

THE RISING SEA

Foundations of Algebraic Geometry

math216.wordpress.com

August 29, 2022 draft

© 2010–2021 by Ravi Vakil.

Note to reader: the index and formatting have yet to be properly dealt with. There remain many issues still to be dealt with in the main part of the notes (including many of your corrections and suggestions).

Contents

Preface	11
0.1. For the reader	12
0.2. For the expert	16
0.3. Background and conventions	17
0.4. $\star\star$ The goals of this book	18
Part I. Preliminaries	23
Chapter 1. Just enough category theory to be dangerous	25
1.1. Motivation	25
1.2. Categories and functors	27
1.3. Universal properties determine an object up to unique isomorphism	33
1.4. Limits and colimits	42
1.5. Adjoint	47
1.6. An introduction to abelian categories	50
1.7. \star Spectral sequences	61
Chapter 2. Sheaves	69
2.1. Motivating example: The sheaf of smooth functions	69
2.2. Definition of sheaf and presheaf	71
2.3. Morphisms of presheaves and sheaves	76
2.4. Properties determined at the level of stalks, and sheafification	80
2.5. Recovering sheaves from a “sheaf on a base”	84
2.6. Sheaves of abelian groups, and \mathcal{O}_X -modules, form abelian categories	87
2.7. The inverse image sheaf	90
Part II. Schemes	95
Chapter 3. Toward affine schemes: the underlying set, and topological space	97
3.1. Toward schemes	97
3.2. The underlying set of an affine scheme	99
3.3. Visualizing schemes: Generic points	111
3.4. The underlying topological space of an affine scheme	113
3.5. A base of the Zariski topology on $\text{Spec } A$: Distinguished open sets	116
3.6. Topological (and Noetherian) properties	117
3.7. The function $I(\cdot)$, taking subsets of $\text{Spec } A$ to ideals of A	125
Chapter 4. The structure sheaf, and the definition of schemes in general	129
4.1. The structure sheaf of an affine scheme	129
4.2. Visualizing schemes: Nilpotents	133

4.3.	Definition of schemes	137
4.4.	Three examples	140
4.5.	Projective schemes, and the Proj construction	146
Chapter 5.	Some properties of schemes	155
5.1.	Topological properties	155
5.2.	Reducedness and integrality	157
5.3.	The Affine Covering Lemma, and properties of schemes that can be checked “affine-locally”	159
5.4.	Normality and factoriality	163
Chapter 6.	Rings are to modules as schemes are to ...	169
6.1.	Quasicoherent sheaves	169
6.2.	Characterizing quasicoherence using the distinguished affine base	172
6.3.	Quasicoherent sheaves form an abelian category	176
6.4.	Finite type quasicoherent, finitely presented, and coherent sheaves	177
6.5.	Visualizing schemes: Associated points and zerodivisors	180
6.6.	★★ Coherent modules over non-Noetherian rings	191
Part III.	Morphisms of schemes	195
Chapter 7.	Morphisms of schemes	197
7.1.	Introduction	197
7.2.	Morphisms of ringed spaces	198
7.3.	From locally ringed spaces to morphisms of schemes	200
7.4.	Maps of graded rings and maps of projective schemes	207
7.5.	Rational maps from reduced schemes	209
7.6.	★ Representable functors and group schemes	216
7.7.	★★ The Grassmannian: First construction	220
Chapter 8.	Useful classes of morphisms of schemes	223
8.1.	“Reasonable” classes of morphisms (such as open embeddings)	223
8.2.	Algebraic interlude: Lying Over and Nakayama	225
8.3.	A gazillion finiteness conditions on morphisms	230
8.4.	Images of morphisms: Chevalley’s Theorem and elimination theory	239
Chapter 9.	Closed embeddings and related notions	247
9.1.	Closed embeddings and closed subschemes	247
9.2.	Locally closed embeddings and locally closed subschemes	250
9.3.	Important examples from projective geometry	252
9.4.	The (closed sub)scheme-theoretic image	259
9.5.	Slicing by effective Cartier divisors, regular sequences and regular embeddings	263
Chapter 10.	Fibered products of schemes, and base change	269
10.1.	They exist	269
10.2.	Computing fibered products in practice	275
10.3.	Interpretations: Pulling back families, and fibers of morphisms	278
10.4.	Properties preserved by base change	284
10.5.	★ Properties not preserved by base change, and how to fix them	286

10.6.	Products of projective schemes: The Segre embedding	294
10.7.	Normalization	296
Chapter 11.	Separated and proper morphisms, and (finally!) varieties	303
11.1.	Motivation	303
11.2.	Fun with diagonal morphisms, and quasiseparatedness made easy	304
11.3.	Separatedness, and varieties	306
11.4.	The locus where two morphisms from X to Y agree, and the “Rational-to-Separated” Theorem	313
11.5.	Proper morphisms	317
Part IV.	“Geometric” properties of schemes	325
Chapter 12.	Dimension	327
12.1.	Dimension and codimension	327
12.2.	Dimension, transcendence degree, and Noether normalization	331
12.3.	Codimension one miracles: Krull and Hartogs	339
12.4.	Dimensions of fibers of morphisms of varieties	347
Chapter 13.	Regularity and smoothness	353
13.1.	The Zariski tangent space	353
13.2.	Regularity, and smoothness over a field	359
13.3.	Examples	365
13.4.	Bertini’s Theorem	368
13.5.	Another (co)dimension one miracle: Discrete valuation rings	371
13.6.	Smooth (and étale) morphisms: First definition	376
13.7.	★ Valuative criteria for separatedness and properness	380
13.8.	★ More sophisticated facts about regular local rings	384
13.9.	★ Filtered rings and modules, and the Artin-Rees Lemma	386
Part V.	Quasicoherent sheaves on schemes, and their uses	389
Chapter 14.	More on quasicoherent and coherent sheaves	391
14.1.	Introduction	391
14.2.	Vector bundles “=” locally free sheaves	391
14.3.	More pleasant properties of finite type and coherent sheaves	400
Chapter 15.	Line bundles: Invertible sheaves and divisors	405
15.1.	Some line bundles on projective space	405
15.2.	Line bundles and Weil divisors	408
15.3.	★ Effective Cartier divisors “=” invertible ideal sheaves	416
Chapter 16.	Quasicoherent sheaves in projective geometry	419
16.1.	The quasicoherent sheaf corresponding to a graded module	419
16.2.	Line bundles on projective A -schemes	420
16.3.	Globally generated and base-point-free line bundles	421
16.4.	Quasicoherent sheaves and graded modules	425
Chapter 17.	Pushforwards and pullbacks of quasicoherent sheaves	429
17.1.	Introduction	429
17.2.	Pushforwards of quasicoherent sheaves	429

17.3.	Pullbacks of quasicoherent sheaves	430
17.4.	Line bundles and maps to projective schemes	437
17.5.	The Curve-to-Projective Extension Theorem	443
17.6.	Ample and very ample line bundles	445
17.7.	★ The Grassmannian as a moduli space	450
Chapter 18.	Relative versions of Spec and Proj, and projective morphisms	455
18.1.	Relative Spec of a (quasicoherent) sheaf of algebras	455
18.2.	Relative Proj of a sheaf of graded algebras	458
18.3.	Projective morphisms	461
18.4.	Applications to curves	467
Chapter 19.	Čech cohomology of quasicoherent sheaves	475
19.1.	(Desired) properties of cohomology	475
19.2.	Definitions and proofs of key properties	480
19.3.	Cohomology of line bundles on projective space	485
19.4.	Riemann-Roch, and arithmetic genus	488
19.5.	A first glimpse of Serre duality	495
19.6.	Hilbert functions, Hilbert polynomials, and genus	498
19.7.	★ Serre's cohomological characterization of ampleness	504
19.8.	Higher pushforward (or direct image) sheaves	507
19.9.	★ From projective to proper hypotheses: Chow's Lemma and Grothendieck's Coherence Theorem	511
Chapter 20.	Application: Curves	517
20.1.	A criterion for a morphism to be a closed embedding	517
20.2.	A series of crucial tools	519
20.3.	Curves of genus 0	522
20.4.	Classical geometry arising from curves of positive genus	523
20.5.	Hyperelliptic curves	526
20.6.	Curves of genus 2	530
20.7.	Curves of genus 3	531
20.8.	Curves of genus 4 and 5	533
20.9.	Curves of genus 1	536
20.10.	Elliptic curves are group varieties	544
20.11.	Counterexamples and pathologies using elliptic curves	550
Chapter 21.	Application: A glimpse of intersection theory	555
21.1.	Intersecting n line bundles with an n -dimensional variety	555
21.2.	Intersection theory on a surface	560
21.3.	The Grothendieck group of coherent sheaves, and an algebraic version of homology	566
21.4.	★★ The Nakai-Moishezon and Kleiman criteria for ampleness	568
Chapter 22.	Differentials	573
22.1.	Motivation and game plan	573
22.2.	Definitions and first properties	574
22.3.	Smoothness of varieties revisited	587
22.4.	Examples	590
22.5.	Understanding smooth varieties using their cotangent bundles	595

22.6.	Unramified morphisms	600
22.7.	The Riemann-Hurwitz Formula	602
Chapter 23.	★ Blowing up	609
23.1.	Motivating example: Blowing up the origin in the plane	609
23.2.	Blowing up, by universal property	611
23.3.	The blow-up exists, and is projective	615
23.4.	Examples and computations	620
Part VI.	More cohomological tools	631
Chapter 24.	Derived functors	633
24.1.	The Tor functors	633
24.2.	Derived functors in general	637
24.3.	Derived functors and spectral sequences	641
24.4.	Derived functor cohomology of \mathcal{O} -modules	646
24.5.	Čech cohomology and derived functor cohomology agree	649
Chapter 25.	Flatness	657
25.1.	Introduction	657
25.2.	Easier facts	659
25.3.	Flatness through Tor	664
25.4.	Ideal-theoretic criteria for flatness	666
25.5.	Topological aspects of flatness	673
25.6.	Local criteria for flatness	678
25.7.	Flatness implies constant Euler characteristic	682
Chapter 26.	Smooth and étale morphisms, and flatness	687
26.1.	Some motivation	687
26.2.	Different characterizations of smooth and étale morphisms	689
26.3.	Generic smoothness and the Kleiman-Bertini Theorem	695
Chapter 27.	Twenty-seven lines	699
27.1.	Introduction	699
27.2.	Preliminary facts	700
27.3.	Every smooth cubic surface (over \bar{k}) contains 27 lines	702
27.4.	Every smooth cubic surface (over \bar{k}) is a blown up plane	706
Chapter 28.	Cohomology and base change theorems	709
28.1.	Statements and applications	709
28.2.	★ Proofs of cohomology and base change theorems	715
28.3.	Applying cohomology and base change to moduli problems	721
Chapter 29.	Depth and Cohen-Macaulayness	727
29.1.	Depth	727
29.2.	Cohen-Macaulay rings and schemes	730
29.3.	★ Serre's $R_1 + S_2$ criterion for normality	733
Chapter 30.	Power series and the Theorem on Formal Functions	739
30.1.	Introduction	739
30.2.	Algebraic preliminaries	739

30.3.	Defining types of singularities	743
30.4.	The Theorem on Formal Functions	746
30.5.	Zariski's Connectedness Lemma and Stein Factorization	748
30.6.	Zariski's Main Theorem	751
30.7.	Castelnuovo's Criterion for contracting (-1) -curves	755
30.8.	★★ Proof of the Theorem on Formal Functions 30.4.2	759
Chapter 31.	★ Proof of Serre duality	763
31.1.	Introduction	763
31.2.	Ext groups and Ext sheaves for \mathcal{O} -modules	768
31.3.	Serre duality for projective k -schemes	772
31.4.	The adjunction formula for the dualizing sheaf, and $\omega_X = \mathcal{K}_X$	777
Bibliography		783

Je pourrais illustrer la ... approche, en gardant l'image de la noix qu'il s'agit d'ouvrir. La première parabole qui m'est venue à l'esprit tantôt, c'est qu'on plonge la noix dans un liquide émollient, de l'eau simplement pourquoi pas, de temps en temps on frotte pour qu'elle pénètre mieux, pour le reste on laisse faire le temps. La coque s'assouplit au fil des semaines et des mois — quand le temps est mûr, une pression de la main suffit, la coque s'ouvre comme celle d'un avocat mûr à point! ...

L'image qui m'était venue il y a quelques semaines était différente encore, la chose inconnue qu'il s'agit de connaître m'apparaissait comme quelque étendue de terre ou de marnes compactes, réticente à se laisser pénétrer. ... La mer s'avance insensiblement et sans bruit, rien ne semble se casser rien ne bouge l'eau est si loin on l'entend à peine... Pourtant elle finit par entourer la substance rétive...

I can illustrate the ... approach with the ... image of a nut to be opened. The first analogy that came to my mind is of immersing the nut in some softening liquid, and why not simply water? From time to time you rub so the liquid penetrates better, and otherwise you let time pass. The shell becomes more flexible through weeks and months — when the time is ripe, hand pressure is enough, the shell opens like a perfectly ripened avocado! ...

A different image came to me a few weeks ago. The unknown thing to be known appeared to me as some stretch of earth or hard marl, resisting penetration ... the sea advances insensibly in silence, nothing seems to happen, nothing moves, the water is so far off you hardly hear it ... yet finally it surrounds the resistant substance.

— A. Grothendieck [Gr6, p. 552-3], translation by C. McLarty [Mc, p. 1]

Preface

This book is intended to give a serious and reasonably complete introduction to algebraic geometry, not just for (future) experts in the field. The exposition serves a narrow set of goals (see §0.4), and necessarily takes a particular point of view on the subject.

It has now been many decades since David Mumford wrote that algebraic geometry “seems to have acquired the reputation of being esoteric, exclusive, and very abstract, with adherents who are secretly plotting to take over all the rest of mathematics! In one respect this last point is accurate ...” ([Mu4, preface] and [Mu7, p. 227]). The revolution has now fully come to pass, and has fundamentally changed how we think about many fields of pure mathematics. A remarkable number of celebrated advances rely in some way on the insights and ideas first articulated by Alexander Grothendieck, Jean-Pierre Serre, and others.

For a number of reasons, algebraic geometry has earned a reputation of being inaccessible. The power of the subject comes from rather abstract heavy machinery, and it is easy to lose sight of the intuitive nature of the objects and methods. Many in nearby fields have only a vague sense of the fundamental ideas of the subject. Algebraic geometry itself has fractured into many parts, and even within algebraic geometry, new researchers are often unaware of the basic ideas in subfields removed from their own.

But there is another more optimistic perspective to be taken. The ideas that allow algebraic geometry to connect several parts of mathematics are fundamental, and well-motivated. Many people in nearby fields would find it useful to develop a working knowledge of the foundations of the subject, and not just at a superficial level. Within algebraic geometry itself, there is a canon (at least for those approaching the subject from this particular direction), that everyone in the field can and should be familiar with. The rough edges of scheme theory have been sanded down over the past half century, although there remains an inescapable need to understand the subject on its own terms.

0.0.1. *The importance of exercises.* This book has a lot of exercises. I have found that unless I have some problems I can think through, ideas don’t get fixed in my mind. Some exercises are trivial — some experts find this offensive, but I find this desirable. A very few necessary ones may be hard, but the reader should have been given the background to deal with them — they are not just an excuse to push hard material out of the text. The exercises are interspersed with the exposition, not left to the end. Most have been extensively field-tested. The point of view here is one I explored with Kedlaya and Poonen in [KPV], a book that was ostensibly about problems, but secretly a case for how one should learn and do and think about mathematics. Most people learn by doing, rather than just passively reading.

Judiciously chosen problems can be the best way of guiding the learner toward enlightenment.

0.0.2. Structure. You will quickly notice that everything is numbered by chapter and section, and everything is numbered the same way after that (for ease of reference), except exercises are indicated by letters (and are sprinkled throughout the text, rather than at the end of sections). Individual paragraphs often get numbers for ease of reference, or to indicate a new topic. Definitions are in bold, and are sometimes given in passing.

0.0.3. Acknowledgments.

This one is going to be really hard, so I'll mostly write this later. Mike Stay is the author of Jokes 1.3.12 and 22.5.2. The phrase “The Rising Sea” is due to Grothendieck [Gr6, p. 552-3], with this particular translation by McLarty [Mc, p. 1], and popularized as the title of Daniel Murfet’s excellent blog [Mur].

0.1 For the reader

This is your last chance. After this, there is no turning back. You take the blue pill, the story ends, you wake up in your bed and believe whatever you want to believe. You take the red pill, you stay in Wonderland and I show you how deep the rabbit-hole goes.
— Morpheus

The contents of this book are intended to be a collection of communal wisdom, necessarily distilled through an imperfect filter. I wish to say a few words on how you might use it, although it is not clear to me if you should or will follow this advice.

Before discussing details, I want to say clearly at the outset: the wonderful machine of modern algebraic geometry was created to understand basic and naive questions about geometry (broadly construed). The purpose of this book is to give you a thorough foundation in these powerful ideas. *Do not be seduced by the lotus-eaters into infatuation with untethered abstraction.* Hold tight to your geometric motivation as you learn the formal structures which have proved to be so effective in studying fundamental questions. When introduced to a new idea, always ask why you should care. Do not expect an answer right away, but demand an answer eventually. Try at least to apply any new abstraction to some concrete example you can understand well.

Understanding algebraic geometry is often thought to be hard because it consists of large complicated pieces of machinery. In fact the opposite is true; to switch metaphors, rather than being narrow and deep, algebraic geometry is shallow but extremely broad. It is built out of a large number of very small parts, in keeping with Grothendieck’s vision of mathematics. It is a challenge to hold the entire organic structure, with its messy interconnections, in your head.

A reasonable place to start is with the idea of “affine complex varieties”: subsets of \mathbb{C}^n cut out by some polynomial equations. Your geometric intuition can immediately come into play — you may already have some ideas or questions about dimension, or smoothness, or solutions over subfields such as \mathbb{R} or \mathbb{Q} . Wiser heads

would counsel spending time understanding complex varieties in some detail before learning about schemes. Instead, I encourage you to learn about schemes immediately, learning about affine complex varieties as the central (but not exclusive) example. This is not ideal, but can save time, and is surprisingly workable. An alternative is to learn about varieties elsewhere, and then come back later.

The intuition for schemes can be built on the intuition for affine complex varieties. Allen Knutson and Terry Tao have pointed out that this involves three different simultaneous generalizations, which can be interpreted as three large themes in mathematics. (i) We allow nilpotents in the ring of functions, which is basically analysis (looking at near-solutions of equations instead of exact solutions). (ii) We glue these affine schemes together, which is what we do in differential geometry (looking at manifolds instead of coordinate patches). (iii) Instead of working over \mathbb{C} (or another algebraically closed field), we work more generally over a ring that isn't an algebraically closed field, or even a field at all, which is basically number theory (solving equations over number fields, rings of integers, etc.).

Because our goal is to be comprehensive, and to understand everything one should know after a first course, it will necessarily take longer to get to interesting sample applications. You may be misled into thinking that one has to work this hard to get to these applications — it is not true! You should deliberately keep an eye out for examples you would have cared about before. This will take some time and patience.

As you learn algebraic geometry, you should pay attention to crucial stepping stones. Of course, the steps get bigger the farther you go.

Chapter 1. Category theory is only language, but it is language with an embedded logic. Category theory is much easier once you realize that it is designed to formalize and abstract things you already know. The initial chapter on category theory prepares you to think cleanly. For example, when someone names something a “cokernel” or a “product”, you should want to know why it deserves that name, and what the name really should mean. The conceptual advantages of thinking this way will gradually become apparent over time. Yoneda’s Lemma — and more generally, the idea of understanding an object through the maps to it — will play an important role.

Chapter 2. The theory of sheaves again abstracts something you already understand well (see the motivating example of §2.1), and what is difficult is understanding how one best packages and works with the information of a sheaf (stalks, sheafification, sheaves on a base, etc.).

Chapters 1 and 2 are a risky gamble, and they attempt a delicate balance. Attempts to explain algebraic geometry often leave such background to the reader, refer to other sources the reader won’t read, or punt it to a telegraphic appendix. Instead, this book attempts to explain everything necessary, but as little as possible, and tries to get across how you should think about (and work with) these fundamental ideas, and why they are more grounded than you might fear.

Chapters 3–5. Armed with this background, you will be able to think cleanly about various sorts of “spaces” studied in different parts of geometry (including real manifolds, topological spaces, and complex manifolds), as ringed spaces that locally are of a certain form. A scheme is just another kind of “geometric space”,

and we are then ready to transport lots of intuition from “classical geometry” to this new setting. (This also will set you up to later think about other geometric kinds of spaces in algebraic geometry, such as complex analytic spaces, algebraic spaces, orbifolds, stacks, rigid analytic spaces, and formal schemes.) The ways in which schemes differ from your geometric intuition can be internalized, and your intuition can be expanded to accommodate them. There are many properties you will realize you will want, as well as other properties that will later prove important. These all deserve names. Take your time becoming familiar with them.

Chapters 7–11. Thinking categorically will lead you to ask about morphisms of schemes (and other spaces in geometry). One of Grothendieck’s fundamental lessons is that the morphisms are central. Important geometric properties should really be understood as properties of morphisms. There are many classes of morphisms with special names, and in each case you should think through why that class deserves a name.

Chapters 12–13. You will then be in a good position to think about fundamental geometric properties of schemes: dimension and smoothness. You may be surprised that these are subtle ideas, but you should keep in mind that they are subtle everywhere in mathematics.

Chapters 14–22. Vector bundles are ubiquitous tools in geometry, and algebraic geometry is no exception. They lead us to the more general notion of quasicoherent sheaves, much as free modules over a ring lead us to modules more generally. We study their properties next, including cohomology. Chapter 20, applying these ideas to study curves, may help make clear how useful they are.

Chapters 24–31. With this in hand, you are ready to learn more advanced tools widely used in the subject. Many examples of what you can do are given, and the classical story of the 27 lines on a smooth cubic surface (Chapter 27) is a good opportunity to see many ideas come together.

The rough logical dependencies among the chapters are shown in Figure 0.1. (Caution: this should be taken with a grain of salt. For example, you can avoid using much of Chapter 19 on curves in later chapters, but it is a crucial source of examples, and a great way to consolidate your understanding. And Chapter 29 on completions uses Chapters 19, 20 and 22 only in the discussion of Castelnuovo’s Criterion 30.7.1.)

In general, I like having as few hypotheses as possible. Certainly a hypothesis that isn’t necessary to the proof is a red herring. But if a reasonable hypothesis can make the proof cleaner and more memorable, I am willing to include it.

In particular, Noetherian hypotheses are handy when necessary, but are otherwise misleading. Even Noetherian-minded readers (normal human beings) are better off having the right hypotheses, as they will make clearer why things are true.

We often state results particular to varieties, especially when there are techniques unique to this situation that one should know. But restricting to algebraically closed fields is useful surprisingly rarely. Geometers needn’t be afraid of arithmetic examples or of algebraic examples; a central insight of algebraic geometry is that the same formalism applies without change.




FIGURE 0.1. Important logical dependences among chapters (or more precisely, a directed graph showing which chapter should be read before which other chapter)

Pathological examples are useful to know. On mountain highways, there are tall sticks on the sides of the road designed for bad weather. In winter, you cannot see the road clearly, and the sticks serve as warning signs: if you cross this line, you will die! Pathologies and (counter)examples serve a similar goal. They also serve as a reality check, when confronting a new statement, theorem, or conjecture, whose veracity you may doubt. (See, for example, §4.1.8.)

When working through a book in algebraic geometry, it is particularly helpful to have other algebraic geometry books at hand, to see different approaches and to have alternate expositions when things become difficult. This book may serve as a good secondary book. If it is your primary source, then two other excellent books with what I consider a similar philosophy are [Liu] and [GW]. De Jong's encyclopedic online reference [Stacks] is peerless. There are many other outstanding sources out there, perhaps one for each approach to the subject; you should browse around and find one you find sympathetic.

If you are looking for a correct or complete history of the subject, you have come to the wrong place. This book is not intended to be a complete guide to the literature, and many important sources are ignored or left out, due to my own ignorance and laziness.

Finally, if you attempt to read this without working through a significant number of exercises (see §0.0.1), I will come to your house and pummel you with [Gr-EGA] until you beg for mercy. It is important to not just have a vague sense of what is true, but to be able to actually get your hands dirty. To quote Mark Kisin: “You can wave your hands all you want, but it still won’t make you fly.” Note: The hints may help you, but sometimes they may not.

0.2 For the expert

If you use this book for a course, you should of course adapt it to your own point of view and your own interests. In particular, you should think about an application or theorem you want to reach at the end of the course (which may well not be in this book), and then work toward it. You should feel no compulsion to sprint to the end; I advise instead taking more time, and ending at the right place for your students. (Figure 0.1, showing large-scale dependencies among the chapters, may help you map out a course.) I have found that the theory of curves (Chapter 20) and the 27 lines on the cubic surface (Chapter 27) have served this purpose well at the end of winter and spring quarters. This was true even if some of the needed background was not covered, and had to be taken by students as some sort of black box.

Faithfulness to the goals of §0.4 required a brutal triage, and I have made a number of decisions you may wish to reverse. I will briefly describe some choices made that may be controversial.

Decisions on how to describe things were made for the sake of the learners. If there were two approaches, and one was “correct” from an advanced point of view, and one was direct and natural from a naive point of view, I went with the latter.

On the other hand, the theory of varieties (over an algebraically closed field, say) was *not* done first and separately. This choice brought me close to tears, but in the end I am convinced that it can work well, if done in the right spirit.

Instead of spending the first part of the course on varieties, I spent the time in a different way. It is tempting to assume that students will either arrive with great comfort and experience with category theory and sheaf theory, or that they should pick up these ideas on their own time. I would love to live in that world. I encourage you to not skimp on these foundational issues. I have found that although these first lectures felt painfully slow to me, they were revelatory to a number of the students, and those with more experience were not bored and did not waste their time. This investment paid off in spades when I was able to rely on their ability to think cleanly and to use these tools in practice. Furthermore, if they left the course with nothing more than hands-on experience with these ideas, the world was still better off for it.

For the most part, we will state results in the maximal generality that the proof justifies, but we will not give a much harder proof if the generality of the stronger result will not be used. There are a few cases where we work harder to prove a somewhat more general result that many readers may not appreciate. For example, we prove a number of theorems for proper morphisms, not just projective

morphisms. But in such cases, readers are invited or encouraged to ignore the subtleties required for the greater generality.

I consider line bundles (and maps to projective space) more fundamental than divisors. General Cartier divisors are not discussed (although *effective* Cartier divisors play an essential role).

Cohomology is done first using the Čech approach (as Serre first did), and derived functor cohomology is introduced only later. I am well aware that Grothendieck thinks of the fact that the agreement of Čech cohomology with derived functor cohomology “should be considered as an accidental phenomenon”, and that “it is important for technical reasons not to take as *definition* of cohomology the Čech cohomology”, [Gr4, p. 108]. But I am convinced that this is the right way for most people to see this kind of cohomology *for the first time*. (It is certainly true that many topics in algebraic geometry are best understood in the language of derived functors. But this is a view from the mountaintop, looking down, and not the best way to explore the forests. In order to appreciate derived functors appropriately, one must understand the homological algebra behind it, and not just take it as a black box.)

We restrict to the Noetherian case only when it is necessary, or (rarely) when it really saves effort. In this way, non-Noetherian people will clearly see where they should be careful, and Noetherian people will realize that non-Noetherian things are not so terrible. Moreover, even if you are interested primarily in Noetherian schemes, it helps to see “Noetherian” in the hypotheses of theorems only when necessary, as it will help you remember how and when this property gets used.

There are some cases where Noetherian readers will suffer a little more than they would otherwise. As an inflammatory example, instead of using Noetherian hypotheses, the notion of quasiseparatedness comes up early and often. The cost is that one extra word has to be remembered, on top of an overwhelming number of other words. But once that is done, it is not hard to remember that essentially every scheme anyone cares about is quasiseparated. Furthermore, whenever the hypotheses “quasicompact and quasiseparated” turn up, the reader will immediately guess a key idea of the proof. As another example, coherent sheaves and finite type (quasicohherent) sheaves are the same in the Noetherian situation, but are still worth distinguishing in statements of the theorems and exercises, for the same reason: to be clearer on what is used in the proof.

Many important topics are not discussed. Valuable criteria are not proved (see §13.7), and their statement is relegated to an optional section. Completely omitted: dévissage, formal schemes, and cohomology with supports. Sorry!

0.3 Background and conventions

“Should you just be an algebraist or a geometer?” is like saying “Would you rather be deaf or blind?”

— M. Atiyah, [At2, p. 659]

All rings are assumed to be commutative unless explicitly stated otherwise. All rings are assumed to contain a unit, denoted 1. Maps of rings must send 1 to 1. We don’t require that $0 \neq 1$; in other words, the “0-ring” (with one element)

is a ring. (There is a ring map from any ring to the 0-ring; the 0-ring only maps to itself. The 0-ring is the final object in the category of rings.) The definition of “integral domain” includes $1 \neq 0$, so the 0-ring is not an integral domain. We accept the Axiom of Choice. In particular, any proper ideal in a ring is contained in a maximal ideal. (The Axiom of Choice also arises in the argument that the category of A -modules has enough injectives, see Exercise 24.2.G.)

The reader should be familiar with some basic notions in commutative ring theory, in particular the notion of ideals (including prime and maximal ideals), various types of rings (including integral domains, principal ideal domains, unique factorization domains, and local rings), localization, and modules. Tensor products and exact sequences of A -modules will be important. We will use the notation (A, \mathfrak{m}) or (A, \mathfrak{m}, k) for local rings (rings with a unique maximal ideal) — A is the ring, \mathfrak{m} its maximal ideal, and $k = A/\mathfrak{m}$ its residue field. We will use the structure theorem for finitely generated modules over a principal ideal domain A : any such module can be written as the direct sum of principal modules $A/(a)$. Some experience with field theory will be important from time to time.

Manifolds will be brought up periodically as examples, but for the most part, they are meant for motivation, so we will often not specify whether they are topological (real) manifolds, differentiable (C^∞ , i.e., smooth) real manifolds, analytic (real) manifolds, or complex (holomorphic) manifolds. Nonetheless, we will define all four (Definition 4.3.9).

0.3.1. Caution about foundational issues. We will not concern ourselves with subtle foundational issues (set-theoretic issues, universes, etc.). It is true that some people should be careful about these issues. But is that really how you want to live your life? (If you are one of these rare people, a good start is [KS2, §1.1].)

0.3.2. Further background. It may be helpful to have books on other subjects at hand that you can dip into for specific facts, rather than reading them in advance. In commutative algebra, [E] is good for this. Other popular choices are [AtM] and [Mat2]. The book [AI] takes a point of view useful to algebraic geometry. For homological algebra, [Weib] is simultaneously detailed and readable.

Background from other parts of mathematics (topology, geometry, complex analysis, number theory, ...) will of course be helpful for intuition and grounding. Some previous exposure to topology is certainly essential.

0.3.3. Nonmathematical conventions. “Unimportant” means “unimportant for the current exposition”, *not* necessarily unimportant in the larger scheme of things. Other words may be used idiosyncratically as well.

There are optional starred sections of topics worth knowing on a second or third (but not first) reading. They are marked with a star: \star . Starred sections are not necessarily harder, merely unimportant. You should not read double-starred sections ($\star\star$) unless you really really want to, but you should be aware of their existence. (It may be strange to have parts of a book that should *not* be read!)

Let’s now find out if you are taking my advice about double-starred sections.

0.4 $\star\star$ The goals of this book

There are a number of possible introductions to the field of algebraic geometry: Riemann surfaces; complex geometry; the theory of varieties; a nonrigorous examples-based introduction; algebraic geometry for number theorists; an abstract functorial approach; and more. All have their place. Different approaches suit different students (and different advisors). This book takes only one route.

Our intent is to cover a canon completely and rigorously, with enough examples and calculations to help develop intuition for the machinery. This is often the content of a second course in algebraic geometry, and in an ideal world, people would learn this material over many years, after having background courses in commutative algebra, algebraic topology, differential geometry, complex analysis, homological algebra, number theory, and French literature. We do not live in an ideal world. For this reason, the book is written as a first introduction, but a challenging one.

This book seeks to do a very few things, but to try to do them well. Our goals and premises are as follows.

The core of the material should be digestible over a single year. After a year of blood, sweat, and tears, readers should have a broad familiarity with the foundations of the subject, and be ready to attend seminars, and learn more advanced material. They should not just have a vague intuitive understanding of the ideas of the subject; they should know interesting examples, know why they are interesting, and be able to work through their details. Readers in other fields of mathematics should know enough to understand the algebro-geometric ideas that arise in their area of interest.

This means that this book is not encyclopedic, and even beyond that, hard choices have to be made. (In particular, analytic aspects are essentially ignored, and are at best dealt with in passing without proof. This is a book about *algebraic* algebraic geometry.)

This book is usable (and has been used) for a course, but the course should (as always) take on the personality of the instructor. With a good course, people should be able to leave early and still get something useful from the experience. With this book, it is possible to leave without regret after learning about category theory, or about sheaves, or about geometric spaces, having become a better person.

The book is also usable (and has been used) for learning on your own. But most mortals cannot learn algebraic geometry fully on their own; ideally you should read in a group, and even if not, you should have someone you can ask questions to (both stupid and smart questions).

There is certainly more than a year's material here, but I have tried to make clear which topics are essential, and which are not. Those teaching a class will choose which "inessential" things are important for the point they wish to get across, and use them.

There is a canon (at least for this particular approach to algebraic geometry). I have been repeatedly surprised at how much people in different parts of algebraic geometry agree on what every civilized algebraic geometer should know after a first (serious) year. (There are of course different canons for different parts of the subject, e.g., complex algebraic geometry, combinatorial algebraic geometry, computational algebraic geometry, etc.) There are extra bells and whistles that different

instructors might add on, to prepare students for their particular part of the field or their own point of view, but the core of the subject remains unified, despite the diversity and richness of the subject. There are some serious and painful compromises to be made to reconcile this goal with the previous one.

Algebraic geometry is for everyone (with the appropriate definition of “everyone”). Algebraic geometry courses tend to require a lot of background, which makes them inaccessible to all but those who know they will go deeply into the subject. Algebraic geometry is too important for that; it is essential that many of those in nearby fields develop some serious familiarity with the foundational ideas and tools of the subject, and not just at a superficial level. (Similarly, algebraic geometers uninterested in any nearby field are necessarily arid, narrow thinkers. Do not be such a person!)

For this reason, this book attempts to require as little background as possible. The background required will, in a technical sense, be surprisingly minimal — ideally just some commutative ring theory and point-set topology, and some comfort with things like prime ideals and localization. This is misleading of course — the more you know, the better. And the less background you have, the harder you will have to work — this is not a light read. On a related note...

The book is intended to be as self-contained as possible. I have tried to follow the motto: “if you use it, you must prove it”. I have noticed that most students are human beings: if you tell them that some algebraic fact is in some late chapter of a book in commutative algebra, they will not immediately go and read it. Surprisingly often, what we need can be developed quickly from scratch, and even if people do not read it, they can see what is involved. The cost is that the book is much denser, and that significant sophistication and maturity is demanded of the reader. The benefit is that more people can follow it; they are less likely to reach a point where they get thrown. On the other hand, people who already have some familiarity with algebraic geometry, but want to understand the foundations more completely, should not be bored, and can focus on more subtle issues.

As just one example, Krull’s Principal Ideal Theorem [12.3.3](#) is an important tool. I have included an essentially standard proof. I do not want people to read it (unless they really really want to), and signal this by a double-star in the title: $\star\star$. Instead, I want people to skim it and realize that they *could* read it, and that it is not seriously hard.

This is an important goal because it is important not just to know what is true, but to know why things are true, and what is hard, and what is not hard. Also, this helps the previous goal, by reducing the number of prerequisites.

The book is intended to build intuition for the formidable machinery of algebraic geometry. The exercises are central for this (see [§0.0.1](#)). Informal language can sometimes be helpful. Many examples are given. Learning how to think cleanly (and in particular categorically) is essential. The advantages of *appropriate* generality should be made clear by example, and not by intimidation. The motivation is more local than global. For example, there is no introductory chapter explaining why one might be interested in algebraic geometry, and instead there is an introductory chapter explaining why you should want to think categorically (and how to actually do this).

Balancing the above goals is already impossible. We must thus give up any hope of achieving any other desiderata. **There are no other goals.**

Part I

Preliminaries

CHAPTER 1

Just enough category theory to be dangerous

Was mich nicht umbringt macht mich stärker.

That which does not kill me, makes me stronger.

— F. Nietzsche [N, aphorism number 8]

1.1 Motivation

The introduction of the digit 0 or the group concept was general nonsense too, and mathematics was more or less stagnating for thousands of years because nobody was around to take such childish steps...

— A. Grothendieck, [BP, p. 4–5]

Before we get to any interesting geometry, we need to develop a language to discuss things cleanly and effectively. This is best done in the language of categories. There is not much to know about categories to get started; it is just a very useful language. Like all mathematical languages, category theory comes with an embedded logic, which allows us to abstract intuitions in settings we know well to far more general situations.

Our motivation is as follows. We will be creating some new mathematical objects (such as schemes, and certain kinds of sheaves), and we expect them to act like objects we have seen before. We could try to nail down precisely what we mean by “act like”, and what minimal set of things we have to check in order to verify that they act the way we expect. Fortunately, we don’t have to — other people have done this before us, by defining key notions, such as *abelian categories*, which behave like modules over a ring.


Our general approach will be as follows. I will try to tell you what you need to know, and no more. (This I promise: if I use the word “topoi”, you can shoot me.) I will begin by telling you things you already know, and describing what is essential about the examples, in a way that we can abstract a more general definition. We will then see this definition in less familiar settings, and get comfortable with using it to solve problems and prove theorems.

For example, we will define the notion of *product* of schemes. We could just give a definition of product, but then you should want to know why this precise definition deserves the name of “product”. As a motivation, we revisit the notion of product in a situation we know well: (the category of) sets. One way to define the product of sets U and V is as the set of ordered pairs $\{(u, v) : u \in U, v \in V\}$. But someone from a different mathematical culture might reasonably define it as


the set of symbols $\{\frac{u}{v} : u \in U, v \in V\}$. These notions are “obviously the same”. Better: there is “an obvious bijection between the two”.

This can be made precise by giving a better definition of product, in terms of a *universal property*. Given two sets M and N , a product is a set P , along with maps $\mu: P \rightarrow M$ and $\nu: P \rightarrow N$, such that for *any set* P' with maps $\mu': P' \rightarrow M$ and $\nu': P' \rightarrow N$, these maps must factor *uniquely* through P :

(1.1.0.1)




(The symbol \exists means “there exists”, and the symbol $!$ means “unique”.) Thus a **product** is a *diagram*




and not just a set P , although the maps μ and ν are often left implicit.

This definition agrees with the traditional definition, with one twist: there isn’t just a single product; but any two products come with a *unique* isomorphism between them. In other words, the product is unique up to unique isomorphism. Here is why: if you have a product




and I have a product



then by the universal property of my product (letting (P_2, μ_2, ν_2) play the role of (P, μ, ν) , and (P_1, μ_1, ν_1) play the role of (P', μ', ν') in (1.1.0.1)), there is a unique map $f: P_1 \rightarrow P_2$ making the appropriate diagram commute (i.e., $\mu_1 = \mu_2 \circ f$ and $\nu_1 = \nu_2 \circ f$). Similarly by the universal property of your product, there is a unique map $g: P_2 \rightarrow P_1$ making the appropriate diagram commute. Now consider the universal property of my product, this time letting (P_2, μ_2, ν_2) play the role of both (P, μ, ν) and (P', μ', ν') in (1.1.0.1). There is a unique map $h: P_2 \rightarrow P_2$ such

that



commutes. However, I can name two such maps: the identity map id_{P_2} , and $f \circ g$. Thus $f \circ g = \text{id}_{P_2}$. Similarly, $g \circ f = \text{id}_{P_1}$. Thus the maps f and g arising from the universal property are bijections. In short, there is a unique bijection between P_1 and P_2 preserving the “product structure” (the maps to M and N). This gives us the right to name any such product $M \times N$, since any two such products are uniquely identified.

This definition has the advantage that it works in many circumstances, and once we define categories, we will soon see that the above argument applies verbatim in any category to show that products, if they exist, are unique up to unique isomorphism. Even if you haven’t seen the definition of category before, you can verify that this agrees with your notion of product in some category that you have seen before (such as the category of vector spaces, or the category of manifolds).

This is handy even in cases that you understand. For example, one way of defining the product of two manifolds M and N is to cut them both up into charts, then take products of charts, then glue them together. But if I cut up the manifolds in one way, and you cut them up in another, how do we know our resulting manifolds are the “same”? We could wave our hands, or make an annoying argument about refining covers, but instead, we should just show that they are “categorical products” and hence canonically the “same” (i.e., isomorphic). We will formalize this argument in §1.3.

Another set of notions we will abstract are categories that “behave like modules”. We will want to define kernels and cokernels for new notions, and we should make sure that these notions behave the way we expect them to. This leads us to the definition of *abelian categories*, first defined by Grothendieck in his Tôhoku paper [Gr1].

In this chapter, we will give an informal introduction to these and related notions, in the hope of giving just enough familiarity to comfortably use them in practice.

1.2 Categories and functors

Before functoriality, people lived in caves. — B. Conrad

We begin with an informal definition of categories and functors.

1.2.1. Categories.

A **category** consists of a collection of **objects**, and for each pair of objects, a set of **morphisms** (or **arrows**) between them. (For experts: technically, this is the definition of a *locally small category*. In the correct definition, the morphisms need

only form a class, not necessarily a set, but see Caution 0.3.1.) Morphisms are often informally called **maps**. The collection of objects of a category \mathcal{C} is often denoted $\text{obj}(\mathcal{C})$, but we will usually denote the collection also by \mathcal{C} . If $A, B \in \mathcal{C}$, then the set of morphisms from A to B is denoted $\text{Mor}(A, B)$. A morphism is often written $f: A \rightarrow B$, and A is said to be the **source** of f , and B the **target** of f . (Of course, $\text{Mor}(A, B)$ is taken to be disjoint from $\text{Mor}(A', B')$ unless $A = A'$ and $B = B'$.)

Morphisms compose as expected: there is a composition $\text{Mor}(B, C) \times \text{Mor}(A, B) \rightarrow \text{Mor}(A, C)$, and if $f \in \text{Mor}(A, B)$ and $g \in \text{Mor}(B, C)$, then their composition is denoted $g \circ f$. Composition is associative: if $f \in \text{Mor}(A, B)$, $g \in \text{Mor}(B, C)$, and $h \in \text{Mor}(C, D)$, then $h \circ (g \circ f) = (h \circ g) \circ f$. For each object $A \in \mathcal{C}$, there is always an **identity morphism** $\text{id}_A: A \rightarrow A$, such that when you (left- or right-)compose a morphism with the identity, you get the same morphism. More precisely, for any morphisms $f: A \rightarrow B$ and $g: B \rightarrow C$, $\text{id}_B \circ f = f$ and $g \circ \text{id}_B = g$. (If you wish, you may check that “identity morphisms are unique”: there is only one morphism deserving the name id_A .) This ends the definition of a category.

We have a notion of **isomorphism** between two objects of a category (a morphism $f: A \rightarrow B$ such that there exists some — necessarily unique — morphism $g: B \rightarrow A$, where $f \circ g$ and $g \circ f$ are the identity on B and A respectively), and a notion of **automorphism** of an object (an isomorphism of the object with itself).

1.2.2. Example. The prototypical example to keep in mind is the category of sets, denoted *Sets*. The objects are sets, and the morphisms are maps of sets. (Because Russell’s paradox shows that there is no set of all sets, we did not say earlier that there is a set of all objects. But as stated in §0.3, we are deliberately omitting all set-theoretic issues.)

1.2.3. Example. Another good example is the category Vec_k of vector spaces over a given field k . The objects are k -vector spaces, and the morphisms are linear transformations. (What are the isomorphisms?)

1.2.A. UNIMPORTANT EXERCISE. A category in which each morphism is an isomorphism is called a **groupoid**. (This notion is not important in what we will discuss. The point of this exercise is to give you some practice with categories, by relating them to an object you know well.)

(a) A perverse definition of a *group* is: a groupoid with one object. Make sense of this. (Similarly, in case you care: a perverse definition of a **monoid** is: a category with one object.)

(b) Describe a groupoid that is not a group.

1.2.B. EXERCISE. If A is an object in a category \mathcal{C} , show that the invertible elements of $\text{Mor}(A, A)$ form a group (called the **automorphism group of A** , denoted $\text{Aut}(A)$). What are the automorphism groups of the objects in Examples 1.2.2 and 1.2.3? Show that two isomorphic objects have isomorphic automorphism groups. (For readers with a topological background: if X is a topological space, then the fundamental groupoid is the category where the objects are points of X , and the morphisms $x \rightarrow y$ are paths from x to y , up to homotopy. Then the automorphism group of x_0 is the (pointed) fundamental group $\pi_1(X, x_0)$. In the case where X is connected, and $\pi_1(X)$ is not abelian, this illustrates the fact that for

a connected groupoid — whose definition you can guess — the automorphism groups of the objects are all isomorphic, but not canonically isomorphic.)

1.2.4. Example: abelian groups. The abelian groups, along with group homomorphisms, form a category Ab .

1.2.5. Important Example: Modules over a ring. If A is a ring, then the A -modules form a category Mod_A . (This category has additional structure; it will be the prototypical example of an *abelian category*, see §1.6.) Taking $A = k$, we obtain Example 1.2.3; taking $A = \mathbb{Z}$, we obtain Example 1.2.4.

1.2.6. Example: rings. There is a category $Rings$, where the objects are rings, and the morphisms are maps of rings in the usual sense (maps of sets which respect addition and multiplication, and which send 1 to 1 by our conventions, §0.3).

1.2.7. Example: topological spaces. The topological spaces, along with continuous maps, form a category Top . The isomorphisms are homeomorphisms.

In all of the above examples, the objects of the categories were in obvious ways sets with additional structure (a **concrete category**, although we won't use this terminology). This needn't be the case, as the next example shows.


1.2.8. Example: partially ordered sets. A **partially ordered set**, (or **poset**), is a set S along with a binary relation \geq on S satisfying:

- (i) $x \geq x$ (reflexivity),
- (ii) $x \geq y$ and $y \geq z$ imply $x \geq z$ (transitivity), and
- (iii) if $x \geq y$ and $y \geq x$ then $x = y$ (antisymmetry).

A partially ordered set (S, \geq) can be interpreted as a category whose objects are the elements of S , and with a single morphism from x to y if and only if $x \geq y$ (and no morphism otherwise).


A trivial example is (S, \geq) where $x \geq y$ if and only if $x = y$. Another example is

(1.2.8.1)




Here there are three objects. The identity morphisms are omitted for convenience, and the two non-identity morphisms are depicted. A third example is

(1.2.8.2)



Here the “obvious” morphisms are again omitted: the identity morphisms, and the morphism from the upper left to the lower right. Similarly,



depicts a partially ordered set, where again, only the “generating morphisms” are depicted.

1.2.9. Example: *the category of subsets of a set, and the category of open subsets of a topological space.* If X is a set, then the subsets form a partially ordered set, where the order is given by inclusion. Informally, if $U \subset V$, then we have exactly one morphism $U \rightarrow V$ in the category (and otherwise none). Similarly, if X is a topological space, then the *open* sets form a partially ordered set, where the order is given by inclusion.

1.2.10. Definition. A **subcategory** \mathcal{A} of a category \mathcal{B} has as its objects some of the objects of \mathcal{B} , and some of the morphisms, such that the morphisms of \mathcal{A} include the identity morphisms of the objects of \mathcal{A} , and are closed under composition. (For example, (1.2.8.1) is in an obvious way a subcategory of (1.2.8.2). Also, we have an obvious “inclusion functor” $i: \mathcal{A} \rightarrow \mathcal{B}$.)

1.2.11. Functors.

A **covariant functor** F from a category \mathcal{A} to a category \mathcal{B} , denoted $F: \mathcal{A} \rightarrow \mathcal{B}$, is the following data. It is a map of objects $F: \text{obj}(\mathcal{A}) \rightarrow \text{obj}(\mathcal{B})$, and for each $A_1, A_2 \in \mathcal{A}$, and morphism $m: A_1 \rightarrow A_2$, a morphism $F(m): F(A_1) \rightarrow F(A_2)$ in \mathcal{B} . We require that F preserves identity morphisms (for $A \in \mathcal{A}$, $F(\text{id}_A) = \text{id}_{F(A)}$), and that F preserves composition ($F(m_2 \circ m_1) = F(m_2) \circ F(m_1)$). (You may wish to verify that covariant functors send isomorphisms to isomorphisms.) A trivial example is the **identity functor** $\text{id}: \mathcal{A} \rightarrow \mathcal{A}$, whose definition you can guess. Here are some less trivial examples.

1.2.12. Example: a forgetful functor. Consider the functor from the category of vector spaces (over a field k) Vec_k to *Sets*, that associates to each vector space its underlying set. The functor sends a linear transformation to its underlying map of sets. This is an example of a **forgetful functor**, where some additional structure is forgotten. Another example of a forgetful functor is $\text{Mod}_A \rightarrow \text{Ab}$ from A -modules to abelian groups, remembering only the abelian group structure of the A -module.

1.2.13. Topological examples. Examples of covariant functors include the fundamental group functor π_1 , which sends a topological space X with choice of a point $x_0 \in X$ to a group $\pi_1(X, x_0)$ (what are the objects and morphisms of the source category?), and the i th homology functor $\text{Top} \rightarrow \text{Ab}$, which sends a topological space X to its i th homology group $H_i(X, \mathbb{Z})$. The covariance corresponds to the fact that a (continuous) morphism of pointed topological spaces $\phi: X \rightarrow Y$ with $\phi(x_0) = y_0$ induces a map of fundamental groups $\pi_1(X, x_0) \rightarrow \pi_1(Y, y_0)$, and similarly for homology groups.

1.2.14. Example. Suppose A is an object in a category \mathcal{C} . Then there is a functor $h^A: \mathcal{C} \rightarrow \text{Sets}$ sending $B \in \mathcal{C}$ to $\text{Mor}(A, B)$, and sending $f: B_1 \rightarrow B_2$ to $\text{Mor}(A, B_1) \rightarrow \text{Mor}(A, B_2)$ described by

$$[g: A \rightarrow B_1] \mapsto [f \circ g: A \rightarrow B_1 \rightarrow B_2].$$

This seemingly silly functor ends up surprisingly being an important concept.

1.2.15. Definitions. If $F: \mathcal{A} \rightarrow \mathcal{B}$ and $G: \mathcal{B} \rightarrow \mathcal{C}$ are covariant functors, then we define a functor $G \circ F: \mathcal{A} \rightarrow \mathcal{C}$ (the **composition** of G and F) in the obvious way. Composition of functors is associative in an evident sense.

A covariant functor $F: \mathcal{A} \rightarrow \mathcal{B}$ is **faithful** if for all $A, A' \in \mathcal{A}$, the map $\text{Mor}_{\mathcal{A}}(A, A') \rightarrow \text{Mor}_{\mathcal{B}}(F(A), F(A'))$ is injective, and **full** if it is surjective. A functor that is full and faithful is **fully faithful**. A subcategory $i: \mathcal{A}' \rightarrow \mathcal{A}$ is a **full subcategory** if i is full. (Inclusions are always faithful, so there is no need for the phrase “faithful subcategory”.) Thus a subcategory \mathcal{A}' of \mathcal{A} is full if and only if for all $A, B \in \text{obj}(\mathcal{A}')$, $\text{Mor}_{\mathcal{A}'}(A, B) = \text{Mor}_{\mathcal{A}}(A, B)$. For example, the forgetful functor $\text{Vec}_k \rightarrow \text{Sets}$ is faithful, but not full; and if A is a ring, the category of finitely generated A -modules is a full subcategory of the category Mod_A of A -modules.

1.2.16. Definition. A **contravariant functor** is defined in the same way as a covariant functor, except the arrows switch directions: in the above language, $F(A_1 \rightarrow A_2)$ is now an arrow from $F(A_2)$ to $F(A_1)$. (Thus $F(m_2 \circ m_1) = F(m_1) \circ F(m_2)$, not $F(m_2) \circ F(m_1)$.)

It is wise to state whether a functor is covariant or contravariant, unless the context makes it very clear. If it is not stated (and the context does not make it clear), the functor is often assumed to be covariant.

Sometimes people describe a contravariant functor $\mathcal{C} \rightarrow \mathcal{D}$ as a covariant functor $\mathcal{C}^{\text{opp}} \rightarrow \mathcal{D}$, where \mathcal{C}^{opp} is the same category as \mathcal{C} except that the arrows go in the opposite direction. Here \mathcal{C}^{opp} is said to be the **opposite category** to \mathcal{C} .

One can define fullness, etc. for contravariant functors, and you should do so.

1.2.17. Linear algebra example. If Vec_k is the category of k -vector spaces (introduced in Example 1.2.3), then taking duals gives a contravariant functor $(\cdot)^{\vee}: \text{Vec}_k \rightarrow \text{Vec}_k$. Indeed, to each linear transformation $f: V \rightarrow W$, we have a dual transformation $f^{\vee}: W^{\vee} \rightarrow V^{\vee}$, and $(f \circ g)^{\vee} = g^{\vee} \circ f^{\vee}$.

1.2.18. Topological example (cf. Example 1.2.13) *for those who have seen cohomology.* The i th cohomology functor $H^i(\cdot, \mathbb{Z}): \text{Top} \rightarrow \text{Ab}$ is a contravariant functor.

1.2.19. Example. There is a contravariant functor $\text{Top} \rightarrow \text{Rings}$ taking a topological space X to the ring of real-valued continuous functions on X . A morphism of topological spaces $X \rightarrow Y$ (a continuous map) induces the pullback map from functions on Y to functions on X .

1.2.20. Example (the functor of points, cf. Example 1.2.14). Suppose A is an object of a category \mathcal{C} . Then there is a contravariant functor $h_A: \mathcal{C} \rightarrow \text{Sets}$ sending $B \in \mathcal{C}$ to $\text{Mor}(B, A)$, and sending the morphism $f: B_1 \rightarrow B_2$ to the morphism $\text{Mor}(B_2, A) \rightarrow \text{Mor}(B_1, A)$ via

$$[g: B_2 \rightarrow A] \mapsto [g \circ f: B_1 \rightarrow B_2 \rightarrow A].$$

This example initially looks weird and different, but Examples 1.2.17 and 1.2.19 may be interpreted as special cases; do you see how? What is A in each case? This functor might reasonably be called the *functor of maps* (to A), but is actually known as the **functor of points**. We will meet this functor again in §1.3.11 and (in the category of schemes) in Definition 7.3.10.

1.2.21. ★ Natural transformations (and natural isomorphisms) of covariant functors, and equivalences of categories.

(This notion won’t come up in an essential way until at least Chapter 7, so you shouldn’t read this section until then.) Suppose F and G are two covariant functors

from \mathcal{A} to \mathcal{B} . A **natural transformation of covariant functors** $F \rightarrow G$ is the data of a morphism $m_A: F(A) \rightarrow G(A)$ for each $A \in \mathcal{A}$ such that for each $f: A \rightarrow A'$ in \mathcal{A} , the diagram

$$\begin{array}{ccc} F(A) & \xrightarrow{F(f)} & F(A') \\ m_A \downarrow & & \downarrow m_{A'} \\ G(A) & \xrightarrow{G(f)} & G(A') \end{array}$$

commutes. A **natural isomorphism** of functors is a natural transformation such that each m_A is an isomorphism. (We make analogous definitions when F and G are both contravariant.)

The data of functors $F: \mathcal{A} \rightarrow \mathcal{B}$ and $F': \mathcal{B} \rightarrow \mathcal{A}$ such that $F \circ F'$ is naturally isomorphic to the identity functor $\text{id}_{\mathcal{B}}$ on \mathcal{B} and $F' \circ F$ is naturally isomorphic to $\text{id}_{\mathcal{A}}$ is said to be an **equivalence of categories**. The right notion of when two categories are “essentially the same” is not *isomorphism* (a functor giving bijections of objects and morphisms) but *equivalence*. Exercises 1.2.C and 1.2.D might give you some vague sense of this. Later exercises (for example, that “rings” and “affine schemes” are essentially the same, once arrows are reversed, Exercise 7.3.E) may help too.

Two examples might make this strange concept more comprehensible. The double dual of a finite-dimensional vector space V is *not* V , but we learn early to say that it is canonically isomorphic to V . We can make that precise as follows. Let $f.d.\text{Vec}_k$ be the category of finite-dimensional vector spaces over k . Note that this category contains oodles of vector spaces of each dimension.

1.2.C. EXERCISE. Let $(\cdot)^{\vee\vee}: f.d.\text{Vec}_k \rightarrow f.d.\text{Vec}_k$ be the double dual functor from the category of finite-dimensional vector spaces over k to itself. Show that $(\cdot)^{\vee\vee}$ is naturally isomorphic to the identity functor on $f.d.\text{Vec}_k$. (Without the finite-dimensionality hypothesis, we only get a natural transformation of functors from id to $(\cdot)^{\vee\vee}$.)

Let \mathcal{V} be the category whose objects are the k -vector spaces k^n for each $n \geq 0$ (there is one vector space for each n), and whose morphisms are linear transformations. The objects of \mathcal{V} can be thought of as vector spaces with bases, and the morphisms as matrices. There is an obvious functor $\mathcal{V} \rightarrow f.d.\text{Vec}_k$, as each k^n is a finite-dimensional vector space.

1.2.D. EXERCISE. Show that $\mathcal{V} \rightarrow f.d.\text{Vec}_k$ gives an equivalence of categories, by describing an “inverse” functor. (Recall that we are being cavalier about set-theoretic assumptions, see Caution 0.3.1, so feel free to simultaneously choose bases for each vector space in $f.d.\text{Vec}_k$. To make this precise, you will need to use Gödel-Bernays set theory or else replace $f.d.\text{Vec}_k$ with a very similar small category, but we won’t worry about this.)

1.2.22. ★★ *Aside for experts.* Your argument for Exercise 1.2.D will show that (modulo set-theoretic issues) this definition of equivalence of categories is the same as another one commonly given: a covariant functor $F: \mathcal{A} \rightarrow \mathcal{B}$ is an equivalence of categories if it is fully faithful and every object of \mathcal{B} is isomorphic to an object of the form $F(A)$ for some $A \in \mathcal{A}$ (F is **essentially surjective**, a term we will not

need). Indeed, one can show that such a functor has a **quasiinverse** (another term we will not use later), i.e., a functor $G: \mathcal{B} \rightarrow \mathcal{A}$ (necessarily also an equivalence and unique up to unique isomorphism) for which $G \circ F \cong \text{id}_{\mathcal{A}}$ and $F \circ G \cong \text{id}_{\mathcal{B}}$, and conversely, any functor that has a quasiinverse is an equivalence.

1.3 Universal properties determine an object up to unique isomorphism

Given some category that we come up with, we often will have ways of producing new objects from old. In good circumstances, such a definition can be made using the notion of a *universal property*. Informally, we wish that there were an object with some property. We first show that if it exists, then it is essentially unique, or more precisely, is unique up to unique isomorphism. Then we go about constructing an example of such an object to show existence.

Explicit constructions are sometimes easier to work with than universal properties, but with a little practice, universal properties are useful in proving things quickly and slickly. Indeed, when learning the subject, people often find explicit constructions more appealing, and use them more often in proofs, but as they become more experienced, they find universal property arguments more elegant and insightful.

1.3.1. Products were defined by a universal property. We have seen one important example of a universal property argument already in §1.1: products. You should go back and verify that our discussion there gives a notion of product in any category, and shows that products, *if they exist*, are unique up to unique isomorphism.

1.3.2. Initial, final, and zero objects. Here are some simple but useful concepts that will give you practice with universal property arguments. An object of a category \mathcal{C} is an **initial object** if it has precisely one map to every object. It is a **final object** if it has precisely one map from every object. It is a **zero object** if it is both an initial object and a final object.

1.3.A. EXERCISE. Show that any two initial objects are uniquely isomorphic. Show that any two final objects are uniquely isomorphic.

In other words, *if* an initial object exists, it is unique up to unique isomorphism, and similarly for final objects. This (partially) justifies the phrase “*the* initial object” rather than “*an* initial object”, and similarly for “*the* final object” and “*the* zero object”. (Convention: we often say “*the*”, not “*a*”, for anything defined up to unique isomorphism.)

1.3.B. EXERCISE. What are the initial and final objects in *Sets*, *Rings*, and *Top* (if they exist)? How about in the two examples of §1.2.9?

1.3.3. Localization of rings and modules. Another important example of a definition by universal property is the notion of *localization* of a ring. We first review a constructive definition, and then reinterpret the notion in terms of universal property. A **multiplicative subset** S of a ring A is a subset closed under multiplication

containing 1. We define a ring $S^{-1}A$. The elements of $S^{-1}A$ are of the form a/s where $a \in A$ and $s \in S$, and where $a_1/s_1 = a_2/s_2$ if (and only if) *for some* $s \in S$, $s(s_2a_1 - s_1a_2) = 0$. We define $(a_1/s_1) + (a_2/s_2) = (s_2a_1 + s_1a_2)/(s_1s_2)$, and $(a_1/s_1) \times (a_2/s_2) = (a_1a_2)/(s_1s_2)$. (If you wish, you may check that this equality of fractions really is an equivalence relation and the two binary operations on fractions are well-defined on equivalence classes and make $S^{-1}A$ into a ring.) We have a canonical ring map

$$(1.3.3.1) \quad A \rightarrow S^{-1}A$$

given by $a \mapsto a/1$. Note that if $0 \in S$, $S^{-1}A$ is the 0-ring.

There are two particularly important flavors of multiplicative subsets. The first is $\{1, f, f^2, \dots\}$, where $f \in A$. This localization is denoted A_f . (Can you describe an isomorphism $A_f \xrightarrow{\sim} A[t]/(tf - 1)$?) The second is $A - \mathfrak{p}$, where \mathfrak{p} is a prime ideal. This localization $S^{-1}A$ is denoted $A_{\mathfrak{p}}$. (Notational warning: If \mathfrak{p} is a prime ideal, then $A_{\mathfrak{p}}$ means you're allowed to divide by elements not in \mathfrak{p} . However, if $f \in A$, A_f means you're allowed to divide by f . This can be confusing. For example, if (f) is a prime ideal, then $A_f \neq A_{(f)}$.)

Warning: sometimes localization is first introduced in the special case where A is an integral domain and $0 \notin S$. In that case, $A \hookrightarrow S^{-1}A$, but this isn't always true, as shown by the following exercise. (But we will see that noninjective localizations needn't be pathological, and we can sometimes understand them geometrically, see Exercise 3.2.L.)

1.3.C. EXERCISE. Show that $A \rightarrow S^{-1}A$ is injective if and only if S contains no zerodivisors. (A **zerodivisor** of a ring A is an element a such that there is a nonzero element b with $ab = 0$. The other elements of A are called **non-zerodivisors**. For example, an invertible element is never a zerodivisor. Counter-intuitively, 0 is a zerodivisor in every ring but the 0-ring. More generally, if M is an A -module, then $a \in A$ is a **zerodivisor for** M if there is a nonzero $m \in M$ with $am = 0$. The other elements of A are called **non-zerodivisors for** M . Equivalently, and *very* usefully, $a \in A$ is a non-zerodivisor for M if and only if $\times a : M \rightarrow M$ is an injection, or equivalently in the language of §1.6, if

$$0 \longrightarrow M \xrightarrow{\times a} M$$

is exact.)

If A is an integral domain and $S = A - \{0\}$, then $S^{-1}A$ is called the **fraction field** of A , which we denote $K(A)$. The previous exercise shows that A is a subring of its fraction field $K(A)$. We now return to the case where A is a general (commutative) ring.

1.3.D. EXERCISE. Verify that $A \rightarrow S^{-1}A$ satisfies the following universal property: $S^{-1}A$ is initial among A -algebras B where every element of S is sent to an invertible element in B . (Recall: the data of “an A -algebra B ” and “a ring map $A \rightarrow B$ ” are the same.) Translation: any map $A \rightarrow B$ where every element of S is sent to an invertible element must factor uniquely through $A \rightarrow S^{-1}A$. Another translation: a ring map out of $S^{-1}A$ is the same thing as a ring map from A that sends every element of S to an invertible element. Furthermore, an $S^{-1}A$ -module is the same

thing as an A -module for which $s \times \cdot : M \rightarrow M$ is an A -module isomorphism for all $s \in S$.

In fact, it is cleaner to *define* $A \rightarrow S^{-1}A$ by the universal property, and to show that it exists, and to use the universal property to check various properties $S^{-1}A$ has. Let's get some practice with this by *defining* localizations of modules by universal property. Suppose M is an A -module. We define the A -module map $\phi: M \rightarrow S^{-1}M$ as being initial among A -module maps $M \rightarrow N$ such that elements of S are invertible in N ($s \times \cdot : N \rightarrow N$ is an isomorphism for all $s \in S$). More precisely, any such map $\alpha: M \rightarrow N$ factors uniquely through ϕ :

$$\begin{array}{ccc} M & \xrightarrow{\phi} & S^{-1}M \\ & \searrow \alpha & \downarrow \exists! \\ & & N \end{array}$$

(Translation: $M \rightarrow S^{-1}M$ is universal (initial) among A -module maps from M to modules that are actually $S^{-1}A$ -modules. Can you make this precise by defining clearly the objects and morphisms in this category?)

Notice: (i) this determines $\phi: M \rightarrow S^{-1}M$ up to unique isomorphism (you should think through what this means); (ii) we are defining not only $S^{-1}M$, but also the map ϕ at the same time; and (iii) essentially by definition the A -module structure on $S^{-1}M$ extends to an $S^{-1}A$ -module structure.

1.3.E. EXERCISE. Show that $\phi: M \rightarrow S^{-1}M$ exists, by constructing something satisfying the universal property. Hint: define elements of $S^{-1}M$ to be of the form m/s where $m \in M$ and $s \in S$, and $m_1/s_1 = m_2/s_2$ if and only if for some $s \in S$, $s(s_2m_1 - s_1m_2) = 0$. Define the additive structure by $(m_1/s_1) + (m_2/s_2) = (s_2m_1 + s_1m_2)/(s_1s_2)$, and the $S^{-1}A$ -module structure (and hence the A -module structure) is given by $(a_1/s_1) \cdot (m_2/s_2) = (a_1m_2)/(s_1s_2)$.

1.3.F. EXERCISE.

(a) Show that localization commutes with finite products, or equivalently, with finite direct sums. In other words, if M_1, \dots, M_n are A -modules, describe an isomorphism (of A -modules, and of $S^{-1}A$ -modules) $S^{-1}(M_1 \times \dots \times M_n) \rightarrow S^{-1}M_1 \times \dots \times S^{-1}M_n$.

(b) Show that localization commutes with *arbitrary* direct sums.

(c) Show that “localization does not necessarily commute with infinite products”: the obvious map $S^{-1}(\prod_i M_i) \rightarrow \prod_i S^{-1}M_i$ induced by the universal property of localization is not always an isomorphism. (Hint: $(1, 1/2, 1/3, 1/4, \dots) \in \mathbb{Q} \times \mathbb{Q} \times \dots$.)

1.3.4. Remark. Localization does not always commute with Hom , see Example 1.6.8. But Exercise 1.6.G will show that in good situations (if the first argument of Hom is *finitely presented*), localization *does* commute with Hom .

1.3.5. Tensor products. Another important example of a universal property construction is the notion of a **tensor product** of A -modules

$$\otimes_A : \quad \text{obj}(\text{Mod}_A) \times \text{obj}(\text{Mod}_A) \longrightarrow \text{obj}(\text{Mod}_A)$$

$$(M, N) \longmapsto M \otimes_A N$$

The subscript A is often suppressed when it is clear from context. The tensor product is often defined as follows. Suppose you have two A -modules M and N . Then elements of the tensor product $M \otimes_A N$ are finite A -linear combinations of symbols $m \otimes n$ ($m \in M, n \in N$), subject to relations $(m_1 + m_2) \otimes n = m_1 \otimes n + m_2 \otimes n$, $m \otimes (n_1 + n_2) = m \otimes n_1 + m \otimes n_2$, $a(m \otimes n) = (am) \otimes n = m \otimes (an)$ (where $a \in A, m_1, m_2 \in M, n_1, n_2 \in N$). More formally, $M \otimes_A N$ is the free A -module generated by $M \times N$, quotiented by the submodule generated by $(m_1 + m_2, n) - (m_1, n) - (m_2, n)$, $(m, n_1 + n_2) - (m, n_1) - (m, n_2)$, $a(m, n) - (am, n)$, and $a(m, n) - (m, an)$ for $a \in A, m, m_1, m_2 \in M, n, n_1, n_2 \in N$. The image of (m, n) in this quotient is $m \otimes n$.

If A is a field k , we recover the tensor product of vector spaces.

1.3.G. EXERCISE (IF YOU HAVEN'T SEEN TENSOR PRODUCTS BEFORE). Show that $\mathbb{Z}/(10) \otimes_{\mathbb{Z}} \mathbb{Z}/(12) \cong \mathbb{Z}/(2)$. (This exercise is intended to give some hands-on practice with tensor products.)

1.3.H. IMPORTANT EXERCISE: RIGHT-EXACTNESS OF $(\cdot) \otimes_A N$. Show that $(\cdot) \otimes_A N$ gives a covariant functor $\text{Mod}_A \rightarrow \text{Mod}_A$. Show that $(\cdot) \otimes_A N$ is a **right-exact functor**, i.e., if

$$M' \rightarrow M \rightarrow M'' \rightarrow 0$$

is an exact sequence of A -modules (which means $f: M \rightarrow M''$ is surjective, and M' surjects onto the kernel of f ; see §1.6), then the induced sequence

$$M' \otimes_A N \rightarrow M \otimes_A N \rightarrow M'' \otimes_A N \rightarrow 0$$

is also exact. This exercise is repeated in Exercise 1.6.F, but you may get a lot out of doing it now. (You will be reminded of the definition of right-exactness in §1.6.5.)

In contrast, you can quickly check that tensor product is not left-exact: tensor the exact sequence of \mathbb{Z} -modules

$$0 \longrightarrow \mathbb{Z} \xrightarrow{\times 2} \mathbb{Z} \longrightarrow \mathbb{Z}/(2) \longrightarrow 0$$

with $\mathbb{Z}/(2)$.

The constructive definition \otimes is a weird definition, and really the “wrong” definition. To motivate a better one: notice that there is a natural A -bilinear map $M \times N \rightarrow M \otimes_A N$. (If $M, N, P \in \text{Mod}_A$, a map $f: M \times N \rightarrow P$ is **A -bilinear** if $f(m_1 + m_2, n) = f(m_1, n) + f(m_2, n)$, $f(m, n_1 + n_2) = f(m, n_1) + f(m, n_2)$, and $f(am, n) = f(m, an) = af(m, n)$.) Any A -bilinear map $M \times N \rightarrow P$ factors through the tensor product uniquely: $M \times N \rightarrow M \otimes_A N \rightarrow P$. (Think this through!)

We can take this as the *definition* of the tensor product as follows. It is an A -module T along with an A -bilinear map $t: M \times N \rightarrow T$, such that given any

A -bilinear map $t': M \times N \rightarrow T'$, there is a unique A -linear map $f: T \rightarrow T'$ such that $t' = f \circ t$.

$$\begin{array}{ccc} M \times N & \xrightarrow{t} & T \\ & \searrow t' & \swarrow \exists! f \\ & T' & \end{array}$$

1.3.I. EXERCISE. Show that $(T, t: M \times N \rightarrow T)$ is unique up to unique isomorphism. Hint: first figure out what “unique up to unique isomorphism” means for such pairs, using a category of pairs (T, t) . Then follow the analogous argument for the product.

In short: given M and N , there is an A -bilinear map $t: M \times N \rightarrow M \otimes_A N$, unique up to unique isomorphism, defined by the following universal property: for any A -bilinear map $t': M \times N \rightarrow T'$ there is a unique A -linear map $f: M \otimes_A N \rightarrow T'$ such that $t' = f \circ t$.

As with all universal property arguments, this argument shows uniqueness *assuming existence*. To show existence, we need an explicit construction.

1.3.J. EXERCISE. Show that the construction of §1.3.5 satisfies the universal property of tensor product.

The three exercises below are useful facts about tensor products with which you should be familiar.

1.3.K. IMPORTANT EXERCISE.

(a) If M is an A -module and $A \rightarrow B$ is a morphism of rings, give $B \otimes_A M$ the structure of a B -module (this is part of the exercise). Show that this describes a functor $\text{Mod}_A \rightarrow \text{Mod}_B$.

(b) If further $A \rightarrow C$ is another morphism of rings, show that $B \otimes_A C$ has a natural structure of a ring. Hint: multiplication will be given by $(b_1 \otimes c_1)(b_2 \otimes c_2) = (b_1 b_2) \otimes (c_1 c_2)$. (Exercise 1.3.U will interpret this construction as a fibered coproduct.)


1.3.L. IMPORTANT EXERCISE. If S is a multiplicative subset of A and M is an A -module, describe a natural isomorphism $(S^{-1}A) \otimes_A M \xrightarrow{\sim} S^{-1}M$ (as $S^{-1}A$ -modules *and* as A -modules).

1.3.M. EXERCISE (\otimes COMMUTES WITH \oplus). Show that tensor products commute with arbitrary direct sums: if M and $\{N_i\}_{i \in I}$ are all A -modules, describe an isomorphism

$$M \otimes (\oplus_{i \in I} N_i) \xrightarrow{\sim} \oplus_{i \in I} (M \otimes N_i).$$

1.3.6. Essential Example: Fibered products. Suppose we have morphisms $\alpha: X \rightarrow Z$ and $\beta: Y \rightarrow Z$ (in *any* category). Then the **fibered product** (or *fibered product*) is an object $X \times_Z Y$ along with morphisms $\text{pr}_X: X \times_Z Y \rightarrow X$ and $\text{pr}_Y: X \times_Z Y \rightarrow Y$, where the two compositions $\alpha \circ \text{pr}_X, \beta \circ \text{pr}_Y: X \times_Z Y \rightarrow Z$ agree, such that given any object W with maps to X and Y (whose compositions to Z

agree), these maps factor through some unique $W \rightarrow X \times_Z Y$:



(Warning: the definition of the fibered product depends on α and β , even though they are omitted from the notation $X \times_Z Y$.)

By the usual universal property argument, if it exists, it is unique up to unique isomorphism. (You should think this through until it is clear to you.) Thus the use of the phrase “the fibered product” (rather than “a fibered product”) is reasonable, and we should reasonably be allowed to give it the name $X \times_Z Y$. We know what maps to it are: they are precisely maps to X and maps to Y that agree as maps to Z .

1.3.7. Definition. As an example, if $\pi: X \rightarrow Y$ is a morphism, and the fibered product $X \times_Y X$ exists, then this determines a **diagonal morphism** $\delta_\pi: X \rightarrow X \times_Y X$. The diagonal morphism will turn out to be a very useful notion.

Depending on your religion, the diagram

$$\begin{array}{ccc}
 X \times_Z Y & \xrightarrow{\text{pr}_Y} & Y \\
 \downarrow \text{pr}_X & & \downarrow \beta \\
 X & \xrightarrow{\alpha} & Z
 \end{array}$$

is called a **fibered/pullback/Cartesian diagram/square** (six possibilities — even more are possible if you prefer “fibred” to “fibered”).

The right way to interpret the notion of fibered product is first to think about what it means in the category of sets.

1.3.N. EXERCISE (FIBERED PRODUCTS OF SETS). Show that in *Sets*,

$$X \times_Z Y = \{(x, y) \in X \times Y : \alpha(x) = \beta(y)\}.$$

More precisely, show that the right side, equipped with its evident maps to X and Y , satisfies the universal property of the fibered product. (This will help you build intuition for fibered products.)

1.3.O. EXERCISE. If X is a topological space, show that fibered products always exist in the category of open sets of X , by describing what a fibered product is. (Hint: it has a one-word description.)

1.3.P. EXERCISE. If Z is the final object in a category \mathcal{C} , and $X, Y \in \mathcal{C}$, show that “ $X \times_Z Y = X \times Y$ ”: “the” fibered product over Z is uniquely isomorphic to “the” product. Assume all relevant (fibered) products exist. (This is an exercise about unwinding the definition.)

1.3.Q. USEFUL EXERCISE: TOWERS OF CARTESIAN DIAGRAMS ARE CARTESIAN DIAGRAMS. If the two squares in the following commutative diagram are Cartesian

diagrams, show that the “outside rectangle” (involving U , V , Y , and Z) is also a Cartesian diagram.

$$\begin{array}{ccc} U & \longrightarrow & V \\ \downarrow & & \downarrow \\ W & \longrightarrow & X \\ \downarrow & & \downarrow \\ Y & \longrightarrow & Z \end{array}$$

1.3.R. EXERCISE. Given morphisms $X_1 \rightarrow Y$, $X_2 \rightarrow Y$, and $Y \rightarrow Z$, show that there is a natural morphism $X_1 \times_Y X_2 \rightarrow X_1 \times_Z X_2$, assuming that both fibered products exist. (This is trivial once you figure out what it is saying. The point of this exercise is to see why it is trivial.)

1.3.S. IMPORTANT EXERCISE: THE DIAGONAL-BASE-CHANGE DIAGRAM. Suppose we are given morphisms $X_1, X_2 \rightarrow Y$ and $Y \rightarrow Z$. Show that the following diagram is a Cartesian square.

$$\boxed{\begin{array}{ccc} X_1 \times_Y X_2 & \longrightarrow & X_1 \times_Z X_2 \\ \downarrow & & \downarrow \\ Y & \longrightarrow & Y \times_Z Y \end{array}}$$

Assume all relevant (fibered) products exist. (If this exercise too hard now, you can try it again at Exercise 1.4.B.) You will appreciate how useful this diagram is when you repeatedly use the diagonal morphism in proofs and constructions; for this reason some call it the *magic diagram*.

If you liked this problem, you may enjoy Exercise 11.2.C.

1.3.8. Coproducts. Define **coproduct** in a category by reversing all the arrows in the definition of product. Define **fibered coproduct** in a category by reversing all the arrows in the definition of fibered product.

1.3.T. EXERCISE. Show that coproduct for *Sets* is disjoint union. This is why we use the notation \coprod for disjoint union.

1.3.U. EXERCISE. Suppose $A \rightarrow B$ and $A \rightarrow C$ are two ring morphisms, so in particular B and C are A -modules. Recall (Exercise 1.3.K) that $B \otimes_A C$ has a ring structure. Show that there is a natural morphism $B \rightarrow B \otimes_A C$ given by $b \mapsto b \otimes 1$. (This is not necessarily an inclusion; see Exercise 1.3.G.) Similarly, there is a natural morphism $C \rightarrow B \otimes_A C$. Show that this gives a fibered coproduct on rings, i.e., that

$$\begin{array}{ccc} B \otimes_A C & \longleftarrow & C \\ \uparrow & & \uparrow \\ B & \longleftarrow & A \end{array}$$

satisfies the universal property of fibered coproduct.

1.3.9. Monomorphisms and epimorphisms.

1.3.10. Definition. A morphism $\pi: X \rightarrow Y$ is a **monomorphism** if any two morphisms $\mu_1: Z \rightarrow X$ and $\mu_2: Z \rightarrow X$ such that $\pi \circ \mu_1 = \pi \circ \mu_2$ must satisfy $\mu_1 = \mu_2$. In other words, there is at most one way of filling in the dotted arrow so that the diagram

$$\begin{array}{ccc} & Z & \\ \downarrow \leq 1 & \searrow & \\ X & \xrightarrow{\pi} & Y \end{array}$$

commutes — for any object Z , the natural map $\text{Mor}(Z, X) \rightarrow \text{Mor}(Z, Y)$ is an injection. Intuitively, it is the categorical version of an injective map, and indeed this notion generalizes the familiar notion of injective maps of sets. (The reason we don't use the word "injective" is that in some contexts, "injective" will have an intuitive meaning which may not agree with "monomorphism". One example: in the category of divisible groups, the map $\mathbb{Q} \rightarrow \mathbb{Q}/\mathbb{Z}$ is a monomorphism but not injective. This is also the case with "epimorphism" (to be defined shortly) vs. "surjective".)

1.3.V. EXERCISE. Show that the composition of two monomorphisms is a monomorphism.

1.3.W. EXERCISE. Prove that a morphism $\pi: X \rightarrow Y$ is a monomorphism if and only if the fibered product $X \times_Y X$ exists, and the induced diagonal morphism $\delta_\pi: X \rightarrow X \times_Y X$ (Definition 1.3.7) is an isomorphism. We may then take this as the definition of monomorphism. (Monomorphisms aren't central to future discussions, although they will come up again. This exercise is just good practice.)

1.3.X. EASY EXERCISE. We use the notation of Exercise 1.3.R. Show that if $Y \rightarrow Z$ is a monomorphism, then the morphism $X_1 \times_Y X_2 \rightarrow X_1 \times_Z X_2$ you described in Exercise 1.3.R is an isomorphism. (Hint: for any object V , give a natural bijection between maps from V to the first and maps from V to the second. It is also possible to use the Diagonal-Base-Change diagram, Exercise 1.3.S.)

The notion of an **epimorphism** is "dual" to the definition of monomorphism, where all the arrows are reversed. This concept will not be central for us, although it turns up in the definition of an abelian category. Intuitively, it is the categorical version of a surjective map. (But be careful when working with categories of objects that are sets with additional structure, as epimorphisms need not be surjective. Example: in the category *Rings*, $\mathbb{Z} \rightarrow \mathbb{Q}$ is an epimorphism, but obviously not surjective.)

1.3.11. Representable functors and Yoneda's Lemma. Much of our discussion about universal properties can be cleanly expressed in terms of representable functors, under the rubric of "Yoneda's Lemma". Yoneda's lemma is an easy fact stated in a complicated way. Informally speaking, you can essentially recover an object in a category by knowing the maps into it. For example, we have seen that the data of maps to $X \times Y$ are naturally (canonically) the data of maps to X and to Y . Indeed, we have now taken this as the *definition* of $X \times Y$.

Recall Example 1.2.20. Suppose A is an object of category \mathcal{C} . For any object $C \in \mathcal{C}$, we have a set of morphisms $\text{Mor}(C, A)$. If we have a morphism $f: B \rightarrow C$,

we get a map of sets

$$(1.3.11.1) \quad \text{Mor}(C, A) \rightarrow \text{Mor}(B, A),$$

by composition: given a map from C to A , we get a map from B to A by precomposing with $f: B \rightarrow C$. Hence this gives a contravariant functor $h_A: \mathcal{C} \rightarrow \text{Sets}$. Yoneda's Lemma states that the functor h_A determines A up to unique isomorphism. More precisely:

1.3.Y. IMPORTANT EXERCISE THAT YOU SHOULD DO ONCE IN YOUR LIFE (YONEDA'S LEMMA).

(a) Suppose you have two objects A and A' in a category \mathcal{C} , and morphisms

$$(1.3.11.2) \quad i_C: \text{Mor}(C, A) \rightarrow \text{Mor}(C, A')$$

that commute with the maps (1.3.11.1). Show that the i_C (as C ranges over the objects of \mathcal{C}) are induced from a unique morphism $g: A \rightarrow A'$. More precisely, show that there is a unique morphism $g: A \rightarrow A'$ such that for all $C \in \mathcal{C}$, i_C is $u \mapsto g \circ u$. (b) If furthermore the i_C are all bijections, show that the resulting g is an isomorphism. (Hint for both: This is much easier than it looks. This statement is so general that there are really only a couple of things that you could possibly try. For example, if you're hoping to find a morphism $A \rightarrow A'$, where will you find it? Well, you are looking for an element $\text{Mor}(A, A')$. So just plug in $C = A$ to (1.3.11.2), and see where the identity goes.)

There is an analogous statement with the arrows reversed, where instead of maps into A , you think of maps *from* A . The role of the contravariant functor h_A of Example 1.2.20 is played by the covariant functor h^A of Example 1.2.14. Because the proof is the same (with the arrows reversed), you needn't think it through.

The phrase "Yoneda's Lemma" properly refers to a more general statement. Although it looks more complicated, it is no harder to prove.

1.3.Z. ★ EXERCISE.

(a) Suppose A and B are objects in a category \mathcal{C} . Give a bijection between the natural transformations $h^A \rightarrow h^B$ of covariant functors $\mathcal{C} \rightarrow \text{Sets}$ (see Example 1.2.14 for the definition) and the morphisms $B \rightarrow A$.

(b) State and prove the corresponding fact for contravariant functors h_A (see Example 1.2.20). Remark: A contravariant functor F from \mathcal{C} to Sets is said to be **representable** if there is a natural isomorphism

$$\xi: F \xrightarrow{\sim} h_A.$$

Thus the representing object A is determined up to unique isomorphism by the pair (F, ξ) . There is a similar definition for covariant functors. (We will revisit this in §7.6, and this problem will appear again as Exercise 7.6.C. The element $\xi^{-1}(\text{id}_A) \in F(A)$ is often called the "universal object"; do you see why?)

(c) **Yoneda's Lemma.** Suppose F is a covariant functor $\mathcal{C} \rightarrow \text{Sets}$, and $A \in \mathcal{C}$. Give a bijection between the natural transformations $h^A \rightarrow F$ and $F(A)$. (The corresponding fact for contravariant functors is essentially Exercise 10.1.B.)

In fancy terms, Yoneda's lemma states the following. Given a category \mathcal{C} , we can produce a new category, called the *functor category* of \mathcal{C} , where the objects are contravariant functors $\mathcal{C} \rightarrow \text{Sets}$, and the morphisms are natural transformations


of such functors. We have a functor (which we can usefully call h) from \mathcal{C} to its functor category, which sends A to h_A . Yoneda's Lemma states that this is a fully faithful functor, called the *Yoneda embedding*. (Fully faithful functors were defined in §1.2.15.)

1.3.12. Joke. The Yoda embedding, contravariant it is.

1.4 Limits and colimits

Limits and colimits are two important definitions determined by universal properties. They generalize a number of familiar constructions. I will give the definition first, and then show you why it is familiar. For example, fractions will be motivating examples of colimits (Exercise 1.4.D(a)), and the p -adic integers (Example 1.4.4) will be motivating examples of limits.

1.4.1. Limits. We say that a category is a **small category** if the objects form a set and the morphisms form a set. (This is a technical condition intended only for experts.) Suppose \mathcal{I} is any small category, and \mathcal{C} is any category. Then a functor $F: \mathcal{I} \rightarrow \mathcal{C}$ (i.e., with an object $A_i \in \mathcal{C}$ for each element $i \in \mathcal{I}$, and appropriate commuting morphisms dictated by \mathcal{I}) is said to be a **diagram indexed by \mathcal{I}** . We call \mathcal{I} an **index category**. Our index categories will usually be partially ordered sets (Example 1.2.8), in which in particular there is at most one morphism between any two objects. (But other examples are sometimes useful.) For example, if \square is the category



and \mathcal{A} is a category, then a functor $\square \rightarrow \mathcal{A}$ is precisely the data of a commuting square in \mathcal{A} .

Then the **limit of the diagram** is an object $\lim_{\mathcal{I}} A_i$ (or $\varprojlim_{\mathcal{I}} A_i$) of \mathcal{C} along with morphisms $f_j: \lim_{\mathcal{I}} A_i \rightarrow A_j$ for each $j \in \mathcal{I}$, such that if $m: j \rightarrow k$ is a morphism in \mathcal{I} , then

(1.4.1.1)


$$\begin{array}{ccc} \lim_{\mathcal{I}} A_i & & \\ f_j \downarrow & \searrow f_k & \\ A_j & \xrightarrow{F(m)} & A_k \end{array}$$

commutes, and this object and maps to each A_i are universal (final) with respect to this property. More precisely, given any other object W along with maps $g_i: W \rightarrow A_i$ commuting with the $F(m)$ (if $m: j \rightarrow k$ is a morphism in \mathcal{I} , then $g_k = F(m) \circ g_j$), then there is a unique map

$$g: W \rightarrow \lim_{\mathcal{I}} A_i$$

so that $g_i = f_i \circ g$ for all i . (In some cases, the limit is sometimes called the **inverse limit** or **projective limit**. We won't use this language.) By the usual universal property argument, if the limit exists, it is unique up to unique isomorphism.

1.4.2. Examples: products. For example, if \mathcal{I} is the partially ordered set



we obtain the fibered product.

If \mathcal{I} is




we obtain the product.

If \mathcal{I} is a set (i.e., the only morphisms are the identity maps), then the limit is called the **product** of the A_i , and is denoted $\prod_i A_i$. The special case where \mathcal{I} has two elements is the example of the previous paragraph.

1.4.A. EXERCISE (REALITY CHECK). Suppose that the partially ordered set \mathcal{I} has an initial object e . Show that the limit of any diagram indexed by \mathcal{I} exists.


1.4.B. EXERCISE: THE DIAGONAL-BASE-CHANGE DIAGRAM, AGAIN. Solve 1.3.S again by identifying both $X_1 \times_Y X_2$ and $Y \times_{(Y \times_Z Y)} (X_1 \times_Z X_2)$ as the limit of the diagram



1.4.3. Example: formal power series. For a prime number A , the **formal power series**, $A[[x]]$, are often described informally (and somewhat unnaturally) as being the ring

$$A[[x]] = a_0 + a_1x + a_2x^2 + \dots$$

(where $a_i \in A$, and the ring operations are the “obvious” ones). They are an example of a limit in the category of rings:



The universal property of limits yields a natural ring morphism $A[x] \rightarrow A[[x]]$. If $A = \mathbb{R}$ or \mathbb{C} , this map factors through the ring of *convergent power series*.

1.4.4. Example: the p -adic integers. For a prime number p , the **p -adic integers** (or more informally, **p -adics**), \mathbb{Z}_p , are often described informally (and somewhat

unnaturally) as being of the form

$$\mathbb{Z}_p = a_0 + a_1p + a_2p^2 + \cdots$$

(where $0 \leq a_i < p$). They are an example of a limit in the category of rings:

$$\begin{array}{ccccccc} & & \mathbb{Z}_p & & & & \\ & \searrow & & \searrow & \searrow & \searrow & \\ \cdots & \longrightarrow & \mathbb{Z}/(p^3) & \longrightarrow & \mathbb{Z}/(p^2) & \longrightarrow & \mathbb{Z}/(p). \end{array}$$

(Warning: \mathbb{Z}_p is sometimes used to denote the integers modulo p , but $\mathbb{Z}/(p)$ or $\mathbb{Z}/p\mathbb{Z}$ is better to use for this, to avoid confusion. Worse: by §1.3.3, \mathbb{Z}_p also denotes those rationals whose denominators are a power of p . Hopefully the meaning of \mathbb{Z}_p will be clear from the context.)

The similarity of Examples 1.4.3 and 1.4.4 is no coincidence. Formal power series and the p -adic integers are examples of *completions*, the topic of Chapter 30.

Limits do not always exist for any index category \mathcal{J} . However, you can often easily check that limits exist if the objects of your category can be interpreted as sets with additional structure, and arbitrary products exist (respecting the set-like structure).

1.4.C. IMPORTANT EXERCISE. Show that in the category *Sets*,

$$\left\{ (a_i)_{i \in \mathcal{J}} \in \prod_i A_i : F(m)(a_j) = a_k \text{ for all } m \in \text{Mor}_{\mathcal{J}}(j, k) \in \text{Mor}(\mathcal{J}) \right\},$$

along with the obvious projection maps to each A_i , is the limit $\lim_{\mathcal{J}} A_i$.

This clearly also works in the category Mod_A of A -modules (in particular Vec_k and Ab), as well as *Rings*.

From this point of view, $2 + 3p + 2p^2 + \cdots \in \mathbb{Z}_p$ can be understood as the sequence $(2, 2 + 3p, 2 + 3p + 2p^2, \dots)$.

1.4.5. Colimits. More immediately relevant for us will be the dual (arrow-reversed version) of the notion of limit (or inverse limit). We just flip the arrows f_i in (1.4.1.1), and get the notion of a **colimit**, which is denoted $\text{colim}_{\mathcal{J}} A_i$ (or $\varinjlim_{\mathcal{J}} A_i$). (You should draw the corresponding diagram.) Again, if it exists, it is unique up to unique isomorphism. (In some cases, the colimit is sometimes called the **direct limit**, **inductive limit**, or **injective limit**. We won't use this language. I prefer using limit/colimit in analogy with kernel/cokernel and product/coproduct. This is more than analogy, as kernels and products may be interpreted as limits, and similarly with cokernels and coproducts. Also, I remember that kernels “map to”, and cokernels are “mapped to”, which reminds me that a limit maps *to* all the objects in the big commutative diagram indexed by \mathcal{J} ; and a colimit has a map *from* all the objects.)

1.4.6. Joke. A comathematician is a device for turning cotheorems into ffee.

Even though we have just flipped the arrows, colimits behave quite differently from limits.

1.4.7. Example. The abelian group $5^{-\infty}\mathbb{Z}$ of rational numbers whose denominators are powers of 5 is a colimit $\operatorname{colim}_{i \in \mathbb{Z}^+} 5^{-i}\mathbb{Z}$. More precisely, $5^{-\infty}\mathbb{Z}$ is the colimit of the diagram

$$\mathbb{Z} \longrightarrow 5^{-1}\mathbb{Z} \longrightarrow 5^{-2}\mathbb{Z} \longrightarrow \dots$$

in the category of abelian groups.

The colimit over an index set I is called the **coproduct**, denoted $\coprod_i A_i$, and is the dual (arrow-reversed) notion to the product.

1.4.D. EXERCISE.

(a) Interpret the statement “ $\mathbb{Q} = \operatorname{colim} \frac{1}{n}\mathbb{Z}$ ”.

(b) Interpret the union of some subsets of a given set as a colimit. (Dually, the intersection can be interpreted as a limit.) The objects of the category in question are the subsets of the given set.

Colimits do not always exist, but there are two useful large classes of examples for which they do.

1.4.8. Definition. A nonempty partially ordered set (S, \geq) is **filtered** (or is said to be a **filtered set**) if for each $x, y \in S$, there is a z such that $x \geq z$ and $y \geq z$. More generally (see Figure 1.1, a nonempty category \mathcal{J} is **filtered** if:

- (i) for each $x, y \in \mathcal{J}$, there is a $z \in \mathcal{J}$ and arrows $x \rightarrow z$ and $y \rightarrow z$, and
- (ii) for every two arrows $u: x \rightarrow y$ and $v: x \rightarrow y$, there is an arrow $w: y \rightarrow z$ such that $w \circ u = w \circ v$.

(Other terminologies are also commonly used, such as “directed partially ordered set” and “filtered index category”, respectively.)




FIGURE 1.1. A filtered category (pictorial definition)

1.4.E. EXERCISE. Suppose \mathcal{J} is filtered. (We will almost exclusively use the case where \mathcal{J} is a filtered set.) Recall the symbol \coprod for disjoint union of sets. Show that any diagram in *Sets* indexed by \mathcal{J} has the following, with the obvious maps to it, as a colimit:

$$\left\{ (a_i, i) \in \coprod_{i \in \mathcal{J}} A_i \right\} / \left((a_i, i) \sim (a_j, j) \text{ if and only if there are } f: A_i \rightarrow A_k \text{ and } g: A_j \rightarrow A_k \text{ in the diagram for which } f(a_i) = g(a_j) \text{ in } A_k \right)$$

(You will see that the “filtered” hypothesis is there is to ensure that \sim is an equivalence relation.)

For example, in Example 1.4.7, each element of the colimit is an element of something upstairs, but you can't say in advance what it is an element of. For instance, $17/125$ is an element of the $5^{-3}\mathbb{Z}$ (or $5^{-4}\mathbb{Z}$, or later ones), but not $5^{-2}\mathbb{Z}$.

This idea applies to many categories whose objects can be interpreted as sets with additional structure (such as abelian groups, A -modules, groups, etc.). For example, the colimit $\text{colim } M_i$ in the category of A -modules Mod_A can be described as follows. The set underlying $\text{colim } M_i$ is defined as in Exercise 1.4.E. To add the elements $m_i \in M_i$ and $m_j \in M_j$, choose an $\ell \in \mathcal{I}$ with arrows $u: i \rightarrow \ell$ and $v: j \rightarrow \ell$, and then define the sum of m_i and m_j to be $F(u)(m_i) + F(v)(m_j) \in M_\ell$. The element $m_i \in M_i$ is 0 if and only if there is some arrow $u: i \rightarrow k$ for which $F(u)(m_i) = 0$, i.e., if it becomes 0 "later in the diagram". Last, multiplication by an element of A is defined in the obvious way.

1.4.F. EXERCISE. Verify that the A -module described above is indeed the colimit. (Make sure you verify that addition is well-defined, i.e., is independent of the choice of representatives m_i and m_j , the choice of ℓ , and the choice of arrows u and v . Similarly, make sure that scalar multiplication is well-defined.)

1.4.G. USEFUL EXERCISE (LOCALIZATION AS A COLIMIT). Generalize Exercise 1.4.D(a) to interpret localization of an integral domain as a colimit over a filtered set: suppose S is a multiplicative set of A , and interpret $S^{-1}A = \text{colim } \frac{1}{s}A$ where the limit is over $s \in S$, and in the category of A -modules. (Aside: Can you make some version of this work even if A isn't an integral domain, e.g., $S^{-1}A = \text{colim } A_s$? This will work in the category of A -algebras.)

A variant of this construction works without the filtered condition, if you have another means of "connecting elements in different objects of your diagram". For example:

1.4.H. EXERCISE: COLIMITS OF A -MODULES WITHOUT THE FILTERED CONDITION. Suppose you are given a diagram of A -modules indexed by $\mathcal{I}: F: \mathcal{I} \rightarrow \text{Mod}_A$, where we let $M_i := F(i)$. Show that the colimit is $\oplus_{i \in \mathcal{I}} M_i$ modulo the relations $m_i - F(n)(m_i)$ for every $n: i \rightarrow j$ in \mathcal{I} (i.e., for every arrow in the diagram). (Somewhat more precisely: "modulo" means "quotiented by the submodule generated by".)

1.4.9. Summary. One useful thing to informally keep in mind is the following. In a category where the objects are "set-like", an element of a limit can be thought of as a family of elements of each object in the diagram, that are "compatible" (Exercise 1.4.C). And an element of a colimit can be thought of ("has a representative that is") an element of a single object in the diagram (Exercise 1.4.E). Even though the definitions of limit and colimit are the same, just with arrows reversed, these interpretations are quite different.

1.4.10. Small remark. In fact, colimits exist in the category of sets for all reasonable ("small") index categories (see for example [E, Thm. A6.1]), but that won't matter to us.

1.4.11. Joke. What do you call someone who reads a paper on category theory? A coauthor!

1.5 Adjoints

We next come to a very useful notion closely related to universal properties. Just as a universal property “essentially” (up to unique isomorphism) determines an object in a category (assuming such an object exists), “adjoints” essentially determine a functor (again, assuming it exists). Two *covariant* functors $F: \mathcal{A} \rightarrow \mathcal{B}$ and $G: \mathcal{B} \rightarrow \mathcal{A}$ are **adjoint** if there is a natural bijection for all $A \in \mathcal{A}$ and $B \in \mathcal{B}$

$$(1.5.0.1) \quad \tau_{AB}: \text{Mor}_{\mathcal{B}}(F(A), B) \rightarrow \text{Mor}_{\mathcal{A}}(A, G(B)).$$

We say that (F, G) form an **adjoint pair**, and that F is **left-adjoint** to G (and G is **right-adjoint** to F). We say F is a **left adjoint** (and G is a **right adjoint**). By “natural” we mean the following. For all $f: A \rightarrow A'$ in \mathcal{A} , we require

$$(1.5.0.2) \quad \begin{array}{ccc} \text{Mor}_{\mathcal{B}}(F(A'), B) & \xrightarrow{Ff^*} & \text{Mor}_{\mathcal{B}}(F(A), B) \\ \downarrow \tau_{A'B} & & \downarrow \tau_{AB} \\ \text{Mor}_{\mathcal{A}}(A', G(B)) & \xrightarrow{f^*} & \text{Mor}_{\mathcal{A}}(A, G(B)) \end{array}$$

to commute, and for all $g: B \rightarrow B'$ in \mathcal{B} we want a similar commutative diagram to commute. (Here f^* is the map induced by $f: A \rightarrow A'$, and Ff^* is the map induced by $Ff: F(A) \rightarrow F(A')$.)

1.5.A. EXERCISE. Write down what this diagram should be.

1.5.B. EXERCISE. Show that the map τ_{AB} (1.5.0.1) has the following properties. For each A there is a map $\eta_A: A \rightarrow GF(A)$ so that for any $g: F(A) \rightarrow B$, the corresponding $\tau_{AB}(g): A \rightarrow G(B)$ is given by the composition

$$A \xrightarrow{\eta_A} GF(A) \xrightarrow{Gg} G(B).$$

Similarly, there is a map $\epsilon_B: FG(B) \rightarrow B$ for each B so that for any $f: A \rightarrow G(B)$, the corresponding map $\tau_{AB}^{-1}(f): F(A) \rightarrow B$ is given by the composition

$$F(A) \xrightarrow{Ff} FG(B) \xrightarrow{\epsilon_B} B.$$

Here is a key example of an adjoint pair.

1.5.C. EXERCISE. Suppose M, N , and P are A -modules (where A is a ring). Describe a bijection $\text{Hom}_A(M \otimes_A N, P) \leftrightarrow \text{Hom}_A(M, \text{Hom}_A(N, P))$. (Hint: try to use the universal property of \otimes .)

1.5.D. EXERCISE. Show that $(\cdot) \otimes_A N$ and $\text{Hom}_A(N, \cdot)$ are adjoint functors.

1.5.E. EXERCISE. Suppose $B \rightarrow A$ is a morphism of rings. If M is an A -module, you can create a B -module M_B by considering it as a B -module. This gives a functor $\cdot_B: \text{Mod}_A \rightarrow \text{Mod}_B$. Show that this functor is right-adjoint to $\cdot \otimes_B A$. In other words, describe a bijection

$$\text{Hom}_A(N \otimes_B A, M) \cong \text{Hom}_B(N, M_B)$$

functorial in both arguments. (This adjoint pair is very important, and is the key player in Chapter 17.)

1.5.1. ★ *Fancier remarks we won't use.* You can check that the left adjoint determines the right adjoint up to natural isomorphism, and vice versa. The maps η_A and ϵ_B of Exercise 1.5.B are called the *unit* and *counit* of the adjunction. This leads to a different characterization of adjunction. Suppose functors $F: \mathcal{A} \rightarrow \mathcal{B}$ and $G: \mathcal{B} \rightarrow \mathcal{A}$ are given, along with natural transformations $\eta: \text{id}_{\mathcal{A}} \rightarrow GF$ and $\epsilon: FG \rightarrow \text{id}_{\mathcal{B}}$ with the property that $G\epsilon \circ \eta G = \text{id}_G$ (for each $B \in \mathcal{B}$, the composition of $\eta_{G(B)}: G(B) \rightarrow GFG(B)$ and $G(\epsilon_B): GFG(B) \rightarrow G(B)$ is the identity) and $\epsilon F \circ F\eta = \text{id}_F$. Then you can check that F is left-adjoint to G . These facts aren't hard to check, so if you want to use them, you should verify everything for yourself.

1.5.2. *Examples from other fields.* For those familiar with representation theory: Frobenius reciprocity may be understood in terms of adjoints. Suppose V is a finite-dimensional representation of a finite group G , and W is a representation of a subgroup $H < G$. Then induction and restriction are an adjoint pair $(\text{Ind}_H^G, \text{Res}_H^G)$ between the category of G -modules and the category of H -modules.

Topologists' favorite adjoint pair may be the suspension functor and the loop space functor.

1.5.3. Example: groupification of abelian semigroups. Here is another motivating example: getting an abelian group from an abelian semigroup. (An **abelian semigroup** is just like an abelian group, except we don't require an identity or an inverse. Morphisms of abelian semigroups are maps of sets preserving the binary operation. One example is the non-negative integers $\mathbb{Z}^{\geq 0} = \{0, 1, 2, \dots\}$ under addition. Another is the positive integers $1, 2, \dots$ under multiplication. You may enjoy groupifying both.) From an abelian semigroup, you can create an abelian group. In our examples, from the nonnegative $(\mathbb{Z}^{\geq 0}, +)$, we create the integers \mathbb{Z} , and from the positive integers under multiplication $(\mathbb{Z}^{> 0}, \times)$, we create the positive rationals $\mathbb{Q}^{> 0}$. Here is a formalization of that notion. A **groupification** of a semigroup S is a map of abelian semigroups $\pi: S \rightarrow G$ such that G is an abelian group, and any map of abelian semigroups from S to an abelian group G' factors *uniquely* through G :

$$\begin{array}{ccc} S & \xrightarrow{\pi} & G \\ & \searrow & \vdots \exists! \\ & & G' \end{array}$$

(Perhaps “abelian groupification” would be more precise than “groupification”.)

1.5.F. EXERCISE (AN ABELIAN GROUP IS GROUPIFIED BY ITSELF). Show that if an abelian semigroup is *already* a group then the identity morphism is the groupification. (More correct: the identity morphism is *a* groupification.) Note that you don't need to construct groupification (or even know that it exists in general) to solve this exercise.

1.5.G. EXERCISE. Construct the “groupification functor” H from the category of nonempty abelian semigroups to the category of abelian groups. (One possible construction: given an abelian semigroup S , the elements of its groupification $H(S)$ are ordered pairs $(a, b) \in S \times S$, which you may think of as $a - b$, with the equivalence that $(a, b) \sim (c, d)$ if $a + d + e = b + c + e$ for some $e \in S$. Describe addition in this group, and show that it satisfies the properties of an abelian group. Describe

the abelian semigroup map $S \rightarrow H(S)$.) Let F be the forgetful functor from the category of abelian groups Ab to the category of abelian semigroups. Show that H is left-adjoint to F .

(Here is the general idea for experts: We have a full subcategory of a category. We want to “project” from the category to the subcategory. We have

$$\text{Mor}_{\text{category}}(S, H) = \text{Mor}_{\text{subcategory}}(G, H)$$

automatically; thus we are describing the left adjoint to the forgetful functor. How the argument worked: we constructed something which was in the smaller category, which automatically satisfies the universal property.)

1.5.H. EXERCISE (CF. EXERCISE 1.5.E). The purpose of this exercise is to give you more practice with “adjoints of forgetful functors”, the means by which we get groups from semigroups, and sheaves from presheaves. Suppose A is a ring, and S is a multiplicative subset. Then $S^{-1}A$ -modules are a full subcategory (§1.2.15) of the category of A -modules (via the obvious inclusion $\text{Mod}_{S^{-1}A} \hookrightarrow \text{Mod}_A$). Then $\text{Mod}_A \rightarrow \text{Mod}_{S^{-1}A}$ can be interpreted as an adjoint to the forgetful functor $\text{Mod}_{S^{-1}A} \rightarrow \text{Mod}_A$. State and prove the correct statements.

(Here is the larger story. Every $S^{-1}A$ -module is an A -module, and this is an injective map, so we have a covariant forgetful functor $F: \text{Mod}_{S^{-1}A} \rightarrow \text{Mod}_A$. In fact this is a fully faithful functor: it is injective on objects, and the morphisms between any two $S^{-1}A$ -modules *as A -modules* are just the same when they are considered as $S^{-1}A$ -modules. Then there is a functor $G: \text{Mod}_A \rightarrow \text{Mod}_{S^{-1}A}$, which might reasonably be called “localization with respect to S ”, which is left-adjoint to the forgetful functor. Translation: If M is an A -module, and N is an $S^{-1}A$ -module, then $\text{Mor}(GM, N)$ (morphisms as $S^{-1}A$ -modules, which are the same as morphisms as A -modules) are in natural bijection with $\text{Mor}(M, FN)$ (morphisms as A -modules).)

Here is a table of most of the adjoints that will come up for us.

situation	category \mathcal{A}	category \mathcal{B}	left adjoint $F: \mathcal{A} \rightarrow \mathcal{B}$	right adjoint $G: \mathcal{B} \rightarrow \mathcal{A}$
A-modules (Ex. 1.5.D)	Mod_A	Mod_A	$(\cdot) \otimes_A N$	$Hom_A(N, \cdot)$
ring maps $B \rightarrow A$ (Ex. 1.5.E)	Mod_B	Mod_A	$(\cdot) \otimes_B A$ (extension of scalars)	$M \mapsto M_B$ (restriction of scalars)
(pre)sheaves on a topological space X (Ex. 2.4.K)	presheaves on X	sheaves on X	sheafification	forgetful
(semi)groups (§1.5.3)	semigroups	groups	groupification	forgetful
sheaves, $\pi: X \rightarrow Y$ (Ex. 2.7.B)	sheaves on Y	sheaves on X	π^{-1}	π_*
sheaves of abelian groups or \mathcal{O} -modules, open embeddings $\pi: U \hookrightarrow Y$ (Ex. 24.4.G)	sheaves on U	sheaves on Y	$\pi_!$	π^{-1}
quasicoherent sheaves, $\pi: X \rightarrow Y$ (Prop. 17.3.6)	$QCoh_Y$	$QCoh_X$	π^*	π_*
ring maps $B \rightarrow A$ (Ex. 31.3.A)	Mod_A	Mod_B	$M \mapsto M_B$ (restriction of scalars)	$N \mapsto$ $Hom_B(A, N)$
quasicoherent sheaves, affine $\pi: X \rightarrow Y$ (Ex. 31.3.B(b))	$QCoh_X$	$QCoh_Y$	π_*	$\pi_{sh}^!$

Other examples will also come up, such as the adjoint pair (\sim, Γ_\bullet) between graded modules over a graded ring, and quasicoherent sheaves on the corresponding projective scheme (§16.4).

1.5.4. Useful comment for experts. One last comment only for people who have seen adjoints before: If (F, G) is an adjoint pair of functors, then F commutes with colimits, and G commutes with limits. Also, limits commute with limits and colimits commute with colimits. We will prove these facts (and a little more) in §1.6.12.

1.6 An introduction to abelian categories

Ton papier sur l'Algèbre homologique a été lu soigneusement, et a converti tout le monde (même Dieudonné, qui semble complètement functorisé!) à ton point de vue.

Your paper on homological algebra was read carefully and converted everyone (even Dieudonné, who seems to be completely functorised!) to your point of view.

— J.-P. Serre, letter to A. Grothendieck, Jul 13, 1955 [GrS, p. 17-18]

Since learning linear algebra, you have been familiar with the notions and behaviors of kernels, cokernels, etc. Later in your life you saw them in the category of abelian groups, and later still in the category of A -modules.

We will soon define some new categories (certain sheaves) that will have familiar-looking behavior, reminiscent of that of modules over a ring. The notions of kernels, cokernels, images, and more will make sense, and they will behave “the way we expect” from our experience with modules. This can be made precise through the notion of an *abelian category*. Abelian categories are the right general setting in which one can do “homological algebra”, in which notions of kernel, cokernel, and so on are used, and one can work with complexes and exact sequences.

We will see enough to motivate the definitions that we will see in general: monomorphism (and subobject), epimorphism, kernel, cokernel, and image. But in this book we will avoid having to show that they behave “the way we expect” in a general abelian category because the examples we will see are directly interpretable in terms of modules over rings. In particular, it is not worth memorizing the definition of abelian category.

Two central examples of an abelian category are the category Ab of abelian groups, and the category Mod_A of A -modules. The first is a special case of the second (just take $A = \mathbb{Z}$). As we give the definitions, you should verify that Mod_A is an abelian category.

We first define the notion of *additive category*. We will use it only as a stepping stone to the notion of an abelian category. Two examples you can keep in mind while reading the definition: the category of free A -modules (where A is a ring), and real (or complex) Banach spaces.

1.6.1. Definition. A category \mathcal{C} is said to be **additive** if it satisfies the following properties.

- Ad1. For each $A, B \in \mathcal{C}$, $\text{Mor}(A, B)$ is an abelian group, such that composition of morphisms distributes over addition. (You should think about what this means — it translates to two distinct statements.)
- Ad2. \mathcal{C} has a zero object, denoted 0 . (This is an object that is simultaneously an initial object and a final object, Definition 1.3.2.)
- Ad3. It has products of two objects (a product $A \times B$ for any pair of objects), and hence by induction, products of any finite number of objects.


In an additive category, the morphisms are often called **homomorphisms**, and Mor is denoted by Hom . In fact, this notation Hom is a good indication that you’re working in an additive category. A functor between additive categories preserving the additive structure of Hom , is called an **additive functor**.

1.6.2. Remarks. It is a consequence of the definition of additive category that finite direct products are also finite direct sums (coproducts) — the details don’t matter to us. The symbol \oplus is used for this notion. Also, it is quick to show that additive functors send zero objects to zero objects (show that Z is a 0-object if and only if $\text{id}_Z = 0_Z$; additive functors preserve both id and 0), and preserve products.

One motivation for the name 0-object is that the 0-morphism in the abelian group $\text{Hom}(A, B)$ is the composition $A \rightarrow 0 \rightarrow B$. (We also remark that the notion of 0-morphism thus makes sense in any category with a 0-object.)

The category of A -modules Mod_A is clearly an additive category, but it has even more structure, which we now formalize as an example of an abelian category.

1.6.3. Definition. Let \mathcal{C} be a category with a 0-object (and thus 0-morphisms). A **kernel** of a morphism $f: B \rightarrow C$ is a map $i: A \rightarrow B$ such that $f \circ i = 0$, and that is universal with respect to this property. Diagrammatically:



(Note that the kernel is not just an object; it is a morphism of an object to B.) Hence it is unique up to unique isomorphism by universal property nonsense. The kernel is written $\ker f \rightarrow B$. A **cokernel** (denoted $\operatorname{coker} f$) is defined dually by reversing the arrows — do this yourself. The kernel of $f: B \rightarrow C$ is the limit (§1.4) of the diagram

(1.6.3.1)

$$\begin{array}{ccc} & 0 & \\ & \downarrow & \\ B & \xrightarrow{f} & C \end{array}$$

and similarly the cokernel is a colimit (see (2.6.0.1)).

If $i: A \rightarrow B$ is a monomorphism, then we say that A is a **subobject** of B , where the map i is implicit. There is also the notion of **quotient object**, defined dually to subobject.

An **abelian category** is an additive category satisfying three additional properties.

- (1) Every map has a kernel and cokernel.
- (2) Every monomorphism is the kernel of its cokernel.
- (3) Every epimorphism is the cokernel of its kernel.

It is a nonobvious (and imprecisely stated) fact that every property you want to be true about kernels, cokernels, etc. follows from these three. (Warning: in part of the literature, additional hypotheses are imposed as part of the definition.)

The **image** of a morphism $f: A \rightarrow B$ is defined as $\operatorname{im}(f) = \ker(\operatorname{coker} f)$ whenever it exists (e.g., in every abelian category). The morphism $f: A \rightarrow B$ factors uniquely through $\operatorname{im} f \rightarrow B$ whenever $\operatorname{im} f$ exists, and $A \rightarrow \operatorname{im} f$ is an epimorphism and a cokernel of $\ker f \rightarrow A$ in every abelian category. The reader may want to verify this as a (hard!) exercise.

The cokernel of a monomorphism is called the **quotient**. The quotient of a monomorphism $A \rightarrow B$ is often denoted B/A (with the map from B implicit).

We will leave the foundations of abelian categories untouched. The key thing to remember is that if you understand kernels, cokernels, images and so on in the category of modules over a given ring, you can manipulate objects in any abelian category. This is made precise by Freyd-Mitchell Embedding Theorem (Remark 1.6.4).

However, the abelian categories we will come across will obviously be related to modules, and our intuition will clearly carry over, so we needn't invoke a theorem whose proof we haven't read. For example, we will show that sheaves of abelian groups on a topological space X form an abelian category (§2.6), and the

interpretation in terms of “compatible germs” will connect notions of kernels, cokernels etc. of sheaves of abelian groups to the corresponding notions of abelian groups.

1.6.4. Small remark on chasing diagrams. It is useful to prove facts (and solve exercises) about abelian categories by chasing elements. Unfortunately, some commonly used abelian categories, such as the category of complexes (to be defined in Exercise 1.6.C), do not have “elements” — they are not naturally “sets with additional structure” in any obvious way. Nonetheless, proof by element-chasing can be justified by the Freyd-Mitchell Embedding Theorem: If \mathcal{C} is an abelian category whose objects form a set, such that $\text{Hom}(X, Y)$ is a set for all $X, Y \in \mathcal{C}$, then there is a ring A and an exact, fully faithful functor from \mathcal{C} into Mod_A , which embeds \mathcal{C} as a full subcategory. (Unfortunately, the ring A need not be commutative.) A proof is sketched in [Weib, §1.6], and references to a complete proof are given there. A proof is also given in [KS1, §9.7]. The upshot is that to prove something about a diagram in some abelian category, we may assume that it is a diagram of modules over some ring, and we may then “diagram-chase” elements. Moreover, any fact about kernels, cokernels, and so on that holds in Mod_A holds in any abelian category.

If invoking a theorem whose proof you haven’t read bothers you, a short alternative is Mac Lane’s “elementary rules for chasing diagrams”, [Mac, Thm. 3, p. 200]; [Mac, Lem. 4, p. 201] gives a proof of the Five Lemma (Exercise 1.7.6) as an example.

But in any case, do what you need to do to put your mind at ease, so you can move forward. Do as little as your conscience will allow.

1.6.5. Complexes, exactness, and homology.

(In this entire discussion, we assume we are working in an abelian category.) We say a sequence

$$(1.6.5.1) \quad \cdots \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow \cdots$$

is a **complex at B** if $g \circ f = 0$, and is **exact at B** if $\ker g = \text{im } f$. (More specifically, g has a kernel that is an image of f . Exactness at B implies being a complex at B — do you see why?) A sequence is a **complex** (resp. **exact**) if it is a complex (resp. exact) at each (internal) term. A **short exact sequence** is an exact sequence with five terms, the first and last of which are zeros — in other words, an exact sequence of the form

$$0 \longrightarrow A \longrightarrow B \longrightarrow C \longrightarrow 0.$$

For example, $0 \longrightarrow A \longrightarrow 0$ is exact if and only if $A = 0$;

$$0 \longrightarrow A \xrightarrow{f} B$$

is exact if and only if f is a monomorphism (with a similar statement for $A \xrightarrow{f} B \longrightarrow 0$);

$$0 \longrightarrow A \xrightarrow{f} B \longrightarrow 0$$

is exact if and only if f is an isomorphism; and

$$0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C$$

is exact if and only if f is a kernel of g (with a similar statement for $A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0$). To show some of these facts it may be helpful to prove that (1.6.5.1) is exact at B if and only if the cokernel of f is a cokernel of the kernel of g .

If you would like practice in playing with these notions before thinking about homology, you can prove the Snake Lemma (stated in Example 1.7.5, with a stronger version in Exercise 1.7.B), or the Five Lemma (stated in Example 1.7.6, with a stronger version in Exercise 1.7.C). (I would do this in the category of A -modules, but see [KS1, Lem. 12.1.1, Lem. 8.3.13] for proofs in general.)

If (1.6.5.1) is a complex at B , then its **homology at B** (often denoted by H) is $\ker g / \operatorname{im} f$. (More precisely, there is some monomorphism $\operatorname{im} f \hookrightarrow \ker g$, and that H is the cokernel of this monomorphism.) Therefore, (1.6.5.1) is exact at B if and only if its homology at B is 0. We say that elements of $\ker g$ (assuming the objects of the category are sets with some additional structure) are the **cycles**, and elements of $\operatorname{im} f$ are the **boundaries** (so homology is “cycles mod boundaries”). If the complex is indexed in decreasing order, the indices are often written as subscripts, and H_i is the homology at $A_{i+1} \rightarrow A_i \rightarrow A_{i-1}$. If the complex is indexed in increasing order, the indices are often written as superscripts, and the homology H^i at $A^{i-1} \rightarrow A^i \rightarrow A^{i+1}$ is often called **cohomology**.

An exact sequence

$$(1.6.5.2) \quad A^\bullet: \quad \dots \longrightarrow A^{i-1} \xrightarrow{f^{i-1}} A^i \xrightarrow{f^i} A^{i+1} \xrightarrow{f^{i+1}} \dots$$

can be “factored” into short exact sequences

$$0 \longrightarrow \ker f^i \longrightarrow A^i \longrightarrow \ker f^{i+1} \longrightarrow 0$$

which is helpful in proving facts about long exact sequences by reducing them to facts about short exact sequences.

More generally, if (1.6.5.2) is assumed only to be a complex, then it can be “factored” into short exact sequences.

$$(1.6.5.3) \quad 0 \longrightarrow \ker f^i \longrightarrow A^i \longrightarrow \operatorname{im} f^i \longrightarrow 0$$

$$0 \longrightarrow \operatorname{im} f^{i-1} \longrightarrow \ker f^i \longrightarrow H^i(A^\bullet) \longrightarrow 0$$

1.6.A. EXERCISE. Describe exact sequences

$$(1.6.5.4) \quad 0 \longrightarrow \operatorname{im} f^i \longrightarrow A^{i+1} \longrightarrow \operatorname{coker} f^i \longrightarrow 0$$

$$0 \longrightarrow H^i(A^\bullet) \longrightarrow \operatorname{coker} f^{i-1} \longrightarrow \operatorname{im} f^i \longrightarrow 0$$

(These are somehow dual to (1.6.5.3). In fact in some mirror universe this might have been given as the standard definition of homology.) Assume the category is that of modules over a fixed ring for convenience, but be aware that the result is true for any abelian category.

1.6.B. EXERCISE AND IMPORTANT DEFINITION. Suppose

$$0 \xrightarrow{d^0} A^1 \xrightarrow{d^1} \dots \xrightarrow{d^{n-1}} A^n \xrightarrow{d^n} 0$$

is a complex of finite-dimensional k -vector spaces (often called A^\bullet for short). Define $h^i(A^\bullet) := \dim H^i(A^\bullet)$. Show that $\sum (-1)^i \dim A^i = \sum (-1)^i h^i(A^\bullet)$. In particular, if A^\bullet is exact, then $\sum (-1)^i \dim A^i = 0$. (If you haven't dealt much with cohomology, this will give you some practice.)

1.6.C. IMPORTANT EXERCISE. Suppose \mathcal{C} is an abelian category. Define the **category $Com_{\mathcal{C}}$ of complexes** as follows. The objects are infinite complexes

$$A^\bullet: \quad \dots \longrightarrow A^{i-1} \xrightarrow{f^{i-1}} A^i \xrightarrow{f^i} A^{i+1} \xrightarrow{f^{i+1}} \dots$$

in \mathcal{C} , and the morphisms $A^\bullet \rightarrow B^\bullet$ are commuting diagrams

$$(1.6.5.5) \quad \begin{array}{ccccccc} \dots & \longrightarrow & A^{i-1} & \xrightarrow{f^{i-1}} & A^i & \xrightarrow{f^i} & A^{i+1} \xrightarrow{f^{i+1}} \dots \\ & & \downarrow & & \downarrow & & \downarrow \\ \dots & \longrightarrow & B^{i-1} & \xrightarrow{g^{i-1}} & B^i & \xrightarrow{g^i} & B^{i+1} \xrightarrow{g^{i+1}} \dots \end{array}$$

Show that $Com_{\mathcal{C}}$ is an abelian category. Feel free to deal with the special case of modules over a fixed ring. (Remark for experts: Essentially the same argument shows that $\mathcal{C}^{\mathcal{I}}$ is an abelian category for any small category \mathcal{I} and any abelian category \mathcal{C} . This immediately implies that the category of presheaves on a topological space X with values in an abelian category \mathcal{C} is automatically an abelian category, cf. §2.3.5.)

1.6.D. IMPORTANT EXERCISE. Show that (1.6.5.5) induces a map of homology $H^i(A^\bullet) \rightarrow H^i(B^\bullet)$. Show furthermore that H^i is a covariant functor $Com_{\mathcal{C}} \rightarrow \mathcal{C}$. (Again, feel free to deal with the special case Mod_A .)

We will later define when two maps of complexes are *homotopic* (§24.1), and show that homotopic maps induce the same map on cohomology (Exercise 24.1.A), but we won't need that any time soon.

1.6.6. Theorem (Long exact sequences). — *A short exact sequence of complexes*

$$\begin{array}{ccccccc} 0^\bullet & : & \dots & \longrightarrow & 0 & \longrightarrow & 0 \longrightarrow 0 \longrightarrow \dots \\ \downarrow & & & & \downarrow & & \downarrow \\ A^\bullet & : & \dots & \longrightarrow & A^{i-1} & \xrightarrow{f^{i-1}} & A^i \xrightarrow{f^i} A^{i+1} \xrightarrow{f^{i+1}} \dots \\ \downarrow & & & & \downarrow & & \downarrow \\ B^\bullet & : & \dots & \longrightarrow & B^{i-1} & \xrightarrow{g^{i-1}} & B^i \xrightarrow{g^i} B^{i+1} \xrightarrow{g^{i+1}} \dots \\ \downarrow & & & & \downarrow & & \downarrow \\ C^\bullet & : & \dots & \longrightarrow & C^{i-1} & \xrightarrow{h^{i-1}} & C^i \xrightarrow{h^i} C^{i+1} \xrightarrow{h^{i+1}} \dots \\ \downarrow & & & & \downarrow & & \downarrow \\ 0^\bullet & : & \dots & \longrightarrow & 0 & \longrightarrow & 0 \longrightarrow 0 \longrightarrow \dots \end{array}$$

induces a **long exact sequence in cohomology**

$$\begin{array}{ccccccc} & & & & \dots & \longrightarrow & H^{i-1}(C^\bullet) \longrightarrow \\ & & & & & & \\ H^i(A^\bullet) & \longrightarrow & H^i(B^\bullet) & \longrightarrow & H^i(C^\bullet) & \longrightarrow & \\ & & & & & & \\ H^{i+1}(A^\bullet) & \longrightarrow & \dots & & & & \end{array}$$

(This requires a definition of the **connecting homomorphism** $H^{i-1}(C^\bullet) \rightarrow H^i(A^\bullet)$, which is “natural” in an appropriate sense.) In the category of modules over a ring, Theorem 1.6.6 will come out of our discussion of spectral sequences, see Exercise 1.7.F, but this is a somewhat perverse way of proving it. For a proof in general, see [KS1, Theorem 12.3.3]. You may want to prove it yourself, by first proving a weaker version of the Snake Lemma (Example 1.7.5), where in the hypotheses (1.7.5.1), the 0’s in the bottom left and top right are removed, and in the conclusion (1.7.5.2), the first and last 0’s are removed.

1.6.7. Exactness of functors. If $F: \mathcal{A} \rightarrow \mathcal{B}$ is an additive covariant functor from one abelian category to another, we say that F is **right-exact** if the exactness of

$$A' \longrightarrow A \longrightarrow A'' \longrightarrow 0,$$

in \mathcal{A} implies that

$$F(A') \longrightarrow F(A) \longrightarrow F(A'') \longrightarrow 0$$

is also exact. Dually, we say that F is **left-exact** if the exactness of

$$0 \longrightarrow A' \longrightarrow A \longrightarrow A'' \quad \text{implies}$$

$$0 \longrightarrow F(A') \longrightarrow F(A) \longrightarrow F(A'') \quad \text{is exact.}$$

An additive contravariant functor is **left-exact** if the exactness of

$$A' \longrightarrow A \longrightarrow A'' \longrightarrow 0 \quad \text{implies}$$

$$0 \longrightarrow F(A'') \longrightarrow F(A) \longrightarrow F(A') \quad \text{is exact.}$$

The reader should be able to deduce what it means for a contravariant functor to be **right-exact**.

An additive covariant or contravariant functor is **exact** if it is both left-exact and right-exact.

1.6.E. EXERCISE. Suppose F is an exact functor. Show that applying F to an exact sequence preserves exactness. For example, if F is covariant, and $A' \rightarrow A \rightarrow A''$ is exact, then $FA' \rightarrow FA \rightarrow FA''$ is exact. (This will be generalized in Exercise 1.6.H(c).)

1.6.F. EXERCISE. Suppose A is a ring, $S \subset A$ is a multiplicative subset, and M is an A -module.

(a) Show that localization of A -modules $Mod_A \rightarrow Mod_{S^{-1}A}$ is an exact covariant functor.

(b) Show that $(\cdot) \otimes_A M$ is a right-exact covariant functor $Mod_A \rightarrow Mod_A$. (This is a repeat of Exercise 1.3.H.)

(c) Show that $Hom(M, \cdot)$ is a left-exact covariant functor $Mod_A \rightarrow Mod_A$. If \mathcal{C} is any abelian category, and $C \in \mathcal{C}$, show that $Hom(C, \cdot)$ is a left-exact covariant functor $\mathcal{C} \rightarrow Ab$.

(d) Show that $Hom(\cdot, M)$ is a left-exact contravariant functor $Mod_A \rightarrow Mod_A$. If \mathcal{C} is any abelian category, and $C \in \mathcal{C}$, show that $Hom(\cdot, C)$ is a left-exact contravariant functor $\mathcal{C} \rightarrow Ab$.

1.6.G. EXERCISE. Suppose M is a **finitely presented A -module**: M has a finite number of generators, and with these generators it has a finite number of relations; or usefully equivalently, fits in an exact sequence

$$(1.6.7.1) \quad A^{\oplus q} \rightarrow A^{\oplus p} \rightarrow M \rightarrow 0$$

Use (1.6.7.1) and the left-exactness of Hom to describe an isomorphism

$$S^{-1} Hom_A(M, N) \xrightarrow{\sim} Hom_{S^{-1}A}(S^{-1}M, S^{-1}N).$$

(You might be able to interpret this in light of a variant of Exercise 1.6.H below, for left-exact contravariant functors rather than right-exact covariant functors.)

1.6.8. Example: *Hom doesn't always commute with localization.* In the language of Exercise 1.6.G, take $A = \mathbb{N} = \mathbb{Z}$, $M = \mathbb{Q}$, and $S = \mathbb{Z} \setminus \{0\}$.

1.6.9. ★ Two useful facts in homological algebra.

We now come to two (sets of) facts I wish I had learned as a child, as they would have saved me lots of grief. They encapsulate what is best and worst of abstract nonsense. The statements are so general as to be nonintuitive. The proofs are very short. They generalize some specific behavior that is easy to prove on an ad hoc basis. Once they are second nature to you, many subtle facts will become obvious to you as special cases. And you will see that they will get used (implicitly or explicitly) repeatedly.

1.6.10. ★ Interaction of homology and (right/left-)exact functors.

You might wait to prove this until you learn about cohomology in Chapter 19, when it will first be used in a serious way.

1.6.H. IMPORTANT EXERCISE (THE FHHF THEOREM). This result can take you far, and perhaps for that reason it has sometimes been called the Fernbahnhof (FernbaHnHoF) Theorem, notably in [Vak1, Exer. 1.6.H]. Suppose $F: \mathcal{A} \rightarrow \mathcal{B}$ is a covariant functor of abelian categories, and C^\bullet is a complex in \mathcal{A} .

- (*F right-exact yields* $FH^\bullet \longrightarrow H^\bullet F$) If F is right-exact, describe a natural morphism $FH^\bullet \rightarrow H^\bullet F$. (More precisely, for each i , the left side is F applied to the cohomology at piece i of C^\bullet , while the right side is the cohomology at piece i of FC^\bullet .)
- (*F left-exact yields* $FH^\bullet \longleftarrow H^\bullet F$) If F is left-exact, describe a natural morphism $H^\bullet F \rightarrow FH^\bullet$.

- (c) (F exact yields $FH^\bullet \xleftarrow{\sim} H^\bullet F$) If F is exact, show that the morphisms of (a) and (b) are inverses and thus isomorphisms.

Hint for (a): use $C^i \xrightarrow{d^i} C^{i+1} \longrightarrow \text{coker } d^i \longrightarrow 0$ to give an isomorphism $F \text{coker } d^i \xleftarrow{\sim} \text{coker } Fd^i$. Then use the first line of (1.6.5.4) to give an epimorphism $F \text{im } d^i \twoheadrightarrow \text{im } Fd^i$. Then use the second line of (1.6.5.4) to give the desired map $FH^i C^\bullet \longrightarrow H^i F C^\bullet$. While you are at it, you may as well describe a map for the fourth member of the quartet $\{\text{coker}, \text{im}, H, \ker\}$: $F \ker d^i \longrightarrow \ker Fd^i$.

1.6.11. If this makes your head spin, you may prefer to think of it in the following specific case, where both \mathcal{A} and \mathcal{B} are the category of A -modules, and F is $(\cdot) \otimes N$ for some fixed N -module. Your argument in this case will translate without change to yield a solution to Exercise 1.6.H(a) and (c) in general. If $\otimes N$ is exact, then N is called a **flat** A -module. (The notion of flatness will turn out to be very important, and is discussed in detail in Chapter 25.)

For example, localization is exact (Exercise 1.6.F(a)), so $S^{-1}A$ is a *flat* A -algebra for all multiplicative sets S . Thus taking cohomology of a complex of A -modules commutes with localization — something you could verify directly.

1.6.12. *Interaction of adjoints, (co)limits, and (left- and right-) exactness.*

A surprising number of arguments boil down to the statement:

Limits commute with limits and right adjoints. In particular, in an abelian category, because kernels are limits, both limits and right adjoints are left-exact.

as well as its dual:

Colimits commute with colimits and left adjoints. In particular, because cokernels are colimits, both colimits and left adjoints are right-exact.

These statements were promised in §1.5.4, and will be proved below. The latter has a useful extension:

In Mod_A , colimits over filtered index categories are exact. “Filtered” was defined in §1.4.8.

1.6.13. *★ ★ Caution.* It is not true that in abelian categories in general, colimits over filtered index categories are exact. (Grothendieck realized the desirability of such colimits being exact, and formalized this as his “AB5” axiom, see for example [Stacks, tag 079A].) Here is a counterexample. Because the axioms of abelian categories are self-dual, it suffices to give an example in which a filtered *limit* fails to be exact, and we do this. Fix a prime p . In the category Ab of abelian groups, for each positive integer n , we have an exact sequence $\mathbb{Z} \rightarrow \mathbb{Z}/(p^n) \rightarrow 0$. Taking the limit over all n in the obvious way, we obtain $\mathbb{Z} \rightarrow \mathbb{Z}_p \rightarrow 0$, which is certainly not exact.)

See Unimportant Remark 1.6.16 for another dashed hope.

1.6.14. If you want to use these statements (for example, later in this book), you will have to prove them. Let’s now make them precise.

1.6.I. EXERCISE (KERNELS COMMUTE WITH LIMITS). Suppose \mathcal{C} is an abelian category, and $a: \mathcal{I} \rightarrow \mathcal{C}$ and $b: \mathcal{I} \rightarrow \mathcal{C}$ are two diagrams in \mathcal{C} indexed by \mathcal{I} . For convenience, let $A_i = a(i)$ and $B_i = b(i)$ be the objects in those two diagrams. Let $h_i: A_i \rightarrow B_i$ be maps commuting with the maps in the diagram. (Translation:

h is a natural transformation of functors $a \rightarrow b$, see §1.2.21.) Then the $\ker h_i$ form another diagram in \mathcal{C} indexed by \mathcal{I} . Describe a canonical isomorphism $\lim \ker h_i \xrightarrow{\sim} \ker(\lim A_i \rightarrow \lim B_i)$, assuming the limits exist.

Implicit in the previous exercise is the idea that limits should somehow be understood as functors.

1.6.J. EXERCISE. Make sense of the statement that “limits commute with limits” in a general category, and prove it. (Hint: recall that kernels are limits. The previous exercise should be a corollary of this one.)

1.6.15. Proposition (right adjoints commute with limits). — Suppose $(F: \mathcal{C} \rightarrow \mathcal{D}, G: \mathcal{D} \rightarrow \mathcal{C})$ is a pair of adjoint functors. If $A = \lim A_i$ is a limit in \mathcal{D} of a diagram indexed by \mathcal{I} , then $GA = \lim GA_i$ (with the corresponding maps $GA \rightarrow GA_i$) is a limit in \mathcal{C} .

Proof. We must show that $GA \rightarrow GA_i$ satisfies the universal property of limits. Suppose we have maps $W \rightarrow GA_i$ commuting with the maps of \mathcal{I} . We wish to show that there exists a unique $W \rightarrow GA$ extending the $W \rightarrow GA_i$. By adjointness of F and G , we can restate this as: Suppose we have maps $FW \rightarrow A_i$ commuting with the maps of \mathcal{I} . We wish to show that there exists a unique $FW \rightarrow A$ extending the $FW \rightarrow A_i$. But this is precisely the universal property of the limit. \square

Of course, the dual statements to Exercise 1.6.J and Proposition 1.6.15 hold by the dual arguments.

If F and G are additive functors between abelian categories, and (F, G) is an adjoint pair, then (as kernels are limits and cokernels are colimits) G is left-exact and F is right-exact.

1.6.K. EXERCISE. Show that in Mod_A , colimits over filtered index categories are exact. (Your argument will apply without change to any abelian category whose objects can be interpreted as “sets with additional structure”.) Right-exactness follows from the above discussion, so the issue is left-exactness. (Possible hint: After you show that localization is exact, Exercise 1.6.F(a), or stalkification is exact, Exercise 2.6.D, in a hands-on way, you will be easily able to prove this. Conversely, if you do this exercise, those two will be easy.)

1.6.L. EXERCISE. Show that filtered colimits commute with homology in Mod_A . Hint: use the FHHF Theorem (Exercise 1.6.H), and the previous Exercise.

In light of Exercise 1.6.L, you may want to think about how limits (and colimits) commute with homology in general, and which way maps go. The statement of the FHHF Theorem should suggest the answer. (Are limits analogous to left-exact functors, or right-exact functors?) We won’t directly use this insight, but see §19.1 (vii) for an example.

Just as colimits are exact (not just right-exact) in especially good circumstances, limits are exact (not just left-exact) too. The following will be used twice in Chapter 30.

1.6.M. EXERCISE. Suppose

$$\begin{array}{ccccccc}
 & \vdots & & \vdots & & \vdots & \\
 & \downarrow & & \downarrow & & \downarrow & \\
 0 & \longrightarrow & A_{n+1} & \longrightarrow & B_{n+1} & \longrightarrow & C_{n+1} \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & A_n & \longrightarrow & B_n & \longrightarrow & C_n \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 & \vdots & & \vdots & & \vdots & \\
 & \downarrow & & \downarrow & & \downarrow & \\
 0 & \longrightarrow & A_0 & \longrightarrow & B_0 & \longrightarrow & C_0 \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 & & 0 & & 0 & & 0
 \end{array}$$

is an inverse system of exact sequences of modules over a ring, such that the maps $A_{n+1} \rightarrow A_n$ are surjective. (We say: “transition maps of the left term are surjective”.) Show that the limit

$$(1.6.15.1) \quad 0 \longrightarrow \lim A_n \longrightarrow \lim B_n \longrightarrow \lim C_n \longrightarrow 0$$

is also exact. (You will need to define the maps in (1.6.15.1).)

1.6.16. Unimportant Remark. Based on these ideas, you may suspect that right-exact functors always commute with colimits. The fact that tensor product commutes with infinite direct sums (Exercise 1.3.M) may reinforce this idea. Unfortunately, it is not true — “double dual” $\cdot^{\vee\vee} : \text{Vec}_k \rightarrow \text{Vec}_k$ is covariant and right exact (in fact, exact), but does not commute with infinite direct sums, as $\bigoplus_{i=1}^{\infty} (k^{\vee\vee})$ is not isomorphic to $(\bigoplus_{i=1}^{\infty} k)^{\vee\vee}$.

1.6.17. ★ Dreaming of derived functors. When you see a left-exact functor, you should always dream that you are seeing the end of a long exact sequence. If

$$0 \rightarrow M' \rightarrow M \rightarrow M'' \rightarrow 0$$

is an exact sequence in abelian category \mathcal{A} , and $F: \mathcal{A} \rightarrow \mathcal{B}$ is a left-exact functor, then

$$0 \rightarrow FM' \rightarrow FM \rightarrow FM''$$

is exact, and you should always dream that it should continue in some natural way. For example, the next term should depend only on M' , call it R^1FM' , and if it is zero, then $FM \rightarrow FM''$ is an epimorphism. This remark holds true for left-exact and contravariant functors too. In good cases, such a continuation exists, and is incredibly useful. We will discuss this in Chapter 24.

1.7 ★ Spectral sequences

Je suis quelque peu affolé par ce déluge de cohomologie, mais j'ai courageusement tenu le coup. Ta suite spectrale me paraît raisonnable (je croyais, sur un cas particulier, l'avoir mise en défaut, mais je m'étais trompé, et cela marche au contraire admirablement bien).

I am a bit panic-stricken by this flood of cohomology, but have borne up courageously. Your spectral sequence seems reasonable to me (I thought I had shown that it was wrong in a special case, but I was mistaken, on the contrary it works remarkably well).

— J.-P. Serre, letter to A. Grothendieck, March 14, 1956 [GrS, p. 38]

Spectral sequences are a powerful book-keeping tool for proving things involving complicated commutative diagrams. They were introduced by Leray in the 1940's at the same time as he introduced sheaves. They have a reputation for being abstruse and difficult. It has been suggested that the name 'spectral' was given because, like spectres, spectral sequences are terrifying, evil, and dangerous. I have heard no one disagree with this interpretation, which is perhaps not surprising since I just made it up.

Nonetheless, the goal of this section is to tell you enough that you can use spectral sequences without hesitation or fear, and why you shouldn't be frightened when they come up in a seminar. What is perhaps different in this presentation is that we will use spectral sequences to prove things that you may have already seen, and that you can prove easily in other ways. This will allow you to get some hands-on experience for how to use them. We will also see them only in the special case of double complexes (the version by far the most often used in algebraic geometry), and not in the general form usually presented (filtered complexes, exact couples, etc.). See [Weib, Ch. 5] for more detailed information if you wish.

You should *not* read this section when you are reading the rest of Chapter 1. Instead, you should read it just before you need it for the first time. When you finally *do* read this section, you *must* do the exercises up to Exercise 1.7.F.

For concreteness, we work in the category Mod_A of module over a ring A . However, everything we say will apply in any abelian category. (And if it helps you feel secure, work instead in the category Vec_k of vector spaces over a field k .)

1.7.1. Double complexes.

A **double complex** is a collection of A -modules $E^{p,q}$ ($p, q \in \mathbb{Z}$), and “rightward” morphisms $d_{\rightarrow}^{p,q}: E^{p,q} \rightarrow E^{p+1,q}$ and “upward” morphisms $d_{\uparrow}^{p,q}: E^{p,q} \rightarrow E^{p,q+1}$. In the superscript, the first entry denotes the column number (the “ x -coordinate”), and the second entry denotes the row number (the “ y -coordinate”). (Warning: this is opposite to the convention for matrices.) The subscript is meant to suggest the direction of the arrows. We will always write these as d_{\rightarrow} and d_{\uparrow} and ignore the superscripts. We require that d_{\rightarrow} and d_{\uparrow} satisfy (a) $d_{\rightarrow}^2 = 0$, (b) $d_{\uparrow}^2 = 0$, and one more condition: (c) either $d_{\rightarrow}d_{\uparrow} = d_{\uparrow}d_{\rightarrow}$ (all the squares commute) or $d_{\rightarrow}d_{\uparrow} + d_{\uparrow}d_{\rightarrow} = 0$ (they all anticommute). Both come up in nature, and you can switch from one to the other by replacing $d_{\uparrow}^{p,q}$ with $(-1)^p d_{\uparrow}^{p,q}$. So I will assume that all the squares anticommute, but that you know how to turn the commuting case into this one. (You will see that there is no difference in the recipe, basically because the image and kernel of a homomorphism f equal the image and

kernel respectively of $-f$.)

$$\begin{array}{ccc}
 E^{p,q+1} & \xrightarrow{d_{\rightarrow}^{p,q+1}} & E^{p+1,q+1} \\
 \uparrow d_{\uparrow}^{p,q} & \text{anticommutes} & \uparrow d_{\uparrow}^{p+1,q} \\
 E^{p,q} & \xrightarrow{d_{\rightarrow}^{p,q}} & E^{p+1,q}
 \end{array}$$

There are variations on this definition, where for example the vertical arrows go downwards, or some subset of the $E^{p,q}$ is required to be zero.

From the double complex we construct a corresponding (single) complex E^\bullet with $E^k = \bigoplus_i E^{i,k-i}$, with $d = d_{\rightarrow} + d_{\uparrow}$. In other words, when there is a *single* superscript k , we mean a sum of the k th antidiagonal of the double complex. The single complex is sometimes called the **total complex**. Note that $d^2 = (d_{\rightarrow} + d_{\uparrow})^2 = d_{\rightarrow}^2 + (d_{\rightarrow}d_{\uparrow} + d_{\uparrow}d_{\rightarrow}) + d_{\uparrow}^2 = 0$, so E^\bullet is indeed a complex.

The cohomology of the single complex is sometimes called the **hypercohomology** of the double complex. We will instead use the phrase **cohomology of the double complex**.

Our initial goal will be to find the cohomology of the double complex. You will see later that we secretly also have other goals.

A spectral sequence is a recipe for computing some information about the cohomology of the double complex. I won't yet give the full recipe. Surprisingly, this fragmentary bit of information is sufficient to prove lots of things.

1.7.2. Approximate Definition. A **spectral sequence with rightward orientation** is a sequence of tables or **pages** $\rightarrow E_0^{p,q}, \rightarrow E_1^{p,q}, \rightarrow E_2^{p,q}, \dots$ ($p, q \in \mathbb{Z}$), where $\rightarrow E_0^{p,q} = E^{p,q}$, along with a differential


$$\rightarrow d_r^{p,q}: \rightarrow E_r^{p,q} \rightarrow \rightarrow E_r^{p-r+1, q+r}$$

with $\rightarrow d_r^{p,q} \circ \rightarrow d_r^{p+r-1, q-r} = 0$, and with an isomorphism of the cohomology of $\rightarrow d_r$ at $\rightarrow E_r^{p,q}$ (i.e., $\ker \rightarrow d_r^{p,q} / \text{im } \rightarrow d_r^{p+r-1, q-r}$) with $\rightarrow E_{r+1}^{p,q}$.

The orientation indicates that our 0th differential is the rightward one: $d_0 = d_{\rightarrow}$. The left subscript “ \rightarrow ” is usually omitted.

The order of the morphisms is best understood visually:

(1.7.2.1)



(the morphisms each apply to different pages). Notice that the map always is “degree 1” in terms of the grading of the single complex E^\bullet . (You should figure out what this informal statement really means.)

The actual definition describes what $E_r^{\bullet,\bullet}$ and $d_r^{\bullet,\bullet}$ really are, in terms of $E^{\bullet,\bullet}$. We will describe d_0 , d_1 , and d_2 below, and you should for now take on faith that this sequence continues in some natural way.

Note that $E_r^{p,q}$ is always a subquotient of the corresponding term on the i th page $E_i^{p,q}$ for all $i < r$. In particular, if $E^{p,q} = 0$, then $E_r^{p,q} = 0$ for all r .

Suppose now that $E^{\bullet,\bullet}$ is a **first quadrant double complex**, i.e., $E^{p,q} = 0$ for $p < 0$ or $q < 0$ (so $E_r^{p,q} = 0$ for all r unless $p, q \in \mathbb{Z}^{\geq 0}$). Then for any fixed p, q , once r is sufficiently large, $E_{r+1}^{p,q}$ is computed from $(E_r^{\bullet,\bullet}, d_r)$ using the complex

$$\begin{array}{ccc} 0 & & \\ & \nearrow d_r^{p,q} & \\ & E_r^{p,q} & \\ & \nwarrow d_r^{p+r-1, q-r} & \\ 0 & & \end{array}$$

and thus we have canonical isomorphisms

$$E_r^{p,q} \cong E_{r+1}^{p,q} \cong E_{r+2}^{p,q} \cong \dots$$

We denote this module $E_\infty^{p,q}$. The same idea works in other circumstances, for example if the double complex is only nonzero in a finite number of rows — $E^{p,q} = 0$ unless $q_0 < q < q_1$. This will come up for example in the mapping cone and long exact sequence discussion (Exercises 1.7.F and 1.7.E below).


We now describe the first few pages of the spectral sequence explicitly. As stated above, the differential d_0 on $E_0^{\bullet,\bullet} = E^{\bullet,\bullet}$ is defined to be d_\rightarrow . The rows are complexes:

$$\begin{array}{ccc} \bullet & \longrightarrow & \bullet \longrightarrow \bullet \\ \text{The 0th page } E_0: & & \bullet \longrightarrow \bullet \longrightarrow \bullet \\ & & \bullet \longrightarrow \bullet \longrightarrow \bullet \end{array}$$

and so E_1 is just the table of cohomologies of the rows. You should check that there are now vertical maps $d_1^{p,q}: E_1^{p,q} \rightarrow E_1^{p,q+1}$ of the row cohomology groups, induced by d_\uparrow , and that these make the columns into complexes. (This is essentially the fact that a map of complexes induces a map on homology.) We have


“used up the horizontal morphisms”, but “the vertical differentials live on”.

The 1st page E_1 :



We take cohomology of d_1 on E_1 , giving us a new table, $E_2^{p,q}$. It turns out that there are natural morphisms from each entry to the entry two above and one to the left, and that the composition of these two is 0. (It is a very worthwhile exercise to work out how this natural morphism d_2 should be defined. Your argument may be reminiscent of the connecting homomorphism in the Snake Lemma 1.7.5 or in the long exact sequence in cohomology arising from a short exact sequence of complexes, Exercise 1.6.C. This is no coincidence.)

The 2nd page E_2 :



This is the beginning of a pattern.

Then it is a theorem that there is a filtration of $H^k(E^\bullet)$ by $E_\infty^{p,q}$ where $p + q = k$. (We can't yet state it as an official **Theorem** because we haven't precisely defined the pages and differentials in the spectral sequence.) More precisely, there is a filtration

$$(1.7.2.2) \quad E_\infty^{0,k} \xrightarrow{E_\infty^{1,k-1}} ? \xrightarrow{E_\infty^{2,k-2}} \dots \xrightarrow{E_\infty^{k,0}} H^k(E^\bullet)$$

where the quotients are displayed above each inclusion. (Here is a tip for remember which way the quotients are supposed to go. The differentials on later and later pages point deeper and deeper into the filtration. Thus the entries in the direction of the later arrowheads are the subobjects, and the entries in the direction of the later “arrowtails” are quotients. This tip has the advantage of being independent of the details of the spectral sequence, e.g., the “quadrant” or the orientation.)

We say that the spectral sequence $\rightarrow E_\infty^{\bullet,\bullet}$ **converges** to $H^\bullet(E^\bullet)$. We often say that $\rightarrow E_2^{\bullet,\bullet}$ (or any other page) **abuts** to $H^\bullet(E^\bullet)$.

Although the filtration gives only partial information about $H^\bullet(E^\bullet)$, sometimes one can find $H^\bullet(E^\bullet)$ precisely. One example is if all $E_\infty^{i,k-i}$ are zero, or if all but one of them are zero (e.g., if $E_r^{\bullet,\bullet}$ has precisely one nonzero row or column, in which case one says that the spectral sequence **collapses** at the r th step, although we will not use this term). Another example is in the category of vector spaces over a field, in which case we can find the dimension of $H^k(E^\bullet)$. Also, in lucky circumstances, E_2 (or some other small page) already equals E_∞ .

1.7.A. EXERCISE: INFORMATION FROM THE SECOND PAGE. Show that $H^0(E^\bullet) = E_\infty^{0,0} = E_2^{0,0}$ and


$$0 \longrightarrow E_2^{0,1} \longrightarrow H^1(E^\bullet) \longrightarrow E_2^{1,0} \xrightarrow{d_2^{1,0}} E_2^{0,2} \longrightarrow H^2(E^\bullet)$$

is exact.

1.7.3. The other orientation.

You may have observed that we could as well have done everything in the opposite direction, i.e., reversing the roles of horizontal and vertical morphisms. Then the sequences of arrows giving the spectral sequence would look like this (compare to (1.7.2.1)).

(1.7.3.1)



This spectral sequence is denoted $\uparrow E_{\bullet,\bullet}^\bullet$ (“with the **upward orientation**”). Then we would again get pieces of a filtration of $H^\bullet(E^\bullet)$ (where we have to be a bit careful with the order with which $\uparrow E_\infty^{p,q}$ corresponds to the subquotients — it is the opposite order to that of (1.7.2.2) for $\rightarrow E_\infty^{p,q}$). Warning: in general there is no isomorphism between $\rightarrow E_\infty^{p,q}$ and $\uparrow E_\infty^{p,q}$.

In fact, this observation that we can start with either the horizontal or vertical maps was our secret goal all along. Both algorithms compute information about the same thing ($H^\bullet(E^\bullet)$), and usually we don’t care about the final answer — we often care about the answer we get in one way, and we get at it by doing the spectral sequence in the *other* way.

1.7.4. Examples.

We are now ready to see how this is useful. The moral of these examples is the following. In the past, you may have proved various facts involving various sorts of diagrams, by chasing elements around. Now, you will just plug them into a spectral sequence, and let the spectral sequence machinery do your chasing for you.

1.7.5. Example: Proving the Snake Lemma. Consider the diagram

$$(1.7.5.1) \quad \begin{array}{ccccccc} 0 & \longrightarrow & D & \longrightarrow & E & \longrightarrow & F \longrightarrow 0 \\ & & \uparrow \alpha & & \uparrow \beta & & \uparrow \gamma \\ 0 & \longrightarrow & A & \longrightarrow & B & \longrightarrow & C \longrightarrow 0 \end{array}$$

where the rows are exact in the middle (at A, B, C, D, E, F) and the squares commute. (Normally the Snake Lemma is described with the vertical arrows pointing downwards, but I want to fit this into my spectral sequence conventions.) We wish

to show that there is an exact sequence

$$(1.7.5.2) \quad 0 \rightarrow \ker \alpha \rightarrow \ker \beta \rightarrow \ker \gamma \rightarrow \operatorname{coker} \alpha \rightarrow \operatorname{coker} \beta \rightarrow \operatorname{coker} \gamma \rightarrow 0.$$

We plug this into our spectral sequence machinery. We first compute the cohomology using the rightward orientation, i.e., using the order (1.7.2.1). Then because the rows are exact, $E_1^{p,q} = 0$, so the spectral sequence has already converged: $E_\infty^{p,q} = 0$.

We next compute this “0” in another way, by computing the spectral sequence using the upward orientation. Then $\uparrow E_1^{\bullet,\bullet}$ (with its differentials) is:

$$0 \longrightarrow \operatorname{coker} \alpha \longrightarrow \operatorname{coker} \beta \longrightarrow \operatorname{coker} \gamma \longrightarrow 0$$

$$0 \longrightarrow \ker \alpha \longrightarrow \ker \beta \longrightarrow \ker \gamma \longrightarrow 0.$$

Then $\uparrow E_2^{\bullet,\bullet}$ is of the form:

$$\begin{array}{ccccccc}
 & & 0 & & & & \\
 & & \searrow & & & & \\
 0 & & & 0 & & & \\
 & \searrow & & \searrow & & & \\
 & & 0 & & ?? & & ? & & 0 \\
 & & \searrow & & \searrow & & \searrow & & \searrow \\
 & & & 0 & & ? & & ?? & & 0 \\
 & & & \searrow & & \searrow & & \searrow & & \searrow \\
 & & & & 0 & & ? & & ?? & & 0 \\
 & & & & & & \searrow & & \searrow & & \searrow \\
 & & & & & & & 0 & & 0 & & 0
 \end{array}$$

We see that after $\uparrow E_2$, all the terms will stabilize except for the double question marks — all maps to and from the single question marks are to and from 0-entries. And after $\uparrow E_3$, even these two double-question-mark terms will stabilize. But in the end our complex must be the 0 complex. This means that in $\uparrow E_2$, all the entries must be zero, except for the two double question marks, and these two must be isomorphic. This means that $0 \rightarrow \ker \alpha \rightarrow \ker \beta \rightarrow \ker \gamma$ and $\operatorname{coker} \alpha \rightarrow \operatorname{coker} \beta \rightarrow \operatorname{coker} \gamma \rightarrow 0$ are both exact (that comes from the vanishing of the single question marks), and

$$\operatorname{coker}(\ker \beta \rightarrow \ker \gamma) \cong \ker(\operatorname{coker} \alpha \rightarrow \operatorname{coker} \beta)$$

is an isomorphism (that comes from the equality of the double question marks). Taken together, we have proved the exactness of (1.7.5.2), and hence the Snake Lemma! (Notice: in the end we didn’t really care about the double complex. We just used it as a prop to prove the Snake Lemma.)

Spectral sequences make it easy to see how to generalize results further. For example, if $A \rightarrow B$ is no longer assumed to be injective, how would the conclusion change?

1.7.B. UNIMPORTANT EXERCISE (GRAFTING EXACT SEQUENCES, A VARIANT OF THE SNAKE LEMMA). Extend the Snake Lemma as follows. Suppose we have a

commuting diagram

$$\begin{array}{ccccccccc}
 0 & \longrightarrow & X' & \longrightarrow & Y' & \longrightarrow & Z' & \longrightarrow & A' & \longrightarrow & \dots \\
 \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \\
 \dots & \longrightarrow & W & \longrightarrow & X & \longrightarrow & Y & \longrightarrow & Z & \longrightarrow & 0.
 \end{array}$$

α b c

where the top and bottom rows are exact. Show that the top and bottom rows can be “grafted together” to an exact sequence

$$\dots \longrightarrow W \longrightarrow \ker a \longrightarrow \ker b \longrightarrow \ker c$$

$$\longrightarrow \operatorname{coker} a \longrightarrow \operatorname{coker} b \longrightarrow \operatorname{coker} c \longrightarrow A' \longrightarrow \dots$$

1.7.6. Example: the Five Lemma. Suppose

$$\begin{array}{ccccccccc}
 F & \longrightarrow & G & \longrightarrow & H & \longrightarrow & I & \longrightarrow & J \\
 \uparrow \alpha & & \uparrow \beta & & \uparrow \gamma & & \uparrow \delta & & \uparrow \epsilon \\
 A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & D & \longrightarrow & E
 \end{array}$$

where the rows are exact and the squares commute.

Suppose $\alpha, \beta, \delta, \epsilon$ are isomorphisms. We will show that γ is an isomorphism.

We first compute the cohomology of the total complex using the rightward orientation (1.7.2.1). We choose this because we see that we will get lots of zeros. Then $\rightarrow E_1^\bullet$ looks like this:

$$\begin{array}{ccccc}
 ? & 0 & 0 & 0 & ? \\
 \uparrow & \uparrow & \uparrow & \uparrow & \uparrow \\
 ? & 0 & 0 & 0 & ?
 \end{array}$$

Then $\rightarrow E_2$ looks similar, and the sequence will converge by E_2 , as we will never get any arrows between two nonzero entries in a table thereafter. We can’t conclude that the cohomology of the total complex vanishes, but we can note that it vanishes in all but four degrees — and most important, it vanishes in the two degrees corresponding to the entries C and H (the source and target of γ).

We next compute this using the upward orientation (1.7.3.1). Then $\uparrow E_1$ looks like this:

$$0 \longrightarrow 0 \longrightarrow ? \longrightarrow 0 \longrightarrow 0$$

$$0 \longrightarrow 0 \longrightarrow ? \longrightarrow 0 \longrightarrow 0$$

and the spectral sequence converges at this step. We wish to show that those two question marks are zero. But they are precisely the cohomology groups of the total complex that we just showed *were* zero — so we are done!

The best way to become comfortable with this sort of argument is to try it out yourself several times, and realize that it really is easy. So you should do the following exercises! Many can readily be done directly, but you should deliberately try to use this spectral sequence machinery in order to get practice and develop confidence.

1.7.C. EXERCISE: A SUBTLER FIVE LEMMA. By looking at the spectral sequence proof of the Five Lemma above, prove a subtler version of the Five Lemma, where one of the isomorphisms can instead just be required to be an injection, and another can instead just be required to be a surjection. (I am deliberately not telling you which ones, so you can see how the spectral sequence is telling you how to improve the result.)

1.7.D. EXERCISE: ANOTHER SUBTLE VERSION OF THE FIVE LEMMA. If β and δ (in (1.7.6.1)) are injective, and α is surjective, show that γ is injective. Give the dual statement (whose proof is of course essentially the same).

The next two exercises no longer involve first quadrant double complexes. You will have to think a little to realize why there is no reason for confusion or alarm.

1.7.E. EXERCISE (THE MAPPING CONE). Suppose $\mu: A^\bullet \rightarrow B^\bullet$ is a morphism of complexes. Suppose C^\bullet is the single complex associated to the double complex $A^\bullet \rightarrow B^\bullet$. (C^\bullet is called the **mapping cone** of μ .) Show that there is a long exact sequence of complexes:

$$\cdots \rightarrow H^{i-1}(C^\bullet) \rightarrow H^i(A^\bullet) \rightarrow H^i(B^\bullet) \rightarrow H^i(C^\bullet) \rightarrow H^{i+1}(A^\bullet) \rightarrow \cdots.$$

(There is a slight notational ambiguity here; depending on how you index your double complex, your long exact sequence might look slightly different.) In particular, we will use the fact that μ induces an isomorphism on cohomology if and only if the mapping cone is exact. (We won't use it until the proof of Theorem 19.2.4.)

1.7.F. EXERCISE. Use spectral sequences to show that a short exact sequence of complexes gives a long exact sequence in cohomology (Theorem 1.6.6). (This is a generalization of Exercise 1.7.E.)

The Grothendieck composition-of-functors spectral sequence (Theorem 24.3.5) will be an important example of a spectral sequence that specializes in a number of useful ways.

You are now ready to go out into the world and use spectral sequences to your heart's content!

1.7.7. Complete definition of spectral sequences, and proof.

You should most definitely not read the precise definition of a spectral sequence, and the proof that they work as advertised, any time soon after reading the introduction to spectral sequences above. But after a suitable interval, you should at least flip through a construction and proof to convince yourself that nothing fancy is involved. The idea is not as bad as you might think, see [Vak2].

It is useful to notice that the proof implies that spectral sequences are functorial in the 0th page: the spectral sequence formalism has good functorial properties in the double complex. Unfortunately, Grothendieck's terminology "spectral functor" [Gr1, §2.4] did not catch on.

CHAPTER 2

Sheaves

It is perhaps surprising that geometric spaces are often best understood in terms of (nice) functions on them. For example, a differentiable manifold that is a subset of \mathbb{R}^n can be studied in terms of its smooth (C^∞) functions. Because “geometric spaces” can have few (everywhere-defined) functions, a more precise version of this insight is that the structure of the space can be well understood by considering all functions on all open subsets of the space. This information is encoded in something called a *sheaf*. Sheaves were introduced by Leray in the 1940’s, and Serre introduced them to algebraic geometry. (The reason for the name will be somewhat explained in Remark 2.4.3.) We will define sheaves and describe useful facts about them. We will begin with a motivating example to convince you that the notion is not so foreign.

One reason sheaves are slippery to work with is that they keep track of a huge amount of information, and there are some subtle local-to-global issues. There are also three different ways of getting a hold of them:

- in terms of open sets (the definition §2.2) — intuitive but in some ways the least helpful;
- in terms of stalks (see §2.4.1); and
- in terms of a base of a topology (§2.5).

(Some people strongly prefer the espace étalé interpretation, §2.2.11, as well.) Knowing which to use requires experience, so it is essential to do a number of exercises on different aspects of sheaves in order to truly understand the concept.

2.1 Motivating example: The sheaf of smooth functions

Consider smooth (C^∞) functions on the topological space $X = \mathbb{R}^n$ (or more generally on a manifold X). The sheaf of smooth functions on X is the data of all smooth functions on all open subsets on X . We will see how to manage these data, and observe some of their properties. On each open set $U \subset X$, we have a ring of smooth functions. We denote this ring of functions $\mathcal{O}(U)$.

Given a smooth function on an open set, you can restrict it to a smaller open set, obtaining a smooth function there. In other words, if $U \subset V$ is an inclusion of open sets, we have a “restriction map” $\text{res}_{V,U}: \mathcal{O}(V) \rightarrow \mathcal{O}(U)$.

Take a smooth function on a big open set, and restrict it to a medium open set, and then restrict that to a small open set. The result is the same as if you restrict the smooth function on the big open set directly to the small open set. In other

words, if $U \hookrightarrow V \hookrightarrow W$, then the following diagram commutes:

$$\begin{array}{ccc} \mathcal{O}(W) & \xrightarrow{\text{res}_{W,V}} & \mathcal{O}(V) \\ & \searrow \text{res}_{W,U} \quad \swarrow \text{res}_{V,U} & \\ & \mathcal{O}(U) & \end{array}$$

Next take two smooth functions f_1 and f_2 on a big open set U , and an open cover of U by some collection of open subsets $\{U_i\}$. (We say $\{U_i\}$ **covers** U , or is an **open cover** of U , if $U = \bigcup U_i$.) Suppose that f_1 and f_2 agree on each of these U_i . Then they must have been the same function to begin with. In other words, if $\{U_i\}_{i \in I}$ is a cover of U , and $f_1, f_2 \in \mathcal{O}(U)$, and $\text{res}_{U,U_i} f_1 = \text{res}_{U,U_i} f_2$, then $f_1 = f_2$. Thus we can *identify* functions on an open set by looking at them on a covering by small open sets.

Finally, suppose you are given the same U and cover $\{U_i\}$, take a smooth function on each of the U_i — a function f_1 on U_1 , a function f_2 on U_2 , and so on — and assume they agree on the pairwise overlaps. Then they can be “glued together” to make one smooth function on all of U . In other words, given $f_i \in \mathcal{O}(U_i)$ for all i , such that $\text{res}_{U_i, U_i \cap U_j} f_i = \text{res}_{U_j, U_i \cap U_j} f_j$ for all i and j , then there is some $f \in \mathcal{O}(U)$ such that $\text{res}_{U,U_i} f = f_i$ for all i .

The entire example above would have worked just as well with continuous functions, or real-analytic functions, or just plain real-valued functions. Thus all of these classes of “nice” functions share some common properties. We will soon formalize these properties in the notion of a sheaf.

2.1.1. The germ of a smooth function. Before we do, we first give another definition, that of the germ of a smooth function at a point $p \in X$. Intuitively, it is a “shred” of a smooth function at p . Germs are objects of the form

$$\{(f, \text{open set } U) : p \in U, f \in \mathcal{O}(U)\}$$

modulo the relation that $(f, U) \sim (g, V)$ if there is some open set $W \subset U, V$ containing p where $f|_W = g|_W$ (i.e., $\text{res}_{U,W} f = \text{res}_{V,W} g$). In other words, two functions that are the same in an open neighborhood of p (but may differ elsewhere) have the same germ. We call this set of germs the stalk at p , and denote it \mathcal{O}_p . Notice that the stalk is a ring: you can add two germs, and get another germ: if you have a function f defined on U , and a function g defined on V , then $f + g$ is defined on $U \cap V$. Moreover, $f + g$ is well-defined: if \tilde{f} has the same germ as f , meaning that there is some open set W containing p on which they agree, and \tilde{g} has the same germ as g , meaning they agree on some open W' containing p , then $\tilde{f} + \tilde{g}$ is the same function as $f + g$ on $U \cap V \cap W \cap W'$.

Notice also that if $p \in U$, you get a map $\mathcal{O}(U) \rightarrow \mathcal{O}_p$. Experts may already see that we are talking about germs as colimits.

We can see that \mathcal{O}_p is a local ring as follows. Consider those germs vanishing at p , which we denote $\mathfrak{m}_p \subset \mathcal{O}_p$. They certainly form an ideal: \mathfrak{m}_p is closed under addition, and when you multiply something vanishing at p by any function, the result also vanishes at p . We check that this ideal is maximal by showing that the quotient ring is a field:

$$(2.1.1.1) \quad 0 \longrightarrow \mathfrak{m}_p := \text{ideal of germs vanishing at } p \longrightarrow \mathcal{O}_p \xrightarrow{f \mapsto f(p)} \mathbb{R} \longrightarrow 0$$

2.1.A. EXERCISE. Show that this is the only maximal ideal of \mathcal{O}_p . (Hint: show that every element of $\mathcal{O}_p \setminus \mathfrak{m}_p$ is invertible.)

Note that we can interpret the value of a function at a point, or the value of a germ at a point, as an element of the local ring modulo the maximal ideal. (We will see that this doesn't work for more general sheaves, but *does* work for things behaving like sheaves of functions. This will be formalized in the notion of a *locally ringed space*, which we will see, briefly, in §7.3.)

2.1.2. Aside. Notice that $\mathfrak{m}_p/\mathfrak{m}_p^2$ is a module over $\mathcal{O}_p/\mathfrak{m}_p \cong \mathbb{R}$, i.e., it is a real vector space. It turns out to be naturally (whatever that means) the cotangent space to the differentiable or analytic manifold at p . This insight will prove handy later, when we define tangent and cotangent spaces of schemes.

2.1.B. ★ EXERCISE FOR THOSE WITH DIFFERENTIAL GEOMETRIC BACKGROUND. Prove this. (Rhetorical question for experts: what goes wrong if the sheaf of continuous functions is substituted for the sheaf of smooth functions? What goes wrong if you use the sheaf of C^1 functions?)

2.2 Definition of sheaf and presheaf

We now formalize these notions, by defining presheaves and sheaves. Presheaves are simpler to define, and notions such as kernel and cokernel are straightforward. Sheaves are more complicated to define, and some notions such as cokernel require more thought. But sheaves are more useful because they are in some vague sense more geometric; you can get information about a sheaf locally.

2.2.1. Definition of sheaf and presheaf on a topological space X .

To be concrete, we will define sheaves of sets. However, in the definition the category *Sets* can be replaced by any category, and other important examples are abelian groups Ab , k -vector spaces Vec_k , rings $Rings$, modules over a ring Mod_A , and more. (You may have to think more when dealing with a category of objects that aren't "sets with additional structure", but there aren't any new complications. In any case, this won't be relevant for us, although people who want to do this should start by solving Exercise 2.2.C.) Sheaves (and presheaves) are often written in calligraphic font. The fact that \mathcal{F} is a sheaf on a topological space X is often written as

$$\begin{array}{c} \mathcal{F} \\ \downarrow \\ X \end{array}$$

2.2.2. Definition: Presheaf. A **presheaf** \mathcal{F} on a topological space X is the following data.

- To each open set $U \subset X$, we have a set $\mathcal{F}(U)$ (e.g., the set of differentiable functions in our motivating example). (Notational warning: Several notations are in use, for various good reasons: $\mathcal{F}(U) = \Gamma(U, \mathcal{F}) = H^0(U, \mathcal{F})$. We will use them all.) The elements of $\mathcal{F}(U)$ are called **sections of \mathcal{F} over U** . (§2.2.11 combined

with Exercise 2.2.G gives a motivation for this terminology, although this isn't so important for us.)

By convention, if the “ U ” is omitted, it is implicitly taken to be X : “**sections of \mathcal{F}** ” means “sections of \mathcal{F} over X ”. These are also called **global sections**.

- For each inclusion $U \hookrightarrow V$ of open sets, we have a **restriction map**

$$\text{res}_{V,U}: \mathcal{F}(V) \rightarrow \mathcal{F}(U)$$

(just as we did for differentiable functions). If $f \in \mathcal{F}(V)$, we often write

$$f|_U$$

for $\text{res}_{V,U}(f)$.

The data is required to satisfy the following two conditions.

- The map $\text{res}_{U,U}$ is the identity: $\text{res}_{U,U} = \text{id}_{\mathcal{F}(U)}$.
- If $U \hookrightarrow V \hookrightarrow W$ are inclusions of open sets, then the restriction maps commute, i.e.,

$$\begin{array}{ccc} \mathcal{F}(W) & \xrightarrow{\text{res}_{W,V}} & \mathcal{F}(V) \\ & \searrow \text{res}_{W,U} \quad \swarrow \text{res}_{V,U} & \\ & \mathcal{F}(U) & \end{array}$$

commutes.

2.2.A. EXERCISE FOR CATEGORY-LOVERS: “A PRESHEAF IS THE SAME AS A CONTRAVARIANT FUNCTOR”. Given any topological space X , we have a “category of open sets” (Example 1.2.9), where the objects are the open sets and the morphisms are inclusions. Verify that the data of a presheaf is precisely the data of a contravariant functor from the category of open sets of X to the category of sets. (This interpretation is surprisingly useful.)

2.2.3. Definition: Stalks and germs. We define the stalk of a presheaf at a point in two equivalent ways. One will be hands-on, and the other will be as a colimit.

2.2.4. Define the **stalk** of a presheaf \mathcal{F} at a point p to be the set of **germs** of \mathcal{F} at p , denoted \mathcal{F}_p , as in the example of §2.1.1. Germs correspond to sections over some open set containing p , and two of these sections are considered the same if they agree on some smaller open set. More precisely: the stalk is

$$\{(f, \text{open } U) : p \in U, f \in \mathcal{F}(U)\}$$

modulo the relation that $(f, U) \sim (g, V)$ if there is some open set $W \subset U, V$ where $p \in W$ and $\text{res}_{U,W} f = \text{res}_{V,W} g$. (To explain the agricultural terminology: the French name “germe” is meant to suggest a tiny shoot sprouting from a seed, cf. “germinate”.)

2.2.5. A useful equivalent definition of a stalk is as a colimit of all $\mathcal{F}(U)$ over all open sets U containing p :

$$\mathcal{F}_p = \text{colim } \mathcal{F}(U).$$

The index category is a filtered set (given any two such open sets, there is a third such set contained in both), so these two definitions are the same by Exercise 1.4.E. Hence we can define stalks for sheaves of sets, groups, rings, and other things for

which colimits exist for directed sets. It is very helpful to keep both definitions of stalk in mind at the same time.

If $p \in U$, and $f \in \mathcal{F}(U)$, then the image of f in \mathcal{F}_p is called the **germ of f at p** . (Warning: unlike the example of §2.1.1, in general, the value of a section at a point doesn't make sense.)

2.2.6. Definition: Sheaf. A presheaf is a **sheaf** if it satisfies two more axioms. Notice that these axioms use the additional information of when some open sets cover another.

Identity axiom. If $\{U_i\}_{i \in I}$ is an open cover of U , and $f_1, f_2 \in \mathcal{F}(U)$, and $f_1|_{U_i} = f_2|_{U_i}$ for all i , then $f_1 = f_2$.

(A presheaf satisfying the identity axiom is called a **separated presheaf**, but we will not use that notation in any essential way.)

Gluability axiom. If $\{U_i\}_{i \in I}$ is an open cover of U , then given $f_i \in \mathcal{F}(U_i)$ for all i , such that $f_i|_{U_i \cap U_j} = f_j|_{U_i \cap U_j}$ for all i, j , then there is some $f \in \mathcal{F}(U)$ such that $\text{res}_{U, U_i} f = f_i$ for all i .

In mathematics, definitions often come paired: “at most one” and “at least one”. In this case, identity means there is at most one way to glue, and gluability means that there is at least one way to glue.

(For experts and scholars of the empty set only: an additional axiom sometimes included is that $\mathcal{F}(\emptyset)$ is a one-element set, and in general, for a sheaf with values in a category, $\mathcal{F}(\emptyset)$ is required to be the final object in the category. This actually follows from the above definitions, assuming that the empty product is appropriately defined as the final object.)

Example. If U and V are disjoint, then $\mathcal{F}(U \cup V) = \mathcal{F}(U) \times \mathcal{F}(V)$. Here we use the fact that $\mathcal{F}(\emptyset)$ is the final object.

The **stalk of a sheaf** at a point is just its stalk as a presheaf — the same definition applies — and similarly for the **germs of a section of a sheaf**.

2.2.B. UNIMPORTANT EXERCISE: PRESHEAVES THAT ARE NOT SHEAVES. Show that the following are presheaves on \mathbb{C} (with the classical topology), but not sheaves: (a) bounded functions, (b) holomorphic functions admitting a holomorphic square root.

Both of the presheaves in the previous Exercise satisfy the identity axiom. A “natural” example failing even the identity axiom is implicit in Remark 2.5.5.

We now make a couple of points intended only for category-lovers.

2.2.7. Interpretation in terms of the equalizer exact sequence. The two axioms for a presheaf to be a sheaf can be interpreted as “exactness” of the “equalizer exact sequence”: $\cdot \longrightarrow \mathcal{F}(U) \longrightarrow \prod \mathcal{F}(U_i) \rightrightarrows \prod \mathcal{F}(U_i \cap U_j)$. Identity is exactness at $\mathcal{F}(U)$, and gluability is exactness at $\prod \mathcal{F}(U_i)$. I won't make this precise, or even explain what the double right arrow means. (What is an exact sequence of sets?!) But you may be able to figure it out from the context.

2.2.C. EXERCISE. The identity and gluability axioms may be interpreted as saying that $\mathcal{F}(\cup_{i \in I} U_i)$ is a certain limit. What is that limit?

Here are a number of examples of sheaves.

2.2.D. EXERCISE.

- (a) Verify that the examples of §2.1 are indeed sheaves (of smooth functions, or continuous functions, or real-analytic functions, or plain real-valued functions, on a manifold or \mathbb{R}^n).
- (b) Show that real-valued continuous functions on (open sets of) a topological space X form a sheaf.

2.2.8. Important Example: Restriction of a sheaf. Suppose \mathcal{F} is a sheaf on X , and U is an open subset of X . Define the **restriction of \mathcal{F} to U** , denoted $\mathcal{F}|_U$, to be the collection $\mathcal{F}|_U(V) = \mathcal{F}(V)$ for all open subsets $V \subset U$. Clearly this is a sheaf on U . (Unimportant but fun fact: §2.7 will tell us how to restrict sheaves to *arbitrary* subsets.)

2.2.9. Important Example: the skyscraper sheaf. Suppose X is a topological space, with $p \in X$, and S is a set. Let $i_p: p \rightarrow X$ be the inclusion. Then $i_{p,*}S$ defined by

$$i_{p,*}S(U) = \begin{cases} S & \text{if } p \in U, \text{ and} \\ \{e\} & \text{if } p \notin U \end{cases}$$

forms a sheaf. Here $\{e\}$ is any one-element set. (Check this if it isn't clear to you — what are the restriction maps?) This is called a **skyscraper sheaf** supported at p , because the informal picture of it looks like a skyscraper at p . (Mild caution: this informal picture suggests that the only nontrivial stalk of a skyscraper sheaf is at p , which isn't the case. Exercise 6.1.B(b) gives an example, although it certainly isn't the simplest one.) There is an analogous definition for sheaves of abelian groups, except $i_{p,*}(S)(U) = \{0\}$ if $p \notin U$; and for sheaves with values in a category more generally, $i_{p,*}S(U)$ should be a final object.

(This notation is admittedly hideous, and the alternative $(i_p)_*S$ is equally bad. In §2.2.12 we explain this notation.)

2.2.10. Constant presheaves and constant sheaves. Let X be a topological space, and S a set. Define $\underline{S}_{\text{pre}}(U) = S$ for all open sets U . You will readily verify that $\underline{S}_{\text{pre}}$ forms a presheaf (with restriction maps the identity). This is called the **constant presheaf associated to S** . This isn't (in general) a sheaf. (It may be distracting to say why. Lovers of the empty set will insist that the sheaf axioms force the sections over the empty set to be the final object in the category, i.e., a one-element set. But even if we patch the definition by setting $\underline{S}_{\text{pre}}(\emptyset) = \{e\}$, if S has more than one element, and X is the two-point space with the **discrete topology**, i.e., where every subset is open, you can check that $\underline{S}_{\text{pre}}$ fails gluing.)

2.2.E. EXERCISE (CONSTANT SHEAVES). Now let $\mathcal{F}(U)$ be the maps to S that are *locally constant*, i.e., for any point p in U , there is an open neighborhood of p where the function is constant. Show that this is a *sheaf*. (A better description is this: endow S with the discrete topology, and let $\mathcal{F}(U)$ be the continuous maps $U \rightarrow S$.) This is called the **constant sheaf** (with values in S); do not confuse it with the constant presheaf. We denote this sheaf \underline{S} .

2.2.F. EXERCISE (“MORPHISMS GLUE”). Suppose Y is a topological space. Show that “continuous maps to Y ” form a sheaf of sets on X . More precisely, to each open set U of X , we associate the set of continuous maps of U to Y . Show that this

forms a sheaf. (Exercise 2.2.D(b), with $Y = \mathbb{R}$, and Exercise 2.2.E, with $Y = S$ with the discrete topology, are both special cases.)

2.2.G. EXERCISE. This is a fancier version of the previous exercise.

(a) (sheaf of sections of a map) Suppose we are given a continuous map $\mu: Y \rightarrow X$. Show that “sections of μ ” form a sheaf. More precisely, to each open set U of X , associate the set of continuous maps $s: U \rightarrow Y$ such that $\mu \circ s = \text{id}_U$. Show that this forms a sheaf. (For those who have heard of vector bundles, these are a good example.) This is motivation for the phrase “section of a sheaf”.

(b) (This exercise is for those who know what a topological group is. If you don’t know what a topological group is, you might be able to guess.) Suppose that Y is a topological group. Show that continuous maps to Y form a sheaf of groups.

2.2.11. ★ The space of sections (espace étalé) of a (pre)sheaf. Depending on your background, you may prefer the following perspective on sheaves. Suppose \mathcal{F} is a presheaf (e.g., a sheaf) on a topological space X . Construct a topological space F along with a continuous map $\pi: F \rightarrow X$ as follows: as a set, F is the disjoint union of all the stalks of \mathcal{F} . This naturally gives a map of sets $\pi: F \rightarrow X$. Topologize F as follows. Each s in $\mathcal{F}(U)$ determines a subset $\{(x, s_x) : x \in U\}$ of F . The topology on F is the weakest topology such that these subsets are open. (These subsets form a base of the topology. For each $y \in F$, there is an open neighborhood V of y and an open neighborhood U of $\pi(y)$ such that $\pi|_V$ is a homeomorphism from V to U . Do you see why these facts are true?) The topological space F could be thought of as the **space of sections** of \mathcal{F} (and in French is called the **espace étalé** of \mathcal{F}). We will not discuss this construction at any length, but it can have some advantages: (a) It is always better to know as many ways as possible of thinking about a concept. (b) Pullback has a natural interpretation in this language (mentioned briefly in Exercise 2.7.C). (c) Sheafification has a natural interpretation in this language (see Remark 2.4.7).

2.2.H. IMPORTANT EXERCISE / DEFINITION: THE PUSHFORWARD SHEAF OR DIRECT IMAGE SHEAF. Suppose $\pi: X \rightarrow Y$ is a continuous map, and \mathcal{F} is a presheaf on X . Then define $\pi_*\mathcal{F}$ by $\pi_*\mathcal{F}(V) = \mathcal{F}(\pi^{-1}(V))$, where V is an open subset of Y . Show that $\pi_*\mathcal{F}$ is a presheaf on Y , and is a sheaf if \mathcal{F} is. This is called the **pushforward** or **direct image** of \mathcal{F} . More precisely, $\pi_*\mathcal{F}$ is called the **pushforward of \mathcal{F} by π** .

2.2.12. As the notation suggests, the skyscraper sheaf (Example 2.2.9) can be interpreted as the pushforward of the constant sheaf \underline{S} on a one-point space p , under the inclusion morphism $i_p: \{p\} \rightarrow X$.

Once we endow sheaves with the structure of a category, we will see that the pushforward is a functor from sheaves on X to sheaves on Y (Exercise 2.3.B).

2.2.I. EXERCISE (PUSHFORWARD INDUCES MAPS OF STALKS). Suppose $\pi: X \rightarrow Y$ is a continuous map, and \mathcal{F} is a sheaf of sets (or rings or A -modules) on X . If $\pi(p) = q$, describe the natural morphism of stalks $(\pi_*\mathcal{F})_q \rightarrow \mathcal{F}_p$. (You can use the explicit definition of stalk using representatives, §2.2.4, or the universal property, §2.2.5. If you prefer one way, you should try the other.)

2.2.13. Important Example: Ringed spaces, and \mathcal{O}_X -modules. Suppose \mathcal{O}_X is a sheaf of rings on a topological space X (i.e., a sheaf on X with values in the

category of *Rings*). Then (X, \mathcal{O}_X) is called a **ringed space**. The sheaf of rings is often denoted by \mathcal{O}_X , pronounced “oh-X”. This sheaf is called the **structure sheaf** of the ringed space. Sections of the structure sheaf \mathcal{O}_X over an open subset U are called **functions on U** . Functions on X are called **global functions**, or just **functions**. (Caution: what we call “functions”, others sometimes call “regular functions”. Furthermore, we will later define “rational functions” on schemes in §5.2.1 and §6.5.34, which are not precisely functions in this sense; they are a particular type of “partially-defined function”.)

The symbol \mathcal{O}_X will always refer to the structure sheaf of a ringed space X . The restriction $\mathcal{O}_X|_U$ of \mathcal{O}_X to an open subset $U \subset X$ is denoted \mathcal{O}_U . (We will later call $(U, \mathcal{O}_U) \rightarrow (X, \mathcal{O}_X)$ an *open embedding* of ringed spaces, see Definition 7.2.1.) The stalk of \mathcal{O}_X at a point p is written “ $\mathcal{O}_{X,p}$ ”, because this looks less hideous than “ \mathcal{O}_{X_p} ”.

Just as we have modules over a ring, we have \mathcal{O}_X -modules over a sheaf of rings \mathcal{O}_X . There is only one possible definition that could go with the name \mathcal{O}_X -**module** (or often **\mathcal{O} -module**) — a sheaf of abelian groups \mathcal{F} with the following additional structure. For each U , $\mathcal{F}(U)$ is an $\mathcal{O}_X(U)$ -module. Furthermore, this structure should behave well with respect to restriction maps: if $U \subset V$, then

$$(2.2.13.1) \quad \begin{array}{ccc} \mathcal{O}_X(V) \times \mathcal{F}(V) & \xrightarrow{\text{action}} & \mathcal{F}(V) \\ \text{res}_{V,U} \times \text{res}_{V,U} \downarrow & & \downarrow \text{res}_{V,U} \\ \mathcal{O}_X(U) \times \mathcal{F}(U) & \xrightarrow{\text{action}} & \mathcal{F}(U) \end{array}$$

commutes. (You should convince yourself that I haven’t forgotten anything.)

Recall that the notion of A -module generalizes the notion of abelian group, because an abelian group is the same thing as a \mathbb{Z} -module. Similarly, the notion of \mathcal{O}_X -module generalizes the notion of sheaf of abelian groups, because the latter is the same thing as a \mathbb{Z} -module. Hence when we are proving things about \mathcal{O}_X -modules, we are also proving things about sheaves of abelian groups.

2.2.J. EXERCISE. If (X, \mathcal{O}_X) is a ringed space, and \mathcal{F} is an \mathcal{O}_X -module, describe how for each $p \in X$, \mathcal{F}_p is an $\mathcal{O}_{X,p}$ -module.

2.2.14. *For those who know about vector bundles.* The motivating example of \mathcal{O}_X -modules is the sheaf of sections of a vector bundle. If (X, \mathcal{O}_X) is a differentiable manifold (so \mathcal{O}_X is the sheaf of smooth functions), and $\pi: V \rightarrow X$ is a vector bundle over X , then the sheaf of smooth sections $\sigma: X \rightarrow V$ is an \mathcal{O}_X -module. Indeed, given a section s of π over an open subset $U \subset X$, and a function f on U , we can multiply s by f to get a new section fs of π over U . Moreover, if U' is a smaller subset, then we could multiply f by s and then restrict to U' , or we could restrict both f and s to U' and then multiply, and we would get the same answer. That is precisely the commutativity of (2.2.13.1).

2.3 Morphisms of presheaves and sheaves

2.3.1. Definitions. Whenever one defines a new mathematical object, category theory teaches to try to understand maps between them. We now define morphisms of presheaves, and similarly for sheaves. In other words, we will describe the *category of presheaves* (of sets, abelian groups, etc.) and the *category of sheaves*.

A **morphism of presheaves** of sets (or indeed of presheaves with values in any category) on X , $\phi: \mathcal{F} \rightarrow \mathcal{G}$, is the data of maps $\phi(U): \mathcal{F}(U) \rightarrow \mathcal{G}(U)$ for all U behaving well with respect to restriction: if $U \hookrightarrow V$ then

$$\begin{array}{ccc} \mathcal{F}(V) & \xrightarrow{\phi(V)} & \mathcal{G}(V) \\ \text{res}_{V,U} \downarrow & & \downarrow \text{res}_{V,U} \\ \mathcal{F}(U) & \xrightarrow{\phi(U)} & \mathcal{G}(U) \end{array}$$

commutes. (Notice: the underlying space of both \mathcal{F} and \mathcal{G} is X .)

Morphisms of sheaves are defined identically: the morphisms from a sheaf \mathcal{F} to a sheaf \mathcal{G} are precisely the morphisms from \mathcal{F} to \mathcal{G} as presheaves. (Translation: The category of sheaves on X is a full subcategory of the category of presheaves on X .) If (X, \mathcal{O}_X) is a ringed space, then morphisms of \mathcal{O}_X -modules have the obvious definition. (Can you write it down?)

An example of a morphism of sheaves is the map from the sheaf of smooth functions on \mathbb{R} to the sheaf of continuous functions. This is a “forgetful map”: we are forgetting that these functions are differentiable, and remembering only that they are continuous.

2.3.2. Notation. We may as well set some notation: let Sets_X , Ab_X , etc. denote the category of sheaves of sets, abelian groups, etc. on a topological space X . Let $\text{Mod}_{\mathcal{O}_X}$ denote the category of \mathcal{O}_X -modules on a ringed space (X, \mathcal{O}_X) . Let $\text{Sets}_X^{\text{pre}}$, etc. denote the category of presheaves of sets, etc. on X .

2.3.3. Aside for category-lovers. If you interpret a presheaf on X as a contravariant functor (from the category of open sets), a morphism of presheaves on X is a natural transformation of functors (§1.2.21).

2.3.A. EXERCISE: MORPHISMS OF (PRE)SHEAVES INDUCE MORPHISMS OF STALKS. If $\phi: \mathcal{F} \rightarrow \mathcal{G}$ is a morphism of presheaves on X , and $p \in X$, describe an induced morphism of stalks $\phi_p: \mathcal{F}_p \rightarrow \mathcal{G}_p$. Translation: taking the stalk at p induces a functor $\text{Sets}_X \rightarrow \text{Sets}$. (Your proof will extend in obvious ways. For example, if ϕ is a morphism of \mathcal{O}_X -modules, then ϕ_p is a map of $\mathcal{O}_{X,p}$ -modules.)

2.3.B. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a continuous map of topological spaces (i.e., a morphism in the category of topological spaces). Show that pushforward gives a functor $\pi_*: \text{Sets}_X \rightarrow \text{Sets}_Y$. Here Sets can be replaced by other categories. (Watch out for some possible confusion: a presheaf is a functor, and presheaves form a category. It may be best to forget that presheaves are functors for now.)

2.3.C. IMPORTANT EXERCISE AND DEFINITION: “SHEAF Hom”. Suppose \mathcal{F} and \mathcal{G} are two sheaves of sets on X . (In fact, it will suffice that \mathcal{F} is a presheaf.) Let $\text{Hom}(\mathcal{F}, \mathcal{G})$ be the collection of data

$$\text{Hom}(\mathcal{F}, \mathcal{G})(U) := \text{Mor}(\mathcal{F}|_U, \mathcal{G}|_U).$$

(Recall the notation $\mathcal{F}|_U$, the restriction of the sheaf to the open set U , Example 2.2.8.) Show that this is a sheaf of sets on X . (To avoid a common confusion: the right side does *not* say $\text{Mor}(\mathcal{F}(U), \mathcal{G}(U))$.) This sheaf is called “sheaf Hom ”. (Strictly speaking, we should reserve Hom for when we are in an additive category, so this should possibly be called “sheaf Mor ”. But the terminology “sheaf Hom ” is too established to uproot.) It will be clear from your construction that, like Hom , Hom is a contravariant functor in its first argument and a covariant functor in its second argument.

Warning: Hom does not commute with taking stalks. More precisely: it is not true that $\text{Hom}(\mathcal{F}, \mathcal{G})_p$ is isomorphic to $\text{Hom}(\mathcal{F}_p, \mathcal{G}_p)$. (Can you think of a counterexample? There is at least a map from one of these to the other — in which direction?)

2.3.4. We will use many variants of the definition of Hom . For example, if \mathcal{F} and \mathcal{G} are sheaves of abelian groups on X , then $\text{Hom}_{\text{Ab}_X}(\mathcal{F}, \mathcal{G})$ is defined by taking $\text{Hom}_{\text{Ab}_X}(\mathcal{F}, \mathcal{G})(U)$ to be the maps *as sheaves of abelian groups* $\mathcal{F}|_U \rightarrow \mathcal{G}|_U$. (Note that $\text{Hom}_{\text{Ab}_X}(\mathcal{F}, \mathcal{G})$ has the structure of a sheaf of abelian groups in a natural way.) Similarly, if \mathcal{F} and \mathcal{G} are \mathcal{O}_X -modules, we define $\text{Hom}_{\text{Mod}_{\mathcal{O}_X}}(\mathcal{F}, \mathcal{G})$ in the analogous way (and it is an \mathcal{O}_X -module). Obnoxiously, the subscripts Ab_X and $\text{Mod}_{\mathcal{O}_X}$ are often dropped (here and in the literature), so be careful which category you are working in! We call $\text{Hom}_{\text{Mod}_{\mathcal{O}_X}}(\mathcal{F}, \mathcal{O}_X)$ the **dual** of the \mathcal{O}_X -module \mathcal{F} , and denote it \mathcal{F}^\vee .

2.3.D. UNIMPORTANT EXERCISE (REALITY CHECK).

- (a) If \mathcal{F} is a sheaf of sets on X , then show that $\text{Hom}(\{p\}, \mathcal{F}) \cong \mathcal{F}$, where $\{p\}$ is the constant sheaf “with values in the one element set $\{p\}$ ”.
- (b) If \mathcal{F} is a sheaf of abelian groups on X , then show that $\text{Hom}_{\text{Ab}_X}(\mathbb{Z}, \mathcal{F}) \cong \mathcal{F}$ (an isomorphism of sheaves of abelian groups).
- (c) If \mathcal{F} is an \mathcal{O}_X -module, then show that $\text{Hom}_{\text{Mod}_{\mathcal{O}_X}}(\mathcal{O}_X, \mathcal{F}) \cong \mathcal{F}$ (an isomorphism of \mathcal{O}_X -modules).

A key idea in (b) and (c) is that 1 “generates” (in some sense) \mathbb{Z} (in (b)) and \mathcal{O}_X (in (c)).

2.3.5. Presheaves of abelian groups (and even “presheaf \mathcal{O}_X -modules”) form an abelian category.

We can make module-like constructions using presheaves of abelian groups on a topological space X . (Throughout this section, all (pre)sheaves are of abelian groups.) For example, we can clearly add maps of presheaves and get another map of presheaves: if $\phi, \psi: \mathcal{F} \rightarrow \mathcal{G}$, then we define the map $\phi + \psi$ by $(\phi + \psi)(V) = \phi(V) + \psi(V)$. (There is something small to check here: that the result is indeed a map of presheaves.) In this way, presheaves of abelian groups form an additive category (Definition 1.6.1: the morphisms between any two presheaves of abelian groups form an abelian group; there is a 0-object; and one can take finite products). For exactly the same reasons, sheaves of abelian groups also form an additive category.

If $\phi: \mathcal{F} \rightarrow \mathcal{G}$ is a morphism of presheaves, define the **presheaf kernel** $\ker_{\text{pre}} \phi$ by $(\ker_{\text{pre}} \phi)(U) := \ker \phi(U)$.

2.3.E. EXERCISE. Show that $\ker_{\text{pre}} \phi$ is a presheaf. (Hint: if $U \hookrightarrow V$, define the restriction map by chasing the following diagram:

$$\begin{array}{ccccccc}
 0 & \longrightarrow & \ker_{\text{pre}} \phi(V) & \longrightarrow & \mathcal{F}(V) & \longrightarrow & \mathcal{G}(V) \\
 & & \downarrow \exists! & & \downarrow \text{res}_{V,U} & & \downarrow \text{res}_{V,U} \\
 0 & \longrightarrow & \ker_{\text{pre}} \phi(U) & \longrightarrow & \mathcal{F}(U) & \longrightarrow & \mathcal{G}(U)
 \end{array}$$

You should check that the restriction maps compose as desired.)

Define the **presheaf cokernel** $\text{coker}_{\text{pre}} \phi$ similarly. It is a presheaf by essentially the same (dual) argument.

2.3.F. EXERCISE: THE COKERNEL DESERVES ITS NAME. Show that the presheaf cokernel satisfies the universal property of cokernels (Definition 1.6.3) in the category of presheaves.

Similarly, $\ker_{\text{pre}} \phi \rightarrow \mathcal{F}$ satisfies the universal property for kernels in the category of presheaves.

It is not too tedious to verify that presheaves of abelian groups form an abelian category, and the reader is free to do so. The key idea is that all abelian-categorical notions may be defined and verified “open set by open set”. We needn’t worry about restriction maps — they “come along for the ride”. Hence we can define terms such as **subpresheaf**, **image presheaf** (or **presheaf image**), and **quotient presheaf** (or **presheaf quotient**), and they behave as you would expect. You construct kernels, quotients, cokernels, and images open set by open set. Homological algebra (exact sequences and so forth) works, and also “works open set by open set”. In particular:

2.3.G. EASY EXERCISE. Show (or observe) that for a topological space X with open set U , $\mathcal{F} \mapsto \mathcal{F}(U)$ gives a functor from presheaves of abelian groups on X , Ab_X^{pre} , to abelian groups, Ab . Then show that this functor is exact.

2.3.H. EXERCISE. Show that a sequence of presheaves $0 \rightarrow \mathcal{F}_1 \rightarrow \mathcal{F}_2 \rightarrow \cdots \rightarrow \mathcal{F}_n \rightarrow 0$ is exact if and only if $0 \rightarrow \mathcal{F}_1(U) \rightarrow \mathcal{F}_2(U) \rightarrow \cdots \rightarrow \mathcal{F}_n(U) \rightarrow 0$ is exact for all U .

The above discussion essentially carries over without change to presheaves with values in an abelian category. (Think this through if you wish.)

However, we are interested in more geometric objects, sheaves, where things can be understood in terms of their local behavior, thanks to the identity and gluing axioms. We will soon see that sheaves of abelian groups also form an abelian category, but a complication will arise that will force the notion of *sheafification* on us. Sheafification will be the answer to many of our prayers. We just haven’t yet realized what we should be praying for.

To begin with, sheaves Ab_X form an additive category, as described in the first paragraph of §2.3.5.

Kernels work just as with presheaves:

2.3.I. IMPORTANT EXERCISE. Suppose $\phi: \mathcal{F} \rightarrow \mathcal{G}$ is a morphism of *sheaves*. Show that the presheaf kernel $\ker_{\text{pre}} \phi$ is in fact a sheaf. Show that it satisfies the universal property of kernels (Definition 1.6.3). (Hint: the second question follows immediately from the fact that $\ker_{\text{pre}} \phi$ satisfies the universal property in the category of *presheaves*.)

Thus if ϕ is a morphism of sheaves, we define

$$\ker \phi := \ker_{\text{pre}} \phi.$$

The problem arises with the cokernel.

2.3.J. IMPORTANT EXERCISE. Let X be \mathbb{C} with the classical topology, let \mathcal{O}_X be the sheaf of holomorphic functions, and let \mathcal{F} be the *presheaf* of functions admitting a holomorphic logarithm. Describe an exact sequence of presheaves on X :

$$0 \longrightarrow \mathbb{Z} \longrightarrow \mathcal{O}_X \longrightarrow \mathcal{F} \longrightarrow 0$$

where $\mathbb{Z} \rightarrow \mathcal{O}_X$ is the natural inclusion and $\mathcal{O}_X \rightarrow \mathcal{F}$ is given by $f \mapsto \exp(2\pi i f)$. (Be sure to verify exactness.) Show that \mathcal{F} is *not* a sheaf. (Hint: \mathcal{F} does not satisfy the gluability axiom. The problem is that there are functions that don't have a logarithm but locally have a logarithm.) This will come up again in Example 2.4.9.

We will have to put our hopes for understanding cokernels of sheaves on hold for a while. We will first learn to understand sheaves using stalks.

2.4 Properties determined at the level of stalks, and sheafification

2.4.1. Properties determined by stalks. We now come to the second way of understanding sheaves mentioned at the start of the chapter. In this section, we will see that lots of facts about sheaves can be checked “at the level of stalks”. We call any property determined at the level of stalks **stalk-local**. This isn't true for presheaves, and reflects the local nature of sheaves. We will see that sections and morphisms are determined “by their stalks”, and the property of a morphism being an isomorphism may be checked at stalks. (The last one is the trickiest.)

2.4.A. IMPORTANT EASY EXERCISE (sections are determined by germs). Prove that a section of a sheaf of sets is determined by its germs, i.e., the natural map

$$(2.4.1.1) \quad \mathcal{F}(U) \rightarrow \prod_{p \in U} \mathcal{F}_p$$

is injective. Hint 1: you won't use the gluability axiom, so this is true for separated presheaves. Hint 2: it is false for presheaves in general, see Exercise 2.4.E, so you *will* use the identity axiom. (Your proof will also apply to sheaves of groups, rings, etc. — to categories of “sets with additional structure”. The same is true of many exercises in this section.)

Exercise 2.4.A suggests a question: which elements of the right side of (2.4.1.1) are in the image of the left side?

2.4.2. Important definition. We say that an element $(s_p)_{p \in U}$ of the right side $\prod_{p \in U} \mathcal{F}_p$ of (2.4.1.1) consists of **compatible germs** if for all $p \in U$, there is some representative

$$(\tilde{s}_p \in \mathcal{F}(U_p), U_p \text{ open in } U)$$

for s_p (where $p \in U_p \subset U$) such that the germ of \tilde{s}_p at all $q \in U_p$ is s_q . Equivalently, there is an open cover $\{U_i\}$ of U , and sections $f_i \in \mathcal{F}(U_i)$, such that if $p \in U_i$, then s_p is the germ of f_i at p . Clearly any section s of \mathcal{F} over U gives a choice of compatible germs for U .

2.4.B. IMPORTANT EXERCISE. Prove that any choice of compatible germs for a sheaf of sets \mathcal{F} over U is the image of a section of \mathcal{F} over U . (Hint: you will use gluability.)

We have thus completely described the image of (2.4.1.1), in a way that will prove useful.

2.4.3. Remark. This perspective motivates the agricultural terminology “sheaf”: a sheaf is (the data of) a bunch of stalks, bundled together appropriately.

Now we throw morphisms into the mix. Recall Exercise 2.3.A: morphisms of (pre)sheaves induce morphisms of stalks.

2.4.C. EXERCISE (morphisms are determined by stalks). If ϕ_1 and ϕ_2 are morphisms from a presheaf of sets \mathcal{F} to a sheaf of sets \mathcal{G} that induce the same maps on each stalk, show that $\phi_1 = \phi_2$. Hint: consider the following diagram.

$$(2.4.3.1) \quad \begin{array}{ccc} \mathcal{F}(U) & \longrightarrow & \mathcal{G}(U) \\ \downarrow & & \downarrow \\ \prod_{p \in U} \mathcal{F}_p & \longrightarrow & \prod_{p \in U} \mathcal{G}_p \end{array}$$

2.4.D. TRICKY EXERCISE (isomorphisms are determined by stalks). Show that a morphism of sheaves of sets is an isomorphism if and only if it induces an isomorphism of all stalks. Hint: Use (2.4.3.1). Once you have injectivity, show surjectivity, perhaps using Exercise 2.4.B, or gluability in some other way; this is more subtle. *Warning:* This exercise does *not* say that if two sheaves have isomorphic stalks, then they are isomorphic.

2.4.E. EXERCISE.

(a) Show that Exercise 2.4.A is false for general presheaves.

(b) Show that Exercise 2.4.C is false for general presheaves.

(c) Show that Exercise 2.4.D is false for general presheaves.

(General hint for finding counterexamples of this sort: consider a 2-point space with the discrete topology.)

2.4.4. Sheafification.

Every sheaf is a presheaf (and indeed by definition sheaves on X form a full subcategory of the category of presheaves on X). Just as groupification (§1.5.3)

gives an abelian group that best approximates an abelian semigroup, sheafification gives the sheaf that best approximates a presheaf, with an analogous universal property. (One possible example to keep in mind is the sheafification of the presheaf of holomorphic functions admitting a square root on \mathbb{C} with the classical topology. See also the exponential exact sequence, Example 2.4.9.)

2.4.5. Definition. If \mathcal{F} is a presheaf on X , then a morphism of presheaves $\text{sh}: \mathcal{F} \rightarrow \mathcal{F}^{\text{sh}}$ on X is a **sheafification of \mathcal{F}** if \mathcal{F}^{sh} is a sheaf, and for any sheaf \mathcal{G} , and any presheaf morphism $g: \mathcal{F} \rightarrow \mathcal{G}$, there exists a unique morphism of sheaves $f: \mathcal{F}^{\text{sh}} \rightarrow \mathcal{G}$ making the diagram

$$\begin{array}{ccc} \mathcal{F} & \xrightarrow{\text{sh}} & \mathcal{F}^{\text{sh}} \\ & \searrow g & \downarrow f \\ & & \mathcal{G} \end{array}$$

commute.

We still have to show that it exists. The following two exercises require existence (which we will show shortly), but not the details of the construction.

2.4.F. EXERCISE. Show that sheafification is unique up to unique isomorphism, assuming it exists. Show that if \mathcal{F} is a sheaf, then the sheafification is $\text{id}: \mathcal{F} \longrightarrow \mathcal{F}$. (This should be second nature by now.)

2.4.G. EASY EXERCISE (SHEAFIFICATION IS A FUNCTOR). Assume for now that sheafification exists. Use the universal property to show that for any morphism of presheaves $\phi: \mathcal{F} \rightarrow \mathcal{G}$, we get a natural induced morphism of sheaves $\phi^{\text{sh}}: \mathcal{F}^{\text{sh}} \rightarrow \mathcal{G}^{\text{sh}}$. Show that sheafification is a functor from presheaves on X to sheaves on X .

2.4.6. Construction. We next show that any presheaf of sets (or groups, rings, etc.) has a sheafification. Suppose \mathcal{F} is a *presheaf*. Define \mathcal{F}^{sh} by defining $\mathcal{F}^{\text{sh}}(\mathcal{U})$ as the set of “compatible germs” of the presheaf \mathcal{F} over \mathcal{U} . Explicitly:

$$\mathcal{F}^{\text{sh}}(\mathcal{U}) := \{(f_p \in \mathcal{F}_p)_{p \in \mathcal{U}} : \text{for all } p \in \mathcal{U}, \text{ there exists an open neighborhood } V \text{ of } p, \text{ contained in } \mathcal{U}, \text{ and } s \in \mathcal{F}(V) \text{ with } s_q = f_q \text{ for all } q \in V\}.$$

Here s_q means the germ of s at q — the image of s in the stalk \mathcal{F}_q .

2.4.H. EASY EXERCISE. Show that \mathcal{F}^{sh} (using the tautological restriction maps) forms a sheaf.

2.4.I. EASY EXERCISE. Describe a natural map of presheaves $\text{sh}: \mathcal{F} \rightarrow \mathcal{F}^{\text{sh}}$.

2.4.J. EXERCISE. Show that the map sh satisfies the universal property of sheafification (Definition 2.4.5). (This is easier than you might fear.)

2.4.K. USEFUL EXERCISE, NOT JUST FOR CATEGORY-LOVERS. Show that the sheafification functor is left-adjoint to the forgetful functor from sheaves on X to presheaves on X . This is not difficult — it is largely a restatement of the universal property. But it lets you use results from §1.6.12, and can “explain” why you don’t need to sheafify when taking kernel (why the presheaf kernel is already the sheaf kernel),

and why you need to sheafify when taking cokernel and (soon, in Exercise 2.6.J) \otimes .

2.4.L. EXERCISE. Show $\mathcal{F} \rightarrow \mathcal{F}^{\text{sh}}$ induces an isomorphism of stalks. (Possible hint: Use the concrete description of the stalks. Another possibility once you read Remark 2.7.3: judicious use of adjoints.)

As a reality check, you may want to verify that “the sheafification of a constant presheaf is the corresponding constant sheaf” (see §2.2.10): if X is a topological space and S is a set, then $(\underline{S}_{\text{pre}})^{\text{sh}}$ may be naturally identified with \underline{S} .

2.4.7. ★ Remark. The “space of sections” (or “espace étalé”) construction (§2.2.11) yields a different-sounding description of sheafification which may be preferred by some readers. The main idea is identical: if \mathcal{F} is a presheaf, let F be the space of sections (or espace étalé) of \mathcal{F} . You may wish to show that if \mathcal{F} is a presheaf, the sheaf of sections of $F \rightarrow X$ (defined in Exercise 2.2.G(a)) is the sheafification of \mathcal{F} . Exercise 2.2.E may be interpreted as an example of this construction. The “space of sections” construction of the sheafification is essentially the same as Construction 2.4.6.

2.4.8. Subsheaves and quotient sheaves.

We now discuss subsheaves and quotient sheaves from the perspective of stalks.

2.4.M. EXERCISE. Suppose $\phi: \mathcal{F} \rightarrow \mathcal{G}$ is a morphism of sheaves of sets on a topological space X . Show that the following are equivalent.

- (a) ϕ is a monomorphism in the category of sheaves.
- (b) ϕ is injective on the level of stalks: $\phi_p: \mathcal{F}_p \rightarrow \mathcal{G}_p$ is injective for all $p \in X$.
- (c) ϕ is injective on the level of open sets: $\phi(U): \mathcal{F}(U) \rightarrow \mathcal{G}(U)$ is injective for all open $U \subset X$.

(Possible hints: for (b) implies (a), recall that morphisms are determined by stalks, Exercise 2.4.C. For (a) implies (c), use the “indicator sheaf” with one section over every open set contained in U , and no section over any other open set.) If these conditions hold, we say that \mathcal{F} is a **subsheaf** of \mathcal{G} (where the “inclusion” ϕ is sometimes left implicit).

(You may later wish to extend your solution to Exercise 2.4.M to show that for any morphism of *presheaves*, if all maps of sections are injective, then all stalk maps are injective. And furthermore, if $\phi: \mathcal{F} \rightarrow \mathcal{G}$ is a morphism from a separated presheaf to an arbitrary presheaf, then injectivity on the level of stalks implies that ϕ is a monomorphism in the category of presheaves. This is useful in some approaches to Exercise 2.6.C.)

2.4.N. EXERCISE. Continuing the notation of the previous exercise, show that the following are equivalent.

- (a) ϕ is an epimorphism in the category of sheaves.
- (b) ϕ is surjective on the level of stalks: $\phi_p: \mathcal{F}_p \rightarrow \mathcal{G}_p$ is surjective for all $p \in X$.

(Possible hint: use a skyscraper sheaf.)

If these conditions hold, we say that \mathcal{G} is a **quotient sheaf** of \mathcal{F} .

Thus *monomorphisms and epimorphisms — subsheafiness and quotient sheafiness — can be checked at the level of stalks.*

Both exercises generalize readily to sheaves with values in any reasonable category, where “injective” is replaced by “monomorphism” and “surjective” is replaced by “epimorphism”.

Notice that there was no part (c) to Exercise 2.4.N, and Example 2.4.9 shows why. (But there is a version of (c) that *implies* (a) and (b): surjectivity on all open sets in the base of a topology implies that the corresponding map of sheaves is an epimorphism, Exercise 2.5.D.)

2.4.9. Example (cf. Exercise 2.3.J). Let $X = \mathbb{C}$ with the classical topology, and define \mathcal{O}_X to be the sheaf of holomorphic functions, and \mathcal{O}_X^* to be the sheaf of invertible (nowhere zero) holomorphic functions. This is a sheaf of abelian groups under multiplication. We have maps of sheaves

$$(2.4.9.1) \quad 0 \longrightarrow \mathbb{Z} \xrightarrow{\times 2\pi i} \mathcal{O}_X \xrightarrow{\exp} \mathcal{O}_X^* \longrightarrow 1.$$

(You can figure out what the sheaves 0 and 1 mean; they are isomorphic, and are written in this way for reasons that may be clear.) We will soon interpret this as an exact sequence of sheaves of abelian groups (the *exponential exact sequence*, see Exercise 2.6.E), although we don’t yet have the language to do so.

2.4.O. ENLIGHTENING EXERCISE. Show that $\exp: \mathcal{O}_X \longrightarrow \mathcal{O}_X^*$ describes \mathcal{O}_X^* as a quotient sheaf of \mathcal{O}_X . Find an open set on which this map is not surjective.

This is a great example to get a sense of what “surjectivity” means for sheaves: nowhere vanishing holomorphic functions (such as the function x away from the origin) have logarithms locally, but they need not have logarithms globally.

2.5 Recovering sheaves from a “sheaf on a base”

Sheaves are natural things to want to think about, but hard to get our hands on. We like the identity and gluability axioms, but they make proving things trickier than for presheaves. We have discussed how we can understand sheaves using stalks (using “compatible germs”). We now introduce a second way of getting a hold of sheaves, by introducing the notion of a *sheaf on a base*. Warning: this way of understanding an entire sheaf from limited information is confusing. It may help to keep sight of the central insight that this partial information is enough to understand germs, and the notion of when they are compatible (with nearby germs).

First, we define the notion of a **base of a topology**. Suppose we have a topological space X , i.e., we know which subsets U_i of X are open. Then a base of a topology is a subcollection of the open sets $\{B_j\} \subset \{U_i\}$, such that each U_i is a union of the B_j . Here is one example that you have seen early in your mathematical life. Suppose $X = \mathbb{R}^n$. Then the way the classical topology is often first defined is by defining *open balls* $B_r(x) = \{y \in \mathbb{R}^n : |y - x| < r\}$, and declaring that any union of open balls is open. So the balls form a base of the classical topology — we say they *generate* the classical topology. As an application of how we use them, to

check continuity of some map $\pi: X \rightarrow \mathbb{R}^n$, you need only think about the pullback of balls on \mathbb{R}^n — part of the traditional δ - ϵ definition of continuity.

Now suppose we have a sheaf \mathcal{F} on a topological space X , and a base $\{B_i\}$ of open sets on X . Then consider the information

$$(\{\mathcal{F}(B_i)\}, \{\text{res}_{B_i, B_j} : \mathcal{F}(B_i) \rightarrow \mathcal{F}(B_j)\}),$$

which is a subset of the information contained in the sheaf — we are only paying attention to the information involving elements of the base, not all open sets.

We can recover the entire sheaf from this information. This is because we can determine the stalks from this information, and we can determine when germs are compatible.

2.5.A. IMPORTANT EXERCISE. Make this precise. How can you recover a sheaf \mathcal{F} from this partial information?

This suggests a notion, called a **sheaf on a base**. A sheaf of sets (or abelian groups, rings, ...) on a base $\{B_i\}$ is the following. For each B_i in the base, we have a set $F(B_i)$. If $B_i \subset B_j$, we have maps $\text{res}_{B_i, B_j} : F(B_j) \rightarrow F(B_i)$, with $\text{res}_{B_i, B_i} = \text{id}_{F(B_i)}$. (Things called “ B ” are always assumed to be in the base.) If $B_i \subset B_j \subset B_k$, then $\text{res}_{B_i, B_k} = \text{res}_{B_j, B_k} \circ \text{res}_{B_i, B_j}$. So far we have defined a **presheaf on a base** $\{B_i\}$.

We also require the **base identity** axiom: If $B = \cup B_i$, then if $f, g \in F(B)$ are such that $\text{res}_{B, B_i} f = \text{res}_{B, B_i} g$ for all i , then $f = g$.

We require the **base gluability** axiom too: If $B = \cup B_i$, and we have $f_i \in F(B_i)$ such that f_i agrees with f_j on any basic open set contained in $B_i \cap B_j$ (i.e., $\text{res}_{B_i, B_k} f_i = \text{res}_{B_j, B_k} f_j$ for all $B_k \subset B_i \cap B_j$) then there exists $f \in F(B)$ such that $\text{res}_{B, B_i} f = f_i$ for all i .

2.5.1. Theorem. — Suppose $\{B_i\}$ is a base on X , and F is a sheaf of sets on this base. Then there is a sheaf \mathcal{F} extending F (with isomorphisms $\mathcal{F}(B_i) \xrightarrow{\sim} F(B_i)$ agreeing with the restriction maps). This sheaf \mathcal{F} is unique up to unique isomorphism.

Proof. We will define \mathcal{F} as the sheaf of compatible germs of F .

Define the **stalk** of a presheaf F on a base at $p \in X$ by

$$F_p = \text{colim } F(B_i)$$

where the colimit is over all B_i (in the base) containing p .

We will say a family of germs in an open set U is compatible near p if there is a section s of F over some B_i containing p such that the germs over B_i are precisely the germs of s . More formally, define

$$\begin{aligned} \mathcal{F}(U) := \{ (f_p \in F_p)_{p \in U} : & \text{for all } p \in U, \text{ there exists } B \text{ with } p \in B \subset U, s \in F(B), \\ & \text{with } s_q = f_q \text{ for all } q \in B \} \end{aligned}$$

where each B is in our base.

This is a sheaf (for the same reasons that the sheaf of compatible germs was, cf. Exercise 2.4.H).

I next claim that if B is in our base, the natural map $F(B) \rightarrow \mathcal{F}(B)$ is an isomorphism.

2.5.B. EXERCISE. Verify that $F(B) \rightarrow \mathcal{F}(B)$ is an isomorphism, likely by showing that it is injective and surjective (or else by describing the inverse map and verifying that it is indeed inverse). Possible hint: elements of $\mathcal{F}(B)$ are determined by stalks, as are elements of $F(B)$.

It will be clear from your solution to Exercise 2.5.B that the restriction maps for F are the same as the restriction maps of \mathcal{F} (for elements of the base).

Finally, you should verify to your satisfaction that \mathcal{F} is indeed unique up to unique isomorphism. (First be sure that you understand what this means!) \square

Theorem 2.5.1 shows that sheaves on X can be recovered from their “restriction to a base”. It is clear from the argument (and in particular the solution to the Exercise 2.5.B) that if \mathcal{F} is a sheaf and F is the corresponding sheaf on the base B , then for any $p \in X$, \mathcal{F}_p is naturally isomorphic to F_p .

Theorem 2.5.1 is a statement about *objects* in a category, so we should hope for a similar statement about *morphisms*.

2.5.C. IMPORTANT EXERCISE: MORPHISMS OF SHEAVES CORRESPOND TO MORPHISMS OF SHEAVES ON A BASE. Suppose $\{B_i\}$ is a base for the topology of X . A morphism $F \rightarrow G$ of sheaves on the base is a collection of maps $F(B_k) \rightarrow G(B_k)$ such that the diagram

$$\begin{array}{ccc} F(B_i) & \longrightarrow & G(B_i) \\ \text{res}_{B_i, B_j} \downarrow & & \downarrow \text{res}_{B_i, B_j} \\ F(B_j) & \longrightarrow & G(B_j) \end{array}$$

commutes for all $B_j \hookrightarrow B_i$.

(a) Verify that a morphism of sheaves is determined by the induced morphism of sheaves on the base.

(b) Show that a morphism of sheaves on the base gives a morphism of the induced sheaves. (Possible hint: compatible stalks.)

2.5.2. Remark. The above constructions and arguments describe an equivalence of categories (§1.2.21) between sheaves on X and sheaves on a given base of X . There is no new content to this statement, but you may wish to think through what it means. What are the functors in each direction? Why aren’t their compositions the identity?

2.5.3. Remark. It will be useful to extend these notions to \mathcal{O}_X -modules (see for example Exercise 6.2.C). You will readily be able to verify that there is a correspondence (really, equivalence of categories) between \mathcal{O}_X -modules and “ \mathcal{O}_X -modules on a base”. Rather than working out the details, you should just informally think through the main points: what is an “ \mathcal{O}_X -module on a base”? Given an \mathcal{O}_X -module on a base, why is the corresponding sheaf naturally an \mathcal{O}_X -module? Later, if you are forced at gunpoint to fill in details, you will be able to.

2.5.D. UNIMPORTANT EXERCISE. Suppose a morphism of sheaves $F \rightarrow G$ on a base $\{B_i\}$ is surjective for all B_i (i.e., $F(B_i) \rightarrow G(B_i)$ is surjective for all i). Show that the corresponding morphism of sheaves (*not* on the base) is surjective (or more precisely: an epimorphism). The converse is not true, unlike the case for injectivity. This gives a useful sufficient criterion for “surjectivity”: a morphism of sheaves is

an epimorphism (“surjective”) if it is surjective for sections on a base. You may enjoy trying this out with Example 2.4.9 (dealing with holomorphic functions in the classical topology on $X = \mathbb{C}$), showing that the exponential map $\exp: \mathcal{O}_X \rightarrow \mathcal{O}_X^*$ is surjective, using the base of contractible open sets.

2.5.4. Gluing sheaves.

We will repeatedly see the theme of constructing some object by gluing, in many different contexts. Keep an eye out for it! In each case, we carefully consider what information we need in order to glue.

2.5.E. IMPORTANT EXERCISE. Suppose $X = \bigcup U_i$ is an open cover of X , and we have sheaves \mathcal{F}_i on U_i along with isomorphisms

$$\phi_{ij}: \mathcal{F}_i|_{U_i \cap U_j} \xrightarrow{\sim} \mathcal{F}_j|_{U_i \cap U_j}$$

(with ϕ_{ii} the identity) that agree on triple overlaps, i.e.,

$$(2.5.4.1) \quad \phi_{jk} \circ \phi_{ij} = \phi_{ik}$$

on $U_i \cap U_j \cap U_k$ (this is called the **cocycle condition**, for reasons we ignore). Show that these sheaves can be glued together into a sheaf \mathcal{F} on X (unique up to unique isomorphism), with isomorphisms $\mathcal{F}|_{U_i} \xrightarrow{\sim} \mathcal{F}_i$, and the isomorphisms over $U_i \cap U_j$ are the obvious ones. *Warning:* we are not assuming this is a finite cover, so you cannot use induction. For this reason this exercise can be perplexing. (You can use the ideas of this section to solve this problem, but you don’t necessarily need to. Hint: As the base, take those open sets contained in *some* U_i .)

Thus we can “glue sheaves together”, using limited patching information. Small observation: the hypothesis ϕ_{ii} is the identity is extraneous, as it follows from the cocycle condition.

2.5.5. Remark for experts. Exercise 2.5.E almost says that the “set” of sheaves forms a sheaf itself, but not quite. Making this precise leads one to the notion of a *stack*.

2.6 Sheaves of abelian groups, and \mathcal{O}_X -modules, form abelian categories

We are now ready to see that sheaves of abelian groups, and their cousins, \mathcal{O}_X -modules, form abelian categories. In other words, we may treat them similarly to vector spaces, and modules over a ring. In the process of doing this, we will see that this is much stronger than an analogy; kernels, cokernels, exactness, and so forth can be understood at the level of stalks (which are just abelian groups), and the compatibility of the germs will come for free.

The category of sheaves of abelian groups on a topological space X is clearly an additive category (Definition 1.6.1). In order to show that it is an abelian category, we must begin by showing that any morphism $\phi: \mathcal{F} \rightarrow \mathcal{G}$ has a kernel and a cokernel. We have already seen that ϕ has a kernel (Exercise 2.3.I): the presheaf kernel is a sheaf, and is a kernel in the category of sheaves.

2.6.A. EXERCISE. Show that the stalk of the kernel is the kernel of the stalks: for all $p \in X$, there is a natural isomorphism

$$(\ker(\mathcal{F} \rightarrow \mathcal{G}))_p \xrightarrow{\sim} \ker(\mathcal{F}_p \rightarrow \mathcal{G}_p).$$

We next address the issue of the cokernel. Now $\phi: \mathcal{F} \rightarrow \mathcal{G}$ has a cokernel in the category of presheaves; call it \mathcal{H}_{pre} (where the subscript is meant to remind us that this is a presheaf). Let $\text{sh}: \mathcal{H}_{\text{pre}} \rightarrow \mathcal{H}$ be its sheafification. Recall that the cokernel is defined using a universal property: it is the colimit of the diagram

$$(2.6.0.1) \quad \begin{array}{ccc} \mathcal{F} & \xrightarrow{\phi} & \mathcal{G} \\ \downarrow & & \\ 0 & & \end{array}$$

in the category of presheaves (cf. (1.6.3.1) and the comment thereafter).

2.6.1. Proposition. — *The composition $\mathcal{G} \rightarrow \mathcal{H}$ is the cokernel of ϕ in the category of sheaves.*

Proof. We show that it satisfies the universal property. Given any sheaf \mathcal{E} and a commutative diagram

$$\begin{array}{ccc} \mathcal{F} & \xrightarrow{\phi} & \mathcal{G} \\ \downarrow & & \downarrow \\ 0 & \longrightarrow & \mathcal{E} \end{array}$$

We construct

$$\begin{array}{ccccc} \mathcal{F} & \xrightarrow{\phi} & \mathcal{G} & & \\ \downarrow & & \downarrow & \searrow & \\ 0 & \longrightarrow & \mathcal{H}_{\text{pre}} & \xrightarrow{\text{sh}} & \mathcal{H} \\ & & & & \downarrow \\ & & & & \mathcal{E} \end{array}$$

(A curved arrow also goes from \mathcal{G} to \mathcal{E} .)

We show that there is a unique morphism $\mathcal{H} \rightarrow \mathcal{E}$ making the diagram commute. As \mathcal{H}_{pre} is the cokernel in the category of presheaves, there is a unique morphism of presheaves $\mathcal{H}_{\text{pre}} \rightarrow \mathcal{E}$ making the diagram commute. But then by the universal property of sheafification (Definition 2.4.5), there is a unique morphism of *sheaves* $\mathcal{H} \rightarrow \mathcal{E}$ making the diagram commute. \square

2.6.B. EXERCISE. Show that the stalk of the cokernel is naturally isomorphic to the cokernel of the stalk.

We have now defined the notions of kernel and cokernel, and verified that they may be checked at the level of stalks. We have also verified that the properties of a morphism being a monomorphism or epimorphism are also determined at the level of stalks (Exercises 2.4.M and 2.4.N). Hence we have proved the following:

2.6.2. Theorem. — *Sheaves of abelian groups on a topological space X form an abelian category.*

That's all there is to it — what needs to be proved has been shifted to the stalks, where everything works because stalks are abelian groups!

And we see more: all structures coming from the abelian nature of this category may be checked at the level of stalks. For example:

2.6.C. EXERCISE. Suppose $\phi: \mathcal{F} \rightarrow \mathcal{G}$ is a morphism of sheaves of abelian groups. Show that the image sheaf $\text{im } \phi$ is the sheafification of the image presheaf. (You must use the definition of image in an abelian category. In fact, this gives the accepted definition of image sheaf for a morphism of sheaves of sets.) Show that the stalk of the image is the image of the stalk.

As a consequence, **exactness of a sequence of sheaves may be checked at the level of stalks**. In particular:

2.6.D. IMPORTANT EXERCISE (CF. EXERCISE 2.3.A). Show that taking the stalk of a sheaf of abelian groups is an exact functor. More precisely, if X is a topological space and $p \in X$ is a point, show that taking the stalk at p defines an exact functor $Ab_X \rightarrow Ab$.

2.6.E. EXERCISE. Check that the exponential exact sequence (2.4.9.1) is exact.

2.6.F. EXERCISE: LEFT-EXACTNESS OF THE FUNCTOR OF “SECTIONS OVER U ”. Suppose $U \subset X$ is an open set, and $0 \rightarrow \mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{H}$ is an exact sequence of sheaves of abelian groups. Show that

$$0 \rightarrow \mathcal{F}(U) \rightarrow \mathcal{G}(U) \rightarrow \mathcal{H}(U)$$

is exact. (You should do this “by hand”, even if you realize there is a very fast proof using the left-exactness of the “forgetful” right adjoint to the sheafification functor.) Show that the section functor need not be exact: show that if $0 \rightarrow \mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{H} \rightarrow 0$ is an exact sequence of sheaves of abelian groups, then

$$0 \rightarrow \mathcal{F}(U) \rightarrow \mathcal{G}(U) \rightarrow \mathcal{H}(U) \rightarrow 0$$

need not be exact. (Hint: the exponential exact sequence (2.4.9.1). But feel free to make up a different example.)

2.6.G. EXERCISE: LEFT-EXACTNESS OF PUSHFORWARD. Suppose $0 \rightarrow \mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{H}$ is an exact sequence of sheaves of abelian groups on X . If $\pi: X \rightarrow Y$ is a continuous map, show that

$$0 \rightarrow \pi_* \mathcal{F} \rightarrow \pi_* \mathcal{G} \rightarrow \pi_* \mathcal{H}$$

is exact. (The previous exercise, dealing with the left-exactness of the global section functor can be interpreted as a special case of this, in the case where Y is a point.)

2.6.H. EXERCISE: LEFT-EXACTNESS OF Hom (CF. EXERCISE 1.6.F(C) AND (D)). Suppose \mathcal{F} is a sheaf of abelian groups on a topological space X . Show that $\text{Hom}(\mathcal{F}, \cdot)$ is a left-exact covariant functor $Ab_X \rightarrow Ab_X$. Show that $\text{Hom}(\cdot, \mathcal{F})$ is a left-exact contravariant functor $Ab_X \rightarrow Ab_X$.

2.6.3. \mathcal{O}_X -modules.

2.6.I. EXERCISE. Show that if (X, \mathcal{O}_X) is a ringed space, then \mathcal{O}_X -modules form an abelian category. (There is a fair bit to check, but there aren't many new ideas.)

2.6.4. Many facts about sheaves of abelian groups carry over to \mathcal{O}_X -modules without change, because a sequence of \mathcal{O}_X -modules is exact if and only if the underlying sequence of sheaves of abelian groups is exact. You should be able to easily check that all of the statements of the earlier exercises in §2.6 also hold for \mathcal{O}_X -modules, when modified appropriately. For example (Exercise 2.6.H), $\text{Hom}_{\mathcal{O}_X}(\cdot, \cdot)$ is a left-exact contravariant functor in its first argument and a left-exact covariant functor in its second argument.

We end with a useful construction using some of the ideas in this section.

2.6.J. IMPORTANT EXERCISE: TENSOR PRODUCTS OF \mathcal{O}_X -MODULES.

(a) Suppose \mathcal{O}_X is a sheaf of rings on X . Define (categorically) what we should mean by **tensor product of two \mathcal{O}_X -modules**. Give an explicit construction, and show that it satisfies your categorical definition. *Hint:* take the “presheaf tensor product” — which needs to be defined — and sheafify. Note: $\otimes_{\mathcal{O}_X}$ is often written \otimes when the subscript is clear from the context. (An example showing sheafification is necessary will arise in Example 15.1.1.)

(b) Show that the tensor product of stalks is the stalk of the tensor product. (If you can show this, you may be able to make sense of the phrase “colimits commute with tensor products”.)

2.6.5. Conclusion. Just as presheaves of abelian groups on a topological space form an abelian category because all abelian-categorical notions make sense open set by open set, sheaves of abelian groups on a topological space form an abelian category because all abelian-categorical notions make sense stalk by stalk.

2.7 The inverse image sheaf

We next describe a notion that is fundamental, but rather intricate. We will not need it for some time, so this may be best left for a second reading. Suppose we have a continuous map $\pi: X \rightarrow Y$. If \mathcal{F} is a sheaf on X , we have defined the pushforward or direct image sheaf $\pi_*\mathcal{F}$, which is a sheaf on Y . There is also a notion of inverse image sheaf. (We will not call it the pullback sheaf, reserving that name for a later construction for quasicoherent sheaves, §17.3.) This is a covariant functor π^{-1} from sheaves on Y to sheaves on X . If the sheaves on Y have some additional structure (e.g., group or ring), then this structure is respected by π^{-1} .

2.7.1. Definition by adjoint: elegant but abstract. We define the inverse image π^{-1} as the left adjoint to π_* .

This isn’t really a definition; we need a construction to show that the adjoint exists. Note that we then get canonical maps $\pi^{-1}\pi_*\mathcal{F} \rightarrow \mathcal{F}$ (associated to the identity in $\text{Mor}_Y(\pi_*\mathcal{F}, \pi_*\mathcal{F})$) and $\mathcal{G} \rightarrow \pi_*\pi^{-1}\mathcal{G}$ (associated to the identity in

$$\text{Mor}_X(\pi^{-1}\mathcal{G}, \pi^{-1}\mathcal{F}).$$

$$\begin{array}{ccc} & \pi^{-1}\mathcal{G} & \longrightarrow \mathcal{F} \\ & \nearrow & \nearrow \\ X & & \\ \downarrow \pi & \mathcal{G} & \longrightarrow \pi_*\mathcal{F} \\ Y & & \end{array}$$

2.7.2. Construction: concrete but ugly. Define the temporary notation

$$\pi_{\text{pre}}^{-1}\mathcal{G}(\mathcal{U}) = \text{colim}_{V \supset \pi(\mathcal{U})} \mathcal{G}(V).$$

(Recall the explicit description of colimit: sections of π_{pre}^{-1} over \mathcal{U} are sections on open sets containing $\pi(\mathcal{U})$, with an equivalence relation. Note that $\pi(\mathcal{U})$ won't be an open set in general.)

2.7.A. EXERCISE. Show that this defines a presheaf on X . Show that it needn't form a sheaf. (Hint: map 2 points to 1 point.)

Now define the **inverse image** of \mathcal{G} by $\pi^{-1}\mathcal{G} := (\pi_{\text{pre}}^{-1}\mathcal{G})^{\text{sh}}$. Note that π^{-1} is a functor from sheaves on Y to sheaves on X . The next exercise shows that π^{-1} is indeed left-adjoint to π_* . But you may wish to try the later exercises first, and come back to Exercise 2.7.B at another time. (For the later exercises, try to give two proofs, one using the universal property, and the other using the explicit description.)

2.7.B. IMPORTANT TRICKY EXERCISE ((π^{-1}, π_*) ARE ADJOINT). If $\pi: X \rightarrow Y$ is a continuous map, and \mathcal{F} is a sheaf on X and \mathcal{G} is a sheaf on Y , describe a bijection

$$\text{Mor}_X(\pi^{-1}\mathcal{G}, \mathcal{F}) \leftrightarrow \text{Mor}_Y(\mathcal{G}, \pi_*\mathcal{F}).$$

Observe that your bijection is “natural” in the sense of the definition of adjoints (i.e., functorial in both \mathcal{F} and \mathcal{G}). Thus Construction 2.7.2 satisfies the universal property of Definition 2.7.1. Possible hint: Show that both sides agree with the following third construction, which we denote $\text{Mor}_{YX}(\mathcal{G}, \mathcal{F})$. A collection of maps $\phi_{V\mathcal{U}}: \mathcal{G}(V) \rightarrow \mathcal{F}(\mathcal{U})$ (as \mathcal{U} runs through all open sets of X , and V runs through all open sets of Y containing $\pi(\mathcal{U})$) is said to be *compatible* if for all open $\mathcal{U}' \subset \mathcal{U} \subset X$ and all open $V' \subset V \subset Y$ with $\pi(\mathcal{U}) \subset V, \pi(\mathcal{U}') \subset V'$, the diagram

$$(2.7.2.1) \quad \begin{array}{ccc} \mathcal{G}(V) & \xrightarrow{\phi_{V\mathcal{U}}} & \mathcal{F}(\mathcal{U}) \\ \text{res}_{V,V'} \downarrow & & \downarrow \text{res}_{\mathcal{U},\mathcal{U}'} \\ \mathcal{G}(V') & \xrightarrow{\phi_{V'\mathcal{U}'}} & \mathcal{F}(\mathcal{U}') \end{array}$$

commutes. Define $\text{Mor}_{YX}(\mathcal{G}, \mathcal{F})$ to be the set of all compatible collections $\phi = \{\phi_{V\mathcal{U}}\}$.

2.7.3. Remark (“stalk and skyscraper are an adjoint pair”). As a special case, if X is a point $p \in Y$, we see that $\pi^{-1}\mathcal{G}$ is the stalk \mathcal{G}_p of \mathcal{G} , and maps from the stalk \mathcal{G}_p to a set S are the same as maps of sheaves on Y from \mathcal{G} to the skyscraper sheaf with

set S supported at p . You may prefer to prove this special case by hand directly before solving Exercise 2.7.B, as it is enlightening. (It can also be useful — can you use it to solve Exercises 2.4.L and 2.4.N?)

2.7.C. EXERCISE. Show that the stalks of $\pi^{-1}\mathcal{G}$ are the same as the stalks of \mathcal{G} . More precisely, if $\pi(p) = q$, describe a natural isomorphism $\mathcal{G}_q \xrightarrow{\sim} (\pi^{-1}\mathcal{G})_p$. (Possible hint: use the concrete description of the stalk, as a colimit. Recall that stalks are preserved by sheafification, Exercise 2.4.L. Alternatively, use adjointness.)

Exercise 2.7.C, along with the notion of compatible germs, may give you a simple way of thinking about (and perhaps visualizing) inverse image sheaves. Closely related: you can think of sections of the inverse image sheaf as, locally, pullbacks of sections on the target. (Those preferring the “espace étalé” or “space of sections” perspective, §2.2.11, can check that the pullback of the “space of sections” is the “space of sections” of the pullback.)

2.7.D. EXERCISE (EASY BUT USEFUL). If U is an open subset of Y , $i: U \rightarrow Y$ is the inclusion, and \mathcal{G} is a sheaf on Y , show that $i^{-1}\mathcal{G}$ is naturally isomorphic to $\mathcal{G}|_U$ (the restriction of \mathcal{G} to U , §2.2.8).

2.7.E. EXERCISE. Show that π^{-1} is an exact functor from sheaves of abelian groups on Y to sheaves of abelian groups on X (cf. Exercise 2.6.D). (Hint: exactness can be checked on stalks, and by Exercise 2.7.C, the stalks are the same.) Essentially the same argument will show that π^{-1} is an exact functor from \mathcal{O}_Y -modules (on Y) to $(\pi^{-1}\mathcal{O}_Y)$ -modules (on X), but don't bother writing that down. (Remark for experts: π^{-1} is a left adjoint, hence right-exact by abstract nonsense, as discussed in §1.6.12. Left-exactness holds because colimits of abelian groups over filtered index sets are exact, Exercise 1.6.K.)

2.7.4. Definition: The push-pull map.

Suppose

$$(2.7.4.1) \quad \begin{array}{ccc} W & \xrightarrow{\beta'} & X \\ \alpha' \downarrow & & \downarrow \alpha \\ Y & \xrightarrow{\beta} & Z \end{array}$$

is a commutative (not necessarily Cartesian!) diagram, and \mathcal{F} is a sheaf on X . Define the **push-pull map**

$$(2.7.4.2) \quad \beta^* \alpha_* \mathcal{F} \longrightarrow \alpha'_*(\beta')^* \mathcal{F}$$

of sheaves on Y as follows. Start with the identity $(\beta')^* \mathcal{F} \xrightarrow{\sim} (\beta')^* \mathcal{F}$ on W . By adjointness of $((\beta')^*, \beta'_*)$, this is the same as the data of a morphism $\mathcal{F} \rightarrow (\beta'_*)(\beta')^* \mathcal{F}$ on X . Apply α_* to get a map $\alpha_* \mathcal{F} \rightarrow \alpha_*(\beta'_*)(\beta')^* \mathcal{F}$ on Z . By the commutativity of (2.7.4.1), this is the map $\alpha_* \mathcal{F} \rightarrow \beta_*(\alpha'_*)(\beta')^* \mathcal{F}$ on Z . By adjointness of (β^*, β_*) , this yields a map (2.7.4.2).

We observe that this entire construction is functorial in \mathcal{F} (i.e., given a map $\mathcal{F} \rightarrow \mathcal{G}$ of sheaves on X , we get a certain commutative diagram of sheaves on Y — what is it?). (We will later extend this to \mathcal{O} -modules, quasicoherent sheaves, and cohomology, see Exercises 17.3.G and 19.8.B.)

2.7.F. EXERCISE. We could have defined the push-pull map in a “dual way” starting with the identity $\alpha_* \mathcal{F} \rightarrow \alpha_* \mathcal{F}$ on Z , then using adjointness of (α^*, α_*) , and continuing from there. Why does this give the *same* definition of the push-pull map?

2.7.5. The support of a sheaf, and the support of a section of a sheaf.

Exercise 2.7.H below gives us an excuse to introduce the notion of *support*, which we use repeatedly later.

2.7.6. Definition. Suppose \mathcal{F} is a sheaf (or indeed separated presheaf) of abelian groups on X , and s is a global section of \mathcal{F} . Define the **support of the section s** , denoted $\text{Supp } s$, to be the set of points p of X where s has a nonzero germ:

$$\text{Supp } s := \{p \in X : s_p \neq 0 \text{ in } \mathcal{F}_p\}.$$

We think of this as the subset of X where “the section s lives” — the complement is the locus where s is the 0-section. (Unimportant: We could define this even if \mathcal{F} is a presheaf, but without the inclusion $\mathcal{F}(U) \hookrightarrow \prod_{p \in U} \mathcal{F}_p$ of Exercise 2.4.A, we could have the strange situation where we have a nonzero section that “lives nowhere”, because it is 0 “near every point”, i.e., is 0 in every stalk.) Define the **support of a sheaf \mathcal{G}** of sets, denoted $\text{Supp } \mathcal{G}$, as the locus where the stalks are nonzero:

$$\text{Supp } \mathcal{G} := \{p \in X : |\mathcal{G}_p| \neq 0\}.$$

Equivalently, $\text{Supp } \mathcal{G}$ is the union of supports of sections over all open sets. Clearly support is a “stalk-local notion”, and hence “commutes” with restriction to open sets.

More generally, if the sheaf has value in some category, the support consists of points where the stalk is not the initial object.

2.7.G. EXERCISE (THE SUPPORT OF A SECTION IS CLOSED). Show that $\text{Supp } s$ is a closed subset of X .

2.7.7. Caution: the locus where a continuous function is nonzero is open; the locus where the germ of a function is nonzero is closed. Basically by the definition of continuity, the locus where the *value* of a continuous function is nonzero is *open*. (More generally, the locus where the value of a function on a locally ringed space is nonzero is open, see Exercise 4.3.F(a).) In contrast, Exercise 2.7.G shows that the locus where the *germ* of a function is nonzero is *closed*. We will try to avoid misunderstanding by using phrases like “ f is 0 at p ” (the value of f is zero, i.e., $f(p) = 0$) and “ f is 0 near p ” (the germ of f is zero, i.e., $f = 0$ in $\mathcal{O}_{X,p}$, or equivalently, f is zero in some neighborhood of p).

2.7.H. EXERCISE.

- (a) Suppose $Z \subset Y$ is a closed subset, and $i: Z \hookrightarrow Y$ is the inclusion. If \mathcal{F} is a sheaf of sets on Z , then show that the stalk $(i_* \mathcal{F})_q$ is \emptyset if $q \notin Z$, and \mathcal{F}_q if $q \in Z$.
- (b) Suppose $\text{Supp } \mathcal{G} \subset Z$ where Z is closed. Show that the natural map $\mathcal{G} \rightarrow i_* i^{-1} \mathcal{G}$ is an isomorphism. Thus a sheaf supported on a closed subset can be considered a sheaf on that closed subset.

2.7.8. Extension by zero, an occasional left adjoint to the inverse image functor. In addition to always being a left adjoint, π^{-1} can sometimes be a right adjoint,

when π is an inclusion of an open subset. We discuss this when we need it, in §24.4.7.

Part II

Schemes

L'idée même de schéma est d'une simplicité enfantine — si simple, si humble, que personne avant moi n'avait songé à se pencher si bas. Si "bébête" même, pour tout dire, que pendant des années encore et en dépit de l'évidence, pour beaucoup de mes savants collègues, ça faisait vraiment "pas sérieux"!

The very idea of scheme is of infantile simplicity — so simple, so humble, that no one before me thought of stooping so low. So childish, in short, that for years, despite all the evidence, for many of my erudite colleagues, it was really "not serious"!

— A. Grothendieck [**Gr6**, p. P32], translated by C. McLarty [**Mc**, p. 313]

CHAPTER 3

Toward affine schemes: the underlying set, and topological space

*There is no serious historical question of how Grothendieck found his definition of schemes. It was in the air. Serre has well said that no one invented schemes... The question is, what made Grothendieck believe he should use this definition to simplify an 80 page paper by Serre into some 1000 pages of *Éléments de Géométrie Algébrique*?*

— C. McLarty [Mc, p. 313]

3.1 Toward schemes

We are now ready to consider the notion of a *scheme*, which is the type of geometric space central to algebraic geometry. We should first think through what we mean by “geometric space”. You have likely seen the notion of a manifold, and we wish to abstract this notion so that it can be generalized to other settings, notably so that we can deal with nonsmooth and arithmetic objects.

The key insight behind this generalization is the following: we can understand a geometric space (such as a manifold) well by understanding the functions on this space. More precisely, we will understand it through the sheaf of functions on the space. If we are interested in differentiable manifolds, we will consider smooth functions; if we are interested in analytic manifolds, we will consider real analytic functions; and so on.

Thus we will define a scheme to be the following data

- *The set*: the points of the scheme
- *The topology*: the open sets of the scheme
- *The structure sheaf*: the sheaf of “algebraic functions” (a sheaf of rings) on the scheme.

Recall that a topological space with a sheaf of rings is called a *ringed space* (§2.2.13).

We will try to draw pictures throughout. Pictures can help develop geometric intuition, which can guide the algebraic development (and, eventually, vice versa). Some people find pictures very helpful, while others are repulsed or confused.

We will try to make all three notions as intuitive as possible. For the set, in the key example of complex (affine) varieties (roughly, things cut out in \mathbb{C}^n by polynomials), we will see that the points are the “traditional points” (n -tuples of complex numbers), plus some extra points that will be handy to have around. For the topology, we will require that “the subset where an algebraic function vanishes must be closed”, and require nothing else. For the sheaf of algebraic functions (the structure sheaf), we will expect that in the complex plane, $(3x^2 + y^2)/(2x + 4xy + 1)$

should be an algebraic function on the open set consisting of points where the denominator doesn't vanish, and this will largely motivate our definition.

3.1.1. Example: Differentiable manifolds. As motivation, we return to our example of differentiable manifolds, reinterpreting them in this light. We will be quite informal in this discussion. Suppose X is a differentiable manifold. It is a topological space, and has a *sheaf of smooth* (C^∞) *functions* \mathcal{O}_X (see §2.1). This gives X the structure of a ringed space. We have observed that evaluation at a point $p \in X$ gives a surjective map from the stalk to \mathbb{R}

$$\mathcal{O}_{X,p} \twoheadrightarrow \mathbb{R},$$

so the kernel, the (germs of) functions vanishing at p , is a maximal ideal $\mathfrak{m}_{X,p}$ (see §2.1.1).

We could *define* a differentiable real manifold as a topological space X with a sheaf of rings (see Definition 4.3.9). We would require that there is a cover of X by open sets such that on each open set the ringed space is isomorphic to a ball around the origin in \mathbb{R}^n (with the sheaf of smooth functions on that ball). With this definition, the ball is the basic patch, and a general manifold is obtained by gluing these patches together. (Admittedly, a great deal of geometry comes from how one chooses to patch the balls together!) In the algebraic setting, the basic patch is the notion of an *affine scheme*, which we will discuss soon. (In the definition of manifold, there is an additional requirement that the topological space be Hausdorff and second-countable, to avoid pathologies. Schemes are often required to be “separated” to avoid essentially the same pathologies. Separatedness will be discussed in Chapter 11.)

Functions are determined by their values at points. This is an obvious statement, but won't be true for schemes in general. We will see an example in Exercise 3.2.A(a), and discuss this behavior further in §3.2.12.

Morphisms of manifolds. How can we describe maps of differentiable manifolds $\pi: X \rightarrow Y$? They are certainly continuous maps — but which ones? We can pull back functions along continuous maps. Smooth functions pull back to smooth functions. More formally, we have a map $\pi^{-1}\mathcal{O}_Y \rightarrow \mathcal{O}_X$. (The inverse image sheaf π^{-1} was defined in §2.7.) Inverse image is left-adjoint to pushforward, so we also get a map $\pi^\sharp: \mathcal{O}_Y \rightarrow \pi_*\mathcal{O}_X$.

Certainly given a map of differentiable manifolds, smooth functions pull back to smooth functions. It is less obvious that *this is a sufficient condition for a continuous map to be smooth*.

3.1.A. IMPORTANT EXERCISE FOR THOSE WITH A LITTLE EXPERIENCE WITH MANIFOLDS. Suppose that $\pi: X \rightarrow Y$ is a continuous map of differentiable manifolds (as topological spaces — not a priori smooth). Show that π is smooth if smooth functions pull back to smooth functions, i.e., if pullback by π gives a map $\mathcal{O}_Y \rightarrow \pi_*\mathcal{O}_X$. (Hint: check this on small patches. Once you figure out what you are trying to show, you will realize that the result is immediate.)

3.1.B. EXERCISE. Show that a morphism of differentiable manifolds $\pi: X \rightarrow Y$ with $\pi(p) = q$ induces a morphism of stalks $\pi^\sharp: \mathcal{O}_{Y,q} \rightarrow \mathcal{O}_{X,p}$. Show that $\pi^\sharp(\mathfrak{m}_{Y,q}) \subset \mathfrak{m}_{X,p}$. In other words, if you pull back a function that vanishes at

q , you get a function that vanishes at p — not a huge surprise. (In §7.3, we formalize this by saying that maps of manifolds are maps of locally ringed spaces.)

3.1.2. *Aside.* Here is a little more for experts: Notice that π induces a map on tangent spaces (see Aside 2.1.2)

$$(\mathfrak{m}_{X,p}/\mathfrak{m}_{X,p}^2)^\vee \rightarrow (\mathfrak{m}_{Y,q}/\mathfrak{m}_{Y,q}^2)^\vee.$$

This is the tangent map you would geometrically expect. Again, it is interesting that the cotangent map $\mathfrak{m}_{Y,q}/\mathfrak{m}_{Y,q}^2 \rightarrow \mathfrak{m}_{X,p}/\mathfrak{m}_{X,p}^2$ is algebraically more natural than the tangent map (there are no “duals”).

Experts are now free to try to interpret other differential-geometric information using only the map of topological spaces and map of sheaves. For example: how can one check if π is a submersion of manifolds? How can one check if f is an immersion? (We will see that the algebro-geometric version of these notions are *smooth morphism* and *unramified morphism*; see Chapter 26, although they will be defined earlier, in Definition 13.6.2 and §22.6, respectively.)

3.1.3. *Side Remark.* Manifolds are covered by disks that are all isomorphic. This isn’t true for schemes (even for “smooth complex varieties”). There are examples of two “smooth complex curves” (the algebraic version of Riemann surfaces) X and Y so that no nonempty open subset of X is isomorphic to a nonempty open subset of Y (see Exercise 7.5.K). And there is a Riemann surface X such that no two open subsets of X are isomorphic (see Exercise 20.7.F). Informally, this is because in the Zariski topology on schemes, all nonempty open sets are “huge” and have more “structure”.

3.1.4. Other examples. If you are interested in differential geometry, you might be interested in differentiable manifolds, on which the functions under consideration are smooth functions. Similarly, if you are interested in topology, you will be interested in topological spaces, on which you will consider the continuous functions. If you are interested in complex geometry, you will be interested in complex manifolds (or possibly “complex analytic varieties”), on which the functions are holomorphic functions. In each of these cases of interesting “geometric spaces”, the topological space and sheaf of functions is clear. The notion of scheme fits naturally into this family.

3.2 The underlying set of an affine scheme

For any ring A , we are going to define something called $\text{Spec } A$, the **spectrum** of A . In this section, we will define it as a set, but we will soon endow it with a topology, and later we will define a sheaf of rings on it (the structure sheaf). Such an object is called an *affine scheme*. Later $\text{Spec } A$ will denote the set along with the topology, and a sheaf of functions. But for now, as there is no possibility of confusion, $\text{Spec } A$ will just be the set.

3.2.1. The set $\text{Spec } A$ is the set of prime ideals of A . The prime ideal \mathfrak{p} of A when considered as an element of $\text{Spec } A$ will be denoted $[\mathfrak{p}]$, to avoid confusion. Elements $a \in A$ will be called **functions on $\text{Spec } A$** , and their **value** at the point $[\mathfrak{p}]$ will be $a \pmod{\mathfrak{p}}$. *This is weird: a function can take values in different rings at different points — the function 5 on $\text{Spec } \mathbb{Z}$ takes the value 1 (mod 2) at $[(2)]$ and 2 (mod 3) at $[(3)]$.*

“An element a of the ring lying in a prime ideal \mathfrak{p} ” translates to “a function a that is 0 at the point $[\mathfrak{p}]$ ” or “a function a vanishing at the point $[\mathfrak{p}]$ ”, and we will use these phrases interchangeably. Notice that if you add or multiply two functions, you add or multiply their values at all points; this is a translation of the fact that $A \rightarrow A/\mathfrak{p}$ is a ring morphism. These translations are important — make sure you are very comfortable with them! They should become second nature.

If A is generated over a base field (or base ring) by elements x_1, \dots, x_r , the elements x_1, \dots, x_r are often called **coordinates**, because we will later be able to reinterpret them as restrictions of “coordinates on r -space”, via the idea of §3.2.10, made precise in Exercise 7.2.E.

3.2.2. *Beginning of a grand dictionary between algebra and geometry.* We are building up the beginning of a grand dictionary; see the first part of §3.7.2.

3.2.3. *Glimpses of the future.* In §4.1.3: we will interpret functions on $\text{Spec } A$ as global sections of the “structure sheaf”, i.e., as a function on a ringed space, in the sense of §2.2.13. We repeat a caution from §2.2.13: what we will call “functions”, others may call “regular functions”. And we will later define “rational functions” (§6.5.34), which are not precisely functions in this sense; they are a particular type of “partially-defined function”.

The notion of “value of a function” will be later interpreted as a value of a function on a particular locally ringed space, see Definition 4.3.7.

3.2.4. We now give some examples.

Example 1 (the complex affine line): $\mathbb{A}_{\mathbb{C}}^1 := \text{Spec } \mathbb{C}[x]$. Let’s find the prime ideals of $\mathbb{C}[x]$. As $\mathbb{C}[x]$ is an integral domain, 0 is prime. Also, $(x - a)$ is prime, for any $a \in \mathbb{C}$: it is even a maximal ideal, as the quotient by this ideal is a field:

$$0 \longrightarrow (x - a) \longrightarrow \mathbb{C}[x] \xrightarrow{f \mapsto f(a)} \mathbb{C} \longrightarrow 0$$

(This exact sequence may remind you of (2.1.1.1) in our motivating example of manifolds.)

We now show that there are no other prime ideals. We use the fact that $\mathbb{C}[x]$ has a division algorithm, and is a unique factorization domain. Suppose \mathfrak{p} is a prime ideal. If $\mathfrak{p} \neq (0)$, then suppose $f(x) \in \mathfrak{p}$ is a nonzero element of smallest degree. It is not constant, as prime ideals can’t contain 1. If $f(x)$ is not linear, then factor $f(x) = g(x)h(x)$, where $g(x)$ and $h(x)$ have positive degree. (Here we use that \mathbb{C} is algebraically closed.) Then $g(x) \in \mathfrak{p}$ or $h(x) \in \mathfrak{p}$, contradicting the minimality of the degree of f . Hence there is a linear element $x - a$ of \mathfrak{p} . Then I claim that $\mathfrak{p} = (x - a)$. Suppose $f(x) \in \mathfrak{p}$. Then the division algorithm would give $f(x) = g(x)(x - a) + m$ where $m \in \mathbb{C}$. Then $m = f(x) - g(x)(x - a) \in \mathfrak{p}$. If $m \neq 0$, then $1 \in \mathfrak{p}$, giving a contradiction.

Thus we can and should (and must!) make a picture of $\mathbb{A}_{\mathbb{C}}^1 = \text{Spec } \mathbb{C}[x]$ (see Figure 3.1). This is just the first illustration of a point of view of Sophie Germain [Ge]: “L’algèbre n’est qu’une géométrie écrite; la géométrie n’est qu’une algèbre figurée.” (Algebra is but written geometry; geometry is but drawn algebra.)

There is one “traditional” point for each complex number, plus one extra (“bonus”) point $[(0)]$. We can mostly picture $\mathbb{A}_{\mathbb{C}}^1$ as \mathbb{C} : the point $[(x - a)]$ we will reasonably associate to $a \in \mathbb{C}$. Where should we picture the point $[(0)]$? The best way of thinking about it is somewhat zen. It is somewhere on the complex line, but nowhere in particular. Because (0) is contained in all of these prime ideals, we will somehow associate it with this line passing through all the other points. This new point $[(0)]$ is called the “generic point” of the line. (We will formally define “generic point” in §3.6.) It is “generically on the line” but you can’t pin it down any further than that. It is not at any particular place on the line. (This is misleading too — we will see in Easy Exercise 3.6.N that it is “near” every point. So it is near everything, but located nowhere precisely.) We will place it far to the right for lack of anywhere better to put it. You will notice that we sketch $\mathbb{A}_{\mathbb{C}}^1$ as one-(real)-dimensional (even though we picture it as an enhanced version of \mathbb{C}); this is to later remind ourselves that this will be a one-dimensional space, where dimensions are defined in an algebraic (or complex-geometric) sense. (Dimension will be defined in Chapter 12.)

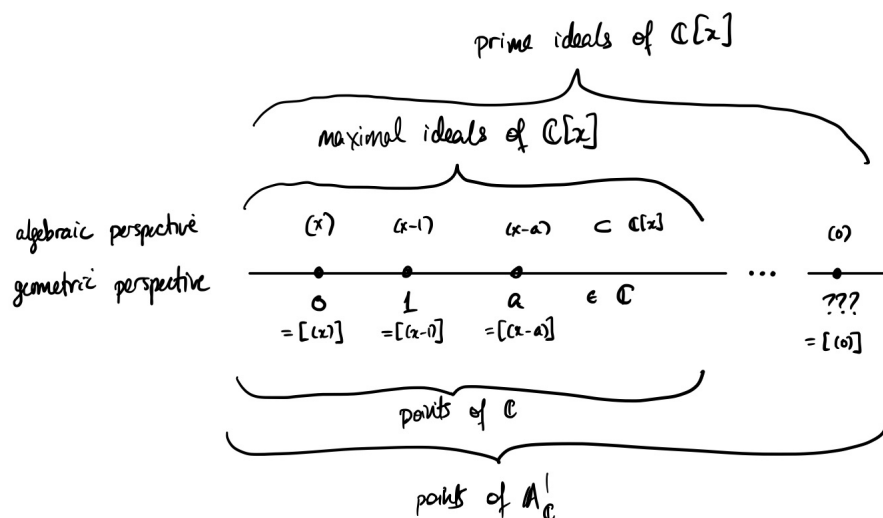


FIGURE 3.1. A picture of $\mathbb{A}_{\mathbb{C}}^1 = \text{Spec } \mathbb{C}[x]$

To give you some feeling for this space, we make some statements that are currently undefined, but suggestive. The functions on $\mathbb{A}_{\mathbb{C}}^1$ are the polynomials. So $f(x) = x^2 - 3x + 1$ is a function. What is its value at $[(x - 1)]$, which we think of as the point $1 \in \mathbb{C}$? Answer: $f(1)$! Or equivalently, we can evaluate $f(x)$ modulo $x - 1$ — this is the same thing by the division algorithm. (What is its value at $[(0)]$? It is $f(x) \pmod{0}$, which is just $f(x)$.)

Here is a more complicated example: $g(x) = (x - 3)^3 / (x - 2)$ is a “rational function”. It is defined everywhere but $x = 2$. (When we know what the structure

sheaf is, we will be able to say that it is an element of the structure sheaf on the open set $\mathbb{A}_{\mathbb{C}}^1 \setminus \{2\}$.) We want to say that $g(x)$ has a triple zero at 3, and a single pole at 2, and we will be able to after §13.5.

Example 2 (the affine line over $k = \bar{k}$): $\mathbb{A}_k^1 := \text{Spec } k[x]$ where k is an algebraically closed field. This is called the affine line over k . All of our discussion in the previous example carries over without change. We will use the same picture, which is after all intended to just be a metaphor.

Example 3: $\text{Spec } \mathbb{Z}$. An amazing fact is that from our perspective, this will look a lot like the affine line \mathbb{A}_k^1 . The integers, like $\bar{k}[x]$, form a unique factorization domain, with a division algorithm. The prime ideals are: (0) , and (p) where p is prime. Thus everything from Example 1 carries over without change, even the picture. Our picture of $\text{Spec } \mathbb{Z}$ is shown in Figure 3.2.




FIGURE 3.2. A “picture” of $\text{Spec } \mathbb{Z}$, which looks suspiciously like Figure 3.1

Let’s blithely carry over our discussion of functions to this space. 100 is a function on $\text{Spec } \mathbb{Z}$. Its value at (3) is “1 (mod 3)”. Its value at (2) is “0 (mod 2)”, and in fact it has a double zero. $27/4$ is a “rational function” on $\text{Spec } \mathbb{Z}$, defined away from (2) . We want to say that it has a double pole at (2) , and a triple zero at (3) . Its value at (5) is

$$27 \times 4^{-1} \equiv 2 \times (-1) \equiv 3 \pmod{5}.$$

(We will gradually make this discussion precise over time.)

Example 4: silly but important examples, and the German word for bacon. The set $\text{Spec } k$ where k is any field is boring: one point. $\text{Spec } 0$, where 0 is the zero-ring, is the empty set, as 0 has no prime ideals.

3.2.A. A SMALL EXERCISE ABOUT SMALL SCHEMES.

(a) Describe the set $\text{Spec } k[\epsilon]/(\epsilon^2)$. The ring $k[\epsilon]/(\epsilon^2)$ is called the ring of **dual numbers**, and will turn out to be quite useful. You should think of ϵ as a very small number, so small that its square is 0 (although it itself is not 0). It is a nonzero function whose value at all points is zero, thus giving our first example of functions not being determined by their values at points. We will discuss this phenomenon further in §3.2.12.

(b) Describe the set $\text{Spec } k[x]_{(x)}$ (see §1.3.3 for a discussion of localization). We will see this scheme again repeatedly, starting with §3.2.9 and Exercise 3.4.K. You might later think of it as a shred of a particularly nice “smooth curve”.

In Example 2, we restricted to the case of algebraically closed fields for a reason: things are more subtle if the field is not algebraically closed.

Example 5 (the affine line over \mathbb{R}): $\mathbb{A}_{\mathbb{R}}^1 = \text{Spec } \mathbb{R}[x]$. Using the fact that $\mathbb{R}[x]$ is a Euclidean domain, similar arguments to those of Examples 1–3 show that the

prime ideals are (0) , $(x - a)$ where $a \in \mathbb{R}$, and $(x^2 + ax + b)$ where $x^2 + ax + b$ is an irreducible quadratic. The latter two are maximal ideals, i.e., their quotients are fields. For example: $\mathbb{R}[x]/(x - 3) \cong \mathbb{R}$, $\mathbb{R}[x]/(x^2 + 1) \cong \mathbb{C}$.

3.2.B. UNIMPORTANT EXERCISE. Show that for the last type of prime, of the form $(x^2 + ax + b)$, the quotient is *always* isomorphic to \mathbb{C} .

So we have the points that we would normally expect to see on the real line, corresponding to real numbers; the generic point 0 ; and new points which we may interpret as *conjugate pairs* of complex numbers (the roots of the quadratic). This last type of point should be seen as more akin to the real numbers than to the generic point. You can picture $\mathbb{A}_{\mathbb{R}}^1$ as the complex plane, folded along the real axis. But the key point is that Galois-conjugate points (such as i and $-i$) are considered glued.

Let's explore functions on this space. Consider the function $f(x) = x^3 - 1$. Its value at the point $[(x - 2)]$ is $f(x) = 7$, or perhaps better, $7 \pmod{x - 2}$. How about at $(x^2 + 1)$? We get

$$x^3 - 1 \equiv -x - 1 \pmod{x^2 + 1},$$

which may be profitably interpreted as $-i - 1$.

One moral of this example is that we can work over a non-algebraically closed field if we wish. It is more complicated, but we can recover much of the information we care about.

3.2.C. IMPORTANT EXERCISE. Describe the set $\mathbb{A}_{\mathbb{Q}}^1$. (This is harder to picture in a way analogous to $\mathbb{A}_{\mathbb{R}}^1$. But the rough cartoon of points on a line, as in Figure 3.1, remains a reasonable sketch.)

Example 6 (the affine line over \mathbb{F}_p): $\mathbb{A}_{\mathbb{F}_p}^1 = \text{Spec } \mathbb{F}_p[x]$. As in the previous examples, $\mathbb{F}_p[x]$ is a Euclidean domain, so the prime ideals are of the form (0) or $(f(x))$ where $f(x) \in \mathbb{F}_p[x]$ is an irreducible polynomial, which can be of any degree. Irreducible polynomials correspond to sets of Galois conjugates in $\overline{\mathbb{F}_p}$.

Note that $\text{Spec } \mathbb{F}_p[x]$ has p points corresponding to the elements of \mathbb{F}_p , but also many more (infinitely more, see Exercise 3.2.D). This makes this space much richer than simply p points. For example, a polynomial $f(x)$ is not determined by its values at the p elements of \mathbb{F}_p , but it *is* determined by its values at the points of $\text{Spec } \mathbb{F}_p[x]$. (As we have mentioned before, this is not true for all schemes.)

You should think about this, even if you are a geometric person — this intuition will later turn up in geometric situations. Even if you think you are interested only in working over an algebraically closed field (such as \mathbb{C}), you will have non-algebraically closed fields (such as $\mathbb{C}(x)$) forced upon you.

3.2.D. EXERCISE. If k is a field, show that $\text{Spec } k[x]$ has infinitely many points. (Hint: Euclid's proof of the infinitude of primes of \mathbb{Z} .)

Example 7 (the complex affine plane): $\mathbb{A}_{\mathbb{C}}^2 = \text{Spec } \mathbb{C}[x, y]$. (As with Examples 1 and 2, our discussion will apply with \mathbb{C} replaced by *any* algebraically closed field.) Sadly, $\mathbb{C}[x, y]$ is not a principal ideal domain: (x, y) is not a principal ideal. We can quickly name *some* prime ideals. One is (0) , which has the same flavor as the (0) ideals in the previous examples. $(x - 2, y - 3)$ is prime, and indeed maximal, because $\mathbb{C}[x, y]/(x - 2, y - 3) \cong \mathbb{C}$, where this isomorphism is via $f(x, y) \mapsto f(2, 3)$.

More generally, $(x - a, y - b)$ is prime for any $(a, b) \in \mathbb{C}^2$. Also, if $f(x, y)$ is an irreducible polynomial (e.g., $y - x^2$ or $y^2 - x^3$) then $(f(x, y))$ is prime.

3.2.E. EXERCISE. Show that we have identified all the prime ideals of $\mathbb{C}[x, y]$. Hint: Suppose \mathfrak{p} is a prime ideal that is not principal. Show you can find $f(x, y), g(x, y) \in \mathfrak{p}$ with no common factor. By considering the Euclidean algorithm in the Euclidean domain $\mathbb{C}(x)[y]$, show that you can find a nonzero $h(x) \in (f(x, y), g(x, y)) \subset \mathfrak{p}$. Using primality, show that one of the linear factors of $h(x)$, say $(x - a)$, is in \mathfrak{p} . Similarly show there is some $(y - b) \in \mathfrak{p}$.




FIGURE 3.3. Picturing $\mathbb{A}_{\mathbb{C}}^2 = \text{Spec } \mathbb{C}[x, y]$

We now attempt to draw a picture of $\mathbb{A}_{\mathbb{C}}^2$ (see Figure 3.3). The maximal prime ideals of $\mathbb{C}[x, y]$ correspond to the traditional points in \mathbb{C}^2 : $[(x - a, y - b)]$ corresponds to $(a, b) \in \mathbb{C}^2$. We now have to visualize the “bonus points”. $[(0)]$ somehow lives behind all of the traditional points; it is somewhere on the plane, but nowhere in particular. So for example, it does not lie on the parabola $y = x^2$. The point $[(y - x^2)]$ lies on the parabola $y = x^2$, but nowhere in particular on it. (Figure 3.3 is a bit misleading. For example, the point $[(0)]$ isn’t in the fourth quadrant; it is somehow near every other point, which is why it is depicted as a somewhat diffuse large dot.) You can see from this picture that we already are implicitly thinking about “dimension”. The prime ideals $(x - a, y - b)$ are somehow of dimension 0, the prime ideals $(f(x, y))$ are of dimension 1, and (0) is of dimension 2. (All of our dimensions here are *complex* or *algebraic* dimensions. The complex plane \mathbb{C}^2 has real dimension 4, but complex dimension 2. Complex dimensions are in general half of real dimensions.) We won’t define dimension precisely until Chapter 12, but you should feel free to keep it in mind before then.

Note too that maximal ideals correspond to the “smallest” points. Smaller ideals correspond to “bigger” points. “One prime ideal contains another” means that the points “have the opposite containment.” All of this will be made precise once we have a topology. This order-reversal is a little confusing, and will remain so even once we have made the notions precise.

We now come to the obvious generalization of Example 7.

Example 8 (complex affine n -space — important!): Let $\mathbb{A}_{\mathbb{C}}^n := \operatorname{Spec} \mathbb{C}[x_1, \dots, x_n]$. (Important definition: More generally, \mathbb{A}_A^n is defined to be $\operatorname{Spec} A[x_1, \dots, x_n]$, where A is an arbitrary ring. When the base ring is clear from context, the subscript A is often omitted. For pedants: the notation \mathbb{A}_A^n implicitly includes the data of the n “coordinate functions” x_1, \dots, x_n .) For concreteness, let’s consider $n = 3$. We now have an interesting question in what at first appears to be pure algebra: What are the prime ideals of $\mathbb{C}[x, y, z]$?

Analogously to before, $(x - a, y - b, z - c)$ is a prime ideal. This is a maximal ideal, because its residue ring is a field (\mathbb{C}); we think of these as “0-dimensional points”. We will often write (a, b, c) for $[(x - a, y - b, z - c)]$ because of our geometric interpretation of these ideals. There are no more maximal ideals, by Hilbert’s Weak Nullstellensatz.

3.2.5. Hilbert’s Weak Nullstellensatz. — *If k is an algebraically closed field, then the maximal ideals of $k[x_1, \dots, x_n]$ are precisely those ideals of the form $(x_1 - a_1, \dots, x_n - a_n)$, where $a_i \in k$.*

We may as well state a slightly stronger version now.

3.2.6. Hilbert’s Nullstellensatz. — *If k is any field, every maximal ideal of $k[x_1, \dots, x_n]$ has residue field a finite extension of k . Translation: any field extension of k that is finitely generated as a ring is necessarily also finitely generated as a module (i.e., is a finite extension of fields).*

This statement is also often called *Zariski’s Lemma*.

3.2.F. EXERCISE. Show that the Nullstellensatz 3.2.6 implies the Weak Nullstellensatz 3.2.5.

We will prove the Nullstellensatz in §8.4.3, and again in Exercise 12.2.B.

The following fact is a useful accompaniment to the Nullstellensatz.

3.2.G. EXERCISE (NOT REQUIRING THE NULLSTELLENSATZ). Any integral domain A which is a finite k -algebra (i.e., a k -algebra that is a finite-dimensional vector space over k) must be a field. Hint: for any nonzero $x \in A$, show $\times x: A \rightarrow A$ is an isomorphism. (Thus, in combination the Nullstellensatz 3.2.6, we see that *prime ideals of $k[x_1, \dots, x_n]$ with finite residue ring are the same as maximal ideals of $k[x_1, \dots, x_n]$. This is worth remembering.*)

There are other prime ideals of $\mathbb{C}[x, y, z]$ too. We have (0) , which corresponds to a “3-dimensional point”. We have $(f(x, y, z))$, where f is irreducible. To this we associate the “hypersurface” $f = 0$, so this is “2-dimensional” in nature. But we have not found them all! One clue: we have prime ideals of “dimension” 0, 2, and 3 — we are missing “dimension 1”. Here is one such prime ideal: (x, y) . We picture this as the locus where $x = y = 0$, which is the z -axis. This is a prime ideal, as the corresponding quotient $\mathbb{C}[x, y, z]/(x, y) \cong \mathbb{C}[z]$ is an integral domain (and should be interpreted as the functions on the z -axis). There are lots of “1-dimensional prime ideals”, and it is not possible to classify them in a reasonable way. It will turn out that they correspond to things that we think of as irreducible curves. Thus remarkably the answer to the purely algebraic question (“what are the prime ideals of $\mathbb{C}[x, y, z]$ ”) is fundamentally geometric!

The fact that the points of $\mathbb{A}_{\mathbb{Q}}^1$ corresponding to maximal ideals of the ring $\mathbb{Q}[x]$ (what we will soon call “closed points”, see Definition 3.6.8) can be interpreted as points of $\overline{\mathbb{Q}}$ where Galois-conjugates are glued together (Exercise 3.2.C) extends to $\mathbb{A}_{\mathbb{Q}}^n$. For example, in $\mathbb{A}_{\mathbb{Q}}^2$, $(\sqrt{2}, \sqrt{2})$ is glued to $(-\sqrt{2}, -\sqrt{2})$ but not to $(\sqrt{2}, -\sqrt{2})$. The following exercise will give you some idea of how this works.

3.2.H. EXERCISE. Describe the maximal ideal of $\mathbb{Q}[x, y]$ corresponding to $(\sqrt{2}, \sqrt{2})$ and $(-\sqrt{2}, -\sqrt{2})$. Describe the maximal ideal of $\mathbb{Q}[x, y]$ corresponding to $(\sqrt{2}, -\sqrt{2})$ and $(-\sqrt{2}, \sqrt{2})$. What are the residue fields in each case?

The description of “closed points” of $\mathbb{A}_{\mathbb{Q}}^2$ (those points corresponding to maximal ideals of the ring $\mathbb{Q}[x, y]$) as Galois-orbits of points in $\overline{\mathbb{Q}}^2$ can even be extended to other “nonclosed” points, as follows.

3.2.I. UNIMPORTANT AND TRICKY BUT FUN EXERCISE. Consider the map of sets $\phi: \mathbb{C}^2 \rightarrow \mathbb{A}_{\mathbb{Q}}^2$ defined as follows. (z_1, z_2) is sent to the prime ideal of $\mathbb{Q}[x, y]$ consisting of polynomials vanishing at (z_1, z_2) .

(a) What is the image of (π, π^2) ?


★ (b) Show that ϕ is surjective. (Warning: You will need some ideas we haven’t discussed in order to solve this. Once we define the Zariski topology on $\mathbb{A}_{\mathbb{Q}}^2$, you will be able to check that ϕ is continuous, where we give \mathbb{C}^2 the classical topology. This example generalizes. For example, you may later be able to generalize this to arbitrary dimension.)

3.2.7. Quotients and localizations. Two natural ways of getting new rings from old — quotients and localizations — have interpretations in terms of spectra.

3.2.8. Quotients: $\text{Spec } A/I$ as a subset of $\text{Spec } A$. It is an important fact that the prime ideals of A/I are in bijection with the prime ideals of A containing I .

3.2.J. ESSENTIAL ALGEBRA EXERCISE (MANDATORY IF YOU HAVEN’T SEEN IT BEFORE). Suppose A is a ring, and I an ideal of A . Let $\phi: A \rightarrow A/I$. Show that ϕ^{-1} gives an inclusion-preserving bijection between prime ideals of A/I and prime ideals of A containing I . Thus we can picture $\text{Spec } A/I$ as a subset of $\text{Spec } A$.

As an important motivational special case, you now have a picture of *affine complex varieties*. Suppose A is a finitely generated \mathbb{C} -algebra, generated by x_1, \dots, x_n , with relations $f_1(x_1, \dots, x_n) = \dots = f_r(x_1, \dots, x_n) = 0$. Then this description in terms of generators and relations naturally gives us an interpretation of $\text{Spec } A$ as a subset of $\mathbb{A}_{\mathbb{C}}^n$, which we think of as “traditional points” (n -tuples of complex numbers) along with some “bonus” points we haven’t yet fully described. To see which of the traditional points are in $\text{Spec } A$, we simply solve the equations $f_1 = \dots = f_r = 0$. For example, $\text{Spec } \mathbb{C}[x, y, z]/(x^2 + y^2 - z^2)$ may be pictured as shown in Figure 3.4. (Admittedly this is just a “sketch of the \mathbb{R} -points”, but we will still find it helpful later.) This entire picture carries over (along with the Nullstellensatz) with \mathbb{C} replaced by any algebraically closed field. Indeed, the picture of Figure 3.4 can be said to depict $k[x, y, z]/(x^2 + y^2 - z^2)$ for most algebraically closed fields k (although it is misleading in characteristic 2, because of the coincidence $x^2 + y^2 - z^2 = (x + y + z)^2$).

FIGURE 3.4. A “picture” of $\text{Spec } \mathbb{C}[x, y, z]/(x^2 + y^2 - z^2)$

3.2.9. Localizations: $\text{Spec } S^{-1}A$ as a subset of $\text{Spec } A$. The following exercise shows how prime ideals behave under localization.

3.2.K. ESSENTIAL ALGEBRA EXERCISE (MANDATORY IF YOU HAVEN’T SEEN IT BEFORE). Suppose S is a multiplicative subset of A . Describe an order-preserving bijection of the prime ideals of $S^{-1}A$ with the prime ideals of A that *don’t meet* the multiplicative set S .

Recall from §1.3.3 that there are two important flavors of localization. The first is $A_f = \{1, f, f^2, \dots\}^{-1}A$ where $f \in A$. A motivating example is $A = \mathbb{C}[x, y]$, $f = y - x^2$. The second is $A_{\mathfrak{p}} = (A - \mathfrak{p})^{-1}A$, where \mathfrak{p} is a prime ideal. A motivating example is $A = \mathbb{C}[x, y]$, $S = A - (x, y)$.

If $S = \{1, f, f^2, \dots\}$, the prime ideals of $S^{-1}A$ are just those prime ideals not containing f — the points where “ f doesn’t vanish”. (In §3.5, we will call this a *distinguished open set*, once we know what open sets are.) So to picture $\text{Spec } \mathbb{C}[x, y]_{y-x^2}$, we picture the affine plane, and throw out those points on the parabola $y - x^2$ — the points (a, a^2) for $a \in \mathbb{C}$ (by which we mean $[(x - a, y - a^2)]$), as well as the “new kind of point” $[(y - x^2)]$.

It can be initially confusing to think about localization in the case where zerodivisors are inverted, because localization $A \rightarrow S^{-1}A$ is not injective (Exercise 1.3.C). Geometric intuition can help. Consider the case $A = \mathbb{C}[x, y]/(xy)$ and $f = x$. What is the localization A_f ? The space $\text{Spec } \mathbb{C}[x, y]/(xy)$ “is” the union of the two axes in the plane. Localizing means throwing out the locus where x vanishes. So we are left with the x -axis, minus the origin, so we expect $\text{Spec } \mathbb{C}[x]_x$. So there should be some natural isomorphism

$$(\mathbb{C}[x, y]/(xy))_x \xrightarrow{\sim} \mathbb{C}[x]_x.$$

3.2.L. EXERCISE. Show that these two rings are isomorphic. (You will see that y on the left goes to 0 on the right.)

If $S = A - \mathfrak{p}$, the prime ideals of $S^{-1}A$ are just the prime ideals of A contained in \mathfrak{p} . In our example $A = \mathbb{C}[x, y]$, $\mathfrak{p} = (x, y)$, we keep all those points corresponding to “things through the origin”, i.e., the 0-dimensional point (x, y) , the 2-dimensional point (0) , and those 1-dimensional points $(f(x, y))$ where $f(0, 0) = 0$, i.e., those

“irreducible curves through the origin”. You can think of this being a shred of the plane near the origin; anything not actually “visible” at the origin is discarded (see Figure 3.5).




FIGURE 3.5. Picturing $\text{Spec } \mathbb{C}[x, y]_{(x, y)}$ as a “shred of $\mathbb{A}^2_{\mathbb{C}}$ ”. Only those points near the origin remain.

Another example is when $A = k[x]$, and $\mathfrak{p} = (x)$ (or more generally when \mathfrak{p} is any maximal ideal). Then $A_{\mathfrak{p}}$ has only two prime ideals (Exercise 3.2.A(b)). You should see this as the germ of a “smooth curve”, where one point is the “classical point”, and the other is the “generic point of the curve”. This is an example of a discrete valuation ring, and indeed all discrete valuation rings should be visualized in such a way. We will discuss discrete valuation rings in §13.5. By then we will have justified the use of the words “smooth” and “curve”. (Reality check: try to picture Spec of \mathbb{Z} localized at (2) and at (0) . How do the two pictures differ?)

3.2.10. Important fact: Maps of rings induce maps of spectra (as sets). We now make an observation that will later grow up to be the notion of morphisms of schemes.

3.2.M. IMPORTANT EASY EXERCISE. If $\phi: B \rightarrow A$ is a map of rings, and \mathfrak{p} is a prime ideal of A , show that $\phi^{-1}(\mathfrak{p})$ is a prime ideal of B .

Hence a map of rings $\phi: B \rightarrow A$ induces a map of sets $\text{Spec } A \rightarrow \text{Spec } B$ “in the opposite direction”. This gives a contravariant functor from the category of rings to the category of sets: the composition of two maps of rings induces the composition of the corresponding maps of spectra.

3.2.N. EASY EXERCISE (REALITY CHECK). Let B be a ring.

(a) Suppose $I \subset B$ is an ideal. Show that the map $\text{Spec } B/I \rightarrow \text{Spec } B$ is the inclusion of §3.2.8.

(b) Suppose $S \subset B$ is a multiplicative set. Show that the map $\text{Spec } S^{-1}B \rightarrow \text{Spec } B$ is the inclusion of §3.2.9.

3.2.11. An explicit example. In the case of “affine complex varieties” (or indeed affine varieties over any algebraically closed field), the translation between maps given by explicit formulas and maps of rings is quite direct. For example, consider a map from the parabola in \mathbb{C}^2 (with coordinates a and b) given by $b = a^2$, to the

“curve” in \mathbb{C}^3 (with coordinates x , y , and z) cut out by the equations $y = x^2$ and $z = y^2$. Suppose the map sends the point $(a, b) \in \mathbb{C}^2$ to the point $(a, b, b^2) \in \mathbb{C}^3$. In our new language, we have a map

$$\operatorname{Spec} \mathbb{C}[a, b]/(b - a^2) \longrightarrow \operatorname{Spec} \mathbb{C}[x, y, z]/(y - x^2, z - y^2)$$

given by

$$\mathbb{C}[a, b]/(b - a^2) \longleftarrow \mathbb{C}[x, y, z]/(y - x^2, z - y^2)$$

$$(a, b, b^2) \longleftarrow (x, y, z),$$

i.e., $x \mapsto a$, $y \mapsto b$, and $z \mapsto b^2$. If the idea is not yet clear, the following two exercises are very much worth doing — they can be very confusing the first time you see them, and very enlightening (and finally, trivial) when you finally figure them out.




FIGURE 3.6. The map $\mathbb{C} \rightarrow \mathbb{C}$ given by $x \mapsto y = x^2$

3.2.O. IMPORTANT EXERCISE (SPECIAL CASE). Consider the map of complex manifolds sending $\mathbb{C} \rightarrow \mathbb{C}$ via $x \mapsto y = x^2$. We interpret the “source” \mathbb{C} as the “ x -line”, and the “target” \mathbb{C} the “ y -line”. You can picture it as the projection of the parabola $y = x^2$ in the xy -plane to the y -axis (see Figure 3.6). Interpret the corresponding map of rings as given by $\mathbb{C}[y] \rightarrow \mathbb{C}[x]$ by $y \mapsto x^2$. Verify that the preimage (the fiber) above the point $a \in \mathbb{C}$ is the point(s) $\pm\sqrt{a} \in \mathbb{C}$, using the definition given above. (A more sophisticated version of this example appears in Example 10.3.4. Warning: the roles of x and y are swapped there, in order to picture double covers in a certain way.)

3.2.P. IMPORTANT EXERCISE (GENERALIZING EXAMPLE 3.2.11). Suppose k is a field, and $f_1, \dots, f_n \in k[x_1, \dots, x_m]$ are given. Let $\phi: k[y_1, \dots, y_n] \rightarrow k[x_1, \dots, x_m]$ be the ring morphism defined by $y_i \mapsto f_i$.

(a) Show that ϕ induces a map of sets $\operatorname{Spec} k[x_1, \dots, x_m]/I \rightarrow \operatorname{Spec} k[y_1, \dots, y_n]/J$ for any ideals $I \subset k[x_1, \dots, x_m]$ and $J \subset k[y_1, \dots, y_n]$ such that $\phi(J) \subset I$. (You may wish to consider the case $I = 0$ and $J = 0$ first. In fact, part (a) has nothing to do with k -algebras; you may wish to prove the statement when the rings $k[x_1, \dots, x_m]$

and $k[y_1, \dots, y_n]$ are replaced by general rings A and B .)

(b) Show that the map of part (a) sends the point $(a_1, \dots, a_m) \in k^m$ (or more precisely, $[(x_1 - a_1, \dots, x_m - a_m)] \in \text{Spec } k[x_1, \dots, x_m]$) to

$$(f_1(a_1, \dots, a_m), \dots, f_n(a_1, \dots, a_m)) \in k^n.$$

3.2.Q. EXERCISE: PICTURING $\mathbb{A}_{\mathbb{Z}}^n$. Consider the map of sets $\pi: \mathbb{A}_{\mathbb{Z}}^n \rightarrow \text{Spec } \mathbb{Z}$, given by the ring map $\mathbb{Z} \rightarrow \mathbb{Z}[x_1, \dots, x_n]$. If p is prime, describe a bijection between the fiber $\pi^{-1}([(p)])$ and $\mathbb{A}_{\mathbb{F}_p}^n$. (You won't need to describe either set! Which is good because you can't.) This exercise may give you a sense of how to picture maps (see Figure 3.7), and in particular why you can think of $\mathbb{A}_{\mathbb{Z}}^n$ as an " \mathbb{A}^n -bundle" over $\text{Spec } \mathbb{Z}$. (Can you interpret the fiber over $[(0)]$ as \mathbb{A}_k^n for some field k ?)




FIGURE 3.7. A picture of $\mathbb{A}_{\mathbb{Z}}^n \rightarrow \text{Spec } \mathbb{Z}$ as a “family of \mathbb{A}^n ’s”, or an “ \mathbb{A}^n -bundle over $\text{Spec } \mathbb{Z}$ ”. What is the field k ? How should you “geometrically” think of the three points indicated?

3.2.12. Functions are not determined by their values at points: the fault of nilpotents. We conclude this section by describing some strange behavior. We are developing machinery that will let us bring our geometric intuition to algebra. There is one serious serious point where your intuition will be false, so you should know now, and adjust your intuition appropriately. As noted by Mumford ([Mu2, p. 12]), “it is this aspect of schemes which was most scandalous when Grothendieck defined them.”

Suppose we have a function (ring element) vanishing at all points. Then it is not necessarily the zero function! The translation of this question is: is the intersection of all prime ideals necessarily just 0? The answer is no, as is shown by the example of the ring of dual numbers $k[\epsilon]/(\epsilon^2)$: $\epsilon \neq 0$, but $\epsilon^2 = 0$. (We saw this ring in Exercise 3.2.A(a).) Any function whose power is zero certainly lies in the intersection of all prime ideals.

3.2.R. EXERCISE. Ring elements that have a power that is 0 are called **nilpotents**.
 (a) Show that if I is an ideal of nilpotents, then the inclusion $\text{Spec } B/I \rightarrow \text{Spec } B$ of Exercise 3.2.J is a bijection. Thus nilpotents don't affect the underlying set. (We will soon see in §3.4.7 that they won't affect the topology either — the difference will be in the structure sheaf.)
 (b) Show that the nilpotents of a ring B form an ideal. This ideal is called the **nilradical**, and is denoted $\mathfrak{N} = \mathfrak{N}(B)$.

Thus the nilradical is contained in the intersection of all the prime ideals. The converse is also true:

3.2.13. Theorem. — *The nilradical $\mathfrak{N}(A)$ is the intersection of all the prime ideals of A . Geometrically: a function on $\text{Spec } A$ vanishes at every point if and only if it is nilpotent.*

3.2.S. EXERCISE. If you don't know this theorem, then look it up, or better yet, prove it yourself. (Hint: Use the fact that any proper ideal of A is contained in a maximal ideal, which requires Zorn's Lemma. Possible further hint: Suppose $x \notin \mathfrak{N}(A)$. We wish to show that there is a prime ideal not containing x . Show that A_x is not the 0-ring, by showing that $1 \neq 0$.)

3.2.14. In particular, although it is upsetting that functions are not determined by their values at points, we have precisely specified what the failure of this intuition is: two functions have the same values at points if and only if they differ by a nilpotent. You should think of this geometrically: a function vanishes at every point of the spectrum of a ring if and only if it has a power that is zero. And if there are no nonzero nilpotents — if $\mathfrak{N} = (0)$ — then functions *are* determined by their values at points. If a ring has no nonzero nilpotents, we say that it is **reduced**.

3.2.T. FUN UNIMPORTANT EXERCISE: DERIVATIVES WITHOUT DELTAS AND EPSILONS (OR AT LEAST WITHOUT DELTAS). Suppose we have a polynomial $f(x) \in k[x]$. Instead, we work in $k[x, \epsilon]/(\epsilon^2)$. What then is $f(x + \epsilon)$? (Do a couple of examples, then prove the pattern you observe.) This is a hint that nilpotents will be important in defining differential information (Chapter 22).

3.3 Visualizing schemes: Generic points

A heavy warning used to be given that pictures are not rigorous; this has never had its bluff called and has permanently frightened its victims into playing for safety. Some pictures, of course, are not rigorous, but I should say most are (and I use them whenever possible myself).

— J. E. Littlewood, [Lit, p. 54]

We all know that Art is not truth. Art is a lie that makes us realize truth, at least the truth that is given us to understand. The artist must know the manner whereby to convince others of the truthfulness of his lies.

— P. Picasso [Pi, p. 315]

For years, you have been able to picture $x^2 + y^2 = 1$ in the plane, and you now have an idea of how to picture $\text{Spec } \mathbb{Z}$. If we are claiming to understand rings as geometric objects (through the Spec functor), then we should wish to develop

geometric insight into them. To develop geometric intuition about schemes, it is helpful to have pictures in your mind, extending your intuition about geometric spaces you are already familiar with. As we go along, we will empirically develop some idea of what schemes should look like. This section summarizes what we have gleaned so far.

Some mathematicians prefer to think completely algebraically, and never think in terms of pictures. Others will be disturbed by the fact that this is an art, not a science. And finally, this hand-waving will necessarily never be used in the rigorous development of the theory. For these reasons, you may wish to skip these sections. However, having the right picture in your mind can greatly help understanding what facts should be true, and how to prove them.

Our starting point is the example of “affine complex varieties” (things cut out by equations involving a finite number of variables over \mathbb{C}), and more generally similar examples over arbitrary algebraically closed fields. We begin with notions that are intuitive (“traditional” points behaving the way you expect them to), and then add in the two features which are new and disturbing, generic points and nonreduced behavior. You can then extend this notion to seemingly different spaces, such as $\text{Spec } \mathbb{Z}$.

Hilbert’s Weak Nullstellensatz 3.2.5 shows that the “traditional points” are present as points of the scheme, and this carries over to any algebraically closed field. If the field is not algebraically closed, the traditional points are glued together into clumps by Galois conjugation, as in Examples 5 (the real affine line) and 6 (the affine line over \mathbb{F}_p) in §3.2. This is a geometric interpretation of Hilbert’s Nullstellensatz 3.2.6.

But we have some additional points to add to the picture. You should remember that they “correspond” to “irreducible” “closed” (algebraic) subsets. As motivation, consider the case of the complex affine plane (Example 7): we had one for each irreducible polynomial, plus one corresponding to the entire plane. We will make “closed” precise when we define the Zariski topology (in the next section). You may already have an idea of what “irreducible” should mean; we make that precise at the start of §3.6. By “correspond” we mean that each closed irreducible subset has a corresponding point sitting on it, called its *generic point* (defined in §3.6). It is a new point, distinct from all the other points in the subset. (The correspondence is described in Exercise 3.7.F for $\text{Spec } A$, and in Exercise 5.1.B for schemes in general.) We don’t know precisely where to draw the generic point, so we may stick it arbitrarily anywhere, but you should think of it as being “almost everywhere”, and in particular, near every other point in the subset.

In §3.2.8, we saw how the points of $\text{Spec } A/I$ should be interpreted as a subset of $\text{Spec } A$. So for example, when you see $\text{Spec } \mathbb{C}[x, y]/(x + y)$, you should picture this not just as a line, but as a line in the xy -plane; the choice of generators x and y of the algebra $\mathbb{C}[x, y]$ implies an inclusion into affine space.

In §3.2.9, we saw how the points of $\text{Spec } S^{-1}A$ should be interpreted as subsets of $\text{Spec } A$. The two most important cases were discussed. The points of $\text{Spec } A_f$ correspond to the points of $\text{Spec } A$ where f doesn’t vanish; we will later (§3.5) interpret this as a distinguished open set.

If \mathfrak{p} is a prime ideal, then $\text{Spec } A_{\mathfrak{p}}$ should be seen as a “shred of the space $\text{Spec } A$ near the subset corresponding to \mathfrak{p} ”. The simplest nontrivial case of this

is $A = k[x]$ and $\mathfrak{p} = (x) \subset A$ (see Exercise 3.2.A, which we discuss again in Exercise 3.4.K).

The test of a first-rate intelligence is the ability to hold two opposed ideas in the mind at the same time, and still retain the ability to function.

— F. Scott Fitzgerald, *The Crack-Up* [Fi, p. 41]

3.4 The underlying topological space of an affine scheme

We next introduce the *Zariski topology* on the spectrum of a ring. When you first hear the definition, it seems odd, but with a little experience it becomes reasonable. As motivation, consider $\mathbb{A}_{\mathbb{C}}^2 = \operatorname{Spec} \mathbb{C}[x, y]$, the complex plane (with a few extra points). In algebraic geometry, we will only be allowed to consider algebraic functions, i.e., polynomials in x and y . The locus where a polynomial vanishes should reasonably be a closed set, and the Zariski topology is defined by saying that the only sets we should consider closed should be these sets, and other sets forced to be closed by these. In other words, it is the coarsest topology where these sets are closed.

3.4.1. In particular, although topologies are often described using open subsets, it will be more convenient for us to define this topology in terms of closed subsets. If S is a subset of a ring A , define the **Vanishing set** of S by

$$V(S) := \{[\mathfrak{p}] \in \operatorname{Spec} A : S \subset \mathfrak{p}\}.$$

It is the set of points on which all elements of S are zero. (It should now be second nature to equate “vanishing at a point” with “contained in a prime”.) We declare that these — and no others — are the closed subsets.

For example, consider $V(xy, yz) \subset \mathbb{A}_{\mathbb{C}}^3 = \operatorname{Spec} \mathbb{C}[x, y, z]$. Which points are contained in this locus? We think of this as solving $xy = yz = 0$. Of the “traditional” points (interpreted as ordered triples of complex numbers, thanks to the Hilbert’s Nullstellensatz 3.2.5), we have the points where $y = 0$ or $x = z = 0$: the xz -plane and the y -axis respectively. Of the “new” points, we have the generic point of the xz -plane (also known as the point $[(y)]$), and the generic point of the y -axis (also known as the point $[(x, z)]$). You might imagine that we also have a number of “one-dimensional” points contained in the xz -plane.

3.4.A. EASY EXERCISE. Check that the x -axis is contained in $V(xy, yz)$. (The x -axis is defined by $y = z = 0$, and the y -axis and z -axis are defined analogously.)

Let’s return to the general situation. The following exercise lets us restrict attention to vanishing sets of *ideals*.

3.4.B. EASY EXERCISE. Show that if (S) is the ideal generated by S , then $V(S) = V((S))$.

3.4.2. Definition. We define the **Zariski topology** by declaring that $V(S)$ is closed for all S . (We may as well state here that **Zariski closure** means closure in the Zariski topology.) Let’s check that the Zariski topology is indeed a topology:

3.4.C. EXERCISE.

- (a) Show that \emptyset and $\text{Spec } A$ are both open subsets of $\text{Spec } A$.
- (b) If I_i is a collection of ideals (as i runs over some index set), show that $\cap_i V(I_i) = V(\sum_i I_i)$. Hence the union of any collection of open sets is open.
- (c) Show that $V(I_1) \cup V(I_2) = V(I_1 I_2)$. (The **product of two ideals** I_1 and I_2 of A are finite A -linear combinations of products of elements of I_1 and I_2 , i.e., elements of the form $\sum_{j=1}^n i_{1,j} i_{2,j}$, where $i_{k,j} \in I_k$. Equivalently, it is the ideal generated by products of elements of I_1 and I_2 . You should quickly check that this is an ideal, and that products are associative, i.e., $(I_1 I_2) I_3 = I_1 (I_2 I_3)$.) Hence the intersection of any finite number of open sets is open.

3.4.3. Properties of the “vanishing set” function $V(\cdot)$. The function $V(\cdot)$ is obviously inclusion-reversing: If $S_1 \subset S_2$, then $V(S_2) \subset V(S_1)$. Warning: We could have equality in the second inclusion without equality in the first, as the next exercise shows.

3.4.D. EXERCISE/DEFINITION. If $I \subset A$ is an ideal, then define its **radical** by

$$\sqrt{I} := \{r \in A : r^n \in I \text{ for some } n \in \mathbb{Z}^{>0}\}.$$

For example, the nilradical \mathfrak{N} (§3.2.R) is $\sqrt{(0)}$. Show that \sqrt{I} is an ideal (cf. Exercise 3.2.R(b)). Show that $V(\sqrt{I}) = V(I)$. We say *an ideal is radical* if it equals its own radical. Show that $\sqrt{\sqrt{I}} = \sqrt{I}$, and that prime ideals are radical.

Here are two useful consequences. As $(I \cap J)^2 \subset IJ \subset I \cap J$ (products of ideals were defined in Exercise 3.4.C), we have that $V(IJ) = V(I \cap J) (= V(I) \cup V(J)$ by Exercise 3.4.C(c)). Also, combining this with Exercise 3.4.B, we see $V(S) = V((S)) = V(\sqrt{(S)})$.

3.4.E. EXERCISE (RADICALS COMMUTE WITH FINITE INTERSECTIONS). If I_1, \dots, I_n are ideals of a ring A , show that $\sqrt{\cap_{i=1}^n I_i} = \cap_{i=1}^n \sqrt{I_i}$. We will use this property repeatedly without referring back to this exercise.

3.4.F. EXERCISE FOR LATER USE. Show that \sqrt{I} is the intersection of all the prime ideals containing I . (Hint: Use Theorem 3.2.13 on an appropriate ring.)

3.4.4. Examples. Let’s see how this meshes with our examples from the previous section.

Recall that $\mathbb{A}_{\mathbb{C}}^1$, as a set, was just the “traditional” points (corresponding to maximal ideals, in bijection with $a \in \mathbb{C}$), and one “new” point $[(0)]$. The Zariski topology on $\mathbb{A}_{\mathbb{C}}^1$ is not that exciting: the open sets are the empty set, and $\mathbb{A}_{\mathbb{C}}^1$ minus a finite number of maximal ideals. (It “almost” has the cofinite topology. Notice that the open sets are determined by their intersections with the “traditional points”. The “new” point $[(0)]$ comes along for the ride, which is a good sign that it is harmless. Ignoring the “new” point, observe that the topology on $\mathbb{A}_{\mathbb{C}}^1$ is a coarser topology than the classical topology on \mathbb{C} .)

3.4.G. EXERCISE. Describe the topological space \mathbb{A}_k^1 (cf. Exercise 3.2.D). (Notice that the strange new point $[(0)]$ is “near every other point” — every neighborhood of every point contains $[(0)]$. This is typical of these new points, see Easy Exercise 3.6.N.)

The case of $\text{Spec } \mathbb{Z}$ is similar. The topology is “almost” the cofinite topology in the same way. The open sets are the empty set, and $\text{Spec } \mathbb{Z}$ minus a finite number of “ordinary” $((p))$ where p is prime.

3.4.5. Closed subsets of $\mathbb{A}_{\mathbb{C}}^2$. The case $\mathbb{A}_{\mathbb{C}}^2$ is more interesting. You should think through where the “one-dimensional prime ideals” fit into the picture. In Exercise 3.2.E, we identified all the prime ideals of $\mathbb{C}[x, y]$ (i.e., the points of $\mathbb{A}_{\mathbb{C}}^2$) as the maximal ideals $[(x - a, y - b)]$ (where $a, b \in \mathbb{C}$ — “zero-dimensional points”), the “one-dimensional points” $[(f(x, y))]$ (where $f(x, y)$ is irreducible), and the “two-dimensional point” $[(0)]$.

Then the closed subsets are of the following form:

- (a) the entire space (the closure of the “two-dimensional point” $[(0)]$),
- (b) a finite number (possibly none) of “curves” (each the closure of a “one-dimensional point” — the “one-dimensional point” along with the “zero-dimensional points” “lying on it”) and a finite number (possibly none) of “zero-dimensional points” (what we will soon call “closed points”, see Definition 3.6.8).

We will soon know enough to verify this using general theory, but you can prove it yourself now, using ideas in Exercise 3.2.E. (The key idea: if $f(x, y)$ and $g(x, y)$ are irreducible polynomials that are not multiples of each other, why do their zero sets intersect in a finite number of points?)

3.4.6. Important fact: Maps of rings induce continuous maps of topological spaces. We saw in §3.2.10 that a map of rings $\phi: B \rightarrow A$ induces a map of sets $\pi: \text{Spec } A \rightarrow \text{Spec } B$.

3.4.H. IMPORTANT EASY EXERCISE. By showing that closed sets pull back to closed sets, show that π is a *continuous* map. Interpret Spec as a contravariant functor $\text{Rings} \rightarrow \text{Top}$.

Not all continuous maps arise in this way. Consider for example the continuous map on $\mathbb{A}_{\mathbb{C}}^1$ that is the identity except 0 and 1 (i.e., $[(x)]$ and $[(x - 1)]$ are swapped; no polynomial can manage this marvellous feat.

In §3.2.10, we saw that $\text{Spec } B/I$ and $\text{Spec } S^{-1}B$ are naturally *subsets* of $\text{Spec } B$. It is natural to ask if the Zariski topology behaves well with respect to these inclusions, and indeed it does.

3.4.I. IMPORTANT EXERCISE (CF. EXERCISE 3.2.N). Suppose that $I, S \subset B$ are an ideal and multiplicative subset respectively.

- (a) Show that $\text{Spec } B/I$ is naturally a *closed* subset of $\text{Spec } B$. If $S = \{1, f, f^2, \dots\}$ ($f \in B$), show that $\text{Spec } S^{-1}B$ is naturally an *open* subset of $\text{Spec } B$. Show that for arbitrary S , $\text{Spec } S^{-1}B$ need not be open or closed. (Hint: $\text{Spec } \mathbb{Q} \subset \text{Spec } \mathbb{Z}$, or possibly Figure 3.5.)

- (b) Show that the Zariski topology on $\text{Spec } B/I$ (resp. $\text{Spec } S^{-1}B$) is the subspace topology induced by inclusion in $\text{Spec } B$. (Hint: compare closed subsets.)

3.4.7. In particular, if $I \subset \mathfrak{N}$ is an ideal of nilpotents, the bijection $\text{Spec } B/I \rightarrow \text{Spec } B$ (Exercise 3.2.R) is a homeomorphism. Thus nilpotents don’t affect the topological space. (The difference will be in the structure sheaf.)

3.4.J. USEFUL EXERCISE FOR LATER. Suppose $I \subset B$ is an ideal. Show that f vanishes on $V(I)$ if and only if $f \in \sqrt{I}$ (i.e., $f^n \in I$ for some $n \geq 1$). (Hint: Exercise 3.4.F. If you are stuck, you will get another hint when you see Exercise 3.5.E.)

3.4.K. EASY EXERCISE (CF. EXERCISE 3.2.A). Describe the topological space $\text{Spec } k[x]_{(x)}$.

3.5 A base of the Zariski topology on $\text{Spec } A$: Distinguished open sets

If $f \in A$, define the **distinguished open set**

$$\begin{aligned} D(f) &:= \{[p] \in \text{Spec } A : f \notin p\} \\ &= \{[p] \in \text{Spec } A : f([p]) \neq 0\}. \end{aligned}$$

It is the locus where f doesn't vanish. (I often privately write this as $D(f \neq 0)$ to remind myself of this. I also privately call this the “Doesn't-vanish set” of f in analogy with $V(f)$ being the Vanishing set of f .) We have already seen this set when discussing $\text{Spec } A_f$ as a subset of $\text{Spec } A$. For example, we have observed that the Zariski-topology on the distinguished open set $D(f) \subset \text{Spec } A$ coincides with the Zariski topology on $\text{Spec } A_f$ (Exercise 3.4.I).

The reason these sets are important is that they form a particularly nice base for the (Zariski) topology:

3.5.A. EASY EXERCISE. Show that the distinguished open sets form a base for the (Zariski) topology. (Hint: Given a subset $S \subset A$, show that the complement of $V(S)$ is $\cup_{f \in S} D(f)$.)

Here are some important but not difficult exercises to give you a feel for this concept.

3.5.B. NIFTY EXERCISE. Suppose $f_i \in A$ as i runs over some index set J . Show that $\cup_{i \in J} D(f_i) = \text{Spec } A$ if and only if $(\{f_i\}_{i \in J}) = A$, or equivalently and very usefully, if there are a_i ($i \in J$), all but finitely many 0, such that $\sum_{i \in J} a_i f_i = 1$. (One of the directions will use the fact that any proper ideal of A is contained in some maximal ideal.)

3.5.C. EXERCISE. Show that if $\text{Spec } A$ is an infinite union of distinguished open sets $\cup_{j \in J} D(f_j)$, then in fact it is a union of a finite number of these, i.e., there is a finite subset J' so that $\text{Spec } A = \cup_{j \in J'} D(f_j)$. (Hint: Exercise 3.5.B.)

3.5.D. EASY EXERCISE. Show that $D(f) \cap D(g) = D(fg)$.

3.5.E. IMPORTANT EXERCISE (CF. EXERCISE 3.4.J). Show that $D(f) \subset D(g)$ if and only if $f^n \in (g)$ for some $n \geq 1$, if and only if g is an invertible element of A_f .

We will use Exercise 3.5.E often. You can solve it thinking purely algebraically, but the following geometric interpretation may be helpful. (You should try to draw your own picture to go with this discussion.) Inside $\text{Spec } A$, we have the closed subset $V(g) = \text{Spec } A/(g)$, where g vanishes, and its complement $D(g)$, where g doesn't vanish. Then f is a function on this closed subset $V(g)$ (or more

precisely, on $\text{Spec } A/(g)$, and by assumption it vanishes at all points of the closed subset. Now any function vanishing at every point of the spectrum of a ring must be nilpotent (Theorem 3.2.13). In other words, there is some n such that $f^n = 0$ in $A/(g)$, i.e., $f^n \equiv 0 \pmod{g}$ in A , i.e., $f^n \in (g)$.

3.5.F. EASY EXERCISE. Show that $D(f) = \emptyset$ if and only if $f \in \mathfrak{N}$.

3.5.G. UNIMPORTANT EXERCISE (INJECTIVE MAPS OF RINGS INDUCE MAPS OF SPECTRA WITH DENSE IMAGE). Show that if $B \hookrightarrow A$, then the induced map of topological spaces $\text{Spec } A \rightarrow \text{Spec } B$ has dense image.

3.6 Topological (and Noetherian) properties

Many topological notions are useful when applied to the topological space $\text{Spec } A$, and later, to schemes.

3.6.1. Possible topological attributes of $\text{Spec } A$: connectedness, irreducibility, quasicompactness.

3.6.2. Connectedness.

A topological space X is **connected** if it cannot be written as the disjoint union of two nonempty open sets. Exercise 3.6.A below gives an example of a nonconnected $\text{Spec } A$, and the subsequent remark explains that all examples are of this form.

3.6.A. EXERCISE. If $A = A_1 \times A_2 \times \cdots \times A_n$, describe a homeomorphism $\text{Spec } A_1 \amalg \text{Spec } A_2 \amalg \cdots \amalg \text{Spec } A_n \rightarrow \text{Spec } A$ for which each $\text{Spec } A_i$ is mapped onto a distinguished open subset $D(f_i)$ of $\text{Spec } A$. Thus $\text{Spec } \prod_{i=1}^n A_i = \bigsqcup_{i=1}^n \text{Spec } A_i$ as topological spaces. (Hint: reduce to $n = 2$ for convenience. Let $f_1 = (1, 0)$ and $f_2 = (0, 1)$.)

3.6.3. Remark. An extension of Exercise 3.6.A (that you can prove if you wish) is that $\text{Spec } A$ is not connected if and only if A is isomorphic to the product of nonzero rings A_1 and A_2 . The key idea is to show that both conditions are equivalent to there existing nonzero $a_1, a_2 \in A$ for which $a_1^2 = a_1$, $a_2^2 = a_2$, $a_1 + a_2 = 1$, and hence $a_1 a_2 = 0$. An element $a \in A$ satisfying $a^2 = a$ is called an **idempotent**. This will appear as Exercise 10.5.L.

3.6.4. Irreducibility.

A topological space is said to be **irreducible** if it is nonempty, and it is not the union of two proper closed subsets. In other words, a nonempty topological space X is irreducible if whenever $X = Y \cup Z$ with Y and Z closed in X , we have $Y = X$ or $Z = X$. Equivalently (and helpfully): any two nonempty open subsets of X intersect. This is a less useful notion in classical geometry — \mathbb{C}^2 is *reducible* (i.e., not irreducible), but we will see that $\mathbb{A}_{\mathbb{C}}^2$ is irreducible (Exercise 3.6.C).

3.6.B. EASY EXERCISE.

(a) Show that in an irreducible topological space, any nonempty open set is dense. (For this reason, you will see that unlike in the classical topology, in the Zariski

topology, nonempty open sets are all “huge”.)

(b) If X is a topological space, and Z is a subset (with the subspace topology), show that Z is irreducible if and only if the closure \overline{Z} in X is irreducible.

3.6.C. EASY EXERCISE. If A is an integral domain, show that $\text{Spec } A$ is irreducible. (Hint: pay attention to the generic point $[(0)]$.) We will generalize this in Exercise 3.7.G.

3.6.D. EXERCISE. Show that an irreducible topological space is connected.

3.6.E. EXERCISE. Give (with proof!) an example of a ring A where $\text{Spec } A$ is connected but reducible. (Possible hint: a picture may help. The symbol “ \times ” has two “pieces” yet is connected.)

3.6.F. TRICKY EXERCISE.

(a) Suppose $I = (wz - xy, wy - x^2, xz - y^2) \subset k[w, x, y, z]$. Show that $\text{Spec } k[w, x, y, z]/I$ is irreducible, by showing that $k[w, x, y, z]/I$ is an integral domain. (This is hard, so here is one of several possible hints: Show that $k[w, x, y, z]/I$ is isomorphic to the subring of $k[a, b]$ generated by monomials of degree divisible by 3. There are other approaches as well, some of which we will see later. This is an example of a hard question: how do you tell if an ideal is prime?) We will later see this as the cone over the *twisted cubic curve* (the twisted cubic curve is defined in Exercise 9.3.A, and is a special case of a Veronese embedding, §9.3.6).

★ (b) Note that the generators of the ideal of part (a) may be rewritten as the equations ensuring that

$$\text{rank} \begin{pmatrix} w & x & y \\ x & y & z \end{pmatrix} \leq 1,$$

i.e., as the determinants of the 2×2 submatrices. Generalize part (a) to the ideal of rank one $2 \times n$ matrices. This notion will correspond to the cone (§9.3.11) over the *degree n rational normal curve* (Exercise 9.3.I).

3.6.5. Quasicompactness.

A topological space X is **quasicompact** if given any cover $X = \cup_{i \in I} U_i$ by open sets, there is a finite subset S of the index set I such that $X = \cup_{i \in S} U_i$. Informally: every open cover has a finite subcover. We will like this condition, because we are afraid of infinity. Depending on your definition of “compactness”, this is the definition of compactness, minus possibly a Hausdorff condition. However, this isn’t really the algebro-geometric analog of “compact” (we certainly wouldn’t want $\mathbb{A}_{\mathbb{C}}^1$ to be compact) — the right analog is “properness” (§11.5).

3.6.G. EXERCISE.

(a) Show that $\text{Spec } A$ is quasicompact. (Hint: Exercise 3.5.C.)

★ (b) (less important) Show that in general $\text{Spec } A$ can have nonquasicompact open sets. Possible hint: let $A = k[x_1, x_2, x_3, \dots]$ and $\mathfrak{m} = (x_1, x_2, \dots) \subset A$, and consider the complement of $V(\mathfrak{m})$. This example will be useful to construct other “counterexamples” later, e.g., Exercises 8.1.F and 5.1.K. In Exercise 3.6.U, we will see that such weird behavior doesn’t happen for “suitably nice” (Noetherian) rings.

3.6.H. EXERCISE.

(a) If X is a topological space that is a finite union of quasicompact spaces, show

that X is quasicompact.

(b) Show that every closed subset of a quasicompact topological space is quasicompact.

3.6.6. ★★ Fun but irrelevant remark. Exercise 3.6.A shows that $\coprod_{i=1}^n \operatorname{Spec} A_i \cong \operatorname{Spec} \prod_{i=1}^n A_i$, but this *never* holds if “ n is infinite” and all A_i are nonzero, as Spec of any ring is quasicompact (Exercise 3.6.G(a)). This leads to an interesting phenomenon. We show that $\operatorname{Spec} \prod_{i=1}^{\infty} A_i$ is “strictly bigger” than $\coprod_{i=1}^{\infty} \operatorname{Spec} A_i$ where each A_i is isomorphic to the field k . First, we have an inclusion of sets $\coprod_{i=1}^{\infty} \operatorname{Spec} A_i \hookrightarrow \operatorname{Spec} \prod_{i=1}^{\infty} A_i$, as there is a maximal ideal of $\prod A_i$ corresponding to each i (precisely, those elements 0 in the i th component.) But there are other maximal ideals of $\prod A_i$. Hint: describe a proper ideal not contained in any of these maximal ideals. (One idea: consider elements $\prod a_i$ that are “eventually zero”, i.e., $a_i = 0$ for $i \gg 0$.) This leads to the notion of *ultrafilters*, which are very useful, but irrelevant to our current discussion.

3.6.7. Possible topological properties of points of $\operatorname{Spec} A$.

3.6.8. Definition. A point of a topological space $p \in X$ is said to be a **closed point** if $\{p\}$ is a closed subset. In the classical topology on \mathbb{C}^n , all points are closed. In $\operatorname{Spec} \mathbb{Z}$ and $\operatorname{Spec} k[t]$, all the points are closed except for $[(0)]$.

3.6.I. EXERCISE. Show that the closed points of $\operatorname{Spec} A$ correspond to the maximal ideals. (In particular, nonempty affine schemes *have* closed points, as nonzero rings have maximal ideals, §0.3.)

3.6.9. Connection to the classical theory of varieties. Hilbert’s Nullstellensatz lets us interpret the closed points of $\mathbb{A}_{\mathbb{C}}^n$ as the n -tuples of complex numbers. More generally, the closed points of $\operatorname{Spec} \bar{k}[x_1, \dots, x_n]/(f_1, \dots, f_r)$ are naturally interpreted as those points in \bar{k}^n satisfying the equations $f_1 = \dots = f_r = 0$ (see Exercises 3.2.J and 3.2.N(a) for example). Hence from now on we will say “closed point” instead of “traditional point” and “non-closed point” instead of “bonus” point when discussing subsets of $\mathbb{A}_{\bar{k}}^n$. The following exercise is the reason that on algebraic varieties, closed points are dense (cf. Exercise 5.3.F and Definition 11.3.9).

3.6.J. EXERCISE.

(a) Suppose that k is a field, and A is a finitely generated k -algebra. Show that closed points of $\operatorname{Spec} A$ are dense, by showing that if $f \in A$, and $D(f)$ is a nonempty (distinguished) open subset of $\operatorname{Spec} A$, then $D(f)$ contains a closed point of $\operatorname{Spec} A$. Hint: note that A_f is *also* a finitely generated k -algebra. Use the Nullstellensatz 3.2.6 to recognize closed points of Spec of a finitely generated k -algebra B as those for which the residue field is a finite extension of k . Apply this to both $B = A$ and $B = A_f$.

(b) Show that if A is a k -algebra that is not finitely generated, the closed points need not be dense. (Hint: Exercise 3.4.K.)

3.6.K. EXERCISE. Suppose k is an algebraically closed field, and $A = k[x_1, \dots, x_n]/I$ is a finitely generated k -algebra with $\mathfrak{N}(A) = \{0\}$ (so the discussion of §3.2.14 applies). Consider the set $X = \operatorname{Spec} A$ as a subset of \mathbb{A}_k^n . The space \mathbb{A}_k^n contains

the “classical” points k^n . Show that functions on X are determined by their values on the closed points (by the weak Nullstellensatz 3.2.5, the “classical” points $k^n \cap \text{Spec } A$ of $\text{Spec } A$). Hint: if f and g are different functions on X , then $f - g$ is nowhere zero on an open subset of X . Use Exercise 3.6.J(a).

Once we know what a variety is (Definition 11.3.9), this will immediately imply that *a function on a variety over an algebraically closed field is determined by its values on the “classical points”*. (Before the advent of scheme theory, functions on varieties — over algebraically closed fields — were thought of as functions on “classical” points, and Exercise 3.6.K basically shows that there is no harm in thinking of “traditional” varieties as a particular flavor of schemes.)

3.6.10. Specialization and generization. Given two points x, y of a topological space X , we say that x is a **specialization** of y , and y is a **generization** of x , if $x \in \overline{\{y\}}$. This (and Exercise 3.6.L) now makes precise our hand-waving about “one point containing another”. It is of course nonsense for a point to contain another. But it is not nonsense to say that the closure of a point contains another. For example, in $\mathbb{A}_{\mathbb{C}}^2 = \text{Spec } \mathbb{C}[x, y]$, $[(y - x^2)]$ is a generization of $[(x - 2, y - 4)] = (2, 4) \in \mathbb{C}^2$, and $(2, 4)$ is a specialization of $[(y - x^2)]$ (see Figure 3.8).




FIGURE 3.8. $(2, 4) = [(x - 2, y - 4)]$ is a specialization of $[(y - x^2)]$. $[(y - x^2)]$ is a generization of $(2, 4)$.

3.6.L. EXERCISE. If $X = \text{Spec } A$, show that $[q]$ is a specialization of $[p]$ if and only if $p \subset q$. Hence show that $V(p) = \overline{\{[p]\}}$.

3.6.11. Definition. We say that a point $p \in X$ is a **generic point** for a closed subset K if $\overline{\{p\}} = K$.

This important notion predates Grothendieck. The early twentieth-century Italian algebraic geometers had a notion of “generic points” of a variety, by which they meant points with no special properties, so that anything proved of “a generic point” was true of “almost all” the points on that variety. The modern “generic point” has the same intuitive meaning. If something is “generically” or “mostly” true for the points of an irreducible subset, in the sense of being true for a dense open subset (for “almost all points”), then it is true for the generic point, and vice

versa. (This is a statement of principle, not of fact. An interesting case is “reducedness”, for which this principle does not hold in general, but *does* hold for “reasonable” schemes such as varieties, see Remark 5.2.2.) For example, a function has value zero at the generic point of an irreducible scheme if and only if it has the value zero at all points. You should keep an eye out for other examples of this.

The phrase **general point** does not have the same meaning. The phrase “the general point of X satisfies such-and-such a property” means “every point of some dense open subset of X satisfies such-and-such a property”. Be careful not to confuse “general” and “generic”. But be warned that accepted terminology does not always follow this convention; witness “generic freeness” (see Lemma 25.5.12), “generic flatness” (§25.5.11), and “generic smoothness” (§26.3, for example).

3.6.M. EXERCISE. Verify that $[(y - x^2)] \in \mathbb{A}_{\mathbb{C}}^2$ is a generic point for $V(y - x^2)$.

As more motivation for the terminology “generic”: we think of $[(y - x^2)]$ as being some nonspecific point on the parabola (with the closed points $(a, a^2) \in \mathbb{C}^2$, i.e., $(x - a, y - a^2)$ for $a \in \mathbb{C}$, being “specific points”); it is “generic” in the conventional sense of the word. We might “specialize it” to a specific point of the parabola; hence for example $(2, 4)$ is a specialization of $[(y - x^2)]$. (Again, see Figure 3.8.) To make this somewhat more precise:

3.6.N. EASY EXERCISE. Suppose p is a generic point for the closed subset K . Show that it is “near every point q of K ” (every neighborhood of q contains p), and “not near any point r not in K ” (there is a neighborhood of r not containing p). (This idea was mentioned in §3.2.4 Example 1 and in Exercise 3.4.G.)

We will soon see (Exercise 3.7.F) that there is a natural bijection between points of $\text{Spec } A$ and irreducible closed subsets of $\text{Spec } A$, sending each point to its closure, and each irreducible closed subset to its (unique) generic point. You can prove this now, but we will wait until we have developed some convenient terminology.

3.6.12. Irreducible and connected components, and Noetherian conditions.

An **irreducible component** of a topological space is a maximal irreducible subset (an irreducible subset not contained in any larger irreducible subset). Irreducible components are closed (as the closure of irreducible subsets are irreducible, Exercise 3.6.B(b)), and it can be helpful to think of irreducible components of a topological space X as maximal among the irreducible *closed* subsets of X . We think of these as the “pieces of X ” (see Figure 3.9).

Similarly, a subset Y of a topological space X is a **connected component** if it is a maximal connected subset (a connected subset not contained in any larger connected subset).

3.6.O. EXERCISE (EVERY TOPOLOGICAL SPACE IS THE UNION OF IRREDUCIBLE COMPONENTS). Show that every point p of a topological space X is contained in an irreducible component of X . Hint: Zorn’s Lemma. More precisely, consider the partially ordered set \mathcal{S} of irreducible closed subsets of X containing p . Show that there exists a maximal totally ordered subset $\{Z_{\alpha}\}$ of \mathcal{S} . Show that $\bigcup Z_{\alpha}$ is irreducible.

3.6.13. Remark. Every point is contained in a connected component, and connected components are always closed. You can prove this now, but we deliberately




FIGURE 3.9. This closed subset of $\mathbb{A}_{\mathbb{C}}^2$ has six irreducible components

postpone this until we need it, in an optional starred section (Exercise 10.5.J). On the other hand, connected components need not be open, see [Stacks, tag 004T]. An example of an affine scheme with connected components that are not open is $\text{Spec}(\prod_1^\infty \mathbb{F}_2)$.

3.6.14. In the examples we have considered, the spaces have naturally broken up into a finite number of irreducible components. For example, the locus $xy = 0$ in $\mathbb{A}_{\mathbb{C}}^2$ we think of as having two “pieces” — the two axes. The reason for this is that their underlying topological spaces (as we shall soon establish) are *Noetherian*. A topological space X is called **Noetherian** if it satisfies the **descending chain condition** for closed subsets: any sequence $Z_1 \supset Z_2 \supset \cdots \supset Z_n \supset \cdots$ of closed subsets eventually stabilizes: there is an r such that $Z_r = Z_{r+1} = \cdots$. Here is a first example (which you should work out explicitly, not using Noetherian rings).

3.6.P. EXERCISE. Show that $\mathbb{A}_{\mathbb{C}}^2$ is a Noetherian topological space: any decreasing sequence of closed subsets of $\mathbb{A}_{\mathbb{C}}^2 = \text{Spec } \mathbb{C}[x, y]$ must eventually stabilize. Note that it can take arbitrarily long to stabilize. (The closed subsets of $\mathbb{A}_{\mathbb{C}}^2$ were described in §3.4.5.) Show that \mathbb{C}^2 with the classical topology is *not* a Noetherian topological space.

3.6.15. Proposition. — Suppose X is a Noetherian topological space. Then every nonempty closed subset Z can be expressed uniquely as a finite union $Z = Z_1 \cup \cdots \cup Z_n$ of irreducible closed subsets, none contained in any other.

Translation: any closed subset Z has a finite number of “pieces”.

Proof. The following technique is called **Noetherian induction**, for reasons that will be clear. We will use it again, many times.

Consider the collection of nonempty closed subsets of X that *cannot* be expressed as a finite union of irreducible closed subsets. We will show that this collection is empty. Otherwise, let Y_1 be one such. If Y_1 properly contains another

such, then choose one, and call it Y_2 . If Y_2 properly contains another such, then choose one, and call it Y_3 , and so on. By the descending chain condition, this must eventually stop, and we must have some Y_r that cannot be written as a finite union of irreducible closed subsets, but every closed subset properly contained in it can be so written. But then Y_r is not itself irreducible, so we can write $Y_r = Y' \cup Y''$ where Y' and Y'' are both proper closed subsets. Both of these by hypothesis can be written as the union of a finite number of irreducible subsets, and hence so can Y_r , yielding a contradiction. Thus each closed subset can be written as a finite union of irreducible closed subsets. We can assume that none of these irreducible closed subsets contain any others, by discarding some of them.

We now show uniqueness. Suppose

$$Z = Z_1 \cup Z_2 \cup \cdots \cup Z_r = Z'_1 \cup Z'_2 \cup \cdots \cup Z'_s$$

are two such representations. Then $Z'_1 \subset Z_1 \cup Z_2 \cup \cdots \cup Z_r$, so $Z'_1 = (Z_1 \cap Z'_1) \cup \cdots \cup (Z_r \cap Z'_1)$. Now Z'_1 is irreducible, so one of these is Z'_1 itself, say (without loss of generality) $Z_1 \cap Z'_1$. Thus $Z'_1 \subset Z_1$. Similarly, $Z_1 \subset Z'_a$ for some a ; but because $Z'_1 \subset Z_1 \subset Z'_a$, and Z'_1 is contained in no other Z'_i , we must have $a = 1$, and $Z'_1 = Z_1$. Thus each element of the list of Z 's is in the list of Z' 's, and vice versa, so they must be the same list. \square

3.6.Q. EXERCISE. Show that every connected component of a topological space X is the union of irreducible components of X . Show that any subset of X that is simultaneously open and closed must be the union of some of the connected components of X . If X is a *Noetherian* topological space, show that the union of any subset of the connected components of X is always open and closed in X . (In particular, connected components of Noetherian topological spaces are always open, which is not true for more general topological spaces, see Remark 3.6.13.)

3.6.16. Noetherian rings. It turns out that all of the spectra we have considered (except in starred Exercise 3.6.G(b)) are Noetherian topological spaces, but that isn't true of the spectra of all rings. The key characteristic all of our examples have had in common is that the rings were *Noetherian*. A ring is **Noetherian** if every ascending sequence $I_1 \subset I_2 \subset \cdots$ of ideals eventually stabilizes: there is an r such that $I_r = I_{r+1} = \cdots$. (This is called the **ascending chain condition** on ideals.)

Here are some quick facts about Noetherian rings. You should be able to prove them all.

- Fields are Noetherian. \mathbb{Z} is Noetherian.
- (*Noetherianness is preserved by quotients.*) If A is Noetherian, and I is any ideal of A , then A/I is Noetherian.
- (*Noetherianness is preserved by localization.*) If A is Noetherian, and S is any multiplicative set, then $S^{-1}A$ is Noetherian.

3.6.R. EXERCISE. Show that principal ideal domains are Noetherian rings.

3.6.S. IMPORTANT EXERCISE. Show that a ring A is Noetherian if and only if every ideal of A is finitely generated.

The next fact is nontrivial.

3.6.17. The Hilbert Basis Theorem. — If A is Noetherian, then so is $A[x]$.

Hilbert proved this in the epochal paper [Hil] where he also proved the Hilbert Syzygy Theorem (§16.3.2), and defined Hilbert functions and showed that they are eventually polynomial (§19.6).

3.6.18. By the results described above, any polynomial ring over any field, or over the integers, is Noetherian — and also any quotient or localization thereof. Hence for example any finitely generated algebra over k or \mathbb{Z} , or any localization thereof, is Noetherian. Most “nice” rings are Noetherian, but not all rings are Noetherian: $k[x_1, x_2, \dots]$ is not, because $(x_1) \subset (x_1, x_2) \subset (x_1, x_2, x_3) \subset \dots$ is a strictly ascending chain of ideals (cf. Exercise 3.6.G(b)).

3.6.19. *Proof of the Hilbert Basis Theorem 3.6.17.* We show that any ideal $I \subset A[x]$ is finitely generated. We inductively produce a set of generators f_1, \dots as follows. For $n \geq 0$, if $I \neq (f_1, \dots, f_n)$, let f_{n+1} be any nonzero element of $I - (f_1, \dots, f_n)$ of lowest degree. Thus f_1 is any element of I of lowest degree, assuming $I \neq (0)$. If this procedure terminates, we are done. Otherwise, let $a_n \in A$ be the initial coefficient of f_n for all $n > 0$. As A is Noetherian, $(a_1, a_2, \dots) = (a_1, \dots, a_N)$ for some N . Say $a_{N+1} = \sum_{i=1}^N b_i a_i$. Then

$$f_{N+1} - \sum_{i=1}^N b_i f_i x^{\deg f_{N+1} - \deg f_i}$$

is an element of I that is not in (f_1, \dots, f_N) (as $f_{N+1} \notin (f_1, \dots, f_N)$), and of lower degree than f_{N+1} , contradicting how we were supposed to have chosen f_{N+1} . \square

We now connect Noetherian rings and Noetherian topological spaces.

3.6.T. EXERCISE. If A is Noetherian, show that $\text{Spec } A$ is a Noetherian topological space. Describe a ring A such that $\text{Spec } A$ is not a Noetherian topological space. (Aside: if $\text{Spec } A$ is a Noetherian topological space, A need not be Noetherian. One example is $A = k[x_1, x_2, x_3, \dots]/(x_1, x_2^2, x_3^3, \dots)$. Then $\text{Spec } A$ has one point, so is Noetherian. But A is not Noetherian as $(x_1) \subsetneq (x_1, x_2) \subsetneq (x_1, x_2, x_3) \subsetneq \dots$ in A .)

3.6.U. EXERCISE (PROMISED IN EXERCISE 3.6.G(b)). Show that every open subset of a Noetherian topological space is quasicompact. Hence if A is Noetherian, every open subset of $\text{Spec } A$ is quasicompact. (If you prefer, show the result for *any* subset, with the induced topology.)

3.6.20. For future use: Noetherian conditions for modules. An important related notion is that of a Noetherian *module*. Although we won’t use this notion for some time (§10.7.3), we will develop their most important properties, while Noetherian ideas are still fresh in your mind. If A is any ring, not necessarily Noetherian, we say an A -**module is Noetherian** if it satisfies the ascending chain condition for submodules. Thus for example a ring A is Noetherian if and only if it is a Noetherian A -module.

3.6.V. EXERCISE (CF. IMPORTANT EXERCISE 3.6.S). Show that if M is a Noetherian A -module, then any submodule of M is a finitely generated A -module. (In fact the converse is true, and you are welcome to prove that too.)

3.6.W. EXERCISE. If $0 \rightarrow M' \rightarrow M \rightarrow M'' \rightarrow 0$ is exact, show that M' and M'' are Noetherian if and only if M is Noetherian. (Hint: Given an ascending chain in M , we get two simultaneous ascending chains in M' and M'' . Possible further hint: prove that if

$$M' \longrightarrow M \xrightarrow{\phi} M''$$

is exact, and $N \subset N' \subset M$, and $N \cap M' = N' \cap M'$ and $\phi(N) = \phi(N')$, then $N = N'$.)

3.6.X. EXERCISE. Show that if A is a Noetherian ring, then $A^{\oplus n}$ is a Noetherian A -module.

3.6.Y. EXERCISE. Show that if A is a Noetherian ring and M is a finitely generated A -module, then M is a Noetherian module. Hence by Exercise 3.6.V, any submodule of a finitely generated module over a Noetherian ring is finitely generated.

3.6.21. Why you should not worry about Noetherian hypotheses. Should you work hard to eliminate Noetherian hypotheses? Should you worry about Noetherian hypotheses? Should you stay up at night thinking about non-Noetherian rings? For the most part, the answer to these questions is “no”. Most people will never need to worry about non-Noetherian rings, but there are reasons to be open to them. First, they can actually come up. For example, fibered products of Noetherian schemes over Noetherian schemes (and even fibered products of Noetherian points over Noetherian points!) can be non-Noetherian (Warning 10.1.5), and the normalization of Noetherian rings can be non-Noetherian (Warning 10.7.4). You can either work hard to show that the rings or schemes you care about don’t have this pathology, or you can just relax and not worry about it. Second, there is often no harm in working with schemes in general. Knowing when Noetherian conditions are needed will help you remember why results are true, because you will have some sense of where Noetherian conditions enter into arguments. Finally, for some people, non-Noetherian rings naturally come up. For example, adèles are not Noetherian. And many valuation rings that naturally arise in arithmetic and tropical geometry are not Noetherian.

3.7 The function $I(\cdot)$, taking subsets of $\text{Spec } A$ to ideals of A

We now introduce a notion that is in some sense “inverse” to the vanishing set function $V(\cdot)$. Given a subset $S \subset \text{Spec } A$, $I(S)$ is the set of functions vanishing on S . In other words, $I(S) = \bigcap_{\mathfrak{p} \in S} \mathfrak{p} \subset A$ (at least when S is nonempty).

We make three quick observations. (Do you see why they are true?)

- $I(S)$ is clearly an ideal of A .
- $I(\cdot)$ is inclusion-reversing: if $S_1 \subset S_2$, then $I(S_2) \subset I(S_1)$.
- $I(\bar{S}) = I(S)$.

3.7.A. EXERCISE. Let $A = k[x, y]$. If $S = \{[(y)], [(x, y - 1)]]\}$ (see Figure 3.10), then $I(S)$ consists of those polynomials vanishing on the x -axis, and at the point $(0, 1)$. Give generators for this ideal.




FIGURE 3.10. The set S of Exercise 3.7.A, pictured as a subset of \mathbb{A}^2

3.7.B. EXERCISE. Suppose $S \subset \mathbb{A}_{\mathbb{C}}^3$ is the union of the three axes. Give generators for the ideal $I(S)$. Be sure to prove it! We will see in Exercise 13.1.F that this ideal is not generated by less than three elements.

3.7.C. EXERCISE. Show that $V(I(S)) = \bar{S}$. Hence $V(I(S)) = S$ for a closed set S .

Note that $I(S)$ is always a radical ideal — if $f \in \sqrt{I(S)}$, then f^n vanishes on S for some $n > 0$, so then f vanishes on S , so $f \in I(S)$.

3.7.D. EASY EXERCISE. Prove that if $J \subset A$ is an ideal, then $I(V(J)) = \sqrt{J}$. (Huge hint: Exercise 3.4.J.)

Exercises 3.7.C and 3.7.D show that V and I are “almost” inverse. More precisely:

3.7.1. Theorem. — $V(\cdot)$ and $I(\cdot)$ give an inclusion-reversing bijection between closed subsets of $\text{Spec } A$ and radical ideals of A (where a closed subset gives a radical ideal by $I(\cdot)$, and a radical ideal gives a closed subset by $V(\cdot)$).

Theorem 3.7.1 is sometimes called Hilbert’s Nullstellensatz, but we reserve that name for Theorem 3.2.6.

3.7.E. EXERCISE. Let $J = (x^2 + y^2 - 1, y - 1)$. Find, with proof, an element of $I(V(J)) \setminus J$.

3.7.F. IMPORTANT EXERCISE (CF. EXERCISE 3.7.G). Show that $V(\cdot)$ and $I(\cdot)$ give a bijection between *irreducible closed subsets* of $\text{Spec } A$ and *prime ideals* of A . From this conclude that in $\text{Spec } A$ there is a bijection between points of $\text{Spec } A$ and irreducible closed subsets of $\text{Spec } A$ (where a point determines an irreducible closed subset by taking the closure). Hence *each irreducible closed subset of $\text{Spec } A$ has precisely one generic point* — any irreducible closed subset Z can be written uniquely as $\overline{\{z\}}$.

3.7.G. EXERCISE/DEFINITION. A prime ideal of a ring A is a **minimal prime ideal** (or more simply, **minimal prime**) if it is minimal with respect to inclusion. (For example, the only minimal prime of $k[x, y]$ is (0) .) If A is any ring, show that the irreducible components of $\text{Spec } A$ are in bijection with the minimal prime ideals of A . In particular, $\text{Spec } A$ is irreducible if and only if A has only one minimal prime ideal; this generalizes Exercise 3.6.C.

Proposition 3.6.15, Exercise 3.6.T, and Exercise 3.7.G imply that every Noetherian ring has a finite number of minimal prime ideals. An algebraic fact is now revealed to be really a “geometric” fact!

3.7.H. EXERCISE (REALITY CHECK). In $\mathbb{A}^n = \operatorname{Spec} k[x_1, \dots, x_n]$, the subset cut out by $f(x_1, \dots, x_n) \in k[x_1, \dots, x_n]$ should certainly have irreducible components corresponding to the distinct irreducible factors of f . Prove this. (What property of the ring $k[x_1, \dots, x_n]$ makes this work?)

3.7.I. EXERCISE. What are the minimal prime ideals of $k[x, y]/(xy)$ (where k is a field)?

3.7.2. Beginning of a grand dictionary between algebra and geometry. We are now well on our way to building a grand dictionary between algebra and geometry. You should add more entries to the table as your understanding deepens.

algebra	geometry	reference
ring A	affine scheme $\operatorname{Spec} A$	§3.2.1
$\mathfrak{p} = I(\mathfrak{p})$ prime ideal of A	$\mathfrak{p} = [\mathfrak{p}]$ point of $\operatorname{Spec} A$	§3.2.1, §3.7
element $f \in A$	function f on $\operatorname{Spec} A$	§3.2.1
$f \pmod{\mathfrak{p}}$	$f(\mathfrak{p})$; value of the function f at \mathfrak{p}	§3.2.1
$f \in \mathfrak{p}$	the function f vanishes (is zero) at \mathfrak{p}	§3.2.1
nilradical ideal $\mathfrak{N} = \sqrt{(0)} = \cap \mathfrak{p}$ of A	functions vanishing at every point	Thm. 3.2.13
maximal ideal of A	irreducible closed subset of $\operatorname{Spec} A$ that is a point (= closed point of $\operatorname{Spec} A$)	Ex. 3.6.I
$\mathfrak{p} = I(Z)$ prime ideal of A (a repeat!)	irreducible closed subset $Z = V(\mathfrak{p})$ of $\operatorname{Spec} A$	Ex. 3.7.F
minimal prime ideal of A	irreducible component of $\operatorname{Spec} A$	§3.7.G
radical ideal $I = \sqrt{I} = I(S)$ of A	closed subset $S = V(I)$ of $\operatorname{Spec} A$	Thm. 3.7.1

The structure sheaf, and the definition of schemes in general

4.1 The structure sheaf of an affine scheme

The final ingredient in the definition of an affine scheme is the *structure sheaf* $\mathcal{O}_{\text{Spec } A}$, which we think of as the “sheaf of algebraic functions”. You should keep in your mind the example of “algebraic functions” on \mathbb{C}^n , which you understand well. For example, in \mathbb{A}^2 , we expect that on the open set $D(xy)$ (away from the two axes), $(3x^4 + y + 4)/x^7y^3$ should be an algebraic function.

These functions will have values at points, but won’t be determined by their values at points. But like all sections of sheaves, they will be determined by their germs (see §4.3.6).

It suffices to describe the structure sheaf as a sheaf (of rings) on the base of distinguished open sets (Theorem 2.5.1 and Exercise 3.5.A).

4.1.1. Definition. Define $\mathcal{O}_{\text{Spec } A}(D(f))$ to be the localization of A at the multiplicative set of all functions that do not vanish outside of $V(f)$ (i.e., those $g \in A$ such that $V(g) \subset V(f)$, or equivalently $D(f) \subset D(g)$, cf. Exercise 3.5.E). This depends only on $D(f)$, and not on f itself. (Scholars of the empty set might notice that by Exercise 3.5.F, we have that $\mathcal{O}_{\text{Spec } A}(\emptyset) = \{0\}$.)

4.1.A. GREAT EXERCISE. Show that the natural map $A_f \rightarrow \mathcal{O}_{\text{Spec } A}(D(f))$ is an isomorphism. (Possible hint: Exercise 3.5.E.)

If $D(f') \subset D(f)$, define the restriction map

$$\text{res}_{D(f), D(f')}: \mathcal{O}_{\text{Spec } A}(D(f)) \longrightarrow \mathcal{O}_{\text{Spec } A}(D(f'))$$

in the obvious way: the latter ring is a further localization of the former ring. The restriction maps obviously compose: this is a “presheaf on the distinguished base”.

4.1.2. Theorem. — *The data just described give a sheaf on the distinguished base, and hence determine a sheaf on the topological space $\text{Spec } A$.*

4.1.3. This sheaf is called the **structure sheaf**, and will be denoted $\mathcal{O}_{\text{Spec } A}$, or sometimes \mathcal{O} if the subscript is clear from the context. The notation $\text{Spec } A$ will hereafter denote the data of a topological space with a structure sheaf. Such a topological space, with structure sheaf, will be called an **affine scheme** (Definition 4.3.1). We

continue to use the language of ringed spaces (§2.2.13): functions on open sets, global functions, and so forth. An important lesson of Theorem 4.1.2 is not just that $\mathcal{O}_{\text{Spec } A}$ is a sheaf, but also that the distinguished base provides a good way of working with $\mathcal{O}_{\text{Spec } A}$. Notice also that we have justified interpreting elements of A as functions on $\text{Spec } A$.

Proof. We must show the base identity and base gluability axioms hold (§2.5). We show that they both hold for the open set that is the entire space $\text{Spec } A$, and leave to you the trick which extends them to arbitrary distinguished open sets (Exercises 4.1.B and 4.1.C). Suppose $\text{Spec } A = \cup_{i \in I} D(f_i)$, or equivalently (Exercise 3.5.B) the ideal generated by the f_i is the entire ring A .

(Aside: experts familiar with the equalizer exact sequence of §2.2.7 will realize that we are showing exactness of

$$(4.1.3.1) \quad 0 \longrightarrow A \longrightarrow \prod_{i \in I} A_{f_i} \longrightarrow \prod_{i \neq j \in I} A_{f_i f_j}$$

where $\{f_i\}_{i \in I}$ is a set of functions with $(f_i)_{i \in I} = A$. Signs are involved in the right-hand map: the map $A_{f_i} \rightarrow A_{f_i f_j}$ is the localization map, and the map $A_{f_j} \rightarrow A_{f_i f_j}$ is the *negative* of the localization map. Base identity corresponds to injectivity at A , and gluability corresponds to exactness at $\prod_i A_{f_i}$.)

We check identity on the base. Suppose that $\text{Spec } A = \cup_{i \in I} D(f_i)$ where i runs over some index set I . Then there is some finite subset of I , which we name $\{1, \dots, n\}$, such that $\text{Spec } A = \cup_{i=1}^n D(f_i)$, i.e., $(f_1, \dots, f_n) = A$ (quasicompactness of $\text{Spec } A$, Exercise 3.5.C). Suppose we are given $s \in A$ such that $\text{res}_{\text{Spec } A, D(f_i)} s = 0$ in A_{f_i} for all i . We wish to show that $s = 0$. The fact that $\text{res}_{\text{Spec } A, D(f_i)} s = 0$ in A_{f_i} implies that there is some m such that for each $i \in \{1, \dots, n\}$, $f_i^m s = 0$. Now $(f_1^m, \dots, f_n^m) = A$ (for example, from $\text{Spec } A = \cup D(f_i) = \cup D(f_i^m)$), so there are $r_i \in A$ with $\sum_{i=1}^n r_i f_i^m = 1$ in A , from which

$$s = \left(\sum r_i f_i^m \right) s = \sum r_i (f_i^m s) = 0.$$

Thus we have checked the “base identity” axiom for $\text{Spec } A$.

4.1.B. EXERCISE. Make tiny changes to the above argument to show base identity for any distinguished open $D(f)$. (Hint: judiciously replace A by A_f in the above argument.)

We next show base gluability. (Serre has described this as a “partition of unity” argument, and if you look at it in the right way, his insight is very enlightening.) Suppose again $\cup_{i \in I} D(f_i) = \text{Spec } A$, where I is an index set (possibly horribly infinite). Suppose we are given elements in each A_{f_i} that agree on the overlaps $A_{f_i f_j}$. Note that intersections of distinguished open sets are also distinguished open sets.

Assume first that I is finite, say $I = \{1, \dots, n\}$. We have elements $a_i/f_i^{l_i} \in A_{f_i}$ agreeing on overlaps $A_{f_i f_j}$ (see Figure 4.1(a)). Letting $g_i = f_i^{l_i}$, using $D(f_i) = D(g_i)$, we can simplify notation by considering our elements as of the form $a_i/g_i \in A_{g_i}$ (Figure 4.1(b)).

The fact that a_i/g_i and a_j/g_j “agree on the overlap” (i.e., in $A_{g_i g_j}$) means that for some m_{ij} ,

$$(g_i g_j)^{m_{ij}} (g_j a_i - g_i a_j) = 0$$




FIGURE 4.1. Base gluability of the structure sheaf

in A . By taking $m = \max m_{ij}$ (here we use the finiteness of I), we can simplify notation:

$$(g_i g_j)^m (g_j a_i - g_i a_j) = 0$$

for all i, j (Figure 4.1(c)). Let $b_i = a_i g_i^m$ for all i , and $h_i = g_i^{m+1}$ (so $D(h_i) = D(g_i)$). Then we can simplify notation even more (Figure 4.1(d)): on each $D(h_i)$, we have a function b_i/h_i , and the overlap condition is

$$(4.1.3.2) \quad h_j b_i = h_i b_j.$$

Now $\cup_i D(h_i) = \text{Spec } A$, implying that $1 = \sum_{i=1}^n r_i h_i$ for some $r_i \in A$. Define

$$(4.1.3.3) \quad r = \sum r_i b_i.$$

This will be the element of A that restricts to each b_j/h_j . Indeed, from the overlap condition (4.1.3.2),

$$r h_j = \sum_i r_i b_i h_j = \sum_i r_i h_i b_j = b_j.$$

We next deal with the case where I is infinite. Choose a finite subset $\{1, \dots, n\} \subset I$ with $(f_1, \dots, f_n) = A$ (or equivalently, use quasicompactness of $\text{Spec } A$ to choose a finite subcover by $D(f_i)$). Construct r as above, using (4.1.3.3). We will show that for any $z \in I - \{1, \dots, n\}$, r restricts to the desired element $a_z/f_z^{l_z}$ of A_{f_z} . Repeat the entire process above with $\{1, \dots, n, z\}$ in place of $\{1, \dots, n\}$, to obtain $r' \in A$ which restricts to a_i for $i \in \{1, \dots, n, z\}$. Then by the base identity axiom, $r' = r$. (Note that we use base identity to *prove* base gluability. This is an example of how the identity axiom is somehow “prior” to the gluability axiom.) Hence r restricts to the desired element $a_z/f_z^{l_z}$ of A_{f_z} .

4.1.C. EXERCISE. Alter this argument appropriately to show base gluability for any distinguished open $D(f)$.

We have now completed the proof of Theorem 4.1.2. \square

4.1.4. ★★ Remark. Definition 4.1.1 and Theorem 4.1.2 suggests a potentially slick way of describing sections of $\mathcal{O}_{\text{Spec } A}$ over *any* open subset: perhaps $\mathcal{O}_{\text{Spec } A}(U)$ is the localization of A at the multiplicative set of all functions that do not vanish at

any point of U . This is not true. A counterexample (that you will later be able to make precise): let $\text{Spec } A$ be two copies of \mathbb{A}_k^2 glued together at their origins (see (25.4.8.1) for explicit equations) and let U be the complement of the origin(s). Then the function which is 1 on the first copy of $\mathbb{A}_k^2 \setminus \{(0, 0)\}$ and 0 on the second copy of $\mathbb{A}_k^2 \setminus \{(0, 0)\}$ is not of this form. (Follow-up question: why would this discussion not work for two copies of \mathbb{A}_k^1 glued at a point?)

4.1.5. Important notion for future use: $\mathcal{O}_{\text{Spec } A}$ -modules coming from A -modules. The following generalization of Theorem 4.1.2 will be essential in the definition of a *quasicoherent sheaf* in Chapter 6. You can leave these exercises for then, but they may be easiest to solve now, while the ideas are fresh in your mind.

4.1.D. IMPORTANT EXERCISE/DEFINITION. Suppose M is an A -module. Show that the following construction describes a sheaf \widetilde{M} on the distinguished base. Define $\widetilde{M}(D(f))$ to be the localization of M at the multiplicative set of all functions that do not vanish outside of $V(f)$. Define restriction maps $\text{res}_{D(f), D(g)}$ in the analogous way to $\mathcal{O}_{\text{Spec } A}$. Show that this defines a sheaf on the distinguished base, and hence a sheaf on $\text{Spec } A$. Then show that this is an $\mathcal{O}_{\text{Spec } A}$ -module.

4.1.6. Important Remark. In the course of answering the previous exercise, you will show that if $(f_i)_{i \in I} = A$,

$$(4.1.6.1) \quad 0 \longrightarrow M \longrightarrow \prod_{i \in I} M_{f_i} \longrightarrow \prod_{i \neq j \in I} M_{f_i f_j}$$

(cf. (4.1.3.1)) is exact. In particular, M can be identified with a specific submodule of $M_{f_1} \times \cdots \times M_{f_r}$. Even though $M \rightarrow M_{f_i}$ may not be an inclusion for any f_i , $M \rightarrow M_{f_1} \times \cdots \times M_{f_r}$ is an inclusion. This will be useful later: we will want to show that if M has some nice property, then M_f does too, which will be easy. We will also want to show that if $(f_1, \dots, f_n) = A$, and the M_{f_i} have this property, then M does too. This idea will be made precise in the Affine Communication Lemma 5.3.2.

4.1.E. EXERCISE. Suppose \mathfrak{p} is a prime ideal of A . Describe a canonical isomorphism $\widetilde{M}_{[\mathfrak{p}]} \xrightarrow{\sim} M_{\mathfrak{p}}$ between the stalk of \widetilde{M} and the localization of M .

The following exercise shows how an important result can be understood quite differently from algebraic and geometric perspectives.

4.1.F. EXERCISE. Suppose M is an A -module. Show that the natural map

$$M \hookrightarrow \prod_{\mathfrak{p} \in \text{Spec } A} M_{\mathfrak{p}}$$

is an injection in two ways: (a) by considering the kernel of the map, and show that any element of it must be 0, and (b) by applying Exercise 2.4.A (sections of a sheaf are determined by germs) to the sheaf \widetilde{M} . (Thus an A -module is zero if and only if all its localizations at prime ideals are zero, and we see this as a simultaneously (a) algebraic and (b) geometric fact.)

4.1.G. EXERCISE. Describe a bijection between maps of A -modules $M \rightarrow N$, and maps of $\mathcal{O}_{\text{Spec } A}$ -modules $\widetilde{M} \rightarrow \widetilde{N}$. (Fancy translation: you will have described Mod_A as a full subcategory of $\text{Mod}_{\mathcal{O}_{\text{Spec } A}}$.)

4.1.7. Definition: Support of a module. Motivated by Exercise 4.1.E, and the notion of support of a sheaf (Definition 2.7.6), define the **support** of $m \in M$ by

$$\text{Supp } m := \{[p] \in \text{Spec } A : m_p \neq 0\} \subset \text{Spec } A,$$

which is the support of m considered as a section of \widetilde{M} . By Exercise 2.7.G, $\text{Supp } m$ is a closed subset of $\text{Spec } A$.

Similarly (also following Definition 2.7.6), define the **support** of M by

$$\text{Supp } M := \text{Supp } \widetilde{M} = \{[p] \in \text{Spec } A : M_p \neq 0\} \subset \text{Spec } A.$$

These notions will come up repeatedly. We will discuss support in more detail in §6.5.3.

4.1.8. Recurring (counter)examples. The example of two planes meeting at a point (§4.1.4) will appear many times as an example or a counterexample. It is important to have such examples (or counterexamples) at the back of your mind, as “boundary markers” reminding you of what is reasonable, and what to watch out for when you hear a new statement. The appearances of the following are listed in the index (under “recurring (counter)examples”), so you can see how often they come up, and where. You do not have the language to understand many of these yet, but you can come back to this list later.

- two planes meeting at a point
- affine space minus the origin, and its inclusion in affine space
- affine space with doubled origin
- (the cone over) the quadric surface, $\text{Spec } A$ and $\text{Proj } A$ where $A = k[w, x, y, z]/(wz - xy)$
- an embedded point on a line, $k[x, y]/(y^2, xy)$
- $x^2 = 0$ in the projective plane
- infinite disjoint unions of schemes (especially $\coprod \text{Spec } k[x]/(x^n)$)
- the morphism $\text{Spec } \overline{\mathbb{Q}} \rightarrow \text{Spec } \mathbb{Q}$
- $k(x) \otimes_k k(y)$

4.2 Visualizing schemes: Nilpotents

The price of metaphor is eternal vigilance.

— A. Rosenbluth and N. Wiener (attribution by [Lew, p. 4])

In §3.3, we discussed how to visualize the underlying set of schemes, adding in generic points to our previous intuition of “classical” (or closed) points. Our later discussion of the Zariski topology fit well with that picture. In our definition of the “affine scheme” $(\text{Spec } A, \mathcal{O}_{\text{Spec } A})$, we have the additional information of nilpotents, which are invisible on the level of points (§3.2.12), so now we figure out to picture them. We will then readily be able to glue them together to picture schemes in general, once we have made the appropriate definitions. As we are building intuition, we cannot be rigorous or precise.

As motivation, note that we have incidence-reversing bijections
maximal ideals of $A \longleftrightarrow$ closed points of $\text{Spec } A$ (Exercise 3.6.I)

prime ideals of $A \longleftrightarrow$ irreducible closed subsets of $\text{Spec } A$ (Exercise 3.7.F)

radical ideals of $A \longleftrightarrow$ closed subsets of $\text{Spec } A$ (Theorem 3.7.1)

If we take the things on the right as “pictures”, our goal is to figure out how to picture ideals that are not radical:

ideals of $A \longleftrightarrow$???

(We will later fill this in rigorously in a different way with the notion of a *closed subscheme*, the scheme-theoretic version of closed subsets, §9.1. But our goal now is to create a picture.)

As motivation, when we see the expression, $\text{Spec } \mathbb{C}[x]/(x(x-1)(x-2))$, we immediately interpret it as a closed subset of $\mathbb{A}_{\mathbb{C}}^1$, namely $\{0, 1, 2\}$. In particular, the map $\mathbb{C}[x] \rightarrow \mathbb{C}[x]/(x(x-1)(x-2))$ can be interpreted (via the Chinese Remainder Theorem) as: take a function on \mathbb{A}^1 , and restrict it to the three points 0, 1, and 2.

This will guide us in how to visualize a nonradical ideal. The simplest example to consider is $\text{Spec } \mathbb{C}[x]/(x^2)$ (Exercise 3.2.A(a)). As a subset of \mathbb{A}^1 , it is just the origin $0 = [(x)]$, which we are used to thinking of as $\text{Spec } \mathbb{C}[x]/(x)$ (i.e., corresponding to the ideal (x) , not (x^2)). We want to enrich this picture in some way. We should picture $\mathbb{C}[x]/(x^2)$ in terms of the information the quotient remembers. The image of a polynomial $f(x)$ is the information of its value at 0, and its derivative (cf. Exercise 3.2.T). We thus picture this as being the point, plus a little bit more — a little bit of infinitesimal “fuzz” on the point (see Figure 4.2). The sequence of restrictions $\mathbb{C}[x] \rightarrow \mathbb{C}[x]/(x^2) \rightarrow \mathbb{C}[x]/(x)$ should be interpreted as nested pictures.

$$\mathbb{C}[x] \twoheadrightarrow \mathbb{C}[x]/(x^2) \twoheadrightarrow \mathbb{C}[x]/(x)$$

$$f(x) \mapsto f(0),$$

Similarly, $\mathbb{C}[x]/(x^3)$ remembers even more information — the second derivative as well. Thus we picture this as the point 0 with even more fuzz.

More subtleties arise in two dimensions (see Figure 4.3). Consider


$$\text{Spec } \mathbb{C}[x, y]/(x, y)^2,$$

which is sandwiched between two rings we know well:


$$\mathbb{C}[x, y] \twoheadrightarrow \mathbb{C}[x, y]/(x, y)^2 \twoheadrightarrow \mathbb{C}[x, y]/(x, y)$$

$$f(x, y) \mapsto f(0, 0).$$

Again, taking the quotient by $(x, y)^2$ remembers the first derivative, “in all directions”. We picture this as fuzz around the point, in the shape of a circle (no

FIGURE 4.2. Picturing quotients of $\mathbb{C}[x]$

direction is privileged). Similarly, $(x, y)^3$ remembers the second derivative “in all directions” — bigger circular fuzz.

FIGURE 4.3. Picturing quotients of $\mathbb{C}[x, y]$

Consider instead the ideal (x^2, y) . What it remembers is the derivative only in the x direction — given a polynomial, we remember its value at 0, and the coefficient of x . We remember this by picturing the fuzz only in the x direction.

This gives us some handle on picturing more things of this sort, but now it becomes more an art than a science. For example, $\text{Spec } \mathbb{C}[x, y]/(x^2, y^2)$ we might picture as a fuzzy square around the origin. (Could you believe that this square is circumscribed by the circular fuzz $\text{Spec } \mathbb{C}[x, y]/(x, y)^3$, and inscribed by the circular fuzz $\text{Spec } \mathbb{C}[x, y]/(x, y)^2$?) One feature of this example is that given two ideals I and J of a ring A (such as $\mathbb{C}[x, y]$), your fuzzy picture of $\text{Spec } A/(I, J)$ should be the “intersection” of your picture of $\text{Spec } A/I$ and $\text{Spec } A/J$ in $\text{Spec } A$. (You will make this precise in Exercise 9.1.I(a).) For example, $\text{Spec } \mathbb{C}[x, y]/(x^2, y^2)$ should be the intersection of two thickened lines. (How would you picture $\text{Spec } \mathbb{C}[x, y]/(x^5, y^3)$? $\text{Spec } \mathbb{C}[x, y, z]/(x^3, y^4, z^5, (x + y + z)^2)$? $\text{Spec } \mathbb{C}[x, y]/((x, y)^5, y^3)$?)

One final example that will motivate us in §6.5 is $\text{Spec } \mathbb{C}[x, y]/(y^2, xy)$. Knowing what a polynomial in $\mathbb{C}[x, y]$ is modulo (y^2, xy) is the same as knowing its value on the x -axis, as well as first-order differential information around the origin. This is worth thinking through carefully: do you see how this information is captured (however imperfectly) in Figure 4.4?



FIGURE 4.4. A picture of the scheme $\text{Spec } k[x, y]/(y^2, xy)$. The fuzz at the origin indicates where “there is nonreducedness”.

Our pictures capture useful information that you already have some intuition for. For example, consider the intersection of the parabola $y = x^2$ and the x -axis (in the xy -plane), see Figure 4.5. You already have a sense that the intersection has multiplicity two. In terms of this visualization, we interpret this as intersecting (in $\text{Spec } \mathbb{C}[x, y]$):

$$\begin{aligned} \text{Spec } \mathbb{C}[x, y]/(y - x^2) \cap \text{Spec } \mathbb{C}[x, y]/(y) &= \text{Spec } \mathbb{C}[x, y]/(y - x^2, y) \\ &= \text{Spec } \mathbb{C}[x, y]/(y, x^2) \end{aligned}$$

which we interpret as the fact that the parabola and line not just meet with multiplicity two, but that the “multiplicity 2” part is in the direction of the x -axis. You will make this example precise in Exercise 9.1.I(b).




FIGURE 4.5. The “scheme-theoretic” intersection of the parabola $y = x^2$ and the x -axis is a nonreduced scheme (with fuzz in the x -direction)

4.2.1. We will later make the location of the fuzz more precise when we discuss associated points. We will see that in reasonable circumstances, the fuzz is concentrated on closed subsets (Exercise 6.5.N).

On a bien souvent répété que la Géométrie est l’art de bien raisonner sur des figures mal faites; encore ces figures, pour ne pas nous tromper, doivent-elles satisfaire à certaines conditions; les proportions peuvent être grossièrement altérées, mais les positions relatives des diverses parties ne doivent pas être bouleversées.

It is often said that geometry is the art of reasoning well from badly made figures; however, these figures, if they are not to deceive us, must satisfy certain conditions; the proportions may be grossly altered, but the relative positions of the different parts must not be upset.

— H. Poincaré, [Po1, p. 2] (see J. Stillwell’s translation [Po2, p. ix])

4.3 Definition of schemes

4.3.1. Definitions. We can now define *scheme* in general. First, define an **isomorphism of ringed spaces** $(X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ as (i) a homeomorphism $\pi: X \rightarrow Y$, and (ii) an isomorphism of sheaves \mathcal{O}_X and \mathcal{O}_Y , considered to be on the same space via π . (Part (ii), more precisely, is an isomorphism $\mathcal{O}_Y \xrightarrow{\sim} \pi_* \mathcal{O}_X$ of sheaves on Y , or equivalently, an isomorphism $\pi^{-1} \mathcal{O}_Y \xrightarrow{\sim} \mathcal{O}_X$ of sheaves on X .) In other words, we have a “correspondence” of sets, topologies, and structure sheaves. Every isomorphism $\alpha: (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ clearly has an “inverse isomorphism” $\alpha^{-1}: (Y, \mathcal{O}_Y) \rightarrow (X, \mathcal{O}_X)$.

An **affine scheme** is a ringed space that is isomorphic to $(\text{Spec } A, \mathcal{O}_{\text{Spec } A})$ for some A . A **scheme** (X, \mathcal{O}_X) is a ringed space such that any point of X has an open neighborhood U such that $(U, \mathcal{O}_X|_U)$ is an affine scheme. The topology on a scheme is called the **Zariski topology**. The scheme can be denoted (X, \mathcal{O}_X) , although it is often denoted X , with the structure sheaf \mathcal{O}_X left implicit.

An **isomorphism of schemes** $(X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ is an isomorphism as ringed spaces. Recall the definition of $\Gamma(\cdot, \cdot)$ in §2.2.2. If $U \subset X$ is an open subset, then the elements of $\Gamma(U, \mathcal{O}_X)$ are said to be the **functions on U** ; this generalizes the definition of functions on an affine scheme, §3.2.1.

4.3.2. Remark. From the definition of the structure sheaf on an affine scheme, several things are clear. First of all, if we are told that (X, \mathcal{O}_X) is an affine scheme, we may recover its ring (i.e., find the ring A such that $\text{Spec } A = X$) by taking the ring of global sections, as $X = D(1)$, so:

$$\begin{aligned} \Gamma(X, \mathcal{O}_X) &= \Gamma(D(1), \mathcal{O}_{\text{Spec } A}) \quad \text{as } D(1) = \text{Spec } A \\ &= A. \end{aligned}$$

(You can verify that we get more, and can “recognize X as the scheme $\text{Spec } A$ ”: we get an isomorphism $\pi: (\text{Spec } \Gamma(X, \mathcal{O}_X), \mathcal{O}_{\text{Spec } \Gamma(X, \mathcal{O}_X)}) \xrightarrow{\sim} (X, \mathcal{O}_X)$. For example, if \mathfrak{m} is a maximal ideal of $\Gamma(X, \mathcal{O}_X)$, then $\{\pi([\mathfrak{m}])\} = V(\mathfrak{m})$.) The following exercise will make these ideas rigorous — they are subtler than they first appear.

4.3.3. Caution. It is not part of the definition that the overlap of two affine open subsets is also affine (see Exercise 4.4.C for an example), although this is true in most cases of interest (see Proposition 11.3.11).

4.3.A. ENLIGHTENING EXERCISE (WHICH CAN BE STRANGELY CONFUSING). Describe a bijection between the isomorphisms $\text{Spec } A \xrightarrow{\sim} \text{Spec } A'$ and the ring isomorphisms $A' \xrightarrow{\sim} A$. Hint: the hardest part is to show that if an isomorphism $\pi: \text{Spec } A \xrightarrow{\sim} \text{Spec } A'$ induces an isomorphism $\pi^\#: A' \xrightarrow{\sim} A$, which in turn induces an isomorphism $\rho: \text{Spec } A \xrightarrow{\sim} \text{Spec } A'$, then $\pi = \rho$. First show this on the level of points; this is (surprisingly) the trickiest part. Then show $\pi = \rho$ as maps of topological spaces. Finally, to show $\pi = \rho$ on the level of structure sheaves, use the distinguished base. Feel free to use insights from later in this section, but be careful to avoid circular arguments. Even struggling with this exercise and failing (until reading later sections) will be helpful.

More generally, given $f \in A$, $\Gamma(D(f), \mathcal{O}_{\text{Spec } A}) \cong A_f$. Thus under the natural inclusion of sets $\text{Spec } A_f \hookrightarrow \text{Spec } A$, the Zariski topology on $\text{Spec } A$ restricts to give the Zariski topology on $\text{Spec } A_f$ (Exercise 3.4.I), and the structure sheaf of $\text{Spec } A$ restricts to the structure sheaf of $\text{Spec } A_f$, as the next exercise shows.

4.3.B. IMPORTANT BUT EASY EXERCISE. Suppose $f \in A$. Show that under the identification of $D(f)$ in $\text{Spec } A$ with $\text{Spec } A_f$ (§3.5), there is a natural isomorphism of ringed spaces $(D(f), \mathcal{O}_{\text{Spec } A}|_{D(f)}) \xrightarrow{\sim} (\text{Spec } A_f, \mathcal{O}_{\text{Spec } A_f})$. Hint: notice that distinguished open sets of $\text{Spec } A_f$ are already distinguished open sets in $\text{Spec } A$.

4.3.C. EASY EXERCISE. If X is a scheme, and U is *any* open subset, prove that $(U, \mathcal{O}_X|_U)$ is also a scheme.

4.3.4. Definitions. We say $(U, \mathcal{O}_X|_U)$ is an **open subscheme** of X . If U is also an affine scheme, we often say U is an **affine open subset**, or an **affine open subscheme**, or sometimes informally just an **affine open**. For example, $D(f)$ is an affine open subscheme of $\text{Spec } A$. (Fun fact: it is not true that every affine open subscheme of $\text{Spec } A$ is of the form $D(f)$; see §20.11.10 for an example.)

4.3.D. EASY EXERCISE. Show that if X is a scheme, then the affine open sets form a base for the Zariski topology.

4.3.E. EASY EXERCISE. The **disjoint union of schemes** is defined as you would expect: it is the disjoint union of sets, with the expected topology (thus it is the disjoint union of topological spaces), with the expected sheaf. We use the symbol $X \coprod Y$ for the disjoint union of X and Y , because once we know what morphisms of schemes are, this construction will turn out to be the coproduct in the category of schemes (Exercise 7.3.K).

(a) Show that the disjoint union of a *finite* number of affine schemes is also an affine scheme. (Hint: Exercise 3.6.A.)

(b) (*a first example of a non-affine scheme*) Show that an infinite disjoint union of (nonempty) affine schemes is not an affine scheme. (Hint: affine schemes are quasicompact, Exercise 3.6.G(a). This is basically answered in Remark 3.6.6.)

4.3.5. Remark: A first glimpse of closed subschemes. Open subsets of a scheme come with a natural scheme structure (Definition 4.3.4). For comparison, closed subsets can have many “natural” scheme structures. We will discuss this later (in §9.1), but for now, it suffices for you to know that a closed subscheme of X is, informally, a particular kind of scheme structure on a closed subset of X . As an example: if $I \subset A$ is an ideal, then $\text{Spec } A/I$ endows the closed subset $V(I) \subset \text{Spec } A$ with a scheme structure; but note that there can be different ideals with the same vanishing set (for example (x) and (x^2) in $k[x]$).

4.3.6. Stalks of the structure sheaf: germs, values at a point, and the residue field of a point. Like every sheaf, the structure sheaf has stalks, and unsurprisingly, they are interesting from an algebraic point of view. In fact, we have seen them before.

4.3.7. Definition. We say a ringed space is a **locally ringed space** if its stalks are local rings. By Exercise 4.1.F, the stalk of $\mathcal{O}_{\text{Spec } A}$ at the point $[p]$ is the local ring

A_p . Hence schemes are locally ringed spaces. Manifolds are another example of locally ringed spaces, see §2.1.1. In both cases, taking the quotient by the maximal ideal may be interpreted as evaluating at the point. The maximal ideal of the local ring $\mathcal{O}_{X,p}$ is denoted $\mathfrak{m}_{X,p}$ or \mathfrak{m}_p , and the **residue field** $\mathcal{O}_{X,p}/\mathfrak{m}_p$ is denoted $\kappa(p)$. Functions on an open subset U of a locally ringed space have **values** at each point of U . The value at p of such a function lies in $\kappa(p)$. As usual, we say that a function **vanishes** at a point p if its value at p is 0. (This generalizes our notion of the value of a function on $\text{Spec } A$, defined in §3.2.1.) Notice that we can't even make sense of the phrase of "function vanishing" on ringed spaces in general.

4.3.F. USEFUL EXERCISE.

- (a) If f is a function on a locally ringed space X , show that the subset of X where f vanishes is closed. Hint: if f is a function on a *ringed space* X , show that the subset of X where the germ of f is invertible is open. (Thus the locus on a locally ringed space where a function has nonzero *value* is open. In contrast, the locus it has nonzero *germ* is *closed*, as mentioned in Caution 2.7.7.)
- (b) Show that if f is a function on a locally ringed space that vanishes nowhere, then f is invertible.

Consider a point $[p]$ of an affine scheme $\text{Spec } A$. (Of course, any point of a scheme can be interpreted in this way, as each point has an affine open neighborhood.) The residue field at $[p]$ is A_p/pA_p , which is isomorphic to $K(A/p)$, the fraction field of the quotient. It is useful to note that localization at p and taking the quotient by p "commute", i.e., the following diagram commutes.

$$(4.3.7.1) \quad \begin{array}{ccc} & A_p & \\ \text{localize} \nearrow & & \searrow \text{quotient} \\ A & & A_p/pA_p = K(A/p) \\ \text{quotient} \searrow & & \nearrow \text{localize, i.e., } K(\cdot) \\ & A/p & \end{array}$$

For example, consider the scheme $\mathbb{A}_k^2 = \text{Spec } k[x, y]$, where k is a field of characteristic not 2. Then $(x^2 + y^2)/x(y^2 - x^5)$ is a function away from the y -axis and the curve $y^2 - x^5$. Its value at $(2, 4)$ (by which we mean $[(x - 2, y - 4)]$) is $(2^2 + 4^2)/(2(4^2 - 2^5))$, as

$$\frac{x^2 + y^2}{x(y^2 - x^5)} \equiv \frac{2^2 + 4^2}{2(4^2 - 2^5)}$$

in the residue field — check this if it seems mysterious. And its value at $[(y)]$, the generic point of the x -axis, is $\frac{x^2}{-x^6} = -1/x^4$, which we see by setting y to 0. This is indeed an element of the fraction field of $k[x, y]/(y)$, i.e., $k(x)$. (If you think you care only about algebraically closed fields, let this example be a first warning: A_p/pA_p won't be algebraically closed in general, even if A is a finitely generated \mathbb{C} -algebra!)

If anything makes you nervous, you should make up an example to make you feel better. Here is one: $27/4$ is a function on $\text{Spec } \mathbb{Z} - \{[(2)], [(7)]\}$ or indeed on an even bigger open set. What is its value at $[(5)]$? Answer: $2/(-1) \equiv -2 \pmod{5}$.

What is its value at the generic point $[(0)]$? Answer: $27/4$. Where does it vanish? At $[(3)]$.

4.3.8. Definition: the fiber and rank of an \mathcal{O}_X -module at a point. If \mathcal{F} is an \mathcal{O}_X -module on a scheme X (or more generally, a locally ringed space), define the **fiber** (or *fibre*) of \mathcal{F} at a point $p \in X$ by

$$\mathcal{F}|_p := \mathcal{F}_p \otimes_{\mathcal{O}_{X,p}} \kappa(p).$$

The fiber is a vector space over $\kappa(p)$. The dimension of this vector space $\dim_{\kappa(p)} \mathcal{F}|_p$ is called the **rank of \mathcal{F} at p** , and is denoted $\text{rank}_p \mathcal{F}$. For example, $\mathcal{O}_X|_p$ is $\kappa(p)$, and $\text{rank}_p \mathcal{O}_X = 1$.

4.3.9. \star Definition: Manifolds of various sorts. In the same way, we can cleanly and concisely define (various notions of) manifolds. A **topological** (respectively, **differentiable**, C^∞ = **smooth**) (real) **manifold** is a ringed space (X, \mathcal{O}_X) , whose underlying topological space X is Hausdorff and second countable, such that any point of X has an open neighborhood U such that $(U, \mathcal{O}_X|_U)$ is isomorphic to an open ball in \mathbb{R}^n with the sheaf of continuous (respectively, differentiable, C^∞ = smooth) real-valued functions on it. The phrase **manifold** (without adjective) often means **smooth manifold**. A **complex manifold** is a ringed space (X, \mathcal{O}_X) , whose underlying topological space X is Hausdorff and second countable, such that any point of X has an open neighborhood U such that $(U, \mathcal{O}_X|_U)$ is isomorphic to an open ball in \mathbb{C}^n with the sheaf of holomorphic functions on it.

4.4 Three examples

We now give three extended examples. Our short-term goal is to see that we can really work with the structure sheaf, and can compute the ring of sections of interesting open sets that aren't just distinguished open sets of affine schemes. Our long-term goal is to meet interesting examples that will come up repeatedly in the future.

4.4.1. First example: The plane minus the origin. This example will show you that the distinguished base is something that you can work with. Let $A = k[x, y]$, so $\text{Spec } A = \mathbb{A}_k^2$. Let's work out the space of functions on the open set $U = \mathbb{A}^2 - \{(0, 0)\} = \mathbb{A}^2 - \{[(x, y)]\}$.

It is not immediately obvious whether this is a distinguished open set. (In fact it is not — you may be able to figure out why within a few paragraphs, if you can't right now. It is not enough to show that (x, y) is not a principal ideal.) But in any case, we can describe it as the union of two things which *are* distinguished open sets: $U = D(x) \cup D(y)$. We will find the functions on U by gluing together functions on $D(x)$ and $D(y)$.

The functions on $D(x)$ are, by Definition 4.1.1, $A_x = k[x, y, 1/x]$. The functions on $D(y)$ are $A_y = k[x, y, 1/y]$. Note that A injects into its localizations (if 0 is not inverted), as it is an integral domain (Exercise 1.3.C), so A injects into both A_x and A_y , and both inject into A_{xy} (and indeed into $k(x, y) = K(A)$). So we are looking for functions on $D(x)$ and $D(y)$ that agree on $D(x) \cap D(y) = D(xy)$, i.e., we are interpreting $A_x \cap A_y$ in A_{xy} (or in $k(x, y)$). Clearly those rational functions with only powers of x in the denominator, and also with only powers of y in the

denominator, are the polynomials. Translation: $A_x \cap A_y = A$. Thus we conclude:

$$(4.4.1.1) \quad \Gamma(U, \mathcal{O}_{\mathbb{A}^2}) \equiv k[x, y].$$

In other words, we get no extra functions by removing the origin. Notice how easy that was to calculate!

4.4.2. Aside. Notice that any function on $\mathbb{A}^2 - \{(0, 0)\}$ extends over all of \mathbb{A}^2 . This is an analog of *Hartogs's Lemma* in complex geometry: you can extend a holomorphic function defined on the complement of a set of codimension at least two on a complex manifold over the missing set. This will work more generally in the algebraic setting: you can extend over points in codimension at least 2 not only if they are “smooth”, but also if they are mildly singular — what we will call *normal*. We will make this precise in §12.3.15. This fact will be very useful for us.

4.4.3. We now show an interesting fact: $(U, \mathcal{O}_{\mathbb{A}^2}|_U)$ is a scheme, but it is not an affine scheme. (This is confusing, so you will have to pay attention.) Here's why: otherwise, if $(U, \mathcal{O}_{\mathbb{A}^2}|_U) \cong (\text{Spec } A, \mathcal{O}_{\text{Spec } A})$, then we can recover A by taking global sections:

$$A = \Gamma(U, \mathcal{O}_{\mathbb{A}^2}|_U),$$

which we have already identified in (4.4.1.1) as $k[x, y]$. So if U is affine, then $U \cong \mathbb{A}_k^2$. But this bijection between prime ideals in a ring and points of the spectrum is more constructive than that: *given the prime ideal I , you can recover the point as the generic point of the closed subset cut out by I , i.e., $V(I)$, and given the point p , you can recover the ideal as those functions vanishing at p , i.e., $I(p)$* . In particular, the prime ideal (x, y) of A should cut out a point of $\text{Spec } A$. But on U , $V(x) \cap V(y) = \emptyset$. Conclusion: U is *not* an affine scheme. (If you are ever looking for a counterexample to something, and you are expecting one involving a non-affine scheme, keep this example in mind!)

4.4.4. Gluing two copies of \mathbb{A}^1 together in two different ways. We have now seen two examples of non-affine schemes: an infinite disjoint union of nonempty schemes (Exercise 4.3.E) and $\mathbb{A}^2 - \{(0, 0)\}$. I want to give you two more examples. They are important because they are the first examples of fundamental behavior, the first pathological, and the second central.

First, I need to tell you how to glue two schemes together. Before that, you should review how to glue topological spaces together along isomorphic open sets. Given two topological spaces X and Y , and open subsets $U \subset X$ and $V \subset Y$ along with a homeomorphism $U \xrightarrow{\sim} V$, we can create a new topological space W , that we think of as gluing X and Y together along $U \xrightarrow{\sim} V$. It is the quotient of the disjoint union $X \sqcup Y$ by the equivalence relation $U \sim V$, where the quotient is given the quotient topology. Then X and Y are naturally (identified with) open subsets of W , and indeed cover W . Can you restate this cleanly with an arbitrary (not necessarily finite) number of topological spaces?

Now that we have discussed gluing topological spaces, let's glue schemes together. (This applies without change more generally to ringed spaces.) Suppose you have two schemes (X, \mathcal{O}_X) and (Y, \mathcal{O}_Y) , and open subsets $U \subset X$ and $V \subset Y$, along with a homeomorphism $\pi: U \xrightarrow{\sim} V$, and an isomorphism of structure sheaves $\mathcal{O}_V \xrightarrow{\sim} \pi_* \mathcal{O}_U$ (i.e., an isomorphism of schemes $(U, \mathcal{O}_X|_U) \xrightarrow{\sim}$

$(V, \mathcal{O}_V|_V)$). Then we can glue these together to get a single scheme. Reason: let W be X and Y glued together using the isomorphism $U \xrightarrow{\sim} V$. Then Exercise 2.5.E shows that the structure sheaves can be glued together to get a sheaf of rings. Note that this is indeed a scheme: any point has an open neighborhood that is an affine scheme. (Do you see why?)

4.4.A. ESSENTIAL EXERCISE (CF. EXERCISE 2.5.E). Show that you can glue an arbitrary collection of schemes together. Suppose we are given:

- schemes X_i (as i runs over some index set I , not necessarily finite),
- open subschemes $X_{ij} \subset X_i$ with $X_{ii} = X_i$,
- isomorphisms $f_{ij}: X_{ij} \xrightarrow{\sim} X_{ji}$ with f_{ii} the identity

such that

- (the cocycle condition) the isomorphisms “agree on triple intersections”, i.e., $f_{ik}|_{X_{ij} \cap X_{ik}} = f_{jk}|_{X_{ji} \cap X_{jk}} \circ f_{ij}|_{X_{ij} \cap X_{ik}}$ (so implicitly, to make sense of the right side, $f_{ij}(X_{ik} \cap X_{ij}) \subset X_{jk}$).

(The cocycle condition ensures that f_{ij} and f_{ji} are inverses. In fact, the hypothesis that f_{ii} is the identity also follows from the cocycle condition.) Show that there is a unique scheme X (up to unique isomorphism) along with open subsets isomorphic to the X_i respecting this gluing data in the obvious sense. (Hint: what is X as a set? What is the topology on this set? In terms of your description of the open sets of X , what are the sections of this sheaf over each open set?)

I will now give you two non-affine schemes. Both are handy to know. In both cases, I will glue together two copies of the affine line \mathbb{A}_k^1 . Let $X = \operatorname{Spec} k[t]$, and $Y = \operatorname{Spec} k[u]$. Let $U = D(t) = \operatorname{Spec} k[t, 1/t] \subset X$ and $V = D(u) = \operatorname{Spec} k[u, 1/u] \subset Y$. We will get both examples by gluing X and Y together along U and V . The difference will be in how we glue.

4.4.5. Second example: the affine line with the doubled origin. Consider the isomorphism $U \xrightarrow{\sim} V$ via the isomorphism $k[t, 1/t] \xrightarrow{\sim} k[u, 1/u]$ given by $t \longleftrightarrow u$ (cf. Exercise 4.3.A). The resulting scheme is called the **affine line with doubled origin**. Figure 4.6 is a picture of it.




FIGURE 4.6. The affine line with doubled origin

As the picture suggests, intuitively this is an analog of a failure of Hausdorffness. Now \mathbb{A}^1 itself is not Hausdorff, so we can't say that it is a failure of Hausdorffness. We see this as weird and bad, so we will want to make a definition that will prevent this from happening. This will be the notion of *separatedness* (to be discussed in Chapter 11). (In fact, in some sense, separatedness is the *right* definition of Hausdorffness, see Exercise 11.3.B!) This will answer other of our prayers as well. For example, on a separated scheme, the “affine base of the Zariski topology” is nice — the intersection of two affine open sets will be affine (Proposition 11.3.11).

4.4.B. EXERCISE. Show that the affine line with doubled origin is not affine. Hint: calculate the ring of global sections, and look back at the argument for $\mathbb{A}^2 - \{(0, 0)\}$.

4.4.C. EASY EXERCISE. Do the same construction with \mathbb{A}^1 replaced by \mathbb{A}^2 . You will have defined the **affine plane with doubled origin**. Describe two affine open subsets of this scheme whose intersection is not an affine open subset. (This example was promised in Caution 4.3.3. Also, an “infinite-dimensional” version comes up in Exercise 5.1.K.)

4.4.6. Third example: the projective line. Consider the isomorphism $U \cong V$ via the isomorphism $k[t, 1/t] \cong k[u, 1/u]$ given by $t \leftrightarrow 1/u$. Figure 4.7 is a suggestive picture of this gluing. The resulting scheme is called the **projective line over the field k** , and is denoted \mathbb{P}_k^1 .




FIGURE 4.7. Gluing two affine lines together to get \mathbb{P}^1

Notice how the points glue. Let me assume that k is algebraically closed for convenience. (You can think about how this changes otherwise.) On the first affine line, we have the closed (“traditional”) points $[(t - a)]$, which we think of as “ a on the t -line”, and we have the generic point $[(0)]$. On the second affine line, we have closed points that are “ b on the u -line”, and the generic point. Then a on the t -line is glued to $1/a$ on the u -line (if $a \neq 0$ of course), and the generic point is glued to the generic point (the ideal (0) of $k[t]$ becomes the ideal (0) of $k[t, 1/t]$ upon localization, and the ideal (0) of $k[u]$ becomes the ideal (0) of $k[u, 1/u]$. And (0) in $k[t, 1/t]$ is (0) in $k[u, 1/u]$ under the isomorphism $t \xrightarrow{\sim} 1/u$.

4.4.7. If k is algebraically closed, we can interpret the closed points of \mathbb{P}_k^1 in the following way, which may make this sound closer to the way you have seen projective space defined earlier. The points are of the form $[a, b]$, where a and b are not both zero, and $[a, b]$ is identified with $[ac, bc]$ where $c \in k^\times$. Then if $b \neq 0$, this is identified with a/b on the t -line, and if $a \neq 0$, this is identified with b/a on the u -line.

4.4.8. Proposition. — \mathbb{P}_k^1 is not affine.

Proof. We do this by calculating the ring of global sections. The global sections correspond to sections over X and sections over Y that agree on the overlap. A section on X is a polynomial $f(t)$. A section on Y is a polynomial $g(u)$. If we restrict

$f(t)$ to the overlap, we get something we can still call $f(t)$; and similarly for $g(u)$. Now we want them to be equal: $f(t) = g(1/t)$. But the only polynomials in t that are at the same time polynomials in $1/t$ are the constants k . Thus $\Gamma(\mathbb{P}^1, \mathcal{O}_{\mathbb{P}^1}) = k$. If \mathbb{P}^1 were affine, then it would be $\text{Spec } \Gamma(\mathbb{P}^1, \mathcal{O}_{\mathbb{P}^1}) = \text{Spec } k$, i.e., one point. But it isn't — it has lots of points. \square

We have proved an analog of an important theorem: the only holomorphic functions on \mathbb{CP}^1 are the constants! (See §11.5.7 for a serious yet easy generalization.)

4.4.9. Important Example: Projective space. We now make a preliminary definition of **projective n -space over a field k** , denoted \mathbb{P}_k^n , by gluing together $n + 1$ open sets each isomorphic to \mathbb{A}_k^n . Judicious choice of notation for these open sets will make our life easier. Our motivation is as follows. In the construction of \mathbb{P}^1 above, we thought of points of projective space as $[x_0, x_1]$, where (x_0, x_1) are only determined up to scalars, i.e., (x_0, x_1) is considered the same as $(\lambda x_0, \lambda x_1)$. Then the first patch can be interpreted by taking the locus where $x_0 \neq 0$, and then we consider the points $[1, t]$, and we think of t as x_1/x_0 ; even though x_0 and x_1 are not well-defined, x_1/x_0 is. The second corresponds to where $x_1 \neq 0$, and we consider the points $[u, 1]$, and we think of u as x_0/x_1 . It will be useful to instead use the notation $x_{1/0}$ for t and $x_{0/1}$ for u .

For \mathbb{P}^n , we glue together $n + 1$ open sets, one for each of $i = 0, \dots, n$. The i th open set U_i will have coordinates $x_{0/i}, \dots, x_{(i-1)/i}, x_{(i+1)/i}, \dots, x_{n/i}$. It will be convenient to write this as

$$(4.4.9.1) \quad \text{Spec } k[x_{0/i}, x_{1/i}, \dots, x_{n/i}] / (x_{i/i} - 1)$$

(so we have introduced a “dummy variable” $x_{i/i}$ which we immediately set to 1). We glue the distinguished open set $D(x_{j/i})$ of U_i to the distinguished open set $D(x_{i/j})$ of U_j , by identifying these two schemes by describing the identification of rings

$$\begin{aligned} \text{Spec } k[x_{0/i}, x_{1/i}, \dots, x_{n/i}, 1/x_{j/i}] / (x_{i/i} - 1) &\cong \\ \text{Spec } k[x_{0/j}, x_{1/j}, \dots, x_{n/j}, 1/x_{i/j}] / (x_{j/j} - 1) \end{aligned}$$

via $x_{k/i} = x_{k/j} / x_{i/j}$ and $x_{k/j} = x_{k/i} / x_{j/i}$ (which implies $x_{i/j} x_{j/i} = 1$). We need to check that this gluing information agrees over triple overlaps.

4.4.D. EXERCISE. Check this, as painlessly as possible. (Possible hint: the triple intersection is affine; describe the corresponding ring.)

4.4.10. Definition: \mathbb{P}_A^n . Note that our definition does not use the fact that k is a field. Hence we may as well define \mathbb{P}_A^n for any *ring* A . This will be useful later.

4.4.E. EXERCISE. Show that the only functions on \mathbb{P}_k^n are constants ($\Gamma(\mathbb{P}_k^n, \mathcal{O}) \cong k$), and hence that \mathbb{P}_k^n is not affine if $n > 0$. Hint: you might fear that you will need some delicate interplay among all of your affine open sets, but you will only need two of your open sets to see this. There is even some geometric intuition behind this: the complement of the union of two open sets has codimension 2. But “Algebraic Hartogs’s Lemma” (discussed informally in §4.4.2, and to be stated rigorously in Theorem 12.3.15) says that any function defined on this union extends to be a function on all of projective space. Because we are expecting to see only

constants as functions on all of projective space, we should already see this for this union of our two affine open sets.

4.4.F. EXERCISE (GENERALIZING §4.4.7). Show that if k is algebraically closed, the closed points of \mathbb{P}_k^n may be interpreted in the traditional way: the points are of the form $[a_0, \dots, a_n]$, where the a_i are not all zero, and $[a_0, \dots, a_n]$ is identified with $[\lambda a_0, \dots, \lambda a_n]$ where $\lambda \in k^\times$.

Helpful translation: we think of the closed points of \mathbb{P}_k^n (where $k = \bar{k}$) as $[x_0, x_1, \dots, x_n]$, with $x_i \in k$ (not all x_i zero), and we identify $[x_0, x_1, \dots, x_n]$ with $[\lambda x_0, \lambda x_1, \dots, \lambda x_n]$ for all $\lambda \in k^\times$. Then $[x_0, x_1, \dots, x_n]$ corresponds to a line through the origin, and $(x_0/i, x_1/i, \dots, x_n/i)$ is where this line meets the hyperplane $x_i = 1$.

We will later give other definitions of projective space (Definition 4.5.8, §17.4.2), and the x_i will be called *projective coordinates* there. Our first definition here will often be handy for computing things. But there is something unnatural about it — projective space is highly symmetric, and that isn't clear from our current definition. Furthermore, as noted by Herman Weyl [Wey, p. 90], “The introduction of numbers as coordinates is an act of violence.”

4.4.11. Fun aside: The Chinese Remainder Theorem is a geometric fact. The Chinese Remainder Theorem is embedded in what we have done. We will see this in a single example, but you should then figure out the general statement. The Chinese Remainder Theorem says that knowing an integer modulo 60 is the same as knowing an integer modulo 3, 4, and 5. Here is how to see this in the language of schemes. What is $\text{Spec } \mathbb{Z}/(60)$? What are the prime ideals of this ring? Answer: those prime ideals containing (60) , i.e., those primes dividing 60, i.e., (2) , (3) , and (5) . Figure 4.8 is a sketch of $\text{Spec } \mathbb{Z}/(60)$. They are all closed points, as these are all maximal ideals, so the topology is the discrete topology. What are the stalks? You can check that they are $\mathbb{Z}/(4)$, $\mathbb{Z}/(3)$, and $\mathbb{Z}/(5)$. The nilpotents “at (2) ” are indicated by the “fuzz” on that point. (We discussed visualizing nilpotents with “infinitesimal fuzz” in §4.2.) So what are global sections on this scheme? They are sections on this open set (2) , this other open set (3) , and this third open set (5) . In other words, we have a natural isomorphism of rings

$$\mathbb{Z}/(60) \xrightarrow{\sim} \mathbb{Z}/(2^2) \times \mathbb{Z}/(3) \times \mathbb{Z}/(5).$$




FIGURE 4.8. A picture of the scheme $\text{Spec } \mathbb{Z}/(60)$

4.4.12. ★★ Example. Here is an example of a function on an open subset of a scheme with some surprising behavior. On $X = \text{Spec } k[w, x, y, z]/(wz - xy)$, consider the open subset $D(y) \cup D(w)$. Clearly the function z/y on $D(y)$ agrees with x/w on $D(w)$ on the overlap $D(y) \cap D(w)$. Hence they glue together to give a section. You may have seen this before when thinking about analytic continuation in complex

geometry — we have a “holomorphic” function which has the description z/y on an open set, and this description breaks down elsewhere, but you can still “analytically continue” it by giving the function a different definition on different parts of the space.

Follow-up for curious experts. This function has no “single description” as a well-defined expression in terms of w, x, y, z ! There is a lot of interesting geometry here, and this scheme will be a constant source of (counter)examples for us (look in the index under “cone over smooth quadric surface”). Here is a glimpse, in terms of words we have not yet defined. The space $\text{Spec } k[w, x, y, z]$ is \mathbb{A}^4 , and is, not surprisingly, 4-dimensional. We are working with the subset X , which is a hypersurface, and is 3-dimensional. It is a cone over a smooth quadric surface in \mathbb{P}^3 (flip to Figure 9.2). The open subset $D(y) \subset X$ is X minus some hypersurface, so we are throwing away a codimension 1 locus. $D(w)$ involves throwing away another codimension 1 locus. You might think that the intersection of these two discarded loci is then codimension 2, and that failure of extending this weird function to a global polynomial comes because of a failure of our Hartogs’s Lemma-type theorem. But that’s not true — $V(y) \cap V(w)$ is in fact codimension 1. Here is what is actually going on. The space $V(y)$ is obtained by throwing away the (cone over the) union of two lines ℓ and m_1 , one in each “ruling” of the surface, and $V(w)$ also involves throwing away the (cone over the) union of two lines ℓ and m_2 . The intersection is the (cone over the) line ℓ , which is a codimension 1 set. Remarkably, despite being “pure codimension 1” the cone over ℓ is not cut out even set-theoretically by a single equation (see Exercise 15.2.R). This means that any expression $f(w, x, y, z)/g(w, x, y, z)$ for our function cannot correctly describe our function on $D(y) \cup D(w)$ — at some point of $D(y) \cup D(w)$ it must be $0/0$. Here’s why. Our function can’t be defined on $V(y) \cap V(w)$, so g must vanish here. But g can’t vanish just on the cone over ℓ — it must vanish elsewhere too.

4.5 Projective schemes, and the Proj construction

Projective schemes are important for a number of reasons. Here are a few. Schemes that were of “classical interest” in geometry — and those that you would have cared about before knowing about schemes — are all projective or an open subset thereof (basically, *quasiprojective*, see §4.5.9). Moreover, most “schemes of interest” tend to be projective or quasiprojective. In fact, it is very hard to even give an example of a scheme satisfying basic properties — for example, finite type and “Hausdorff” (“separated”) over a field — that is provably not quasiprojective. For complex geometers: it is hard to find a compact complex variety that is provably not projective (see Remark 11.5.6), and it is quite hard to come up with a complex variety that is provably not an open subset of a projective variety. So projective schemes are really ubiquitous. Also, the notion of “projective k -scheme” is a good approximation of the algebro-geometric version of compactness (“properness”, see §11.5).

Finally, although projective schemes may be obtained by gluing together affine schemes, and we know that keeping track of gluing can be annoying, there is a simple means of dealing with them without worrying about gluing. Just as there is a rough dictionary between rings and affine schemes, we will have a (slightly

looser) analogous dictionary between graded rings and projective schemes. Just as one can work with affine schemes by instead working with rings, one can work with projective schemes by instead working with graded rings.

4.5.1. Motivation from classical geometry.

For geometric intuition, we recall how one thinks of projective space “classically” (in the classical topology, over the real numbers). \mathbb{P}^n can be interpreted as the lines through the origin in \mathbb{R}^{n+1} . Thus subsets of \mathbb{P}^n correspond to unions of lines through the origin of \mathbb{R}^{n+1} , and closed subsets correspond to such unions which are closed. (The same is not true with “closed” replaced by “open”!)

One often pictures \mathbb{P}^n as being the “points at infinite distance” in \mathbb{R}^{n+1} , where the points infinitely far in one direction are associated with the points infinitely far in the opposite direction. We can make this more precise using the decomposition

$$\mathbb{P}^{n+1} = \mathbb{R}^{n+1} \amalg \mathbb{P}^n$$

by which we mean that there is an open subset in \mathbb{P}^{n+1} identified with \mathbb{R}^{n+1} (the points with last “projective coordinate” nonzero), and the complementary closed subset identified with \mathbb{P}^n (the points with last “projective coordinate” zero). (The phrase “projective coordinate” will be formally defined in §4.5.8, but we will use it even before then, in Exercises 4.5.B and 4.5.E.)

Then for example any equation cutting out some set V of points in \mathbb{P}^n will also cut out some set of points in \mathbb{R}^{n+1} that will be a closed union of lines. We call this the *affine cone* of V . These equations will cut out some union of \mathbb{P}^1 's in \mathbb{P}^{n+1} , and we call this the *projective cone* of V . The projective cone is the disjoint union of the affine cone and V . For example, the affine cone over $x^2 + y^2 = z^2$ in \mathbb{P}^2 is just the “classical” picture of a cone in \mathbb{R}^3 , see Figure 4.9. We will make this analogy precise in our algebraic setting in §9.3.11.




FIGURE 4.9. The affine and projective cone of $x^2 + y^2 = z^2$ in classical geometry

4.5.2. Projective schemes, a first description.

We now describe a construction of projective schemes, which will help motivate the Proj construction. We begin by giving an algebraic interpretation of the cone just described, and more generally, getting some practice with transforming between projective coordinates, and coordinates on the standard affine open subsets. We switch coordinates from x, y, z to x_0, x_1, x_2 in order to use the notation of

§4.4.9. For the next questions, in \mathbb{P}_k^2 , we take projective coordinates x_0, x_1, x_2 . (The terminology “projective coordinate” will not be formally defined until §4.5.8, but you should be able to solve these exercises anyway.) The big open set (or “coordinate chart”) $U_0 = \{[x_0, x_1, x_2] : x_0 \neq 0\}$ has coordinates $x_{1/0}$ and $x_{2/0}$, which we interpret as x_1/x_0 and x_2/x_0 . We have similar definitions for U_1 and U_2 . It will be convenient to define $x_{0/0}$ as 1.

4.5.A. EXERCISE. Describe $(x_{0/1}, x_{2/1})$ in terms of $(x_{1/0}, x_{2/0})$.

4.5.B. EXERCISE. Explain how to define a scheme that should be interpreted as $x_0^2 + x_1^2 - x_2^2 = 0$ “in \mathbb{P}_k^2 ”. Hint: in the affine open subset corresponding to $x_2 \neq 0$, it should (in the language of 4.4.9) be cut out by $x_{0/2}^2 + x_{1/2}^2 - 1 = 0$, i.e., it should “be” the scheme $\text{Spec } k[x_{0/2}, x_{1/2}]/(x_{0/2}^2 + x_{1/2}^2 - 1)$. You can similarly guess what it should be on the other two standard open sets, and show that the three schemes glue together.

4.5.C. EXERCISE. (a) Describe *all* homogeneous polynomials in x_0, x_1, x_2 that can be “dehomogenized” to $x_{1/0}^2 + x_{2/0}^5 = 1$. Explain which one is “best”. (b) In what points do the vanishing sets of these polynomials meet the “line at infinity” $x_0 = 0$ in \mathbb{P}_k^2 ?

4.5.D. EXERCISE. Consider the parabola $x_{2/0} = x_{1/0}^2$ (or, if you prefer, $y = x^2$). How does it meet the line at infinity? (What does this mean?) What is its description in the different coordinate patches U_1 and U_2 ?

4.5.3. Remark: Degree d hypersurfaces in \mathbb{P}^n . The degree d homogeneous polynomials in $n + 1$ variables over a field form a vector space of dimension $\binom{n+d}{d}$. (This is essentially the content of Exercise 9.3.J and Exercise 15.1.C.) It is almost true that two polynomials cut out the same subset of \mathbb{P}_k^n if one is a nonzero scalar multiple of the other. Unfortunately, the examples of $x^2y = 0$ and $xy^2 = 0$ show that this isn’t quite right. You will later be able to check (for example, using Exercise 9.5.C) that two polynomials cut out the same *closed subscheme* (whatever that means) if and only if one is a nonzero multiple of the other. (This is some evidence that the notion of a “closed subscheme” is better than that of a “closed subset”.) The zero polynomial doesn’t really cut out a hypersurface in any reasonable sense of the word. Thus we informally imagine that “degree d hypersurfaces in \mathbb{P}^n are parametrized by $\mathbb{P}^{\binom{n+d}{d}-1}$ ”. This intuition will come up repeatedly (in special cases), and we will give it some precise meaning in §28.3.5. (We will properly define *hypersurfaces* in §9.3.2, once we have the language of closed subschemes. At that time we will also define line, hyperplane, quadric hypersurfaces, conic curves, and other wondrous notions.)

4.5.E. EXERCISE (GENERALIZING EXERCISE 4.5.B). Consider \mathbb{P}_A^n , with projective coordinates x_0, \dots, x_n . Given a collection of homogeneous polynomials $f_i \in A[x_0, \dots, x_n]$, make sense of the scheme “cut out in \mathbb{P}_A^n by the f_i .” (This will later be made precise as an example of a “vanishing scheme”, see Exercise 4.5.S.)

This could be taken as the definition of a *projective A -scheme*, but we will wait until §4.5.9 to state it a little better.

4.5.4. Preliminaries on graded rings.

The Proj construction produces a scheme out of a graded ring. We now give some background on graded rings.

4.5.5. Definition: \mathbb{Z} -graded rings. A **\mathbb{Z} -graded ring** is a ring $S_\bullet = \bigoplus_{n \in \mathbb{Z}} S_n$ (the subscript is called the **grading**), where multiplication respects the grading, i.e., sends $S_m \times S_n$ to S_{m+n} . Clearly S_0 is a subring, each S_n is an S_0 -module, and S_\bullet is a S_0 -algebra. Suppose for the remainder of §4.5.5 that S_\bullet is a \mathbb{Z} -graded ring. Those elements of some S_n are called **homogeneous elements** of S_\bullet . An ideal I of S_\bullet is a **homogeneous ideal** (or a **graded ideal**, although we won't use this terminology) if it is generated by homogeneous elements. Nonzero homogeneous elements have an obvious **degree**. An element of S_d is called a **form of degree d** .

4.5.F. EXERCISE.

- (a) Show that an ideal I is homogeneous if and only if it contains the degree n piece of each of its elements for each n . (Hence I can be decomposed into homogeneous pieces, $I = \bigoplus I_n$, and S_\bullet/I has a natural \mathbb{Z} -grading. This is the reason for the name *homogeneous ideal*.)
- (b) Show that the set of homogeneous ideals of a given \mathbb{Z} -graded ring S_\bullet is closed under sum, product, intersection, and radical.
- (c) Show that a homogeneous ideal $I \subset S_\bullet$ is prime if $I \neq S_\bullet$, and if for any *homogeneous* $a, b \in S_\bullet$, if $ab \in I$, then $a \in I$ or $b \in I$.

If T is a multiplicative subset of S_\bullet containing only homogeneous elements, then $T^{-1}S_\bullet$ has a natural structure as a \mathbb{Z} -graded ring.

4.5.6. $\mathbb{Z}^{\geq 0}$ -graded rings, graded ring over A , and finitely generated graded rings. A **$\mathbb{Z}^{\geq 0}$ -graded ring** is a \mathbb{Z} -graded ring with no elements of negative degree.

For the remainder of the book, graded ring will refer to a $\mathbb{Z}^{\geq 0}$ -graded ring. Warning: this convention is nonstandard (for good reason).

From now on, unless otherwise stated, S_\bullet is assumed to be a graded ring. Fix a ring A , which we call the **base ring**. If $S_0 = A$, we say that S_\bullet is a **graded ring over A** . A key example is $A[x_0, \dots, x_n]$, or more generally $A[x_0, \dots, x_n]/I$ where I is a homogeneous ideal with $I \cap S_0 = 0$ (cf. Exercise 4.5.E). Here we take the conventional grading on $A[x_0, \dots, x_n]$, where each x_i has weight 1.

The subset $S_+ := \bigoplus_{i > 0} S_i \subset S_\bullet$ is an ideal, called the **irrelevant ideal**. The reason for the name “irrelevant” will be clearer in a few paragraphs. If the irrelevant ideal S_+ is a finitely generated ideal, we say that S_\bullet is a **finitely generated graded ring over A** . If S_\bullet is generated by S_1 as an A -algebra, we say that S_\bullet is **generated in degree 1**. (We will later find it useful to interpret “ S_\bullet is generated in degree 1” as “the natural map $\text{Sym}^\bullet S_1 \rightarrow S_\bullet$ is a surjection”. The *symmetric algebra* construction will be briefly discussed in §14.2.13.)

4.5.G. EXERCISE.

- (a) Show that a graded ring S_\bullet over A is a finitely generated graded ring (over A) if and only if S_\bullet is a finitely generated graded A -algebra, i.e., generated over $A = S_0$ by a finite number of homogeneous elements of positive degree. (Hint for the forward implication: show that the generators of S_+ as an ideal are also generators of S_\bullet as an algebra.)

(b) Show that a graded ring S_\bullet over A is Noetherian if and only if $A = S_0$ is Noetherian and S_\bullet is a finitely generated graded ring.

4.5.7. The Proj construction.

We now define a scheme $\text{Proj } S_\bullet$, where S_\bullet is a $(\mathbb{Z}^{\geq 0})$ -graded ring. Here are two examples, to provide a light at the end of the tunnel. If $S_\bullet = A[x_0, \dots, x_n]$, we will recover \mathbb{P}_A^n ; and if $S_\bullet = A[x_0, \dots, x_n]/(f(x_0, \dots, x_n))$ where f is homogeneous, we will construct something “cut out in \mathbb{P}_A^n by the equation $f = 0$ ” (cf. Exercise 4.5.E).

As we did with Spec of a ring, we will build $\text{Proj } S_\bullet$ first as a set, then as a topological space, and finally as a ringed space. In our preliminary definition of \mathbb{P}_A^n , we glued together $n + 1$ well-chosen affine pieces, but we don’t want to make any choices, so we do this by simultaneously considering “all possible” affine open sets. Our affine building blocks will be as follows. For each homogeneous $f \in S_+$, note that the localization $(S_\bullet)_f$ is naturally a \mathbb{Z} -graded ring, where $\deg(1/f) = -\deg f$. Consider

$$(4.5.7.1) \quad \text{Spec}((S_\bullet)_f)_0.$$

where $((S_\bullet)_f)_0$ means the 0-graded piece of the graded ring $(S_\bullet)_f$. (These will be our affine building blocks, as f varies over the homogeneous elements of S_+ .) The notation $((S_\bullet)_f)_0$ is admittedly horrible — the first and third subscripts refer to the grading, and the second refers to localization. As motivation for considering this construction: applying this to $S_\bullet = k[x_0, \dots, x_n]$, with $f = x_i$, we obtain the ring appearing in (4.4.9.1):

$$k[x_{0/i}, x_{1/i}, \dots, x_{n/i}]/(x_{i/i} - 1).$$

(Before we begin the construction: another possible way of defining $\text{Proj } S_\bullet$ is by gluing together affines of this form, by jumping straight to Exercises 4.5.N and 4.5.O. If you prefer that, by all means do so.)

The *points* of $\text{Proj } S_\bullet$ are the set of homogeneous prime ideals of S_\bullet not containing the irrelevant ideal S_+ (the “relevant prime ideals”).

4.5.H. IMPORTANT AND TRICKY EXERCISE. Suppose $f \in S_+$ is homogeneous.

(a) Give a bijection between the prime ideals of $((S_\bullet)_f)_0$ and the homogeneous prime ideals of $(S_\bullet)_f$. Hint: Avoid notational confusion by proving instead that if A is a \mathbb{Z} -graded ring with a homogeneous invertible element f in positive degree, then there is a bijection between prime ideals of A_0 and homogeneous prime ideals of A . Using the ring map $A_0 \rightarrow A$, from each homogeneous prime ideal of A we find a prime ideal of A_0 . The reverse direction is the harder one. Given a prime ideal $P_0 \subset A_0$, define $P \subset A$ (a priori only a subset) as $\bigoplus Q_i$, where $Q_i \subset A_i$, and $a \in Q_i$ if and only if $a^{\deg f}/f^i \in P_0$. Note that $Q_0 = P_0$. Show that $a \in Q_i$ if and only if $a^2 \in Q_{2i}$; show that if $a_1, a_2 \in Q_i$ then $a_1^2 + 2a_1a_2 + a_2^2 \in Q_{2i}$ and hence $a_1 + a_2 \in Q_i$; then show that P is a homogeneous ideal of A ; then show that P is prime.

(b) Interpret the set of prime ideals of $((S_\bullet)_f)_0$ as a subset of $\text{Proj } S_\bullet$.

The correspondence of the points of $\text{Proj } S_\bullet$ with homogeneous prime ideals helps us picture $\text{Proj } S_\bullet$. For example, if $S_\bullet = k[x, y, z]$ with the usual grading, then we picture the homogeneous prime ideal $(z^2 - x^2 - y^2)$ first as a subset of $\text{Spec } S_\bullet$; it is a cone (see Figure 4.9). As in §4.5.1, we picture \mathbb{P}_k^2 as the “plane at infinity”.

Thus we picture this equation as cutting out a conic “at infinity” (in $\text{Proj } S_\bullet$). We will make this intuition somewhat more precise in §9.3.11.

Motivated by the affine case, if T is a set of homogeneous elements of S_\bullet , define the (projective) **vanishing set** of T , $V(T) \subset \text{Proj } S_\bullet$, to be those homogeneous prime ideals containing T but not S_+ . Define $V(f)$ if f is a homogeneous element of *positive* degree in the obvious way. Let $D_+(f) := \text{Proj } S_\bullet \setminus V(f)$ (the **projective distinguished open set**) be the complement of $V(f)$. (The subscript is intended to distinguish this notation from the similar $D(f)$.) Once we define a scheme structure on $\text{Proj } S_\bullet$, we will (without comment) use $D_+(f)$ to refer to the open *subscheme*, not just the open subset. (These definitions can certainly be extended to remove the positive degree hypotheses. For example, the definition of $V(T)$ makes sense for any subset T of S_\bullet , and the definition of $D_+(f)$ makes sense even if f has degree 0. In what follows, we deliberately stick to these narrower definitions. For example, we will want the $D_+(f)$ to form an affine cover, and if f has degree 0, then $D_+(f)$ needn’t be affine.)

4.5.I. EXERCISE. Show that $D_+(f)$ “is” (or more precisely, “corresponds to”) the subset $\text{Spec}((S_\bullet)_f)_0$ you described in Exercise 4.5.H(b). For example, in §4.4.9, the $D_+(x_i)$ are the standard open sets covering projective space.

As in the affine case, the $V(I)$ ’s satisfy the axioms of the closed sets of a topology, and we call this the **Zariski topology** on $\text{Proj } S_\bullet$. (Other definitions given in the literature may look superficially different, but can be easily shown to be the same.) Many statements about the Zariski topology on Spec of a ring carry over to this situation with little extra work. Clearly $D_+(f) \cap D_+(g) = D_+(fg)$, by the same immediate argument as in the affine case (Exercise 3.5.D).

4.5.J. EASY EXERCISE. Verify that the projective distinguished open sets $D_+(f)$ (as f runs through the homogeneous elements of S_+) form a base of the Zariski topology.

4.5.K. EXERCISE. Fix a graded ring S_\bullet .

- (a) Suppose I is any homogeneous ideal of S_\bullet contained in S_+ , and f is a homogeneous element of positive degree. Show that f vanishes on $V(I)$ (i.e., $V(I) \subset V(f)$) if and only if $f^n \in I$ for some n . (Hint: Mimic the affine case; see Exercise 3.4.J.) In particular, as in the affine case (Exercise 3.5.E), if $D_+(f) \subset D_+(g)$, then $f^n \in (g)$ for some n , and vice versa. (Here g is also homogeneous of positive degree.)
- (b) If $Z \subset \text{Proj } S_\bullet$, define $I(Z) \subset S_+$. Show that it is a homogeneous ideal of S_\bullet . For any two subsets, show that $I(Z_1 \cup Z_2) = I(Z_1) \cap I(Z_2)$.
- (c) For any subset $Z \subset \text{Proj } S_\bullet$, show that $V(I(Z)) = \bar{Z}$.

4.5.L. EXERCISE (CF. EXERCISE 3.5.B). Fix a graded ring S_\bullet , and a homogeneous ideal $I \subseteq S_+$. Show that the following are equivalent.

- (a) $V(I) = \emptyset$.
- (b) For any f_i (as i runs through some index set) generating I , $\cup D_+(f_i) = \text{Proj } S_\bullet$.
- (c) $\sqrt{I} = S_+$.

This is more motivation for the ideal S_+ being “irrelevant”: any ideal whose radical contains it is “geometrically irrelevant”.

We now construct $\text{Proj } S_\bullet$ as a *scheme*.

4.5.M. EXERCISE. Suppose some homogeneous $f \in S_+$ is given. Via the inclusion

$$D_+(f) = \text{Spec}((S_\bullet)_f)_0 \hookrightarrow \text{Proj } S_\bullet$$

of Exercise 4.5.I, show that the Zariski topology on $\text{Proj } S_\bullet$ restricts to the Zariski topology on $\text{Spec}((S_\bullet)_f)_0$.

Now that we have defined $\text{Proj } S_\bullet$ as a topological space, we are ready to define the structure sheaf. On $D_+(f)$, we wish it to be the structure sheaf of $\text{Spec}((S_\bullet)_f)_0$. We will glue these sheaves together using Exercise 2.5.E on gluing sheaves.

4.5.N. EXERCISE. If $f, g \in S_+$ are homogeneous and nonzero, describe an isomorphism between $\text{Spec}((S_\bullet)_{fg})_0$ and the distinguished open subset $D(g^{\deg f} / f^{\deg g})$ of $\text{Spec}((S_\bullet)_f)_0$.

Similarly, $\text{Spec}((S_\bullet)_g)_0$ is identified with a distinguished open subset of $\text{Spec}((S_\bullet)_f)_0$. We then glue the various $\text{Spec}((S_\bullet)_f)_0$ (as f varies) altogether, using these pairwise gluings.

4.5.O. EXERCISE. By checking that these gluings behave well on triple overlaps (see Exercise 2.5.E), finish the definition of the scheme $\text{Proj } S_\bullet$.

4.5.P. EXERCISE. (Some will find this essential, others will prefer to ignore it.) (Re)interpret the structure sheaf of $\text{Proj } S_\bullet$ in terms of compatible germs.

4.5.8. Definition. We (re)define **projective space** (over a ring A) by

$$\mathbb{P}_A^n := \text{Proj } A[x_0, \dots, x_n].$$

This definition involves no messy gluing, or special choice of patches. Note that the variables x_0, \dots, x_n , which we call the **projective coordinates** on \mathbb{P}_A^n , are part of the definition. (They may have other names than x 's, depending on the context.)

4.5.Q. EXERCISE. Check that this agrees with our earlier construction of \mathbb{P}_A^n (§4.4.9). (How do you know that the $D_+(x_i)$ cover $\text{Proj } A[x_0, \dots, x_n]$?)

Notice that with our old definition of projective space, it would have been a nontrivial exercise to show that $D_+(x^2 + y^2 - z^2) \subset \mathbb{P}_k^2$ (the complement of a plane conic) is affine; with our new perspective, it is immediate — it is $\text{Spec}(k[x, y, z]_{(x^2 + y^2 - z^2)})_0$.

4.5.R. EXERCISE. Suppose that k is an algebraically closed field. We know from Exercise 4.4.F that the closed points of \mathbb{P}_k^n , as defined in §4.4.9, are in bijection with the points of “classical” projective space. By Exercise 4.5.Q, the scheme \mathbb{P}_k^n as defined in §4.4.9 is isomorphic to $\text{Proj } k[x_0, \dots, x_n]$. Therefore, each point $[a_0, \dots, a_n]$ of classical projective space corresponds to a homogeneous prime ideal of $k[x_0, \dots, x_n]$. Which homogeneous prime ideal is it?

We now figure out the “right definition” of the vanishing scheme, in analogy with the vanishing set $V(\cdot)$ (§3.4.1). You will be defining a *closed subscheme* (mentioned in Remark 4.3.5, and to be properly defined in §9.1).

4.5.S. EXERCISE. If S_\bullet is generated in degree 1, and $f \in S_+$ is homogeneous, explain how to define $V(f)$ “in” $\text{Proj } S_\bullet$, the **vanishing scheme** of f . This is the same notation as the “vanishing set” $V(\cdot)$, but now means something richer than a mere set. (Exercise 9.1.K is similar. Warning: in general, f isn’t a function on $\text{Proj } S_\bullet$. We will later interpret it as something close: a section of a line bundle, see for example §15.1.2.) Hence define $V(I)$ for any homogeneous ideal I of S_+ . (Another solution in more general circumstances will be given in Exercise 14.2.N.)

4.5.9. Projective and quasiprojective schemes.

We call a scheme of the form (i.e., isomorphic to) $\text{Proj } S_\bullet$, where S_\bullet is a *finitely generated* graded ring over A , a **projective scheme over A** , or a **projective A -scheme**. A **quasiprojective A -scheme** is a quasicompact open subscheme of a projective A -scheme. The “ A ” is omitted if it is clear from the context; often A is a field.

4.5.10. Unimportant remarks. (i) Note that $\text{Proj } S_\bullet$ makes sense even when S_\bullet is not finitely generated. This can be useful. For example, you will later be able to do Exercise 7.4.D without worrying about Exercise 7.4.H.

(ii) The quasicompact requirement in the definition of quasiprojectivity is of course redundant in the Noetherian case (cf. Exercise 3.6.U), which is all that matters to most.

4.5.11. Silly example. Note that $\mathbb{P}_A^0 = \text{Proj } A[T] \cong \text{Spec } A$. Thus “ $\text{Spec } A$ is a projective A -scheme”.

4.5.12. Example: projectivization of a vector space $\mathbb{P}V$. We can make this definition of projective space even more choice-free as follows. Let V be an $(n+1)$ -dimensional vector space over k . (Here k can be replaced by any ring A as usual.) Define

$$\text{Sym}^\bullet V^\vee = k \oplus V^\vee \oplus \text{Sym}^2 V^\vee \oplus \cdots .$$

(The reason for the dual is explained by the next exercise.) If for example V is the dual of the vector space with basis associated to x_0, \dots, x_n , we would have $\text{Sym}^\bullet V^\vee = k[x_0, \dots, x_n]$. Then we can define $\mathbb{P}V := \text{Proj}(\text{Sym}^\bullet V^\vee)$. In this language, we have an interpretation for x_0, \dots, x_n : they are linear functionals on the underlying vector space V . (Warning: some authors use the definition $\mathbb{P}V = \text{Proj}(\text{Sym}^\bullet V)$, so be cautious.)

4.5.T. UNIMPORTANT EXERCISE. Suppose k is algebraically closed. Describe a natural bijection between one-dimensional subspaces of V and the closed points of $\mathbb{P}V$. Thus this construction canonically (in a basis-free manner) describes the one-dimensional subspaces of the vector space V .

Unimportant remark: you may be surprised at the appearance of the dual in the definition of $\mathbb{P}V$. This is partially explained by the previous exercise. Most normal (traditional) people define the projectivization of a vector space V to be the space of one-dimensional subspaces of V . Grothendieck considered the projectivization to be the space of one-dimensional *quotients*. One motivation for this is that it gets rid of the annoying dual in the definition above. There are better reasons, that we won’t go into here. In a nutshell, quotients tend to be better-behaved

than subobjects for coherent sheaves, which generalize the notion of vector bundle. (Coherent sheaves are discussed in Chapter 14.)

On another note related to Exercise 4.5.T: you can also describe a natural bijection between points of V and the closed points of $\mathrm{Spec}(\mathrm{Sym}^\bullet V^\vee)$. This construction respects the affine/projective cone picture of §9.3.11.

4.5.13. The Grassmannian. At this point, we could describe the fundamental geometric object known as the *Grassmannian*, and give the “wrong” (but correct) definition of it. We will instead wait until §7.7 to give the wrong definition, when we will know enough to sense that something is amiss. The right definition will be given in §17.7.

Some properties of schemes

5.1 Topological properties

We will now define some useful properties of schemes. As you see each example, you should try these out in specific examples of your choice, such as particular schemes of the form $\operatorname{Spec} \mathbb{C}[x_1, \dots, x_n]/(f_1, \dots, f_r)$.

The definitions of *connected*, *connected component*, *(ir)reducible*, *irreducible component*, *quasicompact*, *generization*, *specialization*, *generic point*, *Noetherian topological space*, and *closed point*, were given in §3.6. You should have pictures in your mind of each of these notions.

Exercise 3.6.C shows that \mathbb{A}^n is irreducible (it was easy). This argument “behaves well under gluing”, yielding:

5.1.A. EASY EXERCISE. Show that \mathbb{P}_k^n is irreducible.

5.1.B. EXERCISE. Exercise 3.7.F showed that there is a bijection between irreducible closed subsets and points for affine schemes (the map sending a point p to the closed subset $\overline{\{p\}}$ is a bijection). Show that this is true of schemes in general.

5.1.C. EASY EXERCISE. Prove that if X is a scheme that has a finite cover $X = \bigcup_{i=1}^n \operatorname{Spec} A_i$ where A_i is Noetherian, then X is a Noetherian topological space (§3.6.14). (We will soon call a scheme with such a cover a *Noetherian scheme*, §5.3.4.) Hint: show that a topological space that is a finite union of Noetherian subspaces is itself Noetherian.

Thus \mathbb{P}_k^n and $\mathbb{P}_{\mathbb{Z}}^n$ are Noetherian topological spaces: we built them by gluing together a finite number of spectra of Noetherian rings.

5.1.D. EASY EXERCISE. Show that a scheme X is quasicompact if and only if it can be written as a finite union of affine open subschemes. (Hence \mathbb{P}_A^n is quasicompact for any ring A .)

5.1.E. EXERCISE: QUASICOMPACT SCHEMES HAVE CLOSED POINTS. Show that if X is a quasicompact scheme, then every point has a closed point in its closure. Show that every nonempty closed subset of X contains a closed point of X . In particular, every nonempty quasicompact scheme has a closed point. (Warning: there exist nonempty schemes with no closed points — see for example [Liu, Exer. 3.27], or [Schw], or Hochster’s thesis [Ho1, Ho2] — so your argument had better use the quasicompactness hypothesis!)

This exercise will often be used in the following way. If there is some property P of points of a scheme that is “open” (if a point p has P , then there is some open neighborhood U of p such that all the points in U have P), then to check if *all* points of a quasicompact scheme have P , it suffices to check only the closed points. (A first example of this philosophy is Exercise 5.2.E.) This provides a connection between schemes and the classical theory of varieties — the points of traditional varieties are the *closed* points of the corresponding schemes (essentially by the Nullstellensatz, see §3.6.9 and Exercise 5.3.F). In many good situations, the closed points are dense (such as for varieties, see §3.6.9 and Exercise 5.3.F again), but this is not true in some fundamental cases (see Exercise 3.6.J(b)).

5.1.1. Quasiseparated schemes. Quasiseparatedness is a weird notion that comes in handy for certain people. (Warning: we will later realize that this is really a property of *morphisms*, not of schemes, §8.3.1.) Most people, however, can ignore this notion, as the schemes they will encounter in real life will all have this property. A topological space is **quasiseparated** if the intersection of any two quasicompact open subsets is quasicompact. (The motivation for the “separatedness” in the name is that it is a weakened version of *separated*, for which the intersection of any two *affine* open sets is affine, see Proposition 11.3.11.)

5.1.F. SHORT EXERCISE. Show that a scheme is quasiseparated if and only if the intersection of any two affine open subsets is a finite union of affine open subsets.

5.1.G. EXERCISE. Show that affine schemes are quasiseparated. (Possible hint: Exercise 5.1.F.)

5.1.H. EXERCISE. Show that a scheme X is quasiseparated if and only if there *exists* a cover of X by affine open subsets, any two of which have intersection also covered by a *finite* number of affine open subsets.

We will see that quasiseparatedness will be a useful hypothesis in theorems (in conjunction with quasicompactness), and that various interesting kinds of schemes (affine, locally Noetherian, separated, see Exercises 5.1.G, 5.3.A, and 11.3.D respectively) are quasiseparated, and this will allow us to state theorems more succinctly (e.g., “if X is quasicompact and quasiseparated” rather than “if X is quasicompact, and either this or that or the other thing hold”).

“Quasicompact and quasiseparated” means something concrete:

5.1.I. EXERCISE. Show that a scheme X is quasicompact and quasiseparated if and only if X can be covered by a finite number of affine open subsets, any two of which have intersection also covered by a finite number of affine open subsets.

5.1.2. Qcqs. So when you see “quasicompact and quasiseparated” as hypotheses in a theorem, you should take this as a clue that you will use this interpretation, and that finiteness will be used in an essential way. This combination of hypotheses is so common that people often describe it a “qcqs” for short (see for example the Qcqs Lemma 6.2.8).

5.1.J. EASY EXERCISE. Show that all projective A -schemes are quasicompact and quasiseparated. (Hint: use the fact that the graded ring in the definition is finitely generated — those finite number of generators will lead you to a covering set.)

5.1.K. EXERCISE (A NONQUASISEPARATED SCHEME). Let $X = \operatorname{Spec} k[x_1, x_2, \dots]$, and let U be $X - [m]$ where m is the maximal ideal (x_1, x_2, \dots) . Take two copies of X , glued along U (“affine ∞ -space with a doubled origin”, see Example 4.4.5 and Exercise 4.4.C for “finite-dimensional” versions). Show that the result is not quasiseparated. Hint: This “open embedding” $U \subset X$ came up earlier in Exercise 3.6.G(b) as an example of a nonquasicompact open subset of an affine scheme.

5.1.3. Dimension. One very important topological notion is *dimension*. (It is amazing that this is a *topological* idea.) But despite being intuitively fundamental, it is more difficult, so we postpone it until Chapter 12.

5.2 Reducedness and integrality

Recall that one of the alarming things about schemes is that functions are not determined by their values at points, and that was because of the presence of nilpotents (§3.2.12).

5.2.1. Definition. We say that a scheme X is **reduced** if its stalks $\mathcal{O}_{X,p}$ have no nonzero nilpotents (for all p).

5.2.A. EXERCISE. Show that if f and g are two functions (global sections of \mathcal{O}_X) on a reduced scheme that agree at all points, then $f = g$. (Two hints: $\mathcal{O}_X(U) \hookrightarrow \prod_{p \in U} \mathcal{O}_{X,p}$ from Exercise 2.4.A, and the nilradical is intersection of all prime ideals from Theorem 3.2.13.)

5.2.B. EXERCISE. Recall that a ring is said to be *reduced* if it has no nonzero nilpotents (§3.2.14). Show that A is a reduced ring if and only if $\operatorname{Spec} A$ is reduced. Show that a scheme X is reduced if $\mathcal{O}_X(U)$ is reduced for every open set U of X .

5.2.C. EASY EXERCISE. Show that \mathbb{A}_k^n and \mathbb{P}_k^n are reduced.

The scheme $\operatorname{Spec} k[x, y]/(y^2, xy)$ is nonreduced. When we sketched it in Figure 4.4, we indicated that the fuzz represented nonreducedness at the origin. The following exercise is a first stab at making this precise.

5.2.D. EXERCISE. Show that $(k[x, y]/(y^2, xy))_x$ has no nonzero nilpotent elements. (Possible hint: show that it is isomorphic to another ring, by considering the geometric picture. Exercise 3.2.L may give another hint.) Show that the only point of $\operatorname{Spec} k[x, y]/(y^2, xy)$ with a nonreduced stalk is the origin.

5.2.E. UNIMPORTANT EXERCISE. If X is a quasicompact scheme, show that it suffices to check reducedness at closed points. Hint: Do not try to show that reducedness is an open condition (see Remark 5.2.2). Instead show that any nonreduced point has a nonreduced closed point in its closure, using Exercise 5.1.E. (This result is interesting, but we won’t use it.)

5.2.2. Remark. We will see in Exercise 6.5.N that reducedness is an open condition for locally Noetherian schemes. But in general, reducedness is not an open

condition. You may be able to identify the underlying topological space of

$$X = \operatorname{Spec} \mathbb{C}[x, y_1, y_2, \dots] / (y_1^2, y_2^2, y_3^2, \dots, (x-1)y_1, (x-2)y_2, (x-3)y_3, \dots)$$

with that of $\operatorname{Spec} \mathbb{C}[x]$, and then to show that the nonreduced points of X are precisely the closed points corresponding to the positive integers. The complement of this set is not Zariski open.

5.2.3. ★ Warning for experts. If a scheme X is reduced, then from the definition of reducedness, its ring of global sections is reduced. However, the converse is not true. You already know enough to verify that if X is the scheme cut out by $x^2 = 0$ in \mathbb{P}_k^2 , then $\Gamma(X, \mathcal{O}_X) \cong k$, and that X is nonreduced. (This example will come up again in §11.5.7 and §19.6.U.)

5.2.F. EXERCISE. Suppose X is quasicompact, and f is a function that vanishes at all points of X . Show that there is some n such that $f^n = 0$. Show that this may fail if X is not quasicompact. (This exercise is less important, but shows why we like quasicompactness, and gives a standard pathology when quasicompactness doesn't hold.) Hint: take an infinite disjoint union of $\operatorname{Spec} A_n$ with $A_n := k[\epsilon]/(\epsilon^n)$. This scheme arises again in §9.4.2 (see Figure 9.5 for a picture) and in Caution/Example 9.4.11.

5.2.4. Definition. A scheme X is **integral** if it is nonempty, and $\mathcal{O}_X(U)$ is an integral domain for every nonempty open subset U of X .

5.2.G. IMPORTANT EXERCISE. Show that a scheme X is integral if and only if it is irreducible and reduced. (Thus we picture integral schemes as: “one piece, no fuzz”. Possible hint: Exercise 4.3.F.)

5.2.H. EXERCISE. Show that an affine scheme $\operatorname{Spec} A$ is integral if and only if A is an integral domain.

5.2.I. EXERCISE. Suppose X is an integral scheme. Then X (being irreducible) has a generic point η . Suppose $\operatorname{Spec} A$ is any nonempty affine open subset of X . Show that $\mathcal{O}_{X,\eta}$ (the stalk of \mathcal{O}_X at η) is naturally identified with $K(A)$, the fraction field of A . This is called the **function field** $K(X)$ of X . It can be computed on any nonempty open set of X , as any such open set contains the generic point. The elements of $K(X)$ are called **rational functions on X** (to be generalized further in Definition 6.5.34).

5.2.J. EXERCISE. Suppose X is an integral scheme. Show that the restriction maps $\operatorname{res}_{U,V}: \mathcal{O}_X(U) \rightarrow \mathcal{O}_X(V)$ are inclusions so long as $V \neq \emptyset$. Suppose $\operatorname{Spec} A$ is any nonempty affine open subset of X (so A is an integral domain). Show that the natural map $\mathcal{O}_X(U) \rightarrow \mathcal{O}_{X,\eta} = K(A)$ (where U is any nonempty open subset) is an inclusion.

Thus irreducible varieties (an important example of integral schemes defined later) have the convenient property that sections over different open sets can be considered subsets of the same ring. In particular, restriction maps (except to the empty set) are always inclusions, and gluing is easy: functions f_i on a cover U_i of U (as i runs over an index set) glue if and only if they are the same element of

$K(X)$. This is one reason why (irreducible) varieties are usually introduced before schemes.

5.2.5. Caution. Integrality is not stalk-local (the disjoint union of two integral schemes is not integral, as $\text{Spec } A \amalg \text{Spec } B = \text{Spec}(A \times B)$ by Exercise 3.6.A), but it almost is, see Exercise 5.3.C.

5.3 The Affine Covering Lemma, and properties of schemes that can be checked “affine-locally”

This section is intended to address something tricky in the definition of schemes. We have defined a scheme as a topological space with a sheaf of rings, that can be covered by affine schemes. Hence we have all of the affine open sets in the cover, but we don’t know how to communicate between any two of them. Somewhat more explicitly, if I have an affine cover, and you have an affine cover, and we want to compare them, and I calculate something on my cover, there should be some way of us getting together, and figuring out how to translate my calculation over to your cover. The Affine Communication Lemma 5.3.2 will provide a convenient machine for doing this.

Thanks to this lemma, we can define a host of important properties of schemes. All of these are “affine-local”, in that they can be checked on any affine cover, i.e., a covering by open affine sets. (We state this more formally after the statement of the Affine Communication Lemma.) We like such properties because we can check them using any affine cover we like. If the scheme in question is quasicompact, then we need only check a finite number of affine open sets.

5.3.1. Proposition. — *Suppose $\text{Spec } A$ and $\text{Spec } B$ are affine open subschemes of a scheme X . Then $\text{Spec } A \cap \text{Spec } B$ is the union of open sets that are simultaneously distinguished open subschemes of $\text{Spec } A$ and $\text{Spec } B$.*




FIGURE 5.1. A trick to show that the intersection of two affine open sets may be covered by open sets that are simultaneously distinguished in both affine open sets

Proof. (See Figure 5.1.) Given any point $p \in \text{Spec } A \cap \text{Spec } B$, we produce an open neighborhood of p in $\text{Spec } A \cap \text{Spec } B$ that is simultaneously distinguished in both $\text{Spec } A$ and $\text{Spec } B$. Let $\text{Spec } A_f$ be a distinguished open subset of $\text{Spec } A$ contained

in $\text{Spec } A \cap \text{Spec } B$ and containing p . Let $\text{Spec } B_g$ be a distinguished open subset of $\text{Spec } B$ contained in $\text{Spec } A_f$ and containing p . Then $g \in \Gamma(\text{Spec } B, \mathcal{O}_X)$ restricts to an element $g' \in \Gamma(\text{Spec } A_f, \mathcal{O}_X) = A_f$. The points of $\text{Spec } A_f$ where g vanishes are precisely the points of $\text{Spec } A_f$ where g' vanishes, so

$$\begin{aligned} \text{Spec } B_g &= \text{Spec } A_f \setminus \{[p] : g' \in \mathfrak{p}\} \\ &= \text{Spec}(A_f)_{g'}. \end{aligned}$$

If $g' = g''/f^n$ ($g'' \in A$) then $\text{Spec}(A_f)_{g'} = \text{Spec } A_{fg''}$, and we are done. \square

The following easy result will be crucial for us.

5.3.2. Affine Communication Lemma. — *Let P be some property enjoyed by some affine open subsets of a scheme X , such that for any affine open subset $\text{Spec } A \hookrightarrow X$,*

- (i) *if $\text{Spec } A \hookrightarrow X$ has property P then for any $f \in A$, $\text{Spec } A_f \hookrightarrow X$ does too.*
- (ii) *if $(f_1, \dots, f_n) = A$ (i.e., the $\text{Spec } A_{f_i}$ cover $\text{Spec } A$, see Exercise 3.5.B), and $\text{Spec } A_{f_i} \hookrightarrow X$ has P for all i , then so does $\text{Spec } A \hookrightarrow X$.*

Suppose that $X = \cup_{i \in I} \text{Spec } A_i$ where $\text{Spec } A_i$ has property P . Then every affine open subset of X has P too.

We say such a property is **affine-local**: a property is *affine-local* if we can check it on any affine cover. (For example, reducedness is an affine-local property. Do you understand why?) Note that if U is an open subscheme of X , then U inherits any affine-local property of X . Note also that any property that is stalk-local (a scheme has property P if and only if all its stalks have property Q) is necessarily affine-local (a scheme has property P if and only if all of its affine open sets have property R , where an affine scheme has property R if and only if all its stalks have property Q). But it is sometimes not so obvious what the right definition of Q is; see for example the discussion of normality in the next section.

Proof. Let $\text{Spec } A$ be an affine subscheme of X . Cover $\text{Spec } A$ with a finite number of distinguished open sets $\text{Spec } A_{g_j}$, each of which is distinguished in some $\text{Spec } A_i$. This is possible by Proposition 5.3.1 and the quasicompactness of $\text{Spec } A$ (Exercise 3.6.G(a)). By (i), each $\text{Spec } A_{g_j}$ has P . By (ii), $\text{Spec } A$ has P . \square

By choosing the property P appropriately, we define some important properties of schemes.

5.3.3. Proposition. — *Suppose A is a ring, and $(f_1, \dots, f_n) = A$.*

- (a) *If A is a Noetherian ring, then so is A_{f_i} . If each A_{f_i} is Noetherian, then so is A .*
- (b) *Suppose B is a ring, and A is a B -algebra. (Hence A_g is a B -algebra for all $g \in A$.) If A is a finitely generated B -algebra, then so is A_{f_i} . If each A_{f_i} is a finitely generated B -algebra, then so is A .*

We will prove these shortly (§5.3.9). But let's first motivate you to read the proof by giving some interesting definitions and results *assuming* Proposition 5.3.3 is true.

5.3.4. Important Definition. Suppose X is a scheme. If X can be covered by affine open sets $\text{Spec } A$ where A is Noetherian, we say that X is a **locally Noetherian scheme**. If in addition X is quasicompact, or equivalently can be covered by finitely many such affine open sets, we say that X is a **Noetherian scheme**. (We will see a

number of definitions of the form “if X has this property, we say that it is locally Q ; if further X is quasicompact, we say that it is Q .”) By Exercise 5.1.C, the underlying topological space of a Noetherian scheme is Noetherian. Hence by Exercise 3.6.U, all open subsets of a Noetherian scheme are quasicompact.

5.3.A. EXERCISE. Show that locally Noetherian schemes are quasiseparated.

5.3.B. EXERCISE. Show that a Noetherian scheme has a finite number of irreducible components. (Hint: Proposition 3.6.15.) Show that a Noetherian scheme has a finite number of connected components, each a finite union of irreducible components.

5.3.C. EXERCISE. Show that a Noetherian scheme X is integral if and only if X is nonempty and connected and all stalks $\mathcal{O}_{X,p}$ are integral domains. Thus in “good situations”, integrality is the union of local (stalks are integral domains) and global (connected) conditions. Hint: Recall that integral = irreducible + reduced (Exercise 5.2.G). If a scheme’s stalks are integral domains, then it is reduced (reducedness is a stalk-local condition, Exercise 5.2.A). If a scheme X has underlying topological space that is Noetherian, then X has finitely many irreducible components (by the previous exercise); if two of them meet at a point p , then $\mathcal{O}_{X,p}$ is not an integral domain. (You can readily extend this from Noetherian schemes to locally Noetherian schemes, by showing that a connected scheme is irreducible if and only if it is nonempty and has a cover by open irreducible subsets. But some Noetherian hypotheses are necessary, see [Stacks, tag 0568].

5.3.5. Unimportant caution. The ring of sections of a Noetherian scheme need not be Noetherian, see Exercise 20.11.H.

5.3.6. Schemes over a given field k , or more generally over a given ring A (A -schemes). You may be particularly interested in working over a particular field, such as \mathbb{C} or \mathbb{Q} , or over a ring such as \mathbb{Z} . Motivated by this, we define the notion of A -scheme, or **scheme over A** , where A is a ring, as a scheme where all the rings of sections of the structure sheaf (over all open sets) are A -algebras, and all restriction maps are maps of A -algebras. (Like some earlier notions such as quasiseparatedness, this will later in Exercise 7.3.H be properly understood as a “relative notion”; it is the data of a morphism $X \rightarrow \operatorname{Spec} A$.) Suppose now X is an A -scheme. If X can be covered by affine open sets $\operatorname{Spec} B_i$ where each B_i is a *finitely generated* A -algebra, we say that X is **locally of finite type over A** , or that it is a **locally finite type A -scheme**. (This is admittedly cumbersome terminology; it will make more sense later, once we know about morphisms in §8.3.9.) If furthermore X is quasicompact, X is (of) **finite type over A** , or a **finite type A -scheme**. Note that a scheme locally of finite type over k or \mathbb{Z} (or indeed any Noetherian ring) is locally Noetherian, and similarly a scheme of finite type over any Noetherian ring is Noetherian. As our key “geometric” examples: (i) $\operatorname{Spec} \mathbb{C}[x_1, \dots, x_n]/I$ is a finite type \mathbb{C} -scheme; and (ii) $\mathbb{P}_{\mathbb{C}}^n$ is a finite type \mathbb{C} -scheme. (The field \mathbb{C} may be replaced by an arbitrary ring A .)

5.3.D. EXERCISE.

(a) (*quasiprojective implies finite type*) If X is a quasiprojective A -scheme (Definition 4.5.9), show that X is of finite type over A . If A is furthermore assumed to

be Noetherian, show that X is a Noetherian scheme, and hence has a finite number of irreducible components.

(b) Suppose U is an open subscheme of a projective A -scheme. Show that U is locally of finite type over A . If A is Noetherian, show that U is quasicompact, and hence quasiprojective over A , and hence by (a) of finite type over A . Show this need not be true if A is not Noetherian. Better: give an example of an open subscheme of a projective A -scheme that is not quasicompact, necessarily for some non-Noetherian A . (Hint: Silly example 4.5.11.)

5.3.7. Varieties. We now make a connection to the classical language of varieties. An affine scheme that is reduced and of finite type over k is said to be an **affine variety (over k)**, or an **affine k -variety**. A reduced (quasi)projective k -scheme is a **(quasi)projective variety (over k)**, or a **(quasi)projective k -variety**. (Warning: in the literature, it is sometimes also assumed in the definition of variety that the scheme is irreducible, or that k is algebraically closed.)

5.3.E. EXERCISE.

(a) Show that $\text{Spec } k[x_1, \dots, x_n]/I$ is an affine k -variety if and only if $I \subset k[x_1, \dots, x_n]$ is a radical ideal.

(b) Suppose $I \subset k[x_0, \dots, x_n]$ is a radical graded ideal. Show that $\text{Proj } k[x_0, \dots, x_n]/I$ is a projective k -variety. (Caution: The example of $I = (x_0^2, x_0x_1, \dots, x_0x_n)$ shows that $\text{Proj } k[x_0, \dots, x_n]/I$ can be a projective k -variety without I being radical.)

We will not define varieties in general until §11.3.9; we will need the notion of separatedness first, to exclude abominations like the line with the doubled origin (Example 4.4.5). But many of the statements we will make in this section about affine k -varieties will automatically apply more generally to k -varieties.

5.3.F. EXERCISE. Show that a point of a locally finite type k -scheme is a closed point if and only if the residue field of the stalk of the structure sheaf at that point is a finite extension of k . Show that on a locally finite type k -scheme, the closed points are dense. Hint: §3.6.9. (Warning: closed points need not be dense even on quite reasonable schemes, see Exercise 3.6.J(b).)

5.3.8. Definition. The **degree of a closed point** p of a locally finite type k -scheme (e.g., a variety over k) is the degree of the field extension $\kappa(p)/k$. For example, in $\mathbb{A}_k^1 = \text{Spec } k[t]$, the point $[[p(t)]]$ ($p(t) \in k[t]$ irreducible) is $\deg p(t)$. If k is algebraically closed, the degree of every closed point is 1.

5.3.9. Proof of Proposition 5.3.3. We divide each part into (i) and (ii) following the statement of the Affine Communication Lemma 5.3.2.

Proof of Proposition 5.3.3(a). (i) If $I_1 \subsetneq I_2 \subsetneq I_3 \subsetneq \dots$ is a strictly increasing chain of ideals of A_f , then we can verify that $J_1 \subsetneq J_2 \subsetneq J_3 \subsetneq \dots$ is a strictly increasing chain of ideals of A , where

$$J_j = \{r \in A : r \in I_j\}$$

where $r \in I_j$ means “the image of r in A_f lies in I_j ”. (We think of this as $I_j \cap A$, except in general A needn’t inject into A_{f_i} .) Clearly J_j is an ideal of A . If $x/f^n \in I_{j+1} \setminus I_j$ where $x \in A$, then $x \in J_{j+1}$, and $x \notin J_j$ (or else $x(1/f)^n \in I_j$ as well).

(ii) Suppose $I_1 \subsetneq I_2 \subsetneq I_3 \subsetneq \cdots$ is a strictly increasing chain of ideals of A . Then for each $1 \leq i \leq n$,

$$I_{i,1} \subset I_{i,2} \subset I_{i,3} \subset \cdots$$

is an increasing chain of ideals in A_{f_i} , where $I_{i,j} = I_j \otimes_A A_{f_i}$. It remains to show that for each j , $I_{i,j} \subsetneq I_{i,j+1}$ for some i ; the result will then follow.

5.3.G. EXERCISE. Finish the proof of Proposition 5.3.3(a). (Hint: if $M := I_j/I_{j-1} \neq 0$, then from $M \hookrightarrow \prod M_{f_i}$ (4.1.6.1), one of the M_{f_i} is nonzero.)

Proof of Proposition 5.3.3(b). Part (i) is clear: if A is generated over B by r_1, \dots, r_n , then A_f is generated over B by $r_1, \dots, r_n, 1/f$.

(ii) Here is the idea. As the f_i generate A , we can write $1 = \sum c_i f_i$ for $c_i \in A$. We have generators of A_{f_i} : $r_{ij}/f_i^{k_{ij}}$, where $r_{ij} \in A$. I claim that $\{f_i\}_i \cup \{c_i\} \cup \{r_{ij}\}_{ij}$ generate A as a B -algebra. Here is why. Suppose you have any $r \in A$. Then in A_{f_i} , we can write r as some polynomial in the r_{ij} 's and f_i , divided by some huge power of f_i . So “in each A_{f_i} , we have described r in the desired way”, except for this annoying denominator. Now use a “partition of unity” type argument as in the proof of Theorem 4.1.2 to combine all of these into a single expression, killing the denominator. Show that the resulting expression you build still agrees with r in each of the A_{f_i} . Thus it is indeed r (by the identity axiom for the structure sheaf).

5.3.H. EXERCISE. Make this argument precise.

This concludes the proof of Proposition 5.3.3. □

5.4 Normality and factoriality

5.4.1. Normality.

We can now define a property of schemes that says that they are “not too far from smooth”, called *normality*, which will come in very handy. We will see later that “locally Noetherian normal schemes satisfy Hartogs’s Lemma” (Algebraic Hartogs’s Lemma 12.3.15 for Noetherian normal schemes): functions defined away from a set of codimension 2 extend over that set. (We saw a first glimpse of this in §4.4.2.) As a consequence, rational functions (defined in Exercise 5.2.I) that have no poles (certain sets of codimension one where the function isn’t defined) are defined everywhere. We need definitions of dimension and poles to make this precise. See §13.8.7 and §29.3.5 for the fact that “smoothness” (really, “regularity”) implies normality.

Recall that an integral domain A is **integrally closed** if the only zeros in $K(A)$ to any monic polynomial in $A[x]$ must lie in A itself. The basic example is \mathbb{Z} (see Exercise 5.4.F for a reason). We say a scheme X is **normal** if all of its stalks $\mathcal{O}_{X,p}$ are normal, i.e., are integral domains, and integrally closed in their fraction fields. (So by definition, normality is a *stalk-local* property.) As reducedness is a stalk-local property (Exercise 5.2.A), normal schemes are reduced. (One might say that a ring A is a **normal ring** if $\text{Spec } A$ is normal, thereby extending the notion of “integral closure” to rings that are not integral domains.)

5.4.A. EXERCISE. Show that integrally closed domains behave well under localization: if A is an integrally closed domain, and S is a multiplicative subset not containing 0, show that $S^{-1}A$ is an integrally closed domain. (Hint: assume that $x^n + a_{n-1}x^{n-1} + \cdots + a_0 = 0$ where $a_i \in S^{-1}A$ has a root in the fraction field. Turn this into another equation in $A[x]$ that also has a root in the fraction field.)

It is no fun checking normality at every single point of a scheme. Thanks to this exercise, we know that if A is an integrally closed domain, then $\text{Spec } A$ is normal. Also, for quasicompact or locally finite type schemes, normality can be checked at closed points, thanks to this exercise, and the fact that for such schemes, any point is a generalization of a closed point (Exercises 5.1.E and 5.3.F).

It is not true that normal schemes are integral. For example, the disjoint union of two normal schemes is normal. Thus $\text{Spec } k \amalg \text{Spec } k \cong \text{Spec}(k \times k) \cong \text{Spec } k[x]/(x(x-1))$ is normal, but its ring of global sections is not an integral domain.

5.4.B. EXERCISE. Show that a Noetherian scheme is normal if and only if it is the finite disjoint union of integral Noetherian normal schemes. (Hint: Exercise 5.3.C.)

We are close to proving a useful result in commutative algebra, so we may as well go all the way.

5.4.2. Proposition. — *If A is an integral domain, then the following are equivalent.*

- (i) A is integrally closed.
- (ii) $A_{\mathfrak{p}}$ is integrally closed for all prime ideals $\mathfrak{p} \subset A$.
- (iii) $A_{\mathfrak{m}}$ is integrally closed for all maximal ideals $\mathfrak{m} \subset A$.

Proof. Exercise 5.4.A shows that integral closedness is preserved by localization, so (i) implies (ii). Clearly (ii) implies (iii).

It remains to show that (iii) implies (i). This argument involves a pretty construction that we will use again. Suppose A is not integrally closed. We show that there is some \mathfrak{m} such that $A_{\mathfrak{m}}$ is also not integrally closed. Suppose

$$(5.4.2.1) \quad x^n + a_{n-1}x^{n-1} + \cdots + a_0 = 0$$

(with $a_i \in A$) has a solution s in $K(A) \setminus A$. Let I be the **ideal of denominators** of s :

$$(5.4.2.2) \quad I := \{r \in A : rs \in A\}.$$

(Note that I is clearly an ideal of A .) Now $I \neq A$, as $1 \notin I$. Thus there is some maximal ideal \mathfrak{m} containing I . Then $s \notin A_{\mathfrak{m}}$, so equation (5.4.2.1) in $A_{\mathfrak{m}}[x]$ shows that $A_{\mathfrak{m}}$ is not integrally closed as well, as desired. \square

5.4.C. UNIMPORTANT EXERCISE. If A is an integral domain, show that $A = \bigcap A_{\mathfrak{m}}$, where the intersection runs over all maximal ideals of A . (We won't use this exercise, but it gives good practice with the ideal of denominators.)

5.4.D. UNIMPORTANT EXERCISE RELATING TO THE IDEAL OF DENOMINATORS. One might naively hope from experience with unique factorization domains that the ideal of denominators is principal. This is not true. As a counterexample,

consider our new friend $A = k[w, x, y, z]/(wz - xy)$ (which we first met in Example 4.4.12, and which we will later recognize as the cone over the quadric surface), and $w/y = x/z \in K(A)$. Show that the ideal of denominators of this element of $K(A)$ is (y, z) .

We will see that the I in the above exercise is not principal (Exercise 13.1.D — you may be able to show it directly, using the fact that I is a homogeneous ideal of a graded ring). But we will also see that in good situations (Noetherian, normal), the ideal of denominators is “pure codimension 1” — this is the content of Algebraic Hartogs’s Lemma 12.3.15. In its proof, §12.3.17, we give a geometric interpretation of the ideal of denominators.

5.4.3. Factoriality.

If all the stalks of a scheme X are unique factorization domains, we say that X is **factorial**. (Unimportant remark: the locus of points on an affine variety over an algebraically closed field that are factorial is an open subset, [BGS, p. 1].)

5.4.E. EXERCISE. Show that any nonzero localization of a unique factorization domain is a unique factorization domain.

5.4.4. Thus if A is a unique factorization domain, then $\text{Spec } A$ is factorial. The converse need not hold, see Exercise 5.4.N. In fact, we will see that elliptic curves are factorial, yet *no* affine open set is the Spec of a unique factorization domain, §20.11.1. Hence one can show factoriality by finding an appropriate affine cover, but there need not *be* such a cover of a factorial scheme. (One might reasonably call a ring A such that $\text{Spec } A$ is factorial, a *factorial ring*; this is a strictly weaker notion than unique factorization domain. We won’t need this terminology.)

5.4.5. Remark: *How to check if a ring is a unique factorization domain.* There are very few means of checking that a Noetherian integral domain is a unique factorization domain. Some useful ones are: (0) elementary means: rings with a Euclidean algorithm such as \mathbb{Z} , $k[t]$, and $\mathbb{Z}[i]$; polynomial rings over a unique factorization domain, by Gauss’s Lemma (see e.g., [Lan, IV.2.3]). (1) the localization of a unique factorization domain is also a unique factorization domain (Exercise 5.4.E). (2) height 1 prime ideals are principal (Proposition 12.3.5). (3) normal and $\text{Cl} = 0$ (Exercise 15.2.T). (4) Nagata’s Lemma (Exercise 15.2.U). (Caution: even if A is a unique factorization domain, $A[[t]]$ need not be; see [Mat2, p. 165]. The first example, due to P. Salmon, was $A = k(u)[[x, y, z]]/(z^2 + x^3 + uz^6)$, see [Sa, Dan].)

5.4.6. Factoriality implies normality. One reason we like factoriality is that it implies normality.

5.4.F. IMPORTANT EXERCISE. Show that unique factorization domains are integrally closed. Hence factorial schemes are normal, and if A is a unique factorization domain, then $\text{Spec } A$ is normal. (However, rings can be integrally closed without being unique factorization domains, as we will see in Exercise 5.4.L. Another example is given without proof in Exercise 5.4.N; in that example, Spec of the ring is factorial. A variation on Exercise 5.4.L will show that schemes can be normal without being factorial, see Exercise 13.1.E.)

5.4.7. Examples.

5.4.G. EASY EXERCISE. Show that the following schemes are normal: \mathbb{A}_k^n , \mathbb{P}_k^n , $\text{Spec } \mathbb{Z}$. (As usual, k is a field. Although it is true that if A is integrally closed then $A[x]$ is as well — see [Bo, Ch. 5, §1, no. 3, Cor. 2] or [E, Ex. 4.18] — this is not an easy fact, so do not use it here.)

5.4.H. HANDY EXERCISE (YIELDING MANY ENLIGHTENING EXAMPLES LATER). Suppose A is a unique factorization domain with 2 invertible, and $z^2 - f$ is irreducible in $A[z]$.

(a) Show that if $f \in A$ has no repeated prime factors, then $\text{Spec } A[z]/(z^2 - f)$ is normal. Hint: $B := A[z]/(z^2 - f)$ is an integral domain, as $(z^2 - f)$ is prime in $A[z]$. Suppose we have monic $F(T) \in B[T]$ so that $F(T) = 0$ has a solution α in $K(B) \setminus K(A)$. Then by replacing $F(T)$ by $\bar{F}(T)F(T)$, we can assume $F(T) \in A[T]$. Also, $\alpha = g + hz$ where $g, h \in K(A)$. Now α is the solution of $Q(T) = 0$ for monic $Q(T) = T^2 - 2gT + (g^2 - h^2f) \in K(A)[T]$, so we can factor $F(T) = P(T)Q(T)$ in $K(A)[T]$. By Gauss's lemma, $2g, g^2 - h^2f \in A$. Say $g = r/2, h = s/t$ (s and t have no common factors, $r, s, t \in A$). Then $g^2 - h^2f = (r^2t^2 - 4s^2f)/4t^2$. Then t is invertible.

(b) Show that if $f \in A$ has repeated prime factors, then $\text{Spec } A[z]/(z^2 - f)$ is *not* normal.

5.4.I. EXERCISE. Show that the following schemes are normal:

- (a) $\text{Spec } \mathbb{Z}[x]/(x^2 - n)$ where n is a square-free integer congruent to 3 modulo 4. Caution: the hypotheses of Exercise 5.4.H do not apply, so you will have to do this directly. (Your argument may also show the result when 3 is replaced by 2. A similar argument shows that $\mathbb{Z}[(1+\sqrt{n})/2]$ is integrally closed if $n \equiv 1 \pmod{4}$ is square-free.)
- (b) $\text{Spec } k[x_1, \dots, x_n]/(x_1^2 + x_2^2 + \dots + x_m^2)$ where $\text{char } k \neq 2, n \geq m \geq 3$.
- (c) $\text{Spec } k[w, x, y, z]/(wz - xy)$ where $\text{char } k \neq 2$. This is our cone over a quadric surface example from Example 4.4.12 and Exercise 5.4.D. Hint: Exercise 5.4.J may help. (The result also holds for $\text{char } k = 2$, but don't worry about this.)

5.4.J. EXERCISE (DIAGONALIZING QUADRATIC FORMS). Suppose k is an algebraically closed field of characteristic not 2. (The hypothesis that k is algebraically closed is not necessary, so feel free to deal with this more general case.)

(a) Show that any quadratic form in n variables can be “diagonalized” by changing coordinates to be a sum of at most n squares (e.g., $uw - v^2 = ((u+w)/2)^2 + (i(u-w)/2)^2 + (iv)^2$), where the linear forms appearing in the squares are linearly independent. (Hint: use induction on the number of variables, by “completing the square” at each step.)

(b) Show that the number of squares appearing depends only on the quadratic. For example, $x^2 + y^2 + z^2$ cannot be written as a sum of two squares. (Possible approach: given a basis x_1, \dots, x_n of the linear forms, write the quadratic form as

$$\begin{pmatrix} x_1 & \cdots & x_n \end{pmatrix} M \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}$$

where M is a symmetric matrix. Determine how M transforms under a change of basis, and show that the rank of M is independent of the choice of basis.)

The **rank** of the quadratic form is the number of (“linearly independent”) squares needed. If the number of squares equals the number of variables, the quadratic form is said to be **full rank** or (of) **maximal rank**.

5.4.K. EASY EXERCISE (RINGS CAN BE INTEGRALLY CLOSED BUT NOT UNIQUE FACTORIZATION DOMAINS, ARITHMETIC VERSION). Show that $\mathbb{Z}[\sqrt{-5}]$ is integrally closed but not a unique factorization domain. (Hints: Exercise 5.4.I(a) and $2 \times 3 = (1 + \sqrt{-5})(1 - \sqrt{-5})$.)

5.4.L. EASY EXERCISE (RINGS CAN BE INTEGRALLY CLOSED BUT NOT UNIQUE FACTORIZATION DOMAINS, GEOMETRIC VERSION). Suppose $\text{char } k \neq 2$. Let $A = k[w, x, y, z]/(wz - xy)$, so $\text{Spec } A$ is the cone over the smooth quadric surface (cf. Exercises 4.4.12 and 5.4.D).

- (a) Show that A is integrally closed. (Hint: Exercises 5.4.I(c) and 5.4.J.)
- (b) Show that A is not a unique factorization domain. (Clearly $wz = xy$. But why are w, x, y , and z irreducible? Hint: A is a graded integral domain. Show that if a homogeneous element factors, the factors must be homogeneous.)

The previous two exercises look similar, but there is a difference. The cone over the quadric surface is normal (by Exercise 5.4.L) but not factorial; see Exercise 13.1.E. On the other hand, $\text{Spec } \mathbb{Z}[\sqrt{-5}]$ is factorial — all of its stalks are unique factorization domains. (You will later be able to show this by showing that $\mathbb{Z}[\sqrt{-5}]$ is a Dedekind domain, §13.5.15, whose stalks are necessarily unique factorization domains by Theorem 13.5.8(f).)

5.4.M. EXERCISE. Suppose A is a k -algebra, and l/k is a finite extension of fields. (Your proof might not use finiteness; this hypothesis is included to avoid distraction by infinite-dimensional vector spaces.) Show that if $A \otimes_k l$ is a normal integral domain, then A is a normal integral domain as well. (Although we won’t need this, a version of the converse is true if l/k is separable, [Gr-EGA, IV₂.6.14.2].) Hint: fix a k -basis for l , $b_1 = 1, \dots, b_d$. Explain why $1 \otimes b_1, \dots, 1 \otimes b_d$ forms a free A -basis for $A \otimes_k l$. Explain why we have injections

$$\begin{array}{ccc} A & \longrightarrow & K(A) \\ \downarrow & & \downarrow \\ A \otimes_k l & \longrightarrow & K(A) \otimes_k l. \end{array}$$

Show that $K(A) \otimes_k l = K(A \otimes_k l)$. (Idea: $A \otimes_k l \subset K(A) \otimes_k l \subset K(A \otimes_k l)$. Why is $K(A) \otimes_k l$ a field?) Show that $(A \otimes_k l) \cap K(A) = A$. Now assume $P(T) \in A[T]$ is monic and has a root $\alpha \in K(A)$, and proceed from there.

5.4.N. EXERCISE (UFD-NESS IS NOT AFFINE-LOCAL). Let $A = (\mathbb{Q}[x, y]_{x^2+y^2})_0$ denote the homogeneous degree 0 part of the ring $\mathbb{Q}[x, y]_{x^2+y^2}$. In other words, it consists of quotients $f(x, y)/(x^2 + y^2)^n$, where f has pure degree $2n$. Show that the distinguished open sets $D(\frac{x^2}{x^2+y^2})$ and $D(\frac{y^2}{x^2+y^2})$ cover $\text{Spec } A$. (Hint: the sum of those two fractions is 1.) Show that $A_{\frac{x^2}{x^2+y^2}}$ and $A_{\frac{y^2}{x^2+y^2}}$ are unique factorization

domains. (Hint: show that both rings are isomorphic to $\mathbb{Q}[t]_{t^2+1}$; this is a localization of the unique factorization domain $\mathbb{Q}[t]$.) Finally, show that A is not a unique factorization domain. Possible hint:

$$\left(\frac{xy}{x^2+y^2}\right)^2 = \left(\frac{x^2}{x^2+y^2}\right)\left(\frac{y^2}{x^2+y^2}\right).$$

(This is generalized in Exercise 15.2.L. It is also related to Exercise 15.2.Q.)

Number theorists may prefer the example of Exercise 5.4.K: $\mathbb{Z}[\sqrt{-5}]$ is not a unique factorization domain, but it turns out that you can cover it with two affine open subsets $D(2)$ and $D(3)$, each corresponding to unique factorization domains. (For number theorists: to show that $\mathbb{Z}[\sqrt{-5}]_2$ and $\mathbb{Z}[\sqrt{-5}]_3$ are unique factorization domains, first show that the class group of $\mathbb{Z}[\sqrt{-5}]$ is $\mathbb{Z}/2$ using the geometry of numbers, as in [Ar4, Ch. 11, Thm. 7.9]. Then show that the ideals $(1 + \sqrt{-5}, 2)$ and $(1 + \sqrt{-5}, 3)$ are not principal, using the usual norm in \mathbb{C} .) The ring $\mathbb{Z}[\sqrt{-5}]$ is an example of a Dedekind domain, as we will discuss in §13.5.15.

5.4.8. Remark. For an example of k -algebra A that is not a unique factorization domain, but becomes one after a particular field extension, see Exercise 15.2.N.

Rings are to modules as schemes are to ...

6.1 Quasicoherent sheaves

When considering the notion of rings for the first time, one is quickly led to define the notion of module. For example, any ring morphism $R \rightarrow S$ expresses S as an R -module. An ideal I of a ring R is also a module, as is the quotient R/I ; and the exact sequence $0 \rightarrow I \rightarrow R \rightarrow R/I \rightarrow 0$ is a first example of the utility of the abelian category structure of Mod_R .

With a scheme X , you may think that we already have the right analog of modules: the category of \mathcal{O}_X -modules Mod_X , which also forms an abelian category (§2.6). But the right analog for a scheme X turns out to be a better-behaved subset (subcategory!) of Mod_X , called *quasicoherent \mathcal{O}_X -modules*, or *quasicoherent sheaves*.

Given an A -module M , we defined an \mathcal{O} -module \widetilde{M} on $\text{Spec } A$ in §4.1.5 — the sections over $D(f)$ were M_f . (Now is a good time to complete the exercises in §4.1.5 if you have not done so already.) These are our “local models” for quasicoherent sheaves. In the same way that a scheme is defined by “gluing together rings”, a quasicoherent sheaf over that scheme is obtained by “gluing together modules over those rings”.

6.1.1. Definition: *the category of quasicoherent sheaves.* If X is a scheme, then an \mathcal{O}_X -module \mathcal{F} is a **quasicoherent sheaf** if for every affine open subset $\text{Spec } A \subset X$, $\mathcal{F}|_{\text{Spec } A} \cong \widetilde{M}$ for some A -module M . We make quasicoherent sheaves on a scheme X into category QCoh_X by taking the morphisms to be morphisms as \mathcal{O}_X -modules.

6.1.2. Theorem. — *Let X be a scheme, and \mathcal{F} an \mathcal{O}_X -module. Suppose P is the property of affine open subschemes $\text{Spec } A \subset X$ that $\mathcal{F}|_{\text{Spec } A} \cong \widetilde{M}$ for some A -module M . Then P satisfies the two hypotheses of the Affine Communication Lemma 5.3.2.*

Thus to check quasicoherence of an \mathcal{O} -module, it suffices to check a collection of affine open sets covering X . For example, \widetilde{M} is a quasicoherent sheaf on $\text{Spec } A$. Exercise 4.1.G shows that the category Mod_A and the category $\text{QCoh}_{\text{Spec } A}$ are essentially the same (we have described an equivalence $\text{Mod}_A \xrightarrow{\sim} \text{QCoh}_{\text{Spec } A}$).

Proof. As usual, the first hypothesis of the Affine Communication Lemma 5.3.2 is easier: if $\text{Spec } A$ has property P , then so does the distinguished open $\text{Spec } A_f$: if M is an A -module, then $\widetilde{M}|_{\text{Spec } A_f} \cong \widetilde{M_f}$ as sheaves of $\mathcal{O}_{\text{Spec } A_f}$ -modules (both sides agree on the level of distinguished open sets and their restriction maps).

We next show the second hypothesis of the Affine Communication Lemma 5.3.2. We are given:

- an $\mathcal{O}_{\text{Spec } A}$ -module \mathcal{F} ,
- elements f_1, \dots, f_n generating A ,
- A_{f_i} -modules M_i ,
- and isomorphisms $\phi_i : \mathcal{F}|_{D(f_i)} \xrightarrow{\sim} \widetilde{M}_i$.

The isomorphisms

$$\begin{array}{ccc} & \mathcal{F}|_{D(f_i) \cap D(f_j)} & \\ \swarrow \sim & & \searrow \sim \\ \widetilde{M}_i|_{D(f_i) \cap D(f_j)} & & \widetilde{M}_j|_{D(f_i) \cap D(f_j)} \end{array}$$

yield (by taking sections over $D(f_i) \cap D(f_j)$) isomorphisms $\phi_{ij} : (M_i)_{f_j} \xrightarrow{\sim} (M_j)_{f_i}$ of $A_{f_i f_j}$ -modules, satisfying the cocycle condition (2.5.4.1). As suggestive notation, let $M_{ij} := (M_i)_{f_j}$ (identified with $(M_j)_{f_i}$ via ϕ_{ij}).

We seek an A -module M with an isomorphism $\widetilde{M} \xleftarrow{\sim} \mathcal{F}$, and (4.1.6.1) suggests that we should define M by

$$(6.1.2.1) \quad 0 \longrightarrow M \longrightarrow M_1 \times \cdots \times M_n \xrightarrow{\gamma} M_{12} \times M_{13} \times \cdots \times M_{(n-1)n}.$$

(Equation (4.1.6.1) shows how the map γ should be defined, with appropriate signs.) Translation: $M = \ker(\gamma)$. If we can describe an isomorphism $\widetilde{M}|_{D(f_i)} \xrightarrow{\sim} \widetilde{M}_i$ on each $D(f_i)$ (or equivalently, an isomorphism $M_{f_i} \xrightarrow{\sim} M_i$ of A_{f_i} -modules), such that triangle

$$\begin{array}{ccc} & \widetilde{M}|_{D(f_i) \cap D(f_j)} & \\ \swarrow \sim & & \searrow \sim \\ \widetilde{M}_i|_{D(f_i) \cap D(f_j)} & \xrightarrow[\sim]{\widetilde{\phi}_{ij}} & \widetilde{M}_j|_{D(f_i) \cap D(f_j)} \end{array}$$

of sheaves on $\text{Spec } A_{f_i f_j}$ commutes (for all i, j), then Exercise 2.5.E (on gluing sheaves) gives our desired isomorphism $\widetilde{M} \xleftarrow{\sim} \mathcal{F}$.

By Exercise 4.1.G, we need only describe an isomorphism $M_{f_i} \xrightarrow{\sim} M_i$ such that the triangles

$$(6.1.2.2) \quad \begin{array}{ccc} & M_{f_i f_j} & \\ \swarrow \sim & & \searrow \sim \\ (M_i)_{f_j} & \xrightarrow{\phi_{ij}} & (M_j)_{f_i} \end{array}$$

commute.

For notational convenience, we assume $i = 1$. Because localization is exact, from (6.1.2.1) we have an exact sequence

$$0 \longrightarrow M_{f_1} \longrightarrow (M_1)_{f_1} \times \cdots \times (M_n)_{f_1} \xrightarrow{\gamma_{f_1}} (M_{12})_{f_1} \times (M_{13})_{f_1} \times \cdots \times (M_{(n-1)n})_{f_1}$$

We will show that

(6.1.2.3)

$$0 \longrightarrow M_1 \xrightarrow{\beta} (M_1)_{f_1} \times \cdots \times (M_n)_{f_1} \xrightarrow{\gamma_{f_1}} (M_{12})_{f_1} \times (M_{13})_{f_1} \times \cdots \times (M_{(n-1)n})_{f_1}$$

is an exact sequence (do you see how β should be defined?), yielding our desired isomorphism $M_{f_1} \xrightarrow{\sim} M_1$. We rewrite (6.1.2.3) as

(6.1.2.4)

$$0 \longrightarrow M_1 \longrightarrow (M_1)_{f_1} \times (M_1)_{f_2} \times \cdots \times (M_1)_{f_n} \xrightarrow{\alpha} (M_1)_{f_1 f_2} \times \cdots \times (M_1)_{f_1 f_n} \times (M_1)_{f_2 f_3} \times \cdots \times (M_1)_{f_{n-1} f_n}.$$

But this is precisely the exact sequence (4.1.6.1), except the ring A is replaced by A_{f_1} , and the module M is replaced by M_1 .

6.1.A. EXERCISE. Tie up the last loose end: Why does the triangle (6.1.2.2) commute?

□

★ *Remark for experts:* The proof of Theorem 6.1.2, phrased slightly more carefully, shows that quasicohherent sheaves satisfy faithfully flat descent.

6.1.3. Stray concluding topics.

6.1.B. UNIMPORTANT EXERCISE (NOT EVERY \mathcal{O}_X -MODULE IS A QUASICOHERENT SHEAF).

(a) Suppose $X = \operatorname{Spec} k[t]$. Let \mathcal{F} be the skyscraper sheaf supported at the origin $[(t)]$, with group $k(t)$ and the usual $k[t]$ -module structure. Show that this is an \mathcal{O}_X -module that is not a quasicohherent sheaf. (More generally, if X is an integral scheme, and $p \in X$ is not the generic point, we could take the skyscraper sheaf at p with group the function field of X . Except in silly circumstances, this sheaf won't be quasicohherent.) See Exercises 9.1.G and 6.2.H for more (pathological) examples of \mathcal{O}_X -modules that are not quasicohherent.

(b) Suppose $X = \operatorname{Spec} k[t]$. Let \mathcal{F} be the skyscraper sheaf supported at the generic point $[(0)]$, with group $k(t)$. Give this the structure of an \mathcal{O}_X -module. Show that this is a quasicohherent sheaf. Describe the restriction maps in the distinguished topology of X . (Remark: Your argument will apply more generally, for example when X is an integral scheme with generic point η , and \mathcal{F} is the skyscraper sheaf $i_{\eta,*}K(X)$.)

6.1.4. Torsion-free sheaves (a stalk-local condition) and torsion sheaves. An A -module M is said to be **torsion-free** if $am = 0$ implies that either a is a zerodivisor in A or $m = 0$.

In the case where A is an integral domain, which is basically the only context in which we will use this concept, the definition of torsion-freeness can be restated as $am = 0$ only if $a = 0$ or $m = 0$. In this case, the **torsion submodule** of M , denoted M_{tors} , consists of those elements of M annihilated by some nonzero element of A . (If A is not an integral domain, this construction needn't yield an A -module.) Clearly M is torsion-free if and only if $M_{\text{tors}} = 0$. We say a module M over an integral domain A is **torsion** if $M = M_{\text{tors}}$; this is equivalent to $M \otimes_A K(A) = 0$.

If X is a scheme, then an \mathcal{O}_X -module \mathcal{F} is said to be **torsion-free** if \mathcal{F}_p is a torsion-free $\mathcal{O}_{X,p}$ -module for all p .

6.1.5. Definition: torsion quasicohherent sheaves on reduced schemes. Motivated by the definition of M_{tors} above, we say that a quasicohherent sheaf on a *reduced* scheme is **torsion** if its stalk at the generic point of every irreducible component is 0. We will mainly use this for coherent sheaves on regular curves, where this notion is very simple indeed (see Exercise 14.3.G(b)), but in the literature it comes up in more general situations.

6.2 Characterizing quasicohherence using the distinguished affine base

Because quasicohherent sheaves are locally of a very special form, in order to “know” a quasicohherent sheaf, we need only know what the sections are over every affine open set, and how to restrict sections from an affine open set U to a *distinguished* affine open subset of U . We make this precise by defining what we will call the *distinguished affine base* of the Zariski topology — not a base in the usual sense. This is a refinement of the notion of a *sheaf on base*. The point of this discussion is to give a useful characterization of quasicohherence, but you may wish to just jump to §6.2.3.

The open sets of the distinguished affine base are the affine open subsets of X . We have already observed that this forms a base. But forget that fact. We like distinguished open sets $\text{Spec } A_f \hookrightarrow \text{Spec } A$, and we don’t really understand open embeddings of one random affine open subset in another. So we just remember the “nice” inclusions.

6.2.1. Definition. The **distinguished affine base** of a scheme X is the data of the affine open sets and the distinguished inclusions.

In other words, we remember only some of the open sets (the affine open sets), and *only some of the morphisms between them* (the distinguished morphisms). For experts: if you think of a topology as a category (the category of open sets), we have described a subcategory.

We can define a sheaf on the distinguished affine base in the obvious way: we have a set (or abelian group, or ring) for each affine open set, and we know how to restrict to distinguished open sets. (You should think through the statement of the identity and gluability axioms yourself.)

Given a sheaf \mathcal{F} on X , we get a sheaf on the distinguished affine base. You can guess where we are going: we will show that all the information of the sheaf is contained in the information of the sheaf on the distinguished affine base.

As a warm-up, we can recover stalks as follows. (We will be implicitly using only the following fact. We have a collection of open subsets, and *some* inclusions among these subsets, such that if we have any $p \in U, V$ where U and V are in our collection of open sets, there is some W containing p , and contained in U and V such that $W \hookrightarrow U$ and $W \hookrightarrow V$ are both in our collection of inclusions. In the case we are considering here, this is the key Proposition 5.3.1 that given any two affine open sets $\text{Spec } A, \text{Spec } B$ in X , $\text{Spec } A \cap \text{Spec } B$ could be covered by affine open sets that were simultaneously distinguished in $\text{Spec } A$ and $\text{Spec } B$. In fancy language: the category of affine open sets, and distinguished inclusions, forms a filtered set.)

The stalk \mathcal{F}_p is the colimit $\text{colim}(f \in \mathcal{F}(U))$ where the colimit is over all open sets contained in X . We compare this to $\text{colim}(f \in \mathcal{F}(U))$ where the colimit is over all affine open sets, and all distinguished inclusions. You can check that the elements of one correspond to elements of the other. (Think carefully about this!)

6.2.A. EXERCISE. Show that a section of a sheaf on the distinguished affine base is determined by the section's germs.

6.2.2. Theorem. —

- (a) *A sheaf on the distinguished affine base \mathcal{F}^b determines a unique (up to unique isomorphism) sheaf \mathcal{F} which when restricted to the affine base is \mathcal{F}^b . (Hence if you start with a sheaf, and take the sheaf on the distinguished affine base, and then take the induced sheaf, you get the sheaf you started with.)*
- (b) *A morphism of sheaves on a distinguished affine base uniquely determines a morphism of sheaves.*
- (c) *An \mathcal{O}_X -module “on the distinguished affine base” yields an \mathcal{O}_X -module.*

This proof is identical to our argument of §2.5 showing that sheaves are (essentially) the same as sheaves on a base, using the “sheaf of compatible germs” construction. The main reason for repeating it is to let you see that all that is needed is for the open sets to form a filtered set (or in the current case, that the category of open sets and distinguished inclusions is filtered).

For experts: (a) and (b) are describing an equivalence of categories between sheaves on the Zariski topology of X and sheaves on the distinguished affine base of X .

Proof. (a) Suppose \mathcal{F}^b is a sheaf on the distinguished affine base. Then we can define stalks.

For any open set U of X , define the sheaf of compatible germs

$$\mathcal{F}(U) := \{(f_p \in \mathcal{F}_p^b)_{p \in U} : \text{for all } p \in U, \text{ there exists } U_p \text{ with } p \in U_p \subset U, \\ s \in \mathcal{F}^b(U_p) \text{ such that } s_q = f_q \text{ for all } q \in U_p\}$$

where each U_p is in our base, and s_q means “the germ of s at q ”. (As usual, those who want to worry about the empty set are welcome to.)

This really is a sheaf: convince yourself that we have restriction maps, identity, and gluability, really quite easily.

I next claim that if U is in our base, that $\mathcal{F}(U) = \mathcal{F}^b(U)$. We clearly have a map $\mathcal{F}^b(U) \rightarrow \mathcal{F}(U)$. This is an isomorphism on stalks, and hence an isomorphism by Exercise 2.4.D.

6.2.B. EXERCISE. Prove (b) (cf. Exercise 2.5.C).

6.2.C. EXERCISE. Prove (c) (cf. Remark 2.5.3).

The proof of Theorem 6.2.2 is now complete. \square

6.2.3. A characterization of quasicoherent sheaves in terms of distinguished inclusions. We use this perspective to give a useful characterization of quasicoherent sheaves among \mathcal{O}_X -modules. Suppose \mathcal{F} is an \mathcal{O}_X -module, and $\text{Spec } A_f \hookrightarrow \text{Spec } A \subset X$ is a distinguished open subscheme of an affine open subscheme of

X. Let $\phi: \Gamma(\operatorname{Spec} A, \mathcal{F}) \rightarrow \Gamma(\operatorname{Spec} A_f, \mathcal{F})$ be the restriction map. The source of ϕ is an A -module, and the target is an A_f -module, so by the universal property of localization (Exercise 1.3.D), ϕ naturally factors as:

$$\begin{array}{ccc} \Gamma(\operatorname{Spec} A, \mathcal{F}) & \xrightarrow{\quad \phi \quad} & \Gamma(\operatorname{Spec} A_f, \mathcal{F}) \\ & \searrow \otimes_A A_f \quad \nearrow \alpha & \\ & \Gamma(\operatorname{Spec} A, \mathcal{F})_f & \end{array}$$

6.2.D. VERY IMPORTANT EXERCISE. Show that an \mathcal{O}_X -module \mathcal{F} is quasicohherent if and only if for each such distinguished $\operatorname{Spec} A_f \hookrightarrow \operatorname{Spec} A$, α is an isomorphism.

Thus a quasicohherent sheaf is (equivalent to) the data of one module for each affine open subset (a module over the corresponding ring), such that the module over a distinguished open set $\operatorname{Spec} A_f$ is given by localizing the module over $\operatorname{Spec} A$.

6.2.E. EXERCISE (GOOD PRACTICE: THE SHEAF OF NILPOTENTS). If A is a ring, and $f \in A$, show that $\mathfrak{N}(A_f) \cong \mathfrak{N}(A)_f$. Use this to define/construct the quasicohherent **sheaf of nilpotents** on any scheme X . This is an example of an ideal sheaf (of \mathcal{O}_X).

6.2.4. Tensor products. Tensor products can also be cleanly described in terms of affine open sets.

6.2.F. EXERCISE. If \mathcal{F} and \mathcal{G} are quasicohherent sheaves on a scheme X , show that $\mathcal{F} \otimes \mathcal{G}$ is a quasicohherent sheaf described by the following information: If $\operatorname{Spec} A \subset X$ is an affine open subset, and $\Gamma(\operatorname{Spec} A, \mathcal{F}) = M$ and $\Gamma(\operatorname{Spec} A, \mathcal{G}) = N$, then $\Gamma(\operatorname{Spec} A, \mathcal{F} \otimes \mathcal{G}) = M \otimes_A N$, and the restriction map $\Gamma(\operatorname{Spec} A, \mathcal{F} \otimes \mathcal{G}) \rightarrow \Gamma(\operatorname{Spec} A_f, \mathcal{F} \otimes \mathcal{G})$ is precisely the localization map $M \otimes_A N \rightarrow (M \otimes_A N)_f \cong M_f \otimes_{A_f} N_f$. (We are using the algebraic fact that $(M \otimes_A N)_f \cong M_f \otimes_{A_f} N_f$. You can prove this by universal property if you want, or by using the explicit construction.)

Note that thanks to the machinery behind the distinguished affine base, sheafification is taken care of. This is a feature we will use often: constructions involving quasicohherent sheaves that involve sheafification for general sheaves don't require sheafification when considered on the distinguished affine base. Along with the fact that injectivity, surjectivity, kernels, tensor products and so on may be computed on affine opens, this is the reason that it is particularly convenient to think about quasicohherent sheaves in terms of affine open sets. (There is a slight caveat in the case of $\mathcal{H}om$, see Exercise 14.3.A.)

6.2.5. Definition. Given a section s of \mathcal{F} and a section t of \mathcal{G} , we have a section $s \otimes t$ of $\mathcal{F} \otimes \mathcal{G}$. If \mathcal{F} is an invertible sheaf, this section is often denoted st .

6.2.6. X_f and the Qcqs lemma. The next result will be useful in the future in showing certain constructions yield quasicohherent sheaves.

6.2.7. Definition: X_f . If X is a scheme, and $f \in \Gamma(X, \mathcal{O}_X)$ is a function, let $X_f \subset X$ be the subset on which f doesn't vanish. By Exercise 4.3.F(a), X_f is open. (We avoid the notation $D(f)$ to avoid any suggestion that X is affine.)

6.2.8. The Qcqs Lemma. — Suppose X is a quasicompact and quasiseparated scheme, \mathcal{F} is a quasicoherent sheaf on X , and $f \in \Gamma(X, \mathcal{O}_X)$ is a function on X . Then the restriction map

$$\mathrm{res}_{X_f \subset X}: \Gamma(X, \mathcal{F}) \longrightarrow \Gamma(X_f, \mathcal{F})$$

is precisely localization: there is an isomorphism $\Gamma(X, \mathcal{F})_f \xrightarrow{\sim} \Gamma(X_f, \mathcal{F})$ making the following diagram commute.

$$\begin{array}{ccc} \Gamma(X, \mathcal{F}) & \xrightarrow{\mathrm{res}_{X_f \subset X}} & \Gamma(X_f, \mathcal{F}) \\ & \searrow \otimes_{\Gamma(X, \mathcal{O}_X)} (\Gamma(X, \mathcal{O}_X)_f) & \nearrow \sim \\ & \Gamma(X, \mathcal{F})_f & \end{array}$$

6.2.G. IMPORTANT EXERCISE. Prove the Qcqs Lemma 6.2.8. Hint: The “quasicompact quasiseparated” hypothesis translates precisely to “covered by a finite number of affine open sets, the pairwise intersection of which is also covered by a finite number of affine open sets”. Apply the exact functor $\otimes_{\Gamma(X, \mathcal{O}_X)} \Gamma(X, \mathcal{O}_X)_f$ to the exact sequence

$$0 \longrightarrow \Gamma(X, \mathcal{F}) \longrightarrow \oplus_i \Gamma(U_i, \mathcal{F}) \longrightarrow \oplus \Gamma(U_{ijk}, \mathcal{F}),$$

where the U_i form a finite affine cover of X and U_{ijk} form a finite affine cover of $U_i \cap U_j$.

6.2.H. LESS IMPORTANT EXERCISE. Give a counterexample to show that the Qcqs Lemma 6.2.8 is false without the quasicompactness hypothesis. (Possible hint: take an infinite disjoint union of affine schemes. The key idea is that infinite direct products do not commute with localization.)

6.2.9. ★★ Grothendieck topologies. The distinguished affine base isn't a topology in the usual sense — the union of two affine sets isn't necessarily affine, for example. It is however a first new example of a generalization of a topology — the notion of a **site** or a **Grothendieck topology**. We give the definition to satisfy the curious, but we certainly won't use this notion. (For a clean statement, see [Stacks, tag 00VH]; our discussion here is intended only as motivation.) The idea is that we should abstract away only those notions we need to define sheaves. We need the notion of open set, but it turns out that we won't even need an underlying set, i.e., we won't even need the notion of points! Let's think through how little we need. For our discussion of sheaves to work, we needed to know what the open sets were, and what the (allowed) inclusions were, and these should “behave well”, and in particular the data of the open sets and inclusions should form a category. (For example, the composition of an allowed inclusion with another allowed inclusion should be an allowed inclusion — in the distinguished affine base, a distinguished open set of a distinguished open set is a distinguished open set.) So we just require the data of *this category*. At this point, we can already

define presheaf (as just a contravariant functor from this category of “open sets”). We saw this idea earlier in Exercise 2.2.A.

In order to extend this definition to that of a sheaf, we need to know more information. We want two open subsets of an open set to intersect in an open set, so *we want the category to be closed under fibered products* (cf. Exercise 1.3.O). For the identity and gluability axioms, we need to know *when some open sets cover another*, so we also remember this as part of the data of a Grothendieck topology. The data of the coverings satisfy some obvious properties. Every open set covers itself (i.e., *the identity map in the category of open sets is a covering*). Coverings pull back: *if we have a map $Y \rightarrow X$, then any cover of X pulls back to a cover of Y* . Finally, *a cover of a cover should be a cover*. Such data (satisfying these axioms) is called a *Grothendieck topology* or a *site*. (There are useful variants of this definition in the literature. Again, we are following [Stacks].) We can define the notion of a sheaf on a Grothendieck topology in the usual way, with no change. A **topos** is a scary name for a category of sheaves of sets on a Grothendieck topology.

Grothendieck topologies are used in a wide variety of contexts in and near algebraic geometry. Étale cohomology (using the étale topology), a generalization of Galois cohomology, is a central tool, as are more general flat topologies, such as the smooth topology. The definition of a Deligne-Mumford or Artin stack uses the étale and smooth topology, respectively. Tate developed a good theory of non-archimedean analytic geometry over totally disconnected ground fields such as \mathbb{Q}_p using a suitable Grothendieck topology. Work in K-theory (related for example to Voevodsky’s work) uses exotic topologies.

6.3 Quasicoherent sheaves form an abelian category

Morphisms from one quasicoherent sheaf on a scheme X to another are defined to be just morphisms as \mathcal{O}_X -modules. In this way, the quasicoherent sheaves on a scheme X form a category, denoted $QCoh_X$. (By definition it is a full subcategory of $Mod_{\mathcal{O}_X}$.) We now show that quasicoherent sheaves on X form an *abelian* category.

When you show that something is an abelian category, you have to check many things, because the definition has many parts. However, if the objects you are considering lie in some ambient abelian category, then it is much easier. You have seen this idea before: there are several things you have to do to check that something is a group. But if you have a subset of group elements, it is much easier to check that it forms a subgroup.

You can look at back at the definition of an abelian category, and you will see that in order to check that a subcategory is an abelian subcategory, it suffices to check only the following:

- (i) 0 is in the subcategory
- (ii) the subcategory is closed under finite sums
- (iii) the subcategory is closed under kernels and cokernels

In our case of $QCoh_X \subset Mod_{\mathcal{O}_X}$, the first two are cheap: 0 is certainly quasicoherent, and the subcategory is closed under finite sums: if \mathcal{F} and \mathcal{G} are sheaves

on X , and over $\text{Spec } A$, $\mathcal{F} \cong \widetilde{M}$ and $\mathcal{G} \cong \widetilde{N}$, then $\mathcal{F} \oplus \mathcal{G} \cong \widetilde{M \oplus N}$ (Do you see why?), so $\mathcal{F} \oplus \mathcal{G}$ is a quasicoherent sheaf.

We now check (iii), using the characterization of Important Exercise 6.2.D. Suppose $\alpha: \mathcal{F} \rightarrow \mathcal{G}$ is a morphism of quasicoherent sheaves. Then on any affine open set U , where the morphism is given by $\beta: M \rightarrow N$, define $(\ker \alpha)(U) = \ker \beta$ and $(\text{coker } \alpha)(U) = \text{coker } \beta$. Then these behave well under inversion of a single element: if

$$0 \longrightarrow K \longrightarrow M \longrightarrow N \longrightarrow P \longrightarrow 0$$

is exact, then so is

$$0 \longrightarrow K_f \longrightarrow M_f \longrightarrow N_f \longrightarrow P_f \longrightarrow 0,$$

from which $(\ker \beta)_f \cong \ker(\beta_f)$ and $(\text{coker } \beta)_f \cong \text{coker}(\beta_f)$. Thus both of these define quasicoherent sheaves. Moreover, by checking stalks, they are indeed the kernel and cokernel of α (exactness can be checked stalk-locally). Thus the quasicoherent sheaves indeed form an abelian category.

6.3.A. EXERCISE. Show that a sequence of quasicoherent sheaves $\mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{H}$ on X is exact if and only if it is exact on every open set in any given affine cover of X . (In particular, taking sections over an affine open $\text{Spec } A$ is an exact functor from the category of quasicoherent sheaves on X to the category of A -modules. Recall that taking sections is only left-exact in general, see §2.6.F.) Thus we may check injectivity or surjectivity of a morphism of quasicoherent sheaves by checking on an affine cover of our choice.

Caution: If $0 \rightarrow \mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{H} \rightarrow 0$ is an exact sequence of quasicoherent sheaves, then for any open set

$$0 \longrightarrow \mathcal{F}(U) \longrightarrow \mathcal{G}(U) \longrightarrow \mathcal{H}(U)$$

is exact, and exactness on the right is guaranteed to hold only if U is affine. (To set you up for cohomology: whenever you see left-exactness, you expect to eventually interpret this as a start of a long exact sequence. So we are expecting H^1 's on the right, and now we expect that $H^1(\text{Spec } A, \mathcal{F}) = 0$. This will indeed be the case.)

6.3.B. LESS IMPORTANT EXERCISE (CONNECTION TO ANOTHER DEFINITION, AND QUASICOHERENT SHEAVES ON RINGED SPACES IN GENERAL). Show that an \mathcal{O}_X -module \mathcal{F} on a scheme X is quasicoherent if and only if there exists an open cover by U_i such that on each U_i , $\mathcal{F}|_{U_i}$ is isomorphic to the cokernel of a map of two free sheaves:

$$\mathcal{O}_{U_i}^{\oplus I} \longrightarrow \mathcal{O}_{U_i}^{\oplus J} \longrightarrow \mathcal{F}|_{U_i} \longrightarrow 0$$

is exact. We have thus connected our definitions to the definition given at the very start of the chapter. This is the definition of a quasicoherent sheaf on a ringed space in general. It is useful in many circumstances, for example in complex analytic geometry.

6.4 Finite type quasicoherent, finitely presented, and coherent sheaves

Here are three natural finiteness conditions on an A -module M . If A is a Noetherian ring, which is the case that almost all of you will ever care about, they are all the same.

The first is the simplest: a module could be **finitely generated**. In other words, there is a surjection $A^{\oplus p} \rightarrow M \rightarrow 0$.

The second is reasonable too. It could be finitely presented — it could have a finite number of generators with a finite number of relations. Translation: there exists a **finite presentation**, i.e., an exact sequence

$$A^{\oplus q} \longrightarrow A^{\oplus p} \longrightarrow M \longrightarrow 0.$$

6.4.1. The third notion is frankly a bit surprising. We say that an A -module M is **coherent** if (i) it is finitely generated, and (ii) for any map $A^{\oplus p} \rightarrow M$ (not necessarily surjective!), the kernel is finitely generated.

6.4.2. Proposition. — *If A is Noetherian, then these three definitions are the same.*

Proof. Clearly coherent implies finitely presented, which in turn implies finitely generated. So suppose M is finitely generated. Take any $\alpha: A^{\oplus p} \rightarrow M$. Then $\ker \alpha$ is a submodule of a finitely generated module over A , and is thus finitely generated by Exercise 3.6.Y. Thus M is coherent. \square

Hence most people can think of these three notions as the same.

6.4.3. Proposition. — *The coherent A -modules form an abelian subcategory of the category of A -modules.*

The proof in general is a series of short exercises in §6.6. You should try them only if you are particularly curious.

Proof if the ring A is Noetherian. Recall from our discussion at the start of §6.3 that we must check three things:

- (i) The 0-module is coherent.
- (ii) The category of coherent modules is closed under finite sums.
- (iii) The category of coherent modules is closed under kernels and cokernels.

The first two are clear. For (iii), suppose that $f: M \rightarrow N$ is a map of finitely generated modules. Then $\operatorname{coker} f$ is finitely generated (it is the image of N), and $\ker f$ is too (it is a submodule of a finitely generated module over a Noetherian ring, Exercise 3.6.Y). \square

We now extend these notions to quasicohherent sheaves.

6.4.A. EXERCISE (FINITE GENERATION SATISFIES THE HYPOTHESES OF THE AFFINE COMMUNICATION LEMMA 5.3.2).

- (a) If $f \in A$, and M is a finitely generated A -module, show that M_f is a finitely generated A_f -module.
- (b) If $(f_1, \dots, f_n) = A$, and M_{f_i} is a finitely generated A_{f_i} -module for all i , show that M is a finitely generated A -module.

For the finite presentation case, the following result is handy, and versions apply in other situations (see Lemma 8.3.R for the “algebra” version). For the longest time I didn’t even suspect such a thing could be true.

6.4.4. Lemma: Finitely presented implies always finitely presented (module version). — Suppose M is a finitely presented A -module (for **some** surjection from a finite free module, the kernel is finitely generated), and $\phi: A^{\oplus p} \rightarrow M$ is **any** surjection. Then $\ker \phi$ is finitely generated.

Proof. Choose a finite presentation of M :

$$\bigoplus_{k=1}^r Af_k \xrightarrow{\beta} \bigoplus_{i=1}^n Ax_i \xrightarrow{\alpha} M \longrightarrow 0$$

so $\alpha(x_1), \dots, \alpha(x_n)$ generate M . Write the surjection of the hypothesis as $\phi: \bigoplus_{j=1}^p Ay_j \rightarrow M$. Choose lifts $g_j \in \bigoplus Ax_i$ of $\phi(y_j) \in M$. Then we can write

$$\begin{aligned} M &= \bigoplus Ax_i / (\beta(f_1), \dots, \beta(f_r)) \\ &= (\bigoplus Ax_i \oplus Ay_j) / (\beta(f_1), \dots, \beta(f_r), y_1 - g_1, \dots, y_p - g_p) \end{aligned}$$

As the $\{\phi(y_j)\}$ generate M , for each $i = 1, \dots, n$ we can choose $h_i \in \bigoplus_{j=1}^p Ay_j$ (i.e., a map $h: \bigoplus_{i=1}^n Ax_i \rightarrow \bigoplus_{j=1}^p Ay_j$) so that $\alpha(x_i) = \phi(h_i)$. Thus

$$\begin{aligned} M &= (\bigoplus Ax_i \oplus Ay_j) / (\beta(f_1), \dots, \beta(f_r), y_1 - g_1, \dots, y_p - g_p, x_1 - h_1, \dots, x_n - h_n) \\ &= \bigoplus Ay_j / (h(\beta(f_1)), \dots, h(\beta(f_r)), y_1 - h(g_1), \dots, y_p - h(g_p)) \end{aligned}$$

□

6.4.B. EXERCISE (FINITE PRESENTATION SATISFIES THE HYPOTHESES OF THE AFFINE COMMUNICATION LEMMA 5.3.2).

(a) If $f \in A$, and if M is a finitely presented A -module, show that M_f is a finitely presented A_f -module.

(b) If $(f_1, \dots, f_n) = A$, and M_{f_i} is a finitely presented A_{f_i} -module for all i , show that M is a finitely presented A -module. (Hint: use Exercise 6.4.A(b) to show first that M is finitely generated, so we may write $M = \operatorname{coker}(\alpha: N \hookrightarrow A^{\oplus n})$ for some submodule $N \subset A^{\oplus n}$. We wish to show that N is finitely-generated. Localize at the f_i , use Lemma 6.4.4 to show N_{f_i} is a finitely generated A_{f_i} module, then apply Exercise 6.4.A(b) to N .)

6.4.C. EXERCISE (COHERENCE SATISFIES THE HYPOTHESES OF THE AFFINE COMMUNICATION LEMMA 5.3.2).

(a) If $f \in A$, and M is a coherent A -module, show that M_f is a coherent A_f -module.

(b) If $(f_1, \dots, f_n) = A$, and M_{f_i} is a coherent A_{f_i} -module for all i , show that M is a coherent A -module. (Hint: if $\phi: A^{\oplus p} \rightarrow M$, then $(\ker \phi)_{f_i} = \ker(\phi_{f_i})$, which is finitely generated for all i . Then apply Exercise 6.4.A(b).)

6.4.5. Definition. A quasicoherent sheaf \mathcal{F} is **finite type** (resp. **finitely presented**, **coherent**) if for every affine open $\operatorname{Spec} A$, $\Gamma(\operatorname{Spec} A, \mathcal{F})$ is a finitely generated (resp. finitely presented, coherent) A -module. Note that coherent sheaves are always finite type, and that on a locally Noetherian scheme, all three notions are the same (by Proposition 6.4.2). Proposition 6.4.3 implies that the coherent sheaves on X form an abelian category, which we denote Coh_X .

By the Affine Communication Lemma 5.3.2, and Exercises 6.4.A, 6.4.B, and 6.4.C, it suffices to check “finite typeness” (resp. finite presentation, coherence) on the open sets in a single affine cover.

6.4.6. Warning. It is not uncommon in the later literature to incorrectly define coherent as finitely generated. Please only use the correct definition, as the wrong definition causes confusion. Besides doing this for the reason of honesty, it will also help you see what hypotheses are actually necessary to prove things. And that always helps you remember what the proofs are — and hence why things are true.

6.4.7. Why coherence? Proposition 6.4.3 is a good motivation for the definition of coherence: it gives a small (in a non-technical sense) abelian category in which we can think about vector bundles.

There are two sorts of people who should care about the details of this definition, rather than living in a Noetherian world where coherent means finite type. Complex geometers should care. They consider complex-analytic spaces with the classical topology. One can define the notion of coherent \mathcal{O}_X -module in a way analogous to this (see [Se1, Def. 2]). Then Oka's Theorem states that the structure sheaf of \mathbb{C}^n (hence of any complex manifold) is coherent, and this is very hard (see [GR, §2.5] or [Rem, §7.2]).

The second sort of people who should care are the sort of arithmetic people who may need to work with non-Noetherian rings, see §3.6.21, or work in non-archimedean analytic geometry.

6.4.8. Remark: Quasicoherent and coherent sheaves on ringed spaces. We will discuss quasicoherent and coherent sheaves on schemes, but they can be defined more generally (see Exercise 6.3.B for quasicoherent sheaves, and [Se1, Def. 2] for coherent sheaves). Many of the results we state will hold in greater generality, but because the proofs look slightly different, we restrict ourselves to schemes to avoid distraction.

6.4.9. ★★ Coherence is not a good notion in smooth geometry. The following example from B. Conrad shows that in quite reasonable (but less “rigid”) situations, the structure sheaf is not coherent over itself. Consider the ring \mathcal{O}_0 of germs of smooth (C^∞) functions at $0 \in \mathbb{R}$, with coordinate x . Now \mathcal{O}_0 is a local ring. Its maximal ideal \mathfrak{m} is generated by x . (Key idea: suppose $f \in \mathfrak{m}$, and suppose f has a representative defined on (ϵ, ϵ) . Then for $t \in (-\epsilon, \epsilon)$, $f(t) = \int_0^t f'(u) du = t \int_0^1 f'(tv) dv$. By “differentiating under the integral sign” repeatedly, we may check that $\int_0^1 f'(tv) dv$ is smooth. We deal with the case $t = 0$ separately as usual.)

Let $\phi \in \mathcal{O}_0$ be the germ of a smooth function that is 0 for $x \leq 0$, and positive for $x > 0$ (such as $\phi(x) = e^{-1/x^2}$ for $x > 0$). Consider the map $\times \phi: \mathcal{O}_0 \rightarrow \mathcal{O}_0$. The kernel is the ideal I_ϕ of functions vanishing for $x \geq 0$. Clearly I_ϕ is nonzero (for example, $\phi(-x) \in I_\phi$), but as $\mathfrak{m} = (x)$, $I_\phi = xI_\phi$, so I_ϕ cannot be finitely generated or else Nakayama's Lemma 8.2.9 would be contradicted. (Essentially the same argument shows that the sheaf of smooth functions on \mathbb{R} is not coherent.) This is why coherence has no useful meaning for smooth manifolds.

6.5 Visualizing schemes: Associated points and zerodivisors

The theory of *associated points* of a module refines the notion of support (§4.1.7). They will help us understand and visualize nilpotents, and generalize the notion of “rational functions” to non-integral schemes. They are useful in ways we won’t use, for example through their connection to primary decomposition. They might be most useful for us in helping us understand and visualize (non-)zerodivisors, which will come up repeatedly, through effective Cartier divisors and line bundles, regular sequences, depth and Cohen-Macaulayness, and more.

There is no particular reason to discuss associated points now, and this section can be read independently, at leisure. We will not need this material in any essential way for some time, but it is a good opportunity to practice visualizing geometry, and to learn some useful algebra.

6.5.1. Motivation. Figure 6.1 is a sketch of a scheme X . We see two connected components, and three irreducible components. The irreducible components of X have dimensions 2, 1, and 1, although we won’t be able to make sense of “dimension” until Chapter 12. Both connected components are nonreduced.

We see a little more in this picture, which we will make precise in this section, in terms of “associated points”. The reducible connected component seems to have different amounts of nonreduced behavior on different loci. The scheme X has six associated points, which are the generic points of the irreducible subsets “visible” in the picture. A function on X is a zerodivisor if its zero locus contains any of these six irreducible subvarieties.




FIGURE 6.1. This scheme has six associated points, of which three are embedded points. A function is a zerodivisor if it vanishes at any of these six points.

Suppose M is a finitely generated module over a Noetherian ring A . For example, M could be A itself. Then there are some special points of $\text{Spec } A$ that are particularly crucial to understanding M . These are the *associated points* of M (or equivalently, the *associated prime ideals* of M — we will use these terms interchangeably). As motivation, we give a zillion properties of associated points, and leave it to you to verify them from the theory developed in the rest of this section.

As you read this section, you may wish to keep in mind

$$M = A = k[x, y]/(y^2, xy)$$

(Figure 4.4) as a running example.

6.5.2. A zillion properties of associated points. Here are some of the properties of associated points that we will prove.

There are finitely many associated points $\text{Ass}_A M \subset \text{Spec } A$.

The support of M is the closure of the associated points of M : $\text{Supp } M = \overline{\text{Ass}_A M}$. The support of any submodule of M is the closure of some subset of the associated points of M . The support of any element of M is the closure of some subset of the associated points.

The associated points of M are precisely the generic points of irreducible components of $\text{Supp } m$ for all $m \in M$. The associated points of M are precisely the generic points of those $\text{Supp } m$ which are irreducible (as m runs over the elements of M). The associated primes of M are precisely those prime ideals that are annihilators of some element of M .

Taking “associated points” commutes with localization. Hence this notion is “geometric in nature”, which will (in §6.5.2) allow us to extend the notion to coherent sheaves on locally Noetherian schemes.

Associated points behave *fairly* well in exact sequences. For example, the associated points of a submodule are a subset of the associated points of the module.

If $I \subset A$ is an ideal, the associated primes \mathfrak{p} of A/I are precisely those \mathfrak{p} such that a \mathfrak{p} -primary ideal appears in the primary decomposition of I .

We will repeatedly use the fact that *an element of A is a zerodivisor if and only if it vanishes at an associated point*.

An element of A is a unit if and only if it vanishes at no associated point. An element of A is nilpotent if and only if it vanishes at every associated point.

The locus of points $[\mathfrak{p}]$ of $\text{Spec } A$ where the stalk $A_{\mathfrak{p}}$ is nonreduced is the closure of some subset of the associated points.

An associated point that is in the closure of another associated point is said to be an *embedded point*. If A is reduced, then $\text{Spec } A$ has no embedded points. Hypersurfaces in \mathbb{A}_k^n have no embedded points. We will later see that complete intersections have no embedded points (§29.2.7).

Elements of M are determined by their localization at the associated points. Sections of the corresponding sheaf \widetilde{M} (Exercise/Definition 4.1.D) are determined by their germs at the associated points.

This discussion immediately implies a notion of **associated point** for a coherent sheaf on a locally Noetherian scheme, with all the good properties described here. The phrase **associated point of a locally Noetherian scheme** X (without explicit mention of a coherent sheaf) means “associated point of \mathcal{O}_X ”, and similarly for **embedded points**.

We now establish these zillion facts.

6.5.3. More on the notion of support.

The notion of *associated points* of an A -module M refines the notion of *support* (in the case where M is finitely generated over a Noetherian ring A). (In what follows, we make no assumptions that A is Noetherian or that M is finitely generated until we need to.) To set this up, recall (§4.1.7) that the *support* of $m \in M$,

$$\text{Supp } m = \{[\mathfrak{p}] \in \text{Spec } A : m_{\mathfrak{p}} \neq 0\},$$

is a closed subset, and thus of the form $V(I)$ for some I . Exercise 6.5.A gives the “best such” I . Define the **annihilator ideal** $\text{Ann}_A m \subset A$ of an element m of an A -module M by:

$$\text{Ann}_A m := \{a \in A : am = 0\} = \ker(A \xrightarrow{\times m} M).$$

The subscript A is omitted if it is clear from context.

6.5.A. EASY IMPORTANT EXERCISE. Show that $\text{Supp } m = V(\text{Ann } m)$.

Recall (Definition 2.7.6) that

$$\text{Supp } \widetilde{M} = \{p \in \text{Spec } A : \widetilde{M}_p \neq 0\},$$

and the analogous Definition 4.1.7 of the *support of the module* M ,

$$(6.5.3.1) \quad \text{Supp } M := \{p \in \text{Spec } A : M_p \neq 0\},$$

so $\text{Supp } M = \text{Supp } \widetilde{M}$. If M is a principal module generated by $m \in M$, then

$$\text{Supp } M = \text{Supp } Am = \text{Supp } m = V(\text{Ann } m).$$

The notions of support and associated points behave well in exact sequences, and under localization. We begin to explain this now.

6.5.4. *The notion of support behaves well in exact sequences.*

6.5.B. EXERCISE. Suppose that $0 \rightarrow M' \rightarrow M \rightarrow M'' \rightarrow 0$ is a short exact sequence of A -modules.

- (a) Show that $\text{Supp } M = \text{Supp } M' \cup \text{Supp } M''$.
- (b) Show that if M is a finitely generated module, then $\text{Supp } M$ is a closed subset of $\text{Spec } A$. (Hint: induction on the number of generators.)

Warning: $\text{Supp } M$ need not be closed in general; consider $A = \mathbb{Z}$ and $M = \bigoplus_{p \text{ prime}} \mathbb{Z}/(p)$.

6.5.C. EXERCISE. Suppose M is a finitely generated A -module, and $x \in A$ has value 0 at all the points of $\text{Supp } M$, i.e., x is contained in all of the primes where M is supported. Show that some power x^n of x annihilates every element of M . (Hint: annihilate a generating set.)

6.5.5. Definition: Associated points and associated primes.

Define the **associated prime ideals** of an A -module M to be those *prime* ideals of A of the form $\text{Ann}_A(m)$ for some $m \in M$. Define the **associated points** of M to be the corresponding points of $\text{Spec } A$; we use the terminology “associated points” and “associated primes” interchangeably. The set of associated points is denoted $\text{Ass}_A M \subset \text{Spec } A$. The subscript A is dropped if it is clear from the context. (To help remember the definition and the notation, some call these the *assassins*, as they are the primes that can ruthlessly annihilate elements of the module, without remorse. But we will not use this term.)

6.5.6. The **associated primes of a ring** A are the associated primes of A considered as an A -module (i.e., $M = A$).

6.5.D. EASY EXERCISE (ASSOCIATED POINTS OF INTEGRAL DOMAINS). If A is an integral domain, show that $\text{Ass } A = \{(0)\}$ — the zero ideal is the only associated prime.

6.5.E. EXERCISE (ASSOCIATED POINTS OF HYPERSURFACES). Given $f \in k[x_1, \dots, x_n]$, show that the associated primes of $k[x_1, \dots, x_n]/(f)$ are those principal ideals generated by the prime factors of f . (Your argument will apply more generally to any $f \in A$ where A is a Unique Factorization Domain.)

6.5.7. *The observation that $[p] \in \text{Ass}_A(M)$ if and only if there is an injection $A/p \hookrightarrow M$ of A -modules will be essential. The next two exercises might drive this home.*

6.5.F. EXERCISE. Suppose $M' \subset M$. Show that $\text{Ass } M' \subset \text{Ass } M$.

The corresponding statement for “support” is implicit in Exercise 6.5.B(a).

6.5.G. EXERCISE. Show that $\text{Ass } M \subset \text{Supp } M$.

If M is finitely generated, then $\text{Supp } M$ is closed (Exercise 6.5.B), so $\overline{\text{Ass } M} \subset \text{Supp } M$. Equality will be shown in Proposition 6.5.24, when A is Noetherian.

6.5.8. Nonzero modules over Noetherian rings have associated points.

Suppose m is a nonzero element of an A -module M . Observe that for any nonzero multiple xm of m , $\text{Ann } m \subseteq \text{Ann } xm \subsetneq A$.

6.5.H. EXERCISE. Suppose A is Noetherian. Show that there is some multiple $n = xm$ such that *any* nonzero multiple $yn \neq 0$ of n satisfies $\text{Ann } yn = \text{Ann } n$.

6.5.9. Proposition. — *Continuing the notation of the previous exercise, we have that $\text{Ann } n$ is a prime ideal.*

Proof. Suppose $ab \in \text{Ann } n$, so $abn = 0$. Then either $bn = 0$ (in which case $b \in \text{Ann } n$), or else $a \in \text{Ann } bn = \text{Ann } n$. \square

We have thus proved the following.

6.5.10. Proposition (nonzero modules over Noetherian rings have associated primes). — *If M is a nonzero module over a Noetherian ring A , then $\text{Ass}_A M$ is nonempty. More precisely, for any $m \neq 0$ in M , there is an associated prime \mathfrak{p} containing $\text{Ann } m$, and $\mathfrak{p} = \text{Ann } xm$ for some $x \in A$.*

6.5.11. Localizations at the associated primes.

Recall the useful fact that $M \rightarrow \prod_{\mathfrak{p} \in \text{Spec } A} M_{\mathfrak{p}}$ is an injection (Exercise 4.1.F). Our current situation is much better: we can take the product over only the localization at *associated* primes.

6.5.I. EXERCISE. Suppose M is a module over a Noetherian ring A . Show that the natural map

$$(6.5.11.1) \quad M \longrightarrow \prod_{\mathfrak{p} \in \text{Ass } M} M_{\mathfrak{p}}$$

is an injection. Hint: if the kernel K is nonzero, then K has an associated prime \mathfrak{p} , which is the annihilator of some $m \in K \subset M$, and m is nonzero in $M_{\mathfrak{p}}$.

Clearly we need only the maximal among the associated primes in (6.5.11.1).

6.5.12. Zerodivisors = elements of associated primes.

6.5.13. Proposition. — Suppose $f \in A$, with A Noetherian. Then f is a zerodivisor on M if and only if f vanishes at an associated point of M . Translation: the set of zerodivisors is the union of the associated prime ideals.

Again, we need only the maximal among the associated primes. For example, if (A, \mathfrak{m}) is a local ring, then \mathfrak{m} is an associated prime if and only if every element of \mathfrak{m} is a zerodivisor.

Proof. Suppose f vanishes at an associated point $[\mathfrak{p}]$ of M . Choose m with $\text{Ann } m = \mathfrak{p}$. Then $fm = 0$ while $m \neq 0$, so f is a zerodivisor.

Conversely, if f vanishes at no associated point, consider the commuting diagram

$$\begin{array}{ccc} M & \longrightarrow & \prod_{\mathfrak{p} \in \text{Ass } M} M_{\mathfrak{p}} \\ \downarrow \times f & & \downarrow \times f \\ M & \longrightarrow & \prod_{\mathfrak{p} \in \text{Ass } M} M_{\mathfrak{p}}, \end{array}$$

where the rows are the maps of (6.5.11.1). The vertical arrow on the right (multiplication by f) is an injection by hypothesis, so the vertical arrow on the left must be an injection as well. \square

6.5.14. Associated points behave fairly well in exact sequences.

6.5.15. Proposition. — Suppose

$$(6.5.15.1) \quad 0 \longrightarrow M' \longrightarrow M \longrightarrow M'' \longrightarrow 0$$

is a short exact sequence of A -modules. Then

$$(6.5.15.2) \quad \text{Ass } M' \subset \text{Ass } M \subset \text{Ass } M' \cup \text{Ass } M''.$$

We come to our first complicated proof of the section.

Proof. The first inclusion of (6.5.15.2) was shown in Exercise 6.5.F.

Suppose next that $[\mathfrak{p}] \in \text{Ass } M$, so there is some $m \in M$ with $\text{Ann } m = \mathfrak{p}$. We wish to find a submodule of M' or M'' isomorphic to A/\mathfrak{p} . If this proposition were true, we would expect to find such a submodule in the “part of (6.5.15.1) spanned by m ”. So we consider instead the exact sequence

$$0 \longrightarrow Am \cap M' \longrightarrow Am \longrightarrow Am/(Am \cap M') \longrightarrow 0,$$

noting that the three modules appearing here are submodules of the corresponding modules in (6.5.15.1). So by Exercise 6.5.F it suffices to prove the result in this “special case”, which can be rewritten as

$$0 \longrightarrow I/\mathfrak{p} \longrightarrow A/\mathfrak{p} \longrightarrow A/I \longrightarrow 0$$

where I is the annihilator of m considered as an element of the module $Am/(Am \cap M')$. For convenience, let $B = A/\mathfrak{p}$ (an integral domain), so we rewrite the exact

sequence further as the top row of

$$\begin{array}{ccccccc}
 0 & \longrightarrow & J & \longrightarrow & B & \longrightarrow & B/J \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & M' & \longrightarrow & M & \longrightarrow & M'' \longrightarrow 0
 \end{array}$$

Now localize the top row of B -modules at $(0) \subset B$, so it becomes an exact sequence of vector spaces over the fraction field $K(B)$, and the central element is one-dimensional:

$$0 \longrightarrow J \otimes K(B) \longrightarrow K(B) \longrightarrow (B/J) \otimes K(B) \longrightarrow 0.$$

Thus one of the outside terms $J \otimes K(B)$ and $(B/J) \otimes K(B)$ has a nonzero element, which (tracing our argument backwards) gives an element of M' or M'' whose annihilator is precisely \mathfrak{p} . \square

6.5.16. Cautionary example. The short exact sequence of \mathbb{Z} -modules

$$0 \longrightarrow \mathbb{Z} \xrightarrow{\times 2} \mathbb{Z} \longrightarrow \mathbb{Z}/2 \longrightarrow 0$$

(and Easy Exercise 6.5.D) shows it is not always true that $\text{Ass } M = \text{Ass } M' \cup \text{Ass } M''$. However, sometimes we can still ensure some associated primes of M'' lift to associated primes of M , as we shall see in §6.5.20.

6.5.17. Finitely generated modules over Noetherian rings have finitely many associated points/primes.

6.5.J. IMPORTANT EXERCISE. Suppose that M is a finitely generated module over a Noetherian ring A . Show that M has a (finite) filtration

(6.5.17.1)

$$0 = M_0 \subset M_1 \subset \cdots \subset M_n = M \quad \text{where } M_{i+1}/M_i \cong A/\mathfrak{p}_i \text{ for some prime } \mathfrak{p}_i.$$

Hint: Build (6.5.17.1) inductively from left to right, using Proposition 6.5.10, and show the process terminates.

6.5.K. EXERCISE. Suppose an A -module M has a finite filtration (6.5.17.1), with no assumptions of finite generation or Noetherianity. Show that every associated prime of M appears as one of the \mathfrak{p}_i . In particular, under this hypothesis (for example, if M is finitely generated over a Noetherian ring) M has finitely many associated points/primes. (Hint: Exercise 6.5.D and Proposition 6.5.15.)

6.5.18. Caution: Non-associated prime ideals may unavoidably appear among the quotients in (6.5.17.1). Example 6.5.16 shows that the non-associated prime ideals may be among the quotients in (6.5.17.1), although in that case it is because the filtration was chosen unwisely. But better choices will not always remedy the problem:

6.5.L. EXERCISE. Consider the module $M = (x, y) \subset A = k[x, y]$ over the ring A . Show that any filtration (6.5.17.1) of M contains a quotient A/\mathfrak{p}_i where \mathfrak{p}_i is not an associated prime.

6.5.19. Remark. However, not all is lost: Exercise 6.5.P will show that for any quotient A/\mathfrak{p}_i in any filtration (6.5.17.1) of M , \mathfrak{p}_i must *contain* an associated prime of M .

6.5.20. Associated points behave fairly well in exact sequences, continued.

6.5.21. Proposition. — We continue to consider the short exact sequence

$$0 \longrightarrow M' \longrightarrow M \longrightarrow M'' \longrightarrow 0$$

of A -modules. Suppose $\mathfrak{p} \in \text{Ass } M''$, but $\mathfrak{p} \notin \text{Supp } M'$ (a stronger hypothesis than $\mathfrak{p} \notin \text{Ass } M'$). Then $\mathfrak{p} \in \text{Ass } M$.

We come to our second complicated proof of the section.

Proof. Choose $m'' \in M''$ with $\mathfrak{p} = \text{Ann } m''$ in M'' . Choose a lift of $m \in M$ of $m'' \in M''$. We apply a strategy similar to that of our proof of Proposition 6.5.15. Consider the “inclusion of short exact sequences”

$$\begin{array}{ccccccc} 0 & \longrightarrow & \ker(Am \rightarrow Am'') & \longrightarrow & Am & \longrightarrow & Am'' \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & M' & \longrightarrow & M & \longrightarrow & M'' \longrightarrow 0. \end{array}$$

As $\text{Supp}(\ker(Am \rightarrow Am'')) \subset \text{Supp } M'$, $\text{Ass}(Am) \subset \text{Ass } M$, and $\text{Ass}(Am'') \subset \text{Ass } M''$ (Exercises 6.5.B(a) and 6.5.F), we have reduced to considering the top row instead of the bottom row. The top row can be conveniently rewritten as

$$0 \longrightarrow \mathfrak{p}/I \longrightarrow A/I \longrightarrow A/\mathfrak{p} \longrightarrow 0$$

(here $I = \text{Ann}(m)$) where our hypothesis translates to $[\mathfrak{p}] \notin \text{Supp}(\mathfrak{p}/I)$. For convenience, let $B = A/I$ so we may now consider the sequence

$$0 \longrightarrow \mathfrak{q} \longrightarrow B \longrightarrow B/\mathfrak{q} \longrightarrow 0,$$

where \mathfrak{q} is prime, with the confusion-inducing hypothesis $[\mathfrak{q}] \notin \text{Supp } \mathfrak{q}$.

From the confusing hypothesis, there is an element b of B that vanishes on $\text{Supp } \mathfrak{q}$ but doesn't vanish at $[\mathfrak{q}]$. Translation: (i) b lies in all the primes of $\text{Supp } \mathfrak{q}$, but (ii) $b \notin \mathfrak{q}$. Then from (i) there is some power b^n of b that annihilates all elements of \mathfrak{q} (Exercise 6.5.C). From (ii), $b^n \notin \mathfrak{q}$.

Then $\text{Ann}(b^n)$ contains \mathfrak{q} from (i), but does not contain any element of $B \setminus \mathfrak{q}$ from (ii), so $\text{Ann}(b^n) = \mathfrak{q}$. hence \mathfrak{q} is an associated prime of B , which (unwinding our argument) shows that \mathfrak{p} is an associated prime of M . \square

6.5.22. Minimal primes are associated.

6.5.M. EXERCISE: MINIMAL PRIMES (“IRREDUCIBLE COMPONENTS”) ARE ASSOCIATED. Suppose M is a finitely generated module over Noetherian A , and $\mathfrak{p} \subset A$ is a prime ideal corresponding to an irreducible component of $\text{Supp } M \subset \text{Spec } A$. Show that $[\mathfrak{p}] \in \text{Ass } M$. Hint: Exercise 6.5.J and Proposition 6.5.21.

6.5.23. Non-Noetherian Remark. By combining Proposition 6.5.13 with Exercise 6.5.M, we see that if A is a Noetherian ring, then any element of any minimal prime \mathfrak{p} is a zerodivisor. This is true without Noetherian hypotheses: suppose $s \in \mathfrak{p}$. Then

by minimality of \mathfrak{p} , $\mathfrak{p}A_{\mathfrak{p}}$ is the unique prime ideal in $A_{\mathfrak{p}}$, so the element $s/1$ of $A_{\mathfrak{p}}$ is nilpotent (because it is contained in all prime ideals of $A_{\mathfrak{p}}$, Theorem 3.2.13). Thus for some $t \in A \setminus \mathfrak{p}$, $ts^n = 0$, so s is a zerodivisor in A . We will use this in Exercise 12.1.G.

6.5.24. Proposition. — *Suppose M is a finitely generated module over a Noetherian ring A . Then $\text{Supp } M = \overline{\text{Ass } M}$.*

Proof. Combine Exercises 6.5.G and 6.5.M. □

6.5.N. EXERCISE. Suppose A is a Noetherian ring. Show that the locus of points $[\mathfrak{p}]$ where $A_{\mathfrak{p}}$ is nonreduced is the support of the nilradical $\text{Supp } \mathfrak{N}$. Hence show that the “reduced locus” of a locally Noetherian scheme is open.

The following justifies a simple way of thinking about associated primes of a ring.

6.5.O. EXERCISE. Show that a prime ideal $\mathfrak{p} \subset A$ is an associated prime of A if and only if there is $f \in A$ such that $\text{Supp } f = V(\mathfrak{p}) = \overline{[\mathfrak{p}]}$.

6.5.P. EXERCISE, PROMISED IN REMARK 6.5.19. Show that for each quotient in the filtration (6.5.17.1) of M , $\text{Supp } A/\mathfrak{p}_i = \overline{[\mathfrak{p}]} \subset \text{Supp } M$, and that every \mathfrak{p}_i contains a minimal prime, and hence an associated prime.

6.5.25. “Support” and “associated points” commute with localization.

Suppose S is a multiplicative subset of A , and $\mathfrak{p} \subset A$ is a prime ideal not meeting S , so (abusing notation slightly) $[\mathfrak{p}] \in \text{Spec } S^{-1}A \subset \text{Spec } A$ (§3.2.9).

6.5.Q. EXERCISE (Supp COMMUTES WITH LOCALIZATION). Show that for any A -module M , $\text{Supp}_A M \cap \text{Spec } S^{-1}A = \text{Supp}_{S^{-1}A} S^{-1}M$.

6.5.26. Proposition (Ass commutes with localization). — *For an A -module M , we have $\text{Ass}_A M \cap \text{Spec } S^{-1}A = \text{Ass}_{S^{-1}A} S^{-1}M$.*

Proof. We first show that $\text{Ass}_A M \cap \text{Spec } S^{-1}A \subset \text{Ass}_{S^{-1}A} S^{-1}M$. If $\mathfrak{p} \in \text{Ass}_A M$, then we have an injection $A/\mathfrak{p} \hookrightarrow M$. Localizing by S (which preserves injectivity), we have $(S^{-1}A)/(S^{-1}\mathfrak{p}) \hookrightarrow M$. (Where did we use $\mathfrak{p} \in \text{Spec } S^{-1}A$?)

We next show that $\text{Ass}_{S^{-1}A} S^{-1}M \subset \text{Ass}_A M \cap \text{Spec } S^{-1}A$. Suppose $\mathfrak{q} := S^{-1}\mathfrak{p} \in \text{Ass}_{S^{-1}A} S^{-1}M$, so $\mathfrak{q} = \text{Ann}_{S^{-1}A}(m/s)$, for some $s \in S$, and $m \in M$. Since the elements of S are units in $S^{-1}A$, we have that $\mathfrak{q} = \text{Ann}_{S^{-1}A} m$. As support commutes with localization, $V(\mathfrak{p})$ must be an irreducible component of $\text{Supp } m$ (as $\text{Supp } m \cap \text{Spec } S^{-1}A$ contains $[\mathfrak{q}]$, but no generizations of $[\mathfrak{q}]$). Then by Exercise 6.5.M, \mathfrak{p} is an associated prime of M . □

6.5.27. Embedded points/primes.

6.5.28. Definition. The associated points that are *not* the generic points of irreducible components of $\text{Supp } M$ are called **embedded points**, and their corresponding primes are called **embedded primes**. The motivation for this language is their geometric incarnation as embedded points (§6.5.32 shortly). For example,

the scheme of Figure 6.1 has three embedded points (the corresponding ring has three embedded primes).

6.5.29. Remark. Exercise 6.5.E translates to “hypersurfaces in \mathbb{A}_k^n have no embedded points”. More generally, if A is a unique factorization domain, then $\operatorname{Spec} A/(f)$ has no embedded points for any $f \in A$. Generalizing in a different direction, we will see that “complete intersections have no embedded points” in §29.2.7.

6.5.R. EXERCISE. Suppose A is a reduced ring (i.e., A has no nonzero nilpotents). Show that $\operatorname{Spec} A$ has no embedded primes. (Hints: if $\mathfrak{p} = \operatorname{Ann} a$ is an embedded prime, show that you can find an element b of \mathfrak{p} not contained in any of the minimal primes of A . From $ab = 0$, show that a is contained in all the minimal primes. Show that a vanishes at all points of $\operatorname{Spec} A$, and hence by Theorem 3.2.13 is nilpotent.)

Thus *reduced rings have no embedded primes*. Even better: *the only elements of a ring that an embedded prime can annihilate are nilpotents*.

6.5.30. Remark. The converse to Exercise 6.5.R is false. Rings without embedded primes can still have nilpotents — witness $k[x]/(x^2)$.

6.5.S. EXERCISE. Show that if \mathfrak{p} is an embedded prime of a ring A , then $A_{\mathfrak{p}}$ is nonreduced.

6.5.31. Get your hands dirty: Explicit algebraic exercises.

6.5.T. EXERCISE. (See Figure 4.4.) Suppose f is a function on $\operatorname{Spec} k[x, y]/(y^2, xy)$, i.e., $f \in k[x, y]/(y^2, xy)$. Show that $\operatorname{Supp} f$ is either the empty set, or the origin, or the entire space. Hence find the associated points of $\operatorname{Spec} k[x, y]/(y^2, xy)$.

6.5.U. EXERCISE (CONTINUING THE PREVIOUS EXERCISE). Show explicitly by hand that $f \in k[x, y]/(y^2, xy)$ is a zerodivisor if and only if $f(0, 0) = 0$.

6.5.V. EXERCISE (PRACTICE WITH FUZZY PICTURES). Suppose $X = \operatorname{Spec} \mathbb{C}[x, y]/I$, and that the associated points of X are $[(y - x^2)]$, $[(x - 1, y - 1)]$, and $[(x - 2, y - 2)]$. (Exercise 6.5.W will verify that such an X actually exists.)

(a) Sketch X as a subset of $\mathbb{A}_{\mathbb{C}}^2 = \operatorname{Spec} \mathbb{C}[x, y]$, including fuzz.

(b) Do you have enough information to know if X is reduced?

(c) Do you have enough information to know if $x + y - 2$ is a zerodivisor? How about $x + y - 3$? How about $y - x^2$?

6.5.W. EXERCISE. Let $I = (y - x^2)^3 \cap (x - 1, y - 1)^{15} \cap (x - 2, y - 2)$. Show that $X = \operatorname{Spec} \mathbb{C}[x, y]/I$ satisfies the hypotheses of Exercise 6.5.V. (Rhetorical question: Is there a “smaller” example? Is there a “smallest”?)

6.5.32. Geometric definitions.

6.5.X. EXERCISE/DEFINITION. Define the **associated points** and **embedded points of a coherent sheaf on a locally Noetherian scheme**. (Idea/hint: do it affinely.)

6.5.Y. EXERCISE. Suppose X is a locally Noetherian scheme, and $U \subset X$ is an open subscheme. Show that the natural map

$$(6.5.32.1) \quad \Gamma(U, \mathcal{O}_X) \longrightarrow \prod_{p \in \text{Ass } X \cap U} \mathcal{O}_{X,p}$$

(cf. (6.5.11.1)) is an injection.

6.5.33. Generalizing the fraction field: the total fraction ring.

6.5.34. Definitions: Rational functions on locally Noetherian schemes. A **rational function** on a locally Noetherian scheme is an element of the image of $\Gamma(U, \mathcal{O}_U)$ in (6.5.32.1) for some U containing all the associated points. Equivalently, the set of rational functions is the colimit of $\mathcal{O}_X(U)$ over all open sets containing the associated points.

For example, on $\text{Spec } k[x, y]/(y^2, xy)$ (Figure 4.4), $\frac{x-2}{(x-1)(x-3)}$ is a rational function, but $\frac{x-2}{x(x-1)}$ is not.

A rational function has a maximal **domain of definition**, because any two actual functions on an open set (i.e., sections of the structure sheaf over that open set) that agree as “rational functions” (i.e., on small enough open sets containing associated points) must be the same function, by the injectivity of (6.5.32.1). We say that a rational function f is **regular at a point** p if p is contained in this maximal domain of definition (or equivalently, if there is some open set containing p where f is defined). For example, on $\text{Spec } k[x, y]/(y^2, xy)$, the rational function $\frac{x-2}{(x-1)(x-3)}$ has domain of definition consisting of everything but 1 and 3 (i.e., $[(x-1)]$ and $[(x-3)]$), and is regular away from those two points. A rational function is **regular** if it is regular at all points. (Unfortunately, “regular” is a regularly overused word in mathematics, and in algebraic geometry in particular.)

The complement of the domain of definition of a rational function f is called the **indeterminacy locus** of f (a phrase we’ll see again in §11.4.3).

The rational functions form a ring, called the **total fraction ring** or **total quotient ring** of X . If $X = \text{Spec } A$ is affine, then this ring is called the **total fraction (or quotient) ring** of A . If X is integral, the total fraction ring is the function field $K(X)$ — the stalk at the generic point — so this extends our earlier Definition 5.2.1 of $K(\cdot)$.

6.5.Z. EXERCISE. Show that the ring of rational functions on a locally Noetherian scheme is the colimit of the functions over all open sets containing the associated points:

$$\text{colim}_U : \text{Ass } X \subset U \mathcal{O}(U).$$

Slightly better (but slightly different): show that a rational function is the data of a function f defined on an open set U containing all associated points, where two such data (U, f) and (U', f') define the same rational function if and only if $f_{U \cap U'} = f'_{U \cap U'}$ (cf. §1.4.9). If X is reduced, show that this is the same as requiring that they are defined on an open set of each of the irreducible components.

6.5.35. Remark: Associated points of integral schemes. In order for some of our discussion elsewhere to make sense in non-Noetherian settings, we note that the notion of associated points for integral schemes works perfectly, because it works for integral domains — only the generic point is associated. In particular, the definition

above of rational functions on an integral scheme X agrees with Definition 5.2.I, as precisely the elements of the function field $K(X)$.

6.5.36. ★★ Aside: Primary ideals and associated primes. Primary decomposition was introduced by the world chess champion E. Lasker, and axiomatized by world math champion E. Noether. We won't need it, but here is the beginning of the story, for the curious reader. An ideal $I \subset A$ in a ring is **primary** if $I \neq A$, and $xy \in I$ implies either $x \in I$ or $y^n \in I$ for some $n > 0$. In this case, \sqrt{I} is prime, and I is said to be \mathfrak{p} -primary. Equivalently, if $I \subset A$ then I is \mathfrak{p} -primary if and only if A/I has only one associated prime \mathfrak{p} . If I is an ideal of a Noetherian ring A , then the associated prime ideals A/I turn out to be precisely the radicals of ideals in a primary decomposition. See [E, §3.3], for example, for more of this important story.

6.6 ★★ Coherent modules over non-Noetherian rings

This section is intended for people who might work with non-Noetherian rings, or who otherwise might want to understand coherent sheaves in a more general setting. Even these people should read this much later. Read this only if you really want to!

Suppose A is a ring. Recall the definition of when an A -module M is finitely generated, finitely presented, and coherent (§6.4). The reason we like coherence is that coherent modules form an abelian category. Here are some accessible exercises working out why these notions behave well. Some repeat earlier discussion in order to keep this section self-contained.

The notion of coherence of a module is only interesting in the case that a ring is coherent over itself. Similarly, coherent sheaves on a scheme X will be interesting only when \mathcal{O}_X is coherent ("over itself"). In this case, coherence is clearly the same as finite presentation. An example where non-Noetherian coherence comes up is the ring $R\langle x_1, \dots, x_n \rangle$ of "restricted power series" over a valuation ring R of a non-discretely valued K (for example, a completion of the algebraic closure of \mathbb{Q}_p). This is relevant to Tate's theory of non-archimedean analytic geometry over K (which you can read about in [BCDKT], for example). The importance of the coherence of the structure sheaf underlines the importance of Oka's Theorem in complex geometry (stated in §6.4.6).

6.6.A. EXERCISE. Show that coherent implies finitely presented implies finitely generated. (This was discussed at the start of §6.4.)

6.6.B. EXERCISE. Show that 0 is coherent.

Suppose for problems 6.6.C–6.6.I that

$$(6.6.0.1) \quad 0 \longrightarrow M \longrightarrow N \longrightarrow P \longrightarrow 0$$

is an exact sequence of A -modules. In this series of problems, we will show that if two of $\{M, N, P\}$ are coherent, the third is as well, which will prove very useful.

6.6.1. Hint \dagger . The following hint applies to several of the problems: try to write

$$\begin{array}{ccccccc}
 0 & \longrightarrow & A^{\oplus p} & \longrightarrow & A^{\oplus(p+q)} & \longrightarrow & A^{\oplus q} \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & M & \longrightarrow & N & \longrightarrow & P \longrightarrow 0
 \end{array}$$

and possibly use the Snake Lemma 1.7.5.

6.6.C. EXERCISE. Show that N finitely generated implies P finitely generated. (You will only need right-exactness of (6.6.0.1).)

6.6.D. EXERCISE. Show that M, P finitely generated implies N finitely generated. Possible hint: \dagger . (You will only need right-exactness of (6.6.0.1).)

6.6.E. EXERCISE. Show that N, P finitely generated need not imply M finitely generated. (Hint: if I is an ideal, we have $0 \rightarrow I \rightarrow A \rightarrow A/I \rightarrow 0$.)

6.6.F. EXERCISE. Show that N coherent, M finitely generated implies M coherent. (You will only need left-exactness of (6.6.0.1).)

6.6.G. EXERCISE. Show that N, P coherent implies M coherent. Hint for (i) in the definition of coherence (§6.4.1):

$$\begin{array}{ccccccc}
 & & A^{\oplus q} & & & & \\
 & & \downarrow & \searrow & & & \\
 & & & A^{\oplus p} & & & \\
 & & & \downarrow & \searrow & & \\
 0 & \longrightarrow & M & \longrightarrow & N & \longrightarrow & P \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \searrow \\
 & & 0 & & 0 & & 0
 \end{array}$$

(You will only need left-exactness of (6.6.0.1).)

6.6.H. EXERCISE. Show that M finitely generated and N coherent implies P coherent. Hint for (ii) in the definition of coherence (§6.4.1): \dagger .

6.6.I. EXERCISE. Show that M, P coherent implies N coherent. (Hint: \dagger .)

6.6.J. EXERCISE. Show that a finite direct sum of coherent modules is coherent.

6.6.K. EXERCISE. Suppose M is finitely generated, N coherent. Then if $\phi: M \rightarrow N$ is any map, then show that $\text{Im } \phi$ is coherent.

6.6.L. EXERCISE. Show that the kernel and cokernel of maps of coherent modules are coherent.

At this point, we have verified that coherent A -modules form an abelian subcategory of the category of A -modules. (Things you have to check: 0 should be

in this set; it should be closed under finite sums; and it should be closed under taking kernels and cokernels.)

6.6.M. EXERCISE. Suppose M and N are coherent submodules of the coherent module P . Show that $M + N$ and $M \cap N$ are coherent. (Hint: consider the right map $M \oplus N \rightarrow P$.)

6.6.N. EXERCISE. Show that if A is coherent (as an A -module) then finitely presented modules are coherent. (Of course, if finitely presented modules are coherent, then A is coherent, as A is finitely presented!)

6.6.O. EXERCISE. If M is finitely presented and N is coherent, show that $\text{Hom}(M, N)$ is coherent. (Hint: Hom is left-exact in its first argument.)

6.6.P. EXERCISE. If M is finitely presented, and N is coherent, show that $M \otimes_A N$ is coherent.

Part III

Morphisms of schemes

Morphisms of schemes

7.1 Introduction

We now describe the morphisms between schemes. We will define some easy-to-state properties of morphisms, but leave more subtle properties for later.

Recall that a scheme is (i) a set, (ii) with a topology, (iii) and a (structure) sheaf of rings, and that it is sometimes helpful to think of the definition as having three steps. In the same way, the notion of morphism of schemes $X \rightarrow Y$ may be defined (i) as a map of sets, (ii) that is continuous, and (iii) with some further information involving the sheaves of functions. In the case of affine schemes, we have already seen the map of sets (§3.2.10) and later saw that this map is continuous (Exercise 3.4.H).

Here are two motivations for how morphisms should behave. The first is algebraic, and the second is geometric.

7.1.1. Algebraic motivation. We will want morphisms of affine schemes $\text{Spec } A \rightarrow \text{Spec } B$ to be precisely the ring maps $B \rightarrow A$. We have already seen that ring maps $B \rightarrow A$ induce maps of topological spaces in the opposite direction (Exercise 3.4.H); the main new ingredient will be to see how to add the structure sheaf of functions into the mix. Then a morphism of schemes should be something that “on the level of affine open sets, looks like this”.

7.1.2. Geometric motivation. Motivated by the theory of manifolds (§3.1.1), which like schemes are ringed spaces, we want morphisms of schemes at the very least to be morphisms of ringed spaces; we now motivate what these are. (We will formalize this in the next section.) Notice that if $\pi: X \rightarrow Y$ is a map of differentiable manifolds, then a smooth function on Y pulls back to a smooth function on X . More precisely, given an open subset $U \subset Y$, there is a natural map $\Gamma(U, \mathcal{O}_Y) \rightarrow \Gamma(\pi^{-1}(U), \mathcal{O}_X)$. This behaves well with respect to restriction (restricting a function to a smaller open set and pulling back yields the same result as pulling back and then restricting), so in fact we have a map of sheaves on Y : $\mathcal{O}_Y \rightarrow \pi_* \mathcal{O}_X$. Similarly a morphism of schemes $\pi: X \rightarrow Y$ should induce a map $\mathcal{O}_Y \rightarrow \pi_* \mathcal{O}_X$. But in fact in the category of differentiable manifolds a continuous map $\pi: X \rightarrow Y$ is a map of differentiable manifolds precisely when smooth functions on Y pull back to smooth functions on X (i.e., the pullback map from smooth functions on Y to *functions* on X in fact lies in the subset of *smooth functions*, i.e., the continuous map π induces a pullback of smooth functions, which can be interpreted as a map $\mathcal{O}_Y \rightarrow \pi_* \mathcal{O}_X$), so this map of sheaves *characterizes* morphisms in

the differentiable category. So we could use this as the *definition* of morphism in the differentiable category (see Exercise 3.1.A).

But how do we apply this to the category of schemes? In the category of differentiable manifolds, a continuous map $\pi: X \rightarrow Y$ induces a pullback of (the sheaf of) functions, and we can ask when this induces a pullback of *smooth* functions. However, functions are odder on schemes, and we can't recover the pullback map just from the map of topological spaces. The right patch is to hardwire this into the definition of morphism, i.e., to have a continuous map $\pi: X \rightarrow Y$, along with a pullback map $\pi^\sharp: \mathcal{O}_Y \rightarrow \pi_* \mathcal{O}_X$. This leads to the definition of the *category* of ringed spaces.

One might hope to define morphisms of schemes as morphisms of ringed spaces. This isn't quite right, as then Motivation 7.1.1 isn't satisfied: as desired, to each morphism $A \rightarrow B$ there is a morphism $\text{Spec } B \rightarrow \text{Spec } A$, but there can be additional morphisms of ringed spaces $\text{Spec } B \rightarrow \text{Spec } A$ not arising in this way (see Exercise 7.2.F). A revised definition as morphisms of ringed spaces that locally look of this form will work, but to avoid the awkwardness of sorting out the details, we take a different approach, using locally ringed spaces, which corresponds to asking not just that functions pull back, but also that *values of functions* pull back. However, we will check that our eventual definition actually is equivalent to this (Exercise 7.3.D).

We begin by formally defining morphisms of ringed spaces.

7.2 Morphisms of ringed spaces

7.2.1. Definition. A **morphism of ringed spaces** $\pi: X \rightarrow Y$ is a continuous map of topological spaces (which we unfortunately also call π) along with a map $\mathcal{O}_Y \rightarrow \pi_* \mathcal{O}_X$, which we think of as a “pullback map”. By adjointness (§2.7.1), this is the same as a map $\pi^{-1} \mathcal{O}_Y \rightarrow \mathcal{O}_X$. (It can be convenient to package this information as in the diagram (2.7.2.1).) There is an obvious notion of composition of morphisms, so ringed spaces form a category. Hence we have notion of automorphisms and isomorphisms. You can easily verify that an isomorphism of ringed spaces means the same thing as it did before (Definition 4.3.1).

If $U \subset Y$ is an open subset, then there is a natural morphism of ringed spaces $(U, \mathcal{O}_Y|_U) \rightarrow (Y, \mathcal{O}_Y)$ (which implicitly appeared earlier in §2.2.13). More precisely, if $U \rightarrow Y$ is an isomorphism of U with an open subset V of Y , and we are given an isomorphism $(U, \mathcal{O}_U) \xrightarrow{\sim} (V, \mathcal{O}_Y|_V)$ (via the isomorphism $U \xrightarrow{\sim} V$), then the resulting map of ringed spaces is called an **open embedding** (or **open immersion**) of ringed spaces, and the morphism $U \rightarrow Y$ is often written $U \hookrightarrow Y$.

7.2.A. EXERCISE (MORPHISMS OF RINGED SPACES GLUE). Suppose (X, \mathcal{O}_X) and (Y, \mathcal{O}_Y) are ringed spaces, $X = \cup_i U_i$ is an open cover of X , and we have morphisms of ringed spaces $\pi_i: U_i \rightarrow Y$ that “agree on the overlaps”, i.e., $\pi_i|_{U_i \cap U_j} = \pi_j|_{U_i \cap U_j}$. Show that there is a unique morphism of ringed spaces $\pi: X \rightarrow Y$ such that $\pi|_{U_i} = \pi_i$. (Exercise 2.2.F essentially showed this for topological spaces.)

7.2.B. EXERCISE. Show that open embeddings of ringed spaces are monomorphisms (in the category of ringed spaces).

7.2.C. EASY IMPORTANT EXERCISE: \mathcal{O} -MODULES PUSH FORWARD. Given a morphism of ringed spaces $\pi: X \rightarrow Y$, show that sheaf pushforward induces a functor $\text{Mod}_{\mathcal{O}_X} \rightarrow \text{Mod}_{\mathcal{O}_Y}$.

7.2.D. EASY IMPORTANT EXERCISE. Given a morphism of ringed spaces $\pi: X \rightarrow Y$ with $\pi(p) = q$, show that there is a map of stalks $(\mathcal{O}_Y)_q \rightarrow (\mathcal{O}_X)_p$.

7.2.E. KEY EXERCISE. Suppose $\pi^\sharp: B \rightarrow A$ is a morphism of rings. Define a morphism of ringed spaces $\pi: \text{Spec } A \rightarrow \text{Spec } B$ as follows. The map of topological spaces was given in Exercise 3.4.H. To describe a morphism of sheaves $\mathcal{O}_{\text{Spec } B} \rightarrow \pi_* \mathcal{O}_{\text{Spec } A}$ on $\text{Spec } B$, it suffices to describe a morphism of sheaves on the distinguished base of $\text{Spec } B$. On $D(g) \subset \text{Spec } B$, we define

$$\mathcal{O}_{\text{Spec } B}(D(g)) \rightarrow \mathcal{O}_{\text{Spec } A}(\pi^{-1}D(g)) = \mathcal{O}_{\text{Spec } A}(D(\pi^\sharp g))$$

by $B_g \rightarrow A_{\pi^\sharp g}$. Verify that this makes sense (e.g., is independent of g), and that this describes a morphism of sheaves on the distinguished base. (This is the third in a series of exercises. We saw that a morphism of rings induces a map of sets in §3.2.10, a map of topological spaces in Exercise 3.4.H, and now a map of ringed spaces here.)

The map of ringed spaces of Key Exercise 7.2.E is really not complicated. Here is an example. Consider the ring map $\mathbb{C}[y] \rightarrow \mathbb{C}[x]$ given by $y \mapsto x^2$ (see Figure 3.6). We are mapping the affine line with coordinate x to the affine line with coordinate y . The map is (on closed points) $a \mapsto a^2$. For example, where does $[(x-3)]$ go to? Answer: $[(y-9)]$, i.e., $3 \mapsto 9$. What is the preimage of $[(y-4)]$? Answer: those prime ideals in $\mathbb{C}[x]$ containing $[(x^2-4)]$, i.e., $[(x-2)]$ and $[(x+2)]$, so the preimage of 4 is indeed ± 2 . This is just about the map of sets, which is old news (§3.2.10), so let's now think about functions pulling back. What is the pullback of the function $3/(y-4)$ on $D(y-4) = \mathbb{A}^1 - \{4\}$? Of course it is $3/(x^2-4)$ on $\mathbb{A}^1 - \{-2, 2\}$.

The construction of Key Exercise 7.2.E will soon be an example of a morphism of schemes. In fact we could make that definition right now. Before we do, we point out (via the next exercise) that not every morphism of ringed spaces between affine schemes is of the form of Key Exercise 7.2.E. (In the language of §7.3, this morphism of ringed spaces is not a morphism of locally ringed spaces.) We won't use this exercise, but it answers a question that everyone will have.

7.2.F. EXERCISE. Recall (Exercise 3.4.K) that $\text{Spec } k[y]_{(y)}$ has two points, $[(0)]$ and $[(y)]$, where the second point is closed, and the first is not. Describe a map of ringed spaces $\text{Spec } k(x) \rightarrow \text{Spec } k[y]_{(y)}$ sending the unique point of $\text{Spec } k(x)$ to the closed point $[(y)]$, where the pullback map on global sections sends k to k by the identity, and sends y to x . Show that this map of ringed spaces is not of the form described in Key Exercise 7.2.E.

7.2.2. Tentative Definition we won't use (cf. Motivation 7.1.1 in §7.1). A morphism of schemes $\pi: (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ is a morphism of ringed spaces that “locally looks like” the maps of affine schemes described in Key Exercise 7.2.E. Precisely, for each choice of affine open sets $\text{Spec } A \subset X$, $\text{Spec } B \subset Y$, such that $\pi(\text{Spec } A) \subset \text{Spec } B$, the induced map of ringed spaces should be of the form shown in Key Exercise 7.2.E.

We would like this definition to be checkable on any affine cover, and we might hope to use the Affine Communication Lemma to develop the theory in this way. This works, but it will be more convenient to use a clever trick: in the next section, we will use the notion of locally ringed spaces, and then once we have used it, we will discard it like yesterday's garbage.

7.3 From locally ringed spaces to morphisms of schemes

In order to prove that morphisms behave in a way we hope, we will use the notion of a *locally ringed space*. It will not be used later, although it is useful elsewhere in geometry. The notion of locally ringed spaces (and maps between them) is inspired by what we know about manifolds (see Exercise 3.1.B). If $\pi: X \rightarrow Y$ is a morphism of manifolds, with $\pi(p) = q$, and f is a function on Y vanishing at q , then the pulled back function $\pi^{\sharp}(f)$ on X should vanish on p . Put differently: germs of functions (at $q \in Y$) vanishing at q should pull back to germs of functions (at $p \in X$) vanishing at p .

7.3.1. Definition. Recall (Definition 4.3.7) that a *locally ringed space* is a ringed space (X, \mathcal{O}_X) such that the stalks $\mathcal{O}_{X,p}$ are all local rings. A **morphism of locally ringed spaces** $\pi: X \rightarrow Y$ is a morphism of ringed spaces such that the induced map of stalks $\pi^{\sharp}: \mathcal{O}_{Y,q} \rightarrow \mathcal{O}_{X,p}$ (Exercise 7.2.D) sends the maximal ideal $\mathfrak{m}_{Y,q}$ of the former into the maximal ideal $\mathfrak{m}_{X,p}$ of the latter (a “**morphism of local rings**”). (Note in particular that locally ringed spaces form a category.) This means something rather concrete and intuitive: “if $p \mapsto q$, and g is a function vanishing at q , then it will pull back to a function vanishing at p .” You would also want: “if $p \mapsto q$, and g is a function *not* vanishing at q , then it will pull back to a function *not* vanishing at p .” This is a consequence of the following exercise.

7.3.A. EXERCISE (VALUES OF FUNCTIONS ON LOCALLY RINGED SPACES PULL BACK). If $\pi: X \rightarrow Y$ is a morphism of locally ringed spaces, and $\pi(p) = q$, show that π induces an inclusion $\kappa(q) \hookrightarrow \kappa(p)$ of residue fields. This captures the fact pull-back sends the locus where a function is zero (respectively, nonzero) to the locus where the pulled back function is zero (respectively, nonzero). In particular, by Exercise 4.3.F(b), the pullback of invertible functions are invertible functions.

To summarize: we use the notion of locally ringed space only to define morphisms of schemes, and to show that morphisms have reasonable properties. The main things you need to remember about locally ringed spaces are (i) that the functions have values at points, and (ii) that given a map of locally ringed spaces, values of functions “pull back”, and in particular, zeros of functions “pull back”.

7.3.B. EXERCISE. Show that morphisms of locally ringed spaces glue (cf. Exercise 7.2.A). (Hint: your solution to Exercise 7.2.A may work without change.)

7.3.C. EASY IMPORTANT EXERCISE.

Recall (§4.3.7 that $\text{Spec } A$ is a locally ringed space. Show that the morphism of ringed spaces $\pi: \text{Spec } A \rightarrow \text{Spec } B$ defined by a ring morphism $\pi^{\sharp}: B \rightarrow A$ (Exercise 7.2.E) is a morphism of locally ringed spaces.

7.3.2. Key Proposition. — *If $\pi: \text{Spec } A \rightarrow \text{Spec } B$ is a morphism of locally ringed spaces then it is the morphism of locally ringed spaces induced by the map $\pi^\sharp: B = \Gamma(\text{Spec } B, \mathcal{O}_{\text{Spec } B}) \rightarrow \Gamma(\text{Spec } A, \mathcal{O}_{\text{Spec } A}) = A$ as in Exercise 7.3.C.*

(Aside: Exercise 4.3.A is a special case of Key Proposition 7.3.2. You should look back at your solution to Exercise 4.3.A, and see where you implicitly used ideas about locally ringed spaces.)

Proof. Suppose $\pi: \text{Spec } A \rightarrow \text{Spec } B$ is a morphism of locally ringed spaces. We wish to show that it is determined by its map on global sections $\pi^\sharp: B \rightarrow A$. We first need to check that the map of points is determined by global sections. Now a point p of $\text{Spec } A$ can be identified with the prime ideal of global functions vanishing on it. The image point $\pi(p)$ in $\text{Spec } B$ can be interpreted as the unique point q of $\text{Spec } B$, where the functions vanishing at q are precisely those that pull back to functions vanishing at p . (Here we use the fact that π is a map of locally ringed spaces.) This is precisely the way in which the map of sets $\text{Spec } A \rightarrow \text{Spec } B$ induced by a ring map $B \rightarrow A$ was defined (§3.2.10).

Note in particular that if $b \in B$, $\pi^{-1}(D(b)) = D(\pi^\sharp b)$, again using the hypothesis that π is a morphism of locally ringed spaces.

It remains to show that $\pi^\sharp: \mathcal{O}_{\text{Spec } B} \rightarrow \pi_* \mathcal{O}_{\text{Spec } A}$ is the morphism of sheaves given by Exercise 7.2.E (cf. Exercise 7.3.C). It suffices to check this on the distinguished base (Exercise 2.5.C(a)). We now want to check that for any map of locally ringed spaces inducing the map of sheaves $\mathcal{O}_{\text{Spec } B} \rightarrow \pi_* \mathcal{O}_{\text{Spec } A}$, the map of sections on any distinguished open set $D(g) \subset \text{Spec } B$ ($g \in B$) is determined by the map of global sections $B \rightarrow A$.

Consider the commutative diagram

$$\begin{array}{ccccc}
 B & \xlongequal{\quad} & \Gamma(\text{Spec } B, \mathcal{O}_{\text{Spec } B}) & \xrightarrow{\pi^\sharp_{\text{Spec } B}} & \Gamma(\text{Spec } A, \mathcal{O}_{\text{Spec } A}) & \xlongequal{\quad} & A \\
 & & \downarrow \text{res}_{\text{Spec } B, D(b)} & & \downarrow \text{res}_{\text{Spec } A, D(\pi^\sharp b)} & & \\
 B_g & \xlongequal{\quad} & \Gamma(D(g), \mathcal{O}_{\text{Spec } B}) & \xrightarrow{\pi^\sharp_{D(g)}} & \Gamma(D(\pi^\sharp g), \mathcal{O}_{\text{Spec } A}) & \xlongequal{\quad} & A_{\pi^\sharp g} = A \otimes_B B_g.
 \end{array}$$

The vertical arrows (restrictions to distinguished open sets) are localizations by g , so the lower horizontal map $\pi^\sharp_{D(g)}$ is determined by the upper map (it is just localization by g). \square

We are ready for our definition.

7.3.3. Definition. If X and Y are schemes, then a morphism $\pi: X \rightarrow Y$ as locally ringed spaces is called a **morphism of schemes**. We have thus defined the **category of schemes**, which we denote Sch . (We then have notions of **isomorphism** — just the same as before, §4.3.7 — and **automorphism**. The *target* Y of π is sometimes called the **base scheme** or the **base**, when we are interpreting π as a “family of schemes parametrized by Y ” — this may become clearer once we have defined the fibers of morphisms in §10.3.2.)

The definition in terms of locally ringed spaces easily implies Tentative Definition 7.2.2:

7.3.D. IMPORTANT EXERCISE. Show that a morphism of schemes $\pi: X \rightarrow Y$ is a morphism of ringed spaces that looks locally like morphisms of affine schemes.

Precisely, if $\text{Spec } A$ is an affine open subset of X and $\text{Spec } B$ is an affine open subset of Y , and $\pi(\text{Spec } A) \subset \text{Spec } B$, then the induced morphism of ringed spaces is a morphism of affine schemes. (In case it helps, note: if $W \subset X$ and $Y \subset Z$ are both open embeddings of ringed spaces, then any morphism of ringed spaces $X \rightarrow Y$ induces a morphism of ringed spaces $W \rightarrow Z$, by composition $W \rightarrow X \rightarrow Y \rightarrow Z$.) Show that it suffices to check on a set $(\text{Spec } A_i, \text{Spec } B_i)$ where the $\text{Spec } A_i$ form an open cover of X and the $\text{Spec } B_i$ form an open cover of Y .

In practice, we will use the affine cover interpretation, and forget completely about locally ringed spaces. In particular, put imprecisely, the category of affine schemes is the category of rings with the arrows reversed. More precisely:

7.3.E. EXERCISE. Show that the category of rings and the opposite category of affine schemes are equivalent (see §1.2.16 and §1.2.21 to read about opposite and equivalent categories, respectively).

In particular, there can be many different maps from one point to another! For example, here are two different maps from the point $\text{Spec } \mathbb{C}$ to the point $\text{Spec } \mathbb{C}$: the identity (corresponding to the identity $\mathbb{C} \rightarrow \mathbb{C}$), and complex conjugation. (There are even more such maps!)

It is clear (from the corresponding facts about locally ringed spaces) that morphisms of schemes glue (Exercise 7.3.B), and the composition of two morphisms is a morphism. Isomorphisms in this category are precisely what we defined them to be earlier (§4.3.7).

7.3.F. ENLIGHTENING EXERCISE. (This exercise can give you some practice with understanding morphisms of schemes by cutting up into affine open sets.) Make sense of the following sentence: “ $\mathbb{A}_k^{n+1} \setminus \{\vec{0}\} \rightarrow \mathbb{P}_k^n$ given by

$$(x_0, x_1, \dots, x_n) \mapsto [x_0, x_1, \dots, x_n]$$

is a morphism of schemes.” Caution: you can’t just say where points go; you have to say where functions go. So you may have to divide these up into affines, and describe the maps, and check that they glue. (Can you generalize to the case where k is replaced by a general ring B ? See Exercise 7.3.O for an answer.)

7.3.4. Morphisms to affine schemes.

The following result shows that it is easy to describe morphisms to an affine scheme without working hard to cover the source with affine open sets.

7.3.G. ESSENTIAL EXERCISE. Show that morphisms $X \rightarrow \text{Spec } A$ are in natural bijection with ring morphisms $A \rightarrow \Gamma(X, \mathcal{O}_X)$. Hint: Show that this is true when X is affine. Use the fact that morphisms glue, Exercise 7.3.B. (This is even true in the category of locally ringed spaces. You are free to prove it in this generality, but it is easier in the category of schemes.)

In particular, there is a canonical morphism from a scheme to Spec of its ring of global sections. (Warning: Even if X is a finite type k -scheme, the ring of global sections might be nasty! In particular, it might not be finitely generated, see 20.11.13.) The canonical morphism $X \rightarrow \text{Spec } \Gamma(X, \mathcal{O}_X)$ is an isomorphism if and only if X is affine (i.e., isomorphic to $\text{Spec } A$ for some ring A), and in this case it is the isomorphism hinted at in Remark 4.3.2.

7.3.H. EASY EXERCISE. Show that this definition of A -scheme given in §7.3.7 agrees with the earlier definition of §5.3.6.

7.3.I. EASY EXERCISE. If S_\bullet is a finitely generated graded A -algebra, describe a natural “structure morphism” $\text{Proj } S_\bullet \rightarrow \text{Spec } A$.

7.3.J. EASY EXERCISE. Show that $\text{Spec } \mathbb{Z}$ is the final object in the category of schemes. In other words, if X is any scheme, there exists a unique morphism to $\text{Spec } \mathbb{Z}$. (Hence the category of schemes is isomorphic to the category of \mathbb{Z} -schemes.) If A is any ring (for example, a field k), show that $\text{Spec } A$ is the final object in the category of A -schemes.

7.3.5. Remark. From Essential Exercise 7.3.G, it is one small step to show that some products of schemes exist: if A and B are rings, then $\text{Spec } A \times \text{Spec } B = \text{Spec}(A \otimes_{\mathbb{Z}} B)$; and if A and B are C -algebras, then $\text{Spec } A \times_{\text{Spec } C} \text{Spec } B = \text{Spec}(A \otimes_C B)$. But we are in no hurry, so we wait until Exercise 10.1.A to discuss this properly.

7.3.6. ★★ Side fact for experts: Γ and Spec are adjoints. We have a contravariant functor Spec from rings to locally ringed spaces, and a contravariant functor Γ from locally ringed spaces to rings. In fact (Γ, Spec) is an adjoint pair! (Caution: we have only discussed adjoints for covariant functors; if you care, you will have to figure out how to define adjoints for contravariant functors.) Thus we could have defined Spec by requiring it to be right-adjoint to Γ . (Fun but irrelevant side question: if you used ringed spaces rather than locally ringed spaces, Γ again has a right adjoint. What is it?)

Our ability to easily describe morphisms to affine schemes will allow us to revisit our earlier discussions of schemes over a given field or ring (§5.3.6).

7.3.7. The category of complex schemes (or more generally the category of k -schemes where k is a field, or more generally the category of A -schemes where A is a ring, or more generally the category of S -schemes where S is a scheme). The category of S -schemes Sch_S (where S is a scheme) is defined as follows. The objects (S -schemes) are morphisms of the form

$$\begin{array}{c} X \\ \downarrow \\ S \end{array}$$

(The morphism to S is called the **structure morphism**. A motivation for this terminology is the fact that if $S = \text{Spec } A$, the structure morphism gives the functions on each open set of X the structure of an A -algebra, cf. §5.3.6.) The morphisms in the category of S -schemes are defined to be commutative diagrams

$$\begin{array}{ccc} X & \longrightarrow & Y \\ \downarrow & & \downarrow \\ S & \xrightarrow{=} & S \end{array}$$

which is more conveniently written as a commutative diagram

$$\begin{array}{ccc} X & \xrightarrow{\quad} & Y \\ & \searrow & \swarrow \\ & S & \end{array}$$

When there is no confusion (if the base scheme is clear), simply the top row of the diagram is given. In the case where $S = \operatorname{Spec} A$, where A is a ring, we get the notion of an A -scheme, which is the same as the same definition as in §5.3.6 (Exercise 7.3.H), but in a more satisfactory form. For example, complex geometers may consider the category of \mathbb{C} -schemes.

7.3.8. Coproducts of schemes.

7.3.K. EASY EXERCISE. Suppose X and Y are schemes. Show that their disjoint union $X \coprod Y$ (defined in Exercise 4.3.E) is the coproduct of X and Y in the category of schemes (justifying the use of the coproduct symbol \coprod).

7.3.9. Morphisms from (some) affine schemes.

Morphisms *from* affine schemes are not quite as simple as morphisms *to* affine schemes, but some cases are worth pointing out.

7.3.L. EXERCISE.

- (a) Suppose p is a point of a scheme X . Describe a canonical (choice-free) morphism $\operatorname{Spec} \mathcal{O}_{X,p} \rightarrow X$. (Hint: do this for affine X first. But then for general X be sure to show that your morphism is independent of choice.)
- (b) Define a canonical morphism $\operatorname{Spec} \kappa(p) \rightarrow X$. (This is often written $p \rightarrow X$; one gives p the obvious interpretation as a scheme, $\operatorname{Spec} \kappa(p)$.)

7.3.M. EXERCISE (MORPHISMS FROM Spec OF A LOCAL RING TO X). Suppose X is a scheme, and (A, \mathfrak{m}) is a local ring. Suppose we have a scheme morphism $\pi: \operatorname{Spec} A \rightarrow X$ sending $[\mathfrak{m}]$ to p . Show that any open set containing p contains the image of π . Show that there is a bijection between $\operatorname{Mor}(\operatorname{Spec} A, X)$ and

$$\{(p \in X, \text{local homomorphisms } \mathcal{O}_{X,p} \rightarrow A)\}.$$

(Possible hint: Exercise 7.3.L(a).)

These exercises lead us to the notion of field-valued points, and more generally ring-valued points, and more generally still, scheme-valued points.

7.3.10. Definition: The functor of points, and scheme-valued points (and ring-valued points, and field-valued points) of a scheme. If Z is a scheme, then Z -valued points of a scheme X , denoted $X(Z)$, are defined to be maps $Z \rightarrow X$. If A is a ring, then A -valued points of a scheme X , denoted $X(A)$, are defined to be the $(\operatorname{Spec} A)$ -valued points of the scheme. (The most common case of this is when A is a field.)

If you are working over a base scheme B — for example, complex algebraic geometers will consider only schemes and morphisms over $B = \operatorname{Spec} \mathbb{C}$ — then in the above definition, there is an implicit structure map $Z \rightarrow B$ (or $\operatorname{Spec} A \rightarrow B$ in the case of $X(A)$). For example, for a complex geometer, if X is a scheme over \mathbb{C} ,

the $\mathbb{C}(t)$ -valued points of X correspond to commutative diagrams of the form

$$\begin{array}{ccc} \mathrm{Spec} \mathbb{C}(t) & \xrightarrow{\quad} & X \\ & \searrow \xi & \swarrow \pi \\ & \mathrm{Spec} \mathbb{C} & \end{array}$$

where $\pi: X \rightarrow \mathrm{Spec} \mathbb{C}$ is the structure map for X , and ξ corresponds to the obvious inclusion of rings $\mathbb{C} \rightarrow \mathbb{C}(t)$. (Warning: a k -valued point of a k -scheme X is sometimes called a “rational point” of X , which is dangerous, as for most of the world, “rational” refers to \mathbb{Q} . We will use the safer phrase “ k -valued point” of X . Another safe choice is “ k -rational point” of X or k -point.)

The terminology “ Z -valued point” (and A -valued point) is unfortunate, because we earlier defined the notion of points of a scheme, and Z -valued points (and A -valued points) are not (necessarily) points! But these usages are well-established in the literature. (Look in the index under “point” to see even more inconsistent use of adjectives that modify this word.)

7.3.N. EXERCISE.

(a) (easy) Show that a morphism of schemes $X \rightarrow Y$ induces a map of Z -valued points $X(Z) \rightarrow Y(Z)$.

(b) Note that morphisms of schemes $X \rightarrow Y$ are not determined by their “underlying” map of points. (What is an example?) Show that they *are* determined by their induced maps of Z -valued points, as Z varies over all schemes. (Hint: pick $Z = X$. In the course of doing this exercise, you will largely prove Yoneda’s Lemma in the guise of Exercise 10.1.B.)

7.3.11. Furthermore, we will see that “products of Z -valued points” behave as you might hope (§10.1.3). A related reason this language is suggestive: the notation $X(Z)$ suggests the interpretation of X as a (contravariant) functor h_X from schemes to sets — the **functor of (scheme-valued) points** of the scheme X (cf. Example 1.2.20).

Here is a low-brow reason A -valued points are a useful notion: *the A -valued points of an affine scheme $\mathrm{Spec} \mathbb{Z}[x_1, \dots, x_n]/(f_1, \dots, f_r)$ (where $f_i \in \mathbb{Z}[x_1, \dots, x_n]$ are relations) are precisely the solutions to the equations*

$$f_1(x_1, \dots, x_n) = \dots = f_r(x_1, \dots, x_n) = 0$$

in the ring A . For example, the rational solutions to $x^2 + y^2 = 16$ are precisely the \mathbb{Q} -valued points of $\mathrm{Spec} \mathbb{Z}[x, y]/(x^2 + y^2 - 16)$. The integral solutions are precisely the \mathbb{Z} -valued points. So A -valued points of an affine scheme (finite type over \mathbb{Z}) can be interpreted simply. In the special case where A is local, A -valued points of a general scheme have a good interpretation too (Exercises 7.3.L and 7.3.M).

On the other hand, Z -valued points of projective space can be subtle. There are some maps we can write down easily, as shown by applying the next exercise in the case $X = \mathrm{Spec} A$, where A is a B -algebra.

7.3.O. EASY (BUT SURPRISINGLY ENLIGHTENING) EXERCISE (CF. EXERCISE 7.3.F).

(a) Suppose B is a ring. If X is a B -scheme, and f_0, \dots, f_n are $n + 1$ functions on X with no common zeros, then show that $[f_0, \dots, f_n]$ gives a morphism of B -schemes $X \rightarrow \mathbb{P}_B^n$.

(b) Suppose g is a nowhere vanishing function on X , and f_i are as in part (a). Show that the morphisms $[f_0, \dots, f_n]$ and $[gf_0, \dots, gf_n]$ to \mathbb{P}_B^n are the same.

7.3.12. Example (Exercise 7.3.F revisited). Consider the $n + 1$ functions x_0, \dots, x_n on \mathbb{A}^{n+1} (otherwise known as $n + 1$ sections of the trivial bundle). They have no common zeros on $\mathbb{A}^{n+1} - 0$. Hence they determine a morphism $\mathbb{A}^{n+1} - 0 \rightarrow \mathbb{P}^n$.

7.3.13. You might hope that Exercise 7.3.O(a) gives all morphisms to projective space (over B). But this isn't the case. Indeed, even the identity morphism $X = \mathbb{P}_k^1 \rightarrow \mathbb{P}_k^1$ isn't of this form, as the source \mathbb{P}^1 has no nonconstant global functions with which to build this map. (There are similar examples with an affine source.) However, there is a correct generalization (characterizing *all* maps from schemes to projective schemes) in Theorem 17.4.1. This result roughly states that this works, so long as the f_i are not quite functions, but sections of a line bundle. Our desire to understand maps to projective schemes in a clean way will be one important motivation for understanding line bundles.

We will see more ways to describe maps to projective space in the next section. A different description directly generalizing Exercise 7.3.O(a) will be given in Exercise 16.3.G, which will turn out (in Theorem 17.4.1) to be a “universal” description.

Incidentally, before Grothendieck, it was considered a real problem to figure out the right way to interpret points of projective space with “coordinates” in a ring. These difficulties were due to a lack of functorial reasoning. And the clues to the right answer already existed (the same problems arise for maps from a manifold to \mathbb{RP}^n) — if you ask such a geometric question (for projective space is geometric), the answer is necessarily geometric, not purely algebraic!

7.3.14. Visualizing morphisms: Picturing maps of schemes when nilpotents are present. You now know how to visualize the points of schemes (§3.3), and nilpotents (§4.2 and §6.5). The following imprecise exercise will give you some sense of how to visualize maps of schemes when nilpotents are involved. Suppose $a \in \mathbb{C}$. Consider the map of rings $\mathbb{C}[x] \rightarrow \mathbb{C}[\epsilon]/(\epsilon^2)$ given by $x \mapsto a\epsilon$. Recall that $\text{Spec } \mathbb{C}[\epsilon]/(\epsilon^2)$ may be pictured as a point with a tangent vector (§4.2). How would you picture this map if $a \neq 0$? How does your picture change if $a = 0$? (The tangent vector should be “crushed” in this case.)

Exercise 13.1.I will extend this considerably. You may enjoy reading its statement now.

7.3.15. ★★ For readers with appropriate background: Analytification of complex algebraic varieties. Warning: Any discussion of analytification is only for readers who are familiar with the notion of complex analytic varieties, or willing to develop it on their own in parallel with our development of schemes.

7.3.P. EXERCISE (ANALYTIFYING COMPLEX ALGEBRAIC VARIETIES). Suppose X is a reduced, finite type \mathbb{C} -scheme. Define the corresponding complex analytic prevariety X_{an} . (The definition of an analytic prevariety is the same as the definition of a variety without the Hausdorff condition.) Caution: your definition should not depend on a choice of an affine cover of X . (Hint: First explain how to analytify reduced finite type affine \mathbb{C} -schemes. Then glue.) Give a bijection between the closed points of X and the points of X_{an} , using the weak Nullstellensatz 3.2.5.

(In fact one may construct a continuous map of sets $X_{\text{an}} \rightarrow X$ generalizing Exercise 3.2.I.)

7.3.Q. EXERCISE (THE ANALYTIFICATION FUNCTOR). Recall the analytification construction of Exercise 7.3.16. For each morphism of reduced finite type \mathbb{C} -schemes $\pi: X \rightarrow Y$ (over \mathbb{C}), define a morphism of complex analytic prevarieties $\pi_{\text{an}}: X_{\text{an}} \rightarrow Y_{\text{an}}$ (the **analytification** of π). Show that analytification gives a functor from the category of reduced finite type \mathbb{C} -schemes to the category of complex analytic prevarieties.

7.3.16. Remark. Two nonisomorphic varieties can have isomorphic analytifications. For example, Serre described two different algebraic structures on the complex manifold $\mathbb{C}^* \times \mathbb{C}^*$, see [Ha2, p. 232] and [MO68421]; one is “the obvious one”, and the other is a \mathbb{P}^1 -bundle over an elliptic curve, with a section removed. For an example of a smooth complex surface with infinitely many algebraic structures, see §20.11.3. On the other hand, a compact complex variety can have only one algebraic structure (see [Se3, §19]).

7.3.17. Further facts about analytification. For more on analytification, see Exercises 11.3.L (Hausdorffness), 14.2.I, 19.4.H, 19.4.U (line bundles), 19.6.4, 21.1.7 (degree), 22.3.B (manifolds). For background on complex analytic spaces, see [GR].

7.4 Maps of graded rings and maps of projective schemes

As maps of rings correspond to maps of affine schemes in the opposite direction, maps of graded rings (over a base ring A) sometimes give maps of projective schemes in the opposite direction. This is an imperfect generalization: not every map of graded rings gives a map of projective schemes (§7.4.2); not every map of projective schemes comes from a map of graded rings (§20.11.9); and different maps of graded rings can yield the same map of schemes (Exercise 7.4.C).

You may find it helpful to think through Examples 7.4.1 and 7.4.2 while working through the following exercise.

7.4.A. ESSENTIAL EXERCISE. Suppose that $\phi: S_{\bullet} \longrightarrow R_{\bullet}$ is a morphism of $(\mathbb{Z}^{\geq 0})$ -graded rings. (By **map of graded rings**, we mean a map of rings that preserves the grading as a map of “graded semigroups” (or “graded monoids”). In other words, there is a $d > 0$ such that S_n maps to R_{dn} for all n .) Show that this induces a morphism of schemes

$$(\text{Proj } R_{\bullet}) \setminus V(\phi(S_+)) \rightarrow \text{Proj } S_{\bullet}.$$

(Hint: Suppose f is a homogeneous element of S_+ . Define a map $D(\phi(f)) \rightarrow D(f)$. Show that they glue together (as f runs over all homogeneous elements of S_+). Show that this defines a map from all of $\text{Proj } R_{\bullet} \setminus V(\phi(S_+))$. In particular, if

$$(7.4.0.1) \quad V(\phi(S_+)) = \emptyset,$$

then we have a morphism $\text{Proj } R_{\bullet} \rightarrow \text{Proj } S_{\bullet}$. From your solution, it will be clear that if ϕ is furthermore a morphism of A -algebras, then the induced morphism $\text{Proj } R_{\bullet} \setminus V(\phi(S_+)) \rightarrow \text{Proj } S_{\bullet}$ is a morphism of A -schemes.

7.4.1. Example. Let's see Exercise 7.4.A in action. We will scheme-theoretically interpret the map of complex projective manifolds \mathbb{CP}^1 to \mathbb{CP}^2 given by

$$\mathbb{CP}^1 \longrightarrow \mathbb{CP}^2$$

$$[s, t] \longmapsto [s^{20}, s^9 t^{11}, t^{20}]$$

Notice first that this is well-defined: $[\lambda s, \lambda t]$ is sent to the same point of \mathbb{CP}^2 as $[s, t]$. The reason for it to be well-defined is that the three polynomials s^{20} , $s^9 t^{11}$, and t^{20} are all homogeneous of degree 20.

Algebraically, this corresponds to a map of graded rings in the opposite direction

$$\mathbb{C}[x, y, z] \rightarrow \mathbb{C}[s, t]$$

given by $x \mapsto s^{20}$, $y \mapsto s^9 t^{11}$, $z \mapsto t^{20}$. You should interpret this in light of your solution to Exercise 7.4.A, and compare this to the affine example of §3.2.11.

7.4.2. Example. Notice that there is no map of complex manifolds $\mathbb{CP}^2 \rightarrow \mathbb{CP}^1$ given by $[x, y, z] \mapsto [x, y]$, because the map is not defined when $x = y = 0$. This corresponds to the fact that the map of graded rings $\mathbb{C}[s, t] \rightarrow \mathbb{C}[x, y, z]$ given by $s \mapsto x$ and $t \mapsto y$, doesn't satisfy hypothesis (7.4.0.1).

7.4.B. EXERCISE. Show that if $\phi: S_\bullet \rightarrow R_\bullet$ satisfies $\sqrt{(\phi(S_+))} = R_+$, then hypothesis (7.4.0.1) is satisfied. (Hint: Exercise 4.5.L.) This algebraic formulation of the more geometric hypothesis can sometimes be easier to verify.

7.4.C. UNIMPORTANT EXERCISE. This exercise shows that different maps of graded rings can give the same map of schemes. Let $R_\bullet = k[x, y, z]/(xz, yz, z^2)$ and $S_\bullet = k[a, b, c]/(ac, bc, c^2)$, where every variable has degree 1. Show that $\text{Proj } R_\bullet \cong \text{Proj } S_\bullet \cong \mathbb{P}_k^1$. Show that the maps $S_\bullet \rightarrow R_\bullet$ given by $(a, b, c) \mapsto (x, y, z)$ and $(a, b, c) \mapsto (x, y, 0)$ give the same (iso)morphism $\text{Proj } R_\bullet \rightarrow \text{Proj } S_\bullet$. (The real reason is that all of these constructions are insensitive to what happens in a finite number of degrees. This will be made precise in a number of ways later, most immediately in Exercise 7.4.F.)

7.4.3. Unimportant remark. Exercise 17.4.I shows that not every morphism of schemes $\text{Proj } R_\bullet \rightarrow \text{Proj } S_\bullet$ comes from a map of graded rings $S_\bullet \rightarrow R_\bullet$, even in quite reasonable circumstances.

7.4.4. Veronese subrings.

Here is a useful construction. Suppose S_\bullet is a finitely generated graded ring. Define the n th **Veronese subring** of S_\bullet by $S_{n\bullet} = \bigoplus_{j=0}^{\infty} S_{nj}$. (The “old degree” n is “new degree” 1.) The geometric interpretation of this construction is the important *Veronese embedding*, discussed in §9.3.6.

7.4.D. EXERCISE. Show that the map of graded rings $S_{n\bullet} \hookrightarrow S_\bullet$ induces an isomorphism $\text{Proj } S_\bullet \xrightarrow{\sim} \text{Proj } S_{n\bullet}$. (Hint: if $f \in S_+$ is homogeneous of degree divisible by n , identify $D(f)$ on $\text{Proj } S_\bullet$ with $D(f)$ on $\text{Proj } S_{n\bullet}$. Why do such distinguished open sets cover $\text{Proj } S_\bullet$?)

7.4.E. EXERCISE. If S_\bullet is generated in degree 1, show that $S_{n\bullet}$ is also generated in degree 1. (You may want to consider the case of the polynomial ring first.)

7.4.F. EXERCISE. Show that if R_\bullet and S_\bullet are the same finitely generated graded rings except in a finite number of nonzero degrees (make this precise!), then $\text{Proj } R_\bullet \cong \text{Proj } S_\bullet$.

7.4.G. EXERCISE. Suppose S_\bullet is generated over S_0 by f_1, \dots, f_n . Find a d such that $S_{d\bullet}$ is finitely generated in “new” degree 1 (= “old” degree d). (This is surprisingly tricky, so here is a hint. Suppose there are generators x_1, \dots, x_n of degrees d_1, \dots, d_n respectively. Show that any monomial $x_1^{a_1} \cdots x_n^{a_n}$ of degree at least $nd_1 \dots d_n$ has $a_i \geq (\prod_j d_j)/d_i$ for some i . Show that the $(nd_1 \dots d_n)$ th Veronese subring is generated by elements in “new” degree 1.)

Exercise 7.4.G, in combination with Exercise 7.4.D, shows that there is little harm in assuming that finitely generated graded rings are generated in degree 1, as after a regrading (or more precisely, keeping only terms of degree a multiple of d , then dividing the degree by d), this is indeed the case. This is handy, as it means that, using Exercise 7.4.D, we can assume that any finitely generated graded ring is generated in degree 1. Exercise 9.3.F will later imply as a consequence that we can embed every Proj in some projective space.

7.4.H. LESS IMPORTANT EXERCISE. Suppose S_\bullet is a finitely generated graded ring. Show that $S_{n\bullet}$ is a finitely generated graded ring. (Possible approach: use the previous exercise, or something similar, to show there is some N such that $S_{nN\bullet}$ is generated in degree 1, so the graded ring $S_{nN\bullet}$ is finitely generated. Then show that for each $0 < j < N$, $S_{nN\bullet+nj}$ is a finitely generated module over $S_{nN\bullet}$.)

7.5 Rational maps from reduced schemes

Informally speaking, a “rational map” is “a morphism defined almost everywhere”, much as a rational function (Definition 5.2.I) is a name for a function defined almost everywhere. We will later see that in good situations, just as with rational functions, where a rational map is defined, it is uniquely defined (the Reduced-to-Separated Theorem 11.4.2), and has a largest “domain of definition” (§11.4.3). For this section only, we assume X to be reduced. A key example will be irreducible varieties (§7.5.6), and the language of rational maps is most often used in this case.

7.5.1. Definition. A **rational map** π from X to Y , denoted $\pi: X \dashrightarrow Y$, is the data of a morphism $\alpha: U \rightarrow Y$ from a dense open set $U \subset X$, with the equivalence relation $(\alpha: U \rightarrow Y) \sim (\beta: V \rightarrow Y)$ if there is a dense open set $Z \subset U \cap V$ such that $\alpha|_Z = \beta|_Z$. (In §11.4.3, we will improve this to: if $\alpha|_{U \cap V} = \beta|_{U \cap V}$ in good circumstances — when Y is separated.) People often use the word “map” for “morphism”, which is quite reasonable, except that a rational map need not be a map. So to avoid confusion, when one means “rational map”, one should never just say “map”.

We will also discuss rational maps of S -schemes for a scheme S . The definition is the same, except now X and Y are S -schemes, and $\alpha: U \rightarrow Y$ is a morphism of S -schemes.

7.5.2. \star Rational maps more generally. Just as with rational functions, Definition 7.5.1 can be extended to where X is not reduced, as is (using the same name, “rational map”), or in a version that imposes some control over what happens over the nonreduced locus (*pseudo-morphisms*, [Stacks, tag 01RX]). We will see in §11.4 that rational maps from reduced schemes to separated schemes behave particularly well, which is why they are usually considered in this context. The reason for the definition of pseudo-morphisms is to extend these results to when X is nonreduced. We will not use the notion of pseudo-morphism.

7.5.3. An obvious example of a rational map is a morphism. Another important example is the projection $\mathbb{P}_A^n \dashrightarrow \mathbb{P}_A^{n-1}$ given by $[x_0, \dots, x_n] \mapsto [x_0, \dots, x_{n-1}]$. (How precisely is this a rational map in the sense of Definition 7.5.1?)

A rational map $\pi: X \dashrightarrow Y$ is **dominant** (or in some sources, *dominating*) if for some (and hence every) representative $U \rightarrow Y$, the image is dense in Y . A morphism is a **dominant morphism** (or *dominating morphism*) if it is dominant as a rational map.

7.5.A. EXERCISE. Show that a rational map $\pi: X \dashrightarrow Y$ of irreducible schemes is dominant if and only if π sends the generic point of X to the generic point of Y .

A little thought will convince you that you can compose (in a well-defined way) a dominant map $\pi: X \dashrightarrow Y$ from an irreducible scheme X to an irreducible scheme Y with a rational map $\rho: Y \dashrightarrow Z$. Furthermore, the composition $\rho \circ \pi$ will be dominant if ρ is dominant. Integral schemes and dominant rational maps between them form a category which is geometrically interesting.

7.5.B. EASY EXERCISE. Show that dominant rational maps of integral schemes give morphisms of function fields in the opposite direction.

In “suitably classical situations” (integral finite type k -schemes — and in particular, irreducible varieties, to be defined in §11.3.9) — this is reversible: dominant rational maps correspond to inclusions of function fields in the opposite direction. We make this precise in §7.5.6 below. But it is not true that morphisms of function fields *always* give dominant rational maps, or even rational maps. For example, $\operatorname{Spec} k[x]$ and $\operatorname{Spec} k(x)$ have the same function field $k(x)$, but there is no corresponding rational map $\operatorname{Spec} k[x] \dashrightarrow \operatorname{Spec} k(x)$ of k -schemes. Reason: such a rational map would correspond to a morphism from an open subset U of $\operatorname{Spec} k[x]$, say $\operatorname{Spec} k[x, 1/f(x)]$, to $\operatorname{Spec} k(x)$. But there is no map of rings $k(x) \rightarrow k[x, 1/f(x)]$ (sending k identically to k and x to x) for any one $f(x)$.

(If you want more evidence that the topologically-defined notion of dominance is simultaneously algebraic, you can show that if $\phi: A \rightarrow B$ is a ring morphism, then the corresponding morphism $\operatorname{Spec} B \rightarrow \operatorname{Spec} A$ is dominant if and only if ϕ has kernel contained in the nilradical of A .)

7.5.4. Definition. A rational map $\pi: X \dashrightarrow Y$ is said to be **birational** if it is dominant, and there is another rational map (a “rational inverse”) ψ that is also dominant, such that $\pi \circ \psi$ is (in the same equivalence class as) the identity on Y , and $\psi \circ \pi$ is (in the same equivalence class as) the identity on X . This is the notion of isomorphism in the category of integral schemes and dominant rational maps. (In the differentiable category, this is not particularly interesting.) A *morphism* is a

birational morphism if it is birational as a rational map. (Note that the “inverse” to a birational morphism may be only a rational map, not a morphism.)

We say X and Y are **birational** (to each other) if there exists a birational map $X \dashrightarrow Y$. If X and Y are irreducible, then birational maps induce isomorphisms of function fields. The fact that maps of function fields correspond to rational maps in the opposite direction for integral finite type k -schemes, to be proved in Proposition 7.5.7, shows that a map between integral finite type k -schemes that induces an isomorphism of function fields is birational. An integral finite type k -scheme is said to be **rational** if it is birational to \mathbb{A}_k^n for some n .

7.5.5. Proposition. — *Suppose X and Y are reduced schemes. Then X and Y are birational if and only if there is a dense open subscheme U of X and a dense open subscheme V of Y such that $U \cong V$.*

(We have “reduced” hypotheses because we have only defined rational maps in this setting. If you wish, you can extend this appropriately once you know about associated points, §6.5.)

Proposition 7.5.5 tells you how to think of birational maps. Just as a rational map is a “mostly-defined function”, two birational reduced schemes are “mostly isomorphic”. For example, a reduced finite type k -scheme (such as an affine k -variety) is rational if and only if it has a dense open subscheme isomorphic to an open subscheme of \mathbb{A}_k^n .

Proof. The “if” direction is immediate, so we prove the “only if” direction. We basically follow our nose, and use what we are given (see Figure 7.1).




FIGURE 7.1. Outline of the proof of Proposition 7.5.5

We have inverse rational maps $F : X \dashrightarrow Y$ and $G : Y \dashrightarrow X$. Choose representative *morphisms* $f : X_1 \rightarrow Y$ (where $X_1 \subset X$) and $g : Y_1 \rightarrow X$ (where $Y_1 \subset Y$) for the rational maps F and G respectively. If $X_2 = (f|_{X_1})^{-1}(Y_1) \subset X_1$, then $g \circ f$ is a *morphism* from X_2 to X that is the identity *as a rational map*. Thus there is a dense open subset $X_3 \subset X_2$ such that the morphism $g \circ f : X_3 \rightarrow X$ is the identity morphism onto X_3 (or more precisely, it is the open embedding $X_3 \hookrightarrow X$).

Similarly, let $Y_2 = (g|_{Y_1})^{-1}(X_3) \subset Y_1$, so $f \circ g$ is a morphism from Y_2 to Y that is the identity as a rational map. Then let $Y_3 \subset Y_2$ be a dense open subset such that $f \circ g : Y_3 \rightarrow Y$ is the inclusion (the “identity”).

Finally, if $X_4 = (f|_{X_3})^{-1}(Y_3) \subset X_3$, then $(g \circ f)|_{X_4}$ is the identity morphism on X_4 (by way of Y_3), and $(f \circ g)|_{Y_3}$ is the identity morphism on Y_3 (by way of X_4), so we have found our isomorphism of open sets that is a “representative” for our birational map. \square

7.5.6. Rational maps of irreducible varieties.

7.5.7. Proposition. — *Suppose X is an integral k -scheme and Y is an integral finite type k -scheme, and we are given an extension of function fields $\phi^\sharp : K(Y) \hookrightarrow K(X)$ preserving k . Then there exists a dominant rational map of k -schemes $\phi : X \dashrightarrow Y$ inducing ϕ^\sharp .*

Proof. By replacing Y with an open subset, we may assume that Y is affine, say $\text{Spec } B$, where B is generated over k by finitely many elements y_1, \dots, y_n . Since we only need to define ϕ on an open subset of X , we may similarly assume that $X = \text{Spec } A$ is affine. Then ϕ^\sharp gives an inclusion $\phi^\sharp : B \hookrightarrow K(A)$. Write $\phi^\sharp(y_i)$ as f_i/g_i ($f_i, g_i \in A$), and let $g := \prod g_i$. Then ϕ^\sharp further induces an inclusion $B \hookrightarrow A_g$. Therefore $\phi : \text{Spec } A_g \rightarrow \text{Spec } B$ induces ϕ^\sharp . The morphism ϕ is dominant because the inverse image of the zero ideal under the inclusion $B \hookrightarrow A_g$ is the zero ideal, so ϕ takes the generic point of X to the generic point of Y . \square

7.5.C. EXERCISE. Let K be a finitely generated field extension of k . (Recall that a field extension K over k is **finitely generated** if there is a finite “generating set” x_1, \dots, x_n in K such that every element of K can be written as a rational function in x_1, \dots, x_n with coefficients in k .) Show that there exists an irreducible affine k -variety with function field K . (Hint: Consider the map $k[t_1, \dots, t_n] \rightarrow K$ given by $t_i \mapsto x_i$, and show that the kernel is a prime ideal \mathfrak{p} , and that $k[t_1, \dots, t_n]/\mathfrak{p}$ has fraction field K . Interpreted geometrically: consider the map $\text{Spec } K \rightarrow \text{Spec } k[t_1, \dots, t_n]$ given by the ring map $t_i \mapsto x_i$, and take the closure of the one-point image.)

7.5.D. EXERCISE. Describe equivalences of categories among the following.

- (a) the category with objects “integral *affine* k -varieties”, and morphisms “dominant rational maps defined over k ”; and
- (b) the opposite (“arrows-reversed”) category with objects “finitely generated field extensions of k ”, and morphisms “inclusions extending the identity on k ”.

(Once we define varieties in general in §11.3.9, you can add in: (c) the category with objects “integral k -varieties”, and morphisms “dominant rational maps defined over k ”.)

In particular, an integral affine k -variety X is rational if its function field $K(X)$ is a purely transcendental extension of k , i.e., $K(X) \cong k(x_1, \dots, x_n)$ for some n . (This needs to be said more precisely: the map $k \hookrightarrow K(X)$ induced by $X \rightarrow \operatorname{Spec} k$ should agree with the “obvious” map $k \hookrightarrow k(x_1, \dots, x_n)$ under this isomorphism.)

7.5.8. More examples of rational maps.

A recurring theme in these examples is that domains of definition of rational maps to projective schemes extend over regular codimension one points. We will make this precise in the Curve-to-Projective Extension Theorem 17.5.1, when we discuss curves.




FIGURE 7.2. Finding primitive Pythagorean triples using geometry

The first example is the classical formula for Pythagorean triples, and its derivation by “stereographic projection”. Suppose you are looking for rational points on the circle C given by $x^2 + y^2 = 1$ (Figure 7.2). One rational point is $p = (1, 0)$. If q is another rational point, then pq is a line of rational (non-infinite) slope. This gives a rational map from the conic C (now interpreted as $\operatorname{Spec} \mathbb{Q}[x, y]/(x^2 + y^2 - 1)$) to $\mathbb{A}_{\mathbb{Q}}^1$, given by $(x, y) \mapsto y/(x - 1)$. (Something subtle just happened: we were talking about \mathbb{Q} -points on a circle, and ended up with a rational map of schemes.) Conversely, given a line of slope m through p , where m is rational, we can recover q by solving the equations $y = m(x - 1)$, $x^2 + y^2 = 1$. We substitute the first equation into the second, to get a quadratic equation in x . We know that we will have a solution $x = 1$ (because the line meets the circle at $(x, y) = (1, 0)$), so we expect to be able to factor this out, and find the other factor. This indeed works:

$$\begin{aligned} x^2 + (m(x - 1))^2 &= 1 \\ \implies (m^2 + 1)x^2 + (-2m^2)x + (m^2 - 1) &= 0 \\ \implies (x - 1)((m^2 + 1)x - (m^2 - 1)) &= 0 \end{aligned}$$

The other solution is $x = (m^2 - 1)/(m^2 + 1)$, which gives $y = -2m/(m^2 + 1)$. Thus we get a birational map between the conic C and \mathbb{A}^1 with coordinate m , given by $f: (x, y) \mapsto y/(x - 1)$ (which is defined for $x \neq 1$), and with inverse rational map given by $m \mapsto ((m^2 - 1)/(m^2 + 1), -2m/(m^2 + 1))$ (which is defined away from $m^2 + 1 = 0$).

We can extend this to a rational map $C \dashrightarrow \mathbb{P}_{\mathbb{Q}}^1$ via the “inclusion” $\mathbb{A}_{\mathbb{Q}}^1 \rightarrow \mathbb{P}_{\mathbb{Q}}^1$ (which we later call an open embedding). Then f is given by $(x, y) \mapsto [y, x - 1]$. We then have an interesting question: what is the domain of definition of f ? It appears to be defined everywhere except for where $y = x - 1 = 0$, i.e., everywhere but p . But in fact it can be extended over p ! Note that $(x, y) \mapsto [x + 1, -y]$ (where $(x, y) \neq (-1, 0)$) agrees with f on their common domains of definition, as $[x + 1, -y] = [y, x - 1]$. Hence this rational map can be extended farther than we at first thought. This will be a special case of the Curve-to-Projective Extension Theorem 17.5.1.

7.5.E. EXERCISE. Use the above to find a “formula” yielding all Pythagorean triples.

7.5.F. EXERCISE. Show that the conic $x^2 + y^2 = z^2$ in \mathbb{P}_k^2 is isomorphic to \mathbb{P}_k^1 for any field k of characteristic not 2. (See Exercise 4.5.S for the definition of “ $x^2 + y^2 = z^2$ in \mathbb{P}_k^2 ” if you need it.)

7.5.9. In fact, any conic in \mathbb{P}_k^2 with a k -valued point (i.e., a point with residue field k) of rank 3 (after base change to \bar{k} , so “rank” makes sense, see Exercise 5.4.J) is isomorphic to \mathbb{P}_k^1 . (The hypothesis of having a k -valued point is certainly necessary: $x^2 + y^2 + z^2 = 0$ over $k = \mathbb{R}$ is a conic that is not isomorphic to \mathbb{P}_k^1 .)

7.5.G. EXERCISE. Find all rational solutions to $y^2 = x^3 + x^2$, by finding a birational map to $\mathbb{A}_{\mathbb{Q}}^1$, mimicking what worked with the conic. Hint: what point should you project from? (In Exercise 20.10.H, we will see that these points basically form a group, and that this is a degenerate elliptic curve.)

You will obtain a rational map to $\mathbb{P}_{\mathbb{Q}}^1$ that is not defined over the node $x = y = 0$, and *cannot* be extended over this codimension 1 set. This is an example of the limits of our future result, the Curve-to-Projective Extension Theorem 17.5.1, showing how to extend rational maps to projective space over codimension 1 sets: the codimension 1 sets have to be regular.

7.5.H. EXERCISE. Use a similar idea to find a birational map from the quadric surface $Q = \{x^2 + y^2 = w^2 + z^2\} \subset \mathbb{P}_{\mathbb{Q}}^3$ to $\mathbb{P}_{\mathbb{Q}}^2$. Use this to find all rational points on Q . (This illustrates a good way of solving Diophantine equations. You will find a dense open subset of Q that is isomorphic to a dense open subset of \mathbb{P}^2 , where you can easily find all the rational points. There will be a closed subset of Q where the rational map is not defined, or not an isomorphism, but you can deal with this subset in an ad hoc fashion.)

7.5.I. EXERCISE (THE CREMONA TRANSFORMATION, A USEFUL CLASSICAL CONSTRUCTION). Consider the rational map $\mathbb{P}_k^2 \dashrightarrow \mathbb{P}_k^2$, given by $[x, y, z] \mapsto [1/x, 1/y, 1/z]$. What is the domain of definition? (It is bigger than the locus where $xyz \neq 0$!) You will observe that you can extend it over “codimension 1 sets” (ignoring the fact that we don’t yet know what codimension means). This again foreshadows the Curve-to-Projective Extension Theorem 17.5.1.

7.5.10. ★ Complex curves that are not rational (fun but inessential).

We now describe two examples of curves C that do not admit a nonconstant rational map from $\mathbb{P}_{\mathbb{C}}^1$. (Admittedly, we do not yet know what “curve” means, but

no matter.) Both proofs are by Fermat's method of *infinite descent*. These results can be interpreted (as you will later be able to check using Theorem 18.4.3) as the fact that these curves have no “nontrivial” $\mathbb{C}(t)$ -valued points, where by this, we mean that any $\mathbb{C}(t)$ -valued point is secretly a \mathbb{C} -valued point. You may notice that if you consider the same examples with $\mathbb{C}(t)$ replaced by \mathbb{Q} (and where C is a curve over \mathbb{Q} rather than \mathbb{C}), you get two fundamental questions in number theory and geometry. The analog of Exercise 7.5.L is the question of rational points on elliptic curves, and you may realize that the analog of Exercise 7.5.J is even more famous. Also, the arithmetic analog of Exercise 7.5.L(a) is the “Four Squares Theorem” (there are not four integer squares in arithmetic progression), first stated by Fermat. These examples will give you a glimpse of how and why facts over number fields are often paralleled by facts over function fields of curves. This parallelism is a recurring deep theme in the subject.

7.5.J. EXERCISE. If $n > 2$, show that $\mathbb{P}_{\mathbb{C}}^1$ has no dominant rational maps to the “Fermat curve” $x^n + y^n = z^n$ in $\mathbb{P}_{\mathbb{C}}^2$. Hint: reduce this to showing that there is no “nonconstant” solution $(f(t), g(t), h(t))$ to $f(t)^n + g(t)^n = h(t)^n$, where $f(t)$, $g(t)$, and $h(t)$ are rational functions in t (that is, elements of $\mathbb{C}(t)$). By clearing denominators, reduce this to showing that there is no nonconstant solution where $f(t)$, $g(t)$, and $h(t)$ are relatively prime polynomials. For this, assume there is a solution, and consider one of the lowest positive degree. Then use the fact that $\mathbb{C}[t]$ is a unique factorization domain, and $h(t)^n - g(t)^n = \prod_{i=1}^n (h(t) - \zeta^i g(t))$, where ζ is a primitive n th root of unity. Argue that each $h(t) - \zeta^i g(t)$ is an n th power. Then use

$$(h(t) - g(t)) + \alpha(h(t) - \zeta g(t)) = \beta(h(t) - \zeta^2 g(t))$$

for suitably chosen α and β to get a solution of smaller degree. (How does this argument fail for $n = 2$?)

7.5.K. EXERCISE. Give two smooth complex curves X and Y so that no nonempty open subset of X is isomorphic to a nonempty open subset of Y . (Try not to be bothered by the fact that we have not yet defined “smoothness”.) Hint: Exercise 7.5.J.

7.5.L. EXERCISE. Suppose a , b , and c are distinct complex numbers. By the following steps, show that if $x(t)$ and $y(t)$ are two rational functions of t such that

$$(7.5.10.1) \quad y(t)^2 = (x(t) - a)(x(t) - b)(x(t) - c),$$

then $x(t)$ and $y(t)$ are constants ($x(t), y(t) \in \mathbb{C}$). (Here \mathbb{C} may be replaced by any field K of characteristic not 2; slight extra care is needed if K is not algebraically closed.)

(a) Suppose $P, Q \in \mathbb{C}[t]$ are relatively prime polynomials such that four linear combinations of them are perfect squares, no two of which are constant multiples of each other. Show that P and Q are constant (i.e., $P, Q \in \mathbb{C}$). Hint: By renaming P and Q , show that you may assume that the perfect squares are P , Q , $P - Q$, $P - \lambda Q$ (for some $\lambda \in \mathbb{C}$). Define u and v to be square roots of P and Q respectively. Show that $u - v$, $u + v$, $u - \sqrt{\lambda}v$, $u + \sqrt{\lambda}v$ are perfect squares, and that u and v are relatively prime. If P and Q are not both constant, note that $0 < \max(\deg u, \deg v) < \max(\deg P, \deg Q)$. Assume from the start that P and Q

were chosen as a counterexample with minimal $\max(\deg P, \deg Q)$ to obtain a contradiction. (Aside: It is possible to have *three* distinct linear combinations that are perfect squares. Such examples essentially correspond to primitive Pythagorean triples in $\mathbb{C}(t)$ — can you see how?)

(b) Suppose $(x, y) = (p/q, r/s)$ is a solution to (7.5.10.1), where $p, q, r, s \in \mathbb{C}[t]$, and p/q and r/s are in lowest terms. Clear denominators to show that $r^2 q^3 = s^2(p - aq)(p - bq)(p - cq)$. Show that $s^2 | q^3$ and $q^3 | s^2$, and hence that $s^2 = \delta q^3$ for some $\delta \in \mathbb{C}$. From $r^2 = \delta(p - aq)(p - bq)(p - cq)$, show that $(p - aq)$, $(p - bq)$, $(p - cq)$ are perfect squares. Show that q is also a perfect square, and then apply part (a).

A much better geometric approach to Exercises 7.5.J and 7.5.L is given in Exercise 22.7.H.

7.6 ★ Representable functors and group schemes

7.6.1. Maps to \mathbb{A}^1 correspond to functions. If X is a scheme, there is a bijection between the maps $X \rightarrow \mathbb{A}^1$ and global sections of the structure sheaf: by Exercise 7.3.G, maps $\pi: X \rightarrow \mathbb{A}_{\mathbb{Z}}^1$ correspond to maps of rings $\pi^\sharp: \mathbb{Z}[t] \rightarrow \Gamma(X, \mathcal{O}_X)$, and $\pi^\sharp(t)$ is a function on X ; this is reversible.

This map is very natural in an informal sense: you can even picture this map to \mathbb{A}^1 as being *given* by the function. (By analogy, a function on a manifold is a map to \mathbb{R} .) But it is natural in a more precise sense: this bijection is functorial in X . We will ponder this example at length, and see that it leads us to two important sophisticated notions: representable functors and group schemes.

7.6.A. EASY EXERCISE. Suppose X is a \mathbb{C} -scheme. Verify that there is a natural bijection between maps $X \rightarrow \mathbb{A}_{\mathbb{C}}^1$ in the category of \mathbb{C} -schemes and functions on X . (Here the base ring \mathbb{C} can be replaced by any ring A .)

This interpretation can be extended to rational maps, as follows.

7.6.B. UNIMPORTANT EXERCISE. Interpret rational functions on an integral scheme (Exercise 5.2.I) as rational maps to $\mathbb{A}_{\mathbb{Z}}^1$.

(If you wish, you can extend your argument to rational maps on locally Noetherian schemes, see Definition 6.5.34.)

7.6.2. Representable functors. We restate the bijection of §7.6.1 as follows. We have two different contravariant functors from Sch to $Sets$: maps to \mathbb{A}^1 (i.e., $H: X \mapsto \text{Mor}(X, \mathbb{A}_{\mathbb{Z}}^1)$), and functions on X ($F: X \mapsto \Gamma(X, \mathcal{O}_X)$). The “naturality” of the bijection — the functoriality in X — is precisely the statement that the bijection gives a natural isomorphism of functors (§1.2.21): given any $\pi: X \rightarrow X'$, the diagram

$$\begin{array}{ccc} H(X') & \xrightarrow{H(\pi)} & H(X) \\ \downarrow & & \downarrow \\ F(X') & \xrightarrow{F(\pi)} & F(X) \end{array}$$

(where the vertical maps are the bijections given in §7.6.1) commutes.

More generally, if Y is an element of a category \mathcal{C} (we care about the special case $\mathcal{C} = \text{Sch}$), recall the contravariant functor $h_Y: \mathcal{C} \rightarrow \text{Sets}$ defined by $h_Y(X) = \text{Mor}(X, Y)$ (Example 1.2.20). We say a contravariant functor from \mathcal{C} to Sets is **represented by Y** if it is naturally isomorphic to the functor h_Y . We say it is **representable** if it is represented by *some* Y .

The bijection of §7.6.1 may now be restated as: *the global section functor is represented by \mathbb{A}^1 .*

7.6.C. IMPORTANT EASY EXERCISE (REPRESENTING OBJECTS ARE UNIQUE UP TO UNIQUE ISOMORPHISM). Show that if a contravariant functor F is represented by Y and by Z , then we have a unique isomorphism $Y \xrightarrow{\sim} Z$ induced by the natural isomorphism of functors $h_Y \xrightarrow{\sim} h_Z$. Hint: this is a version of the universal property arguments of §1.3: once again, we are recognizing an object (up to unique isomorphism) by maps to that object. This exercise is essentially Exercise 1.3.Z(b). (This extends readily to Yoneda's Lemma in this setting, Exercise 10.1.B. You are welcome to try that now.)

You have implicitly seen this notion before: you can interpret the existence of products and fibered products in a category as examples of representable functors. (You may wish to work out how a natural isomorphism $h_{Y \times Z} \xrightarrow{\sim} h_Y \times h_Z$ induces the projection maps $Y \times Z \rightarrow Y$ and $Y \times Z \rightarrow Z$.)

7.6.D. EXERCISE (WARM-UP). Suppose F is the contravariant functor $\text{Sch} \rightarrow \text{Sets}$ defined by $F(X) = \{\text{Grothendieck}\}$ for all schemes X . Show that F is representable. (What is it representable by?)

7.6.E. EXERCISE. In this exercise, \mathbb{Z} may be replaced by any ring.

(a) (*Affine n -space represents the functor of n functions.*) Show that the contravariant functor from $(\mathbb{Z}\text{-})$ schemes to Sets

$$X \longmapsto \{(f_1, \dots, f_n) : f_i \in \Gamma(X, \mathcal{O}_X)\}$$

is represented by $\mathbb{A}_{\mathbb{Z}}^n$. Show that $\mathbb{A}_{\mathbb{Z}}^1 \times_{\mathbb{Z}} \mathbb{A}_{\mathbb{Z}}^1 \cong \mathbb{A}_{\mathbb{Z}}^2$, i.e., that \mathbb{A}^2 satisfies the universal property of $\mathbb{A}^1 \times \mathbb{A}^1$. (You will undoubtedly be able to immediately show that $\prod \mathbb{A}_{\mathbb{Z}}^{m_i} \xrightarrow{\sim} \mathbb{A}_{\mathbb{Z}}^{\sum m_i}$.)

(b) (*The functor of invertible functions is representable.*) Show that the contravariant functor from $(\mathbb{Z}\text{-})$ schemes to Sets taking X to invertible functions on X is representable by $\text{Spec } \mathbb{Z}[t, t^{-1}]$.

7.6.3. Definition. The scheme defined in Exercise 7.6.E(b) is called the **multiplicative group** \mathbb{G}_m . “ \mathbb{G}_m over a field k ” (“the multiplicative group over k ”) means $\text{Spec } k[t, t^{-1}]$, with the same group operations. Better: it represents the group of invertible functions in the category of k -schemes. We can similarly define \mathbb{G}_m over an arbitrary ring or even arbitrary scheme.

7.6.F. LESS IMPORTANT EXERCISE. Fix a ring A . Consider the functor H from the category of locally ringed spaces to Sets given by $H(X) = \{A \rightarrow \Gamma(X, \mathcal{O}_X)\}$. Show that this functor is representable (by $\text{Spec } A$). This gives another (admittedly odd) motivation for the definition of $\text{Spec } A$, closely related to that of §7.3.6.

7.6.4. ** Group schemes (or more generally, group objects in a category).

(The rest of §7.6 should be read only for entertainment.) We return again to Example 7.6.1. Functions on X are better than a set: they form a group. (Indeed they even form a ring, but we will worry about this later.) Given a morphism $X \rightarrow Y$, pullback of functions $\Gamma(Y, \mathcal{O}_Y) \rightarrow \Gamma(X, \mathcal{O}_X)$ is a group homomorphism. So we should expect \mathbb{A}^1 to have some group-like structure. This leads us to the notion of *group scheme*, or more generally a *group object* in a category, which we now define.

Suppose \mathcal{C} is a category with a final object Z and with products. (We know that Sch has a final object $Z = \text{Spec } \mathbb{Z}$, by Exercise 7.3.J. We will later see that it has products, §10.1. But in Exercise 7.6.K we will give an alternative characterization of group objects that applies in any category, so we won't worry about this.)

A **group object** in \mathcal{C} is an element X along with three morphisms:

- *multiplication*: $m: X \times X \rightarrow X$
- *inverse*: $i: X \rightarrow X$
- *identity element*: $e: Z \rightarrow X$ (not the identity map)

These morphisms are required to satisfy several conditions.

(i) associativity axiom:

$$\begin{array}{ccc} X \times X \times X & \xrightarrow{(m, \text{id})} & X \times X \\ \downarrow (\text{id}, m) & & \downarrow m \\ X \times X & \xrightarrow{m} & X \end{array}$$

commutes. (Here id means the equality $X \rightarrow X$.)

(ii) identity axiom:

$$X \xrightarrow{\sim} Z \times X \xrightarrow{e \times \text{id}} X \times X \xrightarrow{m} X$$

and

$$X \xrightarrow{\sim} X \times Z \xrightarrow{\text{id} \times e} X \times X \xrightarrow{m} X$$

are both the identity map $X \rightarrow X$. (This corresponds to the group axiom: “multiplication by the identity element is the identity map”.)

(iii) inverse axiom: $X \xrightarrow{i, \text{id}} X \times X \xrightarrow{m} X$ and $X \xrightarrow{\text{id}, i} X \times X \xrightarrow{m} X$ are both the map that is the composition $X \xrightarrow{\sim} Z \xrightarrow{e} X$.

As motivation, you can check that a group object in the category of sets is in fact the same thing as a group. (This is symptomatic of how you take some notion and make it categorical. You write down its axioms in a categorical way, and if all goes well, if you specialize to the category of sets, you get your original notion. You can apply this to the notion of “rings” in an exercise below.)

A **group scheme** is defined to be a group object in the category of schemes. A **group scheme** over a ring A (or a scheme S) is defined to be a group object in the category of A -schemes (or S -schemes).

7.6.G. EXERCISE. Give $\mathbb{A}_{\mathbb{Z}}^1$ the structure of a group scheme, by describing the three structural morphisms, and showing that they satisfy the axioms. (Hint: the morphisms should not be surprising. For example, inverse is given by $t \mapsto -t$. Note that we know that the product $\mathbb{A}_{\mathbb{Z}}^1 \times \mathbb{A}_{\mathbb{Z}}^1$ exists, by Exercise 7.6.E(a).)

7.6.H. EXERCISE. Show that if G is a group object in a category \mathcal{C} , then for any $X \in \mathcal{C}$, $\text{Mor}(X, G)$ has the structure of a group, and the group structure is preserved by pullback (i.e., $\text{Mor}(\cdot, G)$ is a contravariant functor to *Groups*).

7.6.I. EXERCISE. Show that the group structure described by the previous exercise translates the group scheme structure on $\mathbb{A}_{\mathbb{Z}}^1$ to the group structure on $\Gamma(X, \mathcal{O}_X)$, via the bijection of §7.6.1.

7.6.J. EXERCISE. Define the notion of **abelian group scheme**, and **ring scheme**. (You will undoubtedly at the same time figure out how to define the notion of abelian group object and ring object in any category \mathcal{C} . You may discover a more efficient approach to such questions after reading §7.6.5.)

7.6.5. Group schemes, more functorially. There was something unsatisfactory about our discussion of the “group-respecting” nature of the bijection in §7.6.1: we observed that the right side (functions on X) formed a group, then we developed the axioms of a group scheme, then we cleverly figured out the maps that made $\mathbb{A}_{\mathbb{Z}}^1$ into a group scheme, then we showed that this induced a group structure on the left side of the bijection ($\text{Mor}(X, \mathbb{A}^1)$) that precisely corresponded to the group structure on the right side (functions on X).

The picture is more cleanly explained as follows. The language of scheme-valued points (Definition 7.3.10) has the following advantage: notice that the *points* of a group scheme need not themselves form a group (consider $\mathbb{A}_{\mathbb{Z}}^1$). But Exercise 7.6.H shows that the *Z-valued points* of a group scheme (where Z is any given scheme) indeed form a group.

7.6.K. EXERCISE. Suppose we have a contravariant functor F from *Sch* (or indeed any category) to *Groups*. Suppose further that F composed with the forgetful functor *Groups* \rightarrow *Sets* is represented by an object Y . Show that the group operations on $F(X)$ (as X varies through *Sch*) uniquely determine $m: Y \times Y \rightarrow Y$, $i: Y \rightarrow Y$, $e: Z \rightarrow Y$ satisfying the axioms defining a group scheme, such that the group operation on $\text{Mor}(X, Y)$ is the same as that on $F(X)$.

In particular, the definition of a group object in a category was forced upon us by the definition of group. More generally, you should expect that any class of objects that can be interpreted as sets with additional structure should fit into this picture.

You should apply this exercise to $\mathbb{A}_{\mathbb{Z}}^1$, and see how the explicit formulas you found in Exercise 7.6.G are forced on you.

7.6.L. EXERCISE. Work out the maps m , i , and e in the group schemes of Exercise 7.6.E.

7.6.M. EXERCISE. Explain why the product of group objects in a category can be naturally interpreted as a group object in that category.

7.6.N. EXERCISE.

(a) Define **morphism of group schemes**.

(b) Recall that if A is a ring, then $\text{GL}_n(A)$ (the **general linear group over A**) is the group of invertible $n \times n$ matrices with entries in the ring A . Figure out the right definition of **the group scheme GL_n (over a ring A)**, and describe the **determinant**

map $\det: \mathrm{GL}_n \rightarrow \mathbb{G}_m$.

(c) Make sense of the statement: “ $(\cdot^n): \mathbb{G}_m \rightarrow \mathbb{G}_m$ given by $t \mapsto t^n$ is a morphism of group schemes.”

The language of Exercise 7.6.N(a) suggests that group schemes form a category; feel free to prove this if you want. In fact, the category of group schemes has a zero object. What is it?

7.6.O. EXERCISE (KERNELS OF MAPS OF GROUP SCHEMES). Suppose $F: G_1 \rightarrow G_2$ is a morphism of group schemes. Consider the contravariant functor $Sch \rightarrow Groups$ given by $X \mapsto \ker(\mathrm{Mor}(X, G_1) \rightarrow \mathrm{Mor}(X, G_2))$. If this is representable, by a group scheme G_0 , say, show that $G_0 \rightarrow G_1$ is the kernel of F in the category of group schemes.

7.6.P. EXERCISE. Show that the kernel of (\cdot^n) (Exercise 7.6.N) is representable. If $n > 0$, show that over a field k of characteristic p dividing n , this group scheme is nonreduced. (This group scheme, denoted μ_n , is important, although we will not use it.)

7.6.Q. EXERCISE. Show that the kernel of $\det: \mathrm{GL}_n \rightarrow \mathbb{G}_m$ is representable. This is the group scheme SL_n . (You can do this over \mathbb{Z} , or over a field k , or even over an arbitrary ring A ; the algebra is the same.)

7.6.R. EXERCISE. Show (as easily as possible) that \mathbb{A}_k^1 is a ring k -scheme. (Here k can be replaced by any ring.)

7.6.S. EXERCISE.

(a) Define the notion of a (left) **group scheme action** (of a group scheme on a scheme).

(b) Suppose A is a ring. Show that specifying an integer-valued grading on A is equivalent to specifying an action of \mathbb{G}_m on $\mathrm{Spec} A$. (This interpretation of a grading is surprisingly enlightening. Caution: there are two possible choices of the integer-valued grading, and there are reasons for both. Both are used in the literature.)

7.6.6. Aside: Hopf algebras. Here is a notion that we won’t use, but it is easy enough to define now. Suppose $G = \mathrm{Spec} A$ is an affine group scheme, i.e., a group scheme that is an affine scheme. The categorical definition of group scheme can be restated in terms of the ring A . (This requires thinking through Remark 7.3.5; see Exercise 10.1.A.) Then these axioms define a **Hopf algebra**. For example, we have a “comultiplication map” $A \rightarrow A \otimes A$.

7.6.T. EXERCISE. As $\mathbb{A}_{\mathbb{Z}}^1$ is a group scheme, $\mathbb{Z}[t]$ has a Hopf algebra structure. Describe the comultiplication map $\mathbb{Z}[t] \rightarrow \mathbb{Z}[t] \otimes_{\mathbb{Z}} \mathbb{Z}[t]$.

7.7 ★★ The Grassmannian: First construction

The Grassmannian is a useful geometric construction that is “the geometric object underlying linear algebra”. In (classical) geometry over a field $K = \mathbb{R}$ or

\mathbb{C} , just as projective space parametrizes one-dimensional subspaces of a given n -dimensional vector space, the Grassmannian parametrizes k -dimensional subspaces of n -dimensional space. The Grassmannian $G(k, n)$ is a manifold of dimension $k(n - k)$ (over the field). The manifold structure is given as follows. Given a basis (v_1, \dots, v_n) of n -space, “most” k -planes can be described as the span of the k vectors

$$(7.7.0.1) \quad \left\langle v_1 + \sum_{i=k+1}^n a_{1i}v_i, v_2 + \sum_{i=k+1}^n a_{2i}v_i, \dots, v_k + \sum_{i=k+1}^n a_{ki}v_i \right\rangle.$$

(Can you describe which k -planes are *not* of this form? Hint: row reduced echelon form. Aside: the stratification of $G(k, n)$ by normal form is the decomposition of the Grassmannian into *Schubert cells*. You may be able to show using the normal form that each Schubert cell is isomorphic to an affine space.) Any k -plane of this form can be described in such a way uniquely. We use this to identify those k -planes of this form with the manifold $K^{k(n-k)}$ (with coordinates a_{ji}). This is a large affine patch on the Grassmannian (called the “open Schubert cell” with respect to this basis). As the v_i vary, these patches cover the Grassmannian (why?), and the manifold structures agree (a harder fact).

We now *define* the Grassmannian in algebraic geometry, over a ring A . Suppose $v = (v_1, \dots, v_n)$ is a basis for $A^{\oplus n}$. More precisely: $v_i \in A^{\oplus n}$, and the map $A^{\oplus n} \rightarrow A^{\oplus n}$ given by $(a_1, \dots, a_n) \mapsto a_1v_1 + \dots + a_nv_n$ is an isomorphism.

7.7.A. EXERCISE. Show that any two bases are related by an invertible $n \times n$ matrix over A — a matrix with entries in A whose determinant is an invertible element of A .

For each such basis v , we consider the scheme $U_v \cong \mathbb{A}_A^{k(n-k)}$, with coordinates a_{ji} ($k+1 \leq i \leq n$, $1 \leq j \leq k$), which we imagine as corresponding to the k -plane spanned by the vectors (7.7.0.1).

7.7.B. EXERCISE. Given two bases v and w , explain how to glue U_v to U_w along appropriate open sets. You may find it convenient to work with coordinates a_{ji} where i runs from 1 to n , not just $k+1$ to n , but imposing $a_{ji} = \delta_{ji}$ (i.e., 1 when $i = j$ and 0 otherwise) when $i \leq k$. This convention is analogous to coordinates $x_{i/j}$ on the patches of projective space (§4.4.9). Hint: the relevant open subset of U_v will be where a certain determinant doesn’t vanish.

7.7.C. EXERCISE/DEFINITION. By checking triple intersections, verify that these patches (over all possible bases) glue together to a single scheme (Exercise 4.4.A). This is the **Grassmannian** $G(k, n)$ over the ring A . Because it can be interpreted as a space of linear “ \mathbb{P}_A^{k-1} ’s” in \mathbb{P}_A^{n-1} , it is often also written $G(k-1, n-1)$. (You will see that this is wise notation in Exercise 12.2.L, for example.)

Although this definition is pleasantly explicit (it is immediate that the Grassmannian is covered by $\mathbb{A}^{k(n-k)}$ ’s), and perhaps more “natural” than our original definition of projective space in §4.4.9 (we aren’t making a choice of basis; we use *all* bases), there are several things unsatisfactory about this definition of the Grassmannian. In fact the Grassmannian is projective; this isn’t obvious with this definition. Furthermore, the Grassmannian comes with a natural morphism to

$\mathbb{P}^{\binom{n}{k}-1}$ (which is an example of a *closed embedding*, a notion we will meet in Definition 9.1.1), called the **Plücker embedding**. Finally, there is an action of GL_n on the space of k -planes in n -space, so we should be able to see this in our algebraic incarnation. We will address these issues in §17.7, by giving a better description, as a moduli space.

7.7.1. (Partial) flag varieties. Just as the Grassmannian “parametrizes” k -planes in n -space, the **flag variety** parametrizes “flags”: nested sequences of subspaces of n -space

$$F_0 \subset F_1 \subset \cdots \subset F_n$$

where $\dim F_i = i$. Generalizing both of these is the notion of a **partial flag variety** associated to some data $0 \leq a_1 < \cdots < a_\ell \leq n$, which parametrizes nested sequences of subspaces of n -space

$$F_{a_1} \subset \cdots \subset F_{a_\ell}$$

where $\dim F_{a_i} = a_i$. You should be able to generalize all of the discussion in §7.7 to this setting.

Useful classes of morphisms of schemes

We now define an excessive number of types of morphisms. Some (often finiteness properties) are useful because *most* “reasonable” morphism has such properties, and they will be used in proofs in obvious ways. Others correspond to geometrically meaningful behavior, and you should have a picture of what each means.

8.0.1. Change of perspective: morphisms are more fundamental than objects. One of Grothendieck’s lessons is that things that we often think of as properties of *objects* are better understood as properties of *morphisms*. One way of turning properties of objects into properties of morphisms is as follows. If P is a property of schemes, we often (but not always) say that a *morphism* $\pi: X \rightarrow Y$ has P if for every affine open subset $U \subset Y$, $\pi^{-1}(U)$ has P . We will see this for $P =$ quasicompact, quasiseparated, affine, and more. (As you might hope, in good circumstances, P will satisfy the hypotheses of the Affine Communication Lemma 5.3.2, so we don’t have to check *every* affine open subset.) Informally, you can think of such a morphism as one where all the fibers have P . (You can quickly define the fiber of a morphism as a topological space, but once we define fibered product, we will define the *scheme-theoretic* fiber, and then this discussion will make sense.) But it means more than that: it means that “being P ” is really not just fiber-by-fiber, but behaves well as the fiber varies. (For comparison, “submersion of manifolds” means more than that the fibers are smooth.)

8.1 “Reasonable” classes of morphisms (such as open embeddings)

8.1.1. You will notice that almost all classes of morphisms that are useful have many properties in common, which follow from the following three basic properties. We call any class of morphisms satisfying these properties a “**reasonable**” class of morphisms of schemes.

- (i) The class is **preserved by composition**: if $\pi: X \rightarrow Y$ and $\rho: Y \rightarrow Z$ are both in this class, then so is $\rho \circ \pi$.
- (ii) The class is **preserved by “base change”** (or “pullback” or “fibered product”). Precisely: if $\pi: X \rightarrow Y$ is in this class, then for any $Y' \rightarrow Y$, the induced map $X \times_Y Y' \rightarrow Y'$ is also in this class. (Implicit in this statement is that the fibered product $X \times_Y Y'$ exists; but we will soon see that all fibered products of schemes exist, §10.1.)

- (iii) The class is **local on the target**. In other words, (a) if $\pi: X \rightarrow Y$ is in the class, then for any open subset V of Y , the restricted morphism $\pi^{-1}(V) \rightarrow V$ is in the class; and (b) for a morphism $\pi: X \rightarrow Y$, if there is an open cover $\{V_i\}$ of Y for which each restricted morphism $\pi^{-1}(V_i) \rightarrow V_i$ is in the class, then π is in the class. In particular, as schemes are built out of affine schemes, properties are often easy to verify on any affine cover (the properties are “affine-local” on the target), as described in §8.0.1. (Stalk-local properties are automatically local on the target.)

Properties (i) and (ii) imply a useful additional property — (iv) “reasonable” classes of morphisms are **preserved by product**:

8.1.A. EXERCISE. Suppose P is a property of morphisms preserved by composition and base change ((i) and (ii) above), and $X \rightarrow Y$ and $X' \rightarrow Y'$ are two morphisms of S -schemes with property P . Assume that $X \times_S X'$ and $Y \times_S Y'$ exist (which is indeed true, §10.1). Show that $X \times_S X' \rightarrow Y \times_S Y'$ has property P as well. Hint:

$$\begin{array}{ccccc}
 X \times_S X' & \longrightarrow & X \times_S Y' & \longrightarrow & Y \times_S Y' \\
 \swarrow & & \swarrow & & \swarrow \\
 X' & \longrightarrow & Y' & & X \longrightarrow Y
 \end{array}$$

- (v) Another extremely important consequence of (i) and (ii) is the *Cancellation Theorem* for such morphisms, Theorem 11.2.1.

When you learn of a new class of morphisms, you should immediately ask whether these properties (i)–(iii) (and hence (iv) and (v)) hold. The answer will almost always be “yes”; and if it is “no”, then you should exercise caution, and possibly figure out how to make the definition better (see for example §10.5).

To prepare you think about these properties for schemes, you may want to verify that the analogous properties to (i)–(iii) hold in the category of topological spaces with the classes “injections”, “surjections”, “open embeddings” (open embedding of topological spaces = isomorphism with an open subset), and “closed embeddings” (closed embedding of topological spaces = isomorphism with a closed subset).

As a first example in the category of schemes:

8.1.B. EASY EXERCISE. Show that “isomorphisms” of schemes are “reasonable” in this sense, i.e., that they satisfy the three properties (i)–(iii) above.

We are now ready to introduce a fundamental class of morphisms of schemes.

8.1.2. Definition. A morphism $\pi: X \rightarrow Y$ of schemes is an **open embedding** (or **open immersion**) if it is an open embedding as ringed spaces (§7.2.1). In other words, a morphism $\pi: (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ of schemes is an open embedding if π factors as

$$(X, \mathcal{O}_X) \xrightarrow{\sim} (U, \mathcal{O}_Y|_U) \xhookrightarrow{\tau} (Y, \mathcal{O}_Y)$$

where ρ is an isomorphism, and $\tau: U \hookrightarrow Y$ is an inclusion of an open set. It is immediate that isomorphisms are open embeddings. The symbol \hookrightarrow is often used to indicate that a morphism is an open embedding (or more generally, a locally closed embedding, see §9.2).

If X is actually a subset of Y (and π is the inclusion, i.e., ρ is the identity), then we say (X, \mathcal{O}_X) is an **open subscheme** of (Y, \mathcal{O}_Y) . The difference between open embeddings and open subschemes is a bit confusing, and not too important: at the level of sets, open subschemes *are* subsets, while open embeddings are *bijections onto* subsets.

“Open subschemes” are scheme-theoretic analogs of open subsets. (“Closed subschemes” are scheme-theoretic analogs of closed subsets, but they have a surprisingly different flavor, as we will see in §9.1.)

The next two Exercises verify that the class of open embeddings is “reasonable” in the sense described above.

8.1.C. EXERCISE. Verify that the class of open embeddings satisfies properties (i) and (iii) of §8.1.

8.1.D. IMPORTANT BUT EASY EXERCISE: FIBERED PRODUCTS WITH OPEN EMBEDDINGS EXIST. Verify that the class of open embeddings satisfies property (ii) of §8.1. More specifically: suppose $i: U \rightarrow Z$ is an open embedding, and $\rho: Y \rightarrow Z$ is any morphism. Show that $U \times_Z Y$ exists and $U \times_Z Y \rightarrow Y$ is an open embedding. (Hint: I’ll even tell you what $U \times_Z Y$ is: $(\rho^{-1}(U), \mathcal{O}_Y|_{\rho^{-1}(U)})$.) In particular, if $U \hookrightarrow Z$ and $V \hookrightarrow Z$ are open embeddings, $U \times_Z V \cong U \cap V$: “intersection of open embeddings is fiber product”.

8.1.E. EASY EXERCISE. Show that open embeddings of schemes are monomorphisms (in the category of schemes). Hint: Exercise 7.2.B.

8.1.F. EASY EXERCISE. Suppose $\pi: X \rightarrow Y$ is an open embedding. Show that if Y is locally Noetherian, then X is too. Show that if Y is Noetherian, then X is too. However, show that if Y is quasicompact, X need not be. (Hint: let Y be affine but not Noetherian, see Exercise 3.6.G(b).)

8.1.3. Definition. In analogy with “local on the target” (§8.1.1), we define what it means for a property P of morphisms to be **local on the source**: to check if a morphism $\pi: X \rightarrow Y$ has P , it suffices to check on any open cover $\{U_i\}$ of X . We then define **affine-local on the source** (and **affine-local on the target**) in the obvious way: it suffices to check on any affine open cover of the source (resp. target). The use of “affine-local” rather than “local” is to emphasize that the criterion on affine schemes is simple to describe. (Clearly the Affine Communication Lemma 5.3.2 will be very handy.)

8.1.G. EXERCISE (PRACTICE WITH THE CONCEPT). Show that the notion of “open embedding” is not local on the source.

8.2 Algebraic interlude: Lying Over and Nakayama

Algebra is the offer made by the devil to the mathematician. The devil says: I will give you this powerful machine, it will answer any question you like. All you need to do is give me your soul: give up geometry and you will have this marvelous machine.

— M. Atiyah, [At2, p. 659]; but see the Atiyah quote at the start of §0.3

To set up our discussion in the next section on integral morphisms, we develop some algebraic preliminaries. A clever trick we use can also be used to show Nakayama's Lemma, so we discuss this as well.

Suppose $\phi: B \rightarrow A$ is a ring morphism. We say $a \in A$ is **integral** over B if a satisfies some monic polynomial

$$a^n + ?a^{n-1} + \cdots + ? = 0$$

where the coefficients lie in $\phi(B)$. A ring *morphism* $\phi: B \rightarrow A$ is **integral** if every element of A is integral over $\phi(B)$. An integral ring morphism ϕ is an **integral extension** if ϕ is an *inclusion* of rings. You should think of integral morphisms and integral extensions as ring-theoretic generalizations of the notion of algebraic extensions of fields.

8.2.A. EXERCISE. Show that if $\phi: B \rightarrow A$ is a ring morphism, $(b_1, \dots, b_n) = 1$ in B , and $B_{b_i} \rightarrow A_{\phi(b_i)}$ is integral for all i , then ϕ is integral. Hint: replace B by $\phi(B)$ to reduce to the case where B is a subring of A . Suppose $a \in A$. Show that there is some t and m such that $b_i^t a^m \in B + Ba + Ba^2 + \cdots + Ba^{m-1}$ for some t and m independent of i . Use a "partition of unity" argument as in the proof of Theorem 4.1.2 to show that $a^m \in B + Ba + Ba^2 + \cdots + Ba^{m-1}$.

8.2.B. EXERCISE.

- (a) Show that the property of a *morphism* $\phi: B \rightarrow A$ being integral is always preserved by localization and quotient of B , and quotient of A , but not localization of A . More precisely: suppose ϕ is integral. Show that the induced maps $T^{-1}B \rightarrow \phi(T)^{-1}A$, $B/J \rightarrow A/\phi(J)A$, and $B \rightarrow A/I$ are integral (where T is a multiplicative subset of B , J is an ideal of B , and I is an ideal of A), but $B \rightarrow S^{-1}A$ need not be integral (where S is a multiplicative subset of A). (Hint for the latter: show that $k[t] \rightarrow k[t]$ is an integral ring morphism, but $k[t] \rightarrow k[t]_{(t)}$ is not.)
- (b) Show that the property of ϕ being an integral *extension* is preserved by localization of B , but not localization or quotient of A . (Hint for the latter: $k[t] \rightarrow k[t]$ is an integral extension, but $k[t] \rightarrow k[t]/(t)$ is not.)
- (c) In fact the property of ϕ being an integral *extension* (as opposed to integral *morphism*) is not preserved by taking quotients of B either. (Let $B = k[x, y]/(y^2)$ and $A = k[x, y, z]/(z^2, xz - y)$. Then B injects into A , but $B/(x)$ doesn't inject into $A/(x)$.) But it is in some cases. Suppose $\phi: B \rightarrow A$ is an integral extension, and $J \subset B$ is the restriction of an ideal $I \subset A$. (Side Remark: you can show that this holds if J is prime.) Show that the induced map $B/J \rightarrow A/JA$ is an integral extension. (Hint: show that the composition $B/J \rightarrow A/JA \rightarrow A/I$ is an injection.)

The following lemma uses a useful but sneaky trick.

8.2.1. Lemma. — Suppose $\phi: B \rightarrow A$ is a ring morphism. Then $a \in A$ is integral over B if and only if it is contained in a subalgebra of A that is a finitely generated B -module.

Proof. If a satisfies a monic polynomial equation of degree n , then the B -submodule of A generated by $1, a, \dots, a^{n-1}$ is closed under multiplication, and hence a subalgebra of A .

Assume conversely that a is contained in a subalgebra A' of A that is a finitely generated B -module. Choose a finite generating set m_1, \dots, m_n of A' (as a B -module). Then $am_i = \sum b_{ij}m_j$, for some $b_{ij} \in B$. Thus

$$(8.2.1.1) \quad (a\text{Id}_{n \times n} - [b_{ij}]_{ij}) \begin{pmatrix} m_1 \\ \vdots \\ m_n \end{pmatrix} = \begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix},$$

where Id_n is the $n \times n$ identity matrix in A . We can't invert the matrix $(a\text{Id}_{n \times n} - [b_{ij}]_{ij})$, but we almost can. Recall that an $n \times n$ matrix M has an *adjugate matrix* $\text{adj}(M)$ such that $\text{adj}(M)M = \det(M)\text{Id}_n$. (The (i, j) th entry of $\text{adj}(M)$ is the determinant of the matrix obtained from M by deleting the i th column and j th row, times $(-1)^{i+j}$. You have likely seen this in the form of a formula for M^{-1} when there is an inverse; see for example [DF, p. 440].) The coefficients of $\text{adj}(M)$ are polynomials in the coefficients of M . Multiplying (8.2.1.1) by $\text{adj}(a\text{Id}_{n \times n} - [b_{ij}]_{ij})$, we get

$$\det(a\text{Id}_{n \times n} - [b_{ij}]_{ij}) \begin{pmatrix} m_1 \\ \vdots \\ m_n \end{pmatrix} = \begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix}.$$

So $\det(aI - [b_{ij}])$ annihilates the generating elements m_i , and hence every element of A' , i.e., $\det(aI - [b_{ij}]) = 0$. But expanding the determinant yields an integral equation for a with coefficients in B . \square

8.2.2. Corollary (finite implies integral). — *If A is a finite B -algebra (a finitely generated B -module), then ϕ is an integral ring morphism.*

The converse is false: integral does not imply finite, as $\mathbb{Q} \hookrightarrow \overline{\mathbb{Q}}$ is an integral ring morphism, but $\overline{\mathbb{Q}}$ is not a finite \mathbb{Q} -module. (A field extension is integral if it is algebraic.)

8.2.C. EXERCISE. Show that if $C \rightarrow B$ and $B \rightarrow A$ are both integral ring morphisms, then so is their composition.

8.2.D. EXERCISE. Suppose $\phi: B \rightarrow A$ is a ring morphism. Show that the elements of A integral over B form a subalgebra of A .

8.2.3. Remark: Transcendence theory. These ideas lead to the main facts about transcendence theory we will need for a discussion of dimension of varieties, see Exercise/Definition 12.2.A.

8.2.4. The Lying Over and Going-Up Theorems. The Lying Over Theorem is a useful property of integral extensions.

8.2.5. The Lying Over Theorem. — *Suppose $\phi: B \rightarrow A$ is an integral extension. Then for any prime ideal $\mathfrak{q} \subset B$, there is a prime ideal $\mathfrak{p} \subset A$ such that $\mathfrak{p} \cap B = \mathfrak{q}$.*

To be clear on how weak the hypotheses are: B need not be Noetherian, and A need not be finitely generated over B .

8.2.6. Geometric translation: *if $B \rightarrow A$ is an integral extension, then $\text{Spec } A \rightarrow \text{Spec } B$ is surjective.* Although the Lying Over Theorem 8.2.5 is a theorem in algebra, the

name can be interpreted geometrically: the theorem asserts that the corresponding morphism of schemes is **surjective**. (A map of schemes is **surjective** if the underlying map of sets is surjective.) Translation: “above” every prime \mathfrak{q} “downstairs”, there is a prime \mathfrak{p} “upstairs”, see Figure 8.1. (For this reason, it is often said that \mathfrak{p} “lies over” \mathfrak{q} if $\mathfrak{p} \cap B = \mathfrak{q}$.)




FIGURE 8.1. A picture of the Lying Over Theorem 8.2.5: if $\phi: B \rightarrow A$ is an integral extension, then $\text{Spec } A \rightarrow \text{Spec } B$ is surjective

The following exercise sets up the proof.

8.2.E. ★ EXERCISE. Show that the special case where A is a field translates to: if $B \subset A$ is a subring with A integral over B , then B is a field. Prove this. (Hint: you must show that all nonzero elements in B have inverses in B . Here is the start: If $b \in B$, then $1/b \in A$, and this satisfies some integral equation over B .)

★ *Proof of the Lying Over Theorem 8.2.5.* We first make a reduction: by localizing at \mathfrak{q} (preserving integrality by Exercise 8.2.B(b)), we can assume that (B, \mathfrak{q}) is a local ring. Then let \mathfrak{p} be any *maximal* ideal of A . Consider the following diagram.

$$\begin{array}{ccc} A & \twoheadrightarrow & A/\mathfrak{p} & \text{field} \\ \uparrow & & \uparrow & \\ B & \twoheadrightarrow & B/(\mathfrak{p} \cap B) & \end{array}$$

The right vertical arrow is an integral extension by Exercise 8.2.B(c). By Exercise 8.2.E, $B/(\mathfrak{p} \cap B)$ is a field too, so $\mathfrak{p} \cap B$ is a maximal ideal, hence it is \mathfrak{q} . \square

8.2.F. IMPORTANT EXERCISE (THE Going-Up Theorem).

(a) Suppose $\phi: B \rightarrow A$ is an integral ring *morphism* (not necessarily an integral extension). Show that if $\mathfrak{q}_1 \subset \mathfrak{q}_2 \subset \cdots \subset \mathfrak{q}_n$ is a chain of prime ideals of B , and $\mathfrak{p}_1 \subset \cdots \subset \mathfrak{p}_m$ is a chain of prime ideals of A such that \mathfrak{p}_i “lies over” \mathfrak{q}_i (and $1 \leq m < n$), then the second chain can be extended to $\mathfrak{p}_1 \subset \cdots \subset \mathfrak{p}_n$ so that this remains true. (Hint: reduce to the case $m = 1, n = 2$; reduce to the case where $\mathfrak{q}_1 = (0)$ and $\mathfrak{p}_1 = (0)$; use the Lying Over Theorem 8.2.5.)

(b) Draw a picture of this theorem (akin to Figure 8.1).

There are analogous “Going-Down” results (requiring quite different hypotheses); see for example Theorem 12.2.12 and Exercise 25.5.E.

8.2.7. Nakayama’s Lemma.

The trick in the proof of Lemma 8.2.1 can be used to quickly prove Nakayama’s Lemma, which we will use repeatedly in the future. This name is used for several different but related results, which we discuss here. (A geometric interpretation will be given in Exercise 14.3.E.) We may as well prove it while the trick is fresh in our minds.

8.2.8. Nakayama’s Lemma version 1. — *Suppose A is a ring, I is an ideal of A , and M is a finitely generated A -module, such that $M = IM$. Then there exists an $a \in A$ with $a \equiv 1 \pmod{I}$ with $aM = 0$. (Equivalently, there is some $i \in I$ for which multiplication by i induces the identity on M : $im = m$ for all $m \in M$.)*

Proof. Say M is generated by m_1, \dots, m_n . Then as $M = IM$, we have $m_i = \sum_j a_{ij} m_j$ for some $a_{ij} \in I$. Thus

$$(8.2.8.1) \quad (\text{Id}_n - Z) \begin{pmatrix} m_1 \\ \vdots \\ m_n \end{pmatrix} = 0$$

where $Z = (a_{ij})$. Multiplying both sides of (8.2.8.1) on the left by $\text{adj}(\text{Id}_n - Z)$, we obtain

$$\det(\text{Id}_n - Z) \begin{pmatrix} m_1 \\ \vdots \\ m_n \end{pmatrix} = 0.$$

But when you expand out $\det(\text{Id}_n - Z)$, as Z has entries in I , you get something that is $1 \pmod{I}$. \square

Here is why you care. Suppose I is contained in all maximal ideals of A . (The intersection of all the maximal ideals is called the **Jacobson radical**, but we won’t use this phrase. For comparison, recall that the nilradical was the intersection of the *prime ideals* of A .) Then any $a \equiv 1 \pmod{I}$ is invertible. (We are not using Nakayama yet!) Reason: otherwise $(a) \neq A$, so the ideal (a) is contained in some maximal ideal \mathfrak{m} — but $a \equiv 1 \pmod{\mathfrak{m}}$, contradiction. As a is invertible, we have the following.

8.2.9. Nakayama’s Lemma version 2. — *Suppose A is a ring, I is an ideal of A contained in all maximal ideals, and M is a finitely generated A -module. (The most interesting case is when A is a local ring, and I is the maximal ideal.) Suppose $M = IM$. Then $M = 0$.*

8.2.G. EXERCISE (NAKAYAMA’S LEMMA VERSION 3). Suppose A is a ring, and I is an ideal of A contained in all maximal ideals. Suppose M is a finitely generated A -module, and $N \subset M$ is a submodule. If $N/IN \rightarrow M/IM$ is surjective, then $M = N$.

8.2.H. IMPORTANT EXERCISE (NAKAYAMA’S LEMMA VERSION 4: GENERATORS OF $M/\mathfrak{m}M$ LIFT TO GENERATORS OF M). Suppose (A, \mathfrak{m}) is a local ring. Suppose M is a finitely generated A -module, and $f_1, \dots, f_n \in M$, with (the images of) f_1, \dots, f_n

generating M/mM . Then f_1, \dots, f_n generate M . (In particular, taking $M = m$, if we have generators of m/m^2 , they also generate m .)

8.2.I. IMPORTANT EXERCISE GENERALIZING LEMMA 8.2.1. Recall that a B -module N is said to be **faithful** if the only element of B acting on N by the identity is 1 (or equivalently, if the only element of B acting as the 0-map on N is 0). Suppose S is a subring of a ring A , and $r \in A$. Suppose there is a faithful $S[r]$ -module M that is finitely generated as an S -module. Show that r is integral over S . (Hint: change a few words in the proof of version 1 of Nakayama, Lemma 8.2.8.)

8.2.J. EXERCISE. Suppose A is an integral domain, and \tilde{A} is the **integral closure** of A in $K(A)$, i.e., those elements of $K(A)$ integral over A , which form a subalgebra by Exercise 8.2.D. Show that \tilde{A} is integrally closed in $K(\tilde{A}) = K(A)$.

8.3 A gazillion finiteness conditions on morphisms

By the end of this section, you will have seen the following types of morphisms: quasicompact, quasiseparated, affine, finite, integral, closed, (locally) of finite type, quasifinite — and possibly, (locally) of finite presentation.

8.3.1. Quasicompact and quasiseparated morphisms.

A morphism $\pi: X \rightarrow Y$ of schemes is **quasicompact** if for every open affine subset U of Y , $\pi^{-1}(U)$ is quasicompact. (Equivalently, the preimage of any quasicompact open subset is quasicompact. This is the right definition in other parts of geometry.)

We will like this notion because (i) finite sets have advantages over infinite sets (e.g., a finite set of integers has a maximum; also, things can be proved inductively), and (ii) most reasonable schemes will be quasicompact.

Along with quasicompactness comes the weird notion of quasiseparatedness. A morphism $\pi: X \rightarrow Y$ is **quasiseparated** if for every open affine subset U of Y , $\pi^{-1}(U)$ is a quasiseparated scheme (§5.1.1). (Equivalently, the preimage of any quasicompact open subset is quasiseparated, although we won't worry about proving this. This is the definition that extends to other parts of geometry.) This will be a useful hypothesis in theorems, usually in conjunction with quasicompactness. (For this reason, “quasicompact and quasiseparated” is often abbreviated as **qcqs**, cf. §5.1.2.) Various interesting kinds of morphisms (locally Noetherian source, affine, separated, see Exercises 8.3.B(b), 8.3.D, and 11.3.D resp.) are quasiseparated, and having the word “quasiseparated” will allow us to state theorems more succinctly. *Important remark:* we will give an equivalent definition of quasiseparatedness in §11.2.4 (“quasiseparated = quasicompact diagonal”), which will be much simpler to use.

8.3.A. EASY EXERCISE. Show that the composition of two quasicompact morphisms is quasicompact. (It is also true — but not easy — that the composition of two quasiseparated morphisms is quasiseparated. You are free to show this directly, but will in any case follow easily once we understand it in a more sophisticated way, see §11.2.5.)

Following Grothendieck's philosophy of thinking that the important notions are properties of morphisms, not of objects (§8.0.1), we can restate the definition of a quasicompact (resp. quasiseparated) scheme as a scheme that is quasicompact (resp. quasiseparated) over the final object $\text{Spec } \mathbb{Z}$ in the category of schemes (Exercise 7.3.J).

8.3.B. EASY EXERCISE.

- (a) Show that any morphism from a Noetherian scheme is quasicompact.
- (b) Show that any morphism from a quasiseparated scheme is quasiseparated. Thus by Exercise 5.3.A, any morphism from a locally Noetherian scheme is quasiseparated. Thus readers working only with locally Noetherian schemes may take quasiseparatedness as a standing hypothesis.

8.3.2. Caution. The two parts of the Exercise 8.3.B may lead you to suspect that any morphism $\pi: X \rightarrow Y$ with quasicompact source and target is necessarily quasicompact. This is false, and you may verify that the following is a counterexample. Let Z be the nonquasiseparated scheme constructed in Exercise 5.1.K, and let $X = \text{Spec } k[x_1, x_2, \dots]$ as in Exercise 5.1.K. The obvious open embedding $\pi: X \rightarrow Z$ (identifying X with one of the two pieces glued together to get Z) is not quasicompact. (But once you see the Cancellation Theorem 11.2.1, you will quickly see that any morphism from a quasicompact source to a *quasiseparated* target is necessarily quasicompact.)

8.3.C. EXERCISE. (Obvious hint for both parts: the Affine Communication Lemma 5.3.2.)

- (a) (*quasicompactness is affine-local on the target*) Show that a morphism $\pi: X \rightarrow Y$ is quasicompact if there is a cover of Y by affine open sets U_i such that $\pi^{-1}(U_i)$ is quasicompact.
- (b) (*quasiseparatedness is affine-local on the target*) Show that a morphism $\pi: X \rightarrow Y$ is quasiseparated if there is a cover of Y by affine open sets U_i such that $\pi^{-1}(U_i)$ is quasiseparated.

8.3.3. Affine morphisms.

A morphism $\pi: X \rightarrow Y$ is **affine** if for every affine open set U of Y , $\pi^{-1}(U)$ (interpreted as an open subscheme of X) is an affine scheme. Trivially, the composition of two affine morphisms is affine.

8.3.D. FAST EXERCISE. Show that affine morphisms are quasicompact and quasiseparated. (Hint for the second: Exercise 5.1.G.)

8.3.4. Proposition (the property of “affineness” of a morphism is affine-local on the target). — *A morphism $\pi: X \rightarrow Y$ is affine if there is a cover of Y by affine open sets U such that $\pi^{-1}(U)$ is affine.*

Proof. We apply the Affine Communication Lemma 5.3.2 to the condition “ π is affine over”. We check our two criteria. First, suppose $\pi: X \rightarrow Y$ is affine over $\text{Spec } B$, i.e., $\pi^{-1}(\text{Spec } B) = \text{Spec } A$. Then for any $s \in B$, $\pi^{-1}(\text{Spec } B_s) = \text{Spec } A_{\pi^\# s}$. (Do you see why?)

Second, suppose we are given $\pi: X \rightarrow \text{Spec } B$ and $(s_1, \dots, s_n) = B$ with $X_{\pi^\# s_i}$ affine ($\text{Spec } A_i$, say). (Recall from Definition 6.2.7 that $X_{\pi^\# s_i}$ is the open subset of X where $\pi^\# s_i$ doesn't vanish.) We wish to show that X is affine too. Let $A = \Gamma(X, \mathcal{O}_X)$.

Then $X \rightarrow \operatorname{Spec} B$ factors through the tautological map $\alpha: X \rightarrow \operatorname{Spec} A$ (arising from the (iso)morphism $A \rightarrow \Gamma(X, \mathcal{O}_X)$, Exercise 7.3.G).

$$\begin{array}{ccc} \bigcup_i X_{\pi^{-1}(s_i)} = X & \xrightarrow{\alpha} & \operatorname{Spec} A \\ & \searrow \pi \quad \swarrow \beta & \\ & \bigcup_i D(s_i) = \operatorname{Spec} B & \end{array}$$

We want to show that α is an isomorphism. Now $\beta^{-1}(D(s_i)) = D(\beta^\# s_i) \cong \operatorname{Spec} A_{\beta^\# s_i}$ (the preimage of a distinguished open set is a distinguished open set), and $\pi^{-1}(D(s_i)) = \operatorname{Spec} A_i$. Furthermore, X is quasicompact and quasiseparated by the affine-locality of these notions (Exercise 8.3.C), so the hypotheses of the Qcqs Lemma 6.2.8 are satisfied. Hence by the Qcqs Lemma 6.2.8, we have an induced isomorphism of $A_{\beta^\# s_i} = \Gamma(X, \mathcal{O}_X)_{\beta^\# s_i} \cong \Gamma(X_{\beta^\# s_i}, \mathcal{O}_X) = A_i$. Thus α induces an isomorphism $\operatorname{Spec} A_i \xrightarrow{\sim} \operatorname{Spec} A_{\beta^\# s_i}$ (an isomorphism of rings induces an isomorphism of affine schemes, Exercise 4.3.A). Thus α is an isomorphism over each $\operatorname{Spec} A_{\beta^\# s_i}$, which cover $\operatorname{Spec} A$, and hence α is an isomorphism. Therefore $X \rightarrow \operatorname{Spec} A$ is an isomorphism, so X is affine as desired. \square

The affine-locality of affine morphisms (Proposition 8.3.4) has some nonobvious consequences, as shown in the next exercise.

8.3.E. USEFUL EXERCISE. Suppose Z is a closed subset of an affine scheme $\operatorname{Spec} A$ locally cut out by one equation. (In other words, $\operatorname{Spec} A$ can be covered by smaller open sets, and on each such set Z is cut out by one equation.) Show that the complement Y of Z is affine. (This is clear if Z is globally cut out by one equation f , even set-theoretically; then $Y = \operatorname{Spec} A_f$. However, Z is not always of this form, see §20.11.10.)

8.3.5. Finite and integral morphisms.

Before defining finite and integral morphisms, we give an example to keep in mind. If L/K is a field extension, then $\operatorname{Spec} L \rightarrow \operatorname{Spec} K$ (i) is always affine; (ii) is integral if L/K is algebraic; and (iii) is finite if L/K is finite.

Recall that if we have a ring morphism $B \rightarrow A$ such that A is a finitely generated B -module then we say that A is a **finite** B -algebra. This is stronger than being a finitely generated B -algebra. (The similarity of the terminology “finite” and “finitely-generated” B -algebra is unfortunate.) We say a morphism $\pi: X \rightarrow Y$ is **finite** if for every affine open set $\operatorname{Spec} B$ of Y , $\pi^{-1}(\operatorname{Spec} B)$ is the spectrum of a *finite* B -algebra. By definition, finite morphisms are affine.

8.3.F. EXERCISE (THE PROPERTY OF FINITENESS IS AFFINE-LOCAL ON THE TARGET). Show that a morphism $\pi: X \rightarrow Y$ is finite if there is a cover of Y by affine open sets $\operatorname{Spec} A$ such that $\pi^{-1}(\operatorname{Spec} A)$ is the spectrum of a finite A -algebra.

The following four examples will give you some feeling for finite morphisms. In each example, you will notice two things. In each case, the maps are always finite-to-one (as maps of sets). We will verify this in general in Exercise 8.3.J. You will also notice that the morphisms are **closed** as maps of topological spaces, i.e., the images of closed sets are closed. We will show that finite morphisms are always closed in Exercise 8.3.L (and give a second proof in §9.3.5). Intuitively, you should

think of finite as being closed plus finite fibers, although this isn't quite true. We will make this precise in Theorem 30.6.2.

Example 1: Branched covers. Consider the morphism $\text{Spec } k[t] \rightarrow \text{Spec } k[u]$ given by $u \mapsto p(t)$, where $p(t) \in k[t]$ is a degree n polynomial (see Figure 8.2). This is finite: $k[t]$ is generated as a $k[u]$ -module by $1, t, t^2, \dots, t^{n-1}$.




FIGURE 8.2. The “branched cover” $\mathbb{A}_k^1 \rightarrow \mathbb{A}_k^1$ of the “u-line” by the “t-line” given by $u \mapsto p(t)$ is finite

Example 2: Closed embeddings (to be defined soon, in §9.1.1). If I is an ideal of a ring A , consider the morphism $\text{Spec } A/I \rightarrow \text{Spec } A$ given by the obvious map $A \rightarrow A/I$ (see Figure 8.3 for an example, with $A = k[t]$, $I = (t)$). This is a finite morphism (A/I is generated as a A -module by the element $1 \in A/I$).




FIGURE 8.3. The “closed embedding” $\text{Spec } k \rightarrow \text{Spec } k[t]$ given by $t \mapsto 0$ is finite

Example 3: Normalization (to be defined in §10.7). Consider the morphism $\text{Spec } k[t] \rightarrow \text{Spec } k[x, y]/(y^2 - x^2 - x^3)$ corresponding to $k[x, y]/(y^2 - x^2 - x^3) \rightarrow k[t]$ given by $x \mapsto t^2 - 1$, $y \mapsto t^3 - t$ (check that this is a well-defined ring map!), see Figure 8.4. This is a finite morphism, as $k[t]$ is generated as a $(k[x, y]/(y^2 - x^2 - x^3))$ -module by 1 and t . (The figure suggests that this is an isomorphism away from the “node” of the target. You can verify this, by checking that it induces an isomorphism between $D(t^2 - 1)$ in the source and $D(x)$ in the target. We will meet this example again repeatedly!)

8.3.G. IMPORTANT EXERCISE (EXAMPLE 4, FINITE MORPHISMS TO $\text{Spec } k$). Show that if $X \rightarrow \text{Spec } k$ is a finite morphism, then X is a finite union of points with the discrete topology, each point with residue field a finite extension of k , see




FIGURE 8.4. The “normalization” of the nodal cubic $\text{Spec } k[t] \rightarrow \text{Spec } k[x, y]/(y^2 - x^2 - x^3)$ given by $(x, y) \mapsto (t^2 - 1, t^3 - t)$ is finite

Figure 8.5. (An example is $\text{Spec}(\mathbb{F}_8 \times \mathbb{F}_4[x, y]/(x^2, y^4) \times \mathbb{F}_4[t]/(t^9) \times \mathbb{F}_2) \rightarrow \text{Spec } \mathbb{F}_2$.) Do *not* just quote some fancy theorem! Possible approach: By Exercise 3.2.G, any integral domain which is a finite k -algebra must be a field. If $X = \text{Spec } A$, show that every prime \mathfrak{p} of A is maximal. Show that the irreducible components of $\text{Spec } A$ are closed points. Show $\text{Spec } A$ is discrete and hence finite. Show that the residue fields $K(A/\mathfrak{p})$ of A are finite extensions of k . (See Exercise 8.4.D for an extension to quasifinite morphisms.)



FIGURE 8.5. A picture of a finite morphism to $\text{Spec } k$. Bigger fields are depicted as bigger points.

8.3.H. EASY EXERCISE (CF. EXERCISE 8.2.C). Show that the composition of two finite morphisms is also finite.

8.3.I. EXERCISE (“FINITE MORPHISMS TO $\text{Spec } A$ ARE PROJECTIVE”). If R is an A -algebra, define a graded ring S_\bullet by $S_0 = A$, and $S_n = R$ for $n > 0$. (What is the multiplicative structure? Hint: you know how to multiply elements of R together, and how to multiply elements of A with elements of R .) Describe an isomorphism $\text{Proj } S_\bullet \xrightarrow{\sim} \text{Spec } R$. Show that if R is a *finite* A -algebra (finitely generated as an A -module) then S_\bullet is a finitely generated graded ring over A , and hence that $\text{Spec } R$ is a projective A -scheme (§4.5.9).

Hence finite morphisms are affine (by definition) and projective. The converse is also true, see Corollary 19.1.6.

8.3.J. IMPORTANT EXERCISE. Show that finite morphisms have finite fibers. (This is a useful exercise, because you will have to figure out how to get at points in a fiber of a morphism: given $\pi: X \rightarrow Y$, and $q \in Y$, what are the points of $\pi^{-1}(q)$? This will be easier to do once we discuss fibers in greater detail, see Remark 10.3.5, but it will be enlightening to do it now.) Hint: if $X = \text{Spec } A$ and $Y = \text{Spec } B$ are both affine, and $q = [q]$, then we can throw out everything in B outside \bar{q} by modding out by q ; show that the preimage is $\text{Spec}(A/\pi^\# qA)$. Then you have reduced to the case where Y is the Spec of an integral domain B , and $[q] = [(0)]$ is the generic point. We can throw out the rest of the points of B by localizing at (0) . Show that the preimage is Spec of A localized at $\pi^\#(B \setminus \{0\})$. Show that the condition of finiteness is preserved by the constructions you have done, and thus reduce the problem to Exercise 8.3.G.

The previous two exercises show that finite morphisms are projective and have finite fibers. The converse is also true, see Theorem 19.1.7.

There is more to finiteness than finite fibers, and three examples to keep in mind are Example 8.3.6, Exercise 8.3.K, and Figure 20.1 (a variant of Figure 8.4 showing a morphism that is affine, closed, one-to-one on points and even a monomorphism, but not finite).

8.3.6. Example. The open embedding $\mathbb{A}_k^2 - \{(0, 0)\} \rightarrow \mathbb{A}_k^2$ has finite fibers, but is not affine (as $\mathbb{A}_k^2 - \{(0, 0)\}$ isn't affine, §4.4.1) and hence not finite.

8.3.K. EASY EXERCISE. Show that the open embedding $\mathbb{A}_\mathbb{C}^1 - \{0\} \rightarrow \mathbb{A}_\mathbb{C}^1$ has finite fibers and is affine, but is not finite.

8.3.7. Definition. A morphism $\pi: X \rightarrow Y$ of schemes is **integral** if π is affine, and for every affine open subset $\text{Spec } B \subset Y$, with $\pi^{-1}(\text{Spec } B) = \text{Spec } A$, the induced map $B \rightarrow A$ is an integral ring morphism. This is an affine-local condition by Exercises 8.2.A and 8.2.B, and the Affine Communication Lemma 5.3.2. It is closed under composition by Exercise 8.2.C. Integral morphisms are mostly useful because finite morphisms are integral by Corollary 8.2.2. Note that the converse implication doesn't hold (witness $\text{Spec } \bar{\mathbb{Q}} \rightarrow \text{Spec } \mathbb{Q}$, as discussed after the statement of Corollary 8.2.2).

8.3.L. EXERCISE. Prove that integral morphisms are closed, i.e., that the image of closed subsets are closed. (Hence finite morphisms are closed. A second proof will be given in §9.3.5.) Hint: Reduce to the affine case. If $\pi^\#: B \rightarrow A$ is a ring map, inducing the integral morphism $\pi: \text{Spec } A \rightarrow \text{Spec } B$, then suppose $I \subset A$ cuts out a closed set of $\text{Spec } A$, and $J = (\pi^\#)^{-1}(I)$, then note that $B/J \subset A/I$, and apply the Lying Over Theorem 8.2.5 here.

8.3.M. UNIMPORTANT EXERCISE. Suppose $B \rightarrow A$ is integral. Show that for any ring morphism $B \rightarrow C$, the induced map $C \rightarrow A \otimes_B C$ is integral. (Hint: We wish to show that any $\sum_{i=1}^n a_i \otimes c_i \in A \otimes_B C$ is integral over C . Use the fact that each of the finitely many a_i are integral over B , and then Exercise 8.2.D.) Once we

know what “base change” is, this will imply that the property of integrality of a morphism is preserved by base change, Exercise 10.4.C(e).

8.3.8. Fibers of integral morphisms. Unlike finite morphisms (Exercise 8.3.J), integral morphisms don’t always have finite fibers. (Can you think of an example?) However, once we make sense of fibers as topological spaces (or even schemes) in §10.3.2, you can check (Exercise 12.1.D) that the fibers have the property that no point is in the closure of any other point.

8.3.9. Morphisms (locally) of finite type.

A morphism $\pi: X \rightarrow Y$ is **locally of finite type** if for every affine open set $\text{Spec } B$ of Y , and every affine open subset $\text{Spec } A$ of $\pi^{-1}(\text{Spec } B)$, the induced morphism $B \rightarrow A$ expresses A as a finitely generated B -algebra. By the affine-locality of finite-typeness of B -schemes (Proposition 5.3.3(b)), this is equivalent to: $\pi^{-1}(\text{Spec } B)$ can be covered by affine open subsets $\text{Spec } A_i$ so that each A_i is a finitely generated B -algebra.

A morphism π is **of finite type** if it is locally of finite type and quasicompact. Translation: for every affine open set $\text{Spec } B$ of Y , $\pi^{-1}(\text{Spec } B)$ can be covered with *a finite number of* open sets $\text{Spec } A_i$ so that the induced morphism $B \rightarrow A_i$ expresses A_i as a finitely generated B -algebra.

8.3.10. Linguistic curiosity. It is a common practice to name properties as follows: P equals locally P plus quasicompact. Two exceptions are “ringed space” (§7.3) and “finite presentation” (§8.3.13).

8.3.N. EXERCISE (THE NOTIONS “LOCALLY FINITE TYPE” AND “FINITE TYPE” ARE AFFINE-LOCAL ON THE TARGET). Show that a morphism $\pi: X \rightarrow Y$ is locally of finite type if there is a cover of Y by affine open sets $\text{Spec } B_i$ such that $\pi^{-1}(\text{Spec } B_i)$ is locally finite type over B_i .

Example: the “structure morphism” $\mathbb{P}_A^n \rightarrow \text{Spec } A$ is of finite type, as \mathbb{P}_A^n is covered by $n + 1$ open sets of the form $\text{Spec } A[x_1, \dots, x_n]$.

Our earlier definition of schemes of “finite type over k ” (or “finite type k -schemes”) from §5.3.6 is now a special case of this more general notion: the phrase “a scheme X is of finite type over k ” means that we are given a morphism $X \rightarrow \text{Spec } k$ (the “structure morphism”) that is of finite type.

Here are some properties enjoyed by morphisms of finite type.

8.3.O. EXERCISE (FINITE = INTEGRAL + FINITE TYPE).

- (a) (easier) Show that finite morphisms are of finite type.
- (b) Show that a morphism is finite if and only if it is integral and of finite type.

8.3.P. EXERCISES (NOT HARD, BUT IMPORTANT).

- (a) Show that every open embedding is locally of finite type, and hence that every quasicompact open embedding is of finite type. Show that every open embedding into a locally Noetherian scheme is of finite type.
- (b) Show that the composition of two morphisms locally of finite type is locally of finite type. (Hence as the composition of two quasicompact morphisms is quasicompact, Easy Exercise 8.3.A, the composition of two morphisms of finite type is of finite type.)

(c) Suppose $\pi: X \rightarrow Y$ is locally of finite type, and Y is locally Noetherian. Show that X is also locally Noetherian. If $\pi: X \rightarrow Y$ is a morphism of finite type, and Y is Noetherian, show that X is Noetherian.

8.3.Q. EXERCISE. Suppose X is a scheme over \mathbb{F}_p . Explain how to define (without choice) an endomorphism $F: X \rightarrow X$ such that for each affine open subset $\text{Spec } A \subset X$, F corresponds to the map $A \rightarrow A$ given by $f \mapsto f^p$ for all $f \in A$. (The morphism F is called the **absolute Frobenius morphism**.)

8.3.11. Less important definition. A morphism $\pi: X \rightarrow Y$ is **quasifinite** if it is of finite type, and for all $q \in Y$, $\pi^{-1}(q)$ is a finite set. The main point of this definition is the “finite fiber” part; the “finite type” hypothesis will ensure that this notion is “preserved by fibered product,” Exercise 10.4.E.

Combining Exercise 8.3.J with Exercise 8.3.O(a), we see that finite morphisms are quasifinite. There are quasifinite morphisms which are not finite, such as $\mathbb{A}_k^2 - \{(0, 0)\} \rightarrow \mathbb{A}_k^2$ (Example 8.3.6). However, we will soon see that quasifinite morphisms to $\text{Spec } k$ are finite (Exercise 8.4.D). A key example of a morphism with finite fibers that is not quasifinite is $\text{Spec } \mathbb{C}(t) \rightarrow \text{Spec } \mathbb{C}$. Another is $\text{Spec } \overline{\mathbb{Q}} \rightarrow \text{Spec } \mathbb{Q}$. (For interesting behavior caused by the fact that $\text{Spec } \overline{\mathbb{Q}} \rightarrow \text{Spec } \mathbb{Q}$ is not of finite type, see Warning 10.1.5.)

8.3.12. How to picture quasifinite morphisms. If $\pi: X \rightarrow Y$ is a finite morphism of locally Noetherian schemes, then for any quasicompact open subset $U \subset X$, the induced morphism $\pi|_U: U \rightarrow Y$ is quasifinite. In fact *every* reasonable quasifinite morphism arises in this way. (This simple-sounding statement is in fact a deep and important result — a form of Zariski’s Main Theorem, see Exercise 30.6.E.) Thus the right way to visualize quasifiniteness is as a finite map with some (closed locus of) points removed.

8.3.13. $\star\star$ (Locally) finitely presented morphisms. The following variant of “locally of finite type” is useful in non-Noetherian situations.

8.3.14. Definition. A ring A is a **finitely presented** B -algebra (or $B \rightarrow A$ is **finitely presented**) if “ A has a finite number of generators and a finite number of relations over B ”:

$$A \cong B[x_1, \dots, x_n]/(f_1(x_1, \dots, x_n), \dots, f_r(x_1, \dots, x_n)).$$

If B is Noetherian, then finitely presented is the same as finite type (“finite number of relations” comes for free, Exercise 3.6.S), so normal people will not care.

8.3.R. EXERCISE: FINITELY PRESENTED IMPLIES ALWAYS FINITELY PRESENTED (ALGEBRA VERSION). Suppose $A = B[x_1, \dots, x_n]/(f_1, \dots, f_r)$, where $f_i \in B[x_1, \dots, x_n]$. Suppose also that Y_1, \dots, Y_m generate A as a B -algebra, so we have a surjection $\phi: B[y_1, \dots, y_m] \twoheadrightarrow A$ where $y_j \mapsto Y_j$. Show that the ideal $\ker(\phi) \subset B[y_1, \dots, y_m]$ is finitely generated. Hint: follow the proof of Lemma 6.4.4.

8.3.15. Lemma (finite presentation of A over B is an affine-local property in A). — The condition of a B -algebra being finitely presented satisfies the hypotheses of the Affine Communication Lemma 5.3.2.

Proof. (i) The first hypothesis of the Affine Communication Lemma is the content of the following Exercise.

8.3.S. EXERCISE. Suppose $A = B[x_1, \dots, x_n]/(f_1, \dots, f_r)$, and $g \in A$. Show that A_g is a finitely presented B -algebra. (Hint: it has one more generator and one more relation.)

(ii) (This argument is the hardest part of our discussion on finite presentation.) We suppose that A is a B -algebra, and $g_1, \dots, g_m \in A$ with $(g_1, \dots, g_m) = A$, and A_{g_i} finitely presented over B . We wish to show that A is finitely presented over B . Now the A_{g_i} are finitely generated over B , so by Proposition 5.3.3(b), A is finitely generated over B , say $\phi: B[x_1, \dots, x_n] \twoheadrightarrow A$. We will show that $I := \ker(\phi)$ is a finitely generated ideal of $B[x_1, \dots, x_n]$.

As $(g_1, \dots, g_m) = A$, we may choose $h_j \in A$ so that $1 = h_1 g_1 + \dots + h_m g_m$. Choose lifts H_j and G_j of the h_j and g_j respectively to $B[x_1, \dots, x_n]$. Define $G_0 = 1 - H_1 G_1 - \dots - H_m G_m$, so $G_0 \in I$, and $(G_0, G_1, \dots, G_m) = 1$ in $B[x_1, \dots, x_n]$.

8.3.T. EXERCISE. Show that I_{G_j} is a finitely generated module over $B[x_1, \dots, x_n]_{G_j} = B[x_1, \dots, x_n, y_j]/(y_j G_j - 1)$. (The case $j = 0$ is different but easy — $I_{G_0} = (1)$!) Hint: Lemma 8.3.R.

Then by Exercise 6.4.A (“finite generation of modules is an affine-local condition”), I is a finitely generated $B[x_1, \dots, x_n]$ -module, i.e., a finitely generated ideal of $B[x_1, \dots, x_n]$, as desired. \square

8.3.16. Definition. We say a B -scheme $X \rightarrow \operatorname{Spec} B$ is **locally finitely presented over B** if there exists an affine cover $X = \cup \operatorname{Spec} A_i$ such that A_i is finitely presented over B . This is equivalent (by Lemma 8.3.15) to A being finitely presented over B for *all* affine open subsets $\operatorname{Spec} A$ of X .

8.3.U. EXERCISE (THE NOTION OF “LOCALLY FINITELY PRESENTED MORPHISM” IS AFFINE-LOCAL IN THE TARGET. Suppose $\pi: X \rightarrow Y$ is a morphism of schemes. Show as $\operatorname{Spec} B \subset Y$ runs through the affine open subsets of Y , the condition that $\pi: \pi^{-1}(\operatorname{Spec} B) \rightarrow \operatorname{Spec} B$ is locally finitely presented over $\operatorname{Spec} B$ satisfies the hypotheses of the Affine Communication Lemma. (This is notably easier than Lemma 8.3.15.)

8.3.17. Definition. A morphism $\pi: X \rightarrow Y$ is **locally finitely presented** (or **locally of finite presentation**) if for every affine open set $\operatorname{Spec} B$ of Y , $\pi^{-1}(\operatorname{Spec} B) \rightarrow \operatorname{Spec} B$ is locally finitely presented over $\operatorname{Spec} B$. By Exercise 8.3.U, it suffices to check a single affine cover of Y . (Hence this subsumes Definition 8.3.16.)

8.3.V. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a morphism of schemes.

(a) Show that if π is locally of finite presentation, then for *any* open subscheme $\operatorname{Spec} B$ of Y and *any* affine open subscheme $\operatorname{Spec} A$ of X with $\pi(\operatorname{Spec} A) \subset \operatorname{Spec} B$, A is a finitely presented B -algebra.

(b) Suppose $\cup \operatorname{Spec} A_i = X$ is an affine open cover of X , and $\operatorname{Spec} B_i$ are affine open subsets of Y with $\pi(\operatorname{Spec} A_i) \subset \operatorname{Spec} B_i$. Suppose A_i is finitely presented over B_i . Show that π is locally finitely presented.

8.3.W. EXERCISE. Show that open embeddings are locally finitely presented.

8.3.X. EXERCISE. Show that the composition of two locally finitely presented morphisms is locally finitely presented.

8.3.18. ★★ Remark. A morphism $\pi: X \rightarrow Y$ is locally finitely presented if and only if for every projective system of Y -schemes $\{S_\lambda\}_{\lambda \in I}$ with each S_λ an affine scheme, the natural map

$$\operatorname{colim}_\lambda \operatorname{Hom}_Y(S_\lambda, X) \rightarrow \operatorname{Hom}_Y(\lim_\lambda S_\lambda, X)$$

is a bijection (see [Gr-EGA, IV₃.8.14.2]). This characterization of locally finitely presented morphisms as “limit-preserving” can be useful.

8.3.19. Definition. A morphism is **finitely presented** (or **of finite presentation**) if it is locally finitely presented, quasicompact, and quasiseparated.

This definition violates the principle that erasing “locally” is the same as adding “quasicompact and” (Remark 8.3.10). But it is well motivated: finite presentation means “finite in all possible ways” (the algebra corresponding to each affine open set has a finite number of generators, and a finite number of relations, and a finite number of such affine open sets cover, and their intersections are also covered by a finite number affine open sets) — it is all you would hope for in a scheme without it actually being Noetherian. Exercise 10.3.H makes this precise, and explains how this notion often arises in practice.

If Y is locally Noetherian, then locally of finite presentation is the same as locally of finite type, and finite presentation is the same as finite type. So if you are a Noetherian person, you don’t need to worry about this notion.

8.4 Images of morphisms: Chevalley’s Theorem and elimination theory

In this section, we will answer a question that you may have wondered about long before hearing the phrase “algebraic geometry”. If you have a number of polynomial equations in a number of variables with indeterminate coefficients, you would reasonably ask what conditions there are on the coefficients for a (common) solution to exist. Given the algebraic nature of the problem, you might hope that the answer should be purely algebraic in nature — it shouldn’t be “random”, or involve bizarre functions like exponentials or cosines. You should expect the answer to be given by “algebraic conditions”. This is indeed the case, and it can be profitably interpreted as a question about images of maps of varieties or schemes, in which guise it is answered by Chevalley’s Theorem 8.4.2 (see §8.4.5 for a more precise proof). As a consequence we get an immediate proof of the Nullstellensatz 3.2.6 (§8.4.3).

In certain cases, the image is nicer still. For example, the image of a finite morphism is always closed (Exercise 8.3.L). We will prove a classical result, the Fundamental Theorem of Elimination Theory 8.4.7, which essentially generalizes this (as explained in §9.3.5) to maps from projective space.

In a different direction, in the distant future we will see that in certain good circumstances (“flat” plus a bit more, see Exercise 25.5.G), morphisms are open

(the image of open subsets is open); one example (which you can try to show directly) is $\mathbb{A}_B^n \rightarrow \operatorname{Spec} B$.

8.4.1. Chevalley's Theorem.

If $\pi: X \rightarrow Y$ is a morphism of schemes, the notion of the image of π as *sets* is clear: we just take the points in Y that are the image of points in X . We know that the image can be open (open embeddings), and we have seen examples where it is closed, and more generally, locally closed. But it can be weirder still: consider the morphism $\mathbb{A}_k^2 \rightarrow \mathbb{A}_k^2$ given by $(x, y) \mapsto (x, xy)$. The image is the plane, with the y -axis removed, but the origin put back in (Figure 8.6). We make a definition to capture this phenomenon. A **constructible subset** of a Noetherian topological space is a subset which belongs to the smallest family of subsets such that (i) every open set is in the family, (ii) a finite intersection of family members is in the family, and (iii) the complement of a family member is also in the family. For example, the image of $(x, y) \mapsto (x, xy)$ is constructible.




FIGURE 8.6. The image of $(x, y) \mapsto (x, xy) = (u, v)$.

8.4.A. EXERCISE: CONSTRUCTIBLE SUBSETS ARE FINITE DISJOINT UNIONS OF LOCALLY CLOSED SUBSETS. Recall that a subset of a topological space X is **locally closed** if it is the intersection of an open subset and a closed subset. (Equivalently, it is an open subset of a closed subset, or a closed subset of an open subset. We will later have trouble extending this to open and closed and locally closed subschemes, see Exercise 9.2.B.) Show that a subset of a Noetherian topological space X is constructible if and only if it is the finite disjoint union of locally closed subsets.

As a consequence, if $X \rightarrow Y$ is a continuous map of Noetherian topological spaces, then the preimage of a constructible set is a constructible set. (Important remark: the only reason for the hypothesis of the topological space in question being Noetherian is because this is the only setting in which we have defined constructible sets. An extension of the notion of constructibility to more general topological spaces is mentioned in Exercise 10.3.I.)

8.4.B. EASY EXERCISE (USED IN §8.4.3). Show that the generic point of \mathbb{A}_k^1 does not form a constructible subset of \mathbb{A}_k^1 (where k is a field).

8.4.C. EXERCISE (USED IN EXERCISE 25.5.G).

(a) Show that a constructible subset of a Noetherian scheme is closed if and only if it is “stable under specialization”. More precisely, if Z is a constructible subset of a Noetherian scheme X , then Z is closed if and only if for every pair of points y_1 and y_2 with $y_1 \in \overline{y_2}$, if $y_2 \in Z$, then $y_1 \in Z$. Hint for the “if” implication: show that Z can be written as $\coprod_{i=1}^n U_i \cap Z_i$ where $U_i \subset X$ is open and $Z_i \subset X$ is closed. Show that Z can be written as $\coprod_{i=1}^n U_i \cap Z_i$ (with possibly different n , U_i , Z_i) where each Z_i is irreducible and meets U_i . Now use “stability under specialization” and the generic point of Z_i to show that $Z_i \subset Z$ for all i , so $Z = \cup Z_i$.

(b) Show that a constructible subset of a Noetherian scheme is open if and only if it is “stable under generization”. (Hint: this follows in one line from (a).)

The image of a morphism of schemes can be stranger than a constructible set. Indeed if S is *any* subset of a scheme Y , it can be the image of a morphism: let X be the disjoint union of spectra of the residue fields of all the points of S , and let $\pi: X \rightarrow Y$ be the natural map. This is quite pathological, but in any reasonable situation, the image is essentially no worse than what arose in the previous example of $(x, y) \mapsto (x, xy)$. This is made precise by Chevalley’s Theorem.

8.4.2. Chevalley’s Theorem. — *If $\pi: X \rightarrow Y$ is a finite type morphism of Noetherian schemes, the image of any constructible set is constructible. In particular, the image of π is constructible.*

(For the minority who might care: see §10.3.8 for an extension to locally finitely presented morphisms.) We discuss the proof after giving some important consequences that may seem surprising, in that they are algebraic corollaries of a seemingly quite geometric and topological theorem. The first is a proof of the Nullstellensatz.

8.4.3. Proof of the Nullstellensatz 3.2.6. We wish to show that if K is a field extension of k that is finitely generated as a k -algebra, say by x_1, \dots, x_n , then it is a finite extension of fields. It suffices to show that each x_i is algebraic over k . But if x_i is not algebraic over k , then we have an inclusion of rings $k[x_i] \rightarrow K$, corresponding to a dominant morphism $\pi: \text{Spec } K \rightarrow \mathbb{A}_k^1$ of finite type k -schemes. Of course $\text{Spec } K$ is a single point, so the image of π is one point. By Chevalley’s Theorem 8.4.2 and Exercise 8.4.B, the image of π is not the generic point of \mathbb{A}_k^1 , so $\text{im}(\pi)$ is a closed point of \mathbb{A}_k^1 , and thus π is not dominant. \square

A similar idea can be used in the following exercise.

8.4.D. EXERCISE (QUASIFINITE MORPHISMS TO A FIELD ARE FINITE). Suppose $\pi: X \rightarrow \text{Spec } k$ is a quasifinite morphism. Show that π is finite. (Hint: deal first with the affine case, $X = \text{Spec } A$, where A is finitely generated over k . Suppose A contains an element x that is not algebraic over k , i.e., we have an inclusion $k[x] \hookrightarrow A$. Exercise 8.3.G may help.)

8.4.E. EXERCISE. Suppose X is a scheme of finite type over $\text{Spec } \mathbb{Z}$, and p is a closed point of X . Show that the residue field $\kappa(p)$ is a finite field.

8.4.F. EXERCISE (FOR MAPS OF VARIETIES, SURJECTIVITY CAN BE CHECKED ON CLOSED POINTS). Assume Chevalley’s Theorem 8.4.2. Show that a morphism

of affine k -varieties $\pi: X \rightarrow Y$ is surjective if and only if it is surjective on closed points (i.e., if every closed point of Y is the image of a closed point of X). (Once we define varieties in general, in Definition 11.3.9, you will see that your argument works without change with the adjective “affine” removed.)

8.4.4. Proof of Chevalley’s Theorem 8.4.2.

8.4.G. EXERCISE.

- (a) Reduce the proof of Chevalley’s Theorem to showing just the last sentence: that if $\pi: X \rightarrow Y$ is a finite type morphism of Noetherian schemes, the image of π is constructible.
- (b) Reduce further to the case where Y is affine, say $Y = \text{Spec } B$.
- (b) Reduce further to the case where X is affine.
- (c) Reduce further to the case where $X = \mathbb{A}_B^n = \text{Spec } B[t_1, \dots, t_n]$.
- (d) Reduce further to the case where $X = \mathbb{A}_B^1 = \text{Spec } B[t]$.

We have thus reduced to the case of a morphism of affine schemes corresponding to the ring map

$$B \longrightarrow B[t]/(f_1(t), \dots, f_r(t))$$

where $f_i(t) \in B[t]$. Let $d_i = \deg f_i(t)$, and write

$$f_i(t) = a_{i,d_i} t^{d_i} + \dots + a_{i,1} t + a_{i,0}$$

where $a_{ij} \in B$.

Now $\text{Spec } B$ is the finite disjoint union of locally closed subsets $Y_{\vec{e}}$ indexed by $\vec{e} = (e_1, \dots, e_r)$, where $e_i \in \mathbb{Z}$, $-1 \leq e_i \leq d_i$, and $Y_{\vec{e}}$ is the locus where $f_1(t), \dots, f_r(t)$ have “true degree” e_1, \dots, e_r respectively: the coefficient of t^{e_i} in $f_i(t)$ is nonzero, and the “higher” coefficients are zero (at the points of $Y_{\vec{e}}$). More precisely: $Y_{\vec{e}}$ is the locus where $a_{i,e_i} \neq 0$, and $a_{ij} = 0$ for $j > e_i$. (The zero polynomial has “true degree” -1 under this convention.) Translation: $Y_{\vec{e}} := \text{Spec } B_{\vec{e}}$, where

$$B_{\vec{e}} := \frac{B_{\prod_i a_{i,e_i}}}{(a_{ij})_{j>e_i}}, \quad \text{so} \quad Y_{\vec{e}} = (\cap_i D(a_{i,e_i})) \cap \left(\bigcap_{\substack{i \\ j>e_i}} V(a_{ij}) \right).$$

If $\vec{e} = (-1, \dots, -1)$, then

$$(8.4.4.1) \quad \pi^{-1}(Y_{\vec{e}}) \xrightarrow{\pi} Y_{\vec{e}}$$

is surjective (the map (8.4.4.1) is $\text{Spec } B_{(-1,\dots,-1)}[t] \rightarrow \text{Spec } B_{(-1,\dots,-1)}$), and if $\vec{e} \neq (-1, \dots, -1)$, (8.4.4.1) is a finite morphism corresponding to the map of rings

$$B_{\vec{e}} \longrightarrow B_{\vec{e}}[t]/\left(t^{e_1} + \frac{a_{1,e_1-1}}{a_{1,e_1}} t^{e_1-1} + \dots + \frac{a_{1,0}}{a_{1,e_1}}, \dots, t^{e_r} + \frac{a_{r,e_r-1}}{a_{r,e_r}} t^{e_r-1} + \dots + \frac{a_{r,0}}{a_{r,e_r}}\right).$$

(The ring on the right is generated as a $B_{\vec{e}}$ -module by $1, t, \dots, t^{\max(e_i)}$.) Because the image of finite morphisms are closed (Exercise 8.3.L), the image of X intersect $Y_{\vec{e}}$ (the image of (8.4.4.1)) is a locally closed subset of Y , so the proof is complete. \square

8.4.5. ★ Elimination of quantifiers. A basic sort of question that arises in any number of contexts is when a system of equations has a solution. Suppose for example you have some polynomials in variables x_1, \dots, x_n over an algebraically

closed field \bar{k} , some of which you set to be zero, and some of which you set to be nonzero. (This question is of fundamental interest even before you know any scheme theory!) Then there is an algebraic condition on the coefficients which will tell you if there is a solution. Define the **Zariski topology on \bar{k}^n** in the obvious way: closed subsets are cut out by equations. (A mild generalization of this appears in Exercise 12.2.J.)

8.4.H. EXERCISE (ELIMINATION OF QUANTIFIERS, OVER AN ALGEBRAICALLY CLOSED FIELD). Fix an algebraically closed field \bar{k} . Suppose

$$f_1, \dots, f_p, g_1, \dots, g_q \in \bar{k}[W_1, \dots, W_m, X_1, \dots, X_n]$$

are given. Show that there is a (Zariski-)constructible subset Y of \bar{k}^m such that

$$(8.4.5.1) \quad f_1(w_1, \dots, w_m, X_1, \dots, X_n) = \dots = f_p(w_1, \dots, w_m, X_1, \dots, X_n) = 0$$

and

$$(8.4.5.2) \quad g_1(w_1, \dots, w_m, X_1, \dots, X_n) \neq 0 \quad \dots \quad g_q(w_1, \dots, w_m, X_1, \dots, X_n) \neq 0$$

has a solution $(X_1, \dots, X_n) = (x_1, \dots, x_n) \in \bar{k}^n$ if and only if $(w_1, \dots, w_m) \in Y$. Hints: if Z is a finite type scheme over \bar{k} , and the closed points are denoted Z^{cl} (“cl” is for either “closed” or “classical”), then under the inclusion of topological spaces $Z^{\text{cl}} \hookrightarrow Z$, the Zariski topology on Z induces the Zariski topology on Z^{cl} . Note that we can identify $(\mathbb{A}_{\bar{k}}^p)^{\text{cl}}$ with \bar{k}^p by the Nullstellensatz (Exercise 5.3.F). If X is the locally closed subset of \mathbb{A}^{m+n} cut out by the equalities and inequalities (8.4.5.1) and (8.4.5.2), we have the diagram

$$\begin{array}{ccc} X^{\text{cl}} & \hookrightarrow & X \xrightarrow{\text{loc. cl.}} \mathbb{A}^{m+n} \\ \pi^{\text{cl}} \downarrow & & \downarrow \pi \\ \bar{k}^m & \hookrightarrow & \mathbb{A}^m \end{array}$$

where $Y = \text{im } \pi^{\text{cl}}$. By Chevalley’s Theorem 8.4.2, $\text{im } \pi$ is constructible, and hence so is $(\text{im } \pi) \cap \bar{k}^m$. It remains to show that $(\text{im } \pi) \cap \bar{k}^m = Y (= \text{im } \pi^{\text{cl}})$. You might use the Nullstellensatz.

This is called “elimination of quantifiers” because it gets rid of the quantifier “there exists a solution”. The analogous statement for real numbers, where inequalities are also allowed, is a special case of Tarski’s celebrated theorem of elimination of quantifiers for real closed fields (see, for example, [Ta]).

8.4.6. The Fundamental Theorem of Elimination Theory.

In the case of projective space (and later, projective morphisms), one can do better than Chevalley (at least in describing images of *closed* subsets).

8.4.7. Theorem (Fundamental Theorem of Elimination Theory). — *The morphism $\pi: \mathbb{P}_{\bar{A}}^n \rightarrow \text{Spec } \bar{A}$ is closed (sends closed sets to closed sets).*

Note that *no* Noetherian hypotheses are needed.

A great deal of classical algebra and geometry is contained in this theorem as special cases. Here are some examples.

First, let $A = k[a, b, c, \dots, i]$, and consider the closed subset of $\mathbb{P}_{\bar{A}}^2$ (taken with coordinates x, y, z) corresponding to $ax + by + cz = 0$, $dx + ey + fz = 0$, $gx + hy +$

$iz = 0$. Then we are looking for the locus in $\text{Spec } A$ where these equations have a nontrivial solution. This indeed corresponds to a Zariski-closed set — where

$$\det \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix} = 0.$$

Thus the idea of the determinant is embedded in elimination theory.

As a second example, let $A = k[a_0, a_1, \dots, a_m, b_0, b_1, \dots, b_n]$. Now consider the closed subset of \mathbb{P}_A^1 (taken with coordinates x and y) corresponding to $a_0x^m + a_1x^{m-1}y + \dots + a_my^m = 0$ and $b_0x^n + b_1x^{n-1}y + \dots + b_ny^n = 0$. Then there is a polynomial in the coefficients a_0, \dots, b_n (an element of A) which vanishes if and only if these two polynomials have a common nonzero root — this polynomial is called the *resultant*.

More generally, given a number of homogeneous equations in $n + 1$ variables with indeterminate coefficients, Theorem 8.4.7 implies that one can write down equations in the coefficients that precisely determine when the equations have a nontrivial solution.

8.4.8. Proof of the Fundamental Theorem of Elimination Theory 8.4.7. Suppose $Z \hookrightarrow \mathbb{P}_A^n$ is a closed subset. We wish to show that $\pi(Z)$ is closed.

By the definition of the Zariski topology on $\text{Proj } A[x_0, \dots, x_n]$ (§4.5.7), Z is cut out (set-theoretically) by some homogeneous elements $f_1, f_2, \dots \in A[x_0, \dots, x_n]$. We wish to show that the points $p \in \text{Spec } A$ that are in $\pi(Z)$ form a closed subset. Equivalently, we want to show that those p for which f_1, f_2, \dots have a common zero in $\text{Proj } \kappa(p)[x_0, \dots, x_n]$ form a closed subset of $\text{Spec } A$.

To motivate our argument, we consider a related question. Suppose that $S_\bullet := k[x_0, \dots, x_n]$, and that $g_1, g_2, \dots \in k[x_0, \dots, x_n]$ are homogeneous polynomials. How can we tell if g_1, g_2, \dots have a common zero in $\text{Proj } S_\bullet = \mathbb{P}_k^n$?

They would have a common zero if and only if in

$$\mathbb{A}_k^{n+1} = \text{Spec } S_\bullet,$$

they cut out (set-theoretically) more than the origin, i.e., if set-theoretically

$$V(g_1, g_2, \dots) \not\subset V(x_0, \dots, x_n).$$

By the inclusion-reversing bijection between closed subsets and radical ideals (Theorem 3.7.1), this is true if and only if

$$\sqrt{(g_1, g_2, \dots)} \not\supset \sqrt{(x_0, \dots, x_n)}.$$

But (x_0, \dots, x_n) is radical, so this is true if and only if

$$\sqrt{(g_1, g_2, \dots)} \not\supset (x_0, \dots, x_n).$$

This is true if and only if not all of the generators x_0, \dots, x_n of the ideal (x_0, \dots, x_n) are in $\sqrt{(g_1, g_2, \dots)}$, which is true if only if there is some i such that no power of x_i is in (g_1, g_2, \dots) . This is true if and only if

$$(x_0, \dots, x_n)^N \not\subset (g_1, g_2, \dots) \quad \text{for all } N.$$

(Do you see why this implication goes both ways?) This is equivalent to

$$S_N \not\subset (g_1, g_2, \dots) \quad \text{for all } N,$$

which may be rewritten as

$$(8.4.8.1) \quad S_N \not\subset g_1 S_{N-\deg g_1} \oplus g_2 S_{N-\deg g_2} \oplus \cdots$$

for all N . In other words, this is equivalent to the statement that the k -linear map

$$S_{N-\deg g_1} \oplus S_{N-\deg g_2} \oplus \cdots \rightarrow S_N$$

is not surjective. This map is given by a matrix with $\dim S_N$ rows. (It may have an infinite number of columns, but this will not bother us.) To check that this linear map is not surjective, we need only check that all the “maximal” ($\dim S_N \times \dim S_N$) determinants are zero. (Of course, we need to check this for all N .) Thus the condition that g_1, g_2, \dots have no common zeros in \mathbb{P}_k^n is the same as checking some (admittedly infinite) number of equations. In other words, it is a Zariski-closed condition on the coefficients of the polynomials g_1, g_2, \dots .

8.4.I. EXERCISE. Complete the proof of the Fundamental Theorem of Elimination Theory 8.4.7. (Hint: follow precisely the same argument, with k replaced by A , and the g_i replaced by the f_i . How and why does this prove the theorem?) \square

Notice that projectivity was crucial to the proof: we used graded rings in an essential way. Notice also that the proof is essentially just linear algebra.

Closed embeddings and related notions

9.1 Closed embeddings and closed subschemes

The scheme-theoretic analog of closed subsets has a surprisingly different flavor from the analog of open sets (open subschemes/embeddings). However, just as open subschemes (the scheme-theoretic version of open set) are locally modeled on open sets $U \subset Y$, the analog of closed subsets also has a local model. This was foreshadowed by our understanding of closed subsets of $\text{Spec } B$ as roughly corresponding to ideals. If $I \subset B$ is an ideal, then $\text{Spec } B/I \hookrightarrow \text{Spec } B$ is a morphism of schemes, and we have checked that on the level of topological spaces, this describes $\text{Spec } B/I$ as a closed subset of $\text{Spec } B$, with the subspace topology (Exercise 3.4.I). This morphism is our “local model” of a closed subscheme/embedding.

9.1.1. Definition. A morphism $\pi: X \rightarrow Y$ is a **closed embedding** (or **closed immersion**) if it is an affine morphism, and for every affine open subset $\text{Spec } B \subset Y$, with $\pi^{-1}(\text{Spec } B) \cong \text{Spec } A$, the map $B \rightarrow A$ is surjective (i.e., of the form $B \rightarrow B/I$, our desired local model). (“Embedding” is preferable to “immersion”, because the differential geometric notion of immersion is closer to what algebraic geometers call unramified, which we will define in §22.6.) The symbol \hookrightarrow often is used to indicate that a morphism is a closed embedding (or more generally, a locally closed embedding, §9.2).

If X is a *subset* of Y (and π on the level of sets is the inclusion), we say that X is a **closed subscheme** of Y . A closed embedding is the same thing as an isomorphism with a closed subscheme.

9.1.A. EXERCISE. Show that a closed embedding identifies the topological space of X with a closed subset of the topological space of Y . (Caution: The closed embeddings $\text{Spec } k[x]/(x) \hookrightarrow \text{Spec } k[x]$ and $\text{Spec } k[x]/(x^2) \hookrightarrow \text{Spec } k[x]$ show that the closed subset does not determine the closed subscheme. The “infinitesimal” information, or “fuzz”, is lost.)

9.1.B. EASY EXERCISE. Show that closed embeddings are finite morphisms, hence of finite type.

9.1.C. EASY EXERCISE. Show that the composition of two closed embeddings is a closed embedding.

9.1.D. EXERCISE. Show that the property of being a closed embedding is affine-local on the target.

In particular, if $B \rightarrow A$ is a surjection of rings, then the induced morphism $\text{Spec } A \rightarrow \text{Spec } B$ is a closed embedding.

9.1.2. Important: Closed subschemes correspond to quasicoherent sheaves of ideals. A closed subscheme $\pi: X \hookrightarrow Y$ gives a surjection of \mathcal{O}_Y -modules

$\mathcal{O}_Y \twoheadrightarrow \pi_* \mathcal{O}_X$. The kernel $\mathcal{I}_{X/Y}$ of this map is a sub- \mathcal{O}_Y -module of \mathcal{O}_Y , which we quite reasonably call an **ideal sheaf** (or *sheaf of ideals*) on Y . On each open subset U , it gives an ideal $\mathcal{I}_{X/Y}(U)$ of the ring $\mathcal{O}_Y(U)$; on the affine open subset $\text{Spec } B$, it gives the ideal I of B determining the surjection $B \rightarrow A = B/I$ in the definition of “closed subscheme”. We call the exact sequence (of \mathcal{O}_Y -modules)

$$(9.1.2.1) \quad 0 \longrightarrow \mathcal{I}_{X/Y} \longrightarrow \mathcal{O}_Y \longrightarrow \pi_* \mathcal{O}_X \longrightarrow 0$$

the **closed subscheme exact sequence**. It is actually an exact sequence of *quasicoherent sheaves* on Y :

9.1.E. EXERCISE. Show that $\pi_* \mathcal{O}_X$ is a quasicoherent sheaf on Y . (One way of proceeding is to use Exercise 6.2.D.)

Quasicoherent sheaves form an abelian category, so $\mathcal{I}_{X/Y} = \ker(\mathcal{O}_Y \rightarrow \pi_* \mathcal{O}_X)$ is also a quasicoherent sheaf on Y .

We recover X (and its map to Y) from the ideal sheaf $\mathcal{I}_{X/Y}$ (9.1.2.1): for each affine open subset $\text{Spec } B \subset Y$, the sections over $\text{Spec } B$ give us an exact sequence

$$0 \longrightarrow I \longrightarrow B \longrightarrow B/I \longrightarrow 0, ,$$

and $\text{Spec } B/I$ is $\pi^{-1}(\text{Spec } B)$.

The converse holds:

9.1.F. IMPORTANT EXERCISE. Prove that every quasicoherent sheaf of ideals on Y produces a closed subscheme on Y in this way.

This is a hard exercise, so as a hint, here are two different ways of proceeding; some combination of them may work for you. *Approach 1.* For each affine open $\text{Spec } B$, we have a closed subscheme $\text{Spec } B/I(B) \hookrightarrow \text{Spec } B$. (i) For any two affine open subschemes $\text{Spec } A$ and $\text{Spec } B$, show that the two closed subschemes $\text{Spec } A/I(A) \hookrightarrow \text{Spec } A$ and $\text{Spec } B/I(B) \hookrightarrow \text{Spec } B$ restrict to the *same* closed subscheme of their intersection. (Hint: cover their intersection with open sets simultaneously distinguished in both affine open sets, Proposition 5.3.1.) Thus for example we can glue these two closed subschemes together to get a closed subscheme of $\text{Spec } A \cup \text{Spec } B$. (ii) Use Exercise 4.4.A on gluing schemes (or the ideas therein) to glue together the closed embeddings in all affine open subschemes simultaneously. You will only need to worry about triple intersections. *Approach 2.* (i) Describe X first as a subset of Y . (ii) Check that X is closed. (iii) Define the sheaf of functions \mathcal{O}_X on this subset, perhaps using compatible germs. (iv) Check that this resulting ringed space is indeed locally the closed subscheme given by $\text{Spec } B/I \hookrightarrow \text{Spec } B$.

9.1.G. UNIMPORTANT EXERCISE: A SHEAF OF IDEALS (AS \mathcal{O} -MODULES) THAT IS NOT QUASICOHERENT. Let $X = \text{Spec } k[x]_{(x)}$, the “germ of the affine line at the origin”, which has two points, the closed point and the generic point η . Define $\mathcal{I}(X) = \{0\} \subset \mathcal{O}_X(X) = k[x]_{(x)}$, and $\mathcal{I}(\eta) = k(x) = \mathcal{O}_X(\eta)$. Show that this sheaf of ideals does not correspond to a closed subscheme.

9.1.H. ★★ HARD (AND FRANKLY UNIMPORTANT) EXERCISE (NOT USED LATER). In the literature, the usual definition of a closed embedding is a morphism $\pi: X \rightarrow Y$ such that π induces a homeomorphism of the underlying topological space of X onto a closed subset of the topological space of Y , and the induced map $\pi^\sharp: \mathcal{O}_Y \rightarrow \pi_* \mathcal{O}_X$ of sheaves on Y is surjective. (By “surjective” we mean that the ring morphism on stalks is surjective.) Show that this definition agrees with the one given above. (To show that our definition involving surjectivity on the level of affine open sets implies this definition, you can use the fact that surjectivity of a morphism of sheaves can be checked on a suitably chosen base, Exercise 2.5.D.)

We have now defined the analog of open subsets and closed subsets in the land of schemes. Their definition is slightly less “symmetric” than in the classical topological setting: the “complement” of a closed subscheme is a unique open subscheme, but there are many “complementary” closed subschemes to a given open subscheme in general. (We will soon define one that is “best”, that has a reduced structure, §9.4.9.)

9.1.3. Examples and properties.

9.1.I. IMPORTANT EXERCISE.

(a) In analogy with closed subsets, define the notion of a **finite union of closed subschemes** of X , and an arbitrary (not necessarily finite) **intersection of closed subschemes** of X . (Exercise 9.1.F may help.) Hint: If X is affine, then you might expect that the union of closed subschemes corresponding to I_1 and I_2 would be the closed subscheme corresponding to either $I_1 \cap I_2$ or $I_1 I_2$ — but which one? We would want the union of a closed subscheme with itself to *be* itself, so the right choice is $I_1 \cap I_2$.

(b) Describe the scheme-theoretic intersection of $V(y - x^2)$ and $V(y)$ in \mathbb{A}^2 . See Figure 4.5 for a picture. (For example, explain informally how this corresponds to two curves meeting at a single point with multiplicity 2 — notice how the 2 is visible in your answer. Alternatively, what is the nonreducedness telling you — both its “size” and its “direction”?) Describe their scheme-theoretic union.

(c) Show that the underlying set of a finite union of closed subschemes is the finite union of the underlying sets, and similarly for arbitrary intersections.

(d) Describe the scheme-theoretic intersection of $V(y^2 - x^2)$ and $V(y)$ in \mathbb{A}^2 . Draw a picture. (Did you expect the intersection to have multiplicity one or multiplicity two?) Hence show that if X , Y , and Z are closed subschemes of W , then $(X \cap Z) \cup (Y \cap Z) \neq (X \cup Y) \cap Z$ in general. In particular, not all properties of intersection and union carry over from sets to schemes.

9.1.J. EXERCISE. Show that closed embeddings (like open embeddings, Easy Exercise 8.1.E) are monomorphisms.

In many cases, loci which are a priori closed sets often have enriched versions as closed subschemes. The remaining exercises are examples of this. (We often indicate that we are considering scheme-theoretic enrichments by adding the adjective **scheme-theoretic** — for example, we say “scheme-theoretic intersection”, “scheme-theoretic union”, and so forth.)

9.1.K. IMPORTANT EXERCISE/DEFINITION: THE VANISHING SCHEME (CF. EXERCISE 4.5.S).

(a) Suppose Y is a scheme, and $s \in \Gamma(Y, \mathcal{O}_Y)$. Define the closed subscheme **cut out by s** . We call this the **vanishing scheme** $V(s)$ of s , as it is the scheme-theoretic version of our earlier (set-theoretical) version of $V(s)$ (§3.4.1). (Hint: on affine open $\text{Spec } B$, we just take $\text{Spec } B/(s_B)$, where s_B is the restriction of s to $\text{Spec } B$. Use Exercise 9.1.F to show that this yields a well-defined closed subscheme.)

(b) If u is an invertible function, show that $V(s) = V(su)$.

(c) If S is a set of functions, define $V(S)$.

If M is an A -module, we call the closed subscheme $\text{Spec}(A/\text{Ann } m)$ of $\text{Spec } A$ the **scheme-theoretic support** of m . (This is often described as the “smallest closed subscheme on which m is defined”. Can you make this precise?)

9.1.L. EXERCISE. Show that the underlying set of the scheme-theoretic support of m is the usual “set-theoretic” support, $\text{Supp } m$.

9.1.M. EXERCISE. Define the **scheme-theoretic support of a finitely generated A -module M** to be the scheme-theoretic intersection of all closed subschemes $\text{Spec } A/I \hookrightarrow \text{Spec } A$ for which M is an (A/I) -module. Show that it is the scheme-theoretic union of the scheme-theoretic supports of any finite generating set of M .

9.1.N. EXERCISE (GLOBALIZING THE PREVIOUS EXERCISE). Suppose X is a scheme, and \mathcal{F} is a finite type quasicoherent sheaf on X . How will you define the **scheme-theoretic support of \mathcal{F}** ?

9.2 Locally closed embeddings and locally closed subschemes

Now that we have defined analogs of open and closed subsets, it is natural to define the analog of locally closed subsets. Recall that locally closed subsets are intersections of open subsets and closed subsets. Hence they are closed subsets of open subsets, or equivalently open subsets of closed subsets. The analog of these equivalences will be a little problematic in the land of schemes. (Similarly, the notion of subgroup and quotient group are not quite complementary. Rhetorical question: should “subquotient” of a group G be defined as a subgroup of a quotient of G , or a quotient of a subgroup of G , or are these the same? We will confront the same issue when defining locally closed embeddings.)

We say a morphism $\pi: X \rightarrow Y$ is a **locally closed embedding** (or **locally closed immersion**) if π can be factored into

$$X \xrightarrow[\text{closed}]{\rho} Z \xrightarrow[\text{open}]{\tau} Y$$

where (as indicated) ρ is a closed embedding and τ is an open embedding. For example, the morphism $\text{Spec } k[t, t^{-1}] \rightarrow \text{Spec } k[x, y]$ given by $(x, y) \mapsto (t, 0)$ is a locally closed embedding (Figure 9.1). The symbol \hookrightarrow is often used to indicate that a morphism is a locally closed embedding. Because closed embeddings and open embeddings are monomorphisms (Exercises 9.1.J and 8.1.E respectively), and monomorphisms are preserved by composition, locally closed embeddings are also monomorphisms.




FIGURE 9.1. The locally closed embedding $\text{Spec } k[t, t^{-1}] \rightarrow \text{Spec } k[x, y]$ ($t \mapsto (t, 0) = (x, y)$, i.e., $(x, y) \rightarrow (t, 0)$)

(Warning: The term *immersion* is often used instead of *locally closed embedding* or *locally closed immersion*, but this is unwise terminology, for reasons that already arose for closed embeddings in Definition 9.1.1: The differential geometric notion of immersion is closer to what algebraic geometers call unramified, which we will define in §22.6. Also, the naked term *embedding* should be avoided, because it is needlessly imprecise.)

9.2.1. If X is a subset of Y (and π on the level of sets is the inclusion), we say X is a **locally closed subscheme** of Y . (Warning: the naked term **subscheme** is often used to mean *locally closed subscheme*, but we will avoid this unhappy usage

9.2.A. EASY EXERCISE. Show that locally closed embeddings are locally of finite type.

At this point, you could define the intersection of two locally closed embeddings in a scheme X (which will also be a locally closed embedding in X). But it would be awkward, as you would have to show that your construction is independent of the factorizations of each locally closed embedding into a closed embedding and an open embedding. Instead, we wait until Exercise 10.2.C, when recognizing the intersection as a fibered product will make this easier.

Clearly an open subscheme U of a closed subscheme V of X can be interpreted as a closed subscheme of an open subscheme: as the topology on V is induced from the topology on X , the underlying set of U is the intersection of some open subset U' on X with V . We can take $V' = V \cap U'$, and then $V' \rightarrow U'$ is a closed embedding, and $U' \rightarrow X$ is an open embedding.


It is not clear that a closed subscheme V' of an open subscheme U' can be expressed as an open subscheme of a closed subscheme V . In the category of topological spaces, we would take V as the closure of V' , so we are now motivated to define the analogous construction, which will give us an excuse to introduce several related ideas, in §9.4. We will then resolve this issue in good cases (e.g., if X is Noetherian) in Exercise 9.4.C.

We formalize our discussion in an exercise.

9.2.B. EXERCISE. Suppose $V \rightarrow X$ is a morphism. Consider three conditions:

- (i) V is the intersection of an open subscheme of X and a closed subscheme of X (you will have to define the meaning of “intersection” here, see Exercise 8.1.D, or else see the hint below).
- (ii) V is an open subscheme of a closed subscheme of X , i.e., it factors into an open embedding followed by a closed embedding.
- (iii) V is a closed subscheme of an open subscheme of X , i.e., V is a locally closed embedding.

Show that (i) and (ii) are equivalent, and both imply (iii). (Remark: (iii) does *not* always imply (i) and (ii), see the pathological example [Stacks, tag 01QW].) Hint: It may be helpful to think of the problem as follows. You might hope to think of a locally closed embedding as a fibered diagram.



(We already know that fibered products with open embeddings exist, by Exercise 8.1.D.) Interpret (i) as the existence of the diagram. Interpret (ii) as this diagram minus the lower left corner. Interpret (iii) as the diagram minus the upper right corner.

9.2.C. EXERCISE. Show that the composition of two locally closed embeddings is a locally closed embedding. (Hint: you might use (ii) implies (iii) in the previous exercise.)

9.2.2. Unimportant but fun remark. It may feel odd that in the definition of a locally closed embedding, we had to make a choice (as a composition of a closed embedding followed by an open embedding, rather than vice versa), but this type of issue comes up earlier: a subquotient of a group can be defined as the quotient of a subgroup, or a subgroup of a quotient. Which is the right definition? Or are they the same? (Hint: compositions of two subquotients should certainly be a subquotient, cf. Exercise 9.2.C.)

9.3 Important examples from projective geometry

We now interpret closed embeddings in terms of graded rings. Don’t worry; most of the annoying foundational discussion of graded rings is complete, and we now just take advantage of our earlier work.

9.3.1. Example: Closed embeddings in projective space \mathbb{P}_Λ^n . Recall the definition of projective space \mathbb{P}_Λ^n given in §4.4.10 (and the terminology defined there). Any *homogeneous* polynomial f in x_0, \dots, x_n defines a closed subscheme. (Thus even if f doesn’t make sense as a function, its vanishing scheme still makes sense.) On the open set U_i , the closed subscheme is $V(f(x_{0/i}, \dots, x_{n/i}))$, which we privately

think of as $V(f(x_0, \dots, x_n)/x_i^{\deg f})$. On the overlap

$$U_i \cap U_j = \text{Spec } A[x_{0/i}, \dots, x_{n/i}, x_{j/i}^{-1}]/(x_{i/i} - 1),$$

these functions on U_i and U_j don't exactly agree, but they agree up to a nonvanishing scalar, and hence cut out the same closed subscheme of $U_i \cap U_j$ (Exercise 9.1.K(b)):

$$f(x_{0/i}, \dots, x_{n/i}) = x_{j/i}^{\deg f} f(x_{0/j}, \dots, x_{n/j}).$$

Similarly, a collection of homogeneous polynomials in $A[x_0, \dots, x_n]$ cuts out a closed subscheme of \mathbb{P}_A^n . (Exercise 16.4.H will show that all closed subschemes of \mathbb{P}_A^n are of this form.)

9.3.2. Definition. As usual, let k be a field. A closed subscheme of \mathbb{P}_k^n cut out by a single (nonzero, homogeneous) equation is called a **hypersurface** (in \mathbb{P}_k^n). (Be careful: the word “hypersurface” can be used to mean slightly different things in algebraic geometry, for entirely reasonable reasons, see for example §12.1.5.) Of course, a hypersurface is not cut out by a single global function on \mathbb{P}_k^n : there are no nonconstant global functions (Exercise 4.4.E). The **degree of a hypersurface** is the degree of the polynomial. (Implicit in this is that this notion can be determined from the subscheme itself. You may have the tools to prove this now, but we won't formally prove it until Exercise 19.6.H.) A hypersurface of degree 1 (resp. degree 2, 3, ...) is called a **hyperplane** (resp. **quadric**, **cubic**, **quartic**, **quintic**, **sextic**, **septic**, **octic**, ... **hypersurface**). If $n = 2$, hypersurfaces are called **curves**; a degree 1 hypersurface is called a **line**, and a degree 2 hypersurface is called a **conic curve**, or a **conic** for short. If $n = 3$, a hypersurface is called a **surface**. (In Chapter 12, we will justify the terms *curve* and *surface*.)

9.3.A. EXERCISE.

(a) Show that $wz = xy, x^2 = wy, y^2 = xz$ describes an irreducible subscheme in \mathbb{P}_k^3 . In fact it is a curve, a notion we will define once we know what dimension is. This curve is called the **twisted cubic**. (The twisted cubic is a good nontrivial example of many things, so you should make friends with it as soon as possible. It implicitly appeared earlier in Exercise 3.6.F.)

(b) Show that the twisted cubic is isomorphic to \mathbb{P}_k^1 .

We now extend this discussion to projective schemes in general.

9.3.B. EXERCISE. Suppose that $S_\bullet \twoheadrightarrow R_\bullet$ is a surjection of graded rings. Show that the domain of the induced morphism (Exercise 7.4.A) is $\text{Proj } R_\bullet$, and that the induced morphism $\text{Proj } R_\bullet \rightarrow \text{Proj } S_\bullet$ is a closed embedding.

Exercise 16.4.H will show the converse to Exercise 9.3.B: every closed embedding $X \hookrightarrow \text{Proj } S_\bullet$ into a projective A -scheme (S_\bullet is a finitely generated graded A -algebra) is of the form $\text{Proj}(S_\bullet/I)$, where I is a homogeneous ideal, of “projective functions” vanishing on X . (You may be able to prove this now.)

9.3.C. EXERCISE. Show that an injective linear map of k -vector spaces $V \hookrightarrow W$ induces a closed embedding $\mathbb{P}V \hookrightarrow \mathbb{P}W$. (This is another justification for the definition of $\mathbb{P}V$ in Example 4.5.12 in terms of the *dual* of V .)

9.3.3. Definition. The closed subscheme defined in Exercise 9.3.C is called a **linear space**. Once we know about dimension, we will call this closed subscheme a linear space of dimension $\dim V - 1 = \dim \mathbb{P}V$. More explicitly, a linear space of dimension n in \mathbb{P}^N is any closed subscheme cut out by $N - n$ k -linearly independent homogeneous linear polynomials in x_0, \dots, x_N . A linear space of dimension 1 (resp. 2, n , $\dim \mathbb{P}W - 1$) is called a **line** (resp. **plane**, **n -plane**, **hyperplane**). (If the linear map in the previous exercise is not injective, then the hypothesis (7.4.0.1) of Exercise 7.4.A fails.)

9.3.D. EXERCISE (A SPECIAL CASE OF BÉZOUT'S THEOREM). Suppose $X \subset \mathbb{P}_k^n$ is a degree d hypersurface cut out by $f = 0$, and ℓ is a line not contained in X . A very special case of Bézout's Theorem (Exercise 19.6.K) implies that X and ℓ meet with multiplicity d , "counted correctly". Make sense of this, by restricting the homogeneous degree d polynomial f to the line ℓ , and using the fact that a degree d polynomial in $k[x]$ has d roots, counted properly. (If it makes you feel better, assume $k = \bar{k}$.)

9.3.E. EXERCISE. Show that the map of graded rings $k[w, x, y, z] \rightarrow k[s, t]$ given by $(w, x, y, z) \mapsto (s^3, s^2t, st^2, t^3)$ induces a closed embedding $\mathbb{P}_k^1 \hookrightarrow \mathbb{P}_k^3$, which yields an isomorphism of \mathbb{P}_k^1 with the twisted cubic (defined in Exercise 9.3.A — in fact, this will solve Exercise 9.3.A(b)). Doing this in a hands-on way will set you up well for the general Veronese construction of §9.3.6; see Exercise 9.3.I for a generalization.

9.3.4. A particularly nice case: when S_\bullet is generated in degree 1.

Suppose S_\bullet is a finitely generated graded ring generated in degree 1. Then S_1 is a finitely generated S_0 -module, and the irrelevant ideal S_+ is generated in degree 1 (cf. Exercise 4.5.G(a)).

9.3.F. EXERCISE. Show that if S_\bullet is generated (as an A -algebra) in degree 1 by $n+1$ elements x_0, \dots, x_n , then $\text{Proj } S_\bullet$ may be described as a closed subscheme of \mathbb{P}_A^n as follows. Consider $A^{\oplus(n+1)}$ as a free module with generators t_0, \dots, t_n associated to x_0, \dots, x_n . The surjection of

$$\text{Sym}^\bullet(A^{\oplus(n+1)}) = A[t_0, t_1, \dots, t_n] \twoheadrightarrow S_\bullet$$

$$t_i \longmapsto x_i$$

implies $S_\bullet = A[t_0, t_1, \dots, t_n]/I$, where I is a homogeneous ideal. (In particular, $\text{Proj } S_\bullet$ can always be interpreted as a closed subscheme of some \mathbb{P}_A^n if S_\bullet is finitely generated in degree 1. Then using Exercises 7.4.D and 7.4.G, you can remove the hypothesis of generation in degree 1.)

This is analogous to the fact that if R is a finitely generated A -algebra, then choosing n generators of R as an algebra is the same as describing $\text{Spec } R$ as a closed subscheme of \mathbb{A}_A^n . In the affine case this is "choosing coordinates"; in the projective case this is "choosing projective coordinates".

Recall (Exercise 4.4.F) that if k is algebraically closed, then we can interpret the closed points of \mathbb{P}^n as the lines through the origin in $(n+1)$ -space. The following exercise states this more generally.

9.3.G. EXERCISE. Suppose S_\bullet is a finitely generated graded ring over an algebraically closed field k , generated in degree 1 by x_0, \dots, x_n , inducing closed embeddings $\text{Proj } S_\bullet \hookrightarrow \mathbb{P}^n$ and $\text{Spec } S_\bullet \hookrightarrow \mathbb{A}^{n+1}$. Give a bijection between the closed points of $\text{Proj } S_\bullet$ and the “lines through the origin” in $\text{Spec } S_\bullet \subset \mathbb{A}^{n+1}$.

9.3.5. A second proof that finite morphisms are closed. This interpretation of $\text{Proj } S_\bullet$ as a closed subscheme of projective space (when it is generated in degree 1) yields the following second proof of the fact (shown in Exercise 8.3.L) that finite morphisms are closed. Suppose $\pi: X \rightarrow Y$ is a finite morphism. The question is local on the target, so it suffices to consider the affine case $Y = \text{Spec } B$. It suffices to show that $\pi(X)$ is closed. Then by Exercise 8.3.I, X is a projective B -scheme, and hence by the Fundamental Theorem of Elimination Theory 8.4.7, its image is closed.

9.3.6. Important classical construction: The Veronese embedding.

Suppose $S_\bullet = k[x, y]$, so $\text{Proj } S_\bullet = \mathbb{P}_k^1$. Then $S_{2\bullet} = k[x^2, xy, y^2] \subset k[x, y]$ (see §7.4.4 on the Veronese subring). We identify this subring as follows.

9.3.H. EXERCISE. Let $u = x^2, v = xy, w = y^2$. Show that $S_{2\bullet} \cong k[u, v, w]/(uw - v^2)$, by mapping u, v, w to x^2, xy, y^2 , respectively.

We have a graded ring generated by three elements in degree 1. Thus we think of it as sitting “in” \mathbb{P}^2 , via the construction of §9.3.F. This can be interpreted as “ \mathbb{P}^1 as a conic in \mathbb{P}^2 ”.

9.3.7. Thus if k is algebraically closed of characteristic not 2, using the fact that we can diagonalize quadratics (Exercise 5.4.J), the conics in \mathbb{P}^2 , up to change of coordinates, come in only a few flavors: sums of 3 squares (e.g., our conic of the previous exercise), sums of 2 squares (e.g., $y^2 - x^2 = 0$, the union of 2 lines), a single square (e.g., $x^2 = 0$, which looks set-theoretically like a line, and is nonreduced), and 0 (perhaps not a conic at all). Thus we have proved: any plane conic (over an algebraically closed field of characteristic not 2) that can be written as the sum of three nonzero squares is isomorphic to \mathbb{P}^1 . (See Exercise 7.5.F for a closely related fact.)

We now soup up this example.

9.3.I. EXERCISE. We continue to take $S_\bullet = k[x, y]$. Show that $\text{Proj } S_{d\bullet}$ is given by the equations that

$$\begin{pmatrix} y_0 & y_1 & \cdots & y_{d-1} \\ y_1 & y_2 & \cdots & y_d \end{pmatrix}$$

is rank 1 (i.e., that all the 2×2 minors vanish). This is called the **degree d rational normal curve** “in” \mathbb{P}^d . You did the *twisted cubic* case $d = 3$ in Exercises 9.3.A and 9.3.E.

9.3.8. Definition. More generally, if $S_\bullet = k[x_0, \dots, x_n]$, then $v_d : \text{Proj } S_\bullet \rightarrow \mathbb{P}^{N-1}$ (where N is the dimension of the vector space of homogeneous degree d polynomials in x_0, \dots, x_n) is called the **d-uple embedding** or the (d-uple) **Veronese embedding**. (Combining Exercise 7.4.E with Exercise 9.3.F shows that v_d is a closed embedding.)

9.3.J. COMBINATORIAL EXERCISE (CF. REMARK 4.5.3). Show that $N = \binom{n+d}{d}$.

9.3.K. UNIMPORTANT EXERCISE. Find six linearly independent quadratic equations vanishing on the **Veronese surface** $\text{Proj } S_{2\bullet}$ where $S_\bullet = k[x_0, x_1, x_2]$, which sits naturally in \mathbb{P}^5 . (You needn't show that these equations generate all the equations cutting out the Veronese surface, although this is in fact true.) Possible hint: use the identity

$$\det \begin{pmatrix} x_0x_0 & x_0x_1 & x_0x_2 \\ x_1x_0 & x_1x_1 & x_1x_2 \\ x_2x_0 & x_2x_1 & x_2x_2 \end{pmatrix} = 0.$$

9.3.9. Rulings on the quadric surface. We return to rulings on the quadric surface, which first appeared in the optional (starred) section §4.4.12.

9.3.L. USEFUL GEOMETRIC EXERCISE: THE RULINGS ON THE QUADRIC SURFACE $wz = xy$. This exercise is about the lines on the quadric surface X given by $wz - xy = 0$ in \mathbb{P}_k^3 (where the projective coordinates on \mathbb{P}_k^3 are ordered w, x, y, z). This construction arises all over the place in nature.

(a) Suppose a_0 and b_0 are given elements of k , not both zero. Make sense of the statement: as $[c, d]$ varies in \mathbb{P}^1 , $[a_0c, b_0c, a_0d, b_0d]$ is a line in the quadric surface. (This describes “a family of lines parametrized by \mathbb{P}^1 ”, although we can't yet make this precise.) Find another family of lines. These are the two **rulings** of the smooth quadric surface.

(b) Show that through every k -valued point of the quadric surface X , there passes one line from each ruling.

(c) Show there are no other lines. (There are many ways of proceeding. At risk of predisposing you to one approach, here is a germ of an idea. Suppose L is a line on the quadric surface, and $[1, x, y, z]$ and $[1, x', y', z']$ are distinct points on it. Because they are both on the quadric, $z = xy$ and $z' = x'y'$. Because all of L is on the quadric, $(1+t)(z+tz') - (x+tx')(y+ty') = 0$ for all t . After some algebraic manipulation, this translates into $(x-x')(y-y') = 0$. How can this be made watertight? Another possible approach uses Bézout's Theorem, in the form of Exercise 9.3.D.)

Hence by Exercise 5.4.J, if we are working over an algebraically closed field of characteristic not 2, we have shown that all rank 4 quadric surfaces have two rulings of lines. (In Example 10.6.2, we will recognize this quadric as $\mathbb{P}^1 \times \mathbb{P}^1$.)

9.3.10. Side Remark: Rulings on quadrics. The existence of these two rulings is the first chapter of a number of important and beautiful stories. The second chapter is often the following. If k is an algebraically closed field, then a maximal rank (Exercise 5.4.J) quadric hypersurface X of dimension m contains no linear spaces of dimension greater than $m/2$. (We will see in Exercise 13.3.D that the maximal rank quadric hypersurfaces are the “smooth” quadrics.) If $m = 2a + 1$, then X contains




FIGURE 9.2. One of the two rulings on the quadric surface $V(wz - xy) \subset \mathbb{P}^3$. One ruling contains the line $V(w, x)$, and the other contains the line $V(w, y)$.

an irreducible $\binom{a+2}{2}$ -dimensional family of a -planes. If $m = 2a$, then X contains two irreducible $\binom{a+1}{2}$ -dimensional families of a -planes, and furthermore two a -planes Λ and Λ' are in the same family if and only if $\dim(\Lambda \cap \Lambda') \equiv a \pmod{2}$. These families of linear spaces are also called **rulings**. (For more information, see [GH1, §6.1, p. 735, Prop.].) You already know enough to think through the examples of $m = 0, 1$, and 2 . (The case $m = 2$, rulings on a quadric surface, turn up in §4.4.12, Figure 9.2, Exercise 9.3.L, §10.6.2, §12.2.18, §13.2.14, Exercise 15.2.U, and Exercise 20.10.L.) The case $m = 3$ is discussed in Exercise 17.7.K.

9.3.11. Affine and projective cones (Figure 9.3).

If S_\bullet is a finitely generated graded ring, then the **affine cone** of $\text{Proj } S_\bullet$ is $\text{Spec } S_\bullet$. Caution: this terminology is not ideal, as this construction depends on S_\bullet , not just on $\text{Proj } S_\bullet$. As motivation, consider the graded ring $S_\bullet = \mathbb{C}[x, y, z]/(z^2 - x^2 - y^2)$. Figure 9.4 is a sketch of $\text{Spec } S_\bullet$. (Here we draw the “real picture” of $z^2 = x^2 + y^2$ in \mathbb{R}^3 .) It is a cone in the traditional sense; the origin $(0, 0, 0)$ is the “cone point”.




FIGURE 9.3. The affine and projective cones of $\text{Proj } S_\bullet$.

This gives a useful way of picturing Proj (even over arbitrary rings, not just \mathbb{C}). Intuitively, you could imagine that if you discarded the origin, you would get something that would project onto $\text{Proj } S_\bullet$. The following exercise makes that precise.

9.3.M. EXERCISE. If $\text{Proj } S_\bullet$ is a projective scheme over a field k , describe a natural morphism $\text{Spec } S_\bullet \setminus V(S_+) \rightarrow \text{Proj } S_\bullet$. (Can you see why $V(S_+)$ is a single point, and should reasonably be called the origin?)

This readily generalizes to the following exercise, which again motivates the terminology “irrelevant”.

9.3.N. EASY EXERCISE. If S_\bullet is a finitely generated graded ring over a base ring A , describe a natural morphism $\text{Spec } S_\bullet \setminus V(S_+) \rightarrow \text{Proj } S_\bullet$.

In fact, it can be made precise that $\text{Proj } S_\bullet$ is the quotient (by the multiplicative group of scalars) of the affine cone minus the origin.

9.3.12. Definition. The **projective cone** of $\text{Proj } S_\bullet$ is $\text{Proj } S_\bullet[T]$, where T is a new variable of degree 1. For example, the cone corresponding to the conic $\text{Proj } k[x, y, z]/(z^2 - x^2 - y^2)$ is $\text{Proj } k[x, y, z, T]/(z^2 - x^2 - y^2)$. The projective cone is sometimes called the **projective completion** of $\text{Spec } S_\bullet$. (Note: this depends on S_\bullet , not just on $\text{Proj } S_\bullet$.)

9.3.O. LESS IMPORTANT EXERCISE (CF. §4.5.1). Show that the “projective cone” $\text{Proj } S_\bullet[T]$ of $\text{Proj } S_\bullet$ has a closed subscheme isomorphic to $\text{Proj } S_\bullet$ (informally, corresponding to $T = 0$), whose complement (the distinguished open set $D(T)$) is isomorphic to the affine cone $\text{Spec } S_\bullet$.


This construction can be usefully pictured as the affine cone union some points “at infinity”, and the points at infinity form the Proj . (The reader may wish to start with Figure 9.4, and try to visualize the conic curve “at infinity”, and then compare this visualization to Figure 4.9.)

We have thus completely described the algebraic analog of the classical picture of §4.5.1.

9.3.13. Definition. If we put a nonstandard weighting on the variables of $k[x_1, \dots, x_n]$ — say we give x_i degree d_i — then $\text{Proj } k[x_1, \dots, x_n]$ is called **weighted projective space** $\mathbb{P}(d_1, d_2, \dots, d_n)$.

9.3.P. EXERCISE. (a) Show that $\mathbb{P}(m, n)$ is isomorphic to \mathbb{P}^1 . (b) Show that $\mathbb{P}(1, 1, 2) \cong \text{Proj } k[u, v, w, z]/(uw - v^2)$. Hint: do this by looking at the even-graded parts of $k[x_0, x_1, x_2]$, cf. Exercise 7.4.D. (This is a projective cone over a conic curve. Over a field of characteristic not 2, it is isomorphic to the traditional cone $x^2 + y^2 = z^2$ in \mathbb{P}^3 , see Figure 9.4.)

9.3.Q. EXERCISE (COMBINING A NUMBER OF THE ABOVE CONCEPTS). Show that the $\mathbb{P}(1, 1, n)$ is isomorphic to the projective cone over the degree n rational normal curve.

FIGURE 9.4. The cone $x^2 + y^2 = z^2$.

9.4 The (closed sub)scheme-theoretic image

We now define a series of notions that are all of the form “the smallest closed subscheme such that something or other is true”. One example will be the notion of scheme-theoretic closure of a locally closed embedding, which will allow us to interpret locally closed embeddings in three equivalent ways (open subscheme intersect closed subscheme; open subscheme of closed subscheme; and closed subscheme of open subscheme — cf. Exercise 9.2.B).

9.4.1. Scheme-theoretic image.

We start with the notion of scheme-theoretic image. Set-theoretic images are badly behaved in general (§8.4.1), and even with reasonable hypotheses such as those in Chevalley’s Theorem 8.4.2, things can be confusing. For example, there is no reasonable way to impose a scheme structure on the image of $\mathbb{A}_k^2 \rightarrow \mathbb{A}_k^2$ given by $(x, y) \mapsto (x, xy)$. It will be useful (e.g., Exercise 9.4.C) to define a notion of a closed subscheme of the target that “best approximates” the image. This will incorporate the notion that the image of something with nonreduced structure (“fuzz”) can also have nonreduced structure. As usual, we will need to impose reasonable hypotheses to make this notion behave well (see Theorem 9.4.4 and Corollary 9.4.5).

9.4.2. Definition. Suppose $i: Z \hookrightarrow Y$ is a closed subscheme, giving an exact sequence $0 \rightarrow \mathcal{I}_{Z/Y} \rightarrow \mathcal{O}_Y \rightarrow i_*\mathcal{O}_Z \rightarrow 0$. We say that **the image of $\pi: X \rightarrow Y$ lies in Z** if the composition $\mathcal{I}_{Z/Y} \rightarrow \mathcal{O}_Y \rightarrow \pi_*\mathcal{O}_X$ is zero. Informally, locally, functions vanishing on Z pull back to the zero function on X . If the image of π lies in some subschemes Z_j (as j runs over some index set), it clearly lies in their intersection (cf. Exercise 9.1.I(a) on intersections of closed subschemes). We then define the **scheme-theoretic image of π** , a closed subscheme of Y , as the “smallest closed subscheme containing the image”, i.e., the intersection of all closed subschemes containing the image. In particular (and in our first examples), if Y is affine, the scheme-theoretic image is cut out by functions on Y that are 0 when pulled back to X .

Other reasonable names for this concept are the *closed subscheme-theoretic image*, or *image closed subscheme*. They have the advantage that they remind the reader

that the scheme-theoretic image is a *closed* subscheme, but we won't use these phrases.

Example 1. Consider $\pi: \operatorname{Spec} k[\epsilon]/(\epsilon^2) \rightarrow \operatorname{Spec} k[x] = \mathbb{A}_k^1$ given by $x \mapsto \epsilon$. Then the scheme-theoretic image of π is given by $\operatorname{Spec} k[x]/(x^2)$ (the polynomials pulling back to 0 are precisely multiples of x^2). Thus the image of the fuzzy point still has some fuzz.

Example 2. Consider $\pi: \operatorname{Spec} k[\epsilon]/(\epsilon^2) \rightarrow \operatorname{Spec} k[x] = \mathbb{A}_k^1$ given by $x \mapsto 0$. Then the scheme-theoretic image is given by $\operatorname{Spec} k[x]/(x)$: the image is reduced. In this picture, the fuzz is “collapsed” by π .

Example 3. Consider $\pi: \operatorname{Spec} k[t, t^{-1}] = \mathbb{A}^1 - \{0\} \rightarrow \mathbb{A}^1 = \operatorname{Spec} k[u]$ given by $u \mapsto t$. Any function $g(u)$ which pulls back to 0 as a function of t must be the zero-function. Thus the scheme-theoretic image is everything. The set-theoretic image, on the other hand, is the distinguished open set $\mathbb{A}^1 - \{0\}$. Thus in not-too-pathological cases, the underlying set of the scheme-theoretic image is not the set-theoretic image. But the situation isn't terrible: the underlying set of the scheme-theoretic image must be closed, and indeed it is the closure of the set-theoretic image. We might imagine that in reasonable cases this will be true, and in even nicer cases, the underlying set of the scheme-theoretic image will be set-theoretic image. We will later see that this is indeed the case (§9.4.6).

But sadly pathologies can sometimes happen in, well, pathological situations.

Example 4 (Figure 9.5). Let $X = \coprod \operatorname{Spec} k[\epsilon_n]/((\epsilon_n)^n)$ (a scheme which appeared in the hint to Exercise 5.2.F) and $Y = \operatorname{Spec} k[x]$, and define $X \rightarrow Y$ by $x \mapsto \epsilon_n$ on the n th component of X . If a function $g(x)$ on Y pulls back to 0 on X , then its Taylor expansion is 0 to order n (by examining the pullback to the n th component of X) for all n , so $g(x)$ must be 0. (This argument will be vastly generalized in Exercise 13.9.A(b).) Thus the scheme-theoretic image is $V(0)$ on Y , i.e., Y itself, while the set-theoretic image is easily seen to be just the origin (the closed point 0). (This morphism implicitly arises in Caution/Example 9.4.11.)




FIGURE 9.5. Yuck (Example 4 of Definition 9.4.2)

9.4.3. Criteria for computing scheme-theoretic images affine-locally. Example 4 clearly is weird though, and we can show that in “reasonable circumstances” such pathology doesn't occur.

In the special case where the target Y is affine (pay attention to how this argument works in Examples 1 through 3), say $Y = \operatorname{Spec} B$, almost by definition the

scheme-theoretic image of $\pi: X \rightarrow Y$ is cut out by the ideal $I \subset B$ of functions on Y which pull back to zero on X .

It would be great to use this to compute the scheme-theoretic image affine by affine (affine-locally). On the affine open set $\text{Spec } B \subset Y$, define the ideal $I(B) \subset B$ of functions which pull back to 0 on X . Formally, $I(B) := \ker(B \rightarrow \Gamma(\text{Spec } B, \pi_*(\mathcal{O}_X)))$. If for each such B , and each $g \in B$, the map $\phi: I(B)_g \rightarrow I(B_g)$ is an isomorphism, then we will have defined the scheme-theoretic image as a closed subscheme (see Exercise 9.1.F). Injectivity of ϕ_g is straightforward: if $r \in B$ and r/g^n pulls back to 0 on $\pi^{-1}(\text{Spec } B_g)$, then as g is nonvanishing at all points of $\text{Spec } B_g = D(g)$, we have that $\pi^{-1}(r)$ is 0 on $\pi^{-1}(\text{Spec } B_g)$ as well.

So the question is the surjectivity of ϕ , which translates to: given a function r/g^n on $D(g)$ that pulls back to zero on $\pi^{-1}(D(g))$, is it true that for some m , $rg^m = 0$ when pulled back to $\pi^{-1}(\text{Spec } B)$?

(i) We first show that the answer is yes if $\pi^{-1}(\text{Spec } B)$ is reduced. In this case we can take advantage of Exercise 5.2.A, that a function on a reduced scheme is zero if it has value zero at every point. We may then take $m = 1$ (as $\pi^{-1}(r)$ vanishes on $\pi^{-1}(D(g))$ and g vanishes on $V(g)$, so $\pi^{-1}(rg)$ vanishes on $\pi^{-1}(\text{Spec } B) = \pi^{-1}(D(g)) \cup \pi^{-1}(V(g))$).

(ii) The answer is also yes if $\pi^{-1}(\text{Spec } B)$ is affine, say $\text{Spec } A$: if $r' = \pi^*r$ and $g' = \pi^*g$ in A , then if $r' = 0$ on $D(g')$, then there is an m such that $r'(g')^m = 0$ (as the statement $r' = 0$ in $D(g')$ means precisely this fact — the functions on $D(g')$ are $A_{g'}$).

(ii)' More generally, the answer is yes if $\pi^{-1}(\text{Spec } B)$ is quasicompact: cover $\pi^{-1}(\text{Spec } B)$ with finitely many affine open sets. For each one there will be some m_i so that $rg^{m_i} = 0$ when pulled back to this open set. Then let $m = \max(m_i)$. (We see again that quasicompactness is our friend!)

In conclusion, we have proved the following (subtle) theorem.

9.4.4. Theorem. — *Suppose $\pi: X \rightarrow Y$ is a morphism of schemes. If X is reduced or π is quasicompact, then the scheme-theoretic image of π may be computed affine-locally: on $\text{Spec } A \subset Y$, it is cut out by the functions (elements of A) that pull back to the function 0 (on $\pi^{-1}(\text{Spec } A)$).*

9.4.5. Corollary. — *Under the hypotheses of Theorem 9.4.4, the closure of the set-theoretic image of π is the underlying set of the scheme-theoretic image.*

(Example 4 above shows how we need these hypotheses.)

9.4.6. In particular, if the set-theoretic image is closed (e.g., if π is finite or projective), the set-theoretic image is the underlying set of the scheme-theoretic image, as promised in Example 3 above.

9.4.7. Proof of Corollary 9.4.5. The set-theoretic image is in the underlying set of the scheme-theoretic image. (Check this!) The underlying set of the scheme-theoretic image is closed, so the closure of the set-theoretic image is contained in the underlying set of the scheme-theoretic image. On the other hand, if U is the complement of the closure of the set-theoretic image, $\pi^{-1}(U) = \emptyset$. Under these hypotheses, the scheme theoretic image can be computed locally, so the scheme-theoretic image is the empty set on U . \square

We conclude with a few stray remarks.

9.4.A. EASY EXERCISE. If X is reduced, show that the scheme-theoretic image of $\pi: X \rightarrow Y$ is also reduced.

More generally, you might expect there to be no unnecessary nonreduced structure on the image not forced by nonreduced structure on the source. The following makes this precise in the locally Noetherian case (when we can talk about associated points).

9.4.B. ★ UNIMPORTANT EXERCISE. If $\pi: X \rightarrow Y$ is a *quasicompact* morphism of locally Noetherian schemes, show that the associated points of the image subscheme are a subset of the image of the associated points of X . (The example of $\coprod_{a \in \mathbb{C}} \operatorname{Spec} \mathbb{C}[t]/(t - a) \rightarrow \operatorname{Spec} \mathbb{C}[t]$ shows what can go wrong if you give up quasicompactness — note that reducedness of the source doesn't help.) Hint: reduce to the case where X and Y are affine. (Can you develop your geometric intuition so that this becomes plausible to you?)

9.4.8. Scheme-theoretic closure of a locally closed subscheme.

We define the **scheme-theoretic closure** of a locally closed embedding $\pi: X \hookrightarrow Y$ as the scheme-theoretic image of π . (A shorter phrase for this is *schematic closure*, although this more elegant nomenclature has not caught on.)

9.4.C. EXERCISE. If a locally closed embedding $V \hookrightarrow X$ is quasicompact (e.g., if V is Noetherian, Exercise 8.3.B(a)), or if V is reduced, show that (iii) implies (i) and (ii) in Exercise 9.2.B. Thus in this fortunate situation, a locally closed embedding can be thought of in three different ways, whichever is convenient. (Hint: Corollary 9.4.5.)

9.4.D. UNIMPORTANT EXERCISE, USEFUL FOR INTUITION. If $\pi: X \hookrightarrow Y$ is a locally closed embedding into a locally Noetherian scheme (so X is also locally Noetherian), then the associated points of the scheme-theoretic closure are (naturally in bijection with) the associated points of X . (Hint: Exercise 9.4.B.) Informally, the only nonreduced structure on the scheme-theoretic closure is that “forced by” nonreduced structure on X .

9.4.9. The (reduced) subscheme structure on a closed subset.

Suppose X^{set} is a closed subset of a scheme Y . Then we can define a canonical scheme structure X on X^{set} that is reduced. We could describe it as being cut out by those functions whose values are zero at all the points of X^{set} . On the affine open set $\operatorname{Spec} B$ of Y , if the set X^{set} corresponds to the radical ideal $I = I(X^{\text{set}} \cap \operatorname{Spec} B)$ (recall the $I(\cdot)$ function from §3.7), the scheme X corresponds to $\operatorname{Spec} B/I$. You can quickly check that this behaves well with respect to any distinguished inclusion $\operatorname{Spec} B_f \hookrightarrow \operatorname{Spec} B$. We could also consider this construction as an example of a scheme-theoretic image in the following crazy way: let W be the scheme that is a disjoint union of all the points of X^{set} , where the point corresponding to p in X^{set} is Spec of the residue field of $\mathcal{O}_{Y,p}$. Let $\rho: W \rightarrow Y$ be the “canonical” map sending “ p to p ”, and giving an isomorphism on residue fields. Then the scheme structure on X is the scheme-theoretic image of ρ . A third definition: it is the smallest closed subscheme whose underlying set contains X^{set} .

9.4.E. EXERCISE. Show that all three definitions are the same.

This construction is called the (induced) **reduced subscheme structure** on the closed subset X^{set} . (Vague exercise: Make a definition of the reduced subscheme structure precise and rigorous to your satisfaction.)

9.4.F. EXERCISE. Show that the underlying set of the induced reduced subscheme $X \hookrightarrow Y$ is indeed the closed subset X^{set} . Show that X is reduced.

9.4.10. Reduced version of a scheme.

In the main interesting case where X^{set} is all of Y , we obtain a *reduced closed subscheme* $Y^{\text{red}} \rightarrow Y$, called the **reduction** of Y . On the affine open subset $\text{Spec } B \hookrightarrow Y$, $Y^{\text{red}} \hookrightarrow Y$ corresponds to the nilradical $\mathfrak{N}(B)$ of B . The reduction of a scheme is the “reduced version” of the scheme, and informally corresponds to “shearing off the fuzz”.

An alternative equivalent definition: on the affine open subset $\text{Spec } B \hookrightarrow Y$, the reduction of Y corresponds to the ideal $\mathfrak{N}(B) \subset B$ of nilpotents. As for any $f \in B$, $\mathfrak{N}(B)_f = \mathfrak{N}(B_f)$, by Exercise 9.1.F this defines a closed subscheme.

9.4.G. EXERCISE. Show that this alternative definition is indeed equivalent to the actual one.

9.4.11. ★ Caution/Example. It is not true that for *every* open subset $U \subset Y$, $\Gamma(U, \mathcal{O}_{Y^{\text{red}}})$ is $\Gamma(U, \mathcal{O}_Y)$ modulo its nilpotents. For example, on $Y = \coprod \text{Spec } k[x]/(x^n)$, the function x is not nilpotent, but is 0 on Y^{red} , as it is “locally nilpotent”. This may remind you of Example 4 after Definition 9.4.2.

9.5 Slicing by effective Cartier divisors, regular sequences and regular embeddings

We now introduce regular embeddings, an important class of locally closed embeddings. Locally closed embeddings of regular schemes in regular schemes are one important example of regular embeddings (Exercise 13.2.L(b)). Effective Cartier divisors, basically the codimension 1 case, will turn out to be repeatedly useful as well, with repeated “slicing by effective Cartier divisors” providing the inductive step in many arguments.

9.5.1. Locally principal closed subschemes, and effective Cartier divisors. A closed subscheme is **locally principal** if on each open set in a small enough open cover it is cut out by a single equation (i.e., by a principal ideal, hence the terminology). More specifically, a locally principal closed subscheme is a closed subscheme $\pi: X \hookrightarrow Y$ for which there is an open cover $\{U_i\}$ of Y for which for each i , $\pi^{-1}(U_i) \hookrightarrow U_i$ the closed subscheme $V(s_i)$ of U_i for some $s_i \in \Gamma(U_i, \mathcal{O}_X)$. (These open cover can clearly be chosen to be by affine open subsets.)

For example, hypersurfaces in \mathbb{P}_k^n (Definition 9.3.2) are locally principal: each homogeneous polynomial in $n + 1$ variables defines a locally principal closed subscheme of \mathbb{P}_k^n . (Warning: unlike “local principality”, “principality” is not an

affine-local condition, see §20.11.10! Also, the example of a projective hypersurface, §9.3.1, shows that a locally principal closed subscheme need not be cut out by a globally-defined function.)

If the ideal sheaf is locally generated near every point by a function that is not a zerodivisor, we call the closed subscheme an *effective Cartier divisor*. More precisely: if $\pi: X \rightarrow Y$ is a closed embedding, and there is a cover Y by *affine* open subsets $\text{Spec } A_i \subset Y$, and there exist non-zerodivisors $t_i \in A_i$ with $V(t_i) = X|_{\text{Spec } A_i}$ (scheme-theoretically — i.e., the ideal sheaf of X over $\text{Spec } A_i$ is generated by t_i), then we say that X is an **effective Cartier divisor** on Y . (We will not explain the origin of the phrase, as it is not relevant for this point of view.) We will often use the evocative phrase **slicing by an effective Cartier divisor** when we mean “restrict to an effective Cartier divisor”.

The notion of an effective Cartier divisor is central. The notion of a locally principal closed subscheme is much less so.

9.5.A. EXERCISE. Suppose $t \in A$ is a non-zerodivisor. Show that t is a non-zerodivisor in A_p for each prime p .

9.5.2. Caution. If D is an effective Cartier divisor on an affine scheme $\text{Spec } A$, it is not necessarily true that $D = V(t)$ for some $t \in A$ (see Exercise 15.2.M — §20.11.10 gives a different flavor of example). In other words, the condition of a closed subscheme being an effective Cartier divisor can be verified on *an* affine cover, but cannot be checked on an *arbitrary* affine cover — it is not an affine-local condition in this obvious a way.

9.5.B. EXERCISE. Suppose X is a locally Noetherian scheme. If $t \in \Gamma(X, \mathcal{O}_X)$ is a function on X , show that t (or more precisely the closed subscheme $V(t)$) is an effective Cartier divisor if and only if it doesn’t vanish on any associated point of X . Show that a locally principal closed subscheme of X is an effective Cartier divisor if and only if it doesn’t contain any associated points of X .

9.5.C. EXERCISE. Suppose $V(t) = V(t') \hookrightarrow \text{Spec } A$ is an effective Cartier divisor, with t and t' non-zerodivisors in A . Show that t is an invertible function times t' .

The idea of effective Cartier divisors leads us to the notion of regular sequences. (We will close the loop in Exercise 9.5.H, where we will interpret effective Cartier divisors on any reasonable scheme as regular embeddings of codimension 1.)

9.5.3. Regular sequences.

The notion of “regular sequence”, like so much else, is due to Serre, [Se2]. The definition of regular sequence is the algebraic version of the following geometric idea: locally, we take an effective Cartier divisor (a non-zerodivisor); then an effective Cartier divisor on that; then an effective Cartier divisor on that; and so on, a finite number of times. A little care is necessary; for example, we might want this to be independent of the order of the equations imposed, and this is true only when we say this in the right way.

We make the definition of regular sequence for a ring A , and more generally for an A -module M .

9.5.4. Definition. If M is an A -module, a sequence $x_1, \dots, x_r \in A$ is called an **M -regular sequence** (or a **regular sequence for M**) if the following two conditions are satisfied.

- (i) For each i , x_i is not a zerodivisor for $M/(x_1, \dots, x_{i-1})M$. (The case $i = 1$ should be interpreted as: “ x_1 is not a zerodivisor of M .” Recall that zerodivisors were defined in Exercise 1.3.C.)
- (ii) We have a proper inclusion $(x_1, \dots, x_r)M \subsetneq M$.

In the case most relevant to us, when $M = A$, this should be seen as a reasonable approximation of a “complete intersection”, and indeed we will use this as the definition (§9.5.7).

We say r is the **length of the regular sequence** $x_1, \dots, x_r \in A$. An A -regular sequence is just called a **regular sequence**.

9.5.D. EXERCISE. If M is an A -module, show that an M -regular sequence continues to satisfy condition (i) of the Definition 9.5.4 of regular sequence upon any localization. (Once you know what flatness is, you will see that your argument shows that condition (i) is preserved by any flat base change.)

9.5.E. EXERCISE. If x, y is an M -regular sequence, show that x^N, y is an M -regular sequence. Hint: the difficult part is showing that y is not a zerodivisor of $M/(x^N M)$. Show this by induction on N . If y is a zerodivisor of $M/(x^N M)$, then $ym \equiv 0 \pmod{x^N}$ for some $m \in M \setminus x^N M$. Hence $ym = x^N k$ for some $k \in M$. Use the fact that x, y is a regular sequence to show that m is a multiple of x . (Your argument will easily extend to show more generally that if x_1, \dots, x_n is a regular sequence, and $a_1, \dots, a_n \in \mathbb{Z}^+$, then $x_1^{a_1}, \dots, x_n^{a_n}$ is a regular sequence.)

9.5.5. Regularity of a sequence depends on its order. We now give an example ([E, Example 17.3]) showing that the order of a regular sequence matters. Suppose $A = k[x, y, z]/(x - 1)z$, so $X = \text{Spec } A$ is the union of the $z = 0$ plane and the $x = 1$ plane — X is reduced and has two components (see Figure 9.6). You can readily verify that x is a non-zerodivisor of A ($x = 0$ misses one component of X , and doesn’t vanish entirely on the other), and that the corresponding effective Cartier divisor $X' = \text{Spec } k[x, y, z]/(x, z)$ (on X) is integral. Then $(x - 1)y$ gives an effective Cartier divisor on X' (it doesn’t vanish entirely on X'), so $x, (x - 1)y$ is a regular sequence for A . However, $(x - 1)y$ is *not* a non-zerodivisor of A , as it *does* vanish entirely on one of the two components. Thus $(x - 1)y, x$ is *not* a regular sequence. The reason that reordering the regular sequence $x, (x - 1)y$ ruins regularity is clear: there is a locus on which $(x - 1)y$ isn’t effective Cartier, but it disappears if we enforce $x = 0$ first. The problem is one of “nonlocality” — “near” $x = y = z = 0$ there is no problem. This may motivate the fact that in the (Noetherian) local situation, this problem disappears. We now make this precise.

9.5.6. Theorem. — Suppose (A, \mathfrak{m}) is a Noetherian local ring, and M is a finitely generated A -module. Then any M -regular sequence (x_1, \dots, x_r) in \mathfrak{m} remains a regular sequence upon any reordering.

(In [Di], Dieudonné gives a half-page example showing that Noetherian hypotheses are necessary in Theorem 9.5.6.)

Before proving Theorem 9.5.6 (in Exercise 9.5.F), we prove the first nontrivial case, when $r = 2$. This discussion is secretly a baby case of the Koszul complex.




FIGURE 9.6. Order matters in a regular sequence (in the “non-local” situation), for silly reasons

Suppose $x, y \in \mathfrak{m}$, and x, y is an M -regular sequence. Translation: $x \in \mathfrak{m}$ is a non-zero-divisor on M , and $y \in \mathfrak{m}$ is a non-zero-divisor on M/xM .

Consider the double complex

$$(9.5.6.1) \quad \begin{array}{ccc} M & \xrightarrow{\times(-x)} & M \\ \times y \uparrow & & \uparrow \times y \\ M & \xrightarrow{\times x} & M \end{array}$$

where the bottom left is considered to be in position $(0, 0)$. (The only reason for the minus sign in the top row is solely our arbitrary preference for anti-commuting rather than commuting squares in §1.7.1, but it really doesn't matter.)

We compute the cohomology of the total complex using a (simple) spectral sequence, beginning with the rightward orientation. (The gratuitous use of spectral sequences here, as in many of our other applications, is overkill; we do this partially in order to get practice with the machine.) On the first page, we have

$$\begin{array}{ccc} \ker(M \xrightarrow{\times x} M) & & M/xM \\ \times y \uparrow & & \uparrow \times y \\ \ker(M \xrightarrow{\times x} M) & & M/xM \end{array}$$

The entries $\ker(M \xrightarrow{\times x} M)$ in the first column are 0, as x is a non-zero-divisor on M . Taking homology in the vertical direction to obtain the second page, we find

$$(9.5.6.2) \quad \begin{array}{ccc} 0 & & M/(x, y)M \\ & & \\ 0 & & 0 \end{array}$$

using the fact that y is a non-zero-divisor on M/xM . The sequence clearly converges here. Thus the original double complex (9.5.6.1) only has nonzero cohomology in degree 2, where it is $M/(x, y)M$.

Now we run the spectral sequence on (9.5.6.1) using the upward orientation. The first page of the sequence is:

$$\begin{array}{ccc} M/yM & \xrightarrow{\times(-x)} & M/yM \\ \ker(M \xrightarrow{\times y} M) & \xrightarrow{\times x} & \ker(M \xrightarrow{\times y} M) \end{array}$$

The sequence must converge to (9.5.6.2) after the next step. From the top row, we see that multiplication by x must be injective on M/yM , so x is a non-zerodivisor on M/yM . From the bottom row, multiplication by x gives an isomorphism of $\ker(M \xrightarrow{\times y} M)$ with itself. As $x \in \mathfrak{m}$, by version 2 of Nakayama's Lemma (Lemma 8.2.9), this implies that $\ker(M \xrightarrow{\times y} M) = 0$, so y is a non-zerodivisor on M . Thus we have shown that y, x is a regular sequence on M — the $r = 2$ case of Theorem 9.5.6.

9.5.F. EASY EXERCISE. Prove Theorem 9.5.6. (Hint: show it first in the case of a reordering where only two adjacent x_i are swapped, using the $r = 2$ case just discussed.) Where are the Noetherian hypotheses used?

9.5.7. Regular embeddings.

Suppose $\pi: X \hookrightarrow Y$ is a locally closed embedding. We say that π is a **regular embedding (of codimension r) at a point $p \in X$** if in the local ring $\mathcal{O}_{Y,p}$, the ideal of X is generated by a regular sequence (of length r). We say that π is a **regular embedding (of codimension r)** if it is a regular embedding (of codimension r) at all $p \in X$. (Another reasonable name for a regular embedding might be “local complete intersection”. Unfortunately, “local complete intersection morphism”, or “lci morphism”, is already used for a related notion, see [Stacks, tag 068E].)

Our terminology uses the word “codimension”, which we have not defined as a word on its own. The reason for using this word will become clearer once you meet Krull's Principal Ideal Theorem 12.3.3 and Krull's Height Theorem 12.3.9.

9.5.G. EXERCISE (THE CONDITION OF A LOCALLY CLOSED EMBEDDING BEING A REGULAR EMBEDDING IS OPEN). Show that if a locally closed embedding $\pi: X \hookrightarrow Y$ of locally Noetherian schemes is a regular embedding at p , then it is a regular embedding in some open neighborhood of p in X . Hint: reduce to the case where π is a closed embedding, and then where Y (hence X) is affine — say $Y = \operatorname{Spec} B$, $X = \operatorname{Spec} B/I$, and $p = [p]$ — and there are f_1, \dots, f_r such that in $\mathcal{O}_{Y,p}$, the images of the f_i are a regular sequence generating I_p . We wish to show that $(f_1, \dots, f_r) = I$ “in an open neighborhood of p ”. Prove the following fact in algebra: if I and J are ideals of a Noetherian ring A , and $\mathfrak{p} \subset A$ is a prime ideal such that $I_{\mathfrak{p}} = J_{\mathfrak{p}}$, show that there exists $a \in A \setminus \mathfrak{p}$ such that $I_a = J_a$ in A_a . To do this, show that it suffices to consider the special case $I \subset J$, by considering $I \cap J$ and J instead of I and J . To show this special case, let $K = J/I$, a finitely generated module, and show that if $K_{\mathfrak{p}} = 0$ then $K_a = 0$ for some $a \in A \setminus \mathfrak{p}$.

Hence if X is locally of finite type over a field (e.g., a variety), then to check that a closed embedding π is a regular embedding it suffices to check at closed points of X (Exercise 5.3.F).

Exercise 13.1.F(b) will show that not all closed embeddings are regular embeddings.

9.5.H. EXERCISE. Show that a closed embedding $X \hookrightarrow Y$ of locally Noetherian schemes is a regular embedding of codimension 1 if and only if X is an effective Cartier divisor on Y . Unimportant remark: the Noetherian hypotheses can be replaced by requiring \mathcal{O}_Y to be coherent, and essentially the same argument applies. It is interesting to note that “effective Cartier divisor” implies “regular embedding of codimension 1” always, but that the converse argument requires Noetherian(-like) assumptions. (See [MO129242] for a counterexample to the converse.)

9.5.8. Definition. A **codimension r complete intersection** in a scheme Y is a closed subscheme X that can be written as the scheme-theoretic intersection of r effective Cartier divisors D_1, \dots, D_r , such that at every point $p \in X$, the equations corresponding to D_1, \dots, D_r form a regular sequence. The phrase **complete intersection** means “codimension r complete intersection for some r ”.

CHAPTER 10

Fibered products of schemes, and base change

10.1 They exist

Before we get to products, we note that coproducts exist in the category of schemes: just as with the category of sets (Exercise 1.3.T), coproduct is disjoint union. The next exercise makes this precise (and directly extends to coproducts of an infinite number of schemes).

We will now construct the fibered product in the category of schemes.

10.1.1. Theorem: Fibered products exist. — Suppose $\alpha: X \rightarrow Z$ and $\beta: Y \rightarrow Z$ are morphisms of schemes. Then the fibered product

$$\begin{array}{ccc} X \times_Z Y & \xrightarrow{\alpha'} & Y \\ \beta' \downarrow & & \downarrow \beta \\ X & \xrightarrow{\alpha} & Z \end{array}$$

exists in the category of schemes.

Note: if A is a ring, people often sloppily write \times_A for $\times_{\text{Spec } A}$. If B is an A -algebra, and X is an A -scheme, people often write X_B or $X \times_A B$ for $X \times_{\text{Spec } A} \text{Spec } B$.

10.1.2. Warning: products of schemes aren't products of sets. Before showing existence, here is a warning: the product of schemes isn't a product of sets (and more generally for fibered products). We have made a big deal about schemes being *sets*, endowed with a *topology*, upon which we have a *structure sheaf*. So you might think that we will construct the product in this order: take the product set, endow it with the product topology, and then figure out what the structure sheaf must be. But we won't, because scheme products behave oddly on the level of sets. You may have checked (Exercise 7.6.E(a)) that the product of two affine lines over your favorite algebraically closed field \bar{k} is the affine plane: $\mathbb{A}_{\bar{k}}^1 \times_{\bar{k}} \mathbb{A}_{\bar{k}}^1 \cong \mathbb{A}_{\bar{k}}^2$. But the underlying set of the latter is *not* the underlying set of the former — we get additional points, corresponding to curves in \mathbb{A}^2 that are not lines parallel to the axes!

10.1.3. On the other hand, W -valued points (where W is a scheme, Definition 7.3.10) *do* behave well under (fibered) products (as mentioned in §7.3.11). This is just the universal property *definition* of fibered product: a W -valued point of a scheme X is defined as an element of $\text{Hom}(W, X)$, and the fibered product is defined (via

Yoneda, see §1.3.1) by

$$(10.1.3.1) \quad \mathrm{Hom}(W, X \times_Z Y) = \mathrm{Hom}(W, X) \times_{\mathrm{Hom}(W, Z)} \mathrm{Hom}(W, Y).$$

This is one justification for making the definition of scheme-valued point. For this reason, those classical people preferring to think only about varieties over an algebraically closed field \bar{k} (or more generally, finite type schemes over \bar{k}), and preferring to understand them through their closed points — or equivalently, the \bar{k} -valued points, by the Nullstellensatz (Exercise 5.3.F) — needn't worry: the closed points of the product of two finite type \bar{k} -schemes over \bar{k} are (naturally identified with) the product of the closed points of the factors. This will follow from the fact that the product is also finite type over \bar{k} , which we verify in Exercise 10.2.E. This is one of the reasons that varieties over algebraically closed fields can be easier to work with. But over a nonalgebraically closed field, things become even more interesting; Example 10.2.3 is a first glimpse.

10.1.4. Fancy aside. You may feel that (i) “products of topological spaces are products on the underlying sets” is natural, while (ii) “products of schemes are not necessarily products on the underlying sets” is weird. But really (i) is the lucky consequence of the fact that the underlying set of a topological space can be interpreted as set of p -valued points, where p is a point, so it is best seen as a consequence of paragraph 10.1.3, which is the “more correct” — i.e., more general — fact.

10.1.5. Warning on Noetherianness. The fibered product of Noetherian schemes need not be Noetherian. You will later be able to verify that Exercise 10.2.F gives an example, i.e., that $A := \mathbb{Q} \otimes_{\mathbb{Q}} \mathbb{Q}$ is not Noetherian, as follows. By Exercise 12.1.G(a), $\dim A = 0$. A Noetherian dimension 0 scheme has a finite number of points (Exercise 12.1.C). But by Exercise 10.2.F, $\mathrm{Spec} A$ has an infinite number of points.

On the other hand, the fibered product of finite type k -schemes over finite type k -schemes is a finite type k -scheme (Exercise 10.2.E), so this pathology does not arise for varieties.

10.1.6. Philosophy behind the proof of Theorem 10.1.1. The proof of Theorem 10.1.1 can be confusing. The following comments may help a little.

We already basically know existence of fibered products in two cases: the case where X, Y , and Z are affine (stated explicitly below), and the case where $\beta: Y \rightarrow Z$ is an open embedding (Exercise 8.1.D).

10.1.A. EXERCISE (PROMISED IN REMARK 7.3.5). Use Exercise 7.3.G ($\mathrm{Hom}_{\mathrm{Sch}}(W, \mathrm{Spec} A) = \mathrm{Hom}_{\mathrm{Rings}}(A, \Gamma(W, \mathcal{O}_W))$) to show that given ring maps $C \rightarrow A$ and $C \rightarrow B$,

$$\mathrm{Spec}(A \otimes_C B) \cong \mathrm{Spec} A \times_{\mathrm{Spec} C} \mathrm{Spec} B.$$

(Interpret tensor product as the “fibered coproduct” in the category of rings.) Hence the fibered product of affine schemes exists (in the category of schemes). (This generalizes the fact that the product of affine lines exist, Exercise 7.6.E(a).)

The main theme of the proof of Theorem 10.1.1 is that because schemes are built by gluing affine schemes along open subsets, these two special cases will be all that we need. The argument will repeatedly use the same ideas — roughly,

that schemes glue (Exercise 4.4.A), and that morphisms of schemes glue (Exercise 7.3.B). This is a sign that something more structural is going on; §10.1.7 describes this for experts.

Proof of Theorem 10.1.1. The key idea is this: we cut everything up into affine open sets, do fibered products there, and show that everything glues nicely. The conceptually difficult part of the proof comes from the gluing, and the realization that we have to check almost nothing. We divide the proof up into a number of bite-sized pieces.

Step 1: fibered products of affine with “almost-affine” over affine. We begin by combining the affine case with the open embedding case as follows. Suppose X and Z are affine, and $\beta: Y \rightarrow Z$ factors as $Y \xrightarrow{\iota} Y' \rightarrow Z$ where ι is an open embedding and Y' is affine. Then $X \times_Z Y$ exists. This is because if the two small squares of

$$\begin{array}{ccc} W & \longrightarrow & Y \\ \text{open} \downarrow & & \downarrow \text{open} \\ W' & \longrightarrow & Y' \\ \downarrow & & \downarrow \\ X & \longrightarrow & Z \end{array}$$

are fibered diagrams, then the “outside rectangle” is also a fibered diagram. (This was Exercise 1.3.Q, although you should be able to see this on the spot.) It will be important to remember (from Important Exercise 8.1.D) that “open embeddings” are “preserved by fibered product”: the fact that $Y \rightarrow Y'$ is an open embedding implies that $W \rightarrow W'$ is an open embedding.

Key Step 2: fibered product of affine with arbitrary over affine exists. We now come to the key part of the argument: if X and Z are affine, and Y is arbitrary. This is confusing when you first see it, so we first deal with a special case, when Y is the union of two affine open sets $Y_1 \cup Y_2$. Let $Y_{12} = Y_1 \cap Y_2$.

Now for $i = 1$ and 2 , $X \times_Z Y_i$ exists by the affine case, Exercise 10.1.A. Call this W_i . Also, $X \times_Z Y_{12}$ exists by Step 1 (call it W_{12}), and comes with *canonical* open embeddings into W_1 and W_2 (by construction of fibered products with open embeddings, see the last sentence of Step 1). Thus we can glue W_1 to W_2 along W_{12} ; call this resulting scheme W .

We check that the result is the fibered product by verifying that it satisfies the universal property. Suppose we have maps $\alpha'': V \rightarrow X$, $\beta'': V \rightarrow Y$ that compose (with α and β respectively) to the same map $V \rightarrow Z$. We need to construct a unique map $\gamma: V \rightarrow W$, so that $\alpha' \circ \gamma = \beta''$ and $\beta' \circ \gamma = \alpha''$.

(10.1.6.1)

$$\begin{array}{ccccc} & & V & & \\ & \searrow & \downarrow \alpha'' & \searrow \beta'' & \\ & & W & \xrightarrow{\alpha'} & Y \\ & \searrow \beta' & \downarrow \beta & & \\ & & X & \xrightarrow{\alpha} & Z \end{array}$$

For $i = 1, 2$, define $V_i := (\beta'')^{-1}(Y_i)$. Define $V_{12} := (\beta'')^{-1}(Y_{12}) = V_1 \cap V_2$. Then there is a unique map $V_i \rightarrow W_i$ such that the composed maps $V_i \rightarrow X$ and $V_i \rightarrow Y_i$ are as desired (by the universal property of the fibered product $X \times_Z Y_i = W_i$), hence a unique map $\gamma_i: V_i \rightarrow W$. Similarly, there is a unique map $\gamma_{12}: V_{12} \rightarrow W$ such that the composed maps $V_{12} \rightarrow X$ and $V_{12} \rightarrow Y$ are as desired. But the restriction of γ_i to V_{12} is one such map, so it must be γ_{12} . Thus the maps γ_1 and γ_2 agree on V_{12} , and glue together to a unique map $\gamma: V \rightarrow W$. We have shown existence and uniqueness of the desired γ .

We have thus shown that if Y is the union of two affine open sets, and X and Z are affine, then $X \times_Z Y$ exists.

We now tackle the general case. (You may prefer to first think through the case where “two” is replaced by “three”.) We now cover Y with open sets Y_i , as i runs over some index set (not necessarily finite!). As before, we define W_i and W_{ij} . We can glue these together to produce a scheme W along with open sets we identify with W_i (Exercise 4.4.A — you should check the triple intersection “cocycle” condition).

As in the two-affine case, we show that W is the fibered product by showing that it satisfies the universal property. Suppose we have maps $\alpha'': V \rightarrow X$, $\beta'': V \rightarrow Y$ that compose to the same map $V \rightarrow Z$. We construct a unique map $\gamma: V \rightarrow W$, so that $\alpha' \circ \gamma = \beta''$ and $\beta' \circ \gamma = \alpha''$. Define $V_i = (\beta'')^{-1}(Y_i)$ and $V_{ij} := (\beta'')^{-1}(Y_{ij}) = V_i \cap V_j$. Then there is a unique map $V_i \rightarrow W_i$ such that the composed maps $V_i \rightarrow X$ and $V_i \rightarrow Y_i$ are as desired, hence a unique map $\gamma_i: V_i \rightarrow W$. Similarly, there is a unique map $\gamma_{ij}: V_{ij} \rightarrow W$ such that the composed maps $V_{ij} \rightarrow X$ and $V_{ij} \rightarrow Y$ are as desired. But the restriction of γ_i to V_{ij} is one such map, so it must be γ_{ij} . Thus the maps γ_i and γ_j agree on V_{ij} . Thus the γ_i glue together to a unique map $\gamma: V \rightarrow W$. We have shown existence and uniqueness of the desired γ , completing this step.

Step 3: Z affine, X and Y arbitrary. We next show that if Z is affine, and X and Y are arbitrary schemes, then $X \times_Z Y$ exists. We just follow Step 2, with the roles of X and Y reversed, using the fact that by the previous step, we can assume that the fibered product of an affine scheme with an arbitrary scheme over an affine scheme exists.

Step 4: Z “almost-affine”, X and Y arbitrary. This is akin to Step 1. Let $Z \hookrightarrow Z'$ be an open embedding into an affine scheme. Then $X \times_Z Y$ satisfies the universal property of $X \times_{Z'} Y$. (If you don’t see this directly, one argument is to use Easy Exercise 1.3.X and that open embeddings are monomorphisms, Exercise 8.1.E.)

Step 5: the general case. We employ the same trick yet again. Suppose $\alpha: X \rightarrow Z$, $\beta: Y \rightarrow Z$ are two morphisms of schemes. Cover Z with affine open subschemes Z_i , and let $X_i = \alpha^{-1}(Z_i)$ and $Y_i = \beta^{-1}(Z_i)$. Define $Z_{ij} := Z_i \cap Z_j$, $X_{ij} := \alpha^{-1}(Z_{ij})$, and $Y_{ij} := \beta^{-1}(Z_{ij})$. Then $W_i := X_i \times_{Z_i} Y_i$ exists for all i (Step 3), and $W_{ij} := X_{ij} \times_{Z_{ij}} Y_{ij}$ exists for all i, j (Step 4), and for each i and j , W_{ij} comes with a canonically open embedding into both W_i and W_j (see the last sentence in Step 1). As W_i satisfies the universal property of $X \times_Z Y_i$ (Do you see why?), we may canonically identify W_i (which we know to exist by Step 3) with $X \times_Z Y_i$. Similarly, we identify W_{ij} with $X \times_Z Y_{ij}$.

We then proceed exactly as in Step 2: the W_i ’s can be glued together along the W_{ij} (the cocycle condition can be readily checked to be satisfied), and W can be checked to satisfy the universal property of $X \times_Z Y$ (again, exactly as in Step 2). \square

10.1.7. ★ Describing the existence of fibered products using the high-falutin' language of representable functors. The proof above can be described more cleanly in the language of representable functors (§7.6). This will be enlightening only after you have absorbed the above argument and meditated on it for a long time. It may be most useful to shed light on representable functors, rather than on the existence of the fibered product.

Until the end of §10.1 only, by functor, we mean contravariant functor from the category Sch of schemes to the category of Sets. For each scheme X , we have a functor h_X , taking a scheme Y to the set $\text{Mor}(Y, X)$ (§1.2.20). Recall (§1.3.11, §7.6) that a functor is *representable* if it is naturally isomorphic to some h_X . If a functor is representable, then the representing scheme is unique up to unique isomorphism (Exercise 7.6.C). This can be usefully extended as follows:

10.1.B. EXERCISE (Yoneda's Lemma).

(a) If X and Y are schemes, describe a bijection between morphisms of schemes $X \rightarrow Y$ and natural transformations of functors $h_X \rightarrow h_Y$. Hence show that the category of schemes is a full subcategory (defined in §1.2.15) of the “functor category” (all contravariant functors $Sch \rightarrow \text{Sets}$).

(b) If X is a scheme and F is a contravariant functor $Sch \rightarrow \text{Sets}$, give a bijection between $F(X)$ and the natural transformation of functors $h_X \rightarrow F$. (You are free to check that your bijection is “functorial” in both F and X .)

Hint: this question has nothing to do with schemes; your argument will work in any category. This is the contravariant version of Exercise 1.3.Z(c).

One of Grothendieck's insights is that we should try to treat such functors as “geometric spaces”, without worrying about representability. Many notions carry over to this more general setting without change, and some notions are easier. For example, fibered products of functors always exist: $h \times_{h''} h'$ may be defined by

$$(h \times_{h''} h')(W) = h(W) \times_{h''(W)} h'(W),$$

where the fibered product on the right is a fibered product of sets, which always exists. (This isn't quite enough to define a functor; we have only described where objects go. You should work out where morphisms go too.) We didn't use anything about schemes; this works with Sch replaced by any category.

Then “ $X \times_Z Y$ exists” translates to “ $h_X \times_{h_Z} h_Y$ is representable”.

10.1.8. Representable functors are Zariski sheaves. Because “morphisms to schemes glue” (Exercise 7.3.B), we have a necessary condition for a functor to be representable. We know that if $\{U_i\}$ is an open cover of Y , a morphism $Y \rightarrow X$ is determined by its restrictions $U_i \rightarrow X$, and given morphisms $U_i \rightarrow X$ that agree on the overlap $U_i \cap U_j \rightarrow X$, we can glue them together to get a morphism $Y \rightarrow X$. In the language of equalizer exact sequences (§2.2.7),

$$\cdot \longrightarrow h_X(Y) \longrightarrow \prod h_X(U_i) \rightrightarrows \prod h_X(U_i \cap U_j)$$

is exact. Thus morphisms to X (i.e., the functor h_X) form a sheaf on every scheme Y . If this holds, we say that the functor is a **Zariski sheaf**. (You can impress your friends by telling them that this is a *sheaf on the big Zariski site*.) We can repeat this discussion with Sch replaced by the category Sch_S of schemes over a given base scheme S . We have proved (or observed) that *in order for a functor to be representable, it is necessary for it to be a Zariski sheaf*.

The fibered product passes this test:

10.1.C. EXERCISE. If $X, Y \rightarrow Z$ are schemes, show that $h_X \times_{h_Z} h_Y$ is a Zariski sheaf. (Do not use the fact that $h_X \times_{h_Z} h_Y$ is representable! The point of this section is to recover representability from a more sophisticated perspective.)

We can make some other definitions that extend notions from schemes to functors. We say that a map (i.e., natural transformation) of functors $h' \rightarrow h$ expresses h' as a **subfunctor** of h if for all schemes X , the induced map of sets $h'(X) \rightarrow h(X)$ is an inclusion. We say that a subfunctor $h' \hookrightarrow h$ is an **open subfunctor** if for all schemes X and maps $h_X \rightarrow h$ (i.e., all elements of $h(X)$), the fibered product $h_X \times_h h'$ is representable, by the scheme U say, and $h_U \rightarrow h_X$ corresponds to an open embedding of schemes $U \rightarrow X$. The following Cartesian square may help.

$$\begin{array}{ccc} h_U & \xrightarrow{\text{open}} & h_X \\ \downarrow & & \downarrow \\ h' & \longrightarrow & h \end{array}$$

10.1.D. EXERCISE. Show that a map of representable functors $h_W \rightarrow h_Z$ is an open subfunctor if and only if $W \rightarrow Z$ is an open embedding, so this indeed extends the notion of open embedding to (contravariant) functors ($Sch \rightarrow Sets$).

10.1.E. EXERCISE (THE GEOMETRIC NATURE OF THE NOTION OF “OPEN SUBFUNCTOR”).

- (a) Show that an open subfunctor of an open subfunctor is also an open subfunctor.
- (b) Suppose $h' \rightarrow h$ and $h'' \rightarrow h$ are two open subfunctors of h . Define the intersection of these two open subfunctors, which should also be an open subfunctor of h .
- (c) Suppose U and V are two open subschemes of a scheme X , so $h_U \rightarrow h_X$ and $h_V \rightarrow h_X$ are open subfunctors. Show that the intersection of these two open subfunctors is, as you would expect, $h_{U \cap V}$.

10.1.F. EXERCISE. Suppose $\alpha: X \rightarrow Z$ and $\beta: Y \rightarrow Z$ are morphisms of schemes, and $U \subset X$, $V \subset Y$, $W \subset Z$ are open embeddings, where U and V map to W . Interpret $h_U \times_{h_W} h_V$ as an open subfunctor of $h_X \times_{h_Z} h_Y$. (Hint: given a map $h_T \rightarrow h_X \times_{h_Z} h_Y$, what open subset of T should correspond to $U \times_W V$?)

A collection h_i of open subfunctors of h is said to **cover** h if for *every* map $h_X \rightarrow h$ from a representable subfunctor, the corresponding open subsets $U_i \hookrightarrow X$ cover X .

Given that functors do not have an obvious underlying set (let alone a topology), it is rather amazing that we are talking about when one is an “open subset” of another, or when some functors “cover” another!

10.1.G. EXERCISE. Suppose $\{Z_i\}_i$ is an affine cover of Z , $\{X_{ij}\}_j$ is an affine cover of the preimage of Z_i in X , and $\{Y_{ik}\}_k$ is an affine cover of the preimage of Z_i in Y . Show that $\{h_{X_{ij}} \times_{h_{Z_i}} h_{Y_{ik}}\}_{ijk}$ is an open cover of the functor $h_X \times_{h_Z} h_Y$. (Hint: consider a map $h_T \rightarrow h_X \times_{h_Z} h_Y$, and extend your solution to Exercise 10.1.F.)

We now come to a key point: a Zariski sheaf that is “locally representable” must be representable:

10.1.H. KEY EXERCISE. If a functor h is a Zariski sheaf that has an open cover by representable functors (“is covered by schemes”), then h is representable. (Hint: use Exercise 4.4.A to glue together the schemes representing the open subfunctors.)

This immediately leads to the existence of fibered products as follows. Exercise 10.1.C shows that $h_X \times_{h_Z} h_Y$ is a Zariski sheaf. But $h_{X_{ij}} \times_{h_{Z_i}} h_{Y_{ik}}$ is representable for each i, j, k (fibered products of affines over an affine exist, Exercise 10.1.A), and these functors are an open cover of $h_X \times_{h_Z} h_Y$ by Exercise 10.1.G, so by Key Exercise 10.1.H we are done.

10.2 Computing fibered products in practice

Before giving some examples, we first see how to compute fibered products in practice. There are four types of morphisms (1)–(4) that it is particularly easy to take fibered products with, and all morphisms can be built from these atomic components. More precisely, (1) will imply that we can compute fibered products locally on the source and target. Thus to understand fibered products in general, it suffices to understand them on the level of affine sets, i.e., to be able to compute $A \otimes_B C$ given ring maps $B \rightarrow A$ and $B \rightarrow C$. Any map $B \rightarrow A$ (and similarly $B \rightarrow C$) may be expressed as $B \rightarrow B[t_1, \dots]/I$, so if we know how to base change by “adding variables” (2) and “taking quotients” (3), we can “compute” any fibered product (at least in theory). The fourth type of morphism (4), corresponding to localization, is useful to understand explicitly as well.

(1) Base change by open embeddings.

We have already done this (Exercise 8.1.D), and we used it in the proof that fibered products of schemes exist.

(2) Adding an extra variable.

10.2.A. EASY ALGEBRA EXERCISE. Show that $A \otimes_B B[t] \cong A[t]$, so the following is a fibered diagram. (Your argument might naturally extend to allow the addition of infinitely many variables, but we won’t need this generality.) Hint: show that $A[t]$ satisfies an appropriate universal property.

$$\begin{array}{ccc} \mathrm{Spec} A[t] & \longrightarrow & \mathrm{Spec} B[t] \\ \downarrow & & \downarrow \\ \mathrm{Spec} A & \longrightarrow & \mathrm{Spec} B \end{array}$$

10.2.1. Definition (affine space over an arbitrary scheme). If X is any scheme, we define \mathbb{A}_X^n as $X \times_{\mathbb{Z}} \mathrm{Spec} \mathbb{Z}[x_1, \dots, x_n]$. Clearly $\mathbb{A}_{\mathrm{Spec} A}^n$ is canonically the same as \mathbb{A}_A^n as defined in Example 8 of §3.2.4.

(3) Base change by closed embeddings.

10.2.B. EXERCISE. Suppose $\phi: B \rightarrow A$ is a ring morphism, and $I \subset B$ is an ideal. Let $I^e := \langle \phi(i) \rangle_{i \in I} \subset A$ be the **extension of I to A** . Describe a natural isomorphism $A/I^e \xrightarrow{\sim} A \otimes_B (B/I)$. (Hint: consider $I \rightarrow B \rightarrow B/I \rightarrow 0$, and use the right-exactness of $\otimes_B A$, Exercise 1.3.H.)

10.2.2. As an immediate consequence: the fibered product with a closed subscheme is a closed subscheme of the fibered product in the obvious way. We say that “closed embeddings are preserved by base change”.

10.2.C. EXERCISE.

- (a) Interpret the intersection of two closed embeddings into X (cf. Exercise 9.1.I) as their fibered product over X .
- (b) Show that “locally closed embeddings” are preserved by base change.
- (c) Define the **intersection of n locally closed embeddings** $X_i \hookrightarrow Z$ ($1 \leq i \leq n$) by the fibered product of the X_i over Z (mapping to Z). Show that the intersection of (a finite number of) locally closed embeddings is also a locally closed embedding.

10.2.D. EXERCISE. Suppose $I \subset A[x_1, \dots, x_m]$ and $J \subset A[y_1, \dots, y_n]$ are ideals. Prove the isomorphism

$$A[x_1, \dots, x_m]/I \otimes_A A[y_1, \dots, y_n]/J \xrightarrow{\sim} A[x_1, \dots, x_m, y_1, \dots, y_n]/(I, J).$$

Hint: Exercises 10.2.A and 10.2.B.

As an application, we can compute tensor products of finitely generated k -algebras over k . For example, we have

$$k[x_1, x_2]/(x_1^2 - x_2) \otimes_k k[y_1, y_2]/(y_1^3 + y_2^3) \cong k[x_1, x_2, y_1, y_2]/(x_1^2 - x_2, y_1^3 + y_2^3).$$

10.2.E. EXERCISE. Suppose X and Y are locally of finite type A -schemes. Show that $X \times_A Y$ is also locally of finite type over A . Prove the same thing with “locally” removed from both the hypothesis and conclusion.

10.2.3. Example. We can use these ideas to compute $\mathbb{C} \otimes_{\mathbb{R}} \mathbb{C}$:

$$\begin{aligned} \mathbb{C} \otimes_{\mathbb{R}} \mathbb{C} &\cong \mathbb{C} \otimes_{\mathbb{R}} (\mathbb{R}[x]/(x^2 + 1)) \\ &\cong (\mathbb{C} \otimes_{\mathbb{R}} \mathbb{R}[x])/(x^2 + 1) && \text{by §10.2(3)} \\ &\cong \mathbb{C}[x]/(x^2 + 1) && \text{by §10.2(2)} \\ &\cong \mathbb{C}[x]/((x - i)(x + i)) \\ &\cong \mathbb{C}[x]/(x - i) \times \mathbb{C}[x]/(x + i) && \text{by the Chinese Remainder Theorem} \\ &\cong \mathbb{C} \times \mathbb{C} \end{aligned}$$

Thus $\text{Spec } \mathbb{C} \times_{\mathbb{R}} \text{Spec } \mathbb{C} \cong \text{Spec } \mathbb{C} \amalg \text{Spec } \mathbb{C}$. This example is the first example of many different behaviors. Notice for example that two points somehow correspond to the Galois group of \mathbb{C} over \mathbb{R} ; for one of them, x (the “ i ” in one of the copies of \mathbb{C}) equals i (the “ i ” in the other copy of \mathbb{C}), and in the other, $x = -i$.

10.2.4. ★ Remark. Here is a clue that there is something deep going on behind Example 10.2.3. If L/K is a (finite) Galois extension with Galois group G , then $L \otimes_K L$ is isomorphic to L^G (the product of $|G|$ copies of L). This turns out to be a restatement of the classical form of linear independence of characters! In the

language of schemes, $\text{Spec } L \times_K \text{Spec } L$ is a union of a number of copies of $\text{Spec } L$ that naturally form a torsor over the Galois group G ; but we will not define torsor here.

10.2.F. ★ HARD BUT FASCINATING EXERCISE FOR THOSE FAMILIAR WITH $\text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$. Show that the points of $\text{Spec } \overline{\mathbb{Q}} \otimes_{\mathbb{Q}} \overline{\mathbb{Q}}$ are in natural bijection with $\text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$, and the Zariski topology on the former agrees with the profinite topology on the latter. (Some hints: first do the case of finite Galois extensions. Relate the topology on Spec of a direct limit of rings to the inverse limit of Spec s. Can you see which point corresponds to the identity of the Galois group?)

At this point, we can compute any $A \otimes_B C$ (where A and C are B -algebras): any map of rings $\phi: B \rightarrow A$ can be interpreted by adding variables (perhaps infinitely many) to B , and then imposing relations. But in practice (4) is useful, as we will see in examples.

(4) Base change of affine schemes by localization.

10.2.G. EXERCISE. Suppose $\phi: B \rightarrow A$ is a ring morphism, and $S \subset B$ is a multiplicative subset of B , which implies that $\phi(S)$ is a multiplicative subset of A . Describe a natural isomorphism $\phi(S)^{-1}A \xrightarrow{\sim} A \otimes_B (S^{-1}B)$.

Informal translation: “the fibered product with a localization is the localization of the fibered product in the obvious way”. We say that “localizations are preserved by base change”. This is handy if the localization is of the form $B \hookrightarrow B_f$ (corresponding to taking distinguished open sets) or (if B is an integral domain) $B \hookrightarrow K(B)$ (corresponding to taking the generic point), and various things in between.

10.2.5. Examples. These four facts let you calculate lots of things in practice, and we will use them freely.

10.2.H. EXERCISE: THE THREE IMPORTANT TYPES OF MONOMORPHISMS OF SCHEMES. Show that the following are monomorphisms (Definition 1.3.10): open embeddings, closed embeddings, and localization of affine schemes. As monomorphisms are closed under composition (Exercise 1.3.V) compositions of the above are also monomorphisms — for example, locally closed embeddings, or maps from “ Spec of stalks at points of X ” to X . (Caution: if p is a point of a scheme X , the natural morphism $\text{Spec } \mathcal{O}_{X,p} \rightarrow X$, cf. Exercise 7.3.M, is a monomorphism but is not in general an open embedding.)

10.2.I. EXERCISE. Recall that $\mathbb{A}_{\mathbb{A}}^n \cong \mathbb{A}_{\mathbb{Z}}^n \times_{\text{Spec } \mathbb{Z}} \text{Spec } \mathbb{A}$ (§10.2.1). Prove similarly that $\mathbb{P}_{\mathbb{A}}^n \cong \mathbb{P}_{\mathbb{Z}}^n \times_{\text{Spec } \mathbb{Z}} \text{Spec } \mathbb{A}$. Thus affine space and projective space are pulled back from their “universal manifestation” over the final object $\text{Spec } \mathbb{Z}$.

10.2.6. Extending the base field. One special case of base change is called **extending the base field**: if X is a k -scheme, and ℓ is a field extension (often ℓ is the algebraic closure of k), then $X \times_{\text{Spec } k} \text{Spec } \ell$ (sometimes informally written $X \times_k \ell$ or X_{ℓ}) is an ℓ -scheme. Often properties of X can be checked by verifying them instead on X_{ℓ} . This is the subject of *descent* — certain properties “descend” from X_{ℓ} to X . We have

already seen that the property of being the Spec of a normal integral domain descends in this way (Exercise 5.4.M). Exercises 10.2.J and 10.2.K give other examples of properties which descend: the property of two morphisms being equal, and the property of a(n affine) morphism being a closed embedding, both descend in this way. Those interested in schemes over non-algebraically closed fields will use this repeatedly, to reduce results to the algebraically closed case.

10.2.J. EXERCISE. Suppose $\pi: X \rightarrow Y$ and $\rho: X \rightarrow Y$ are morphisms of k -schemes, ℓ/k is a field extension, and $\pi_\ell: X \times_{\text{Spec } k} \text{Spec } \ell \rightarrow Y \times_{\text{Spec } k} \text{Spec } \ell$ and $\rho_\ell: X \times_{\text{Spec } k} \text{Spec } \ell \rightarrow Y \times_{\text{Spec } k} \text{Spec } \ell$ are the induced maps of ℓ -schemes. (Be sure you understand what this means!) Show that if $\pi_\ell = \rho_\ell$ then $\pi = \rho$. (Hint: show that π and ρ are the same on the level of sets. To do this, you may use that $X \times_{\text{Spec } k} \text{Spec } \ell \rightarrow X$ is surjective, which we will soon prove in Exercise 10.4.F. Then reduce to the case where X and Y are affine.)

10.2.K. EASY EXERCISE. Suppose $\pi: X \rightarrow Y$ is an affine morphism over k , and ℓ/k is a field extension. Show that π is a closed embedding if and only if $\pi \times_k \ell: X \times_k \ell \rightarrow Y \times_k \ell$ is. (The affine hypothesis is not necessary for this result, but it makes the proof easier, and this is the situation in which we will most need it.)

10.2.L. EXERCISE (SEEMINGLY PATHOLOGICAL BEHAVIOR IN NONPATHOLOGICAL CIRCUMSTANCES). Suppose k is a field, and $A = k(x) \otimes_k k(y)$. Show that A is a (nonzero) localization of $k[x, y]$, hence an integral domain. Thus (0) is the unique minimal prime ideal of A . Show that the remaining prime ideals of A correspond to ideals $(f(x, y)) \subset k[x, y]$, where $f(x, y)$ is an irreducible polynomial in $k[x, y]$ containing *both* the variables x and y . (Hint: (2) and (4) above.)

This example will come up again in Remarks 12.2.16 and 25.5.10. The idea in the solution to Exercise 10.2.L also yields the following.

10.2.M. UNIMPORTANT BUT FUN EXERCISE. Show that $\text{Spec } \mathbb{Q}(t) \otimes_{\mathbb{Q}} \mathbb{C}$ has closed points in natural correspondence with the transcendental complex numbers. This scheme doesn't come up in nature, but it is certainly neat!

10.3 Interpretations: Pulling back families, and fibers of morphisms

10.3.1. Pulling back families.

Before making any definitions, we give a motivating informal example. Consider the “family of curves”

$$y^2 = x^3 + tx$$

in the xy -plane parametrized by t . Translation: consider

$$\text{Spec } k[x, y, t]/(y^2 - x^3 - tx) \longrightarrow \text{Spec } k[t].$$

If we “pull back this family” to the uv -plane via $uv = t$, we get the family

$$y^2 = x^3 + uvx.$$

If instead we set t to 3, we get the curve

$$y^2 = x^3 + 3x,$$

which we interpret as the fiber of the original family above $t = 3$. You may have noticed the fiber products underlying in these constructions:

$$\begin{array}{ccc} \mathrm{Spec} k[x, y, u, v]/(y^2 - x^3 - uvx) & \longrightarrow & \mathrm{Spec} k[x, y, t]/(y^2 - x^3 - tx) \\ \downarrow & & \downarrow \\ \mathrm{Spec} k[u, v] & \xrightarrow{uv \mapsto t} & \mathrm{Spec} k[t] \end{array}$$

and

$$\begin{array}{ccc} \mathrm{Spec} k[x, y]/(y^2 - x^3 - 3x) & \longrightarrow & \mathrm{Spec} k[x, y, t]/(y^2 - x^3 - tx) \\ \downarrow & & \downarrow \\ \mathrm{Spec} k & \xrightarrow{3 \mapsto t} & \mathrm{Spec} k[t] \end{array}$$

We now formalize this.




FIGURE 10.1. A picture of a pulled back family

Suppose $Y \rightarrow Z$ is a morphism. We interpret this as a “family of schemes parametrized by a **base scheme** (or just plain **base**) Z .” Then if we have another morphism $\psi: X \rightarrow Z$, we interpret the induced map $X \times_Z Y \rightarrow X$ as the “pulled back family” (see Figure 10.1).

$$\begin{array}{ccc} X \times_Z Y & \longrightarrow & Y \\ \text{pulled back family} \downarrow & & \downarrow \text{family} \\ X & \xrightarrow{\psi} & Z \end{array}$$

We sometimes say that $X \times_Z Y$ is the **scheme-theoretic preimage** (or *scheme-theoretic pullback*, or *scheme-theoretic inverse image*, or *inverse image scheme*) of Y . (Our forthcoming discussion of fibers may give some motivation for this.) For this reason, fibered product is often called **base change** or **change of base** or **pullback**. In addition to the various names for a Cartesian diagram given in §1.3.7, in algebraic geometry it is often called a **base change diagram** or a **pullback diagram**, and $X \times_Z Y \rightarrow X$ is called the **pullback** of $Y \rightarrow Z$ by ψ , and $X \times_Z Y$ is called the **pullback** of Y by ψ . One often uses the phrase “over X ” or “above X ” when discussing $X \times_Z Y$, especially if X is a locally closed subscheme of Z . (Random side remark: scheme-theoretic preimage always makes sense, while the notion of scheme-theoretic image is somehow problematic, as discussed in §9.4.1.)

10.3.2. Fibers of morphisms.

(If you did Exercise 8.3.J, that finite morphisms have finite fibers, you will not find the following discussion surprising.) A special case of pullback is the notion of a fiber of a morphism. We motivate this with the notion of fiber in the category of topological spaces.

10.3.A. EXERCISE. Show that if $Y \rightarrow Z$ is a continuous map of topological spaces, and X is a point p of Z , then the fiber of Y over p (the set-theoretic fiber, with the induced topology) is naturally identified with $X \times_Z Y$.


More generally, for any $\pi: X \rightarrow Z$, the fiber of $X \times_Z Y \rightarrow X$ over a point p of X is naturally identified with the fiber of $Y \rightarrow Z$ over $\pi(p)$.

10.3.3. Definition. Motivated by topology, we return to the category of schemes. Suppose $p \rightarrow Z$ is the inclusion of a point (not necessarily closed). More precisely, if p is a point with residue field K , consider the map $\text{Spec } K \rightarrow Z$ sending $\text{Spec } K$ to p , with the natural isomorphism of residue fields. Then if $g: Y \rightarrow Z$ is any morphism, the scheme-theoretic preimage of p is called the **scheme-theoretic fiber** (or *fibre*) of g **above** p , and is denoted $g^{-1}(p)$. If Z is irreducible, the scheme-theoretic fiber above the generic point of Z is called the **generic fiber** (of g). In an affine open subscheme $\text{Spec } A$ containing p , p corresponds to some prime ideal \mathfrak{p} , and the morphism $\text{Spec } K \rightarrow Z$ corresponds to the ring map $A \rightarrow A_{\mathfrak{p}}/\mathfrak{p}A_{\mathfrak{p}}$. This is the composition of localization and closed embedding, and thus can be computed by the tricks above. (Note that $p \rightarrow Z$ is a monomorphism, by Exercise 10.2.H.)

10.3.B. EXERCISE. Show that the underlying topological space of the scheme-theoretic fiber of $X \rightarrow Y$ above a point p is naturally identified with the topological fiber of $X \rightarrow Y$ above p .

10.3.C. EXERCISE (ANALOG OF EXERCISE 10.3.A). Suppose that $\pi: Y \rightarrow Z$ and $\tau: X \rightarrow Z$ are morphisms, and $p \in X$ is a point. Show that the fiber of $X \times_Z Y \rightarrow X$ over p is (isomorphic to) the base change to p of the fiber of $\pi: Y \rightarrow Z$ over $\tau(p)$.

10.3.4. Example (enlightening in several ways). Consider the projection of the parabola $y^2 = x$ to the x -axis over \mathbb{Q} , corresponding to the map of rings $\mathbb{Q}[x] \rightarrow \mathbb{Q}[y]$, with $x \mapsto y^2$. If \mathbb{Q} alarms you, replace it with your favorite field and see what happens. (You should look at Figure 10.2, which is a flipped version of the

FIGURE 10.2. The map $\mathbb{C} \rightarrow \mathbb{C}$ given by $y \mapsto y^2$

parabola of Figure 3.6, and figure out how to edit it to reflect what we glean here.) Writing $\mathbb{Q}[y]$ as $\mathbb{Q}[x, y]/(y^2 - x)$ helps us interpret the morphism conveniently.

(i) Then the preimage of 1 is two points:

$$\begin{aligned} \operatorname{Spec} \mathbb{Q}[x, y]/(y^2 - x) \otimes_{\mathbb{Q}[x]} \mathbb{Q}[x]/(x - 1) &\cong \operatorname{Spec} \mathbb{Q}[x, y]/(y^2 - x, x - 1) \\ &\cong \operatorname{Spec} \mathbb{Q}[y]/(y^2 - 1) \\ &\cong \operatorname{Spec} \mathbb{Q}[y]/(y - 1) \coprod \operatorname{Spec} \mathbb{Q}[y]/(y + 1). \end{aligned}$$

(ii) The preimage of 0 is one nonreduced point:

$$\operatorname{Spec} \mathbb{Q}[x, y]/(y^2 - x, x) \cong \operatorname{Spec} \mathbb{Q}[y]/(y^2).$$

(iii) The preimage of -1 is one reduced point, but of “size 2 over the base field”.

$$\begin{aligned} \operatorname{Spec} \mathbb{Q}[x, y]/(y^2 - x, x + 1) &\cong \operatorname{Spec} \mathbb{Q}[y]/(y^2 + 1) \\ &\cong \operatorname{Spec} \mathbb{Q}[i] \\ &= \operatorname{Spec} \mathbb{Q}(i). \end{aligned}$$

(iv) The preimage of the generic point is again one reduced point, but of “size 2 over the residue field”, as we verify now.

$$\operatorname{Spec} \mathbb{Q}[x, y]/(y^2 - x) \otimes_{\mathbb{Q}[x]} \mathbb{Q}(x) \cong \operatorname{Spec} \mathbb{Q}[y] \otimes_{\mathbb{Q}[y^2]} \mathbb{Q}(y^2)$$

i.e., (informally) the Spec of the ring of polynomials in y divided by polynomials in y^2 . A little thought shows you that in this ring you may invert *any* polynomial in y , as if $f(y)$ is any polynomial in y , then

$$\frac{1}{f(y)} = \frac{f(-y)}{f(y)f(-y)},$$

and the latter denominator is a polynomial in y^2 . Thus

$$\mathbb{Q}[x, y]/(y^2 - x) \otimes_{\mathbb{Q}[x]} \mathbb{Q}(x) \cong \mathbb{Q}(y)$$

which is a degree 2 field extension of $\mathbb{Q}(x)$ (note that $\mathbb{Q}(x) = \mathbb{Q}(y^2)$).

(You might want to work through the preimages of other points of $\mathbb{A}_{\mathbb{Q}}^1$, such as $[(x^3 - 4)]$ and $[(x^2 - 2)]$.)

Notice the following interesting fact: in each of the four cases, the number of preimages can be interpreted as 2, where you count to two in several ways: you can count points (as in the case of the preimage of 1); you can get nonreduced

behavior (as in the case of the preimage of 0); or you can have a field extension of degree 2 (as in the case of the preimage of -1 or the generic point). In each case, the fiber is an affine scheme whose dimension as a vector space over the residue field of the point is 2. Number theoretic readers may have seen this behavior before. We will discuss this example again in §18.4.9. This is going to be symptomatic of a very important kind of morphism (a finite flat morphism, see Remark 25.4.8 and §25.4.12).

Try to draw a picture of this morphism if you can, so you can develop a pictorial shorthand for what is going on. A good first approximation is the parabola of Figure 10.2, but you will want to somehow depict the peculiarities of (iii) and (iv).

10.3.5. Remark: *Finite morphisms have finite fibers.* If you haven't done Exercise 8.3.J, that finite morphisms have finite fibers, now would be a good time to do it, as you will find it more straightforward given what you know now.

10.3.D. EXERCISE (IMPORTANT FOR THOSE WITH MORE ARITHMETIC BACKGROUND). What is the scheme-theoretic fiber of $\text{Spec } \mathbb{Z}[i] \rightarrow \text{Spec } \mathbb{Z}$ over the prime (p) ? Your answer will depend on p , and there are four cases, corresponding to the four cases of Example 10.3.4. (Can you draw a picture?)

10.3.E. EXERCISE. (This exercise will give you practice in computing a fibered product over something that is not a field.) Consider the morphism of schemes $X = \text{Spec } k[t] \rightarrow Y = \text{Spec } k[u]$ corresponding to $k[u] \rightarrow k[t]$, $u \mapsto t^2$, where $\text{char } k \neq 2$. Show that $X \times_Y X$ has two irreducible components. (What happens if $\text{char } k = 2$? See Exercise 10.5.A for a clue.)

10.3.6. A first view of a blow-up.

10.3.F. IMPORTANT CONCRETE EXERCISE. (The discussion here immediately generalizes to \mathbb{A}_k^n .) Define a closed subscheme $\text{Bl}_{(0,0)} \mathbb{A}_k^2$ of $\mathbb{A}_k^2 \times_k \mathbb{P}_k^1$ as follows (see Figure 10.3). If the coordinates on \mathbb{A}_k^2 are x, y , and the projective coordinates on \mathbb{P}_k^1 are u, v , this subscheme is cut out in $\mathbb{A}_k^2 \times_k \mathbb{P}_k^1$ by the single equation $xv = yu$. (You may wish to interpret $\text{Bl}_{(0,0)} \mathbb{A}_k^2$ as follows. The \mathbb{P}_k^1 parametrizes lines through the origin. The blow-up corresponds to ordered pairs of (point p , line ℓ) such that $(0,0)$ and p both lie on ℓ .) Describe the fiber of the morphism $\text{Bl}_{(0,0)} \mathbb{A}_k^2 \rightarrow \mathbb{P}_k^1$ over each closed point of \mathbb{P}_k^1 . Show that the morphism $\text{Bl}_{(0,0)} \mathbb{A}_k^2 \rightarrow \mathbb{A}_k^2$ is an isomorphism away from $(0,0) \in \mathbb{A}_k^2$. Show that the fiber over $(0,0)$ is an effective Cartier divisor (§9.5.1, a closed subscheme that is locally cut out by a single equation, which is not a zerodivisor). It is called the **exceptional divisor**. We will discuss blow-ups in Chapter 23. This particular example will come up in the motivating example of §23.1, and in Exercise 21.2.D.

We haven't yet discussed regularity, but here is a hand-waving argument suggesting that the $\text{Bl}_{(0,0)} \mathbb{A}_k^2$ is "smooth": the preimage above either standard open set $U_i \subset \mathbb{P}^1$ is isomorphic to \mathbb{A}^2 . Thus "the blow-up is a surgery that takes the smooth surface \mathbb{A}_k^2 , cuts out a point, and glues back in a \mathbb{P}^1 , in such a way that the outcome is another smooth surface."

10.3.7. General fibers, generic fibers, generically finite morphisms.




FIGURE 10.3. A first example of a blow-up

The phrases “generic fiber” and “general fiber” parallel the phrases “generic point” and “general point” (Definition 3.6.11). Suppose $\pi: X \rightarrow Y$ is a morphism of schemes. When one says the **general fiber** (or *a* general fiber) of π has a certain property, this means that there exists a dense open subset $U \subset Y$ such that the fibers above any point in U have that property.

When one says *the generic fiber* of $\pi: X \rightarrow Y$ (Definition 10.3.3), this implicitly means that Y is irreducible, and the phrase refers to the fiber over the generic point. *General fiber* and *generic fiber* are not the same thing! Clearly if something holds for the general fiber, then it holds for the generic fiber, but the converse is not always true. However, in good circumstances, it can be — properties of the generic fiber extend to an honest open neighborhood. For example, if Y is irreducible and Noetherian, and π is finite type, then if the generic fiber of π is empty (resp. nonempty), then the general fiber is empty (resp. nonempty), by Chevalley’s Theorem 8.4.2.

If $\pi: X \rightarrow Y$ is finite type, we say π is **generically finite** if π is finite after base change to the generic point of each irreducible component (or equivalently, by Exercise 8.4.D, if the preimage of the generic point of each irreducible component of Y is a finite set). (The notion of generic finiteness can be defined in more general circumstances, see [Stacks, tag 073A].)

10.3.G. EXERCISE (“GENERICALLY FINITE” USUALLY MEANS “GENERALLY FINITE”). Suppose $\pi: X \rightarrow Y$ is an affine, finite type, generically finite morphism of locally Noetherian schemes, and Y is reduced. Show that there is an open neighborhood of each generic point of Y over which π is actually finite. (The hypotheses can be weakened considerably, see [Stacks, tag 02NW].) Hint: reduce to the case where Y is $\text{Spec } B$, where B is an integral domain. Then X is affine, say $X = \text{Spec } A$. Write $A = B[x_1, \dots, x_n]/I$. Now $A \otimes_B K(B)$ is a finite $K(B)$ -module (finite-dimensional vector space) by hypothesis, so there are monic polynomials $f_i(t) \in K(B)[t]$ such that $f_i(x_i) = 0$ in $A \otimes_B K(B)$. Let b be the product of the (finite number of) denominators appearing in the coefficients in the $f_i(x)$. By replacing B by B_b , argue that you can assume that $f_i(t) \in B[t]$. Then $f_i(x_i) = 0$ in $A \otimes_B K(B)$, meaning that $f_i(x_i)$ is annihilated by some nonzero element of B . By replacing B by its localization at the product of these n nonzero elements (“shrinking $\text{Spec } B$ further”), argue that $f_i(x_i) = 0$ in A . Then conclude.

10.3.8. ★★ Finitely presented families (morphisms) are locally pullbacks of particularly nice families. If you are macho and are embarrassed by Noetherian rings, the following exercise can be used to extend results from the Noetherian case to locally finitely presented situations. Exercise 10.3.I, an extension of Chevalley's Theorem 8.4.2, is a good example.

10.3.H. EXERCISE. Suppose $\pi: X \rightarrow \operatorname{Spec} B$ is a finitely presented morphism. Show that there exists a base change diagram of the form

$$\begin{array}{ccc} X & \xrightarrow{\quad} & X' \\ \pi \downarrow & & \downarrow \pi' \\ \operatorname{Spec} B & \xrightarrow{\quad \rho \quad} & \operatorname{Spec} \mathbb{Z}[x_1, \dots, x_N] \end{array}$$

where N is some integer, and π' is finitely presented (= finite type as the target is Noetherian, see §8.3.13). Thus each finitely presented morphism is locally (on the base) a pullback of a finite type morphism to a Noetherian scheme. Hence any result proved for Noetherian schemes and stable under base change is automatically proved for locally finitely presented morphisms to arbitrary schemes. Hint: think about the case where X is affine first. If $X = \operatorname{Spec} A$, then $A = B[y_1, \dots, y_n]/(f_1, \dots, f_r)$. Choose one variable x_i for each coefficient of $f_i \in B[y_1, \dots, y_n]$. What is X' in this case? Then consider the case where X is the union of two affine open sets, that intersect in an affine open set. Then consider more general cases until you solve the full problem. You will need to use every part of the definition of finite presentation. (Exercise 28.2.E extends this result.)

10.3.I. EXERCISE (CHEVALLEY'S THEOREM FOR LOCALLY FINITELY PRESENTED MORPHISMS).

(a) Suppose that A is a finitely presented B -algebra (B not necessarily Noetherian), so $A = B[x_1, \dots, x_n]/(f_1, \dots, f_r)$. Show that the image of $\operatorname{Spec} A \rightarrow \operatorname{Spec} B$ is a finite union of locally closed subsets of $\operatorname{Spec} B$. Hint: Exercise 10.3.H (the simpler affine case).

(b) Show that if $\pi: X \rightarrow Y$ is a quasicompact locally finitely presented morphism, and Y is quasicompact, then $\pi(X)$ is a finite union of locally closed subsets. (For hardened experts only: [Gr-EGA, 0_{III}.9.1] gives a definition of *local constructibility*, and of constructibility in more generality. The general form of Chevalley's Constructibility Theorem [Gr-EGA, IV₁.1.8.4] is that the image of a locally constructible subset, under a finitely presented map, is also locally constructible.)

10.4 Properties preserved by base change

All “reasonable” properties of morphisms are preserved under base change (cf. §8.1(ii)). We discuss this, and in §10.5.1 we will explain how to fix those that don't fit this pattern.

We have already shown that the notion of “open embedding” is preserved by base change (Exercise 8.1.D). We did this by explicitly describing what the fibered product of an open embedding is: if $Y \hookrightarrow Z$ is an open embedding, and $\psi: X \rightarrow Z$

is any morphism, then we checked that the open subscheme $\psi^{-1}(Y)$ of X satisfies the universal property of fibered products.

We have also shown that the notion of “closed embedding” is preserved by base change (§10.2 (3)). In other words, given a Cartesian diagram

$$\begin{array}{ccc} W & \longrightarrow & Y \\ \downarrow & & \downarrow \text{cl. emb.} \\ X & \longrightarrow & Z \end{array}$$

where $Y \hookrightarrow Z$ is a closed embedding, $W \rightarrow X$ is as well.

10.4.A. EASY EXERCISE. Show that the closed embeddings and locally closed embeddings are both “reasonable” classes of morphisms in the sense of §8.1.

10.4.B. EASY EXERCISE. Show that locally principal closed subschemes (Definition 9.5.1) pull back to locally principal closed subschemes.

Similarly, other important properties are preserved by base change.

10.4.C. EXERCISE. Show that the following properties of morphisms are preserved by base change.

- (a) quasicompact
- (b) affine morphism
- (c) finite
- (d) integral
- (e) locally of finite type
- (f) finite type
- ★ (g) locally of finite presentation

10.4.D. EXERCISE. Show that the properties of morphisms of the previous exercise are all “reasonable” in the sense of §8.1.

10.4.E. ★ EXERCISE. Show that the notion of “quasifinite morphism” (finite type + finite fibers, Definition 8.3.11) is preserved by base change. (Warning: the notion of “finite fibers” is not preserved by base change. $\text{Spec } \overline{\mathbb{Q}} \rightarrow \text{Spec } \mathbb{Q}$ has finite fibers, but $\text{Spec } \overline{\mathbb{Q}} \otimes_{\mathbb{Q}} \overline{\mathbb{Q}} \rightarrow \text{Spec } \mathbb{Q}$ has one point for each element of $\text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$, see Exercise 10.2.F.) Hint: reduce to the case $\text{Spec } A \rightarrow \text{Spec } B$. Reduce to the case $\phi: \text{Spec } A \rightarrow \text{Spec } k$. By Exercise 8.4.D, such ϕ are actually finite, and finiteness is preserved by base change.

10.4.F. EXERCISE. (a) Show that surjectivity is preserved by base change. (**Surjectivity** has its usual meaning: surjective as a map of sets.) You may end up showing that for any fields k_1 and k_2 containing k_3 , $k_1 \otimes_{k_3} k_2$ is nonzero, and using the Axiom of Choice to find a maximal ideal in $k_1 \otimes_{k_3} k_2$. (b) Show that surjective morphisms form a “reasonable” class (§8.1).

10.4.1. On the other hand, injectivity is not preserved by base change — witness the bijection $\text{Spec } \mathbb{C} \rightarrow \text{Spec } \mathbb{R}$, which loses injectivity upon base change by $\text{Spec } \mathbb{C} \rightarrow \text{Spec } \mathbb{R}$ (see Example 10.2.3). This can be rectified (see §10.5.I).

10.4.G. EXERCISE (CF. EXERCISE 10.2.E). Suppose X and Y are integral finite type \bar{k} -schemes. Show that $X \times_{\bar{k}} Y$ is an integral finite type \bar{k} -scheme. (Once we define “variety”, this will become the important fact that the product of irreducible varieties over an algebraically closed field is an irreducible variety, Exercise 11.3.K. The fact that the base field \bar{k} is algebraically closed is important, see §10.5.1. See Exercise 10.5.O for an improvement.) Hint: reduce to the case where X and Y are both affine, say $X = \operatorname{Spec} A$ and $Y = \operatorname{Spec} B$ with A and B integral domains. You might flip ahead to Easy Exercise 10.5.N to see how to do this. Suppose $(\sum a_i \otimes b_i)(\sum a'_j \otimes b'_j) = 0$ in $A \otimes_{\bar{k}} B$ with $a_i, a'_j \in A$, $b_i, b'_j \in B$, where both $\{b_i\}$ and $\{b'_j\}$ are linearly independent over \bar{k} , and a_1 and a'_1 are nonzero. Show that $D(a_1 a'_1) \subset \operatorname{Spec} A$ is nonempty. By the Weak Nullstellensatz 3.2.5, there is a maximal $\mathfrak{m} \subset A$ in $D(a_1 a'_1)$ with $A/\mathfrak{m} = \bar{k}$. By reducing modulo \mathfrak{m} , deduce $(\sum \bar{a}_i \otimes b_i)(\sum \bar{a}'_j \otimes b'_j) = 0$ in B , where the overline indicates residue modulo \mathfrak{m} . Show that this contradicts the fact that B is an integral domain.

10.5 ★ Properties not preserved by base change, and how to fix them

We saw in the previous section that many useful properties of morphisms are preserved by base change. Indeed, as discussed in §8.1, “preserved by base change” is one of the properties any “reasonable” class of morphisms should have. But some seemingly reasonable notions are not preserved by base change, and in this section we discuss how to modify them appropriately so that they are. We do this not for our own amusement, but because the resulting notions come up in nature, and perhaps in retrospect are the notions we should started with.

The “universal” patch to this problem is as follows. If P is some property of schemes, then a morphism of schemes is said to be **universally P** if it remains P under any base change. Then the class of universally P morphisms is preserved by base change. (Do you see why?)

An important example arising later (§11.5) is the notion of *universally closed morphisms*. Another example that we discuss next is that of universally injective morphisms, which turns out to be related to inseparable field extensions.

One problem with “universally P ” morphisms is that it is a priori hard to determine whether a given morphism is universally P — you seem to have to check every single base change of the morphism. Finding other equivalent criteria is thus essential.

10.5.1. Geometric fibers.

There are some notions that you should reasonably expect to be preserved by pullback based on your geometric intuition. Given a family in the topological category, fibers pull back in reasonable ways. So for example, any pullback of a family in which all the fibers are irreducible will also have this property; ditto for connected. Unfortunately, both of these fail in algebraic geometry, as Example 10.2.3

shows:

$$\begin{array}{ccc} \mathrm{Spec} \mathbb{C} \amalg \mathrm{Spec} \mathbb{C} & \longrightarrow & \mathrm{Spec} \mathbb{C} \\ \downarrow & & \downarrow \\ \mathrm{Spec} \mathbb{C} & \longrightarrow & \mathrm{Spec} \mathbb{R} \end{array}$$

The family on the right (the vertical map) has irreducible and connected fibers, and the one on the left doesn't. The same example shows that the notion of "integral fibers" also doesn't behave well under pullback. And we used it in §10.4.1 to show that injectivity isn't preserved by base change.

10.5.A. EXERCISE. Suppose k is a field of characteristic p , so $k(u)/k(u^p)$ is an inseparable extension. By considering $k(u) \otimes_{k(u^p)} k(u)$, show that the notion of "reduced fibers" does not necessarily behave well under pullback. (We will soon see that this happens only in characteristic p , in the presence of inseparability.)

We rectify this problem as follows.

10.5.2. A **geometric point** of a scheme X is defined to be a morphism $\mathrm{Spec} k \rightarrow X$ where k is an algebraically closed field. Awkwardly, this is now the third kind of "point" of a scheme! There are just plain points, which are elements of the underlying set; there are Z -valued points (Z a scheme), which are maps $Z \rightarrow X$, §7.3.10; and there are geometric points. Geometric points are clearly a flavor of a scheme-valued point, but they are also an enriched version of a (plain) point: they are the data of a point with an inclusion of the residue field of the point in an algebraically closed field.

A **geometric fiber** of a morphism $X \rightarrow Y$ is defined to be the fiber over a geometric point of Y , i.e., the fibered product with the geometric point $\mathrm{Spec} k \rightarrow Y$. A morphism has **connected** (resp. **irreducible**, **integral**, **reduced**) **geometric fibers** if all its geometric fibers are connected (resp. irreducible, integral, reduced). One usually says that the morphism has **geometrically connected** (resp. **geometrically irreducible**, **geometrically integral**, **geometrically reduced**) **fibers**. A k -scheme X is **geometrically connected** (resp. **geometrically irreducible**, **geometrically integral**, **geometrically reduced**) if the structure morphism $X \rightarrow \mathrm{Spec} k$ has geometrically connected (resp. irreducible, integral, reduced) fibers. We will soon see that to check any of these conditions, we need only base change to \bar{k} .

(Warning: in some sources, in the definition of "geometric point", "algebraically closed" is replaced by "separably closed".)

10.5.B. EXERCISE. Show that the notion of "connected (resp. irreducible, integral, reduced) geometric fibers" behaves well under base change.

10.5.C. EXERCISE FOR THE ARITHMETICALLY-MINDED. Show that for the morphism $\mathrm{Spec} \mathbb{C} \rightarrow \mathrm{Spec} \mathbb{R}$, all geometric fibers consist of two reduced points. (Cf. Example 10.2.3.) Thus $\mathrm{Spec} \mathbb{C}$ is a geometrically reduced but not geometrically irreducible \mathbb{R} -scheme.

10.5.D. EASY EXERCISE. Give examples of k -schemes that:

- (a) are reduced but not geometrically reduced;
- (b) are connected but not geometrically connected;

(c) are integral but not geometrically integral.

10.5.E. EXERCISE (TO CONVINCE GEOMETERS WHY GEOMETRIC FIBERS ARE MEANINGFUL). Recall Example 10.3.4, the projection of the parabola $y^2 = x$ to the x -axis, corresponding to the map of rings $\mathbb{Q}[x] \rightarrow \mathbb{Q}[y]$, with $x \mapsto y^2$. Show that the geometric fibers of this map are always two points, except for those geometric fibers “over $0 = [(x)]$ ”. (Note that $\text{Spec } \mathbb{C} \rightarrow \text{Spec } \mathbb{Q}[x]$ and $\text{Spec } \overline{\mathbb{Q}} \rightarrow \text{Spec } \mathbb{Q}[x]$, both corresponding to ring maps with $x \mapsto 0$, are both geometric points “above 0 ”.)

Checking whether a k -scheme is geometrically connected etc. seems annoying: you need to check every single algebraically closed field containing k . However, in each of these four cases, the failure of nice behavior of geometric fibers can already be detected after a finite extension of fields. For example, $\text{Spec } \mathbb{Q}(i) \rightarrow \text{Spec } \mathbb{Q}$ is not geometrically connected, and in fact you only need to base change by $\text{Spec } \mathbb{Q}(i)$ to see this. We make this precise as follows.

Suppose X is a k -scheme. If K/k is a field extension, define $X_K = X \times_k \text{Spec } K$. Consider the following twelve statements of the form “ X_K is [property]” for all fields [condition]’:

- $(C_K), (C_{K=\overline{K}}), (C_{\overline{K}}), (C_{k^s}),$
- $(I_K), (I_{K=\overline{K}}), (I_{\overline{K}}), (I_{k^s}),$
- $(R_K), (R_{K=\overline{K}}), (R_{\overline{K}}), (R_{k^p}).$

Here C means “connected”, I means “irreducible”, and R means “reduced”. Also, k^s is the separable closure of k , and k^p is the perfect closure of k . If $\text{char } k = 0$, then $\overline{k} = k^p = k^s$, so we can ignore $(C_{k^s}), (I_{k^s}),$ and (R_{k^p}) .

Thus (C_K) means “ X_K is connected for all fields K ”, $(I_{\overline{K}})$ means “ $X_{\overline{K}}$ is irreducible”, and (R_{k^p}) means “ X_{k^p} is reduced”. Geometrically connected (resp., geometrically irreducible, geometrically reduced) translates to $(C_{K=\overline{K}})$ (resp., $(I_{K=\overline{K}}), (R_{K=\overline{K}})$).

Trivially (C_K) implies $(C_{K=\overline{K}})$ implies $(C_{\overline{K}})$, and (C_K) implies (C_{k^s}) , and similarly (with appropriate changes) with “connected” replaced by “irreducible” and “reduced”.

10.5.F. EXERCISE.

- (a) Suppose that E/F is a field extension, and A is an F -algebra. Show that A is a subalgebra of $A \otimes_F E$. (Hint: think of these as vector spaces over F .)
- (b) Show that: $(C_{K=\overline{K}})$ implies (C_K) and $(C_{\overline{K}})$ implies (C_{k^s}) .
- (c) Show that: $(I_{K=\overline{K}})$ implies (I_K) and $(I_{\overline{K}})$ implies (I_{k^s}) .
- (d) Show that: $(R_{K=\overline{K}})$ implies (R_K) and $(R_{\overline{K}})$ implies (R_{k^p}) .

Thus for example a k -scheme is geometrically integral if and only if it remains integral under any field extension. Hence “geometrically integral” means “universally, integral fibers”.

10.5.3. Hard fact. In fact, (C_{k^p}) implies (C_K) , and thus (C_K) through (C_{k^s}) are all equivalent, and similarly for the other two properties (irreducibility and reducedness instead of connectedness, with k^p replacing k^s in the case of reducedness). We defer the explanation to the end of this section, in a double-starred discussion (§10.5.5). On a first reading, you may want to read only Corollary 10.5.13

on connectedness, Proposition 10.5.16 on irreducibility, Proposition 10.5.22 on reducedness, and Theorem 11.3.15 on varieties, and then to use them to solve Exercise 10.5.P. You can later come back and read the proofs, which include some useful tricks turning questions about general schemes over a field to questions about finite type schemes.

The following exercise may help even the geometrically-minded reader appreciate the utility of these notions. (There is nothing important about the dimension 2 and the degree 4 in this exercise!)

10.5.G. ★ EXERCISE. Recall from Remark 4.5.3 that the quartic curves in \mathbb{P}_k^2 are parametrized by a \mathbb{P}_k^{14} . (This will be made much more precise in §28.3.5.) Show that the points of \mathbb{P}_k^{14} corresponding to geometrically integral curves form an open subset. Explain the necessity of the modifier “geometrically” (even if k is algebraically closed).

10.5.4. Universally injective (radicial) morphisms.

As remarked in §10.4.1, injectivity is not preserved by base change. A better notion is that of **universally injective** morphisms: morphisms that are injections of sets after any base change. In keeping with the traditional agricultural terminology (sheaves, germs, . . . , cf. Remark 2.4.3), these morphisms were named **radicial** after one of the lesser root vegetables. As a first example: all locally closed embeddings are universally injective (as they are injective, and remain locally closed embeddings upon any base change). If you wish, you can show more generally that all monomorphisms are universally injective. (Hint: show that monomorphisms are injective, and that the class of monomorphisms is preserved by base change.)

Suppose you want to determine whether a given morphism is universally injective. A map of sets is injective if and only if each fiber contains at most one point. Now “the fiber of a base change is the base change of the fiber” (Exercise 10.3.C). Also, the underlying set of the scheme-theoretic fiber is the fiber of the map on underlying sets. Hence the question of universal injectivity turns into one about field theory. We make this precise in Exercise 10.5.I, after first making the field theory question precise in Exercise 10.5.H. En route, we will see why universal injectivity is algebro-geometric generalization of the notion of purely inseparable extensions of fields.

10.5.H. EXERCISE. Suppose E/F is an extension of fields, inducing $\phi : \text{Spec } E \rightarrow \text{Spec } F$.

(a) If E contains an element x transcendental over F , show that ϕ is not universally injective. Hint: Exercise 10.2.L and the fibered squares

$$\begin{array}{ccccc} ? & \longrightarrow & \text{Spec } F(x) \otimes_F F(y) & \longrightarrow & \text{Spec } F(y) \\ \downarrow & & \downarrow & & \downarrow \\ \text{Spec } E & \longrightarrow & \text{Spec } F(x) & \longrightarrow & \text{Spec } F. \end{array}$$

(b) If $E \setminus F$ contains an element x algebraic over F (with minimal polynomial $m_x(t) \in F[t]$, say), and separable over F , show that ϕ is not universally injective.

Hint: the fibered squares

$$\begin{array}{ccccc}
 ? & \longrightarrow & \operatorname{Spec} F[t]/(m_x(t)) \otimes_F \bar{F} & \longrightarrow & \operatorname{Spec} \bar{F} \\
 \downarrow & & \downarrow & & \downarrow \\
 \operatorname{Spec} E & \longrightarrow & \operatorname{Spec} F[t]/(m_x(t)) & \longrightarrow & \operatorname{Spec} F.
 \end{array}$$

(c) Suppose E contains no elements of the above forms — i.e., E/F is a **purely inseparable** extension (all elements of E are algebraic over F , and have some p^N th power in F), and E'/F is any other field extension, show that $\operatorname{Spec} E \otimes_F E'$ contains precisely one point.

10.5.I. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a morphism of schemes. Show that the following are equivalent.

- (i) The morphism π is universally injective.
- (ii) The morphism π is injective, and for each $p \in X$, the field extension $\kappa(p)/\kappa(\pi(p))$ is purely inseparable.

You may already see that the class of universally injective morphisms is also closed under composition, and local on the target, and hence forms a “reasonable” class (§8.1). In any case, we will see this in a different way in §11.3.14.

10.5.5. ★★ Proof of Hard Fact 10.5.3.

(The rest of this section is under construction.)

10.5.6. Proposition. — Suppose A and B are finite type k -algebras. Then $\operatorname{Spec} A \times_k \operatorname{Spec} B \rightarrow \operatorname{Spec} B$ is an open map.

This is the one fact we will not prove here. We could (it isn’t too hard), but instead we leave it until Remark 25.5.4.

10.5.7. Preliminary discussion.

10.5.8. Lemma. — Suppose X is a k -scheme. Then $X \rightarrow \operatorname{Spec} k$ is universally open, i.e., remains open after any base change.

Proof. If S is an arbitrary k -scheme, we wish to show that $X_S \rightarrow S$ is open. It suffices to consider the case $X = \operatorname{Spec} A$ and $S = \operatorname{Spec} B$. To show that $\phi: \operatorname{Spec} A \otimes_k B \rightarrow \operatorname{Spec} B$ is open, it suffices to show that the image of a distinguished open set $D(f)$ ($f \in A \otimes_k B$) is open.

We come to a trick we will use repeatedly, which we will call the tensor-finiteness trick. Write $f = \sum a_i \otimes b_i$, where the sum is *finite*. It suffices to replace A by the subring generated by the a_i . (Reason: if this ring is A' , then factor ϕ through $\operatorname{Spec}(A' \otimes_k B)$.) Thus we may assume A is finitely generated over k . Then use Proposition 10.5.6. \square

10.5.9. Lemma. — Suppose the field extension E/F is purely inseparable. Suppose X is any F -scheme. Then $\phi: X_E \rightarrow X$ is a homeomorphism.

Proof. The morphism ϕ is a bijection, so we may identify the points of X and X_E . (Reason: for any point $p \in X$, the scheme-theoretic fiber $\phi^{-1}(p)$ is a single point,

by the definition of pure inseparability.) The morphism ϕ is continuous (so opens in X are open in X_E), and by Lemma 10.5.8, ϕ is open (so opens in X_E are open in X). \square

10.5.10. Connectedness.

Recall that a connected component of a topological space is a maximal connected subset §3.6.12.

10.5.J. EXERCISE (PROMISED IN REMARK 3.6.13). Show that every point is contained in a connected component, and that connected components are closed. (Hint: see the hint for Exercise 3.6.O.)

10.5.K. TOPOLOGICAL EXERCISE. Suppose $\phi: X \rightarrow Y$ is open, and has nonempty connected fibers. Then ϕ induces a bijection of connected components.

10.5.11. Lemma. — Suppose X is geometrically connected over k . Then for any scheme Y/k , $X \times_k Y \rightarrow Y$ induces a bijection of connected components.

Proof. Combine Lemma 10.5.8 and Exercise 10.5.K. \square

10.5.L. EXERCISE (PROMISED IN REMARK 3.6.3). Show that a scheme X is disconnected if and only if there exists a function $e \in \Gamma(X, \mathcal{O}_X)$ that is an idempotent ($e^2 = e$) distinct from 0 and 1. (Hint: if X is the disjoint union of two open sets X_0 and X_1 , let e be the function that is 0 on X_0 and 1 on X_1 . Conversely, given such an idempotent, define $X_0 = V(e)$ and $X_1 = V(1 - e)$.)

10.5.12. Proposition. — Suppose k is separably closed, and A is a k -algebra with $\text{Spec } A$ connected. Then $\text{Spec } A$ is geometrically connected over k .

Proof. We wish to show that $\text{Spec } A \otimes_k K$ is connected for any field extension K/k . It suffices to assume that K is algebraically closed (as $\text{Spec } A \otimes_k \bar{K} \rightarrow \text{Spec } A \otimes_k K$ is surjective). By choosing an embedding $\bar{k} \hookrightarrow K$ and considering the diagram

$$\begin{array}{ccccc} \text{Spec } A \otimes_k K & \longrightarrow & \text{Spec } A \otimes_k \bar{k} & \xrightarrow[\text{by Lem. 10.5.9}]{\text{homeo.}} & \text{Spec } A \\ \downarrow & & \downarrow & & \downarrow \\ \text{Spec } K & \longrightarrow & \text{Spec } \bar{k} & \longrightarrow & \text{Spec } k \end{array}$$

it suffices to assume k is algebraically closed.

If $\text{Spec } A \otimes_k K$ is disconnected, then $A \otimes_k K$ contains an idempotent $e \neq 0, 1$ (by Exercise 10.5.L). By the tensor-finiteness trick, we may assume that A is a finitely generated algebra over k , and K is a finitely generated field extension. Write $K = K(B)$ for some integral domain B of finite type over k . Then by the tensor-finiteness trick, by considering the finite number of denominators appearing in a representative of e as a sum of decomposable tensors, $e \in A \otimes_k B[1/b]$ for some nonzero $b \in B$, so $\text{Spec } A \otimes_k B[1/b]$ is disconnected, say with open subsets U and V with $U \coprod V = \text{Spec } A \otimes_k B[1/b]$.

Now $\phi: \text{Spec } A \otimes_k B[1/b] \rightarrow \text{Spec } B[1/b]$ is an open map (Proposition 10.5.6), so $\phi(U)$ and $\phi(V)$ are nonempty open sets. As $\text{Spec } B[1/b]$ is connected, the intersection $\phi(U) \cap \phi(V)$ is a nonempty open set, which has a closed point p (with

residue field k , as $k = \bar{k}$). But then $\phi^{-1}(p) \cong \operatorname{Spec} A$, and we have covered $\operatorname{Spec} A$ with two disjoint open sets, yielding a contradiction. \square

10.5.13. Corollary. — *If k is separably closed, and Y is a connected k -scheme, then Y is geometrically connected.*

Proof. We wish to show that for any field extension K/k , Y_K is connected. By Proposition 10.5.12, $\operatorname{Spec} K$ is geometrically connected over k . Then apply Lemma 10.5.11 with $X = \operatorname{Spec} K$. \square

10.5.14. Irreducibility.

10.5.15. Proposition. — *Suppose k is separably closed, A is a k -algebra with $\operatorname{Spec} A$ irreducible, and K/k is a field extension. Then $\operatorname{Spec} A \otimes_k K$ is irreducible.*

Proof. We follow the philosophy of the proof of Proposition 10.5.12. As in the first paragraph of that proof, it suffices to assume that K and k are algebraically closed. If $A \otimes_k K$ is not irreducible, then we can find x and y with $V(x), V(y) \neq \operatorname{Spec} A \otimes_k K$ and $V(x) \cup V(y) = \operatorname{Spec} A \otimes_k K$. As in the second paragraph of the proof of Proposition 10.5.12, we may assume that A is a finitely generated algebra over k , and $K = K(B)$ for an integral domain B of finite type over k , and $x, y \in A \otimes_k B[1/b]$ for some nonzero $b \in B$. Then $D(x)$ and $D(y)$ are nonempty open subsets of $\operatorname{Spec} A \otimes_k B[1/b]$, whose image in $\operatorname{Spec} B[1/b]$ are nonempty opens, and thus their intersection is nonempty and contains a closed point p . But then $\phi^{-1}(p) \cong \operatorname{Spec} A$, and we have covered $\operatorname{Spec} A$ with two proper closed sets (the restrictions of $V(x)$ and $V(y)$), yielding a contradiction. \square

10.5.M. EXERCISE. Suppose k is separably closed, and A and B are k -algebras, both irreducible (with irreducible Spec , i.e., with one minimal prime). Show that $A \otimes_k B$ is irreducible too. (Hint: reduce to the case where A and B are finite type over k . Extend the proof of the previous proposition.)

10.5.N. EASY EXERCISE. Show that a scheme X is irreducible if and only if there exists an open cover $X = \bigcup U_i$ with U_i irreducible for all i , and $U_i \cap U_j \neq \emptyset$ for all i, j .

10.5.16. Proposition. — *Suppose K/k is a field extension of a separably closed field and X_k is irreducible. Then X_K is irreducible.*

Proof. Take an open cover $X_k = \bigcup U_i$ by pairwise intersecting irreducible affine open subsets. The base change of each U_i to K is irreducible by Proposition 10.5.15, and they pairwise intersect. The result then follows from Exercise 10.5.N. \square

10.5.O. EXERCISE. Suppose B is a geometrically integral k -algebra, and A is an integral k -algebra. Show that $A \otimes_k B$ is integral. (Once we define “variety”, this will imply that the product of a geometrically integral variety with an integral variety is an integral variety.) Hint: revisit the proof of Exercise 10.4.G.

10.5.17. Reducedness.

We recall the following fact from field theory, which is a refined version of the basics of transcendence theory developed in Exercise 12.2.A. Because this is a starred section, we content ourselves with a reference rather than a proof.

10.5.18. Algebraic Fact: finitely generated extensions of perfect fields are separably generated, see [E, Cor. 16.17(b)] or [vdW, §19.7]. — Suppose E/F is a finitely generated extension of a perfect field. Then it can be factored into a finite separable part and a purely transcendental part: $E/F(t_1, \dots, t_n)/F$.

10.5.19. Proposition. — Suppose B is a geometrically reduced k -algebra, and A is a reduced k -algebra. Then $A \otimes_k B$ is reduced.

(Compare this to Exercise 10.5.O.)

Proof. Reduce to the case where A is finitely generated over k using the tensor-finiteness trick. (Suppose we have $x \in A \otimes_k B$ with $x^n = 0$. Then $x = \sum a_i \otimes b_i$. Let A' be the finitely generated subring of A generated by the a_i . Then $A' \otimes_k B$ is a subring of $A \otimes_k B$. Replace A by A' .) Then A is a subring of the product $\prod K_i$ of the function fields of its irreducible components. (One can see this as follows: the kernel of $A \rightarrow \prod A/p_i$ is the intersection of the primes p_i , which is the ideal of nilpotents \mathfrak{N} by Theorem 3.2.13; but $\mathfrak{N} = 0$ as A is reduced. And $A/p_i \hookrightarrow K(A/p_i) = K_i$ clearly. Alternatively, this is clear from our discussion of associated points: Exercise ????, see also Exercise 6.5.I). So it suffices to prove it for A a product of fields. Then it suffices to prove it when A is a field. But then we are done, by the definition of geometric reducedness. \square

10.5.20. Proposition. — Suppose A is a reduced k -algebra. Then:

- (a) $A \otimes_k k(t)$ is reduced.
- (b) If E/k is a finite separable extension, then $A \otimes_k E$ is reduced.

Proof. (a) Clearly $A \otimes k[t]$ is reduced, and localization preserves reducedness (as reducedness is stalk-local, Exercise 5.2.A).

(b) Working inductively, we can assume E is generated by a single element, with minimal polynomial $p(t)$. By the tensor-finiteness trick, we can assume A is finitely generated over k . Then by the same trick as in the proof of Proposition 10.5.19, we can replace A by the product of its function fields of its components, and then we can assume A is a field. But then $A[t]/p(t)$ is reduced by the definition of separability of p . \square

10.5.21. Lemma. — Suppose E/k is a field extension of a perfect field, and A is a reduced k -algebra. Then $A \otimes_k E$ is reduced.

Proof. By the tensor-finiteness trick, we may assume E is finitely generated over k . By Algebraic Fact 10.5.18, we can factor E/k into extensions of the forms of Proposition 10.5.20 (a) and (b). We then apply Proposition 10.5.20. \square

10.5.22. Proposition. — Suppose E/k is an extension of a perfect field, and X is a reduced k -scheme. Then X_E is reduced.

Proof. Reduce to the case where X is affine. Use Lemma 10.5.21. \square

10.5.23. Corollary. — Suppose k is perfect, and A and B are reduced k -algebras. Then $A \otimes_k B$ is reduced.

Proof. By Lemma 10.5.21, A is a geometrically reduced k -algebra. Then apply Lemma 10.5.19. \square

10.5.P. EXERCISE (COMPLETING HARD FACT 10.5.3). Show that (R_{k^p}) implies (R_K) , (I_{k^s}) implies (I_K) , and (C_{k^s}) implies (C_K) .

10.5.Q. EXERCISE. Suppose that A and B are two integral domains that are \bar{k} -algebras. Show that $A \otimes_{\bar{k}} B$ is an integral domain.

10.6 Products of projective schemes: The Segre embedding

We next describe products of projective A -schemes over A . (The case of greatest initial interest is if $A = k$.) To do this, we need only describe $\mathbb{P}_A^m \times_A \mathbb{P}_A^n$, because any projective A -scheme has a closed embedding in some \mathbb{P}_A^m , and closed embeddings behave well under base change, so if $X \hookrightarrow \mathbb{P}_A^m$ and $Y \hookrightarrow \mathbb{P}_A^n$ are closed embeddings, then $X \times_A Y \hookrightarrow \mathbb{P}_A^m \times_A \mathbb{P}_A^n$ is also a closed embedding, cut out by the equations of X and Y (§10.2(3)). We will describe $\mathbb{P}_A^m \times_A \mathbb{P}_A^n$, and see that it too is a projective A -scheme. (Hence if X and Y are projective A -schemes, then their product $X \times_A Y$ over A is also a projective A -scheme.)

Before we do this, we will get some motivation from classical projective spaces (nonzero vectors modulo nonzero scalars, Exercise 4.4.F) in a special case. Our map will send $[x_0, x_1, x_2] \times [y_0, y_1]$ to a point in \mathbb{P}^5 , whose coordinates we think of as being entries in the “multiplication table”

$$\begin{bmatrix} x_0 y_0, & x_1 y_0, & x_2 y_0, \\ x_0 y_1, & x_1 y_1, & x_2 y_1 \end{bmatrix}.$$

This is indeed a well-defined map of sets. Notice that the resulting matrix is rank one, and from the matrix, we can read off $[x_0, x_1, x_2]$ and $[y_0, y_1]$ up to scalars. For example, to read off the point $[x_0, x_1, x_2] \in \mathbb{P}^2$, we take the first row, unless it is all zero, in which case we take the second row. (They can't both be all zero.) In conclusion: in classical projective geometry, given a point of \mathbb{P}^m and \mathbb{P}^n , we have produced a point in \mathbb{P}^{mn+m+n} , and from this point in \mathbb{P}^{mn+m+n} , we can recover the points of \mathbb{P}^m and \mathbb{P}^n .

Suitably motivated, we return to algebraic geometry. We define a map

$$\mathbb{P}_A^m \times_A \mathbb{P}_A^n \rightarrow \mathbb{P}_A^{mn+m+n}$$

by

$$\begin{aligned} ([x_0, \dots, x_m], [y_0, \dots, y_n]) &\mapsto [z_{00}, z_{01}, \dots, z_{ij}, \dots, z_{mn}] \\ &= [x_0 y_0, x_0 y_1, \dots, x_i y_j, \dots, x_m y_n]. \end{aligned}$$

More explicitly, we consider the map from the affine open set $U_i \times V_j$ (where $U_i = D(x_i)$ and $V_j = D(y_j)$) to the affine open set $W_{ij} = D(z_{ij})$ by

$$(x_{0/i}, \dots, x_{m/i}, y_{0/j}, \dots, y_{n/j}) \mapsto (x_{0/i} y_{0/j}, \dots, x_{i/i} y_{j/j}, \dots, x_{m/i} y_{n/j})$$

or, in terms of algebras, $z_{ab/ij} \mapsto x_{a/i} y_{b/j}$.

10.6.A. EXERCISE. Check that these maps glue to give a well-defined morphism $\mathbb{P}_A^m \times_A \mathbb{P}_A^n \rightarrow \mathbb{P}_A^{m+n}$.

10.6.1. We next show that this morphism is a closed embedding. We can check this on an open cover of the target (the notion of being a closed embedding is affine-local, Exercise 9.1.D). Let's check this on the open set where $z_{ij} \neq 0$. The preimage of this open set in $\mathbb{P}_A^m \times \mathbb{P}_A^n$ is the locus where $x_i \neq 0$ and $y_j \neq 0$, i.e., $U_i \times V_j$. As described above, the map of rings is given by $z_{ab/ij} \mapsto x_{a/i}y_{b/j}$; this is clearly a surjection, as $z_{aj/ij} \mapsto x_{a/i}$ and $z_{ib/ij} \mapsto y_{b/j}$. (A generalization of this ad hoc description will be given in Exercise 17.4.E.)

This map is called the **Segre morphism** or **Segre embedding**. If A is a field, the image is called the **Segre variety**.

10.6.B. EXERCISE. Show that the Segre scheme (the image of the Segre embedding) is cut out (scheme-theoretically) by the equations corresponding to

$$\text{rank} \begin{pmatrix} a_{00} & \cdots & a_{0n} \\ \vdots & \ddots & \vdots \\ a_{m0} & \cdots & a_{mn} \end{pmatrix} = 1,$$

i.e., that all 2×2 minors vanish. Hint: suppose you have a polynomial in the a_{ij} that becomes zero upon the substitution $a_{ij} = x_i y_j$. Give a recipe for subtracting polynomials of the form “monomial times 2×2 minor” so that the end result is 0. (The analogous question for the Veronese embedding in special cases is the content of Exercises 9.3.I and 9.3.K.)

10.6.2. Important Example. Let's consider the first nontrivial example, when $m = n = 1$. We get $\mathbb{P}_k^1 \times_k \mathbb{P}_k^1 \hookrightarrow \mathbb{P}_k^3$ (where k is a field). We get a single equation

$$\text{rank} \begin{pmatrix} a_{00} & a_{01} \\ a_{10} & a_{11} \end{pmatrix} = 1,$$

i.e., $a_{00}a_{11} - a_{01}a_{10} = 0$. We again meet our old friend, the quadric surface (§9.3.9)! Hence: the smooth quadric surface $wz - xy = 0$ (Figure 9.2) is isomorphic to $\mathbb{P}_k^1 \times_k \mathbb{P}_k^1$. Recall from Exercise 9.3.L that the quadric surface has two families (rulings) of lines. You may wish to check that one family of lines corresponds to the image of $\{x\} \times_k \mathbb{P}_k^1$ as x varies, and the other corresponds to the image $\mathbb{P}_k^1 \times_k \{y\}$ as y varies.

If k is an algebraically closed field of characteristic not 2, then by diagonalizability of quadratics (Exercise 5.4.J), all rank 4 (“full rank”) quadratics in 4 variables are isomorphic, so all rank 4 quadric surfaces over an algebraically closed field of characteristic not 2 are isomorphic to $\mathbb{P}_k^1 \times_k \mathbb{P}_k^1$.

Note that this is not true over a field that is not algebraically closed. For example, over \mathbb{R} , $w^2 + x^2 + y^2 + z^2 = 0$ (in $\mathbb{P}_{\mathbb{R}}^3$) is not isomorphic to $\mathbb{P}_{\mathbb{R}}^1 \times_{\mathbb{R}} \mathbb{P}_{\mathbb{R}}^1$. Reason: the former has no real points, while the latter has lots of real points.

You may wish to do the next two exercises in either order. The second can be used to show the first, but the first may give you insight into the second.

10.6.C. EXERCISE: A COORDINATE-FREE DESCRIPTION OF THE SEGRE EMBEDDING. Show that the Segre embedding can be interpreted as $\mathbb{P}V \times \mathbb{P}W \rightarrow \mathbb{P}(V \otimes W)$ via

the surjective map of graded rings

$$\mathrm{Sym}^\bullet(V^\vee \otimes W^\vee) \twoheadrightarrow \bigoplus_{i=0}^\infty (\mathrm{Sym}^i V^\vee) \otimes (\mathrm{Sym}^i W^\vee)$$

“in the opposite direction”.

10.6.D. EXERCISE: A COORDINATE-FREE DESCRIPTION OF PRODUCTS OF PROJECTIVE A -SCHEMES IN GENERAL. Suppose that S_\bullet and T_\bullet are finitely generated graded rings over A . Describe an isomorphism

$$(\mathrm{Proj} S_\bullet) \times_A (\mathrm{Proj} T_\bullet) \xleftarrow{\sim} \mathrm{Proj} \bigoplus_{n=0}^\infty (S_n \otimes_A T_n)$$

(where hopefully the definition of multiplication in the graded ring $\bigoplus_{n=0}^\infty (S_n \otimes_A T_n)$ is clear).

10.7 Normalization

We discuss normalization now only because the central construction gives practice with the central idea behind the construction in §10.1 of the fibered product (see Exercises 10.7.B and 10.7.I).

Normalization is a means of turning a *reduced* scheme into a normal scheme. (Unimportant remark: reducedness is not a necessary hypothesis, and is included only to avoid distraction. If you care, you can follow through the construction, and realize that the normalization of a scheme factors through the reduction, and is the normalization of the reduction.) A *normalization* of a reduced scheme X is a morphism $\nu: \tilde{X} \rightarrow X$ from a normal scheme, where ν induces a bijection of irreducible components of \tilde{X} and X , and ν gives a birational morphism on each of the irreducible components. It will satisfy a universal property, and hence it will be unique up to unique isomorphism. Figure 8.4 is an example of a normalization.

We begin with the case where X is irreducible, and hence (by Exercise 5.2.G) integral. (We will then deal with a more general case, and also discuss normalization in a function field extension.) In this case of irreducible X , the **normalization** $\nu: \tilde{X} \rightarrow X$ is a dominant morphism (not just dominant rational map!) from an irreducible normal scheme to X , such that any other such morphism factors through ν :

$$(10.7.0.1) \quad \begin{array}{ccccc} & \text{normal} & & & \text{normal} \\ & Y & \xrightarrow{\exists!} & \tilde{X} & \\ & \searrow \text{dominant} & & \swarrow \nu \text{ dominant} & \\ & & X & & \end{array}$$

Thus if the normalization exists, then it is unique up to unique isomorphism. We now have to show that it exists, and we do this in a way that will look familiar. We deal first with the case where X is affine, say $X = \mathrm{Spec} A$, where A is an integral domain. Then let \tilde{A} be the integral closure of A in its fraction field $K(A)$.

10.7.A. EXERCISE. (Recall that A is an integral domain.) Show that $\nu: \mathrm{Spec} \tilde{A} \rightarrow \mathrm{Spec} A$ satisfies the universal property of normalization. (En route, you might

show that the global sections of an irreducible normal scheme are also “normal”, i.e., integrally closed.)

10.7.B. IMPORTANT (BUT SURPRISINGLY EASY) EXERCISE. Show that normalizations of integral schemes exist in general. (Hint: Ideas from the existence of fibered products, §10.1, may help.)

10.7.C. EASY EXERCISE. Show that normalizations are integral and surjective. (Hint for surjectivity: the Lying Over Theorem, see §8.2.6.)

We will soon see that normalization of integral finite type k -schemes is always a birational morphism, in Exercise 10.7.N.

10.7.D. EXERCISE. Explain (by defining a universal property) how to extend the notion of normalization to the case where X is a reduced scheme, with possibly more than one component, but under the hypothesis that every affine open subset of X has finitely many irreducible components. Note that this includes all locally Noetherian schemes. (If you wish, you can show that the normalization exists in this case. See [Stacks, tag 035Q] for more.)

Here are some examples.

10.7.E. EXERCISE (NORMALIZATION OF THE NODAL CUBIC). Suppose that $\text{char } k \neq 2$. Show that $\text{Spec } k[t] \rightarrow \text{Spec } k[x, y]/(y^2 - x^2(x + 1))$ given by $(x, y) \mapsto (t^2 - 1, t(t^2 - 1))$ (see Figure 8.4) is a normalization. The target curve is called the **nodal cubic curve**. (Hint: show that $k[t]$ and $k[x, y]/(y^2 - x^2(x + 1))$ have the same fraction field. Show that $k[t]$ is integrally closed. Show that $k[t]$ is contained in the integral closure of $k[x, y]/(y^2 - x^2(x + 1))$.)

You will see from the previous exercise that once we guess what the normalization is, it isn't hard to verify that it is indeed the normalization. Perhaps a few words are in order as to where the polynomials $t^2 - 1$ and $t(t^2 - 1)$ arose in the previous exercise. The key idea is to guess $t = y/x$. (Then $t^2 = x + 1$ and $y = xt$ quickly.) This idea comes from three possible places. We begin by sketching the curve, and noticing the “node” at the origin. (“Node” will be formally defined in §30.3.1.) (a) The function y/x is well-defined away from the node, and at the node, the two branches have “values” $y/x = 1$ and $y/x = -1$. (b) We can also note that if $t = y/x$, then t^2 is a polynomial, so we will need to adjoin t in order to obtain the normalization. (c) The curve is cubic, so we expect a general line to meet the cubic in three points, counted with multiplicity. (We will make this precise when we discuss Bézout's Theorem, Exercise 19.6.K, but in this case we have already gotten a hint of this in Exercise 7.5.G.) There is a \mathbb{P}^1 parametrizing lines through the origin (with coordinate equal to the slope of the line, y/x), and most such lines meet the curve with multiplicity two at the origin, and hence meet the curve at precisely one other point of the curve. So this “coordinatizes” most of the curve, and we try adding in this coordinate.

10.7.F. EXERCISE. Find the normalization of the *cusp* $y^2 = x^3$ (see Figure 10.4). (“Cusp” will be formally defined in Definition 30.3.3.)




FIGURE 10.4. Normalization of a cusp


10.7.G. EXERCISE. Suppose $\text{char } k \neq 2$. Find the normalization of the *tacnode* $y^2 = x^4$, and draw a picture analogous to Figure 10.4. (“Tacnode” will be formally defined in Definition 30.3.3.)

(Although we haven’t defined “singularity”, “smooth”, “curve”, or “dimension”, you should still read this.) Notice that in the previous examples, normalization “resolves” the singularities (“non-smooth” points) of the curve. In general, it will do so in dimension one (in reasonable Noetherian circumstances, as normal local Noetherian integral domains of dimension one are all discrete valuation rings, §13.5), but won’t do so in higher dimension (the cone $z^2 = x^2 + y^2$ over a field k of characteristic not 2 is normal, Exercise 5.4.I(b)).

10.7.H. EXERCISE. Suppose $X = \text{Spec } \mathbb{Z}[15i]$. Describe the normalization $\tilde{X} \rightarrow X$. (Hint: $\mathbb{Z}[i]$ is a unique factorization domain, §5.4.5(0), and hence is integrally closed by Exercise 5.4.F.) Over what points of X is the normalization not an isomorphism?

Another exercise in a similar vein is the normalization of the “pinched plane”, Exercise 13.5.I.

10.7.I. EXERCISE (NORMALIZATION IN A FUNCTION FIELD EXTENSION, AN IMPORTANT GENERALIZATION). Suppose X is an integral scheme. The **normalization** $\nu: \tilde{X} \rightarrow X$ of X in a given algebraic field extension L of the function field $K(X)$ of X is a dominant morphism from a normal integral scheme \tilde{X} with function field L , such that ν induces the inclusion $K(X) \hookrightarrow L$, and that is universal with respect to this property.



Show that the normalization in a finite extension of fields exists.

The following two examples, one arithmetic and one geometric, show that this is an interesting construction.

10.7.J. EXERCISE. Suppose $X = \operatorname{Spec} \mathbb{Z}$ (with function field \mathbb{Q}). Find its integral closure in the field extension $\mathbb{Q}(i)$. (There is no “geometric” way to do this; it is purely an algebraic problem, although the answer should be understood geometrically.)

10.7.1. Remark: Rings of integers in number fields. A finite extension K of \mathbb{Q} is called a **number field**, and the integral closure of \mathbb{Z} in K the **ring of integers in K** , denoted \mathcal{O}_K . (This notation is awkward given our other use of the symbol \mathcal{O} .)

$$\begin{array}{ccc} \operatorname{Spec} K & \longrightarrow & \operatorname{Spec} \mathcal{O}_K \\ \downarrow & & \downarrow \\ \operatorname{Spec} \mathbb{Q} & \longrightarrow & \operatorname{Spec} \mathbb{Z} \end{array}$$

By the previous exercises, \mathcal{O}_K is a normal integral domain, and we will see soon (Theorem 10.7.3(a)) that it is Noetherian, and later (Exercise 12.1.F) that it has “dimension 1”. This is an example of a *Dedekind domain*, see §13.5.15. We will think of it as a “smooth” curve as soon as we define what “smooth” (really, regular) and “curve” mean.

10.7.K. EXERCISE. Find the ring of integers in $\mathbb{Q}(\sqrt{n})$, where n is square-free and $n \equiv 3 \pmod{4}$. (Hint: Exercise 5.4.I(a), where you will also be able to figure out the answer for square-free n in general.)

10.7.L. EXERCISE. Suppose $\operatorname{char} k \neq 2$ for convenience (although it isn’t necessary).

(a) Suppose $X = \operatorname{Spec} k[x]$ (with function field $k(x)$). Find its normalization in the field extension $k(y)$, where $y^2 = x^2 + x$. (Again we get a Dedekind domain.) Hint: this can be done without too much pain. Show that $\operatorname{Spec} k[x, y]/(x^2 + x - y^2)$ is normal, possibly by identifying it as an open subset of \mathbb{P}_k^1 , or possibly using Exercise 5.4.H.

(b) Suppose $X = \mathbb{P}^1$, with distinguished open subscheme $\operatorname{Spec} k[x]$. Find its normalization in the field extension $k(x, y)$, where $y^2 = x^2 + x$. (Part (a) involves computing the normalization over one affine open set; now figure out what happens over the “other” affine open set, and how to glue.)

10.7.2. Fancy fact: Finiteness of integral closure.

The following fact is useful.

10.7.3. Theorem (finiteness of integral closure). — Suppose A is a Noetherian integral domain, $K = K(A)$, L/K is a finite extension of fields, and B is the integral closure of A in L (“the integral closure of A in the field extension L/K ”, i.e., those elements of L integral over A).

- (a) If A is integrally closed and L/K is separable, then B is a finitely generated A -module.
 (b) (E. Noether) If A is a finitely generated k -algebra, then B is a finitely generated A -module.

Eisenbud gives a proof in a page and a half: (a) is [E, Prop. 13.14] and (b) is [E, Cor. 13.13]. A sketch is given in §10.7.5.

10.7.4. Warning. Part (b) does *not* hold for Noetherian A in general. In fact, the integral closure of a Noetherian ring need not be Noetherian (see [E, p. 299] for

some discussion). This is alarming. The existence of such an example is a sign that Theorem 10.7.3 is not easy.

In the following four exercises, you can replace “integral finite type k -scheme” by “integral variety” once you know what a variety is.

10.7.M. EXERCISE (FINITENESS OF NORMALIZATION).

(a) Show that if X is an integral finite type k -scheme, then its normalization $v: \tilde{X} \rightarrow X$ is a finite morphism.

(b) Suppose X is a locally Noetherian integral scheme. Show that if X is normal, then the normalization in a *finite* separable field extension is a finite morphism. Show that if X is an integral finite type k -scheme, then the normalization in a finite extension of fields is a finite morphism. In particular, once we define “variety” (Definition 11.3.9), you will see that this implies that the normalization of a variety (including in a finite extension of the function field) is a variety.

10.7.N. EXERCISE. Show that if X is an integral finite type k -scheme, then the normalization morphism is birational. (Hint: Proposition 7.5.7; or solve Exercise 10.7.O first.)

10.7.O. EXERCISE. Suppose that X is an integral finite type k -scheme. Show that the normalization map of X is an isomorphism on an open dense subset of X . Hint: Proposition 7.5.5.

10.7.P. EXERCISE. Suppose $\rho: Z \rightarrow X$ is a finite birational morphism from an integral finite type k -scheme to a *normal* integral finite type k -scheme. Show that ρ is an isomorphism.

10.7.5. ★★ Sketch of proof of finiteness of integral closure, Theorem 10.7.3. Here is a sketch to show the structure of the argument. It uses commutative algebra ideas from Chapter 12, so you should only glance at this to see that nothing fancy is going on. *Part (a):* reduce to the case where L/K is Galois, with group $\{\sigma_1, \dots, \sigma_n\}$. Choose $b_1, \dots, b_n \in B$ forming a K -vector space basis of L . Let M be the matrix (familiar from Galois theory) with ij th entry $\sigma_i b_j$, and let $d = \det M$. Show that the entries of M lie in B , and that $d^2 \in K$ (as d^2 is Galois-fixed). Show that $d \neq 0$ using linear independence of characters. Then complete the proof by showing that $B \subset d^{-2}(Ab_1 + \dots + Ab_n)$ (submodules of finitely generated modules over Noetherian rings are also Noetherian, Exercise 3.6.Y) as follows. Suppose $b \in B$, and write $b = \sum c_i b_i$ ($c_i \in K$). If c is the column vector with entries c_i , show that the i th entry of the column vector Mc is $\sigma_i b \in B$. Multiplying Mc on the left by $\text{adj } M$ (see the trick of the proof of Lemma 8.2.1), show that $dc_i \in B$. Thus $d^2 c_i \in B \cap K = A$ (as A is integrally closed), as desired.

For (b), use the Noether Normalization Lemma 12.2.4 to reduce to the case $A = k[x_1, \dots, x_n]$. Reduce to the case where L is normally closed over K . Let L' be the subextension of L/K so that L/L' is Galois and L'/K is purely inseparable. Use part (a) to reduce to the case $L = L'$. If $L' \neq K$, then for some q , L' is generated over K by the q th root of a finite set of rational functions. Reduce to the case $L' = k'(x_1^{1/q}, \dots, x_n^{1/q})$ where k'/k is a finite purely inseparable extension. In this case, show that $B = k'[x_1^{1/q}, \dots, x_n^{1/q}]$, which is indeed finite over $k[x_1, \dots, x_n]$. \square

10.7.6. Reduction of schemes, revisited. It may be enlightening to revisit the “reduction” of a scheme using this point of view. The argument is clean and simple, and by working it out for yourself, you will better appreciate the various moving parts. We now define “reduction” from scratch, and then show that it agrees with the definition of §9.4.10.

10.7.7. Definition. Suppose X is a scheme. In analogy with (10.7.0.1), we say $\rho : X^{\text{red}} \rightarrow X$ is “the” **reduction of X** if any morphism $W \rightarrow X$ from a reduced scheme W factors uniquely through ρ :

$$\begin{array}{ccc} \text{reduced} & & \\ & W \xrightarrow{\exists!} X^{\text{red}} & \\ & \searrow \quad \downarrow \rho & \\ & & X \end{array}$$

The “the” is in quotes because this is a slight abuse of notation. In analogy with fibered products and normalization, we now check that such an X^{red} exists:

10.7.Q. EXERCISE.

- In analogy with Exercise 10.7.A, show that $\rho : \text{Spec } A/\mathfrak{N}(A) \rightarrow \text{Spec } A$ satisfies the universal property of reduction.
- In analogy with Exercise 10.7.B, show that the reduction of schemes exists in general.
- To make sure we haven’t contradicted our earlier selves: Show that the construction of §9.4.10 satisfies this universal property.

10.7.8. Definition. By the universal property of reduction, any morphism of schemes $\pi : X \rightarrow Y$ induces a morphism $X^{\text{red}} \rightarrow Y^{\text{red}}$, the **reduction of π** which we quite reasonably name π^{red} . (You likely see that “red” gives a functor from the category of schemes to the category of reduced schemes.)

Separated and proper morphisms, and (finally!) varieties

11.1 Motivation

Separatedness is a fundamental notion. It is the analog of the Hausdorff condition for manifolds (see Exercise 11.3.B), and as with Hausdorffness, this geometrically intuitive notion ends up being just the right hypothesis to make theorems work. Although the definition initially looks odd, in retrospect it is just perfect.

Let's review why we like Hausdorffness. Recall that a topological space is **Hausdorff** if for every two points x and y , there are disjoint open neighborhoods of x and y . The real line is Hausdorff, but the "real line with doubled origin" (of which Figure 4.6 may be taken as a sketch) is not. Many proofs and results about manifolds use Hausdorffness in an essential way. For example, the classification of compact one-dimensional manifolds is very simple, but if the Hausdorff condition were removed, we would have a very wild set.

So once armed with this definition, we can cheerfully exclude the line with doubled origin from civilized discussion, and we can (finally) define the notion of a *variety*, in a way that corresponds to the classical definition.

With our motivation from manifolds, we shouldn't be surprised that all of our affine and projective schemes are separated. Subsets of Hausdorff topological spaces are automatically Hausdorff. Hence if Y is a manifold, and X is a subset that satisfies all the hypotheses of a manifold except possibly Hausdorffness, X is a manifold. Similarly, we will see that locally closed embeddings in something separated are also separated (combine Exercise 11.3.C and Proposition 11.3.3).

Grothendieck taught us that one should try to define properties of morphisms, not of objects; then we can say that an object has that property if its morphism to the final object has that property. We discussed this briefly at the start of Chapter 8. In this spirit, separatedness will be a property of morphisms, not schemes.

As an unexpected bonus, a separated morphism to an affine scheme has the property that the intersection of two affine open sets in the source is affine (Proposition 11.3.11). This will make Čech cohomology work very easily on (quasicompact) schemes (Chapter 19). You might consider this an analog of the fact that in \mathbb{R}^n , the intersection of two convex sets is also convex. As affine schemes are trivial from the point of view of quasicoherent cohomology, just as convex sets in \mathbb{R}^n have no cohomology, this metaphor is apt.

A lesson arising from the construction is the importance of the diagonal morphism (Definition 1.3.7). More precisely, given a morphism $\pi: X \rightarrow Y$, good

consequences can be leveraged from good behavior of the **diagonal morphism** $\delta_\pi: X \rightarrow X \times_Y X$ (the product of the identity morphism $X \rightarrow X$ with itself), usually through fun diagram chases. This lesson applies across many fields of geometry. (Another nice gift of the diagonal morphism: it will give us a good algebraic definition of differentials, in Chapter 22.)

So before we define separatedness, we make some observations about diagonal morphisms.

11.2 Fun with diagonal morphisms, and quasiseparatedness made easy

Recall that most classes of morphisms are “reasonable” in the sense of §8.1: they are (i) preserved by composition, (ii) preserved by base change, (iii) local on the target, and as a consequence they are (iv) preserved by product (Exercise 8.1.A).

Another useful consequence is that (v) “reasonable” classes of morphisms have a “Cancellation Property”. This involves morphisms whose *diagonal* lies in P , i.e., those $\pi: X \rightarrow Y$ such that $\delta_\pi: X \rightarrow X \times_Y X$ lies in P . We call this auxiliary class of morphisms $P\delta$ (“ P -diagonal”).

The Cancellation Theorem is a strange-looking, but very useful, result. Like the Diagonal-Base-Change diagram 1.3.S, this result is unexpectedly ubiquitous.

11.2.1. The Cancellation Theorem for a Property P of Morphisms. — *Let P be a class of morphisms that is preserved by base change and composition. (Every “reasonable” class of morphisms satisfies this, §8.1.) Suppose*

$$\begin{array}{ccc} X & \xrightarrow{\pi} & Y \\ & \searrow \tau & \swarrow \rho \\ & Z & \end{array}$$


is a commuting diagram of schemes. Suppose ρ is in $P\delta$ and τ is in P . Then π is in P .

The following diagram summarizes this important theorem:

$$\begin{array}{ccc} X & \xrightarrow{\therefore \in P} & Y \\ & \searrow \in P & \swarrow \in P\delta \\ & Z & \end{array}$$

When you plug in different P , you get very different-looking (and nonobvious) consequences, the first of which are given in Exercise 11.3.E. (Here are some facts you can prove easily, but which can be interpreted as applications of the Cancellation Theorem in *Sets*, and which may thus shed light on how the Cancellation Theorem works. If $\pi: X \rightarrow Y$ and $\rho: Y \rightarrow Z$ are maps of sets, and $\rho \circ \pi$ is injective, then so is π ; and if $\rho \circ \pi$ is surjective and ρ is injective, then π is surjective.)

11.2.A. EXERCISE. Prove the Cancellation Theorem 11.2.1. Hint: Interpret and make use of the following diagram.




(vi) If P is a “reasonable” class, so is $P\delta$:


11.2.2. Theorem. — *If P is a “reasonable” class of morphisms, then the class $P\delta$ of morphisms whose diagonal is in P is also a “reasonable” class.*

Proof. We verify that $P\delta$ satisfies the “three axioms of reasonableness”.


11.2.B. EXERCISE. Prove that “ $P\delta$ is preserved by composition”. Hint: Interpret and make use of the following diagram.



11.2.C. EXERCISE. Prove that “ $P\delta$ is preserved by base change”. Hint: Suppose (a) below is Cartesian. Show that (b) is also Cartesian.



11.2.D. EXERCISE. Prove that “ $P\delta$ is local on the target”. Hint: Suppose we have a cover $\{U_i \rightarrow Y\}$ of Y so that we have a Cartesian diagram (a). Then we have Cartesian diagrams (b).



□

(vii) Finally, if you have two morphisms from X that are of certain form, the next exercise shows what more you need to ensure that the map to the product is also of that form. This time you can find the “proof by diagram” by yourself.

11.2.E. EXERCISE. Suppose $\rho : X \rightarrow Y$ and $\sigma : X \rightarrow Z$ are morphisms of S -schemes in class \mathcal{P} , and $X \rightarrow S$ is in $\mathcal{P}\delta$. Show that $(\rho, \sigma) : X \rightarrow Y \times_S Z$ is in class \mathcal{P} as well. (See Exercises 11.3.A and ?? for applications.)

11.2.F. ★ EXERCISE. Show that finitely presented morphisms form a “reasonable” class of morphisms.

11.2.3. Example: Monomorphisms. Recall that a morphism $\pi : X \rightarrow Y$ is a monomorphism if and only if the diagonal $\delta_\pi : X \rightarrow X \times_Y X$ is an isomorphism (Exercise 1.3.W). Because isomorphisms form a “reasonable” class of morphisms, we have immediately shown that monomorphisms form a “reasonable” class as well (and accidentally solving Exercise 1.3.V, and proving that monomorphisms are preserved by base change).

11.2.4. Redefinition: Quasiseparated morphisms.

We (now) say a morphism $\pi : X \rightarrow Y$ is **quasiseparated** if the diagonal morphism $\delta_\pi : X \rightarrow X \times_Y X$ is quasicompact.

11.2.G. EXERCISE. Show that this agrees with our earlier definition of quasiseparated (§8.3.1): show that $\pi : X \rightarrow Y$ is quasiseparated if and only if for any affine open $\text{Spec } A$ of Y , and two affine open subsets U and V of X mapping to $\text{Spec } A$, $U \cap V$ is a *finite* union of affine open sets. (Possible hint: compare this to Proposition 11.3.11. Another possible hint: the Diagonal-Base-Change diagram, Exercise 1.3.S.)

11.2.5. Because quasicompact morphisms are a “reasonable” class (§10.4.D), Theorem 11.2.2 implies that quasiseparated morphisms are a reasonable class! In particular, quasiseparated morphisms are preserved by composition and base change, and the notion is local on the target.

11.2.6. All of this discussion was remarkably removed from geometry. To bring in the geometry, we now look more closely at the diagonal, which isn’t just any old kind of morphism.

11.3 Separatedness, and varieties

11.3.1. Proposition. — *Let $\pi : X \rightarrow Y$ be a morphism of schemes.*

(a) *If X and Y are affine, then the diagonal morphism $\delta_\pi : X \rightarrow X \times_Y X$ is a closed embedding.*

(b) *In general, the diagonal morphism $\delta_\pi : X \rightarrow X \times_Y X$ is a locally closed embedding.*

The corresponding locally closed subscheme of $X \times_Y X$ — the **diagonal (locally closed) subscheme** — will be denoted Δ .

Proof. (a) Suppose $X = \text{Spec } A$ and $Y = \text{Spec } B$. Then the diagonal morphism corresponds to the natural ring map $A \otimes_B A \rightarrow A$ given by $a_1 \otimes a_2 \mapsto a_1 a_2$, which is obviously surjective.

(b) We will describe a union of open subsets of $X \times_Y X$ covering the image of X , such that the image of X is a closed embedding in this union.

Say Y is covered with affine open sets V_i and X is covered with affine open sets U_{ij} , with $\pi: U_{ij} \rightarrow V_i$. Note that $U_{ij} \times_{V_i} U_{ij}$ is an affine open subscheme of the product $X \times_Y X$ (basically this is how we constructed the product, by gluing together affine building blocks). Then the diagonal is covered by these affine open subsets $U_{ij} \times_{V_i} U_{ij}$. (Any point $p \in X$ lies in some U_{ij} ; then $\delta(p) \in U_{ij} \times_{V_i} U_{ij}$. Figure 11.1 may be helpful.) Note that $\delta^{-1}(U_{ij} \times_{V_i} U_{ij}) = U_{ij}$: clearly $U_{ij} \subset \delta^{-1}(U_{ij} \times_{V_i} U_{ij})$, and because $\text{pr}_1 \circ \delta = \text{id}_X$ (where pr_1 is the first projection), $\delta^{-1}(U_{ij} \times_{V_i} U_{ij}) \subset U_{ij}$. Finally, by part (a), $U_{ij} \rightarrow U_{ij} \times_{V_i} U_{ij}$ is a closed embedding.




FIGURE 11.1. An open neighborhood of the diagonal is covered by $U_{ij} \times_{V_i} U_{ij}$

The open subsets we described may not cover $X \times_Y X$, so we have not shown that δ is a closed embedding.

11.3.A. EASY EXERCISE. Suppose $X \rightarrow Y$ and $X \rightarrow Z$ are both locally closed embeddings. Prove that $X \rightarrow Y \times Z$ is a locally closed embedding. Hint: Exercise 11.2.E.

11.3.2. Definition. A morphism $\pi: X \rightarrow Y$ is **separated** if the diagonal morphism $\delta_\pi: X \rightarrow X \times_Y X$ is a closed embedding. An A -scheme X is said to be **separated over A** if the structure morphism $X \rightarrow \text{Spec } A$ is separated. When people say that a scheme (rather than a morphism) X is separated, they mean implicitly that some “structure morphism” is separated. For example, if they are talking about A -schemes, they mean that X is separated over A .

Thanks to Proposition 11.3.1(b) (and once you show that a locally closed embedding whose image is closed is actually a closed embedding), a morphism is separated if and only if the diagonal Δ is a closed subset — a purely topological condition on the diagonal. This is reminiscent of a definition of Hausdorff, as the next exercise shows.

11.3.B. EXERCISE (ONLY FOR THOSE SEEKING TOPOLOGICAL MOTIVATION). Show that a topological space X is Hausdorff if and only if the diagonal is a closed subset of $X \times X$. (The reason separatedness of schemes doesn't give Hausdorffness — i.e., that for any two open points x and y there aren't necessarily disjoint open neighborhoods — is that in the category of schemes, the topological space $X \times X$ is not in general the product of the topological space X with itself, see §10.1.2.)

11.3.3. Proposition. — *Separated morphisms form a “reasonable” class in the sense of 8.1 — they are by preserved by composition, base change, and the notion is local on the target.*

Proof. The class of locally closed embeddings is “reasonable” in this sense (Exercise 10.4.A), so by Theorem 11.2.2 the same is true of separated morphisms. \square

11.3.C. EXERCISE. Show that monomorphisms are separated. Hence locally closed embeddings (and in particular open and closed embeddings) are separated.

11.3.D. EXERCISE. Show that separated morphisms are quasiseparated. (Hint: closed embeddings are affine, hence quasicompact.)

11.3.4. Proposition. — *Affine morphisms are separated. In particular, finite morphisms are separated.*

Proof. By Proposition 11.3.1(a), every morphism of affine schemes is separated. The notion of separatedness is local on the target, so we are done. \square

Many of the following exercises will drive home the importance of the Cancellation Theorem 11.2.1.

11.3.E. EXERCISE. Show that an A -scheme is separated (over A) if and only if it is separated over \mathbb{Z} . In particular, a complex scheme is separated over \mathbb{C} if and only if it is separated over \mathbb{Z} , so complex geometers and arithmetic geometers can discuss separated schemes without confusion.

11.3.F. EASY EXERCISE. Suppose we have morphisms $X \xrightarrow{\pi} Y \xrightarrow{\rho} Z$.

(a) Show that if $\rho \circ \pi$ is a locally closed embedding (resp. locally of finite type, separated), then so is π .

(b) If $\rho \circ \pi$ is quasicompact, and Y is Noetherian, show that π is quasicompact. Hint: Exercise 8.3.B(a).

(c) If $\rho \circ \pi$ is quasiseparated, show that π is quasiseparated. Hint: Exercise 8.3.B(b).

11.3.G. EXERCISE. Suppose $\mu: Z \rightarrow X$ is a morphism, and $\sigma: X \rightarrow Z$ is a section of μ , i.e., $\mu \circ \sigma$ is the identity on X .

$$\begin{array}{c} Z \\ \mu \downarrow \uparrow \sigma \\ X \end{array}$$

Show that σ is a locally closed embedding. Show that if μ is separated, then σ is a closed embedding. Show that if μ is quasiseparated, then σ is quasicompact. Give an example to show that σ need not be a closed embedding if μ is not separated.

11.3.5. The graph morphism. Suppose $\pi: X \rightarrow Y$ is a morphism of Z -schemes. The morphism $\gamma_\pi: X \rightarrow X \times_Z Y$ given by $\gamma_\pi = (\text{id}, \pi)$ is called the **graph morphism** (see Figure 11.2). Then π factors as $\text{pr}_2 \circ \gamma_\pi$, where pr_2 is the second projection. The diagram of Figure 11.2 is often called the **graph of a morphism** π . (We will discuss graphs of rational maps in §11.4.4.) Exercise 11.3.G applied to $X \times_Z Y \rightarrow X$ yields the following.




FIGURE 11.2. The graph morphism

11.3.6. Proposition. — *The graph morphism γ_π is always a locally closed embedding. If Y is a separated Z -scheme (i.e., $Y \rightarrow Z$ is separated), then γ_π is a closed embedding. If Y is a quasiseparated Z -scheme, then γ_π is quasicompact.*

The image of γ_π (a locally closed subscheme of $X \times_Z Y$) is often also called the **graph of π** ; we might denote it Γ_π .

11.3.H. EXERCISE WE WON'T USE. Suppose P is a class of morphisms such that closed embeddings are in P , and P is preserved by fibered product and composition. Recall the Definition 10.7.8 of the “reduction” of a morphism. Show that if $\pi: X \rightarrow Y$ is in P then $\pi^{\text{red}}: X^{\text{red}} \rightarrow Y^{\text{red}}$ is in P . (Two examples are the classes of separated morphisms and quasiseparated morphisms.) Hint:

$$\begin{array}{ccc} X^{\text{red}} & \xrightarrow{\pi^{\text{red}}} & Y^{\text{red}} \\ \downarrow \in P & \searrow & \downarrow \in P \delta \\ X & \xrightarrow{\pi \in P} & Y \end{array}$$

11.3.7. Important examples.

11.3.I. EXERCISE. Show that the line with doubled origin (Example 4.4.5) is not separated, by verifying that the image of the diagonal morphism is not closed. (Another argument is given below, in Exercise 11.4.C. A fancy argument is given in Exercise 13.7.C.)

We next come to our first example of something separated but not affine. The following single calculation will imply that all quasiprojective A -schemes are separated (once we know that the composition of separated morphisms are separated, Proposition 11.3.3).

11.3.8. Proposition. — *The morphism $\mathbb{P}_A^n \rightarrow \text{Spec } A$ is separated.*

We give two proofs. The first is by direct calculation. The second requires no calculation, and just requires that you remember some classical constructions described earlier.

Proof 1: Direct calculation. We cover $\mathbb{P}_A^n \times_A \mathbb{P}_A^n$ with open sets of the form $U_i \times_A U_j$, where U_0, \dots, U_n form the “usual” affine open cover. The case $i = j$ was taken care of before, in the proof of Proposition 11.3.1(a). If $i \neq j$ then

$$U_i \times_A U_j \cong \operatorname{Spec} A[x_{0/i}, \dots, x_{n/i}, y_{0/j}, \dots, y_{n/j}] / (x_{i/i} - 1, y_{j/j} - 1).$$

Now the restriction of the diagonal Δ is contained in U_i (as the diagonal morphism composed with projection to the first factor is the identity), and similarly is contained in U_j . Thus the diagonal morphism over $U_i \times_A U_j$ is $U_i \cap U_j \rightarrow U_i \times_A U_j$. This is a closed embedding, as the corresponding map of rings

$$A[x_{0/i}, \dots, x_{n/i}, y_{0/j}, \dots, y_{n/j}] \rightarrow A[x_{0/i}, \dots, x_{n/i}, x_{j/i}^{-1}] / (x_{i/i} - 1)$$

(given by $x_{k/i} \mapsto x_{k/i}$, $y_{k/j} \mapsto x_{k/i}/x_{j/i}$) is clearly a surjection (as each generator of the ring on the right is clearly in the image — note that $x_{j/i}^{-1}$ is the image of $y_{i/j}$). \square

Proof 2: Classical geometry. Note that the diagonal morphism $\delta: \mathbb{P}_A^n \rightarrow \mathbb{P}_A^n \times_A \mathbb{P}_A^n$ followed by the Segre embedding $S: \mathbb{P}_A^n \times_A \mathbb{P}_A^n \rightarrow \mathbb{P}^{n^2+2n}$ (§10.6, a closed embedding) can also be factored as the second Veronese embedding $v_2: \mathbb{P}_A^n \rightarrow \mathbb{P}^{\binom{n+2}{2}-1}$ (§9.3.6) followed by a linear map $L: \mathbb{P}^{\binom{n+2}{2}-1} \rightarrow \mathbb{P}^{n^2+2n}$, both of which are closed embeddings (the linear map L by Exercise 9.3.C).

$$\begin{array}{ccccc}
 & & \mathbb{P}_A^n \times_A \mathbb{P}_A^n & & \\
 & \nearrow \delta & & \searrow S & \\
 \mathbb{P}_A^n & & & & \mathbb{P}^{n^2+2n} \\
 & \searrow v_2 & & \nearrow L & \\
 & & \mathbb{P}^{\binom{n+2}{2}-1} & &
 \end{array}$$

δ is labeled “1. (?) cl. emb.”, S is labeled “cl. emb.”, v_2 is labeled “cl. emb.”, and L is labeled “cl. emb.”.

Informally, in coordinates:

$$\begin{array}{ccc}
 ([x_0, x_1, \dots, x_n], [x_0, x_1, \dots, x_n]) & & \\
 \nearrow \delta & & \searrow S \\
 [x_0, x_1, \dots, x_n] & & \begin{bmatrix} x_0^2 & x_0 x_1 & \cdots & x_0 x_n \\ x_1 x_0 & x_1^2 & \cdots & x_1 x_n \\ \vdots & \vdots & \ddots & \vdots \\ x_n x_0 & x_n x_1 & \cdots & x_n^2 \end{bmatrix} \\
 \searrow v_2 & & \nearrow L \\
 [x_0^2, x_0 x_1, \dots, x_{n-1} x_n, x_n^2] & &
 \end{array}$$

The composed map $\mathbb{P}_A^n \rightarrow \mathbb{P}^{n^2+2n}$ may be written as

$$[x_0, \dots, x_n] \mapsto [x_0^2, x_0 x_1, x_0 x_2, \dots, x_n^2],$$

where the subscripts on the right run over all ordered pairs (i, j) where $0 \leq i, j \leq n$.) This forces δ to send closed sets to closed sets (or else $S \circ \delta$ won't, but $L \circ v_2$ does). \square

11.3.9. Important Definition (finally!). A **variety** over a field k , or k -**variety**, is a reduced, separated scheme of finite type over k . For example, a reduced finite type affine k -scheme is a variety. (Do you see why separatedness holds?) The notion of variety generalizes our earlier notion of affine variety (§5.3.7) and projective variety (§5.3.E, see Proposition 11.3.13). By Exercise 5.3.F, the closed points of a variety are dense. (Notational caution: In some sources, the additional condition of irreducibility is imposed. Also, it is often assumed that k is algebraically closed.)

Varieties over k form a category: morphisms of k -varieties are just morphisms as k -schemes. We will shortly see (in Exercise 11.3.J) that such morphisms are automatically finite type and separated. (Of course, the category of varieties over an algebraically closed field k long predates modern scheme theory.)

A **subvariety** of a variety X is a *reduced* locally closed subscheme of X (which you can quickly check is a variety itself). An open subvariety of X is an open subscheme of X . (Reducedness is automatic in this case.) A closed subvariety of X is a reduced closed subscheme of X .

11.3.J. EASY EXERCISE. Show that morphisms of k -varieties are finite type and separated.

11.3.K. EXERCISE (PRODUCTS OF IRREDUCIBLE VARIETIES OVER \bar{k} ARE IRREDUCIBLE VARIETIES). Use Exercise 10.4.G and properties of separatedness to show that the product of two irreducible \bar{k} -varieties is an irreducible \bar{k} -variety.

11.3.L. ★★ EXERCISE (COMPLEX ALGEBRAIC VARIETIES YIELD COMPLEX ANALYTIC VARIETIES; FOR THOSE WITH SUFFICIENT BACKGROUND). Show that the analytification (Exercise 7.3.16) of a complex algebraic variety is a complex analytic variety.

If you have read the double-starred section on group objects in a category (§7.6.4), you will automatically have the notion of a **group variety**. We will discuss this a bit more in §11.5.9.

11.3.10. Here is a very handy consequence of separatedness.

11.3.11. Proposition. — Suppose $X \rightarrow \operatorname{Spec} A$ is a separated morphism to an affine scheme, and U and V are affine open subsets of X . Then $U \cap V$ is an affine open subset of X .

Before proving this, we state a consequence that is otherwise not clear. If $X = \operatorname{Spec} A$, then the intersection of any two affine open subsets of A is an affine open subset. This is certainly not an obvious fact! We know the intersection of two distinguished affine open sets is affine (from $D(f) \cap D(g) = D(fg)$), but we have little handle on affine open sets in general.

Warning: this property does not characterize separatedness. (The correct property is “affine diagonal”, which we will not discuss, but which is useful in the theory of stacks.) For example, if A is a field k and X is the line with doubled origin over k , then X also has this property.

In the proof of Proposition 11.3.11, we will use the following nifty result.

11.3.M. EASY EXERCISE. Suppose U and V are open subsets of an A -scheme X . Show that $\Delta \cap (U \times_A V) \cong U \cap V$. (The intersection of two locally closed subschemes was defined in Exercise 10.2.C.) Hints: (i) this was basically proved in the course of Proof 1 of Proposition 11.3.8. (ii) See Figure 11.3. (iii) You can interpret this statement using the Diagonal-Base-Change diagram, Exercise 1.3.S:




FIGURE 11.3. $\Delta \cap (U \times_A V) \cong U \cap V$ (Exercise 11.3.M)

11.3.12. Proof of Proposition 11.3.11. By Exercise 11.3.M, $(U \times_A V) \cap \Delta \cong U \cap V$, where Δ is the diagonal. But $U \times_A V$ is affine (the fibered product of two affine schemes over an affine scheme is affine, Step 1 of our construction of fibered products, Theorem 10.1.1), and $(U \times_A V) \cap \Delta$ is a closed subscheme of the affine scheme $U \times_A V$, and hence $U \cap V$ is affine. \square

11.3.13. Corollary. — *Every quasiprojective A -scheme is separated over A . In particular, every reduced quasiprojective k -scheme is a k -variety.*

Proof. Suppose $X \rightarrow \operatorname{Spec} A$ is a quasiprojective A -scheme. The structure morphism can be factored into an open embedding composed with a closed embedding followed by $\mathbb{P}_A^n \rightarrow \operatorname{Spec} A$. Open embeddings and closed embeddings are separated (Exercise 11.3.C), and $\mathbb{P}_A^n \rightarrow \operatorname{Spec} A$ is separated (Proposition 11.3.8). Compositions of separated morphisms are separated (Proposition 11.3.3), so we are done. \square

11.3.14. $\star\star$ Universally injective morphisms and the diagonal.

11.3.N. EXERCISE. Show that $\pi: X \rightarrow Y$ is universally injective if and only if the diagonal morphism $\delta_\pi: X \rightarrow X \times_Y X$ is surjective. (We will soon see that δ_π is *always* injective, by Proposition 11.3.1(b).)

Because surjective morphisms form a “reasonable” class (Exercise 10.4.F), we see that universally injective morphisms also form a “reasonable” class.

11.3.O. EASY EXERCISE. If $\pi: X \rightarrow Y$ and $\rho: Y \rightarrow Z$ are morphisms, and $\rho \circ \pi$ is universally injective, show that π is universally injective.

11.3.P. EXERCISE.

- (a) Show that universally injective morphisms are separated.
- (b) Show that a map between finite type schemes over an algebraically closed field \bar{k} is universally injective if and only if it is injective on closed points.

11.3.15. Theorem. —

- (a) If k is perfect, the product of affine k -varieties (over $\text{Spec } k$) is an affine k -variety.
- (b) If k is separably closed, the product of irreducible affine k -varieties is an irreducible affine k -variety.
- (c) If k is separably closed, the product of connected affine k -varieties is a connected affine k -variety.

Once we define varieties in general, in Definition 11.3.9, you will be able to remove the adjective “affine” throughout this statement of Theorem 11.3.15. (We also remark that (b) was proven in Exercise 10.4.G under the stronger hypothesis that k is *algebraically* closed.)

Proof. (a) The finite type and separated statements are straightforward, as both properties are preserved by base change and composition. For reducedness, reduce to the affine case, then use Corollary 10.5.23.

(b) It only remains to show irreducibility. Reduce to the affine case using Exercise 10.5.N (as in the proof of Proposition 10.5.16). Then use Exercise 10.5.M.

(c) This follows from Corollary 10.5.13. \square

11.4 The locus where two morphisms from X to Y agree, and the “Rational-to-Separated” Theorem

When we introduced rational maps in §7.5, we promised that in good circumstances, a rational map has a “largest domain of definition”. We are now ready to make precise what “good circumstances” means, in the Reduced-to-Separated Theorem 11.4.2. We first introduce an important result making sense of locus where two morphisms with the same source and target “agree”.

11.4.A. USEFUL EXERCISE: THE LOCUS WHERE TWO MORPHISMS AGREE. Suppose $\pi: X \rightarrow Y$ and $\pi': X \rightarrow Y$ are two morphisms over some scheme Z .

$$\begin{array}{ccc} X & \begin{array}{c} \xrightarrow{\pi} \\ \xrightarrow{\pi'} \end{array} & Y \\ & \searrow & \swarrow \\ & Z & \end{array}$$

We can now give meaning to the phrase ‘the locus where π and π' agree’, and that in particular there is a largest locally closed subscheme where they agree — which is closed if Y is separated over Z . Suppose $\mu: W \rightarrow X$ is some morphism (not assumed to be a locally closed embedding). We say that π and π' agree on μ if $\pi \circ \mu = \pi' \circ \mu$. Show that there is a locally closed subscheme $i: V \hookrightarrow X$ on which π and π' agree, such that any morphism $\mu: W \rightarrow X$ on which π and π' agree factors uniquely through i , i.e., there is a unique $j: W \rightarrow V$ such that $\mu = i \circ j$. Show further that if $Y \rightarrow Z$ is separated, then $i: V \hookrightarrow X$ is a closed embedding. Hint: define V to be the following fibered product:

$$\begin{array}{ccc} V & \longrightarrow & Y \\ \downarrow & & \downarrow \delta \\ X & \xrightarrow{(\pi, \pi')} & Y \times_Z Y. \end{array}$$

As δ is a locally closed embedding, $V \rightarrow X$ is too. Then if $\mu: W \rightarrow X$ is any morphism such that $\pi \circ \mu = \pi' \circ \mu$, then μ factors through V .

The fact that the locus where two maps agree can be nonreduced should not come as a surprise: consider two maps from \mathbb{A}_k^1 to itself, $\pi(x) = 0$ and $\pi'(x) = x^2$. They agree when $x = 0$, but the situation is (epsilononically) better than that — they should agree even on $\text{Spec } k[x]/(x^2)$.

11.4.1. Minor Remarks.

(i) In the previous exercise, we describe $V \hookrightarrow X$ by a universal property. Taking this as the definition, it is not a priori clear that V is a locally closed subscheme of X , or even that it exists.

(ii) Warning: consider two maps from $\text{Spec } \mathbb{C}$ to itself over $\text{Spec } \mathbb{R}$, the identity and complex conjugation. These are both maps from a point to a point, yet they do not agree despite agreeing as maps of sets. (If you are not convinced that they disagree as morphisms, this might help: after base change $\text{Spec } \mathbb{C} \rightarrow \text{Spec } \mathbb{R}$, they do not agree even as maps of sets.)

(iii) More generally, the (set-theoretic) locus where π and π' agree can be interpreted as follows: π and π' agree at p if $\pi(p) = \pi'(p)$ and the two maps of residue fields are the same. (You may enjoy thinking this through as an exercise.)

11.4.B. EXERCISE: MAPS OF \bar{k} -VARIETIES ARE DETERMINED BY THE MAPS ON CLOSED POINTS. Suppose $\pi: X \rightarrow Y$ and $\pi': X \rightarrow Y$ are two morphisms of \bar{k} -varieties that are the same at the level of closed points (i.e., for each closed point $p \in X$, $\pi(p) = \pi'(p)$). Show that $\pi = \pi'$. (This implies that the functor from the category of “classical varieties over \bar{k} ”, which we won’t define here, to the category of \bar{k} -schemes, is fully faithful. Can you generalize this appropriately to non-algebraically closed fields?)

11.4.C. LESS IMPORTANT EXERCISE. Show that the line with doubled origin X (Example 4.4.5) is not separated, by finding two morphisms $\pi: W \rightarrow X, \pi': W \rightarrow X$ whose domain of agreement is not a closed subscheme (cf. Proposition 11.3.1(b)). (Another argument was given above, in Exercise 11.3.I. A fancy argument will be given in Exercise 13.7.C.)

We now come to the central result of this section.

11.4.2. Reduced-to-Separated Theorem. — *Two S -morphisms $\pi: U \rightarrow Z, \pi': U \rightarrow Z$ from a reduced scheme to a separated S -scheme agreeing on a dense open subset of U are the same.*

Proof. Let V be the locus where π and π' agree. It is a closed subscheme of U (by Exercise 11.4.A), which contains a dense open set. But the only closed subscheme of a reduced scheme U whose underlying set is dense is all of U . (Do you see why this last statement is true? Can you give a counterexample if the reducedness hypothesis is removed?) \square

11.4.3. Consequence 1. Hence (as X is reduced and Y is separated) if we have two morphisms from open subsets of X to Y , say $\pi: U \rightarrow Y$ and $\pi': V \rightarrow Y$, and they agree on a dense open subset Z of $U \cap V$, then they necessarily agree on $U \cap V$.

Consequence 2. A rational map has a largest **domain of definition** on which $\pi: U \dashrightarrow Y$ is a morphism, which is the union of all the domains of definition. In particular, a rational function on a reduced scheme has a largest domain of definition. For example, the domain of definition of $\mathbb{A}_k^2 \dashrightarrow \mathbb{P}_k^1$ given by $(x, y) \mapsto [x, y]$ has domain of definition $\mathbb{A}_k^2 \setminus \{(0, 0)\}$ (cf. §7.5.3). This partially extends the definition of the domain of a rational function on a locally Noetherian scheme (Definition 6.5.34). The complement of the domain of definition is called the **indeterminacy locus** (a phrase used earlier in §6.5.34), and its points are sometimes called **fundamental points** of the rational map, although we won't use these phrases. (We will see in Exercise 23.4.L that a rational map to a projective scheme can be upgraded to an honest morphism by “blowing up” a scheme-theoretic version of the locus of indeterminacy.)

11.4.D. EXERCISE. Show that the Reduced-to-Separated Theorem 11.4.2 is false if we give up reducedness of the source or separatedness of the target. Here are some possibilities. For the first, consider the two maps from $\operatorname{Spec} k[x, y]/(y^2, xy)$ to $\operatorname{Spec} k[t]$, where we take π given by $t \mapsto x$ and π' given by $t \mapsto x + y$; f_1 and f_2 agree on the distinguished open set $D(x)$, see Figure 11.4. For the second, consider the two maps from $\operatorname{Spec} k[t]$ to the line with the doubled origin, one of which maps to the “upper half”, and one of which maps to the “lower half”. These two morphisms agree on the dense open set $D(t)$, see Figure 11.4.

11.4.4. Graphs of rational maps. (Graphs of *morphisms* were defined in §11.3.5.) If X is reduced and Y is separated, define the **graph Γ_π of a rational map $\pi: X \dashrightarrow Y$** as follows. Let (U, π') be any representative of this rational map (so $\pi': U \rightarrow Y$ is a morphism). Let Γ_π be the scheme-theoretic closure of $\Gamma_{\pi'} \hookrightarrow U \times Y \hookrightarrow X \times Y$, where the first map is a closed embedding (Proposition 11.3.6), and the second is an open embedding. The product here should be taken in the category you are working in.




FIGURE 11.4. Two different maps to a nonseparated scheme agreeing on a dense open set (see Exercise 11.4.D)

For example, if you are working with k -schemes, the fibered product should be taken over k .

11.4.E. EXERCISE. Show that the graph of a rational map $\pi: X \dashrightarrow Y$ is independent of the choice of representative of π . Hint: Suppose $\xi': U \rightarrow Y$ and $\xi: V \rightarrow Y$ are two representatives of π . Reduce to the case where V is the domain of definition of π (§11.4.3), and $\xi' = \xi|_U$. Reduce to the case $V = X$. Show an isomorphism $\Gamma_\pi \cong X$, and $\Gamma_{\xi|_U} \cong U$. (Remark: the separatedness of Y is not necessary.)

In analogy with graphs of morphisms, the following diagram of a graph of a rational map π can be useful (cf. Figure 11.2).

$$\begin{array}{ccc}
 \Gamma_\pi & \xrightarrow{\text{cl. emb.}} & X \times Y \\
 \uparrow \gamma_\pi & \searrow \text{pr}_1 & \searrow \text{pr}_2 \\
 X & \xrightarrow{\pi} & Y
 \end{array}$$

11.4.F. EXERCISE (THE BLOW-UP OF THE PLANE AS THE GRAPH OF A RATIONAL MAP). Consider the rational map $\mathbb{A}_k^2 \dashrightarrow \mathbb{P}_k^1$ given by $(x, y) \mapsto [x, y]$. Show that this rational map cannot be extended over the origin. (A similar argument arises in Exercise 7.5.I on the Cremona transformation.) Show that the graph of the rational map is the morphism (the blow-up) described in Exercise 10.3.F. (When we define blow-ups in general, we will see that they are often graphs of rational maps, see Exercise 23.4.M.)

11.4.5. Variations.

Variations of the short proof of Theorem 11.4.2 yield other useful results. Exercise 11.4.B is one example. Here is another.

11.4.G. EXERCISE (MAPS TO A SEPARATED SCHEME CAN BE EXTENDED OVER AN EFFECTIVE CARTIER DIVISOR IN AT MOST ONE WAY). Suppose X is a Z -scheme (not necessarily reduced!), and Y is a separated Z -scheme. Suppose further that D is an effective Cartier divisor on X . Show that any Z -morphism $X \setminus D \rightarrow Y$ can be extended in at most one way to a Z -morphism $X \rightarrow Y$. (Hint: reduce to the case where $X = \text{Spec } A$, and D is the vanishing scheme of $t \in A$. Reduce to showing that the scheme-theoretic image of $D(t)$ in X is all of X . Show this by showing that $A \rightarrow A_t$ is an inclusion.)

As noted in §7.5.2, rational maps can be defined from any X that has associated points to any Y . The Reduced-to-Separated Theorem 11.4.2 can be extended to this setting, as follows.

11.4.H. EXERCISE (THE “ASSOCIATED-TO-SEPARATED THEOREM”). Prove that two S -morphisms π and π' from a locally Noetherian scheme U to a separated S -scheme Z , agreeing on an open subset containing the associated points of U , are the same.

11.5 Proper morphisms

A map of topological spaces (also known as a continuous map!) is said to be *proper* if the preimage of any compact set is compact. *Properness* of morphisms is an analogous property. For example, a variety over \mathbb{C} will be proper if it is compact in the classical topology (see [Se3, §7]). Alternatively, we will see that projective A -schemes are proper over A — so this is a nice property satisfied by projective schemes, which also is convenient to work with.

Recall (§8.3.5) that a (continuous) map of topological spaces $\pi: X \rightarrow Y$ is *closed* if for each closed subset $S \subset X$, $\pi(S)$ is also closed. A morphism of schemes is closed if the underlying continuous map is closed. We say that a morphism of schemes $\pi: X \rightarrow Y$ is **universally closed** if for every morphism $Z \rightarrow Y$, the induced morphism $Z \times_Y X \rightarrow Z$ is closed. In other words, a morphism is universally closed if it remains closed under any base change. By now you will automatically expect the following exercise.

11.5.A. EASY EXERCISE. Show that the class of universally closed morphisms is “reasonable” in the sense of §8.1.

To motivate the definition of properness for schemes, we remark that a continuous map $\pi: X \rightarrow Y$ of locally compact Hausdorff spaces which have countable bases for their topologies is universally closed if and only if it is proper (i.e., preimages of compact subsets are compact). You are welcome to prove this as an exercise.

11.5.1. Definition. A morphism $\pi: X \rightarrow Y$ is **proper** if it is separated, finite type, and universally closed. A scheme X is often said to be proper if some implicit structure morphism is proper. For example, a k -scheme X is often described as proper if $X \rightarrow \operatorname{Spec} k$ is proper. (A k -scheme is often said to be **complete** if it is proper. We will not use this terminology.) If A is a ring, one often says that an A -scheme is **proper over A** if it is proper over $\operatorname{Spec} A$.

Let’s try this idea out in practice. We expect that $\mathbb{A}_{\mathbb{C}}^1 \rightarrow \operatorname{Spec} \mathbb{C}$ is not proper, because the complex manifold corresponding to $\mathbb{A}_{\mathbb{C}}^1$ is not compact. However, note that this map is separated (it is a map of affine schemes), finite type, and (trivially) closed. So the “universally” is what matters here.

11.5.B. EASY EXERCISE. Show that $\mathbb{A}_{\mathbb{C}}^1 \rightarrow \operatorname{Spec} \mathbb{C}$ is not proper, by finding a base change that turns this into a non-closed map. Possible hint: Consider a well-chosen map such as $\mathbb{A}_{\mathbb{C}}^1 \times_{\mathbb{C}} \mathbb{A}_{\mathbb{C}}^1 \rightarrow \mathbb{A}_{\mathbb{C}}^1$ or $\mathbb{A}_{\mathbb{C}}^1 \times_{\mathbb{C}} \mathbb{P}_{\mathbb{C}}^1 \rightarrow \mathbb{P}_{\mathbb{C}}^1$. (See Figure 20.1 for

another finite type, separated, closed morphism that is not proper. Showing *that* morphism is not proper requires more creativity.)

11.5.2. Example. As a first example: closed embeddings are proper. They are clearly separated, as affine morphisms are separated, Proposition 11.3.4. They are finite type. After base change, they remain closed embeddings (§10.2.2), and closed embeddings are always closed. This easily extends further as follows.

11.5.3. Proposition. — *Finite morphisms are proper.*

Proof. Finite morphisms are separated (as they are affine by definition, and affine morphisms are separated, Proposition 11.3.4), and finite type (basically because finite modules over a ring are automatically finitely generated). To show that finite morphisms are closed after any base change, we note that they remain finite after any base change (finiteness is preserved by base change, Exercise 10.4.C(d)), and finite morphisms are closed (Exercise 8.3.L). \square

11.5.4. Proposition. —

- (a) The notion of “proper morphism” is stable under base change.
- (b) The notion of “proper morphism” is local on the target (i.e., $\pi: X \rightarrow Y$ is proper if and only if for any affine open cover $U_i \rightarrow Y$, $\pi^{-1}(U_i) \rightarrow U_i$ is proper). Note that the “only if” direction follows from (a) — consider base change by $U_i \hookrightarrow Y$.
- (c) The notion of “proper morphism” is closed under composition.
- (d) The product of two proper morphisms is proper: if $\pi: X \rightarrow Y$ and $\pi': X' \rightarrow Y'$ are proper, where all morphisms are morphisms of Z -schemes, then $\pi \times \pi': X \times_Z X' \rightarrow Y \times_Z Y'$ is proper.
- (e) Suppose

(11.5.4.1)

$$\begin{array}{ccc} X & \xrightarrow{\pi} & Y \\ & \searrow \tau & \swarrow \rho \\ & Z & \end{array}$$

is a commutative diagram, and τ is proper, and ρ is separated. Then π is proper.

A sample application of (e): a morphism (over $\text{Spec } k$) from a proper k -scheme to a separated k -scheme is always proper.

Proof. Parts (a) through (c) say that proper morphisms form a “reasonable” class of morphisms in the sense of §8.1. So they follow because separated morphisms, finite type morphisms, and universal closed morphisms are all reasonable classes. Then (d) follows from (a) and (c) by Exercise 8.1.A. Closed embeddings are proper (Example 11.5.2), so (e) follows from the Cancellation Theorem 11.2.1 for proper morphisms. \square

11.5.C. UNIMPORTANT EXERCISE (“IMAGES OF PROPER SCHEMES ARE PROPER”). Suppose in the diagram

$$\begin{array}{ccc} X & \xrightarrow{\pi} & Y \\ & \searrow \tau & \swarrow \rho \\ & Z & \end{array}$$

that τ is proper and ρ is separated and finite type. Show that the scheme-theoretic image of X under π is a proper Z -scheme. (We won’t use this fact, but it reassures us that properness in algebraic geometry behaves like properness in topology.)

We now come to the most important example of proper morphisms.

11.5.5. Theorem. — *Projective A -schemes are proper over A .*

(As finite morphisms to $\operatorname{Spec} A$ are projective A -schemes, Exercise 8.3.I, Theorem 11.5.5 can be used to give a second proof that finite morphisms are proper, Proposition 11.5.3.)

Proof. The structure morphism of a projective A -scheme $X \rightarrow \operatorname{Spec} A$ factors as a closed embedding followed by $\mathbb{P}_A^n \rightarrow \operatorname{Spec} A$. Closed embeddings are proper (Example 11.5.2), and compositions of proper morphisms are proper (Proposition 11.5.4), so it suffices to show that $\mathbb{P}_A^n \rightarrow \operatorname{Spec} A$ is proper. We have already seen that this morphism is finite type (Easy Exercise 5.3.D) and separated (Proposition 11.3.8), so it suffices to show that $\mathbb{P}_A^n \rightarrow \operatorname{Spec} A$ is universally closed. As $\mathbb{P}_A^n = \mathbb{P}_{\mathbb{Z}}^n \times_{\mathbb{Z}} \operatorname{Spec} A$, it suffices to show that $\mathbb{P}_X^n := \mathbb{P}_{\mathbb{Z}}^n \times_{\mathbb{Z}} X \rightarrow X$ is closed for any scheme X . But the property of being closed is local on the target on X , so by covering X with affine open subsets, it suffices to show that $\mathbb{P}_B^n \rightarrow \operatorname{Spec} B$ is closed for all rings B . This is the Fundamental Theorem of Elimination Theory (Theorem 8.4.7). \square

11.5.6. Remark: “Reasonable” proper schemes are projective. It is not easy to come up with an example of an A -scheme that is proper but not projective! Over a field, all proper curves are projective (see Remark 19.7.2), and all smooth surfaces are projective (see [Ba, Thm. 1.28] for a proof of this theorem of Zariski; smoothness of course is not yet defined). We will meet a first example of a proper but not projective variety (a singular threefold) in §17.4.11. We will later see an example of a proper nonprojective surface in §20.11.11, and a simpler one in Exercise 21.2.G. Once we know about flatness, we will see Hironaka’s example of a proper nonprojective irreducible smooth threefold over \mathbb{C} (§25.7.6).

11.5.7. Functions on connected reduced proper \bar{k} -schemes must be constant.

As an enlightening application of these ideas, we show that if X is a connected reduced proper k -scheme where $k = \bar{k}$, then $\Gamma(X, \mathcal{O}_X) = k$. The analogous fact for connected compact complex manifolds uses the maximum principle. We saw this in the special case $X = \mathbb{P}^n$ in Exercise 4.4.E; it will be vastly generalized by Grothendieck’s Coherence Theorem 19.9.1.

Suppose $f \in \Gamma(X, \mathcal{O}_X)$ (f is a function on X). This is the same as a map $\pi: X \rightarrow \mathbb{A}_k^1$ (Exercise 7.3.G, discussed further in §7.6.1). Let π' be the composition of π with the open embedding $\mathbb{A}^1 \hookrightarrow \mathbb{P}^1$. By Proposition 11.5.4(e), π' is proper, and

in particular closed. As X is connected, the image of π' is as well. Thus the set-theoretic image of π' must be either a closed point, or all of \mathbb{P}^1 . But the set-theoretic image of π' lies in \mathbb{A}^1 , so it must be a closed point p (which we identify with an element of k).

By Corollary 9.4.5, the support of the *scheme-theoretic* image of π is the closed point p . By Exercise 9.4.A, the scheme-theoretic image is precisely p (with the reduced structure). Thus π can be interpreted as the structure map to $\text{Spec } k$, followed by a closed embedding to \mathbb{A}^1 identifying $\text{Spec } k$ with p . You should be able to verify that this is the map to \mathbb{A}^1 corresponding to the constant function $f = p$.

(What are counterexamples if different hypotheses are relaxed? Interesting fact: $x^2 = 0$ in \mathbb{P}_k^2 shows that even nonreduced proper k -schemes can have only constants as functions; see §5.2.3 and Exercise 19.6.U.)

11.5.8. Facts (not yet proved) that may help you correctly think about finiteness.

The following facts may shed some light on the notion of finiteness. We will prove them later.

A morphism of locally Noetherian schemes is finite if and only if it is proper and affine, if and only if it is proper and quasifinite (Theorem 30.6.2). We have proved parts of this statement, but we will only finish the proof once we know Zariski's Main Theorem, cf. §8.3.12).

As an application: quasifinite morphisms from proper schemes to separated schemes are finite. Here is why: suppose $\pi: X \rightarrow Y$ is a quasifinite morphism over Z , where X is proper over Z . Then by the Cancellation Theorem 11.2.1 for proper morphisms, $X \rightarrow Y$ is proper. Hence as π is quasifinite and proper, π is finite.

As an explicit example, consider the map $\pi: \mathbb{P}_k^1 \rightarrow \mathbb{P}_k^1$ given by $[x, y] \mapsto [f(x, y), g(x, y)]$, where f and g are homogeneous polynomials of the same degree with no common roots in \mathbb{P}^1 . The fibers are finite, and π is proper (from the Cancellation Theorem 11.2.1 for proper morphisms, as discussed after the statement of Theorem 11.5.4), so π is finite. This could be checked directly as well, but now we can save ourselves the annoyance.

11.5.9. ★★ Group varieties. The rest of §11.5 is double-starred, and we will also throughout work in the category of *varieties over a field* k . We briefly discuss group varieties, mainly because we can, not because we have anything profound to say. We discuss them right now only because properness gives an unexpected consequence, that proper group varieties are abelian (thanks to the surprisingly simple Rigidity Lemma 11.5.12).

As discussed in §11.3.9, now that we have the category of varieties over k , we immediately have the notion of a **group variety** (thanks to §7.6.4 on group objects in any category). An **algebraic group** (over k) is a smooth group variety. Examples include GL_n (Exercise 7.6.N(b)), which includes \mathbb{G}_m as a special case, and SL_n (Exercise 7.6.Q).

11.5.10. Side Remarks (that we won't prove or use). Group varieties are automatically smooth (hence algebraic groups) in characteristic 0. Group varieties are not necessarily smooth in positive characteristic. (Let k be an imperfect field of characteristic p , and choose $t \in k$ without a p th root in k . Consider the "closed sub-group scheme" $x^p = ty^p$ of the "additive group" $\mathbb{G}_a^2 = \text{Spec } k[x, y]$. This group scheme

is reduced, but not geometrically reduced; and not smooth, or even regular.) Algebraic groups are automatically quasiprojective. An algebraic group G is affine (as a scheme) if and only if it admits a closed immersion into GL_n for some $n \geq 0$, such that $G \rightarrow GL_n$ is a homomorphism of group schemes (Exercise 7.6.N(a)). For further discussion on these and other issues, see [P2, §5].

11.5.11. Definition: Abelian varieties. We can now define one of the most important classes of varieties: abelian varieties. The bad news is that we have no real example yet. We will later see elliptic curves as an example (§20.10), but we will not meet anything more exotic than that. Still, it is pleasant to know that we can make the definition this early.

An **abelian variety** over a field k is an algebraic group that is geometrically integral and projective (over k). It turns out that the analytification of any abelian variety over \mathbb{C} is a complex torus, i.e., \mathbb{C}^n/Λ , where Λ is a lattice (of rank $2n$). (The key idea: connected compact complex Lie groups are commutative, see Remark 11.5.14; the universal cover is a simply connected commutative complex Lie group, and thus \mathbb{C}^n . See [BL, Lem. 1.1.1].)


We now show that abelian varieties are abelian, i.e., abelian varieties are commutative algebraic group varieties. (This is less tautological than it sounds! The adjective “abelian” is used in two different senses; what they have in common is that they are derived from the name of Niels Henrik Abel.) As algebraic groups need not be commutative (witness GL_n), it is somehow surprising that commutativity could be forced by “compactness”.

The key result is the following fact, beautiful in its own right.

11.5.12. Rigidity Lemma. — *Let X , Y , and Z be varieties, where X is proper and geometrically integral, with a rational point p , and Y is irreducible. Let $\text{pr}: X \times Y \rightarrow Y$ be the projection, and let $\alpha: X \times Y \rightarrow Z$ be any morphism. Suppose $q \in Y$ is such that α is constant on $X \times \{q\}$; say $\alpha(X \times \{q\}) = r \in Z$.*

(a) *Then there is a morphism $\psi: Y \rightarrow Z$ such that $\alpha = \psi \circ \text{pr}$. (In particular, for every $q' \in Y$, α is constant on $X \times \{q'\}$.)*

(b) *If α is also constant on $\{p\} \times Y$, then α is a constant function on all of $X \times Y$.*



The hypotheses can be relaxed in a number of ways, but this version suffices for our purposes. If you have not read §10.5.1, where geometric integrality is discussed, you need not worry too much; this is automatic if $k = \bar{k}$ and X is integral (Hard Fact 10.5.3). Even if you *have* read §10.5.1, it is helpful to first read this proof under the assumption that $k = \bar{k}$, to avoid being distracted from the main idea by geometric points. But of course all of §11.5.9 is double-starred, so you shouldn't be reading this anyway.


11.5.D. EXERCISE. Show that the properness hypothesis on X is necessary. (Can you make this accord with your intuition?)

Proof of the Rigidity Lemma 11.5.12. Define $\beta: X \times Y \rightarrow Z$ by $\beta(x, y) = \alpha(p, y)$. (More precisely, if σ_p is the section of $X \times Y \rightarrow Y$ pulled back from $\text{Spec } k \mapsto p$, then $\beta = \alpha \circ \sigma_p \circ \text{pr}$.) We will show that $\beta(x, y) = \alpha(x, y)$. This will imply (a) (take $\psi = \alpha \circ \sigma_p$), which in turn immediately implies (b).

Proposition 10.5.19 (geometrically reduced times reduced is reduced) implies that $X \times Y$ is reduced. Exercise 10.5.O (geometrically irreducible times irreducible is irreducible) implies that $X \times Y$ is irreducible. As Z is separated, by Exercise 11.4.A, it suffices to show that $\alpha = \beta$ on a nonempty (hence dense) open subset of $X \times Y$.

Let $U \subset Z$ be an affine open neighborhood of r . Then $\alpha^{-1}(Z \setminus U)$ is a closed subset of $X \times Y$. As the projection pr is proper (using properness of X , and preservation of properness under base change, Proposition 11.5.4(a)), we have $\text{pr}(\alpha^{-1}(Z \setminus U))$ is closed. Its complement is open, and contains q ; let $V \subset Y$ be an open neighborhood of q disjoint from $\text{pr}(\alpha^{-1}(Z \setminus U))$.

11.5.E. EXERCISE. Suppose $\gamma: \text{Spec } \bar{K} \rightarrow V$ is a geometric point of V . Show that the fiber $X_{\bar{K}}$ over $\text{Spec } \bar{K}$ (or more precisely, over γ) is mapped to a point in Z by the restriction $\alpha_{\bar{K}}$ of α to $X_{\bar{K}}$. Hint: review §11.5.7.



By Exercise 11.4.B (maps of varieties over an algebraically closed field are determined by the map of closed points), $\alpha_{\bar{K}} = \beta_{\bar{K}}$.

11.5.F. EXERCISE. Use Exercise 11.4.A (the fact that the locus where $\alpha = \beta$ is a closed subscheme — and represents a certain functor) to show that $\alpha = \beta$ on $\text{pr}^{-1}(V)$. □

11.5.13. Corollary. — Suppose A is an abelian variety, G is a group variety, and $\phi: A \rightarrow G$ is any morphism of varieties. Then ϕ is the composition of a translation and a homomorphism.

Proof. Let e_A and e_G be the identity points of A and G respectively. Composing ϕ with a translation, we may assume that $\phi(e_A) = e_G$. Consider the morphism $\alpha: A \times A \rightarrow G$ given by $\alpha(a_1, a_2) = \phi(a_1 a_2) \phi(a_1)^{-1} \phi(a_2)^{-1}$. Then $\alpha(\{e_A\} \times A) = \alpha(A \times \{e_A\}) = \{e_G\}$, so by the Rigidity Lemma 11.5.12, α is a constant, and sends $A \times A$ to e_G . Thus $\phi(a_1 a_2) = \phi(a_1) \phi(a_2)$, so ϕ is a homomorphism. □

11.5.G. EXERCISE (ABELIAN VARIETIES ARE ABELIAN). Show that an abelian variety is an abelian group variety. Hint: apply Corollary 11.5.13 to the inversion morphism $i: A \rightarrow A$.

11.5.14. Remark. A similar idea is used to show that connected compact complex Lie groups are abelian, see for example [BL, Lem. 1.1.1].

Part IV

“Geometric” properties of schemes

CHAPTER 12

Dimension

12.1 Dimension and codimension

Everyone knows what a curve is, until he has studied enough mathematics to become confused ...

— F. Klein, [RS, p. 90]

At this point, you know a fair bit about schemes, but there are some fundamental notions you cannot yet define. In particular, you cannot use the phrase “smooth surface”, as it involves the notion of dimension and of smoothness. You may be surprised that we have gotten so far without using these ideas. You may also be disturbed to find that these notions can be subtle, but you should keep in mind that they are subtle in all parts of mathematics.

In this chapter, we will address the first notion, that of dimension of schemes. This should agree with, and generalize, our geometric intuition. Although we think of dimension as a basic notion in geometry, it is a slippery concept, as it is throughout mathematics. Even in linear algebra, the definition of dimension of a vector space is surprising the first time you see it, even though it quickly becomes second nature. The definition of dimension for manifolds is equally nontrivial. For example, how do we know that there isn’t an isomorphism between some 2-dimensional manifold and some 3-dimensional manifold? Your answer will likely use topology, and hence you should not be surprised that the notion of dimension is often quite topological in nature.

A caution for those thinking over the complex numbers: our dimensions will be algebraic, and hence half that of the “real” picture. For example, we will see very shortly that $\mathbb{A}_{\mathbb{C}}^1$, which you may picture as the complex numbers (plus one generic point), has dimension 1.

12.1.1. Definition(s): dimension. Surprisingly, the right definition is purely topological — it just depends on the topological space, and not on the structure sheaf. We define the **dimension** of a topological space X (denoted $\dim X$) as the supremum of lengths of chains of closed irreducible sets, starting the indexing of the closed irreducible sets with 0. (The dimension may be infinite.) Scholars of the empty set can take the dimension of the empty set to be $-\infty$. (An analogy from linear algebra: the dimension of a vector space is the supremum of the length of chains of subspaces.) Define the **dimension** of a ring as the supremum of the lengths of the chains of nested prime ideals (where indexing starts at zero). These two definitions of dimension are sometimes called **Krull dimension**. (You might think a

Noetherian ring has finite dimension because all chains of prime ideals are finite, but this isn't necessarily true — see Exercise 12.1.K.)

12.1.A. EASY EXERCISE. Show that $\dim \operatorname{Spec} A = \dim A$. (Hint: Exercise 3.7.F gives a bijection between irreducible closed subsets of $\operatorname{Spec} A$ and prime ideals of A . It is “inclusion-reversing”.)

The homeomorphism between $\operatorname{Spec} A$ and $\operatorname{Spec} A/\mathfrak{N}(A)$ (§3.4.7: the Zariski topology disregards nilpotents) implies that $\dim \operatorname{Spec} A = \dim \operatorname{Spec} A/\mathfrak{N}(A)$.

12.1.2. Examples. We have identified all the prime ideals of $k[t]$ (they are 0, and $(f(t))$ for irreducible polynomials $f(t)$), \mathbb{Z} ((0) and (p)), k (only (0)), and $k[x]/(x^2)$ (only (x)), so we can quickly check that $\dim \mathbb{A}_k^1 = \dim \operatorname{Spec} \mathbb{Z} = 1$, $\dim \operatorname{Spec} k = 0$, $\dim \operatorname{Spec} k[x]/(x^2) = 0$.

12.1.3. We must be careful with the notion of dimension for reducible spaces. If Z is the union of two closed subsets X and Y , then $\dim Z = \max(\dim X, \dim Y)$. Thus dimension is not a “local” characteristic of a space. This sometimes bothers us, so we try to only talk about dimensions of irreducible topological spaces. We say a topological space is **equidimensional** or **pure dimensional** (resp. **equidimensional of dimension n** or **pure dimension n**) if each of its irreducible components has the same dimension (resp. they are all of dimension n). An equidimensional dimension 1 (resp. 2, n) topological space is said to be a **curve** (resp. **surface**, **n -fold**).

12.1.B. USEFUL EXERCISE. Show that a scheme has dimension n if and only if it admits an open cover by affine open subsets of dimension at most n , where equality is achieved for some affine open subset. Hint: You may find it helpful, here and later, to show the following. For any topological space X and open subset $U \subset X$, there is a bijection between irreducible closed subsets of U , and irreducible closed subsets of X that meet U .

12.1.C. EASY EXERCISE. Show that a Noetherian scheme of dimension 0 has a finite number of points.

12.1.D. EXERCISE (FIBERS OF INTEGRAL MORPHISMS, PROMISED IN §8.3.8). Suppose $\pi: X \rightarrow Y$ is an integral morphism. Show that every (nonempty) fiber of π has dimension 0. Hint: As integral morphisms are preserved by base change, we assume that $Y = \operatorname{Spec} k$. Hence we must show that if $\phi: k \rightarrow A$ is an integral extension, then $\dim A = 0$. Outline of proof: Suppose $\mathfrak{p} \subset \mathfrak{m}$ are two prime ideals of A . Mod out by \mathfrak{p} , so we can assume that A is an integral domain. I claim that any nonzero element is invertible: Say $x \in A$, and $x \neq 0$. Then the minimal monic polynomial for x has nonzero constant term. But then x is invertible — recall the coefficients are in a field.

12.1.E. IMPORTANT EXERCISE. Show that if $\pi: \operatorname{Spec} A \rightarrow \operatorname{Spec} B$ corresponds to an integral extension of rings, then $\dim \operatorname{Spec} A = \dim \operatorname{Spec} B$. Hint: show that a chain of prime ideals downstairs gives a chain upstairs of the same length, by the Going-Up Theorem (Exercise 8.2.F). Conversely, a chain upstairs gives a chain downstairs. Use Exercise 12.1.D to show that no two elements of the chain upstairs go to the same element $[q] \in \operatorname{Spec} B$ of the chain downstairs.

12.1.F. EXERCISE. Show that if $v: \tilde{X} \rightarrow X$ is the normalization of a scheme (possibly in a finite extension of the function field of X), then $\dim \tilde{X} = \dim X$. Feel free to assume that X is integral for convenience.

12.1.G. EXERCISE. Suppose X is an affine k -scheme, and K/k is an *algebraic* field extension.

(a) Suppose X has pure dimension n . Show that $X_K := X \times_k K$ has pure dimension n . (See Exercise 25.5.F for a generalization, which for example removes the affine hypothesis. Also, see Exercise 12.2.K and Remark 12.2.16 for the fate of possible generalizations to arbitrary field extensions.) Hint: If $X = \operatorname{Spec} A$, reduce to the case where A is an integral domain. An irreducible component of X_K corresponds to a minimal prime \mathfrak{p} of $A_K := A \otimes_k K$. Suppose $\alpha \in \ker(A \rightarrow A_K/\mathfrak{p})$. Show that $\alpha = 0$, using the fact that α lies in a minimal prime \mathfrak{p} of A_K (and is hence a zerodivisor, by §6.5.12 combined with §6.5.22), and A_K is a free A -module (so multiplication in A_K by $\alpha \in A$ is injective if α is nonzero). Thus $A \rightarrow A_K/\mathfrak{p}$ is injective. Then use Exercise 12.1.E.

(b) Prove the converse to (a): show that if X_K has pure dimension n , then X has pure dimension n .

12.1.H. EXERCISE. Show that $\dim \mathbb{Z}[x] = 2$. (Hint: The prime ideals of $\mathbb{Z}[x]$ were implicitly determined in Exercise 3.2.Q.)

12.1.4. Codimension. Because dimension behaves oddly for disjoint unions, we need some care when defining codimension, and in using the phrase. For example, if Y is a closed subset of X , we might define the codimension to be $\dim X - \dim Y$, but this behaves badly. For example, if X is the disjoint union of a point Y and a curve Z , then $\dim X - \dim Y = 1$, but this has nothing to do with the local behavior of X near Y .

A better definition is as follows. In order to avoid excessive pathology, we define the codimension of Y in X *only when Y is irreducible*. (Use extreme caution when using this word in any other setting.) Define the **codimension of an irreducible subset** $Y \subset X$ of a topological space as the supremum of lengths of *increasing* chains of irreducible closed subsets starting with \bar{Y} (where indexing starts at 0 — recall that the closure of an irreducible set is irreducible, Exercise 3.6.B(b)). In particular, the **codimension of a point** is the codimension of its closure. The codimension of Y in X is denoted by $\operatorname{codim}_X Y$.

We say that a prime ideal \mathfrak{p} in a ring has **codimension** equal to the supremum of lengths of the chains of decreasing prime ideals starting at \mathfrak{p} , with indexing starting at 0. Thus in an integral domain, the ideal (0) has codimension 0; and in \mathbb{Z} , the ideal (23) has codimension 1. Note that the codimension of the prime ideal \mathfrak{p} in A is $\dim A_{\mathfrak{p}}$ (see §3.2.9). (This notion is often called **height**.) Thus the codimension of \mathfrak{p} in A is the codimension of $[\mathfrak{p}]$ in $\operatorname{Spec} A$.

(Continuing an analogy with linear algebra: the codimension of a vector subspace $Y \subset X$ is the supremum of lengths of increasing chains of subspaces starting with Y . This is a better definition than $\dim X - \dim Y$, because it works even when $\dim X = \infty$. You might prefer to define $\operatorname{codim}_X Y$ as $\dim(X/Y)$; that is analogous to defining the codimension of \mathfrak{p} in A as the dimension of $A_{\mathfrak{p}}$ — see the previous paragraph.)

12.1.I. EXERCISE. Show that if Y is an irreducible closed subset of a scheme X , and η is the generic point of Y , then the codimension of Y is the dimension of the local ring $\mathcal{O}_{X,\eta}$.

12.1.5. Notice that Y is codimension 0 in X if it is an irreducible component of X . Similarly, Y is codimension 1 if it is not an irreducible component, and for every irreducible component Y' it is contained in, there is no irreducible closed subset strictly between Y and Y' . (See Figure 12.1 for examples.) A closed subset all of whose irreducible components are codimension 1 in some ambient space X is said to be a **hypersurface** in X . (Be careful: the word “hypersurface” can be used to mean slightly different things in algebraic geometry, for entirely reasonable reasons, see for example Definition 9.3.2.)




FIGURE 12.1. Behavior of codimension

12.1.J. EASY EXERCISE. If Y is an irreducible closed subset of a scheme X , show that

$$(12.1.5.1) \quad \text{codim}_X Y + \dim Y \leq \dim X.$$

We will soon see that equality always holds if X and Y are irreducible varieties (Theorem 12.2.9), but equality doesn't hold in general (§12.3.19).

Warning. The notion of codimension still can behave slightly oddly. For example, consider Figure 12.1. (You should think of this as an intuitive sketch.) Here the total space X has dimension 2, but point p is dimension 0, and codimension 1. We also have an example of a codimension 2 subset q contained in a codimension 0 subset C with no codimension 1 subset “in between”.

Worse things can happen; we will soon see an example of a closed point in an *irreducible* surface that is nonetheless codimension 1, not 2, in §12.3.19. However, for irreducible *varieties* this can't happen, and inequality (12.1.5.1) must be an equality (Theorem 12.2.9).

12.1.6. In unique factorization domains, codimension 1 prime ideals are principal. We make a short observation, for future reference.

12.1.7. Lemma. — *In a unique factorization domain A , all codimension 1 prime ideals are principal.*

This is a first glimpse of the fact that codimension one is rather special — this theme will continue in §12.3. We will see that the converse of Lemma 12.1.7 holds as well (when A is a Noetherian integral domain, Proposition 12.3.5).

Proof. Suppose \mathfrak{p} is a codimension 1 prime. Choose any $f \neq 0$ in \mathfrak{p} , and let g be any irreducible/prime factor of f that is in \mathfrak{p} (there is at least one). Then (g) is a nonzero prime ideal contained in \mathfrak{p} , so $(0) \subset (g) \subset \mathfrak{p}$. As \mathfrak{p} is codimension 1, we must have $\mathfrak{p} = (g)$, and thus \mathfrak{p} is principal. \square

12.1.8. A fun but unimportant counterexample. We end this introductory section with a fun pathology. As a Noetherian ring has no infinite chain of prime ideals, you may think that Noetherian rings must have finite dimension. Nagata, the master of counterexamples, shows you otherwise with the following example.

12.1.K. ★★ EXERCISE: AN INFINITE-DIMENSIONAL NOETHERIAN RING. Let $A = k[x_1, x_2, \dots]$. Choose an increasing sequence of positive integers m_1, m_2, \dots whose differences are also increasing ($m_{i+1} - m_i > m_i - m_{i-1}$). Let $\mathfrak{p}_i = (x_{m_i+1}, \dots, x_{m_{i+1}})$ and $S = A - \cup_i \mathfrak{p}_i$.

- (a) Show that S is a multiplicative set.
- (b) Show that $S^{-1}\mathfrak{p}_i$ in $S^{-1}A$ is the largest prime ideal in a chain of prime ideals of length $m_{i+1} - m_i$. Hence conclude that $\dim S^{-1}A = \infty$.
- (c) Suppose B is a ring such that (i) for every maximal ideal \mathfrak{m} , $B_{\mathfrak{m}}$ is Noetherian, and (ii) every nonzero $b \in B$ is contained in finitely many maximal ideals. Show that B is Noetherian. (One possible approach: show that for any x_1, x_2, \dots , (x_1, x_2, \dots) is finitely generated.)
- (d) Use (c) to show that $S^{-1}A$ is Noetherian.

12.1.9. Remark: Noetherian local rings have finite dimension. However, we shall see in Exercise 12.3.H(a) that Noetherian *local* rings always have finite dimension. (This requires a surprisingly hard fact, Krull's Height Theorem 12.3.9.) Thus points of locally Noetherian schemes always have finite codimension.

12.2 Dimension, transcendence degree, and Noether normalization

We now give a powerful alternative interpretation for dimension for irreducible varieties, in terms of transcendence degree. The proof will involve a classical construction, *Noether normalization*, which will be useful in other ways as well. In case you haven't seen transcendence theory, here is a lightning introduction.

12.2.A. EXERCISE/DEFINITION. Recall that an *element* of a field extension E/F is *algebraic over F* if it is integral over F . Recall also that a *field extension* E/F is an *algebraic extension* if it is an integral extension (if all elements of E are algebraic

over F , cf. §8.2). The composition of two algebraic extensions is algebraic, by Exercise 8.2.C. If E/F is a field extension, and F' and F'' are two intermediate field extensions, then we write $F' \sim F''$ if $F'F''$ is algebraic over both F' and F'' . Here $F'F''$ is the *compositum* of F' and F'' , the smallest field extension in E containing F' and F'' .

(a) Show that \sim is an equivalence relation on subextensions of E/F . A **transcendence basis** of E/F is a set of elements $\{x_i\}$ that are algebraically independent over F (there is no nontrivial polynomial relation among the x_i with coefficients in F) such that $F(\{x_i\}) \sim E$.

(b) Show that if E/F has two transcendence bases, and one has cardinality n , then both have cardinality n . (Hint: show that you can substitute elements from the one basis into the other one at a time.) The size of any transcendence basis is called the **transcendence degree** (which may be ∞), and is denoted trdeg . Any finitely generated field extension necessarily has finite transcendence degree. (Remark: A related result was mentioned in Algebraic Fact 10.5.18.)

12.2.1. Theorem (dimension = transcendence degree). — Suppose A is a **finitely generated domain over a field k** (i.e., a finitely generated k -algebra that is an integral domain). Then $\dim \text{Spec } A = \text{trdeg } K(A)/k$. Hence if X is an irreducible k -variety, then $\dim X = \text{trdeg } K(X)/k$.

We will prove Theorem 12.2.1 shortly (§12.2.7). We first show that it is useful by giving some immediate consequences. We seem to have immediately $\dim \mathbb{A}_k^n = n$. However, our proof of Theorem 12.2.1 will go *through* this fact, so it isn't really a consequence.

A more substantive consequence is the following. If X is an irreducible k -variety, then $\dim X$ is the transcendence degree of the function field (the stalk at the generic point) $\mathcal{O}_{X,\eta}$ over k . Thus (as the generic point lies in all nonempty open sets) the dimension can be computed in any open set of X . (Warning: this is false without the finite type hypothesis, even in quite reasonable circumstances: let X be the two-point space $\text{Spec } k[x]_{(x)}$, and let U consist of only the generic point, see Exercise 3.4.K.)

Another consequence is a second proof of the Nullstellensatz 3.2.6.

12.2.B. EXERCISE: THE NULLSTELLENSATZ FROM DIMENSION THEORY. Suppose $A = k[x_1, \dots, x_n]/I$. Show that the residue field of any maximal ideal of A is a finite extension of k . (Hint: the maximal ideals correspond to dimension 0 points, which correspond to transcendence degree 0 finitely generated extensions of k , i.e., finite extensions of k .)

Yet another consequence is geometrically believable.

12.2.C. EXERCISE. If $\pi: X \rightarrow Y$ is a dominant morphism of irreducible k -varieties, then $\dim X \geq \dim Y$. (This is false more generally: consider the inclusion of the generic point into an irreducible curve.)

12.2.D. EXERCISE. Suppose $f_1(x_1, \dots, x_n), \dots, f_m(x_1, \dots, x_n) \in k[x_1, \dots, x_n]$ are m polynomials in n variables over a field, with $m > n$. Show that the polynomials f_1, \dots, f_m are algebraically dependent.

12.2.E. EXERCISE (PRACTICE WITH THE CONCEPT). Show that the three equations

$$wz - xy = 0, \quad wy - x^2 = 0, \quad xz - y^2 = 0$$

cut out an integral *surface* S in \mathbb{A}_k^4 . (You may recognize these equations from Exercises 3.6.F and 9.3.A.) You might expect S to be a curve, because it is cut out by three equations in four-space. One of many ways to proceed: cut S into pieces. For example, show that $D(w) \cong \text{Spec } k[x, w]_w$. (You may recognize S as the affine cone over the twisted cubic. The twisted cubic was defined in Exercise 9.3.A.) It turns out that you need three equations to cut out this surface. The first equation cuts out a threefold in \mathbb{A}_k^4 (by Krull's Principal Ideal Theorem 12.3.3, which we will meet soon). The second equation cuts out a surface: our surface, along with another surface. The third equation cuts out our surface, and removes the “extraneous component”. One last aside: notice once again that the cone over the quadric surface $k[w, x, y, z]/(wz - xy)$ makes an appearance.)

12.2.2. Definition: degree of a rational map of irreducible varieties. If $\pi: X \dashrightarrow Y$ is a dominant rational map of irreducible (hence integral) k -varieties of the same dimension, the degree of the field extension is called the **degree** of the rational map. This readily extends if X is reducible: we add up the degrees on each of the components of X . If π is a rational map of integral affine k -varieties of the same dimension that is *not* dominant we say the degree is 0. We will interpret this degree in terms of counting preimages of general points of Y in §25.5.13. Note that degree is multiplicative under composition: if $\rho: Y \dashrightarrow Z$ is a rational map of integral k -varieties of the same dimension, then $\deg(\rho \circ \pi) = \deg(\rho) \deg(\pi)$, as degrees of field extensions are multiplicative in towers.

12.2.3. Noether normalization.

Our proof of Theorem 12.2.1 will use another important classical notion, Noether normalization.

12.2.4. Noether Normalization Lemma. — Suppose A is an integral domain, finitely generated over a field k . If $\text{trdeg } K(A)/k = n$, then there are elements $x_1, \dots, x_n \in A$, algebraically independent over k , such that A is a finite (hence integral by Corollary 8.2.2) extension of $k[x_1, \dots, x_n]$.

The geometric content behind this result is that given any integral finite type affine k -scheme X , we can find a surjective finite morphism $X \rightarrow \mathbb{A}_k^n$, where n is the transcendence degree of the function field of X (over k). Surjectivity follows from the Lying Over Theorem 8.2.5, in particular Exercise 12.1.E. This interpretation is sometimes called *geometric Noether normalization*. (See Remark 31.3.7 for a projective version of Noether normalization.)

12.2.5. Nagata's proof of the Noether Normalization Lemma 12.2.4. Suppose we can write $A = k[y_1, \dots, y_m]/\mathfrak{p}$, i.e., that A can be chosen to have m generators. Note that $m \geq n$. We show the result by induction on m . The base case $m = n$ is immediate.

Assume now that $m > n$, and that we have proved the result for smaller m . We will find $m - 1$ elements z_1, \dots, z_{m-1} of A such that A is finite over $k[z_1, \dots, z_{m-1}]/\mathfrak{q}$ (i.e., the subring of A generated by z_1, \dots, z_{m-1} , where the z_i satisfy the relations given by the ideal \mathfrak{q}). Then by the inductive hypothesis,

$k[z_1, \dots, z_{m-1}]/q$ is finite over some $k[x_1, \dots, x_n]$, and A is finite over $k[z_1, \dots, z_{m-1}]$, so by Exercise 8.3.H, A is finite over $k[x_1, \dots, x_n]$.

$$\begin{array}{c}
 A = k[y_1, \dots, y_m]/p \\
 \left| \text{finite} \right. \\
 k[z_1, \dots, z_{m-1}]/q \\
 \left| \text{finite} \right. \\
 k[x_1, \dots, x_n]
 \end{array}$$

As y_1, \dots, y_m are algebraically dependent in A , there is some nonzero algebraic relation $f(y_1, \dots, y_m) = 0$ among them (where f is a polynomial in m variables over k).

Let $z_1 = y_1 - y_m^{r_1}$, $z_2 = y_2 - y_m^{r_2}$, \dots , $z_{m-1} = y_{m-1} - y_m^{r_{m-1}}$, where r_1, \dots, r_{m-1} are positive integers to be chosen shortly. Then

$$f(z_1 + y_m^{r_1}, z_2 + y_m^{r_2}, \dots, z_{m-1} + y_m^{r_{m-1}}, y_m) = 0.$$

Then upon expanding this out, each monomial in f (as a polynomial in m variables) will yield a single term that is a constant times a power of y_m (with no z_i factors). By choosing the r_i so that $0 \ll r_1 \ll r_2 \ll \dots \ll r_{m-1}$, we can ensure that the powers of y_m appearing are all distinct, and so that in particular there is a leading term y_m^N , and all other terms (including those with factors of z_i) are of smaller degree in y_m . Thus we have described an integral dependence of y_m on z_1, \dots, z_{m-1} as desired. \square

12.2.6. The geometry behind Nagata's proof. Here is the geometric intuition behind Nagata's argument. Suppose we have an n -dimensional variety in \mathbb{A}_k^m with $n < m$, for example $xy = 1$ in \mathbb{A}^2 . One approach is to hope the projection to a hyperplane is a finite morphism. In the case of $xy = 1$, if we projected to the x -axis, it wouldn't be finite, roughly speaking because the asymptote $x = 0$ prevents the map from being closed (cf. Exercise 8.3.K). If we instead projected to a random line, we might hope that we would get rid of this problem, and indeed we usually can: this problem arises for only a finite number of directions. But we might have a problem if the field were finite: perhaps the finite number of directions in which to project each have a problem. (You can show that if k is an infinite field, then the substitution in the above proof $z_i = y_i - y_m^{r_i}$ can be replaced by the linear substitution $z_i = y_i - a_i y_m$ where $a_i \in k$, and that for a nonempty Zariski-open choice of a_i , we indeed obtain a finite morphism.) Nagata's trick in general is to "jiggle" the variables in a nonlinear way, and this jiggling kills the nonfiniteness of the map.

12.2.F. EXERCISE (DIMENSION IS ADDITIVE FOR PRODUCTS OF FINITE TYPE k -SCHEMES).

If X and Y are irreducible k -varieties, show that $\dim X \times_k Y = \dim X + \dim Y$. Hint: If we had surjective finite morphisms $X \rightarrow \mathbb{A}_k^{\dim X}$ and $Y \rightarrow \mathbb{A}_k^{\dim Y}$, we could construct a surjective finite morphism $X \times_k Y \rightarrow \mathbb{A}_k^{\dim X + \dim Y}$. (Warning: this can be false without the finite type assumption, see Remark 12.2.16.)

In particular, we have $\dim A[x] = \dim A + 1$ when A is finitely generated over a field. (This is true more generally when A is Noetherian, although we will not prove this.)

12.2.7. Proof of Theorem 12.2.1 on dimension and transcendence degree. Suppose X is an integral affine k -scheme of finite type. We show that $\dim X$ equals the transcendence degree n of its function field, by induction on n . (The idea is that we reduce from X to \mathbb{A}^n to a hypersurface in \mathbb{A}^n to \mathbb{A}^{n-1} .) Assume the result is known for all transcendence degrees less than n .

By the Noether Normalization Lemma 12.2.4, there exists a surjective finite morphism $X \rightarrow \mathbb{A}_k^n$. By Exercise 12.1.E, $\dim X = \dim \mathbb{A}_k^n$. If $n = 0$, we are done, as $\dim \mathbb{A}_k^0 = 0$.

We now show that $\dim \mathbb{A}_k^n = n$ for $n > 0$, by induction. Clearly $\dim \mathbb{A}_k^n \geq n$, as we can describe a chain of irreducible subsets of length n : if x_1, \dots, x_n are coordinates on \mathbb{A}^n , consider the chain of ideals

$$(0) \subset (x_1) \subset \cdots \subset (x_1, \dots, x_n)$$

in $k[x_1, \dots, x_n]$. Suppose we have a chain of prime ideals of length at least $n + 1$:

$$(0) = \mathfrak{p}_0 \subset \cdots \subset \mathfrak{p}_m.$$

Choose any nonzero element g of \mathfrak{p}_1 , and let f be any irreducible factor of g lying in \mathfrak{p}_1 . Then replace \mathfrak{p}_1 by (f) . (Of course, \mathfrak{p}_1 may have been (f) to begin with.) Then $K(k[x_1, \dots, x_n]/(f(x_1, \dots, x_n)))$ has transcendence degree $n - 1$, so by induction,

$$\dim k[x_1, \dots, x_n]/(f) = n - 1,$$

so $\dim k[x_1, \dots, x_n] \leq n$, completing the proof. \square

12.2.8. Codimension is the difference of dimensions for irreducible varieties.

Noether normalization will help us show that codimension is the difference of dimensions for irreducible varieties, i.e., that the inequality (12.1.5.1) is always an equality.

12.2.9. Theorem. — Suppose X is a pure-dimensional k -scheme locally of finite type (for example, an irreducible k -variety), Y is an irreducible closed subset, and η is the generic point of Y . Then $\dim Y + \dim \mathcal{O}_{X,\eta} = \dim X$. Hence by Exercise 12.1.I, $\dim Y + \text{codim}_X Y = \dim X$ — inequality (12.1.5.1) is always an equality.

Proving this will give us an excuse to introduce some useful notions, such as the Going-Down Theorem for finite extensions of integrally closed domains (Theorem 12.2.12). Before we begin the proof, we give an algebraic translation.

12.2.G. EXERCISE. A ring A is called **catenary** if for every nested pair of prime ideals $\mathfrak{p} \subset \mathfrak{q} \subset A$, all maximal chains of prime ideals between \mathfrak{p} and \mathfrak{q} have the same length. (We will not use this term in any serious way later.) Show that if A is a localization of a finitely generated ring over a field k , then A is catenary.

12.2.10. Remark. Most rings arising naturally in algebraic geometry are catenary. Important examples include: localizations of finitely generated \mathbb{Z} -algebras; complete Noetherian local rings; Dedekind domains; and Cohen-Macaulay rings (see §29.2.13). It is hard to give an example of a noncatenary ring; see for example [Stacks, tag 02JE] or [He].

12.2.11. Proof of Theorem 12.2.9.

12.2.H. EXERCISE. Reduce the proof of Theorem 12.2.9 to the following problem. If X is an irreducible affine k -variety and Z is a closed irreducible subset maximal among those smaller than X (the only larger closed irreducible subset is X), then $\dim Z = \dim X - 1$.

Let $d = \dim X$ for convenience. By Noether normalization 12.2.4, we have a finite morphism $\pi: X \rightarrow \mathbb{A}^d$ corresponding to a finite extension of rings. Then $\pi(Z)$ is an irreducible closed subset of \mathbb{A}^d (finite morphisms are closed, Exercise 8.3.L).

12.2.I. EXERCISE. Show that it suffices to show that $\pi(Z)$ is a hypersurface. (Hint: the dimension of any hypersurface is $d - 1$ by Theorem 12.2.1 on dimension and transcendence degree. Exercise 12.1.E implies that $\dim \pi^{-1}(\pi(Z)) = \dim \pi(Z)$. But be careful: Z is not $\pi^{-1}(\pi(Z))$ in general.)

Now if $\pi(Z)$ is not a hypersurface, then it is properly contained in an irreducible hypersurface H , so by the Going-Down Theorem 12.2.12 for finite extensions of integrally closed domains (which we shall now prove), there is some closed irreducible subset Z' of X properly containing Z , contradicting the maximality of Z . \square

12.2.12. Theorem (Going-Down Theorem for finite extensions of integrally closed domains). — Suppose $\phi: B \hookrightarrow A$ is a finite extension of rings (so A is a finite B -module), B is an integrally closed domain, and A is an integral domain. Then given nested prime ideals $\mathfrak{q} \subset \mathfrak{q}'$ of B , and a prime \mathfrak{p}' of A lying over \mathfrak{q}' (i.e., $\mathfrak{p}' \cap B = \mathfrak{q}'$), there exists a prime \mathfrak{p} of A contained in \mathfrak{p}' , lying over \mathfrak{q} .

As usual, you should sketch a geometric picture of the statement. This theorem is usually stated about extending a chain of ideals, in the same way as the Going-Up Theorem (Exercise 8.2.F), and you may want to think this through. (Another Going-Down Theorem, for flat morphisms, will be given in Exercise 25.5.E.)

This theorem is true more generally with “finite” replaced by “integral”; see [E, p. 291] (“Completion of the proof of 13.9”) for the extension of Theorem 12.2.12, or else see [AtM, Thm. 5.16] or [Mat1, Thm. 5(v)] for completely different proofs. See [E, Fig. 10.4] for an example (in the form of a picture) of why the “integrally closed” hypothesis on B cannot be removed.


In the course of the proof, we will use the following easy fact. We could prove it now, but we leave it until Exercise 12.3.C, because the proof uses a trick arising in Exercise 12.3.B.

12.2.13. Proposition (prime avoidance). — Suppose $\mathfrak{p}_1, \dots, \mathfrak{p}_n$ are prime ideals of a ring A , and I is another ideal of A not contained in any \mathfrak{p}_i . Then I is not contained in $\cup \mathfrak{p}_i$: there is an element $f \in I$ not in any of the \mathfrak{p}_i .

(Can you give a geometric interpretation of this result? Can you figure out why it is called “prime avoidance”?)

12.2.14. Proof of Theorem 12.2.12 (Going-Down Theorem for finite extensions of integrally closed domains). The proof uses Galois theory. Let L be the normal closure of $K(A)/K(B)$ (the smallest subfield of $\overline{K(B)}$ containing $K(A)$, and that is mapped to

itself by any automorphism over $\overline{K(B)}/K(B)$). Let C be the integral closure of B in L (discussed in Exercise 10.7.I). Because $A \hookrightarrow C$ is an integral extension, there is a prime Q of C lying over $q \subset B$ (by the Lying Over Theorem 8.2.5), and a prime Q' of C containing Q lying over q' (by the Going-Up Theorem, Exercise 8.2.F). Similarly, there is a prime P' of C lying over $p' \subset A$ (and thus over $q' \subset B$). We would be done if $P' = Q'$ (just take $p = Q \cap A$), but this needn't be the case. However, Lemma 12.2.15 below shows there is an automorphism σ of C over B , that sends Q' to P' , and then the image of $\sigma(Q)$ in A will do the trick, completing the proof. The following diagram, in geometric terms, may help.



12.2.15. Lemma. — Suppose B is an integrally closed domain, $L/K(B)$ is a finite normal field extension, and C is the integral closure of B in L . If q' is a prime ideal of B , then automorphisms of $L/K(B)$ act transitively on the prime ideals of C lying over q' .

This result is often first seen in number theory, with $B = \mathbb{Z}$ and L a Galois extension of \mathbb{Q} .

Proof. Let P and Q_1 be two prime ideals of C lying over q' , and let Q_2, \dots, Q_n be the prime ideals of C conjugate to Q_1 (the image of Q_1 under $\text{Aut}(L/K(B))$). If P is not one of the Q_i , then P is not contained in any of the Q_i . (Do you see why? Hint: Exercise 12.1.D or 12.1.E.) Hence by prime avoidance (Proposition 12.2.13), P is not contained in their union, so there is some $a \in P$ not contained in any Q_i . Thus no conjugate of a can be contained in Q_1 , so the norm $N_{L/K(B)}(a) \in B$ is not contained in $Q_1 \cap B = q'$. (Recall: If $L/K(B)$ is separable, the norm is just the product of the conjugates. But even if $L/K(B)$ is not separable, the norm is the product of conjugates to the appropriate power, because we can factor $L/K(B)$ as a separable extension followed by a purely inseparable extension, or because we can take an explicit basis for the extension and calculate the norm of a . See [Lan, §VI.5] for more on norms.) But since $a \in P$, its norm lies in P , but also in B , and hence in $P \cap B = q'$, yielding a contradiction. \square

We end with three exercises which may give you practice and enlightenment.

12.2.J. EXERCISE/DEFINITION. Suppose p is a closed point of a locally finite type k -scheme X . Show that the following three integers are the same:

- (a) the largest dimension of an irreducible component of X containing p ;
- (b) $\dim \mathcal{O}_{X,p}$;
- (c) $\text{codim}_X p$.

We call this the **dimension of X at p** . Use (a) to show that this function on X is an upper semicontinuous function on the *closed points* of X . (The **Zariski topology on the closed points of a finite type k -scheme** in the obvious way: closed subsets are cut out by equations. A case of this already arose in §8.4.5.)

12.2.K. EXERCISE (THE DIMENSION OF A LOCALLY FINITE TYPE k -SCHEME IS PRESERVED BY ANY FIELD EXTENSION, CF. EXERCISE 12.1.G(A)). Suppose X is a locally finite type k -scheme of pure dimension n , and K/k is a field extension (not necessarily algebraic). Show that X_K has pure dimension n . Hint: Reduce to the case where X is affine, so say $X = \text{Spec } A$. Reduce to the case where A is an integral domain. Show (using the axiom of choice) that K/k can be written as an algebraic extension of a purely transcendental extension. Hence by Exercise 12.1.G(a), it suffices to deal with the case where K/k is purely transcendental, say with transcendence basis $\{e_i\}_{i \in I}$ (possibly infinite). Show that $A_K := A \otimes_k K$ is an integral domain, by interpreting it as a certain localization of the domain $A[[e_i]]$. If t_1, \dots, t_d is a transcendence basis for $K(A)/k$, show that $\{e_i\} \cup \{t_j\}$ is a transcendence basis for $K(A_K)/k$. Show that $\{t_j\}$ is a transcendence basis for $K(A_K)/K$.

Exercise 12.2.K is conceptually very useful. For example, if X is described by some equations with \mathbb{Q} -coefficients, the dimension of X doesn't depend on whether we consider it as a \mathbb{Q} -scheme or as a \mathbb{C} -scheme.

12.2.16. Remark. Unlike Exercise 12.1.G, Exercise 12.2.K has finite type hypotheses on X . It is not true that if X is an arbitrary k -scheme of pure dimension n , and K/k is an arbitrary extension, then X_K necessarily has pure dimension n . For example, Exercise 10.2.L shows that if k is a field, then $\dim k(x) \otimes_k k(y) = 1$.

12.2.17. ★ Lines on hypersurfaces, part 1. Notice: although dimension theory is not central to the following statement, it is essential to the proof.

12.2.L. ENLIGHTENING STRENUOUS EXERCISE: MOST SURFACES IN THREE-SPACE OF DEGREE $d > 3$ HAVE NO LINES. In this exercise, we work over an algebraically closed field \bar{k} . For any $d > 3$, show that most degree d surfaces in \mathbb{P}^3 contain no lines. Here, “most” means “all closed points of a Zariski-open subset of the parameter space for degree d homogeneous polynomials in 4 variables, up to scalars”. As there are $\binom{d+3}{3}$ such monomials, the degree d hypersurfaces are parametrized by $\mathbb{P}^{\binom{d+3}{3}-1}$ (see Remark 4.5.3). Hint: Construct an “incidence” correspondence

$$X = \{(\ell, H) : [\ell] \in \mathbb{G}(1, 3), [H] \in \mathbb{P}^{\binom{d+3}{3}-1}, \ell \subset H\},$$

parametrizing lines in \mathbb{P}^3 contained in a hypersurface: define a closed subscheme X of $\mathbb{P}^{\binom{d+3}{3}-1} \times \mathbb{G}(1, 3)$ that makes this notion precise. (Recall that $\mathbb{G}(1, 3)$ is a Grassmannian.) Show that X is a $\mathbb{P}^{\binom{d+3}{3}-1-(d+1)}$ -bundle over $\mathbb{G}(1, 3)$. (Possible hint for this: how many degree d hypersurfaces contain the line $x = y = 0$?) Show that $\dim \mathbb{G}(1, 3) = 4$ (see §7.7: $\mathbb{G}(1, 3)$ has an open cover by \mathbb{A}^4 's). Show that $\dim X = \binom{d+3}{3} - 1 - (d+1) + 4$. Show that the image of the projection $X \rightarrow \mathbb{P}^{\binom{d+3}{3}-1}$

must lie in a proper closed subset. The following diagram may help.

$$\dim \binom{d+3}{3} - 1 - (d+1) + 4$$

$$\begin{array}{ccc} & X & \\ \swarrow & & \searrow^{\mathbb{P} \binom{d+3}{3} - 1 - (d+1)} \\ \mathbb{P} \binom{d+3}{3} - 1 & & \mathbb{G}(1,3) \quad \dim 4 \end{array}$$

(The argument readily generalizes to show that if $d > 2n-3$, then “most” degree d hypersurfaces in \mathbb{P}^n have no lines. The case $n = 1$ and $n = 2$ are trivial but worth thinking through.)

12.2.18. Side Remark. If you do the previous Exercise, your dimension count will suggest the true facts that degree 1 hypersurfaces — i.e., hyperplanes — have 2-dimensional families of lines, and that most degree 2 hypersurfaces have 1-dimensional families (rulings) of lines, as shown in Exercise 9.3.L. They will also suggest that most degree 3 hypersurfaces contain a finite number of lines, which reflects the celebrated fact that regular cubic surfaces over an algebraically closed field always contain 27 lines (Theorem 27.1.1), and we will use this “incidence correspondence” or “incidence variety” to prove it (§27.4). The statement about quartic surfaces generalizes to the Noether-Lefschetz Theorem implying that a very general surface of degree d at least 4 contains no curves that are not the intersection of the surface with a hypersurface, see [Lef, GH2]. “**Very general**” means that in the parameter space (in this case, the projective space parametrizing surfaces of degree d), the statement is true away from a countable union of proper Zariski-closed subsets. Like “general”, (which was defined in §10.3.7), “very general” is a weaker version of the phrase “almost every”.

12.3 Codimension one miracles: Krull and Hartogs

In this section, we will explore a number of results related to codimension one. We introduce two results that apply in more general situations, and link functions and the codimension one points where they vanish: Krull’s Principal Ideal Theorem 12.3.3, and Algebraic Hartogs’s Lemma 12.3.15. We will find these two theorems very useful. For example, Krull’s Principal Ideal Theorem will help us compute codimensions, and will show us that codimension can behave oddly, and Algebraic Hartogs’s Lemma will give us a useful characterization of unique factorization domains (Proposition 12.3.5). The results in this section will require (locally) Noetherian hypotheses. They are harder, in that the proofs are technical, and don’t shed much light on the uses of the results. Thus it is more important to understand how to use these results than to be familiar with their proofs.

12.3.1. Krull’s Principal Ideal Theorem. In a vector space, a single linear equation always cuts out a subspace of codimension 0 or 1 (and codimension 0 occurs only when the equation is 0). The Principal Ideal Theorem, due to W. Krull, generalizes this linear algebra fact.

12.3.2. Krull's Principal Ideal Theorem (geometric version). — Suppose X is a locally Noetherian scheme, and f is a function. The irreducible components of $V(f)$ are codimension 0 or 1.

This is clearly a consequence of the following algebraic statement. The full proof is technical, and included in [REFERENCE TO BE ADDED LATER.] only to show you that it isn't excessively long.

12.3.3. Krull's Principal Ideal Theorem (algebraic version). — Suppose A is a Noetherian ring, and $f \in A$. Then every prime \mathfrak{p} minimal among those containing f has codimension at most 1. If furthermore f is not a zerodivisor, then every such prime \mathfrak{p} containing f has codimension precisely 1.

For example, locally principal closed subschemes have “codimension 0 or 1”, and effective Cartier divisors have “pure codimension 1”. Here is another example, that you could certainly prove directly, without the Principal Ideal Theorem.

12.3.A. EXERCISE. Show that an irreducible homogeneous polynomial in $n + 1$ variables over a field k describes an integral scheme of dimension $n - 1$ in \mathbb{P}_k^n .

12.3.B. IMPORTANT EXERCISE (TO BE USED REPEATEDLY). This is a cool argument. (a) (*Hypersurfaces meet everything of dimension at least 1 in projective space, unlike in affine space.*) Suppose X is a closed subset of \mathbb{P}_k^n of dimension at least 1, and H is a nonempty hypersurface in \mathbb{P}_k^n . Show that H meets X . (Hint: note that the affine cone over H contains the origin in \mathbb{A}_k^{n+1} . Apply Krull's Principal Ideal Theorem 12.3.3 to the cone over X .)

(b) Suppose $X \hookrightarrow \mathbb{P}_k^n$ is a closed subset of dimension r . Show that any codimension r linear space meets X . Hint: Refine your argument in (a). (Exercise 12.3.E generalizes this to show that any two things in projective space that you would expect to meet for dimensional reasons do in fact meet.)

(c) Show further that there is an intersection of $r + 1$ nonempty hypersurfaces missing X . (The key step: show that there is a hypersurface of sufficiently high degree that doesn't contain any generic point of X . Show this by induction on the number of generic points. To get from m to $m + 1$: take a hypersurface not vanishing on p_1, \dots, p_m . If it doesn't vanish on p_{m+1} , we are done. Otherwise, call this hypersurface f_{m+1} . Do something similar with $m + 1$ replaced by i for each $1 \leq i \leq m$. Then consider $\sum_i f_1 \cdots \hat{f}_i \cdots f_{m+1}$.) If k is infinite, show that there is a codimension $r + 1$ linear subspace missing X . (The key step: show that there is a hyperplane not containing any generic point of a component of X .)

(d) If k is an infinite field, show that there is an intersection of r hyperplanes meeting X in a finite number of points. (We will see in Exercise 13.4.C that if $k = \bar{k}$, for “most” choices of these r hyperplanes, this intersection is reduced, and in Exercise 19.6.N that the number of points is the “degree” of X . But first of course we must define “degree”.)

The proof of prime avoidance (the following exercise) is similar to that of Exercise 12.3.B(c), and you may wish to solve it first. We will find it useful shortly.

12.3.C. EXERCISE. Prove Proposition 12.2.13 (prime avoidance). Hint: by induction on n . Don't look in the literature — you might find a much longer argument.

12.3.D. EXERCISE (BOUND ON CODIMENSION OF INTERSECTIONS IN \mathbb{A}_k^n). Let k be a field. Suppose X and Y are equidimensional closed subvarieties (possibly singular) of codimension m and n respectively in \mathbb{A}_k^d . Show that every component of $X \cap Y$ has codimension at most $m + n$ in \mathbb{A}_k^d as follows. Show that the diagonal $\mathbb{A}_k^d \cong \Delta \subset \mathbb{A}_k^d \times_k \mathbb{A}_k^d$ is a regular embedding of codimension d . (You will quickly guess the d equations for Δ .) Figure out how to identify the intersection of X and Y in \mathbb{A}_k^d with the intersection of $X \times Y$ with Δ in $\mathbb{A}_k^d \times_k \mathbb{A}_k^d$. Then show that locally, $X \cap Y$ is cut out in $X \times Y$ by d equations. Use Krull's Principal Ideal Theorem 12.3.3. You will also need Exercise 12.2.F. (See Exercise 13.2.M for a generalization.)

12.3.E. FUN EXERCISE (GENERALIZING EXERCISE 12.3.B(B)). Suppose X and Y are equidimensional closed subvarieties of \mathbb{P}^n of codimensions d and e respectively, and $d + e \leq n$. Show that X and Y intersect. Hint: apply Exercise 12.3.D to the affine cones of X and Y . Recall the argument you used in Exercise 12.3.B(a) or (b).

12.3.F. USEFUL EXERCISE. Suppose f is an element of a Noetherian ring A , contained in no codimension zero or one prime ideals. Show that f is invertible. (Hint: if a function vanishes nowhere, it is invertible, by Exercise 4.3.F(b).)

12.3.4. A useful characterization of unique factorization domains.

We can use Krull's Principal Ideal Theorem to prove one of the four useful criteria for unique factorization domains, promised in §5.4.5.

12.3.5. Proposition. — *Suppose that A is a Noetherian integral domain. Then A is a unique factorization domain if and only if all codimension 1 prime ideals are principal.*

This contains Lemma 12.1.7 and (in some sense) its converse. (In fact this result is true even without Noetherian hypotheses, without much more work; see [Ka, Thm. 5, p. 4].)

Proof. One direction is Lemma 12.1.7: if A is a unique factorization domain, then all codimension 1 prime ideals are principal. So we assume conversely that all codimension 1 prime ideals of A are principal, and we will show that A is a unique factorization domain. Now each codimension one prime ideal, being principal, is generated by a prime element, which is unique up to multiplication by a unit. All prime elements of a ring are irreducible, by the usual argument. (Suppose $p \in A$ is prime, and $p = ab$, with $a, b \in A \setminus \{0\}$. Then by primality of the ideal (p) , either a or b is in (p) , say without loss of generality $a = kp$. Then $p = kpb$, from which $1 = bk$, from which b is unit. Thus p is irreducible.) Conversely, if $a \in A$ is irreducible, then by the Principal Ideal Theorem 12.3.3, $V(a)$ contains some codimension one point $[(p)]$, so $a \in (p)$, say $a = kp$. But by irreducibility of a , as p is not a unit, k must be a unit, from which $(a) = (p)$, so a is prime. Thus we have shown that the prime elements are the irreducible elements (and incidentally, that they are, up to multiplicative units, in correspondence with the codimension one prime ideals).

Finally, A is a unique factorization domain by the usual argument: each non-unit nonzero element can be factored into a finite number of irreducibles, and the uniqueness of such a factorization can be shown using the fact that those irreducibles are also prime. \square

12.3.G. EXERCISE. Suppose (A, \mathfrak{m}) is a Noetherian local ring, and $f \in \mathfrak{m}$. Show that $\dim A/(f) \geq \dim A - 1$.

12.3.6. Systems of parameters, and Krull's Height Theorem..

The following result is easily seen to be equivalent to Krull's Principal Ideal Theorem, but I like it better. (Perhaps mention linear algebra...)

12.3.7. Key Lemma. — *Suppose*

$$\mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \cdots \subsetneq \mathfrak{p}_n$$

is a strict chain of prime ideals in a Noetherian ring A , and $f \in \mathfrak{p}_n \setminus \mathfrak{p}_0$. Then there exists a strict chain of prime ideals

$$\mathfrak{p}_0 = \mathfrak{q}_0 \subsetneq \mathfrak{q}_1 \subsetneq \cdots \subsetneq \mathfrak{q}_{n-1} \subsetneq \mathfrak{q}_n = \mathfrak{p}_n.$$

in A , with $f \in \mathfrak{q}_1 \setminus \mathfrak{q}_0$.

(You should feel compelled to immediately sketch of a picture of this lemma!)

Proof. Our plan:

$$\begin{array}{ccccccccc} \mathfrak{p}_0 & \longrightarrow & \mathfrak{p}_1 & \longrightarrow & \cdots & \longrightarrow & \mathfrak{p}_{n-2} & \longrightarrow & \mathfrak{p}_{n-1} & \longrightarrow & \mathfrak{p}_n \\ \parallel & \searrow & & \searrow & & \searrow & & \searrow & & \searrow & \parallel \\ \mathfrak{q}_0 & \longrightarrow & \mathfrak{q}_1 & \longrightarrow & \cdots & \longrightarrow & \mathfrak{q}_{n-2} & \longrightarrow & \mathfrak{q}_{n-1} & \longrightarrow & \mathfrak{q}_n \end{array}$$

Suppose $\mathfrak{p}_{j-2} \subset \mathfrak{p}_{j-1} \subset \mathfrak{p}_j$ with $f \in \mathfrak{p}_j \setminus \mathfrak{p}_{j-1}$. We will find a prime ideal \mathfrak{q}_{j-1} containing f , with $\mathfrak{p}_{j-2} \subsetneq \mathfrak{q}_{j-1} \subsetneq \mathfrak{p}_j$. (This result, by induction, will imply the Lemma.)

Take \mathfrak{q}_{j-1} to be a prime contained in (f, \mathfrak{p}_{j-2}) , and minimal among those contained in \mathfrak{p}_j .

Now \mathfrak{q}_{j-1} is codimension 1 in R/\mathfrak{p}_{j-2} by Krull's Principal Ideal theorem, and \mathfrak{p}_j is codimension at least 2 (we have found a chain of length 2). Thus $\mathfrak{q}_{j-1} \subsetneq \mathfrak{p}_j$. \square

12.3.8. Corollary (Krull for dimension of a Noetherian local ring). — *If (A, \mathfrak{m}) is a Noetherian local ring, and $f \in \mathfrak{m}$, then $\dim A/(f) \geq \dim A - 1$.*

Proof. Easy! Make chain for A/f out of chain for A . \square

The following generalization of Krull's Principal Ideal Theorem looks like it might follow by induction from Krull, but it is more subtle.

12.3.9. Krull's Height Theorem. — *Suppose $X = \operatorname{Spec} A$ where A is Noetherian, and Z is an irreducible component of $V(r_1, \dots, r_\ell)$, where $r_1, \dots, r_\ell \in A$. Then the codimension of Z is at most ℓ .*

(As mentioned in §12.1.4, *height* is another common word for *codimension*.)

Proof. Localize at \mathfrak{p} . Then you have a Noetherian local ring. Take a chain of length d . Then use the r 's in order. \square

Noetherian local ring is finite-dimensional. Proof: $\mathfrak{m} = (f_1, \dots, f_r)$. Why codimension = \dim ? Explain! Cut in from:

12.3.H. IMPORTANT EXERCISE. Suppose (A, \mathfrak{m}) is a Noetherian local ring.

(a) (Noetherian local rings have finite dimension, promised in Remark 12.1.9) Use Krull's Height Theorem 12.3.9 to prove that if there are g_1, \dots, g_ℓ such that $V(g_1, \dots, g_\ell) = \{\mathfrak{m}\}$, then $\dim A \leq \ell$. Hence show that A has finite dimension. (For comparison, Noetherian rings in general may have infinite dimension, see Exercise 12.1.K.)

(b) Let $d = \dim A$. Show that there exist $g_1, \dots, g_d \in A$ such that $V(g_1, \dots, g_d) = \{\mathfrak{m}\}$. Hint: use induction on d . Find an equation g_d knocking the dimension down by 1, i.e., $\dim A/(g_d) = \dim A - 1$. Suppose $\mathfrak{p}_1, \dots, \mathfrak{p}_n$ correspond to the irreducible components of $\text{Spec } A$ of dimension d , and $\mathfrak{q}_i \supset \mathfrak{p}_i$ are prime ideals corresponding to irreducible closed subsets of codimension 1 and dimension $d - 1$. Use prime avoidance (Proposition 12.2.13) to find $h_i \in \mathfrak{q}_i \setminus \bigcup_{j=1}^n \mathfrak{p}_j$. Let $g_d = \prod_{i=1}^n h_i$.

12.3.10. Definition. The geometric translation of the above exercise is the following. Given a d -dimensional “germ of a reasonable space” around a point p , p can be cut out set-theoretically by d equations, and we always need at least d equations. These d elements of A are called a **system of parameters** for the Noetherian local ring (A, \mathfrak{m}) .

12.3.11. Remark. There exists a Noetherian local ring that is not catenary (ref needed). Old example?

12.3.12. Definition. **system of parameters** for a Noetherian local ring, or (essentially equivalently) a point on a locally Noetherian scheme. It is a sequence cutting out precisely (set-theoretically) the closed point (so $\sqrt{f_1, \dots, f_d} = \mathfrak{m}$).

VERSION ONE.

VERSION TWO.

12.3.13. Later remark. Counterexample from before (or after?) is no problem.

12.3.14. Algebraic Hartogs's Lemma for Noetherian normal schemes.

Hartogs's Lemma in several complex variables states (informally) that a holomorphic function defined away from a codimension two set can be extended over that. We now describe an algebraic analog, for Noetherian normal schemes. (It may also be profitably compared to the Second Riemann Extension Theorem.) We will use this repeatedly and relentlessly when connecting line bundles and divisors.

12.3.15. Algebraic Hartogs's Lemma. — Suppose A is a integrally closed Noetherian integral domain. Then

$$A = \bigcap_{\mathfrak{p} \text{ codimension } 1} A_{\mathfrak{p}}.$$

The equality takes place in $K(A)$; recall that any localization of an integral domain A is naturally a subset of $K(A)$ (Exercise 1.3.C). Warning: few people call this Algebraic Hartogs's Lemma. I call it this because it parallels the statement in complex geometry.

One might say that if $f \in K(A)$ does not lie in $A_{\mathfrak{p}}$ where \mathfrak{p} has codimension 1, then f has a pole at $[\mathfrak{p}]$, and if $f \in K(A)$ lies in $\mathfrak{p}A_{\mathfrak{p}}$ where \mathfrak{p} has codimension 1, then f has a zero at $[\mathfrak{p}]$.

12.3.16. It is worth *geometrically* interpreting Algebraic Hartogs’s Lemma as saying that *a rational function on a Noetherian normal scheme with no poles is in fact regular* (an element of A). Informally: “Noetherian normal schemes have the Hartogs property.” (We will properly define zeros and poles in §13.5.7, see also Exercise 13.5.H.)

One can state Algebraic Hartogs’s Lemma more generally in the case that $\text{Spec } A$ is a Noetherian normal scheme, meaning that A is a product of Noetherian normal integral domains; the reader may wish to do so.

Another generalization (and something closer to the “right” statement) is that if A is a subring of a field K , then the integral closure of A in K is the intersection of all valuation rings of K containing A ; see [AtM, Cor. 5.22] for explanation and proof.

12.3.17. ★★ Proof of Algebraic Hartogs’s Lemma 12.3.15. (This proof may be stated completely algebraically, but we state it as geometrically as possible, at the expense of making it longer. Nonetheless, this proof may be best skipped on a first reading — the main thing to notice is that it is short and not frightening. See [Stacks, tag 031T] for another proof using Serre’s criterion for normality, which we prove in §29.3.) The left side is obviously contained in the right, so assume some x lies in every A_p but not in A . As in the proof of Proposition 5.4.2, we measure the failure of x to be a function (an element of $\text{Spec } A$) with the “ideal of denominators” I of x (5.4.2.2):

$$I := \{r \in A : rx \in A\}.$$

(As an important remark not necessary for the proof: it is helpful to interpret the ideal of denominators as scheme-theoretically measuring the failure of x to be regular, or better, giving a scheme-theoretic structure to the locus where x is not regular.) As $1 \notin I$, we have $I \neq A$. Choose a minimal prime q containing I .

Our second step in obtaining a contradiction is to focus near the point $[q]$, i.e., focus attention on A_q rather than A , and as a byproduct notice that $\text{codim } q > 1$. The construction of the ideal of denominators behaves well with respect to localization — if p is any prime, then the ideal of denominators of x in A_p is I_p , and it again measures “the failure of Algebraic Hartogs’s Lemma for x ,” this time in A_p . But Algebraic Hartogs’s Lemma is vacuously true for dimension 1 rings, so no codimension 1 prime contains I . Thus q has codimension at least 2. By localizing at q , we can assume that A is a local ring with maximal ideal q , and that q is the *only* prime containing I .

In the third step, we construct a suitable multiple z of x that is still not a function on $\text{Spec } A$, such that multiplying z by anything vanishing at $[q]$ results in a function. (Translation: $z \notin A$, but $zq \subset A$.) As q is the only prime containing I , $\sqrt{I} = q$ (Exercise 3.4.F), so as q is finitely generated, there is some n with $I \supset q^n$ (Do you see why?). Take the minimal such n , so $I \not\supset q^{n-1}$, and choose any $y \in q^{n-1} - I$. Let $z = yx$. As $y \notin I$, $z \notin A$. On the other hand, $yq \subset q^n \subset I$, so $zq \subset Ix \subset A$, so zq is an ideal of A , completing this step.

Finally, we have two cases: either there is function vanishing on $[q]$ that, when multiplied by z , doesn’t vanish on $[q]$; or else every function vanishing on $[q]$, multiplied by z , still vanishes on $[q]$. Translation: (i) either zq is contained in q , or (ii) it is not.

(i) If $zq \subset q$, then we would have a finitely generated A -module (namely q) with a faithful $A[z]$ -action, forcing z to be integral over A (and hence in A , as A is integrally closed) by Exercise 8.2.I, yielding a contradiction.

(ii) If zq is an ideal of A not contained in the unique maximal ideal q , then it must be A ! Thus $zq = A$ from which $q = A(1/z)$, from which q is principal. But then $\text{codim } q = \dim A \leq \dim_{A/q} q/q^2 \leq 1$ (the inequality coming from principality of q), contradicting $\text{codim } q \geq 2$. \square

12.3.18. Proof of Krull's Principal Ideal Theorem 12.3.3.

We prove Krull's Principal Ideal Theorem 12.3.3 in three exercises.

12.3.I. EXERCISE. Show that it suffices to prove Theorem 12.3.3 in the case where the ring A is an integral domain: if A is a Noetherian integral domain, and $f \in A$ is nonzero, then each irreducible component of $V(f)$ is codimension 1 in $\text{Spec } A$.


12.3.J. EXERCISE. Prove the result (in the form of the reduction, in Exercise 12.3.I) in the case where A is normal, by considering the rational function $1/f \in K(A)$, and using Algebraic Hartogs's Lemma 12.3.15.

12.3.K. EXERCISE. Prove the result (in the form of the reduction, in Exercise 12.3.I) in general, by first considering $1/f \in K(\tilde{A}) = K(A)$, where \tilde{A} is the integral closure of A in its fraction field $K(A)$. Then prove the result for A , using your proof of Exercise 12.1.F (the dimension of an integral scheme is the dimension of its normalization). \square

We conclude with two starred sections.

12.3.19. ★ Pathologies of the notion of “codimension”. We can use Krull's Principal Ideal Theorem to produce the example of pathology in the notion of codimension promised earlier this chapter. Let $A = k[x]_{(x)}[t]$. In other words, elements of A are polynomials in t , whose coefficients are quotients of polynomials in x , where no factors of x appear in the denominator. (Warning: A is not $k[x, t]_{(x)}$.) Clearly, A is an integral domain, so $xt - 1$ is not a zero divisor. You can verify that $A/(xt - 1) \cong k[x]_{(x)}[1/x] \cong k(x)$ — “in $k[x]_{(x)}$, we may divide by everything but x , and now we are allowed to divide by x as well” — so $A/(xt - 1)$ is a field. Thus $(xt - 1)$ is not just prime but also maximal. By Krull's theorem, $(xt - 1)$ is codimension 1. Thus $(0) \subset (xt - 1)$ is a maximal chain. However, A has dimension at least 2: $(0) \subset (t) \subset (x, t)$ is a chain of prime ideals of length 2. (In fact, A has dimension precisely 2, although we don't need this fact in order to observe the pathology.) Thus we have a codimension 1 prime in a dimension 2 ring whose quotient is dimension 0, not 1. Here is a picture of this partially ordered set of

ideals.



This example comes from geometry, and it is enlightening to draw a picture, see Figure 12.2. $\text{Spec } k[x]_{(x)}$ corresponds to a “germ” of \mathbb{A}_k^1 near the origin, and $\text{Spec } k[x]_{(x)}[t]$ corresponds to “this \times the affine line”. You may be able to see from the picture some motivation for this pathology — $V(xt - 1)$ doesn’t meet $V(x)$, so it can’t have any specialization on $V(x)$, and there is nowhere else for $V(xt - 1)$ to specialize. It is disturbing that this misbehavior turns up even in a relatively benign-looking ring.




FIGURE 12.2. Dimension and codimension behave oddly on the surface $\text{Spec } k[x]_{(x)}[t]$

12.3.20. ★ Lines on hypersurfaces, part 2. (Part 1 was §12.2.17.) We now give a geometric application of Krull’s Principal Ideal Theorem 12.3.3, applied through Exercise 12.3.B(a). Throughout, we work over an algebraically closed field \bar{k} .

12.3.L. EXERCISE.

(a) Suppose

$$f(x_0, \dots, x_n) = f_d(x_1, \dots, x_n) + x_0 f_{d-1}(x_1, \dots, x_n) + \cdots + x_0^{d-1} f_1(x_1, \dots, x_n)$$

is a homogeneous degree d polynomial (so $\deg f_i = i$) cutting out a hypersurface X in \mathbb{P}^n containing $p := [1, 0, \dots, 0]$. Show that there is a line through p contained in X if and only if $f_1 = f_2 = \dots = f_d = 0$ has a common zero in $\mathbb{P}^{n-1} = \text{Proj } \bar{k}[x_1, \dots, x_n]$. (Hint: given a common zero $[a_1, \dots, a_n] \in \mathbb{P}^{n-1}$, show that line joining p to $[0, a_1, \dots, a_n]$ is contained in X .)

(b) If $d \leq n - 1$, show that through any point $p \in X$, there is a line contained in X . Hint: Exercise 12.3.B(a).

(c) If $d \geq n$, show that for “most hypersurfaces” X of degree d in \mathbb{P}^n (for all hypersurfaces whose corresponding point in the parameter space $\mathbb{P}^{\binom{n+d}{d}-1}$ — cf. Remark 4.5.3 and Exercise 9.3.J — lies in some nonempty Zariski-open subset), “most points $p \in X$ ” (all points in a nonempty dense Zariski-open subset of X) have no lines in X passing through them. (Hint: first show that there is a single p in a single X contained in no line. Chevalley’s Theorem 8.4.2 may help.)

12.3.21. Remark. A projective (or proper) \bar{k} -variety X is **uniruled** if every point $p \in X$ is contained in some $\mathbb{P}^1 \subset X$. (We won’t use this word beyond this remark.) Part (b) shows that all hypersurfaces of degree at most $n-1$ are uniruled. One can show (using methods beyond what we know now, see [Ko1, Cor. IV.1.11]) that if $\text{char } \bar{k} = 0$, then every smooth hypersurface of degree at least $n+1$ in $\mathbb{P}_{\bar{k}}^n$ is *not* uniruled (thus making the open set in (c) explicit). Furthermore, smooth hypersurfaces of degree n are uniruled, but covered by conics rather than lines. Thus there is a strong difference in how hypersurfaces behave depending on how the degree relates to $n+1$. This is true in many other ways as well. Smooth hypersurfaces of degree less than $n+1$ are examples of *Fano* varieties; smooth hypersurfaces of degree $n+1$ are examples of *Calabi-Yau* varieties (with the possible exception of $n=1$, depending on the definition); and smooth hypersurfaces of degree greater than $n+1$ are examples of *general type* varieties. We define these terms in §22.5.5.

12.4 Dimensions of fibers of morphisms of varieties

In this section, we show that the dimensions of fibers of morphisms of varieties behave in a way you might expect from our geometric intuition. The reason we have waited until now to discuss this is because we will use Theorem 12.2.9 (for varieties, codimension is the difference of dimensions). We discuss generalizations in §12.4.5.

Recall that a function f from a topological space X to \mathbb{R} is **upper semicontinuous** if for each $x \in \mathbb{R}$, $f^{-1}((-\infty, x))$ is open. Informally: the function can jump “up” upon taking limits. Our upper semicontinuous functions will map to \mathbb{Z} , so informally functions jump “up” on closed subsets. Similarly, f is **lower semicontinuous** if $x \in \mathbb{R}$, $f^{-1}((x, \infty))$ is open.

Before we begin, let’s make sure we are on the same page with respect to our intuition. Elimination theory (Theorem 8.4.7) tells us that the projection $\pi: \mathbb{P}_A^n \rightarrow \text{Spec } A$ is closed. We can interpret this as follows. A closed subset X of \mathbb{P}_A^n is cut out by a bunch of homogeneous equations in $n+1$ variables (over A). The image of X is the subset of $\text{Spec } A$ where these equations have a common nontrivial solution. If we try hard enough, we can describe this by saying that the *existence* of a nontrivial solution (or the existence of a preimage of a point under $\pi: X \rightarrow \text{Spec } A$)

is an “upper semicontinuous” fact. More generally, your intuition might tell you that the locus where a number of homogeneous polynomials in $n + 1$ variables over A have a solution space (in \mathbb{P}_A^n) of dimension at least d should be a closed subset of $\text{Spec } A$. (As a special case, consider linear equations. The condition for m linear equations in $n + 1$ variables to have a solution space of dimension at least $d + 1$ is a closed condition on the coefficients — do you see why, using linear algebra?) This intuition will be correct, and will use properness in a fundamental way (Theorem 12.4.2(b)). We will also make sense of upper semicontinuity in fiber dimension on the *source* (Theorem 12.4.2(a)). A useful example to think through is the map from the xy -plane to the xz -plane ($\text{Spec } k[x, y] \rightarrow \text{Spec } k[x, z]$), given by $(x, z) \mapsto (x, xy)$. (This example also came up in §8.4.1.)

We begin our substantive discussion with an inequality that holds more generally in the locally Noetherian setting.

12.4.A. KEY EXERCISE (CODIMENSION BEHAVES AS YOU MIGHT EXPECT FOR A MORPHISM, OR “FIBER DIMENSIONS CAN NEVER BE LOWER THAN EXPECTED”). Suppose $\pi: X \rightarrow Y$ is a morphism of locally Noetherian schemes, and $p \in X$ and $q \in Y$ are points such that $q = \pi(p)$. Show that

$$\text{codim}_X p \leq \text{codim}_Y q + \text{codim}_{\pi^{-1}(q)} p$$

(see Figure 12.3). Hint: take a system of parameters (Definition 12.3.10) for q “in Y ”, and a system of parameters for p “in $\pi^{-1}(q)$ ”, and use them to find $\text{codim}_Y q + \text{codim}_{\pi^{-1}(q)} p$ elements of $\mathcal{O}_{X,p}$ cutting out $p = \{[m]\}$ in $\text{Spec } \mathcal{O}_{X,p}$. Use Exercise 12.3.H.




FIGURE 12.3. Exercise 12.4.A: the codimension of a point in the total space is bounded by the sum of the codimension of the point in the fiber plus the codimension of the image in the target

Does Exercise 12.4.A agree with your geometric intuition? You should be able to come up with enlightening examples where equality holds, and where equality fails. We will see that equality always holds for sufficiently nice — flat — morphisms, see Proposition 25.5.7.

We now show that the inequality of Exercise 12.4.A is actually an equality over “most of Y ” if Y is an irreducible variety.

12.4.1. Theorem. — Suppose $\pi: X \rightarrow Y$ is a (necessarily finite type) morphism of irreducible k -varieties, with $\dim X = m$ and $\dim Y = n$. Then there exists a nonempty open subset $U \subset Y$ such that for all $q \in U$, the fiber over q has pure dimension $m - n$, or is empty.

(Unimportant remark: the empty set technically has pure dimension $m - n$ — every irreducible component of the empty set has dimension $m - n$. This makes the statement a bit cleaner.)

Proof. We begin with three quick reductions. (i) By shrinking Y if necessary, we may assume that Y is affine, say $\text{Spec } B$. (ii) We may also assume that X is affine, say $\text{Spec } A$. (Reason: cover X with a finite number of affine open subsets X_1, \dots, X_a , and take the intersection of the U ’s for each of the $\pi|_{X_i}$.) (iii) If π is not dominant, then we are done, as the image misses a dense open subset U of $\text{Spec } B$. So we assume now that π is dominant.

In order to motivate the rest of the argument, we describe our goal. We will produce a nonempty distinguished open subset U of $\text{Spec } B$ so that $\pi^{-1}(U) \rightarrow U$ factors through $\mathbb{A}_k^{m-n} := \mathbb{A}_k^{m-n} \times_k U$ via a finite surjective morphism:

$$(12.4.1.1) \quad \begin{array}{ccc} \text{Spec } A & \xleftarrow{\text{open emb.}} & \pi^{-1}(U) \\ \downarrow \pi & & \downarrow \text{finite surj.} \\ & & \mathbb{A}_U^{m-n} \\ & & \downarrow \\ \text{Spec } B & \xleftarrow{\text{open emb.}} & U \end{array}$$

12.4.B. EXERCISE. Show that this suffices to prove the Theorem 12.4.1. (Hint: Use Exercise 12.4.A, and Theorem 12.2.9 that codimension is the difference of dimensions for varieties, to show that each component of the fiber over a point of U has dimension at least $m - n$. Show that any irreducible variety mapping finitely to \mathbb{A}_k^{m-n} has dimension at most $m - n$.)

So we now work to build (12.4.1.1). We begin by noting that we have inclusions of B into both A and $K(B)$, and from both A and $K(B)$ into $K(A)$. The maps from A and $K(B)$ into $K(A)$ both factor through $A \otimes_B K(B)$ (whose Spec is the generic fiber of π), so the maps from both A and $K(B)$ to $A \otimes_B K(B)$ must be inclusions.

$$(12.4.1.2) \quad \begin{array}{ccccc} & & & & K(A) \\ & & & \nearrow & \\ & & & A \otimes_B K(B) & \\ \uparrow & \hookrightarrow & & \uparrow & \\ A & & & K(B) & \\ \uparrow & \hookrightarrow & & \uparrow & \\ B & \hookrightarrow & & K(B) & \end{array}$$

Clearly $K(A) \otimes_B K(B) = K(A)$ (as $A \otimes_B K(B)$ can be interpreted as taking A and inverting those nonzero elements of B), and $A \otimes_B K(B)$ is a finitely generated algebra over the field $K(B)$.

By transcendence theory (Exercise 12.2.A), $K(A)$ has transcendence degree $m - n$ over $K(B)$ (as $K(A)$ has transcendence degree m over k , and $K(B)$ has transcendence degree n over k). Thus by Noether normalization 12.2.4, we can find elements $t_1, \dots, t_{m-n} \in A \otimes_B K(B)$, algebraically independent over $K(B)$, such that $A \otimes_B K(B)$ is integral over $K(B)[t_1, \dots, t_{m-n}]$.

Now, we can think of the elements $t_i \in A \otimes_B K(B)$ as fractions, with numerators in A and (nonzero) denominators in B . If f is the product of the denominators appearing for each t_i , then by replacing B by B_f (replacing $\text{Spec } B$ by its distinguished open subset $D(f)$), we may assume that the t_i are all in A . Thus (after sloppily renaming B_f as B , and A_f as A) we can trim and extend (12.4.1.2) to the following.

$$\begin{array}{ccc}
 A & \xrightarrow{\quad} & A \otimes_B K(B) \\
 \uparrow & & \uparrow \text{integral} \\
 B[t_1, \dots, t_{m-n}] & \xrightarrow{\quad} & K(B)[t_1, \dots, t_{m-n}] \\
 \uparrow & & \uparrow \\
 B & \xrightarrow{\quad} & K(B)
 \end{array}$$

Now A is finitely generated over B , and hence over $B[t_1, \dots, t_{m-n}]$, say by u_1, \dots, u_q . Noether normalization implies that each u_i satisfies some monic equation $f_i(u_i) = 0$, where $f_i \in K(B)[t_1, \dots, t_{m-n}][t]$. The coefficients of f_i are a priori fractions in B , but by multiplying by all those denominators, we can assume each $f_i \in B[t_1, \dots, t_{m-n}][t]$. Let $b \in B$ be the product of the leading coefficients of all the f_i . If $U = D(b)$ (the locus where b is invertible), then over U , the f_i (can be taken to) have leading coefficient 1, so the u_i (in A_b) are integral over $B_b[t_1, \dots, t_{m-n}]$. Thus $\text{Spec } A_b \rightarrow \text{Spec } B_b[t_1, \dots, t_{m-n}]$ is finite and surjective (the latter by the Lying Over Theorem 8.2.5).

We have now constructed (12.4.1.1), as desired. \square

There are a couple of things worth pointing out about the proof. First, this result is interesting for classical varieties over a field k . But the proof of it uses the theory of varieties over *another* field, notably the function field $K(B)$. This is an example of how the introduction of generic points to algebraic geometry is useful even for considering more “classical” questions.

Second, the idea of the main part of the argument is that we have a result over the generic point ($\text{Spec } A \otimes_B K(B)$ finite and surjective over affine space over $K(B)$), and we want to “spread it out” to an open neighborhood of the generic point of $\text{Spec } B$. We do this by realizing that “finitely many denominators” appear when correctly describing the problem, and inverting those functions. This “spreading out” idea is a recurring theme.

12.4.C. EXERCISE (USEFUL CRITERION FOR IRREDUCIBILITY). Suppose $\pi: X \rightarrow Y$ is a proper morphism to an irreducible variety, and all the fibers of π are nonempty, and irreducible of the same dimension. Show that X is irreducible.

12.4.2. Theorem (upper semicontinuity of fiber dimension). — Suppose $\pi: X \rightarrow Y$ is a morphism of finite type k -schemes.

(a) (upper semicontinuity on the source) The dimension of the fiber of π at $p \in X$ (the dimension of the largest component of $\pi^{-1}(\pi(p))$ containing p) is an upper semicontinuous function in p (i.e., on X).

(b) (upper semicontinuity on the target) If furthermore π is closed (e.g., if π is proper), then the dimension of the fiber of π over $q \in Y$ is an upper semicontinuous function in q (i.e., on Y).

You should be able to immediately construct a counterexample to part (b) if the closedness hypothesis is dropped. (We also remark that Theorem 12.4.2(b) for projective morphisms is done, in a simple way, in Exercise 19.1.B.)

Proof. (a) Let F_n be the subset of X consisting of points where the fiber dimension is at least n . We wish to show that F_n is a closed subset for all n . We argue by induction on $\dim Y$. The base case $\dim Y = 0$ is trivial. So we fix Y , and assume the result for all smaller-dimensional targets.

12.4.D. EXERCISE. Show that it suffices to prove the result when X and Y are integral, and π is dominant.

Let $r = \dim X - \dim Y$ be the “relative dimension” of π . If $n \leq r$, then $F_n = X$ by Exercise 12.4.A (combined with Theorem 12.2.9, that codimension is the difference of dimensions for varieties).

If $n > r$, then let $U \subset Y$ be the dense open subset of Theorem 12.4.1, where “the fiber dimension is exactly r ”. Then F_n does not meet the preimage of U . By replacing Y with $Y \setminus U$ (and X by $X \setminus \pi^{-1}(U)$), we are done by the inductive hypothesis.

12.4.E. EASY EXERCISE. Prove (b) (using (a)).

□

12.4.3. Proposition (“generically finite implies generally finite”, cf. Exercise 10.3.G).

— Suppose $\pi: X \rightarrow Y$ is a generically finite morphism of irreducible k -varieties of dimension n . Then there is a dense open subset $V \subset Y$ above which π is finite.

(If you wish, you can later relax the irreducibility hypothesis to simply requiring X and Y to be simply of pure dimension n .)

Proof. As in the proof of Theorem 12.4.1, we may assume that Y is affine, and that π is dominant. Suppose that $X = \bigcup_{i=1}^n U_i$, where the U_i are affine open subschemes of X . By Exercise 10.3.G, there are dense open subsets $V_i \subset Y$ over which $\pi|_{U_i}$ is finite. By replacing Y by an affine open subset of $\bigcap V_i$, we may assume that $\pi|_{U_i}$ is finite.

12.4.F. EXERCISE. Show that π is closed. Hint: you will just use that $\pi|_{U_i}$ is closed, and that there are a finite number of U_i .

Then $X \setminus U_1$ is a closed subset, so $\pi(X \setminus U_1)$ is closed.

12.4.G. EXERCISE. Show that this closed subset is not all of Y .

Define $V := Y \setminus \pi(X \setminus U_1)$. Then π is finite above V : it is the restriction of the finite morphism $\pi|_{U_1} : U_1 \rightarrow Y$ to the open subset V of the target Y . \square

12.4.4. Aside: Semicontinuity. Semicontinuity is a recurring theme in algebraic geometry. It is worth keeping an eye out for it. here are some examples.

- (i) fiber dimension (Theorem 12.4.2 above)
- (ii) the rank of a matrix of functions (because rank drops on closed subsets, where various discriminants vanish)
- (iii) the rank of a finite type quasicoherent sheaf (Exercise 14.3.J)
- (iv) degree of a finite morphism, as a function of the target (§14.3.4)
- (v) dimension of tangent space at closed points of a variety over an algebraically closed field (Exercise 22.2.J)
- (vi) rank of cohomology groups of coherent sheaves, in proper flat families (Theorem 28.1.1)

All but (ii) are upper semicontinuous; (ii) is a lower semicontinuous function.

12.4.5. ★ Generalizing results of §12.4 beyond varieties. The above arguments can be extended to more general situations than varieties. We remain in the locally Noetherian situation for safety, until the last sentence of §12.4.6. One fact used repeatedly was that codimension is the difference of dimensions (Theorem 12.2.9). This holds much more generally; see Remark 12.2.10 on catenary rings. Extensions of Theorem 12.4.1 should require that π be finite type (which was automatic in the statement of Theorem 12.4.1, by the Cancellation Theorem 11.2.1 for finite type morphisms). In the proof of Theorem 12.4.1, we use that the dimension of the generic fiber of the morphism $\pi: X \rightarrow Y$ of irreducible schemes is $\dim X - \dim Y$; this can be proved using Proposition 25.5.7. The remaining results then readily follow without change.

12.4.6. We make one particular generalization explicit because we will mention it later. The ring \mathbb{Z} (along with finitely generated algebras over it, and localizations thereof) is excellent in all possible ways, including in the ways needed in the previous paragraph, so the argument for upper semicontinuity on the source (Theorem 12.4.2(a)) applies without change if the target is the Spec of a finitely generated \mathbb{Z} -algebra. Any finite type morphism of locally Noetherian schemes is locally pulled back from a finite type morphism to the Spec of a finitely generated \mathbb{Z} -algebra (this will essentially be made precise and shown in Exercise 28.2.E), so we then have upper semicontinuity of fiber dimension on the source for all finite type morphisms of locally Noetherian schemes, and better yet, for all locally finitely presented morphisms.

Regularity and smoothness

One natural notion we expect to see for geometric spaces is the notion of when an object is “smooth”. In algebraic geometry, this notion, called *regularity*, is easy to define (Definition 13.2.3) but a bit subtle in practice. This will lead us to a different related notion of when a *variety is smooth* (Definition 13.2.6).

This chapter has many moving parts, of which §13.1–13.6 are the important ones. In §13.1, the *Zariski tangent space* is motivated and defined. In §13.2, we define *regularity* and *smoothness over a field*, the central topics of this chapter, and discuss some of their important properties. In §13.3, we give a number of important examples, mostly in the form of exercises. In §13.4, we discuss *Bertini’s Theorem*, a fundamental classical result. In §13.5, we give many characterizations of *discrete valuation rings*, which play a central role in algebraic geometry. Having seen clues that “smoothness” is a “relative” notion rather than an “absolute” one, in §13.6 we define *smooth morphisms* (and in particular, *étale morphisms*), and give some of their properties. (We will revisit this definition in §22.3.1, once we know more.)

The remaining sections are less central. In §13.7, we discuss the valuative criteria for separatedness and properness. In §13.8, we mention some more sophisticated facts about regular local rings (the local rings at regular points). In §13.9, we prove the Artin-Rees Lemma, because it will have been invoked in §13.5 (and will be used later as well).

13.1 The Zariski tangent space

We begin by defining the tangent space of a scheme at a point. It behaves like the tangent space you know and love at “smooth” points, but also makes sense at other points. In other words, geometric intuition at the “smooth” points guides the definition, and then the definition guides the algebra at all points, which in turn lets us refine our geometric intuition.

The definition is short but surprising. The main difficulty is convincing yourself that it deserves to be called the tangent space. This is tricky to explain, because we want to show that it agrees with our intuition, but our intuition is worse than we realize. So you should just accept this definition for now, and later convince yourself that it is reasonable.

13.1.1. Definition. The **Zariski cotangent space** of a local ring (A, \mathfrak{m}) is defined to be $\mathfrak{m}/\mathfrak{m}^2$; it is a vector space over the residue field A/\mathfrak{m} . The dual vector space is the **Zariski tangent space**. If X is a scheme, the **Zariski cotangent space** $T_{X,p}^\vee$ at a point $p \in X$ is defined to be the Zariski cotangent space of the local ring $\mathcal{O}_{X,p}$ (and similarly for the **Zariski tangent space** $T_{X,p}$). Elements of the Zariski cotangent

space are called **cotangent vectors** or **differentials**; elements of the tangent space are called **tangent vectors**.

The cotangent space is more algebraically natural than the tangent space, in that the definition is shorter. There is a moral reason for this: the cotangent space is more naturally determined in terms of functions on a space, and we are very much thinking about schemes in terms of “functions on them”. This will come up later.

Here are two plausibility arguments that this is a reasonable definition. Hopefully one will catch your fancy.

In differential geometry, the tangent space at a point is sometimes defined as the vector space of derivations at that point. A derivation at a point p of a manifold is an \mathbb{R} -linear operation that takes in functions f near p (i.e., elements of \mathcal{O}_p), and outputs elements $f'(p)$ of \mathbb{R} , and satisfies the Leibniz rule

$$(fg)' = f'g + g'f.$$

(We will later define derivations in a more general setting, §22.2.17.) A derivation is the same as a map $\mathfrak{m} \rightarrow \mathbb{R}$, where \mathfrak{m} is the maximal ideal of \mathcal{O}_p . (The map $\mathcal{O}_p \rightarrow \mathbb{R}$ extends this, via the map $\mathcal{O}_p \rightarrow \mathfrak{m}$ given by $f - f(p)$.) But \mathfrak{m}^2 maps to 0, as if $f(p) = g(p) = 0$, then

$$(fg)'(p) = f'(p)g(p) + g'(p)f(p) = 0.$$

Thus a derivation induces a map $\mathfrak{m}/\mathfrak{m}^2 \rightarrow \mathbb{R}$, i.e., an element of $(\mathfrak{m}/\mathfrak{m}^2)^\vee$.

13.1.A. EXERCISE. Check that this is reversible, i.e., that any linear map $\mathfrak{m}/\mathfrak{m}^2 \rightarrow \mathbb{R}$ gives a derivation. In other words, verify that the Leibniz rule holds.

Here is a second, vaguer, motivation that this definition is plausible for the cotangent space of the origin of \mathbb{A}^n . (I prefer this one, as it is more primitive and elementary.) Functions on \mathbb{A}^n should restrict to a linear function on the tangent space. What (linear) function does $x^2 + xy + x + y$ restrict to “near the origin”? You will naturally answer: $x + y$. Thus we “pick off the linear terms”. Hence $\mathfrak{m}/\mathfrak{m}^2$ are the linear functionals on the tangent space, so $\mathfrak{m}/\mathfrak{m}^2$ is the cotangent space. In particular, you should picture functions vanishing at a point (i.e., lying in \mathfrak{m}) as giving functions on the tangent space in this obvious way.

13.1.2. Old-fashioned example. Computing the Zariski tangent space is actually quite hands-on, because you can compute it just as you did when you learned multivariable calculus. In \mathbb{A}^3 , we have a curve cut out by $x + y + z^2 + xyz = 0$ and $x - 2y + z + x^2y^2z^3 = 0$. (You can use Krull’s Principal Ideal Theorem 12.3.3 to check that this is a curve, but it is not important to do so.) What is the tangent line near the origin? (Is it even smooth there?) Answer: the first surface looks like $x + y = 0$ and the second surface looks like $x - 2y + z = 0$. The curve has tangent line cut out by $x + y = 0$ and $x - 2y + z = 0$. It is smooth (in the traditional sense). In multivariable calculus, the students work hard to get the answer, because we aren’t allowed to tell them to just pick out the linear terms.

Let’s make explicit the fact that we are using. If A is a ring, \mathfrak{m} is a maximal ideal, and $f \in \mathfrak{m}$ is a function vanishing at the point $[\mathfrak{m}] \in \text{Spec } A$, then the Zariski tangent space of $\text{Spec } A/(f)$ at \mathfrak{m} is cut out in the Zariski tangent space of $\text{Spec } A$ (at \mathfrak{m}) by the single linear equation $f \pmod{\mathfrak{m}^2}$. The next exercise will force you to think this through.

13.1.B. IMPORTANT EXERCISE (“KRULL’S PRINCIPAL IDEAL THEOREM FOR TANGENT SPACES” — BUT MUCH EASIER THAN KRULL’S PRINCIPAL IDEAL THEOREM 12.3.3!). Suppose A is a ring, and \mathfrak{m} a maximal ideal. If $f \in \mathfrak{m}$, show that the Zariski tangent space of A/f is cut out in the Zariski tangent space of A by $f \pmod{\mathfrak{m}^2}$. (Note: we can quotient by f and localize at \mathfrak{m} in either order, as quotient and localization commute, (4.3.7.1).) Hence the dimension of the Zariski tangent space of $\text{Spec } A/(f)$ at $[\mathfrak{m}]$ is the dimension of the Zariski tangent space of $\text{Spec } A$ at $[\mathfrak{m}]$, or one less. (That last sentence should be suitably interpreted if the dimension is infinite, although it is less interesting in this case.)

Here is another example to see this principle in action, extending Example 13.1.2: $x + y + z^2 = 0$ and $x + y + x^2 + y^4 + z^5 = 0$ cuts out a curve, which obviously passes through the origin. If I asked my multivariable calculus students to calculate the tangent line to the curve at the origin, they would do calculations which would boil down (without them realizing it) to picking off the linear terms. They would end up with the equations $x + y = 0$ and $x + y = 0$, which cut out a plane, not a line. They would be disturbed, and I would explain that this is because the curve isn’t smooth at a point, and their techniques don’t work. We on the other hand bravely declare that the tangent space is cut out by $x + y = 0$, and (will soon) use this pathology as *definition* of what makes a point singular (or non-regular). (Intuitively, the curve near the origin is very close to lying in the plane $x + y = 0$.) Notice: the cotangent space jumped up in dimension from what it was “supposed to be”, not down. We will see that this is not a coincidence, in Theorem 13.2.1.

13.1.C. EXERCISE. Suppose Y and Z are closed subschemes of X , both containing the point $p \in X$.

- (a) Show that $T_{Z,p}$ is naturally a sub- $\kappa(p)$ -vector space of $T_{X,p}$.
- (b) Show that $T_{Y \cap Z, p} = T_{Y,p} \cap T_{Z,p}$, where \cap as usual in this context is scheme-theoretic intersection.
- (c) Show that $T_{Y \cup Z, p}$ contains the span of $T_{Y,p}$ and $T_{Z,p}$, where \cup as usual in this context is scheme-theoretic union.
- (d) Show that $T_{Y \cup Z, p}$ can be strictly larger than the span of $T_{Y,p}$ and $T_{Z,p}$. (Hint: Figure 4.5.)

Here is a pleasant consequence of the notion of Zariski tangent space.

13.1.3. Problem. Consider the ring $A = k[x, y, z]/(xy - z^2)$. Show that (x, z) is not a principal ideal.

As $\dim A = 2$ (why?), and $A/(x, z) \cong k[y]$ has dimension 1, we see that this ideal is codimension 1 (as codimension is the difference of dimensions for irreducible varieties, Theorem 12.2.9). Our geometric picture is that $\text{Spec } A$ is a cone (we can diagonalize the quadric as $xy - z^2 = ((x+y)/2)^2 - ((x-y)/2)^2 - z^2$, at least if $\text{char } k \neq 2$ — see Exercise 5.4.J), and that (x, z) is a line on the cone. (See Figure 13.1 for a sketch.) This suggests that we look at the cone point.

13.1.4. Solution. Let $\mathfrak{m} = (x, y, z)$ be the maximal ideal corresponding to the origin. Then $\text{Spec } A$ has Zariski tangent space of dimension 3 at the origin, and $\text{Spec } A/(x, z)$ has Zariski tangent space of dimension 1 at the origin. But $\text{Spec } A/(f)$




FIGURE 13.1. $V(x, z) \subset \operatorname{Spec} k[x, y, z]/(xy - z^2)$ is a line on a cone

must have Zariski tangent space of dimension at least 2 at the origin by Exercise 13.1.B. (As an added bonus, we have shown that A is not a Unique Factorization Domain, by Lemma 12.1.7.)

13.1.5. ★ Remark. Another approach to solving the problem, not requiring the definition of the Zariski tangent space, is to use the fact that the ring is graded (where x , y , and z each have degree 1), and the ideal (x, z) is a homogeneous ideal. (You may enjoy thinking this through.) The advantage of using the tangent space is that it applies to more general situations where there is no grading. For example, (a) (x, z) is not a principal ideal of $k[x, y, z]/(xy - z^2 - z^3)$. As a different example, (b) (x, z) is not a principal ideal of the local ring $(k[x, y, z]/(xy - z^2))_{(x, y, z)}$ (the “germ of the cone”). However, we remark that the graded case is still very useful. The construction of replacing a filtered ring by its “associated graded” ring can turn more general rings into graded rings (and can be used to turn example (a) into the graded case). The construction of completion can turn local rings into graded local rings (and can be used to turn example (b) into, essentially, the graded case). Filtered rings will come up in §13.9, the associated graded construction will implicitly come up in our discussions of the blow-up in §23.3, and many aspects of completions will be described in Chapter 30.

13.1.D. EXERCISE. Show that $(x, z) \subset k[w, x, y, z]/(wz - xy)$ is a codimension 1 ideal that is not principal, using the method of Solution 13.1.4. (See Figure 13.2 for the projectivization of this situation — a line on a smooth quadric surface.) This example was promised just after Exercise 5.4.D. An improvement is given in Exercise 15.2.R.

13.1.E. EXERCISE. Let $A = k[w, x, y, z]/(wz - xy)$. Show that $\operatorname{Spec} A$ is not factorial. (Exercise 5.4.L shows that A is not a unique factorization domain, but this is not enough — why is the localization of A at the prime (w, x, y, z) not factorial? One possibility is to do this “directly”, by trying to imitate the solution to Exercise 5.4.L, but this is hard. Instead, use the intermediate result that in a unique factorization domain, any codimension 1 prime is principal, Lemma 12.1.7, and considering Exercise 13.1.D.) As A is integrally closed if $\operatorname{char} k \neq 2$ (Exercise 5.4.I(c)), this yields an example of a scheme that is normal but not factorial, as promised in Exercise 5.4.F. A slight generalization will be given in 23.4.N.




FIGURE 13.2. The line $V(x, z)$ on the smooth quadric surface $V(wz - xy) \subset \mathbb{P}^3$.

13.1.F. LESS IMPORTANT EXERCISE (“HIGHER-ORDER DATA”). (This exercise is fun, but won’t be used.)

(a) In Exercise 3.7.B, you computed the equations cutting out the (union of the) three coordinate axes of \mathbb{A}_k^3 . (Call this scheme X .) Your ideal should have had three generators. Show that the ideal cannot be generated by fewer than three elements. (Hint: working modulo $\mathfrak{m} = (x, y, z)$ won’t give any useful information, so work modulo a higher power of \mathfrak{m} .)

(b) Show that the coordinate axes in \mathbb{A}_k^3 are not a regular embedding in \mathbb{A}_k^3 . (This was promised at the end of §9.5.)

13.1.6. Morphisms and tangent spaces. Suppose $\pi: X \rightarrow Y$, and $\pi(p) = q$. Then if we were in the category of differentiable manifolds, we would expect a tangent map, from the tangent space of p to the tangent space at q . Indeed that is the case; we have a map of stalks $\mathcal{O}_{Y,q} \rightarrow \mathcal{O}_{X,p}$, which sends the maximal ideal of the former \mathfrak{n} to the maximal ideal of the latter \mathfrak{m} (we have checked that this is a “local morphism” when we briefly discussed locally ringed spaces, see §7.3.1). Thus \mathfrak{n}^2 maps to \mathfrak{m}^2 , from which we see that $\mathfrak{n}/\mathfrak{n}^2$ maps to $\mathfrak{m}/\mathfrak{m}^2$. If $(\mathcal{O}_{X,p}, \mathfrak{m})$ and $(\mathcal{O}_{Y,q}, \mathfrak{n})$ have the same residue field κ , so $\mathfrak{n}/\mathfrak{n}^2 \rightarrow \mathfrak{m}/\mathfrak{m}^2$ is a linear map of κ -vector spaces, we have a natural map $(\mathfrak{m}/\mathfrak{m}^2)^\vee \rightarrow (\mathfrak{n}/\mathfrak{n}^2)^\vee$. This is the map from the tangent space of p to the tangent space at q that we sought. (Aside: note that the *cotangent* map *always* exists, without requiring p and q to have the same residue field — a sign that cotangent spaces are more natural than tangent spaces in algebraic geometry.)

13.1.7. Practice.

Here are some exercises to give you practice with the Zariski tangent space. If you have some differential geometric background, the first will further convince you that this definition correctly captures the idea of (co)tangent space.

13.1.G. IMPORTANT EXERCISE (THE JACOBIAN COMPUTES THE ZARISKI COTANGENT SPACE). Suppose X is a finite type k -scheme. Then locally it is of the form $\text{Spec } k[x_1, \dots, x_n]/(f_1, \dots, f_r)$. Show that the Zariski cotangent space at a k -valued

point (a closed point with residue field k) is given by the cokernel of the Jacobian map $k^r \rightarrow k^n$ given by the Jacobian matrix

$$(13.1.7.1) \quad J = \begin{pmatrix} \frac{\partial f_1}{\partial x_1}(p) & \cdots & \frac{\partial f_r}{\partial x_1}(p) \\ \vdots & \ddots & \vdots \\ \frac{\partial f_1}{\partial x_n}(p) & \cdots & \frac{\partial f_r}{\partial x_n}(p) \end{pmatrix}.$$

(This makes precise our example of a curve in \mathbb{A}^3 cut out by a couple of equations, where we picked off the linear terms, see Example 13.1.2.) You might be alarmed: what does $\frac{\partial f}{\partial x_i}$ mean? Do you need deltas and epsilons? No! Just define derivatives formally, e.g.,

$$\frac{\partial}{\partial x_1}(x_1^2 + x_1 x_2 + x_2^2) = 2x_1 + x_2.$$

Hint: Do this first when p is the origin, and consider linear terms, just as in Example 13.1.2 and Exercise 13.1.B. For the general case, “translate p to the origin”.

13.1.8. Remark. This result can be extended to closed points of X whose residue field is separable over k (and in particular, to *all* closed points if $\text{char } k = 0$ or if k is finite), see Remark 22.3.10. The fact that we wait until Chapter 22 to show this does not mean that it needs to be so complicated; a more elementary proof is possible.

13.1.9. Warning. It is more common in mathematics (but not universal) to define the Jacobian matrix as the transpose of this. But it will be more convenient for us to follow this minority convention.

13.1.H. EXERCISE (THE CORANK OF THE JACOBIAN IS INDEPENDENT OF THE PRESENTATION). Suppose A is a finitely generated k -algebra, generated by x_1, \dots, x_n , with ideal of relations I generated by f_1, \dots, f_r . Let p be a point of $\text{Spec } A$.

(a) Suppose $g \in I$. Show that appending the column of partials of g to the Jacobian matrix (13.1.7.1) does not change the **corank** (the dimension of the cokernel) at p . Hence show that the corank of the Jacobian matrix at p does not depend on the choice of generators of I .

(b) Suppose $q(x_1, \dots, x_n) \in k[x_1, \dots, x_n]$. Let h be the polynomial $y - q(x_1, \dots, x_n) \in k[x_1, \dots, x_n, y]$. Show that the Jacobian matrix of (f_1, \dots, f_r, h) with respect to the variables (x_1, \dots, x_n, y) has the same corank at p as the Jacobian matrix of (f_1, \dots, f_r) with respect to (x_1, \dots, x_n) . Hence show that the corank of the Jacobian matrix at p is independent of the choice of generators for A .

13.1.I. EXERCISE. Suppose X is a k -scheme. Describe a natural bijection from $\text{Mor}_k(\text{Spec } k[\epsilon]/(\epsilon^2), X)$ to the data of a point p with residue field k (necessarily a closed point) and a tangent vector at p . (This is important, for example in deformation theory.)

13.1.J. EXERCISE. Find the dimension of the Zariski tangent space at the point $[(2, 2i)]$ of $\mathbb{Z}[2i] \cong \mathbb{Z}[x]/(x^2 + 4)$. Find the dimension of the Zariski tangent space at the point $[(2, x)]$ of $\mathbb{Z}[\sqrt{-2}] \cong \mathbb{Z}[x]/(x^2 + 2)$. (If you prefer geometric versions of the same examples, replace \mathbb{Z} by \mathbb{C} , and 2 by y : consider $\mathbb{C}[x, y]/(x^2 + y^2)$ and $\mathbb{C}[x, y]/(x^2 + y)$.)

13.2 Regularity, and smoothness over a field

The key idea in the definition of regularity is contained in the following result, that “the dimension of the Zariski tangent space is at least the dimension of the local ring”.

13.2.1. Theorem. — *For any Noetherian local ring (A, \mathfrak{m}, k) , $\dim A \leq \dim_k \mathfrak{m}/\mathfrak{m}^2$.*

13.2.2. Proof of Theorem 13.2.1. Note that \mathfrak{m} is finitely generated (as A is Noetherian), so $\mathfrak{m}/\mathfrak{m}^2$ is a finitely generated $(A/\mathfrak{m} = k)$ -module, hence finite-dimensional. Say $\dim_k \mathfrak{m}/\mathfrak{m}^2 = n$. Choose a basis of $\mathfrak{m}/\mathfrak{m}^2$, and lift it to elements f_1, \dots, f_n of \mathfrak{m} . Then by Nakayama’s Lemma (version 4, Exercise 8.2.H), $(f_1, \dots, f_n) = \mathfrak{m}$.

Then by Exercise 12.3.H (a consequence of Krull’s Height Theorem 12.3.9), $\dim A \leq n$. \square

We remark that the f_i are a system of parameters for A (Definition 12.3.10), and are “particularly good ones”.

13.2.3. Definition. If equality holds in Theorem 13.2.1, we say that A is a **regular local ring**. In particular, *regular local rings are Noetherian by definition*. (If a Noetherian ring A is regular at all of its prime ideals, i.e., if $A_{\mathfrak{p}}$ is a regular local ring for all prime ideals \mathfrak{p} of A , then A is said to be a **regular ring**. We basically won’t use this terminology. Fact 13.8.2 reassuringly implies that regular local rings are indeed regular rings.) A locally Noetherian scheme X is **regular** at a point p if the local ring $\mathcal{O}_{X,p}$ is regular. (The word **nonsingular** is often used as well, notably in the case when X is finite type over a field, but for the sake of consistency we will use “regular” throughout.) It is **singular** at the point otherwise, and we say that the point is a **singularity**. (The word **nonregular** is also used, but for the sake of consistency, we will use “singular”, despite the fact that this choice is *inconsistent* with our choice of “regular” over “nonsingular”. Possible motivation for the inconsistency: “regular local ring” and “singularity” are both standard terminology, so at some point we are forced to make a choice.) A scheme is **regular** (or **nonsingular**) if it is regular at all points. It is **singular** (or **nonregular**) otherwise (i.e., if it is singular at *at least one* point — if it has a singularity).

13.2.A. EXERCISE. Show that a dimension 0 Noetherian local ring is regular if and only if it is a field.

You will hopefully gradually become convinced that this is the right notion of “smoothness” of schemes. Remarkably, Krull introduced the notion of a regular local ring for purely algebraic reasons, some time before Zariski realized that it was a fundamental notion in geometry in 1947.

13.2.B. EXERCISE. Suppose p is a regular closed point of X , so $(\mathcal{O}_{X,p}, \mathfrak{m})$ is a regular local ring, of dimension n , say. Suppose $f \in \mathcal{O}_{X,p}$. Show that $\mathcal{O}_{X,p}/(f)$ is a regular local ring of dimension $n - 1$ if and only if $f \in \mathfrak{m} \setminus \mathfrak{m}^2$. (Hint: Krull’s Principal Ideal Theorem for tangent spaces, Exercise 13.1.B.)

13.2.C. EXERCISE (THE SLICING CRITERION FOR REGULARITY). Suppose X is a Noetherian scheme (such as a variety), D is an effective Cartier divisor on X (Definition 9.5.1), and $p \in D$. Show that if p is a regular point of D then p is a regular point of X . (Hint: Krull's Principal Ideal Theorem for tangent spaces, Exercise 13.1.B.)

13.2.4. The Jacobian criterion.

A finite type k -scheme is locally of the form $\operatorname{Spec} k[x_1, \dots, x_n]/(f_1, \dots, f_r)$. The Jacobian criterion for regularity (Exercise 13.2.D) gives a hands-on method for checking for singularity at closed points, using the equations f_1, \dots, f_r , if $k = \bar{k}$.

13.2.D. IMPORTANT EXERCISE (THE JACOBIAN CRITERION — EASY, GIVEN EXERCISE 13.1.G). Suppose $X = \operatorname{Spec} k[x_1, \dots, x_n]/(f_1, \dots, f_r)$ has pure dimension d . Show that a k -valued point $p \in X$ is regular if and only if the corank of the Jacobian matrix (13.1.7.1) (the dimension of the cokernel) at p is d .

13.2.E. EASY EXERCISE. Suppose $k = \bar{k}$. Show that the singular *closed* points of the hypersurface $f(x_1, \dots, x_n) = 0$ in \mathbb{A}_k^n are given by the equations

$$f = \frac{\partial f}{\partial x_1} = \dots = \frac{\partial f}{\partial x_n} = 0.$$

(Translation: the singular points of $f = 0$ are where the gradient of f vanishes. This is not shocking.)

13.2.5. Smoothness over a field.

There seem to be two serious drawbacks with the Jacobian criterion. For finite type schemes over \bar{k} , the criterion gives a necessary condition for regularity, but it is not obviously sufficient, as we need to check regularity at non-closed points as well. We can prove sufficiency by working hard to show Fact 13.8.2, which implies that the non-closed points must be regular as well. A second failing is that the criterion requires k to be algebraically closed. These problems suggest that old-fashioned ideas of using derivatives and Jacobians are ill-suited to the fancy modern notion of regularity. But that is wrong — the fault is with the concept of regularity. There is a better notion of *smoothness over a field*. Better yet, this idea generalizes to the notion of a smooth morphism of schemes (to be discussed in §13.6, and again in Chapter 26), which behaves well in all possible ways (including in the sense of §8.1). This is another sign that some properties we think of as of objects (“absolute notions”) should really be thought of as properties of morphisms (“relative notions”). We know enough to imperfectly (but correctly) define what it means for a scheme to be **k -smooth**, or **smooth over k** .

13.2.6. Definition. A k -scheme is **k -smooth of dimension d** , or **smooth of dimension d over k** , if it is of pure dimension d , and there exists a cover by affine open sets $\operatorname{Spec} k[x_1, \dots, x_n]/(f_1, \dots, f_r)$ where the Jacobian matrix has corank d at all points. (In particular, it is locally of finite type.) A k -scheme is **smooth** over k if it is smooth of some dimension. The k is often omitted when it is clear from context.

13.2.F. EXERCISE (FIRST EXAMPLES).

(a) Show that \mathbb{A}_k^n is smooth for any n and k . For which characteristics is the curve $y^2z = x^3 - xz^2$ in \mathbb{P}_k^2 smooth (cf. Exercise 13.3.C)?

(b) Suppose $f \in k[x_1, \dots, x_n]$ is a polynomial such that the system of equations

$$f = \frac{\partial f}{\partial x_1} = \dots = \frac{\partial f}{\partial x_n} = 0$$

has no solutions in \bar{k} . Show that the hypersurface $f = 0$ in \mathbb{A}_k^n is smooth. (Compare this to Exercise 13.2.E, which has the additional hypothesis $k = \bar{k}$.)

13.2.G. EXERCISE (SMOOTHNESS OF VARIETIES IS PRESERVED BY EXTENSION OF BASE FIELD). Suppose X is a finite type k -scheme, and $k \subset \ell$ is a field extension. Show that if X is smooth over k then $X \times_{\operatorname{Spec} k} \operatorname{Spec} \ell$ is smooth over ℓ . (The converse will be proved in Exercise 22.3.C.)

The next exercise shows that we need only check closed points, thereby making a connection to classical geometry.

13.2.H. EXERCISE. Show that if the Jacobian matrix for $X = \operatorname{Spec} k[x_1, \dots, x_n]/(f_1, \dots, f_r)$ has corank d at all *closed* points, then it has corank d at *all* points. (Hint: the locus where the Jacobian matrix has corank d can be described in terms of vanishing and nonvanishing of determinants of certain explicit matrices.)

13.2.7. You can check that any open subset of a smooth k -variety is also a smooth k -variety. With what we know now, we could show that this implies that k -smoothness is equivalent to the Jacobian being corank d everywhere for *every* affine open cover (and by *any* choice of generators of the ring corresponding to such an open set). Indeed, you should feel free to do this if you cannot restrain yourself. But the cokernel of the Jacobian matrix is secretly the space of differentials (which might not be surprising if you have experience with differentials in differential geometry), so this will come for free when we give a better version of this definition in Definition 22.3.1. The current imperfect definition will suffice for us to work out examples. And if you don't want to wait until Definition 22.3.1, you can use Exercise 13.2.I below to show that if k is algebraically closed, then smoothness can be checked on any open cover.

13.2.8. In defense of regularity. Having made a spirited case for smoothness, we should be clear that regularity is still very useful. For example, it is the only concept which makes sense in mixed characteristic. In particular, \mathbb{Z} is regular at its points (it is a regular ring), and more generally, discrete valuation rings are incredibly useful examples of regular rings.

13.2.9. Regularity vs. smoothness.

13.2.I. EXERCISE. Suppose X is a locally finite type scheme of pure dimension d over an algebraically closed field $k = \bar{k}$. Show that X is regular at its closed points if and only if it is smooth. (We will soon learn that for locally finite type \bar{k} -schemes, regularity at closed points is the same as regularity everywhere, Theorem 13.8.3.) Hint to show regularity implies smoothness: use the Jacobian criterion to show that the corank of the Jacobian is d at the closed points of X . Then use Exercise 13.2.H.

More generally, if k is perfect (e.g., if $\operatorname{char} k = 0$ or k is a finite field), then smoothness is the same as regularity at closed points (see Exercise 22.3.D). More

generally still, we will later prove the following fact. We mention it now because it will make a number of statements cleaner long before we finally prove it. (There will be no circularity.)

13.2.10. Smoothness-Regularity Comparison Theorem. —

- (a) If k is perfect, every regular finite type k -scheme is smooth over k .
- (b) Every smooth k -scheme is regular (with no hypotheses on perfection).

Part (a) will be proved in Exercise 22.3.D. Part (b) will be proved in §26.2.3. The fact that Theorem 13.2.10 will be proved so far in the future does not mean that we truly need to wait that long. We could prove both parts by the end of this chapter, with some work. We instead postpone the proof until we have machinery that will do much of the work for us. But if you wish some insight right away, here is the outline of the argument for (b), which you can even try to implement after reading this chapter. We will soon show (in Exercise 13.3.O) that \mathbb{A}_k^n is regular for all fields k . Once we know the definition of étale morphism, we will realize that every smooth variety locally admits an étale morphism to \mathbb{A}_k^n (Observation 13.6.4). You can then show that if $\pi: X \rightarrow Y$ is an étale morphism of locally Noetherian schemes, and $p \in X$, then p is regular if $\pi(p)$ is regular. This will be shown in Exercise 26.2.D, but you could reasonably do this after reading the definition of étaleness.

13.2.11. Caution: Regularity does not imply smoothness. If k is not perfect, then regularity does *not* imply smoothness, as demonstrated by the following example. Let $k = \mathbb{F}_p(u)$, and consider the hypersurface $X = \operatorname{Spec} k[x]/(x^p - u)$. Now $k[x]/(x^p - u)$ is a field, hence regular. But if $f(x) = x^p - u$, then $f(u^{1/p}) = \frac{df}{dx}(u^{1/p}) = 0$, so the Jacobian criterion fails — X is not smooth over k . (Never forget that smoothness requires a choice of field — it is a “relative” notion, and we will later define smoothness over an arbitrary scheme, in §13.6.) Technically, this argument is not yet complete: as noted in §13.2.7, we have not shown that it suffices to check the Jacobian on *any* affine cover. But as mentioned in §13.2.7, this will be rectified when we give a better definition of smoothness in Definition 22.3.1.

In case the previous example is too “small” to be enlightening (because the scheme in question *is* smooth over a *different* field, namely $k[x]/(x^p - u)$), here is another. Let $k = \mathbb{F}_p(u)$ as before, with $p > 2$, and consider the curve $\operatorname{Spec} k[x, y]/(y^2 - x^p + u)$. Then the closed point $(y, x^p - u)$ is regular but not smooth.

Thus you should not use “regular” and “smooth” interchangeably.

13.2.12. Regular local rings are integral domains.

You might expect from geometric intuition that a scheme is “locally irreducible” at a “smooth” point. Put algebraically:

13.2.13. Theorem. — Suppose (A, \mathfrak{m}, k) is a regular local ring of dimension n . Then A is an integral domain.

Before proving it, we give some consequences.

13.2.J. EXERCISE. Suppose p is a regular point of a locally Noetherian scheme X . Show that only one irreducible component of X passes through p .

13.2.K. EASY EXERCISE. Show that a nonempty regular Noetherian scheme is irreducible if and only if it is connected.

13.2.L. IMPORTANT EXERCISE (REGULAR SCHEMES IN REGULAR SCHEMES ARE REGULAR EMBEDDINGS).

- (a) Suppose (A, \mathfrak{m}, k) is a regular local ring of dimension n , and $I \subset A$ is an ideal of A cutting out a regular local ring of dimension d . Let $r = n - d$. Show that $\text{Spec } A/I$ is a regular embedding in $\text{Spec } A$. Hint: show that there are elements f_1, \dots, f_r of I spanning the k -vector space $I/(I \cap \mathfrak{m}^2)$. Show that the quotient of A by both (f_1, \dots, f_r) and I yields dimension d regular local rings. Show that a surjection of integral domains of the same dimension must be an isomorphism.
- (b) Suppose $\pi: X \rightarrow Y$ is a closed embedding of regular schemes. Show that π is a regular embedding.

Exercise 13.2.L has the following striking geometric consequence.

13.2.M. EXERCISE (GENERALIZING EXERCISE 12.3.D). Suppose dimension d , and X and Y are equidimensional subvarieties (possibly singular) of \mathbb{A}_k^d codimension m and n respectively. Show that every component of $X \cap Y$ has codimension at most $m + n$ in \mathbb{A}_k^d as follows. Show that the diagonal $\mathbb{A}_k^d \cong \Delta \subset \mathbb{A}_k^d \times_k \mathbb{A}_k^d \cong \mathbb{A}_k^{2d}$ is a regular embedding of codimension d . Then follow the rest of the hint to Exercise 12.3.D.

13.2.14. Remark. The following example shows that the regularity hypotheses in Exercise 13.2.M cannot be (completely) dropped. Let $W = \text{Spec } k[w, x, y, z]/(wz - xy)$ be the cone over the smooth quadric surface, which is an integral threefold. Let X be the surface $w = x = 0$ and Y the surface $y = z = 0$; both lie in W . Then $X \cap Y$ is just the origin, so we have two codimension 1 subvarieties meeting in a codimension 3 subvariety. (It is no coincidence that X and Y are the affine cones over two lines in the same ruling, see Exercise 9.3.L. This example will arise again in Exercise 23.4.N.)

13.2.15. Proof of Theorem 13.2.13 (following [Liu, Prop. 2.11]). We prove the result by induction on n . The case $n = 0$ follows from Exercise 13.2.A (dimension 0 regular local rings are fields, which are certainly integral domains). We now assume $n > 0$, and that we have proved the result for smaller dimension. Fix any $f \in \mathfrak{m} \setminus \mathfrak{m}^2$. By Exercise 13.1.B, the Zariski tangent space of the local ring $A/(f)$ (at its maximal ideal \mathfrak{m}) has dimension $n - 1$. By Exercise 12.3.G, $\dim A/(f) \geq n - 1$. Thus by Theorem 13.2.1, $\dim A/(f) = n - 1$, and $A/(f)$ is a regular local ring. By our inductive hypothesis, $A/(f)$ is an integral domain (of dimension $n - 1$).

Now choose any minimal prime ideal \mathfrak{p} of A with $\dim A/\mathfrak{p} = n$. We wish to show that $\mathfrak{p} = (0)$. The Zariski cotangent space of A/\mathfrak{p} is a quotient of that of A , and thus has dimension at most n . But $\dim A/\mathfrak{p} = n$, so by Theorem 13.2.1, A/\mathfrak{p} is a regular local ring of dimension n . By the argument of the previous paragraph with A replaced by A/\mathfrak{p} , we see that $A/(\mathfrak{p} + (f))$ is an integral domain of dimension $n - 1$. But it is a quotient of $A/(f)$. The only way one integral domain can be the quotient of another of the same dimension is if the quotient is an isomorphism (as discussed in the hint to Exercise 13.2.L(a)).

Thus $\mathfrak{p} + (f) = (f)$, i.e., $\mathfrak{p} \subset fA$. Hence each element u of \mathfrak{p} can be written as fv for some $v \in A$. As $f \notin \mathfrak{p}$ (as $\dim A/(\mathfrak{p} + (f)) = n - 1 < n = \dim A/\mathfrak{p}$), we have $v \in \mathfrak{p}$, and so $\mathfrak{p} \subset f\mathfrak{p}$.

Clearly $f\mathfrak{p} \subset \mathfrak{p}$, so $\mathfrak{p} = f\mathfrak{p}$. Then by the second version of Nakayama's Lemma 8.2.9, $\mathfrak{p} = (0)$ as desired. \square

13.2.16. $\star\star$ Checking regularity of k -schemes at closed points by base changing to \bar{k} .

(We revisit these ideas using a different approach in 22.3.E, so you should read this only if you are particularly curious.) The Jacobian criterion is a great criterion for checking regularity of finite type k -schemes at k -valued points. The following result extends its applicability to more general closed points.

Suppose X is a finite type k -scheme of pure dimension n , and $p \in X$ is a closed point with residue field k' . By the Nullstellensatz 3.2.6, k'/k is a finite extension of fields; suppose that it is separable. Define $\pi: X_{\bar{k}} := X \times_k \bar{k} \rightarrow X$ by base change from $\text{Spec } \bar{k} \rightarrow \text{Spec } k$.

13.2.N. EXERCISE.

- (a) Suppose $f(x) \in k[x]$ is a separable polynomial (i.e., f has distinct roots in \bar{k}), and irreducible, so $k'' := k[x]/(f(x))$ is a field extension of k . Show that $k'' \otimes_k \bar{k}$ is, as a ring, $\bar{k} \times \cdots \times \bar{k}$, where there are $\deg f = \deg k''/k$ factors.
(b) Show that $\pi^{-1}(p)$ consists of $\deg(k'/k)$ *reduced* points.

13.2.O. EXERCISE. Suppose p is a closed point of X , with residue field k' that is separable over k of degree d . Show that $X_{\bar{k}}$ is regular at all the preimages p_1, \dots, p_d of p if and only if X is regular at p as follows.

- (a) Reduce to the case $X = \text{Spec } A$.
(b) Let $\mathfrak{m} \subset A$ be the maximal ideal corresponding to p . By tensoring the exact sequence $0 \rightarrow \mathfrak{m} \rightarrow A \rightarrow k' \rightarrow 0$ with \bar{k} (field extensions preserve exactness of sequences of vector spaces), we have the exact sequence

$$0 \rightarrow \mathfrak{m} \otimes_k \bar{k} \rightarrow A \otimes_k \bar{k} \rightarrow k' \otimes_k \bar{k} \rightarrow 0.$$

Show that $\mathfrak{m} \otimes_k \bar{k} \subset A \otimes_k \bar{k}$ is the ideal corresponding to the pullback of p to $\text{Spec } A \otimes_k \bar{k}$. Verify that $(\mathfrak{m} \otimes_k \bar{k})^2 = \mathfrak{m}^2 \otimes_k \bar{k}$.

- (c) By tensoring the short exact sequence of k -vector spaces $0 \rightarrow \mathfrak{m}^2 \rightarrow \mathfrak{m} \rightarrow \mathfrak{m}/\mathfrak{m}^2 \rightarrow 0$ with \bar{k} , show that

$$\sum_{i=1}^d \dim_{\bar{k}} T_{X_{\bar{k}}, p_i} = d \dim_{k'} T_{X, p}.$$

- (d) Use Exercise 12.1.G(a) and the inequalities $\dim_{\bar{k}} T_{X_{\bar{k}}, p_i} \geq \dim X_{\bar{k}}$ and $\dim_{k'} T_{X, p} \geq \dim X$ (Theorem 13.2.1) to conclude.

13.2.17. Remark. In fact, regularity at a single p_i is enough to conclude regularity at p . You can show this by following up on Exercise 13.2.O; first deal with the case when k'/k is Galois, and obtain some transitive group action of $\text{Gal}(k'/k)$ on $\{p_1, \dots, p_d\}$. Another approach is given in Exercise 22.3.E.

13.2.18. Remark. In Exercise 22.3.D, we will use this to prove that a variety over a perfect field is smooth if and only if it is regular at all closed points (cf. the Smoothness-Regularity Comparison Theorem 13.2.10).

13.3 Examples

We now explore regularity in practice through a series of examples and exercises. Much of this discussion is secretly about smoothness rather than regularity. In particular, in order to use the Jacobian criterion, we will usually work over an algebraically closed field.

13.3.1. Geometric examples.

13.3.A. EASY EXERCISE. Suppose k is a field. Show that \mathbb{A}_k^1 and \mathbb{A}_k^2 are regular, by directly checking the regularity of all points. Show that \mathbb{P}_k^1 and \mathbb{P}_k^2 are regular. (The generalization to arbitrary dimension is harder, so we leave it to Exercise 13.3.O.)

13.3.B. EXERCISE (THE EULER OR JACOBIAN TEST FOR PROJECTIVE HYPERSURFACES). Suppose $k = \bar{k}$. Show that the singular *closed* points of the hypersurface $f = 0$ in \mathbb{P}_k^n correspond to the locus

$$f = \frac{\partial f}{\partial x_0} = \cdots = \frac{\partial f}{\partial x_n} = 0.$$

If the degree of the hypersurface is not divisible by $\text{char } k$ (e.g., if $\text{char } k = 0$), show that it suffices to check $\frac{\partial f}{\partial x_0} = \cdots = \frac{\partial f}{\partial x_n} = 0$. Hint: show that $(\deg f)f = \sum_i x_i \frac{\partial f}{\partial x_i}$. (In fact, this will give the singular points in general, not just the singular closed points, cf. §13.2.5. We won't use this, so we won't prove it.)

13.3.C. EXERCISE. Suppose that $k = \bar{k}$ does not have characteristic 2. Show that $y^2z = x^3 - xz^2$ in \mathbb{P}_k^2 is an irreducible regular curve. (Eisenstein's Criterion gives one way of showing irreducibility. Warning: we didn't specify $\text{char } k \neq 3$, so be careful when using the Euler test.)

13.3.D. EXERCISE. Suppose $k = \bar{k}$ has characteristic not 2. Show that a quadric hypersurface in \mathbb{P}^n is regular if and only if it is maximal rank. ("Maximal rank" was defined in Exercise 5.4.J.)

13.3.E. EXERCISE. Suppose $k = \bar{k}$ has characteristic 0. Show that there exists a regular (projective) plane curve of degree d . Hint: try a "Fermat curve" $x^d + y^d + z^d = 0$. (Feel free to weaken the hypotheses. Bertini's Theorem 13.4.2 will give another means of showing existence.)

13.3.F. EXERCISE (SEE FIGURE 13.3). Find all the singular closed points of the following plane curves. Here we work over $k = \bar{k}$ of characteristic 0 to avoid distractions.

- (a) $y^2 = x^2 + x^3$. This is an example of a *node*.
- (b) $y^2 = x^3$. This is called a *cusp*; we met it earlier in Exercise 10.7.F.
- (c) $y^2 = x^4$. This is called a *tacnode*; we met it earlier in Exercise 10.7.G.

(A precise definition of a node etc. will be given in Definition 30.3.1.)



FIGURE 13.3. Plane curve singularities

13.3.G. EXERCISE. Suppose $k = \bar{k}$. Use the Jacobian criterion appropriately to show that the twisted cubic $\text{Proj } k[w, x, y, z]/(wz - xy, wy - x^2, xz - y^2)$ is regular. (You can do this, without any hypotheses on k , using the fact that it is isomorphic to \mathbb{P}^1 . But do this with the explicit equations, for the sake of practice. The twisted cubic was defined in Exercise 9.3.A.)

13.3.2. Tangent planes and tangent lines.

Suppose a scheme $X \subset \mathbb{A}^n$ is cut out by equations f_1, \dots, f_r , and X is regular of dimension d at the k -valued point $a = (a_1, \dots, a_n)$. Then the **tangent d -plane to X at p** (sometimes denoted $T_p X$) is given by the r equations

$$\left(\frac{\partial f_i}{\partial x_1}(a) \right) (x_1 - a_1) + \dots + \left(\frac{\partial f_i}{\partial x_n}(a) \right) (x_n - a_n) = 0,$$

where (as in (13.1.7.1)) $\frac{\partial f_i}{\partial x_1}(a)$ is the evaluation of $\frac{\partial f_i}{\partial x_1}$ at $a = (a_1, \dots, a_n)$.

13.3.H. EXERCISE. Why is this independent of the *choice* of defining equations f_1, \dots, f_r of X ?

The Jacobian criterion (Exercise 13.2.D) ensures that these r equations indeed cut out a d -plane. (If $d = 1$, this is called the **tangent line**.) This can be readily shown to be notion of tangent plane that we see in multivariable calculus, but note that here this is the *definition*, and thus don't have to worry about δ 's and ϵ 's. Instead we will have to just be careful that it behaves the way we want it to.

13.3.I. EXERCISE. Compute the tangent line to the curve of Exercise 13.3.F(b) at $(1, 1)$.

13.3.J. EXERCISE. Suppose $X \subset \mathbb{P}_k^n$ (k as usual a field) is cut out by homogeneous equations f_1, \dots, f_r , and $p \in X$ is a k -valued point that is regular of dimension d . Define the (projective) tangent d -plane to X at p . (Definition 9.3.3 gives the definition of a d -plane in \mathbb{P}_k^n , but you shouldn't need to refer there.)

13.3.3. Side Remark to help you think cleanly. We would want the definition of tangent k -plane to be natural in the sense that for any automorphism σ of \mathbb{A}_k^n (or, in the case of the previous Exercise, \mathbb{P}_k^n), $\sigma(T_p X) = T_{\sigma(p)} \sigma(X)$. You could verify this by hand, but you can also say this in a cleaner way, by interpreting the equations cutting out the tangent space in a coordinate free manner. Informally speaking, we are using the canonical identification of n -space with the tangent space to n -space

at p , and using the fact that the Jacobian “linear transformation” cuts out $T_p X$ in $T_p \mathbb{A}^n$ in a way independent of choice of coordinates on \mathbb{A}^n or defining equations of X . Your solution to Exercise 13.3.H will help you start to think in this way.

13.3.K. EXERCISE. Suppose $X \subset \mathbb{P}_k^n$ is a degree d hypersurface cut out by $f = 0$, and L is a line not contained in X . Exercise 9.3.D (a case of Bézout’s Theorem) showed that X and L meet at d points, counted “with multiplicity”. Suppose L meets X “with multiplicity at least 2” at a k -valued point $p \in L \cap X$, and that p is a regular point of X . Show that L is contained in the tangent plane to X at p . (Do you have a picture of this in your mind?)

13.3.4. Arithmetic examples.

13.3.L. EASY EXERCISE. Show that $\text{Spec } \mathbb{Z}$ is a regular curve.

13.3.M. EXERCISE. (This tricky exercise is for those who know about the primes of the Gaussian integers $\mathbb{Z}[i]$.) There are several ways of showing that $\mathbb{Z}[i]$ is dimension 1. (For example: (i) it is a principal ideal domain; (ii) it is the normalization of \mathbb{Z} in the field extension $\mathbb{Q}(i)/\mathbb{Q}$; (iii) using Krull’s Principal Ideal Theorem 12.3.3 and the fact that $\dim \mathbb{Z}[x] = 2$ by Exercise 12.1.H.) Show that $\text{Spec } \mathbb{Z}[i]$ is a regular curve. (There are several ways to proceed. You could use Exercise 13.1.B. As an example to work through first, consider the prime $(1 + i)$, which is cut out by the equation $1 + x$ in $\text{Spec } \mathbb{Z}[x]/(x^2 + 1)$.) We will later (§13.5.10) have a simpler approach once we discuss discrete valuation rings.

13.3.N. EXERCISE. Show that $[(5, 5i)]$ is the unique singular point of $\text{Spec } \mathbb{Z}[5i]$. (Hint: $\mathbb{Z}[i]_5 \cong \mathbb{Z}[5i]_5$. Use the previous exercise.)

13.3.5. Back to geometry: \mathbb{A}_k^n is regular.

The key step to showing that \mathbb{A}_k^n is regular (where k is a field) is the following.

13.3.6. Proposition. — Suppose (B, \mathfrak{n}, k) is a regular local ring of dimension d . Let $\phi: B \rightarrow B[x]$. Suppose \mathfrak{p} is a prime ideal of $A := B[x]$ such that $\mathfrak{n}B[x] \subset \mathfrak{p}$. Then $A_{\mathfrak{p}}$ is a regular local ring.

Proof. Geometrically: we have a morphism $\pi: X = \text{Spec } B[x] \rightarrow Y = \text{Spec } B$, and $\pi(\overline{[\mathfrak{p}]}) = [\mathfrak{n}]$. The fiber $\pi^{-1}([\mathfrak{n}])$ is $\text{Spec}(B[x]/\mathfrak{n}B[x]) = \text{Spec } k[x]$. Thus either (i) $\overline{[\mathfrak{p}]}$ is the fiber \mathbb{A}_k^1 above $[\mathfrak{n}]$, or (ii) $\overline{[\mathfrak{p}]}$ is a closed point of the fiber \mathbb{A}_k^1 . (See Figure 13.4.)




FIGURE 13.4. The geometry behind the proof of Proposition 13.3.6

Before considering these two cases, we make two remarks. As (B, \mathfrak{n}) is a regular local ring of dimension d , \mathfrak{n} is generated by d elements of B , say f_1, \dots, f_d (as discussed in the proof of Theorem 13.2.1), and there is a chain of prime ideals

$$(13.3.6.1) \quad 0 = \mathfrak{q}_0 \subset \mathfrak{q}_1 \subset \dots \subset \mathfrak{q}_d = \mathfrak{n}.$$

Case (i): $\overline{[p]}$ is the fiber \mathbb{A}_k^1 . In this case, $\mathfrak{p} = \mathfrak{n}B[x]$. Thus by Krull's Height Theorem 12.3.9, the codimension of $\mathfrak{p} = \mathfrak{n}B[x]$ is at most d , because \mathfrak{p} is generated by the d elements f_1, \dots, f_d . But \mathfrak{p} has codimension at least d : given the chain (13.3.6.1) ending with \mathfrak{n} , we have a corresponding chain of prime ideals of $B[x]$ ending with $\mathfrak{n}B[x]$. Hence the codimension of \mathfrak{p} is precisely d , and $\mathfrak{p}A_{\mathfrak{p}}$ is generated by f_1, \dots, f_d , implying that it is a regular local ring.

Case (ii): $\overline{[p]}$ is a closed point of the fiber \mathbb{A}_k^1 . The closed point \mathfrak{p} of $k[x]$ corresponds to some monic irreducible polynomial $g(x) \in k[x]$. Arbitrarily lift the coefficients of g to B ; we sloppily denote the resulting polynomial in $B[x]$ by $g(x)$ as well. Then $\mathfrak{p} = (f_1, \dots, f_d, g)$, so by Krull's Height Theorem 12.3.9, the codimension of \mathfrak{p} is at most $d + 1$. But \mathfrak{p} has codimension at least $d + 1$: given the chain (13.3.6.1) ending with \mathfrak{n} , we have a corresponding chain in $B[x]$ ending with $\mathfrak{n}B[x]$, and we can extend it by appending \mathfrak{p} . Hence the codimension of \mathfrak{p} is precisely $d + 1$, and $\mathfrak{p}A_{\mathfrak{p}}$ is generated by f_1, \dots, f_d, g , implying that it is a regular local ring. \square

13.3.O. EXERCISE. Use Proposition 13.3.6 to show that if X is a regular (locally Noetherian) scheme, then so is $X \times \mathbb{A}^1$. In particular, show that \mathbb{A}_k^n is regular. (Hint: think about fibers.)

13.4 Bertini's Theorem

We now discuss Bertini's Theorem, a fundamental classical result.

13.4.1. Definition: dual projective space. The **dual** (or **dual projective space**) to \mathbb{P}_k^n (with coordinates x_0, \dots, x_n), is informally the space of hyperplanes in \mathbb{P}_k^n . Somewhat more precisely, it is a projective space \mathbb{P}_k^n with coordinates a_0, \dots, a_n (which we denote $\mathbb{P}_k^{n\vee}$ with the futile intent of preventing confusion), along with the data of the "incidence variety" or "incidence correspondence" $I \subset \mathbb{P}^n \times \mathbb{P}^{n\vee}$ cut out by the equation $a_0x_0 + \dots + a_nx_n = 0$. Note that the k -valued points of $\mathbb{P}^{n\vee}$ indeed correspond to hyperplanes in \mathbb{P}^n defined over k , and this is also clearly a duality relation (there is a symmetry in the definition between the x -variables and the a -variables). So this is concrete enough to use in practice, and extends over an arbitrary base (notably $\text{Spec } \mathbb{Z}$). (But if you have a delicate and refined sensibility, you may want to come up with a coordinate-free definition.)

13.4.2. Bertini's Theorem. — Suppose X is a smooth subvariety of \mathbb{P}_k^n . Then there is a nonempty (=dense) open subset U of dual projective space $\mathbb{P}_k^{n\vee}$ such that for any closed point $[H] \in U$, H doesn't contain any component of X , and the scheme $H \cap X$ is k -smooth.

The hypothesis that k is algebraically closed is not essential (even in the proof below): see Exercise 13.4.A.

We remark that any theorem of this flavor is often called a “Bertini Theorem”. One example is the Kleiman-Bertini Theorem 26.3.7, which was *not* proved jointly by Kleiman and Bertini.

As an application of Bertini’s Theorem 13.4.2, a general degree $d > 0$ hypersurface in \mathbb{P}_k^n intersects X in a regular subvariety of codimension 1 in X : replace $X \hookrightarrow \mathbb{P}^n$ with the composition

$$X \longrightarrow \mathbb{P}^n \xrightarrow{\nu_d} \mathbb{P}^N$$

where ν_d is the d th Veronese embedding (9.3.8). Here “general” has its usual meaning in algebraic geometry, see §10.3.7, except that we are considering only *closed* points of $U \subset \mathbb{P}_k^{n\vee}$. A useful consequence of this, taking $X = \mathbb{P}^n$: we immediately see that there exists a smooth degree d hypersurface in \mathbb{P}^n (over an algebraically closed field \bar{k}), without any messing around with specific equations (as in the special case Exercise 13.3.E).

Exercise 26.3.D gives a useful improvement of Bertini’s Theorem in characteristic 0 (see Exercise 26.3.E).

Proof. In order to keep the proof as clean as possible, we assume X is irreducible, but essentially the same proof applies in general.

The central idea of the proof is quite naive. We will describe the hyperplanes that are “bad”, and show that they form a closed subset of dimension at most $n - 1$ of $\mathbb{P}_k^{n\vee}$, and hence that the complement is a dense open subset. Somewhat more precisely, we will define a projective variety $Z \subset X \times \mathbb{P}_k^{n\vee}$ that can informally be described as:

$$(13.4.2.1) \quad Z = \{(p \in X, H \subset \mathbb{P}_k^n) : p \in H, \\ \text{and either } p \text{ is a singular point of } H \cap X, \text{ or } X \subset H\}$$

We will see that the projection $\pi: Z \rightarrow X$ has fibers at closed points that are projective spaces of dimension $n - 1 - \dim X$, and use this to show that $\dim Z \leq n - 1$. Thus the image of Z in $\mathbb{P}_k^{n\vee}$ will be a closed subset (Theorem 8.4.7), of dimension of at most $n - 1$, so its complement will be open and nonempty. We now put this strategy into action.

13.4.3. We first define Z more precisely, in terms of equations on $\mathbb{P}^n \times \mathbb{P}^{n\vee}$, where the coordinates on \mathbb{P}^n are x_0, \dots, x_n , and the dual coordinates on $\mathbb{P}^{n\vee}$ are a_0, \dots, a_n . Suppose X is cut out by f_1, \dots, f_r . Then we take these equations as the first of the defining equations of Z . (So far we have defined the subscheme $X \times \mathbb{P}^{n\vee}$.) We also add the equation $a_0 x_0 + \dots + a_n x_n = 0$. (So far we have described the subscheme of $\mathbb{P}^n \times \mathbb{P}^{n\vee}$ corresponding to points (p, H) where $p \in X$ and $p \in H$.) Note that the Jacobian matrix ((13.1.7.1), except with variables starting with x_0 rather than x_1)

$$\begin{pmatrix} \frac{\partial f_1}{\partial x_0}(p) & \cdots & \frac{\partial f_r}{\partial x_0}(p) \\ \vdots & \ddots & \vdots \\ \frac{\partial f_1}{\partial x_n}(p) & \cdots & \frac{\partial f_r}{\partial x_n}(p) \end{pmatrix}$$

has corank equal to $\dim X$ at all closed points of X — this is precisely the Jacobian criterion for regularity (Exercise 13.2.D). We then require that the Jacobian matrix

with a new column

$$\begin{pmatrix} a_0 \\ \vdots \\ a_n \end{pmatrix}$$

appended has corank $\geq \dim X$ (hence $= \dim X$). This is cut out by equations (the determinants of certain minors). By the Jacobian description of the Zariski tangent space, this condition encodes the requirement that the Zariski tangent space of $H \cap X$ at p has dimension precisely $\dim X$, which is $\dim H \cap X + 1$ (i.e., $H \cap X$ is singular at p) if H does not contain X , or if H contains X . This is precisely the notion that we wished to capture. (Remark 13.4.4 works through an example, which may help clarify how this works.)

We next show that $\dim Z \leq n - 1$. For each closed point $p \in X$, let W_p be the locus of hyperplanes containing p , such that $H \cap X$ is singular at p , or else contains all of X ; what is the dimension of W_p ? Suppose $\dim X = d$. Then the restrictions on the hyperplanes in the definition of W_p correspond to $d + 1$ linear conditions. (Do you see why?) This means that W_p is a codimension $d + 1$, or dimension $n - d - 1$, projective space. Thus the fiber of $\pi: Z \rightarrow X$ over each closed point has pure dimension $n - d - 1$. By Key Exercise 12.4.A, this implies that $\dim Z \leq n - 1$. (If you wish, you can use Exercise 12.4.C to show that $\dim Z = n - 1$, and you can later show that Z is a projective bundle over X , once you know what a projective bundle is (Definition 18.2.3). But we don't need this for the proof.) \square

13.4.4. Remark. Here is an example that may help convince you that the algebra of paragraph 13.4.3 is describing the geometry we desire. Consider the plane conic $x_0^2 - x_1^2 - x_2^2 = 0$ over a field of characteristic not 2, which you might picture as the circle $x^2 + y^2 = 1$ from the real picture in the chart U_0 . Consider the point $[1, 1, 0]$, corresponding to $(1, 0)$ on the circle. We expect the tangent line in the affine plane to be $x = 1$, which corresponds to $x_0 - x_1 = 0$. Let's see what the algebra gives us. The Jacobian matrix (13.1.7.1) is

$$\begin{pmatrix} 2x_0 \\ -2x_1 \\ -2x_2 \end{pmatrix} = \begin{pmatrix} 2 \\ -2 \\ 0 \end{pmatrix},$$

which indeed has rank 1 as expected. Our recipe asks that the matrix

$$\begin{pmatrix} 2 & a_0 \\ -2 & a_1 \\ 0 & a_2 \end{pmatrix}$$

have rank 1 (i.e., $a_0 = -a_1$ and $a_2 = 0$), and also that $a_0x_0 + a_1x_1 + a_2x_2 = 0$, which (you should check) is precisely what we wanted.

13.4.A. EXERCISE. Reword the proof of Bertini's Theorem so as to remove the $k = \bar{k}$ hypothesis.

13.4.B. EASY EXERCISE.

- (a) Prove Bertini's Theorem with the weaker hypothesis that X has finitely many singular points.
- (b) Prove Bertini's Theorem with the weaker hypothesis that $X \rightarrow \mathbb{P}_k^n$ is a locally closed embedding.

13.4.C. EXERCISE. Continue to assume $k = \bar{k}$. Show that if X is a projective variety of dimension n in \mathbb{P}^m , then the intersection of X with n general hyperplanes consists of a finite number of reduced points. More precisely: if $\mathbb{P}^{m\vee}$ is the dual projective space, then there is a nonempty Zariski-open subset $U \subset (\mathbb{P}^{m\vee})^n$ such that for each closed point (H_1, \dots, H_n) of U , the scheme-theoretic intersection $H_1 \cap \dots \cap H_n \cap X$ consists of a finite number of reduced points. (The number of such points, counted correctly, is called the *degree* of the variety, see Exercise 19.6.N.)

13.4.5. Dual varieties. (We continue to assume $k = \bar{k}$ for convenience, although this can be relaxed.) This gives us an excuse to mention a classical construction. The image of Z (see (13.4.2.1)) in $\mathbb{P}^{n\vee}$ is called the **dual variety** of X . As $\dim Z = n - 1$, we “expect” the dual of X to be a hypersurface of $\mathbb{P}^{n\vee}$. It is a nonobvious fact that this in fact is a duality: the dual of the dual of X is X itself (see [H2, Thm. 15.24]). The following exercise will give you some sense of the dual variety.

13.4.D. EXERCISE. Show that the dual of a hyperplane in \mathbb{P}^n is the corresponding point of the dual space $\mathbb{P}^{n\vee}$. In this way, the duality between \mathbb{P}^n and $\mathbb{P}^{n\vee}$ is a special case of duality between projective varieties.

13.4.E. EXERCISE. Suppose $C \subset \mathbb{P}^2$ is a smooth conic over an algebraically closed field of characteristic not 2. Show that the dual variety to C is also a smooth conic. Thus for example, through a general point in the plane (if $k = \bar{k}$), there are two tangents to C . (The points on a line in the dual plane corresponds to those lines through a point of the original plane.)

13.4.F. ★ EXERCISE (THERE IS ONE SMOOTH CONIC TANGENT TO FIVE GENERAL LINES, AND GENERALIZATIONS). Continuing the notation of the previous problem, show that the number of smooth conics C containing i generally chosen points and tangent to $5 - i$ generally chosen lines is 1, 2, 4, 4, 2, 1 respectively for $i = 0, 1, 2, 3, 4, 5$. You might interpret the symmetry of the sequence in terms of the duality between the conic and the dual conic. This fact was likely known in the paleolithic era.

13.5 Another (co)dimension one miracle: Discrete valuation rings

The case of (co)dimension 1 is important, because if you understand how prime ideals behave that are separated by dimension 1, then you can use induction to prove facts in arbitrary dimension. This is one reason why Krull’s Principal Ideal Theorem 12.3.3 is so useful.

A dimension 1 regular local ring can be thought of as a “germ of a smooth curve” (see Figure 13.5). Two examples to keep in mind are $k[x]_{(x)} = \{f(x)/g(x) : x \nmid g(x)\}$ and $\mathbb{Z}_{(5)} = \{a/b : 5 \nmid b\}$. The first example is “geometric” and the second is “arithmetic”, but hopefully it is clear that they have something fundamental in common.

The purpose of this section is to give a long series of equivalent definitions of these rings. Before beginning, we quickly sketch these seven definitions. There are a number of ways a Noetherian local ring can be “nice”. It can be regular, or a



FIGURE 13.5. A germ of a curve

principal ideal domain, or a unique factorization domain, or normal. In dimension 1, these are the same. Also equivalent are nice properties of ideals: if \mathfrak{m} is principal; or if *all* ideals are either powers of the maximal ideal, or 0. Finally, the ring can have a *discrete valuation*, a measure of “size” of elements that behaves particularly well.

13.5.1. Theorem. — *Suppose (A, \mathfrak{m}) is a Noetherian local ring of dimension 1. Then the following are equivalent.*

- (a) (A, \mathfrak{m}) is regular.
- (b) \mathfrak{m} is principal.

Proof. Here is why (a) implies (b). If A is regular, then $\mathfrak{m}/\mathfrak{m}^2$ is one-dimensional. Choose any element $t \in \mathfrak{m} - \mathfrak{m}^2$. Then t generates $\mathfrak{m}/\mathfrak{m}^2$, so generates \mathfrak{m} by Nakayama’s Lemma 8.2.H (\mathfrak{m} is finitely generated by the Noetherian hypothesis). We call such an element a **uniformizer**.

Conversely, if \mathfrak{m} is generated by one element t over A , then $\mathfrak{m}/\mathfrak{m}^2$ is generated by one element t over $A/\mathfrak{m} = k$. Since $\dim_k \mathfrak{m}/\mathfrak{m}^2 \geq 1$ by Theorem 13.2.1, we have $\dim_k \mathfrak{m}/\mathfrak{m}^2 = 1$, so (A, \mathfrak{m}) is regular. \square

We will soon use a useful fact, which is geometrically motivated, and is a special case of an important result, the Artin-Rees Lemma 13.9.3. We will prove it in §13.9.

13.5.2. Proposition. — *If (A, \mathfrak{m}) is a Noetherian local ring, then $\cap_i \mathfrak{m}^i = 0$.*

13.5.3. The geometric intuition for this is that any function that is analytically zero at a point (vanishes to all orders) actually vanishes in an open neighborhood of that point. (Exercise 13.9.B will make this precise.) The geometric intuition also suggests an example showing that Noetherianness is necessary: consider e^{-1/x^2} in the ring of germs of C^∞ -functions on \mathbb{R} at the origin. (In particular, this implies that the ring of C^∞ -functions on \mathbb{R} , localized at the origin, is not Noetherian!)

It is tempting to argue that

$$(13.5.3.1) \quad \mathfrak{m}(\cap_i \mathfrak{m}^i) = \cap_i \mathfrak{m}^i,$$

and then to use Nakayama’s Lemma 8.2.9 to argue that $\cap_i \mathfrak{m}^i = 0$. Unfortunately, it is not obvious that this first equality is true: product does not commute with infinite descending intersections in general. (Aside: product also doesn’t commute with finite intersections in general, as for example $k[x, y, z]/(xz - yz)$, $(z)((x) \cap (y)) \neq (xz) \cap (yz)$.) We will establish Proposition 13.5.2 in Exercise 13.9.A(b). (We could do it directly right now without too much effort.)

13.5.4. Theorem. — Suppose (A, \mathfrak{m}) is a Noetherian local ring of dimension 1. Then (a) and (b) are equivalent to:

(c) A is an integral domain, and all ideals are of the form \mathfrak{m}^n (for $n \geq 0$) or (0) .

Proof. Assume (a): suppose (A, \mathfrak{m}, k) is a regular local ring of dimension 1. Then I claim that $\mathfrak{m}^n \neq \mathfrak{m}^{n+1}$ for any n . Otherwise, by Nakayama's Lemma, $\mathfrak{m}^n = 0$, from which $\mathfrak{t}^n = 0$. But A is an integral domain (by Theorem 13.2.13), so $\mathfrak{t} = 0$, from which $A = A/\mathfrak{m}$ is a field, which doesn't have dimension 1, contradiction.

I next claim that $\mathfrak{m}^n/\mathfrak{m}^{n+1}$ is dimension 1. Reason: $\mathfrak{m}^n = (\mathfrak{t}^n)$. So \mathfrak{m}^n is generated as an A -module by one element, and $\mathfrak{m}^n/(\mathfrak{m}\mathfrak{m}^n)$ is generated as a $(A/\mathfrak{m} = k)$ -module by 1 element (nonzero by the previous paragraph), so it is a one-dimensional vector space.

So we have a chain of ideals $A \supset \mathfrak{m} \supset \mathfrak{m}^2 \supset \mathfrak{m}^3 \supset \cdots$ with $\bigcap \mathfrak{m}^i = (0)$ (Proposition 13.5.2). We want to say that there is no room for any ideal besides these, because each pair is “separated by dimension 1”, and there is “no room at the end”. Proof: suppose $I \subset A$ is an ideal. If $I \neq (0)$, then there is some n such that $I \subset \mathfrak{m}^n$ but $I \not\subset \mathfrak{m}^{n+1}$. Choose some $u \in I - \mathfrak{m}^{n+1}$. Then $(u) \subset I$. But u generates $\mathfrak{m}^n/\mathfrak{m}^{n+1}$, hence by Nakayama it generates \mathfrak{m}^n , so we have $\mathfrak{m}^n \subset I \subset \mathfrak{m}^n$, so we are done: (c) holds.

We now show that (c) implies (a). Assume (a) does not hold: suppose we have a dimension 1 Noetherian local integral domain that is not regular, so $\mathfrak{m}/\mathfrak{m}^2$ has dimension at least 2. Choose any $u \in \mathfrak{m} - \mathfrak{m}^2$. Then (u, \mathfrak{m}^2) is an ideal, but $\mathfrak{m}^2 \subsetneq (u, \mathfrak{m}^2) \subsetneq \mathfrak{m}$. \square

13.5.A. EASY EXERCISE. Suppose (A, \mathfrak{m}) is a Noetherian dimension 1 local ring. Show that (a)–(c) above are equivalent to:

(d) A is a principal ideal domain.

13.5.5. Discrete valuation rings. We next define the notion of a discrete valuation ring. Suppose K is a field. A **discrete valuation** on K is a surjective homomorphism $v: K^\times \rightarrow \mathbb{Z}$ (in particular, $v(xy) = v(x) + v(y)$) satisfying

$$v(x + y) \geq \min(v(x), v(y))$$

except if $x + y = 0$ (in which case the left side is undefined). (Such a valuation is called *non-archimedean*, although we will not use that term.) It is often convenient to say $v(0) = \infty$. More generally, a **valuation** is a surjective homomorphism $v: K^\times \rightarrow G$ to a totally ordered abelian group G , although this isn't so important to us.

Here are three key examples.

- (i) (the 5-adic valuation) $K = \mathbb{Q}$, $v(r)$ is the “power of 5 appearing in r ”, e.g., $v(35/2) = 1$, $v(27/125) = -3$.
- (ii) $K = k(x)$, $v(f)$ is the “power of x appearing in f .”
- (iii) $K = k(x)$, $v(f)$ is the negative of the degree. This is really the same as (ii), with x replaced by $1/x$.

Then $0 \cup \{x \in K^\times : v(x) \geq 0\}$ is a ring, which we denote \mathcal{O}_v . It is called the **valuation ring** of v . (Not every valuation is discrete. Consider the ring of *Puiseux series* over a field k , $K = \bigcup_{n \geq 1} k((x^{1/n}))$, with $v: K^\times \rightarrow \mathbb{Q}$ given by $v(x^q) = q$.)

13.5.B. EXERCISE. Describe the valuation rings in the three examples (i)–(iii) above. (You will notice that they are familiar-looking dimension 1 Noetherian local rings. What a coincidence!)

13.5.C. EXERCISE. Show that $\{0\} \cup \{x \in K^\times : v(x) > 0\}$ is the unique maximal ideal of the valuation ring. (Hint: show that everything in the complement is invertible.) Thus the valuation ring is a local ring.

An integral domain A is called a **discrete valuation ring** (or **DVR**) if there exists a discrete valuation v on its fraction field $K = K(A)$ for which $\mathcal{O}_v = A$. Similarly, A is a **valuation ring** if there exists a valuation v on K for which $\mathcal{O}_v = A$.

Now if A is a regular local ring of dimension 1, and t is a uniformizer (a generator of \mathfrak{m} as an ideal, or equivalently of $\mathfrak{m}/\mathfrak{m}^2$ as a k -vector space) then any nonzero element r of A lies in some $\mathfrak{m}^n - \mathfrak{m}^{n+1}$, so $r = t^n u$ where u is invertible (as t^n generates \mathfrak{m}^n by Nakayama, and so does r), so $K(A) = A_t = A[1/t]$. So any element of $K(A)^\times$ can be written uniquely as ut^n where u is invertible and $n \in \mathbb{Z}$. Thus we can define a valuation v by $v(ut^n) = n$.

13.5.D. EXERCISE. Show that v is a discrete valuation.

13.5.E. EXERCISE. Conversely, suppose (A, \mathfrak{m}) is a discrete valuation ring. Show that (A, \mathfrak{m}) is a regular local ring of dimension 1. (Hint: Show that the ideals are all of the form (0) or $I_n = \{r \in A : v(r) \geq n\}$, and (0) and I_1 are the only prime ideals. Thus we have Noetherianness, and dimension 1. Show that I_1/I_2 is generated by the image of any element of $I_1 - I_2$.)

Hence we have proved:

13.5.6. Theorem. — *An integral domain A is a Noetherian local ring of dimension 1 satisfying (a)–(d) if and only if*

(e) *A is a discrete valuation ring.*

13.5.F. EXERCISE. Show that there is only one discrete valuation on a discrete valuation ring.

13.5.7. Definition. Thus any regular local ring of dimension 1 comes with a unique valuation on its fraction field. If the valuation of an element is $n > 0$, we say that the element has a **zero of order** n . If the valuation is $-n < 0$, we say that the element has a **pole of order** n . We will come back to this shortly, after dealing with (f) and (g).

13.5.8. Theorem. — *Suppose (A, \mathfrak{m}) is a Noetherian local ring of dimension 1. Then (a)–(e) are equivalent to:*

(f) *A is a unique factorization domain, and*

(g) *A is an integral domain, integrally closed in its fraction field $K = K(A)$.*

Proof. (a)–(e) clearly imply (f), because we have the following stupid unique factorization: each nonzero element of A can be written uniquely as ut^n where $n \in \mathbb{Z}^{\geq 0}$ and u is invertible.

Now (f) implies (g), because unique factorization domains are integrally closed in their fraction fields (Exercise 5.4.F).

It remains to check that (g) implies (a)–(e). We will show that (g) implies (b).

Suppose (A, \mathfrak{m}) is a Noetherian local integral domain of dimension 1, integrally closed in its fraction field $K = K(A)$. Choose any nonzero $r \in \mathfrak{m}$. Then $S = A/(r)$ is a Noetherian local ring of dimension 0 — its only prime is the image of \mathfrak{m} , which we denote \mathfrak{n} to avoid confusion. Then \mathfrak{n} is finitely generated, and each generator is nilpotent (the intersection of all the prime ideals in any ring are the nilpotents, Theorem 3.2.13). Then $\mathfrak{n}^N = 0$, where N is sufficiently large. Hence there is some n such that $\mathfrak{n}^n = 0$ but $\mathfrak{n}^{n-1} \neq 0$.

Now comes the crux of the argument. Thus in A , $\mathfrak{m}^n \subset (r)$ but $\mathfrak{m}^{n-1} \not\subset (r)$. Choose $s \in \mathfrak{m}^{n-1} - (r)$. Consider $s/r \in K(A)$. As $s \notin (r)$, $s/r \notin A$, so as A is integrally closed, s/r is not integral over A .

Now (inside $K(A)$) $\frac{s}{r}\mathfrak{m} \not\subset \mathfrak{m}$ (or else $\frac{s}{r}\mathfrak{m} \subset \mathfrak{m}$ would imply that \mathfrak{m} is a faithful $A[\frac{s}{r}]$ -module, contradicting Exercise 8.2.I). But $s\mathfrak{m} \subset \mathfrak{m}^n \subset rA$, so $\frac{s}{r}\mathfrak{m} \subset A$. Thus $\frac{s}{r}\mathfrak{m} = A$, from which $\mathfrak{m} = \frac{r}{s}A$, so \mathfrak{m} is principal. \square

13.5.9. Geometry of normal Noetherian schemes. We can finally make precise (and generalize) the fact that the function $(x-2)^2x/(x-3)^4$ on $\mathbb{A}_{\mathbb{C}}^1$ has a double zero at $x=2$ and a quadruple pole at $x=3$. Furthermore, we can say that $75/34$ has a double zero at 5, and a single pole at 2. (What are the zeros and poles of $x^3(x+y)/(x^2+xy)^3$ on \mathbb{A}^2 ?) Suppose X is a locally Noetherian scheme. Then for any regular codimension 1 point p (i.e., any point p where $\mathcal{O}_{X,p}$ is a regular local ring of dimension 1), we have a discrete valuation v . If f is any nonzero element of the fraction field of $\mathcal{O}_{X,p}$ (e.g., if X is integral, and f is a nonzero element of the function field of X), then if $v(f) > 0$, we say that the element has a **zero of order** $v(f)$ **at** p , and if $v(f) < 0$, we say that the element has a **pole of order** $-v(f)$ **at** p . (We are not yet allowed to discuss order of vanishing at a point that is not regular and codimension 1. One can make a definition, but it doesn't behave as well as it does when you have a discrete valuation.)

13.5.G. EXERCISE (FINITENESS OF ZEROS AND POLES ON NOETHERIAN SCHEMES). Suppose X is an integral Noetherian scheme, and $f \in K(X)^\times$ is a nonzero element of its function field. Show that f has a finite number of zeros and poles. (Hint: reduce to $X = \text{Spec } A$. If $f = f_1/f_2$, where $f_i \in A$, prove the result for f_i .)

Suppose A is a Noetherian integrally closed domain. Then it is **regular in codimension 1** (translation: its points of codimension at most 1 are regular). If A is dimension 1, then obviously A is regular.

13.5.H. EXERCISE. If f is a nonzero rational function on a locally Noetherian normal scheme, and f has no poles, show that f is regular. (Hint: Algebraic Hartogs's Lemma 12.3.15.)

13.5.10. For example (cf. Exercise 13.3.M), $\text{Spec } \mathbb{Z}[i]$ is regular, because it is dimension 1, and $\mathbb{Z}[i]$ is a unique factorization domain. Hence $\mathbb{Z}[i]$ is normal, so all its closed (codimension 1) points are regular. Its generic point is also regular, as $\mathbb{Z}[i]$ is an integral domain.

13.5.11. Remark. A (Noetherian) scheme can be singular in codimension 2 and still be normal. For example, you have shown that the cone $x^2 + y^2 = z^2$ in \mathbb{A}^3 in

characteristic not 2 is normal (Exercise 5.4.I(b)), but it is singular at the origin (the Zariski tangent space is visibly three-dimensional).

But singularities of normal schemes are not so bad in some ways: we have Algebraic Hartogs's Lemma 12.3.15 for Noetherian normal schemes, which states that you can extend functions over codimension 2 sets.

13.5.12. Remark. We know that for Noetherian rings we have implications

unique factorization domain \implies integrally closed \implies regular in codimension 1.

Hence for locally Noetherian schemes, we have similar implications:

factorial \implies normal \implies regular in codimension 1.

Here are two examples to show you that these inclusions are strict.

13.5.I. EXERCISE (THE PINCHED PLANE). Let A be the subring $k[x^3, x^2, xy, y]$ of $k[x, y]$. (Informally, we allow all polynomials that don't include a nonzero multiple of the monomial x .) Show that $\text{Spec } k[x, y] \rightarrow \text{Spec } A$ is a normalization. Show that A is not integrally closed. Show that $\text{Spec } A$ is regular in codimension 1. (Hint for the last part: show it is dimension 2, and when you throw out the origin you get something regular, by inverting x^2 and y respectively, and considering A_{x^2} and A_y .)

13.5.13. Example (the quadric surface once again). Suppose k is algebraically closed of characteristic not 2. Then $k[w, x, y, z]/(wz - xy)$ is integrally closed, but not a unique factorization domain, see Exercise 5.4.L (and Exercise 13.1.E).

13.5.14. Important remark: Finitely generated modules over a principal ideal domain (e.g., a discrete valuation ring). We record a useful fact for future reference. Recall that finitely generated modules over a principal ideal domain are finite direct sums of cyclic modules (see for example [DF, §12.1, Thm. 5]). Hence any finitely generated module over a discrete valuation ring A with uniformizer t is a finite direct sum of terms A and $A/(t^r)$ (for various r). See Proposition 14.3.2 for an immediate consequence.

13.5.15. Aside: Dedekind domains. A **Dedekind domain** is a Noetherian integral domain of dimension at most one that is normal (integrally closed in its fraction field). The localization of a Dedekind domain at any prime but (0) (i.e., a codimension one prime) is hence a discrete valuation ring. This is an important notion, but we won't use it much. Rings of integers of number fields are examples, see §10.7.1. In particular, if n is a square-free integer congruent to 3 (mod 4), then $\mathbb{Z}[\sqrt{n}]$ is a Dedekind domain, by Exercise 5.4.I(a). If you wish you can prove *unique factorization of ideals in a Dedekind domain*: any nonzero ideal in a Dedekind domain can be uniquely factored into prime ideals.

13.6 Smooth (and étale) morphisms: First definition

In §13.2.5, we defined smoothness over a field. It is an imperfect definition (see §13.2.7), and we will improve it when we know more (in §22.3.1).

We know enough to define *smooth morphisms* (and, as a special case, *étale morphisms*) in general. Our definition will be imperfect in a number of ways. For example, it will look a little surprising. As another sign, it will not be obvious that it actually generalizes smoothness over a field. As a third sign, like Definition 13.2.6, because it is of the form “there exist open covers satisfying some property”, it is poorly designed to be used to show that some morphism is *not* smooth. But it has a major advantage that we can give the definition right now, and the basic properties of smooth morphisms will be straightforward to show.

13.6.1. Differential geometric motivation. The notion of a smooth morphism is motivated by the following idea from differential geometry. For the purposes of this discussion, we say a map $\pi: X \rightarrow Y$ of manifolds (or real analytic spaces, or any variant thereof) is *smooth of relative dimension n* if locally (on X) it looks like $Y' \times \mathbb{R}^n \rightarrow Y'$. Translation: it is “locally on the source a smooth fibration”. (If you have not heard the word fibration before, don’t worry. Informally speaking, a *fibration* in some “category of geometric objects” is a map $\pi: X \rightarrow Y$ that locally on the source is a product $V \times W \rightarrow V$; in some sort of category of manifolds, a fibration is a “locally on the source a smooth fibration” if the W in the local description can be taken to be \mathbb{R}^a for some a .)

In particular, “smooth of relative dimension 0” is the same as the important notion of “local isomorphism” (“isomorphism locally on the source” — not quite the same as “covering space”). Carrying this idea too naively into algebraic geometry leads to the “wrong” definition (see Exercise 13.6.E). An explanation of how this differential geometric intuition leads to a good definition of smoothness will be given in §26.1, when we will be ready to *give* a good definition of smoothness. For now we will settle for a *satisfactory* (and correct) definition.

13.6.2. Definition. A morphism $\pi: X \rightarrow Y$ is **smooth of relative dimension n** if there exist open covers $\{U_i\}$ of X and $\{V_i\}$ of Y , with $\pi(U_i) \subset V_i$, such that for every i we have a commutative diagram

$$\begin{array}{ccc} U_i & \xrightarrow{\sim} & W \\ \pi|_{U_i} \downarrow & & \downarrow \rho|_W \\ V_i & \xrightarrow{\sim} & \text{Spec } B \end{array}$$

where there is a morphism $\rho: \text{Spec } B[x_1, \dots, x_{n+r}]/(f_1, \dots, f_r) \rightarrow \text{Spec } B$ (induced by the obvious map of rings in the opposite direction), and W is an open subscheme of $\text{Spec } B[x_1, \dots, x_{n+r}]/(f_1, \dots, f_r)$, such that the determinant of the Jacobian matrix of the f_i ’s with respect to the *first* r x_i ’s,

$$(13.6.2.1) \quad \det \left(\frac{\partial f_j}{\partial x_i} \right)_{i,j \leq r},$$

is an invertible (= nowhere zero) function on W .

Étale means *smooth of relative dimension 0*.

13.6.3. Quick observations. From the definition, smooth morphisms are locally of finite presentation (hence locally of finite type). Also, $\mathbb{A}_B^n \rightarrow \text{Spec } B$ is immediately seen to be smooth of relative dimension n . From the definition, the locus on X

where π is smooth of relative dimension n is open. In particular, the locus where π is étale is open.

As easy examples, open embeddings are étale, and the projection $\mathbb{A}^n \times Y \rightarrow Y$ is smooth of relative dimension n . You may already have some sense of what makes this definition imperfect. For example, how do you know a morphism is smooth if it isn't given to you in a form where the "right" variables x_i and "right" equations are clear?

13.6.A. MOTIVATING EXERCISE (FOR THOSE WITH DIFFERENTIAL-GEOMETRIC BACKGROUND). Show how Definition 13.6.2 gives the definition in differential geometry, as described in §13.6.1. (Your argument will use the implicit function theorem, which "doesn't work" in algebraic geometry, see Exercise 13.6.E.)

13.6.B. EXERCISE. Show that the notion of smoothness of relative dimension n is local on both the source and target.

We can thus make sense of the phrase " $\pi: X \rightarrow Y$ is smooth of relative dimension n at $p \in X$ ": it means that there is an open neighborhood U of p such that $\pi|_U$ is smooth of relative dimension n .

The phrase **smooth morphism** (without reference to relative dimension n) often informally means "smooth morphism of some relative dimension n ", but sometimes can mean "smooth of some relative dimension in the neighborhood of every point".

13.6.C. EASY EXERCISE. Show that the notion of smoothness of relative dimension n is preserved by base change.

13.6.D. EXERCISE. Suppose $\pi: X \rightarrow Y$ is smooth of relative dimension m and $\rho: Y \rightarrow Z$ is smooth of relative dimension n . Show that $\rho \circ \pi$ is smooth of relative dimension $m + n$. (In particular, the composition of étale morphisms is étale. Furthermore, this implies that smooth morphisms are closed under product, by Exercise 8.1.A.)

Exercises 13.6.B–13.6.D imply that smooth morphisms (and étale morphisms) form a "reasonable" class in the sense of §8.1.

13.6.4. Observation. Suppose $\pi: X \rightarrow Y$ is smooth of relative dimension n . Then locally on X , π can be described as an étale cover of \mathbb{A}_Y^n . More precisely, for every $p \in X$, there is an open neighborhood U of p , such that $\pi|_U$ can be factored into

$$U \xrightarrow{\alpha} \mathbb{A}_Y^n \xrightarrow{\beta} Y$$

where α is étale and β is the obvious projection.

13.6.E. EXERCISE. Suppose k is a field of characteristic not 2. Let $Y = \operatorname{Spec} k[t]$, and $X = \operatorname{Spec} k[u, 1/u]$. Show that the morphism $\pi: X \rightarrow Y$ induced by $t \mapsto u^2$ is étale. (Sketch π !) Show that there is no nonempty open subset U of X on which π is an isomorphism. (This shows that étale cannot be defined as "local isomorphism", i.e., "isomorphism locally on the source". In particular, the naive extension of the differential intuition of §13.6.1 does not give the right definition of étaleness in algebraic geometry.)

We now show that our new Definition 13.6.2 specializes to Definition 13.2.6 when the target is a field. This is the hardest fact in this section. (It is also a consequence of Theorem 26.2.2, so you could skip the proof of Theorem 13.6.5 and take it on faith until Chapter 26. The crux of the proof of Theorem 26.2.2 is similar to the crux of the proof of Theorem 13.6.5.)

13.6.5. Theorem. — *Suppose X is a k -scheme. Then the following are equivalent.*

- (i) X is smooth of relative dimension n over $\operatorname{Spec} k$ (Definition 13.6.2).
- (ii) X has pure dimension n , and is smooth over k (in the sense of Definition 13.2.6).

This allows us to use the phrases “smooth over k ” and “smooth over $\operatorname{Spec} k$ ” interchangeably, as we would expect. More generally, “**smooth over a ring A** ” means “smooth over $\operatorname{Spec} A$ ”.

★ *Proof.* (i) implies (ii). Suppose X is smooth of relative dimension n over $\operatorname{Spec} k$ (Definition 13.6.2). We will show that X has pure dimension n , and is smooth over k (Definition 13.2.6). Our goal is a Zariski-local statement, so we may assume that X is an open subset

$$W \subset \operatorname{Spec} k[x_1, \dots, x_{n+r}]/(f_1, \dots, f_r)$$

as described in Definition 13.6.2.

The central issue is showing that X has “pure dimension n ”, as the latter half of the statement would then follow because the corank of the Jacobian matrix is then n .

Let $X_{\bar{k}} := X \times_k \bar{k}$. By Exercise 12.1.G, $X_{\bar{k}}$ is pure dimension n if and only if X has pure dimension n . Let p be a closed point of the (finite type scheme) $X_{\bar{k}}$. Then X has dimension at least n at p , by Krull’s Principal Ideal Theorem 12.3.3, combined with Theorem 12.2.9 (codimension is the difference of dimensions for varieties). But the Zariski tangent space at a closed point of a finite type scheme over \bar{k} is cut out by the Jacobian matrix (Exercise 13.1.G), and thus has dimension exactly n (by the corank hypothesis in the definition of k -smooth). As the dimension of the Zariski tangent space bounds the dimension of the scheme (Theorem 13.2.1), $X_{\bar{k}}$ has pure dimension n as desired.

(ii) implies (i). We will prove the result near any given closed point p of X . (Here we use that any union of open neighborhoods of all the closed points is all of X , by Exercise 5.3.F.) We thus may reduce to the following problem. Suppose we are given $X := \operatorname{Spec} k[x_1, \dots, x_{n+r}]/(f_1, \dots, f_r)$, and a point $p \in X$, and an open neighborhood U of p in X such that U is pure dimension n , and the Jacobian matrix of the f_j ’s with respect to the x_i ’s has corank n everywhere on U . We wish to show that the structure morphism $\pi: X \rightarrow \operatorname{Spec} k$ is smooth of relative dimension n at p (in the sense of Definition 13.6.2). We may as well replace X by U . At p , there exists an $r \times r$ submatrix of the Jacobian matrix that is invertible (by the corank condition). Rearrange the x_i ’s and f_j ’s so that the Jacobian of the first r of the f_j ’s with respect to the first r of the x_i ’s has nonzero determinant. Let $I = (f_1, \dots, f_r)$. We will show that $I = (f_1, \dots, f_r)$ “near p ”, or more precisely, that there is some $g \in k[x_1, \dots, x_{n+r}]$ with $g(p) \neq 0$ such that $I_g = (f_1, \dots, f_r)_g$ (so $I = (f_1, \dots, f_r)$ “in $D(g)$ ”).

Choose a point $\bar{p} \in X_{\bar{k}}$ above p , i.e., such that if ψ is the morphism $X_{\bar{k}} \rightarrow X$ induced by $\operatorname{Spec} \bar{k} \rightarrow \operatorname{Spec} k$, then $\psi(\bar{p}) = p$. (Why does such a \bar{p} exist?) Now $X_{\bar{k}}$ is \bar{k} -smooth (in the sense of Definition 13.2.6, by Exercise 13.2.G) at \bar{p} , hence regular

(Exercise 13.2.I). Its Zariski tangent space at \bar{p} is cut out in the tangent space of \mathbb{A}_k^{n+r} by the Jacobian conditions (Exercise 13.1.G), and hence cut out by f_1, \dots, f_r .

13.6.F. EXERCISE. Suppose \bar{m} is the ideal of $\bar{k}[x_1, \dots, x_{n+r}]$ corresponding to \bar{p} . Show that f_1, \dots, f_r generate the ideal (f_1, \dots, f_N) in $\bar{k}[x_1, \dots, x_{n+r}]_{\bar{m}}$. Hint: look at your argument for Exercise 13.2.L(a).

For notational compactness, define $A := k[x_1, \dots, x_{n+r}]/(f_1, \dots, f_r)$, and $\bar{A} := \bar{k}[x_1, \dots, x_{n+r}]/(f_1, \dots, f_r) = A \otimes_k \bar{k}$. Let I (resp. \bar{I}) be the ideal of A (resp. \bar{A}) generated by f_{r+1}, \dots, f_N (so $\bar{I} = I \otimes_k \bar{k}$). We wish to show that $I = 0$ “near p ”. Here $p = [m]$, where m is a maximal ideal (recall p is a closed point).

Exercise 13.6.F means that $\bar{I}_{\bar{m}} = 0$, so taking the quotient by the maximal ideal \bar{m} of \bar{A} yields

$$I \otimes_A (\bar{A}/\bar{m}) = 0$$

which implies

$$(I \otimes_A (A/m)) \otimes_{A/m} (\bar{A}/\bar{m}) = 0.$$

But $\bar{k} = \bar{A}/\bar{m}$ is a field extension of $\ell := A/m$, and an ℓ -vector space V is 0 if (and only if) $V \otimes_{\ell} \bar{k} = 0$, so $I \otimes_A (A/m) = 0$, i.e., $I_m/mI_m = 0$.

13.6.G. EXERCISE. Conclude that there is some $g \in A \setminus m$ such that $gI = 0$. Hint: Nakayama, and the finite generation of I . (Be sure you understand why the proof is now complete.)

□

13.7 ★ Valuative criteria for separatedness and properness

(The only reason this section is placed here is that we need the theory of discrete valuation rings.)

In reasonable circumstances, it is possible to verify separatedness by checking only maps from spectra of valuation rings. There are four reasons you might like this (even if you never use it). First, it gives useful intuition for what separated morphisms look like. Second, given that we understand schemes by maps to them (the Yoneda philosophy), we might expect to understand morphisms by mapping certain maps of schemes to them, and this is how you can interpret the diagram appearing in the valuative criterion. And the third concrete reason is that one of the two directions in the statement is much easier (a special case of the Reduced-to-Separated Theorem 11.4.2, see Exercise 13.7.A), and this is the direction we will repeatedly use. Finally, the criterion is very useful!

Similarly, there is a valuative criterion for properness.

In this section, we will meet the valuative criteria, but aside from outlining the proof of one result (the DVR version of the valuative criterion of separatedness), we will not give proofs, and satisfy ourselves with references. There are two reasons for this controversial decision. First, the proofs require the development of some commutative algebra involving valuation rings that we will not otherwise need. Second, we will not use these results in any essential way later in this book.

We begin with a valuative criterion for separatedness that applies in a case that will suffice for the interests of most people, that of finite type morphisms of Noetherian schemes. We will then give a more general version for more general readers.

13.7.1. Theorem (valuative criterion for separatedness, DVR version; see Figure 13.6). — Suppose $\pi: X \rightarrow Y$ is a morphism of finite type of locally Noetherian schemes. Then π is separated if and only if the following condition holds: for any discrete valuation ring A , and any diagram of the form

$$(13.7.1.1) \quad \begin{array}{ccc} \mathrm{Spec} K(A) & \longrightarrow & X \\ \text{open emb.} \downarrow & & \downarrow \pi \\ \mathrm{Spec} A & \longrightarrow & Y \end{array}$$

(where the vertical morphism on the left corresponds to the inclusion $A \hookrightarrow K(A)$), there is at most one morphism $\mathrm{Spec} A \rightarrow X$ such that the diagram

$$(13.7.1.2) \quad \begin{array}{ccc} \mathrm{Spec} K(A) & \longrightarrow & X \\ \text{open emb.} \downarrow & \nearrow \leq 1 & \downarrow \pi \\ \mathrm{Spec} A & \longrightarrow & Y \end{array}$$

commutes.




FIGURE 13.6. The valuative criterion for separatedness

The idea behind the proof is explained in §13.7.3. We can show one direction right away, in the next exercise.

13.7.A. EXERCISE (THE EASY DIRECTION). Use the Reduced-to-Separated Theorem 11.4.2 to prove one direction of the theorem: that if π is separated, then the valuative criterion holds.

13.7.B. EXERCISE. Suppose X is an integral Noetherian separated curve. If $p \in X$ is a regular closed point, then $\mathcal{O}_{X,p}$ is a discrete valuation ring, so each regular point yields a discrete valuation on $K(X)$. Use the previous exercise to show that distinct points yield distinct discrete valuations.

Here is the intuition behind the valuative criterion (see Figure 13.6). We think of Spec of a discrete valuation ring A as a “germ of a curve”, and $\mathrm{Spec} K(A)$ as

the “germ minus the origin” (even though it is just a point!). Then the valuative criterion says that if we have a map from a germ of a curve to Y , and have a lift of the map away from the origin to X , then there is at most one way to lift the map from the entire germ. In the case where Y is the spectrum of a field, you can think of this as saying that limits of one-parameter families are unique (if they exist).

For example, this captures the idea of what is wrong with the map of the line with the doubled origin over k (Figure 13.7): we take $\text{Spec } A$ to be the germ of the affine line at the origin, and consider the map of the germ minus the origin to the line with doubled origin. Then we have two choices for how the map can extend over the origin.



FIGURE 13.7. The line with the doubled origin fails the valuative criterion for separatedness. (You may notice Figure 11.4 embedded in this diagram.)

13.7.C. EXERCISE. Make this precise: show that map of the line with doubled origin over k to $\text{Spec } k$ fails the valuative criterion for separatedness. (Earlier arguments were given in Exercises 11.3.I and 11.4.C.)

13.7.2. ★★ Remark for experts: moduli spaces and the valuative criterion for separatedness. If $Y = \text{Spec } k$, and X is a (fine) moduli space (a term we won’t define here) of some type of object, then the question of the separatedness of X (over $\text{Spec } k$) has a natural interpretation: given a family of your objects parametrized by a “punctured discrete valuation ring”, is there always at most one way of extending it over the closed point?

13.7.3. Idea behind the proof of Theorem 13.7.1 (the valuative criterion for separatedness, DVR version). (One direction was done in Exercise 13.7.A.) If π is *not* separated, our goal is to produce a diagram (13.7.1.1) that can be completed to (13.7.1.2) in more than one way. If π is not separated, then $\delta: X \rightarrow X \times_Y X$ is a locally closed embedding that is not a closed embedding.

13.7.D. EXERCISE. Show that you can find points p not in the diagonal Δ of $X \times_Y X$ and q in Δ such that $p \in \overline{q}$, and there are no points “between p and q ” (no points r distinct from p and q with $p \in \overline{r}$ and $r \in \overline{q}$). (Exercise 8.4.C may shed some light.)

Let Q be the scheme obtained by giving the induced reduced subscheme structure to \overline{q} . Let $B = \mathcal{O}_{Q,p}$ be the local ring of Q at p .

13.7.E. EXERCISE. Show that B is a Noetherian local integral domain of dimension 1.

If B were regular, then we would be done: composing the inclusion morphism $Q \rightarrow X \times_Y X$ with the two projections induces the same morphism $q \rightarrow X$ (i.e., $\text{Spec } \kappa(q) \rightarrow X$) but different extensions to B precisely because p is not in the diagonal. To complete the proof, one shows that the normalization of B is Noetherian; then localizing at any prime above p (there is one by the Lying Over Theorem 8.2.5) yields the desired discrete valuation ring A .

For an actual proof, see [Stacks, tag 0207] or [Gr-EGA, II.72.3].

With a more powerful invocation of commutative algebra, we can prove a valuative criterion with much less restrictive hypotheses.

13.7.4. Theorem (valuative criterion for separatedness). — Suppose $\pi: X \rightarrow Y$ is a quasiseparated morphism. Then π is separated if and only if for any valuation ring A with function field K , and any diagram of the form (13.7.1.1), there is at most one morphism $\text{Spec } A \rightarrow X$ such that the diagram (13.7.1.2) commutes.

Because I have already failed to completely prove the DVR version, I feel no urge to prove this harder fact. The proof of one direction, that π separated implies that the criterion holds, follows from the identical argument as in Exercise 13.7.A. For a complete proof, see [Stacks, tags 01KY and 01KZ] or [Gr-EGA, II.7.2.3].

13.7.5. Valuative criteria for (universal closedness and) properness.

There is a valuative criterion for properness too. It is philosophically useful, and sometimes directly useful, although we won't need it. It naturally comes from the valuative criterion for separatedness combined with a valuative criterion for universal closedness.

13.7.6. Theorem (valuative criterion for universal closedness and properness, DVR version). — Suppose $\pi: X \rightarrow Y$ is a morphism of finite type of locally Noetherian schemes. Then π is universally closed (resp. proper) if and only if for any discrete valuation ring A and any diagram (13.7.1.1), there is at least one (resp. exactly one) morphism $\text{Spec } A \rightarrow X$ such that the diagram (13.7.1.2) commutes.

See [Gr-EGA, II.7.3.8], [Ha1, Thm. II.4.7], or [Gr-EGA, I.5.5] for proofs. A comparison with Theorem 13.7.1 will convince you that these three criteria naturally form a family.

In the case where Y is a field, you can think of the valuative criterion of properness as saying that limits of one-parameter families in proper varieties always exist, and are unique. This is a useful intuition for the notion of properness.

13.7.F. EASY EXERCISE. Use the valuative criterion for properness to prove that $\mathbb{P}_A^n \rightarrow \text{Spec } A$ is proper if A is Noetherian. (Don't be fooled: Because this requires the valuative criterion, this is a difficult way to prove a fact that we already showed in Theorem 11.5.5.)

13.7.G. EXERCISE (CF. EXERCISE 13.7.B). Suppose X is an irreducible regular (Noetherian) curve, proper over a field k (respectively, over \mathbb{Z}). Describe a bijection between the discrete valuations on $K(X)$ for which the elements of k (respectively \mathbb{Z}) have valuation 0, and the closed points of X .

13.7.7. Remarks for experts. There is a moduli-theoretic interpretation similar to that for separatedness (Remark 13.7.2): X is proper if and only if there is always precisely one way of filling in a family over a punctured spectrum of a discrete valuation ring.

13.7.8. Finally, here is a fancier version of the valuative criterion for universal closedness and properness.

13.7.9. Theorem (valuative criterion for universal closedness and properness). — Suppose $\pi: X \rightarrow Y$ is a quasiseparated, finite type (hence quasicompact) morphism. Then π is universally closed (resp. proper) if and only if the following condition holds. For any valuation ring A and any diagram of the form (13.7.1.1), there is at least one (resp. exactly one) morphism $\text{Spec } A \rightarrow X$ such that the diagram (13.7.1.2) commutes.

Clearly the valuative criterion for properness is a consequence of the valuative criterion for separatedness (Theorem 13.7.4) and the valuative criterion for universal closedness. For proofs, see [Stacks, tag 01KF] or [Gr-EGA, II.7.3.8].

13.7.10. On the importance of valuation rings in general. Although we have only discussed discrete valuation rings in depth, general valuation rings should not be thought of as an afterthought. Serre makes the case to Grothendieck in a letter, [GrS, p. 125]:

Tu es bien sévère pour les Valuations! Je persiste pourtant à les garder, pour plusieurs raisons: d'abord une pratique: n rédacteurs ont sué dessus, il n'y a rien à reprocher au fourbi, on ne doit pas le vider sans raisons très sérieuses (et tu n'en as pas). ... Même un noethérien impénitent a besoin des valuations discrètes, et de leurs extension; en fait, Tate, Dwork, et tous les gen p-adiques te diront qu'on ne peut se limiter au cas discret et que le cas de rang 1 est indispensable; à ce moment, les méthodes noethériennes deviennent un vrai carcan, et on comprend beaucoup mieux en rédigeant le cas général que uniquement le cas de rang 1. ... Il n'y a pas lieu d'en faire un plat, bien sûr, et c'est pourquoi j'avais vivement combattu le plan Weil initial qui en faisait le théorème central de l'Algèbre Commutative, mais d'autre part il faut le garder.

You are very harsh on Valuations! I persist nonetheless in keeping them, for several reasons, of which the first is practical: n people have sweated over them, there is nothing wrong with the result, and it should not be thrown out without very serious reasons (which you do not have). ... Even an unrepentant Noetherian needs discrete valuations and their extensions; in fact, Tate, Dwork and all the p-adic people will tell you that one cannot restrict oneself to the discrete case and the rank 1 case is indispensable; Noetherian methods then become a burden, and one understands much better if one considers the general case and not only the rank 1 case. ... It is not worth making a mountain out of it, of course, which is why I energetically fought Weil's original plan to make it the central theorem of Commutative Algebra, but on the other hand it must be kept.

13.8 ★ More sophisticated facts about regular local rings

Regular local rings have essentially every good property you could want, but some of them require hard work. We now discuss a few fancier facts that may help you sleep well at night.

13.8.1. Localizations of regular local rings are regular local rings.

13.8.2. Fact ([E, Cor. 19.14], [Mat2, Thm. 19.3]). — *If (A, \mathfrak{m}) is a regular local ring, then any localization of A at a prime is also a regular local ring.*

(We will not need this, and hence will not prove it.) This major theorem was an open problem in commutative algebra for a long time until settled by Serre and Auslander-Buchsbaum using homological methods.

Hence to check if $\text{Spec } A$ is regular (A Noetherian), it suffices to check at closed points (at maximal ideals). Assuming Fact 13.8.2 (and using Exercise 5.1.E), you can check regularity of a Noetherian scheme by checking at closed points.

We will prove two important cases of Fact 13.8.2. The first you can do right now.

13.8.A. EXERCISE. Suppose X is a Noetherian dimension 1 scheme that is regular at its closed points. Show that X is reduced. Hence show (without invoking Fact 13.8.2) that X is regular.

The second important case will be proved in §22.3.12:

13.8.3. Theorem. — *If X is a finite type scheme over a perfect field k that is regular at its closed points, then X is regular.*

More generally, Exercise 22.3.G will show that Fact 13.8.2 holds if A is the localization of a finite type algebra over a perfect field.

13.8.B. EXERCISE (GENERALIZING EXERCISE 13.3.E). Suppose k is an algebraically closed field of characteristic 0. Assuming Theorem 13.8.3, show that there exists a regular hypersurface of every positive degree d in \mathbb{P}^n . (As in Exercise 13.3.E, feel free to weaken the hypotheses.)

13.8.4. Regular local rings are unique factorization domains, integrally closed, and Cohen-Macaulay.

13.8.5. Fact (Auslander-Buchsbaum Theorem). — *Regular local rings are unique factorization domains.*

(This is a hard theorem, so we will not prove it, and will therefore not use it. For a proof, see [E, Thm. 19.19], [Mat2, Thm. 20.3] or [Gr-EGA, IV₄.21.11.1], or [Mu7, §III.7] in the special case of varieties.) Thus regular schemes are factorial, and hence normal by Exercise 5.4.F.

13.8.6. Remark: Factoriality is weaker than regularity. The implication “regular implies factorial” is strict. Here is an example showing this. Suppose k is an algebraically closed field of characteristic not 2. Let $A = k[x_1, \dots, x_n]/(x_1^2 + \dots + x_n^2)$ (cf. Exercise 5.4.I). Note that $\text{Spec } A$ is clearly singular at the origin. In Exercise 15.2.V, we will show that A is a unique factorization domain when $n \geq 5$, so $\text{Spec } A$ is factorial. In particular, $A_{(x_1, \dots, x_n)}$ is a Noetherian local ring that is a unique factorization domain, but not regular. (More generally, it is a consequence of

Grothendieck's proof of a conjecture of P. Samuel that a Noetherian local ring that is a complete intersection — in particular a hypersurface — that is factorial in codimension at most 3 must be factorial, [SGA2, Exp. XI, Cor. 3.14]. For a shorter and somewhat more elementary proof, see [CL].)

13.8.7. Regular local rings are integrally closed. The Auslander-Buchsbaum Theorem (Fact 13.8.5) implies that regular local rings are integrally closed (by Exercise 5.4.F). We will prove this (without appealing to Fact 13.8.5) in the most important geometric cases in §29.3.5.

13.8.8. Regular local rings are Cohen-Macaulay. In §29.2.5, we will show that regular local rings are Cohen-Macaulay (a notion defined in Chapter 29).

13.9 ★ Filtered rings and modules, and the Artin-Rees Lemma

We conclude Chapter 13 by discussing the Artin-Rees Lemma 13.9.3, which was used to prove Proposition 13.5.2. The Artin-Rees Lemma generalizes the intuition behind Proposition 13.5.2, that any function that is analytically zero at a point actually vanishes in an open neighborhood of that point (§13.5.3). Because we will use it later (proving the Cohomology and Base Change Theorem 28.1.6), and because it is useful to recognize it in other contexts, we discuss it in some detail.

13.9.1. Definitions. Suppose I is an ideal of a ring A . A descending filtration of an A -module M

$$(13.9.1.1) \quad M = M_0 \supset M_1 \supset M_2 \supset \cdots$$

is called an **I -filtration** if $I^d M_n \subset M_{n+d}$ for all $d, n \geq 0$. An example is the **I -adic filtration** where $M_k = I^k M$. We say an I -filtration is **I -stable** if for some s and all $d \geq 0$, $I^d M_s = M_{d+s}$. For example, the I -adic filtration is I -stable.

Let $A_\bullet(I)$ be the graded ring $\bigoplus_{n \geq 0} I^n$. This is called the **Rees algebra** of the ideal I in A , although we will not need this terminology. Define $M_\bullet(I) := \bigoplus M_n$. It is naturally a *graded* module over $A_\bullet(I)$.

13.9.2. Proposition. *If A is Noetherian, M is a finitely generated A -module, and (13.9.1.1) is an I -filtration, then $M_\bullet(I)$ is a finitely generated $A_\bullet(I)$ -module if and only if the filtration (13.9.1.1) is I -stable.*

Proof. Note that $A_\bullet(I)$ is Noetherian (by Exercise 4.5.G(b), as A is Noetherian, and I is a finitely generated A -module).

Assume first that $M_\bullet(I)$ is finitely generated over the Noetherian ring $A_\bullet(I)$, and hence Noetherian. Consider the increasing chain of $A_\bullet(I)$ -submodules whose k th element L_k is

$$M \oplus M_1 \oplus M_2 \oplus \cdots \oplus M_k \oplus IM_k \oplus I^2 M_k \oplus \cdots$$

(which agrees with $M_\bullet(I)$ up until M_k , and then “ I -stabilizes”). This chain must stabilize by Noetherianness. But $\bigcup L_k = M_\bullet(I)$, so for some $s \in \mathbb{Z}^{\geq 0}$, $L_s = M_\bullet(I)$, so $I^d M_s = M_{s+d}$ for all $d \geq 0$ — (13.9.1.1) is I -stable.

For the other direction, assume that $M_{d+s} = I^d M_s$ for a fixed s and all $d \geq 0$. Then $M_\bullet(I)$ is generated over $A_\bullet(I)$ by $M \oplus M_1 \oplus \cdots \oplus M_s$. But each M_j is finitely generated, so $M_\bullet(I)$ is indeed a finitely generated $A_\bullet(I)$ -module. \square

13.9.3. Artin-Rees Lemma. — Suppose A is a Noetherian ring, and (13.9.1.1) is an I -stable filtration of a finitely generated A -module M . Suppose that $L \subset M$ is a submodule, and let $L_n := L \cap M_n$. Then

$$L = L_0 \supset L_1 \supset L_2 \supset \cdots$$

is an I -stable filtration of L .

Proof. Note that L_\bullet is an I -filtration, as $IL_n \subset IL \cap IM_n \subset L \cap M_{n+1} = L_{n+1}$. Also, $L_\bullet(I)$ is an $A_\bullet(I)$ -submodule of the finitely generated $A_\bullet(I)$ -module $M_\bullet(I)$, and hence finitely generated by Exercise 3.6.Y (as $A_\bullet(I)$ is Noetherian, see the proof of Proposition 13.9.2). \square

An important special case is the following.

13.9.4. Corollary. — Suppose $I \subset A$ is an ideal of a Noetherian ring, and M is a finitely generated A -module, and L is a submodule. Then for some integer s , $I^d(L \cap I^s M) = L \cap I^{d+s} M$ for all $d \geq 0$.

Warning: it need not be true that $I^d L = L \cap I^d M$ for all d . (Can you think of a counterexample to this statement?)

Proof. Apply the Artin-Rees Lemma 13.9.3 to the filtration $M_n = I^n M$. \square

13.9.A. EXERCISE (KRULL INTERSECTION THEOREM).

- (a) Suppose I is an ideal of a Noetherian ring A , and M is a finitely generated A -module. Show that there is some $a \equiv 1 \pmod{I}$ such that $a \cap_{j=1}^\infty I^j M = 0$. Hint: Apply the Artin-Rees Lemma 13.9.3 with $L = \cap_{j=1}^\infty I^j M$ and $M_n = I^n M$. Show that $L = IL$, and apply the first version of Nakayama (Lemma 8.2.8).
- (b) Show that if A is a Noetherian integral domain or a Noetherian local ring, and I is a proper ideal, then $\cap_{j=1}^\infty I^j = 0$. In particular, you will have proved Proposition 13.5.2: if (A, \mathfrak{m}) is a Noetherian local ring, then $\cap_i \mathfrak{m}^i = 0$.

13.9.B. EXERCISE. Make the following precise, and prove it (thereby justifying the intuition in §13.5.3): if X is a locally Noetherian scheme, and f is a function on X that is analytically zero at a point $p \in X$, then f vanishes in a (Zariski) open neighborhood of p .

Part V

Quasicoherent sheaves on schemes, and their uses

CHAPTER 14

More on quasicoherent and coherent sheaves

14.1 Introduction

In Chapter 6, we discussed how *algebraically*, quasicoherent and coherent sheaves generalized modules over a ring. In this chapter, we explore how *geometrically*, quasicoherent and coherent sheaves generalize vector bundles. Informally, a vector bundle V on a geometric space X (such as a manifold) is a family of vector spaces continuously parametrized by points of X . In other words, for each point p of X , there is a vector space, and these vector spaces are glued into a space V so that, as p varies, the vector space above p varies continuously. Nontrivial examples to keep in mind are the tangent bundle to a differentiable manifold, and the Möbius strip over a circle (interpreted as a line bundle). We will make this somewhat more precise in §14.2, but if you have not seen this idea before, you shouldn't be concerned; the main thing we will use this notion for is motivation.

A **free sheaf** on a ringed space X is an \mathcal{O}_X -module isomorphic to $\mathcal{O}_X^{\oplus I}$ where the sum is over some index set I . The **rank** of the free sheaf is the cardinality of I . (This agrees its rank as an \mathcal{O} -module at a point, see Definition 4.3.8.) A **locally free sheaf** on a ringed space X is an \mathcal{O}_X -module locally isomorphic to a free sheaf. This corresponds to the notion of a vector bundle (§14.2). Quasicoherent sheaves form a convenient abelian category containing the locally free sheaves (i.e., vector bundles) that is much smaller than the full category of \mathcal{O} -modules. Quasicoherent sheaves generalize locally free sheaves in much the way that modules generalize free modules. Coherent sheaves are, roughly speaking, a finite rank version of quasicoherent sheaves, which form a well-behaved abelian category containing finite rank locally free sheaves (or equivalently, finite rank vector bundles). Just as the notion of free modules leads us to the notion of modules in general, and finitely generated modules, the notion of free sheaves will lead us inevitably to the notion of quasicoherent sheaves and coherent sheaves. (There is a slight fib in comparing finitely generated modules to coherent sheaves, see §6.4.)

14.2 Vector bundles “=” locally free sheaves

We recall somewhat more precisely the notion of vector bundles on manifolds. Arithmetically-minded readers shouldn't tune out: for example, fractional ideals of the ring of integers in a number field (defined in §10.7.1) turn out to be an example of a “line bundle on a smooth curve” (Exercise 14.2.J).

A **rank n vector bundle on a manifold M** is a map $\pi: V \rightarrow M$ with the structure of an n -dimensional real vector space on $\pi^{-1}(p)$ for each point $p \in M$, such that for every $p \in M$, there is an open neighborhood U and a homeomorphism

$$\phi: U \times \mathbb{R}^n \rightarrow \pi^{-1}(U)$$

over U (so that the diagram

$$(14.2.0.1) \quad \begin{array}{ccc} \pi^{-1}(U) & \xleftarrow{\cong} & U \times \mathbb{R}^n \\ & \searrow \pi|_{\pi^{-1}(U)} & \swarrow \text{projection to first factor} \\ & U & \end{array}$$

commutes) that is an isomorphism of vector spaces over each $p \in U$. An isomorphism (14.2.0.1) is called a **trivialization over U** . If $U = M$, we say that V is a **trivial bundle**.

We call n the **rank** of the vector bundle. A rank 1 vector bundle is called a **line bundle**. (It can also be convenient to be agnostic about the rank of the vector bundle, so it can have different ranks on different connected components. It is also sometimes convenient to consider infinite-rank vector bundles.)

14.2.1. Transition functions. Given trivializations over U_1 and U_2 , over their intersection, the two trivializations must be related by an element T_{12} of GL_n with entries consisting of functions on $U_1 \cap U_2$. If $\{U_i\}$ is a cover of M , and we are given trivializations over each U_i , then the $\{T_{ij}\}$ must satisfy the **cocycle condition**:

$$(14.2.1.1) \quad T_{jk}|_{U_i \cap U_j \cap U_k} \circ T_{ij}|_{U_i \cap U_j \cap U_k} = T_{ik}|_{U_i \cap U_j \cap U_k}.$$

(This implies $T_{ij} = T_{ji}^{-1}$.) The data of the T_{ij} are called **transition functions** (or *transition matrices*) for the trivialization.

This is reversible: given the data of a cover $\{U_i\}$ and transition functions T_{ij} , we can recover the vector bundle (up to unique isomorphism) by “gluing together the various $U_i \times \mathbb{R}^n$ along $U_i \cap U_j$ using T_{ij} ”.

14.2.2. The sheaf of sections. Fix a rank n vector bundle $V \rightarrow M$. The sheaf of sections \mathcal{F} of V (Exercise 2.2.G) is an \mathcal{O}_M -module — given any open set U , we can multiply a section over U by a function on U and get another section.

Moreover, given a trivialization over U , the sections over U are naturally identified with n -tuples of functions of U .

$$\begin{array}{c} U \times \mathbb{R}^n \\ \downarrow \pi \\ U \end{array} \quad \begin{array}{c} \updownarrow \\ \text{n-tuple of functions} \end{array}$$

Thus given a trivialization, over each open set U_i , we have an isomorphism $\mathcal{F}|_{U_i} \xrightarrow{\sim} \mathcal{O}_{U_i}^{\oplus n}$. As mentioned in §14.1, we say that such an \mathcal{F} is a **locally free sheaf of rank n** .

14.2.3. Transition functions for the sheaf of sections. Suppose we have a vector bundle on M , along with a trivialization over an open cover U_i . Suppose we have a section of the vector bundle over M . (This discussion will apply with M replaced

by any open subset.) Then over each U_i , the section corresponds to an n -tuple of functions over U_i , say \vec{s}^i .

14.2.A. EXERCISE. Show that over $U_i \cap U_j$, the vector-valued function \vec{s}^i is related to \vec{s}^j by the (same) transition functions: $T_{ij}\vec{s}^i = \vec{s}^j$. (Don't do this too quickly — make sure your i 's and j 's are on the correct side.)

Given a locally free sheaf \mathcal{F} with rank n , and a trivializing open neighborhood of \mathcal{F} (an open cover $\{U_i\}$ such that over each U_i , $\mathcal{F}|_{U_i} \cong \mathcal{O}_{U_i}^{\oplus n}$ as \mathcal{O} -modules), we have transition functions $T_{ij} \in \mathrm{GL}_n(\mathcal{O}(U_i \cap U_j))$ satisfying the cocycle condition (14.2.1.1). Thus the data of a locally free sheaf of rank n is *equivalent* to the data of a vector bundle of rank n . This change of perspective is useful, and is similar to an earlier change of perspective when we introduced ringed spaces: understanding spaces is the same as understanding (sheaves of) functions on the spaces, and understanding vector bundles (a type of “space over M ”) is the same as understanding (sheaves of) sections.

14.2.4. Definition. A rank 1 locally free sheaf is called an **invertible sheaf**. (Unimportant aside: “invertible sheaf” is a heinous term for something that is essentially a line bundle. The motivation is that if X is a locally ringed space, and \mathcal{F} and \mathcal{G} are \mathcal{O}_X -modules with $\mathcal{F} \otimes_{\mathcal{O}_X} \mathcal{G} \cong \mathcal{O}_X$, then \mathcal{F} and \mathcal{G} are invertible sheaves [MO33489]. Thus in the monoid of \mathcal{O}_X -modules under tensor product, invertible sheaves are the invertible elements. We will never use this fact. People often informally use the phrase “line bundle” when they mean “invertible sheaf”. The phrase *line sheaf* has been proposed but has not caught on.)

14.2.5. Locally free sheaves more generally.

We can generalize the notion of locally free sheaves to schemes (or more generally, ringed spaces) without change. A **locally free sheaf of rank n on a scheme X** is defined as an \mathcal{O}_X -module \mathcal{F} that is locally a free sheaf of rank n . Precisely, there is an open cover $\{U_i\}$ of X such that for each U_i , $\mathcal{F}|_{U_i} \cong \mathcal{O}_{U_i}^{\oplus n}$. This open cover determines transition functions — the data of a cover $\{U_i\}$ of X , and functions $T_{ij} \in \mathrm{GL}_n(\mathcal{O}(U_i \cap U_j))$ satisfying the cocycle condition (14.2.1.1) — which in turn determine the locally free sheaf. As before, given these data, we can find the sections over any open set U . Informally, they are sections of the free sheaves over each $U \cap U_i$ that agree on overlaps. More formally, for each i , they are

$$\vec{s}^i = \begin{pmatrix} s_1^i \\ \vdots \\ s_n^i \end{pmatrix} \in \Gamma(U \cap U_i, \mathcal{O}_X)^n,$$

satisfying $T_{ij}\vec{s}^i = \vec{s}^j$ on $U \cap U_i \cap U_j$.

You should think of these as vector bundles, but just keep in mind that they are not the “same”, just equivalent notions. We will later (Definition 18.1.4) define the “total space” of the vector bundle $V \rightarrow X$ (a scheme over X) in terms of the sheaf version of Spec (or more precisely, $\mathrm{Spec}(\mathrm{Sym} V^\bullet)$). But the locally free sheaf perspective will prove to be more useful. As one example: the definition of a locally free sheaf is much shorter than that of a vector bundle.

As in our motivating discussion, it is sometimes convenient to let the rank vary among connected components, or to consider infinite rank locally free sheaves.

14.2.6. The problem with locally free sheaves / vector bundles (as opposed to vector spaces).

Recall that \mathcal{O}_X -modules form an abelian category: we can talk about kernels, cokernels, and so forth, and we can do homological algebra. Similarly, vector spaces form an abelian category. But locally free sheaves (i.e., vector bundles), along with reasonably natural maps between them (those that arise as maps of \mathcal{O}_X -modules), don't form an abelian category. As a motivating example in the category of manifolds (cf. Warning 14.2.10), consider the map of the trivial line bundle on \mathbb{R} (with coordinate t) to itself, corresponding to multiplying by the coordinate t . Then this map jumps rank, and if you try to define a kernel or cokernel you will get confused.

This problem is resolved by enlarging our notion of nice \mathcal{O}_X -modules in a natural way, to quasicoherent sheaves.

$$\begin{array}{ccccc} \mathcal{O}_X\text{-modules} & \supset & \text{quasicoherent sheaves} & \supset & \text{locally free sheaves} \\ (\text{abelian category}) & & (\text{abelian category}) & & (\text{not an abelian category}) \end{array}$$

You can turn this into a *definition* of quasicoherent sheaves, equivalent to those we gave in §6.1. We want a notion that is local on X of course. So we ask for the smallest abelian subcategory of $\text{Mod}_{\mathcal{O}_X}$ that is “local” and includes vector bundles. It turns out that the main obstruction to vector bundles to be an abelian category is the failure of cokernels of maps of locally free sheaves — as \mathcal{O}_X -modules — to be locally free; we could define quasicoherent sheaves to be those \mathcal{O}_X -modules that are locally cokernels, yielding a description that works more generally on ringed spaces, as described in Exercise 6.3.B. You may wish to later check that our definitions are equivalent to these.

Similarly, in the locally Noetherian setting, finite rank locally free sheaves will sit in a nice smaller abelian category, that of *coherent sheaves*.

$$\begin{array}{ccccc} \text{quasicoherent sheaves} & \supset & \text{coherent sheaves} & \supset & \text{finite rank locally free sheaves} \\ (\text{abelian category}) & & (\text{abelian category}) & & (\text{not an abelian category}) \end{array}$$

14.2.7. Useful constructions on locally free sheaves.

We now give some useful constructions in the form of a series of exercises about locally free sheaves on a scheme. They are useful, important, and surprisingly nontrivial! Two hints: Exercises 14.2.B–14.2.H will apply for ringed spaces in general, so you shouldn't use special properties of schemes. Furthermore, they are all local on X , so you can reduce to the case where the locally free sheaves in question are actually free.

14.2.B. EXERCISE. Suppose \mathcal{F} and \mathcal{G} are locally free sheaves on X of rank m and n respectively. Show that $\text{Hom}_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G})$ is a locally free sheaf of rank mn .

14.2.C. EXERCISE. If \mathcal{E} is a locally free sheaf on X of (finite) rank n , Exercise 14.2.B implies that $\mathcal{E}^\vee := \text{Hom}(\mathcal{E}, \mathcal{O}_X)$ is also a locally free sheaf of rank n . This is called the **dual** of \mathcal{E} (cf. §2.3.4). Given transition functions for \mathcal{E} , describe transition functions for \mathcal{E}^\vee . (Note that if \mathcal{E} is rank 1, i.e., invertible, the transition functions of the dual are the inverse of the transition functions of the original.) Show that $\mathcal{E} \cong \mathcal{E}^{\vee\vee}$. (Caution: your argument showing that there is a canonical isomorphism $\mathcal{F} \xrightarrow{\sim} (\mathcal{F}^\vee)^\vee$ better not also show that there is an isomorphism

$\mathcal{F} \xrightarrow{\sim} \mathcal{F}^\vee$! We will see an example in §15.1 of a locally free \mathcal{F} that is not isomorphic to its dual: the invertible sheaf $\mathcal{O}(1)$ on \mathbb{P}^n .)

14.2.D. EXERCISE. If \mathcal{F} and \mathcal{G} are locally free sheaves, show that $\mathcal{F} \otimes \mathcal{G}$ is a locally free sheaf. (Here \otimes is tensor product as \mathcal{O}_X -modules, defined in Exercise 2.6.J.) If \mathcal{F} is an invertible sheaf, show that $\mathcal{F} \otimes \mathcal{F}^\vee \cong \mathcal{O}_X$.

14.2.E. EXERCISE. Recall that tensor products tend to be only right-exact in general (Exercise 1.3.H). Show that tensoring by a locally free sheaf is exact. More precisely, if \mathcal{F} is a locally free sheaf, and $\mathcal{G}' \rightarrow \mathcal{G} \rightarrow \mathcal{G}''$ is an exact sequence of \mathcal{O}_X -modules, then so is $\mathcal{G}' \otimes \mathcal{F} \rightarrow \mathcal{G} \otimes \mathcal{F} \rightarrow \mathcal{G}'' \otimes \mathcal{F}$. (Possible hint: it may help to check exactness by checking exactness at stalks. Recall that the tensor product of stalks can be identified with the stalk of the tensor product, so for example there is a “natural” isomorphism $(\mathcal{G} \otimes_{\mathcal{O}_X} \mathcal{F})_p \xrightarrow{\sim} \mathcal{G}_p \otimes_{\mathcal{O}_{X,p}} \mathcal{F}_p$, Exercise 2.6.J(b).)

14.2.F. EXERCISE. If \mathcal{E} is a locally free sheaf of finite rank, and \mathcal{F} and \mathcal{G} are \mathcal{O}_X -modules, show that $\text{Hom}(\mathcal{F}, \mathcal{G} \otimes \mathcal{E}) \cong \text{Hom}(\mathcal{F} \otimes \mathcal{E}^\vee, \mathcal{G})$. (Possible hint: first consider the case where \mathcal{E} is free.)

14.2.8. The Picard group.

14.2.G. EXERCISE AND IMPORTANT DEFINITION. Show that the invertible sheaves on X , up to isomorphism, form an abelian group under tensor product. This is called the **Picard group** of X , and is denoted $\text{Pic } X$.

14.2.H. EXERCISE. If $\pi: X \rightarrow Y$ is a morphism of ringed spaces, and \mathcal{G} is a locally free sheaf of rank n on Y , figure out how to define a “pulled back” locally free sheaf of rank n on X , denoted $\pi^*\mathcal{G}$. Then show that Pic is a contravariant functor from the category of ringed spaces to the category Ab of abelian groups.

14.2.I. ★ EXERCISE (FOR THOSE WITH SUFFICIENT COMPLEX-ANALYTIC BACKGROUND). Recall the analytification functor (Exercises 7.3.16 and 11.3.L), that takes a complex finite type reduced scheme and produces a complex analytic space.

(a) If \mathcal{L} is an invertible sheaf on a complex (algebraic) variety X , define (up to unique isomorphism) the corresponding invertible sheaf on the complex variety X_{an} .

(b) Show that the induced map $\text{Pic } X \rightarrow \text{Pic } X_{\text{an}}$ is a group homomorphism.

(c) Show that this construction is functorial: if $\pi: X \rightarrow Y$ is a morphism of complex varieties, the following diagram commutes:

$$\begin{array}{ccc} \text{Pic } Y & \xrightarrow{\pi^*} & \text{Pic } X \\ \downarrow & & \downarrow \\ \text{Pic } Y_{\text{an}} & \xrightarrow{\pi_{\text{an}}^*} & \text{Pic } X_{\text{an}} \end{array}$$

where the vertical maps are the ones you have defined.

14.2.J. ★ EXERCISE (FOR THOSE WITH SUFFICIENT ARITHMETIC BACKGROUND; SEE ALSO PROPOSITION 15.2.10 AND §15.2.13). Recall the definition of the ring of integers \mathcal{O}_K in a number field K , Remark 10.7.1. A **fractional ideal** \mathfrak{a} of \mathcal{O}_K is a nonzero \mathcal{O}_K -submodule of K such that there is a nonzero $a \in \mathcal{O}_K$ such that

$\mathfrak{a}\mathfrak{a} \subset \mathcal{O}_K$. Products of fractional ideals are defined analogously to products of ideals in a ring (defined in Exercise 3.4.C): $\mathfrak{a}\mathfrak{b}$ consists of (finite) \mathcal{O}_K -linear combinations of products of elements of \mathfrak{a} and elements of \mathfrak{b} . Thus fractional ideals form a semigroup under multiplication, with \mathcal{O}_K as the identity. In fact fractional ideals of \mathcal{O}_K form a group.

- (a) Explain how a fractional ideal on a ring of integers in a number field yields an invertible sheaf. (Although we won't need this, it is worth noting that a fractional ideal is the same as an invertible sheaf *with a trivialization at the generic point*.)
- (b) A fractional ideal is **principal** if it is of the form $r\mathcal{O}_K$ for some $r \in K^\times$. Show that any two that differ by a principal ideal yield the same invertible sheaf.
- (c) Show that two fractional ideals that yield the same invertible sheaf differ by a principal ideal.
- (d) The **class group** is defined to be the group of fractional ideals modulo the principal ideals (i.e., modulo K^\times). Give an isomorphism of the class group with the Picard group of \mathcal{O}_K .

(This discussion applies to any Dedekind domain. See Exercise 15.2.S for a follow-up.)

14.2.9. Locally free sheaves on schemes. The next exercises are specific to schemes.

14.2.K. EASY EXERCISE. Prove that locally free sheaves on a scheme are quasicoherent.

Similarly, finite rank locally free sheaves are always finite type, and if \mathcal{O}_X is coherent, finite rank locally free sheaves on X are coherent. (Remember: If \mathcal{O}_X is not coherent, then coherence is a pretty useless notion on X .)

14.2.L. UNIMPORTANT EXERCISE. Use the example of Exercise 6.1.B(b) to show that not every quasicoherent sheaf is locally free.

14.2.10. Warning for those already familiar with vector bundles. Morphisms of vector bundles are more restrictive than morphisms of quasicoherent sheaves that happen to be locally free sheaves. A locally free subsheaf of a locally free sheaf does not always yield a “subvector bundle” of a vector bundle. The archetypal example is the exact sequence of quasicoherent sheaves on $\mathbb{A}_k^1 = \text{Spec } k[t]$ corresponding to the following exact sequence of $k[t]$ -modules:

$$(14.2.10.1) \quad 0 \longrightarrow tk[t] \longrightarrow k[t] \longrightarrow k \longrightarrow 0.$$

The locally free sheaf $\widehat{tk[t]}$ is a subsheaf of $\widehat{k[t]}$, but it does not correspond to a “subvector bundle”; the cokernel is not a vector bundle. (This example came up in §14.2.6.)

14.2.M. EXERCISE. Show that every (finite) rank n vector bundle on \mathbb{A}_k^1 is trivial of rank n . Hint: finitely generated modules over a principal ideal domain are finite direct sums of cyclic modules, as mentioned in Remark 13.5.14. (Be careful: Before invoking this, you need to be sure that the module in question is finitely generated!) See the aside in §15.2.8 for the difficult generalization to \mathbb{A}_k^n .

14.2.11. Definition. We define **rational (and regular) sections of a locally free sheaf** on a scheme X just as we did rational (and regular) functions (see for example §6.5 and §7.5).

14.2.N. EXERCISE. Suppose s is a section of a locally free sheaf \mathcal{F} on a scheme X . Define the notion of the **subscheme cut out by $s = 0$** , denoted (for obvious reasons) $V(s)$. Be sure to check that your definition is independent of choices! (This exercise gives a new solution to Exercise 4.5.S.) Hint: given a trivialization over an open set U , s corresponds to a number of functions f_1, \dots on U ; on U , take the scheme cut out by these functions. Alternate hint that avoids coordinates: figure out how to define it as the largest closed subscheme on which s restricts to 0.

14.2.O. EASY EXERCISE. Show that locally free sheaves on locally Noetherian normal schemes satisfy “Hartogs’s Lemma”: sections defined away from a set of codimension at least 2 extend over that set. (Algebraic Hartogs’s Lemma for Noetherian normal schemes is Theorem 12.3.15.)

14.2.P. EASY EXERCISE. Suppose s is a nonzero rational section of an invertible sheaf on a locally Noetherian normal scheme. Show that if s has no poles, then s is regular. (Hint: Exercise 13.5.H.)

14.2.12. Locally free sheaves in short exact sequences.

14.2.Q. EXERCISE (POSSIBLE HELP FOR LATER PROBLEMS).

(a) Suppose

$$(14.2.12.1) \quad 0 \longrightarrow \mathcal{F}' \longrightarrow \mathcal{F} \longrightarrow \mathcal{F}'' \longrightarrow 0$$

is a short exact sequence of quasicoherent sheaves on X . Suppose $U = \operatorname{Spec} A$ is an affine open set where \mathcal{F}' , \mathcal{F}'' are free, say $\mathcal{F}'|_{\operatorname{Spec} A} = \tilde{A}^{\oplus a}$, $\mathcal{F}''|_{\operatorname{Spec} A} = \tilde{A}^{\oplus b}$. (Here a and b are assumed to be finite for convenience, but this is not necessary, so feel free to generalize to the infinite rank case.) Show that \mathcal{F} is also free on $\operatorname{Spec} A$, and that $0 \rightarrow \mathcal{F}' \rightarrow \mathcal{F} \rightarrow \mathcal{F}'' \rightarrow 0$ can be interpreted as coming from the tautological exact sequence $0 \rightarrow A^{\oplus a} \rightarrow A^{\oplus(a+b)} \rightarrow A^{\oplus b} \rightarrow 0$. (As a consequence, given an exact sequence of quasicoherent sheaves (14.2.12.1) where \mathcal{F}' and \mathcal{F}'' are locally free, \mathcal{F} must also be locally free.)

(b) In the finite rank case, show that given an open covering by trivializing affine open sets (of the form described in (a)), the transition functions (really, matrices) of \mathcal{F} may be interpreted as block upper triangular matrices, where the top left $a \times a$ blocks are transition functions for \mathcal{F}' , and the bottom $b \times b$ blocks are transition functions for \mathcal{F}'' .

14.2.R. EXERCISE. Suppose (14.2.12.1) is a short exact sequence of quasicoherent sheaves on X . By Exercise 14.2.Q(a), if \mathcal{F}' and \mathcal{F}'' are locally free, then \mathcal{F} is too.

(a) If \mathcal{F} and \mathcal{F}'' are locally free of finite rank, show that \mathcal{F}' is too. Hint: Reduce to the case $X = \operatorname{Spec} A$ and \mathcal{F} and \mathcal{F}'' free. Interpret the map $\phi: \mathcal{F} \rightarrow \mathcal{F}''$ as an $n \times m$ matrix M with values in A , with m the rank of \mathcal{F} and n the rank of \mathcal{F}'' . For each point p of X , show that there exist n columns $\{c_1, \dots, c_n\}$ of M that are linearly independent at p and hence near p (as linear independence is given by nonvanishing of the appropriate $n \times n$ determinant). Thus X can be covered by distinguished open subsets in bijection with the choices of n columns of M .

Restricting to one subset and renaming columns, reduce to the case where the determinant of the first n columns of M is invertible. Then change coordinates on $A^{\oplus m} = \mathcal{F}(\text{Spec } A)$ so that M with respect to the new coordinates is the identity matrix in the first n columns, and 0 thereafter. Finally, in this case interpret \mathcal{F}' as $\widetilde{A^{\oplus(m-n)}}$.

(b) If \mathcal{F}' and \mathcal{F} are both locally free, show that \mathcal{F}'' need not be. (Hint: consider (14.2.10.1), which we will soon interpret as the closed subscheme exact sequence (9.1.2.1) for a point on \mathbb{A}^1 .)

14.2.13. Tensor algebra constructions.

For the next exercises, recall the following. If M is an A -module, then the **tensor algebra** $T^\bullet(M)$ is a noncommutative algebra, graded by $\mathbb{Z}^{\geq 0}$, defined as follows. $T^0(M) = A$, $T^n(M) = M \otimes_A \cdots \otimes_A M$ (where n terms appear in the product), and multiplication is what you expect.

The **symmetric algebra** $\text{Sym}^\bullet M$ is a symmetric algebra, graded by $\mathbb{Z}^{\geq 0}$, defined as the quotient of $T^\bullet(M)$ by the (two-sided) ideal generated by all elements of the form $x \otimes y - y \otimes x$ for all $x, y \in M$. Thus $\text{Sym}^n M$ is the quotient of $M \otimes \cdots \otimes M$ by the relations of the form $m_1 \otimes \cdots \otimes m_n - m'_1 \otimes \cdots \otimes m'_n$ where (m'_1, \dots, m'_n) is a rearrangement of (m_1, \dots, m_n) .

The **exterior algebra** $\wedge^\bullet M$ is defined to be the quotient of $T^\bullet M$ by the (two-sided) ideal generated by all elements of the form $x \otimes x$ for all $x \in M$. Expanding $(a+b) \otimes (a+b)$, we see that $a \otimes b = -b \otimes a$ in $\wedge^2 M$. This implies that if 2 is invertible in A (e.g., if A is a field of characteristic not 2), $\wedge^n M$ is the quotient of $M \otimes \cdots \otimes M$ by the relations of the form $m_1 \otimes \cdots \otimes m_n - (-1)^{\text{sgn}(\sigma)} m_{\sigma(1)} \otimes \cdots \otimes m_{\sigma(n)}$ where σ is a permutation of $\{1, \dots, n\}$. The exterior algebra is a “skew-commutative” A -algebra.

Better: both Sym and \wedge can be defined by universal properties. For example, any map of A -modules $M^{\otimes n} \rightarrow N$ that is symmetric in the n entries factors uniquely through $\text{Sym}_A^n(M)$.

It is most correct to write $T_A^\bullet(M)$, $\text{Sym}_A^\bullet(M)$, and $\wedge_A^\bullet(M)$, but the “base ring” A is usually omitted for convenience.

14.2.S. EXERCISE. If \mathcal{F} is locally free of rank m , show that $T^n \mathcal{F}$, $\text{Sym}^n \mathcal{F}$, and $\wedge^n \mathcal{F}$ are locally free, and find their ranks.

(Remark: These constructions can be defined for \mathcal{O} -modules on an arbitrary ringed space.) We note that in this case, $\wedge^{\text{rank } \mathcal{F}} \mathcal{F}$ is denoted $\det \mathcal{F}$, and is called the **determinant (line) bundle** or (both better and worse) the **determinant locally free sheaf**.

You can also define the sheaf of noncommutative algebras $T^\bullet \mathcal{F}$, the sheaf of commutative algebras $\text{Sym}^\bullet \mathcal{F}$, and the sheaf of skew-commutative algebras $\wedge^\bullet \mathcal{F}$.

14.2.T. EXERCISE. Suppose $0 \rightarrow \mathcal{F}' \rightarrow \mathcal{F} \rightarrow \mathcal{F}'' \rightarrow 0$ is an exact sequence of locally free sheaves. Show that for any r , there is a filtration of $\text{Sym}^r \mathcal{F}$

$$\text{Sym}^r \mathcal{F} = \mathcal{G}^0 \supset \mathcal{G}^1 \supset \cdots \supset \mathcal{G}^r \supset \mathcal{G}^{r+1} = 0$$

with subquotients

$$\mathcal{G}^p / \mathcal{G}^{p+1} \cong (\mathrm{Sym}^p \mathcal{F}') \otimes (\mathrm{Sym}^{r-p} \mathcal{F}'').$$

(Here are two different possible hints for this and Exercise 14.2.U: (1) Interpret the transition matrices for \mathcal{F} as block upper triangular, with two blocks, where one diagonal block gives the transition matrices for \mathcal{F}' , and the other gives the transition matrices for \mathcal{F}'' (cf. Exercise 14.2.12.1(b)). Then appropriately interpret the transition matrices for $\mathrm{Sym}^r \mathcal{F}$ as block upper triangular, with $r+1$ blocks. (2) It suffices to consider a small enough affine open set $\mathrm{Spec} A$, where \mathcal{F}' , \mathcal{F} , \mathcal{F}'' are free, and to show that your construction behaves well with respect to localization at an element $f \in A$. In such an open set, the sequence is $0 \rightarrow A^{\oplus p} \rightarrow A^{\oplus(p+q)} \rightarrow A^{\oplus q} \rightarrow 0$ by the Exercise 14.2.Q. Let e_1, \dots, e_p be the standard basis of $A^{\oplus p}$, and f_1, \dots, f_q be the standard basis of $A^{\oplus q}$. Let e'_1, \dots, e'_p be denote the images of e_1, \dots, e_p in $A^{\oplus(p+q)}$. Let f'_1, \dots, f'_q be any lifts of f_1, \dots, f_q to $A^{\oplus(p+q)}$. Note that f'_i is well-defined modulo e'_1, \dots, e'_p . Note that

$$\mathrm{Sym}^r \mathcal{F}|_{\mathrm{Spec} A} \cong \bigoplus_{i=0}^r \mathrm{Sym}^i \mathcal{F}'|_{\mathrm{Spec} A} \otimes_{\mathcal{O}_{\mathrm{Spec} A}} \mathrm{Sym}^{r-i} \mathcal{F}''|_{\mathrm{Spec} A}.$$

Show that $\mathcal{G}^p := \bigoplus_{i=p}^r \mathrm{Sym}^i \mathcal{F}'|_{\mathrm{Spec} A} \otimes_{\mathcal{O}_{\mathrm{Spec} A}} \mathrm{Sym}^{r-i} \mathcal{F}''|_{\mathrm{Spec} A}$ gives a well-defined (locally free) subsheaf that is independent of the choices made, e.g., of the basis $e_1, \dots, e_p, f_1, \dots, f_q$, and the lifts f'_1, \dots, f'_q .)

14.2.U. USEFUL EXERCISE. Suppose $0 \rightarrow \mathcal{F}' \rightarrow \mathcal{F} \rightarrow \mathcal{F}'' \rightarrow 0$ is an exact sequence of locally free sheaves. Show that for any r , there is a filtration of $\wedge^r \mathcal{F}$:

$$\wedge^r \mathcal{F} = \mathcal{G}^0 \supset \mathcal{G}^1 \supset \dots \supset \mathcal{G}^r \supset \mathcal{G}^{r+1} = 0$$

with subquotients

$$\mathcal{G}^p / \mathcal{G}^{p+1} \cong (\wedge^p \mathcal{F}') \otimes (\wedge^{r-p} \mathcal{F}'')$$

for each p . In particular, if the sheaves have finite rank, then $\det \mathcal{F} = (\det \mathcal{F}') \otimes (\det \mathcal{F}'')$.

14.2.V. EXERCISE. Suppose \mathcal{F} is locally free of rank n . Describe a map $\wedge^r \mathcal{F} \times \wedge^{n-r} \mathcal{F} \rightarrow \wedge^n \mathcal{F}$ that induces an isomorphism $\wedge^r \mathcal{F} \xrightarrow{\sim} (\wedge^{n-r} \mathcal{F})^\vee \otimes \wedge^n \mathcal{F}$. This is called a **perfect pairing of vector bundles**. (If you know about perfect pairings of vector spaces, do you see why this is a generalization?) You might use this later in showing duality of Hodge numbers of regular varieties over algebraically closed fields, Exercise 22.5.L.

14.2.W. EXERCISE (DETERMINANT LINE BUNDLES BEHAVE WELL IN EXACT SEQUENCES).

Suppose $0 \rightarrow \mathcal{F}_1 \rightarrow \dots \rightarrow \mathcal{F}_n \rightarrow 0$ is an exact sequence of finite rank locally free sheaves on X . Show that “the alternating product of determinant bundles is trivial”:

$$\det(\mathcal{F}_1) \otimes \det(\mathcal{F}_2)^\vee \otimes \det(\mathcal{F}_3) \otimes \det(\mathcal{F}_4)^\vee \otimes \dots \otimes \det(\mathcal{F}_n)^{(-1)^n} \cong \mathcal{O}_X.$$

(Hint: break the exact sequence into short exact sequences. Use Exercise 14.2.R(a) to show that they are short exact sequences of *finite rank locally free sheaves*. Then use Exercise 14.2.U.)

14.2.X. EXERCISE. Suppose \mathcal{F} is a quasicoherent sheaf. Define the quasicoherent sheaves $T^n \mathcal{F}$, $\text{Sym}^n \mathcal{F}$, and $\wedge^n \mathcal{F}$. (One possibility: describe them on each affine open set, and use the characterization of Important Exercise 6.2.D.)

14.3 More pleasant properties of finite type and coherent sheaves

We now describe a number of important facts about quasicoherent sheaves when we have some finiteness conditions.

14.3.A. EXERCISE (*Hom* BEHAVES REASONABLY IF THE SOURCE IS FINITELY PRESENTED).

(a) (“*Hom*(*fpr*, *qcoh*) is *qcoh*”) Suppose \mathcal{F} is a finitely presented (quasicoherent) sheaf on X , and \mathcal{G} is a quasicoherent sheaf on X . Show that $\text{Hom}(\mathcal{F}, \mathcal{G})$ is a quasicoherent sheaf. Hint: take a “finite presentation” of \mathcal{F} locally: on $U \subset X$, $\mathcal{O}_U^{\oplus m} \rightarrow \mathcal{O}_U^{\oplus n} \rightarrow \mathcal{F} \rightarrow 0$. Apply the left-exact functor (from \mathcal{O}_X -modules to \mathcal{O}_X -modules, §2.6.4) $\text{Hom}_{\mathcal{O}_X}(\cdot, \mathcal{G})$. (Aside: For an example of quasicoherent sheaves \mathcal{F} and \mathcal{G} on a scheme X such that $\text{Hom}(\mathcal{F}, \mathcal{G})$ is not quasicoherent, let A be a discrete valuation ring with uniformizer t , let $X = \text{Spec } A$, let $\mathcal{F} = \widetilde{M}$ and $\mathcal{G} = \widetilde{N}$ with $M = \bigoplus_{i=1}^{\infty} A$ and $N = A$. Then $M_t = \bigoplus_{i=1}^{\infty} A_t$, and of course $N_t = A_t$. Consider the homomorphism $\phi: M_t \rightarrow N_t$ sending 1 in the i th factor of M_t to $1/t^i$. Then ϕ is not the localization of any element of $\text{Hom}_A(M, N)$.)

(b) (“*Hom*(*fpr*, *coh*) is *coh*”) If further \mathcal{G} is coherent and \mathcal{O}_X is coherent, show that $\text{Hom}(\mathcal{F}, \mathcal{G})$ is also coherent.

(c) Suppose \mathcal{F} is a finitely presented sheaf on X , and \mathcal{G} is a quasicoherent sheaf on X . Show that $\text{Hom}(\mathcal{F}, \cdot)$ gives a left-exact covariant functor $\text{QCoh}_X \rightarrow \text{QCoh}_X$ and $\text{Coh}_X \rightarrow \text{Coh}_X$, and that $\text{Hom}(\cdot, \mathcal{G})$ gives a left-exact contravariant functor from finitely presented sheaves on X to QCoh_X (cf. Exercise 2.6.H).

14.3.1. Duals of coherent sheaves. From Exercise 14.3.A(b), if \mathcal{F} is coherent, its dual $\mathcal{F}^\vee := \text{Hom}(\mathcal{F}, \mathcal{O}_X)$ (see §2.3.4) is too. This generalizes the notion of duals of vector bundles in Exercise 14.2.C. Your argument there generalizes to show that there is always a natural “double dual” morphism $\mathcal{F} \rightarrow (\mathcal{F}^\vee)^\vee$. Unlike in the vector bundle case, this is not always an isomorphism. (For an example, let \mathcal{F} be the coherent sheaf associated to $k[t]/(t)$ on $\mathbb{A}^1 = \text{Spec } k[t]$, and show that $\mathcal{F}^\vee = 0$.) Coherent sheaves for which the “double dual” morphism is an isomorphism are called **reflexive sheaves**, but we won’t use this notion. The canonical map $\mathcal{F} \otimes \mathcal{F}^\vee \rightarrow \mathcal{O}_X$ is called the *trace map* — can you see why?

14.3.B. EXERCISE. Suppose \mathcal{F} is a finite rank locally free sheaf, and \mathcal{G} is a quasicoherent sheaf. Describe an isomorphism $\text{Hom}(\mathcal{F}, \mathcal{G}) \xrightarrow{\sim} \mathcal{F}^\vee \otimes \mathcal{G}$. (This holds more generally if \mathcal{G} is an \mathcal{O} -module, but we won’t use that, so you may as well prove the simpler result given in this exercise.)

14.3.C. EXERCISE. Suppose

$$(14.3.1.1) \quad 0 \longrightarrow \mathcal{F} \longrightarrow \mathcal{G} \longrightarrow \mathcal{H} \longrightarrow 0$$

is an exact sequence of quasicoherent sheaves on a scheme X , where \mathcal{H} is a locally free quasicoherent sheaf, and suppose \mathcal{E} is a quasicoherent sheaf. By left-exactness

of $\mathcal{H}om$ (Exercise 2.6.H),

$$0 \longrightarrow \mathcal{H}om(\mathcal{H}, \mathcal{E}) \longrightarrow \mathcal{H}om(\mathcal{G}, \mathcal{E}) \longrightarrow \mathcal{H}om(\mathcal{F}, \mathcal{E}) \longrightarrow 0$$

is exact except possibly on the right. Show that it is also exact on the right. (Hint: this is local, so you can assume that X is affine, say $\text{Spec } A$, and $\mathcal{H} = \widetilde{A^{\oplus I}}$, so (14.3.1.1) can be written as $0 \rightarrow M \rightarrow N \rightarrow A^{\oplus I} \rightarrow 0$. Show that this exact sequence splits, so we can write $N = M \oplus A^{\oplus I}$ in a way that respects the exact sequence.) In particular, if \mathcal{F}, \mathcal{G} , and \mathcal{H} are coherent, and \mathcal{H} is locally free, then we have an exact sequence of coherent sheaves

$$0 \longrightarrow \mathcal{H}^\vee \longrightarrow \mathcal{G}^\vee \longrightarrow \mathcal{F}^\vee \longrightarrow 0.$$

14.3.D. EXERCISE (“Supp(FT \mathcal{F}) IS CLOSED”). Suppose \mathcal{F} is a sheaf of abelian groups. (a) Show that the support of a finite type quasicoherent sheaf on a scheme X is a closed subset. (Hint: Exercise 6.5.B(b).) (b) Show that the support of a quasicoherent sheaf need not be closed. (Hint: If $A = \mathbb{C}[t]$, then $\mathbb{C}[t]/(t - a)$ is an A -module supported at a . Consider $\bigoplus_{a \in \mathbb{C}} \mathbb{C}[t]/(t - a)$. Be careful: this example won’t work if \bigoplus is replaced by \prod .)

We next come to a geometric interpretation of Nakayama’s Lemma, which is why Nakayama’s Lemma should be considered a geometric fact (with an algebraic proof).

14.3.E. USEFUL EXERCISE: GEOMETRIC NAKAYAMA (FOR FINITE TYPE SHEAVES, GENERATORS AT p ARE GENERATORS NEAR p). Suppose X is a scheme, and \mathcal{F} is a finite type quasicoherent sheaf. Show that if $U \subset X$ is an open neighborhood of $p \in X$ and $a_1, \dots, a_n \in \mathcal{F}(U)$ so that their images $a_1|_p, \dots, a_n|_p$ generate the fiber $\mathcal{F}|_p$ (defined as $\mathcal{F}_p \otimes \kappa(p)$, §4.3.8), then there is an affine open neighborhood $p \in \text{Spec } A \subset U$ of p such that “ $a_1|_{\text{Spec } A}, \dots, a_n|_{\text{Spec } A}$ generate $\mathcal{F}|_{\text{Spec } A}$ ” in the following senses:

- (i) $a_1|_{\text{Spec } A}, \dots, a_n|_{\text{Spec } A}$ generate $\mathcal{F}(\text{Spec } A)$ as an A -module;
- (ii) for any $q \in \text{Spec } A$, a_1, \dots, a_n generate the stalk \mathcal{F}_q as an $\mathcal{O}_{X,q}$ -module (and hence for any $q \in \text{Spec } A$, the fibers $a_1|_q, \dots, a_n|_q$ generate the fiber $\mathcal{F}|_q$ as a $\kappa(q)$ -vector space).

In particular, if $\mathcal{F}_p \otimes \kappa(p) = 0$, then there exists an open neighborhood V of p such that $\mathcal{F}|_V = 0$.

14.3.F. USEFUL EXERCISE (FOR FINITELY PRESENTED SHEAVES, LOCALLY FREE = FREE STALKS). Suppose \mathcal{F} is a finitely presented sheaf on a scheme X . Show that if \mathcal{F}_p is a free $\mathcal{O}_{X,p}$ -module for some $p \in X$, then \mathcal{F} is locally free in some open neighborhood of p . Hence \mathcal{F} is locally free if and only if \mathcal{F}_p is a free $\mathcal{O}_{X,p}$ -module for all $p \in X$. Translation: for a finitely presented sheaf, local freeness is a stalk-local property, and free stalks imply local freeness nearby.

Extended hint: Find an open neighborhood U of p , and n elements of $\mathcal{F}(U)$ that generate \mathcal{F}_p . Using Geometric Nakayama, Exercise 14.3.E, show that the sections generate \mathcal{F}_q for all q in some open neighborhood V of p in U . Thus you have described a surjection $\mathcal{O}_V^{\oplus n} \rightarrow \mathcal{F}|_V$. Show that the kernel of this map

is finite type (using Lemma 6.4.4, finitely presented implies always finitely presented), and hence has closed support (say $Z \subset Y$), which does not contain p . Thus $\mathcal{O}_{Y \setminus Z}^{\oplus n} \rightarrow \mathcal{F}|_{Y \setminus Z}$ is an isomorphism.

This is enlightening in a number of ways. It shows that for finitely presented sheaves, local freeness is a stalk-local condition. Furthermore, on an integral scheme, any finitely presented sheaf \mathcal{F} is automatically free over the generic point (do you see why?), so every finitely presented sheaf on an integral scheme is locally free over a dense open subset. And any finitely presented sheaf that is 0 at the generic point of an irreducible scheme is necessarily 0 on a dense open subset. The last two sentences show the utility of generic points; such statements would have been more mysterious in classical algebraic geometry.

14.3.G. EXERCISE. (Torsion-free and torsion sheaves were defined in §6.1.4.)

- (a) Show that torsion-free coherent sheaves on a regular (hence implicitly locally Noetherian) curve are locally free.
- (b) Show that torsion coherent sheaves on a quasicompact regular integral curve are supported at a finite number of closed points.
- (c) Suppose \mathcal{F} is a coherent sheaf on a quasicompact (for convenience) regular curve. Describe a canonical short exact sequence $0 \rightarrow \mathcal{F}_{\text{tors}} \rightarrow \mathcal{F} \rightarrow \mathcal{F}_{\text{lf}} \rightarrow 0$, where \mathcal{F}_{tor} is a torsion sheaf, and \mathcal{F}_{lf} is locally free.

To answer Exercise 14.3.G, use Useful Exercise 14.3.F (local freeness can be checked at stalks) to reduce to the discrete valuation ring case, and recall Remark 13.5.14, the structure theorem for finitely generated modules over a principal ideal domain A : any such module can be written as the direct sum of principal modules $A/(a)$. Hence:

14.3.2. Proposition. — *If M is a finitely generated module over a principal ideal domain, then M is torsion-free if and only if M is free.*

(Exercise 25.2.B is closely related.)

Proposition 14.3.2 is false without the finite generation hypothesis: consider $M = K(A)$ for a suitably general ring A . It is also false if we give up the principal ideal domain hypothesis: consider $(x, y) \subset \mathbb{C}[x, y]$ and $(x, y) \subset \mathbb{C}[x, y]/(xy)$. (These examples require some verification.) These examples show that Exercise 14.3.G(a) is false if we give up the “dimension 1” or “regular” hypothesis.

14.3.3. The fiber and rank of a quasicoherent sheaf at a point.

We now apply the definitions of fiber and rank of an \mathcal{O}_X -module at a point $p \in X$ (§4.3.8) to a quasicoherent sheaf \mathcal{F} on a scheme X . The *fiber* of \mathcal{F} at p is the vector space

$$\mathcal{F}|_p := \mathcal{F}_p / \mathfrak{m}_p \mathcal{F}_p = \mathcal{F}_p \otimes_{\mathcal{O}_{X,p}} \kappa(p),$$

where \mathfrak{m} is the maximal ideal of $\mathcal{O}_{X,p}$, and $\kappa(p)$ is as usual the residue field $\mathcal{O}_{X,p}/\mathfrak{m}$ at p . A section of \mathcal{F} over an open set containing p can be said to take on a value at p , which is an element of this vector space. The *rank* of a quasicoherent sheaf \mathcal{F} at a point p is $\dim_{\kappa(p)} \mathcal{F}_p / \mathfrak{m}_p \mathcal{F}_p$ (possibly infinite). More explicitly, on any affine set $\text{Spec } A$ where $p = [p]$ and $\mathcal{F}(\text{Spec } A) = M$, then the rank is $\dim_{K(A/p)} M_p / \mathfrak{p} M_p$. Note that this definition of rank is consistent with the notion of rank of a locally

free sheaf. In the locally free case, the rank is a (locally) constant function of the point. The converse is sometimes true, see Exercise 14.3.K below.

If X is irreducible, and \mathcal{F} is a quasicoherent (usually coherent) sheaf on X , then $\text{rank } \mathcal{F}$ (with no mention of a point) by convention means at the generic point. (For example, a rank 0 quasicoherent sheaf on an integral scheme is a torsion quasicoherent sheaf, see Definition 6.1.5.)

14.3.H. EXERCISE. Consider the coherent sheaf \mathcal{F} on $\mathbb{A}_k^1 = \text{Spec } k[t]$ corresponding to the module $k[t]/(t)$. Find the rank of \mathcal{F} at every point of \mathbb{A}^1 . Don't forget the generic point!

14.3.I. EXERCISE. Show that at any point, $\text{rank}(\mathcal{F} \oplus \mathcal{G}) = \text{rank}(\mathcal{F}) + \text{rank}(\mathcal{G})$ and $\text{rank}(\mathcal{F} \otimes \mathcal{G}) = \text{rank } \mathcal{F} \text{ rank } \mathcal{G}$. (Hint: Show that direct sums and tensor products commute with ring quotients and localizations, i.e., $(M \oplus N) \otimes_R (R/I) \cong M/IM \oplus N/IN$, $(M \otimes_R N) \otimes_R (R/I) \cong (M \otimes_R R/I) \otimes_{R/I} (N \otimes_R R/I) \cong M/IM \otimes_{R/I} N/IN$, etc.)

If \mathcal{F} is a finite type quasicoherent sheaf, then the rank of \mathcal{F} at p is finite, and by Nakayama's Lemma, the rank is the minimal number of generators of M_p as an A_p -module.

14.3.J. IMPORTANT EXERCISE (RANK IS UPPERSEMICONTINUOUS FOR FINITE TYPE QUASICOHERENT SHEAVES). If \mathcal{F} is a finite type quasicoherent sheaf on X , show that $\text{rank}(\mathcal{F})$ is an upper semicontinuous function on X . Hint: generators at a point p are generators nearby by Geometric Nakayama's Lemma, Exercise 14.3.E. (The example in Exercise 14.3.D shows the necessity of the finite type hypothesis.)

14.3.K. IMPORTANT EXERCISE (FOR FINITE TYPE QUASICOHERENT SHEAVES ON REDUCED SCHEMES, CONSTANT RANK = VECTOR BUNDLE).

(a) If X is reduced, \mathcal{F} is a finite type quasicoherent sheaf on X , and the rank of \mathcal{F} is constant, show that \mathcal{F} is locally free.

Hint: prove the result in an open neighborhood of an arbitrary point p as follows. Suppose $n = \text{rank}_p \mathcal{F}$. Choose n generators of the fiber $\mathcal{F}|_p$ (a basis as a $\kappa(p)$ -vector space). By Geometric Nakayama's Lemma 14.3.E, we can find an affine open neighborhood $\text{Spec } A \subset X$ of p , with $\mathcal{F}|_{\text{Spec } A} = \widetilde{M}$, so that the chosen generators $\mathcal{F}|_p$ lift to generators m_1, \dots, m_n of M . Let $\phi: A^{\oplus n} \twoheadrightarrow M$ $((r_1, \dots, r_n) \mapsto \sum r_i m_i)$ be the corresponding surjection. Suppose $\ker \phi \neq 0$ — say $(r_1, \dots, r_n) \in \ker \phi$, with $r_i \neq 0$. As X is reduced, there is some p where $r_i \notin \mathfrak{p}$. Then r_i is invertible in A_p , so M_p has fewer than n generators, contradicting the constancy of rank.

(b) Then use upper semicontinuity of rank (Exercise 14.3.J) to show that finite type quasicoherent sheaves on a reduced scheme are locally free on a dense open set.

(c) Show that part (a) can be false without the condition of X being reduced. (Hint: Consider $A = k[x]/(x^2)$, $M = k$.)

You can use the notion of rank to help visualize finite type quasicoherent sheaves, or even quasicoherent sheaves. For example, finite type quasicoherent sheaves generalize finite rank vector bundles as follows: to each point there is an associated vector space, and although the ranks can jump, they fit together in families as well as one might hope. You might try to visualize the example of

Example 14.3.H. Nonreducedness can fit into the picture as well — how would you picture the coherent sheaf on $\operatorname{Spec} k[\epsilon]/(\epsilon^2)$ corresponding to $k[\epsilon]/(\epsilon)$ (Exercise 14.3.K(c))? How about $k[\epsilon]/(\epsilon^2) \oplus k[\epsilon]/(\epsilon)$?

14.3.4. Degree of a finite morphism at a point. Suppose $\pi: X \rightarrow Y$ is a finite morphism. Then $\pi_* \mathcal{O}_X$ is a finite type (quasicoherent) sheaf on Y , and the rank of this sheaf at a point p is called the **degree** of the finite morphism at p . By Exercise 14.3.J, the degree of π is an upper semicontinuous function on Y . The degree can jump: consider the closed embedding of a point into a line corresponding to $k[t] \rightarrow k$ given by $t \mapsto 0$. It can also be constant in cases that you might initially find surprising — see Exercise 10.3.4, where the degree is always 2, but the 2 is obtained in a number of different ways.

14.3.L. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a finite morphism. By unwinding the definition, verify that the degree of π at p is the dimension of the space of functions on the scheme-theoretic preimage of p , considered as a vector space over the residue field $\kappa(p)$. In particular, the degree is zero if and only if $\pi^{-1}(p)$ is empty.

Line bundles: Invertible sheaves and divisors

When learning algebraic geometry for the first time, it is surprising to hear that line bundles (invertible sheaves) are of such vital importance. *i)* They are, of course, the simplest vector bundles. *ii)* We can often classify them, using “codimension 1” information. In particular, we can build them and manipulate them. *iii)* Knowing line bundles on a variety X helps us classify X , and understand maps to projective space, and to other varieties. *iv)* And this only the beginning.

In this chapter, we next describe convenient and powerful ways of working with and classifying line bundles. We begin with a fundamental example, the line bundles $\mathcal{O}(n)$ on projective space, §15.1. We then introduce Weil divisors (formal sums of codimension 1 subsets), and use them to determine $\text{Pic } X$ in a number of circumstances, §15.2. We finally discuss sheaves of ideals that happen to be invertible (effective Cartier divisors), §15.3. A recurring theme is that line bundles are closely related to “codimension 1 information”.

Based on your intuition for line bundles on manifolds, you might hope that every point has a “small” open neighborhood on which all invertible sheaves (or locally free sheaves) are trivial. Sadly, this is not the case. We will eventually see (§20.11.1) that for the curve $y^2 - x^3 - x = 0$ in $\mathbb{A}_{\mathbb{C}}^2$, every nonempty open set has nontrivial invertible sheaves. (This will use the fact that it is an open subset of an *elliptic curve*.) This is a persistent problem in algebraic geometry: the Zariski topology is too blunt.

15.1 Some line bundles on projective space

We now describe an important family of invertible sheaves on projective space over a field k .

As a warm-up, we begin with the invertible sheaf $\mathcal{O}_{\mathbb{P}_k^1}(1)$ on $\mathbb{P}_k^1 = \text{Proj } k[x_0, x_1]$. The subscript \mathbb{P}_k^1 refers to the space on which the sheaf lives, and is often omitted when it is clear from the context. We describe the invertible sheaf $\mathcal{O}(1)$ using transition functions. It is trivial on the usual affine open sets $U_0 = D(x_0) = \text{Spec } k[x_1/x_0]$ and $U_1 = D(x_1) = \text{Spec } k[x_0/x_1]$. (We continue to use the convention $x_{i/j}$ for describing coordinates on patches of projective space, see §4.4.9.) Thus the data of a section over U_0 is a polynomial in x_1/x_0 . The transition function from U_0 to U_1 is multiplication by $x_{0/1} = x_{1/0}^{-1}$. The transition function from U_1 to U_0 is hence multiplication by $x_{1/0} = x_{0/1}^{-1}$.

This information is summarized below:

$$\begin{array}{lll}
 \text{open cover} & \mathcal{U}_0 = \operatorname{Spec} k[x_{1/0}] & \mathcal{U}_1 = \operatorname{Spec} k[x_{0/1}] \\
 \\
 \text{trivialization and transition functions} & \begin{array}{ccc} & \xrightarrow{\times x_{0/1} = x_{1/0}^{-1}} & \\ k[x_{1/0}] & \xleftrightarrow{\hspace{1cm}} & k[x_{0/1}] \\ & \xleftarrow{\times x_{1/0} = x_{0/1}^{-1}} & \end{array} &
 \end{array}$$

To test our understanding, let's compute the global sections of $\mathcal{O}(1)$. This will generalize our hands-on calculation that $\Gamma(\mathbb{P}_k^1, \mathcal{O}_{\mathbb{P}_k^1}) \cong k$ (Example 4.4.6). A global section is a polynomial $f(x_{1/0}) \in k[x_{1/0}]$ and a polynomial $g(x_{0/1}) \in k[x_{0/1}]$ such that $f(1/x_{0/1})x_{0/1} = g(x_{0/1})$. A little thought will show that f must be linear: $f(x_{1/0}) = ax_{1/0} + b$, and hence $g(x_{0/1}) = a + bx_{0/1}$. Thus

$$\dim \Gamma(\mathbb{P}_k^1, \mathcal{O}(1)) = 2 \neq 1 = \dim \Gamma(\mathbb{P}_k^1, \mathcal{O}).$$

Thus $\mathcal{O}(1)$ is not isomorphic to \mathcal{O} , and we have constructed our first (proved) example of a nontrivial line bundle!

We next define more generally $\mathcal{O}_{\mathbb{P}_k^1}(n)$ on \mathbb{P}_k^1 . It is defined in the same way, except that the transition functions are the n th powers of those for $\mathcal{O}(1)$.

$$\begin{array}{lll}
 \text{open cover} & \mathcal{U}_0 = \operatorname{Spec} k[x_{1/0}] & \mathcal{U}_1 = \operatorname{Spec} k[x_{0/1}] \\
 \\
 \text{trivialization and transition functions} & \begin{array}{ccc} & \xrightarrow{\times x_{0/1}^n = x_{1/0}^{-n}} & \\ k[x_{1/0}] & \xleftrightarrow{\hspace{1cm}} & k[x_{0/1}] \\ & \xleftarrow{\times x_{1/0}^n = x_{0/1}^{-n}} & \end{array} &
 \end{array}$$

In particular, thanks to the explicit transition functions, we see that $\mathcal{O}(n) = \mathcal{O}(1)^{\otimes n}$ (with the obvious meaning if n is negative: $(\mathcal{O}(1)^{\otimes (-n)})^\vee$). Clearly also $\mathcal{O}(m) \otimes \mathcal{O}(n) = \mathcal{O}(m+n)$.

15.1.A. IMPORTANT EXERCISE. Show that $\dim \Gamma(\mathbb{P}^1, \mathcal{O}(n)) = n+1$ if $n \geq 0$, and 0 otherwise.

15.1.1. Example. Long ago (in Exercise 2.6.J(a)), we warned that sheafification was necessary when tensoring \mathcal{O}_X -modules: if \mathcal{F} and \mathcal{G} are two \mathcal{O}_X -modules on a ringed space, then it is not necessarily true that $\mathcal{F}(X) \otimes_{\mathcal{O}_X(X)} \mathcal{G}(X) \cong (\mathcal{F} \otimes \mathcal{G})(X)$. We now have an example: let $X = \mathbb{P}_k^1$, $\mathcal{F} = \mathcal{O}(1)$, $\mathcal{G} = \mathcal{O}(-1)$, and use the fact that $\mathcal{O}(-1)$ has no nonzero global sections.

15.1.B. EXERCISE. Show that if $m \neq n$, then $\mathcal{O}(m) \not\cong \mathcal{O}(n)$. Hence conclude that we have an injection of groups $\mathbb{Z} \hookrightarrow \operatorname{Pic} \mathbb{P}_k^1$ given by $n \mapsto \mathcal{O}(n)$.

It is useful to identify the global sections of $\mathcal{O}(n)$ with the homogeneous polynomials of degree n in x_0 and x_1 , i.e., with the degree n part of $k[x_0, x_1]$ (cf. §15.1.2 for the generalization to \mathbb{P}^m). Can you see this from your solution to Exercise 15.1.A? We will see that this identification is natural in many ways. For example, you can show that the definition of $\mathcal{O}(n)$ doesn't depend on a choice of affine cover, and this polynomial description is also independent of cover. (For this, see Example 4.5.12; you can later compare this to Exercise 28.1.M.) As an immediate check of the usefulness of this point of view, ask yourself: where does the

section $x_0^3 - x_0x_1^2$ of $\mathcal{O}(3)$ vanish? The section $x_0 + x_1$ of $\mathcal{O}(1)$ can be multiplied by the section x_0^2 of $\mathcal{O}(2)$ to get a section of $\mathcal{O}(3)$. Which one? Where does the rational section $x_0^4(x_1 + x_0)/x_1^7$ of $\mathcal{O}(-2)$ have zeros and poles, and to what order? (We saw the notion of zeros and poles in Definition 13.5.7, and will meet them again in §15.2, but you should intuitively answer these questions already.)

We now define the invertible sheaf $\mathcal{O}_{\mathbb{P}_k^m}(n)$ on the projective space \mathbb{P}_k^m . On the usual affine open set $U_i = \text{Spec } k[x_{0/i}, \dots, x_{m/i}]/(x_{i/i} - 1) = \text{Spec } A_i$, it is trivial, so sections (as an A_i -module) are isomorphic to A_i . The transition function from U_i to U_j is multiplication by $x_{i/j}^n = x_{j/i}^{-n}$.

$$U_i = \text{Spec } k[x_{0/i}, \dots, x_{m/i}]/(x_{i/i} - 1) \quad U_j = \text{Spec } k[x_{0/j}, \dots, x_{m/j}]/(x_{j/j} - 1)$$

$$k[x_{0/i}, \dots, x_{m/i}]/(x_{i/i} - 1) \begin{array}{c} \xrightarrow{\times x_{i/j}^n = x_{j/i}^{-n}} \\ \xleftarrow{\times x_{j/i}^n = x_{i/j}^{-n}} \end{array} k[x_{0/j}, \dots, x_{m/j}]/(x_{j/j} - 1)$$

Note that these transition functions clearly satisfy the cocycle condition.

15.1.C. ESSENTIAL EXERCISE (CF. EXERCISE 9.3.J, THEOREM 19.1.2). Show that $\dim_k \Gamma(\mathbb{P}_k^m, \mathcal{O}_{\mathbb{P}_k^m}(n)) = \binom{m+n}{m}$.

15.1.2. As in the case of \mathbb{P}^1 , sections of $\mathcal{O}(n)$ on \mathbb{P}_k^m are naturally identified with homogeneous degree n polynomials in our $m+1$ variables. (Important question: Do you see why? Can you work out this dictionary?) Thus $x + y + 2z$ is a section of $\mathcal{O}(1)$ on \mathbb{P}^2 . It isn't a function, but we know where this section vanishes — precisely where $x + y + 2z = 0$.

Also, notice that for fixed m , $\binom{m+n}{m}$ is a polynomial in n of degree m for $n \geq 0$ (or better: for $n \geq -m$). This should be telling you that this function “wants to be a polynomial,” but won't succeed without assistance. We will later define $h^0(\mathbb{P}_k^m, \mathcal{O}(n)) := \dim_k \Gamma(\mathbb{P}_k^m, \mathcal{O}(n))$, and later still we will define higher cohomology groups, and we will define the *Euler characteristic* $\chi(\mathbb{P}_k^m, \mathcal{O}(n)) := \sum_{i=0}^{\infty} (-1)^i h^i(\mathbb{P}_k^m, \mathcal{O}(n))$ (cohomology will vanish in degree higher than m). We will discover the moral that the Euler characteristic is better-behaved than h^0 , and so we should now suspect (and later prove, see Theorem 19.1.2, Remark 19.3.1, and Exercise 19.3.B) that this polynomial is in fact the Euler characteristic, and the reason that it agrees with h^0 for $n \geq 0$ is because all the other cohomology groups should vanish.

We finally note that we can define $\mathcal{O}(n)$ on \mathbb{P}_A^m for any ring A : the above definition applies without change.

15.1.3. These are the only line bundles on \mathbb{P}_k^m .

Suppose that k is a field. We will see in §15.2.9 that these $\mathcal{O}(n)$ are the only invertible sheaves on \mathbb{P}_k^m . The next Exercise shows this when $m = 1$, although this approach will soon be trumped.

15.1.D. EXERCISE. Show that every invertible sheaf on \mathbb{P}_k^1 is of the form $\mathcal{O}(n)$ for some n . Hint: use the classification of finitely generated modules over a principal

ideal domain (Remark 13.5.14) to show that all invertible sheaves on \mathbb{A}_k^1 are trivial (a special case of Exercise 14.2.M). Reduce to determining possible transition functions between the two open subsets in the standard cover of \mathbb{P}_k^1 .

Caution: there can exist invertible sheaves on \mathbb{P}_A^1 not of the form $\mathcal{O}(n)$. You may later be able to think of examples. (Hints to find an example for when you know more: what if $\text{Spec } A$ is disconnected? Or if that leads to too silly an example, what if $\text{Spec } A$ has nontrivial invertible sheaves?)

15.2 Line bundles and Weil divisors

The notion of Weil divisors gives a great way of understanding and classifying line bundles, at least on Noetherian normal schemes. Some of what we discuss will apply in more general circumstances, and the expert is invited to consider generalizations by judiciously weakening hypotheses in various statements. Before we get started, you should be warned: this is one of those topics in algebraic geometry that is hard to digest — learning it changes the way in which you think about line bundles. But once you become comfortable with the imperfect dictionary to divisors, it becomes second nature.

For the rest of this section, we consider only *Noetherian schemes*. We do this because we will use finite decomposition into irreducible components (Exercise 5.3.B), and Algebraic Hartogs's Lemma 12.3.15.

Define a **Weil divisor** as a formal \mathbb{Z} -linear combination of codimension 1 irreducible closed subsets of X . In other words, a Weil divisor is defined to be an object of the form

$$\sum_{Y \subset X \text{ codimension } 1} n_Y [Y]$$

where the n_Y are integers, all but a finite number of which are zero. Weil divisors obviously form an abelian group, denoted $\text{Weil } X$. For example, if X is a curve, the Weil divisors are linear combinations of closed points.

We say that $[Y]$ is an **irreducible** (Weil) divisor. A Weil divisor $D = \sum n_Y [Y]$ is said to be **effective** if $n_Y \geq 0$ for all Y . In this case we say $D \geq 0$, and by $D_1 \geq D_2$ we mean $D_1 - D_2 \geq 0$. The **support** of a Weil divisor D , denoted $\text{Supp } D$, is the subset $\cup_{n_Y \neq 0} Y$. If $U \subset X$ is an open set, we define the **restriction map** $\text{Weil } X \rightarrow \text{Weil } U$ by $\sum n_Y [Y] \mapsto \sum_{Y \cap U \neq \emptyset} n_Y [Y \cap U]$.

Suppose now that X is *regular in codimension 1*. We add this hypothesis because we will use properties of discrete valuation rings. Assume also that X is *reduced*. (This is only so we can talk about rational functions without worrying about them being defined at embedded points. Feel free to relax this hypothesis.) Suppose that \mathcal{L} is an invertible sheaf, and s a rational section not vanishing everywhere on any irreducible component of X . (Rational sections are given by a section over a dense open subset of X , with the obvious equivalence, §14.2.11.) Then s determines a Weil divisor

$$\text{div}(s) := \sum_Y \text{val}_Y(s) [Y]$$

where the summation runs over all irreducible divisors Y of X . We call $\text{div}(s)$ the **divisor of zeros and poles** of the rational section s (cf. Definition 13.5.7). To

determine the valuation $\text{val}_Y(s)$ of s along Y , take any open set U containing the generic point of Y where \mathcal{L} is trivializable, along with any trivialization over U ; under this trivialization, s is a nonzero rational function on U , which thus has a valuation. Any two such trivializations differ by an invertible function (transition functions are invertible), so this valuation is well-defined. Note that $\text{val}_Y(s) = 0$ for all but finitely many Y , by Exercise 13.5.G.

Now consider the set $\{(\mathcal{L}, s)\}$ of pairs of line bundles \mathcal{L} with nonzero rational sections s of \mathcal{L} , up to isomorphism. This set (after taking quotient by isomorphism) forms an abelian group under tensor product \otimes , with identity $(\mathcal{O}_X, 1)$. (Tricky question: what is the inverse of (\mathcal{L}, s) in this group?) It is important to notice that if t is an invertible function on X , then multiplication by t gives an isomorphism $(\mathcal{L}, s) \xrightarrow{\sim} (\mathcal{L}, st)$. Similarly, $(\mathcal{L}, s)/(\mathcal{L}, u) = (\mathcal{O}, s/u)$. The map div yields a group homomorphism

$$(15.2.0.1) \quad \text{div}: \{(\mathcal{L}, s)/\text{isomorphism}\} \longrightarrow \text{Weil } X.$$

15.2.A. EASIER EXERCISE.

(a) (*divisors of rational functions*) Verify that on \mathbb{A}_k^1 , $\text{div}(x^3/(x+1)) = 3[(x)] - [(x+1)]$ (“ $= 3[0] - [-1]$ ”).

(b) (*divisor of rational sections of a nontrivial invertible sheaf*) On \mathbb{P}_k^1 , there is a rational section of $\mathcal{O}(1)$ “corresponding to” $x^2/(x+y)$. Figure out what this means, and calculate $\text{div}(x^2/(x+y))$.

The homomorphism (15.2.0.1) will be the key to determining all the line bundles on many X . (Note that any invertible sheaf will have such a rational section. For each irreducible component, take a nonempty open set not meeting any other irreducible component; then shrink it so that \mathcal{L} is trivial; choose a trivialization; then take the union of all these open sets, and choose the section on this union corresponding to 1 under the trivialization.) We will see that in reasonable situations, this map div will be injective, and often an isomorphism. Thus by forgetting the rational section (i.e., taking an appropriate quotient), we will have described the Picard group of all line bundles. Let’s put this strategy into action.

15.2.1. Proposition. — *If X is normal and Noetherian then the map div is injective.*

Proof. Suppose $\text{div}(\mathcal{L}, s) = 0$. Then s has no poles. By Exercise 14.2.P, s is a regular section. We now show that the morphism $\times s: \mathcal{O}_X \rightarrow \mathcal{L}$ is in fact an isomorphism; this will prove the Proposition, as it will give an isomorphism $(\mathcal{O}_X, 1) \xrightarrow{\sim} (\mathcal{L}, s)$.

It suffices to show that $\times s$ is an isomorphism on an open subset U of X where \mathcal{L} is trivial, as X is covered by trivializing open neighborhoods of \mathcal{L} (as \mathcal{L} is locally trivial). Choose an isomorphism $i: \mathcal{L}|_U \xrightarrow{\sim} \mathcal{O}_U$. Composing $\times s$ with i yields a map $\times s': \mathcal{O}_U \rightarrow \mathcal{O}_U$ that is multiplication by a rational function $s' = i(s)$ that has no zeros and no poles. The rational function s' is regular because it has no poles (Exercise 13.5.H), and $1/s'$ is regular for the same reason. Thus s' is an invertible function on U , so $\times s'$ is an isomorphism. Hence $\times s$ is an isomorphism over U . \square

Motivated by this, we try to find an inverse to div , or at least to determine the image of div .

15.2.2. Important Definition. Assume now that X is irreducible (purely to avoid making (15.2.2.1) look uglier — but feel free to relax this, see Exercise 15.2.B). **Assume also that X is normal** — this will be a standing assumption for the rest of this section. Suppose D is a Weil divisor. Define the sheaf $\mathcal{O}_X(D)$ by

$$(15.2.2.1) \quad \Gamma(U, \mathcal{O}_X(D)) := \{t \in K(X)^\times : \operatorname{div}|_U t + D|_U \geq 0\} \cup \{0\}.$$

Here $\operatorname{div}|_U t$ means the divisor of t considered as a rational function on U , i.e., consider just the irreducible divisors of U . (The subscript X in $\mathcal{O}_X(D)$ is omitted when it is clear from context.) The sections of $\mathcal{O}_X(D)$ over U are the rational functions on U that have poles and zeros “constrained by D ”: a positive coefficient in D allows a pole of that order; a negative coefficient demands a zero of that order. Away from the support of D , this is (isomorphic to) the structure sheaf (by Algebraic Hartogs’s Lemma 12.3.15).

15.2.3. Remark. It will be helpful to note that $\mathcal{O}_X(D)$ comes along with a canonical “rational section” corresponding to $1 \in K(X)^\times$. (It is a rational section in the sense that it is a section over a dense open set, namely the complement of $\operatorname{Supp} D$.)

15.2.B. LESS IMPORTANT EXERCISE. Generalize the definition of $\mathcal{O}_X(D)$ to the case when X is not necessarily irreducible. (This is just a question of language. Once you have done this, feel free to drop this hypothesis in the rest of this section.)

15.2.C. EASY EXERCISE. Verify that $\mathcal{O}_X(D)$ is a quasicoherent sheaf. (Hint: the distinguished affine criterion for quasicoherence of Exercise 6.2.D.)

In good situations, $\mathcal{O}_X(D)$ is an invertible sheaf. For example, let $X = \mathbb{A}_k^1$. Consider

$$\mathcal{O}_X(-2[(x)] + [(x-1)] + [(x-2)]),$$

often written $\mathcal{O}(-2[0] + [1] + [2])$ for convenience. Then $3x^3/(x-1)$ is a global section; it has the required two zeros at $x = 0$ (and even one to spare), and takes advantage of the allowed pole at $x = 1$, and doesn’t have a pole at $x = 2$, even though one is allowed. (Unimportant aside: the statement remains true in characteristic 2, although the explanation requires editing.)

15.2.D. EASY EXERCISE. (This is a consequence of later discussion as well, but you should be able to do this by hand.)

(a) Show that any global section of $\mathcal{O}_{\mathbb{A}_k^1}(-2[(x)] + [(x-1)] + [(x-2)])$ is a $k[x]$ -multiple of $x^2/(x-1)(x-2)$.

(b) Extend the argument of (a) to give an isomorphism

$$\mathcal{O}_{\mathbb{A}_k^1}(-2[(x)] + [(x-1)] + [(x-2)]) \xrightarrow{\sim} \mathcal{O}_{\mathbb{A}_k^1}.$$

In the next several exercises, we show that in good circumstances, $\mathcal{O}_X(D)$ is an invertible sheaf. (In fact the $\mathcal{O}_X(D)$ construction can be useful even if $\mathcal{O}_X(D)$ is *not* an invertible sheaf, but this won’t concern us here. An example of an $\mathcal{O}_X(D)$ that is not an invertible sheaf is given in Exercise 15.2.H.)

15.2.E. HARD BUT IMPORTANT EXERCISE. Suppose \mathcal{L} is an invertible sheaf, and s is a nonzero rational section of \mathcal{L} .

- (a) Describe an isomorphism $\mathcal{O}(\operatorname{div} s) \xrightarrow{\sim} \mathcal{L}$. (You will use the normality hypothesis!) Hint: show that those open subsets U for which $\mathcal{O}(\operatorname{div} s)|_U \cong \mathcal{O}_U$ form a base for the Zariski topology. For each such U , define $\phi_U: \mathcal{O}(\operatorname{div} s)(U) \rightarrow \mathcal{L}(U)$ sending a rational function t (with zeros and poles “constrained by $\operatorname{div} s$ ”) to st . Show that ϕ_U is an isomorphism (with the obvious inverse map, division by s). Argue that this map induces an isomorphism of sheaves $\phi: \mathcal{O}(\operatorname{div} s) \xrightarrow{\sim} \mathcal{L}$.
- (b) Let σ be the map from $K(X)$ to the rational sections of \mathcal{L} , where $\sigma(t)$ is the rational section of $\mathcal{O}_X(D) \cong \mathcal{L}$ defined via (15.2.2.1) (as described in Remark 15.2.3). Show that the isomorphism of (a) can be chosen such that $\sigma(1) = s$. (Hint: the map in part (a) sends 1 to s .)

15.2.F. EXERCISE (THE EXAMPLE OF §15.1). Suppose $X = \mathbb{P}_k^n$, $\mathcal{L} = \mathcal{O}(1)$, s is the section of $\mathcal{O}(1)$ corresponding to x_0 , and $D = \operatorname{div} s$. Verify that $\mathcal{O}(mD) \cong \mathcal{O}(m)$, and the canonical rational section of $\mathcal{O}(mD)$ is precisely s^m . (Watch out for possible confusion: 1 has no pole along $x_0 = 0$, but $\sigma(1) = s^m$ *does* have a zero if $m > 0$.) For this reason, $\mathcal{O}(1)$ is sometimes called the **hyperplane class** in $\operatorname{Pic} X$. (Of course, x_0 can be replaced by any linear form.)

15.2.4. Definition. If D is a Weil divisor on (Noetherian normal irreducible) X such that $D = \operatorname{div} f$ for some rational function f , we say that D is **principal**. Principal divisors clearly form a subgroup of $\operatorname{Weil} X$; denote this group of principal divisors $\operatorname{Prin} X$. Note that div induces a group homomorphism $K(X)^\times \rightarrow \operatorname{Prin} X$. If X can be covered with open sets U_i such that on U_i , D is principal, we say that D is **locally principal**. Locally principal divisors form a subgroup of $\operatorname{Weil} X$, which we denote $\operatorname{LocPrin} X$. (This notation is not standard, and we only use it in (15.2.7.1).)

15.2.5. Important observation. As a consequence of Exercise 15.2.E(a) (taking $\mathcal{L} = \mathcal{O}$), if D is principal (and X is normal, a standing hypothesis), then $\mathcal{O}(D) \cong \mathcal{O}$. (Diagram (15.2.7.1) will imply that the converse holds: if $\mathcal{O}(D) \cong \mathcal{O}$, then D is principal.) Thus if D is *locally* principal, $\mathcal{O}_X(D)$ is *locally* isomorphic to \mathcal{O}_X — $\mathcal{O}_X(D)$ is an invertible sheaf.

15.2.G. IMPORTANT EXERCISE. Suppose $\mathcal{O}_X(D)$ is an invertible sheaf.

- (a) Show that $\operatorname{div}(\sigma(1)) = D$, where σ was defined in Exercise 15.2.E(b).
 (b) Show the converse to Observation 15.2.5: show that D is locally principal.

15.2.6. Remark. In definition (15.2.2.1), it may seem cleaner to consider those s such that $\operatorname{div} s \geq D|_U$. The reason for the convention comes from our desire that $\operatorname{div} \sigma(1) = D$. (Taking the “opposite” convention would yield the dual bundle, in the case where D is locally principal.)

15.2.H. LESS IMPORTANT EXERCISE: A WEIL DIVISOR THAT IS NOT LOCALLY PRINCIPAL. Let $X = \operatorname{Spec} k[x, y, z]/(xy - z^2)$, a cone, and let D be the line $z = x = 0$ (see Figure 13.1).

- (a) Show that D is not locally principal. (Hint: consider the stalk at the origin. Use the Zariski tangent space, see Problem 13.1.3.) In particular $\mathcal{O}_X(D)$ is not an invertible sheaf.
 (b) Show that $\operatorname{div}(x) = 2D$. This corresponds to the fact that the plane $x = 0$ is tangent to the cone X along D .

15.2.I. IMPORTANT EXERCISE. If X is Noetherian and factorial, show that for any Weil divisor D , $\mathcal{O}(D)$ is an invertible sheaf. (Hint: It suffices to deal with the case where D is irreducible, say $D = [Y]$, and to cover X by open sets so that on each open set U there is a function whose divisor is $[Y \cap U]$. One open set will be $X - Y$. Next, we find an open set U containing an arbitrary $p \in Y$, and a function on U . As $\mathcal{O}_{X,p}$ is a unique factorization domain, the prime corresponding to Y is codimension 1 and hence principal by Lemma 12.1.7. Let f be a generator of this prime ideal, interpreted as an element of $K(X)$. It is regular at p , it has a finite number of zeros and poles, and through p , $[Y]$ is the “only zero” (the only component of the divisor of zeros). Let U be X minus all the other zeros and poles.)

15.2.7. The class group. We can now get a handle on the Picard group of a normal Noetherian scheme. Define the **class group** of X , $\text{Cl } X$, by $\text{Weil } X / \text{Prin } X$. By taking the quotient of the inclusion (15.2.0.1) by $\text{Prin } X$, we have the inclusion $\text{Pic } X \hookrightarrow \text{Cl } X$. This is summarized in the convenient and enlightening diagram.

$$(15.2.7.1) \quad \begin{array}{ccccc} & & (\mathcal{O}(D), \sigma(1)) \hookrightarrow D & & \\ & \swarrow \sim & \downarrow \text{div} & \searrow & \\ \{(\mathcal{L}, s)\}/\text{iso.} & \xrightarrow{\sim} & \text{LocPrin } X & \xrightarrow{\quad} & \text{Weil } X \\ \downarrow & & \downarrow / \text{Prin } X & & \downarrow / \text{Prin } X \\ \text{Pic } X \equiv \{\mathcal{L}\}/\text{iso.} & \xrightarrow{\sim} & \text{LocPrin } X / \text{Prin } X & \xrightarrow{\quad} & \text{Cl } X \\ & \nwarrow \sim & \uparrow \mathcal{O}(D) \hookrightarrow D & & \end{array}$$

This diagram is very important, and although it is short to state, it takes time to digest.

In particular, if A is a unique factorization domain, then all Weil divisors on $\text{Spec } A$ are principal by Lemma 12.1.7, so $\text{Cl Spec } A = 0$, and hence $\text{Pic Spec } A = 0$.

15.2.8. As $k[x_1, \dots, x_n]$ has unique factorization, $\text{Cl}(\mathbb{A}_k^n) = 0$, so $\boxed{\text{Pic}(\mathbb{A}_k^n) = 0}$. Geometers might find this believable — “ \mathbb{C}^n is a contractible complex manifold, and hence should have no nontrivial line bundles” — even if some caution is in order, as the kinds of line bundles being considered are entirely different: holomorphic vs. topological or C^∞ . (Aside: for this reason, you might expect that \mathbb{A}_k^n also has no nontrivial vector bundles. This is the Quillen-Suslin Theorem, formerly known as Serre’s Conjecture, part of Quillen’s work leading to his 1978 Fields Medal. The case $n = 1$ was Exercise 14.2.M. For a short proof by Vaserstein, see [Lan, p. 850].)

Removing a closed subset of X of codimension greater than 1 doesn’t change the class group, as it doesn’t change the Weil divisor group or the principal divisors. (Warning: it *can* affect the Picard group, see Exercise 15.2.Q.)

Removing a subset of codimension 1 changes the Weil divisor group in a controllable way. Suppose Z is an *irreducible* codimension 1 subset of X . Then we clearly have an exact sequence:

$$0 \longrightarrow \mathbb{Z} \xrightarrow{1 \mapsto [Z]} \text{Weil } X \longrightarrow \text{Weil}(X - Z) \longrightarrow 0.$$

When we take the quotient by principal divisors, taking into account the fact that we may lose exactness on the left, we get an **excision exact sequence for class**

groups:

$$(15.2.8.1) \quad \mathbb{Z} \xrightarrow{1 \mapsto [Z]} \mathrm{Cl} X \longrightarrow \mathrm{Cl}(X - Z) \longrightarrow 0.$$

(Do you see why?)

For example, if U is an open subscheme of $X = \mathbb{A}^n$, $\mathrm{Pic} U = \{0\}$.

As another application, let $X = \mathbb{P}_k^n$, and Z be the hyperplane $x_0 = 0$. We have

$$\mathbb{Z} \longrightarrow \mathrm{Cl} \mathbb{P}_k^n \longrightarrow \mathrm{Cl} \mathbb{A}_k^n \longrightarrow 0$$

from which $\mathrm{Cl} \mathbb{P}_k^n$ is generated by the class $[Z]$, and $\mathrm{Pic} \mathbb{P}_k^n$ is a subgroup of this.

15.2.9. By Exercise 15.2.F, $[Z] \mapsto \mathcal{O}(1)$, and as $\mathcal{O}(n)$ is nontrivial for $n \neq 0$ (Exercise 15.1.B), $[Z]$ is not torsion in $\mathrm{Cl} \mathbb{P}_k^n$. Hence $\mathrm{Pic}(\mathbb{P}_k^n) \hookrightarrow \mathrm{Cl}(\mathbb{P}_k^n)$ is an isomorphism, and $\boxed{\mathrm{Pic}(\mathbb{P}_k^n) \cong \mathbb{Z}}$, with generator $\mathcal{O}(1)$. The **degree** of an invertible sheaf on \mathbb{P}^n is defined using this: define $\deg \mathcal{O}(d)$ to be d .

We have gotten good mileage from the fact that the Picard group of the spectrum of a unique factorization domain is trivial. More generally, Exercise 15.2.I gives us:

15.2.10. Proposition. — *If X is Noetherian and factorial, then for any Weil divisor D , $\mathcal{O}(D)$ is invertible, and hence the map $\mathrm{Pic} X \rightarrow \mathrm{Cl} X$ is an isomorphism.*

This can be used to make the connection to the class group in number theory precise, see Exercise 14.2.J; see also §15.2.13.

15.2.11. Mild but important generalization: twisting line bundles by divisors. The above constructions can be extended, with \mathcal{O}_X replaced by an arbitrary invertible sheaf, as follows. Let \mathcal{L} be an invertible sheaf on a normal Noetherian scheme X . Then define $\mathcal{L}(D)$ by $\mathcal{O}_X(D) \otimes \mathcal{L}$.

15.2.J. EASY EXERCISE.

(a) Assume for convenience that X is irreducible. Show that sections of $\mathcal{L}(D)$ can be interpreted as rational sections of \mathcal{L} with zeros and poles constrained by D , just as in (15.2.2.1):

$$\Gamma(U, \mathcal{L}(D)) := \{t \text{ nonzero rational section of } \mathcal{L} : \mathrm{div}|_U t + D|_U \geq 0\} \cup \{0\}.$$

(b) Suppose D_1 and D_2 are locally principal. Show that

$$(\mathcal{O}(D_1))(D_2) \cong \mathcal{O}(D_1 + D_2).$$

15.2.12. Fun examples.

We can now actually calculate some Picard and class groups. First, a useful observation: notice that you can restrict invertible sheaves on Y to any subscheme X , and this can be a handy way of checking that an invertible sheaf is not trivial. Effective Cartier divisors (§9.5.1) sometimes restrict too: if you have an effective Cartier divisor on Y , then it restricts to a closed subscheme on X , locally cut out scheme-theoretically by one equation. If you are fortunate and this equation doesn't vanish on any associated point of X (§6.5.2), then you get an effective Cartier divisor on X . You can check that the restriction of effective Cartier divisors corresponds to restriction of invertible sheaves (in the sense of Exercise 14.2.H).

15.2.K. EXERCISE: A TORSION PICARD GROUP. Suppose that Y is a hypersurface in \mathbb{P}_k^n corresponding to an irreducible degree d polynomial. Show that $\text{Pic}(\mathbb{P}_k^n - Y) \cong \mathbb{Z}/(d)$. (For differential geometers: this is related to the fact that $\pi_1(\mathbb{P}_k^n - Y) \cong \mathbb{Z}/(d)$.) Hint: (15.2.8.1).

The next two exercises explore consequences of Exercise 15.2.K, and provide us with some examples promised in Exercise 5.4.N.

15.2.L. EXERCISE (GENERALIZING EXERCISE 5.4.N). Keeping the same notation, assume $d > 1$ (so $\text{Pic}(\mathbb{P}^n - Y) \neq 0$), and let H_0, \dots, H_n be the $n + 1$ coordinate hyperplanes on \mathbb{P}^n . Show that $\mathbb{P}^n - Y$ is affine, and $\mathbb{P}^n - Y - H_i$ is a distinguished open subset of it. Show that the $\mathbb{P}^n - Y - H_i$ form an open cover of $\mathbb{P}^n - Y$. Show that $\text{Pic}(\mathbb{P}^n - Y - H_i) = 0$. Then by Exercise 15.2.T, each $\mathbb{P}^n - Y - H_i$ is the Spec of a unique factorization domain, but $\mathbb{P}^n - Y$ is not. Thus the property of being a unique factorization domain is not an affine-local property — it satisfies only one of the two hypotheses of the Affine Communication Lemma 5.3.2.

15.2.M. EXERCISE. Keeping the same notation as the previous exercise, show that on $\mathbb{P}^n - Y$, H_i (restricted to this open set) is an effective Cartier divisor that is not cut out by a single equation. (Hint: Otherwise it would give a trivial element of the class group.)

15.2.N. EXERCISE. Show that $A := \mathbb{R}[x, y]/(x^2 + y^2 - 1)$ is not a unique factorization domain, but $A \otimes_{\mathbb{R}} \mathbb{C}$ is. Hint: Exercise 15.2.L.

15.2.O. EXERCISE: PICARD GROUP OF $\mathbb{P}^1 \times \mathbb{P}^1$. Consider

$$X = \mathbb{P}_k^1 \times_k \mathbb{P}_k^1 \cong \text{Proj } k[w, x, y, z]/(wz - xy),$$

a smooth quadric surface (see Figure 9.2, and Example 10.6.2). Show that $\text{Pic } X \cong \mathbb{Z} \oplus \mathbb{Z}$ as follows: Show that if $L = \{\infty\} \times_k \mathbb{P}^1 \subset X$ and $M = \mathbb{P}^1 \times_k \{\infty\} \subset X$, then $X - L - M \cong \mathbb{A}^2$. This will give you a surjection $\mathbb{Z} \oplus \mathbb{Z} \rightarrow \text{Cl } X$. Show that $\mathcal{O}(L)$ restricts to \mathcal{O} on L and $\mathcal{O}(1)$ on M . Show that $\mathcal{O}(M)$ restricts to \mathcal{O} on M and $\mathcal{O}(1)$ on L . (This exercise takes some time, but is enlightening.)

15.2.P. EXERCISE. Show that irreducible smooth projective surfaces (over k) can be birational but not isomorphic. Hint: show \mathbb{P}^2 is not isomorphic to $\mathbb{P}^1 \times \mathbb{P}^1$ using the Picard group. (Aside: we will see in Exercise 21.2.D that the Picard group of the “blown up plane” is \mathbb{Z}^2 , but in Exercise 21.2.E we will see that the blown up plane is not isomorphic to $\mathbb{P}^1 \times \mathbb{P}^1$, using a little more information in the Picard group.)

This is unlike the case for curves: birational irreducible smooth projective curves (over k) must be isomorphic, as we will see in Theorem 18.4.3. Nonetheless, any two surfaces are related in a simple way: if X and X' are projective, regular, and birational, then X can be sequentially blown up at judiciously chosen points, and X' can too, such that the two results are isomorphic (see [Ha1, Thm. V.5.5]; blowing up will be discussed in Chapter 23).

15.2.Q. EXERCISE: PICARD GROUP OF THE CONE. Let $X = \text{Spec } k[x, y, z]/(xy - z^2)$, a cone, where $\text{char } k \neq 2$. (The characteristic hypothesis is not necessary for the result, but is included so you can use Exercise 5.4.H to show normality of X .) Show

that $\text{Pic } X = 0$, and $\text{Cl } X \cong \mathbb{Z}/2$. Hint: show that the class of $Z = \{x = z = 0\}$ (the “affine cone over a line”) generates $\text{Cl } X$ by showing that its complement $D(x)$ is isomorphic to an open subset of \mathbb{A}_k^2 . Show that $2[Z] = \text{div}(x)$ and hence principal, and that Z is not principal, Exercise 15.2.H. (Remark: You know enough to show that $X - \{(0, 0, 0)\}$ is factorial. So although the class group is insensitive to removing loci of codimension greater than 1, §15.2.8, this is not true of the Picard group.)

A Weil divisor (on a normal scheme) with a nonzero multiple corresponding to a line bundle is called **\mathbb{Q} -Cartier**. (We won’t use this terminology beyond the next exercise.) Exercise 15.2.Q gives an example of a Weil divisor that does not correspond to a line bundle, but is nonetheless \mathbb{Q} -Cartier. We now give an example of a Weil divisor that is *not* \mathbb{Q} -Cartier.

15.2.R. EXERCISE (A NON- \mathbb{Q} -CARTIER DIVISOR). On the cone over the smooth quadric surface $X = \text{Spec } k[w, x, y, z]/(wz - xy)$, let Z be the Weil divisor cut out by $w = x = 0$. Exercise 13.1.D showed that Z is not cut out scheme-theoretically by a single equation. Show more: that if $n \neq 0$, then $n[Z]$ is not locally principal. Hint: show that the complement of an effective Cartier divisor on an affine scheme is also affine, using Proposition 8.3.4. Then if some multiple of Z were locally principal, then the closed subscheme of the complement of Z cut out by $y = z = 0$ would be affine — any closed subscheme of an affine scheme is affine. But this is the scheme $y = z = 0$ (also known as the wx -plane) minus the point $w = x = 0$, which we have seen is non-affine, §4.4.1.

15.2.S. ★ EXERCISE (FOR THOSE WITH SUFFICIENT ARITHMETIC BACKGROUND). Identify the (ideal) class group of the ring of integers \mathcal{O}_K in a number field K , as defined in Exercise 14.2.J, with the class group of $\text{Spec } \mathcal{O}_K$, as defined in this section. In particular, you will recover the common description of the class group as formal sums of prime ideals, modulo an equivalence relation coming from principal fractional ideals.

15.2.13. More on class groups and unique factorization.

As mentioned in §5.4.5, there are few commonly used means of checking that a ring is a unique factorization domain. The next exercise is one of them, and it is useful. For example, it implies the classical fact that for rings of integers in number fields, the class group is the obstruction to unique factorization (see Exercise 14.2.J and Proposition 15.2.10).

15.2.T. EXERCISE. Suppose that A is a Noetherian integral domain. Show that A is a unique factorization domain if and only if A is integrally closed and $\text{Cl Spec } A = 0$. (One direction is easy: we have already shown that unique factorization domains are integrally closed in their fraction fields. Also, Lemma 12.1.7 shows that all codimension 1 prime ideals of a unique factorization domain are principal, so that implies that $\text{Cl Spec } A = 0$. It remains to show that if A is integrally closed and $\text{Cl Spec } A = 0$, then all codimension 1 prime ideals are principal, as this characterizes unique factorization domains (Proposition 12.3.5). Algebraic Hartogs’s Lemma 12.3.15, may arise in your argument.) This is the third important characterization of unique factorization domains promised in §5.4.5.

My final favorite method of checking that a ring is a unique factorization domain (§5.4.5) is Nagata’s Lemma. It is also the least useful.

15.2.U. ★★ EXERCISE (NAGATA’S LEMMA). Suppose A is a Noetherian domain, $x \in A$ an element such that (x) is prime and $A_x = A[1/x]$ is a unique factorization domain. Then A is a unique factorization domain. (Hint: Exercise 15.2.T. Use the short exact sequence

$$\mathbb{Z}[(x)] \rightarrow \text{Cl Spec } A \rightarrow \text{Cl Spec } A_x \rightarrow 0$$

(15.2.8.1) to show that $\text{Cl Spec } A = 0$. Prove that $A[1/x]$ is integrally closed, then show that A is integrally closed as follows. Suppose $T^n + a_{n-1}T^{n-1} + \cdots + a_0 = 0$, where $a_i \in A$, and $T \in K(A)$. Then by integral closedness of A_x , we have that $T = r/x^m$, where if $m > 0$, then $r \notin (x)$. Then we quickly get a contradiction if $m > 0$.)

This leads to a fun algebra fact promised in Remark 13.8.6. Suppose k is an algebraically closed field of characteristic not 2. Let $A = k[x_1, \dots, x_n]/(x_1^2 + \cdots + x_m^2)$ where $m \leq n$. When $m \leq 2$, we get some special behavior. (If $m = 0$, we get affine space; if $m = 1$, we get a nonreduced scheme; if $m = 2$, we get a reducible scheme that is the union of two affine spaces.)

If $m \geq 3$, we have verified that $\text{Spec } A$ is normal, in Exercise 5.4.I(b). In fact, if $m \geq 3$, then A is a unique factorization domain *unless* $m = 3$ or $m = 4$ (Exercise 5.4.L; see also Exercise 13.1.E). For the case $m = 3$:

$$A = k[x, y, z, w_1, \dots, w_{n-3}]/(x^2 + y^2 - z^2)$$

is not a unique factorization domain, as it is has nonzero class group (by essentially the same argument as for Exercise 15.2.Q).

The failure at 4 comes from the geometry of the quadric surface: we have checked that in $\text{Spec } k[w, x, y, z]/(wz - xy)$, there is a codimension 1 irreducible subset — the cone over a line in a ruling — that is not principal.

15.2.V. EXERCISE (THE CASE $m \geq 5$). Suppose that k is algebraically closed of characteristic not 2. Show that if $\ell \geq 3$, then

$$A = k[a, b, x_1, \dots, x_n]/(ab - x_1^2 - \cdots - x_\ell^2)$$

is a unique factorization domain, by using Nagata’s Lemma with $x = a$.

15.3 ★ Effective Cartier divisors “=” invertible ideal sheaves

We now give a different means of describing invertible sheaves on a scheme. One advantage of this over Weil divisors is that it can give line bundles on everywhere nonreduced schemes (such a scheme can’t be regular at any codimension 1 prime). But we won’t use this, so it is less important.

Suppose $D \hookrightarrow X$ is a closed subscheme such that the corresponding ideal sheaf \mathcal{I} is an invertible sheaf. Then \mathcal{I} is locally trivial; suppose U is a trivializing affine open set $\text{Spec } A$. Then the closed subscheme exact sequence (9.1.2.1)

$$0 \rightarrow \mathcal{I} \rightarrow \mathcal{O}_X \rightarrow \mathcal{O}_D \rightarrow 0$$

corresponds to

$$0 \rightarrow I \rightarrow A \rightarrow A/I \rightarrow 0$$

with $I \cong A$ as A -modules. Thus I is generated by a single element, say a , and this exact sequence starts as

$$0 \longrightarrow A \xrightarrow{\times a} A$$

As multiplication by a is injective, a is not a zerodivisor. We conclude that D is locally cut out by a single equation, that is not a zerodivisor. This was the definition of *effective Cartier divisor* given in §9.5.1. This argument is clearly reversible, so we have a quick new definition of effective Cartier divisor (an ideal sheaf \mathcal{I} that is an invertible sheaf — or equivalently, the corresponding closed subscheme).

15.3.A. EASY EXERCISE. Show that a is unique up to multiplication by an invertible function.

In the case where X is locally Noetherian, we can use the language of associated points (§6.5.2), so we can restate this definition as: D is locally cut out by a single equation, not vanishing at any associated point of X .

We now define an invertible sheaf corresponding to D . The seemingly obvious definition would be to take \mathcal{I}_D , but instead we define the invertible sheaf $\mathcal{O}(D)$ corresponding to an effective Cartier divisor to be the *dual*: \mathcal{I}_D^\vee . (The reason for the dual is Exercise 15.3.C.)

15.3.B. EXERCISE. Check that this agrees with our earlier definition of $\mathcal{O}(D)$, Important Definition 15.2.2.

The ideal sheaf \mathcal{I}_D is then $\mathcal{O}(-D)$, so the closed subscheme exact sequence (9.1.2.1) is

$$0 \rightarrow \mathcal{O}(-D) \rightarrow \mathcal{O} \rightarrow \mathcal{O}_D \rightarrow 0.$$

The invertible sheaf $\mathcal{O}(D)$ has a canonical section s_D : Tensoring $0 \rightarrow \mathcal{I} \rightarrow \mathcal{O}$ with \mathcal{I}^\vee gives us $\mathcal{O} \rightarrow \mathcal{I}^\vee$. (Easy unimportant fact: instead of tensoring $\mathcal{I} \rightarrow \mathcal{O}$ with \mathcal{I}^\vee , we could have dualized $\mathcal{I} \rightarrow \mathcal{O}$, and we would get the same section.)

15.3.C. IMPORTANT AND SURPRISINGLY TRICKY EXERCISE. Recall that a section of a locally free sheaf on X cuts out a closed subscheme of X (Exercise 14.2.N). Show that this section s_D cuts out D . (Compare this to Remark 15.2.6.)

This construction is reversible:

15.3.D. EXERCISE. Suppose \mathcal{L} is an invertible sheaf, and s is a section that is locally not a zerodivisor. (Make sense of this! In particular, if X is locally Noetherian, this means “ s does not vanish at an associated point of X ”, see §6.5.2.) Show that $s = 0$ cuts out an effective Cartier divisor D , and $\mathcal{O}(D) \cong \mathcal{L}$.

15.3.E. EXERCISE. Suppose \mathcal{I} and \mathcal{J} are invertible ideal sheaves (hence corresponding to effective Cartier divisors, say D and D' respectively). Show that $\mathcal{I}\mathcal{J}$ is an invertible ideal sheaf. (We define the **product of two quasicoherent ideal sheaves** $\mathcal{I}\mathcal{J}$ as you might expect: on each affine, we take the product of the two corresponding ideals. To make sure this is well-defined, we need only check that if A is a ring, and $f \in A$, and $I, J \subset A$ are two ideals, then $(IJ)_f = I_f J_f$ in A_f .) We define the corresponding Cartier divisor to be $D + D'$. Verify that $\mathcal{O}(D + D') \cong \mathcal{O}(D) \otimes \mathcal{O}(D')$.

We thus have an important correspondence between *effective Cartier divisors* (closed subschemes whose ideal sheaves are invertible, or equivalently locally cut out by one non-zerodivisor, or in the locally Noetherian case, locally cut out by one equation not vanishing at an associated point) and *ordered pairs* (\mathcal{L}, s) where \mathcal{L} is an invertible sheaf, and s is a section that is not locally a zerodivisor (or in the locally Noetherian case, not vanishing at an associated point). The effective Cartier divisors form an abelian semigroup. We have a map of semigroups, from effective Cartier divisors to invertible sheaves with sections not locally zerodivisors (and hence also to the Picard group of invertible sheaves).

We get lots of invertible sheaves, by taking differences of two effective Cartier divisors. In fact we “usually get them all” — it is very hard to describe an invertible sheaf on a finite type k -scheme that is not describable in such a way. For example, there are none if the scheme is regular or even factorial (basically by Proposition 15.2.10 for factoriality; and regular schemes are factorial by the Auslander-Buchsbaum Theorem 13.8.5). Exercise 17.6.F will imply that there are none if the scheme is projective. It holds in all other reasonable circumstances, see [Gr-EGA, IV₄.21.3.4]. However, it does not always hold; the first and best example is due to Kleiman, see [K13].

Quasicoherent sheaves in projective geometry

The first two sections of this chapter are relatively straightforward, and the last two are trickier.

16.1 The quasicoherent sheaf corresponding to a graded module

We now describe quasicoherent sheaves on a projective A -scheme. Recall that a projective A -scheme is produced from the data of a $\mathbb{Z}^{\geq 0}$ -graded ring S_{\bullet} , with $S_0 = A$, and S_+ is a finitely generated ideal (a “finitely generated graded ring over A ”, §4.5.6). The resulting scheme is denoted $\text{Proj } S_{\bullet}$.

Suppose M_{\bullet} is a graded S_{\bullet} -module, *graded by \mathbb{Z}* . (While reading the next section, you may wonder why we don’t grade by $\mathbb{Z}^{\geq 0}$. You will see that it doesn’t matter. A \mathbb{Z} -grading will make things cleaner when we produce an M_{\bullet} from a quasicoherent sheaf on $\text{Proj } S_{\bullet}$.) We define the quasicoherent sheaf \widetilde{M}_{\bullet} as follows. (I will avoid calling it \widetilde{M} , as this might cause confusion with the affine case; but \widetilde{M}_{\bullet} is *not* graded in any way.) For each homogeneous f of *positive degree*, we define a quasicoherent sheaf $\widetilde{M}_{\bullet}(f)$ on the distinguished open $D(f) = \{p : f(p) \neq 0\}$ by $((M_{\bullet})_f)_0$ — note that $((M_{\bullet})_f)_0$ is an $((S_{\bullet})_f)_0$ -module, and recall that $D(f)$ is identified with $\text{Spec}((S_{\bullet})_f)_0$ (Exercise 4.5.I). As in (4.5.7.1), the subscript 0 means “the 0-graded piece”. We have obvious isomorphisms of the restriction of $\widetilde{M}_{\bullet}(f)$ and $\widetilde{M}_{\bullet}(g)$ to $D(fg)$, satisfying the cocycle conditions. (Think through this yourself, to be sure you agree with the word “obvious”!) By Exercise 2.5.E, these sheaves glue together to a single sheaf \widetilde{M}_{\bullet} on $\text{Proj } S_{\bullet}$. We then discard the temporary notation $\widetilde{M}_{\bullet}(f)$.

The \mathcal{O} -module \widetilde{M}_{\bullet} is clearly quasicoherent, because it is quasicoherent on each $D(f)$, and quasicoherence is local.

16.1.A. EXERCISE. Give an isomorphism between the stalk of \widetilde{M}_{\bullet} at a point corresponding to homogeneous prime $\mathfrak{p} \subset S_{\bullet}$ and $((M_{\bullet})_{\mathfrak{p}})_0$. (Remark: You can use this exercise to give an alternate definition of \widetilde{M}_{\bullet} in terms of “compatible stalks”, cf. Exercise 4.5.P.)

Given a map of graded modules $\phi: M_{\bullet} \rightarrow N_{\bullet}$, we get an induced map of sheaves $\widetilde{M}_{\bullet} \rightarrow \widetilde{N}_{\bullet}$. Explicitly, over $D(f)$, the map $M_{\bullet} \rightarrow N_{\bullet}$ induces $(M_{\bullet})_f \rightarrow (N_{\bullet})_f$, which induces $\phi_f: ((M_{\bullet})_f)_0 \rightarrow ((N_{\bullet})_f)_0$; and this behaves well with respect to restriction to smaller distinguished open sets, i.e., the following diagram

commutes.

$$\begin{array}{ccc} ((M_\bullet)_f)_0 & \xrightarrow{\phi_f} & ((N_\bullet)_f)_0 \\ \downarrow & & \downarrow \\ ((M_\bullet)_{fg})_0 & \xrightarrow{\phi_{fg}} & ((N_\bullet)_{fg})_0. \end{array}$$

Thus \sim is a functor from the category of graded S_\bullet -modules to the category of quasicoherent sheaves on $\text{Proj } S_\bullet$.

16.1.B. EASY EXERCISE. Show that \sim is an exact functor. (Hint: everything in the construction is exact.)

16.1.C. EXERCISE. Show that if M_\bullet and M'_\bullet agree in high enough degrees, then $\widetilde{M_\bullet} \cong \widetilde{M'_\bullet}$. Then show that the map from graded S_\bullet -modules (up to isomorphism) to quasicoherent sheaves on $\text{Proj } S_\bullet$ (up to isomorphism) is not a bijection. (Really: show this is not an equivalence of categories.)

Exercise 16.1.C shows that \sim isn't an isomorphism (or equivalence) of categories, but it is close. The relationship is akin to that between presheaves and sheaves, and the sheafification functor (see §16.4).

16.1.D. EXERCISE. Describe a map of S_0 -modules $M_0 \rightarrow \Gamma(\text{Proj } S_\bullet, \widetilde{M_\bullet})$. (This foreshadows the “saturation map” of §16.4.5 that takes a graded module to its saturation, see Exercise 16.4.C.)

16.1.1. Example: Homogeneous ideals of S_\bullet give closed subschemes of $\text{Proj } S_\bullet$. Recall that a homogeneous ideal $I_\bullet \subset S_\bullet$ yields a closed subscheme $\text{Proj } S_\bullet/I_\bullet \hookrightarrow \text{Proj } S_\bullet$. For example, suppose $S_\bullet = k[w, x, y, z]$, so $\text{Proj } S_\bullet \cong \mathbb{P}^3$. The ideal $I_\bullet = (wz - xy, x^2 - wy, y^2 - xz)$ yields our old friend, the twisted cubic (defined in Exercise 9.3.A)

16.1.E. EXERCISE. Show that if the functor \sim is applied to the exact sequence of graded S_\bullet -modules

$$0 \rightarrow I_\bullet \rightarrow S_\bullet \rightarrow S_\bullet/I_\bullet \rightarrow 0$$

we obtain the closed subscheme exact sequence (9.1.2.1) for $\text{Proj } S_\bullet/I_\bullet \hookrightarrow \text{Proj } S_\bullet$.

All closed subschemes of $\text{Proj } S_\bullet$ arise in this way (see Exercise 16.4.H).

16.1.2. Remark. If M_\bullet is finitely generated (resp. finitely presented, coherent), then so is $\widetilde{M_\bullet}$, [Roh, (4.4.11)]. We will not need this fact.

16.2 Line bundles on projective A -schemes

Suppose M_\bullet is a graded S_\bullet -module. Define the graded module $M(n)_\bullet$ by $M(n)_m := M_{n+m}$. Thus the quasicoherent sheaf $\widetilde{M(n)_\bullet}$ satisfies

$$\Gamma(D(f), \widetilde{M(n)_\bullet}) = ((M_\bullet)_f)_n$$

where here the subscript means we take the n th graded piece. (These subscripts are admittedly confusing!)

16.2.A. EXERCISE. If $S_\bullet = A[x_0, \dots, x_m]$, so $\text{Proj } S_\bullet = \mathbb{P}_A^m$, show $\widetilde{S(n)_\bullet} \cong \mathcal{O}(n)$ using transition functions (cf. §15.1). (Recall from §15.1.2 that the global sections of $\mathcal{O}(n)$ should be identified with the homogeneous degree n polynomials in x_0, \dots, x_m . Can you see that in the context of this exercise?)

16.2.1. Definition. Motivated by this, if S_\bullet is a graded ring generated in degree 1, we define $\mathcal{O}_{\text{Proj } S_\bullet}(n)$ (or simply $\mathcal{O}(n)$, where S_\bullet is implicit) by $\widetilde{S(n)_\bullet}$.

16.2.B. IMPORTANT EXERCISE. If S_\bullet is generated in degree 1, show that $\mathcal{O}_{\text{Proj } S_\bullet}(n)$ is an invertible sheaf.

If S_\bullet is generated in degree 1, and \mathcal{F} is a quasicoherent sheaf on $\text{Proj } S_\bullet$, define $\mathcal{F}(n) := \mathcal{F} \otimes \mathcal{O}(n)$. This is often called **twisting \mathcal{F} by $\mathcal{O}(n)$** , or **twisting \mathcal{F} by n** . More generally, if \mathcal{L} is an invertible sheaf, then $\mathcal{F} \otimes \mathcal{L}$ is often called **twisting \mathcal{F} by \mathcal{L}** .

16.2.C. EXERCISE. If S_\bullet is generated in degree 1, show that $\widetilde{M_\bullet(n)} \cong \widetilde{M(n)_\bullet}$. (Hereafter, we can be cavalier with the placement of the “dot” in such situations.)

16.2.D. EXERCISE. If S_\bullet is generated in degree 1, show that $\mathcal{O}(m+n) \cong \mathcal{O}(m) \otimes \mathcal{O}(n)$.

16.2.2. Unimportant remark. Even if S_\bullet is not generated in degree 1, then by Exercise 7.4.G, $S_{d\bullet}$ is generated in degree 1 for some d . In this case, we may define the invertible sheaves $\mathcal{O}(dn)$ for $n \in \mathbb{Z}$. This does *not* mean that we *can't* define $\mathcal{O}(1)$; this depends on S_\bullet . For example, if S_\bullet is the polynomial ring $k[x, y]$ with the usual grading, except without linear terms (so $S_\bullet = k[x^2, xy, y^2, x^3, x^2y, xy^2, y^3]$), then $S_{2\bullet}$ and $S_{3\bullet}$ are both generated in degree 1, meaning that we may define $\mathcal{O}(2)$ and $\mathcal{O}(3)$. There is good reason to call their “difference” $\mathcal{O}(1)$.

16.3 Globally generated and base-point-free line bundles

16.3.A. IMPORTANT EXERCISE (TO BE USED REPEATEDLY). Generalize Exercise 6.2.G as follows. Suppose X is a quasicompact quasiseparated scheme, \mathcal{L} is an invertible sheaf on X with section s , and \mathcal{F} is a quasicoherent sheaf on X . As in Exercise 6.2.G, let X_s be the open subset of X where s doesn't vanish. Show that any section of \mathcal{F} over X_s can be interpreted as the quotient of a global section of $\mathcal{F} \otimes_{\mathcal{O}_X} \mathcal{L}^{\otimes n}$ by s^n . In other words, any section of \mathcal{F} over X_s can be extended over all of X , once you multiply it by a large enough power of s . More precisely: note that $\bigoplus_{n \geq 0} \Gamma(X, \mathcal{L}^{\otimes n})$ is a graded ring, and we interpret s as a degree 1 element of it. Note also that $\bigoplus_{n \geq 0} \Gamma(X, \mathcal{F} \otimes_{\mathcal{O}_X} \mathcal{L}^{\otimes n})$ is a graded module over this ring. Describe a natural map

$$\left(\left(\bigoplus_{n \geq 0} \Gamma(X, \mathcal{F} \otimes_{\mathcal{O}_X} \mathcal{L}^{\otimes n}) \right)_s \right)_0 \longrightarrow \Gamma(X_s, \mathcal{F})$$

and show that it is an isomorphism. (Hint: after showing the existence of the natural map, show it is an isomorphism in the affine case.)

We now come to a topic that is harder, but that will be important. Throughout this section, S_\bullet will be a finitely generated graded ring over A , generated in degree 1. We will prove the following result.

16.3.1. Theorem. — *Any finite type quasicoherent sheaf \mathcal{F} on $\text{Proj } S_\bullet$ can be presented in the form*

$$\bigoplus_{\text{finite}} \mathcal{O}(-n) \rightarrow \mathcal{F} \rightarrow 0.$$

Because we can work with the line bundles $\mathcal{O}(-n)$ in a hands-on way, this result will give us great control over *all* coherent sheaves (and in particular, vector bundles) on $\text{Proj } S_\bullet$. As just a first example, it will allow us to show that every coherent sheaf on a projective k -scheme has a finite-dimensional space of global sections (Corollary 19.1.4). (This fact will grow up to be the fact that the higher pushforwards of coherent sheaves under proper morphisms are also coherent, see Theorem 19.8.1(d) and Grothendieck's Coherence Theorem 19.9.1.)

Rather than proceeding directly to a proof, we use this as an excuse to introduce notions that are useful in wider circumstances (*global generation*, *base-point-freeness*, *ampleness*), and their interrelationships. But first we use it as an excuse to mention an important classical result.

16.3.2. The Hilbert Syzygy Theorem.

Given any coherent sheaf \mathcal{F} on \mathbb{P}_k^n , Theorem 16.3.1 gives the existence of a surjection $\phi: \bigoplus_{\text{finite}} \mathcal{O}(-m) \rightarrow \mathcal{F} \rightarrow 0$. The kernel of the surjection is also coherent, so iterating this construction, we can construct an infinite resolution of \mathcal{F} by a direct sum of line bundles:

$$\cdots \bigoplus_{\text{finite}} \mathcal{O}(m_{2,j}) \longrightarrow \bigoplus_{\text{finite}} \mathcal{O}(m_{1,j}) \longrightarrow \bigoplus_{\text{finite}} \mathcal{O}(m_{0,j}) \longrightarrow \mathcal{F} \longrightarrow 0.$$

The Hilbert Syzygy Theorem states that there is in fact a *finite* resolution, of length at most n . Because we won't use this, we don't give a proof, but [E] (especially [E, Thm. 1.13] and the links thereafter) has an excellent discussion. (Hilbert's original statement, and the discussion in [E], is about graded modules. The connection to the statement here on coherent sheaves is provided by the exactness of the functor \sim , Exercise 16.1.B.) See the comments after Theorem 3.6.17 for the original history of this result.

16.3.3. Globally generated sheaves. Suppose X is a scheme, and \mathcal{F} is an \mathcal{O} -module. The most important definition of this section is the following: \mathcal{F} is **globally generated** (or **generated by global sections**) if it admits a surjection from a free sheaf on X :

$$\mathcal{O}^{\oplus I} \twoheadrightarrow \mathcal{F}.$$

Here I is some index set. The global sections in question are the images of the $|I|$ sections corresponding to 1 in the various summands of $\mathcal{O}_X^{\oplus I}$; those images generate the stalks of \mathcal{F} . We say \mathcal{F} is **finitely globally generated** (or **generated by a finite number of global sections**) if the index set I can be taken to be finite.

More definitions in more detail: we say that \mathcal{F} is **globally generated at a point** p (or sometimes **generated by global sections at p**) if we can find $\phi: \mathcal{O}^{\oplus I} \rightarrow \mathcal{F}$

that is surjective on stalks at p :

$$\mathcal{O}_p^{\oplus I} \xrightarrow{\phi_p} \mathcal{F}_p.$$

(It would be more precise to say that the stalk of \mathcal{F} at p is generated by global sections of \mathcal{F} .) The key insight is that global generation at p means that every germ at p is a linear combination (over the local ring $\mathcal{O}_{X,p}$) of germs of global sections.

Note that \mathcal{F} is *globally generated* if it is globally generated at *all* points p . (Reason: Exercise 2.4.D showed that isomorphisms can be checked on the level of stalks. An easier version of the same argument shows that surjectivity can also be checked on the level of stalks.) Notice that we can take a single index set for all of X , by taking the union of all the index sets for each p .

16.3.B. EASY EXERCISE (REALITY CHECK). Show that every quasicoherent sheaf on every *affine* scheme is globally generated. Show that every finite type quasicoherent sheaf on every affine scheme is generated by a finite number of global sections. (Hint for both: for any A -module M , there is a surjection onto M from a free A -module.)

Clearly if \mathcal{F} and \mathcal{G} are globally generated, then so is $\mathcal{F} \oplus \mathcal{G}$ as well (and similarly for finitely globally generated, generated at a point p , etc.). Similarly for tensor product:

16.3.C. EASY EXERCISE (GLOBALLY GENERATED \otimes GLOBALLY GENERATED IS GLOBALLY GENERATED). Show that if quasicoherent sheaves \mathcal{F} and \mathcal{G} are globally generated at a point p , then so is $\mathcal{F} \otimes \mathcal{G}$.

16.3.D. EASY BUT IMPORTANT EXERCISE. Suppose \mathcal{F} is a finite type quasicoherent sheaf on X .

(a) Show that \mathcal{F} is globally generated at p if and only if “the fiber of \mathcal{F} is generated by global sections at p ”, i.e., the map from global sections to the fiber $\mathcal{F}_p/\mathfrak{m}_p\mathcal{F}_p$ is surjective, where \mathfrak{m} is the maximal ideal of $\mathcal{O}_{X,p}$. (Hint: Geometric Nakayama, Exercise 14.3.E.)

(b) Show that if \mathcal{F} is globally generated at p , then “ \mathcal{F} is globally generated near p ”: there is an open neighborhood U of p such that \mathcal{F} is globally generated at every point of U .

(c) Suppose further that X is a quasicompact scheme. Show that if \mathcal{F} is globally generated at all closed points of X , then \mathcal{F} is globally generated at all points of X . (Note that nonempty quasicompact schemes *have* closed points, Exercise 5.1.E.)

16.3.E. EASY EXERCISE. If \mathcal{F} is a finite type quasicoherent sheaf on X , and X is quasicompact, show that \mathcal{F} is globally generated if and only if it is generated by a *finite number* of global sections.

16.3.F. EASY EXERCISE. An invertible sheaf \mathcal{L} on X is globally generated if and only if for any point $p \in X$, there is a global section of \mathcal{L} not vanishing at p . (See Theorem 17.4.1 for why we care.)

16.3.4. Definitions. If \mathcal{L} is an invertible sheaf on X , then those points where all global sections of \mathcal{L} vanish are called the **base points** of \mathcal{L} , and the set of base

points is called the **base locus** of \mathcal{L} ; it is a closed subset of X . (We can refine this to a closed subscheme: by taking the scheme-theoretic intersection of the vanishing loci of the sections of \mathcal{L} , we obtain the **scheme-theoretic base locus**.) The complement of the **base locus** is the **base-point-free locus**. If \mathcal{L} has no base points, it is **base-point-free**. By the previous discussion, (i) the base-point-free locus is an open subset of X , and (ii) \mathcal{L} is generated by global sections if and only if it is base-point-free. By Exercise 16.3.C, the tensor product of two base-point-free line bundles is base-point-free.

(Remark: We will see in Exercise 19.2.I that if X is a k -scheme, and \mathcal{L} is an invertible sheaf on X , and K/k is any field extension, then \mathcal{L} is base-point-free if and only if it is “base-point-free after base change to K ”. You could reasonably prove this now.)

16.3.5. Base-point-free line bundles and maps to projective space. The main reason we care about the definitions above is the following. Recall Exercise 7.3.O(a), which shows that $n + 1$ functions on a scheme X with no common zeros yield a map to \mathbb{P}^n . This notion generalizes.

16.3.G. EASY EXERCISE (A VITALLY IMPORTANT CONSTRUCTION). Suppose s_0, \dots, s_n are $n + 1$ global sections of an invertible sheaf \mathcal{L} on a scheme X , with no common zero. Define a corresponding map to \mathbb{P}^n :

$$X \xrightarrow{[s_0, \dots, s_n]} \mathbb{P}^n$$

Hint: If U is an open subset on which \mathcal{L} is trivial, choose a trivialization, then translate the s_i into functions using this trivialization, and use Exercise 7.3.O(a) to obtain a morphism $U \rightarrow \mathbb{P}^n$. Then show that all of these maps (for different U and different trivializations) “agree”, using Exercise 7.3.O(b).

(In Theorem 17.4.1, we will see that this yields *all* maps to projective space.) Note that this exercise as written “works over \mathbb{Z} ” (as *all* morphisms are “over” the final object in the category of schemes), although many readers will just work over a particular base such as a given field k . Here is some convenient classical language which is used in this case.

16.3.6. Definitions. A **linear series** on a k -scheme X is a k -vector space V (usually finite-dimensional), an invertible sheaf \mathcal{L} , and a linear map $\lambda: V \rightarrow \Gamma(X, \mathcal{L})$. Such a linear series is often called “ V ”, with the rest of the data left implicit. If the map λ is an isomorphism, it is called a **complete linear series**, and is often written $|\mathcal{L}|$. The language of base points (Definition 16.3.4) readily translates to this situation. For example, given a linear series, any point $p \in X$ on which all elements of the linear series V vanish, we say that p is a **base point** of V . If V has no base points, we say that it is **base-point-free**. The set of base points is called the **base locus** of the linear series. One can similarly define the **scheme-theoretic base locus** of the linear series.

As a reality check, you should understand why an $n + 1$ -dimensional linear series on a k -scheme X with base-point-free locus U defines a morphism $U \rightarrow \mathbb{P}_k^n$.

16.3.7. Serre’s Theorem A. We are now able to state a celebrated result of Serre.

16.3.8. Serre’s Theorem A. — Suppose S_\bullet is generated in degree 1, and finitely generated over $A = S_0$. Let \mathcal{F} be any finite type quasicoherent sheaf on $\text{Proj } S_\bullet$. Then there exists some n_0 such that for all $n \geq n_0$, $\mathcal{F}(n)$ is finitely globally generated.

We could now prove Serre’s Theorem A directly, but will continue to use this as an excuse to introduce more ideas; it will be a consequence of Theorem 17.6.2.

16.3.9. Proof of Theorem 16.3.1 assuming Serre’s Theorem A (Theorem 16.3.8). Suppose we have m global sections s_1, \dots, s_m of $\mathcal{F}(n)$ that generate $\mathcal{F}(n)$. This gives a map

$$\oplus_m \mathcal{O} \longrightarrow \mathcal{F}(n)$$


given by $(f_1, \dots, f_m) \mapsto f_1 s_1 + \dots + f_m s_m$ on any open set. Because these global sections generate $\mathcal{F}(n)$, this is a surjection. Tensoring with $\mathcal{O}(-n)$ (which is exact, as tensoring with any locally free sheaf is exact, Exercise 14.2.E) gives the desired result. \square

16.4 Quasicoherent sheaves and graded modules

(This section answers some fundamental questions, but it is surprisingly tricky. You may wish to skip this section, or at least the proofs, on first reading, unless you have a particular need for them.)

Throughout this section, S_\bullet is a finitely generated graded algebra generated in degree 1, so in particular we have the invertible sheaf $\mathcal{O}(n)$ for all n by Exercise 16.2.B. Also, throughout, M_\bullet is a graded S_\bullet -module, and \mathcal{F} is a quasicoherent sheaf on $\text{Proj } S_\bullet$.

We know how to get quasicoherent sheaves on $\text{Proj } S_\bullet$ from graded S_\bullet -modules. We will now see that we can get them all in this way. We will define a functor Γ_\bullet from (the category of) quasicoherent sheaves on $\text{Proj } S_\bullet$ to (the category of) graded S_\bullet -modules that will attempt to reverse the \sim construction. They are not quite inverses, as \sim can turn two different graded modules into the same quasicoherent sheaf (see for example Exercise 16.1.C). But we will see a natural isomorphism $\widetilde{\Gamma_\bullet(\mathcal{F})} \xrightarrow{\sim} \mathcal{F}$. In fact $\Gamma_\bullet(\widetilde{M_\bullet})$ is a better (“saturated”) version of M_\bullet , and there is a saturation functor $M_\bullet \rightarrow \Gamma_\bullet(\widetilde{M_\bullet})$ that is akin to groupification and sheafification — it is adjoint to the forgetful functor from saturated graded modules to graded modules. And thus we come to the fundamental relationship between \sim and Γ_\bullet : they are an adjoint pair.



We now make some of this precise, but as little as possible to move forward. In particular, we will show that every quasicoherent sheaf on a projective A -scheme arises from a graded module (Corollary 16.4.3), and we will see that every closed subscheme of $\text{Proj } S_\bullet$ arises from a homogeneous ideal $I_\bullet \subset S_\bullet$ (Exercise 16.4.H).

16.4.1. Definition of Γ_\bullet . When you do Essential Exercise 15.1.C (on global sections of $\mathcal{O}_{\mathbb{P}^n_k}(n)$), you will suspect that in good situations,

$$M_n \cong \Gamma(\text{Proj } S_\bullet, \widetilde{M}(n)_\bullet).$$

Motivated by this, we define

$$\Gamma_n(\mathcal{F}) := \Gamma(\text{Proj } S_\bullet, \mathcal{F}(n)).$$

16.4.A. EXERCISE. Describe a morphism of S_0 -modules $M_n \rightarrow \Gamma(\text{Proj } S_\bullet, \widetilde{M}(n)_\bullet)$, extending the $n = 0$ case of Exercise 16.1.D.

16.4.B. EXERCISE. Show that $\Gamma_\bullet(\mathcal{F})$ is a graded S_\bullet -module. (Hint: consider $S_n \rightarrow \Gamma(\text{Proj } S_\bullet, \mathcal{O}(n))$.)

16.4.C. EXERCISE. Show that the map $M_\bullet \rightarrow \Gamma_\bullet(\widetilde{M}_\bullet)$ arising from the previous two exercises is a map of S_\bullet -modules. We call this the **saturation map**.

16.4.D. EXERCISE. Show that the saturation map need not be injective, nor need it be surjective. (Hint: $S_\bullet = k[x]$, $M_\bullet = k[x]/(x^2)$ or $M_\bullet = xk[x]$.)

16.4.E. EXERCISE. Show that Γ_\bullet is a functor from $QCoh_{\text{Proj } S_\bullet}$ to the category of graded S_\bullet -modules. In other words, if $\mathcal{F} \rightarrow \mathcal{G}$ is a morphism of quasicoherent sheaves on $\text{Proj } S_\bullet$, describe the natural map $\Gamma_\bullet \mathcal{F} \rightarrow \Gamma_\bullet \mathcal{G}$, and show that such maps respect the identity and composition. Thus the saturation map can be better called the **saturation functor**.

16.4.2. ★ The reverse map. Now that we have defined the saturation map $M_\bullet \rightarrow \Gamma_\bullet \widetilde{M}_\bullet$, we will describe a map $\widetilde{\Gamma_\bullet \mathcal{F}} \rightarrow \mathcal{F}$. While subtler to define, it will have the advantage of being an isomorphism.

16.4.F. EXERCISE. Define the natural map $\widetilde{\Gamma_\bullet \mathcal{F}} \rightarrow \mathcal{F}$ as follows. First describe the map on sections over $D(f)$. Note that sections of the left side are of the form m/f^n where $m \in \Gamma_{n \deg f}(\mathcal{F})$, and $m/f^n = m'/f^{n'}$ if there is some N with $f^N(f^{n'}m - f^n m') = 0$. Sections on the right are implicitly described in Exercise 16.3.A. Show that your map behaves well on overlaps $D(f) \cap D(g) = D(fg)$.

16.4.G. TRICKY EXERCISE. Show that the natural map $\widetilde{\Gamma_\bullet \mathcal{F}} \rightarrow \mathcal{F}$ is an isomorphism, by showing that it is an isomorphism of sections over $D(f)$ for any f . First show surjectivity, using Exercise 16.3.A to show that any section of \mathcal{F} over $D(f)$ is of the form m/f^n where $m \in \Gamma_{n \deg f}(\mathcal{F})$. Then verify that it is injective.

16.4.3. Corollary. — *Every quasicoherent sheaf on a projective A -scheme arises from the \sim construction.*

16.4.H. EXERCISE (CONVERSE TO EXERCISE 9.3.B). Show that each closed subscheme of $\text{Proj } S_\bullet$ arises from a homogeneous ideal $I_\bullet \subset S_\bullet$. (Hint: Suppose Z is

a closed subscheme of $\text{Proj } S_\bullet$. Consider the exact sequence $0 \rightarrow \mathcal{I}_Z \rightarrow \mathcal{O}_{\text{Proj } S_\bullet} \rightarrow \mathcal{O}_Z \rightarrow 0$. Apply Γ_\bullet , and then \sim . Be careful: Γ_\bullet is left-exact, but not necessarily exact.) This fulfills promises made in Exercises 9.3.B and 16.1.E.

For the first time, we see that every closed subscheme of a projective scheme is cut out by homogeneous equations. This is the analog of the fact that every closed subscheme of an affine scheme is cut out by equations. It is disturbing that it is so hard to prove this fact. (For comparison, this was easy on the level of the Zariski topology — see Exercise 4.5.K(c).)

16.4.I. ★★ EXERCISE (Γ_\bullet AND \sim ARE ADJOINT FUNCTORS). Describe a natural bijection $\text{Hom}(M_\bullet, \Gamma_\bullet \mathcal{F}) \cong \text{Hom}(\widetilde{M_\bullet}, \mathcal{F})$, as follows.

- (a) Show that maps $M_\bullet \rightarrow \Gamma_\bullet \mathcal{F}$ are the “same” as maps $((M_\bullet)_f)_0 \rightarrow ((\Gamma_\bullet \mathcal{F})_f)_0$ as f varies through S_+ , that are “compatible” as f varies, i.e., if $D(g) \subset D(f)$, there is a commutative diagram

$$\begin{array}{ccc} ((M_\bullet)_f)_0 & \longrightarrow & ((\Gamma_\bullet \mathcal{F})_f)_0 \\ \downarrow & & \downarrow \\ ((M_\bullet)_g)_0 & \longrightarrow & ((\Gamma_\bullet \mathcal{F})_g)_0 \end{array}$$

More precisely, give a bijection between $\text{Hom}(M_\bullet, \Gamma_\bullet \mathcal{F})$ and the set of compatible maps

$$\left(\text{Hom}((M_\bullet)_f)_0 \rightarrow ((\Gamma_\bullet \mathcal{F})_f)_0 \right)_{f \in S_+}.$$

- (b) Describe a bijection between the set of compatible maps $(\text{Hom}((M_\bullet)_f)_0 \rightarrow ((\Gamma_\bullet \mathcal{F})_f)_0)_{f \in S_+}$ and the set of compatible maps $\Gamma(D(f), \widetilde{M_\bullet}) \rightarrow \Gamma(D(f), \mathcal{F})$.

16.4.4. The special case $M_\bullet = S_\bullet$. We have a saturation map $S_\bullet \rightarrow \Gamma_\bullet \widetilde{S_\bullet}$, which is a map of S_\bullet -modules. But $\Gamma_\bullet \widetilde{S_\bullet}$ has the structure of a graded ring (basically because we can multiply sections of $\mathcal{O}(m)$ by sections of $\mathcal{O}(n)$ to get sections of $\mathcal{O}(m+n)$, see Exercise 16.2.D).

16.4.J. EXERCISE. Show that the map of graded rings $S_\bullet \rightarrow \Gamma_\bullet \widetilde{S_\bullet}$ induces (via the construction of Essential Exercise 7.4.A) an isomorphism $\text{Proj } \Gamma_\bullet \widetilde{S_\bullet} \xrightarrow{\sim} \text{Proj } S_\bullet$, and under this isomorphism, the respective $\mathcal{O}(1)$ ’s are identified.

This addresses the following question: to what extent can we recover S_\bullet from $(\text{Proj } S_\bullet, \mathcal{O}(1))$? The answer is: we cannot recover S_\bullet , but we can recover its “saturation”. And better yet: given a projective A -scheme $\pi: X \rightarrow \text{Spec } A$, along with $\mathcal{O}(1)$, we obtain it as a Proj of a graded algebra in a canonical way, via

$$X \cong \text{Proj} \left(\bigoplus_{n \geq 0} \Gamma(X, \mathcal{O}(n)) \right).$$

16.4.5. ★★ Saturated S_\bullet -modules. We end with a remark: different graded S_\bullet -modules give the same quasicoherent sheaf on $\text{Proj } S_\bullet$, but the results of this section show that there is a “best” (= saturated) graded module for each quasicoherent sheaf, and there is a map from each graded module to its “best” version, $M_\bullet \rightarrow \Gamma_\bullet \widetilde{M_\bullet}$. A module for which this is an isomorphism (a “best” module) is called **saturated**. We won’t use this term later.

This “saturation” map $M_{\bullet} \rightarrow \Gamma_{\bullet} \widetilde{M}_{\bullet}$ is analogous to the sheafification map, taking presheaves to sheaves. For example, the saturation of the saturation equals the saturation.

There is a bijection between saturated quasicoherent sheaves of ideals on $\text{Proj } S_{\bullet}$ and closed subschemes of $\text{Proj } S_{\bullet}$.

Pushforwards and pullbacks of quasicoherent sheaves

17.1 Introduction

This chapter is devoted to pushforward and pullbacks of quasicoherent sheaves, their properties, and some applications.

Suppose $B \rightarrow A$ is a morphism of rings. Recall (from Exercise 1.5.E) that $(\cdot \otimes_B A, \cdot_B)$ is an adjoint pair between the categories of A -modules and B -modules: we have a bijection

$$\mathrm{Hom}_A(N \otimes_B A, M) \cong \mathrm{Hom}_B(N, M_B)$$

functorial in both arguments. These constructions behave well with respect to localization (in an appropriate sense), and hence work (often) in the category of quasicoherent sheaves on schemes (and indeed always in the category of \mathcal{O} -modules on ringed spaces, see Definition 17.3.5, although we won't particularly care). The easier construction $(M \mapsto M_B)$ will turn into our old friend pushforward. The other $(N \mapsto A \otimes_B N)$ will be a relative of pullback, whom I'm reluctant to call an "old friend".

17.2 Pushforwards of quasicoherent sheaves

We begin with the pushforwards, for which we have already done much of the work.

The main moral of this section is that in "reasonable" situations, the pushforward of a quasicoherent sheaf is quasicoherent, and that this can be understood in terms of one of the module constructions defined above. We begin with a motivating example:

17.2.A. EXERCISE. Let $\pi: \mathrm{Spec} A \rightarrow \mathrm{Spec} B$ be a morphism of affine schemes, and suppose M is an A -module, so \widetilde{M} is a quasicoherent sheaf on $\mathrm{Spec} A$. Give an isomorphism $\pi_* \widetilde{M} \xrightarrow{\sim} \widetilde{M_B}$. (Hint: There is only one reasonable way to proceed: look at distinguished open sets.)

In particular, $\pi_* \widetilde{M}$ is quasicoherent. Perhaps more important, this implies that the pushforward of a quasicoherent sheaf under an affine morphism is also quasicoherent.

17.2.B. EXERCISE. If $\pi: X \rightarrow Y$ is an affine morphism, show that π_* is an exact functor $QCoh_X \rightarrow QCoh_Y$.

The following result generalizes the fact that the pushforward of a quasicoherent sheaf under an affine morphism is also quasicoherent.

17.2.C. ESSENTIAL EXERCISE (COROLLARY TO EXERCISE 6.2.G: “ $\pi_* : QCoh_X \rightarrow QCoh_Y$ FOR QCQS π ”). Suppose $\pi: X \rightarrow Y$ is a quasicompact quasiseparated morphism, and \mathcal{F} is a quasicoherent sheaf on X . Show that $\pi_*\mathcal{F}$ is a quasicoherent sheaf on Y .

Coherent sheaves do not always push forward to coherent sheaves. For example, consider the structure morphism $\pi: \mathbb{A}_k^1 \rightarrow \text{Spec } k$, corresponding to $k \rightarrow k[t]$. Then $\pi_*\mathcal{O}_{\mathbb{A}_k^1}$ is the quasicoherent sheaf corresponding to $k[t]$, which is not a finitely generated k -module. But in good situations, coherent sheaves do push forward. For example:

17.2.D. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a finite morphism of locally Noetherian schemes. If \mathcal{F} is a coherent sheaf on X , show that $\pi_*\mathcal{F}$ is a coherent sheaf. Hint: Show first that $\pi_*\mathcal{O}_X$ is finite type.

Once we define cohomology of quasicoherent sheaves, we will quickly prove that if \mathcal{F} is a coherent sheaf on \mathbb{P}_k^n , then $\Gamma(\mathbb{P}_k^n, \mathcal{F})$ is a finite-dimensional k -module, and more generally if \mathcal{F} is a coherent sheaf on $\text{Proj } S_\bullet$, then $\Gamma(\text{Proj } S_\bullet, \mathcal{F})$ is a coherent A -module (where $S_0 = A$). This is a special case of the fact that the “pushforwards of coherent sheaves by projective morphisms are also coherent sheaves” (Theorem 19.8.1(d)). (The notion of projective morphism, a relative version of $\text{Proj } S_\bullet \rightarrow \text{Spec } A$, will be defined in §18.3.)

More generally, given Noetherian hypotheses, pushforwards of coherent sheaves by proper morphisms are also coherent sheaves (Grothendieck’s Coherence Theorem 19.9.1).

17.3 Pullbacks of quasicoherent sheaves

We next discuss the pullback of a quasicoherent sheaf: if $\pi: X \rightarrow Y$ is a morphism of schemes, π^* is a covariant functor $QCoh_Y \rightarrow QCoh_X$. The notion of the pullback of a quasicoherent sheaf can be confusing on first (and second) glance. (For example, it is *not* the inverse image sheaf, although we will see that it is related.)

Here are three contexts in which you have seen the pullback, or can understand it quickly. It may be helpful to keep these in mind, to keep you anchored in the long discussion that follows. Suppose \mathcal{G} is a quasicoherent sheaf on a scheme Y .

(i) (*restriction to open subsets*) If $i: U \hookrightarrow Y$ is an open embedding, then $i^*\mathcal{G}$ is $\mathcal{G}|_U$, the restriction of \mathcal{G} to U (Example 2.2.8).

(ii) (*restriction to points*) If $i: p \hookrightarrow Y$ is the “inclusion” of a point p in Y (for example, the closed embedding of a closed point; see Exercise 7.3.L(b)), then $i^*\mathcal{G}$ is $\mathcal{G}|_p$, the fiber of \mathcal{G} at p (Definition 4.3.8).

The similarity of the notation $\mathcal{G}|_U$ and $\mathcal{G}|_p$ is precisely because both are pullbacks. Pullbacks (especially to locally closed subschemes or generic points) are often called *restriction*. For this reason, if $\pi: X \rightarrow Y$ is some sort of inclusion (such as a locally closed embedding, or an “inclusion of a generic point”) then $\pi^*\mathcal{G}$ is often written as $\mathcal{G}|_X$ and called the **restriction of \mathcal{G} to X** , when π can be interpreted as some type of “inclusion”.

(iii) (*pulling back vector bundles*) Suppose \mathcal{G} is a locally free sheaf on Y , $\pi: X \rightarrow Y$ is any morphism. If $\{U_i\}$ are trivializing open neighborhoods for \mathcal{G} , and $T_{ij} \in \mathrm{GL}_r(\mathcal{O}_Y(U_i \cap U_j))$ are transition matrices for \mathcal{G} between U_i and U_j , then $\{\pi^{-1}U_i\}$ are trivializing open neighborhoods for $\pi^*\mathcal{G}$, and π^*T_{ij} are transition matrices for $\pi^*\mathcal{G}$. (This will be established in Theorem 17.3.7(3).)

17.3.1. Strategy. We will see three different ways of thinking about the pullback. Each has significant disadvantages, but together they give a good understanding.

(a) Because we are understanding quasicoherent sheaves in terms of affine open sets, and modules over the corresponding rings, we begin with an interpretation in this vein. This will be very useful for proving and understanding facts. The disadvantage is that it is annoying to make a definition out of this (when the target is not affine), because gluing arguments can be tedious.

(b) As we saw with fibered product, gluing arguments can be made simpler using universal properties, so our second “definition” will be by universal property. This is elegant, but has the disadvantage that it still needs a construction, and because it works in the larger category of \mathcal{O} -modules, it isn’t clear from the universal property that it takes quasicoherent sheaves to quasicoherent sheaves. But if the target is affine, our construction of (a) is easily seen to satisfy universal property. Furthermore, the universal property is “local on the target”: if $\pi: X \rightarrow Y$ is any morphism, $i: U \hookrightarrow Y$ is an open embedding, and \mathcal{G} is a quasicoherent sheaf on Y , then if $\pi^*\mathcal{G}$ exists, then its restriction to $\pi^{-1}(U)$ is (canonically identified with) $(\pi|_U)^*(\mathcal{G}|_U)$. Thus if the pullback exists in general (even as an \mathcal{O} -module), affine-locally on Y it looks like the construction of (a) (and thus is quasicoherent).

(c) The third definition is one that works on ringed spaces in general. It is short, and is easily seen to satisfy the universal property. It doesn’t obviously take quasicoherent sheaves to quasicoherent sheaves (at least in the way that we have defined quasicoherent sheaves) — a priori it takes quasicoherent sheaves to \mathcal{O} -modules. But thanks to the discussion at the end of (b) above, which used (a), this shows that the pullback of a quasicoherent sheaf is indeed quasicoherent.

17.3.2. First attempt at describing the pullback, using affines. Suppose $\pi: X \rightarrow Y$ is a morphism of schemes, and \mathcal{G} is a quasicoherent sheaf on Y . We want to define the pullback quasicoherent sheaf $\pi^*\mathcal{G}$ on X in terms of affine open sets on X and Y . Suppose $\mathrm{Spec} A \subset X$, $\mathrm{Spec} B \subset Y$ are affine open sets, with $\pi(\mathrm{Spec} A) \subset \mathrm{Spec} B$. Suppose $\mathcal{G}|_{\mathrm{Spec} B} \cong \tilde{N}$. Perhaps motivated by the fact that pullback should relate to tensor product, we want

$$(17.3.2.1) \quad \Gamma(\mathrm{Spec} A, \pi^*\mathcal{G}) = N \otimes_B A.$$

More precisely, we would like $\Gamma(\mathrm{Spec} A, \pi^*\mathcal{G})$ and $N \otimes_B A$ to be identified. This could mean that we use this to construct a definition of $\pi^*\mathcal{G}$, by “gluing all this information together” (and showing it is well-defined). Or it could mean that we define $\pi^*\mathcal{G}$ in some other way, and then find a natural identification (17.3.2.1). The

first approach can be made to work (and §17.3.3 is the first step), but we will follow the second.




FIGURE 17.1. The pullback of a quasicoherent sheaf (module-theoretically)

17.3.3. We begin this project by *fixing* an affine open subset $\text{Spec } B \subset Y$. To avoid confusion, let $\phi = \pi|_{\pi^{-1}(\text{Spec } B)} : \pi^{-1}(\text{Spec } B) \rightarrow \text{Spec } B$. We will define a quasicoherent sheaf on $\pi^{-1}(\text{Spec } B)$ that will turn out to be $\phi^*(\mathcal{G}|_{\text{Spec } B})$ (and will also be the restriction of $\pi^*\mathcal{G}$ to $\pi^{-1}(\text{Spec } B)$).

If $\text{Spec } A_f \subset \text{Spec } A$ is a distinguished open set, then

$$\Gamma(\text{Spec } A_f, \phi^*\mathcal{G}) = N \otimes_B A_f = (N \otimes_B A)_f = \Gamma(\text{Spec } A, \phi^*\mathcal{G})_f$$

where “=” means “canonically isomorphic”. Define the restriction map $\Gamma(\text{Spec } A, \phi^*\mathcal{G}) \rightarrow \Gamma(\text{Spec } A_f, \phi^*\mathcal{G})$,

$$(17.3.3.1) \quad \Gamma(\text{Spec } A, \phi^*\mathcal{G}) \longrightarrow \Gamma(\text{Spec } A, \phi^*\mathcal{G}) \otimes_A A_f,$$

by $\alpha \mapsto \alpha \otimes 1$ (of course). Thus $\phi^*\mathcal{G}$ is (or: extends to) a quasicoherent sheaf on $\pi^{-1}(\text{Spec } B)$ (by Exercise 6.2.D).

We have now defined a quasicoherent sheaf on $\pi^{-1}(\text{Spec } B)$, for every affine open subset $\text{Spec } B \subset Y$. We want to show that this construction, as $\text{Spec } B$ varies, glues into a single quasicoherent sheaf on X .

You are welcome to do this gluing appropriately, for example using the distinguished affine base of Y (§6.2.1). This works, but can be confusing, so we take another approach.

17.3.4. Universal property definition of pullback. Suppose $\pi: X \rightarrow Y$ is a morphism of ringed spaces, and \mathcal{G} is an \mathcal{O}_Y -module. (We are of course interested in the case where π is a morphism of schemes, and \mathcal{G} is quasicoherent. Even once we specialize our discussion to schemes, much of our discussion will extend without change to this more general situation.) We “define” the pullback $\pi^*\mathcal{G}$ as an \mathcal{O}_X -module, using the following adjointness universal property: for any \mathcal{O}_X -module \mathcal{F} , there is a bijection $\text{Hom}_{\mathcal{O}_X}(\pi^*\mathcal{G}, \mathcal{F}) \leftrightarrow \text{Hom}_{\mathcal{O}_Y}(\mathcal{G}, \pi_*\mathcal{F})$, and these bijections are functorial in \mathcal{F} . By universal property nonsense, this determines $\pi^*\mathcal{G}$ up to unique isomorphism; we just need to make sure that it exists (which is why the word “define” is in quotes). Notice that we avoid worrying about when the push-forward of a quasicoherent sheaf \mathcal{F} is quasicoherent by working in the larger category of \mathcal{O} -modules.

17.3.A. IMPORTANT EXERCISE. If Y is affine, say $Y = \operatorname{Spec} B$, show that the construction of the quasicoherent sheaf in §17.3.3 satisfies this universal property of pullback of \mathcal{G} . Thus calling this sheaf $\pi^*\mathcal{G}$ is justified. (Hint: Interpret both sides of the alleged bijection explicitly. The adjointness in the ring/module case should turn up.)

17.3.B. IMPORTANT EXERCISE. Suppose $i: U \hookrightarrow X$ is an open embedding of ringed spaces, and \mathcal{F} is an \mathcal{O}_X -module. Show that $\mathcal{F}|_U$ satisfies the universal property of $i^*\mathcal{F}$ (and thus deserves to be called $i^*\mathcal{F}$). In other words, for each \mathcal{O}_U -module \mathcal{E} , describe a bijection

$$\operatorname{Hom}_{\mathcal{O}_U}(\mathcal{F}|_U, \mathcal{E}) \longleftrightarrow \operatorname{Hom}_{\mathcal{O}_X}(\mathcal{F}, i_*\mathcal{E}),$$

functorial in \mathcal{E} .

We next show that if $\pi^*\mathcal{G}$ satisfies the universal property (for the morphism $\pi: X \rightarrow Y$), then if $j: V \hookrightarrow Y$ is any open subset, and $i: U = \pi^{-1}(V) \hookrightarrow X$ (see (17.3.4.1)), then $(\pi^*\mathcal{G})|_U$ satisfies the universal property for $\pi|_U: U \rightarrow V$. Thus $(\pi^*\mathcal{G})|_U$ deserves to be called $\pi|_U^*(\mathcal{G}|_V)$. You will notice that we really need to work with \mathcal{O} -modules, not just with quasicoherent sheaves.

$$(17.3.4.1) \quad \begin{array}{ccc} \pi^{-1}(V) & \xlongequal{\quad} & U \xrightarrow{i} X \\ & \pi|_U \downarrow & \downarrow \pi \\ & V \xrightarrow{j} & Y \end{array}$$

If \mathcal{F}' is an \mathcal{O}_U -module, we have a series of bijections (using Important Exercise 17.3.B and adjointness of pullback and pushforward):

$$\begin{aligned} \operatorname{Hom}_{\mathcal{O}_U}((\pi^*\mathcal{G})|_U, \mathcal{F}') &\cong \operatorname{Hom}_{\mathcal{O}_U}(i^*(\pi^*\mathcal{G}), \mathcal{F}') \\ &\cong \operatorname{Hom}_{\mathcal{O}_X}(\pi^*\mathcal{G}, i_*\mathcal{F}') \\ &\cong \operatorname{Hom}_{\mathcal{O}_Y}(\mathcal{G}, \pi_*i_*\mathcal{F}') \\ &\cong \operatorname{Hom}_{\mathcal{O}_Y}(\mathcal{G}, j_*(\pi|_U)_*\mathcal{F}') \\ &\cong \operatorname{Hom}_{\mathcal{O}_V}(j^*\mathcal{G}, (\pi|_U)_*\mathcal{F}') \\ &\cong \operatorname{Hom}_{\mathcal{O}_V}(\mathcal{G}|_V, (\pi|_U)_*\mathcal{F}'). \end{aligned}$$

We have thus described a bijection

$$\operatorname{Hom}_{\mathcal{O}_U}((\pi^*\mathcal{G})|_U, \mathcal{F}') \longleftrightarrow \operatorname{Hom}_{\mathcal{O}_V}(\mathcal{G}|_V, (\pi|_U)_*\mathcal{F}'),$$

which is clearly (by construction) functorial in \mathcal{F}' . Hence the discussion in the first paragraph of §17.3.4 is justified. For example, thanks to Important Exercise 17.3.A, the pullback π^* exists if Y is an open subset of an affine scheme.

At this point, we could show that the pullback exists, following the idea behind the construction of the fibered product: we would start with the definition when Y is affine, and “glue”. We will instead take another route.

17.3.5. Third definition: pullback of \mathcal{O} -modules via explicit construction. Suppose $\pi: X \rightarrow Y$ is a morphism of ringed spaces, and \mathcal{G} is an \mathcal{O}_Y -module. Of course, our case of interest is if π is a morphism of schemes, and \mathcal{G} is quasicoherent. Now $\pi^{-1}\mathcal{G}$ is a $\pi^{-1}\mathcal{O}_Y$ -module. (Notice that we are using the ringed space $(X, \pi^{-1}\mathcal{O}_Y)$,

not (X, \mathcal{O}_X) . The inverse image construction π^{-1} was discussed in §2.7.) Furthermore, \mathcal{O}_X is also a $\pi^{-1}\mathcal{O}_Y$ -module, via the map $\pi^{-1}\mathcal{O}_Y \rightarrow \mathcal{O}_X$ that is part of the data of the morphism π . Define the **pullback** of \mathcal{G} by π as the \mathcal{O}_X -module

$$(17.3.5.1) \quad \pi^*\mathcal{G} := \pi^{-1}\mathcal{G} \otimes_{\pi^{-1}\mathcal{O}_Y} \mathcal{O}_X.$$

It is immediate that pullback is a covariant functor $\pi^*: \text{Mod}_{\mathcal{O}_Y} \rightarrow \text{Mod}_{\mathcal{O}_X}$.

17.3.C. IMPORTANT EXERCISE. Show that this definition (17.3.5.1) of pullback satisfies the universal property. Thus the pullback exists, at least as a functor $\text{Mod}_{\mathcal{O}_Y} \rightarrow \text{Mod}_{\mathcal{O}_X}$.

17.3.D. IMPORTANT EXERCISE. Show that if $\pi: X \rightarrow Y$ is a morphism of schemes, then π^* gives a covariant functor $\text{QCoh}_Y \rightarrow \text{QCoh}_X$. (You might use §17.3.3, Exercise 17.3.B, Exercise 17.3.A, and/or §17.3.4.)

The following is then immediate from the universal property.

17.3.6. Proposition. — Suppose $\pi: X \rightarrow Y$ is a quasicompact, quasiseparated morphism. Then $(\pi^*: \text{QCoh}_Y \rightarrow \text{QCoh}_X, \pi_*: \text{QCoh}_X \rightarrow \text{QCoh}_Y)$ are an adjoint pair: there is an isomorphism

$$(17.3.6.1) \quad \text{Hom}_{\mathcal{O}_X}(\pi^*\mathcal{G}, \mathcal{F}) \xleftarrow{\sim} \text{Hom}_{\mathcal{O}_Y}(\mathcal{G}, \pi_*\mathcal{F}),$$

functorial in both $\mathcal{F} \in \text{QCoh}_X$ and $\mathcal{G} \in \text{QCoh}_Y$.

The “quasicompact and quasiseparated” hypotheses are solely to ensure that π_* indeed sends QCoh_X to QCoh_Y (Exercise 17.2.C).

We are now ready to show that pullback has all sorts of desired properties.

17.3.7. Theorem. — Suppose $\pi: X \rightarrow Y$ is a morphism of schemes, and \mathcal{G} is a quasicoherent sheaf on Y .

- (1) (pullback preserves the structure sheaf) There is a canonical isomorphism $\pi^*\mathcal{O}_Y \xleftarrow{\sim} \mathcal{O}_X$.
- (2) (pullback preserves finite type quasicoherent sheaves) If \mathcal{G} is a finite type quasicoherent sheaf, so is $\pi^*\mathcal{G}$. Hence if X is locally Noetherian, and \mathcal{G} is coherent, then so is $\pi^*\mathcal{G}$. (It is not always true that the pullback of a coherent sheaf is coherent, and the interested reader can think of a counterexample.)
- (3) (pullback preserves vector bundles, and their transition functions) If \mathcal{G} is locally free sheaf of rank r , then so is $\pi^*\mathcal{G}$. (In particular, the pullback of an invertible sheaf is invertible.) Furthermore, if $\{\mathcal{U}_i\}$ are trivializing open neighborhoods for \mathcal{G} , and $T_{ij} \in \text{GL}_r(\mathcal{O}_Y(\mathcal{U}_i \cap \mathcal{U}_j))$ are transition matrices for \mathcal{G} between \mathcal{U}_i and \mathcal{U}_j , then $\{\pi^{-1}\mathcal{U}_i\}$ are trivializing open neighborhoods for $\pi^*\mathcal{G}$, and π^*T_{ij} are transition matrices for $\pi^*\mathcal{G}$.
- (4) (functoriality in the morphism) If $\xi: W \rightarrow X$ is a morphism of schemes, then there is a canonical isomorphism $\xi^*\pi^*\mathcal{G} \xleftarrow{\sim} (\pi \circ \xi)^*\mathcal{G}$.
- (5) (functoriality in the quasicoherent sheaf) π^* is a functor $\text{QCoh}_Y \rightarrow \text{QCoh}_X$.
- (6) (pulling back a section) Note that a section of \mathcal{G} is the data of a map $\mathcal{O}_Y \rightarrow \mathcal{G}$. By (1) and (5), if $s: \mathcal{O}_Y \rightarrow \mathcal{G}$ is a section of \mathcal{G} then there is a natural section $\pi^*s: \mathcal{O}_X \rightarrow \pi^*\mathcal{G}$ of $\pi^*\mathcal{G}$. The pullback of the locus where s vanishes is the locus where the pulled-back section π^*s vanishes.

(7) (pullback on stalks) If $\pi: X \rightarrow Y$, $\pi(p) = q$, then pullback induces an isomorphism

$$(17.3.7.1) \quad (\pi^*\mathcal{G})_p \xrightarrow{\sim} \mathcal{G}_q \otimes_{\mathcal{O}_{Y,q}} \mathcal{O}_{X,p}.$$

(8) (pullback on fibers of the quasicoherent sheaves) Pullback of fibers are given as follows: if $\pi: X \rightarrow Y$, where $\pi(p) = q$, then the map

$$(\pi^*\mathcal{G})|_p \xrightarrow{\sim} \mathcal{G}|_q \otimes_{\kappa(q)} \kappa(p)$$

induced by (17.3.7.1) is an isomorphism.

(9) (pullback preserves tensor product) $\pi^*(\mathcal{G} \otimes_{\mathcal{O}_Y} \mathcal{G}') = \pi^*\mathcal{G} \otimes_{\mathcal{O}_X} \pi^*\mathcal{G}'$. (Here \mathcal{G}' is also a quasicoherent sheaf on Y .)

(10) Pullback is a right-exact functor.

All of the above are interconnected in obvious ways that you should be able to prove by hand. (As just one example: the germ of a pulled back section, (6), is the expected element of the pulled back stalk, (7).) In fact much more is true, that you should be able to prove on a moment's notice, such as for example that the pullback of the symmetric power of a locally free sheaf is naturally isomorphic to the symmetric power of the pullback, and similarly for wedge powers and tensor powers.

17.3.E. IMPORTANT EXERCISE. Prove Theorem 17.3.7. Possible hints: You may find it convenient to do right-exactness (10) early; it is related to right-exactness of \otimes . For the tensor product fact (9), show that $(M \otimes_B A) \otimes (N \otimes_B A) \cong (M \otimes_B N) \otimes_B A$, and that this behaves well with respect to localization. The proof of the fiber fact (8) is as follows. Given a ring map $B \rightarrow A$ with $[m] \mapsto [n]$, where $m \subset A$ and $n \subset B$ are maximal ideals, show that $(N \otimes_B A) \otimes_A (A/m) \cong (N \otimes_B (B/n)) \otimes_{B/n} (A/m)$ by showing both sides are isomorphic to $N \otimes_B (A/m)$.

17.3.F. UNIMPORTANT EXERCISE. Verify that the following is an example showing that pullback is not left-exact: consider the exact sequence of sheaves on \mathbb{A}_k^1 , where $i: p \hookrightarrow \mathbb{A}_k^1$ is the origin:

$$0 \longrightarrow \mathcal{O}_{\mathbb{A}_k^1}(-p) \longrightarrow \mathcal{O}_{\mathbb{A}_k^1} \longrightarrow i_*\mathcal{O}_p \longrightarrow 0.$$

(This is the closed subscheme exact sequence (9.1.2.1) for $p \in \mathbb{A}_k^1$, and corresponds to the exact sequence of $k[t]$ -modules $0 \rightarrow tk[t] \rightarrow k[t] \rightarrow k \rightarrow 0$. Warning: here $\mathcal{O}|_p$ is not the stalk \mathcal{O}_p ; it is the structure sheaf of the scheme p .) Restrict to p .

17.3.G. EXERCISE.

(a) (the push-pull map for \mathcal{O} -modules) Suppose

$$(17.3.7.2) \quad \begin{array}{ccc} X' & \xrightarrow{\psi'} & X \\ \pi' \downarrow & & \downarrow \pi \\ Y' & \xrightarrow{\psi} & Y \end{array}$$

is a commutative diagram of ringed spaces, and \mathcal{F} is an \mathcal{O}_X -module. Describe a natural map (functorial in \mathcal{F}) of $\mathcal{O}_{Y'}$ -modules,

$$(17.3.7.3) \quad \psi^*\pi_*\mathcal{F} \longrightarrow \pi'_*(\psi')^*\mathcal{F}$$

following the Definition 2.7.4 of the push-pull map for sheaves in general.

(b) (*the push-pull map for quasicoherent sheaves*) Suppose now that (17.3.7.2) is a commutative diagram of schemes, and π and π' are both quasicompact and quasiseparated (so pushforward by both sends quasicoherent sheaves to quasicoherent sheaves). Describe a natural map (functorial in \mathcal{F}) of quasicoherent sheaves on Y' , (17.3.7.3). (See Exercise 19.8.B for a generalization to cohomology.)

By applying the above exercise in the special case where Y' is a point q of Y , we see that there is a natural map from the fiber of the pushforward to the sections over the fiber:

$$(17.3.7.4) \quad \pi_* \mathcal{F} \otimes \kappa(q) \longrightarrow \Gamma(\pi^{-1}(q), \mathcal{F}|_{\pi^{-1}(q)}).$$

One might hope that (17.3.7.4) is an isomorphism, i.e., that $\pi_* \mathcal{F}$ “glues together” the fibers $\Gamma(\pi^{-1}(q), \mathcal{F}|_{\pi^{-1}(q)})$. This is too much to ask, but at least (17.3.7.4) gives a map. (In fact, under just the right circumstances, (17.3.7.4) is an isomorphism, see §28.1.)

17.3.H. IMPORTANT EXERCISE (PROJECTION FORMULA, TO BE GENERALIZED IN EXERCISE 19.8.E). Suppose $\pi: X \rightarrow Y$ is quasicompact and quasiseparated, and \mathcal{F} and \mathcal{G} are quasicoherent sheaves on X and Y respectively.

(a) Describe a natural morphism $(\pi_* \mathcal{F}) \otimes \mathcal{G} \rightarrow \pi_*(\mathcal{F} \otimes \pi^* \mathcal{G})$.

(b) If \mathcal{G} is locally free, show that this natural morphism is an isomorphism. (Hint: what if \mathcal{G} is free?)

(c) If π is affine, then show that this natural morphism is an isomorphism.

17.3.8. Remark: Flatness. Given $\pi: X \rightarrow Y$, if the functor π^* from quasicoherent sheaves on Y to quasicoherent sheaves on X is exact (not just right-exact as given by Theorem 17.3.7(10)), we will say that π is a **flat morphism**. This is an incredibly important notion. We will return to it (and define it in a different-looking, equivalent but better way) in Chapter 25; see Exercise 25.2.L.

17.3.9. Remark: Pulling back ideal sheaves. There is one subtlety in pulling back quasicoherent *ideal* sheaves. Suppose $i: X \hookrightarrow Y$ is a closed embedding, and $\mu: Y' \rightarrow Y$ is an arbitrary morphism. Let $X' := X \times_Y Y'$. As “closed embeddings pull back” (§10.2.2), the pulled back map $i': X' \rightarrow Y'$ is a closed embedding.

$$\begin{array}{ccc} X' & \longrightarrow & X \\ i' \downarrow & & \downarrow i \\ Y' & \xrightarrow{\mu} & Y \end{array}$$

Now μ^* induces canonical isomorphisms $\mu^* \mathcal{O}_Y \xrightarrow{\sim} \mathcal{O}_{Y'}$ and $\mu^*(i_* \mathcal{O}_X) \xrightarrow{\sim} (i'_* \mathcal{O}_{X'})$. (Why is the latter isomorphism true? Hint: check affine-locally.) But it is *not* necessarily true that $\mu^* \mathcal{I}_{X/Y} = \mathcal{I}_{X'/Y'}$. (Exercise 17.3.F yields an example.) This is because the application of μ^* to the closed subscheme exact sequence

$$0 \longrightarrow \mathcal{I}_{X/Y} \longrightarrow \mathcal{O}_Y \longrightarrow i_* \mathcal{O}_X \longrightarrow 0$$

yields something that is a priori only right-exact:

$$\mu^* \mathcal{I}_{X/Y} \longrightarrow \mathcal{O}_{Y'} \longrightarrow i'_* \mathcal{O}_{X'} \longrightarrow 0.$$

Thus, as $\mathcal{I}_{X'/Y'}$ is the kernel of $\mathcal{O}_{Y'} \rightarrow i'_* \mathcal{O}_{X'}$, we see that $\mathcal{I}_{X'/Y'}$ is the image of $\mu^* \mathcal{I}_{X/Y}$ in $\mathcal{O}_{Y'}$. We can also see this explicitly from Exercise 10.2.B: affine-locally, the ideal of the pullback is generated by the pullback of the ideal.

Note also that if μ is flat (Remark 17.3.8), then $\mu^* \mathcal{I}_{X/Y} \rightarrow \mathcal{I}_{X'/Y'}$ is an isomorphism.

17.4 Line bundles and maps to projective schemes

Theorem 17.4.1, the converse or completion to Exercise 16.3.G, will give one reason why line bundles are crucially important: they tell us about maps to projective space, and more generally, to quasiprojective A -schemes. Given that we have had a hard time naming any non-quasiprojective schemes, they tell us about maps to essentially all interesting schemes.

17.4.1. Important Theorem. — *For a fixed scheme X , morphisms $X \rightarrow \mathbb{P}^n$ are in bijection with the data $(\mathcal{L}, s_0, \dots, s_n)$, where \mathcal{L} is an invertible sheaf and s_0, \dots, s_n are sections of \mathcal{L} with no common zeros, up to isomorphism of these data.*

(This works over \mathbb{Z} , or indeed any base.) Informally: morphisms to \mathbb{P}^n correspond to $n + 1$ sections of a line bundle, not all vanishing at any point, modulo global sections of \mathcal{O}_X^* , as multiplication by an invertible function gives an automorphism of \mathcal{L} . This is one of those important theorems in algebraic geometry that is easy to prove, but quite subtle in its effect on how one should think. It takes some time to properly digest. A “coordinate-free” version is given in Exercise 17.4.F.

17.4.2. Theorem 17.4.1 describes all morphisms to projective space, and hence by the Yoneda philosophy, this can be taken as the *definition* of projective space: it defines projective space up to unique isomorphism. *Projective space \mathbb{P}^n (over \mathbb{Z}) is the moduli space of line bundles \mathcal{L} along with $n + 1$ sections of \mathcal{L} with no common zeros. (Can you give an analogous definition of projective space over X , denoted \mathbb{P}_X^n ?)*

Every time you see a map to projective space, you should immediately simultaneously keep in mind the invertible sheaf and sections.

Maps to projective schemes can be described similarly. For example, if $Y \hookrightarrow \mathbb{P}_k^2$ is the curve $x_2^2 x_0 = x_1^3 - x_1 x_0^2$, then maps from a scheme X to Y are given by an invertible sheaf on X along with three sections s_0, s_1, s_2 , with no common zeros, satisfying $s_2^2 s_0 - s_1^3 + s_1 s_0^2 = 0$.

Here more precisely is the correspondence of Theorem 17.4.1. Any $n + 1$ sections of \mathcal{L} with no common zeros determine a morphism to \mathbb{P}^n , by Exercise 16.3.G. Conversely, if you have a map to projective space $\pi: X \rightarrow \mathbb{P}^n$, then we have $n + 1$ sections of $\mathcal{O}_{\mathbb{P}^n}(1)$, corresponding to the hyperplane sections, x_0, \dots, x_n . Then $\pi^* x_0, \dots, \pi^* x_n$ are sections of $\pi^* \mathcal{O}_{\mathbb{P}^n}(1)$, and they have no common zero. (Reminder: it is helpful to think of pulling back invertible sheaves in terms of pulling back transition functions, see Theorem 17.3.7(3).)

So to prove Theorem 17.4.1, we just need to show that these two constructions compose to give the identity in either direction.

Proof of Important Theorem 17.4.1. Suppose we are given $n + 1$ sections s_0, \dots, s_n of an invertible sheaf \mathcal{L} , with no common zeros, which (via Exercise 16.3.G) induce

a morphism $\pi: X \rightarrow \mathbb{P}^n$. For each s_i , we get a trivialization of \mathcal{L} on the open set X_{s_i} where s_i doesn't vanish. The transition functions for \mathcal{L} are precisely s_i/s_j on $X_{s_i} \cap X_{s_j}$. As $\mathcal{O}(1)$ is trivial on the standard affine open sets $D(x_i)$ of \mathbb{P}^n , $\pi^*(\mathcal{O}(1))$ is trivial on $X_{s_i} = \pi^{-1}(D(x_i))$. Moreover, $s_i/s_j = \pi^*(x_i/x_j)$ (directly from the construction of π in Exercise 16.3.G). This gives an isomorphism $\mathcal{L} \xrightarrow{\sim} \pi^*\mathcal{O}(1)$ — the two invertible sheaves have the same transition functions.

17.4.A. EXERCISE. Show that this isomorphism can be chosen so that

$$(\mathcal{L}, s_0, \dots, s_n) \xrightarrow{\sim} (\pi^*\mathcal{O}(1), \pi^*x_0, \dots, \pi^*x_n),$$

thereby completing one of the two implications of the theorem.

For the other direction, suppose we are given a map $\pi: X \rightarrow \mathbb{P}^n$. Let $s_i = \pi^*x_i \in \Gamma(X, \pi^*(\mathcal{O}(1)))$. As the x_i 's have no common zeros on \mathbb{P}^n , the s_i 's have no common zeros on X . The map $[s_0, \dots, s_n]$ is the same as the map π . We see this as follows. The preimage of $D(x_i)$ (by the morphism $[s_0, \dots, s_n]$) is $D(s_i) = D(\pi^*x_i) = \pi^{-1}(D(x_i))$, so “the right open sets go to the right open sets”. To show the two morphisms $D(s_i) \rightarrow D(x_i)$ (induced from (s_1, \dots, s_n) and π) are the same, we use the fact that maps to an affine scheme $D(x_i)$ are determined by their maps of global sections in the opposite direction (Essential Exercise 7.3.G). Both morphisms $D(s_i) \rightarrow D(x_i)$ corresponds to the ring map $\pi^\sharp: x_j/x_i = x_j/x_i \mapsto s_j/s_i$. \square

17.4.3. Remark: Extending Theorem 17.4.1 to rational maps. Suppose s_0, \dots, s_n are sections of an invertible sheaf \mathcal{L} on a scheme X . Then Theorem 17.4.1 yields a morphism $X - V(s_0, \dots, s_n) \rightarrow \mathbb{P}^n$. In particular, if X is integral, and the s_i are not all 0, these data yield a rational map $X \dashrightarrow \mathbb{P}^n$.

17.4.4. Examples and applications.

17.4.B. EXERCISE (AUTOMORPHISMS OF PROJECTIVE SPACE). Show that all the automorphisms of projective space \mathbb{P}_k^n (fixing k) correspond to $(n+1) \times (n+1)$ invertible matrices over k , modulo scalars (also known as $\mathrm{PGL}_{n+1}(k)$). (Hint: Suppose $\pi: \mathbb{P}_k^n \rightarrow \mathbb{P}_k^n$ is an automorphism. Show that this induces an isomorphism $\pi^*\mathcal{O}(1) \xrightarrow{\sim} \mathcal{O}(1)$. Show that $\pi^*: \Gamma(\mathbb{P}^n, \mathcal{O}(1)) \rightarrow \Gamma(\mathbb{P}^n, \mathcal{O}(1))$ is an isomorphism.)

17.4.5. Automorphisms of projective space are often called **projective transformations.** Because of Exercise 17.4.B, in their incarnation of matrices modulo scalars, projective transformations are also called **projective changes of coordinates**. Exercise 17.4.B will be useful later, especially for the case $n = 1$. In this case, these automorphisms are called **fractional linear transformations**. (For experts: why was Exercise 17.4.B not stated over an arbitrary base ring A ? Where does the argument go wrong in that case? For what rings A does the result still work?)

17.4.C. EXERCISE. Show that $\mathrm{Aut}(\mathbb{P}_k^1)$ is strictly three-transitive on k -valued points, i.e., given two triplets (p_1, p_2, p_3) and (q_1, q_2, q_3) each of distinct k -valued points of \mathbb{P}^1 , there is precisely one automorphism of \mathbb{P}^1 sending p_i to q_i ($i = 1, 2, 3$).

17.4.D. EXERCISE. Suppose p_0, \dots, p_{n+1} are $n+2$ distinct points of \mathbb{P}_k^n , no $n+1$ of which lie on a hyperplane. Show that there is a unique projective transformation

taking p_i ($0 \leq i \leq n$) to $[0, \dots, 0, 1, 0, \dots, 0]$ (where the 1 is in the i th position), and taking p_{n+1} to $[1, \dots, 1]$.

Here are more examples of these ideas in action.

17.4.6. *Example: the Veronese embedding is $|\mathcal{O}_{\mathbb{P}^n_k}(d)|$.* Consider the line bundle $\mathcal{O}_{\mathbb{P}^n_k}(d)$ on \mathbb{P}^n_k . We have checked that the sections of this line bundle form a vector space of dimension $\binom{n+d}{d}$, with a basis corresponding to homogeneous degree d polynomials in the projective coordinates for \mathbb{P}^n_k . Also, they have no common zeros (as for example the subset of sections $x_0^d, x_1^d, \dots, x_n^d$ have no common zeros). Thus the complete linear series is base-point-free, and determines a morphism $v_d: \mathbb{P}^n \rightarrow \mathbb{P}^{\binom{n+d}{d}-1}$. This is the Veronese embedding (Definition 9.3.8). For example, if $n = 2$ and $d = 2$, we get a map $\mathbb{P}^2_k \rightarrow \mathbb{P}^5_k$.

In §9.3.8, we saw that this is a closed embedding. The following is a more general method of checking that maps to projective space are closed embeddings.

17.4.E. LESS IMPORTANT EXERCISE. Suppose $\pi: X \rightarrow \mathbb{P}^n_k$ corresponds to an invertible sheaf \mathcal{L} on X , and sections s_0, \dots, s_n . Show that π is a closed embedding if and only if

- (i) each open set X_{s_i} is affine, and
- (ii) for each i , the map of rings $A[y_0, \dots, y_n] \rightarrow \Gamma(X_{s_i}, \mathcal{O})$ given by $y_j \mapsto s_j/s_i$ is surjective.

17.4.7. *Special case of Example 17.4.6.* Recall that the image of the Veronese embedding when $n = 1$ is called a *rational normal curve of degree d* (Exercise 9.3.I). Our map is $\mathbb{P}^1_k \rightarrow \mathbb{P}^d_k$ given by $[x, y] \mapsto [x^d, x^{d-1}y, \dots, xy^{d-1}, y^d]$.

17.4.F. EXERCISE. Show that a (finite-dimensional) base-point-free linear series V on X corresponding to \mathcal{L} induces a morphism to projective space

$$\phi_V: X \longrightarrow \mathbb{P}V^\vee.$$

(This should be seen as a coordinate-free version of Theorem 17.4.1.)

17.4.G. EXERCISE. (For this exercise, we work over a field k .) Suppose we have a morphism $\pi: X \rightarrow \mathbb{P}^n_k$, corresponding to the base-point-free linear series $\Gamma(\mathbb{P}^n, \mathcal{O}(1)) \rightarrow \Gamma(X, \mathcal{L})$ (so $\mathcal{L} = \pi^*\mathcal{O}(1)$). If the scheme-theoretic image of X in \mathbb{P}^n lies in a hyperplane, we say that the linear series (or X itself) is **degenerate** (and otherwise, **nondegenerate**). Show that a base-point-free linear series $V \rightarrow \Gamma(X, \mathcal{L})$ is nondegenerate if and only if the map $V \rightarrow \Gamma(X, \mathcal{L})$ is an inclusion. In particular, a complete linear series is always nondegenerate.

17.4.H. EXERCISE. Suppose we are given a map $\pi: \mathbb{P}^1_k \rightarrow \mathbb{P}^n_k$ where the corresponding invertible sheaf on \mathbb{P}^1_k is $\mathcal{O}(d)$. (This can reasonably be called a *degree d map*, cf. Exercises 19.4.E and 19.6.I.) Show that if $d < n$, then the image is degenerate. Show that if $d = n$ and the image is nondegenerate, then the image is isomorphic (via an automorphism of projective space, Exercise 17.4.B) to a rational normal curve.

17.4.I. EXERCISE. Define the graded rings $R_\bullet = k[u, v, w]/(uw - v^2)$ and $S_\bullet = k[x, y]$ (with all variables having degree 1). By Exercise 9.3.H, we have an isomorphism $\text{Proj } R_\bullet \xrightarrow{\sim} \text{Proj } S_\bullet$ (via the Veronese embedding v_2). Show that this isomorphism is not induced by a map of graded rings $S_\bullet \rightarrow R_\bullet$.

17.4.8. Remark. You may be able to show that after “regrading”, the isomorphism $\text{Proj } R_\bullet \xrightarrow{\sim} \text{Proj } S_\bullet$ does arise from a map of graded rings ($S_{2\bullet} \rightarrow R_\bullet$). Exercise 20.11.B gives an example where it is not possible to patch the lack of maps of graded rings by just regrading.

17.4.J. EXERCISE: AN EARLY LOOK AT INTERSECTION THEORY, RELATED TO BÉZOUT’S THEOREM. A classical definition of the degree of a curve in projective space is as follows: intersect it with a “general” hyperplane, and count the number of points of intersection, with appropriate multiplicity. We interpret this in the case of $\pi: \mathbb{P}_k^1 \rightarrow \mathbb{P}_k^n$. Show that there is a hyperplane H of \mathbb{P}_k^n not containing $\pi(\mathbb{P}_k^1)$. Equivalently, $\pi^*H \in \Gamma(\mathbb{P}^1, \mathcal{O}_{\mathbb{P}^1}(d))$ (interpreted as the pullback of a section of $\mathcal{O}_{\mathbb{P}^n}(1)$) is not 0. Show that the number of zeros of π^*H is precisely d . (You will have to define “appropriate multiplicity”. What does it mean geometrically if π is a closed embedding, and π^*H has a double zero? Aside: Can you make sense of this even if π is not a closed embedding?) Thus this classical notion of degree agrees with the notion of degree in Exercise 17.4.H. (See Exercise 9.3.D for another case of Bézout’s Theorem. Here we intersect a degree d curve with a degree 1 hyperplane; there we intersect a degree 1 curve with a degree d hypersurface. Exercise 19.6.K will give a common generalization.)

17.4.9. Example: The Segre embedding revisited. The Segre embedding can also be interpreted in this way. This is a useful excuse to define some notation. Suppose \mathcal{F} is a quasicoherent sheaf on a Z -scheme X , and \mathcal{G} is a quasicoherent sheaf on a Z -scheme Y . Let π_X, π_Y be the projections from $X \times_Z Y$ to X and Y respectively. Then $\mathcal{F} \boxtimes \mathcal{G}$ (pronounced “ \mathcal{F} box-times \mathcal{G} ”) is defined to be $\pi_X^* \mathcal{F} \otimes \pi_Y^* \mathcal{G}$. In particular, $\mathcal{O}_{\mathbb{P}^m \times \mathbb{P}^n}(a, b)$ is defined to be $\mathcal{O}_{\mathbb{P}^m}(a) \boxtimes \mathcal{O}_{\mathbb{P}^n}(b)$ (over any base Z). The Segre embedding $\mathbb{P}^m \times \mathbb{P}^n \rightarrow \mathbb{P}^{mn+m+n}$ corresponds to the complete linear series for the invertible sheaf $\mathcal{O}(1, 1)$.

When we first saw the Segre embedding in §10.6, we saw (in different language) that this complete linear series is base-point-free. We also checked by hand (§10.6.1) that it is a closed embedding, essentially by Exercise 17.4.E.

Recall that if \mathcal{L} and \mathcal{M} are both base-point-free invertible sheaves on a scheme X , then $\mathcal{L} \otimes \mathcal{M}$ is also base-point-free (Exercise 16.3.C, see also Definition 16.3.4). We may interpret this fact using the Segre embedding (under reasonable hypotheses on X). If $\phi_{\mathcal{L}}: X \rightarrow \mathbb{P}^M$ is a morphism corresponding to a (base-point-free) linear series based on \mathcal{L} , and $\phi_{\mathcal{M}}: X \rightarrow \mathbb{P}^N$ is a morphism corresponding to a linear series on \mathcal{M} , then the Segre embedding yields a morphism $X \rightarrow \mathbb{P}^M \times \mathbb{P}^N \rightarrow \mathbb{P}^{(M+1)(N+1)-1}$, which corresponds to a base-point-free series of sections of $\mathcal{L} \otimes \mathcal{M}$.

17.4.K. FUN EXERCISE. Suppose X is a quasiprojective k -scheme, and $\pi: \mathbb{P}_k^n \rightarrow X$ is any morphism (over k). Show that either the image of π has dimension n , or π contracts \mathbb{P}_k^n to a point. In particular, there are no nonconstant maps from projective space to a smaller-dimensional quasi-projective variety. Hint: if $X \subset \mathbb{P}^N$,

define d by $\pi^* \mathcal{O}_{\mathbb{P}^n}(1) \cong \mathcal{O}_{\mathbb{P}^n_k}(d)$. Try to show that $d = 0$. To do that, show that if $m \leq n$ then m nonempty hypersurfaces in \mathbb{P}^n have nonempty intersection. For this, use the fact that any nonempty hypersurface in \mathbb{P}^n_k has nonempty intersection with any subscheme of dimension at least 1 (Exercise 12.3.B(a)).

17.4.L. EXERCISE. Explain how GL_n acts (nontrivially!) on \mathbb{P}^{n-1} (over \mathbb{Z} , or over a field of your choice). (The group scheme GL_n was defined in Exercise 7.6.N. The *action* of a group scheme appeared earlier in Exercise 7.6.S(a).) Hint: this is much more easily done with the language of functors, §7.6, using our functorial description of projective space (§17.4.2), than with our old description of projective space in terms of patches. (A generalization to the Grassmannian will be given in Exercise 17.7.L.)

17.4.10. Remark. Over an algebraically closed field \bar{k} , GL_n acts transitively on the closed points of $\mathbb{P}^n_{\bar{k}}$, and the stabilizer of the point $[1, 0, \dots, 0]$ consists of the subgroup P of matrices with 0's in the first column below the first row. (Side Remark: the point is better written as a column vector, so the GL_n -action can be interpreted as matrix multiplication in the usual way.) This suggests that $\mathbb{P}^{n-1}_{\bar{k}}$ is the quotient GL_n/P . This is largely true; but we first would have to make sense of the notion of group quotient.

17.4.11. ★★ A proper nonprojective k -variety — and gluing schemes along closed subschemes.

We conclude by using what we have developed to describe an example of a scheme that is proper but not projective (promised in Remark 11.5.6). We use a construction that looks so fundamental that you may be surprised to find that we won't use it in any meaningful way later.

Fix an algebraically closed field k . For $i = 1, 2$, let $X_i \cong \mathbb{P}^3_k$, Z_i be a line in X_i , and Z'_i be a regular conic in X_i disjoint from Z_i (both Z_i and Z'_i isomorphic to \mathbb{P}^1_k). The construction of §17.4.12 will allow us to glue X_1 to X_2 so that Z_1 is identified with Z'_2 and Z'_1 is identified with Z_2 , see Figure 17.2. (You will be able to make this precise after reading §17.4.12.) The result, call it X , is proper, by Exercise 17.4.P.




FIGURE 17.2. Building a proper nonprojective variety

Then X is not projective. For if it were, then it would be embedded in projective space by some invertible sheaf \mathcal{L} . If X is embedded, then X_1 is too, so \mathcal{L} must restrict to an invertible sheaf on X_1 of the form $\mathcal{O}_{X_1}(n_1)$, where $n_1 > 0$. You can check that the restriction of \mathcal{L} to Z_1 is $\mathcal{O}_{Z_1}(n_1)$, and the restriction of \mathcal{L} to Z'_1 is

$\mathcal{O}_{Z'_1}(2n_1)$. Symmetrically, the restriction of \mathcal{L} to Z_2 is $\mathcal{O}_{Z_2}(n_2)$ for some $n_2 > 0$, and the restriction of \mathcal{L} to Z'_2 is $\mathcal{O}_{Z'_2}(2n_2)$. But after gluing, $Z_1 = Z'_2$, and $Z'_1 = Z_2$, so we have $n_1 = 2n_2$ and $2n_1 = n_2$, which is impossible.

17.4.12. *Gluing two schemes together along isomorphic closed subschemes.*

It is straightforward to show that you can glue two schemes along isomorphic open subschemes. (More precisely, if X_1 and X_2 are schemes, with open subschemes U_1 and U_2 respectively, and an isomorphism $U_1 \xrightarrow{\sim} U_2$, you can make sense of gluing X_1 and X_2 along $U_1 \xrightarrow{\sim} U_2$. You should think this through.) You can similarly glue two schemes along isomorphic *closed* subschemes. We now make this precise. Suppose $Z_1 \hookrightarrow X_1$ and $Z_2 \hookrightarrow X_2$ are closed embeddings, and $\phi: Z_1 \xrightarrow{\sim} Z_2$ is an isomorphism. We will explain how to glue X_1 to X_2 along ϕ . The result will be called $X_1 \coprod_{\phi} X_2$.

17.4.13. *Motivating example.* Our motivating example is if $X_i = \operatorname{Spec} A_i$ and $Z_i = \operatorname{Spec} A_i/I_i$, and ϕ corresponds to $\phi^\#: A_2/I_2 \xrightarrow{\sim} A_1/I_1$. Then the result will be $\operatorname{Spec} R$, where R is the ring consisting of ordered pairs $(a_1, a_2) \in A_1 \times A_2$ that “agree via ϕ ”. More precisely, this is a fibered product of rings:

$$R := A_1 \times_{\phi^\#: A_1/I_1 \rightarrow A_2/I_2} A_2.$$

17.4.14. *The general construction, as a locally ringed space.* In our general situation, we might wish to cover X_1 and X_2 by open charts of this form. We would then have to worry about gluing and choices, so to avoid this, we instead first construct $X_1 \coprod_{\phi} X_2$ as a locally ringed space. As a topological space, the definition is clear: we glue the underlying sets together along the underlying sets of $Z_1 \cong Z_2$, and topologize it so that a subset of $X_1 \coprod_{\phi} X_2$ is open if and only if its restrictions to X_1 and X_2 are both open. For convenience, let Z be the image of Z_1 (or equivalently Z_2) in $X_1 \coprod_{\phi} X_2$. We next define the stalk of the structure sheaf at any point $p \in X_1 \coprod_{\phi} X_2$. If $p \in X_i \setminus Z = (X_1 \coprod_{\phi} X_2) \setminus X_{3-i}$ (hopefully the meaning of this is clear), we define the stalk as $\mathcal{O}_{X_i, p}$. If $p \in X_1 \cap X_2$, we define the stalk to consist of elements $(s_1, s_2) \in \mathcal{O}_{X_1, p} \times \mathcal{O}_{X_2, p}$ which agree in $\mathcal{O}_{Z_1, p} \cong \mathcal{O}_{Z_2, p}$. The meaning of everything in this paragraph will be clear to you if you can do the following.

17.4.M. EXERCISE. Define the structure sheaf of $\mathcal{O}_{X_1 \coprod_{\phi} X_2}$ in terms of compatible germs. (What should it mean for germs to be compatible? Hint: for $z \in Z$, suppose we have open subsets U_1 of X_1 and U_2 of X_2 , with $U_1 \cap Z = U_2 \cap Z$, so U_1 and U_2 glue together to give an open subset U of $X_1 \coprod_{\phi} X_2$. Suppose we also have functions f_1 on U_1 and f_2 on U_2 that “agree on $U \cap Z$ ” — what does that mean? Then we declare that the germs of the “function on U obtained by gluing together f_1 and f_2 ” are compatible.) Show that the resulting ringed space is a locally ringed space.

We next want to show that the locally ringed space $X_1 \coprod_{\phi} X_2$ is a scheme. Clearly it is a scheme away from Z . We first verify a special case.

17.4.N. EXERCISE. Show that in Example 17.4.13 the construction of §17.4.14 indeed yields $\operatorname{Spec}(A_1 \times_{\phi^\#} A_2)$.

17.4.O. EXERCISE. In the general case, suppose $x \in Z$. Show that there is an affine open subset $\text{Spec } A_i \subset X_i$ such that $Z \cap \text{Spec } A_1 = Z \cap \text{Spec } A_2$. Then use Exercise 17.4.M to show that $X_1 \coprod_{\phi} X_2$ is a scheme in an open neighborhood of x , and thus a scheme.

17.4.15. Remarks.

(a) As the notation suggests, this is a fibered coproduct in the category of schemes, and indeed in the category of locally ringed spaces. We won't need this fact, but you can prove it if you wish; it isn't hard. Unlike the situation for products, fibered coproducts don't exist in general in the category of schemes. Miraculously (and for reasons that are specific to schemes), the resulting cofibered diagram is *also* a fibered diagram. This has pleasant ramifications. For example, this construction "behaves well with respect to" (or "commutes with") flat base change.

(b) You might hope that if you have a single scheme X with two disjoint closed subschemes W' and W'' , and an isomorphism $W' \xrightarrow{\sim} W''$, then you should be able to glue X to itself along $W' \rightarrow W''$. This construction doesn't work, and indeed it may not be possible. You can still make sense of the quotient as an *algebraic space*, which we will not define here.

17.4.P. EXERCISE. We continue to use the notation X_i , ϕ , etc. Suppose we are working in the category of A -schemes.

(a) If X_1 and X_2 are universally closed, show that $X_1 \coprod_{\phi} X_2$ is as well.

(b) If X_1 and X_2 are separated, show that $X_1 \coprod_{\phi} X_2$ is as well.

(c) If X_1 and X_2 are finite type over a *Noetherian* ring A , show that $X_1 \coprod_{\phi} X_2$ is as well. (Hint: Reduce to the "affine" case of the Motivating Example 17.4.13. Choose generators x_1, \dots, x_n of A_1 , and y_1, \dots, y_n , such that x_i modulo I_1 agrees with y_i modulo I_2 via ϕ . Choose generators g_1, \dots, g_m of I_2 — here use Noetherianness of A . Show that (x_i, y_i) and $(0, g_i)$ generate $R \subset A_1 \times A_2$, as follows. Suppose $(a_1, a_2) \in R$. Then there is some polynomial m such that $a_1 = m(x_1, \dots, x_n)$. Hence $(a_1, a_2) - m((x_1, y_1), \dots, (x_n, y_n)) = (0, a'_2)$ for some $a'_2 \in I_2$. Then a'_2 can be written as $\sum_{i=1}^m \ell_i(y_1, \dots, y_n)g_i$. But then $(0, a'_2) = \sum_{i=1}^m \ell_i((x_1, y_1), \dots, (x_n, y_n))(0, g_i)$.)

Thus if X_1 and X_2 are proper, so is $X_1 \coprod_{\phi} X_2$.

On the other hand, we have just seen (§17.4.11) that if X_1 and X_2 are *projective*, then $X_1 \coprod_{\phi} X_2$ needn't be projective.

17.5 The Curve-to-Projective Extension Theorem

We now use the main theorem of the previous section, Theorem 17.4.1, to prove something useful and concrete. (In fact, we could have proved this far earlier — with a little cleverness you can replace the invocation of Theorem 17.4.1 by Exercise 7.3.M, and prove this as soon as you know about Discrete Valuation Rings.)

17.5.1. The Curve-to-Projective Extension Theorem. — Suppose C is a pure dimension 1 Noetherian scheme over an affine base $S = \text{Spec } A$, and $p \in C$ is a regular closed

point of it. Suppose Y is a projective S -scheme. Then any morphism $C \setminus \{p\} \rightarrow Y$ (of S -schemes) extends to all of C .

In practice, we will use this theorem when $S = \operatorname{Spec} k$, and C is a k -variety. The only reason we assume S is affine is because we won't know the meaning of "projective S -scheme" until we know what a projective morphism is (§18.3). But the proof below extends immediately to general S once we know the meaning of the statement.

Note that if such an extension exists, then it is unique: the nonreduced locus of C is a closed subset (Exercise 6.5.N). Hence by replacing C by an open neighborhood of p that is reduced, we can use the Reduced-to-Separated Theorem 11.4.2 that maps from reduced schemes to separated schemes are determined by their behavior on a dense open set. Alternatively, maps to a separated scheme can be extended over an effective Cartier divisor in at most one way (Exercise 11.4.G).

The following exercise shows that the hypotheses are necessary.

17.5.A. EXERCISE. In each of the following cases, prove that the morphism $C \setminus \{p\} \rightarrow Y$ cannot be extended to a morphism $C \rightarrow Y$.

(a) *Projectivity of Y is necessary.* Suppose $C = \mathbb{A}_k^1$, $p = 0$, $Y = \mathbb{A}_k^1$, and $C \setminus \{p\} \rightarrow Y$ is given by " $t \mapsto 1/t$ ".

(b) *One-dimensionality of C is necessary.* Suppose $C = \mathbb{A}_k^2$, $p = (0, 0)$, $Y = \mathbb{P}_k^1$, and $C \setminus \{p\} \rightarrow Y$ is given by $(x, y) \mapsto [x, y]$.

(c) *Non-singularity of C is necessary.* Suppose $C = \operatorname{Spec} k[x, y]/(y^2 - x^3)$, $p = 0$, $Y = \mathbb{P}_k^1$, and $C \setminus \{p\} \rightarrow Y$ is given by $(x, y) \mapsto [x, y]$.

We remark that by combining this (easy) theorem with the (hard) valuative criterion for properness (Theorem 13.7.6), one obtains a proof of the properness of projective space bypassing the (tricky but not hard) Fundamental Theorem of Elimination Theory 8.4.7 (see Exercise 13.7.F). Fancier remark: the valuative criterion of properness can be used to show that Theorem 17.5.1 remains true if Y is only required to be proper, but it requires some thought.

17.5.2. Central idea of proof. The central idea of the proof may be summarized as "clear denominators", and is illustrated by the following motivating example. Suppose you have a morphism from $\mathbb{A}^1 - \{0\}$ to projective space, and you wanted to extend it to \mathbb{A}^1 . Suppose the map was given by $t \mapsto [t^4 + t^{-3}, t^{-2} + 4t]$. Then of course you would "clear the denominators", and replace the map by $t \mapsto [t^7 + 1, t + 4t^4]$. Similarly, if the map was given by $t \mapsto [t^2 + t^3, t^2 + t^4]$, you would divide by t^2 , to obtain the map $t \mapsto [1 + t, 1 + t^2]$.

Proof. Our plan is to maneuver ourselves into the situation where we can apply the idea of §17.5.2. We begin with some quick reductions. The nonreduced locus of C is closed (Exercise 6.5.N) and doesn't contain p , so by replacing C by an appropriate open neighborhood of p , we may assume that C is reduced and affine.

We next reduce to the case where $Y = \mathbb{P}_A^n$. Choose a closed embedding $Y \rightarrow \mathbb{P}_A^n$. If the result holds for \mathbb{P}^n , and we have a morphism $C \rightarrow \mathbb{P}^n$ with $C \setminus \{p\}$ mapping to Y , then C must map to Y as well. Reason: we can reduce to the case where the source is an affine open subset, and the target is $\mathbb{A}_A^n \subset \mathbb{P}_A^n$ (and hence affine). Then the functions vanishing on $Y \cap \mathbb{A}_A^n$ pull back to functions that vanish at the

generic points of the irreducible components of C and hence vanish everywhere on C (using reducedness of C), i.e., C maps to Y .

Choose a uniformizer $t \in \mathfrak{m} - \mathfrak{m}^2$ in the local ring of C at p . As t is a function in some neighborhood of p , we may assume that t is a function on C (by replacing C by this neighborhood of p). Then $V(t)$ contains p (as $t \in \mathfrak{m}$), but $V(t)$ does not contain the component of C containing p (as $t \notin \mathfrak{m}^2$), so by replacing C by an affine open neighborhood of p in $(C \setminus V(t)) \cup \{p\}$ — note that $(C \setminus V(t)) \cup \{p\} \subset C$ is open because C (almost) has the cofinite topology (§3.4.1) — we may assume that p is the only zero of the function t (and of course t vanishes at p with multiplicity 1).

We have a map $C \setminus \{p\} \rightarrow \mathbb{P}_A^n$, which by Theorem 17.4.1 corresponds to a line bundle \mathcal{L} on $C \setminus \{p\}$ and $n + 1$ sections of it with no common zeros in $C \setminus \{p\}$. Let U be a nonempty open set of $C \setminus \{p\}$ on which $\mathcal{L} \cong \mathcal{O}$. Then by replacing C by $U \cup \{p\}$, we interpret the map to \mathbb{P}^n as $n + 1$ rational functions f_0, \dots, f_n , defined away from p , with no common zeros away from p . Let $N = \min_i (\text{val}_p f_i)$. Then $t^{-N} f_0, \dots, t^{-N} f_n$ are $n + 1$ functions with no common zeros. Thus they determine a morphism $C \rightarrow \mathbb{P}_A^n$ extending $C \setminus \{p\} \rightarrow \mathbb{P}_A^n$ as desired. \square

17.5.B. EXERCISE (USEFUL PRACTICE). Suppose X is a Noetherian k -scheme, and Z is an irreducible codimension 1 subvariety whose generic point is a regular point of X (so the local ring $\mathcal{O}_{X,Z}$ is a discrete valuation ring). Suppose $\pi: X \dashrightarrow Y$ is a rational map to a projective k -scheme. Show that the domain of definition of the rational map includes a dense open subset of Z . In other words, rational maps from Noetherian k -schemes to projective k -schemes can be extended over regular codimension 1 sets. (We have already seen this principle in action — see Exercise 7.5.I on the Cremona transformation.)

17.6 Ample and very ample line bundles

Suppose $\pi: X \rightarrow \text{Spec } A$ is a proper morphism, and \mathcal{L} is an invertible sheaf on X . (The case when A is a field is the one of most immediate interest.) We say that \mathcal{L} is **very ample over A** or **π -very ample**, or **relatively very ample** if $X \cong \text{Proj } S_\bullet$ where S_\bullet is a finitely generated graded ring over A generated in degree 1 (Definition 4.5.6), and $\mathcal{L} \cong \mathcal{O}_{\text{Proj } S_\bullet}(1)$. One often just says **very ample** if the structure morphism is clear from the context. Note that the existence of a very ample line bundle implies that X is projective.

17.6.A. EASY BUT IMPORTANT EXERCISE (EQUIVALENT DEFINITION OF VERY AMPLE OVER A). Suppose $\pi: X \rightarrow \text{Spec } A$ is proper, and \mathcal{L} is an invertible sheaf on X . Show that \mathcal{L} is very ample if and only if there exist a finite number of global sections s_0, \dots, s_n of \mathcal{L} , with no common zeros, such that the morphism

$$[s_0, \dots, s_n]: X \rightarrow \mathbb{P}_A^n$$

is a closed embedding.

17.6.B. EASY EXERCISE (VERY AMPLE IMPLIES BASE-POINT-FREE). Show that a very ample invertible sheaf \mathcal{L} on a proper A -scheme must be base-point-free. In other words, show that for any point p on the scheme, there is a section of \mathcal{L} not vanishing there.

17.6.C. EXERCISE (VERY AMPLE \otimes BASE-POINT-FREE IS VERY AMPLE, HENCE VERY AMPLE \otimes VERY AMPLE IS VERY AMPLE). Suppose \mathcal{L} and \mathcal{M} are invertible sheaves on a proper A -scheme X , and \mathcal{L} is very ample over A and \mathcal{M} is base-point-free, then $\mathcal{L} \otimes \mathcal{M}$ is very ample. (Hint: \mathcal{L} gives a closed embedding $X \hookrightarrow \mathbb{P}^m$, and \mathcal{M} gives a morphism $X \rightarrow \mathbb{P}^n$. Show that the product map $X \rightarrow \mathbb{P}^m \times \mathbb{P}^n$ is a closed embedding, using the Cancellation Theorem 11.2.1 for closed embeddings on $X \rightarrow \mathbb{P}^m \times \mathbb{P}^n \rightarrow \mathbb{P}^m$. Finally, consider the composition $X \hookrightarrow \mathbb{P}^m \times \mathbb{P}^n \hookrightarrow \mathbb{P}^{mn+m+n}$, where the last closed embedding is the Segre embedding.)

17.6.D. EXERCISE (VERY AMPLE \boxtimes VERY AMPLE IS VERY AMPLE, CF. EXAMPLE 17.4.9). Suppose X and Y are proper A -schemes, and \mathcal{L} (resp. \mathcal{M}) is a very ample invertible sheaf on X (resp. Y). If $\pi_X: X \times_A Y \rightarrow X$ and $\pi_Y: X \times_A Y \rightarrow Y$ are the usual projections, show that $\pi_X^* \mathcal{L} \otimes \pi_Y^* \mathcal{M}$ (also known as $\mathcal{L} \boxtimes \mathcal{M}$, see §17.4.9) is very ample on $X \times_A Y$.

17.6.1. Definition. We say an invertible sheaf \mathcal{L} on a proper A -scheme X is **ample over A** or **π -ample** (where $\pi: X \rightarrow \operatorname{Spec} A$ is the structure morphism), or **relatively ample** if one of the following equivalent conditions holds.

17.6.2. Theorem. — Suppose $\pi: X \rightarrow \operatorname{Spec} A$ is proper, and \mathcal{L} is an invertible sheaf on X . The following are equivalent.

- (a) (projective geometry) For some $N > 0$, $\mathcal{L}^{\otimes N}$ is very ample over A .
- (a') For all $n \gg 0$, $\mathcal{L}^{\otimes n}$ is very ample over A .
- (b) (global generation) For all finite type quasicoherent sheaves \mathcal{F} , there is an n_0 such that for $n \geq n_0$, $\mathcal{F} \otimes \mathcal{L}^{\otimes n}$ is globally generated.
- (c) (topology) As f runs over all the global sections of $\mathcal{L}^{\otimes n}$ (over all $n > 0$), the open subsets $X_f = \{p \in X : f(p) \neq 0\}$ form a base for the topology of X .
- (c') As f runs over the global sections of $\mathcal{L}^{\otimes n}$ (over all $n > 0$), those open subsets X_f which are affine form a base for the topology of X .

(Variants of this Theorem 17.6.2 in the “absolute” and “relative” settings will be given in Theorems 17.6.6 and 18.3.9 respectively.)

Properties (a) and (a') relate to projective geometry, and property (b) relates to global generation (stalks). Properties (c) and (c') are somehow more topological, and while they may seem odd, they will provide the connection between (a)/(a') and (b). Note that (b), (c) and (c') make no reference to the structure morphism π . We will later (Theorem 19.7.1) meet a cohomological criterion (due, unsurprisingly, to Serre). The Kodaira Embedding Theorem (see, for example, [GH1, p. 181]) also gives a criterion for ampleness in the complex category: if X is a compact complex manifold, then an invertible sheaf \mathcal{L} on X is ample if and only if it admits a Hermitian metric with positive curvature everywhere.

The different flavor of these conditions gives some indication that ampleness is better-behaved than very ampleness in a number of ways. We mention without proof another property: if $\pi: X \rightarrow T$ is a finitely presented proper morphism, and \mathcal{L} is an invertible sheaf on X , then those points on T where \mathcal{L} is ample on the fiber form an open subset of T . Furthermore, on this open subset, \mathcal{L} is relatively ample over the base. We won't use these facts (proved in [Gr-EGA, IV₃.9.6.4]), but they are good to know.

Before getting to the proof of Theorem 17.6.2, we give some sample applications. We begin by noting that the fact that (a) implies (b) gives Serre's Theorem A (Theorem 16.3.8).

17.6.E. IMPORTANT EXERCISE. Suppose \mathcal{L} and \mathcal{M} are invertible sheaves on a proper A -scheme X , and \mathcal{L} is ample. Show that $\mathcal{L}^{\otimes n} \otimes \mathcal{M}$ is very ample for $n \gg 0$. (Hint: use both (a) and (b) of Theorem 17.6.2, and Exercise 17.6.C.)

17.6.F. IMPORTANT EXERCISE. Show that every line bundle on a projective A -scheme X is the difference of two very ample line bundles. More precisely, for any invertible sheaf \mathcal{L} on X , we can find two very ample invertible sheaves \mathcal{M} and \mathcal{N} such that $\mathcal{L} \cong \mathcal{M} \otimes \mathcal{N}^\vee$. (Hint: use the previous Exercise.)

17.6.G. IMPORTANT EXERCISE (USED REPEATEDLY). Suppose $\pi: X \rightarrow Y$ is a finite morphism of proper A -schemes, and \mathcal{L} is an ample line bundle on Y . Show that $\pi^*\mathcal{L}$ is ample on X . Hint: use the criterion of Theorem 17.6.2(b). Suppose \mathcal{F} is a finite type quasicoherent sheaf on X . We wish to show that $\mathcal{F} \otimes (\pi^*\mathcal{L})^{\otimes n}$ is globally generated for $n \gg 0$. Note that $(\pi_*\mathcal{F}) \otimes \mathcal{L}^{\otimes n}$ is globally generated for $n \gg 0$ by ampleness of \mathcal{L} on Y , i.e., there exists a surjection

$$\mathcal{O}_Y^{\oplus I} \twoheadrightarrow (\pi_*\mathcal{F}) \otimes \mathcal{L}^{\otimes n},$$

where I is some index set. Show that

$$\mathcal{O}_X^{\oplus I} \xleftarrow{\sim} \pi^*(\mathcal{O}_Y^{\oplus I}) \longrightarrow \pi^*(\pi_*\mathcal{F} \otimes \mathcal{L}^{\otimes n})$$

is surjective. Pullback π^* preserves tensor products (Theorem 17.3.7(9)), so we have an isomorphism $\pi^*(\pi_*\mathcal{F} \otimes \mathcal{L}^{\otimes n}) \xrightarrow{\sim} \pi^*(\pi_*\mathcal{F}) \otimes (\pi^*\mathcal{L})^{\otimes n}$. Show (using only affineness of π) that $\pi^*\pi_*\mathcal{F} \rightarrow \mathcal{F}$ is surjective. Connect these pieces together to describe a surjection

$$\mathcal{O}_X^{\oplus I} \twoheadrightarrow \mathcal{F} \otimes (\pi^*\mathcal{L})^{\otimes n}.$$

(Remark for those who have read about ampleness in the absolute setting in §17.6.5: the argument applies in that situation, i.e., with “proper A -schemes” changed to “schemes”, without change. The only additional thing to note is that ampleness of \mathcal{L} on Y implies that Y is quasicompact from the definition, and separated from Theorem 17.6.6(d). A relative version of this result appears in §18.3.8. It can be generalized even further, with “ π finite” replaced by “ π quasiaffine” — to be defined in §18.3.11 — see [Gr-EGA, II.5.1.12].)

17.6.H. EXERCISE (AMPLE \otimes AMPLE IS AMPLE, AMPLE \otimes BASE-POINT-FREE IS AMPLE). Suppose \mathcal{L} and \mathcal{M} are invertible sheaves on a proper A -scheme X , and \mathcal{L} is ample. Show that if \mathcal{M} is ample or base-point-free, then $\mathcal{L} \otimes \mathcal{M}$ is ample.

17.6.I. LESS IMPORTANT EXERCISE (AMPLE \boxtimes AMPLE IS AMPLE). Solve Exercise 17.6.D with “very ample” replaced by “ample”.

17.6.3. Proof of Theorem 17.6.2 in the case X is Noetherian. Noetherian hypotheses are used at only one point in the proof, and we explain how to remove them, and give a reference for the details.

Obviously, (a') implies (a).

Clearly $(c') \text{ implies } (c)$. We now show that $(c) \text{ implies } (c')$. Suppose we have a point p in an open subset U of X . We seek an *affine* X_f containing p and contained in U . By shrinking U , we may assume that U is affine. From (c) , U contains some X_f containing p . But this X_f is affine, as it is the complement of the vanishing locus of a section of a line bundle on an affine scheme (Exercise 8.3.E), so (c') holds. Note for future reference that the equivalence of (c) and (c') did not require the hypothesis of properness.

We next show that $(a) \text{ implies } (c)$. We embed X in projective space by some power of \mathcal{L} . Given a closed subset $Z \subset X$, and a point p of the complement $X \setminus Z$, we seek a section of some $\mathcal{L}^{\otimes N}$ that vanishes on Z and not on p . The existence of such a section follows from the fact that $V(I(Z)) = Z$ (Exercise 4.5.K(c)): there is some element of $I(Z)$ that does not vanish on p .

We next show that $(b) \text{ implies } (c)$. Suppose we have a point p in an open subset U of X . We seek a section of $\mathcal{L}^{\otimes N}$ that doesn't vanish at p , but vanishes on $X \setminus U$. Let \mathcal{I} be the sheaf of ideals of functions vanishing on $X \setminus U$ (the quasicoherent sheaf of ideals cutting out $X \setminus U$, with reduced structure). As X is Noetherian, \mathcal{I} is finite type, so by (b) , $\mathcal{I} \otimes \mathcal{L}^{\otimes N}$ is generated by global sections for some N , so there is some section of it not vanishing at p . (*Noetherian note:* This is the only part of the argument where we use Noetherian hypotheses. They can be removed as follows. Show that for a quasicompact quasiseparated scheme, every ideal sheaf is generated by its finite type subideal sheaves. Indeed, any quasicoherent sheaf on a quasicompact quasiseparated scheme is the union of its finite type quasicoherent subsheaves. This can be proved by hand, and is a suitable but hard exercise, see [GW, Cor. 10.50]. One of these finite type ideal sheaves doesn't vanish at p ; use this as \mathcal{I} instead.)

We now have to start working harder.

We next show that $(c') \text{ implies } (b)$. We wish to show that $\mathcal{F} \otimes \mathcal{L}^{\otimes n}$ is globally generated for $n \gg 0$.

We first show that (c') implies that for some $N > 0$, $\mathcal{L}^{\otimes N}$ is globally generated, as follows. For each closed point $p \in X$, there is some $f \in \Gamma(X, \mathcal{L}^{\otimes N(p)})$ not vanishing at p , so $p \in X_f$. (Don't forget that quasicompact schemes have closed points, Exercise 5.1.E!) As p varies, these X_f cover all of X . Use quasicompactness of X to select a finite number of these X_f that cover X . To set notation, say these are X_{f_1}, \dots, X_{f_n} , where $f_i \in \Gamma(X, \mathcal{L}^{\otimes N_i})$. By replacing f_i with $f_i^{\otimes (\prod_j N_j)/N_i}$, we may assume that they are all sections of the *same* power $\mathcal{L}^{\otimes N}$ of \mathcal{L} ($N = \prod_j N_j$). Then $\mathcal{L}^{\otimes N}$ is generated by these global sections.

We next show that it suffices to show that for all finite type quasicoherent sheaves \mathcal{F} , $\mathcal{F} \otimes \mathcal{L}^{\otimes mN}$ is globally generated for $m \gg 0$. For if we knew this, we could apply it to $\mathcal{F}, \mathcal{F} \otimes \mathcal{L}, \dots, \mathcal{F} \otimes \mathcal{L}^{\otimes (N-1)}$ (a finite number of times), and the result would follow. For this reason, we can replace \mathcal{L} by $\mathcal{L}^{\otimes N}$. In other words, to show that (c') implies (b) , we may also assume the additional hypothesis that \mathcal{L} is globally generated.

For each closed point p , choose an affine open neighborhood of the form X_f , using (c') . Then $\mathcal{F}|_{X_f}$ is generated by a finite number of global sections (Easy Exercise 16.3.B — a finite type quasicoherent sheaf on $\text{Spec } A$ corresponds to a finitely generated A -module). By Exercise 16.3.A, each of these generators can be expressed as a quotient of a section (over X) of $\mathcal{F} \otimes \mathcal{L}^{\otimes M(p)}$ by $f^{M(p)}$. (Note: we

can take a single $M(p)$ for each p .) Then $\mathcal{F} \otimes \mathcal{L}^{\otimes M(p)}$ is globally generated at p by a finite number of global sections. By Exercise 16.3.D(b), $\mathcal{F} \otimes \mathcal{L}^{\otimes M(p)}$ is globally generated at all points in some open neighborhood U_p of p . As \mathcal{L} is also globally generated, this implies that $\mathcal{F} \otimes \mathcal{L}^{\otimes M'}$ is globally generated at all points of U_p for $M' \geq M(p)$ (cf. Easy Exercise 16.3.C). From quasicompactness of X , a finite number of these U_p cover X , so we are done (by taking the maximum of these $M(p)$).

Our penultimate step is to show that (c') implies (a). Our goal is to assume (c'), and to find sections of some $\mathcal{L}^{\otimes N}$ that embeds X into projective space. Choose a cover of (quasicompact) X by n affine open subsets X_{a_1}, \dots, X_{a_n} , where a_1, \dots, a_n are all sections of powers of \mathcal{L} . By replacing each section with a suitable power, we may assume that they are all sections of the *same* power of \mathcal{L} , say $\mathcal{L}^{\otimes N}$. Say $X_{a_i} = \text{Spec } A_i$, where (using that π is finite type) $A_i = \text{Spec } A[a_{i1}, \dots, a_{ij_i}]/I_i$. By Exercise 16.3.A, each a_{ij} is of the form $s_{ij}/a_i^{m_{ij}}$, where $s_{ij} \in \Gamma(X, \mathcal{L}^{\otimes Nm_{ij}})$ (for some m_{ij}). Let $m = \max_{i,j} m_{ij}$. Then for each i, j , $a_{ij} = (s_{ij} a_i^{m-m_{ij}})/a_i^m$. For convenience, let $b_i = a_i^m$, and $b_{ij} = s_{ij} a_i^{m-m_{ij}}$; these are all global sections of $\mathcal{L}^{\otimes mN}$. Now consider the linear series generated by the b_i and b_{ij} . As the $D(b_i) = X_{a_i}$ cover X , this linear series is base-point-free, and hence (by Exercise 16.3.G) gives a morphism to \mathbb{P}^Q (where $Q = \#b_i + \#b_{ij} - 1$). Let $x_1, \dots, x_n, \dots, x_{ij}, \dots$ be the projective coordinates on \mathbb{P}^Q , so $f^*x_i = b_i$, and $f^*x_{ij} = b_{ij}$. Then the morphism of affine schemes $X_{a_i} \rightarrow D(x_i)$ is a closed embedding, as the associated maps of rings is a surjection (the generator a_{ij} of A_i is the image of x_{ij}/x_i).

At this point, we note for future reference that we have shown the following. If $X \rightarrow \text{Spec } A$ is finite type, and \mathcal{L} satisfies (c)=(c'), then X is an open embedding into a projective A -scheme. (We did not use separatedness.) We conclude our proof that (c') implies (a) by using properness to show that the image of this open embedding into a projective A -scheme is in fact closed, so X is a projective A -scheme.

Finally, we note that (a) and (b) together imply (a'): if $\mathcal{L}^{\otimes N}$ is very ample (from (a)), and $\mathcal{L}^{\otimes n}$ is base-point-free for $n \geq n_0$ (from (b)), then $\mathcal{L}^{\otimes n}$ is very ample for $n \geq n_0 + N$ by Exercise 17.6.C. \square

17.6.4. ★★ Semiample line bundles. Just as an invertible sheaf is ample if some tensor power of it is very ample, an invertible sheaf \mathcal{L} is said to be **semiample** if some tensor power of it is base-point-free. (Translation: \mathcal{L} is ample if some power gives a closed embedding into projective space, and \mathcal{L} is semiample if some power gives just a *morphism* to projective space.) We won't use this notion.

17.6.5. ★ Ampleness in the absolute setting. (We will not use this section in any serious way later.) Note that global generation is already an absolute notion, i.e., is defined for a quasicohherent sheaf on a scheme, with no reference to any morphism. An examination of the proof of Theorem 17.6.2 shows that ampleness may similarly be interpreted in an absolute setting. We make this precise. Suppose \mathcal{L} is an invertible sheaf on a *quasicompact* scheme X . We say that \mathcal{L} is **ample** if as f runs over the sections of $\mathcal{L}^{\otimes n}$ ($n > 0$), the open subsets $X_f = \{p \in X : f(p) \neq 0\}$ form a base for the topology of X . (We emphasize that quasicompactness in X is part of the condition of ampleness of \mathcal{L} .) For example, (i) if X is an affine scheme, every invertible sheaf is ample, and (ii) if X is a projective A -scheme, $\mathcal{O}(1)$ is ample.

17.6.J. EASY EXERCISE (PROPERTIES OF ABSOLUTE AMPLENESS).

- (a) Fix a positive integer n . Show that \mathcal{L} is ample if and only if $\mathcal{L}^{\otimes n}$ is ample.
 (b) Show that if $Z \hookrightarrow X$ is a closed embedding, and \mathcal{L} is ample on X , then $\mathcal{L}|_Z$ is ample on Z .

The following result will give you some sense of how ampleness behaves. We will not use it, and hence omit the proof (which is given in [Stacks, tag 01Q3]). However, many parts of the proof are identical to (or generalize) the corresponding arguments in Theorem 17.6.2. The labeling of the statements parallels the labelling of the statements in Theorem 17.6.2.

17.6.6. Theorem (cf. Theorem 17.6.2). — *Suppose \mathcal{L} is an invertible sheaf on a quasicompact scheme X . The following are equivalent.*

- (b) X is quasiseparated, and for every finite type quasicoherent sheaf \mathcal{F} , there is an n_0 such that for $n \geq n_0$, $\mathcal{F} \otimes \mathcal{L}^{\otimes n}$ is globally generated.
- (c) As f runs over the global sections of $\mathcal{L}^{\otimes n}$ ($n > 0$), the open subsets $X_f = \{p \in X : f(p) \neq 0\}$ form a base for the topology of X (i.e., \mathcal{L} is ample).
- (c') As f runs over the global sections of $\mathcal{L}^{\otimes n}$ ($n > 0$), those open subsets X_f which are affine form a base for the topology of X .
- (d) Let S_\bullet be the graded ring $\bigoplus_{n \geq 0} \Gamma(X, \mathcal{L}^{\otimes n})$. Then the open sets X_f (with $f \in S_+$ homogeneous of positive degree) cover X , and the associated map $X \rightarrow \text{Proj } S_\bullet$ is an open embedding. (Warning: S_\bullet need not be finitely generated, and $\text{Proj } S_\bullet$ is not necessarily finite type.)

Part (d) implies that X is separated (and thus quasiseparated).

17.6.7. ★ Transporting global generation, base-point-freeness, and ampleness to the relative situation.

These notions can be “relativized”. We could do this right now, but we wait until §18.3.7, when we will have defined the notion of a projective morphism, and thus a “relatively very ample” line bundle.

17.7 ★ The Grassmannian as a moduli space

We first defined projective space inelegantly in §4.4.9, and in §17.4.2 we gave a clean (if perhaps surprising) functorial definition. Similarly, in §7.7, we gave a preliminary description of the Grassmannian. We are now in a position to give a better definition.

We describe the “Grassmannian functor” (which we also denote $G(k, n)$), then show that it is representable (§7.6.2). The construction works over an arbitrary base scheme, so we work over the final object $\text{Spec } \mathbb{Z}$. (You should think through what to change if you wish to work with, for example, complex schemes.) The functor is defined as follows. To a scheme B , we associate the set of *locally free rank k quotients of the rank n free sheaf*,

$$(17.7.0.1) \quad \mathcal{O}_B^{\oplus n} \twoheadrightarrow \mathcal{Q}$$

up to isomorphism. An isomorphism of two such quotients $\phi: \mathcal{O}_B^{\oplus n} \rightarrow \mathcal{Q} \rightarrow 0$ and $\phi': \mathcal{O}_B^{\oplus n} \rightarrow \mathcal{Q}' \rightarrow 0$ is an isomorphism $\sigma: \mathcal{Q} \xrightarrow{\sim} \mathcal{Q}'$ such that the diagram

$$\begin{array}{ccc} \mathcal{O}_B^{\oplus n} & \xrightarrow{\phi} & \mathcal{Q} \\ & \searrow \phi' & \downarrow \sigma \\ & & \mathcal{Q}' \end{array}$$

commutes. By Exercise 14.2.R(a), $\ker \phi$ is locally free of rank $n - k$. (Thus if you prefer, you can extend (17.7.0.1), and instead consider the functor to take B to short exact sequences

$$(17.7.0.2) \quad 0 \rightarrow \mathcal{L} \rightarrow \mathcal{O}_B^{\oplus n} \rightarrow \mathcal{Q} \rightarrow 0$$

of locally free sheaves over B , of ranks $n - k$, n , and k respectively.)

It may surprise you that we are considering rank k *quotients* of a rank n sheaf, not rank k *subobjects*, given that the Grassmannian should parametrize k -dimensional subspace of an n -dimensional space. This is done for several reasons. One is that the kernel of a surjective map of locally free sheaves must be locally free, while the cokernel of an injective map of locally free sheaves need not be locally free (Exercise 14.2.R(a) and (b) respectively). Another reason: we will see in §28.3.3 that the geometric incarnation of this problem indeed translates to this. We can already see a key example here: if $k = 1$, our definition yields one-dimensional quotients $\mathcal{O}_B^{\oplus n} \rightarrow \mathcal{L} \rightarrow 0$. But this is precisely the data of n sections of \mathcal{L} , with no common zeros, which by Theorem 17.4.1 (the functorial description of projective space) corresponds precisely to maps to \mathbb{P}^{n-1} , so the $k = 1$ case parametrizes what we want.

We now show that the Grassmannian functor is representable for given n and k .

So far we have described what the Grassmannian functor does on *objects* B ; to fully describe a functor, we need to say what it does to *morphisms* $B_1 \rightarrow B_2$. Can you see what the definition must be?

17.7.A. EXERCISE. Show that the Grassmannian functor is a Zariski sheaf (§10.1.8).

Hence by Key Exercise 10.1.H, to show that the Grassmannian functor is representable, we need only cover it with open subfunctors that are representable.

Throughout the rest of this section, a *k-subset* is a subset of $\{1, \dots, n\}$ of size k .

17.7.B. EXERCISE.

(a) Suppose I is a k -subset. Make the following statement precise: there is an open subfunctor $G(k, n)_I$ of $G(k, n)$ where the k sections of \mathcal{Q} corresponding to I (of the n sections of \mathcal{Q} coming from the surjection $\phi: \mathcal{O}_B^{\oplus n} \rightarrow \mathcal{Q}$) are linearly independent. Hint: in a trivializing open neighborhood of \mathcal{Q} , where we can choose an isomorphism $\mathcal{Q} \xrightarrow{\sim} \mathcal{O}_B^{\oplus k}$, ϕ can be interpreted as a $k \times n$ matrix M , and this locus is where the determinant of the $k \times k$ matrix consisting of the I columns of M is nowhere zero. Show that this locus behaves well under transitions between trivializations.

(b) Show that these open subfunctors $G(k, n)_I$ cover the functor $G(k, n)$ (as I runs through the k -subsets).

Hence by Exercise 10.1.H, to show $G(k, n)$ is representable, we need only show that $G(k, n)_I$ is representable for arbitrary I . After renaming the summands of $\mathcal{O}^{\oplus n}$, without loss of generality we may assume $I = \{1, \dots, k\}$.

17.7.C. EXERCISE. Show that $G(k, n)_{\{1, \dots, k\}}$ is represented by $\mathbb{A}^{k(n-k)}$ as follows. (You will have to make this precise.) Given a surjection $\phi: \mathcal{O}^{\oplus n} \rightarrow \mathcal{Q}$, let $\phi_i: \mathcal{O} \rightarrow \mathcal{Q}$ be the map from the i th summand of $\mathcal{O}^{\oplus n}$. (Really, ϕ_i is just a section of \mathcal{Q} .) For the open subfunctor $G(k, n)_I$, show that

$$\phi_1 \oplus \cdots \oplus \phi_k: \mathcal{O}^{\oplus k} \rightarrow \mathcal{Q}$$

is an isomorphism. For a scheme B , the bijection $G(k, n)_I(B) \leftrightarrow \text{Hom}(B, \mathbb{A}^{k(n-k)})$ is given as follows. Given an element $\phi \in G(k, n)_I(B)$, for $j \in \{k+1, \dots, n\}$, $\phi_j = a_{1j}\phi_1 + a_{2j}\phi_2 + \cdots + a_{kj}\phi_k$, where a_{ij} are functions on B . But $k(n-k)$ functions on B is the same as a map to $\mathbb{A}^{k(n-k)}$ (Exercise 7.6.E). Conversely, given $k(n-k)$ functions a_{ij} ($1 \leq i \leq k < j \leq n$), define a surjection $\phi: \mathcal{O}^{\oplus n} \rightarrow \mathcal{O}^{\oplus k}$ as follows: (ϕ_1, \dots, ϕ_k) is the identity, and $\phi_j = a_{1j}\phi_1 + a_{2j}\phi_2 + \cdots + a_{kj}\phi_k$ for $j > k$.

You have now shown that $G(k, n)$ is representable, by covering it with $\binom{n}{k}$ copies of $\mathbb{A}^{k(n-k)}$. (You might wish to relate this to the description you gave in §7.7.) In particular, the Grassmannian over a field is smooth, and irreducible of dimension $k(n-k)$. (The Grassmannian over any base is smooth over that base, because $\mathbb{A}_B^{k(n-k)} \rightarrow B$ is smooth, see §13.6.2.)

17.7.1. The universal exact sequence over the Grassmannian. Note that we have a tautological exact sequence

$$0 \rightarrow \mathcal{S} \rightarrow \mathcal{O}^{\oplus n} \rightarrow \mathcal{Q} \rightarrow 0.$$

17.7.2. The Plücker embedding.

By applying \wedge^k to a surjection $\phi: \mathcal{O}^{\oplus n} \rightarrow \mathcal{Q}$ (over an arbitrary base B), we get a surjection

$$\wedge^k \phi: \mathcal{O}^{\oplus \binom{n}{k}} \rightarrow \det \mathcal{Q}$$

(Exercise 14.2.U). But a surjection from a rank N free sheaf to a line bundle is the same as a map to \mathbb{P}^{N-1} (Theorem 17.4.1).

17.7.D. EXERCISE. Use this to describe a map $P: G(k, n) \rightarrow \mathbb{P}^{\binom{n}{k}-1}$. (This is just a tautology: a natural transformation of functors induces a map of the representing schemes. This is Yoneda's Lemma, although if you didn't do Exercise 1.3.Z, you may wish to do this exercise by hand. But once you do, you may as well go back to prove Yoneda's Lemma and do Exercise 1.3.Z, because the argument is just the same!)

17.7.E. EXERCISE. The projective coordinate x_i on $\mathbb{P}^{\binom{n}{k}-1}$ corresponding to the i th factor of $\mathcal{O}^{\oplus \binom{n}{k}}$ may be interpreted as the determinant of the map $\phi_I: \mathcal{O}^{\oplus k} \rightarrow \mathcal{Q}$, where the $\mathcal{O}^{\oplus k}$ consists of the summands of $\mathcal{O}^{\oplus n}$ corresponding to I . Make this precise.

17.7.F. EXERCISE. Show that the standard open set U_I of $\mathbb{P}^{\binom{n}{k}-1}$ corresponding to k -subset I (i.e., where the corresponding coordinate x_I doesn't vanish) pulls back to the open subscheme $G(k, n)_I \subset G(k, n)$. Denote this map $P_I: G(k, n)_I \rightarrow U_I$.

17.7.G. EXERCISE. Show that P_I is a closed embedding as follows. We may deal with the case $I = \{1, \dots, k\}$. Note that $G(k, n)_I$ is affine — you described it $\text{Spec } \mathbb{Z}[a_{ij}]_{1 \leq i \leq k < j \leq n}$ in Exercise 17.7.C. Also, U_I is affine, with coordinates $x_{I'/I}$, as I' varies over the other k -subsets. (The case $k = 1$ is the useful notation $x_{i/j}$ we used to name coordinates on charts for projective space, §4.4.9.) You want to show that the map

$$P_I^\sharp: \mathbb{Z}[x_{I'/I}]_{I' \subset \{1, \dots, n\}, |I'|=k} / (x_{I/I} - 1) \rightarrow \mathbb{Z}[a_{ij}]_{1 \leq i \leq k < j \leq n}$$

is a surjection. By interpreting the map $\phi: \mathcal{O}^{\oplus n} \rightarrow \mathcal{O}^{\oplus k}$ as a $k \times n$ matrix M whose left k columns are the identity matrix and whose remaining entries are a_{ij} ($1 \leq i \leq k < j \leq n$), interpret P_I^\sharp as taking $x_{I'/I}$ to the determinant of the $k \times k$ submatrix corresponding to the columns in I' . For each (i, j) (with $1 \leq i \leq k < j \leq n$), find some I' so that $x_{I'/I} \mapsto \pm a_{ij}$. (Let $I' = \{1, \dots, i-1, i+1, \dots, k, j\}$.)

Hence $G(k, n) \rightarrow \mathbb{P}^{\binom{n}{k}-1}$ is a closed embedding, so $G(k, n)$ is projective over \mathbb{Z} .

17.7.H. UNIMPORTANT EXERCISE. As an entertaining geometric consequence: if V is a vector space over a field, show that the “pure tensors in $\wedge^k V$ are pure in exactly one way”: if $v_1 \wedge \dots \wedge v_k = w_1 \wedge \dots \wedge w_k \neq 0$ in $\wedge^k V$, show that there is a $k \times k$ matrix of determinant 1 relating the v_i to the w_i .

17.7.3. The Plücker equations. The equations of $G(k, n) \rightarrow \mathbb{P}^{\binom{n}{k}-1}$ are particularly nice. There are quadratic relations among the $k \times k$ minors of a $k \times (n - k)$ matrix, called the Plücker equations. By our construction, they are equations satisfied by $G(k, n)$. It turns out that these equations cut out $G(k, n)$, and in fact generate the homogeneous ideal of $G(k, n)$, but this takes more work (see [F1, §9.1, Lem. 1]). We explore this in one example.

17.7.I. EXERCISE. Suppose v_1, v_2, v_3 , and v_4 are four vectors in a two-dimensional vector space V over some field. Show that

$$(v_1 \wedge v_2)(v_3 \wedge v_4) - (v_1 \wedge v_3)(v_2 \wedge v_4) + (v_1 \wedge v_4)(v_2 \wedge v_3) = 0.$$

The tricky part of this exercise is figuring out how to interpret multiplication (of Plücker coordinates) here.

17.7.J. EXERCISE. Note that the Plücker embedding embeds the Grassmannian $G(2, 4)$ into \mathbb{P}^5 .

(a) Show that $G(2, 4)$ is cut out by the quadratic equation

$$x_{12}x_{34} - x_{13}x_{24} + x_{14}x_{23} = 0.$$

(Hint: Use Exercise 17.7.I to show that the quadratic vanishes on $G(2, 4)$. But that isn't enough.)

(b) Show that every smooth quadric hypersurface in \mathbb{P}^5 over an algebraically closed field \bar{k} of characteristic not 2 is isomorphic to the Grassmannian (over \bar{k}). (For comparison, every smooth quadric hypersurface in $\mathbb{P}_{\bar{k}}^1$ is two points; every smooth

quadric hypersurface in $\mathbb{P}_{\bar{k}}^2$ is isomorphic to $\mathbb{P}_{\bar{k}}^1$, §7.5.9; and every smooth quadric hypersurface in $\mathbb{P}_{\bar{k}}^3$ is isomorphic to $\mathbb{P}_{\bar{k}}^1 \times_{\bar{k}} \mathbb{P}_{\bar{k}}^1$, Example 10.6.2.)

17.7.K. EXERCISE (RULINGS ON QUADRICS IN \mathbb{P}^5). Suppose $k = \bar{k}$. From Remark 9.3.10, we expect two three-dimensional families of planes in $G(2, 4)$ (interpreted as a hypersurface in \mathbb{P}^5 via the Plücker embedding, see Exercise 17.7.I). One of them may be described as follows: for each point $p \in \mathbb{P}_k^3$, we have a two-dimensional family of lines through p ; this is a plane in $G(2, 4)$. There is a three-dimensional family of planes corresponding to the choice of p . This is one of the two rulings. What is the other one? Prove (as rigorously as you can manage, given what you know) that these are both rulings.

17.7.4. Further discussion.

17.7.L. EXERCISE. Show that the group scheme GL_n acts on the Grassmannian $G(k, n)$. (The *action* of a group scheme appeared earlier in Exercise 7.6.S(a).) Hint: this is much more easily done with the language of functors, §7.6, than with the description of the Grassmannian in terms of patches, §7.7. (Exercise 17.4.L was the special case of projective space.)

17.7.M. ★★ EXERCISE (GRASSMANNIAN BUNDLES). Suppose \mathcal{F} is a rank n locally free sheaf on a scheme X . Define the **Grassmannian bundle** $G(k, \mathcal{F})$ over X . Intuitively, if \mathcal{F} is a varying family of n -dimensional vector spaces over X , $G(k, \mathcal{F})$ should parametrize k -dimensional quotients of the fibers. You may want to define the functor first, and then show that it is representable. Your construction will behave well under base change.

17.7.5. (Partial) flag varieties. The discussion here extends without change to partial flag varieties (§7.7.1), and the interested reader should think this through.

CHAPTER 18

Relative versions of Spec and Proj, and projective morphisms

In this chapter, we will use universal properties to define two useful constructions, *Spec* of a sheaf of algebras \mathcal{A} , and *Proj* of a sheaf of graded algebras \mathcal{A}_\bullet on a scheme X . These will both generalize (globalize) our constructions of *Spec* of A -algebras and *Proj* of graded A -algebras. We will see that affine morphisms are precisely those of the form $\text{Spec } \mathcal{A} \rightarrow X$, and so we will *define* projective morphisms to be those of the form $\text{Proj } \mathcal{A}_\bullet \rightarrow X$.

In both cases, our plan is to make a notion we know well over a ring work more generally over a scheme. The main issue is how to glue the constructions over each affine open subset together. The slick way we will proceed is to give a universal property, then show that the affine construction satisfies this universal property, then that the universal property behaves well with respect to open subsets, then to use the idea that lets us glue together the fibered product (or normalization) together to do all the hard gluing work. The most annoying part of this plan is finding the right universal property, especially in the *Proj* case.

18.1 Relative Spec of a (quasicoherent) sheaf of algebras

Given an A -algebra, B , we can take its *Spec* to get an affine scheme over *Spec* A : $\text{Spec } B \rightarrow \text{Spec } A$. We will now see a universal property description of a globalization of that notation. Consider an arbitrary scheme X , and a quasicoherent sheaf of algebras \mathcal{B} on it. We will define how to take *Spec* of this sheaf of algebras, and we will get a scheme $\text{Spec } \mathcal{B} \rightarrow X$ that is “affine over X ”, i.e., the structure morphism is an affine morphism. You can think of this in two ways.

18.1.1. First, and most concretely, for any affine open set $\text{Spec } A \subset X$, $\Gamma(\text{Spec } A, \mathcal{B})$ is some A -algebra; call it B . Then above $\text{Spec } A$, $\text{Spec } \mathcal{B}$ will be $\text{Spec } B$.

18.1.2. Second, it will satisfy a universal property. We could define the A -scheme $\text{Spec } B$ by the fact that morphisms to $\text{Spec } B$ (from an A -scheme W , over $\text{Spec } A$) correspond to maps of A -algebras $B \rightarrow \Gamma(W, \mathcal{O}_W)$ (this is our old friend Exercise 7.3.G). The universal property for $\beta: \text{Spec } \mathcal{B} \rightarrow X$ generalizes this. Given a

morphism $\mu: W \rightarrow X$, the X -morphisms $W \rightarrow \text{Spec } \mathcal{B}$ are in functorial (in W) bijection with morphisms α of sheaves of algebras making

(18.1.2.1)

$$\begin{array}{ccc} & \mathcal{O}_X & \\ \swarrow & & \searrow \\ \mathcal{B} & \xrightarrow{\alpha} & \mu_* \mathcal{O}_W \end{array}$$

commute. Here the map $\mathcal{O}_X \rightarrow \mu_* \mathcal{O}_W$ is that coming from the map of ringed spaces, and the map $\mathcal{O}_X \rightarrow \mathcal{B}$ comes from the \mathcal{O}_X -algebra structure on \mathcal{B} . (For experts: it needn't be true that $\mu_* \mathcal{O}_W$ is quasicoherent, but that doesn't matter. Non-experts should completely ignore this parenthetical comment.)

By universal property nonsense, these data determine $\beta: \text{Spec } \mathcal{B} \rightarrow X$ up to unique isomorphism, assuming that it exists.

Fancy translation: in the category of X -schemes, $\beta: \text{Spec } \mathcal{B} \rightarrow X$ represents the functor

$$(\mu: W \rightarrow X) \longmapsto \{(\alpha: \mathcal{B} \rightarrow \mu_* \mathcal{O}_W)\}.$$

18.1.A. EXERCISE. Show that if X is affine, say $\text{Spec } A$, and $\mathcal{B} = \tilde{B}$, where B is an A -algebra, then $\text{Spec } B \rightarrow \text{Spec } A$ satisfies this universal property. (Hint: Exercise 7.3.G.)

18.1.3. Proposition. — Suppose $\beta: \text{Spec } \mathcal{B} \rightarrow X$ satisfies the universal property for (X, \mathcal{B}) , and $U \hookrightarrow X$ is an open subset. Then $\beta|_U: \text{Spec } \mathcal{B} \times_X U = (\text{Spec } \mathcal{B})|_U \rightarrow U$ satisfies the universal property for $(U, \mathcal{B}|_U)$.

Proof. For convenience, let $V = \text{Spec } \mathcal{B} \times_X U$. A U -morphism $W \rightarrow V$ is the same as an X -morphism $W \rightarrow \text{Spec } \mathcal{B}$ (where by assumption $\mu: W \rightarrow X$ factors through U). By the universal property of $\text{Spec } \mathcal{B}$, this is the same information as a map $\mathcal{B} \rightarrow \mu_* \mathcal{O}_W$, which by the universal property definition of pullback (§17.3.4) is the same as $\mu^* \mathcal{B} \rightarrow \mathcal{O}_W$, which is the same information as $(\mu|_U)^* \mathcal{B} \rightarrow \mathcal{O}_W$. By adjointness again this is the same as $\mathcal{B}|_U \rightarrow (\mu|_U)_* \mathcal{O}_W$. \square

Combining the above Exercise and Proposition, we have shown the existence of $\text{Spec } \mathcal{B}$ in the case that X is an open subscheme of an affine scheme.

18.1.B. EXERCISE. Show the existence of $\text{Spec } \mathcal{B}$ in general, following the philosophy of our construction of the fibered product, normalization, and so forth.

We make some quick observations. First, $\text{Spec } \mathcal{B}$ can be “computed affinely locally on X ”. We also have an isomorphism $\phi: \mathcal{B} \xrightarrow{\sim} \beta_* \mathcal{O}_{\text{Spec } \mathcal{B}}$.

18.1.C. EXERCISE. Given an X -morphism γ :

$$\begin{array}{ccc} W & \xrightarrow{\gamma} & \text{Spec } \mathcal{B} \\ & \searrow \mu & \swarrow \beta \\ & X & \end{array}$$

show that the morphism α (in (18.1.2.1)) is the composition

$$\mathcal{B} \xrightarrow{\phi} \beta_* \mathcal{O}_{\text{Spec } \mathcal{B}} \longrightarrow \beta_* \gamma_* \mathcal{O}_W = \mu_* \mathcal{O}_W.$$

The Spec construction gives an important way to understand affine morphisms. Note that $\text{Spec } \mathcal{B} \rightarrow X$ is an affine morphism. The “converse” is also true:

18.1.D. EXERCISE. Show that if $\mu: Z \rightarrow X$ is an affine morphism, then we have a natural isomorphism $Z \xrightarrow{\sim} \text{Spec } \mu_* \mathcal{O}_Z$ of X -schemes.

Hence we can recover any affine morphism using the Spec construction. More precisely, a morphism is affine if and only if it is of the form $\mu: \text{Spec } \mathcal{B} \rightarrow X$ constructed above.

18.1.E. EXERCISE. Suppose $\mu: \text{Spec } \mathcal{B} \rightarrow X$ is such a morphism. Show that the category of quasicoherent sheaves on $\text{Spec } \mathcal{B}$ is equivalent to the category of quasicoherent sheaves on X with the structure of \mathcal{B} -modules (quasicoherent \mathcal{B} -modules on X).

This is useful if X is quite simple but $\text{Spec } \mathcal{B}$ is complicated. We will use this before long when $X \cong \mathbb{P}^1$, and $\text{Spec } \mathcal{B}$ is a more complicated curve.

18.1.F. EXERCISE (Spec BEHAVES WELL WITH RESPECT TO BASE CHANGE). Suppose $\mu: Z \rightarrow X$ is any morphism, and \mathcal{B} is a quasicoherent sheaf of algebras on X . Show that there is a natural isomorphism $Z \times_X \text{Spec } \mathcal{B} \xrightarrow{\sim} \text{Spec } \mu^* \mathcal{B}$. (If you think about, you will see that this generalizes the statement and possibly proof of Proposition 18.1.3.)

18.1.4. Definition. An important example of the Spec construction is the **total space of a finite rank locally free sheaf** \mathcal{F} , which we define to be $\text{Spec}(\text{Sym}^\bullet \mathcal{F}^\vee)$.

18.1.G. EXERCISE. Suppose \mathcal{F} is a locally free sheaf of rank n . Show that the total space of \mathcal{F} is a rank n *vector bundle*, i.e., that given any point $p \in X$, there is an open neighborhood $U \subset X$ such that

$$\text{Spec}(\text{Sym}^\bullet \mathcal{F}^\vee|_U) \cong \mathbb{A}_U^n.$$

Show that \mathcal{F} is isomorphic to the sheaf of sections of the total space $\text{Spec}(\text{Sym}^\bullet \mathcal{F}^\vee)$. (Possible hint: use transition functions.) For this reason, the total space is also called the **vector bundle associated to a locally free sheaf** \mathcal{F} . (Caution: some authors, e.g., [Stacks, tag 01M2], call $\text{Spec}(\text{Sym}^\bullet \mathcal{F})$, the *dual* of this vector bundle, the vector bundle associated to \mathcal{F} .)

In particular, if $\mathcal{F} = \mathcal{O}_X^{\oplus n}$, then $\text{Spec}(\text{Sym}^\bullet \mathcal{F}^\vee)$ is called \mathbb{A}_X^n , generalizing our earlier notions of \mathbb{A}_λ^n . As the notion of free sheaf behaves well with respect to base change, so does the notion of \mathbb{A}_X^n , i.e., given $X \rightarrow Y$, $\mathbb{A}_Y^n \times_Y X \cong \mathbb{A}_X^n$. (Aside: you may notice that the construction $\text{Spec } \text{Sym}^\bullet$ can be applied to any coherent sheaf \mathcal{F} (without dualizing, i.e., $\text{Spec}(\text{Sym}^\bullet \mathcal{F})$). This is sometimes called the **abelian cone** associated to \mathcal{F} . This concept can be useful, but we won’t need it.)

18.1.H. EXERCISE (THE TAUTOLOGICAL BUNDLE ON \mathbb{P}^n IS $\mathcal{O}(-1)$). Suppose k is a field. Define the subset $X \subset \mathbb{A}_k^{n+1} \times \mathbb{P}_k^n$ corresponding to “points of \mathbb{A}_k^{n+1} on the corresponding line of \mathbb{P}_k^n ”, so that the fiber of the map $\pi: X \rightarrow \mathbb{P}^n$ corresponding to a point $l = [x_0, \dots, x_n]$ is the line in \mathbb{A}_k^{n+1} corresponding to l , i.e., the scalar multiples of (x_0, \dots, x_n) . Show that $\pi: X \rightarrow \mathbb{P}_k^n$ is (the line bundle corresponding to) the invertible sheaf $\mathcal{O}(-1)$. (Possible hint: work first over the usual affine open

sets of \mathbb{P}_k^n , and figure out transition functions.) For this reason, $\mathcal{O}(-1)$ is often called the **tautological bundle** of \mathbb{P}_k^n (even over an arbitrary base, not just a field). (Side Remark: The projection $X \rightarrow \mathbb{A}_k^{n+1}$ is the blow-up of \mathbb{A}_k^{n+1} at the “origin”, see Exercise 10.3.F.)

18.2 Relative Proj of a sheaf of graded algebras

In parallel with the relative version *Spec* of *Spec*, we define a relative version of *Proj*, denoted *Proj* (called “relative Proj” or “sheaf Proj”), of a quasicoherent graded sheaf of algebras (satisfying some hypotheses) on a scheme X . We have already done the case where the base X is affine, in §4.5.7, using the regular Proj construction over a ring A . The elegant way to proceed would be to state the right universal property, and then use this cleverly to glue together the constructions over each affine, just as we did in the constructions of fibered product, normalization, and *Spec*. But because graded rings and graded modules make everything confusing, we do not do this. Instead we guiltily take a more pedestrian approach. (But the universal property can be made to work, see [Stacks, tag 01O0].)

18.2.A. EXERCISE (Proj COMMUTES WITH AFFINE BASE CHANGE). Suppose $A \rightarrow B$ is map of rings, and S_\bullet is a $\mathbb{Z}^{\geq 0}$ -graded ring over A .

(a) Give a canonical isomorphism

$$(18.2.0.1) \quad \alpha: \text{Proj}_B(S_\bullet \otimes_A B) \xrightarrow{\sim} (\text{Proj}_A S_\bullet) \times_{\text{Spec } A} \text{Spec } B$$

(b) (easy) Suppose X is a projective A -scheme (§4.5.9). Show that $X \times_{\text{Spec } A} \text{Spec } B$ is a projective B -scheme.

(c) Suppose S_\bullet is generated in degree 1, so $\mathcal{O}_{\text{Proj}_A S_\bullet}(1)$ is an invertible sheaf (§16.2). Clearly $S_\bullet \otimes_A B$ is generated in degree 1 as a B -algebra. Describe an isomorphism

$$\mathcal{O}_{\text{Proj}_B(S_\bullet \otimes_A B)}(1) \xrightarrow{\sim} \alpha^* \gamma^* \mathcal{O}_{\text{Proj}_A S_\bullet}(1),$$

where γ is the top morphism in the pullback diagram

$$\begin{array}{ccc} (\text{Proj}_A S_\bullet) \times_{\text{Spec } A} \text{Spec } B & \xrightarrow{\gamma} & \text{Proj}_A S_\bullet \\ \downarrow & & \downarrow \\ \text{Spec } B & \longrightarrow & \text{Spec } A \end{array}$$

Possible hint: transition functions.

We now give a general means of constructing schemes over X (from [Stacks, tag 01LH]), if we know what they should be over any affine open set, and how these behave under open embeddings of one affine open set into another.

18.2.B. EXERCISE. Suppose we are given a scheme X , and the following data:

- (i) For each affine open subset $U \subset X$, we are given some morphism $\pi_U: Z_U \rightarrow U$ (a “scheme over U ”).
- (ii) For each (open) inclusion of affine open subsets $V \subset U \subset X$, we are given an open embedding $\rho_V^U: Z_V \hookrightarrow Z_U$.

Assume this data satisfies:

- (a) for each $V \subset U \subset X$, ρ_V^U induces an isomorphism $Z_V \xrightarrow{\sim} \pi_U^{-1}(V)$ of schemes over V , and
- (b) whenever $W \subset V \subset U \subset X$ are three nested affine open subsets, $\rho_W^U = \rho_V^U \circ \rho_W^V$.

Show that there exists an X -scheme $\pi: Z \rightarrow X$, and isomorphisms $i_U: \pi^{-1}(U) \xrightarrow{\sim} Z_U$ for each affine open set U , such that for nested affine open sets $V \subset U$, ρ_V^U agrees with the composition

$$Z_V \xrightarrow{i_V^{-1}} \pi^{-1}(V) \hookrightarrow \pi^{-1}(U) \xrightarrow{i_U} Z_U$$

Hint (cf. Exercise 4.4.A): construct Z first as a set, then as a topological space, then as a scheme. (Your construction will be independent of choices. Your solution will work in more general situations, for example when the category of schemes is replaced by ringed spaces, and when the affine open subsets are replaced by any base of the topology.)

18.2.C. IMPORTANT EXERCISE AND DEFINITION (relative Proj). Suppose $\mathcal{S}_\bullet = \bigoplus_{n \geq 0} \mathcal{S}_n$ is a quasicoherent sheaf of $\mathbb{Z}^{\geq 0}$ -graded algebras on a scheme X . Over each affine open subset $\text{Spec } A \cong U \subset X$, we have an U -scheme $\text{Proj}_A \mathcal{S}_\bullet(U) \rightarrow U$. Show that these can be glued together to form an X -scheme, which we call $\text{Proj}_X \mathcal{S}_\bullet$; we have a “structure morphism” $\beta: \text{Proj}_X \mathcal{S}_\bullet \rightarrow X$. (The structure morphism β is part of the definition.)

By the construction of Exercise 18.2.B, the preimage over any affine open set can be computed using the original Proj construction. (You may enjoy going back and giving constructions of X^{red} , the normalization of X , and Spec of a quasicoherent sheaf of \mathcal{O} -algebras using this idea. But there is a moral price to be paid by giving up the universal property.)

18.2.1. Ongoing (reasonable) hypotheses on \mathcal{S}_\bullet : “finite generation in degree 1”. The Proj construction is most useful when applied to an A -algebra S_\bullet satisfying some reasonable hypotheses (§4.5.6), notably when S_\bullet is a finitely generated $\mathbb{Z}^{\geq 0}$ -graded A -algebra, and ideally if it is generated in degree 1. For this reason, in the rest of the book, we will enforce these assumptions on \mathcal{S}_\bullet , once we make sense of them for quasicoherent sheaves of algebras. (If you later need to relax these hypotheses — for example, to keep the finite generation hypothesis but remove the “generation in degree 1” hypothesis — it will not be too difficult.) Precisely, **we now always require that (i) \mathcal{S}_\bullet is “generated in degree 1”, and (ii) \mathcal{S}_1 is finite type.** The cleanest way to make condition (i) precise is to require the natural map

$$\text{Sym}_{\mathcal{O}_X}^\bullet \mathcal{S}_1 \rightarrow \mathcal{S}_\bullet$$

to be surjective. (Important: implicit in (i) is that \mathcal{S}_0 is a quotient of \mathcal{O}_X .) Because the Sym^\bullet construction may be computed affine-locally (§14.2.13), we can check generation in degree 1 on any affine cover. The motivation for these hypotheses is the following construction.

18.2.D. IMPORTANT EXERCISE: $\mathcal{O}(1)$ ON $\text{Proj } \mathcal{S}_\bullet$. If \mathcal{S}_\bullet is finitely generated in degree 1 (Hypotheses 18.2.1), construct an invertible sheaf $\mathcal{O}_{\text{Proj } \mathcal{S}_\bullet}(1)$ on $\text{Proj } \mathcal{S}_\bullet$ that “restricts to $\mathcal{O}_{\text{Proj } \mathcal{S}_\bullet(\text{Spec } A)}(1)$ over each affine open subset $\text{Spec } A \subset X$ ”.

18.2.E. EXERCISE (“Proj COMMUTES WITH BASE CHANGE”). Suppose \mathcal{S}_\bullet is a quasicoherent sheaf of $\mathbb{Z}^{\geq 0}$ -graded algebras on X . Let $\rho: Z \rightarrow X$ be any morphism. Give a natural isomorphism

$$(\text{Proj } \rho^* \mathcal{S}_\bullet, \mathcal{O}_{\text{Proj } \rho^* \mathcal{S}_\bullet}(1)) \xrightarrow{\sim} (Z \times_X \text{Proj } \mathcal{S}_\bullet, \psi^* \mathcal{O}_{\text{Proj } \mathcal{S}_\bullet}(1))$$

where ψ is the “top” morphism in the base change diagram

$$\begin{array}{ccc} Z \times_X \text{Proj } \mathcal{S}_\bullet & \xrightarrow{\psi} & \text{Proj } \mathcal{S}_\bullet \\ \downarrow & & \downarrow \beta \\ Z & \xrightarrow{\rho} & X. \end{array}$$

18.2.2. Definition: π -very ample. Suppose $\pi: X \rightarrow Y$ is proper. If \mathcal{L} is an invertible sheaf on X , then we say that \mathcal{L} is **very ample (with respect to π)**, **relatively very ample**, or (awkwardly) **π -very ample** if we can write $X = \text{Proj}_Y \mathcal{S}_\bullet$ with $\mathcal{L} \cong \mathcal{O}(1)$, where \mathcal{S}_\bullet is a quasicoherent sheaf of algebras on Y satisfying Hypotheses 18.2.1 (“finite generation in degree 1”). (The notion of very ampleness can be extended to more general situations, see for example [Stacks, tag 01VM]. But this is of interest only to people with esoteric tastes.)

18.2.F. EXERCISE. Suppose \mathcal{S}_\bullet is finitely generated in degree 1 (Hypotheses 18.2.1). Describe a map of graded quasicoherent sheaves $\phi: \mathcal{S}_\bullet \rightarrow \bigoplus_n \beta_* \mathcal{O}(n)$ (β is the structure morphism, see Exercise 18.2.C). Hint: Exercise 16.4.C.

18.2.G. EXERCISE. Suppose \mathcal{L} is an invertible sheaf on X , and \mathcal{S}_\bullet is a quasicoherent sheaf of graded algebras on X generated in degree 1 (Hypotheses 18.2.1). Define $\mathcal{S}'_\bullet = \bigoplus_{n=0}^\infty (\mathcal{S}_n \otimes \mathcal{L}^{\otimes n})$. Then \mathcal{S}'_\bullet has a natural algebra structure inherited from \mathcal{S}_\bullet ; describe it. Give a natural isomorphism of “ X -schemes with line bundles”

$$(\text{Proj } \mathcal{S}'_\bullet, \mathcal{O}_{\text{Proj } \mathcal{S}'_\bullet}(1)) \xrightarrow{\sim} (\text{Proj } \mathcal{S}_\bullet, \mathcal{O}_{\text{Proj } \mathcal{S}_\bullet}(1) \otimes \beta^* \mathcal{L}),$$

where $\beta: \text{Proj } \mathcal{S}_\bullet \rightarrow X$ is the structure morphism. In other words, informally speaking, the Proj is the same, but the $\mathcal{O}(1)$ is twisted by \mathcal{L} .

18.2.3. Definition. If \mathcal{F} is a finite type quasicoherent sheaf on X , then $\text{Proj}(\text{Sym}^\bullet \mathcal{F})$ is called its **projectivization**, and is denoted $\mathbb{P}\mathcal{F}$. You can check that this construction behaves well with respect to base change. Define $\mathbb{P}_X^n := \mathbb{P}(\mathcal{O}_X^{\oplus (n+1)})$. (Then $\mathbb{P}_{\text{Spec } A}^n$ agrees with our earlier definition of \mathbb{P}_A^n , cf. Exercise 4.5.Q, and \mathbb{P}_X^n agrees with our earlier usage, see for example the proof of Theorem 11.5.5.) More generally, if \mathcal{F} is locally free of rank $n+1$, then $\mathbb{P}\mathcal{F}$ is a **projective bundle** or **\mathbb{P}^n -bundle** over X . By Exercise 18.2.G, if \mathcal{F} is a finite rank locally free sheaf on X , there is a canonical isomorphism $\mathbb{P}\mathcal{F} \xrightarrow{\sim} \mathbb{P}(\mathcal{L} \otimes \mathcal{F})$.

18.2.4. Example: ruled surfaces. If C is a regular curve and \mathcal{F} is locally free of rank 2, then $\mathbb{P}\mathcal{F}$ is called a **ruled surface** over C . If C is further isomorphic to

\mathbb{P}^1 , $\mathbb{P}\mathcal{F}$ is called a **Hirzebruch surface**. All vector bundles on \mathbb{P}^1 split as a direct sum of line bundles (see §19.5.5 for a proof), so each Hirzebruch surface is of the form $\mathbb{P}(\mathcal{O}(n_1) \oplus \mathcal{O}(n_2))$. By Exercise 18.2.G, this depends only on $n_2 - n_1$. The Hirzebruch surface $\mathbb{P}(\mathcal{O} \oplus \mathcal{O}(n))$ ($n \geq 0$) is often denoted \mathbb{F}_n . We will discuss the Hirzebruch surfaces in greater length in §21.2.9. We will see that the \mathbb{F}_n are all distinct in Exercise 21.2.Q.

18.2.H. EXERCISE. If \mathcal{S}_\bullet is finitely generated in degree 1 (Hypotheses 18.2.1), describe a canonical closed embedding

$$\begin{array}{ccc} \text{Proj } \mathcal{S}_\bullet & \xrightarrow{i} & \mathbb{P}\mathcal{S}_1 \\ & \searrow \beta & \swarrow \\ & X & \end{array}$$

and an isomorphism $\mathcal{O}_{\text{Proj } \mathcal{S}_\bullet}(1) \xrightarrow{\sim} i^* \mathcal{O}_{\mathbb{P}\mathcal{S}_1}(1)$ arising from the surjection $\text{Sym}^\bullet \mathcal{S}_1 \rightarrow \mathcal{S}_\bullet$.

18.2.I. EXERCISE. Suppose \mathcal{F} is a locally free sheaf of rank $n + 1$ on X . Exhibit a bijection between the set of sections $s: X \rightarrow \mathbb{P}\mathcal{F}$ of $\mathbb{P}\mathcal{F} \rightarrow X$ and the set of surjective homomorphisms $\mathcal{F} \rightarrow \mathcal{L} \rightarrow 0$ of \mathcal{F} onto invertible sheaves on X .

18.2.5. Remark (the relative version of the projective and affine cone). There is a natural morphism from $\text{Spec } \mathcal{S}_\bullet$ minus the zero-section to $\text{Proj } \mathcal{S}_\bullet$ (cf. Exercise 9.3.N). Just as $\text{Proj } \mathcal{S}_\bullet[T]$ contains a closed subscheme identified with $\text{Proj } \mathcal{S}_\bullet$ whose complement can be identified with $\text{Spec } \mathcal{S}_\bullet$ (Exercise 9.3.O), $\text{Proj } \mathcal{S}_\bullet[T]$ contains a closed subscheme identified with $\text{Proj } \mathcal{S}_\bullet$ whose complement can be identified with $\text{Spec } \mathcal{S}_\bullet$. You are welcome to think this through.

18.2.6. Remark. If you wish, you can describe (with proof) a universal property of $\text{Proj } \mathcal{S}_\bullet$. (You may want to describe a universal property of Proj first.) I recommend against it — a universal property should make your life easier, not harder. One possible universal property is given in [Stacks, tag 01NS].

18.3 Projective morphisms

In §18.1, we reinterpreted affine morphisms: $X \rightarrow Y$ is an affine morphism if there is an isomorphism $X \xrightarrow{\sim} \text{Spec } \mathcal{B}$ of Y -schemes for some quasicoherent sheaf of algebras \mathcal{B} on Y . We will *define* the notion of a projective morphism similarly.

You might think that because projectivity is such a classical notion, there should be some obvious definition, that is reasonably behaved. But this is not the case, and there are many possible variant definitions of projective (see [Stacks, tag 01W8]). All are imperfect, including the accepted definition we give here. Although projective morphisms are preserved by base change, we will manage to show that they are preserved by composition only when the target is quasicompact (Exercise 18.3.B), and we will only show that the notion is local on the target when we add the data of a line bundle, and even then only under locally Noetherian hypotheses (§18.3.4).

18.3.1. Definition. A morphism $\pi: X \rightarrow Y$ is **projective** if there is an isomorphism

$$\begin{array}{ccc} X & \xrightarrow{\sim} & \text{Proj } \mathcal{S}_\bullet \\ & \searrow \pi & \swarrow \\ & Y & \end{array}$$

for a quasicoherent sheaf of algebras \mathcal{S}_\bullet on Y (satisfying “finite generation in degree 1”, Hypotheses 18.2.1). We say X is a **projective Y -scheme**, or X is **projective over Y** . Using Exercise 7.4.D, this generalizes the notion of a projective A -scheme.

18.3.2. Warnings. First, notice that $\mathcal{O}(1)$, an important part of the concept of Proj , is not mentioned in the definition. (I would prefer that it be part of the definition, but this isn’t accepted practice.) As a result, the notion of affine morphism is affine-local on the target, but the notion of projectivity of a morphism is not clearly affine-local on the target. (In Noetherian circumstances, with the additional data of the invertible sheaf $\mathcal{O}(1)$, it is, as we will see in §18.3.4. We will also later see an example showing that the property of being projective is *not* local, §25.7.7.)

Second, [Ha1, p. 103] gives a different definition of projective morphism; we follow the more general definition of Grothendieck. These definitions turn out to be the same in nice circumstances. (But finite morphisms are not always projective in the sense of [Ha1], while they *are* projective in our sense.)

18.3.A. EXERCISE.

(a) (*a useful characterization of projective morphisms*) Suppose $\pi: X \rightarrow Y$ is a morphism. Show that π is projective if and only if there exist a finite type quasicoherent sheaf \mathcal{S}_1 on Y , and a closed embedding $i: X \hookrightarrow \mathbb{P}_{\mathcal{S}_1}$ (over Y , i.e., commuting with the maps to Y). Hint: Exercise 18.2.H.

(b) (*a useful characterization of projective morphisms, with line bundle*) Suppose \mathcal{L} is an invertible sheaf on X , and $\pi: X \rightarrow Y$ is a morphism. Show that π is projective, with $\mathcal{O}(1) \cong \mathcal{L}$, if and only if there exist a finite type quasicoherent sheaf \mathcal{S}_1 on Y , a closed embedding $i: X \hookrightarrow \mathbb{P}_{\mathcal{S}_1}$ (over Y , i.e., commuting with the maps to Y), and an isomorphism $i^* \mathcal{O}_{\mathbb{P}_{\mathcal{S}_1}}(1) \xrightarrow{\sim} \mathcal{L}$.

(c) If furthermore Y admits an ample line bundle \mathcal{M} , show that π is projective if and only if there exists a closed embedding $i: X \hookrightarrow \mathbb{P}_Y^n$ (over Y) for some n . (If you wish, assume Y is proper over $\text{Spec } A$, so you can avoid the starred section §17.6.5.) Hint: the harder direction is the forward implication. Use the finite type quasicoherent sheaf \mathcal{S}_1 from (a). Tensor \mathcal{S}_1 with a high enough power of \mathcal{M} so that it is finitely globally generated (Theorem 17.6.6, or Theorem 17.6.2 in the proper setting), to obtain a surjection

$$\mathcal{O}_Y^{\oplus(n+1)} \twoheadrightarrow \mathcal{S}_1 \otimes \mathcal{M}^{\otimes N}.$$

Then use Exercise 18.2.G.

18.3.3. Definition: Quasiprojective morphisms. In analogy with projective and quasiprojective A -schemes (§4.5.9), one may define quasiprojective morphisms. If Y is *quasicompact*, we say that $\pi: X \rightarrow Y$ is a **quasiprojective morphism** if π can be expressed as a quasicompact open embedding into a scheme projective over Y . This is not a great notion, and we will not use it. (The general definition of quasiprojective is slightly delicate — see [Gr-EGA, II.5.3] — but we won’t need it.)

18.3.4. Properties of projective morphisms.

We start to establish a number of properties of projective morphisms. First, the property of a morphism being projective is clearly preserved by base change, as the *Proj* construction behaves well with respect to base change (Exercise 18.2.E). Also, projective morphisms are proper: properness is local on the target (Theorem 11.5.4(b)), and we saw earlier that projective A -schemes are proper over A (Theorem 11.5.5). In particular (by definition of properness), projective morphisms are separated, finite type, and universally closed.

Exercise 18.3.G (in a future optional section) implies that if $\pi: X \rightarrow Y$ is a proper morphism of locally Noetherian schemes, and \mathcal{L} is an invertible sheaf on X , the question of whether π is a projective morphism with \mathcal{L} as $\mathcal{O}(1)$ is local on Y .

18.3.B. IMPORTANT CHALLENGING EXERCISE (THE COMPOSITION OF PROJECTIVE MORPHISMS IS PROJECTIVE, IF THE FINAL TARGET IS QUASICOMPACT). Suppose $\pi: X \rightarrow Y$ and $\rho: Y \rightarrow Z$ are projective morphisms, and Z is quasicompact. Show that $\pi \circ \rho$ is projective. Hint: the criterion for projectivity given in Exercise 18.3.A(b) will be useful. (i) Deal first with the case where Z is affine. Build the following commutative diagram, thereby finding a closed embedding $X \hookrightarrow \mathbb{P}(\mathcal{F}^{\oplus n})$ over Z . In this diagram, all inclusions are closed embeddings, and all script fonts refer to finite type quasicoherent sheaves.

$$\begin{array}{ccccccc}
 X & \xrightarrow{\quad} & \mathbb{P}(\mathcal{E}) & \xrightarrow{(\dagger)} & \mathbb{P}_Z^{n-1} \times_Z Y & \xrightarrow{\quad} & \mathbb{P}_Z^{n-1} \times_Z \mathbb{P}(\mathcal{F}) \xrightarrow[\text{cf. Ex. 10.6.C, 10.6.D}]{\text{Segre}} \mathbb{P}(\mathcal{F}^{\oplus n}) \\
 & \searrow \pi & \downarrow & & \downarrow & & \downarrow \\
 & & Y & \xrightarrow{\quad} & \mathbb{P}(\mathcal{F}) & & \\
 & & \downarrow \rho & & \swarrow & \nearrow & \\
 & & Z & & & &
 \end{array}$$

Construct the closed embedding (\dagger) as follows. Suppose \mathcal{M} is the very ample line bundle on Y over Z . Then \mathcal{M} is ample, and so by Theorem 17.6.2, for $m \gg 0$, $\mathcal{E} \otimes \mathcal{M}^{\otimes m}$ is generated by a finite number of global sections. Suppose $\mathcal{O}_Y^{\oplus n} \twoheadrightarrow \mathcal{E} \otimes \mathcal{M}^{\otimes m}$ is the corresponding surjection. This induces a closed embedding $\mathbb{P}(\mathcal{E} \otimes \mathcal{M}^{\otimes m}) \hookrightarrow \mathbb{P}_Y^{n-1}$. But $\mathbb{P}(\mathcal{E} \otimes \mathcal{M}^{\otimes m}) \cong \mathbb{P}(\mathcal{E})$ (Exercise 18.2.G), and $\mathbb{P}_Y^{n-1} = \mathbb{P}_Z^{n-1} \times_Z Y$. (ii) Unwind this diagram to show that (for Z affine) if \mathcal{L} is π -very ample and \mathcal{M} is ρ -very ample, then for $m \gg 0$, $\mathcal{L} \otimes \pi^* \mathcal{M}^{\otimes m}$ is $(\rho \circ \pi)$ -very ample. Then deal with the general case by covering Z with a finite number of affines.

18.3.5. Finite morphisms are projective.

18.3.C. IMPORTANT EXERCISE: FINITE MORPHISMS ARE PROJECTIVE (CF. EXERCISE 8.3.I). Show that finite morphisms are projective as follows. Suppose $Z \rightarrow X$ is finite, so that $Z \cong \text{Spec } \mathcal{B}$ where \mathcal{B} is a finite type quasicoherent sheaf of algebras on X . Describe a sheaf of graded algebras \mathcal{S}_\bullet where $\mathcal{S}_0 \cong \mathcal{O}_X$ and $\mathcal{S}_n \cong \mathcal{B}$ for $n > 0$. Describe an X -isomorphism $Z \xrightarrow{\sim} \text{Proj } \mathcal{S}_\bullet$.

In particular, closed embeddings are projective. We have the sequence of implications for morphisms

$$\text{closed embedding} \implies \text{finite} \implies \text{projective} \implies \text{proper}.$$

We know that finite morphisms are projective (Exercise 18.3.C), and have finite fibers (Exercise 8.3.J). We will show the converse in Theorem 19.1.7, and state the extension to proper morphisms immediately after (although we won't prove it until Theorem 30.6.2).

18.3.6. Caution: *Consequences of projectivity not being “reasonable” in the sense of §8.1.* Because the property of being projective is preserved by base change (§18.3.4), and composition to *quasicompact targets* (Exercise 18.3.B), the property of being projective is “usually” preserved by products (Exercise 8.1.A): if $\pi: X \rightarrow Y$ and $\pi': X' \rightarrow Y'$ are projective, then so is $\pi \times \pi': X \times X' \rightarrow Y \times Y'$, so long as $Y \times Y'$ is quasicompact. Also, if you follow through the proof of the Cancellation Theorem 11.2.1, you will see that if $\pi: X \rightarrow Y$ is a morphism, $\rho: Y \rightarrow Z$ is a separated morphism (so the diagonal δ_ρ is a closed embedding and hence projective), and $\rho \circ \pi$ is projective, and Y is quasicompact, then π is projective.

18.3.D. EXERCISE. Show that a morphism (over $\text{Spec } k$) from a projective k -scheme to a quasicompact separated k -scheme is always projective. (Hint: the Cancellation Theorem 11.2.1 for projective morphisms; see also Caution 18.3.6.)

18.3.7. ★★ Global generation and (very) ampleness in the relative setting.

We extend the discussion of §16.3 to the relative setting, in order to give ourselves the language of relatively base-point-freeness. We won't use this discussion, so on a first reading you should jump directly to §18.4. But these ideas come up repeatedly in the research literature.

Suppose $\pi: X \rightarrow Y$ is a quasicompact quasiseparated morphism. If \mathcal{F} is a quasicoherent sheaf on X , we say that \mathcal{F} is **relatively globally generated** or **globally generated with respect to π** if the natural map of quasicoherent sheaves $\pi^* \pi_* \mathcal{F} \rightarrow \mathcal{F}$ is surjective. (Quasicompactness and quasiseparatedness are needed ensure that $\pi_* \mathcal{F}$ is a quasicoherent sheaf, Exercise 17.2.C.) But these hypotheses are not very restrictive. Global generation is most useful only in the quasicompact setting, and most people won't be bothered by quasiseparated hypotheses. Unimportant aside: these hypotheses can be relaxed considerably. If $\pi: X \rightarrow Y$ is a morphism of *locally ringed spaces* — not necessarily schemes — with no other hypotheses, and \mathcal{F} is a quasicoherent sheaf on X , then we say that \mathcal{F} is **relatively globally generated** or **globally generated with respect to π** if the natural map $\pi^* \pi_* \mathcal{F} \rightarrow \mathcal{F}$ of \mathcal{O}_X -modules is surjective.)

Thanks to our hypotheses, as the natural map $\pi^* \pi_* \mathcal{F} \rightarrow \mathcal{F}$ is a morphism of quasicoherent sheaves, the condition of being relatively globally generated is affine-local on Y .

Suppose now that \mathcal{L} is a locally free sheaf on X , and $\pi: X \rightarrow Y$ is a morphism. We say that \mathcal{L} is **relatively base-point-free** or **base-point-free with respect to π** if it is relatively globally generated.

18.3.E. EXERCISE. Suppose \mathcal{L} is a finite rank locally free sheaf on X , $\pi: X \rightarrow Y$ is a quasicompact separated morphism, and $\pi_* \mathcal{L}$ is finite type on Y . (We will later show in Grothendieck's Coherence Theorem 19.9.1 that this latter statement

is true if π is proper and Y is Noetherian. This is much easier if π is projective, see Theorem 19.8.1. We could work hard and prove it now, but it isn't worth the trouble.) Describe a canonical morphism $\psi: X \rightarrow \mathbb{P}\pi_*\mathcal{L}$. (Possible hint: this generalizes the fact that base-point-free line bundles give maps to projective space, so generalize that argument, see §16.3.5.)

We say that \mathcal{L} is **relatively ample** or **π -ample** or **relatively ample with respect to π** if for every affine open subset $\text{Spec } B$ of Y , $\mathcal{L}|_{\pi^{-1}(\text{Spec } B)}$ is ample on $\pi^{-1}(\text{Spec } B)$ over B , or equivalently (by §17.6.5), $\mathcal{L}|_{\pi^{-1}(\text{Spec } B)}$ is (absolutely) ample on $\pi^{-1}(\text{Spec } B)$. By the discussion in §17.6.5, if \mathcal{L} is ample then π is necessarily quasicompact, and (by Theorem 17.6.6) separated; if π is affine, then all invertible sheaves are ample; and if π is projective, then the corresponding $\mathcal{O}(1)$ is ample. By Exercise 17.6.J, \mathcal{L} is π -ample if and only if $\mathcal{L}^{\otimes n}$ is π -ample, and if $Z \hookrightarrow X$ is a closed embedding, then $\mathcal{L}|_Z$ is ample over Y .

From Theorem 17.6.6(d), we have a natural open embedding $X \rightarrow \text{Proj}_Y \oplus \pi_*\mathcal{L}^{\otimes d}$. (Do you see what this map is? Also, be careful: $\oplus \pi_*\mathcal{L}^{\otimes d}$ need not be a finitely generated graded sheaf of algebras, so we are using the *Proj* construction where one of the usual hypotheses doesn't hold.)

The notions of relative global generation and relative ampleness are most useful in the proper setting, because of Theorem 17.6.2.

18.3.8. Many statements of §16.3 carry over without change. For example, we have the following. Suppose $\pi: X \rightarrow Y$ is proper, \mathcal{F} and \mathcal{G} are quasicoherent sheaves on X , and \mathcal{L} and \mathcal{M} are invertible sheaves on X . If π is affine, then \mathcal{F} is relatively globally generated (from Easy Exercise 16.3.B). If \mathcal{F} and \mathcal{G} are relatively globally generated, so is $\mathcal{F} \otimes \mathcal{G}$ (Easy Exercise 16.3.C). If \mathcal{L} is π -very ample (Definition 18.2.2), then it is π -base-point-free (Easy Exercise 17.6.B). If \mathcal{L} is π -very ample, and \mathcal{M} is π -base-point-free (if for example it is π -very ample), then $\mathcal{L} \otimes \mathcal{M}$ is π -very ample (Exercise 17.6.C). Exercise 17.6.G extends immediately to show that if

$$\begin{array}{ccc} X & \xrightarrow{\pi} & Y \\ & \searrow \tau & \swarrow \rho \\ & S & \end{array}$$

is a finite morphism of S -schemes, and if \mathcal{L} is a ρ -ample invertible sheaf on Y , then $\pi^*\mathcal{L}$ is τ -ample.

By the nature of the statements, some of the statements of §16.3 require quasicompactness hypotheses on Y , or other patches. For example:

18.3.9. Theorem. — Suppose $\pi: X \rightarrow Y$ is proper, \mathcal{L} is an invertible sheaf on X , and Y is quasicompact. The following are equivalent.

- (a) For some $N > 0$, $\mathcal{L}^{\otimes N}$ is π -very ample.
- (a') For all $n \gg 0$, $\mathcal{L}^{\otimes n}$ is π -very ample.
- (b) For all finite type quasicoherent sheaves \mathcal{F} , there is an n_0 such that for $n \geq n_0$, $\mathcal{F} \otimes \mathcal{L}^{\otimes n}$ is relatively globally generated.
- (c) The invertible sheaf \mathcal{L} is π -ample.

18.3.F. EXERCISE. Prove Theorem 18.3.9 using Theorem 17.6.2. (Unimportant remark: The proof given of Theorem 17.6.2 used Noetherian hypotheses, but as stated there, they can be removed.)

After doing the above Exercise, it will be clear how to adjust the statement of Theorem 18.3.9 if you need to remove the quasicompactness assumption on Y .

18.3.G. EXERCISE (A USEFUL EQUIVALENT DEFINITION OF VERY AMPLENESS UNDER NOETHERIAN HYPOTHESES). Suppose $\pi: X \rightarrow Y$ is a proper morphism, Y is locally Noetherian (hence X is too, as π is finite type), and \mathcal{L} is an invertible sheaf on X . Suppose that you know that in this situation $\pi_*\mathcal{L}$ is finite type. (We will later show this, as described in Exercise 18.3.E.) Show that \mathcal{L} is very ample if and only if (i) \mathcal{L} is relatively base-point-free, and (ii) the canonical Y -morphism $i: X \rightarrow \mathbb{P}(\pi_*\mathcal{L})$ of Exercise 18.3.E is a closed embedding. Conclude that the notion of relative very ampleness is affine-local on Y (it may be checked on *any* affine cover Y), if Y is locally Noetherian and π is proper.

As a consequence, Theorem 18.3.9 implies the notion of relative ampleness is affine-local on Y (if π is proper and Y is locally Noetherian).

18.3.10. ★★ Ample vector bundles. The notion of an **ample vector bundle** is useful in some parts of the literature, so we define it, although we won't use the notion. A locally free sheaf \mathcal{E} on a proper A -scheme X is **ample** if $\mathcal{O}_{\mathbb{P}\mathcal{E}/X}(1)$ is an ample invertible sheaf. In particular, using Exercise 18.2.G, you can verify that an invertible sheaf is ample as a locally free sheaf (this definition) if and only if it is ample as an invertible sheaf (Definition 17.6.1), preventing a notational crisis. (The proper hypotheses can be relaxed; it is included only because Definition 17.6.1 of ampleness is only for proper schemes.)

18.3.11. ★★ Quasiaffine morphisms.

Because we have introduced quasiprojective morphisms (Definition 18.3.3), we briefly introduce quasiaffine morphisms (and quasiaffine schemes), as some readers may have cause to use them. Many of these ideas could have been introduced long before, but because we will never use them, we deal with them all at once.

A scheme X is **quasiaffine** if it admits a quasicompact open embedding into an affine scheme. This implies that X is quasicompact and separated. Note that if X is Noetherian (the most relevant case for most people), then any open embedding is of course automatically quasicompact.

18.3.H. EXERCISE. Show that X is quasiaffine if and only if the canonical map $X \rightarrow \operatorname{Spec} \Gamma(X, \mathcal{O}_X)$ (defined in Exercise 7.3.G and the paragraph following it) is a quasicompact open embedding. Thus a quasiaffine scheme comes with a *canonical* quasicompact open embedding into an affine scheme. Hint: Let $A = \Gamma(X, \mathcal{O}_X)$ for convenience. Suppose $X \hookrightarrow \operatorname{Spec} R$ is a quasicompact open embedding. We wish to show that $X \hookrightarrow \operatorname{Spec} A$ is a quasicompact open embedding. Factor $X \hookrightarrow \operatorname{Spec} R$ through $X \rightarrow \operatorname{Spec} A \rightarrow \operatorname{Spec} R$. Show that $X \rightarrow \operatorname{Spec} A$ is an open embedding in an open neighborhood of any chosen point $p \in X$, as follows. Choose $r \in R$ such that $p \in D(r) \subset X$. Notice that if $X_r = \{q \in X : r(q) \neq 0\}$, then $\Gamma(X_r, \mathcal{O}_X) = \Gamma(X, \mathcal{O}_X)_r$ by Exercise 16.3.A, using the fact that X is quasicompact and quasiseparated. Use this to show that the map $X_r \rightarrow \operatorname{Spec} A_r$ is an isomorphism.

It is not hard to show that X is quasiaffine if and only if \mathcal{O}_X is ample, but we won't use this fact.

A morphism $\pi: X \rightarrow Y$ is **quasiaffine** if the inverse image of every affine open subset of Y is a quasiaffine scheme. By Exercise 18.3.H, this is equivalent to π being quasicompact and separated, and the natural map $X \rightarrow \text{Spec } \pi_* \mathcal{O}_X$ being a quasicompact open embedding. This implies that the notion of quasiaffineness is local on the target (may be checked on an open cover), and also affine-local on a target (one may choose an affine cover, and check that the preimages of these open sets are quasiaffine). Quasiaffine morphisms are preserved by base change: if a morphism $X \hookrightarrow Z$ over Y is a quasicompact open embedding into an affine Y -scheme, then for any $W \rightarrow Y$, $X \times_Y W \hookrightarrow Z \times_Y W$ is a quasicompact open embedding into an affine W -scheme. (Interestingly, Exercise 18.3.H is *not* the right tool to use to show this base change property.)

One may readily check that quasiaffine morphisms are preserved by composition, [Stacks, tag 01SN]. Thus quasicompact locally closed embeddings are quasiaffine. If X is affine, then $X \rightarrow Y$ is quasiaffine if and only if it is quasicompact (as the preimage of any affine open subset of Y is an open subset of an affine scheme, namely X). In particular, from the Cancellation Theorem 11.2.1 for quasicompact morphisms, any morphism from an affine scheme to a quasiseparated scheme is quasiaffine.

18.4 Applications to curves

We now apply what we have learned to curves.

18.4.1. Theorem (every integral curve has a birational model that is regular and projective). — *If C is an integral curve of finite type over a field k , then there exists a regular projective k -curve C' birational to C .*




FIGURE 18.1. Constructing a projective regular model of a curve C over k via a finite cover of \mathbb{P}^1

Proof. We can assume C is affine. By the Noether Normalization Lemma 12.2.4 (or using transcendence theory, see Exercise 12.2.A), we can find some $x \in K(C) \setminus k$ with $K(C)/k(x)$ a finite extension of fields. By identifying a standard open of \mathbb{P}_k^1 with $\text{Spec } k[x]$, and taking the normalization of \mathbb{P}^1 in the function field $K(C)$ (Definition 10.7.I), we obtain a finite morphism $C' \rightarrow \mathbb{P}^1$, where C' is a curve


(finiteness by Exercise 10.7.M, $\dim C' = \dim \mathbb{P}^1$ by Exercise 12.1.F), and regular (it is reduced hence regular at the generic point, and regular at the closed points by the main theorem on discrete valuation rings in §13.5). Also, C' is birational to C as they have isomorphic function fields (Exercise 7.5.D).

Finally, $C' \rightarrow \mathbb{P}_k^1$ is finite hence projective (Exercise 18.3.C), and $\mathbb{P}_k^1 \rightarrow \operatorname{Spec} k$ is projective, so as composition of projective morphisms (to a quasicompact target) are projective (Exercise 18.3.B), $C' \rightarrow \operatorname{Spec} k$ is projective. \square

18.4.2. Theorem. — *If C is an irreducible separated regular curve, finite type over a field k , then there is an open embedding $C \hookrightarrow C'$ into some projective regular curve C' (over k).*

Proof. We first prove the result in the case where C is affine. Then we have a closed embedding $C \hookrightarrow \mathbb{A}^n$, and we consider \mathbb{A}^n as a standard open subset of \mathbb{P}^n . Taking the scheme-theoretic closure of C in \mathbb{P}^n , we obtain a projective integral curve \bar{C} , containing C as an open subset. The normalization $\tilde{\bar{C}}$ of \bar{C} is a finite morphism (finiteness of normalization, Exercise 10.7.M), so $\tilde{\bar{C}}$ is Noetherian, and regular (as normal Noetherian dimension 1 local rings are discrete valuation rings, §13.5). Moreover, by the universal property of normalization, normalization of \bar{C} doesn't affect the normal open set C , so we have an open embedding $C \hookrightarrow \tilde{\bar{C}}$. Finally, $\tilde{\bar{C}} \rightarrow \bar{C}$ is finite hence projective, and $\bar{C} \rightarrow \operatorname{Spec} k$ is projective, so (by Exercise 18.3.B) $\tilde{\bar{C}}$ is projective.


We next consider the case of general C . Let C_1 be any nonempty affine open subset of C . By the discussion in the previous paragraph, we have a regular projective compactification $\tilde{\bar{C}}_1$. The Curve-to-Projective Extension Theorem 17.5.1 (applied successively to the finite number of points $C \setminus C_1$) implies that the morphism $C_1 \hookrightarrow \tilde{\bar{C}}_1$ extends to a birational morphism $C \rightarrow \tilde{\bar{C}}_1$. Because points of a regular curve are determined by their valuation (Exercise 13.7.B), this is an inclusion of sets. Because the topology on curves is stupid (cofinite), it expresses C as an open subset of $\tilde{\bar{C}}_1$.




But why is $C \rightarrow \tilde{\bar{C}}_1$ an open embedding of schemes?

We show it is an open embedding near a point $p \in C$ as follows. Let C_2 be an affine open neighborhood of p in C . We repeat the construction we used on C_1 , to

obtain the following diagram, with open embeddings marked.



By the Curve-to-Projective Extension Theorem 17.5.1, the map $C_1 \rightarrow \tilde{\tilde{C}}_2$ extends to $\pi_{12}: \tilde{\tilde{C}}_1 \rightarrow \tilde{\tilde{C}}_2$, and we similarly have a morphism $\pi_{21}: \tilde{\tilde{C}}_2 \rightarrow \tilde{\tilde{C}}_1$, extending $C_2 \rightarrow \tilde{\tilde{C}}_1$. The composition $\pi_{21} \circ \pi_{12}$ is the identity morphism (as it is the identity rational map, see Theorem 11.4.2). The same is true for $\pi_{12} \circ \pi_{21}$, so π_{12} and π_{21} are isomorphisms. The enhanced diagram



commutes (by Theorem 11.4.2 again, implying that morphisms of reduced separated schemes are determined by their behavior on dense open sets). But $C_2 \rightarrow \tilde{\tilde{C}}_1$ is an open embedding (in particular, at p), so $C \rightarrow \tilde{\tilde{C}}_1$ is an open embedding there as well. \square

18.4.A. EXERCISE. Show that all regular proper curves over k are projective.

18.4.3. Theorem (various categories of curves are the same). — *The following categories are equivalent.*

- (i) *irreducible regular projective curves over k , and surjective k -morphisms.*
- (ii) *irreducible regular projective curves over k , and dominant k -morphisms.*
- (iii) *irreducible regular projective curves over k , and dominant rational maps over k .*
- (iv) *integral curves of finite type over k , and dominant rational maps over k .*
- (v) *the opposite category (defined in 1.2.16) of finitely generated fields of transcendence degree 1 over k , and k -homomorphisms.*

All morphisms and maps in the following discussion are assumed to be defined over k .

(Aside: The interested reader can tweak the proof below to show the following variation of the theorem: in (i)–(iv), consider only geometrically irreducible curves, and in (v), consider only fields K such that $K \cap k^s = k$ in \bar{k} . This variation allows us to exclude “weird” curves we may not want to consider. For example, if $k = \mathbb{R}$, then we are allowing curves such as $\mathbb{P}_{\mathbb{C}}^1$ which are not geometrically irreducible, as $\mathbb{P}_{\mathbb{C}}^1 \times_{\mathbb{R}} \text{Spec } \mathbb{C} \cong \mathbb{P}_{\mathbb{C}}^1 \amalg \mathbb{P}_{\mathbb{C}}^1$.)

Proof. Every surjective morphism is a dominant morphism, and every dominant morphism is a dominant rational map, and each integral regular projective curve is a quasiprojective curve, so we have shown (informally speaking) how to get from (i) to (ii) to (iii) to (iv). To get from (iv) to (i), suppose we have a dominant rational map $C_1 \dashrightarrow C_2$ of integral curves. Replace C_1 by a dense open set so the rational map is a morphism $C_1 \rightarrow C_2$. This induces a map of normalizations $\tilde{C}_1 \rightarrow \tilde{C}_2$ of regular irreducible curves. Let $\tilde{\tilde{C}}_i$ be a regular projective compactification of \tilde{C}_i (for $i = 1, 2$), as in Theorem 18.4.2. Then the morphism $\tilde{C}_1 \rightarrow \tilde{C}_2$ extends to a morphism $\tilde{\tilde{C}}_1 \rightarrow \tilde{\tilde{C}}_2$ by the Curve-to-Projective Extension Theorem 17.5.1. This morphism is surjective (Do you see why?), so we have produced a morphism in category (i).

18.4.B. EXERCISE. Put the above pieces together to describe equivalences of categories (i) through (iv).

It remains to connect (v). This is essentially the content of Exercise 7.5.D; details are left to the reader. \square

Theorem 18.4.3 has a number of implications. For example, each quasiprojective reduced curve is birational to precisely one projective regular curve. Here is another interesting consequence.

18.4.4. Important notion: Degree of a projective morphism from a curve to a regular curve.

You might already have a reasonable sense that a map of compact Riemann surfaces has a well-behaved degree, that the number of preimages of a point of C' is constant, so long as the preimages are counted with appropriate multiplicity. For example, if f locally looks like $z \mapsto z^m = y$, then near $y = 0$ and $z = 0$ (but not at $z = 0$), each point has precisely m preimages, but as y goes to 0, the m preimages coalesce. Enlightening Example 10.3.4 showed this phenomenon in a more complicated context.

We now show the algebraic version of this fact. Suppose $\pi: C \rightarrow C'$ is a surjective (or equivalently, dominant) map of regular projective curves. We will show that π has a well-behaved degree, in a sense that we will now make precise.

First we show that π is finite. Theorem 19.1.7 (finite = projective + finite fibers) implies this, but we haven't proved it yet. So instead we show the finiteness of π as follows. Let C'' be the normalization of C' in the function field of C . Then we have an isomorphism $K(C) \cong K(C'')$ (really, equality) which leads to birational maps $C \dashrightarrow C''$ which extend to morphisms as both C and C'' are regular and projective (by the Curve-to-Projective Extension Theorem 17.5.1). Thus this yields an isomorphism of C and C'' . But $C'' \rightarrow C'$ is a finite morphism by the finiteness of normalization (Exercise 10.7.M).

18.4.5. Proposition. — Suppose that $\pi: C \rightarrow C'$ is a finite morphism, where C is a (pure dimension 1) curve with no embedded points (the most important case: C is reduced), and C' is a regular curve. Then $\pi_* \mathcal{O}_C$ is locally free of finite rank.

The “no embedded points” hypothesis is the same as requiring that every associated point of C maps to a generic point of (some component of) C' . Notice that

if C is reduced, then there are no restrictions on how bad the singularities of C can be!

We will prove Proposition 18.4.5 in §18.4.10, after showing how useful it is. The regularity hypothesis on C' is necessary: the normalization of a nodal curve (Figure 8.4) is an example where most points have one preimage, and one point (the “node”) has two. (We will later see, in Remark 25.4.8 and §25.4.12, that what matters in the hypotheses of Proposition 18.4.5 is that the morphism is finite and flat.)

18.4.6. Definition. If C' is irreducible, the rank of this locally free sheaf is the **degree** of π .

18.4.C. EXERCISE. Recall that the degree of a rational map from one integral curve to another is defined as the degree of the function field extension (Definition 12.2.2). Show that (with the notation of Proposition 18.4.5) if C and C' are integral, the degree of π as a rational map is the same as the rank of $\pi_* \mathcal{O}_C$.

18.4.7. Remark for those with complex-analytic background (algebraic degree = analytic degree). If $C \rightarrow C'$ is a finite map of regular complex algebraic curves, Proposition 18.4.5 establishes that the algebraic degree as defined above is the same as the analytic degree (counting preimages, with multiplicity).

18.4.D. EXERCISE. We continue the notation and hypotheses of Proposition 18.4.5. Suppose p is a closed point of C' . The scheme-theoretic preimage $\pi^{-1}(p)$ of p is a dimension 0 scheme over k .

(a) Suppose C' is finite type over a field k , and n is the dimension of the k -vector space of global sections of the structure sheaf of $\pi^{-1}(p)$. Show that $n = (\deg \pi)(\deg p)$. (The degree of a point was defined in §5.3.8.)

(b) Suppose that C is regular, and $\pi^{-1}(p) = \{p_1, \dots, p_m\}$. Suppose t is a uniformizer of the discrete valuation ring $\mathcal{O}_{C', p}$. Show that

$$\deg \pi = \sum_{i=1}^m (\text{val}_{p_i} \pi^* t) \deg(\kappa(p_i)/\kappa(p)),$$

where $\deg(\kappa(p_i)/\kappa(p))$ denotes the degree of the field extension of the residue fields. (Can you extend (a) to remove the hypotheses of working over a field? If you are a number theorist, can you recognize (b) in terms of splitting prime ideals in extensions of rings of integers in number fields?)

18.4.E. EXERCISE. Suppose that C is an irreducible regular projective curve over k , and s is a nonzero rational function on C . Show that the number of zeros of s (counted with appropriate multiplicity) equals the number of poles. Hint: recognize this as the degree of a morphism $s: C \rightarrow \mathbb{P}_k^1$. (In the complex category, this is an important consequence of the Residue Theorem.)

18.4.8. Remark. In Exercise 19.4.E, we will see that the number of zeros (and poles) in the previous exercise is the *degree* of the line bundle giving the map to \mathbb{P}^1 (via Important Theorem 17.4.1).

18.4.9. Revisiting Example 10.3.4. Proposition 18.4.5 and Exercise 18.4.D make precise what general behavior we observed in Example 10.3.4. Suppose C' is irreducible, and that d is the rank of this allegedly locally free sheaf. Then the fiber over any point of C with residue field K is the Spec of an algebra of dimension d over K . This means that the number of points in the fiber, counted with appropriate multiplicity, is always d .

As a motivating example, we revisit Example 10.3.4, the map $\mathbb{Q}[y] \rightarrow \mathbb{Q}[x]$ given by $x \mapsto y^2$, the projection of the parabola $x = y^2$ to the x -axis. We observed the following.

- (i) The fiber over $x = 1$ is $\mathbb{Q}[y]/(y^2 - 1)$, so we get 2 points.
- (ii) The fiber over $x = 0$ is $\mathbb{Q}[y]/(y^2)$ — we get one point, with multiplicity 2, arising because of the nonreducedness.
- (iii) The fiber over $x = -1$ is $\mathbb{Q}[y]/(y^2 + 1) \cong \mathbb{Q}(i)$ — we get one point, with multiplicity 2, arising because of the field extension.
- (iv) Finally, the fiber over the generic point $\text{Spec } \mathbb{Q}(x)$ is $\text{Spec } \mathbb{Q}(y)$, which is one point, with multiplicity 2, arising again because of the field extension (as $\mathbb{Q}(y)/\mathbb{Q}(x)$ is a degree 2 extension).

We thus see three sorts of behaviors ((iii) and (iv) are really the same). Note that even if you only work with algebraically closed fields, you will still be forced to deal with this third type of behavior, because residue fields at generic points are usually not algebraically closed (witness case (iv) above).

18.4.10. Proof of Proposition 18.4.5. The key idea, useful in other circumstances, is to reduce to a fact about discrete valuation rings.

The question is local on the target, so we may assume that C' is affine. Now C' is normal, so by Exercise 5.4.B, its connected components are irreducible. Replacing C' by each of its irreducible components, we may also assume C' is integral.

By Exercise 14.3.K, if the rank of the finite type quasicoherent sheaf $\pi_* \mathcal{O}_{C'}$ is constant, then (as C' is reduced) $\pi_* \mathcal{O}_{C'}$ is locally free. We will show this by showing the rank at any closed point p of C' is the same as the rank at the generic point. (This is another example of the usefulness of the generic point.)

Suppose $C' = \text{Spec } A'$, where A' is an integral domain, and $p = [\mathfrak{m}]$. As π is an affine morphism, C is an affine scheme as well; say $C = \text{Spec } A$.

We wish to show that (i) $\dim_{A'/\mathfrak{m}}(A/\mathfrak{m})$ (the rank of $\pi_* \mathcal{O}_C$ at p) equals (ii) $\dim_{K(A')}(A'^{\times})^{-1}A$ (the rank of $\pi_* \mathcal{O}_C$ at the generic point). In other words, we take A (considered as an A' -module), and (i) quotient by \mathfrak{m} , and (ii) invert all nonzero elements of A' , and in each case compute the result's dimension over the appropriate field.

Both (i) and (ii) factor through localizing at \mathfrak{m} , so it suffices to show that $A_{\mathfrak{m}}$ is a finite rank free $A'_{\mathfrak{m}}$ -module, of rank d , say, as the answers to both (i) and (ii) will then be d .

Now $A'_{\mathfrak{m}}$ is a discrete valuation ring; let t be its uniformizer. We can assume that $t \in A'$ (as otherwise, we replace A' by a suitable localization, at the “denominators” of t). Then $A_{\mathfrak{m}}$ is a finitely generated $A'_{\mathfrak{m}}$ -module, and hence by Remark 13.5.14 is a finite sum of principal modules, of the form $A'_{\mathfrak{m}}$ or $A'_{\mathfrak{m}}/(t^n)$ (for various n). We wish to show that there are no summands of the latter type. But if there were, then t (interpreted as an element of $A_{\mathfrak{m}}$) would be a zerodivisor of $A_{\mathfrak{m}}$, and thus (interpreted as an element of A) a zerodivisor of A . But then by

Proposition 6.5.13, there is an associated point of C in $\pi^{-1}(p)$, contradicting the hypotheses that C has no embedded points. \square

CHAPTER 19

Čech cohomology of quasicoherent sheaves

This topic is surprisingly simple and elegant. You may think cohomology must be complicated, and that this is why it appears so late in the book. But you will see that we need very little background. After defining schemes, we could have immediately defined quasicoherent sheaves, and then defined cohomology, and verified that it had many useful properties.

19.1 (Desired) properties of cohomology

Rather than immediately defining cohomology of quasicoherent sheaves, we first discuss why we care, and what properties it should have.

As $\Gamma(X, \cdot)$ is a left-exact functor, if $0 \rightarrow \mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{H} \rightarrow 0$ is a short exact sequence of quasicoherent sheaves on a scheme X , then

$$0 \rightarrow \mathcal{F}(X) \rightarrow \mathcal{G}(X) \rightarrow \mathcal{H}(X)$$

is exact. We dream that this sequence continues to the right, giving a long exact sequence. More explicitly, there should be some covariant functors H^i ($i \geq 0$) from quasicoherent sheaves on X to groups such that H^0 is the global section functor Γ , and so that there is a “long exact sequence in cohomology”.

$$(19.1.0.1) \quad 0 \longrightarrow H^0(X, \mathcal{F}) \longrightarrow H^0(X, \mathcal{G}) \longrightarrow H^0(X, \mathcal{H})$$

$$\longrightarrow H^1(X, \mathcal{F}) \longrightarrow H^1(X, \mathcal{G}) \longrightarrow H^1(X, \mathcal{H}) \longrightarrow \dots$$

(In general, whenever we see a left-exact or right-exact functor, we should hope for this, and in good cases our dreams will come true. The machinery behind this usually involves *derived functors*, which we will discuss in Chapter 24.)

Before defining cohomology groups of quasicoherent sheaves explicitly, we first describe their important properties, which are in some ways more important than the formal definition. The boxed properties will be the important ones.

Suppose X is a *separated and quasicompact* A -scheme. For each quasicoherent sheaf \mathcal{F} on X , we will define A -modules $H^i(X, \mathcal{F})$. In particular, if $A = k$, they are k -vector spaces. In this case, we define $h^i(X, \mathcal{F}) = \dim_k H^i(X, \mathcal{F})$ (where k is left implicit on the left side).

- (i) Each H^i is a covariant functor $QCoh_X \rightarrow Mod_A$.

(ii) The functor H^0 is identified with the functor Γ : $H^0(X, \mathcal{F}) = \Gamma(X, \mathcal{F})$, and the covariance of (i) for $i = 0$ is just the usual covariance for Γ ($\mathcal{F} \rightarrow \mathcal{G}$ induces $\Gamma(X, \mathcal{F}) \rightarrow \Gamma(X, \mathcal{G})$).

(iii) If $0 \rightarrow \mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{H} \rightarrow 0$ is a short exact sequence of quasicoherent sheaves on X , then we have a long exact sequence (19.1.0.1). The maps $H^i(X, \mathcal{F}) \rightarrow H^i(X, \mathcal{G})$ come from covariance, and similarly for $H^i(X, \mathcal{G}) \rightarrow H^i(X, \mathcal{H})$. The connecting homomorphisms $H^i(X, \mathcal{H}) \rightarrow H^{i+1}(X, \mathcal{F})$ will have to be defined.

(iv) If $\pi: X \rightarrow Y$ is any morphism of quasicoherent separated A -schemes, and \mathcal{F} is a quasicoherent sheaf on X , then there is a natural morphism $H^i(Y, \pi_* \mathcal{F}) \rightarrow H^i(X, \mathcal{F})$ extending $\Gamma(Y, \pi_* \mathcal{F}) \rightarrow \Gamma(X, \mathcal{F})$. (Note that π is quasicompact and separated by the Cancellation Theorem 11.2.1 for quasicompact and separated morphisms, taking $Z = \text{Spec } A$ in the statement of the Cancellation Theorem, so $\pi_* \mathcal{F}$ is indeed a quasicoherent sheaf by Exercise 17.2.C.) We will later see this as part of a larger story, the *Leray spectral sequence* (Theorem 24.4.5). If \mathcal{G} is a quasicoherent sheaf on Y , then setting $\mathcal{F} := \pi^* \mathcal{G}$ and using the adjunction map $\mathcal{G} \rightarrow \pi_* \pi^* \mathcal{G}$ and covariance of (ii) gives a natural **pullback map** $H^i(Y, \mathcal{G}) \rightarrow H^i(X, \pi^* \mathcal{G})$ (via $H^i(Y, \mathcal{G}) \rightarrow H^i(Y, \pi_* \pi^* \mathcal{G}) \rightarrow H^i(X, \pi^* \mathcal{G})$) extending $\Gamma(Y, \mathcal{G}) \rightarrow \Gamma(X, \pi^* \mathcal{G})$. In this way, H^i is a “contravariant functor in the space”.

(v) If $\pi: X \rightarrow Y$ is an affine morphism, and \mathcal{F} is a quasicoherent sheaf on X , the natural map of (iv) is an isomorphism: $H^i(Y, \pi_* \mathcal{F}) \xrightarrow{\sim} H^i(X, \mathcal{F})$. When π is a closed embedding and $Y = \mathbb{P}_A^n$, this isomorphism translates calculations on arbitrary projective A -schemes to calculations on \mathbb{P}_A^n .

(vi) (*affine cover cohomology vanishing*) If X can be covered by n affine open sets, then $H^i(X, \mathcal{F}) = 0$ for $i \geq n$ for all \mathcal{F} . In particular, on affine schemes, all higher ($i > 0$) quasicoherent cohomology groups vanish. The vanishing of H^1 in this case, along with the long exact sequence (iii) implies that Γ is an exact functor for quasicoherent sheaves on affine schemes, something we already knew (Exercise 6.3.A).

19.1.1. Remarks (not part of property (vi)). It is also true that if $\dim X = n$, then $H^i(X, \mathcal{F}) = 0$ for all $i > n$ and for all \mathcal{F} (**dimensional cohomology vanishing**). We will prove this for projective k -schemes (Theorem 19.2.6) and even quasiprojective k -schemes (Exercise 23.4.T). See §19.2.7 for discussion of the general case.

The converse to (vi) in the case when $n = 1$ is Serre’s *cohomological criterion for affineness*, see §19.7.5.

Let’s get back to our list.

(vii) The functor H^i behaves well under direct sums, and more generally under filtered colimits: $H^i(X, \text{colim } \mathcal{F}_j) = \text{colim } H^i(X, \mathcal{F}_j)$.

(viii) We will also identify the cohomology of all $\mathcal{O}(m)$ on \mathbb{P}_A^n :

19.1.2. Theorem. —

- (a) $H^0(\mathbb{P}_A^n, \mathcal{O}_{\mathbb{P}_A^n}(m))$ is a free A -module of rank $\binom{n+m}{m}$ if $m \geq 0$.
- (b) $H^n(\mathbb{P}_A^n, \mathcal{O}_{\mathbb{P}_A^n}(m))$ is a free A -module of rank $\binom{-m-1}{-n-m-1}$ if $m \leq -n-1$.
- (c) $H^i(\mathbb{P}_A^n, \mathcal{O}_{\mathbb{P}_A^n}(m)) = 0$ otherwise.

We have already shown the first statement in Essential Exercise 15.1.C.

Theorem 19.1.2 has a number of features that will be the first appearances of facts that we will prove later.

- The cohomology of these bundles vanish in degree above n ((vi) above, “affine cover cohomology vanishing”)
- These cohomology groups are always *finitely generated* A -modules. This will be true for all coherent sheaves on projective A -schemes (Theorem 19.1.3(i)), and indeed (with more work) on proper A -schemes (Grothendieck’s Coherence Theorem 19.9.1).
- The top cohomology group vanishes for $m > -n - 1$. We will later see this as an example of *Kodaira vanishing*, see §22.5.7.

If A is a field k :

- The top cohomology group is one-dimensional for $m = -n - 1$. This is the first appearance of the *dualizing sheaf*.
- There is a natural duality

$$H^i(\mathbb{P}^n, \mathcal{O}(m)) \times H^{n-i}(\mathbb{P}^n, \mathcal{O}(-n-1-m)) \rightarrow H^n(\mathbb{P}^n, \mathcal{O}(-n-1))$$

This is the first appearance of *Serre duality*. (For the case $n = 1$, see Example 19.5.4.)

- The alternating sum $\sum (-1)^i h^i(\mathbb{P}^n, \mathcal{O}(m))$ is a polynomial in m . This is a first example of a *Hilbert polynomial*.

Before proving these facts, let’s first use them to prove interesting things, as motivation.

By Theorem 16.3.1, for any coherent sheaf \mathcal{F} on \mathbb{P}_A^n we can find a surjection $\mathcal{O}(m)^{\oplus j} \rightarrow \mathcal{F}$, which yields the exact sequence

$$(19.1.2.1) \quad 0 \longrightarrow \mathcal{G} \longrightarrow \mathcal{O}(m)^{\oplus j} \longrightarrow \mathcal{F} \longrightarrow 0$$

for some coherent sheaf \mathcal{G} . We can use this to prove the following.

19.1.3. Theorem. — (i) For any coherent sheaf \mathcal{F} on a projective A -scheme X where A is Noetherian, $H^i(X, \mathcal{F})$ is a coherent (finitely generated) A -module.
(ii) (Serre vanishing) Furthermore, for $m \gg 0$, $H^i(X, \mathcal{F}(m)) = 0$ for all $i > 0$ (even without Noetherian hypotheses).

(A slightly fancier version of Serre vanishing will be given in Theorem 19.8.G.)

Proof. Because cohomology of a closed subscheme can be computed on the ambient space ((v) above), we may immediately reduce to the case $X = \mathbb{P}_A^n$.

(i) Consider the long exact sequence:

$$\begin{aligned}
 0 &\longrightarrow H^0(\mathbb{P}_A^n, \mathcal{G}) \longrightarrow H^0(\mathbb{P}_A^n, \mathcal{O}(\mathfrak{m})^{\oplus j}) \longrightarrow H^0(\mathbb{P}_A^n, \mathcal{F}) \longrightarrow \\
 &H^1(\mathbb{P}_A^n, \mathcal{G}) \longrightarrow H^1(\mathbb{P}_A^n, \mathcal{O}(\mathfrak{m})^{\oplus j}) \longrightarrow H^1(\mathbb{P}_A^n, \mathcal{F}) \longrightarrow \dots \\
 &\dots \longrightarrow H^{n-1}(\mathbb{P}_A^n, \mathcal{G}) \longrightarrow H^{n-1}(\mathbb{P}_A^n, \mathcal{O}(\mathfrak{m})^{\oplus j}) \longrightarrow H^{n-1}(\mathbb{P}_A^n, \mathcal{F}) \longrightarrow \\
 &H^n(\mathbb{P}_A^n, \mathcal{G}) \longrightarrow H^n(\mathbb{P}_A^n, \mathcal{O}(\mathfrak{m})^{\oplus j}) \longrightarrow H^n(\mathbb{P}_A^n, \mathcal{F}) \longrightarrow 0
 \end{aligned}$$

The exact sequence ends here because \mathbb{P}_A^n is covered by $n + 1$ affine open sets ((vi) above, “affine cover cohomology vanishing”). Then $H^n(\mathbb{P}_A^n, \mathcal{O}(\mathfrak{m})^{\oplus j})$ is finitely generated by Theorem 19.1.2(b), hence $H^n(\mathbb{P}_A^n, \mathcal{F})$ is finitely generated for all coherent sheaves \mathcal{F} . Hence in particular, $H^n(\mathbb{P}_A^n, \mathcal{G})$ is finitely generated. As $H^{n-1}(\mathbb{P}_A^n, \mathcal{O}(\mathfrak{m})^{\oplus j})$ is finitely generated, and $H^n(\mathbb{P}_A^n, \mathcal{G})$ is too, we have that $H^{n-1}(\mathbb{P}_A^n, \mathcal{F})$ is finitely generated for all coherent sheaves \mathcal{F} . We continue inductively downwards.

(ii) Twist (19.1.2.1) by $\mathcal{O}(N)$ for $N \gg 0$. Then

$$H^n(\mathbb{P}_A^n, \mathcal{O}(\mathfrak{m} + N)^{\oplus j}) = H^n(\mathbb{P}_A^n, \mathcal{O}(\mathfrak{m} + N))^{\oplus j} = 0$$

(by (vii) above), so $H^n(\mathbb{P}_A^n, \mathcal{F}(N)) = 0$. Translation: for any coherent sheaf, its top cohomology vanishes once you twist by $\mathcal{O}(N)$ for N sufficiently large. Hence this is true for \mathcal{G} as well. Hence from the long exact sequence, $H^{n-1}(\mathbb{P}_A^n, \mathcal{F}(N)) = 0$ for $N \gg 0$. As in (i), we induct downwards, until we get that $H^1(\mathbb{P}_A^n, \mathcal{F}(N)) = 0$. (The induction stops here, as it is *not* true that $H^0(\mathbb{P}_A^n, \mathcal{O}(\mathfrak{m} + N)^{\oplus j}) = 0$ for large N — quite the opposite.) \square

In particular, we have proved the following, that we would have cared about even before we knew about cohomology.

19.1.4. Corollary. — *Any projective k -scheme has a finite-dimensional space of functions. More generally, if A is Noetherian and \mathcal{F} is a coherent sheaf on a projective A -scheme, then $H^0(X, \mathcal{F})$ is a coherent A -module.*

(We will generalize this in Theorem 19.8.1.) I want to emphasize how remarkable this proof is. It is a question about global sections, i.e., H^0 , which we think of as the most down to earth cohomology group, yet the proof is by downward induction for H^n , starting with n large.

Corollary 19.1.4 is true more generally for proper k -schemes, not just projective k -schemes (Grothendieck’s Coherence Theorem 19.9.1).

Here are some important consequences. They can also be shown directly, without the use of cohomology, but with much more elbow grease.

19.1.5. As a partial converse, if $h^0(X, \mathcal{O}_X) = 1$, then X is connected (why?), but need not be reduced: witness the subscheme in \mathbb{P}^2 cut out by $x^2 = 0$. (For experts: the geometrically connected hypothesis is necessary, as $X = \text{Spec } \mathbb{C}$ is a projective

integral \mathbb{R} -scheme, with $h^0(X, \mathcal{O}_X) = 2$. Similarly, a nontrivial purely inseparable field extension can be used to show that the geometrically reduced hypothesis is also necessary.)

19.1.A. CRUCIAL EXERCISE (PUSHFORWARDS OF COHERENT SHEAVES BY PROJECTIVE MORPHISMS ARE COHERENT). Suppose $\pi: X \rightarrow Y$ is a projective morphism of locally Noetherian schemes. Show that the pushforward of a coherent sheaf on X is a coherent sheaf on Y . (See Grothendieck's Coherence Theorems 19.8.1 and 19.9.1 for generalizations.)

Finite morphisms are affine (from the definition) and projective (18.3.C). We can now show that this is a characterization of finiteness.

19.1.6. Corollary (projective + affine = finite). — Suppose Y is locally Noetherian. Then a morphism $\pi: X \rightarrow Y$ is projective and affine if and only if π is finite.

We will see in Exercise 19.9.A that the projective hypotheses can be relaxed to proper.

Proof. We already know that finite morphisms are affine (by definition) and projective (Exercise 18.3.C), so we show the converse. Suppose π is projective and affine. By Exercise 19.1.A, $\pi_* \mathcal{O}_X$ is coherent and hence finite type. \square

The following result was promised in §18.3.5, and has a number of useful consequences.

19.1.7. Theorem (projective + finite fibers = finite). — Suppose $\pi: X \rightarrow Y$ with Y Noetherian. Then π is projective and has finite fibers if and only if it is finite. Equivalently, π is projective and quasifinite if and only if it is finite.

(Recall that quasifinite = finite fibers + finite type. But projective includes finite type.) It is true more generally that (with Noetherian hypotheses) proper + finite fibers = finite, see Theorem 30.6.2.

Proof. We show π is finite near a point $q \in Y$. Fix an affine open neighborhood $\text{Spec } A$ of q in Y . Pick a hypersurface H in \mathbb{P}_A^n missing the preimage of q , so $H \cap X$ is closed. (See Figure 19.1.) Let $H' = \pi(H \cap X)$, which is closed in $\text{Spec } A$, and doesn't contain q . Let $U = \text{Spec } A - H'$, which is an open set containing q . Then above U , π is projective and affine, so we are done by Corollary 19.1.6. \square

A similar trick was used in the proof of the Rigidity Lemma 11.5.12.

19.1.B. EXERCISE (UPPER SEMICONTINUITY OF FIBER DIMENSION ON THE TARGET, FOR PROJECTIVE MORPHISMS). Use a similar argument as in Theorem 19.1.7 to prove *upper semicontinuity of fiber dimension of projective morphisms*: suppose $\pi: X \rightarrow Y$ is a projective morphism where Y is locally Noetherian (or more generally \mathcal{O}_Y is coherent over itself). Show that $\{q \in Y : \dim \pi^{-1}(q) > k\}$ is a Zariski-closed subset of Y . In other words, the dimension of the fiber “jumps over Zariski-closed subsets” of the target. (You can interpret the case $k = -1$ as the fact that projective morphisms are closed, which is basically the Fundamental Theorem of Elimination Theory 8.4.7, cf. §18.3.4.) This exercise is rather important for having a sense of how projective morphisms behave. (The case of varieties was done earlier, in Theorem 12.4.2(b). This approach is much simpler.)




FIGURE 19.1. The key construction in the proof of Theorem 19.1.7

The final exercise of the section is on a different theme.

19.1.C. EXERCISE. Suppose $0 \rightarrow \mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{H} \rightarrow 0$ is an exact sequence of quasicoherent sheaves on X , projective over A , with \mathcal{F} coherent. Show that for $n \gg 0$,

$$0 \longrightarrow H^0(X, \mathcal{F}(n)) \longrightarrow H^0(X, \mathcal{G}(n)) \longrightarrow H^0(X, \mathcal{H}(n)) \longrightarrow 0$$

is also exact. (Hint: for $n \gg 0$, $H^1(X, \mathcal{F}(n)) = 0$.)

19.2 Definitions and proofs of key properties

This section could be read much later; the facts we will use are all stated in the previous section. However, the arguments are not complicated, so you want to read this right away. As you read this, you should go back and check off all the facts in the previous section, to assure yourself that you understand everything promised.

19.2.1. Čech cohomology. Čech cohomology in general settings is defined using a limit over finer and finer covers of a space. In our algebro-geometric setting, the situation is much cleaner, and we can use a single cover.

Suppose X is quasicompact and separated, which is true for example if X is quasiprojective over A . In particular, X may be covered by a finite number of affine open sets, and the intersection of any two affine open sets is also an affine open set (by separatedness, Proposition 11.3.11). We will use quasicompactness and separatedness only in order to ensure these two nice properties.

Suppose \mathcal{F} is a quasicoherent sheaf, and $\mathcal{U} = \{U_i\}_{i=1}^n$ is a *finite* collection of affine open sets covering X . For $I \subset \{1, \dots, n\}$ define $U_I = \cap_{i \in I} U_i$, which is affine by the separated hypothesis. (Here is a strong analogy for those who have seen

cohomology in other contexts: cover a topological space X with a finite number of open sets U_i , such that all intersections $\cap_{i \in I} U_i$ are contractible.) Consider the **Čech complex**

$$(19.2.1.1) \quad 0 \longrightarrow \prod_{\substack{|I|=1 \\ I \subset \{1, \dots, n\}}} \mathcal{F}(U_I) \longrightarrow \dots \longrightarrow \prod_{\substack{|I|=i \\ I \subset \{1, \dots, n\}}} \mathcal{F}(U_I) \longrightarrow \prod_{\substack{|I|=i+1 \\ I \subset \{1, \dots, n\}}} \mathcal{F}(U_I) \longrightarrow \dots$$

The maps are defined as follows. The map from $\mathcal{F}(U_I) \rightarrow \mathcal{F}(U_J)$ is 0 unless $I \subset J$, i.e., $J = I \cup \{j\}$. If j is the k th element of J , then the map is $(-1)^{k-1}$ times the restriction map res_{U_I, U_J} .

19.2.A. EASY EXERCISE (FOR THOSE WHO HAVEN'T SEEN ANYTHING LIKE THE ČECH COMPLEX BEFORE). Show that the Čech complex is indeed a complex, i.e., that the composition of two consecutive arrows is 0.

Define $H_{\mathcal{U}}^i(X, \mathcal{F})$ to be the i th cohomology group of the complex (19.2.1.1). (The indexing starts with $i = 0$.) Note that if X is an A -scheme, then $H_{\mathcal{U}}^i(X, \mathcal{F})$ is an A -module. We have almost succeeded in defining the Čech cohomology group H^i , except our definition a priori depends on a choice of a cover \mathcal{U} . Note that $H_{\mathcal{U}}^i(X, \cdot)$ is clearly a covariant functor $QCoh_X \rightarrow Mod_A$.

19.2.B. EASY EXERCISE. Identify $H_{\mathcal{U}}^0(X, \mathcal{F})$ with $\Gamma(X, \mathcal{F})$. (Hint: use the sheaf axioms for \mathcal{F} .)

19.2.C. EXERCISE. Suppose

$$(19.2.1.2) \quad 0 \longrightarrow \mathcal{F} \longrightarrow \mathcal{G} \longrightarrow \mathcal{H} \longrightarrow 0$$

is a short exact sequence of sheaves of abelian groups on a topological space, and \mathcal{U} is a finite open cover such that on any intersection U_I of open subsets in \mathcal{U} , the map $\Gamma(U_I, \mathcal{G}) \rightarrow \Gamma(U_I, \mathcal{H})$ is surjective. Show that we get a “long exact sequence of cohomology for $H_{\mathcal{U}}^i$ ” (where we take the same definition of $H_{\mathcal{U}}^i$). In our situation, where X is a quasicompact separated A -scheme, and (19.2.1.2) is a short exact sequence of quasicoherent sheaves on X , show that we get a long exact sequence for the A -modules $H_{\mathcal{U}}^i$.

In the proof of Theorems 19.2.4 and 19.8.1, we will make use of the fact that your construction of the connecting homomorphism will “commute with localization of A ”. More precisely, we will need the following.

19.2.D. EXERCISE. Suppose we are given a short exact sequence (19.2.1.2) of quasicoherent sheaves on a quasicompact separated A -scheme $\pi: X \rightarrow \text{Spec } A$, a cover \mathcal{U} of X by affine open sets, and some $f \in A$. The restriction of the sets of \mathcal{U} to X_f yields an affine open cover \mathcal{U}' of $X_f = \pi^{-1}(D(f))$. Identify the long exact sequence associated to (19.2.1.2) using $H_{\mathcal{U}}^i$, localized at f , with the long exact sequence associated to the restriction of (19.2.1.2) to X_f , using the affine open cover \mathcal{U}' . (First check that the maps such as $H_{\mathcal{U}}^i(\mathcal{F}) \rightarrow H_{\mathcal{U}}^i(\mathcal{G})$ given by covariance “commute with localization”, and then check that the connecting homomorphisms do as well.)

19.2.2. Theorem/Definition. — Our standing assumption is that X is quasicompact and separated. $H_{\mathcal{U}}^i(X, \mathcal{F})$ is independent of the choice of (finite) cover $\{U_i\}$. More precisely, for any two covers $\{U_i\} \subset \{V_i\}$, the maps $H_{\{V_i\}}^i(X, \mathcal{F}) \rightarrow H_{\{U_i\}}^i(X, \mathcal{F})$ induced by the natural map of Čech complexes (19.2.1.1) are isomorphisms. Define the **Čech cohomology group** $H^i(X, \mathcal{F})$ to be this group.

If you are unsure of what the “natural map of Čech complexes” is, by (19.2.3.1) it should become clear.

19.2.3. For experts: maps of complexes inducing isomorphisms on cohomology groups are called **quasiisomorphisms**. We are actually getting a finer invariant than cohomology out of this construction; we are getting an element of the *derived category of A -modules*.

Proof. We need only prove the result when $|\{V_i\}| = |\{U_i\}| + 1$. We will show that if $\{U_i\}_{1 \leq i \leq n}$ is a cover of X , and U_0 is any other affine open set, then the map $H_{\{U_i\}_{0 \leq i \leq n}}^i(X, \mathcal{F}) \rightarrow H_{\{U_i\}_{1 \leq i \leq n}}^i(X, \mathcal{F})$ is an isomorphism. Consider the exact sequence of complexes

(19.2.3.1)

$$\begin{array}{ccccccc}
 & & 0 & & 0 & & 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 \cdots \longrightarrow & \prod_{\substack{|I| = i-1 \\ 0 \in I}} \mathcal{F}(U_I) & \longrightarrow & \prod_{\substack{|I| = i \\ 0 \in I}} \mathcal{F}(U_I) & \longrightarrow & \prod_{\substack{|I| = i+1 \\ 0 \in I}} \mathcal{F}(U_I) & \longrightarrow \cdots \\
 & \downarrow & & \downarrow & & \downarrow & \\
 \cdots \longrightarrow & \prod_{|I| = i-1} \mathcal{F}(U_I) & \longrightarrow & \prod_{|I| = i} \mathcal{F}(U_I) & \longrightarrow & \prod_{|I| = i+1} \mathcal{F}(U_I) & \longrightarrow \cdots \\
 & \downarrow & & \downarrow & & \downarrow & \\
 \cdots \longrightarrow & \prod_{\substack{|I| = i-1 \\ 0 \notin I}} \mathcal{F}(U_I) & \longrightarrow & \prod_{\substack{|I| = i \\ 0 \notin I}} \mathcal{F}(U_I) & \longrightarrow & \prod_{\substack{|I| = i+1 \\ 0 \notin I}} \mathcal{F}(U_I) & \longrightarrow \cdots \\
 & \downarrow & & \downarrow & & \downarrow & \\
 & 0 & & 0 & & 0 &
 \end{array}$$

Throughout, $I \subset \{0, \dots, n\}$. The bottom two rows are Čech complexes with respect to two covers, and the map between them induces the desired map on cohomology. We get a long exact sequence of cohomology from this short exact sequence of complexes (Exercise 1.6.C). Thus we wish to show that the top row is exact and thus has vanishing cohomology. (Note that $U_0 \cap U_j$ is affine by our separatedness hypothesis, Proposition 11.3.11.) But the i th cohomology of the top row is precisely $H_{\{U_j \cap U_0\}_{j>0}}^{i-1}(U_0, \mathcal{F})$ except at step 0, where we get 0 (because the complex starts off $0 \rightarrow \mathcal{F}(U_0) \rightarrow \prod_{j=1}^n \mathcal{F}(U_0 \cap U_j)$). So it suffices to show that higher Čech groups of affine schemes are 0. Hence we are done by the following result. \square

19.2.4. Theorem. — *The higher Čech cohomology $H_{\mathcal{U}}^i(X, \mathcal{F})$ of an affine A -scheme X vanishes (for any affine cover \mathcal{U} , $i > 0$, and quasicoherent \mathcal{F}).*

This is another “partition of unity” argument, in the spirit of the proof of Theorem 4.1.2.

Proof. (The following argument can be made shorter using spectral sequences, but we avoid them for the sake of clarity.) We want to show that the “extended” complex

$$(19.2.4.1) \quad 0 \longrightarrow \mathcal{F}(X) \longrightarrow \prod_{|I|=1} \mathcal{F}(U_I) \longrightarrow \prod_{|I|=2} \mathcal{F}(U_I) \longrightarrow \cdots$$

(where the global sections $\mathcal{F}(X)$ have been appended to the start) has no cohomology, i.e., is exact. We do this with a trick.

Suppose first that some U_i , say U_0 , is X . We deal with this case by sleight of hand. The complex (19.2.4.1) is the middle row of the following short exact sequence of complexes.

$$(19.2.4.2) \quad \begin{array}{ccccccc} 0 & \longrightarrow & 0 & \longrightarrow & \prod_{|I|=1, 0 \in I} \mathcal{F}(U_I) & \longrightarrow & \prod_{|I|=2, 0 \in I} \mathcal{F}(U_I) \longrightarrow \cdots \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & \mathcal{F}(X) & \longrightarrow & \prod_{|I|=1} \mathcal{F}(U_I) & \longrightarrow & \prod_{|I|=2} \mathcal{F}(U_I) \longrightarrow \cdots \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & \mathcal{F}(X) & \longrightarrow & \prod_{|I|=1, 0 \notin I} \mathcal{F}(U_I) & \longrightarrow & \prod_{|I|=2, 0 \notin I} \mathcal{F}(U_I) \longrightarrow \cdots \end{array}$$

The top row is the same as the bottom row, slid over by 1. The corresponding long exact sequence of cohomology shows that the central row has vanishing cohomology. (You should show that the “connecting homomorphism” on cohomology is indeed an isomorphism.) This might remind you of the *mapping cone* construction (Exercise 1.7.E).

We next prove the general case by sleight of hand. Say $X = \operatorname{Spec} R$. We wish to show that the complex of A -modules (19.2.4.1) is exact. It is also a complex of R -modules, so we wish to show that the complex of R -modules (19.2.4.1) is exact. To show that it is exact, it suffices to show that for a cover of $\operatorname{Spec} R$ by distinguished open sets $D(f_i)$ ($1 \leq i \leq r$) (i.e., $(f_1, \dots, f_r) = 1$ in R) the complex is exact. (Translation: exactness of a sequence of sheaves may be checked locally.) We choose a cover so that each $D(f_i)$ is contained in some $U_j = \operatorname{Spec} A_j$. Consider the complex localized at f_i . As

$$\Gamma(\operatorname{Spec} A, \mathcal{F})_{f_i} = \Gamma(\operatorname{Spec}(A_j)_{f_i}, \mathcal{F})$$

(by quasicoherence of \mathcal{F} , Exercise 6.2.D), as $U_j \cap D(f_i) = D(f_i)$, we are in the situation where one of the U_i ’s is X , so we are done. \square

We have now proved properties (i)–(iii) of the previous section. The “affine cover cohomology vanishing” property (vi) is also straightforward: if X is covered by n affine open sets, use these as the cover \mathcal{U} , and notice that the Čech complex ends by the n th step.

19.2.E. EXERCISE (PROPERTY (v)). Suppose $\pi: X \rightarrow Y$ is an affine morphism, and Y is a quasicompact and separated A -scheme (and hence X is too, as affine morphisms are both quasicompact and separated). If \mathcal{F} is a quasicoherent sheaf on X , describe a natural isomorphism $H^i(Y, \pi_* \mathcal{F}) \xrightarrow{\sim} H^i(X, \mathcal{F})$. (Hint: if \mathcal{U} is an affine cover of Y , “ $\pi^{-1}(\mathcal{U})$ ” is an affine cover of X . Use these covers to compute the cohomology of \mathcal{F} .)

19.2.F. EXERCISE (PROPERTY (iv)). Suppose $\pi: X \rightarrow Y$ is any quasicompact separated morphism, \mathcal{F} is a quasicoherent sheaf on X , and Y is a quasicompact separated A -scheme. The hypotheses on π ensure that $\pi_* \mathcal{F}$ is a quasicoherent sheaf on Y . Describe a natural morphism $H^i(Y, \pi_* \mathcal{F}) \rightarrow H^i(X, \mathcal{F})$ extending $\Gamma(Y, \pi_* \mathcal{F}) \rightarrow \Gamma(X, \mathcal{F})$. (Aside: this morphism is an isomorphism for $i = 0$, but need not be an isomorphism for higher i : consider $i = 1$, $X = \mathbb{P}_k^1$, $\mathcal{F} = \mathcal{O}(-2)$, and let Y be a point $\text{Spec } k$.)

19.2.G. EXERCISE. Prove Property (vii) of the previous section, that “cohomology commutes with filtered colimits”. (This can be done by hand. Hint: in the category of modules over a ring, taking the colimit over a filtered sets is an exact functor, §1.6.12.)

We have now proved all of the properties of the previous section, except for (viii), which we will get to in §19.3.

19.2.5. Useful facts about cohomology for k -schemes.

19.2.H. EXERCISE (COHOMOLOGY AND CHANGE OF BASE FIELD). Suppose X is a quasicompact separated k -scheme, and \mathcal{F} is a quasicoherent sheaf on X . Give an isomorphism

$$H^i(X, \mathcal{F}) \otimes_k K \xrightarrow{\sim} H^i(X \times_{\text{Spec } k} \text{Spec } K, \mathcal{F} \otimes_k K)$$

for all i , where K/k is any field extension. Here $\mathcal{F} \otimes_k K$ means the pullback of \mathcal{F} to $X \times_{\text{Spec } k} \text{Spec } K$. Hence $h^i(X, \mathcal{F}) = h^i(X \times_{\text{Spec } k} \text{Spec } K, \mathcal{F} \otimes_k K)$. If $i = 0$ (taking $H^0 = \Gamma$), show the result without the quasicompact and separated hypotheses. (This is useful for relating facts about k -schemes to facts about schemes over algebraically closed fields. Your proof might use vector spaces — i.e., linear algebra — in a fundamental way. If it doesn’t, you may prove something more general, if $k \rightarrow K$ is replaced by a flat ring map $B \rightarrow A$. Recall that $B \rightarrow A$ is flat if $\otimes_B A$ is an exact functor $\text{Mod}_B \rightarrow \text{Mod}_A$. A hint for this harder exercise: the FHHF Theorem, Exercise 1.6.H. See Exercise 19.8.B(b) for the next generalization of this.)

19.2.I. EXERCISE (BASE-POINT-FREENESS IS INDEPENDENT OF EXTENSION OF BASE FIELD). Suppose X is a scheme over a field k , \mathcal{L} is an invertible sheaf on X , and K/k is a field extension. Show that \mathcal{L} is base-point-free if and only if its pullback to $X \times_{\text{Spec } k} \text{Spec } K$ is base-point-free. (Hint: Exercise 19.2.H with $i = 0$ implies that a basis of sections of \mathcal{L} over k becomes, after tensoring with K , a basis of sections of $\mathcal{L} \otimes_k K$.)

19.2.6. Theorem (dimensional cohomology vanishing for quasicoherent sheaves on projective k -schemes). — Suppose X is a projective k -scheme, and \mathcal{F} is a quasicoherent sheaf on X . Then $H^i(X, \mathcal{F}) = 0$ for $i > \dim X$.

In other words, cohomology vanishes above the dimension of X .

Proof. Suppose $X \hookrightarrow \mathbb{P}^N$, and let $n = \dim X$. We show that X may be covered by $n + 1$ affine open sets. Exercise 12.3.B shows that there are $n + 1$ effective Cartier divisors on \mathbb{P}^N such that their complements U_0, \dots, U_n cover X . Then U_i is affine, so $U_i \cap X$ is affine, and thus we have covered X with $n + 1$ affine open sets. \square

(It turns out that $n + 1$ affine open sets are always necessary. One way of proving this is by showing that the complement of a dense affine open subset is always pure codimension 1, see for example [RotV, Lem. 2.3].)

19.2.7. \star Dimensional cohomology vanishing more generally. Using the theory of blowing up, Theorem 19.2.6 can be extended to quasiprojective k -schemes, see §23.4.15. Dimensional cohomology vanishing is even true in much greater generality. To state it, we need to define cohomology with the more general machinery of derived functors (Chapter 24). If X is a Noetherian topological space (§3.6.14) and \mathcal{F} is any sheaf of abelian groups on X , we have $H^i(X, \mathcal{F}) = 0$ for all $i > \dim X$. (Grothendieck sketches his elegant proof in [GrS, p. 29-30]; see [Ha1, Theorem III.2.7] for a more detailed explanation.) In particular, if X is a k -variety of dimension n , we *always* have dimensional cohomology vanishing, even for crazy varieties that can't be covered with $n + 1$ affine open subsets (see §23.4.15).

19.2.8. The Künneth formula.

Suppose X and Y are quasicompact separated k -schemes, and \mathcal{F} and \mathcal{G} are quasicoherent sheaves on X and Y respectively. Let $\pi_X: X \times_k Y \rightarrow X$ and $\pi_Y: X \times_k Y \rightarrow Y$ be the two projections. Recall the definition $\mathcal{F} \boxtimes \mathcal{G} := \pi_X^* \mathcal{F} \otimes \pi_Y^* \mathcal{G}$ (§17.4.9). Then we have an isomorphism

$$H^m(X \times_k Y, \mathcal{F} \boxtimes \mathcal{G}) \xrightarrow{\sim} \bigoplus_{p+q=m} H^p(X, \mathcal{F}) \otimes_k H^q(Y, \mathcal{G}).$$

To show this, choose affine covers of X and Y , and produce the Čech complexes for \mathcal{F} and \mathcal{G} . Show that the tensor product of these two complexes (the total complex associated to the double complex) is the Čech complex for $\mathcal{F} \boxtimes \mathcal{G}$ (with respect to the products of the affine covers of X and Y). Finally, show that the cohomology of the tensor product of two complexes over k is the tensor products of the cohomologies, a result known as the *Eilenberg-Zilber Theorem*.

19.2.9. The cup product. The cup product in Čech cohomology can be defined in a simple way; see [Liu, Exer. 2.17] for a particularly elegant description. We will not need this construction.

19.3 Cohomology of line bundles on projective space

We now finally prove the last promised basic fact about cohomology, property (viii) of §19.1, Theorem 19.1.2, on the cohomology of line bundles on projective space. More correctly, we will do one case and you will do the rest.

We begin with a warm-up that will let you (implicitly) see some of the structure that will arise in the proof. It also gives good practice in computing cohomology groups.

19.3.A. EXERCISE. Compute the cohomology groups $H^i(\mathbb{A}_k^2 \setminus \{(0,0)\}, \mathcal{O})$. (Hint: the case $i = 0$ was done in Example 4.4.1. The case $i > 1$ is clear from “affine cover cohomology vanishing”, property (vi) above.) In particular, show that $H^1(\mathbb{A}_k^2 \setminus \{(0,0)\}, \mathcal{O}) \neq 0$, and thus give another proof (see §4.4.3) of the fact that $\mathbb{A}_k^2 \setminus \{(0,0)\}$ is not affine. (Cf. Serre’s cohomological criterion for affineness, Remark ??.)

19.3.1. Remark. Essential Exercise 15.1.C and the ensuing discussion showed that $H^0(\mathbb{P}_A^n, \mathcal{O}_{\mathbb{P}_A^n}(m))$ should be interpreted as the homogeneous degree m polynomials in x_0, \dots, x_n (with A -coefficients). Similarly, $H^n(\mathbb{P}_A^n, \mathcal{O}_{\mathbb{P}_A^n}(m))$ should be interpreted as the homogeneous degree m Laurent polynomials in x_0, \dots, x_n , where in each monomial, each x_i appears with degree at most -1 .

19.3.2. Proof of Theorem 19.1.2. We take the standard cover $U_0 = D(x_0), \dots, U_n = D(x_n)$ of \mathbb{P}_A^n .

19.3.B. EXERCISE (ESSENTIAL FOR THE PROOF OF THEOREM 19.1.2; CF. REMARK 19.3.1 ABOVE). If $I \subset \{0, \dots, n\}$, then give an isomorphism (of A -modules) of $\Gamma(U_I, \mathcal{O}(m))$ with the homogeneous degree m Laurent monomials (in x_0, \dots, x_n , with coefficients in A) where each x_i for $i \notin I$ appears with non-negative degree. Your construction should be such that the restriction map $\Gamma(U_I, \mathcal{O}(m)) \rightarrow \Gamma(U_J, \mathcal{O}(m))$ ($I \subset J$) corresponds to the natural inclusion: a Laurent polynomial in $\Gamma(U_I, \mathcal{O}(m))$ maps to the *same* Laurent polynomial in $\Gamma(U_J, \mathcal{O}(m))$.

We now give the proof for $n = 2$, because all of the complications present themselves here. You will see how the argument works (without change) for general n (Exercise 19.3.D).

The Čech complex for $\mathcal{O}(m)$ is the degree m part of
(19.3.2.1)

$$\begin{aligned} 0 \longrightarrow & A[x_0, x_1, x_2, x_0^{-1}] \times A[x_0, x_1, x_2, x_1^{-1}] \times A[x_0, x_1, x_2, x_2^{-1}] \longrightarrow \\ & A[x_0, x_1, x_2, x_0^{-1}, x_1^{-1}] \times A[x_0, x_1, x_2, x_1^{-1}, x_2^{-1}] \times A[x_0, x_1, x_2, x_0^{-1}, x_2^{-1}] \\ & \longrightarrow A[x_0, x_1, x_2, x_0^{-1}, x_1^{-1}, x_2^{-1}] \longrightarrow 0. \end{aligned}$$

Rather than consider $\mathcal{O}(m)$ for each m independently, it is notationally simpler to consider them all at once, by considering $\mathcal{F} = \bigoplus_{m \in \mathbb{Z}} \mathcal{O}(m)$: the Čech complex for \mathcal{F} is (19.3.2.1). It is useful to write which U_I corresponds to which factor (see (19.3.2.2) below). The maps (from one factor of one term to one factor of the next) are all natural inclusions, or negative of natural inclusions, and in particular preserve degree.

We extend (19.3.2.1) by replacing the $0 \rightarrow$ on the left by $0 \rightarrow A[x_0, x_1, x_2] \rightarrow$:
(19.3.2.2)

$$H^0 \qquad \qquad U_0 \ U_1 \ U_2 \qquad \qquad U_{012}$$

$$0 \longrightarrow A[x_0, x_1, x_2] \longrightarrow \cdots \longrightarrow \cdots \longrightarrow A[x_0, x_1, x_2, x_0^{-1}, x_1^{-1}, x_2^{-1}] \longrightarrow 0.$$

19.3.C. EXERCISE. Show that if (19.3.2.2) is exact, except that at U_{012} the cohomology/cokernel is

$$x_0^{-1} x_1^{-1} x_2^{-1} A[x_0^{-1}, x_1^{-1}, x_2^{-1}],$$

then Theorem 19.1.2 holds for $n = 2$. (Hint: Remark 19.3.1.)

Because the maps in (19.3.2.2) preserve multidegree (degrees of each x_i independently), we can study exactness of (19.3.2.2) monomial by monomial.

The “3 negative exponents” case. Consider first the monomial $x_0^{a_0} x_1^{a_1} x_2^{a_2}$, where the exponents a_i are all negative. Then (19.3.2.2) in this multidegree is:

$$0 \longrightarrow 0_{H^0} \longrightarrow 0_0 \times 0_1 \times 0_2 \longrightarrow 0_{01} \times 0_{12} \times 0_{02} \longrightarrow A_{012} \longrightarrow 0.$$

Here the subscripts serve only to remind us which “Čech” terms the factors correspond to. (For example, A_{012} corresponds to the coefficient of $x_0^{a_0} x_1^{a_1} x_2^{a_2}$ in $A[x_0, x_1, x_2, x_0^{-1}, x_1^{-1}, x_2^{-1}]$.) Clearly this complex only has (co)homology at the U_{012} spot, as desired.

The “2 negative exponents” case. Consider next the case where *two* of the exponents, say a_0 and a_1 , are negative. Then the complex in this multidegree is

$$0 \longrightarrow 0_{H^0} \longrightarrow 0_0 \times 0_1 \times 0_2 \longrightarrow A_{01} \times 0_{12} \times 0_{02} \longrightarrow A_{012} \longrightarrow 0,$$

which is clearly exact.

The “1 negative exponent” case. We next consider the case where *one* of the exponents, say a_0 , is negative. Then the complex in this multidegree is

$$0 \longrightarrow 0_{H^0} \longrightarrow A_0 \times 0_1 \times 0_2 \longrightarrow A_{01} \times 0_{12} \times A_{02} \longrightarrow A_{012} \longrightarrow 0$$

With a little thought (paying attention to the signs on the arrows $A \rightarrow A$), you will see that it is exact. (The subscripts, by reminding us of the subscripts in the original Čech complex, remind us what signs to take in the maps.)

The “0 negative exponent” case. Finally, consider the case where *none* of the exponents are negative. Then the complex in this multidegree is

$$0 \longrightarrow A_{H^0} \longrightarrow A_0 \times A_1 \times A_2 \longrightarrow A_{01} \times A_{12} \times A_{02} \longrightarrow A_{012} \longrightarrow 0$$

We wish to show that this is exact. We write this complex as the middle of a short exact sequence of complexes:

(19.3.2.3)

$$\begin{array}{ccccccccc} 0 & \longrightarrow & 0 & \longrightarrow & A_2 & \longrightarrow & A_{02} \times A_{12} & \longrightarrow & A_{012} & \longrightarrow & 0 \\ \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & A_{H^0} & \longrightarrow & A_0 \times A_1 \times A_2 & \longrightarrow & A_{01} \times A_{12} \times A_{02} & \longrightarrow & A_{012} & \longrightarrow & 0 \\ \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & A_{H^0} & \longrightarrow & A_0 \times A_1 & \longrightarrow & A_{01} & \longrightarrow & 0 & \longrightarrow & 0 \end{array}$$

Thus we get a long exact sequence in cohomology (Theorem 1.6.6). But the top and bottom rows are exact (basically from the “1-negative” case), i.e., cohomology-free, so the middle row must be exact too.

19.3.D. EXERCISE. Prove Theorem 19.1.2 for general n . (I could of course just have given you the proof for general n , but seeing the argument in action may be enlightening. In particular, your argument may be much shorter. For example, the “2-negative” case could be done in the same way as the “1-negative” case, so you will not need $n + 1$ separate cases if you set things up carefully.)

19.3.3. Remark. This argument is basically the proof that the reduced homology of the boundary of a simplex S (known in some circles as a “sphere”) is 0, unless S is the empty set, in which case it is one-dimensional. The “empty set” case corresponds to the “3-negative” case. \square

19.4 Riemann-Roch, and arithmetic genus

We have seen some powerful uses of Čech cohomology, to prove things about spaces of global sections, and to prove Serre vanishing. We will now see some classical constructions come out very quickly and cheaply.

In this section, we will work over a field k . Suppose \mathcal{F} is a coherent sheaf on a projective k -scheme X . Recall the notation (§19.1) $h^i(X, \mathcal{F}) := \dim_k H^i(X, \mathcal{F})$. By Theorem 19.1.3, $h^i(X, \mathcal{F})$ is finite. (The arguments in this section will extend without change to proper X once we have this finiteness for proper morphisms, by Grothendieck’s Coherence Theorem 19.9.1.) Define the **Euler characteristic** of \mathcal{F} by

$$\chi(X, \mathcal{F}) := \sum_{i=0}^{\dim X} (-1)^i h^i(X, \mathcal{F}).$$

We will see repeatedly here and later that Euler characteristics behave better than individual cohomology groups. As one sign, notice that for fixed n , and $m \geq 0$,

$$h^0(\mathbb{P}_k^n, \mathcal{O}(m)) = \binom{n+m}{m} = \frac{(m+1)(m+2) \cdots (m+n)}{n!}.$$

Notice that the expression on the right is a polynomial in m of degree n . (For later reference, notice also that the leading term is $m^n/n!$.) But it is not true that

$$h^0(\mathbb{P}_k^n, \mathcal{O}(m)) = \frac{(m+1)(m+2) \cdots (m+n)}{n!}$$

for *all* m — it breaks down for $m \leq -n-1$. Still, you can check (using Theorem 19.1.2) that

$$\chi(\mathbb{P}_k^n, \mathcal{O}(m)) = \frac{(m+1)(m+2) \cdots (m+n)}{n!}.$$

So one lesson is this: if one cohomology group (usually the top or bottom) behaves well in a certain range, and then messes up, likely it is because (i) it is actually the Euler characteristic which behaves well *always*, and (ii) the other cohomology groups vanish in that certain range.

In fact, we will see that it is often hard to calculate cohomology groups (even h^0), but it can be easier calculating Euler characteristics. So one important way of getting a hold of cohomology groups is by computing the Euler characteristic, and then showing that all the *other* cohomology groups vanish. Hence the ubiquity and importance of *vanishing theorems*. (A vanishing theorem usually states that a certain cohomology group vanishes under certain conditions.) We will see this in action when discussing curves. (One of the first applications will be (20.2.5.1).)

The following exercise shows another way in which Euler characteristic behaves well: it is *additive in exact sequences*.

19.4.A. EXERCISE. Show that if $0 \rightarrow \mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{H} \rightarrow 0$ is an exact sequence of coherent sheaves on a projective k -scheme X , then $\chi(X, \mathcal{G}) = \chi(X, \mathcal{F}) + \chi(X, \mathcal{H})$. (Hint: consider the long exact sequence in cohomology.) More generally, if

$$0 \rightarrow \mathcal{F}_1 \rightarrow \cdots \rightarrow \mathcal{F}_n \rightarrow 0$$

is an exact sequence of coherent sheaves, show that

$$\sum_{i=1}^n (-1)^i \chi(X, \mathcal{F}_i) = 0.$$

(This exercise both generalizes the “exact” case of Exercise 1.6.B — consider the case where $X = \operatorname{Spec} k$ — and uses it in the proof.)

19.4.1. The Riemann-Roch Theorem for line bundles on a regular projective curve. Suppose $D := \sum_{p \in C} a_p [p]$ is a divisor on a regular projective curve C over a field k (where $a_p \in \mathbb{Z}$, and all but finitely many a_p are 0). Define the **degree of D** by

$$\deg D = \sum a_p \deg p.$$

(The degree of a point p was defined in §5.3.8, as the degree of the field extension of the residue field over k .)

19.4.B. ESSENTIAL EXERCISE: THE RIEMANN-ROCH THEOREM FOR LINE BUNDLES ON A REGULAR PROJECTIVE CURVE. Show that

$$\chi(C, \mathcal{O}_C(D)) = \deg D + \chi(C, \mathcal{O}_C)$$

by induction on $\sum |a_p|$ (where $D = \sum a_p [p]$ as above). Hint: to show that $\chi(C, \mathcal{O}_C(D)) = \deg p + \chi(C, \mathcal{O}_C(D - p))$, tensor the closed subscheme exact sequence (9.1.2.1)

$$0 \rightarrow \mathcal{O}_C(-p) \rightarrow \mathcal{O}_C \rightarrow \mathcal{O}_p \rightarrow 0$$

(where \mathcal{O}_p is the structure sheaf of the scheme p , not the stalk $\mathcal{O}_{C,p}$) by $\mathcal{O}_C(D)$, and use additivity of Euler characteristics in exact sequences (Exercise 19.4.A).

As every invertible sheaf \mathcal{L} is of the form $\mathcal{O}_C(D)$ for some D (see §15.2), this exercise is very powerful.

19.4.2. Important definition: degree of a line bundle on a reduced projective curve (generalized in §19.4.6). Suppose \mathcal{L} is an invertible sheaf on a reduced projective curve C over k . Define the **degree of \mathcal{L}** (denoted $\deg \mathcal{L}$ or $\deg_C \mathcal{L}$) by

$$\deg \mathcal{L} := \chi(C, \mathcal{L}) - \chi(C, \mathcal{O}_C).$$

19.4.C. IMPORTANT EASY EXERCISE: DEGREE OF A DIVISOR = DEGREE OF THE CORRESPONDING LINE BUNDLE. Suppose \mathcal{L} is an invertible sheaf on a *regular* projective curve C over k , and let s be a nonzero rational section of \mathcal{L} on C . Let D be the divisor of zeros and poles of s :

$$D := \sum_{p \in C} v_p(s)[p].$$

Show that $\deg \mathcal{L} = \deg D$. In particular, the degree can be computed by counting zeros and poles of *any* section not vanishing on a component of C .

The next few exercises give useful applications.

19.4.D. EXERCISE. Give a new solution to Exercise 18.4.E (a nonzero rational function on a projective curve has the same number of zeros and poles, counted appropriately) using the ideas above.

19.4.E. EXERCISE. Suppose s_1 and s_2 are two sections of a degree d line bundle \mathcal{L} on an irreducible regular curve C , with no common zeros. Then s_1 and s_2 determine a morphism $\pi: C \rightarrow \mathbb{P}_k^1$. Show that the degree of π is d . (Translation: a two-dimensional base-point-free degree d linear system on C defines a degree d cover of \mathbb{P}^1 .)

19.4.F. EXERCISE. If \mathcal{L} and \mathcal{M} are two line bundles on a regular projective curve C , show that $\deg \mathcal{L} \otimes \mathcal{M} = \deg \mathcal{L} + \deg \mathcal{M}$. (Hint: choose nonzero rational sections of \mathcal{L} and \mathcal{M} .)

19.4.G. EXERCISE. Suppose $\pi: C \rightarrow C'$ is a degree d morphism of integral projective regular curves, and \mathcal{L} is an invertible sheaf on C' . Show that $\deg_C \pi^* \mathcal{L} = d \deg_{C'} \mathcal{L}$. Hint: compute $\deg_C \mathcal{L}$ using any nonzero rational section s of \mathcal{L} , and compute $\deg_C \pi^* \mathcal{L}$ using the rational section $\pi^* s$ of $\pi^* \mathcal{L}$. Note that zeros pull back to zeros, and poles pull back to poles. Depending on your approach, Exercise 18.4.D might help.

19.4.H. ★★ EXERCISE (COMPLEX-ANALYTIC INTERPRETATION OF DEGREE; ONLY FOR THOSE WITH SUFFICIENT ANALYTIC BACKGROUND). Suppose X is a connected regular projective complex curve. Show that the degree map is the composition of group homomorphisms

$$\mathrm{Pic} X \longrightarrow \mathrm{Pic} X_{\mathrm{an}} \xrightarrow{c_1} H^2(X_{\mathrm{an}}, \mathbb{Z}) \xrightarrow{\cap [X_{\mathrm{an}}]} H_0(X_{\mathrm{an}}, \mathbb{Z}) \cong \mathbb{Z}.$$

Hint: show it for a generator $\mathcal{O}(p)$ of the group $\mathrm{Pic} X$, using explicit transition functions. (The first map was discussed in Exercise 14.2.I. The second map takes a line bundle to its first Chern class, and can be interpreted as follows. The transition functions for a line bundle yield a Čech 1-cycle for $\mathcal{O}_{X_{\mathrm{an}}}^*$; this yields a map $\mathrm{Pic} X_{\mathrm{an}} \rightarrow H^1(X_{\mathrm{an}}, \mathcal{O}_{X_{\mathrm{an}}}^*)$. Combining this with the map $H^1(X_{\mathrm{an}}, \mathcal{O}_{X_{\mathrm{an}}}^*) \rightarrow H^2(X_{\mathrm{an}}, \mathbb{Z})$ from the long exact sequence in cohomology corresponding to the exponential exact sequence (2.4.9.1) yields the first Chern class map.)

19.4.3. Arithmetic genus.

Motivated by geometry (Miracle 19.4.4 below), we define the **arithmetic genus** of a scheme X as $1 - \chi(X, \mathcal{O}_X)$. This is sometimes denoted $p_a(X)$. For integral projective curves over an algebraically closed field, $h^0(X, \mathcal{O}_X) = 1$ (§11.5.7), so $p_a(X) = h^1(X, \mathcal{O}_X)$. (In higher dimension, this is a less natural notion.)

We can restate the Riemann-Roch formula for curves (Exercise 19.4.B) as:

$$(19.4.3.1) \quad h^0(C, \mathcal{L}) - h^1(C, \mathcal{L}) = \deg \mathcal{L} - p_a(C) + 1.$$

This is the most common formulation of the Riemann-Roch formula.

19.4.4. Miracle. If C is a regular irreducible projective complex curve, then the corresponding complex-analytic object, a compact *Riemann surface*, has a notion called the *genus* g , which is (informally speaking) the number of holes (see Figure 19.2). Miraculously, $g = p_a$ in this case (see Exercise 22.7.I), and for this reason, we

will often write g for p_a when discussing regular (projective irreducible) curves, over any field. We will discuss genus further in §19.6.6, when we will be able to compute it in many interesting cases. (Warning: the arithmetic genus of $\mathbb{P}_{\mathbb{C}}^1$ as an \mathbb{R} -variety is -1 !)




FIGURE 19.2. A genus 3 Riemann surface

19.4.5. Degree and rank of a coherent sheaf.

Suppose C is an irreducible reduced curve (possibly singular), over a field k . If \mathcal{F} is a coherent sheaf on C , recall (from §14.3.3) that the *rank* of \mathcal{F} , denoted $\text{rank } \mathcal{F}$, is its rank at the generic point of C .

19.4.I. EASY EXERCISE. Show that the rank is additive in exact sequences: if

$$0 \rightarrow \mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{H} \rightarrow 0$$

is an exact sequence of coherent sheaves, show that

$$\text{rank } \mathcal{F} - \text{rank } \mathcal{G} + \text{rank } \mathcal{H} = 0.$$

Hint: localization is exact. (Caution: your argument will use the fact that the rank is at the generic point; the example

$$0 \longrightarrow \widetilde{k[t]} \xrightarrow{\times t} \widetilde{k[t]} \longrightarrow \widetilde{k[t]/(t)} \longrightarrow 0$$

on \mathbb{A}_k^1 shows that rank at a closed point is not additive in exact sequences.)

19.4.6. Definition. If C is *projective*, define the **degree** of \mathcal{F} by

$$(19.4.6.1) \quad \deg \mathcal{F} = \chi(C, \mathcal{F}) - (\text{rank } \mathcal{F}) \cdot \chi(C, \mathcal{O}_C).$$

The statement (19.4.6.1) is often called **Riemann-Roch for coherent sheaves** (or vector bundles) on a projective curve.

If \mathcal{F} is an invertible sheaf (or if more generally the rank is the same on each irreducible component), we can drop the irreducibility hypothesis. This then generalizes the notion of the degree of a line bundle on a reduced curve in (Important Definition 19.4.2).

We now study the behavior of this notion. (In Exercise 22.7.B, you will show that if \mathcal{F} is supported at a finite number of points, the degree of \mathcal{F} splits up into a contribution from each point.)

19.4.J. EASY EXERCISE. Show that degree (as a function of coherent sheaves on a fixed curve C) is additive in exact sequences.

19.4.K. EXERCISE. If \mathcal{F} is a torsion sheaf on C , show that $\deg \mathcal{F} \geq 0$, with equality if and only $\mathcal{F} = 0$.

19.4.L. EXERCISE. Show that the degree of a vector bundle is the degree of its determinant bundle. Hint: Exercise 14.2.W.

19.4.M. EXERCISE. If C is a projective curve, and \mathcal{L} is an ample line bundle on C , show that $\deg \mathcal{L} > 0$. (Hint: show it if \mathcal{L} is *very* ample.)

19.4.N. EXERCISE. Suppose \mathcal{L} is a base-point-free invertible sheaf on a proper variety X , and hence induces some morphism $\phi: X \rightarrow \mathbb{P}^n$. Then \mathcal{L} is ample if and only if ϕ is finite. (Hint: if ϕ is finite, use Exercise 17.6.G. If ϕ is not finite, show that there is a curve C contracted by ϕ , using Theorem 19.1.7. Show that \mathcal{L} has degree 0 on C .)

19.4.7. Extended example: the universal plane conic has no rational sections.

We use the theory of the degree to get an interesting consequence. We work over a fixed field k . We consider the following diagram.

$$\begin{array}{ccc} \mathcal{C} & \xrightarrow{\text{cl. emb.}} & \mathbb{P}^2 \times \mathbb{P}^5 \longrightarrow \mathbb{P}^2 \\ & \searrow & \downarrow \\ & & \mathbb{P}^5 \end{array}$$

If the \mathbb{P}^2 has projective coordinates x_0, x_1, x_2 , then \mathbb{P}^5 has coordinates $a_{00}, a_{01}, a_{11}, a_{02}, a_{12}, a_{22}$, and \mathcal{C} is cut out by the single equation

$$a_{00}x_0^2 + a_{01}x_0x_1 + \cdots + a_{22}x_2^2 = 0.$$

We interpret \mathbb{P}^5 as the parameter space of conics (in \mathbb{P}^2), and \mathcal{C} as the universal conic over \mathbb{P}^5 (parametrizing a conic C along with a point $p \in C$), which comes with a canonical projection $\mathcal{C} \rightarrow \mathbb{P}^5$.

19.4.O. EXERCISE. By interpreting \mathcal{C} as a \mathbb{P}^4 -bundle over \mathbb{P}^5 , show that \mathcal{C} is a smooth sixfold, and that $\text{Pic } \mathcal{C} \cong \mathbb{Z} \times \mathbb{Z}$.

19.4.P. EXERCISE. Fix a line $\ell \subset \mathbb{P}^2$ and a k -valued point $q \in \mathbb{P}^2$. Let D_ℓ be the divisor on \mathcal{C} corresponding to (C, p) with p lying on ℓ . Let D_q be the divisor on \mathcal{C} corresponding to (C, p) with $q \in C$. Using your description of \mathcal{C} as a \mathbb{P}^4 -bundle over \mathbb{P}^5 , show that D_ℓ and D_q generate $\text{Pic } \mathcal{C}$.

19.4.Q. EXERCISE. Suppose K is a fiber of $\mathcal{C} \rightarrow \mathbb{P}^5$ over a point $r \in \mathbb{P}^5$ — i.e., a conic in \mathbb{P}^2 over the field $\kappa(r)$. Suppose further that neither q nor ℓ are contained in K . (This hypothesis is unnecessary, but simplifies the problem.) Show that $\deg_K \mathcal{O}(D_q) = 0$ and $\deg_K \mathcal{O}(D_\ell) = 2$. Hence show that if \mathcal{L} is any invertible sheaf on \mathcal{C} , then $\deg_K(\mathcal{L})$ is even.

19.4.R. EXERCISE. Show that there is no rational section to the projection $\pi: \mathcal{C} \rightarrow \mathbb{P}^5$. Hint: if there were a regular section over an open subset U of \mathbb{P}^5 , it would be a divisor on $\pi^{-1}(U)$; let D be its closure in \mathcal{C} . Show that D meets any fiber of π over U in multiplicity 1. Use Exercise 19.4.Q to obtain a contradiction.

We can restate Exercise 19.4.R in the following dramatic way: there is no way to write down three rational functions X_0, X_1, X_2 in a_{00}, \dots, a_{22} such that

$$a_{00}X_0^2 + a_{01}X_0X_1 + \cdots + a_{22}X_2^2 = 0$$

without $X_0 = X_1 = X_2 = 0$.

The question of a rational point on a conic is one of arithmetic. (Think: $x^2 + y^2 = z^2$.) Our solution was topological. The unification of topology and arithmetic in this example is the beginning of a long and fruitful story in algebraic geometry.

19.4.S. ★ EXERCISE. Suppose C is a projective curve over a field k , \mathcal{L} is a line bundle on C , and \mathcal{F} is a coherent sheaf on C . Show that $\chi(\mathcal{L} \otimes \mathcal{F}) - \chi(\mathcal{F})$ is the sum over the irreducible components C_i of C of the degree of \mathcal{L} on C_i^{red} (the reduction was defined in §9.4.10) times the length of \mathcal{F} at the generic point η_i of C_i (the length of \mathcal{F}_{η_i} as an \mathcal{O}_{η_i} -module). This could reasonably be called **Riemann-Roch for nonreduced curves**. Hints: (1) First reduce to the case where \mathcal{F} is scheme-theoretically supported on C^{red} , by showing that both sides of the alleged equality are additive in short exact sequences, and using the filtration

$$0 = \mathcal{I}^r \mathcal{F} \subset \mathcal{I}^{r-1} \mathcal{F} \subset \cdots \subset \mathcal{I} \mathcal{F} \subset \mathcal{F}$$

of \mathcal{F} , where \mathcal{I} is the ideal sheaf cutting out C^{red} in C . Thus we need only consider the case where C is reduced. (2) As C is projective, we can write $\mathcal{L} \cong \mathcal{O}(\sum n_j p_j)$ where the p_j are regular points distinct from the associated points of \mathcal{F} . Use this avatar of \mathcal{L} , and perhaps induction on the number of p_j .

In fact, all proper curves over k are projective (Remark 19.7.2), so “projective” can be replaced by “proper” in Exercise 19.4.S. In this guise, we will use Exercise 19.4.S when discussing intersection theory in Chapter 21.

19.4.8. ★ Numerical equivalence, the Néron-Severi group, nef line bundles, and the nef and ample cones.

The notion of a degree on a line bundle leads to important and useful notions. Suppose X is a proper k -variety, and \mathcal{L} is an invertible sheaf on X . If $i: C \hookrightarrow X$ is a one-dimensional closed reduced subscheme of X , define the degree of \mathcal{L} on C by $\deg_C \mathcal{L} := \deg_C i^* \mathcal{L}$. If $\deg_C \mathcal{L} = 0$ for all C , we say that \mathcal{L} is **numerically trivial**.

19.4.T. EXERCISE.

- (a) Show that \mathcal{L} is numerically trivial if and only if $\deg_C \mathcal{L} = 0$ for all *integral* curves C in X .
- (b) Show that if $\pi: X \rightarrow Y$ is a proper morphism, and \mathcal{L} is a numerically trivial invertible sheaf on Y , then $\pi^* \mathcal{L}$ is numerically trivial on X .
- (c) Show that \mathcal{L} is numerically trivial if and only if \mathcal{L} is numerically trivial on each of the irreducible components of X .
- (d) Show that if \mathcal{L} and \mathcal{L}' are numerically trivial, then $\mathcal{L} \otimes \mathcal{L}'$ and \mathcal{L}^\vee are both numerically trivial.

19.4.9. Numerical equivalence. By part (d), the numerically trivial invertible sheaves form a subgroup of $\text{Pic } X$, denoted $\text{Pic}^\tau X$. Two line bundles equivalent modulo the subgroup of numerically trivial line bundles are called **numerically equivalent**. A property of invertible sheaves stable under numerical equivalence is said to be a *numerical property*. We will see that “nefness” and ampleness are numerical properties (Definition 19.4.10 and Remark 21.4.2 respectively).

We will later define the *Néron-Severi group* $\text{NS}(X)$ of X as $\text{Pic } X$ modulo algebraic equivalence (Exercise 25.7.5). (We will define algebraic equivalence once we

have discussed flatness.) The highly nontrivial **Néron-Severi Theorem** (or **Theorem of the Base**) states that $\text{NS}(X)$ is a finitely generated group. (The proof is quite difficult; see [K11, p. 334, Prop. 3]. For a simpler proof over \mathbb{C} , see [GH1, p. 462].) The group $\text{Pic } X / \text{Pic}^\tau X$ is denoted $N^1(X)$. We will see (in §25.7.5) that it is a quotient of $\text{NS}(X)$, so it is also finitely generated. As the group $N^1(X)$ is clearly abelian and torsion-free, it is finite free \mathbb{Z} -module (by the classification of finitely generated modules over a principal ideal domain, see §0.3). The rank of $N^1(X)$ is called the **Picard number**, and is denoted $\rho(X)$ (although we won't have need of this notion, except in our discussion of the Hodge Index Theorem in §21.2.10). For example, $\rho(\mathbb{P}^n) = 1$ and $\rho((\mathbb{P}^1)^n) = n$. We define $N_{\mathbb{Q}}^1(X) := N^1(X) \otimes_{\mathbb{Z}} \mathbb{Q}$ (so $\rho(X) = \dim_{\mathbb{Q}} N_{\mathbb{Q}}^1(X)$), and call the elements of this group **\mathbb{Q} -line bundles**, for lack of any common term in the literature.

19.4.U. ★★ EXERCISE (FINITENESS OF PICARD NUMBER IN THE COMPLEX CASE, ONLY FOR THOSE WITH SUFFICIENT BACKGROUND). Show (without the Néron-Severi Theorem) that if X is a complex proper variety, then $\rho(X)$ is finite, by interpreting it as a subquotient of $H^2(X, \mathbb{Z})$. Hint: show that the image of (\mathcal{L}, C) under the map $H^2(X, \mathbb{Z}) \times H_2(X, \mathbb{Z}) \rightarrow H_0(X, \mathbb{Z}) \rightarrow \mathbb{Z}$ is $\deg_C \mathcal{L}$. Hint: figure out how to reduce to the case where C is a smooth projective curve, then use Exercise 19.4.H.

19.4.10. Definition. We say that an invertible sheaf \mathcal{L} is **numerically effective**, or **nef** if for all such C , $\deg_C \mathcal{L} \geq 0$. Clearly nefness is a numerical property.

19.4.V. EASY EXERCISE (CF. EXERCISE 19.4.T).

- (a) Show that \mathcal{L} is nef if and only if $\deg_C \mathcal{L} \geq 0$ for all *integral* curves C in X .
- (b) Show that if $\pi: X \rightarrow Y$ is a proper morphism, and \mathcal{L} is a nef invertible sheaf on Y , then $\pi^* \mathcal{L}$ is nef on X . (Hint: Exercise 19.4.G will be needed.)
- (c) Show that \mathcal{L} is nef if and only if \mathcal{L} is nef on each of the irreducible components of X .
- (d) Show that if \mathcal{L} and \mathcal{L}' are nef, then $\mathcal{L} \otimes \mathcal{L}'$ is nef. Thus the nef elements of $\text{Pic } X$ form a semigroup.
- (e) Show that ample invertible sheaves are nef.
- (f) Suppose $n \in \mathbb{Z}^{>0}$. Show that \mathcal{L} is nef if and only if $\mathcal{L}^{\otimes n}$ is nef.

19.4.W. EXERCISE. Define what it means for a \mathbb{Q} -line bundle to be nef. Show that the nef \mathbb{Q} -line bundles form a closed cone in $N_{\mathbb{Q}}^1(X)$. This is called the **nef cone**.

19.4.X. EXERCISE. Describe the nef cones of \mathbb{P}_k^2 and $\mathbb{P}_k^1 \times_k \mathbb{P}_k^1$. (Notice in the latter case that the two boundaries of the cone correspond to linear series contracting one of the \mathbb{P}^1 's. This is true in general: informally speaking, linear series corresponding to the boundaries of the cone give interesting contractions. Another example will be given in Exercise 21.2.F.)

It is a surprising fact that whether an invertible sheaf \mathcal{L} on X is ample depends only on its class in $N_{\mathbb{Q}}^1(X)$, i.e., on how it intersects the curves in X . Because of this (as for any $n \in \mathbb{Z}^{\geq 0}$, \mathcal{L} is ample if and only if $\mathcal{L}^{\otimes n}$ is ample, see Theorem 17.6.2), it makes sense to define when a \mathbb{Q} -line bundle is ample. Then by Exercise 17.6.H, the ample divisors form a cone in $N_{\mathbb{Q}}^1(X)$, necessarily contained in the nef cone by Exercise 19.4.V(e). It turns out that if X is projective, the ample divisors are precisely

the interior of the nef cone. The new facts in this paragraph are a consequence of Kleiman's numerical criterion for ampleness, Theorem 21.4.6.

19.5 A first glimpse of Serre duality

A common version of Riemann-Roch involves Serre duality, which unlike Riemann-Roch, is *hard*.

19.5.1. Theorem (Serre duality for smooth projective varieties). — Suppose X is a geometrically irreducible smooth projective k -variety, of dimension n . Then there is an invertible sheaf ω_X (or simply ω) on X such that

$$h^i(X, \mathcal{F}) = h^{n-i}(X, \omega_X \otimes \mathcal{F}^\vee)$$

for all $i \in \mathbb{Z}$ and all finite rank locally free sheaves \mathcal{F} .

The invertible sheaf ω_X is an example of a *dualizing sheaf*, which will be formally defined in §31.1.4. We will see in Chapter 31 that Theorem 19.5.1 is a consequence of a perfect pairing

$$(19.5.1.1) \quad H^i(X, \mathcal{F}) \times H^{n-i}(X, \omega_X \otimes \mathcal{F}^\vee) \rightarrow H^n(X, \omega_X) \cong k,$$

and that smoothness can be relaxed somewhat.

19.5.2. Further miracle: the sheaf of algebraic volume forms is Serre-dualizing. The invertible sheaf ω_X turns out to be the *canonical sheaf* (or *canonical bundle*) \mathcal{K}_X , which is defined as the determinant of the cotangent bundle $\Omega_{X/k}$ of X (see §22.5.3); we will define the cotangent bundle in Chapter 22. We connect the dualizing sheaf to the canonical bundle in §31.4; see Desideratum 31.1.1.

19.5.3. Back to Riemann-Roch. For the purposes of restating Riemann-Roch for a curve C , it suffices to note that $h^1(C, \mathcal{L}) = h^0(C, \omega_C \otimes \mathcal{L}^\vee)$. Then the Riemann-Roch formula can be rewritten as

$$h^0(C, \mathcal{L}) - h^0(C, \omega_C \otimes \mathcal{L}^\vee) = \deg \mathcal{L} - p_a(C) + 1.$$

19.5.A. EXERCISE (ASSUMING SERRE DUALITY). Suppose C is a geometrically integral smooth projective curve over k .

(a) Show that $h^0(C, \omega_C)$ is the genus g of C . (The genus g , in the guise of the arithmetic genus p_a , was defined in §19.4.3–19.4.4.)

(b) Show that $\deg \omega_C = 2g - 2$. (Hint: Riemann-Roch for $\mathcal{L} = \omega_C$.)

19.5.4. Example. If $C = \mathbb{P}_k^1$, Exercise 19.5.A implies that $\omega_C \cong \mathcal{O}(-2)$. Moreover, we also have a natural perfect pairing (cf. (19.5.1.1))

$$H^0(\mathbb{P}^1, \mathcal{O}(n)) \times H^1(\mathbb{P}^1, \mathcal{O}(-2-n)) \rightarrow k \cong H^1(\mathbb{P}^1, \mathcal{O}(-2)).$$

We can interpret this pairing as follows. If $n < 0$, both factors on the left are 0, so we assume $n > 0$. Then $H^0(\mathbb{P}^1, \mathcal{O}(n))$ corresponds to homogeneous degree n polynomials in x and y , and $H^1(\mathbb{P}^1, \mathcal{O}(-2-n))$ corresponds to homogeneous degree $-2-n$ Laurent polynomials in x and y so that the degrees of x and y are both at most -1 (see Remark 19.3.1). You can quickly check that the dimension of both vector spaces are $n+1$. The pairing is given as follows: multiply the

polynomial by the Laurent polynomial, to obtain a Laurent polynomial of degree -2 . Read off the coefficient of $x^{-1}y^{-1}$. (This works more generally for \mathbb{P}_A^n ; see the discussion after the statement of Theorem 19.1.2.)

19.5.B. EXERCISE (ASSUMING SERRE DUALITY): AMPLE DIVISORS ON A CONNECTED SMOOTH PROJECTIVE VARIETY ARE CONNECTED. Suppose X is a connected smooth projective \bar{k} -variety of dimension at least 2, and D is an effective ample divisor. Show that D is connected. (Hint: Suppose $D = V(s)$, where s is a section of an ample invertible sheaf \mathcal{L} . Then $V(s^n) = V(s)$ (as vanishing *sets*, not vanishing *schemes*) for all $n > 0$, so we may replace \mathcal{L} with a high power of our choosing. Use the long exact sequence for $0 \rightarrow \mathcal{O}_X(-nD) \rightarrow \mathcal{O}_X \rightarrow \mathcal{O}_{V(s^n)} \rightarrow 0$ to show that for $n \gg 0$, $h^0(X, \mathcal{O}_{V(s^n)}) = 1$.)

Once we know that Serre duality holds for Cohen-Macaulay projective schemes (§31.3), this result will automatically extend to these schemes. (A related result is Exercise 19.6.U, which doesn't use Serre duality.) On the other hand, the result is false if X is the union of two 2-planes in \mathbb{P}^4 meeting at a point (why?), so this will imply that this X is not Cohen-Macaulay. (We will show this in another way in Counterexample 29.2.2.)

19.5.5. ★ Classification of vector bundles on \mathbb{P}_k^1 .

As promised in Example 18.2.4, we classify the vector bundles on \mathbb{P}_k^1 . We discuss this in §19.5 because you might use Serre duality at one step (although you might not).

19.5.6. Theorem. — *If \mathcal{E} is a rank r vector bundle on \mathbb{P}_k^1 , then $\mathcal{E} \cong \mathcal{O}(a_1) \oplus \cdots \oplus \mathcal{O}(a_r)$, for a unique nondecreasing sequence of integers a_1, \dots, a_r .*

This result was proved independently many times (see for example Dedekind and Weber's article [DW, §22]), and is a special case of a theorem of Grothendieck, [Gr2]. It is sometimes called Grothendieck's Theorem, because Grothendieck doesn't have enough theorems named after him.

19.5.7. For \mathbb{P}_k^n more generally, the case $r = 1$ was shown in §15.2.9, but the statement is false for $r > 1$ and $n > 1$. A counterexample is given in Exercise 22.4.G, which shows that it is not even possible to *filter* the rank two bundle $\Omega_{\mathbb{P}^2}$ on \mathbb{P}^2 by line bundles. One true generalization is a theorem of Horrocks, which states that a finite rank locally free sheaf \mathcal{E} on \mathbb{P}_k^n splits precisely when “all the middle cohomology of all of its twists is zero” — when $H^i(\mathbb{P}_k^n, \mathcal{E}(m)) = 0$ for $0 < i < n$ and all $m \in \mathbb{Z}$. This has the surprising consequence that if $n \geq 2$, then a finite rank locally free sheaf on \mathbb{P}_k^n splits if and only if its restriction to some previously chosen 2-plane (\mathbb{P}_k^2) splits — all of the additional complication turns up already in dimension 2. See [OSS, §2.3] for more.

Proof. Note that the classification makes no reference to cohomology, but the proof uses cohomology in an essential way. It is possible to prove Theorem 19.5.6 with no cohomological machinery (see for example [HM] or [GW, p. 314–5]), or with more cohomological machinery (see for example [Ha1, Ex. V.2.6]).

19.5.C. EXERCISE. Suppose that $\mathcal{E} \cong \mathcal{O}(a_1) \oplus \cdots \oplus \mathcal{O}(a_r)$, with $a_1 \leq \cdots \leq a_r$. Show that the a_i can be determined from the numbers $\dim_k \operatorname{Hom}(\mathcal{O}(m), \mathcal{E})$, as m

ranges over the integers. Use this to show the uniqueness part of Theorem 19.5.6.

We now begin the proof, by induction on r . Fix a rank r locally free sheaf \mathcal{E} . The case $r = 1$ was established in §15.2.9, so we assume $r > 1$.

19.5.D. EXERCISE. Show that for $m \ll 0$, $\text{Hom}(\mathcal{O}(m), \mathcal{E}) \neq 0$, and that for $m \gg 0$, $\text{Hom}(\mathcal{O}(m), \mathcal{E}) = 0$. Hint: show that $\mathcal{H}om(\mathcal{O}(m), \mathcal{E}) = \mathcal{E}(-m)$ using Exercise 14.2.F. Use Serre vanishing (Theorem 19.1.3(ii)) for the first part. Feel free to use Serre duality and Serre vanishing for the second part. But you may prefer to come up with an argument without Serre duality, to avoid invoking something we have not yet proved.

Thus there is some a_r for which $\dim \text{Hom}(\mathcal{O}(a_r), \mathcal{E}) > 0$, but for which $\text{Hom}(\mathcal{O}(m), \mathcal{E}) = 0$ for all $m > a_r$. Choose a nonzero map $\phi: \mathcal{O}(a_r) \rightarrow \mathcal{E}$.

19.5.E. EXERCISE. Show that ϕ is an injection. (Hint: $\mathcal{O}(a_r)$ is torsion-free, and thus the kernel is torsion-free.)

19.5.F. EXERCISE. Let \mathcal{F} be the cokernel of ϕ . Show that \mathcal{F} is locally free. Hint: Exercise 14.3.G(c) gives an exact sequence $0 \rightarrow \mathcal{F}_{\text{tors}} \rightarrow \mathcal{F} \rightarrow \mathcal{F}_{\text{lf}} \rightarrow 0$, where \mathcal{F}_{lf} is locally free. Let \mathcal{L} be the kernel of the surjection $\mathcal{E} \rightarrow \mathcal{F}_{\text{lf}}$. Show that \mathcal{L} is locally free, and thus is isomorphic to $\mathcal{O}(N)$ for some N . Show that there is a nonzero map $\phi': \mathcal{O}(a_r) \rightarrow \mathcal{O}(N)$. Show (using the same idea as the previous exercise) that this nonzero map ϕ' must be an injection, so $a_r \leq N$. Show that $N \leq a_r$ because there is a nonzero map $\mathcal{O}(N) \rightarrow \mathcal{E}$ (recall how a_r was chosen). Show that ϕ' is an isomorphism, and thus that $\mathcal{F} = \mathcal{F}_{\text{lf}}$.

By our inductive hypothesis, we have $\mathcal{F} = \mathcal{O}(a_1) \oplus \cdots \oplus \mathcal{O}(a_{r-1})$, where $a_1 \leq \cdots \leq a_{r-1}$, so we have a short exact sequence

$$(19.5.7.1) \quad 0 \longrightarrow \mathcal{O}(a_r) \longrightarrow \mathcal{E} \longrightarrow \mathcal{O}(a_1) \oplus \cdots \oplus \mathcal{O}(a_{r-1}) \longrightarrow 0.$$

We next show that $a_r \geq a_i$ for $i < r$. We are motivated by the fact that for a quasicoherent sheaf \mathcal{G} on \mathbb{P}^1 ,

$$\text{Hom}_{\mathbb{P}^1}(\mathcal{O}(a_r + 1), \mathcal{G}) = H^0(\mathbb{P}^1, \mathcal{G}(-a_r - 1))$$

(Exercise 14.2.F). Tensor (19.5.7.1) with $\mathcal{O}(-a_r - 1)$ (preserving exactness, by Exercise 14.2.E), and take the long exact sequence in cohomology. Part of the long exact sequence is

$$H^0(\mathbb{P}^1, \mathcal{E}(-a_r - 1)) \longrightarrow H^0(\mathbb{P}^1, \mathcal{F}(-a_r - 1)) \longrightarrow H^1(\mathbb{P}^1, \mathcal{O}(-1)).$$

Notice that $H^0(\mathbb{P}^1, \mathcal{E}(-a_r - 1)) = 0$ (as $\text{Hom}(\mathcal{O}(a_r + 1), \mathcal{E}) = 0$, by the definition of a_r), and $H^1(\mathbb{P}^1, \mathcal{O}(-1)) = 0$ (Theorem 19.1.2(b)). Thus

$$\begin{aligned} 0 &= H^0(\mathbb{P}^1, \mathcal{F}(-a_r - 1)) \\ &= H^0(\mathbb{P}^1, \oplus_i \mathcal{O}(a_i - a_r - 1)) \\ &= \oplus_i H^0(\mathbb{P}^1, \mathcal{O}(a_i - a_r - 1)). \end{aligned}$$

Hence $a_i - a_r - 1 < 0$ (by Exercise 15.1.A), so $a_r \geq a_i$ as desired.

Finally, we wish to show that exact sequence (19.5.7.1) expresses \mathcal{E} as a direct sum (of the subsheaf and quotient sheaf). For simplicity, we focus on the case $r = 2$, and return to the general case in Exercise 19.5.I.

19.5.G. EXERCISE. Show that the transition functions for the vector bundle, in appropriate coordinates, are given by

$$\begin{pmatrix} t^{-a_2} & \alpha(t) \\ 0 & t^{-a_1} \end{pmatrix},$$

where $\alpha(t)$ is a Laurent polynomial in t . Hint/reminder: Recall that all vector bundles on \mathbb{A}_k^1 are trivial, Exercise 14.2.M. Transition matrices for extensions of one vector bundle (with known transition matrices) by another were discussed in Exercise 14.2.Q.

19.5.H. EXERCISE. Implicit in the above 2×2 matrix is a choice of a basis of a rank 2 free module M over the ring $k[t]$ corresponding to one of the standard affine open subsets $\text{Spec } k[t]$, and a rank 2 free module M' over the ring $k[1/t]$ over the other standard open subset. Show that by an appropriate “upper-triangular” change of basis of M , you can arrange for $\alpha(t)$ to have no monomials of degree $\geq -a_2$. Show that by an appropriate “upper-triangular” change of basis of M' , you can arrange for $\alpha(t)$ to have no monomials of degree $\leq -a_1$. Thus by choosing bases of M and M' appropriately, we can take $\alpha(t) = 0$. Show that this implies that $\mathcal{E} = \mathcal{O}(a_1) \oplus \mathcal{O}(a_2)$.

19.5.I. EXERCISE. Finish the proof of Theorem 19.5.6 for arbitrary r . Hint: there are no new ideas beyond the case $r = 2$. □

19.6 Hilbert functions, Hilbert polynomials, and genus

If \mathcal{F} is a coherent sheaf on a projective k -scheme $X \subset \mathbb{P}^n$, define the **Hilbert function of \mathcal{F}** by

$$h_{\mathcal{F}}(m) := h^0(X, \mathcal{F}(m)).$$

Note that the Hilbert function depends on the projective embedding of X (by way of its line bundle), although this dependency is suppressed in the notation. The **Hilbert function of X** is the Hilbert function of the structure sheaf. The ancients were aware that the Hilbert function is “eventually polynomial”, i.e., for large enough m , it agrees with some polynomial. This polynomial contains lots of interesting geometric information, as we will soon see. In modern language, we expect that this “eventual polynomiality” arises because the Euler characteristic should be a polynomial, and that for $m \gg 0$, the higher cohomology vanishes. This is indeed the case, as we now verify.

19.6.1. Theorem. — *If \mathcal{F} is a coherent sheaf on a projective k -scheme $X \hookrightarrow \mathbb{P}_k^n$, $\chi(X, \mathcal{F}(m))$ is a polynomial of degree equal to $\dim \text{Supp } \mathcal{F}$. Hence by Serre vanishing (Theorem 19.1.3 (ii)), for $m \gg 0$, $h^0(X, \mathcal{F}(m))$ is a polynomial $p_{\mathcal{F}}(m)$ of degree $\dim \text{Supp } \mathcal{F}$. In particular, for $m \gg 0$, $h^0(X, \mathcal{O}_X(m))$ is polynomial with degree equal to $\dim X$.*

19.6.2. Definition. The polynomial $p_{\mathcal{F}}(m)$ defined in Theorem 19.6.1 is called the **Hilbert polynomial**. If $X \subset \mathbb{P}^n$ is a projective k -scheme, define $p_X(m) := p_{\mathcal{O}_X}(m)$.

In Theorem 19.6.1, $\mathcal{O}_X(m)$ is the restriction or pullback of $\mathcal{O}_{\mathbb{P}^n_k}(m)$. Both the degree of the 0 polynomial and the dimension of the empty set are defined to be -1 . In particular, the only coherent sheaf with Hilbert polynomial 0 is the zero-sheaf.

This argument uses the notion of associated points of a coherent sheaf on a locally Noetherian scheme, §6.5.2. (The resolution given by the Hilbert Syzygy Theorem, §16.3.2, can give a shorter proof; but we haven't proved the Hilbert Syzygy Theorem.)

Proof. Define $p_{\mathcal{F}}(m) = \chi(X, \mathcal{F}(m))$. We will show that $p_{\mathcal{F}}(m)$ is a polynomial of the desired degree.

We first use Exercise 19.2.H to reduce to the case where k is algebraically closed, and in particular infinite. (This is one of those cases where even if you are concerned with potentially arithmetic questions over some non-algebraically closed field like \mathbb{F}_p , you are forced to consider the “geometric” situation where the base field is algebraically closed.)

The coherent sheaf \mathcal{F} has a finite number of associated points (§6.5.17). We show a useful fact that we will use again.

19.6.A. EXERCISE. Suppose X is a projective k -scheme with k infinite, and \mathcal{F} is a coherent sheaf on X . Show that if \mathcal{L} is a very ample invertible sheaf on X , then there is an effective Cartier divisor D on X with $\mathcal{L} \cong \mathcal{O}(D)$, and where D does not meet the associated points of \mathcal{F} . (Hint: show that given any finite set of points of \mathbb{P}^n_k , there is a hyperplane not containing any of them. This is a variant of the key step in Exercise 12.3.B(c).)

Thus there is a hyperplane $x = 0$ ($x \in \Gamma(X, \mathcal{O}(1))$) missing this finite number of points. (This is where we use the infinitude of k .)

Then the map $\mathcal{F}(-1) \xrightarrow{x} \mathcal{F}$ is injective. (On any affine open subset \mathcal{F} corresponds to a module, and x is not a zerodivisor on that module, as it doesn't vanish at any associated point of that module, see Proposition 6.5.13.) Thus we have a short exact sequence

$$(19.6.2.1) \quad 0 \longrightarrow \mathcal{F}(-1) \longrightarrow \mathcal{F} \longrightarrow \mathcal{G} \longrightarrow 0$$

where \mathcal{G} is a coherent sheaf.

19.6.B. EXERCISE. Show that $\text{Supp } \mathcal{G} = (\text{Supp } \mathcal{F}) \cap V(x)$. (Hint: show that $\mathcal{F}(-1) \rightarrow \mathcal{F}$ is an isomorphism away from $V(x)$, and hence $\mathcal{G} = 0$ on this locus. If $p \in V(x)$, show that the $\mathcal{F}(-1)|_p \rightarrow \mathcal{F}|_p$ is the 0 map, and hence $\mathcal{F}|_p \rightarrow \mathcal{G}|_p$ is an isomorphism.)

Hence $V(x)$ meets all positive-dimensional components of $\text{Supp } \mathcal{F}$ (Exercise 12.3.B(a)), so $\dim \text{Supp } \mathcal{G} = \dim \text{Supp } \mathcal{F} - 1$ by Krull's Principal Ideal Theorem 12.3.3 unless $\mathcal{F} = 0$ (in which case we already know the result, so assume this is not the case).

Twisting (19.6.2.1) by $\mathcal{O}(m)$ yields

$$0 \longrightarrow \mathcal{F}(m-1) \longrightarrow \mathcal{F}(m) \longrightarrow \mathcal{G}(m) \longrightarrow 0$$

Euler characteristics are additive in exact sequences (Exercise 19.4.A), from which $p_{\mathcal{F}}(m) - p_{\mathcal{F}}(m-1) = p_{\mathcal{G}}(m)$. Now $p_{\mathcal{G}}(m)$ is a polynomial of degree $\dim \text{Supp } \mathcal{F} - 1$.

The result is then a consequence from the following elementary fact about polynomials in one variable.

19.6.C. EXERCISE. Suppose f and g are functions on the integers, $f(m+1) - f(m) = g(m)$ for all m , and $g(m)$ is a polynomial of degree $d \geq 0$. Show that f is a polynomial of degree $d+1$. □

Example 1. The Hilbert polynomial of projective space is $p_{\mathbb{P}^n}(m) = \binom{m+n}{n}$, where we interpret this as the polynomial $(m+1) \cdots (m+n)/n!$.

Example 2. Suppose H is a degree d hypersurface in \mathbb{P}^n , with $i: H \hookrightarrow \mathbb{P}^n$ the closed embedding. Then from the closed subscheme exact sequence

$$0 \longrightarrow \mathcal{O}_{\mathbb{P}^n}(-d) \longrightarrow \mathcal{O}_{\mathbb{P}^n} \longrightarrow i_*\mathcal{O}_H \longrightarrow 0,$$

we have

$$(19.6.2.2) \quad p_H(m) = p_{\mathbb{P}^n}(m) - p_{\mathbb{P}^n}(m-d) = \binom{m+n}{n} - \binom{m+n-d}{n}.$$

(Implicit in this argument is the fact that $(i_*\mathcal{O}_H) \otimes \mathcal{O}_{\mathbb{P}^n}(m) \cong i_*(\mathcal{O}_H \otimes i^*\mathcal{O}_{\mathbb{P}^n}(m))$. This follows from the projection formula, Exercise 17.3.H(b). You can also show this directly with transition functions.)

19.6.D. EXERCISE. Show that the twisted cubic (\mathbb{P}^1 embedded in \mathbb{P}^3 by the complete linear series $|\mathcal{O}_{\mathbb{P}^1}(3)|$) has Hilbert polynomial $3m+1$. (The twisted cubic was defined in Exercise 9.3.A.)

19.6.E. EXERCISE. More generally, find the Hilbert polynomial for the d th Veronese embedding of \mathbb{P}^n (i.e., the closed embedding of \mathbb{P}^n in a bigger projective space by way of the line bundle $\mathcal{O}(d)$, §9.3.6).

19.6.F. EXERCISE (TO BE USED SEVERAL TIMES IN CHAPTER 20). Suppose

$$X \xhookrightarrow{i} Y \xhookrightarrow{\quad} \mathbb{P}_k^n$$

is a sequence of closed embeddings.

(a) Show that $p_X(m) \leq p_Y(m)$ for $m \gg 0$. Hint: let $\mathcal{I}_{X/Y}$ be the ideal sheaf of X in Y . Consider the exact sequence

$$0 \longrightarrow \mathcal{I}_{X/Y}(m) \longrightarrow \mathcal{O}_Y(m) \longrightarrow i_*\mathcal{O}_X(m) \longrightarrow 0.$$

(b) If $p_X(m) = p_Y(m)$ for $m \gg 0$, show that $X = Y$. Hint: Show that if the Hilbert polynomial of a coherent sheaf \mathcal{F} is 0, then $\mathcal{F} = 0$. (Handy trick: For $m \gg 0$, $\mathcal{F}(m)$ is generated by global sections.) Apply this to $\mathcal{F} = \mathcal{I}_{X/Y}$.

From the Hilbert polynomial, we can extract many invariants, of which two are particularly important. The first is the *degree*, and the second is the arithmetic

genus (§19.6.6). The **degree of a projective k -scheme of dimension n** is defined to be leading coefficient of the Hilbert polynomial (the coefficient of m^n) times $n!$.

Using the examples above, we see that the degree of \mathbb{P}^n in itself is 1. The degree of the twisted cubic is 3.

19.6.G. EXERCISE. Show that the degree is always an integer. Hint: by induction, show that any polynomial in m of degree k taking on only integer values must have coefficient of m^k an integral multiple of $1/k!$. Hint for this: if $f(x)$ takes on only integral values and is of degree k , then $f(x+1) - f(x)$ takes on only integral values and is of degree $k-1$.

19.6.H. EXERCISE. Show that the degree of a degree d hypersurface (Definition 9.3.2) is d (preventing a notational crisis).

19.6.I. EXERCISE (DEGREE=DEGREE). Suppose a curve C is embedded in projective space via an invertible sheaf of degree d (as defined in §19.4.5). In other words, this line bundle determines a closed embedding. Show that the degree of C under this embedding is d , preventing another notational crisis. Hint: Riemann-Roch, Exercise 19.4.B. (A similar notational crisis was averted in Exercise 19.4.E.)

19.6.J. EXERCISE. Show that the degree of the d th Veronese embedding of \mathbb{P}^n is d^n .

19.6.K. EXERCISE (BÉZOUT'S THEOREM, GENERALIZING EXERCISES 9.3.D AND 17.4.J).

Suppose X is a projective scheme of dimension at least 1, and H is a hypersurface not containing any associated points of X . (For example, if X is reduced and thus has no embedded points, we are just requiring H not to contain any irreducible components of X .) Show that $\deg(H \cap X) = (\deg H)(\deg X)$. (As an example, we have *Bézout's Theorem for plane curves*: if C and D are plane curves of degrees m and n respectively, with no common components, then C and D meet at mn points, counted with appropriate multiplicity.

19.6.3. This is a very handy result! For example: if two projective plane curves of degree m and degree n share no irreducible components, then they intersect in mn points, counted with appropriate multiplicity. (To apply Exercise 19.6.K, you need to know that plane curves have no embedded points. You can either do this using Exercise 6.5.E, or save time by assuming that one of the curves is reduced.) The notion of multiplicity of intersection is just the degree of the intersection as a k -scheme.

19.6.L. EXERCISE. Suppose C is a degree 1 curve in \mathbb{P}_k^3 (or more precisely, a degree 1 pure-one-dimensional closed subscheme of \mathbb{P}_k^3 , with no embedded points). Show that C is a line. Hint: reduce to the case $k = \bar{k}$. Suppose p and q are distinct closed points on C . Use Bézout's Theorem (Exercise 19.6.K) to show that any hyperplane containing p and q must contain C , and thus that $C \subset \overline{pq}$.

19.6.M. FUN EXERCISE. Let k be a field, which we assume to be algebraically closed for convenience (although you are free to remove this hypothesis if you wish). Suppose C is a degree d integral curve in \mathbb{P}_k^N with $N \geq d$. Show that C is

contained in a linear $\mathbb{P}_k^d \subset \mathbb{P}_k^N$. (Exercise 19.6.L is not quite a special case of this problem, but the hint may still be helpful.)

19.6.N. EXERCISE (ANOTHER FORM OF BÉZOUT’S THEOREM). Classically, the degree of a complex projective variety of dimension n was defined as follows. We slice the variety with n generally chosen hyperplanes. Then the intersection will be a finite number of reduced points, by Exercise 13.4.C (a consequence of Bertini’s Theorem 13.4.2). The degree is this number of points. Use Bézout’s Theorem to make sense of this in a way that agrees with our definition of degree. You will need to assume that k is infinite.

Thus the classical definition of the degree, which involved making a choice and then showing that the result is independent of choice, has been replaced by making a cohomological definition involving Euler characteristics. This should remind you of how we got around to “correctly” understanding the degree of a line bundle. It was traditionally defined as the degree of a divisor of any nonzero rational section (Important Exercise 19.4.C), and we found a better definition in terms of Euler characteristics (§19.4.5).

19.6.4. ★★ *Aside: Connection to the topological definition of degree (for those with sufficient background).* Another definition of degree of a dimension d complex projective variety $X \subset \mathbb{P}_\mathbb{C}^n$ is as the number d such that $[X]$ is d times the “positive” generator of $H_{2d}(\mathbb{P}_\mathbb{C}^n, \mathbb{Z})$. You can show this by induction on d as follows. Suppose X is a complex projective variety of dimension d . For a generally chosen hyperplane, $H \cap X$ is a complex projective variety of (complex) dimension $d - 1$. (Do you see why?) Show that $[H \cap X] = c_1(\mathcal{O}(1)) \cup [X]$ in $H_{2(d-1)}(X, \mathbb{Z})$, by suitably generalizing the solution to Exercise 19.4.H. (A further generalization is given in §21.1.7.) For this reason, $c_1(\mathcal{O}(1))$ is often called the “hyperplane class”.

19.6.5. *Revisiting an earlier example.* We revisit the enlightening example of Example 10.3.4 and §18.4.9: let $k = \mathbb{Q}$, and consider the parabola $x = y^2$. We intersect it with the four lines, $x = 1$, $x = 0$, $x = -1$, and $x = 2$, and see that we get 2 each time (counted with the same convention as with the last time we saw this example).

If we intersect it with $y = 2$, we only get one point — but that’s because this isn’t a projective curve, and we really should be doing this intersection on \mathbb{P}_k^2 , and in this case, the conic meets the line in two points, one of which is “at ∞ ”.

19.6.O. EXERCISE. Show that the degree of the d -fold Veronese embedding of \mathbb{P}^n is d^n in a different way from Exercise 19.6.J as follows. Let $v_d: \mathbb{P}^n \rightarrow \mathbb{P}^N$ be the Veronese embedding. To find the degree of the image, we intersect it with n hyperplanes in \mathbb{P}^N (scheme-theoretically), and find the number of intersection points (counted with multiplicity). But the pullback of a hyperplane in \mathbb{P}^N to \mathbb{P}^n is a degree d hypersurface. Perform this intersection in \mathbb{P}^n , and use Bézout’s theorem (Exercise 19.6.K).

19.6.P. EXERCISE (DEGREE IS ADDITIVE FOR UNIONS). Suppose X and Y are two d -dimensional closed subschemes of \mathbb{P}_k^n , with no d -dimensional irreducible components in common. Show that $\deg X \cup Y = \deg X + \deg Y$.

19.6.6. Arithmetic genus, again.

There is another central piece of information residing in the Hilbert polynomial. Notice that $1 - p_X(0) = 1 - \chi(X, \mathcal{O}_X)$ is the arithmetic genus (§19.4.3), an *intrinsic* invariant of the scheme X , independent of the projective embedding.

Imagine how amazing this must have seemed to the ancients: they defined the Hilbert function by counting how many “functions of various degrees” there are; then they noticed that when the degree gets large, it agrees with a polynomial; and then when they plugged 0 into the polynomial — extrapolating backwards, to where the Hilbert function and Hilbert polynomials didn’t agree — they found a magic invariant! Furthermore, in the case when X is a complex curve, this invariant was basically the topological genus!

We can now see a large family of curves over an algebraically closed field that is provably not \mathbb{P}^1 ! Note that the Hilbert polynomial of \mathbb{P}^1 is $(m+1)/1 = m+1$, so $\chi(\mathcal{O}_{\mathbb{P}^1}) = 1$. Suppose C is a degree d curve in \mathbb{P}^2 . Then the Hilbert polynomial of C is

$$p_{\mathbb{P}^2}(m) - p_{\mathbb{P}^2}(m-d) = (m+1)(m+2)/2 - (m-d+1)(m-d+2)/2.$$

Plugging in $m=0$ gives us $-(d^2-3d)/2$. Thus when $d > 2$, we have a curve that cannot be isomorphic to \mathbb{P}^1 ! (And it is not hard to show that there exists a *regular* degree d curve, Exercise 13.3.E.)

Now from $0 \rightarrow \mathcal{O}_{\mathbb{P}^2}(-d) \rightarrow \mathcal{O}_{\mathbb{P}^2} \rightarrow \mathcal{O}_C \rightarrow 0$, using $h^1(\mathbb{P}^2, \mathcal{O}_{\mathbb{P}^2}(-d)) = 0$, we have that $h^0(C, \mathcal{O}_C) = 1$. As $h^0 - h^1 = \chi$, we have

$$(19.6.6.1) \quad h^1(C, \mathcal{O}_C) = (d-1)(d-2)/2.$$

We now revisit an interesting question we first saw in §7.5.10. If k is an algebraically closed field, is every finitely generated transcendence degree 1 extension of k isomorphic to $k(x)$? In that section, we found ad hoc (but admittedly beautiful) examples showing that the answer is “no”. But we now have a better answer. The question initially looks like an algebraic question, but we now recognize it as a fundamentally geometric one. There is an integer-valued cohomological invariant of such field extensions that has good geometric meaning: the genus.

Equation (19.6.6.1) yields examples of curves of genus 0, 1, 3, 6, 10, ... (corresponding to degree 1 or 2, 3, 4, 5, ...). This begs some questions, such as: are there curves of other genera? (We will see soon, in §20.5.5, that the answer is yes.) Are there other genus 0 curves? (Not if k is algebraically closed, but sometimes yes otherwise — consider $x^2 + y^2 + z^2 = 0$ in $\mathbb{P}_{\mathbb{R}}^2$, which has no \mathbb{R} -points and hence is not isomorphic to $\mathbb{P}_{\mathbb{R}}^1$ — we will discuss this more in §20.3.) Do we have all the curves of genus 3? (Almost all, but not quite. We will see more in §20.7.) Do we have all the curves of genus 6? (We are missing “most of them”, as will be suggested by §20.8.3.)

Caution: The Euler characteristic of the structure sheaf is an incomplete invariant. It doesn’t always distinguish between isomorphism classes of irreducible smooth projective varieties. For example, \mathbb{P}^2 and $\mathbb{P}^1 \times \mathbb{P}^1$ both have structure sheaf Euler characteristic 1 (see Theorem 19.1.2 and Exercise ??), but are not isomorphic — $\text{Pic } \mathbb{P}^2 \cong \mathbb{Z}$ (§15.2.7) while $\text{Pic}(\mathbb{P}^1 \times \mathbb{P}^1) \cong \mathbb{Z} \oplus \mathbb{Z}$ (Exercise 15.2.O).

19.6.7. Complete intersections.

19.6.Q. EXERCISE. Suppose X is a dimension m complete intersection in \mathbb{P}_k^N . Show that $h^i(X, \mathcal{O}_X(n)) = 0$ for $0 < i < m$.

19.6.R. EXERCISE. Find the genus of the complete intersection of 2 quadric surfaces in \mathbb{P}_k^3 .

19.6.S. EXERCISE. More generally, find the genus of the complete intersection of a degree m surface with a degree n surface in \mathbb{P}_k^3 . (If $m = 2$ and $n = 3$, you should get genus 4. We will see in §20.8 that in some sense most genus 4 curves arise in this way. Note that Bertini's Theorem 13.4.2 ensures that there *are* regular curves of this form.)

19.6.T. EXERCISE. Show that the rational normal curve of degree d in \mathbb{P}^d is *not* a complete intersection if $d > 2$. (Hint: If it *were* the complete intersection of $d - 1$ hypersurfaces, what would the degree of the hypersurfaces be? Why could none of the degrees be 1?)

19.6.U. EXERCISE (POSITIVE-DIMENSIONAL COMPLETE INTERSECTIONS ARE CONNECTED). Show that complete intersections of *positive* dimension are connected. Hint: show that $h^0(X, \mathcal{O}_X) = 1$. (This argument will even show that they are geometrically connected, defined in §10.5.1, as h^0 is preserved by field extension, see Exercise 19.2.H.)

Notice that this statement does not require X to be reduced, so for example this implies that if X is the scheme cut out by $x^2 = 0$ in \mathbb{P}_k^2 , then $h^0(X, \mathcal{O}_X) = 1$ — there are no nonzero nilpotent functions, even though there are local functions that are nonzero and nilpotent! This example came up earlier in §5.2.3 and §11.5.7.

19.6.V. EXERCISE. Show that the union of two planes in \mathbb{P}^4 meeting at a point is not a complete intersection. Hint: it is connected, but you can slice with another hyperplane and get something not connected (see Exercise 19.6.U).

19.7 ★ Serre's cohomological characterization of ampleness

Theorem 17.6.2 gave a number of characterizations of ampleness, in terms of projective geometry, global generation, and the Zariski topology. Here is another characterization, this time cohomological, under Noetherian hypotheses. Because (somewhat surprisingly) we won't use this result, this section is starred.

19.7.1. Theorem (Serre's cohomological criterion for ampleness). — Suppose A is a Noetherian ring, X is a proper A -scheme, and \mathcal{L} is an invertible sheaf on X . Then the following are equivalent.

- (a-c) The invertible sheaf \mathcal{L} is ample on X (over A).
- (e) For all coherent sheaves \mathcal{F} on X , there is an n_0 such that for $n \geq n_0$, $H^i(X, \mathcal{F} \otimes \mathcal{L}^{\otimes n}) = 0$ for all $i > 0$.

The label (a-c) is intended to reflect the statement of Theorem 17.6.2. We avoid the label (d) because it appeared in Theorem 17.6.6. (Aside: the “properness” assumption cannot be removed, as can be shown by the example of Exercise 19.8.F.) Before getting to the proof, we motivate this result by giving some applications. (As a warm-up, you can give a second solution to Exercise 17.6.G in

the Noetherian case, using the affineness of π to show that $H^i(X, \mathcal{F} \otimes (\pi^* \mathcal{L})^{\otimes m}) = H^i(Y, (\pi_* \mathcal{F}) \otimes \mathcal{L}^{\otimes m})$.

19.7.A. EXERCISE. Suppose X is a proper A -scheme, and \mathcal{L} is an invertible sheaf on X . Show that \mathcal{L} is ample on X if and only if $\mathcal{L}|_{X^{\text{red}}}$ is ample on X^{red} . Hint: for the “only if” direction, use Exercise 17.6.G. For the “if” direction, let \mathcal{I} be the ideal sheaf cutting out the closed subscheme X^{red} in X . Filter \mathcal{F} by powers of \mathcal{I} :

$$0 = \mathcal{I}^r \mathcal{F} \subset \mathcal{I}^{r-1} \mathcal{F} \subset \cdots \subset \mathcal{I} \mathcal{F} \subset \mathcal{F}.$$

(Essentially the same filtration appeared in Exercise 19.4.S, for similar reasons.) Show that each quotient $\mathcal{I}^n \mathcal{F} / \mathcal{I}^{n-1} \mathcal{F}$, twisted by a high enough power of \mathcal{L} , has no higher cohomology. Use descending induction on n to show each part $\mathcal{I}^n \mathcal{F}$ of the filtration (and hence in particular \mathcal{F}) has this property as well.

19.7.B. EXERCISE. Suppose X is a proper A -scheme, and \mathcal{L} is an invertible sheaf on X . Show that \mathcal{L} is ample on X if and only if \mathcal{L} is ample on each component. Hint: follow the outline of the solution to the previous exercise, taking instead \mathcal{I} as the ideal sheaf of one component. Perhaps first reduce to the case where $X = X^{\text{red}}$.

19.7.C. EXERCISE. (In Exercise 20.2.E, we will show that on a smooth projective integral curve, an invertible sheaf is ample if and only if it has positive degree. Use that fact in this exercise. There will be no logical circularity.) Show that a line bundle on a projective curve is ample if and only if it has positive degree on each component.

19.7.2. Remark: Proper curves are projective. Serre’s criterion for ampleness is the key ingredient for showing that every proper curve over a field is projective. The steps are as follows. (i) Recall that every regular integral proper curve is projective, Exercise 18.4.A. (ii) The hardest step is showing that every reduced integral proper curve C is projective. This is done by choosing a regular point on each irreducible component of C , and letting \mathcal{L} be the corresponding invertible sheaf. Because of Exercise 19.7.C, we hope that \mathcal{L} will be ample. Show that a line bundle on C is ample if its pullback to the normalization of C is ample (a partial converse to Exercise 17.6.G, see for example [Ha1, Ex. III.5.7(d)]). Thus our \mathcal{L} is ample. (iii) Show that every reduced proper curve is projective using Exercise 19.7.B. (iv) Show that every proper curve C is projective, using Exercise 19.7.A, after first finding an invertible sheaf on C that will be shown to be ample.

19.7.3. Very ample versus ample. The previous exercises don’t work with “ample” replaced by “very ample”, which shows again how the notion of ampleness is better-behaved than very ampleness.

19.7.4. Proof of Theorem 19.7.1. For the fact that (a-c) implies (e), use the fact that $\mathcal{L}^{\otimes N}$ is very ample for some N (Theorem 17.6.2(a)), and apply Serre vanishing (Theorem 19.1.3(ii)) to \mathcal{F} , $\mathcal{F} \otimes \mathcal{L}$, \dots , and $\mathcal{F} \otimes \mathcal{L}^{\otimes (N-1)}$.

So we now assume (e), and show that \mathcal{L} is ample by criterion (b) of Theorem 17.6.2: we will show that for any coherent sheaf \mathcal{F} on X , $\mathcal{F} \otimes \mathcal{L}^{\otimes n}$ is globally generated for $n \gg 0$.

We begin with a special case: we will show that $\mathcal{L}^{\otimes n}$ is globally generated (i.e., base-point-free) for $n \gg 0$. To do this, it suffices to show that every closed point p has an open neighborhood U so that there exists some N_p so that $n \geq N_p$, $\mathcal{L}^{\otimes n}$ is globally generated for all points of U_p . (Reason: by quasicompactness, every closed subset of X contains a closed point, by Exercise 5.1.E. So as p varies over the closed points of X , these U_p cover X . By quasicompactness again, we can cover X by a finite number of these U_p . Let N be the maximum of the corresponding N_p . Then for $n \geq N$, $\mathcal{L}^{\otimes n}$ is globally generated in each of these U_p , and hence on all of X .)

Let p be a closed point of X . For all n , $m_p \otimes \mathcal{L}^{\otimes n}$ is coherent (by our Noetherian hypotheses). By (e), there exists some n_0 so that for $n \geq n_0$, $H^1(X, m_p \otimes \mathcal{L}^{\otimes n}) = 0$. By the long exact sequence arising from the closed subscheme exact sequence

$$0 \rightarrow m_p \otimes \mathcal{L}^{\otimes n} \rightarrow \mathcal{L}^{\otimes n} \rightarrow \mathcal{L}^{\otimes n}|_p \rightarrow 0,$$

we have that $\mathcal{L}^{\otimes n}$ is globally generated at p for $n \geq n_0$. By Exercise 16.3.D(b), there is an open neighborhood V_0 of p such that $\mathcal{L}^{\otimes n_0}$ is globally generated at all points of V_0 . Thus $\mathcal{L}^{\otimes kn_0}$ is globally generated at all points of V_0 for all positive integers k (using Easy Exercise 16.3.C). For each $i \in \{1, \dots, n_0 - 1\}$, there is an open neighborhood V_i of p such that $\mathcal{L}^{\otimes(n_0+i)}$ is globally generated at all points of V_i (again by Exercise 16.3.D(b)). We may take each V_i to be contained in V_0 . By Easy Exercise 16.3.C, $\mathcal{L}^{\otimes(kn_0+n_0+i)}$ is globally generated at every point of V_i (as this is the case for $\mathcal{L}^{\otimes kn_0}$ and $\mathcal{L}^{\otimes(n_0+i)}$). Thus in the open neighborhood $U_p := \bigcap_{i=0}^{n_0-1} V_i$, $\mathcal{L}^{\otimes n}$ is globally generated for $n \geq N_p := 2n_0$.

We have now shown that there exists some N such that for $n \geq N$, $\mathcal{L}^{\otimes n}$ is globally generated. Now suppose \mathcal{F} is a coherent sheaf. To conclude the proof, we will show that $\mathcal{F} \otimes \mathcal{L}^{\otimes n}$ is globally generated for $n \gg 0$. This argument has a similar flavor to what we have done so far, so we give it as an exercise.

19.7.D. EXERCISE. Suppose p is a closed point of X .

(a) Show that for $n \gg 0$, $\mathcal{F} \otimes \mathcal{L}^{\otimes n}$ is globally generated at p .

(b) Show that there exists an open neighborhood U_p of p such that for $n \gg 0$, $\mathcal{F} \otimes \mathcal{L}^{\otimes n}$ is globally generated at every point of U_p . Caution: while it is true that by Exercise 16.3.D(b), for each $n \gg 0$, there is some open neighborhood V_n of p such that $\mathcal{F} \otimes \mathcal{L}^{\otimes n}$ is globally generated there, it need not be true that

$$(19.7.4.1) \quad \bigcap_{n \gg 0} V_n$$

is an open set. You may need to use the fact that $\mathcal{L}^{\otimes n}$ is globally generated for $n \geq N$ to replace (19.7.4.1) by a finite intersection.

19.7.E. EXERCISE. Conclude the proof of Theorem 19.7.1 by showing that $\mathcal{F} \otimes \mathcal{L}^{\otimes n}$ is globally generated for $n \gg 0$. \square

19.7.5. Aside. Serre's characterization of affineness is similar in flavor to Theorem 19.7.1.

19.7.6. Theorem (Serre's cohomological characterization of affineness). — Suppose X is a quasicompact quasiseparated scheme. Then the following are equivalent.

(a) The scheme X is affine.

(b) For any quasicoherent sheaf \mathcal{F} on X , $H^i(X, \mathcal{F}) = 0$ for all $i > 0$.

(c) For any coherent sheaf of ideals \mathcal{I} on X , $H^1(X, \mathcal{I}) = 0$.

Clearly (a) implies (b) implies (c) (the former from “affine cover cohomology vanishing”, Property (vi) of §19.1) without any Noetherian assumptions, so the real substance is in the implication from (c) to (a). Because we won’t use it, we omit the proof (which in any case is given in [GW, Thm. 12.35]).

Serre proved an analogous result in complex analytic geometry: Stein spaces are also characterized by the vanishing of cohomology of coherent sheaves.

19.8 Higher pushforward (or direct image) sheaves

Cohomology groups were defined for $X \rightarrow \operatorname{Spec} A$ where the structure morphism is quasicompact and separated; for any quasicohherent \mathcal{F} on X , we defined $H^i(X, \mathcal{F})$. We will now define a “relative” version of this notion, for quasicompact and separated morphisms $\pi: X \rightarrow Y$: for any quasicohherent \mathcal{F} on X , we will define $R^i\pi_*\mathcal{F}$, a quasicohherent sheaf on Y . (Now would be a good time to do Exercise 1.6.H, the FHHF Theorem, if you haven’t done it before.)

We have many motivations for doing this. In no particular order:

- (1) It “globalizes” what we did before with cohomology.
- (2) If $0 \rightarrow \mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{H} \rightarrow 0$ is a short exact sequence of quasicohherent sheaves on X , then we know that $0 \rightarrow \pi_*\mathcal{F} \rightarrow \pi_*\mathcal{G} \rightarrow \pi_*\mathcal{H} \rightarrow 0$ is exact, and higher pushforwards will extend this to a long exact sequence.
- (3) We will later see that this will show how cohomology groups vary in families, especially in “nice” situations. Intuitively, if we have a nice family of varieties, and a family of sheaves on them, we could hope that the cohomology varies nicely in families, and in fact in “nice” situations, this is true. (As always, “nice” usually means “flat”, whatever that means. We will see that Euler characteristics are locally constant in proper flat families in §25.7, and the Cohomology and Base Change Theorem 28.1.6 will show that in particularly good situations, dimensions of cohomology groups are constant.)

All of the important properties of cohomology described in §19.1 will carry over to this more general situation. Best of all, there will be no extra work required.

In the notation $R^i\pi_*\mathcal{F}$ for higher pushforward sheaves, the “R” stands for “right derived functor”, and corresponds to the fact that we get a long exact sequence in cohomology extending to the right (from the 0th terms). In Chapter 24, we will see that in good circumstances, if we have a left-exact functor, there is a long exact sequence going off to the right, in terms of right derived functors. Similarly, if we have a right-exact functor (e.g., if M is an A -module, then $\otimes_A M$ is a right-exact functor from the category of A -modules to itself), there may be a long exact sequence going off to the left, in terms of left derived functors.

Suppose $\pi: X \rightarrow Y$, and \mathcal{F} is a quasicohherent sheaf on X . For each $\operatorname{Spec} A \subset Y$, we have A -modules $H^i(\pi^{-1}(\operatorname{Spec} A), \mathcal{F})$. We now show that these patch together to form a quasicohherent sheaf, in the sense of §6.2.3. We need check only one fact: that this behaves well with respect to taking distinguished open sets. In other

words, we must check that for each $f \in A$, the natural map

$$H^i(\pi^{-1}(\operatorname{Spec} A), \mathcal{F}) \longrightarrow H^i(\pi^{-1}(\operatorname{Spec} A_f), \mathcal{F})$$

(induced by the map of spaces in the opposite direction — H^i is contravariant in the space) is precisely the localization $\otimes_A A_f$. But this can be verified easily: let $\{U_i\}$ be an affine open cover of $\pi^{-1}(\operatorname{Spec} A)$. We can compute $H^i(\pi^{-1}(\operatorname{Spec} A), \mathcal{F})$ using the Čech complex (19.2.1.1). But this induces a cover of $\pi^{-1}(\operatorname{Spec} A_f)$ in a natural way: if $U_i = \operatorname{Spec} A_i$ is an affine open subset of $\pi^{-1}(\operatorname{Spec} A)$, we define $U'_i := \operatorname{Spec}(A_i)_{\pi^\# f}$. The resulting Čech complex for $\operatorname{Spec} A_f$ is the localization of the Čech complex for $\operatorname{Spec} A$. As taking cohomology of a complex commutes with localization (as discussed in the FHHF Theorem, Exercise 1.6.H), we have defined a quasicoherent sheaf on Y by the characterization of quasicoherent sheaves in §6.2.3.

Define the **i th higher pushforward sheaf** or the **i th higher direct image sheaf** $R^i\pi_*\mathcal{F}$ to be this quasicoherent sheaf.

Caution: Although when $U \subset Y$ is an affine open subset, $H^0(U, R^\bullet\pi_*\mathcal{F}) = H^\bullet(\pi^{-1}(U), \mathcal{F})$, this is not true for open subsets $U \subset Y$ in general. You might later try to construct an example of this behavior.

19.8.1. Theorem. — Suppose $\pi: X \rightarrow Y$ is a quasicompact separated morphism of schemes. Then:

- (a) $R^i\pi_*$ is a covariant functor $QCoh_X \rightarrow QCoh_Y$.
- (b) We can identify $R^0\pi_*$ with π_* .
- (c) (the **long exact sequence of higher pushforward sheaves**) A short exact sequence $0 \rightarrow \mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{H} \rightarrow 0$ of sheaves on X induces a long exact sequence

$$(19.8.1.1) \quad 0 \longrightarrow R^0\pi_*\mathcal{F} \longrightarrow R^0\pi_*\mathcal{G} \longrightarrow R^0\pi_*\mathcal{H} \longrightarrow$$

$$R^1\pi_*\mathcal{F} \longrightarrow R^1\pi_*\mathcal{G} \longrightarrow R^1\pi_*\mathcal{H} \longrightarrow \dots$$

of sheaves on Y .

- (d) (projective pushforwards of coherent are coherent: Grothendieck's Coherence Theorem for projective morphisms) If π is a projective morphism and \mathcal{O}_Y is coherent on Y (this hypothesis is automatic for Y locally Noetherian), and \mathcal{F} is a coherent sheaf on X , then for all i , $R^i\pi_*\mathcal{F}$ is a coherent sheaf on Y .

19.8.2. Unimportant Remark. If X and Y are Noetherian, the hypotheses “quasicompact and separated” on π can be removed; see Unimportant Remark 24.5.8 for the “separated” hypothesis.

19.8.3. Proof of Theorem 19.8.1. We first show covariance: if $\mathcal{F} \rightarrow \mathcal{G}$ is a morphism of quasicoherent sheaves on X , we define a map $R^i\pi_*\mathcal{F} \rightarrow R^i\pi_*\mathcal{G}$. (It will be clear we will have shown that $R^i\pi_*$ is a functor.) It suffices to define this map on the “distinguished affine base” of Y (Definition 6.2.1). Thus it suffices to show the following: if X' is a quasicompact separated A -scheme, and $\mathcal{F} \rightarrow \mathcal{G}$ is a morphism of quasicoherent sheaves on X , then the map $H^i(X', \mathcal{F}) \rightarrow H^i(X', \mathcal{G})$ constructed

in §19.2 (property (i) of §19.1) “commutes with localization at $f \in A$ ”. But this was shown in Exercise 19.2.D.

In a similar way, we construct the connecting homomorphism $R^i\pi_*\mathcal{H} \rightarrow R^{i+1}\pi_*\mathcal{F}$ in the long exact sequence (19.8.1.1), by showing that the construction in the case where $Y = \text{Spec } A$ “commutes with localization at $f \in A$ ”. Again, this was shown in Exercise 19.2.D.

It suffices to check all other parts of this statement on affine open subsets of Y , so they all follow from the analogous statements in Čech cohomology (§19.1). \square

The following result is handy, and essentially immediate from our definition.

19.8.A. EASY EXERCISE. Show that if π is affine, then for $i > 0$, $R^i\pi_*\mathcal{F} = 0$.

19.8.4. How higher pushforwards behave with respect to base change.

19.8.B. EXERCISE (HIGHER PUSHFORWARDS AND BASE CHANGE).

(a) (the push-pull map for cohomology) Suppose $\psi: Y' \rightarrow Y$ is any morphism, and $\pi: X \rightarrow Y$ is quasicompact and separated. Suppose \mathcal{F} is a quasicoherent sheaf on X . Let

$$(19.8.4.1) \quad \begin{array}{ccc} X' & \xrightarrow{\psi'} & X \\ \pi' \downarrow & & \downarrow \pi \\ Y' & \xrightarrow{\psi} & Y \end{array}$$

be a Cartesian diagram. Describe a natural morphism

$$\psi^*(R^i\pi_*\mathcal{F}) \rightarrow R^i\pi'_*(\psi')^*\mathcal{F}$$

of quasicoherent sheaves on Y' . (Hint: the FHHF Theorem, Exercise 1.6.H.) Show that the $i = 0$ case is the push-pull map of Exercise 17.3.G.

(b) (cohomology commutes with affine flat base change) If $\psi: Y' \rightarrow Y$ is an affine morphism, and for a cover $\text{Spec } A_i$ of Y , where $\psi^{-1}(\text{Spec } A_i) = \text{Spec } B_i$, B_i is a flat A_i -algebra (§1.6.11: $\otimes_{A_i} B_i$ is exact), and the diagram in (a) is a Cartesian diagram, show that the natural morphism of (a) is an isomorphism. (Exercise 19.2.H was a special case of this exercise. You can likely generalize this to non-affine morphisms — and thus show that cohomology commutes with flat base change (Theorem 25.2.9) — but we wait until Chapter 25 to discuss flatness at length.)

19.8.C. EXERCISE (CF. EXERCISE 17.3.G). Prove Exercise 19.8.B(a) *without* the hypothesis that (19.8.4.1) is a Cartesian diagram, but adding the requirement that π' is quasicompact and separated (just so our definition of $R^i\pi'_*$ applies). In the course of the proof, you will see a map arising in the Leray spectral sequence (Theorem 24.4.5). (Hint: use Exercise 19.8.B(a).)


A useful special case of Exercise 19.8.B(a) is the following.

19.8.D. EXERCISE. If $q \in Y$, describe a natural morphism $(R^i\pi_*\mathcal{F}) \otimes \kappa(q) \rightarrow H^i(\pi^{-1}(q), \mathcal{F}|_{\pi^{-1}(q)})$. (Be sure you understand the meaning of the left side!) Hint: the FHHF Theorem, Exercise 1.6.H.

Thus the fiber of the pushforward may not be the cohomology of the fiber, but at least it always maps to it. We will later see that in good situations this map is an

isomorphism, and thus the higher pushforward sheaf indeed “patches together” the cohomology on fibers (the Cohomology and Base Change Theorem 28.1.6).

19.8.E. EXERCISE (PROJECTION FORMULA, GENERALIZING EXERCISE 17.3.H). Suppose $\pi: X \rightarrow Y$ is quasicompact and separated, and \mathcal{F} and \mathcal{G} are quasicoherent sheaves on X and Y respectively.



(a) Describe a natural morphism

$$(R^i \pi_* \mathcal{F}) \otimes \mathcal{G} \longrightarrow R^i \pi_* (\mathcal{F} \otimes \pi^* \mathcal{G}).$$

(Hint: the FHHF Theorem, Exercise 1.6.H.)

(b) If \mathcal{G} is locally free, show that this natural morphism is an isomorphism.

The following fact uses the same trick as Theorem 19.1.7 and Exercise 19.1.B.

19.8.5. Theorem (relative dimensional cohomology vanishing). — *If $\pi: X \rightarrow Y$ is a projective morphism and Y is locally Noetherian (or more generally \mathcal{O}_Y is coherent over itself), then the higher pushforwards vanish in degree higher than the maximum dimension of the fibers.*

This is false without the projective hypothesis (see Exercise 19.8.F below). In particular, you might hope that just as dimensional cohomology vanishing generalized from projective varieties to quasiprojective varieties or more general settings (§19.2.7) that relative dimensional cohomology vanishing would generalize from projective morphisms to quasiprojective morphisms, but this is not the case.

19.8.F. EXERCISE. Consider the open embedding $\pi: \mathbb{A}^n - \{0\} \rightarrow \mathbb{A}^n$. By direct calculation, show that $R^{n-1} \pi_* \mathcal{O}_{\mathbb{A}^n - \{0\}} \neq 0$. (This calculation will remind you of the proof of Theorem 19.1.2(b); see also Remark 19.3.1.)

Proof of Theorem 19.8.5. Let m be the maximum dimension of all the fibers.

The question is local on Y , so we will show that the result holds near a point p of Y . We may assume that Y is affine, and hence that $X \hookrightarrow \mathbb{P}_Y^n$.

Let k be the residue field at p . Then $\pi^{-1}(p)$ is a projective k -scheme of dimension at most m . By Exercise 12.3.B we can find affine open sets $D(f_1), \dots, D(f_{m+1})$ that cover $\pi^{-1}(p)$. In other words, the intersection of the $V(f_i)$ does not intersect $\pi^{-1}(p)$.

If $Y = \text{Spec } A$ and $p = [\mathfrak{p}]$ (so $k = A_{\mathfrak{p}}/\mathfrak{p}A_{\mathfrak{p}}$), then arbitrarily lift each f_i from an element of $k[x_0, \dots, x_n]$ to an element f'_i of $A_{\mathfrak{p}}[x_0, \dots, x_n]$. Let F be the product of the denominators of the f'_i ; note that $F \notin \mathfrak{p}$, i.e., $\mathfrak{p} = [\mathfrak{p}] \in D(F)$. Then $f'_i \in A_F[x_0, \dots, x_n]$. The intersection of their zero loci $\cap_i V(f'_i) \subset \mathbb{P}_{A_F}^n$ is a closed subscheme of $\mathbb{P}_{A_F}^n$. Intersect it with X to get another closed subscheme of $\mathbb{P}_{A_F}^n$. Take its image under π ; as projective morphisms are closed, we get a closed subset

of $D(F) = \text{Spec } A_F$. But this closed subset does not include p ; hence we can find an affine open neighborhood $\text{Spec } B$ of p in Y missing the image. But if f'_i are the restrictions of f'_i to $B[x_0, \dots, x_n]$, then $D(f'_i)$ cover $\pi^{-1}(\text{Spec } B)$; in other words, $\pi^{-1}(\text{Spec } B)$ is covered by $m + 1$ affine open sets, so by affine cover cohomology vanishing (Property (vi) of §19.1), its cohomology vanishes in degree at least $m + 1$. But the higher direct image sheaf is computed using these cohomology groups, hence the higher pushforward sheaf $R^i\pi_*\mathcal{F}$ vanishes on $\text{Spec } B$ too. \square

19.8.G. EXERCISE (RELATIVE SERRE VANISHING, CF. THEOREM 19.1.3(II)). Suppose $\pi: X \rightarrow Y$ is a proper morphism of Noetherian schemes, and \mathcal{L} is a π -ample invertible sheaf on X . Show that for any coherent sheaf \mathcal{F} on X , for $m \gg 0$, $R^i\pi_*(\mathcal{F} \otimes \mathcal{L}^{\otimes m}) = 0$ for all $i > 0$.

19.9 ★ From projective to proper hypotheses: Chow's Lemma and Grothendieck's Coherence Theorem

The main proofs in this section are double-starred because the results aren't absolutely necessary in the rest of our discussions, and may not be worth reading right now. But just knowing the statement of Grothendieck's Coherence Theorem 19.9.1, (generalizing Theorem 19.8.1(d)) will allow you to immediately translate many of our arguments about projective schemes and morphisms to proper schemes and morphisms, and Chow's Lemma is a multi-purpose tool to extend results from the projective situation to the proper situation in general.

So you have four choices: (i) read the proof of Grothendieck's Coherence Theorem; (ii) skim the proof and convince yourself that you *could* understand it if you wanted to; (iii) skip the proof and assume it is true; or (iv) just retreat to projective hypotheses whenever Grothendieck's Coherence Theorem is invoked. Life is full of compromises, and I won't judge you.

19.9.1. Grothendieck's Coherence Theorem. — Suppose $\pi: X \rightarrow Y$ is a proper morphism of locally Noetherian schemes. Then for any coherent sheaf \mathcal{F} on X , $R^i\pi_*\mathcal{F}$ is coherent on Y .

The case $i = 0$ has already been mentioned a number of times.

19.9.A. EXERCISE. Recall that finite morphisms are affine (by definition) and proper. Use Theorem 19.9.1 to show that if $\pi: X \rightarrow Y$ is proper and affine and Y is Noetherian, then π is finite. (Hint: mimic the proof of the weaker result where proper is replaced by projective, Corollary 19.1.6.)

The proof of Theorem 19.9.1 requires two sophisticated facts. The first is the Leray spectral sequence (Theorem 24.4.5, which applies in this situation because of Exercise 24.5.F). Suppose $\pi: X \rightarrow Y$ and $\rho: Y \rightarrow Z$ are quasicompact separated morphisms. Then for any quasicohherent sheaf \mathcal{F} on X , there is a spectral sequence with E_2 term given by $R^p\rho_*(R^q\pi_*\mathcal{F})$ converging to $R^{p+q}(\rho \circ \pi)_*\mathcal{F}$. Because this would be a reasonable (but hard) exercise in the case we need it (where Z is affine), we will feel comfortable using it. But because we will later prove it, we won't prove it now.

We will also need Chow's Lemma.

19.9.2. Chow's Lemma (proper = projective - birational projective). — Suppose $\pi: X \rightarrow \text{Spec } A$ is a proper morphism, and A is Noetherian. Then there exists $\mu: X' \rightarrow X$ which is surjective and projective, such that $\pi \circ \mu$ is also projective, and such that μ is an isomorphism on a dense open subset of X :

$$\begin{array}{ccc} X' & \xrightarrow{\mu} & X \\ & \searrow \pi \circ \mu & \swarrow \pi \\ & \text{Spec } A. & \end{array}$$

In particular, if X is a proper k -variety, it admits a projective birational morphism from a projective k -variety.

Many generalizations of results from projective to proper situations go through Chow's Lemma. We will prove this version, and state other versions of Chow's Lemma, in §19.9.4. Assuming these two facts, we now prove Grothendieck's Coherence Theorem 19.9.1 in a series of exercises.

19.9.3. ★★ Proof of Grothendieck's Coherence Theorem 19.9.1, assuming Chow's Lemma 19.9.2.

The question is local on Y , so we may assume Y is affine, say $Y = \text{Spec } A$. We work by Noetherian induction on $\text{Supp } \mathcal{F}$, with the base case when $\text{Supp } \mathcal{F} = \emptyset$, i.e., $\mathcal{F} = 0$, which is obvious. So fix \mathcal{F} , and assume the result is known for all coherent sheaves with strictly smaller support.

19.9.B. EXERCISE. Show that we may assume that $\text{Supp } \mathcal{F} = X$. (Hint: Replace X by the scheme-theoretic support of \mathcal{F} , defined in Exercise 9.1.N.)

We now invoke Chow's Lemma to construct a projective morphism $\mu: X' \rightarrow X$ that is an isomorphism on a dense open subset U of X (so $X \setminus U$ is a closed subset of X , strictly smaller than X), and such that $\pi \circ \mu: X' \rightarrow \text{Spec } A$ is projective.

Then $\mathcal{G} = \mu^* \mathcal{F}$ is a coherent sheaf on X' , $\mu_* \mathcal{G}$ is a coherent sheaf on X (by the projective case, Theorem 19.8.1(d)) and the adjunction map $\mathcal{F} \rightarrow \mu_* \mathcal{G} = \mu_* \mu^* \mathcal{F}$ is an isomorphism on U . The kernel \mathcal{E} and cokernel \mathcal{H} are coherent sheaves on X that have strictly smaller support:

$$0 \longrightarrow \mathcal{E} \longrightarrow \mathcal{F} \longrightarrow \mu_* \mathcal{G} \longrightarrow \mathcal{H} \longrightarrow 0.$$

19.9.C. EXERCISE. By the inductive hypothesis, the higher pushforwards of \mathcal{E} and \mathcal{H} are coherent. Show that if all the higher pushforwards of $\mu_* \mathcal{G}$ are coherent, then the higher pushforwards of \mathcal{F} are coherent.

So we are reduced to showing that the higher pushforwards of $\mu_* \mathcal{G}$ are coherent for any coherent \mathcal{G} on X' .

The Leray spectral sequence (Theorem 24.4.5) for $X' \xrightarrow{\mu} X \xrightarrow{\pi} \text{Spec } A$ has E_2 page given by $R^p \pi_* (R^q \mu_* \mathcal{G})$ converging to $R^{p+q} (\pi \circ \mu)_* \mathcal{G}$. Now $R^q \mu_* \mathcal{G}$ is coherent by Theorem 19.8.1(d). Furthermore, as μ is an isomorphism on a dense open subset U of X , $R^q \mu_* \mathcal{G}$ is zero on U , and is thus supported on the complement of U , whose support is strictly smaller than X . Hence by our inductive hypothesis,

$R^p\pi_*(R^q\mu_*\mathcal{G})$ is coherent for all p , and all $q \geq 1$. The only possibly noncoherent sheaves on the E_2 page are in the row $q = 0$ — precisely the sheaves we are interested in. Also, by Theorem 19.8.1(d) applied to $\pi \circ \mu$, $R^{p+q}(\pi \circ \mu)_*\mathcal{G}$ is coherent.

19.9.D. EXERCISE. Show that $E_n^{p,q}$ is always coherent for any $n \geq 2$, $q > 0$. Show that $E_n^{p,0}$ is coherent for a given $n \geq 2$ if and only if $E_2^{p,0}$ is coherent. Show that $E_\infty^{p,q}$ is coherent, and hence that $E_2^{p,0}$ is coherent, thereby completing the proof of Grothendieck's Coherence Theorem 19.9.1.

□

19.9.4. ★★ Proof of Chow's Lemma 19.9.2.

We use the properness hypothesis on π through each of its three constituent parts: finite type, separated, universally closed. The parts using separatedness are particularly tricky.

As X is Noetherian, it has finitely many irreducible components. Cover X with affine open sets U_1, \dots, U_n . We may assume that each U_i meets each irreducible component. (If some U_i does not meet an irreducible component Z , then take any affine open subset Z' of $Z - \bar{X} - \bar{Z}$, and replace U_i by $U_i \cup Z'$.) Then $U := \cap_i U_i$ is a dense open subset of X . As each U_i is finite type over A , we can choose a closed embedding $U_i \subset \mathbb{A}_A^{n_i}$. Let \bar{U}_i be the (scheme-theoretic) closure of U_i in $\mathbb{P}_A^{n_i}$.

Now we have the diagonal morphism $U \rightarrow X \times_A \prod \bar{U}_i$ (where the product is over $\text{Spec } A$), which is a locally closed embedding (the composition of the closed embedding $U \hookrightarrow U^{n+1}$ with the open embedding $U^{n+1} \hookrightarrow X \times_A \prod \bar{U}_i$). Let X' be the scheme-theoretic closure of U in $X \times_A \prod \bar{U}_i$. Let μ be the composed morphism $X' \rightarrow X \times_A \prod \bar{U}_i \rightarrow X$, so we have a diagram

$$\begin{array}{ccc}
 X' & & \\
 \downarrow \text{cl. emb.} & \searrow \mu & \\
 X \times_A \prod \bar{U}_i & \xrightarrow{\text{proj.}} & X \\
 \downarrow \text{proper} & & \downarrow \text{proper} \\
 \prod \bar{U}_i & \xrightarrow{\text{proj.}} & \text{Spec } A
 \end{array}$$

(where the square is Cartesian). The morphism μ is projective (as it is the composition of two projective morphisms and X is quasicompact, Exercise 18.3.B). We will conclude the argument by showing that $\mu^{-1}(U) = U$ (or more precisely, μ is an isomorphism above U), and that $X' \rightarrow \prod \bar{U}_i$ is a closed embedding (from which the composition

$$X \longrightarrow \prod \bar{U}_i \longrightarrow \text{Spec } A$$

is projective).

19.9.E. EXERCISE. Suppose T_0, \dots, T_n are *separated* schemes over A with isomorphic open sets, which we sloppily call V in each case. Then V is a locally closed subscheme of $T_0 \times_A \cdots \times_A T_n$. Let \bar{V} be the closure of this locally closed subscheme.

Show that

$$\begin{aligned}
 V &\cong \overline{V} \cap (V \times_A T_1 \times_A \cdots \times_A T_n) \\
 &= \overline{V} \cap (T_0 \times_A V \times_A T_2 \times_A \cdots \times_A T_n) \\
 &= \cdots \\
 &= \overline{V} \cap (T_0 \times_A \cdots \times_A T_{n-1} \times_A V).
 \end{aligned}$$

(Hint for the first isomorphism: the graph of the morphism $V \rightarrow T_1 \times_A \cdots \times_A T_n$ is a closed embedding, as $T_1 \times_A \cdots \times_A T_n$ is separated over A , by Proposition 11.3.6. Thus the scheme-theoretic closure of V in $V \times_A T_1 \times_A \cdots \times_A T_n$ is V itself. Finally, the scheme-theoretic closure can be computed locally, essentially by Theorem 9.4.4.)

19.9.F. EXERCISE. Using (the idea behind) the previous exercise, show that $\mu^{-1}(U) = U$.

It remains to show that $X' \rightarrow \prod \overline{U}_i$ is a closed embedding. Now $X' \rightarrow \prod \overline{U}_i$ is closed (it is the composition of two closed maps), so it suffices to show that $X' \rightarrow \prod \overline{U}_i$ is a locally closed embedding.

19.9.G. EXERCISE. Let A_i be the closure of U in

$$B_i := X \times_A \overline{U}_1 \times_A \cdots \times_A U_i \times_A \cdots \times_A \overline{U}_n$$

(only the i th term is missing the bar), and let C_i be the closure of U in

$$D_i := \overline{U}_1 \times_A \cdots \times_A U_i \times_A \cdots \times_A \overline{U}_n.$$

Show that there is an isomorphism $A_i \xrightarrow{\sim} C_i$ induced by the projection $B_i \rightarrow D_i$. Hint: note that the section $D_i \rightarrow B_i$ of the projection $B_i \rightarrow D_i$, given informally by $(t_1, \dots, t_n) \mapsto (t_i, t_1, \dots, t_n)$, is a closed embedding, as it can be interpreted as the graph of a morphism to a separated A -scheme (Proposition 11.3.6). So U can be interpreted as a locally closed subscheme of D_i , which in turn can be interpreted as a closed subscheme of B_i . Thus the closure of U in D_i may be identified with its closure in B_i .

As the U_i cover X , the $\mu^{-1}(U_i)$ cover \overline{X} . But $\mu^{-1}(U_i) = A_i$ (closure can be computed locally — the closure of U in B_i is the intersection of B_i with the closure \overline{X} of U in $X \times_A \overline{U}_1 \times_A \cdots \times_A \overline{U}_n$).

Hence over each U_i , we get a closed embedding of $A_i \hookrightarrow D_i$, and thus $X' \rightarrow \prod \overline{U}_i$ is a locally closed embedding as desired. \square

19.9.5. Other versions of Chow's Lemma. We won't use these versions, but their proofs are similar to what we have already shown.

19.9.6. Remark. Notice first that if X is reduced (resp. irreducible, integral), then X' can be taken to be reduced (resp. irreducible, integral) as well.

19.9.H. EXERCISE. By suitably crossing out lines in the proof above, weaken the hypothesis “ π is proper” to “ π finite type and separated”, at the expense of weakening the conclusion “ $\pi \circ \mu$ is projective” to “ $\pi \circ \mu$ is quasiprojective”.

19.9.7. Remark. The target $\operatorname{Spec} A$ can be generalized to a scheme S that is (i) Noetherian, or (ii) separated and quasicompact with finitely many irreducible components. This can be combined with Remark 19.9.6 and Exercise 19.9.H. See [Gr-EGA, II.5.6.1] for a proof. See also [GW, Thm. 13.100] for a version that is slightly more general.

Application: Curves

We now use what we have developed to study something explicit — curves. Throughout this chapter, we will assume that all curves are projective, geometrically integral, regular curves over a field k . We will sometimes add the hypothesis that k is algebraically closed. Most people are happy with working over algebraically closed fields, and those people should ignore the adverb “geometrically”.

We certainly do not need the massive machinery we have developed in order to understand curves, but with the perspective we have gained, the development is quite clean. The key ingredients we will need are as follows. We use a criterion for a morphism to be a closed embedding, that we prove in §20.1. We use the “black box” of Serre duality (to be proved in Chapter 31). In §20.2, we use this background to observe a very few useful facts, which we will use repeatedly. Finally, in the course of applying them to understand curves of various genera, we develop the theory of hyperelliptic curves in a hands-on way (§20.5), in particular proving a special case of the Riemann-Hurwitz formula.

If you are jumping into this chapter without reading much beforehand, you should skip §20.1, taking Theorem 20.1.1 as a black box. Depending on your background, you may want to skip §20.2 as well, taking the “crucial tools” as a black box.

20.1 A criterion for a morphism to be a closed embedding

We will repeatedly use a criterion for when a morphism is a closed embedding, which is not special to curves. Before stating it, we recall some facts about closed embeddings. Suppose $\pi: X \rightarrow Y$ is a closed embedding. Then π is projective, and it is injective on points. This is not enough to ensure that it is a closed embedding, as the example of the normalization of the cusp shows (Figure 10.4). Another example is the following.

20.1.A. EXERCISE (FROBENIUS, CF. §8.3.Q). Show that $\pi: \mathbb{A}_{\mathbb{F}_p}^1 \rightarrow \mathbb{A}_{\mathbb{F}_p}^1$ given by $x \mapsto x^p$ is a bijection on points, and induces an isomorphism of residue fields on closed points, yet is not a closed embedding.

The additional information you need to ensure that a morphism π is a closed embedding is that the tangent map is injective at all closed points.

20.1.B. EXERCISE. Show (directly, not invoking Theorem 20.1.1) that in the two examples described above (the normalization of a cusp and the Frobenius morphism), the tangent map is *not* injective at all closed points.

20.1.1. Theorem. — Suppose $k = \bar{k}$, and $\pi: X \rightarrow Y$ is a projective morphism of finite type k -schemes that is injective on closed points and injective on tangent vectors at closed points. Then π is a closed embedding.

(Remark: In this case of finite type schemes over an algebraically closed field, “injective on tangent vectors at closed points” is equivalent to “unramified”, see Exercise 22.6.I. We will define unramified in §22.6.)

The example $\text{Spec } \mathbb{C} \rightarrow \text{Spec } \mathbb{R}$ shows that we need the hypothesis that k is algebraically closed in Theorem 20.1.1. Those allergic to algebraically closed fields should still pay attention, as we will use this to prove things about curves over k where k is *not* necessarily algebraically closed (see also Exercises 10.2.K and 20.1.E).

We need the hypothesis that the morphism be projective, as shown by the example of Figure 20.1. It is the normalization of the node, except we erase one of the preimages of the node. We map \mathbb{A}^1 to the plane, so that its image is a curve with one node. We then consider the morphism we get by discarding one of the preimages of the node. Then this morphism is an injection on points, and is also injective on tangent vectors, but it is not a closed embedding. (In the world of differential geometry, this fails to be an embedding because the map doesn’t give a homeomorphism onto its image.)




FIGURE 20.1. (variant of Figure 8.4) The projectivity hypothesis in Theorem 20.1.1 cannot be dropped

Theorem 20.1.1 appears to be fundamentally a statement about varieties, but it isn’t. We will reduce it to the following result.

20.1.2. Theorem. — Suppose $\pi: X \rightarrow Y$ is a finite morphism of Noetherian schemes whose degree at every point of Y (§14.3.4) is 0 or 1. Then π is a closed embedding.

20.1.C. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a finite morphism whose degree at every point of Y is 0 or 1. Show that π is injective on points (easy). If $p \in X$ is any point, show that π induces an isomorphism of residue fields $\kappa(\pi(p)) \xrightarrow{\sim} \kappa(p)$. Show that π induces an injection of tangent spaces. Thus key hypotheses of Theorem 20.1.1 are implicitly in the hypotheses of Theorem 20.1.2.

20.1.3. Reduction of Theorem 20.1.1 to Theorem 20.1.2. The property of being a closed embedding is local on the target, so we may assume that Y is affine, say $\text{Spec } B$.

I next claim that π has finite fibers, not just finite fibers above closed points: the fiber dimension for projective morphisms is upper semicontinuous (Exercise 19.1.B, or Theorem 12.4.2(b)), so the locus where the fiber dimension is at least 1 is a closed subset, so if it is nonempty, it must contain a closed point of Y . Thus the fiber over any point is a dimension 0 finite type scheme over that point, hence a finite set.

Hence π is a projective morphism with finite fibers, thus finite by Theorem 19.1.7.

But the degree of a finite morphism is upper semicontinuous, (§14.3.4), and is at most 1 at closed points of Y (do you see why?), hence is at most 1 at all points (do you see why?).

20.1.4. Proof of Theorem 20.1.2. The problem is local on Y , so we may assume Y is affine, say $Y = \text{Spec } B$. Thus X is affine too, say $\text{Spec } A$, and π corresponds to a ring morphism $B \rightarrow A$. We wish to show that this is a surjection of rings, or (equivalently) of B -modules. Let K be the cokernel of this morphism of B -modules:

$$(20.1.4.1) \quad B \rightarrow A \rightarrow K \rightarrow 0.$$

We wish to show that $K = 0$. It suffices to show that for any maximal ideal \mathfrak{n} of B , $K_{\mathfrak{n}} = 0$. (Do you remember why?) Localizing (20.1.4.1) at \mathfrak{n} , we obtain the exact sequence

$$B_{\mathfrak{n}} \rightarrow A_{\mathfrak{n}} \rightarrow K_{\mathfrak{n}} \rightarrow 0.$$

Applying the right-exact functor $\otimes_{B_{\mathfrak{n}}} (B_{\mathfrak{n}}/\mathfrak{n}B_{\mathfrak{n}})$, we obtain

$$(20.1.4.2) \quad B_{\mathfrak{n}}/\mathfrak{n}B_{\mathfrak{n}} \xrightarrow{\alpha} A_{\mathfrak{n}}/\mathfrak{n}A_{\mathfrak{n}} \longrightarrow K_{\mathfrak{n}}/\mathfrak{n}K_{\mathfrak{n}} \longrightarrow 0.$$

Now $\text{Spec } A_{\mathfrak{n}}/\mathfrak{n}A_{\mathfrak{n}}$ is the scheme theoretic preimage of $[\mathfrak{n}] \in \text{Spec } B$, so by hypothesis, it is either empty, or the map α is an isomorphism (in which case we will see that the map α will turn out to be the map of residue fields). In the first case, $A_{\mathfrak{n}}/\mathfrak{n}A_{\mathfrak{n}} = 0$, from which $K_{\mathfrak{n}}/\mathfrak{n}K_{\mathfrak{n}} = 0$. In the second case, $K_{\mathfrak{n}}/\mathfrak{n}K_{\mathfrak{n}} = 0$ as well. Applying Nakayama's Lemma 8.2.9 (noting that A is a finitely generated B -module, hence K is too, hence $K_{\mathfrak{n}}$ is a finitely generated $B_{\mathfrak{n}}$ -module), $K_{\mathfrak{n}} = 0$ as desired. \square

Exercise 10.2.K can be used to extend Theorem 20.1.1's reach to general fields k , not necessarily algebraically closed. The next exercise is an example of how.

20.1.D. EXERCISE. Use Theorem 20.1.1 to show that the d th Veronese embedding from \mathbb{P}_k^n , corresponding to the complete linear series $|\mathcal{O}_{\mathbb{P}_k^n}(d)|$, is a closed embedding. Do the same for the Segre embedding from $\mathbb{P}_k^m \times_{\text{Spec } k} \mathbb{P}_k^n$. (This is just for practice for using this criterion. It is a weaker result than we had before; we have earlier checked both of these statements over an arbitrary base ring in Remark 9.3.8 and §10.6 respectively, and we are now checking it only over fields. However, see Exercise 20.1.E below.)

20.1.E. LESS IMPORTANT EXERCISE. Using the ideas from this section, prove that the d th Veronese embedding from $\mathbb{P}_{\mathbb{Z}}^n$ (over the integers!), is a closed embedding. (Again, we have done this before. This exercise is simply to show that these methods can easily extend to work more generally.)

20.2 A series of crucial tools

We are now ready to start understanding curves in a hands-on way. We will repeatedly make use of the following series of crucial remarks, and it will be important to have them at the tip of your tongue.

20.2.1. In what follows, C will be a projective, geometrically regular, geometrically integral curve over a field k , and \mathcal{L} is an invertible sheaf on C .

20.2.2. Reminder: Serre duality. Serre duality (Theorem 19.5.1) on a geometrically irreducible regular genus g curve C over k involves an invertible sheaf ω_C (of degree $2g - 2$, with g sections, Exercise 19.5.A), such that for any coherent sheaf \mathcal{F} on C , $h^i(C, \mathcal{F}) = h^{1-i}(C, \omega_C \otimes \mathcal{F}^\vee)$ for $i = 0, 1$. (Better: there is a duality between the two cohomology groups.)

20.2.3. Negative degree line bundles have no nonzero section. If $\deg \mathcal{L} < 0$, then $h^0(C, \mathcal{L}) = 0$. Reason: $\deg \mathcal{L}$ is the number of zeros minus the number of poles (suitably counted) of any rational section (Important Exercise 19.4.C). If there is a regular section (i.e., with no poles), then this is necessarily non-negative. Refining this argument yields the following.

20.2.4. Degree 0 line bundles, and recognizing when they are trivial. If $\deg \mathcal{L} = 0$, then $h^0(C, \mathcal{L}) = 0$ or 1; and if $h^0(C, \mathcal{L}) = 1$ then $\mathcal{L} \cong \mathcal{O}_C$. Reason: if there is a nonzero section s , it has no poles, and hence no zeros, because $\deg \mathcal{L} = 0$. Then $\operatorname{div} s = 0$, so $\mathcal{L} \cong \mathcal{O}_C(\operatorname{div} s) = \mathcal{O}_C$. (Recall how this works, cf. Important Exercise 15.2.E: s gives a trivialization for the invertible sheaf. We have a natural bijection for any open set $\Gamma(U, \mathcal{L}) \leftrightarrow \Gamma(U, \mathcal{O}_U)$, where the map from left to right is $s' \mapsto s'/s$, and the map from right to left is $f \mapsto sf$.) Conversely, for a geometrically integral projective variety, $h^0(\mathcal{O}) = 1$. (§11.5.7 shows this for k algebraically closed — this is where geometric integrality is used — and Exercise 19.2.H shows that cohomology commutes with base field extension.)

Serre duality turns these statements about line bundles of degree at most 0 into statements about line bundles of degree at least $2g - 2$, as follows.

20.2.5. We know $h^0(C, \mathcal{L})$ if the degree is sufficiently high. If $\deg \mathcal{L} > 2g - 2$, then

$$(20.2.5.1) \quad h^0(C, \mathcal{L}) = \deg \mathcal{L} - g + 1.$$

So we know $h^0(C, \mathcal{L})$ if $\deg \mathcal{L} \gg 0$. (This is important — remember this!) Reason: $h^1(C, \mathcal{L}) = h^0(C, \omega_C \otimes \mathcal{L}^\vee)$; but $\omega_C \otimes \mathcal{L}^\vee$ has negative degree (as ω_C has degree $2g - 2$), and thus this invertible sheaf has no sections. The result then follows from the Riemann-Roch Theorem in the guise of (19.4.3.1).

20.2.A. USEFUL EXERCISE (RECOGNIZING ω_C AMONG DEGREE $2g - 2$ LINE BUNDLES). Suppose \mathcal{L} is a degree $2g - 2$ invertible sheaf. Show that it has $g - 1$ or g sections, and it has g sections if and only if $\mathcal{L} \cong \omega_C$.

20.2.6. Twisting \mathcal{L} by a (degree 1) point changes h^0 by at most 1. Suppose p is any closed point of degree 1 (i.e., the residue field of p is k). Then $h^0(C, \mathcal{L}) - h^0(C, \mathcal{L}(-p)) = 0$ or 1. (The twist of \mathcal{L} by a divisor, such as $\mathcal{L}(-p)$, was defined in §15.2.11.) Reason: consider $0 \rightarrow \mathcal{O}_C(-p) \rightarrow \mathcal{O}_C \rightarrow \mathcal{O}_p \rightarrow 0$, tensor with \mathcal{L} (this

is exact as \mathcal{L} is locally free) to get

$$0 \rightarrow \mathcal{L}(-p) \rightarrow \mathcal{L} \rightarrow \mathcal{L}|_p \rightarrow 0.$$

Then $h^0(C, \mathcal{L}|_p) = 1$, so as the long exact sequence of cohomology starts off

$$0 \rightarrow H^0(C, \mathcal{L}(-p)) \rightarrow H^0(C, \mathcal{L}) \rightarrow H^0(C, \mathcal{L}|_p),$$

we are done.

20.2.7. Criterion for \mathcal{L} to be base-point-free. Suppose for this remark that k is algebraically closed, so *all* closed points have degree 1 over k . Then if $h^0(C, \mathcal{L}) - h^0(C, \mathcal{L}(-p)) = 1$ for *all* closed points p , then \mathcal{L} is base-point-free, and hence induces a morphism from C to projective space (Theorem 17.4.1). Reason: given any p , our equality shows that there exists a section of \mathcal{L} that does not vanish at p — so by definition, p is not a base point of \mathcal{L} .

20.2.8. Next, suppose p and q are distinct (closed) points of degree 1. Then $h^0(C, \mathcal{L}) - h^0(C, \mathcal{L}(-p - q)) = 0, 1$, or 2 (by repeating the argument of Remark 20.2.6 twice). If $h^0(C, \mathcal{L}) - h^0(C, \mathcal{L}(-p - q)) = 2$, then necessarily (20.2.8.1)

$$h^0(C, \mathcal{L}) = h^0(C, \mathcal{L}(-p)) + 1 = h^0(C, \mathcal{L}(-q)) + 1 = h^0(C, \mathcal{L}(-p - q)) + 2.$$

Then the linear series \mathcal{L} separates points p and q , i.e., the corresponding map f to projective space satisfies $f(p) \neq f(q)$. Reason: there is a hyperplane of projective space passing through $f(q)$ but not passing through $f(p)$, or equivalently, there is a section of \mathcal{L} vanishing at q but not vanishing at p . This is because of the last equality in (20.2.8.1).

20.2.9. By the same argument as above, if p is a (closed) point of degree 1, then $h^0(C, \mathcal{L}) - h^0(C, \mathcal{L}(-2p)) = 0, 1$, or 2 . I claim that if this is 2 , then the map corresponding to \mathcal{L} (which is already seen to be base-point-free at p from the above) separates the tangent vectors at p . To show this, we need to show that the cotangent map is *surjective*. To show surjectivity onto a one-dimensional vector space, we just need to show that the map is nonzero. So it suffices to find a function on the target vanishing at the image of p that pulls back to a function that vanishes at p to order 1 but not 2. In other words, we want a section of \mathcal{L} vanishing at p to order 1 but not 2. But that is the content of the statement $h^0(C, \mathcal{L}(-p)) - h^0(C, \mathcal{L}(-2p)) = 1$.

20.2.10. Criterion for \mathcal{L} to be very ample. Combining some of our previous comments: suppose C is a curve over an *algebraically closed* field k , and \mathcal{L} is an invertible sheaf such that for *all* closed points p and q , *not necessarily distinct*, $h^0(C, \mathcal{L}) - h^0(C, \mathcal{L}(-p - q)) = 2$, then \mathcal{L} gives a *closed embedding into projective space*, as it separates points and tangent vectors, by Theorem 20.1.1.

20.2.B. EXERCISE. Suppose that k is algebraically closed, so the previous remark applies. Show that $C \setminus \{p\}$ is affine. Hint: Show that if $j \gg 0$, then $\mathcal{O}(jp)$ is base-point-free and has at least two linearly independent sections, one of which has divisor jp . Use these two sections to map to \mathbb{P}^1 so that the set-theoretic preimage of ∞ is p . Argue that the map is finite, and that $C \setminus \{p\}$ is the preimage of \mathbb{A}^1 . (A trivial variation of this argument shows that $C \setminus \{p_1, \dots, p_n\}$ is affine if $n > 0$.)

20.2.11. Powerful conclusion: Numerical criteria for \mathcal{L} to be base-point-free and very ample. We can combine much of the above discussion to give the following useful fact. If k is algebraically closed, then $\deg \mathcal{L} \geq 2g$ implies that \mathcal{L} is base-point-free (and hence determines a morphism to projective space). Also, $\deg \mathcal{L} \geq 2g + 1$ implies that for any basis of sections $s_0, \dots, s_{\deg \mathcal{L} - g}$ of \mathcal{L} , the morphism given by $(s_0, \dots, s_{\deg \mathcal{L} - g})$ is in fact a closed embedding into $\mathbb{P}^{\deg \mathcal{L} - g}$ (so \mathcal{L} is very ample). Remember this!

20.2.C. EXERCISE (ON A PROJECTIVE, REGULAR, INTEGRAL CURVE OVER \bar{k} , AMPLE = POSITIVE DEGREE). Show that an invertible sheaf \mathcal{L} on a projective, regular, integral curve over \bar{k} is ample if and only if $\deg \mathcal{L} > 0$.

(This can be extended to curves over general fields using Exercise 20.2.D below.) Thus there is a blunt purely numerical criterion for ampleness of line bundles on curves. This generalizes to projective varieties of higher dimension; this is called *the Nakai-Moishezon criterion for ampleness*, Theorem 21.4.1.

20.2.D. EXERCISE (EXTENSION TO NON-ALGEBRAICALLY CLOSED FIELDS). Show that the statements in §20.2.11 hold even without the hypothesis that k is algebraically closed. (Hint: to show one of the facts about some curve C and line bundle \mathcal{L} , consider instead $C \times_{\text{Spec } k} \text{Spec } \bar{k}$. Then show that if the pullback of \mathcal{L} here has sections giving you one of the two desired properties, then there are sections downstairs with the same properties. You may want to use facts that we have used, such as the fact that base-point-freeness is independent of extension of base field, Exercise 19.2.I, or that the property of an affine morphism over k being a closed embedding holds if and only if it does after an extension of k , Exercise 10.2.K.)

20.2.E. EXERCISE (EXTENDING EXERCISE 20.2.C). Suppose \mathcal{L} is an invertible sheaf on a projective, geometrically regular, geometrically integral curve C (over a field k). Show that \mathcal{L} is ample if and only if $\deg \mathcal{L} > 0$. Hint: reduce to the case where k is algebraically closed, with the help of Exercise 10.2.K. (This was promised in Exercise 19.7.C.)

20.2.F. EXERCISE. Suppose C is a curve (in the sense of §20.2.1 as usual in this section), and $p \in C$ is a closed point. Show that $C - \{p\}$ is affine. Hint: Exercise 20.2.E.

We are now ready to take these facts and go to the races.

20.3 Curves of genus 0

We are now ready to (in some form) answer the question: what are the curves of genus 0?

In §7.5.9, we saw a genus 0 curve (over a field k) that was *not* isomorphic to \mathbb{P}^1 : $x^2 + y^2 + z^2 = 0$ in $\mathbb{P}_{\mathbb{R}}^2$. (It has genus 0 by (19.6.6.1).) We have already observed that this curve is *not* isomorphic to $\mathbb{P}_{\mathbb{R}}^1$, because it doesn't have an \mathbb{R} -valued point. On the other hand, we haven't seen a genus 0 curve over an algebraically closed field with this property. This is no coincidence: the lack of an existence of a k -valued point is the only obstruction to a genus 0 curve being \mathbb{P}^1 .

20.3.1. Proposition. — Suppose C is genus 0, and C has a k -valued (degree 1) point. Then $C \cong \mathbb{P}_k^1$.

Thus we see that all genus 0 (integral, regular, projective) curves over an algebraically closed field are isomorphic to \mathbb{P}^1 .

Proof. Let p be the point, and consider $\mathcal{L} = \mathcal{O}(p)$. Then $\deg \mathcal{L} = 1$, so we can apply what we know above: first, $h^0(C, \mathcal{L}) = 2$ (Remark 20.2.5), and second, these two sections give a closed embedding into \mathbb{P}_k^1 (Remark 20.2.11). But the only closed embedding of a curve into the integral curve \mathbb{P}_k^1 is an isomorphism! \square

As a bonus, Proposition 20.3.1 implies that $x^2 + y^2 + z^2 = 0$ in $\mathbb{P}_{\mathbb{R}}^2$ has no line bundles of degree 1 over \mathbb{R} ; otherwise, we could just apply the above argument to the corresponding line bundle. This example shows us that over a non-algebraically closed field, there can be genus 0 curves that are not isomorphic to \mathbb{P}_k^1 . The next result lets us get our hands on them as well.

20.3.2. Claim. — All genus 0 curves can be described as conics in \mathbb{P}_k^2 .

Proof. Any genus 0 curve has a degree -2 line bundle — the dualizing sheaf ω_C . Thus any genus 0 curve has a degree 2 line bundle: $\mathcal{L} = \omega_C^\vee$. We apply Remark 20.2.11: $\deg \mathcal{L} = 2 \geq 2g + 1$, so this line bundle gives a closed embedding into \mathbb{P}^2 . (Exercise 19.6.I ensures that the image really is a conic.) \square

20.3.A. EXERCISE. Suppose C is a genus 0 curve (projective, geometrically integral and regular). Show that C has a point of degree at most 2. (The degree of a point was defined in §5.3.8.)

The geometric means of finding Pythagorean triples presented in §7.5.8 looked quite different, but was really the same. There was a genus 0 curve C (a plane conic) with a k -valued point p , and we proved that it was isomorphic to \mathbb{P}_k^1 . The line bundle used to show the isomorphism wasn't the degree 1 line bundle $\mathcal{O}_C(p)$; it was the degree 1 line bundle $\mathcal{O}_{\mathbb{P}^2}(1)|_C \otimes \mathcal{O}_C(-p)$.

20.4 Classical geometry arising from curves of positive genus

We will use the following Proposition and Corollary later, and we take this as an excuse to revisit some very classical geometry from a modern standpoint.

20.4.1. Proposition. — Recall our standing assumptions for this chapter (§20.2.1), that C is a projective, geometrically regular, geometrically integral curve over a field k . Suppose C is not isomorphic to \mathbb{P}_k^1 (with no assumptions on the genus of C), and \mathcal{L} is an invertible sheaf of degree 1. Then $h^0(C, \mathcal{L}) < 2$.

Proof. Otherwise, let s_1 and s_2 be two (independent) sections. As the degree of the divisor of zeros of s_i is the degree of \mathcal{L} , each vanishes at a single point p_i (to order 1).

Now $p_1 \neq p_2$. (Otherwise, if $p_1 = p_2$, then s_1/s_2 has no poles or zeros, and hence is an invertible function by Hartogs' Lemma, see §12.3.16, and a constant

function using §11.5.7 after base change to \bar{k} . This contradicts the linear independence of s_1 and s_2 .)

Thus s_1 and s_2 generate a base-point-free linear series, so we have an induced morphism $C \rightarrow \mathbb{P}^1$. This is a finite degree 1 morphism from one regular curve to another (Exercise 19.4.E), which hence induces a degree 1 extension of function fields, i.e., an isomorphism of function fields, which means that the curves are isomorphic (cf. Theorem 18.4.3). But we assumed that C is not isomorphic to \mathbb{P}_k^1 , so we have a contradiction. \square

20.4.2. Corollary. — *If C is a projective, geometrically regular, geometrically integral curve over k , not isomorphic to \mathbb{P}_k^1 and p and q are degree 1 points, then $\mathcal{O}_C(p) \cong \mathcal{O}_C(q)$ if and only if $p = q$.*

20.4.A. EXERCISE. Show that if k is algebraically closed, then C has genus 0 if and only if all degree 0 line bundles are trivial.

20.4.B. EXERCISE. Suppose C is a regular plane curve of degree $e > 2$, and D_1 and D_2 are two plane curves of the same degree d not containing C . By Bézout's Theorem for plane curves (§19.6.3), D_i meets C at de points, counted “correctly”. Suppose D_1 and D_2 meet C at $de - 1$ of the “same points”, plus one more. Show that the remaining points are the same as well. More precisely, suppose there is a divisor E on C of degree $de - 1$, and degree 1 (k -valued) points p_1 and p_2 such that $D_i \cap C = E + p_i$ (as divisors on C). Show that $p_1 = p_2$. (The case $d = e = 3$ is Chasles' Theorem, and is the first case of the Cayley-Bacharach Theorem, see [EGH].)

As an entertaining application of Exercise 20.4.B, we can prove two classical results. For convenience, in this discussion we will assume k is algebraically closed, although this assumption can be easily removed.

20.4.3. Pappus's Theorem (Pappus of Alexandria). — (See Figure 20.2.) Suppose ℓ and m are distinct lines in the plane, and α, β, γ are distinct points on ℓ , and α', β', γ' are distinct points on m , and all six points are distinct from $\ell \cap m$. Then points $x = \overline{\alpha\beta'} \cap \overline{\alpha'\beta}$, $y = \overline{\alpha\gamma'} \cap \overline{\alpha'\gamma}$, and $z = \overline{\beta\gamma'} \cap \overline{\beta'\gamma}$ are collinear.




FIGURE 20.2. Pappus' Theorem

Pascal's "Mystical Hexagon" Theorem was discovered by Pascal at age 16. (What were you doing at age 16?)

20.4.4. Pascal's "Mystical Hexagon" Theorem. — (See Figure 20.3.) If a hexagon $\alpha\gamma'\beta\alpha'\gamma\beta'$ is inscribed in a smooth conic K , and opposite pairs of sides are extended until they meet, the three intersection points $x = \overline{\alpha\beta'} \cap \overline{\alpha'\beta}$, $y = \overline{\alpha'\gamma} \cap \overline{\alpha\gamma'}$, and $z = \overline{\beta\gamma'} \cap \overline{\beta'\gamma}$ are collinear.




FIGURE 20.3. Pascal's "Mystical Hexagon" Theorem

Pappus's Theorem can be seen as a degeneration of Pascal's Theorem: the conic degenerates into the union of two lines, and the six points degenerate so that α , β , and γ are on one, and α' , β' , and γ' are on the other. We thus prove Pascal's Theorem, and you should check that the proof readily applies to prove Pappus's Theorem.

20.4.C. EXERCISE. Suppose $\alpha, \beta, \gamma, \alpha', \beta'$ are five points in $\mathbb{P}_{\mathbb{R}}^2$, no three on a line. Show that there is a unique conic K passing through the five points. Show that K is regular. Explain how to construct the tangent to K at α using only a straightedge. (Hint for the last part: apply Pascal's Theorem, taking $\gamma' = \alpha$. You may need to first figure out why you can apply Pascal's Theorem in this degenerate case.)

In particular, given an ellipse $K \subset \mathbb{R}^2$ and a point $\alpha \in K$, you will now be able to construct the tangent to K at α using only a straightedge.

20.4.5. Proof of Pascal's Theorem 20.4.4. We wish to show that the line \overline{xy} meets $\overline{\beta\gamma'}$ and $\overline{\beta'\gamma}$ at the same point. Let C be the curve that is the union of K and the line \overline{xy} . (Warning, cf. §20.2.1: for one time in this chapter, we are not assuming C to be regular!)

Let D_1 be the union of the three lines $\overline{\alpha\beta'}$, $\overline{\beta\gamma'}$, and $\overline{\gamma\alpha'}$, and let D_2 be the union of the three lines $\overline{\alpha'\beta}$, $\overline{\beta'\gamma}$, and $\overline{\gamma'\alpha}$. Note that D_1 meets C at the nine points $\alpha, \beta, \gamma, \alpha', \beta', \gamma', x, y$, and $\overline{xy} \cap \overline{\beta\gamma'}$, and D_2 meets C at the same nine points, except $\overline{xy} \cap \overline{\beta\gamma'}$ is replaced by $\overline{xy} \cap \overline{\beta'\gamma}$. Thus if we knew Corollary 20.4.2 so that it applied to our C (which is singular), then we would be done (cf. Exercise 20.4.B).

So we extend Corollary 20.4.2 to our situation. So to do this, we extend Proposition 20.4.1 to our situation. We have a plane cubic C which is the union of a line L and a conic K , and two points $p = \overline{xy} \cap \overline{\beta\gamma'}$ and $q = \overline{xy} \cap \overline{\beta'\gamma}$, with $\mathcal{O}(p) \cong \mathcal{O}(q)$. (Reason: both are the restriction of $\mathcal{O}_{\mathbb{P}^2}(3)$ to C , twisted by $-(\alpha + \beta + \gamma + \alpha' + \beta' + \gamma' + x + y)$.) Call this invertible sheaf \mathcal{L} . Suppose $p \neq q$. Then the two sections of \mathcal{L} with zeros at p and q give a base-point-free linear series, and thus a morphism $\pi: C \rightarrow \mathbb{P}^1$, with $\pi^{-1}(0) = p$ and $\pi^{-1}(\infty) = q$.

By the argument in the proof of Proposition 20.4.1, π gives an isomorphism of L with \mathbb{P}^1 . As π is proper, $\pi(K)$ is closed, and as K is irreducible, $\pi(K)$ is irreducible. As $\pi(K)$ does not contain 0 (or ∞), it can't be all of \mathbb{P}^1 . Hence $\pi(K)$ is a point.

The conic K meets the line L in two points with multiplicity (by Bézout's Theorem, §19.6.3, or by simple algebra). If $K \cap L$ consists of two points a and b , we have a contradiction: K can't be contracted to both $\pi(a)$ and $\pi(b)$. But what happens if K meets L at one point, i.e., if the conic is tangent to the line?

20.4.D. EXERCISE. Finish the proof of Pascal's Theorem by dealing with this case. Hint: nilpotents will come to the rescue. The intuition is as follows: $K \cap L$ is a subscheme of L of length 2, but the scheme theoretic image $\pi(K)$ can be shown to be a reduced closed point. □

20.4.6. Remark. The key motivating fact that makes our argument work, Proposition 20.4.1, is centrally about curves not of genus 0, yet all the curves involved in Pappus's Theorem and Pascal's Theorem have genus 0. The insight to keep in mind is that union of curves of genus 0 need not have genus 0. In our case, it mattered that cubic curves have genus 1, even if they are union of \mathbb{P}^1 's.

20.5 Hyperelliptic curves

We next discuss an important class of curves, the hyperelliptic curves. In this section, *we assume k is algebraically closed of characteristic not 2*. (These hypotheses can be relaxed, at some cost.)

A (projective regular integral) genus g curve C is **hyperelliptic** if it admits a double cover of (i.e., degree 2, necessarily finite, morphism to) \mathbb{P}_k^1 . For convenience, when we say C is hyperelliptic, we will implicitly have in mind a *choice* of double cover $\pi: C \rightarrow \mathbb{P}^1$. (We will later see that if $g \geq 2$, then there is at most one such double cover, Proposition 20.5.7, so this is not a huge assumption.) The map π is called the **hyperelliptic map**.

By Exercise 18.4.D, the preimage of any closed point p of \mathbb{P}^1 consists of either one or two points. If $\pi^{-1}(p)$ is a single point, we say p is a **branch point**, and $\pi^{-1}(p)$ is a **ramification point** of π . (The notion of ramification will be defined more generally in §22.6.)

20.5.1. Theorem (hyperelliptic Riemann-Hurwitz formula). — Suppose $k = \bar{k}$ and $\text{char } k \neq 2$, $\pi: C \rightarrow \mathbb{P}_k^1$ is a double cover by a projective regular irreducible genus g curve over k . Then π has $2g + 2$ branch points.

This is a special case of the Riemann-Hurwitz formula, which we will state and prove in §22.7. You may have already heard about genus 1 complex curves double covering \mathbb{P}^1 , branched over 4 points.

To prove Theorem 20.5.1, we first prove the following.

20.5.2. Proposition. — *Assume $\text{char } k \neq 2$ and $k = \bar{k}$. Given r distinct points $p_1, \dots, p_r \in \mathbb{P}^1$, there is precisely one double cover branched at precisely these points if r is even, and none if r is odd.*

Proof. Pick points 0 and ∞ of \mathbb{P}^1 distinct from the r branch points. All r branch points are in $\mathbb{P}^1 - \infty = \mathbb{A}^1 = \text{Spec } k[x]$. Suppose we have a double cover of \mathbb{A}^1 , $C' \rightarrow \mathbb{A}^1$, where x is the coordinate on \mathbb{A}^1 . This induces a quadratic field extension K over $k(x)$. As $\text{char } k \neq 2$, this extension is Galois. Let $\sigma: K \rightarrow K$ be the Galois involution. Let y be a nonzero element of K such that $\sigma(y) = -y$, so 1 and y form a basis for K over the field $k(x)$, and are eigenvectors of σ . Now $\sigma(y^2) = y^2$, so $y^2 \in k(x)$. We can replace y by an appropriate $k(x)$ -multiple so that y^2 is a polynomial, with no repeated factors, and monic. (This is where we use the hypothesis that k is algebraically closed, to get leading coefficient 1.)

Thus $y^2 = x^N + a_{N-1}x^{N-1} + \dots + a_0$, where the polynomial on the right (call it $f(x)$) has no repeated roots. The Jacobian criterion (in the guise of Exercise 13.2.E) implies that this curve C'_0 in $\mathbb{A}^2 = \text{Spec } k[x, y]$ is regular. Then C'_0 is normal and has the same function field as C' . Thus C'_0 and C' are both normalizations of \mathbb{A}^1 in the finite extension of fields generated by y , and hence are isomorphic. Thus we have identified C' in terms of an explicit equation.

The branch points correspond to those values of x for which there is exactly one value of y , i.e., the roots of $f(x)$. In particular, $N = r$, and

$$f(x) = (x - p_1) \cdots (x - p_r),$$

where the p_i are interpreted as elements of \bar{k} .

Having mastered the situation over \mathbb{A}^1 , we return to the situation over \mathbb{P}^1 . We will examine the branched cover over the affine open set $\mathbb{P}^1 \setminus \{0\} = \text{Spec } k[u]$, where $u = 1/x$. The previous argument applied to $\text{Spec } k[u]$ rather than $\text{Spec } k[x]$ shows that any such double cover must be of the form

$$\begin{aligned} C'' &= \text{Spec } k[z, u] / (z^2 - (u - 1/p_1) \cdots (u - 1/p_r)) \\ &= \text{Spec } k[z, u] / \left(\left(\prod p_i \right) z^2 - (-u)^r f(1/u) \right) \\ &\longrightarrow \text{Spec } k[u] = \mathbb{A}^1. \end{aligned}$$

So if there is a double cover over all of \mathbb{P}^1 , it must be obtained by gluing C'' to C' , “over” the gluing of $\text{Spec } k[x]$ to $\text{Spec } k[u]$ to obtain \mathbb{P}^1 .

Thus in $K(C)$, we must have

$$z^2 = u^r f(1/u) = f(x)/x^r = y^2/x^r$$

from which $z^2 = y^2/x^r$.

If r is even, considering $K(C)$ as generated by y and x , there are two possible values of z : $z = \pm y/x^{r/2}$. After renaming z by $-z$ if necessary, there is a single way of gluing these two patches together (we choose the positive square root).

If r is odd, the result follows from Exercise 20.5.A below.

20.5.A. EXERCISE. Suppose $\text{char } k \neq 2$. Show that x does not have a square root in the field $k(x)[y]/(y^2 - f(x))$, where f is a polynomial with nonzero roots p_1, \dots, p_r . (Possible hint: why is $\sqrt{3} \notin \mathbb{Q}(\sqrt{2})$?) Explain how this proves Proposition 20.5.2 in the case where r is odd. \square

For future reference, we collect here our explicit (two-affine) description of the hyperelliptic cover $C \rightarrow \mathbb{P}^1$.

$$(20.5.2.1) \quad \begin{array}{ccc} \text{Spec } k[x, y]/(y^2 - f(x)) & \xrightarrow[\substack{z=y/x^{r/2} \\ y=z/u^{r/2}}]{} & \text{Spec } k[u, z]/(z^2 - u^r f(1/u)) \\ \downarrow & & \downarrow \\ \text{Spec } k[x] & \xrightarrow[\substack{u=1/x \\ x=1/u}]{} & \text{Spec } k[u] \end{array}$$

20.5.3. If k is not algebraically closed. If k is not algebraically closed (but of characteristic not 2), the above argument shows that if we have a double cover of \mathbb{A}^1 , then it is of the form $y^2 = af(x)$, where f is monic (and $a \neq 0$). Furthermore, if a and a' differ (multiplicatively) by an element of $(k^\times)^2$, then $y^2 = af(x)$ is isomorphic to $y^2 = a'f(x)$. You may be able to use this to show that (assuming that $k^\times \neq (k^\times)^2$) a double cover is *not* determined by its branch points. Moreover, this failure is classified by $k^\times/(k^\times)^2$. Thus we have lots of curves that are not isomorphic over k , but become isomorphic over \bar{k} . These are often called *twists* of each other.

(In particular, once we define elliptic curves, you will be able to show that there exist two elliptic curves over \mathbb{Q} with the same j -invariant, that are not isomorphic, see Exercise 20.10.G.)

20.5.4. Back to proving the hyperelliptic Riemann-Hurwitz formula, Theorem 20.5.1. Our explicit description of the unique double cover of \mathbb{P}^1 branched over r different points will allow us to compute the genus, thereby completing the proof of Theorem 20.5.1.

We continue the notation (20.5.2.1) of the proof of Proposition 20.5.2. Suppose \mathbb{P}^1 has affine cover by $\text{Spec } k[x]$ and $\text{Spec } k[u]$, with $u = 1/x$, as usual. Suppose $C \rightarrow \mathbb{P}^1$ is a double cover, given by $y^2 = f(x)$ over $\text{Spec } k[x]$, where f has degree r , and $z^2 = u^r f(1/u)$. Then C has an affine open cover by $\text{Spec } k[x, y]/(y^2 - f(x))$ and $\text{Spec } k[u, z]/(z^2 - u^r f(1/u))$. The corresponding Čech complex for \mathcal{O}_C is

$$0 \longrightarrow k[x, y]/(y^2 - f(x)) \times k[u, z]/(z^2 - u^r f(1/u)) \xrightarrow{d} (k[x, y]/(y^2 - f(x)))_x \longrightarrow 0.$$

The second (and last) nonzero part of the complex has basis consisting of monomials $x^n y^\epsilon$, where $n \in \mathbb{Z}$ and $\epsilon = 0$ or 1 . To compute the genus $g = h^1(C, \mathcal{O}_C)$, we must compute $\text{coker } d$. We can use the first factor $k[x, y]/(y^2 - f(x))$ to hit the monomials $x^n y^\epsilon$ where $n \in \mathbb{Z}^{\geq 0}$, and $\epsilon = 0$ or 1 . The image of the second factor is generated by elements of the form $u^m z^\epsilon$, where $m \geq 0$ and $\epsilon = 0$ or 1 . But

$u^m z^e = x^{-m} (y/x^{r/2})^e$. By inspection, the cokernel has basis generated by monomials $x^{-1}y, x^{-2}y, \dots, x^{-r/2+1}y$, and thus has dimension $r/2 - 1$. Hence $g = r/2 - 1$, from which Theorem 20.5.1 follows. \square

20.5.5. Curves of every genus. As a consequence of the hyperelliptic Riemann-Hurwitz formula (Theorem 20.5.1), we see that there are curves of every genus $g \geq 0$ over an algebraically closed field of characteristic not 2: to get a curve of genus g , consider the branched cover branched over $2g + 2$ distinct points. The unique genus 0 curve is of this form, and we will see in §20.6.2 that every genus 2 curve is of this form. We will soon see that every genus 1 curve (reminder: over an algebraically closed field!) is too (§20.9.5). But it is too much to hope that all curves are of this form, and we will soon see (§20.7.3) that there are genus 3 curves that are *not* hyperelliptic, and we will even get heuristic evidence that “most” genus 3 curves are not hyperelliptic. We will later give vague evidence (that can be made precise) that “most” genus g curves are not hyperelliptic if $g > 2$ (§20.8.3).

20.5.B. EXERCISE. Verify that a curve C of genus at least 1 admits a degree 2 cover of \mathbb{P}^1 if and only if it admits a degree 2 invertible sheaf \mathcal{L} with $h^0(C, \mathcal{L}) = 2$. Possibly in the course of doing this, verify that if C is a curve of genus at least 1, and C has a degree 2 invertible sheaf \mathcal{L} with at least 2 (linearly independent) sections, then \mathcal{L} has precisely two sections, and that this \mathcal{L} is base-point-free and gives a hyperelliptic map.

20.5.6. Proposition. — Assume $g \geq 2$. If \mathcal{L} corresponds to a hyperelliptic cover $C \rightarrow \mathbb{P}^1$, then $\mathcal{L}^{\otimes(g-1)} \cong \omega_C$.

Proof. Compose the hyperelliptic map with the $(g - 1)$ th Veronese map:

$$(20.5.6.1) \quad C \xrightarrow{|\mathcal{L}|} \mathbb{P}^1 \xrightarrow{|\mathcal{O}_{\mathbb{P}^1}(g-1)|} \mathbb{P}^{g-1}.$$

The composition corresponds to $\mathcal{L}^{\otimes(g-1)}$. This invertible sheaf has degree $2g - 2$ (by Exercise 19.4.G). The pullback $H^0(\mathbb{P}^{g-1}, \mathcal{O}(1)) \rightarrow H^0(C, \mathcal{L}^{\otimes(g-1)})$ is injective because the image of C in \mathbb{P}^{g-1} (a rational normal curve) is nondegenerate: if there were a hyperplane $s \in H^0(\mathbb{P}^{g-1}, \mathcal{O}(1))$ that pulled back to 0 on C , then the image of C would lie in that hyperplane, yet a rational normal curve cannot. Thus $\mathcal{L}^{\otimes(g-1)}$ has at least g sections. But by Exercise 20.2.A, the only invertible sheaf of degree $2g - 2$ with (at least) g sections is the dualizing sheaf. \square

As an added bonus, we see that the composition of (20.5.6.1) is the *complete* linear series $|\mathcal{L}^{\otimes(g-1)}|$ — all sections of $\mathcal{L}^{\otimes(g-1)}$ come up in this way.

20.5.7. Proposition (a genus ≥ 2 curve can be hyperelliptic in only one way). — Any curve C of genus at least 2 admits at most one double cover of \mathbb{P}^1 . More precisely, if \mathcal{L} and \mathcal{M} are two degree two line bundles yielding maps $C \rightarrow \mathbb{P}^1$, then $\mathcal{L} \cong \mathcal{M}$.

Proof. If C is hyperelliptic, then we can recover the hyperelliptic map by considering the canonical linear series given by ω_C (the *canonical map*, which we will use again repeatedly in the next few sections): it is a double cover of a degree $g - 1$ rational normal curve (by the previous proposition), which is isomorphic to \mathbb{P}^1 . This

double cover is the hyperelliptic cover (also by the proof of the previous proposition). Thus we have uniquely recovered the map $C \rightarrow \mathbb{P}^1$, and this map must be induced by both \mathcal{L} and \mathcal{M} , from which we have $\mathcal{L} \cong \mathcal{M}$ (using Theorem 17.4.1, relating maps to projective space and line bundles). \square

20.5.8. The “space of hyperelliptic curves”. Thanks to Proposition 20.5.7, we can now classify hyperelliptic curves of genus at least 2. Hyperelliptic curves of genus $g \geq 2$ correspond to precisely $2g + 2$ distinct points on \mathbb{P}^1 modulo S_{2g+2} , and modulo automorphisms of \mathbb{P}^1 . Thus “the space of hyperelliptic curves” has dimension

$$2g + 2 - \dim \operatorname{Aut} \mathbb{P}^1 = 2g - 1.$$

This is not a well-defined statement, because we haven’t rigorously defined “the space of hyperelliptic curves” — it is an example of a *moduli space*. For now, take this as a plausibility statement. It is also plausible that this space is irreducible and reduced — it is the image of something irreducible and reduced.

20.6 Curves of genus 2

20.6.1. The reason for leaving genus 1 for later. It might make most sense to jump to genus 1 at this point, but the theory of elliptic curves is especially rich and subtle, so we will leave it for §20.9.

In general, curves have quite different behaviors (topologically, arithmetically, geometrically) depending on whether $g = 0$, $g = 1$, or $g \geq 2$. This trichotomy extends to varieties of higher dimension. We already have some inkling of it in the case of curves. Arithmetically, genus 0 curves can have lots and lots of rational points, genus 1 curves can have lots of rational points, and by Faltings’ Theorem (Mordell’s Conjecture) any curve of genus at least 2 has at most finitely many rational points. (Thus even before the proof of important cases of the Taniyama-Shimura conjecture [Wi, TWi], and hence Fermat’s Last Theorem, we knew that $x^n + y^n = z^n$ in \mathbb{P}^2 has at most finitely many rational solutions for $n \geq 4$, as such curves have genus $\binom{n-1}{2} > 1$, see (19.6.6.1).) In the language of differential geometry, Riemann surfaces of genus 0 are positively curved, Riemann surfaces of genus 1 are flat, and Riemann surfaces of genus at least 2 are negatively curved. It is a fact that curves of genus at least 2 have finite automorphism groups (see §22.7.9), while curves of genus 1 have one-dimensional automorphism groups, see Question 20.10.6), and the unique curve of genus 0 over an algebraically closed field has a three-dimensional automorphism group (see Exercises 17.4.B and 17.4.C). (See Exercise 22.5.G for more on this issue.)

20.6.2. Back to curves of genus 2.

Over an algebraically closed field, we saw in §20.3 that there is only one genus 0 curve. In §20.5 we established that there are hyperelliptic curves of genus 2. How can we get a hold of curves of genus 2? For example, are they all hyperelliptic? “How many” are there? We now tackle these questions.

Fix a curve C of genus $g = 2$. Then ω_C is degree $2g - 2 = 2$, and has 2 sections (Exercise 20.2.A). By Exercise 20.5.B, ω_C is base-point-free, and thus induces a

double cover $C \rightarrow \mathbb{P}^1$ (unique up to automorphisms of \mathbb{P}^1 , which we studied in §17.4.5). Conversely, any double cover $C \rightarrow \mathbb{P}^1$ arises from a degree 2 invertible sheaf with at least 2 sections. Hence if $g(C) = 2$, this invertible sheaf must be the dualizing sheaf, just because it is a degree $2g - 2 = 2$ line bundle with $g = 2$ sections (or alternatively, by the easiest case of Proposition 20.5.6).

Hence we have a natural bijection between genus 2 curves and genus 2 double covers of \mathbb{P}^1 (up to automorphisms of \mathbb{P}^1). If the characteristic is not 2, the hyperelliptic Riemann-Hurwitz formula (Theorem 20.5.1) shows that the double cover is branched over $2g + 2 = 6$ geometric points. In particular, we have a “three-dimensional space of genus 2 curves”. This is not rigorous, but we can certainly show that there are an infinite number of genus 2 curves. Precisely:

20.6.A. EXERCISE. Fix an algebraically closed field k of characteristic not 2. Show that there are an infinite number of (pairwise) non-isomorphic genus 2 curves over k .

20.6.B. EXERCISE. Show that every genus 2 curve (over any field of characteristic not 2) has finite automorphism group.

20.7 Curves of genus 3

20.7.1. The canonical embedding.

We begin with two exercises in general genus $g > 1$, then specialize to genus 3.

20.7.A. EXERCISE. Assume $k = \bar{k}$ (purely to avoid distraction — feel free to remove this hypothesis). Suppose C is a genus g curve, with $g > 1$. Show that if C is not hyperelliptic, then the dualizing invertible sheaf gives a closed embedding $C \hookrightarrow \mathbb{P}^{g-1}$. (In the hyperelliptic case, we have already seen that the dualizing sheaf gives us a double cover of a rational normal curve.) Such a curve in \mathbb{P}^{g-1} is called a **canonical curve**, and this closed embedding is called the **canonical embedding** of C .

20.7.B. EXERCISE. Suppose C is a curve of genus $g > 1$, over a field k that is not necessarily algebraically closed. Show that C has a closed point of degree at most $2g - 2$ over the base field. (For comparison: if $g = 1$, for any n , there is a genus 1 curve over \mathbb{Q} with no point of degree less than n , [MO286390]!)

(This exercise gives a second proof that all genus 2 curves are hyperelliptic.)

Suppose now that C is a curve of genus 3. Then ω_C has degree $2g - 2 = 4$, and has $g = 3$ sections.

20.7.2. Claim. — *The invertible sheaf ω_C is base-point-free, and hence gives a map to \mathbb{P}^2 .*

Proof. We check base-point-freeness by working over the algebraic closure \bar{k} (which we can, by Exercise 19.2.I). For any point p , by Riemann-Roch,

$$h^0(C, \omega_C(-p)) - h^0(C, \mathcal{O}(p)) = \deg(\omega_C(-p)) - g + 1 = 3 - 3 + 1 = 1.$$

But $h^0(C, \mathcal{O}(p)) = 1$ by Proposition 20.4.1, so

$$h^0(C, \omega_C(-p)) = 2 = h^0(C, \omega_C) - 1.$$

Thus p is not a base point of ω_C for any p , so by Criterion 20.2.7 ω_C is base-point-free. \square

The next natural question is: Is this a closed embedding? Again, we can check over the algebraic closure. We use our “closed embedding test” of §20.1, in the form of §20.2.10.

Thus we conclude that if (and only if) C is not hyperelliptic, then the canonical map describes C as a degree 4 curve in \mathbb{P}^2 .

Conversely, any quartic plane curve is canonically embedded (Exercise 20.7.A). Reason: the curve has genus 3 (see (19.6.6.1)), and is mapped by an invertible sheaf of degree 4 with 3 sections. But by Exercise 20.2.A, the only invertible sheaf of degree $2g - 2$ with g sections is ω_C .

In particular, each nonhyperelliptic genus 3 curve can be described as a quartic plane curve in only one way (up to automorphisms of \mathbb{P}^2 , which were described in Exercise 17.4.B).

In conclusion, there is a bijection between nonhyperelliptic genus 3 curves, and plane quartic curves up to projective linear transformations (§17.4.5).

20.7.3. Remark. In particular, as there exist regular plane quartics (Exercise 13.3.E), there exist nonhyperelliptic genus 3 curves.

20.7.C. EXERCISE (CF. §20.5.8). Give a heuristic (non-rigorous) argument that the nonhyperelliptic curves of genus 3 form a family of dimension 6. (Hint: Count the dimension of the family of regular quartic plane curves, and quotient by $\text{Aut } \mathbb{P}^2 = \text{PGL}(3)$.)

The genus 3 curves thus seem to come in two families: the hyperelliptic curves (a family of dimension 5), and the nonhyperelliptic curves (a family of dimension 6). This is misleading — they actually come in a single family of dimension 6.

In fact, hyperelliptic curves are naturally limits of nonhyperelliptic curves. We can write down an explicit family. (This explanation necessarily requires some hand-waving, as it involves topics we haven’t seen yet.) Suppose we have a hyperelliptic curve branched over $2g + 2 = 8$ points of \mathbb{P}^1 . Choose an isomorphism of \mathbb{P}^1 with a conic in \mathbb{P}^2 . There is a regular quartic meeting the conic at precisely those 8 points. (This requires a short argument using Bertini’s Theorem 13.4.2, which we omit.) Then if f is the equation of the conic, and g is the equation of the quartic, then $f^2 + t^2g$ is a family of quartics that are smooth for most t (smoothness is an open condition in t , as we will see in Theorem 26.3.3 on generic smoothness). The $t = 0$ case is a double conic. Then it is a fact that if you normalize the total space of the family, the central fiber (above $t = 0$) turns into our hyperelliptic curve. Thus we have expressed our hyperelliptic curve as a limit of nonhyperelliptic curves.

20.7.D. UNIMPORTANT EXERCISE. A (projective) curve (over a field k) admitting a degree 3 cover of \mathbb{P}^1 is called **trigonal**. Show that every nonhyperelliptic genus 3 complex curve is trigonal, by taking the quartic model in \mathbb{P}^2 , and projecting to \mathbb{P}^1 from any point on the curve. Do this by choosing coordinates on \mathbb{P}^2 so that p is

at $[0, 0, 1]$. (The same idea, applied to cubics rather than quartics, will be used in §20.9.9.)

20.7.4. ★ A genus 3 curve with no nontrivial automorphisms.

We have seen that a (smooth projective integral) curve of genus at most 2 always has nontrivial automorphisms. It turns out that there are genus 3 curves with no nontrivial automorphisms.

20.7.E. EXERCISE. Suppose $C' \subset \mathbb{P}^2$ is a smooth plane quartic curve (over any field k). Show that there is bijection between automorphisms of C' and automorphisms of \mathbb{P}^2 preserving C' (as a set).

Thus to find a genus 3 curve with no nontrivial automorphisms, we need only find a smooth quartic plane curve C' such that the only automorphism of \mathbb{P}^2 fixing C' as a set must be the identity. Your intuition may (correctly) tell you that most quartics have this property. But exhibiting a specific C' (with proof) requires rolling up our sleeves and getting to work. Poonen gives automorphism free curves over any field in [P1]; an example in characteristic 0 is

$$y^3z - 3yz^3 = 3x^4 - 4x^3z + z^4.$$

20.7.F. EXERCISE. Suppose C is a smooth projective curve with no nontrivial automorphisms. Show that no two distinct open subsets of C are isomorphic.

20.7.5. Genus 3 curves with nontrivial automorphisms.

Certainly genus 3 curves can have nontrivial automorphisms: witness hyperelliptic curves. More impressive is the Klein quartic

$$x^3y + y^3z + z^3x = 0$$

in $\mathbb{P}_{\mathbb{C}}^2$, which has 168 automorphisms. (Can you find them all?) In fact, the automorphism group of the Klein quartic is the unique finite simple group of order 168 (the second-smallest nonabelian finite simple group). In §22.7.9, we will see that no genus 3 curve can have more automorphisms (assuming the hard fact that curves of genus greater than 1 have finite automorphism groups).

20.8 Curves of genus 4 and 5

We first consider nonhyperelliptic curves C of genus 4. Note that $\deg \omega_C = 6$ and $h^0(C, \omega_C) = 4$, so the canonical map expresses C as a sextic curve in \mathbb{P}^3 . Let $i : C \hookrightarrow \mathbb{P}^3$ be the canonical (closed) embedding. We shall see that all such C are complete intersections of quadric surfaces and cubic surfaces, and conversely all regular complete intersections of quadric surfaces and cubic surfaces are genus 4 nonhyperelliptic curves, canonically embedded.

By (20.2.5.1) (Riemann-Roch and Serre duality),

$$h^0(C, i^* \mathcal{O}_{\mathbb{P}^3}(2)) = h^0(C, \omega_C^{\otimes 2}) = \deg \omega_C^{\otimes 2} - g + 1 = 12 - 4 + 1 = 9.$$

On the other hand, $h^0(\mathbb{P}^3, \mathcal{O}_{\mathbb{P}^3}(2)) = \binom{3+2}{2} = 10$ (Exercise 15.1.C). Thus the restriction map

$$H^0(\mathbb{P}^3, \mathcal{O}_{\mathbb{P}^3}(2)) \longrightarrow H^0(C, i^* \mathcal{O}_{\mathbb{P}^3}(2))$$

must have a nontrivial kernel, so there is at least one quadric surface Q in \mathbb{P}^3 that contains our curve C . Now quadric surfaces are either double planes, or the union of two planes, or cones, or regular quadrics. (They correspond to quadric forms of rank 1, 2, 3, and 4 respectively. The rank of a quadratic form was defined in Exercise 5.4.J.) But C can't lie in a plane, so Q must be a cone or regular. In particular, Q is irreducible.

Now C can't lie on *two* (distinct) such quadrics, say Q and Q' . Otherwise, as Q and Q' have no common components (they are irreducible and not the same!), $Q \cap Q'$ is a curve (not necessarily reduced or irreducible). By Bézout's Theorem (Exercise 19.6.K), $Q \cap Q'$ is a curve of degree 4. Thus our curve C , being of degree 6, cannot be contained in $Q \cap Q'$. (If you don't see why directly, Exercise 19.6.F might help.)

20.8.1. We next consider cubic surfaces. By (20.2.5.1) again,

$$h^0(C, \omega_C^{\otimes 3}) = \deg \omega_C^{\otimes 3} - g + 1 = 18 - 4 + 1 = 15.$$

Now $\text{Sym}^3 \Gamma(C, \omega_C)$ has dimension $\binom{3+3}{3} = 20$. Thus C lies on at least a 5-dimensional vector space of cubics. Now a 4-dimensional subspace comes from multiplying the quadratic form Q by a linear form ($?w + ?x + ?y + ?z$). Hence there is still one cubic K whose underlying form is not divisible by the quadric form Q (i.e., K doesn't contain Q .) Then K and Q share no component, so $K \cap Q$ is a complete intersection containing C as a closed subscheme. Now $K \cap Q$ and C are both degree 6 (the former by Bézout's Theorem, Exercise 19.6.K, and the latter because C is embedded by a degree 6 line bundle, Exercise 19.6.I). Also, $K \cap Q$ and C both have arithmetic genus 4 (the former by Exercise 19.6.S). These two invariants determine the (linear) Hilbert polynomial, so $K \cap Q$ and C have the same Hilbert polynomial. Hence $C = K \cap Q$ by Exercise 19.6.F.

We now show the converse, and that any regular complete intersection C of a quadric surface with a cubic surface is a canonically embedded genus 4 curve. By Exercise 19.6.S, such a complete intersection has genus 4.

20.8.A. EXERCISE. Show that $\mathcal{O}_C(1)$ has at least 4 sections. (Translation: C doesn't lie in a hyperplane.)

The only degree $2g-2$ invertible sheaf with (at least) g sections is the dualizing sheaf (Exercise 20.2.A), so $\mathcal{O}_C(1) \cong \omega_C$, and C is indeed canonically embedded.

20.8.B. EXERCISE. Give a heuristic argument suggesting that the nonhyperelliptic curves of genus 4 "form a family of dimension 9".

On to genus 5!

20.8.C. EXERCISE. Suppose C is a nonhyperelliptic genus 5 curve. Show that the canonical curve is degree 8 in \mathbb{P}^4 . Show that it lies on a three-dimensional vector space of quadrics (i.e., it lies on 3 linearly independent quadrics). Show that a regular complete intersection of 3 quadrics is a canonical(ly embedded) genus 5 curve.

Unfortunately, not all canonical genus 5 curves are the complete intersection of 3 quadrics in \mathbb{P}^4 . But in the same sense that most genus 3 curves can be described as plane quartics, most canonical genus 5 curves are complete intersections of 3

quadrics, and most genus 5 curves are nonhyperelliptic. The correct way to say this is that there is a dense Zariski-open locus in the moduli space of genus 5 curves consisting of nonhyperelliptic curves whose canonical embedding is cut out by 3 quadrics.

(Those nonhyperelliptic genus 5 canonical curves not cut out by a three-dimensional vector space of quadrics are precisely the trigonal curves, see Exercise 20.7.D. The triplets of points mapping to the same point of \mathbb{P}^1 under the trigonal map turn out to lie on a line in the canonical map. Any quadric vanishing along those 3 points must vanish along the line — basically, any quadratic polynomial with three zeros must be the zero polynomial.)

20.8.D. EXERCISE. Assuming the discussion above, count complete intersections of three quadrics to give a heuristic argument suggesting that the curves of genus 5 “form a family of dimension 12”.

20.8.2. We have now understood curves of genus 3 through 5 by thinking of canonical curves as complete intersections. Sadly our luck has run out.

20.8.E. EXERCISE. Show that if $C \subset \mathbb{P}^{g-1}$ is a canonical curve of genus $g \geq 6$, then C is *not* a complete intersection. (Hint: Bézout’s Theorem, Exercise 19.6.K.)

20.8.3. Some discussion on curves of general genus. However, we still have some data. If \mathcal{M}_g is this ill-defined “moduli space of genus g curves”, we have heuristics to find its dimension for low g . In genus 0, over an algebraically closed field, there is only genus 0 curve (Proposition 20.3.1), so it appears that $\dim \mathcal{M}_0 = 0$. In genus 1, over an algebraically closed field, we will soon see that the elliptic curves are classified by the j -invariant (Exercise 20.10.F), so it appears that $\dim \mathcal{M}_1 = 1$. We have also informally computed $\dim \mathcal{M}_2 = 3$, $\dim \mathcal{M}_3 = 6$, $\dim \mathcal{M}_4 = 9$, $\dim \mathcal{M}_5 = 12$. What is the pattern? In fact in some strong sense it was known by Riemann that $\dim \mathcal{M}_g = 3g - 3$ for $g > 1$. What goes wrong in genus 0 and genus 1? As a clue, recall our insight when discussing Hilbert functions (§19.6) that whenever some function is “eventually polynomial”, we should assume that it “wants to be polynomial”, and there is some better function (usually an Euler characteristic) that *is* polynomial, and that cohomology-vanishing ensures that the original function and the better function “eventually agree”. Making sense of this in the case of \mathcal{M}_g is far beyond the scope of our current discussion, so we will content ourselves by observing the following facts. *Every* regular curve of genus greater than 1 has a finite number of automorphisms — a zero-dimensional automorphism group. *Every* regular curve of genus 1 has a one-dimensional automorphism group (see Question 20.10.6). And the only regular curve of genus 0 has a three-dimensional automorphism group (Exercise 17.4.C). (See Aside 22.5.12 for more discussion.) So notice that for all $g \geq 0$,

$$\dim \mathcal{M}_g - \dim \operatorname{Aut} C_g = 3g - 3$$

where $\operatorname{Aut} C_g$ means the automorphism group of any curve of genus g .

In fact, in the language of stacks (or orbifolds), it makes sense to say that the dimension of the moduli space of (smooth projective geometrically irreducible) genus 0 curves is -3 , and the dimension of the moduli space of genus 1 curves is 0.

20.9 Curves of genus 1

Finally, we come to the very rich case of curves of genus 1. We will present the theory by discussing interesting things about line bundles of steadily increasing degree.

20.9.1. Line bundles of degree 0.

Suppose C is a genus 1 curve. Then $\deg \omega_C = 2g - 2 = 0$ and $h^0(C, \omega_C) = g = 1$ (by Exercise 20.2.A). But the only degree 0 invertible sheaf with a section is the structure sheaf (§20.2.4), so we conclude that $\omega_C \cong \mathcal{O}_C$. (If you know that complex genus 1 curves are of the form \mathbb{C} modulo a lattice, and Miracle 19.5.2 that the sheaf of differentials is dualizing, then you might not be surprised.)

We move on to line bundles of higher degree. Next, note that if $\deg \mathcal{L} > 0$, then Riemann-Roch and Serre duality (20.2.5.1) give

$$h^0(C, \mathcal{L}) = \deg \mathcal{L} - g + 1 = \deg \mathcal{L}.$$

20.9.2. Line bundles of degree 1.

Each degree 1 (k -valued) point q determines a line bundle $\mathcal{O}(q)$, and two distinct points determine two distinct line bundles (as a degree 1 line bundle has only one section, up to scalar multiples). Conversely, any degree 1 line bundle \mathcal{L} is of the form $\mathcal{O}(q)$ (as \mathcal{L} has a section — then just take its divisor of zeros).

Thus we have a canonical bijection between degree 1 line bundles and degree 1 (closed) points. (If k is algebraically closed, as all closed points have residue field k , this means that we have a canonical bijection between degree 1 line bundles and closed points.)

Define an **elliptic curve** to be a genus 1 curve E with a choice of k -valued point p . The choice of this point should always be considered part of the definition of an elliptic curve — “elliptic curve” is not a synonym for “genus 1 curve”. (Note: a genus 1 curve need not have any k -valued points at all! For example, the genus 1 curve “compactifying” $y^2 = -x^4 - 1$ in $\mathbb{A}_{\mathbb{Q}}^2$ has no \mathbb{R} -points, and hence no \mathbb{Q} -points. Of course, if $k = \bar{k}$, then any closed point is k -valued, by the Nullstellensatz 3.2.5.) We will often denote elliptic curves by E rather than C .

If (E, p) is an elliptic curve, then there is a canonical bijection between the set of degree 0 invertible sheaves (up to isomorphism) and the set of degree 1 points of E : simply twist the degree 1 line bundles by $\mathcal{O}(-p)$. Explicitly, the bijection is given by

$$\mathcal{L} \longmapsto \operatorname{div}(\mathcal{L}(p))$$

$$\mathcal{O}(q - p) \longleftarrow q$$

But the degree 0 invertible sheaves form a group (under tensor product), so we have proved:

20.9.3. Proposition (the group law on the degree 1 points of an elliptic curve). — *The above bijection defines an abelian group structure on the degree 1 points of an elliptic curve, where p is the identity.*

From now on, we will identify closed points of E with degree 0 invertible sheaves on E without comment.

For those familiar with the complex analytic picture, this is not surprising: E is isomorphic to the complex numbers modulo a lattice: $E \cong \mathbb{C}/\Lambda$. (We haven't shown this, of course.)

Proposition 20.9.3 is currently just a bijection of sets. Given that E has a much richer structure (it has a generic point, and the structure of a variety), this is a sign that there should be a way of defining some *scheme* $\text{Pic}^0(E)$, and that this should be an isomorphism of schemes. We will soon show (Theorem 20.10.4) that this group structure on the degree 1 points of E comes from a group variety structure on E .

20.9.4. Aside: The Mordell-Weil Theorem, group, and rank. This is a good excuse to mention the *Mordell-Weil Theorem*: for any elliptic curve E over \mathbb{Q} , the \mathbb{Q} -points of E form a *finitely generated* abelian group, often called the *Mordell-Weil group*. By the classification of finitely generated abelian groups (a special case of the classification of finitely generated modules over a principal ideal domain, Remark 13.5.14), the \mathbb{Q} -points are a direct sum of a torsion part, and of a free \mathbb{Z} -module. The rank of the \mathbb{Z} -module is called the *Mordell-Weil rank*.

20.9.5. Line bundles of degree 2 (see Figure 20.4).

Note that $\mathcal{O}_E(2p)$ has 2 sections, so E admits a double cover of \mathbb{P}^1 (Exercise 20.5.B). One of the ramification points is p : one of the sections of $\mathcal{O}_E(2p)$ vanishes at p to order 2, so there is a point of \mathbb{P}^1 whose preimage consists of p (with multiplicity 2). Assume now that $k = \bar{k}$ and $\text{char } k \neq 2$, so we can use the hyperelliptic Riemann-Hurwitz formula (Theorem 20.5.1), which implies that E has 4 ramification points (p and three others). Conversely, given 4 points in \mathbb{P}^1 , there exists a unique double cover branched at those 4 points (Proposition 20.5.2). Thus elliptic curves correspond to 4 distinct points in \mathbb{P}^1 , where one is marked P , up to automorphisms of \mathbb{P}^1 . Equivalently, by placing P at ∞ , elliptic curves correspond to 3 points in \mathbb{A}^1 , up to affine maps $x \mapsto ax + b$.




FIGURE 20.4. Elliptic curves yield double covers of the projective line, branched at four points (§20.9.5)

20.9.A. EXERCISE. Show that the other three ramification points are precisely the (non-identity) 2-torsion points in the group law. (Hint: if one of the points is q , show that $\mathcal{O}(2q) \cong \mathcal{O}(2p)$, but $\mathcal{O}(q)$ is not congruent to $\mathcal{O}(p)$.)

Thus (if $\text{char } k \neq 2$ and $k = \bar{k}$) every elliptic curve has precisely four 2-torsion points. If you are familiar with the complex picture $E \cong \mathbb{C}/\Lambda$, this isn't surprising.

20.9.6. Follow-up remark. An elliptic curve with *full level n structure* is an elliptic curve with an isomorphism of its n -torsion points with $(\mathbb{Z}/n)^2$. (This notion has problems if n is divisible by $\text{char } k$.) Thus an elliptic curve with *full level 2 structure* is the same thing as an elliptic curve with an ordering of the three other branch points in its degree 2 cover description. Thus (if $k = \bar{k}$) these objects are parametrized by the λ -line, which we discuss below.

Follow-up to the follow-up. There is a notion of moduli spaces of elliptic curves with full level n structure. Such moduli spaces are smooth curves (where this is interpreted appropriately — they are stacks), and have smooth compactifications. A *weight k level n modular form* is a section of $\omega_M^{\otimes k}$ where ω_M is the dualizing (or canonical) sheaf of this moduli space (“modular curve”).

20.9.7. The cross-ratio and the j -invariant. (We maintain the assumption that k is algebraically closed and characteristic not 2.) If the three other branch points are temporarily labeled $Q_1, Q_2, Q_3 \in \mathbb{P}^1$, there is a unique automorphism of \mathbb{P}^1 taking P, Q_1, Q_2 to $(\infty, 0, 1)$ respectively (as $\text{Aut } \mathbb{P}^1$ is three-transitive, Exercise 17.4.C). Suppose that Q_3 is taken to some number λ under this map, where necessarily $\lambda \neq 0, 1, \infty$.

The value λ is called the **cross-ratio** of the four points (P, Q_1, Q_2, Q_3) of \mathbb{P}^1 (first defined by W. Clifford, but implicitly known since the time of classical Greece).

20.9.B. EXERCISE. Show that isomorphism classes of four ordered distinct (closed) points on \mathbb{P}^1 (over an algebraically closed field k), up to projective equivalence (automorphisms of \mathbb{P}^1), are classified by the cross-ratio.

We have not defined the notion of *moduli space*, but the previous exercise illustrates the fact that $\mathbb{P}^1 - \{0, 1, \infty\}$ (the image of the cross-ratio map) is the moduli space for four ordered distinct points of \mathbb{P}^1 up to projective equivalence.

Notice:

- If we had instead sent P, Q_2, Q_1 to $(\infty, 0, 1)$, then Q_3 would have been sent to $1 - \lambda$.
- If we had instead sent P, Q_1, Q_3 to $(\infty, 0, 1)$, then Q_2 would have been sent to $1/\lambda$.
- If we had instead sent P, Q_3, Q_1 to $(\infty, 0, 1)$, then Q_2 would have been sent to $1 - 1/\lambda = (\lambda - 1)/\lambda$.
- If we had instead sent P, Q_2, Q_3 to $(\infty, 0, 1)$, then Q_1 would have been sent to $1/(1 - \lambda)$.
- If we had instead sent P, Q_3, Q_2 to $(\infty, 0, 1)$, then Q_1 would have been sent to $1 - 1/(1 - \lambda) = \lambda/(\lambda - 1)$.

Thus these six values (which correspond to S_3 , the symmetric group permuting Q_1, Q_2 , and Q_3) yield the same elliptic curve, and this elliptic curve will (upon choosing an ordering of the other 3 branch points) yield one of these six values.

This is fairly satisfactory already. To check if two elliptic curves $(E, p), (E', p')$ over $k = \bar{k}$ are isomorphic, we write both as double covers of \mathbb{P}^1 ramified at p and p' respectively, then order the remaining branch points, then compute their respective λ 's (say λ and λ' respectively), and see if they are related by one of the

six expressions above:

$$(20.9.7.1) \quad \lambda' = \lambda, 1/\lambda, 1 - \lambda, 1/(1 - \lambda), (\lambda - 1)/\lambda, \text{ or } \lambda/(\lambda - 1).$$

It would be far more convenient if, instead of a “six-valued invariant” λ , there were a single invariant (let’s call it j), such that $j(\lambda) = j(\lambda')$ if and only if one of the equalities of (20.9.7.1) holds. This j -function should presumably be algebraic, so it would give a map j from the λ -line $\mathbb{A}^1 - \{0, 1\}$ to the \mathbb{A}^1 . By the Curve-to-Projective Extension Theorem 17.5.1, this would extend to a morphism $j: \mathbb{P}^1 \rightarrow \mathbb{P}^1$. By Exercise 18.4.D, because this is (for most λ) a 6-to-1 map, the degree of this cover is 6 (or more correctly, at least 6).

We can make this dream more precise as follows. The elliptic curves over k correspond to k -valued points of $\mathbb{P}^1 - \{0, 1, \infty\}$, modulo the action of S_3 on λ given above. Consider the subfield K of $k(\lambda)$ fixed by S_3 . Then the extension $k(\lambda)/K$ is necessarily Galois (see for example [DF, §14.2, Thm. 9]), and degree 6. We are hoping that this subfield is of the form $k(j)$, and if so, we would obtain the j -map $\mathbb{P}^1 \rightarrow \mathbb{P}^1$ as described above. One could show that K is finitely generated over k , and then invoke Lüroth’s Theorem, which we will soon prove in Example 22.7.7; but we won’t need this.

Instead, we will just hunt for such a j . Note that λ should satisfy a sextic polynomial over $k(j)$ (or more precisely given what we know right now, a polynomial of degree *at least* six), as for each j -invariant, there are six values of λ in general.

As you are undoubtedly aware, there *is* such a j -invariant. Here is the formula for the j -invariant that everyone uses:

$$(20.9.7.2) \quad j = 2^8 \frac{(\lambda^2 - \lambda + 1)^3}{\lambda^2(\lambda - 1)^2}.$$

You can readily check that $j(\lambda) = j(1/\lambda) = j(1 - \lambda) = \dots$, and that as j (in lowest terms) has a degree 6 numerator and degree < 6 denominator, j indeeds determines a degree 6 map from \mathbb{P}^1 (with coordinate λ) to \mathbb{P}^1 (with coordinate j). But this complicated-looking formula begs the question: where did this formula come from? How did someone think of it? We will largely answer this, but we will ignore the 2^8 . This, as you might imagine, arises from characteristic 2 issues, and in order to invoke the results of §20.5 we have been assuming $\text{char } k \neq 2$. (To see why the 2^8 is forced upon us by characteristic 2, see [De, p. 64]. From a different, complex-analytic, point of view, the 2^8 comes from the fact that the q -expansion of j begins $j = q^{-1} + 744 + \dots$: the q^{-1} term has coefficient 1; see [Se5, p. 90, Rem. 2]. These seemingly different explanations are related by the theory of the Tate curve, see [Si2, Thm. V.3.1].)

Rather than using the formula handed to us, let’s try to guess what j is. We won’t expect to get the same formula as (20.9.7.2), but our answer should differ by an automorphism of the j -line (\mathbb{P}^1) — we will get $j' = (aj + b)/(cj + d)$ for some a, b, c, d (Exercise 17.4.B).

We are looking for some $j'(\lambda)$ such that $j'(\lambda) = j'(1/\lambda) = \dots$. Hence we want some expression in λ that is invariant under this S_3 -action. A first possibility would be to take the product of the six numbers

$$\lambda \cdot (1 - \lambda) \cdot \frac{1}{\lambda} \cdot \frac{\lambda - 1}{\lambda} \cdot \frac{1}{1 - \lambda} \cdot \frac{\lambda}{\lambda - 1}$$

This is silly, as the product is obviously 1.

A better idea is to add them all together:

$$\lambda + (1 - \lambda) + \frac{1}{\lambda} + \frac{\lambda - 1}{\lambda} + \frac{1}{1 - \lambda} + \frac{\lambda}{\lambda - 1}$$

This also doesn't work, as they add to 3 — the six terms come in pairs adding to 1.

(Another reason you might realize this can't work: if you look at the sum, you will realize that you will get something of the form “degree at most 3” divided by “degree at most 2”. Then if $j' = p(\lambda)/q(\lambda)$, then λ is a root of a cubic over j . But we said that λ should satisfy a sextic over j' . The only way we avoid a contradiction is if $j' \in k$.)

But you will undoubtedly have another idea immediately. One good idea is to take the second symmetric function in the six roots. An equivalent one that is easier to do by hand is to add up the squares of the six terms. Even before doing the calculation, we can see that this will work: it will clearly produce a fraction whose numerator and denominator have degree at most 6, and it is not constant, as when λ is some fixed small number (say $1/2$), the sum of squares is some small real number, while when λ is a large real number, the sum of squares will have to be some large real number (different from the value when $\lambda = 1/2$).

When you add up the squares by hand (which is not hard), you will get

$$j' = \frac{2\lambda^6 - 6\lambda^5 + 9\lambda^4 - 8\lambda^3 + 9\lambda^2 - 6\lambda + 2}{\lambda^2(\lambda - 1)^2}.$$

Indeed $k(j) \cong k(j')$: you can check (again by hand) that

$$2j/2^8 = \frac{2\lambda^6 - 6\lambda^5 + 12\lambda^4 - 14\lambda^3 + 12\lambda^2 - 6\lambda + 2}{\lambda^2(\lambda - 1)^2}.$$

Thus $2j/2^8 - j' = 3$.

20.9.8. Line bundles of degree 3.

In the discussion of degree 2 line bundles 20.9.5, we assumed $\text{char } k \neq 2$ and $k = \bar{k}$, in order to invoke the Riemann-Hurwitz formula. In this section, we will start with no assumptions, and add them as we need them. In this way, you will see what partial results hold with weaker assumptions.

Consider the degree 3 invertible sheaf $\mathcal{O}_E(3p)$. By Riemann-Roch (20.2.5.1), $h^0(E, \mathcal{O}_E(3p)) = \deg(3p) - g + 1 = 3$. As $\deg \mathcal{O}_E(3p) > 2g$, this gives a closed embedding (Remark 20.2.11 and Exercise 20.2.D). Thus we have a closed embedding $E \hookrightarrow \mathbb{P}_k^2$ as a cubic curve. Moreover, there is a line in \mathbb{P}_k^2 meeting E at point p with multiplicity 3, corresponding to the section of $\mathcal{O}(3p)$ vanishing precisely at p with multiplicity 3. (A line in the plane meeting a smooth curve with multiplicity at least 2 is a *tangent line*, see Definition 13.3.2. A line in the plane meeting a smooth curve with multiplicity at least 3 is said to be a **flex line**, and that point is a **flex point** or **flex** of the curve.) (See Figure 20.5.)




FIGURE 20.5. Finding the Weierstrass cubic incarnation of an elliptic curve

Choose projective coordinates on \mathbb{P}_k^2 so that p maps to $[0, 1, 0]$, and the flex line is the line at infinity $z = 0$. Then the cubic is of the following form:

$$\begin{aligned}
 & ? x^3 + ? x^2 y + ? x y^2 + ? y^3 \\
 & + ? x^2 z + ? x y z + ? y^2 z = 0 \\
 & + ? x z^2 + ? y z^2 \\
 & + ? z^3
 \end{aligned}$$

The coefficient of x^3 is not 0 (or else this cubic is divisible by z). Dividing the entire equation by this coefficient, we can assume that the coefficient of x^3 is 1. The coefficient of $y^2 z$ is not 0 either (or else this cubic is singular at $x = z = 0$). We can scale z (i.e., replace z by a suitable multiple) so that the coefficient of $y^2 z$ is -1 . If the characteristic of k is not 2, then we can then replace y by $y + ?x + ?z$ so that the coefficients of xyz and yz^2 are 0, and if the characteristic of k is not 3, we can replace x by $x + ?z$ so that the coefficient of $x^2 z$ is also 0. In conclusion, if $\text{char } k \neq 2, 3$, the elliptic curve may be written as

$$(20.9.8.1) \quad y^2 z = x^3 + axz^2 + bz^3.$$

This is called the **Weierstrass normal form** of the curve.

20.9.9. From the Weierstrass cubic to the double cover. We see the hyperelliptic (double cover) description of the curve by setting $z = 1$, or more precisely, by working in the distinguished open set $z \neq 0$ and using inhomogeneous coordinates. Here is the geometric explanation of why the double cover description is visible in the cubic description. Project the cubic from $p = [0, 1, 0]$ (see Figure 20.6). This is a

map $E - p \rightarrow \mathbb{P}^1$ (given by $[x, y, z] \mapsto [x, z]$), and is basically the vertical projection of the cubic to the x -axis. (Figure f:Wdc may help you visualize this.) By the Curve-to-Projective Extension Theorem 17.5.1, the morphism $E - p \rightarrow \mathbb{P}^1$ extends over p . If $\mathcal{L} := \mathcal{O}(3p)$ is the line bundle giving the morphism $E \hookrightarrow \mathbb{P}^2$ (via Theorem 17.4.1 describing maps to projective space in terms of line bundles), then the two sections x and z giving the map to \mathbb{P}^1 vanish at p to order 1 and 3 respectively. To “resolve” the rational map into an honest morphism, we interpret x and z as sections of $\mathcal{L}(-p) = \mathcal{O}(2p)$, and now they generate a base-point-free linear series, and thus a morphism to \mathbb{P}^1 . (You may be able to interpret this as implementing the proof of Theorem 17.4.1 in a different language.) A similar idea, applied to quartics rather than cubics, was used in Exercise 20.7.D.




FIGURE 20.6. From the Weierstrass cubic to the double cover interpretation of an elliptic curve (cf. Figure 20.5)

As a consequence, with a little sweat, we can compute the j -invariant as a function of a and b :

$$j(a, b) = \frac{2^8 3^3 a^3}{4a^3 + 27b^2}.$$

20.9.C. EXERCISE. Show that the flexes of the cubic are the 3-torsion points in the group E . (“Flex” was defined in §20.9.8: it is a point where the tangent line meets the curve with multiplicity at least 3 at that point. In fact, if k is algebraically closed and $\text{char } k \neq 3$, there are nine of them. This won’t be surprising if you are familiar with the complex story, $E \cong \mathbb{C}/\Lambda$.)

20.9.10. The group law, geometrically.

The group law has a beautiful classical description in terms of the Weierstrass form. Consider Figure 20.7. In the Weierstrass coordinates, the origin p is the only point of E meeting the line at infinity ($z = 0$); in fact the line at infinity corresponds to the tautological section of $\mathcal{O}(3p)$. If a line meets E at three points p_1, p_2, p_3 , then

$$\mathcal{O}(p_1 + p_2 + p_3) \cong \mathcal{O}(3p)$$

from which (in the group law) $p_1 + p_2 + p_3 = 0$.

Hence to find the inverse of a point s , we consider the intersection of E with the line \overline{sp} ; $-s$ is the third point of intersection. To find the sum of two points q and r , we consider the intersection of E with the line \overline{qr} , and call the third point s . We then compute $-s$ by connecting s to p , obtaining $q + r$.




FIGURE 20.7. The group law on the elliptic curve, geometrically

We could give this description of a group law on a cubic curve in Weierstrass normal form to anyone familiar with the notion of projective space, and the notion of a group, but we would then have to prove that the construction we are giving indeed defines a group. In particular, we would have to prove associativity, which is not a priori clear. But in this case, we have already established that the degree 1 points form a group, by giving a bijection to $\text{Pic}^0 E$, and we are merely interpreting the group law on $\text{Pic}^0 E$.

Note that this description works even in characteristic 2 and 3; we do not need the cubic to be in Weierstrass normal form, and we need only that $\mathcal{O}(3p)$ gives a closed embedding into \mathbb{P}^2 .

20.9.11. Degree 4 line bundles. You have probably forgotten that we began by studying line bundles degree by degree. The story doesn't stop in degree 3. In the same way that we showed that a canonically embedded nonhyperelliptic curve of genus 4 is the complete intersection in \mathbb{P}^3_k of a quadric surface and a cubic surface (§20.8), we can show the following.

20.9.D. EXERCISE. Show that the complete linear series for $\mathcal{O}(4p)$ embeds E in \mathbb{P}^3 as the complete intersection of two quadrics. Hint: Show the image of E is

contained in at least 2 linearly independent quadrics. Show that neither can be reducible, so they share no components. Use Bézout’s Theorem (Exercise 19.6.K) and Exercise 19.6.F, as in §20.8.1.

The beautiful structure doesn’t stop with degree 4, but it gets more complicated. For example, the degree 5 embedding is not a complete intersection (of hypersurfaces), but is the complete intersection of $G(2, 5)$ under its Plücker embedding with five hyperplanes (or perhaps better, a codimension 5 linear space). In seemingly different terminology, its equations are 4×4 Pfaffians of a general 5×5 skew-symmetric matrix of linear forms, although I won’t say what this means.

20.10 Elliptic curves are group varieties

Ever since 1949, I considered the construction of an algebraic theory of the Picard variety as the task of greatest urgency in abstract algebraic geometry.

— A. Weil [Weil, p. 537]

We initially described the group law on the degree 1 points of an algebraic curve in a rather abstract way. From that definition, it was not clear that over \mathbb{C} the group operations (addition, inverse) are continuous. But the explicit description in terms of the Weierstrass cubic makes this clear. In fact we can observe even more: addition and inverse are algebraic in general. Better yet, elliptic curves are group varieties. (Thus they are abelian varieties, see Definition 11.5.11.)

20.10.1. (This is a clue that $\text{Pic}^0(E)$ really wants to be a scheme, and not just a group. Once the notion of “moduli space of line bundles on a variety” is made precise, this can be shown, see for example [FGIKNV, Part 5].)

We begin with the “group inverse” morphism, as a warm-up.

20.10.2. Proposition. — *If $\text{char } k \neq 2, 3$, there is a morphism of k -varieties $E \rightarrow E$ sending a (degree 1) point to its inverse, and this construction behaves well under field extension of k .*

In other words, the “inverse map” in the group law actually arises from a morphism of schemes — it isn’t just a set map. (You are welcome to think through the two remaining characteristics, and to see that essentially the same proof applies. But the proof of Theorem 20.10.4 will give you a better sense of how to proceed.)

Proof. Consider the map (the hyperelliptic involution) $y \mapsto -y$ of the Weierstrass normal form. \square

The algebraic description of addition would be a big mess if we were to write it down. We will be able to show algebraicity by a trick — not by writing it down explicitly, but by thinking through how we *could* write it down explicitly. The main part of the trick is the following proposition. We give it in some generality just because it can be useful, but you may prefer to assume that $k = \bar{k}$ and C is a regular cubic.

20.10.3. Proposition. — Suppose $C \subset \mathbb{P}_k^2$ is a geometrically integral cubic curve (so in particular C contains no lines). Let C^{reg} be the regular points of C . There is a unique morphism $t: C^{\text{reg}} \times C^{\text{reg}} \rightarrow C^{\text{reg}}$ such that

- (a) if p and q are distinct regular k -valued points of C , then $t(p, q)$ is obtained by intersecting the line \overline{pq} with C , and taking the third “residual” point of intersection with C . More precisely, \overline{pq} will meet C at three points with multiplicity (Exercise 9.3.D), including p and q ; $t(p, q)$ is the third point; and
- (b) this property remains true after extension to \bar{k} .

Furthermore, if p is a k -valued point of C^{reg} , then $t(p, p)$ is where the tangent line ℓ to C at p meets C again. More precisely, ℓ will meet C at three points with multiplicity, which includes p with multiplicity 2; $t(p, p)$ is the third point.

We will need property (b) because C may have few enough k -valued points (perhaps none!) that the morphism t can not be determined by its behavior on them. In the course of the proof, we will see that (b) can be extended to “this property remains true after any field extension of k ”.

Proof. We first show (in this paragraph) that if p and q are distinct regular points, then the third point r of intersection of \overline{pq} with C is also regular. If $r = p$ or $r = q$, we are done. Otherwise, the cubic obtained by restricting C to \overline{pq} has three distinct (hence reduced, i.e., multiplicity 1) roots, p , q , and r . Thus $C \cap \overline{pq}$ is regular at r , so r is a regular point of C by the slicing criterion for regularity, Exercise 13.2.C.

We now assume that $k = \bar{k}$, and leave the general case to the end. Fix p , q , and r , where $p \neq q$, and r is the “third” point of intersection of \overline{pq} with C . We will describe a morphism $t_{p,q}$ in an open neighborhood of $(p, q) \in C^{\text{reg}} \times C^{\text{reg}}$. By Exercise 11.4.B, showing that morphisms of varieties over \bar{k} are determined by their behavior on closed (\bar{k} -valued) points, these morphisms glue together (uniquely) to give a morphism t , completing the proof in the case $k = \bar{k}$.

Choose projective coordinates on \mathbb{P}^2 in such a way that $U_0 \cong \text{Spec } k[x_1, x_2]$ contains p , q , and r , and the line \overline{pq} is not “vertical”. More precisely, in $\text{Spec } k[x_1, x_2]$, say $p = (p_1, p_2)$ (in terms of “classical coordinates” — more pedantically, $p = [(x_1 - p_1, x_2 - p_2)]$), $q = (q_1, q_2)$, $r = (r_1, r_2)$, and $p_1 \neq q_1$. In these coordinates, the curve C is cut out by some cubic, which we also sloppily denote C : $C(x_1, x_2) = 0$.

Now suppose $P = (P_1, P_2)$ and $Q = (Q_1, Q_2)$ are two points of $C \cap U_0$ (not necessarily our p and q). We attempt to compute the third point of intersection of \overline{PQ} with C , in a way that works on an open subset of $C \times C$ that includes (p, q) . To do this explicitly requires ugly high school algebra, but because we know how it looks, we will be able to avoid dealing with any details!

The line \overline{PQ} is given by $x_2 = mx_1 + b$, where $m = \frac{P_2 - Q_2}{P_1 - Q_1}$ and $b = P_2 - mP_1$ are both rational functions of P and Q . Then m and b are defined for all P and Q such that $P_1 \neq Q_1$ (and hence for an open neighborhood of (p, q) , as $p_1 \neq q_1$, and as $P_1 \neq Q_1$ is an open condition).

Now we solve for $C \cap \overline{PQ}$, by substituting $x_2 = mx_1 + b$ into C , to get $C(x_1, mx_1 + b)$. This is a cubic in x_1 , say

$$\gamma(x_1) = Ax_1^3 + Bx_1^2 + Cx_1 + D = 0.$$

The coefficients of γ are rational functions of P_1, P_2, Q_1 , and Q_2 . The cubic γ has 3 roots (with multiplicity) so long as $A \neq 0$, which is a Zariski-open condition on m

and b , and hence a Zariski-open condition on P_1, P_2, Q_1, Q_2 . As $P, Q \in C \cap \overline{PQ} \cap U_0$, P_1 and Q_1 are two of the roots of $\gamma(x_1) = 0$. The sum of the roots of $\gamma(x_1) = 0$ is $-B/A$ (by Viète's formula), so the third root of γ is $R_1 := -B/A - P_1 - Q_1$. Thus if we take $R_2 = mR_1 + b$, we have found the third point of intersection of \overline{PQ} with C (which happily lies in U_0). We have thus described a morphism from the open subset of $(C^{\text{reg}} \cap U_0) \times (C^{\text{reg}} \cap U_0)$, containing (p, q) , that does what we want. (Precisely, the open subset is defined by $A \neq 0$, which can be explicitly unwound.) We have thus completed the proof of Proposition 20.10.3 (except for the last paragraph) for $k = \bar{k}$. (Those who believe they are interested only in algebraically closed fields can skip ahead.)

We extend this to Proposition 20.10.3 for every field k except \mathbb{F}_2 . Suppose $U_0^{[x_1, x_2]} = \text{Spec } k[x_1, x_2]$ is any affine open subset of \mathbb{P}_k^2 , along with choice of coordinates. (The awkward superscript " $[x_1, x_2]$ " is there to emphasize that the particular coordinates are used in the construction.) Then the construction above gives a morphism *defined over* k from an open subset of $(C^{\text{reg}} \cap U_0^{[x_1, x_2]}) \times (C^{\text{reg}} \cap U_0^{[x_1, x_2]})$ (note that all of the hypothetical algebra was done over k), that sends P and Q to the third point of intersection of \overline{PQ} with C . Note that this construction commutes with any field extension, as the construction is insensitive to the field we are working over. Thus after base change to the algebraic closure, the map also has the property that it takes as input two points, and spits out the third point of intersection of the line with the cubic. Furthermore, all of these maps (as $U_0^{[x_1, x_2]}$ varies over all complements U_0 of lines "with k -coefficients", and choices of coordinates on U_0) can be glued together: they agree on their pairwise overlaps (as after base change to \bar{k} they are the same, by our previous discussion, and two maps that are the same after base change to \bar{k} were the same to begin with by Exercise 10.2.J), and this is what is required to glue them together (Exercise 7.2.A).

We can geometrically interpret the open subset $(C^{\text{reg}} \cap U_0^{[x_1, x_2]}) \times (C^{\text{reg}} \cap U_0^{[x_1, x_2]})$ by examining the construction: it is defined in the locus $\{P = (P_1, P_2), Q = (Q_1, Q_2)\}$ where (i) $P_1 \neq Q_1$, and (ii) the third point of intersection R of \overline{PQ} with C also lies in U_0 .

So which points (P, Q) of $C^{\text{reg}} \times C^{\text{reg}}$ are missed? Condition (i) isn't important; if (P, Q) satisfies (ii) but not (i), we can swap the roles of x_1 and x_2 , and (P, Q) will then satisfy (i). The only way (P, Q) can not be covered by one of these open sets is if there is *no* U_0 (a complement of a line defined over k) that includes P, Q , and R .

20.10.A. EXERCISE. Use $|k| > 2$ to show that there is a linear form on \mathbb{P}^2 with coefficients in k that misses P, Q , and R . (This is sadly *not* true if $k = \mathbb{F}_2$ — do you see why?)

20.10.B. EXERCISE. Prove the last statement of Proposition 20.10.3.

20.10.C. ★★ UNIMPORTANT EXERCISE. Complete the proof by dealing with the case $k = \mathbb{F}_2$. Hint: first produce the morphism t over \mathbb{F}_4 . The goal is then to show that this t is really "defined over" \mathbb{F}_2 ("descends to" \mathbb{F}_2). The morphism t is initially described locally by considering the complement of a line defined over \mathbb{F}_4 (and then letting the line vary). Instead, look at the map by looking at the complement of a line and its "conjugate". The complement of the line and

its conjugate is an affine \mathbb{F}_2 -variety. The partially-defined map t on this affine variety is a priori defined over \mathbb{F}_4 , and is preserved by conjugation. Show that this partially defined map is “really” defined over \mathbb{F}_2 . (If you figure out what all of this means, you will have an important initial insight into the theory of “descent”.) \square

We can now use this to define the group variety structure on E .

20.10.4. Theorem. — Suppose (E, p) is an elliptic curve (a regular genus 1 curve over k , with a k -valued point p). Take the Weierstrass embedding of E in \mathbb{P}_k^2 , via the complete linear series $|\mathcal{O}_E(3p)|$. Define the k -morphism $e: \operatorname{Spec} k \rightarrow E$ by sending $\operatorname{Spec} k$ to p . Define the k -morphism $i: E \rightarrow E$ via $q \mapsto t(p, q)$, or more precisely, as the composition

$$E \xrightarrow{(id, e)} E \times E \xrightarrow{t} E.$$

Define the k -morphism $m: E \times E \rightarrow E$ via $(q, r) \mapsto t(p, t(q, r))$. Then (E, e, i, m) is a group variety over k .

By the construction of t , all of these morphisms “commute with arbitrary base extension”.

Proof. We need to check that various pairs of morphisms described in axioms (i)–(iii) of §7.6.4 are equal. For example, in axiom (iii), we need to show that $m \circ (i, id) = m \circ (id, i)$; all of the axioms are clearly of this sort.

Assume first that $k = \bar{k}$. Then each of these pairs of morphisms agree as maps of \bar{k} -points: $\operatorname{Pic} E$ is a group, and under the bijection between $\operatorname{Pic} E$ and E of Proposition 20.9.3, the group operations translate into the maps described in the statement of Theorem 20.10.4 by the discussion of §20.9.10.

But morphisms of \bar{k} -varieties are determined by their maps on the level of \bar{k} -points (Exercise 11.4.B), so each of these pairs of morphisms are the same.

For general k , we note that from the \bar{k} case, these morphisms agree after base change to the algebraic closure. Then by Exercise 10.2.J, they must agree to begin with. \square

20.10.5. Features of this construction. The most common derivation of the properties of an elliptic curve is to describe it as a cubic, and then to describe addition using the explicit construction with lines. Then one has to work to prove that the multiplication described is associative.

Instead, we started with something that was patently a group (the degree 0 line bundles). We interpreted the maps used in the definition of the group (addition and inverse) geometrically using our cubic interpretation of elliptic curves. This allowed us to see that these maps were algebraic.

As a bonus, we see that in some vague sense, the Picard group of an elliptic curve wants to be an algebraic variety.

20.10.D. EXERCISE. Suppose p and q are k -points of a genus 1 curve E . Show that there is an automorphism of E sending p to q .

20.10.E. EXERCISE. Suppose (E, p) is an elliptic curve over an algebraically closed field k of characteristic not 2 or 3. Show that the automorphism group of (E, p) is isomorphic to $\mathbb{Z}/2$, $\mathbb{Z}/4$, or $\mathbb{Z}/6$. (An automorphism of an elliptic curve (E, p)

over $k = \bar{k}$ is an automorphism of E fixing p scheme-theoretically, or equivalently, fixing the k -valued point p by Exercise 11.4.B.) Hint: reduce to the question of automorphisms of \mathbb{P}^1 fixing a point ∞ and a set of distinct three points $\{p_1, p_2, p_3\} \in \mathbb{P}^1 \setminus \{\infty\}$. (The algebraic closure of k is not essential, so feel free to remove this hypothesis, using Exercise 10.2.J. What happens if the characteristic is 3?)

20.10.F. EXERCISE. Explain why genus 1 curves over an algebraically closed field of characteristic not 2 are classified (up to isomorphism) by j -invariant. (Caution: note that the problem says “genus 1”, not “elliptic”.)

20.10.G. EXERCISE. Give (with proof) two elliptic curves over \mathbb{Q} with the same j -invariant that are not isomorphic. (Hint: §20.5.3.)

20.10.6. Vague question. What are the possible automorphism groups of a genus 1 curve over an algebraically closed k of characteristic not 2? You should be able to convince yourself that the group has “dimension 1”.

20.10.H. IMPORTANT EXERCISE: A DEGENERATE ELLIPTIC CURVE. Consider the genus 1 curve $C \subset \mathbb{P}_k^2$ given by $y^2z = x^3 + x^2z$, with the point $p = [0, 1, 0]$. Emulate the above argument to show that $C \setminus \{[0, 0, 1]\}$ is a group variety (where the operation is that given by tensoring line bundles). Show that it is isomorphic to \mathbb{G}_m (the multiplicative group scheme $\text{Spec } k[t, t^{-1}]$, see Exercise 7.6.E) with coordinate $t = y/x$, by showing an isomorphism of schemes, and showing that multiplication and inverse in both group varieties agree under this isomorphism.

20.10.I. EXERCISE: AN EVEN MORE DEGENERATE ELLIPTIC CURVE. Consider the genus 1 curve $C \subset \mathbb{P}_k^2$ given by $y^2z = x^3$, with the point $p = [0, 1, 0]$. Emulate the above argument to show that $C \setminus \{[0, 0, 1]\}$ is a group variety. Show that it is isomorphic to \mathbb{A}^1 (with additive group structure) with coordinate $t = x/y$, by showing an isomorphism of schemes, and showing that multiplication/addition and inverse in both group varieties agree under this isomorphism.

20.10.7. ★ Towards Poncelet’s Porism.

These ideas lead to a beautiful classical fact, Poncelet’s Porism. (A *porism* is an archaic name for a type of mathematical result.)




FIGURE 20.8. Poncelet’s Porism 20.10.8

20.10.8. Poncelet's Porism. — Suppose C and D are two ellipses in \mathbb{R}^2 , with C containing D . Choose any point p_0 on C . Choose one of the two tangents ℓ_1 from p_0 to D . Then ℓ_1 meets C at two points in total: p_0 and another point p_1 . From p_1 , there are two tangents to D , ℓ_1 and another line ℓ_2 . The line ℓ_2 meets C at some other point p_2 . Continue this to get a sequence of points p_0, p_1, p_2, \dots . Suppose this sequence starting with p_0 is periodic, i.e., $p_0 = p_n$ for some n . Then it is periodic with any starting point $p \in C$ (see Figure 20.8).

It is possible to prove Poncelet's Porism in an elementary manner, but a proof involving elliptic curves is quite beautiful, and gives connections to more sophisticated ideas. Rather than proving Poncelet's Porism, we discuss some related facts.

20.10.J. EXERCISE. Suppose E is a smooth degree 3 projective plane curve over an algebraically closed field, $q, r \in E$, and n is a positive integer. Let $F: E \rightarrow E$ be the morphism $x \mapsto t(r, t(q, x))$. Suppose that $F^n(p) = p$ for some closed point $p \in E$. Show that $F^n(x) = x$ for all points $x \in E$.

You should feel the urge to improve this result. (Can you show that if $F^n(p) = p$ for some $p \in E$, then F^n is the identity morphism on E ? Can you extend this to the case where the base field is not algebraically closed?)

20.10.K. EXERCISE. Suppose E is an ellipse in \mathbb{R}^2 . For each point $(x, y) \in E$, there is another point $G(x, y) := (x, y')$ (possibly the same), with the same first coordinate, corresponding to the "other" intersection of the vertical line through (x, y) with E . Similarly, there is another point $H(x, y) := (x', y)$, with the same second coordinate. Let $F = G \circ H$. Show that if $F^{17}(p) = p$ for some point $p \in E$, then $F^{17}(x) = x$ for all points $x \in E$. (See Figure 20.9, where 17 is replaced by 3.) Hint: this might be best done not using any fancy methods from algebraic geometry.

Again, you should wish to improve this. To what extent is this dependent on the real numbers? What if the ellipse was instead a different conic section? (And what does this have to do with elliptic curves? Why is the conic called E ?)

20.10.L. ★★ EXERCISE. Let k be a field, algebraically closed purely for convenience. Suppose Q is a smooth quadric surface in \mathbb{P}_k^3 , and K is a cone (a rank 3 quadric) in \mathbb{P}_k^3 , such that $E = Q \cap K$ is a smooth curve.

- Show that E has genus 1.
- If ℓ is a line on Q , show that E meets ℓ at two points (with multiplicity).
- Fix one of the rulings (family of lines) of Q . Show that there exists a morphism $G: E \rightarrow E$ that takes a closed point p of E , and sends it to the other point of E on the line ℓ in the ruling containing p .
- Define $H: E \rightarrow E$ similarly, using the other ruling of Q . Fix a positive integer n . Let $F = G \circ H$. Show that if $F^n(p) = p$ for one closed point $p \in E$, then $F^n(x) = x$ for all $x \in E$.

We now connect Exercise 20.10.L to Poncelet's Porism 20.10.8. Let v be the cone point of K . Then $v \notin Q$, as otherwise v would be a singular point of $Q \cap K$ (Do you see why?). It is a fact that projection from v gives a morphism $Q \rightarrow \mathbb{P}_k^2$ that is branched over a smooth conic D . (You may be able to make this precise, and prove it, after reading Chapter 22.) The lines on Q project to tangent lines to D in \mathbb{P}_k^2 . The




FIGURE 20.9. Exercise 20.10.K (with 17 replaced by 3)

projection from v contracts $K \setminus \{v\}$ to a conic C . You may be able to interpret the statement of Poncelet's Porism 20.10.8 (for this C and D) in terms of Exercise 20.10.L. (Caution: the field is not algebraically closed in Poncelet's Porism 20.10.8, and is algebraically closed in Exercise 20.10.L.) You may even be able to take an arbitrary C and D as in Poncelet's Porism 20.10.8, and reverse-engineer K and Q as in Exercise 20.10.L, thereby proving Poncelet's Porism (and discovering why there is an elliptic curve hidden in its statement).

20.11 Counterexamples and pathologies using elliptic curves

We now give some fun counterexamples using our understanding of elliptic curves. The main extra juice elliptic curves give us comes from the fact that elliptic curves are the simplest varieties with "continuous Picard groups".

20.11.1. An example of a scheme that is factorial, but such that no affine open neighborhood of any point has ring that is a unique factorization domain.

Suppose E is an elliptic curve over \mathbb{C} (or some other uncountable algebraically closed field). Consider $p \in E$. The local ring $\mathcal{O}_{E,p}$ is a discrete valuation ring and hence a unique factorization domain (Theorem 13.5.8). Then any affine open neighborhood of E is of the form $E - q_1 - \cdots - q_n$ for some points q_1, \dots, q_n . I claim that its Picard group is nontrivial. Recall the exact sequence:

$$\mathbb{Z}^{\oplus n} \xrightarrow{(a_1, \dots, a_n) \mapsto a_1 q_1 + \cdots + a_n q_n} \text{Pic } E \longrightarrow \text{Pic}(E - q_1 - \cdots - q_n) \longrightarrow 0.$$

But the group on the left is countable, and the group in the middle is uncountable, so the group on the right, $\text{Pic}(E - q_1 - \cdots - q_n)$, is nonzero. We have shown that every nonempty open subset of E has nonzero line bundles, as promised in Remark 15.

As $E - q_1 - \cdots - q_n$ is affine, by Exercise 15.2.T the corresponding ring is not a unique factorization domain. To summarize: complex elliptic curves are factorial, but no affine open subset has a ring that has unique factorization, as promised in §5.4.4.

20.11.A. EXERCISE. The above argument shows that over an uncountable field, $\text{Pic } E$ is not a finitely generated group. Show that even over the countable field \mathbb{Q} , $\text{Pic } E$ is not a finitely generated group, as follows. If the elliptic curve E is generated by q_1, \dots, q_n , then there is a finite extension of fields K of \mathbb{Q} over which all q_i are defined — the compositum of the residue fields of the q_i . (Careful — do you see what it means for the q_i to be “defined” over a smaller field than \mathbb{Q} ?) Show that any point in the subgroup of E generated by the q_i must also be defined over K . Show that E has a point not defined over K . Use this to show that $\text{Pic } E$ is not finitely generated. (The same argument works with \mathbb{Q} replaced by $\overline{\mathbb{F}}_p$.)

20.11.2. Remark. In contrast to the above discussion, over \mathbb{Q} , the Mordell-Weil Theorem states that $\text{Pic } E$ is finitely generated (see Aside 20.9.4).

20.11.3. ★★ A complex surface with infinitely many algebraic structures (for those with sufficient complex geometric background).

As remarked in Exercise 7.3.16, a complex manifold may have many algebraic structures. The following example of M. Kim gives an example with *infinitely* many algebraic structures. Suppose E is an elliptic curve over \mathbb{C} with origin p , and let \mathcal{L} be a nontrivial line bundle on $E - p$. Then \mathcal{L} is *analytically* trivial because $E - p$ is a Stein space, so the analytification of the total space is independent of the choice of \mathcal{L} . However, you know enough to show (with work) that there are infinitely many pairwise (algebraically) nonisomorphic \mathcal{L} , and also that their total spaces are likewise pairwise (algebraically) nonisomorphic. This gives a complex surface with infinitely many algebraic structures (indeed a continuum of them). See [MO68421] for more details.

20.11.4. Counterexamples using a non-torsion point.

We next give a number of counterexamples using the existence of a non-torsion point of a complex elliptic curve. We first show the existence of such a point.

20.11.5. We have a “multiplication by n ” map $[n]: E \rightarrow E$, which sends p to np . If $n = 0$, the map $[n]$ has degree 0 (even though $[n]$ isn’t dominant, degree is still defined, see Definition 12.2.2). If $n = 1$, the map $[n]$ has degree 1. Given the complex picture of a torus, you might not be surprised that the degree of $[n]$ is n^2 . If $n = 2$, we have almost shown that it has degree 4, as we have checked that there are precisely 4 points q such that $2p = 2q$ (Exercise 20.9.A). All that really shows (using Exercise 18.4.D(b)) is that the degree is at least 4. (We could check by hand that the degree is 4 is we really wanted to.)

20.11.6. Proposition. — Suppose E is an elliptic curve over a field k of characteristic not 2. For each $n > 0$, the “multiplication by n ” morphism $[n]$ has positive degree, so there are only a finite number of n -torsion points.

Proof. We may assume $k = \overline{k}$, as the degree of a map of curves is independent of field extension.

We prove the result by induction; it is true for $n = 1$ and $n = 2$.

If n is odd ($2k + 1$, say), then assume otherwise that $nq = 0$ for all closed points q . Let r be a non-trivial 2-torsion point, so $2r = 0$. But $nr = 0$ as well, so $r = (n - 2k)r = 0$, contradicting $r \neq 0$.

If n is even, then $[n] = [2] \circ [n/2]$ (degree is multiplicative under composition of rational maps, §12.2.2), and by our inductive hypothesis both $[2]$ and $[n/2]$ have positive degree. \square

In particular, the total number of torsion points on E is countable. If k is an uncountable field, then E has an uncountable number of closed points (consider an open subset of the curve as $y^2 = x^3 + ax + b$; there are uncountably many choices for x , and each of them has 1 or 2 choices for y). Thus we have the following.

20.11.7. Corollary. — *If E is a curve over an uncountable algebraically closed field of characteristic not 2 (e.g., \mathbb{C}), then E has a non-torsion point.*

Proof. For each n , there are only finitely many n -torsion points. Thus there are (at most) countably many torsion points. The curve E has uncountably many closed points. (One argument for this: take a double cover $\pi: E \rightarrow \mathbb{P}^1$. Then \mathbb{P}^1 has uncountably many closed points, and π is surjective on closed points.) \square

20.11.8. Remark. This argument clearly breaks down over countable fields. In fact, over $\overline{\mathbb{F}}_p$, all points of an elliptic curve E are torsion. (Any given point is defined over some finite field \mathbb{F}_{p^r} . The points defined over \mathbb{F}_{p^r} form a subgroup of E , using the explicit geometric construction of the group law, and there are finite number of points over \mathbb{F}_{p^r} — certainly no more than the number of \mathbb{F}_{p^r} -points of \mathbb{P}^2 .) But over \mathbb{Q} , there are elliptic curves with non-torsion points. Even better, there are examples over \mathbb{Q} : $[2, 1, 8]$ is a \mathbb{Q} -point of the elliptic curve $y^2z = x^3 + 4xz^2 - z^3$ that is not torsion. The proof would carry us too far afield, but one method is to use the Nagell-Lutz Theorem (see for example [Sil, Cor. 7.2]).

We now use the existence of a non-torsion point to create some interesting pathologies.

20.11.9. A map of projective varieties not arising from a map of graded rings, even after regrading.

20.11.B. EXERCISE. Suppose $E \hookrightarrow \mathbb{P}_{\mathbb{C}}^2$ is a smooth complex plane cubic (hence genus 1), yielding a graded ring S_{\bullet} . Let $t: E \rightarrow E$ be translation by a non-torsion point. Show that t does not correspond to a map of graded rings $S_{\bullet} \rightarrow S_{\bullet}$, even after regrading (cf. Remark 17.4.8). (If you wish, you can show the following. If $u: E \rightarrow E$ is a translation, show that u corresponds to some map of regraded rings $S_{n\bullet} \rightarrow S_{n\bullet}$ if and only if u is translation by a torsion point.)

20.11.10. An affine open subset of an affine scheme that is not a distinguished open set.

We can use this to construct an example of an affine scheme X and an affine open subset Y that is not distinguished in X . Let $X = E - p$, which is affine (see Exercise 20.2.B, or better, note that the linear series $\mathcal{O}(3p)$ sends E to \mathbb{P}^2 in such a way that the “line at infinity” meets E only at p ; then $E - p$ has a closed embedding into the affine scheme \mathbb{A}^2).

Let q be another point on E so that $q - p$ is non-torsion. Then $E - p - q$ is affine (Exercise 20.2.B). Assume that it is distinguished. Then there is a function f on $E - p$ that vanishes on q (to some positive order d). Thus f is a rational function on E that vanishes at q to order d , and (as the total number of zeros minus poles of f is 0) has a pole at p of order d . But then $d(p - q) = 0$ in $\text{Pic}^0 E$, contradicting our assumption that $p - q$ is non-torsion.

In particular, $E - p$ is an affine scheme, and q is locally cut out by one equation, but it is not globally cut out *even set-theoretically* by one equation. This was promised in Exercise 8.3.E and §9.5.1.

20.11.11. A proper (nonprojective) surface with no nontrivial line bundles.

We next use a non-torsion point p on an elliptic curve E to construct a proper nonprojective surface with no nontrivial line bundles. Let X_1 be \mathbb{P}^2 , and Z_1 a smooth cubic in X_1 , identified with E . Let $Z_2 \subset X_2$ be exactly the same. Glue X_1 to X_2 along the isomorphism $Z_1 \xrightarrow{\sim} Z_2$ given, not by the identity, but by translation by p . (This gluing construction was described in §17.4.12.) Call the result X . By Exercise 17.4.P (and the sentence thereafter), X is proper.

20.11.C. EXERCISE. Show that every line bundle on X is trivial. Hint: Suppose \mathcal{L} is an invertible sheaf on X . Then $\mathcal{L}|_{X_i}$ is $\mathcal{O}(d_i)$ for some d_i (for $i = 1, 2$). The restriction of $\mathcal{L}|_{X_1}$ to Z_1 must agree with the restriction of $\mathcal{L}|_{X_2}$ to Z_2 . Use this to show that $d_1 = d_2 = 0$. Explain why gluing two trivial bundles (on X_1 and X_2 respectively) together in this way yields a trivial bundle on X .

20.11.D. EXERCISE. Show that X is not projective.

See Exercise 21.2.G for a somewhat simpler example of a proper nonprojective surface.

20.11.12. A Picard group that has no chance of being a scheme.

We informally observed that the Picard group of an elliptic curve “wants to be” a scheme (see §20.10.5). This is true of projective (and even proper) varieties in general (see [FGIKNV, Ch. 9]). On the other hand, if we work over \mathbb{C} , the affine scheme $E - p - q$ (in the language of §20.11.10 above) has a Picard group that can be interpreted as \mathbb{C} modulo a lattice modulo a non-torsion point (e.g., $\mathbb{C}/\langle 1, i, \pi \rangle$). This has no reasonable interpretation as a manifold, let alone a variety. So the fact that the Picard group of proper varieties turns out to be a scheme should be seen as quite remarkable.

20.11.13. A variety whose ring of functions is not finitely generated.

We next show an example of a complex variety whose ring of global sections is not finitely generated. (An example over \mathbb{Q} can be constructed in the same way using the curve of Remark 20.11.8. If you wish, you can show that these examples are even quasiaffine.) This is related to Hilbert’s fourteenth problem (see [Mu5, §3]).

20.11.E. PRELIMINARY EXERCISE. Suppose X is a scheme, and L is the total space of a line bundle corresponding to invertible sheaf \mathcal{L} , so $L = \text{Spec } \bigoplus_{n \geq 0} (\mathcal{L}^\vee)^{\otimes n}$. (This construction first appeared in Definition 18.1.4.) Show that if X is quasicompact, then $H^0(L, \mathcal{O}_L) = \bigoplus H^0(X, (\mathcal{L}^\vee)^{\otimes n})$. (Possible hint: choose a trivializing

cover for \mathcal{L} . Rhetorical question: can you figure out the more general statement if \mathcal{L} is a rank r locally free sheaf?)

Let E be an elliptic curve over some ground field k , \mathcal{N} a degree 0 non-torsion invertible sheaf on E , and \mathcal{P} a positive-degree invertible sheaf on E . Then $H^0(E, \mathcal{N}^m \otimes \mathcal{P}^n)$ is nonzero if and only if either (i) $n > 0$, or (ii) $m = n = 0$ (in which case the sections are elements of k).

20.11.F. EASY EXERCISE. Show that the ring $R = \bigoplus_{m,n \geq 0} H^0(E, \mathcal{N}^m \otimes \mathcal{P}^n)$ is not finitely generated.

20.11.G. EXERCISE. Let X be the total space of the vector bundle associated to $(\mathcal{N} \oplus \mathcal{P})^\vee$ over E . Show that the ring of global sections of X is R , and hence is not finitely generated. (Hint: interpret X as a line bundle over a line bundle over E .)

20.11.H. EXERCISE. Show that X (as in the above exercise) is a variety (and in particular, Noetherian) whose ring of global sections is not Noetherian.

Application: A glimpse of intersection theory

This section is intended to show some of the interesting things we can do with what we have built so far. But it can be skipped on a first reading, by young people in a hurry.

The only reason this chapter appears after Chapter 20 is because we will use Exercise 20.2.E in the double-starred proof of the Nakai-Moishezon criterion for ampleness (see §21.4.5).

Throughout this chapter, X will be a k -variety. In most applications, X will be projective.

21.1 Intersecting n line bundles with an n -dimensional variety

The central tool in this chapter is the following.

21.1.1. Definition: intersection product, or intersection number $(\mathcal{L}_1 \cdots \mathcal{L}_n), (\mathcal{L}_1 \cdots \mathcal{L}_n \cdot \mathcal{F}), (\mathcal{L}_1 \cdots \mathcal{L}_n \cdot Y)$.

Suppose \mathcal{F} is a coherent sheaf on X with proper support (automatic if X is proper) of dimension at most n , and $\mathcal{L}_1, \dots, \mathcal{L}_n$ are invertible sheaves on X . Let $(\mathcal{L}_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F})$ be the signed sum over the 2^n subsets of $\{1, \dots, n\}$

$$(21.1.1.1) \quad \sum_{\{i_1, \dots, i_m\} \subset \{1, \dots, n\}} (-1)^m \chi(X, \mathcal{L}_{i_1}^\vee \otimes \cdots \otimes \mathcal{L}_{i_m}^\vee \otimes \mathcal{F}).$$

We call this the **intersection of $\mathcal{L}_1, \dots, \mathcal{L}_n$ with \mathcal{F}** . (Never forget that whenever we write $(\mathcal{L}_1 \cdots \mathcal{L}_n \cdot \mathcal{F})$, we are implicitly assuming that $\dim \text{Supp } \mathcal{F} \leq n$.)

The case we will find most useful is if \mathcal{F} is the structure sheaf of a closed subscheme Y (of dimension at most n). In this case, we may write it $(\mathcal{L}_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot Y)$, and call it **the intersection of $\mathcal{L}_1, \dots, \mathcal{L}_n$ with Y** .

If $Y = X$, we may write it $(\mathcal{L}_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n)$, and call it **the intersection of $\mathcal{L}_1, \dots, \mathcal{L}_n$** (with the “ X ” left implicit).

If the \mathcal{L}_i are all the same, say \mathcal{L} , one often writes $(\mathcal{L}^n \cdot \mathcal{F})$ or $(\mathcal{L}^n \cdot Y)$ or (\mathcal{L}^n) . (Be careful with this confusing notation: \mathcal{L}^n does not mean $\mathcal{L}^{\otimes n}$.)

In some circumstances the convention is to omit the parentheses.

We will prove many things about the intersection product in this chapter. One fact will be left until we study flatness (Exercise 25.7.D): that it is “deformation-invariant” — that it is constant in “nice” families.

21.1.A. EXERCISE (REALITY CHECK). Show that if $\mathcal{L}_1 \cong \mathcal{O}_X$ then $(\mathcal{L}_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F}) = 0$.

The following exercise suggests that the intersection product might be interesting, as it “interpolates” between two useful notions: the degree of a line bundle on a curve, and Bézout’s Theorem.

21.1.B. EXERCISE.

(a) If X is a projective curve, and \mathcal{L} is an invertible sheaf on X , show that $(\mathcal{L} \cdot X) = \deg_X \mathcal{L}$.

(b) Suppose k is an infinite field, $X = \mathbb{P}^N$, and Y is a dimension n subvariety of X . If H_1, \dots, H_n are generally chosen hypersurfaces of degrees d_1, \dots, d_n respectively (so $\dim(H_1 \cap \dots \cap H_n \cap Y) = 0$ by Exercise 12.3.B(d)), then by Bézout’s Theorem (Exercise 19.6.K),

$$\deg(H_1 \cap \dots \cap H_n \cap Y) = d_1 \cdots d_n \deg(Y).$$

Show that

$$(\mathcal{O}_X(H_1) \cdots \mathcal{O}_X(H_n) \cdot Y) = d_1 \cdots d_n \deg(Y).$$

We now describe some of the properties of the intersection product. In the course of proving Exercise 21.1.B(b) you may in effect solve the following exercise.

21.1.C. EXERCISE. Suppose D is an effective Cartier divisor on X that restricts to an effective Cartier divisor $D|_Y = D \cap Y$ on Y (i.e., remains locally not a zerodivisor on Y). Show that

$$(\mathcal{L}_1 \cdots \mathcal{L}_{n-1} \cdot \mathcal{O}(D) \cdot Y) = (\mathcal{L}_1 \cdots \mathcal{L}_{n-1} \cdot D|_Y).$$

More generally, if D is an effective Cartier divisor on X that does not contain any associated point of \mathcal{F} , show that

$$(\mathcal{L}_1 \cdots \mathcal{L}_{n-1} \cdot \mathcal{O}(D) \cdot \mathcal{F}) = (\mathcal{L}_1 \cdots \mathcal{L}_{n-1} \cdot \mathcal{F}|_D).$$

(A similar idea came up in the proof that the Hilbert polynomial is actually polynomial; see the discussion around (19.6.2.1).)

21.1.2. Definition. For this reason, if D is an effective Cartier divisor, in the symbol for the intersection product, we often write D instead of $\mathcal{O}(D)$. We interchangeably think of intersecting effective Cartier divisors rather than line bundles. For example, we will discuss the special case of intersection theory on a surface in §21.2, and when we intersect two curves C and D , we will write the intersection as $(C \cdot D)$ or even $C \cdot D$.

21.1.D. EXERCISE. Show that the intersection product (21.1.1.1) is preserved by field extension of k .

21.1.3. Proposition. — Assume X is projective. For fixed \mathcal{F} , the intersection product $(\mathcal{L}_1 \cdots \mathcal{L}_n \cdot \mathcal{F})$ is a symmetric multilinear function of the $\mathcal{L}_1, \dots, \mathcal{L}_n$.

Proposition 21.1.3 is actually true with “projective” replaced by “proper”, see [K11, Prop. 2] or [Ko1, Prop. VI.2.7]. Unlike most extensions to the proper case, this is not just an application of Chow’s Lemma. It involves a different approach, involving a beautiful trick called *dévissage*.

Proof. Symmetry is clear. By Exercise 21.1.D, we may assume that k is infinite (e.g., algebraically closed). We now prove the result by induction on n .

21.1.E. EXERCISE (BASE CASE). Prove the result when $n = 1$. Hint: Exercise 19.4.S. (In fact, you can take the base case to be $n = 0$, if this doesn't confuse you.)

We now assume the result for smaller values of n . We use a trick. We wish to show that (for arbitrary $\mathcal{L}_1, \mathcal{L}'_1, \mathcal{L}_2, \dots, \mathcal{L}_n$),

$$(21.1.3.1) \quad (\mathcal{L}_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F}) + (\mathcal{L}'_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F}) - ((\mathcal{L}_1 \otimes \mathcal{L}'_1) \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F})$$

is 0.

21.1.F. EXERCISE. Rewrite (21.1.3.1) as

$$(21.1.3.2) \quad (\mathcal{L}_1 \cdot \mathcal{L}'_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F}).$$

(There are now $n + 1$ line bundles appearing in the product, but this does not contradict the definition of the intersection product, as $\dim \text{Supp } \mathcal{F} \leq n < n + 1$.)

21.1.G. EXERCISE. Use the inductive hypothesis to show that (21.1.3.1) is 0 if $\mathcal{L}_n \cong \mathcal{O}(D)$ for D an effective Cartier divisor missing the associated points of \mathcal{F} .

In particular, if \mathcal{L}_n is very ample, then (21.1.3.1) is 0, as Exercise 19.6.A shows that there exists a section of \mathcal{L}_n missing the associated points of \mathcal{F} .

By the symmetry of its incarnation as (21.1.3.2), expression (21.1.3.1) vanishes if \mathcal{L}_1 is very ample. Let \mathcal{A} and \mathcal{B} be any two very ample line bundles on X . Then by substituting $\mathcal{L}_1 = \mathcal{B}$ and $\mathcal{L}'_1 = \mathcal{A} \otimes \mathcal{B}^\vee$, using the vanishing of (21.1.3.1), we have

$$(21.1.3.3) \quad (\mathcal{A} \otimes \mathcal{B}^\vee \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F}) = (\mathcal{A} \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F}) - (\mathcal{B} \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F})$$

Both summands on the right side of (21.1.3.3) are linear in \mathcal{L}_n , so the same is true of the left side. But by Exercise 17.6.F, any invertible sheaf on X may be written in the form $\mathcal{A} \otimes \mathcal{B}^\vee$ ("as the difference of two very amples"), so $(\mathcal{L}_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F})$ is linear in \mathcal{L}_n , and thus (by symmetry) in each of the \mathcal{L}_i . (An interesting feature of this argument is that we intended to show linearity in \mathcal{L}_1 , and ended up showing linearity in \mathcal{L}_n .) \square

We have an added bonus arising from the proof.

21.1.H. EXERCISE. Suppose X is projective. Show that if $\dim \text{Supp } \mathcal{F} < n + 1$, and $\mathcal{L}_1, \mathcal{L}'_1, \mathcal{L}_2, \dots, \mathcal{L}_n$ are invertible sheaves on X , then (21.1.3.2) vanishes. In other words, the intersection product of $n + 1$ invertible sheaves with a coherent sheaf \mathcal{F} vanishes if the $\dim \text{Supp } \mathcal{F} < n + 1$. (In fact, the result holds with "projective" replaced by "proper", as the results it relies on hold in this greater generality.)

21.1.4. Proposition. — Suppose X is projective. The intersection product depends only on the numerical equivalence classes of the \mathcal{L}_i .

(Numerical equivalence was defined in §19.4.9 — two line bundles are linearly equivalent if their degrees are the same on any curve.) Again, the result remains true with "projective" replaced by "proper". But in the proof here, we use the fact that every line bundle is the difference two very ample line bundles in both the proof of Proposition 21.1.3 and in the proof of Proposition 21.1.4 itself. For the proof of the Proposition in the proper case, see [FGIKNV, Prop. B.20].

Proof. Suppose \mathcal{L}_1 is numerically equivalent to \mathcal{L}'_1 , and $\mathcal{L}_2, \dots, \mathcal{L}_n$, and \mathcal{F} are arbitrary. We wish to show that $(\mathcal{L}_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F}) = (\mathcal{L}'_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F})$.

By Exercise 21.1.D, we may assume that k is infinite (e.g., algebraically closed). We proceed by induction on n . The case $n = 1$ follows from Exercise 19.4.S. We assume that $n > 1$, and assume the result for “smaller n ”. By multilinearity of the intersection product, and the fact that each \mathcal{L}_n may be written as the “difference” of two very ample invertible sheaves (Exercise 17.6.F), it suffices to prove the result in the case when \mathcal{L}_n is very ample. We may then write $\mathcal{L}_n = \mathcal{O}(D)$, where D is an effective Cartier divisor missing the associated points of \mathcal{F} (Exercise 19.6.A). Then

$$\begin{aligned} (\mathcal{L}_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F}) &= (\mathcal{L}_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_{n-1} \cdot \mathcal{F}|_D) \quad (\text{Exercise 21.1.C}) \\ &= (\mathcal{L}'_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_{n-1} \cdot \mathcal{F}|_D) \quad (\text{inductive hypothesis}) \\ &= (\mathcal{L}'_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F}) \quad (\text{Exercise 21.1.C}). \end{aligned}$$

□

21.1.5. Asymptotic Riemann-Roch.

If Y is a projective (=proper, see Remark 19.7.2) curve, $\chi(Y, \mathcal{L}^{\otimes m}) = m \deg_Y \mathcal{L} + \chi(Y, \mathcal{O}_Y)$ (Exercise 19.4.S) is a linear polynomial in m , whose leading term is an intersection product. This generalizes as follows.

21.1.I. EXERCISE (ASYMPTOTIC RIEMANN-ROCH). Suppose X is projective. (As usual, the result will remain true with “projective” replaced by “proper”.) Suppose \mathcal{F} is a coherent sheaf on X with $\dim \text{Supp } \mathcal{F} \leq n$, and \mathcal{L} is a line bundle on X . Show that $\chi(X, \mathcal{L}^{\otimes m} \otimes \mathcal{F})$ is a polynomial in m of degree at most n . Show that the coefficient of m^n in this polynomial (the “leading term”) is $(\mathcal{L}^n \cdot \mathcal{F})/n!$. Hint: Exercise 21.1.H implies that $(\mathcal{L}^{n+1} \cdot (\mathcal{L}^{\otimes i} \otimes \mathcal{F})) = 0$. (Careful with this notation: \mathcal{L}^{n+1} doesn’t mean $\mathcal{L}^{\otimes(n+1)}$, it means $\mathcal{L} \cdot \mathcal{L} \cdots \mathcal{L}$ with $n+1$ factors.) Expand this out using (21.1.1.1) to get a recursion for $\chi(X, \mathcal{L}^{\otimes m} \otimes \mathcal{F})$. Your argument may resemble the proof of polynomiality of the Hilbert polynomial, Theorem 19.6.1, so you may find further hints there. Exercise 19.6.C in particular might help.

We know all the coefficients of this polynomial if X is a curve, by Riemann-Roch (see (19.4.6.1)), or basically by definition. We will know/interpret all the coefficients if X is a regular projective surface and \mathcal{F} is an invertible sheaf when we prove Riemann-Roch for surfaces (Exercise 21.2.B(b)). To understand the general case, we need the theory of Chern classes. The result is the Hirzebruch-Riemann-Roch Theorem, which can be further generalized to the celebrated Grothendieck-Riemann-Roch Theorem (see [F2, §15.2]).

21.1.J. EXERCISE (THE PROJECTION FORMULA). Suppose $\pi: X_1 \rightarrow X_2$ is a (projective) morphism of integral projective varieties of the same dimension n , and $\mathcal{L}_1, \dots, \mathcal{L}_n$ are invertible sheaves on X_2 . Show that

$$(\pi^* \mathcal{L}_1 \cdots \pi^* \mathcal{L}_n) = \deg(X_1/X_2)(\mathcal{L}_1 \cdots \mathcal{L}_n).$$

(The first intersection is on X_1 , and the second is on X_2 .) Hint: Let $d = \deg \pi$, and assume $d > 0$. (Deal with the case where π is not dominant separately, so $d = 0$ by convention, using Chevalley’s Theorem 8.4.2.) Argue that by the multilinearity of the intersection product, it suffices to deal with the case where the \mathcal{L}_i are very ample. Then choose sections of each \mathcal{L}_i , all of whose intersections lie in an open subset U where π has “genuine degree d ”. To find U : first use Exercise 10.3.G to find a dense open subset $U' \subset X_2$ over which π is finite. Then use Useful

Exercise 14.3.F to show that there exists a dense open subset $U \subset U'$ on which $\pi_* \mathcal{O}$ is a locally free sheaf of rank d . In the “flatness” language of Chapter 25, you are showing that there is a dense open subset U of X_2 over which π is finite and flat (and hence has “constant degree”, see Remark 25.4.8). (As usual, the result holds with “projective” replaced with “proper”; see [Ko1, Prop. VI.2.11].)

21.1.6. Remark: A more general projection formula. Suppose $\pi: X_1 \rightarrow X_2$ is a proper morphism of proper varieties, and \mathcal{F} is a coherent sheaf on X_1 with $\dim \text{Supp } \mathcal{F} \leq n$ (so $\dim \text{Supp } \pi_* \mathcal{F} \leq n$, using Exercise 12.2.C). Suppose also that $\mathcal{L}_1, \dots, \mathcal{L}_n$ are invertible sheaves on X_2 . Then

$$(\pi^* \mathcal{L}_1 \cdots \pi^* \mathcal{L}_n \cdot \mathcal{F}) = (\mathcal{L}_1 \cdots \mathcal{L}_n \cdot \pi_* \mathcal{F}).$$

This is also called the **projection formula** (and generalizes, in a nonobvious way, Exercise 21.1.J). Proofs are given in [FGIKNV, B.15] and [Ko1, Prop. VI.2.11]. Because we won’t use this version of the projection formula, we do not give a proof here.

21.1.K. EXERCISE (INTERSECTING WITH AMPLE LINE BUNDLES). Suppose X is a projective k -variety, and \mathcal{L} is an ample line bundle on X . Show that for any subvariety Y of X of dimension n , $(\mathcal{L}^n \cdot Y) > 0$. (Hint: use Proposition 21.1.3 and Theorem 17.6.2 to reduce to the case where \mathcal{L} is very ample. Then show that $(\mathcal{L}^n \cdot Y) = \deg Y$ in the embedding into projective space induced by the linear series $|\mathcal{L}|$.)

The Nakai-Moishezon criterion (Theorem 21.4.1) states that this characterizes ampleness.

21.1.7. ★ Cohomological interpretation in the complex projective case, generalizing Exercise 19.4.H and §19.6.4. If $k = \mathbb{C}$, we can interpret $(\mathcal{L}_1 \cdots \mathcal{L}_n \cdot Y)$ as the degree of

$$(21.1.7.1) \quad c_1((\mathcal{L}_1)_{\text{an}}) \cup \cdots \cup c_1((\mathcal{L}_n)_{\text{an}}) \cap [Y_{\text{an}}]$$

in $H_0(Y_{\text{an}}, \mathbb{Z})$. (Recall $c_1((\mathcal{L}_i)_{\text{an}}) \in H^2(X_{\text{an}}, \mathbb{Z})$, as discussed in Exercise 19.4.H.) One way of proving this is to use multilinearity of both the intersection product and (21.1.7.1) to reduce to the case where the \mathcal{L}_n is very ample, so $\mathcal{L}_n \cong \mathcal{O}(D)$, where D restricts to an effective Cartier divisor E on Y . Then show that if \mathcal{L} is an analytic line bundle on Y_{an} with nonzero section E_{an} , then $c_1(\mathcal{L}) \cap [Y_{\text{an}}] = [E_{\text{an}}]$. Finally, use induction on n and Exercise 21.1.C.

21.1.8. Important remark: additive notation for line bundles. There is a standard, useful, but confusing convention suggested by the multilinearity of the intersection product: we write tensor product of invertible sheaves *additively*. This is convenient for working with multilinearity. (Some people try to avoid confusion by using divisors rather than line bundles, as we add divisors when we “multiply” the corresponding line bundles. This is psychologically helpful, but may add more confusion, as one then has to worry about the whether and why and how and when line bundles correspond to divisors.) We will use this, for example, in Exercises 21.2.B–21.2.G, §21.2.10, and §21.4.

21.2 Intersection theory on a surface

We now apply the general machinery of §21.1 to the case of a regular projective surface X . (What matters is that X is Noetherian and factorial, so $\text{Pic } X \rightarrow \text{Cl } X$ is an isomorphism, Proposition 15.2.10. Recall that regular schemes are factorial by the Auslander-Buchsbaum Theorem 13.8.5.)

21.2.A. EXERCISE/DEFINITION. Suppose C and D are effective divisors (i.e., curves) on X .

(a) Show that

$$(21.2.0.1) \quad \deg_C \mathcal{O}_X(D)|_C$$

$$(21.2.0.2) \quad = (\mathcal{O}(C) \cdot \mathcal{O}(D) \cdot X)$$

$$(21.2.0.3) \quad = \deg_D \mathcal{O}_X(C)|_D.$$

We call this the **intersection number** of C and D , and denote it $C \cdot D$.

(b) If C does not contain any associated point of D (so in particular, if C and D have no components in common), show that

$$(21.2.0.4) \quad C \cdot D = h^0(C \cap D, \mathcal{O}_{C \cap D})$$

where $C \cap D$ is the scheme-theoretic intersection of C and D on X .

21.2.1. Remark. In particular, if C and D have no irreducible component in common, then $C \cdot D \geq 0$. This means that we cannot hope to understand intersection numbers using moving lemmas as in topology — if we have some irreducible curve E with $E^2 < 0$, we cannot “move” E to something (call it F) “equivalent” to E , to compute $E \cdot E$ as $E \cdot F$.

21.2.2. Important aside. The hypothesis in Exercise 21.2.A(b), that C not contain any associated point of D , is a red herring. In fact, D can never have any embedded points. We can see this in two ways, by either invoking (i) a fact we have stated but not proved, or (ii) a fact we have not yet stated but *will* prove. ... , as we will see in §29.2.5 when discussing Cohen-Macaulay rings. (More generally, we will see that local complete intersections have no embedded points — this should help motivate you to learn about Cohen-Macaulayness.) The case of \mathbb{A}^2 (and hence \mathbb{P}^2) can be done by hand (Exercise 6.5.E).

Thus the hypothesis in Exercise 21.2.A(b) can be replaced by the more simple “ C and D have no common components”.

21.2.3. Advantages and disadvantages. We thus have four descriptions of the intersection number (21.2.0.1)–(21.2.0.4), each with advantages and disadvantages. The Euler characteristic description (21.2.0.2) is remarkably useful (for example, in the exercises below), but the geometry is obscured. The definition $\deg_C \mathcal{O}_X(D)|_C$, (21.2.0.1), is not obviously symmetric in C and D . The definition $h^0(C \cap D, \mathcal{O}_{C \cap D})$, (21.2.0.4), is clearly local — to each point of $C \cap D$, we have a vector space. For example, we know that in $\mathbb{A}_{\mathbb{K}}^2$, $y - x^2 = 0$ meets the x -axis with multiplicity 2, because h^0 of the scheme-theoretic intersection $(\mathbb{K}[x, y]/(y - x^2, y))$ has dimension 2. (This h^0 is also the *length* of the dimension 0 scheme, whose definition you may be able to figure out given Definition ?? of the length of a module.)

By Proposition 21.1.3, the intersection number induces a bilinear “intersection form”

$$(21.2.3.1) \quad \text{Pic } X \times \text{Pic } X \longrightarrow \mathbb{Z}.$$

By Asymptotic Riemann-Roch (Exercise 21.1.I), $\chi(X, \mathcal{O}(nD))$ is a quadratic polynomial in n .

21.2.4. You can verify that Exercise 21.2.A recovers Bézout’s Theorem for plane curves (see Exercise 19.6.K), using $\chi(\mathbb{P}^2, \mathcal{O}(n)) = (n+2)(n+1)/2$ (from Theorem 19.1.2), and the fact that effective Cartier divisors on \mathbb{P}_k^2 have no embedded points (Exercise 6.5.E).

Before getting to a number of interesting explicit examples, we derive a couple of fundamental theoretical facts.

21.2.B. EXERCISE. Assuming Serre duality for X (Theorem 19.5.1), prove the following for a smooth projective surface X . (We are mixing divisor and invertible sheaf notation, as described in Remark 21.1.8, so be careful. Here K_X is a divisor corresponding to ω_X .)

(a) (sometimes called the adjunction formula) $C \cdot (K_X + C) = 2p_a(C) - 2$ for any curve $C \subset X$. Hint: compute $C \cdot (-K_X - C)$ instead. (See Exercise 22.5.B and §31.4 for other versions of the adjunction formula.)

(b) (Riemann-Roch for surfaces)

$$\chi(X, \mathcal{O}_X(D)) = D \cdot (D - K_X)/2 + \chi(X, \mathcal{O}_X)$$

for any Weil divisor D (cf. Riemann-Roch for curves, Exercise 19.4.B).

21.2.5. Two explicit examples: $\mathbb{P}^1 \times \mathbb{P}^1$ and $\text{Bl}_p \mathbb{P}^2$.

21.2.C. EXERCISE: $X = \mathbb{P}^1 \times \mathbb{P}^1$. Recall from Exercise 15.2.O that $\text{Pic}(\mathbb{P}^1 \times \mathbb{P}^1) = \mathbb{Z}\ell \times \mathbb{Z}m$, where ℓ is the curve $\mathbb{P}^1 \times \{0\}$ and m is the curve $\{0\} \times \mathbb{P}^1$. Show that the intersection form (21.2.3.1) is given by $\ell \cdot \ell = m \cdot m = 0$, $\ell \cdot m = 1$. (Hint: You can compute the cohomology groups of line bundles on $\mathbb{P}^1 \times \mathbb{P}^1$ using Exercise ??, but it is much faster to use Exercise 21.2.A(b).) What is the class of the diagonal in $\mathbb{P}^1 \times \mathbb{P}^1$ in terms of these generators?

21.2.D. EXERCISE: THE BLOWN UP PROJECTIVE PLANE. (You needn’t have read Chapter 23 to do this exercise!) Let $X = \text{Bl}_p \mathbb{P}^2$ be the blow-up of \mathbb{P}_k^2 at a k -valued point (the origin, say) p — see Exercise 10.3.F, which describes the blow-up of \mathbb{A}_k^2 , and “compactify”. Interpret $\text{Pic } X$ as generated (as an abelian group) by ℓ and e , where ℓ is a line not passing through the origin, and e is the exceptional divisor. Show that the intersection form (21.2.3.1) is given by $\ell \cdot \ell = 1$, $e \cdot e = -1$, and $\ell \cdot e = 0$. Hence show that $\text{Pic } X \cong \mathbb{Z}\ell \times \mathbb{Z}e$ (as promised in the aside in Exercise 15.2.P). In particular, the exceptional divisor has negative self-intersection. (This exercise will be generalized in §23.4.13.)

21.2.6. Hint. Here is a possible hint to get the intersection form in Exercise 21.2.D. The scheme-theoretic preimage in $\text{Bl}_p \mathbb{P}^2$ of a line through the origin is the scheme-theoretic union of the exceptional divisor e and the “proper transform” m of the line through the origin. Show that $\ell = e + m$ in $\text{Pic}(\text{Bl}_p \mathbb{P}^2)$ (writing the Picard

group operation additively, cf. Remark 21.1.8). Show that $\ell \cdot m = e \cdot m = 1$ and $m \cdot m = 0$.

21.2.7. Definition: (-1) -curve. Notice that the exceptional divisor e has self-intersection -1 . We will see more generally in Exercise 23.4.O that this is the case for all exceptional divisors of blow-ups of smooth surfaces (over k) blown up at a k -valued point. We give such curves a name. If X is a surface over k , and $C \subset X$ is a curve in X consisting of smooth points, with $C \cong \mathbb{P}_k^1$, and $C \cdot C = -1$, we say that C is a (-1) -curve. (Once we know more, we will be able to restate “ $C \cdot C = -1$ ” as “ C has normal bundle $\mathcal{O}(-1)$ ”: combine Exercise 22.2.H with Exercise 21.2.A(a).)

21.2.E. EXERCISE. Show that the blown up projective plane $\text{Bl}_p \mathbb{P}^2$ in Exercise 21.2.D is not isomorphic to $\mathbb{P}^1 \times \mathbb{P}^1$, perhaps considering their (isomorphic) Picard groups, and identifying which classes are effective (represented by effective divisors). (This is an example of a pair of smooth projective birational surfaces that have isomorphic Picard groups, but which are not isomorphic. This exercise shows that \mathbb{F}_0 is not isomorphic to \mathbb{F}_1 . (The method and result are generalized in Exercise 21.2.Q.)

21.2.F. EXERCISE (CF. EXERCISE 19.4.X). Show that the nef cone (Exercise 19.4.W) of $\text{Bl}_p \mathbb{P}^2$ is generated by ℓ and m . Hint: show that ℓ and m are nef. By intersecting line bundles with the curves e and ℓ , show that nothing outside the cone spanned by ℓ and m is nef. (Side Remark: note that as in Exercise 19.4.X, linear series corresponding to the boundaries of the cone give “interesting contractions”.)

21.2.G. EXERCISE: A PROPER NONPROJECTIVE SURFACE. Show the existence of a proper nonprojective surface over a field as follows, paralleling the construction of a proper nonprojective threefold in §17.4.11. Take two copies of the blown up projective plane $\text{Bl}_p \mathbb{P}^2$, gluing ℓ on the first to e on the second, and e on the second to ℓ on the first. Hint: show that if \mathcal{L} is a line bundle having positive degree on each effective curve, then $\mathcal{L} \cdot \ell > \mathcal{L} \cdot e$, using $\ell = e + m$ from Hint 21.2.6. (See §20.11.11 for another example of a proper nonprojective surface.)

21.2.8. Fibrations.

Suppose $\pi: X \rightarrow B$ is a morphism from a regular projective surface to a regular curve and $b \in B$ is a closed point. Let $F = \pi^{-1}(b)$. Then $\mathcal{O}_X(F) = \pi^* \mathcal{O}_B(b)$, which is isomorphic to \mathcal{O} on F . Thus $F \cdot F = \deg_F \mathcal{O}_X(F) = 0$: “the self-intersection of a fiber is 0”. The same argument works without X being regular, so long as you phrase it properly: $(\pi^* \mathcal{O}_X(b))^2 = 0$.

21.2.H. EXERCISE. Suppose E is an elliptic curve, with origin p . On $E \times E$, let Δ be the diagonal. By considering the “difference” map $E \times E \rightarrow E$, for which $\pi^{-1}(p) = \Delta$, show that $\Delta^2 = 0$. Show that $N_{\mathbb{Q}}^1(E \times E)$ has rank at least 3. Show that in general for schemes X and Y , $\text{Pic } X \times \text{Pic } Y \rightarrow \text{Pic}(X \times Y)$ (defined by pulling back and tensoring) need not be isomorphism; the case of $X = Y = \mathbb{P}^1$ is misleading.

Remark: $\dim_{\mathbb{Q}} N_{\mathbb{Q}}^1(E \times E)$ is always 3 or 4. It is 4 if there is a nontrivial endomorphism from E to itself (i.e., not just multiplication $[n]$ by some n , §20.11.5, followed by a translation); the additional class comes from the graph of this endomorphism. (See [Mu3, §21, App. III] for an introduction to the tools needed to show this.)

Our next goal is to describe the self-intersection of a curve on a ruled surface (Exercise 21.2.J). To set this up, we have a useful preliminary result.

21.2.I. EXERCISE (THE NORMAL BUNDLE TO A SECTION OF Proj OF A RANK 2 VECTOR BUNDLE). Suppose X is a scheme, and \mathcal{V} is a rank 2 locally free sheaf on X . Explain how the short exact sequences

$$(21.2.8.1) \quad 0 \rightarrow \mathcal{S} \rightarrow \mathcal{V} \rightarrow \mathcal{Q} \rightarrow 0$$

on X , where \mathcal{S} and \mathcal{Q} have rank 1, correspond to the sections $\sigma: X \rightarrow \mathbb{P}\mathcal{V}$ to the projection $\mathbb{P}\mathcal{V} \rightarrow X$. Show that the normal bundle to $\sigma(X)$ in $\mathbb{P}\mathcal{V}$ is $\mathcal{Q} \otimes \mathcal{S}^\vee$. (A generalization is stated in §22.4.10.) Hint: (i) For simplicity, it is convenient to assume $\mathcal{S} = \mathcal{O}_X$, by replacing \mathcal{V} by $\mathcal{V} \otimes \mathcal{S}^\vee$, as the statement of the problem respects tensoring by an invertible sheaf (see Exercise 18.2.G). (ii) Assume now (*with loss of generality*) that $\mathcal{Q} \cong \mathcal{O}_X$, and the exact sequence splits. Then describe the section as $\sigma: X \rightarrow \mathbb{P}^1 \times X$, with X mapping to the 0 section. Describe an isomorphism of \mathcal{O}_X with the normal bundle to $\sigma(X) \rightarrow \mathbb{P}^1 \times X$. (Do *not* just say that the normal bundle “is trivial”.) (iii) Now consider the case where \mathcal{Q} is general. Choose trivializing open neighborhoods U_i of \mathcal{Q} , and let g_{ij} be the transition function for \mathcal{Q} . On the overlap between two trivializing open neighborhoods $U_i \cap U_j$, determine how your two isomorphisms of \mathcal{O}_X with $N_{\sigma(X)/\mathbb{P}^1 \times X}$ from (ii) (one for U_i , one for U_j) are related. In particular, show that they differ by g_{ij} .

21.2.J. EXERCISE (SELF-INTERSECTIONS OF SECTIONS OF RULED SURFACES). Suppose C is a regular curve, and \mathcal{V} is a rank 2 locally free sheaf on C . Then $\mathbb{P}\mathcal{V}$ is a ruled surface (Example 18.2.4). Fix a section σ of $\mathbb{P}\mathcal{V}$ corresponding to a filtration (21.2.8.1). Show that $\sigma(C) \cdot \sigma(C) = \deg_C(\mathcal{Q} \otimes \mathcal{S}^\vee)$.

21.2.9. The Hirzebruch surfaces $\mathbb{F}_n = \text{Proj}_{\mathbb{P}^1}(\mathcal{O}_{\mathbb{P}^1} \oplus \mathcal{O}_{\mathbb{P}^1}(n))$.

Recall the definition of the Hirzebruch surface $\mathbb{F}_n = \text{Proj}_{\mathbb{P}^1}(\mathcal{O}_{\mathbb{P}^1} \oplus \mathcal{O}_{\mathbb{P}^1}(n))$ in Example 18.2.4. It is a \mathbb{P}^1 -bundle over \mathbb{P}^1 ; let $\pi: \mathbb{F}_n \rightarrow \mathbb{P}^1$ be the structure morphism. Using Exercise 21.2.J, corresponding to

$$0 \longrightarrow \mathcal{O}(n) \longrightarrow \mathcal{O} \oplus \mathcal{O}(n) \longrightarrow \mathcal{O} \longrightarrow 0,$$

we have a section of π of self-intersection $-n$; call it $E \subset \mathbb{F}_n$. Similarly, corresponding to

$$0 \longrightarrow \mathcal{O} \longrightarrow \mathcal{O} \oplus \mathcal{O}(n) \longrightarrow \mathcal{O}(n) \longrightarrow 0,$$

we have a section $C \subset \mathbb{F}_n$ of self-intersection n . Let p be any k -valued point of \mathbb{P}^1 , and let $F = \pi^{-1}(p)$.

21.2.K. EXERCISE. Show that the line bundle $\mathcal{O}(F)$ is independent of the choice of p .

21.2.L. EXERCISE. Show that $\text{Pic } \mathbb{F}_n$ is generated by E and F . In the course of doing this, you will develop “local charts” for \mathbb{F}_n , which will help you solve later exercises.

21.2.M. EXERCISE. Compute the intersection matrix on $\text{Pic } \mathbb{F}_n$. Show that E and F are independent, and thus $\text{Pic } \mathbb{F}_n \cong \mathbb{Z}E \oplus \mathbb{Z}F$. Calculate C in terms of E and F .

21.2.N. EXERCISE. Show how to identify $\mathbb{F}_n \setminus E$, along with the structure map π , with the total space of the line bundle $\mathcal{O}(n)$ on \mathbb{P}^1 , with C as the 0-section. Similarly show how to identify $\mathbb{F}_n \setminus C$ with the total space of the line bundle $\mathcal{O}(-n)$ on \mathbb{P}^1 ; with E as the 0-section.

21.2.O. EXERCISE. Show that $h^0(\mathbb{F}_n, \mathcal{O}_{\mathbb{F}_n}(C)) > 1$. Hint: As $\mathcal{O}_{\mathbb{F}_n}(C)$ has a section — namely C — we have that $h^0(\mathbb{F}_n, \mathcal{O}_{\mathbb{F}_n}(C)) \geq 1$. One way to proceed is to write down another section using local charts for \mathbb{F}_n .

21.2.P. EXERCISE. Show that every effective curve on \mathbb{F}_n is a non-negative linear combination of E and F . (Conversely, it is clear that for every nonnegative a and b , $\mathcal{O}(aE + bF)$ has a section, corresponding to the effective curve “ $aE + bF$ ”. The extension of this to $N_{\mathbb{Q}}^1$ is called the **effective cone**, and this notion, extended to proper varieties more generally, can be very useful. This exercise shows that E and F generate the effective cone of \mathbb{F}_n .) Hint: show that because “ F moves”, any effective curve must intersect F nonnegatively, and similarly because “ C moves” (Exercise 21.2.O), any effective curve must intersect C nonnegatively. If $\mathcal{O}(aE + bF)$ has a section corresponding to an effective curve D , what does this say about a and b ?

21.2.Q. EXERCISE. By comparing effective cones, and the intersection pairing, show that the \mathbb{F}_n are pairwise nonisomorphic. (This result was promised in Example 18.2.4. Exercise 21.2.E is a special case.)

Exercise 21.2.Q is difficult to do otherwise, and foreshadows the fact that nef and effective cones are useful tools in classifying and understanding varieties in general. In particular, they are central to the minimal model program.

21.2.R. EXERCISE. If $n = 0$, show that there are no curves on \mathbb{F}_n of negative self-intersection. If $n \geq 0$, show that E is the unique curve on \mathbb{F}_n with self-intersection $-n$, and there are no reduced curves on \mathbb{F}_n of smaller self-intersection. This again gives another (related) means of showing that the \mathbb{F}_n are pairwise nonisomorphic.

21.2.S. EXERCISE. Show that the nef cone of \mathbb{F}_n is generated by C and F . (We will soon see that by Kleiman’s criterion for ampleness, Theorem 21.4.7, that the ample cone is the interior of this cone, so we have now identified the ample line bundles on \mathbb{F}_n .)

21.2.T. EXERCISE. We have seen earlier (Exercises 21.2.F and 19.4.X) that the boundary of the nef cone give “interesting contractions”. Show that the map given by the linear series corresponding to $\mathcal{O}(F)$ is the projection of \mathbb{F}_n onto \mathbb{P}^1 , and the map given by the linear series corresponding to $\mathcal{O}(C)$ sends \mathbb{F}_n to the projective cone over the degree n rational normal curve.

After this series of exercises, you may wish to revisit Exercises 21.2.C–21.2.F, and interpret them as special cases: $\mathbb{F}_0 \cong \mathbb{P}^1 \times \mathbb{P}^1$ and $\mathbb{F}_1 \cong \text{Bl}_{\mathbb{P}} \mathbb{P}^2$.

21.2.10. The Hodge Index Theorem.

We use what we have learned to prove the following celebrated result.

21.2.11. The Hodge Index Theorem. — Suppose X is an irreducible smooth projective surface (over a field k) with $\mathcal{L}, \mathcal{H} \in \text{Pic } X$, with $\mathcal{H} \cdot \mathcal{H} > 0$ and $\mathcal{L} \cdot \mathcal{H} = 0$. Then (a) $\mathcal{L} \cdot \mathcal{L} \leq 0$, and (b) equality holds if and only if \mathcal{L} is numerically trivial.

Before reading on, you should review the definitions of $N^1(X)$ and $\rho(X)$ in §19.4.9.

21.2.U. EXERCISE (REASON FOR THE NAME). (We will not use this.) By arguments similar to the classification of quadratic forms over algebraically closed fields (Exercise 5.4.J), one can show that given a symmetric bilinear form on a *real* vector space V of (finite) dimension n , there are integers a_+, a_0, a_- such that then given any orthogonal basis v_1, \dots, v_n , the number of $v_i \cdot v_i$ that are positive (resp. zero, negative) is a_+ (resp. a_0, a_-). The ordered pair (a_+, a_-) is often called the **index**, and is sometimes called the **signature**. (This result is sometimes called *Sylvester's law of inertia*. See, for example, [Lan, §XV.4] for a proof.) Show that the Hodge Index Theorem implies that $N_{\mathbb{R}}^1(X) := N^1(X) \otimes_{\mathbb{Z}} \mathbb{R}$, along with the bilinear form coming from the intersection product, has index $(1, \rho(X) - 1)$. You should think of the positive “term” as coming from the existence of an ample class; the “rest” is negative.

Proof.

21.2.12. We begin with the following fact. If \mathcal{H} is an ample invertible sheaf, \mathcal{M} is any invertible sheaf, and $\mathcal{M} \cdot \mathcal{H} > \omega_X \cdot \mathcal{H}$, then $h^2(X, \mathcal{M}) = 0$. (Translation: for all invertible sheaves “more positive than ω_X ”, there is no top cohomology; compare this to §20.2.5, the corresponding statement in dimension 1.) Reason: by Serre duality (Theorem 19.5.1), we wish to show that $h^0(X, \omega_X \otimes \mathcal{M}^\vee) = 0$. But if otherwise $\omega_X \otimes \mathcal{M}^\vee$ had a nonzero section, then its (effective nonzero) divisor would intersect \mathcal{H} nonnegatively (Exercise 21.1.K), so (using additive notation, see Remark 21.1.8) $(\omega_X - \mathcal{M}) \cdot \mathcal{H} \geq 0$, contradicting $\mathcal{M} \cdot \mathcal{H} > \omega_X \cdot \mathcal{H}$.

21.2.13. We next show that if \mathcal{H} is ample, and \mathcal{L} is any invertible sheaf with $\mathcal{L} \cdot \mathcal{L} > 0$ and $\mathcal{L} \cdot \mathcal{H} > 0$, then for $n \gg 0$, $\mathcal{L}^{\otimes n}$ has a section. Here is why. For $n \gg 0$,

$$\mathcal{L}^{\otimes n} \cdot \mathcal{H} = n\mathcal{L} \cdot \mathcal{H} > \omega_X \cdot \mathcal{H},$$

so $h^2(X, \mathcal{L}^{\otimes n}) = 0$ by §21.2.12 (taking $\mathcal{M} = \mathcal{L}^{\otimes n}$). Then for $n \gg 0$, by Riemann-Roch for surfaces (Exercise 21.2.B(b)),

$$\begin{aligned} h^0(X, \mathcal{L}^{\otimes n}) - h^1(X, \mathcal{L}^{\otimes n}) &= \chi(X, \mathcal{L}^{\otimes n}) \\ &= (n^2/2)\mathcal{L} \cdot \mathcal{L} - (n/2)\mathcal{L} \cdot \omega_X + \chi(X, \mathcal{O}_X). \end{aligned}$$

As $\mathcal{L} \cdot \mathcal{L} > 0$, the right side is positive for $n \gg 0$, so

$$h^0(X, \mathcal{L}^{\otimes n}) \geq h^0(X, \mathcal{L}^{\otimes n}) - h^1(X, \mathcal{L}^{\otimes n}) > 0$$

as desired.

21.2.14. We are now ready to prove the Hodge Index Theorem. Assume first that $\mathcal{L} \cdot \mathcal{L} > 0$ (and $\mathcal{L} \cdot \mathcal{H} = 0$). Then for $n \gg 0$, $\mathcal{H}' := \mathcal{L} \otimes \mathcal{H}^{\otimes n}$ is ample (indeed very ample) by Exercise 17.6.E. Then $\mathcal{L} \cdot \mathcal{H}' = \mathcal{L} \cdot \mathcal{L} > 0$, so by §21.2.13, $\mathcal{L}^{\otimes n}$ is effective, which implies that $\mathcal{L} \cdot \mathcal{H} = (\mathcal{L}^{\otimes n} \cdot \mathcal{H})/n > 0$ (using Exercise 21.1.K again), contradicting our hypothesis.

21.2.15. Assume finally that $\mathcal{L} \cdot \mathcal{H} = 0$, and (i) $\mathcal{L} \cdot \mathcal{L} = 0$, but that (ii) \mathcal{L} is not numerically trivial. By (ii), we can find an invertible sheaf \mathcal{Q} such that $\mathcal{Q} \cdot \mathcal{L} \neq 0$. Then we can find an invertible sheaf \mathcal{R} such that $\mathcal{R} \cdot \mathcal{L} \neq 0$ and $\mathcal{R} \cdot \mathcal{H} = 0$: take

$$\mathcal{R} = (\mathcal{H} \cdot \mathcal{H})\mathcal{Q} - (\mathcal{Q} \cdot \mathcal{H})\mathcal{H}.$$

Then take $\mathcal{L}' = \mathcal{L}^{\otimes n} \otimes \mathcal{R}$, so $\mathcal{L}' \cdot \mathcal{H} = 0$, but because $\mathcal{R} \cdot \mathcal{L} \neq 0$, we can find some $n \in \mathbb{Z}$ so that

$$\mathcal{L}' \cdot \mathcal{L}' = n\mathcal{L} \cdot \mathcal{R} + \mathcal{R} \cdot \mathcal{R} > 0.$$

Then the argument of §21.2.14 applies, with \mathcal{L}' in the place of \mathcal{L} . \square

21.2.16. Generalizations and variations. The hypotheses can be weakened considerably. We used smoothness only because we need Serre duality, with an invertible sheaf ω_X . We will see in §31.4 that we need less than smoothness for this. But we can do much better, and a proof where X is merely required to be geometrically irreducible and proper requires only a minor modification of our argument, see [FGIKNV, Thm. B.27].

21.3 The Grothendieck group of coherent sheaves, and an algebraic version of homology

The construction of the intersection product (21.1.1.1) may leave you hungry for something more, especially in light of the cohomological interpretation of §21.1.7. You may want some sort of homology-like theory which is a repository for cycles of different dimensions, on which (Chern classes of) line bundles can act. We can actually do this easily, given what we know.

21.3.1. Definition. If X is a k -variety, we define the **Grothendieck group of coherent sheaves**, which we denote $K(X)$ (and which is often denoted $K_0(X)$), as the abelian group generated symbols of the form $[\mathcal{F}]$ where \mathcal{F} is a coherent sheaf on X , subject to the relations that $[\mathcal{F}] = [\mathcal{F}']$ if $\mathcal{F} \cong \mathcal{F}'$, and $[\mathcal{F}'] + [\mathcal{F}''] = [\mathcal{F}]$ if there is a short exact sequence

$$0 \rightarrow \mathcal{F}' \rightarrow \mathcal{F} \rightarrow \mathcal{F}'' \rightarrow 0.$$

(Caution: the notation $K(X)$ is now used for two concepts, the Grothendieck group and the function field. The meaning of the symbol should be clear from the context.) By construction, the Grothendieck group is the universal construction of an operator on the category of coherent sheaves that “behaves well in exact sequences”. For example, if X is proper, then:

- (i) We have a map $\chi: \text{Coh}_X \rightarrow \mathbb{Z}$, as the Euler characteristic of a coherent sheaf is finite by Theorem 19.1.3(i) in the projective case, and more generally by Grothendieck’s Coherence Theorem 19.9.1 in the proper case. This map descends to $\chi: K(X) \rightarrow \mathbb{Z}$ by Exercise 19.4.A (which extends without change to the proper case).
- (ii) If X is integral, then the rank function

$$\text{rank}: \text{Coh}_X \rightarrow \mathbb{Z}$$

descends to a function $\text{rank}: K(X) \rightarrow \mathbb{Z}$. (The argument of Exercise 19.4.I applies.)

21.3.2. Definition and nonstandard notation. The Grothendieck group is filtered by dimension: let $K(X)^{\leq d}$ be the subgroup of $K(X)$ generated by coherent sheaves supported in dimension at most d . (This is called the **dimensional filtration** or **coniveau filtration on $K(X)$** , although we won't need these names.) Let $A'_d(X)$ be the d th graded piece of $K(X)$, i.e., $A'_d(X) := K(X)^{\leq d} / K(X)^{\leq d-1}$.

If \mathcal{L} is an invertible sheaf on X , define $\mathcal{L} \cdot: K(X) \rightarrow K(X)$ by $\mathcal{L} \cdot [\mathcal{F}] = [\mathcal{F}] - [\mathcal{L}^\vee \otimes \mathcal{F}]$. (Do you see why this operator is well-defined?)

21.3.A. EXERCISE. If X is projective and k is infinite, show that $\mathcal{L} \cdot$ sends $K(X)^{\leq d}$ to $K(X)^{\leq d-1}$. (Hint: for each fixed \mathcal{F} supported on a subset of dimension at most d , write \mathcal{L} is a difference of two very ample invertible sheaves, and choose sections of those two, missing the associated points of \mathcal{F} .)

21.3.3. Remark. The previous exercise holds true without k being infinite or X being proper; see [K01, Prop. VI.2.5].

21.3.4. For the rest of this section, we assume X is projective and k is infinite. (But in light of Remark 21.3.3, these hypotheses can be removed.) By Exercise 21.3.A, $\mathcal{L} \cdot$ descends to a map $A'_d(X) \rightarrow A'_{d-1}(X)$; we denote this operator $c_1(\mathcal{L}) \cap$. (It is the action of the first Chern class.)

21.3.B. EXERCISE ("c₁ IS ADDITIVE"). Show that $c_1(\mathcal{L} \otimes \mathcal{L}') \cap = (c_1(\mathcal{L}) \cap) + (c_1(\mathcal{L}') \cap)$. Hint: show that $((\mathcal{L} \otimes \mathcal{L}') \cdot) - (\mathcal{L} \cdot) - (\mathcal{L}' \cdot) = -(\mathcal{L} \cdot) \circ (\mathcal{L}' \cdot)$, and thus sends $K(X)^{\leq d}$ to $K(X)^{\leq d-2}$. (This will remind you of the trick in the proof of Proposition 21.1.3, and indeed is the motivation for that trick. Caution: the action of $\text{Pic}(X)$ is not additive on $K(X)$; it only becomes additive once we pass to the associated graded ring $A'_\bullet(X)$.)

If \mathcal{F} has support of dimension at most n , and $\mathcal{L}_1, \dots, \mathcal{L}_n$ are n invertible sheaves, we can now reinterpret the intersection product $(\mathcal{L}_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F})$ as

$$c_1(\mathcal{L}_1) \cap \cdots \cap c_1(\mathcal{L}_n) \cap [\mathcal{F}]$$

where $[\mathcal{F}]$ is interpreted as lying in either $K(X)$ or $A'_n(X)$. You can now go back and read §21.1, and reprove all the results with this starting point, sometimes obtaining interesting generalizations. We have now promoted the intersection product from an integer to a class in $A'_0(X)$, which may have much more information. (This is true even if X is a curve of positive genus over \mathbb{C} , using the theory developed in [F2].)

These $A'_d(X)$ behave like homology groups in a number of ways. If $Y \subset X$ is a closed subvariety of pure dimension d , we have a class $[Y] := [\mathcal{O}_Y] \in A'_d(X)$. The groups $A'_d(X)$ have appropriate functorial properties. For example, if $\pi: X_1 \rightarrow X_2$ is a proper morphism (even if the X_i are not themselves proper), we have a map $\pi_*: K(X_1) \rightarrow K(X_2)$ given by

$$(21.3.4.1) \quad \pi_*[\mathcal{F}] = [\pi_*\mathcal{F}] - [R^1\pi_*\mathcal{F}] + \cdots$$

(using the long exact sequence for higher pushforwards, Theorem 19.8.1 (c)) which descends to a map $A'_d(X_1) \rightarrow A'_d(X_2)$, where the "later terms" on the right side of

(21.3.4.1) disappear, so $\pi_*[\mathcal{F}] = [\pi_*\mathcal{F}]$. This pushforward interacts well with the first Chern class of line bundles, yielding the projection formula of Remark 21.1.6.

If π is instead a *flat morphism* (Remark 17.3.8, soon to be discussed at length in Chapter 25), then π^* is exact, so we have a map $\pi^*: K(X_2) \rightarrow K(X_1)$. If the “relative dimension” of this map is r (to be properly defined in Definition 25.5.8), this yields a map $\pi^*: A'_d(X_2) \rightarrow A'_{d+r}(X_1)$. This interacts well with first Chern classes and proper pushforwards.

If $k = \mathbb{C}$ and X is proper, there is a map $A'_d(X) \rightarrow H_{2d}(X, \mathbb{Q})$, which behaves as you might hope (for example, in its interaction with Chern classes of line bundles). If X is *not* proper (but $k = \mathbb{C}$), then the map is to Borel-Moore homology rather than usual homology.

Our groups $A'_d(X)$ are a good approximation of the theory of Chow groups $A_d(X)$, as developed in [F2]. In fact, there is a surjective map

$$(21.3.4.2) \quad A_d(X) \rightarrow A'_d(X)$$

[F2, Examp. 15.1.5], and this map is an isomorphism once tensored with \mathbb{Q} [F2, Thms. 18.2 and 18.3]. This is the beginning of a long and rich story in algebraic geometry.

21.3.5. Side Remark. The surjection map (21.3.4.2) need not be an isomorphism, see [SGA6, Exp. XIV, §4.5-4.7] (although its kernel must be torsion, as described above). As a tantalizing example, if X is the group E_8 , the kernel has 2-torsion, 3-torsion, and 5-torsion, see [KN, DZ].

21.4 ★★ The Nakai-Moishezon and Kleiman criteria for ampleness

Exercise 21.1.K stated that if X is projective k -variety, and \mathcal{L} is an ample line bundle on X , then for any subvariety Y of X of dimension n , $(\mathcal{L}^n \cdot Y) > 0$. The Nakai-Moishezon criterion states that this is a characterization:

21.4.1. Theorem (Nakai-Moishezon criterion for ampleness). — *If \mathcal{L} is an invertible sheaf on a projective k -scheme X , and for every irreducible subvariety Y of X , $(\mathcal{L}^{\dim Y} \cdot Y) > 0$, then \mathcal{L} is ample.*

21.4.2. Remarks. We note that X need only be proper for this result to hold ([K11, Thm. III.1.1]).

Before proving the Nakai-Moishezon criterion, we point out some consequences related to our discussion of numerical equivalence in §19.4.8. By Proposition 21.1.4, $(\mathcal{L}^{\dim Y} \cdot Y)$ depends only on the numerical equivalence class of \mathcal{L} , so ampleness is a numerical property. As a result, the notion of ampleness makes sense on $N_{\mathbb{Q}}^1(X)$. As the tensor product of two ample invertible sheaves is ample (Exercise 17.6.H), the ample \mathbb{Q} -line bundles in $N_{\mathbb{Q}}^1(X)$ form a cone, called the **ample cone** of X .

21.4.3. Proposition. — *If X is a projective k -scheme, the ample cone is open.*

21.4.4. In the rest of this section, we often use additive notation for the tensor product of invertible sheaves, as described in Remark 21.1.8. This is because we

want to deal with intersections on the \mathbb{Q} -vector space $N_{\mathbb{Q}}^1(X)$. For example by $((a\mathcal{L}_1 + b\mathcal{L}_1') \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F})$ (where $a, b \in \mathbb{Q}$), we mean $a(\mathcal{L}_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F}) + b(\mathcal{L}_1' \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F})$.

Proof. Suppose \mathcal{A} is an ample invertible sheaf on X . We will describe a small open neighborhood of $[\mathcal{A}]$ in $N_{\mathbb{Q}}^1(X)$ consisting of ample \mathbb{Q} -line bundles. Choose invertible sheaves $\mathcal{L}_1, \dots, \mathcal{L}_n$ on X whose classes form a basis of $N_{\mathbb{Q}}^1(X)$. By Exercise 17.6.E, there is some m such that $\mathcal{A}^{\otimes m} \otimes \mathcal{L}_i$ and $\mathcal{A}^{\otimes m} \otimes \mathcal{L}_i^{\vee}$ are both very ample for all i . Thus (in the additive notation of §21.4.4), $\mathcal{A} + \frac{1}{m}\mathcal{L}_i$ and $\mathcal{A} - \frac{1}{m}\mathcal{L}_i$ are both ample. As the ample \mathbb{Q} -line bundles form a cone, it follows that $\mathcal{A} + \epsilon_1\mathcal{L}_1 + \cdots + \epsilon_n\mathcal{L}_n$ is ample for $\sum_i |\epsilon_i| \leq 1/m$. \square

21.4.5. Proof of the Nakai-Moishezon criterion, Theorem 21.4.1. We prove the Nakai-Moishezon criterion in several steps.

21.4.A. UNIMPORTANT EXERCISE. Prove the case where $\dim X = 0$.

Step 1: initial reductions. Suppose \mathcal{L} satisfies the hypotheses of the Theorem; we wish to show that \mathcal{L} is ample. By Exercises 19.7.A and 19.7.B, we may assume that X is integral. Moreover, we can work by induction on dimension, so we can assume that \mathcal{L} is ample on any closed subvariety. The base case is dimension 1, which was done in Exercise 20.2.E.

Step 2: sufficiently high powers of \mathcal{L} have sections. We show that $H^0(X, \mathcal{L}^{\otimes m}) \neq 0$ for $m \gg 0$.

Our plan is as follows. By Asymptotic Riemann-Roch (Exercise 21.1.I), $\chi(X, \mathcal{L}^{\otimes m}) = m^n(\mathcal{L}^n)/n! + \cdots$ grows (as a function of m) without bound. A plausible means of attack is to show that $h^i(X, \mathcal{L}^{\otimes m}) = 0$ for $i > 0$ and $m \gg 0$. We won't do that, but will do something similar.

By Exercise 17.6.C, \mathcal{L} is the difference of two very ample line bundles, say $\mathcal{L} \cong \mathcal{A} \otimes \mathcal{B}^{-1}$ with $\mathcal{A} = \mathcal{O}(A)$ and $\mathcal{B} = \mathcal{O}(B)$. From $0 \rightarrow \mathcal{O}(-A) \rightarrow \mathcal{O} \rightarrow \mathcal{O}_A \rightarrow 0$ we have

$$(21.4.5.1) \quad 0 \rightarrow \mathcal{L}^{\otimes m}(-B) \rightarrow \mathcal{L}^{\otimes(m+1)} \rightarrow \mathcal{L}^{\otimes(m+1)}|_A \rightarrow 0.$$

From $0 \rightarrow \mathcal{O}(-B) \rightarrow \mathcal{O} \rightarrow \mathcal{O}_B \rightarrow 0$, we have

$$(21.4.5.2) \quad 0 \rightarrow \mathcal{L}^{\otimes m}(-B) \rightarrow \mathcal{L}^{\otimes m} \rightarrow \mathcal{L}^{\otimes m}|_B \rightarrow 0.$$

Choose m large enough so that both $\mathcal{L}^{\otimes(m+1)}|_A$ and $\mathcal{L}^{\otimes m}|_B$ have vanishing higher cohomology (i.e., $h^{>0} = 0$ for both; use the inductive hypothesis that \mathcal{L} is ample on all proper closed subvarieties, and Serre vanishing, Theorem 19.1.3(ii)). This implies that for $i \geq 2$,

$$\begin{aligned} H^i(X, \mathcal{L}^{\otimes m}) &\cong H^i(X, \mathcal{L}^{\otimes m}(-B)) \quad (\text{long exact sequence for (21.4.5.2)}) \\ &\cong H^i(X, \mathcal{L}^{\otimes(m+1)}) \quad (\text{long exact sequence for (21.4.5.1)}) \end{aligned}$$

so the higher cohomology stabilizes (is constant) for large m . From

$$\chi(X, \mathcal{L}^{\otimes m}) = h^0(X, \mathcal{L}^{\otimes m}) - h^1(X, \mathcal{L}^{\otimes m}) + \text{constant},$$

$H^0(\mathcal{L}^{\otimes m}) \neq 0$ for $m \gg 0$, completing Step 2.

So by replacing \mathcal{L} by a suitably large multiple (ampleness is independent of taking tensor powers, Theorem 17.6.2), we may assume \mathcal{L} has a section D . We now use D as a crutch.

Step 3: $\mathcal{L}^{\otimes m}$ is globally generated for $m \gg 0$.

As D is effective, $\mathcal{L}^{\otimes m}$ is globally generated on the complement of D : we have a section nonvanishing on that big open set. Thus any base locus must be contained in D . Consider the short exact sequence

$$(21.4.5.3) \quad 0 \rightarrow \mathcal{L}^{\otimes(m-1)} \rightarrow \mathcal{L}^{\otimes m} \rightarrow \mathcal{L}^{\otimes m}|_D \rightarrow 0$$

Now $\mathcal{L}|_D$ is ample by our inductive hypothesis. Choose m so large that $H^1(X, \mathcal{L}^{\otimes m}|_D) = 0$ (Serre vanishing, Theorem 19.1.3(ii)). From the exact sequence associated to (21.4.5.3),

$$\phi_m: H^1(X, \mathcal{L}^{\otimes(m-1)}) \rightarrow H^1(X, \mathcal{L}^{\otimes m})$$

is surjective for $m \gg 0$. Using the fact that the $H^1(X, \mathcal{L}^{\otimes m})$ are finite-dimensional vector spaces, as m grows, $H^1(X, \mathcal{L}^{\otimes m})$ must eventually stabilize, so the ϕ_m are isomorphisms for $m \gg 0$.

Thus for large m , from the long exact sequence in cohomology for (21.4.5.3), $H^0(X, \mathcal{L}^{\otimes m}) \rightarrow H^0(X, \mathcal{L}^{\otimes m}|_D)$ is surjective for $m \gg 0$. But $H^0(X, \mathcal{L}^{\otimes m}|_D)$ has no base points by our inductive hypothesis (applied to D), i.e., for any point p of D there is a section of $\mathcal{L}^{\otimes m}|_D$ not vanishing at p , so $H^0(X, \mathcal{L}^{\otimes m})$ has no base points on D either, completing Step 3.

Step 4. Thus \mathcal{L} is a base-point-free line bundle with positive degree on each curve (by hypothesis of Theorem 21.4.1), so by Exercise 19.4.N we are done. \square

The following result is the key to proving Kleiman's numerical criterion of ampleness, Theorem 21.4.7.

21.4.6. Kleiman's Theorem. — Suppose X is a projective k -scheme. If \mathcal{L} is a nef invertible sheaf on X , then $(\mathcal{L}^k \cdot V) \geq 0$ for every irreducible subvariety $V \subset X$ of dimension k .

As usual, this extends to the proper case, see [K11, Thm. IV.2.1]. And as usual, we postpone the proof until after we appreciate the consequences.

21.4.B. EXERCISE.

(a) (*limit of amples is nef*) If \mathcal{L} and \mathcal{H} are any two invertible sheaves such that $\mathcal{L} + \epsilon\mathcal{H}$ is ample for all sufficiently small $\epsilon > 0$, show that \mathcal{L} is nef. (Hint: $\lim_{\epsilon \rightarrow 0}$. This doesn't require Kleiman's Theorem.)

(b) (*nef + ample = ample*) Suppose X is a projective k -scheme, \mathcal{H} is ample, and \mathcal{L} is nef. Show that $\mathcal{L} + \epsilon\mathcal{H}$ is ample for all $\epsilon \in \mathbb{Q}^{\geq 0}$. (Hint: use the Nakai-Moishezon criterion: $((\mathcal{L} + \epsilon\mathcal{H})^k \cdot V) > 0$. This may help you appreciate the additive notation.)

21.4.7. Theorem (Kleiman's numerical criterion for ampleness). — Suppose X is a projective k -scheme.

- (a) The nef cone is the closure of the ample cone.
- (b) The ample cone is the interior of the nef cone.

Side remark: of course (a) is false if “projective” is relaxed to “proper” (as the ample cone of a proper nonprojective variety is empty, but 0 is nef so the nef cone is nonempty). However, part (b) is true if X is proper and factorial (see [K11,

Thm. IV.2.2] for a proof and a more general statement). Hence if X is smooth, proper, and nonprojective, then the interior of the nef cone is empty.

21.4.C. EXERCISE. See [K11, p. 326, Ex. 2] for Mumford's example of a non-ample line bundle on a smooth projective surface that meets every curve positively. Doesn't this contradict Theorem 21.4.7?

Proof. (a) Ample invertible sheaves are nef (Exercise 19.4.V(e)), and the nef cone is closed (Exercise 19.4.W), so the closure of the ample cone is contained in the cone. Conversely, each nef element of $N_{\mathbb{Q}}^1(X)$ is the limit of ample classes by Exercise 21.4.B(a), so the nef cone is contained in the closure of the ample cone.

(b) As the ample cone is open (Proposition 21.4.3), the ample cone is contained in the interior of the nef cone. Conversely, suppose \mathcal{L} is in the interior of the nef cone, and \mathcal{H} is any ample class. Then $\mathcal{L} - \epsilon\mathcal{H}$ is nef for all small enough positive ϵ . Then by Exercise 21.4.B(b), $\mathcal{L} = (\mathcal{L} - \epsilon\mathcal{H}) + \epsilon\mathcal{H}$ is ample. \square

Suitably motivated, we prove Kleiman's Theorem 21.4.6.

Proof. We may immediately reduce to the case where X is irreducible and reduced. We work by induction on $n := \dim X$. The base case $n = 1$ is obvious. So we assume that $(\mathcal{L}^{\dim V} \cdot V) \geq 0$ for all irreducible V not equal to X . We need only show that $(\mathcal{L}^n \cdot X) \geq 0$.

Fix some very ample invertible sheaf \mathcal{H} on X .

21.4.D. EXERCISE. Show that $(\mathcal{L}^k \cdot \mathcal{H}^{n-k} \cdot X) \geq 0$ for all $k < n$. (Hint: use the inductive hypothesis.)

Consider $P(t) := ((\mathcal{L} + t\mathcal{H})^n \cdot X) \in \mathbb{Q}[t]$. We wish to show that $P(0) \geq 0$. Assume otherwise that $P(0) < 0$. Now for $t \gg 0$, $\mathcal{L} + t\mathcal{H}$ is ample, so $P(t)$ is positive for large t . Thus $P(t)$ has positive real roots. Let t_0 be the largest positive real root of t . (In fact there is only one positive root, as Exercise 21.4.D shows that all the nonconstant coefficients of $P(t)$ are nonnegative.)

21.4.E. EXERCISE. Show that for (rational) $t > t_0$, $\mathcal{L} + t\mathcal{H}$ is ample. Hint: use the Nakai-Moishezon criterion (Theorem 21.4.1); if $V \neq X$ is an irreducible subvariety, show that $((\mathcal{L} + t\mathcal{H})^{\dim V} \cdot V) > 0$ by expanding $(\mathcal{L} + t\mathcal{H})^{\dim V}$.

Let $Q(t) := (\mathcal{L} \cdot (\mathcal{L} + t\mathcal{H})^{n-1} \cdot X)$ and $R(t) := (t\mathcal{H} \cdot (\mathcal{L} + t\mathcal{H})^{n-1} \cdot X)$, so $P(t) = Q(t) + R(t)$.

21.4.F. EXERCISE. Show that $Q(t) \geq 0$ for all rational $t > t_0$. Hint (which you will have to make sense of): $(\mathcal{L} + t\mathcal{H})$ is ample by Exercise 21.4.E, so for N sufficiently large, $N(\mathcal{L} + t\mathcal{H})$ is very ample. Use the idea of the proof of Proposition 21.1.4 to intersect X with $n - 1$ divisors in the class of $N(\mathcal{L} + t\mathcal{H})$ so that $((N(\mathcal{L} + t\mathcal{H}))^{n-1} \cdot X)$ is an effective curve C . Then $(\mathcal{L} \cdot C) \geq 0$ as \mathcal{L} is nef. Thus $Q(t_0) \geq 0$.

21.4.G. EXERCISE. Show that $R(t_0) > 0$. (Hint: Exercise 21.4.D.)

Thus $P(t_0) > 0$ as desired. \square

Differentials

22.1 Motivation and game plan

Differentials are an intuitive geometric notion, and we are going to figure out the right description of them algebraically. The algebraic manifestation is somewhat non-intuitive, so it is helpful to understand differentials first in terms of geometry. Also, although the algebraic statements are odd, none of the proofs are hard or long. You will notice that this topic could have been done as soon as we knew about morphisms and quasicoherent sheaves. We have usually introduced new ideas through a number of examples, but in this case we will spend a fair amount of time discussing theory, and only then get to examples.

Suppose X is a “smooth” k -variety. We would like to define a tangent bundle. We will see that the right way to do this will easily apply in much more general circumstances.

- We will see that the cotangent sheaf is more “natural” for schemes than tangent bundle/sheaf. This is similar to the fact that the Zariski *cotangent space* is more natural than the *tangent space* (i.e., if A is a ring and \mathfrak{m} is a maximal ideal, then $\mathfrak{m}/\mathfrak{m}^2$ is “more natural” than $(\mathfrak{m}/\mathfrak{m}^2)^\vee$), as we have repeatedly discussed since §13.1. In both cases this is because we are understanding “spaces” via their (sheaf of) functions on them, which is somehow dual to the geometric pictures you have of spaces in your mind.

So we will define the cotangent sheaf first. An element of the (co)tangent space will be called a **(co)tangent vector**.

- Our undergraduate intuition will continue to work — we can calculate using the Jacobian matrix, the answers will be the same in cases where our undergraduate calculations apply, and we will even have the advantage of avoiding δ ’s and ϵ ’s.
- Our construction will automatically apply for general X , even if X is not “smooth” (or even at all nice, e.g., finite type). The cotangent sheaf will not necessarily be locally free, but it will still be a quasicoherent sheaf.
- Better yet, this construction will naturally work “relatively”. For *any* $\pi: X \rightarrow Y$, we will define $\Omega_\pi = \Omega_{X/Y}$, a quasicoherent sheaf on X , the sheaf of *relative differentials*. The fiber of this sheaf at a point will be the cotangent vectors of the fiber of the map. This will specialize to the earlier case by taking $Y = \operatorname{Spec} k$. The idea is that this glues together the cotangent sheaves of the fibers of the family. Figure 22.1 is a sketch of the relative tangent space of a map $X \rightarrow Y$ at a point $p \in X$ — it is the tangent to the fiber. (The tangent space is easier to draw than the cotangent space!) An element of the relative (co)tangent space is called a **vertical or relative (co)tangent vector**.




FIGURE 22.1. The relative tangent space of a morphism $X \rightarrow Y$ at a point p

Thus the central concept of this chapter is the cotangent sheaf $\Omega_\pi = \Omega_{X/Y}$ for a morphism $\pi: X \rightarrow Y$ of schemes. A good picture to have in your mind is the following. If $\pi: X \rightarrow Y$ is a submersion of manifolds (a map inducing a surjection on tangent spaces), you might hope that the tangent spaces to the fibers at each point $p \in X$ might fit together to form a vector bundle. This is the relative tangent bundle (of π), and its dual is $\Omega_{X/Y}$ (see Figure 22.1). Even if you are not geometrically minded, you will find this useful. (For an arithmetic example, see Exercise 22.2.F.)

22.2 Definitions and first properties

22.2.1. The affine case: three definitions.

We first study the affine case. Suppose A is a B -algebra, so we have a morphism of rings $\phi: B \rightarrow A$ and a morphism of schemes $\text{Spec } A \rightarrow \text{Spec } B$. I will define an A -module $\Omega_{A/B}$ in three ways. This is called the **module of relative differentials** or the **module of Kähler differentials**. The module of differentials will be defined to be this module, as well as a map $d: A \rightarrow \Omega_{A/B}$ satisfying three properties. (Caution: although d sends an A -module to an A -module, it is not in general A -linear. A priori we take it as a homomorphism of abelian groups, but we will momentarily make it a homomorphism of B -modules, Exercise 22.2.A.)

- (i) *additivity*: $da + da' = d(a + a')$.
- (ii) *Leibniz*: $d(aa') = a da' + a' da$.
- (iii) *triviality on pullbacks*: $db = 0$ for $b \in \phi(B)$.

These properties will not be surprising if you have seen differentials in any other context.

22.2.A. TRIVIAL EXERCISE. Show that d is B -linear.

22.2.B. EXERCISE. Prove the quotient rule: if $a' = as$, then $s^2 da = s da' - a' ds$.

22.2.C. EXERCISE. State and prove the chain rule for $d(f(g))$ where f is a polynomial with B -coefficients, and $g \in A$. (As motivation, think of the case $B = k$. So for example, $da^n = na^{n-1}da$, and more generally, if f is a polynomial in one variable, $df(a) = f'(a)da$, where f' is defined formally: if $f = \sum c_i x^i$ then $f' = \sum c_i i x^{i-1}$.)

We will now see three definitions of the module of Kähler differentials, which will soon “sheafify” to the sheaf of relative differentials. The first definition is a concrete hands-on definition. The second is by universal property. And the third will globalize well, and will allow us to define $\Omega_{X/Y}$ conveniently in general.

22.2.2. First definition of differentials: explicit description. We define $\Omega_{A/B}$ to be finite A -linear combinations of symbols “ da ” for $a \in A$, subject to the three rules (i)–(iii) above. For example, take $A = k[x, y]$, $B = k$. Then a sample differential is $3x^2 dy + 4 dx \in \Omega_{A/B}$. We have identities such as $d(3xy^2) = 3y^2 dx + 6xy dy$.

22.2.3. Key fact. Note that if A is generated over B (as an algebra) by $x_i \in A$ (where i lies in some index set, possibly infinite), subject to some relations r_j (where j lies in some index set, and each is a polynomial in the x_i), then the A -module $\Omega_{A/B}$ is generated by the dx_i , subject to the relations $dr_j = 0$. In short, we needn’t take every single element of A ; we can take a generating set. And we needn’t take every single relation among these generating elements; we can take generators of the relations.

22.2.D. EXERCISE. Verify Key fact 22.2.3. (If you wish, use the affine conormal exact sequence, Theorem 22.2.12, to verify it; different people prefer to work through the theory in different orders. Just take care not to make circular arguments.)

In particular:

22.2.4. Proposition. — *If A is a finitely generated B -algebra, then $\Omega_{A/B}$ is a finite type (i.e., finitely generated) A -module. If A is a finitely presented B -algebra, then $\Omega_{A/B}$ is a finitely presented A -module.*

Recall (§8.3.13) that a ring A is *finitely presented* over another ring B if it can be expressed with finite number of generators and finite number of relations:

$$A = B[x_1, \dots, x_n] / (r_1(x_1, \dots, x_n), \dots, r_j(x_1, \dots, x_n)).$$

If A is Noetherian, then finitely presented is the same as finite type, as the “finite number of relations” comes for free, so most of you will not care.

Let’s now see some examples. Among these examples are three particularly important building blocks for ring maps: adding free variables; localizing; and taking quotients. If we know how to deal with these, we know (at least in theory) how to deal with any ring map. (They were similarly useful in understanding the fibered product in practice, in §10.2.)

22.2.5. Example: taking a quotient. If $A = B/I$, then $\Omega_{A/B} = 0$: $da = 0$ for all $a \in A$, as each such a is the image of an element of B . This should be believable; in this case, there are no “vertical tangent vectors”.

22.2.6. Example: adding variables. If $A = B[x_1, \dots, x_n]$, then $\Omega_{A/B} = A dx_1 \oplus \dots \oplus A dx_n$. (Note that this argument applies even if we add an arbitrarily infinite

number of indeterminates.) The intuitive geometry behind this makes the answer very reasonable. The cotangent bundle of affine n -space should indeed be free of rank n .

22.2.7. Explicit example: an affine plane curve. Consider the plane curve $y^2 = x^3 - x$ in \mathbb{A}_k^2 , where the characteristic of k is not 2 (see Figure 20.7 for a picture). Let $A = k[x, y]/(y^2 - x^3 + x)$ and $B = k$. By Key fact 22.2.3, the module of differentials $\Omega_{A/B}$ is generated by dx and dy , subject to the relation

$$2y \, dy = (3x^2 - 1) \, dx.$$

Thus in the locus where $y \neq 0$, dx is a generator (as dy can be expressed in terms of dx). We conclude that where $y \neq 0$, $\tilde{\Omega}_{A/B}$ is isomorphic to the trivial line bundle (invertible sheaf). Similarly, in the locus where $3x^2 - 1 \neq 0$, dy is a generator. These two loci cover the entire curve, as solving $y = 0$ gives $x^3 - x = 0$, i.e., $x = 0$ or ± 1 , and in each of these cases $3x^2 - 1 \neq 0$. We have shown that $\tilde{\Omega}_{A/B}$ is an invertible sheaf.

We can interpret dx and dy geometrically. Where does the differential dx vanish? The previous paragraph shows that it doesn't vanish on the patch where $2y \neq 0$. On the patch where $3x^2 - 1 \neq 0$, where dy is a generator, $dx = (2y/(3x^2 - 1)) \, dy$ from which we see that dx vanishes precisely where $y = 0$. You should find this believable from the picture. We have shown that $dx = 0$ precisely where the curve has a vertical tangent vector. Once we can pull back differentials (Exercise 22.2.K(a) or Theorem 22.2.27), we can interpret dx as the pullback of a differential on the x -axis to $\text{Spec } A$ (pulling back along the projection to the x -axis). When we do that, using the fact that dx doesn't vanish on the x -axis, we can interpret the locus where $dx = 0$ as the locus where the projection map branches. (Can you compute where $dy = 0$, and interpret it geometrically?)

This discussion applies to plane curves more generally. Suppose $A = k[x, y]/f(x, y)$, where for convenience $k = \bar{k}$. Then the same argument as the one given above shows that $\tilde{\Omega}_{A/k}$ is free of rank 1 on the open set $D(\partial f/\partial x)$, and also on $D(\partial f/\partial y)$. If $\text{Spec } A$ is a regular curve, then these two sets cover all of $\text{Spec } A$. (Exercise 13.2.E — basically the Jacobian criterion — gives regularity at the closed points. Furthermore, the curve must be reduced, or else as the nonreduced locus is closed, it would be nonreduced at a closed point, contradicting regularity. Finally, reducedness at a generic point is equivalent to regularity — a scheme whose underlying set is a point is reduced if and only if it is regular. Alternatively, we could invoke a big result, Fact 13.8.2, to get regularity at the generic point from regularity at the closed points.)

If, on the other hand, the plane curve is singular, then Ω is *not* locally free of rank one. For example, consider the plane curve $\text{Spec } A$ where $A = \mathbb{C}[x, y]/(y^2 - x^3)$, so

$$\Omega_{A/\mathbb{C}} = (A \, dx \oplus A \, dy)/(2y \, dy - 3x^2 \, dx).$$

Then the fiber of $\Omega_{A/\mathbb{C}}$ over the origin (computed by setting $x = y = 0$) is rank 2, as it is generated by dx and dy , with no relation.

Implicit in the above discussion is the following exercise, showing that Ω can be computed using the Jacobian matrix.

22.2.E. IMPORTANT BUT EASY EXERCISE (JACOBIAN DESCRIPTION OF $\Omega_{A/B}$). Suppose $A = B[x_1, \dots, x_n]/(f_1, \dots, f_r)$. Then

$$\Omega_{A/B} = (\oplus_i (A \, dx_i)) / (\oplus_j df_j)$$

may be interpreted as the cokernel of the Jacobian matrix (13.1.7.1)

$$J: A^{\oplus r} \rightarrow A^{\oplus n}.$$

(You might notice that if J has full rank everywhere, so A is smooth over B , then $\Omega_{A/B}$ is locally free; see Exercise 22.2.Q for more.)

22.2.8. Example: localization. If T is a multiplicative subset of B , and $A = T^{-1}B$, then $\Omega_{A/B} = 0$. Reason: by the quotient rule (Exercise 22.2.B), if $a = b/t$, then $da = (t \, db - b \, dt)/t^2 = 0$. If $A = B_f$, this is intuitively believable; then $\text{Spec } A$ is an open subset of $\text{Spec } B$, so there should be no vertical (co)tangent vectors.

22.2.F. IMPORTANT EXERCISE (FIELD EXTENSIONS). This notion of relative differentials is interesting even for finite extensions of fields. In other words, even when you map a reduced point to a reduced point, something interesting can happen with differentials.

(a) Suppose K/k is a separable algebraic extension. Show that $\Omega_{K/k} = 0$. Do not assume that K/k is a finite extension! (Hint: for any $\alpha \in K$, there is a polynomial $f(x)$ such that $f(\alpha) = 0$ and $f'(\alpha) \neq 0$.)

(b) Suppose k is a field of characteristic p , $K = k(t^p)$, $L = k(t)$. Compute $\Omega_{L/K}$ (where $K \hookrightarrow L$ is the “obvious” inclusion).

We now delve a little deeper, and discuss two useful and geometrically motivated exact sequences.

22.2.9. Theorem (cotangent exact sequence, affine version). — Suppose $C \rightarrow B \rightarrow A$ are ring morphisms. Then there is a natural exact sequence of A -modules

$$A \otimes_B \Omega_{B/C} \xrightarrow{a \otimes db \mapsto a \, db} \Omega_{A/C} \xrightarrow{da \mapsto da} \Omega_{A/B} \longrightarrow 0.$$

The proof will be quite straightforward algebraically, but the statement comes fundamentally from geometry, and that is how best to remember it. Figure 22.2 is a sketch of a map $\pi: X \rightarrow Y$. Here X should be interpreted as $\text{Spec } A$, Y as $\text{Spec } B$, and $\text{Spec } C$ (we will soon call it Z) is a point. (If you would like a picture with a higher-dimensional $\text{Spec } C$, just “take the product of Figure 22.2 with a curve”.) In the Figure, Y is “smooth”, and X is “smooth over Y ” — which means roughly that all the fibers are smooth. Suppose p is a point of X . Then the tangent space of the fiber of π at p is certainly a subspace of the tangent space of the total space of X at p . The cokernel is naturally the pullback of the tangent space of Y at $\pi(p)$. This short exact sequence for each p should be part of a short exact sequence of “relative tangent sheaves”

$$0 \rightarrow \mathcal{T}_{X/Y} \rightarrow \mathcal{T}_{X/Z} \rightarrow \pi^* \mathcal{T}_{Y/Z} \rightarrow 0$$

on X . (We will formally define “relative tangent sheaf” in §22.2.20.) Dualizing this yields

$$0 \rightarrow \pi^* \Omega_{Y/Z} \rightarrow \Omega_{X/Z} \rightarrow \Omega_{X/Y} \rightarrow 0.$$




FIGURE 22.2. A sketch of the geometry behind the cotangent exact sequence

This is precisely the statement of Theorem 22.2.9, except we also have left-exactness. This discrepancy is because the statement of the theorem is more general; we will see in Exercise 22.2.S that in the “smooth” case, we indeed have left-exactness.

22.2.10. Intriguing Remark. As always, whenever you see something right-exact, you should suspect that there should be some sort of (co)homology theory so that this is the end of a long exact sequence. This is indeed the case, and this exact sequence involves *André-Quillen homology* (see [E, p. 386] for more). You should expect that the next term to the left should be the first homology corresponding to A/B , and in particular shouldn’t involve C . So if you already suspect that you have exactness on the left in the case where A/B and B/C are “smooth” (whatever that means), and the intuition of Figure 22.2 applies, then you should expect further that all that is necessary is that A/B be “smooth”, and that this would imply that the first André-Quillen homology should be zero. Even though you wouldn’t precisely know what all the words meant, you would be completely correct! You would also be developing a vague inkling about the *cotangent complex*. We will see examples when left-exactness holds in a sufficiently “smooth” situation in Proposition 22.7.2 and Exercise 22.2.S. For a more general statement, see [Stacks, tag 06B6]. See also [Liu, Cor. 6.3.22] and [Liu, Prop. 6.3.11] for clean discussions of the next two terms in the sequence: “cotangent exact sequence”, “conormal exact sequence”, . . .

22.2.11. Proof of Theorem 22.2.9 (the cotangent exact sequence, affine version). First, note that surjectivity of $\Omega_{A/C} \rightarrow \Omega_{A/B}$ is clear, as this map is given by $da \mapsto da$ (where $a \in A$).

Next, the composition over the middle term is clearly 0, as this composition is given by $a \otimes db \mapsto adb \mapsto 0$.

Finally, we wish to identify $\Omega_{A/B}$ as the cokernel of $A \otimes_B \Omega_{B/C} \rightarrow \Omega_{A/C}$. Now $\Omega_{A/B}$ is exactly the same as $\Omega_{A/C}$, except we have extra relations: $db = 0$ for $b \in B$. These are precisely the images of $1 \otimes db$ on the left. \square

22.2.12. Theorem (conormal exact sequence, affine version). — Suppose B is a C -algebra, I is an ideal of B , and $A = B/I$. Then there is a natural exact sequence of A -modules

$$I/I^2 \xrightarrow{\delta} A \otimes_B \Omega_{B/C} \xrightarrow{a \otimes db \mapsto a \, db} \Omega_{A/C} \longrightarrow 0.$$

Here δ is, informally, $i \mapsto 1 \otimes di$, or more formally, $1 \otimes d: B/I \otimes_B I \rightarrow B/I \otimes_B \Omega_{B/C}$.

(You will recognize the map $A \otimes_B \Omega_{B/C} \rightarrow \Omega_{A/C}$ from the cotangent exact sequence, Theorem 22.2.9.) The proof is algebraic, so the geometric discussion thereafter may help clarify how you should really think of it.

Proof. We will identify the cokernel of $\delta: I/I^2 \rightarrow A \otimes_B \Omega_{B/C}$ with $\Omega_{A/C}$. Consider $A \otimes_B \Omega_{B/C}$. As an A -module, it is generated by db (where $b \in B$), subject to three types of relations: $dc = 0$ for $c \in \phi(C)$ (where $\phi: C \rightarrow B$ describes B as a C -algebra), additivity, and the Leibniz rule. Given any relation in B , d of that relation is 0.

Now $\Omega_{A/C}$ is defined similarly, except there are more relations in A ; these are precisely the elements of $I \subset B$. Thus we obtain $\Omega_{A/C}$ by starting out with $A \otimes_B \Omega_{B/C}$, and adding the additional relations di where $i \in I$. But this is precisely the image of δ ! (Be sure that you see how the identification of the cokernel of δ with $\Omega_{A/C}$ is precisely via the map $a \otimes db \mapsto a \, db$.) \square

We now give a geometric interpretation of the conormal exact sequence, and in particular define conormal modules/sheaves/bundles.

As with the cotangent exact sequence (Theorem 22.2.9), the conormal exact sequence is fundamentally about geometry. To motivate it, consider the sketch of Figure 22.3. In the sketch, everything is “smooth”, X is one-dimensional, Y is two-dimensional, j is the inclusion $j: X \hookrightarrow Y$, and Z is a point. Then at a point $p \in X$, the tangent space $\mathcal{T}_X|_p$ clearly injects into the tangent space of $j(p)$ in Y , and the cokernel is the normal vector space to X in Y at p . This should give an exact sequence of bundles on X :

$$0 \rightarrow \mathcal{T}_X \rightarrow j^* \mathcal{T}_Y \rightarrow \mathcal{N}_{X/Y} \rightarrow 0.$$

Dualizing this should give

$$0 \rightarrow \mathcal{N}_{X/Y}^\vee \rightarrow j^* \Omega_{Y/Z} \rightarrow \Omega_{X/Z} \rightarrow 0.$$

This is precisely what appears in the statement of the Theorem, except (i) the exact sequence in algebraic geometry is not necessarily exact on the left, and (ii) we see I/I^2 instead of $\mathcal{N}_{\text{Spec } A / \text{Spec } B}^\vee$.

22.2.13. We resolve the first issue (i) by expecting that the sequence of Theorem 22.2.12 is exact on the left in appropriately “smooth” situations, and this is indeed the case, see Theorem 22.3.7 and Remark 22.3.8. (If you enjoyed Remark 22.2.10, you might correctly guess several things. The next term on the left should be the first André-Quillen homology of A/C , so we should only need that




FIGURE 22.3. A sketch of the geometry behind the conormal exact sequence

A/C is smooth, and B should be irrelevant. Also, if $A = B/I$, then we should expect that I/I^2 is the first André-Quillen homology of A/B .)

22.2.14. Conormal modules and conormal sheaves. We resolve the second issue (ii) by declaring I/I^2 to be the **conormal module**, and in Definition 22.2.15 we will define the obvious analog as the *conormal sheaf*.

Here is some geometric intuition as to why we might want to call (the sheaf associated to) I/I^2 the conormal sheaf, which will likely confuse you, but may offer some enlightenment. First, if $\text{Spec } A$ is a closed point of $\text{Spec } B$, we expect the conormal space to be precisely the cotangent space. And indeed if $A = B/\mathfrak{m}$, the Zariski cotangent space is $\mathfrak{m}/\mathfrak{m}^2$. (We made this subtle connection in §13.1.) In particular, at some point you will develop a sense of why the conormal (=cotangent) space to the origin in $\mathbb{A}_k^2 = \text{Spec } k[x, y]$ is naturally the space of linear forms $\alpha x + \beta y$. But then consider the z -axis in $\text{Spec } k[x, y, z] = \mathbb{A}_k^3$, cut out by $I = (x, y)$. Elements of I/I^2 may be written as $\alpha(z)x + \beta(z)y$, where $\alpha(z)$ and $\beta(z)$ are polynomial. This reasonably should be the conormal space to the z -axis: as z varies, the coefficients of x and y vary. More generally, the same idea suggests that the conormal module/sheaf to any coordinate k -plane inside n -space corresponds to I/I^2 . Now consider a k -dimensional manifold X inside an n -dimensional manifold Y , with the classical topology. We can apply the same construction: if \mathcal{I} is the ideal sheaf of X in Y , then $\mathcal{I}/\mathcal{I}^2$ can be identified with the conormal sheaf (essentially the conormal vector bundle), because analytically locally $X \hookrightarrow Y$ can be identified with $\mathbb{R}^k \hookrightarrow \mathbb{R}^n$. For this reason, you might hope that in algebraic geometry, if $\text{Spec } A \hookrightarrow \text{Spec } B$ is an inclusion of something “smooth” in something “smooth”, I/I^2 should be the conormal module (or, after applying the functor \sim , the conormal sheaf). Motivated by this, we define the conormal module as I/I^2 *always*, and then notice that it has good properties (such as Theorem 22.2.12), but

take care to learn what unexpected behavior it might have when we are not in the “smooth” situation, by working out examples such as that of §22.2.7.

22.2.15. Definition. Suppose $i: X \hookrightarrow Y$ is a closed embedding of schemes cut out by ideal sheaf \mathcal{I} . Define the **conormal sheaf for a closed embedding** by $\mathcal{I}/\mathcal{I}^2$, denoted by $\mathcal{N}_{X/Y}^\vee$. (The product of quasicoherent ideal sheaves was defined in Exercise 15.3.E.) Important: We interpret $\mathcal{N}_{X/Y}^\vee$ as a quasicoherent sheaf on X , **not** (just) Y — be sure you understand why we may do so. (Exercise 18.1.E is one reason, but it may confuse the issue more than help.)

Define the **normal sheaf** as its dual $\mathcal{N}_{X/Y} := \mathcal{H}om(\mathcal{N}_{X/Y}^\vee, \mathcal{O}_X)$. This is imperfect notation, because it suggests that the dual of \mathcal{N} is always \mathcal{N}^\vee . This is not always true, as for A -modules, the natural morphism from a module to its double-dual is not always an isomorphism. (Modules for which this is true are called **reflexive**, but we won’t use this notion.)

If the normal sheaf is locally free of finite rank, then we call the associated vector bundle the **normal bundle** of X in Y .

22.2.G. EASY EXERCISE. Define the **conormal sheaf** $\mathcal{N}_{X/Y}^\vee$ (and hence the normal sheaf $\mathcal{N}_{X/Y}$) for a locally closed embedding $i: X \hookrightarrow Y$ of schemes, a quasicoherent sheaf on X . (Make sure your definition is well-defined!)

In the good situation of a regular embedding, the conormal sheaf (and hence the normal sheaf) is locally free of finite rank (or more informally, a vector bundle). In particular, the dual of \mathcal{N} is indeed \mathcal{N}^\vee . As a warm-up we deal with the important codimension 1 case.

22.2.H. EXERCISE: NORMAL BUNDLES TO EFFECTIVE CARTIER DIVISORS. Suppose $D \subset X$ is an effective Cartier divisor (§9.5.1). Show that the conormal sheaf $\mathcal{N}_{D/X}^\vee$ is $\mathcal{O}(-D)|_D$ (and in particular is an invertible sheaf), and hence that the normal sheaf is $\mathcal{O}(D)|_D$. It may be surprising that the normal sheaf should be locally free if $X \cong \mathbb{A}^2$ and D is the union of the two axes (and more generally if X is regular but D is singular), because you may be used to thinking that a “tubular neighborhood” is isomorphic to the normal bundle.

We now treat the general case.

22.2.16. Proposition. —

(a) Suppose A is Noetherian. If I is generated by a regular sequence x_1, \dots, x_r , then the map $\gamma: (A/I)^{\oplus r} \rightarrow I/I^2$ given by $(a_1, \dots, a_r) \mapsto a_1x_1 + \dots + a_rx_r$ describes I/I^2 as a free module of rank r over A/I with basis x_1, \dots, x_r .

(b) The (co)normal sheaf of a codimension r regular embedding is locally free of rank r .

Proof. (a) Clearly γ is surjective. We now show that it is injective. It suffices to show that it is injective upon localization to all points of $\text{Spec } A/I$, because injectivity is a stalk-local condition. Then in such a localization, the regular sequence in A remains regular (Exercise 9.5.D, plus the fact that the strict inclusion $(x_1, \dots, x_r) \subsetneq A$ remains strict upon localization at a prime containing I). We are thus reduced to the local question. (Caution: Make sure you see why $(I/I^2)_m = I_m/(I_m)^2$.)

Consider the a_i ’s instead as elements of A . Suppose now that $(a_1, \dots, a_r) \in \ker \gamma$; we will show that each a_i is in I . Now because x_r is not a zerodivisor on

$A/(x_1, \dots, x_{r-1})$, since $a_r x_r = 0$ in $A/(x_1, \dots, x_{r-1})$, we have $a_r \in (x_1, \dots, x_{r-1}) \subset I$. Because the roles of the a_i 's are symmetric by Theorem 9.5.6, we are done.

(b) follows immediately from (a). \square

We will soon meet a related but harder fact, that if \mathcal{I} is the ideal sheaf of a regular embedding of codimension r , then $\mathcal{I}^n/\mathcal{I}^{n+1}$ is locally free, because it is $\text{Sym}^n(\mathcal{I}/\mathcal{I}^2)$ (see Theorem 23.3.8).

22.2.17. Second definition: universal property. Here is a second definition that is important philosophically, by universal property. Of course, it is a characterization rather than a definition: by universal property nonsense, it shows that if the module exists (with the d map), then it is unique up to unique isomorphism, and then one still has to construct it to make sure that it exists.

Suppose A is a B -algebra, and M is a A -module. A **B -linear derivation of A into M** is a map $d: A \rightarrow M$ of B -modules (not necessarily a map of A -modules) satisfying the Leibniz rule: $d(fg) = f dg + g df$. As an example, suppose $B = k$, and $A = k[x]$, and $M = A$. Then d/dx is a k -linear derivation. As a second example, if $B = k$, $A = k[x]$, and $M = k$, then $(d/dx)|_0$ (the operator “evaluate the derivative at 0”) is a k -linear derivation.

A third example is $d: A \rightarrow \Omega_{A/B}$, and indeed $d: A \rightarrow \Omega_{A/B}$ is the *universal B -linear derivation of A* . Precisely, the map $d: A \rightarrow \Omega_{A/B}$ is defined by the following universal property: any other B -linear derivation $d': A \rightarrow M$ factors uniquely through d :

$$\begin{array}{ccc} A & \xrightarrow{d'} & M \\ & \searrow d \quad \nearrow f & \\ & \Omega_{A/B} & \end{array}$$

Here f is a map of A -modules. (Note again that d and d' are not necessarily maps of A -modules — they are only B -linear.) By universal property nonsense, if it exists, it is unique up to unique isomorphism. The map $d: A \rightarrow \Omega_{A/B}$ clearly satisfies this universal property, essentially by definition.

The next result connects the cotangent module $\Omega_{A/B}$ to the cotangent space at a (rational) point.

22.2.18. Proposition (the fiber of Ω at a rational point is the cotangent space). — Suppose B is a k -algebra, and $\mathfrak{m} \subset B$ is a maximal ideal with residue field k . Then there is an isomorphism of k -vector spaces $\delta: \mathfrak{m}/\mathfrak{m}^2 \xrightarrow{\sim} \Omega_{B/k} \otimes_B k$ (where the k on the right is a B -module via the isomorphism $k \xleftarrow{\sim} B/\mathfrak{m}$).

Corollary 22.3.9 will give a quite different proof, and generalize it to the case where B/\mathfrak{m} is a separable extension of k .

Proof. We instead show an isomorphism of dual vector spaces

$$\text{Hom}_k(\Omega_{B/k} \otimes_B k, k) \longrightarrow \text{Hom}_k(\mathfrak{m}/\mathfrak{m}^2, k).$$

We have (canonical) isomorphisms

$$\begin{aligned}\mathrm{Hom}_k(\Omega_{B/k} \otimes_B k, k) &= \mathrm{Hom}_B(\Omega_{B/k} \otimes_B k, k) \\ &= \mathrm{Hom}_B(\Omega_{B/k}, \mathrm{Hom}_B(k, k)) \\ &= \mathrm{Hom}_B(\Omega_{B/k}, \mathrm{Hom}_k(k, k)) \\ &= \mathrm{Hom}_B(\Omega_{B/k}, k),\end{aligned}$$

where in the right argument of $\mathrm{Hom}_B(\Omega_{B/k}, k)$, k is a B -module via its manifestation as B/\mathfrak{m} . By the universal property of $\Omega_{B/k}$ (§22.2.17), $\mathrm{Hom}_B(\Omega_{B/k}, k)$ corresponds to the k -derivations of B into $B/\mathfrak{m} \cong k$.

22.2.I. EXERCISE. Show that these are precisely the elements of $\mathrm{Hom}_k(\mathfrak{m}/\mathfrak{m}^2, k)$. (The algebra involved is essentially the same as that of Exercise 13.1.A.) \square

You can verify that this δ is the one appearing in the conormal exact sequence, Theorem 22.2.12, with $I = \mathfrak{m}$ and $A = C = k$. In fact, from the conormal exact sequence, we can immediately see that δ is a surjection, as $\Omega_{k/k} = 0$.

22.2.19. Remark. Proposition 22.2.18, in combination with the Jacobian exercise 22.2.E above, gives a second proof of Exercise 13.1.G, the Jacobian method for computing the Zariski tangent space at a k -valued point of a finite type k -scheme. Corollary 22.3.9 will extend this to the case of a separable closed point.

22.2.J. EXERCISE. Suppose X is a finite type scheme over an algebraically closed field. Show that the function from the closed points of X to $\mathbb{Z}^{\geq 0}$ given by $p \mapsto \dim T_p X$ is upper semicontinuous in the Zariski topology. (Be clear on what the Zariski topology is on the set of closed points.) Corollary 22.3.9 will allow you to extend this to all fields of characteristic 0 and all finite fields.

Depending on how your brain works, you may prefer using the first (constructive) or second (universal property) definition to do the next two exercises.

22.2.K. EXERCISE.

(a) (*pullback of differentials*) If

$$\begin{array}{ccc} A' & \longleftarrow & A \\ \uparrow & & \uparrow \\ B' & \longleftarrow & B \end{array}$$

is a commutative diagram, describe a natural homomorphism of A' -modules $A' \otimes_A \Omega_{A/B} \rightarrow \Omega_{A'/B'}$. An important special case is $B = B'$.

(b) (*differentials behave well with respect to base extension, affine case*) If furthermore the above diagram is a “tensor diagram” (i.e., $A' \cong B' \otimes_B A$, so the diagram is “co-Cartesian”) then show that $A' \otimes_A \Omega_{A/B} \rightarrow \Omega_{A'/B'}$ is an isomorphism. (Depending on how you proceed, this may be trickier than you expect.)

22.2.L. EXERCISE: LOCALIZATION (STRONGER FORM; CF. EXAMPLE 22.2.8). Suppose $\phi: B \rightarrow A$ is a map of rings, S is a multiplicative subset of A , and T is a multiplicative subset of B with $\phi(T) \subset S$, so we have the following commutative

diagram.

$$\begin{array}{ccc} S^{-1}A & \longleftarrow & A \\ \uparrow & & \uparrow \\ T^{-1}B & \longleftarrow & B \end{array}$$

Show that the pullback of differentials $S^{-1}\Omega_{A/B} \rightarrow \Omega_{S^{-1}A/T^{-1}B}$ of Exercise 22.2.K(a) is an isomorphism. (This should be believable from the intuitive picture of “vertical cotangent vectors”.) An important case is when $T = \{1\}$.

22.2.M. EXERCISE (FIELD EXTENSIONS CONTINUED, CF. EXERCISE 22.2.F).

(a) Compute $\Omega_{k(t)/k}$. (Hint: §22.2.6 followed by Exercise 22.2.L.)

(b) If K/k is **separably generated** by $t_1, \dots, t_n \in K$ (i.e., t_1, \dots, t_n form a transcendence basis, and $K/k(t_1, \dots, t_n)$ is algebraic and separable), show that $\Omega_{K/k}$ is a free K -module (i.e., vector space) with basis dt_1, \dots, dt_n . Hint: use the cotangent exact sequence (Theorem 22.2.9) for $k \hookrightarrow k(t_1, \dots, t_n) \hookrightarrow K$ to show that the dt_i span $\Omega_{K/k}$ as a K -vector space. The tricky part is showing that the dt_i are linearly independent. Do this by showing that there exists a unique map $\Omega_{K/k} \rightarrow K$ sending dt_1 to 1, and dt_i to 0 for $i > 1$. Do this first for the case where $K/k(t_1, \dots, t_n)$ is generated by one element; then where it is finitely generated; then the general case by defining it on all finitely generated subextensions, and using uniqueness to show that they “all agree”.

22.2.20. Third definition: global. We now want to globalize this definition for an arbitrary morphism of schemes $\pi: X \rightarrow Y$. We could do this “affine by affine”; we just need to make sure that the above notion behaves well with respect to “change of affine sets”. Thus a relative differential on X would be the data of, for every affine $U \subset X$, a differential of the form $\sum a_i db_i$, and on the intersection of two affine open sets $U \cap U'$, with representatives $\sum a_i db_i$ on U and $\sum a'_i db'_i$ on the second, an equality on the overlap. Instead, we take a different approach. I will give the (seemingly unintuitive) definition, then tell you how to think about it, and then get back to the definition.

Let $\pi: X \rightarrow Y$ be any morphism of schemes. Recall that $\delta: X \rightarrow X \times_Y X$ is a locally closed embedding (Proposition 11.3.1(b)). Define the **relative cotangent sheaf** $\Omega_{X/Y}$ (or Ω_π) as the conormal sheaf $\mathcal{N}_{X/X \times_Y X}^\vee$ of the diagonal (see §22.2.14 — and if $X \rightarrow Y$ is separated you needn’t even worry about Exercise 22.2.G). (Now is also as good a time as any to define the **relative tangent sheaf** $\mathcal{T}_{X/Y}$ as the dual $\mathcal{H}om(\Omega_{X/Y}, \mathcal{O}_X)$ to the relative cotangent sheaf. If we are working in the category of k -schemes, then $\Omega_{X/k}$ and $\mathcal{T}_{X/k}$ are often called the **cotangent sheaf** and **tangent sheaf** of X respectively.)

We now define $d: \mathcal{O}_X \rightarrow \Omega_{X/Y}$. Let $\text{pr}_1: X \times_Y X \rightarrow X$ and $\text{pr}_2: X \times_Y X \rightarrow X$ be the two projections. Then define $d: \mathcal{O}_X \rightarrow \Omega_{X/Y}$ on the open set U as follows:

$$df = \text{pr}_2^* f - \text{pr}_1^* f.$$

(Warning: this is not a morphism of quasicoherent sheaves on X , although it is \mathcal{O}_Y -linear in the only possible meaning of that phrase.) We will soon see that d is indeed a derivation of the sheaf \mathcal{O}_X (in the only possible meaning of the phrase), and at the same time see that our new notion of differentials agrees with our old definition on affine open sets, and hence globalizes the definition. Note that for

any open subset $U \subset X$, d induces a map

$$(22.2.20.1) \quad \Gamma(U, \mathcal{O}_X) \rightarrow \Gamma(U, \Omega_{X/Y}),$$

which we also call d , and interpret as “taking the derivative”.

22.2.21. Motivation. Before connecting this to our other definitions, let me try to convince you that this is a reasonable definition to make. (This discussion is informal and nonrigorous.) Say for example that Y is a point, and X a manifold. Then the tangent bundle $T_{X \times X}$ on $X \times X$ is $\text{pr}_1^* T_X \oplus \text{pr}_2^* T_X$, where pr_1 and pr_2 are the projections from $X \times X$ onto its two factors. Restrict this to the diagonal Δ , and look at the normal bundle exact sequence:

$$0 \rightarrow T_\Delta \rightarrow T_{X \times X}|_\Delta \rightarrow N_{\Delta/X} \rightarrow 0.$$

Now the left morphism sends v to (v, v) , so the cokernel can be interpreted as $(v, -v)$. Thus $N_{\Delta/X}$ is isomorphic to T_X . Thus we can turn this on its head: we know how to find the normal bundle (or more precisely the conormal sheaf), and we can use this to define the tangent bundle (or more precisely the cotangent sheaf).

22.2.22. Testing this out in the affine case. Let’s now see how this works for the special case $\text{Spec } A \rightarrow \text{Spec } B$. Then the diagonal morphism $\delta: \text{Spec } A \hookrightarrow \text{Spec } A \otimes_B A$ corresponds to the ideal I of $A \otimes_B A$ that is the kernel of the ring map

$$\alpha: \sum x_i \otimes y_i \rightarrow \sum x_i y_i.$$

22.2.23. The ideal I of $A \otimes_B A$ is generated by the elements of the form $1 \otimes a - a \otimes 1$. Reason: if $\alpha(\sum x_i \otimes y_i) = 0$, i.e., $\sum x_i y_i = 0$, then

$$\sum x_i \otimes y_i = \sum (x_i \otimes y_i - x_i y_i \otimes 1) = \sum x_i (1 \otimes y_i - y_i \otimes 1).$$

The derivation is $d: A \rightarrow A \otimes_B A$, $a \mapsto (1 \otimes a - a \otimes 1)$ (an element of I , taken modulo I^2). (We shouldn’t really call this “ d ” until we have verified that it agrees with our earlier definition, but we irresponsibly will anyway.)

Let’s check that d is indeed a derivation (§22.2.17). Clearly d is B -linear, so we check the Leibniz rule:

$$\begin{aligned} d(aa') - a da' - a' da &= 1 \otimes aa' - aa' \otimes 1 - a \otimes a' + aa' \otimes 1 - a' \otimes a + a' a \otimes 1 \\ &= -a \otimes a' - a' \otimes a + a' a \otimes 1 + 1 \otimes aa' \\ &= (1 \otimes a - a \otimes 1)(1 \otimes a' - a' \otimes 1) \\ &\in I^2. \end{aligned}$$

Thus by the universal property of $\Omega_{A/B}$, we have a natural morphism $\Omega_{A/B} \rightarrow I/I^2$ of A -modules.

22.2.24. Theorem. — *The natural morphism $f: \Omega_{A/B} \rightarrow I/I^2$ induced by the universal property of $\Omega_{A/B}$ is an isomorphism.*

Proof. We will show this as follows. (i) We will show that f is surjective, and (ii) we will describe $g: I/I^2 \rightarrow \Omega_{A/B}$ such that $g \circ f: \Omega_{A/B} \rightarrow \Omega_{A/B}$ is the identity (showing that f is injective).

(i) The map f sends da to $1 \otimes a - a \otimes 1$, and such elements generate I (§22.2.23), so f is surjective.

(ii) Consider the map $A \otimes_B A \rightarrow \Omega_{A/B}$ defined by $x \otimes y \mapsto x \, dy$. (This is a well-defined map, by the universal property of \otimes , see §1.3.5.) Define $g: I/I^2 \rightarrow \Omega_{A/B}$ as the restriction of this map to I . We need to check that this is well-defined, i.e., that elements of I^2 are sent to 0, i.e., we need that

$$\left(\sum_i x_i \otimes y_i \right) \left(\sum_j x'_j \otimes y'_j \right) = \sum_{i,j} x_i x'_j \otimes y_i y'_j \mapsto 0$$

when $\sum_i x_i y_i = \sum_j x'_j y'_j = 0$. But by the Leibniz rule,

$$\begin{aligned} \sum_{i,j} x_i x'_j \, d(y_i y'_j) &= \sum_{i,j} x_i x'_j y_i \, dy'_j + \sum_{i,j} x_i x'_j y'_j \, dy_i \\ &= \left(\sum_i x_i y_i \right) \left(\sum_j x'_j \, dy'_j \right) + \left(\sum_i x_i \, dy_i \right) \left(\sum_j x'_j y'_j \right) \\ &= 0. \end{aligned}$$

Then $g \circ f$ is indeed the identity, as

$$da \xrightarrow{f} 1 \otimes a - a \otimes 1 \xrightarrow{g} 1 \, da - a \, d1 = da$$

as desired. \square

22.2.N. EASY EXERCISE. Suppose $\pi: X \rightarrow Y$ is a morphism of schemes, with open subschemes $\text{Spec } A \subset X$ and $\text{Spec } B \subset Y$, with $\text{Spec } A \subset \pi^{-1}(\text{Spec } B)$. Identify $\Omega_{X/Y}|_{\text{Spec } A}$ with $\tilde{\Omega}_{A/B}$, and identify $d: \Gamma(\text{Spec } A, \mathcal{O}_X) \rightarrow \Gamma(\text{Spec } A, \Omega_{X/Y})$ with $d: A \rightarrow \Omega_{A/B}$. Thus the global construction indeed naturally “glues together” the affine construction.

We can now use our understanding of how Ω works on affine open sets to generalize previous statements to non-affine settings.

22.2.O. EXERCISE. If $U \subset X$ is an open subset, show that the map (22.2.20.1) is a derivation.

22.2.P. EXERCISE. Suppose $\pi: X \rightarrow Y$ is locally of finite type, and Y (and hence X) is locally Noetherian. Show that $\Omega_{X/Y}$ is a coherent sheaf on X . (Feel free to weaken the Noetherian hypotheses for weaker conclusions.)

22.2.Q. EXERCISE. Suppose $\pi: X \rightarrow Y$ is smooth of relative dimension n . Prove that $\Omega_{X/Y}$ is locally free of rank n . Hint: the Jacobian description of Ω , Exercise 22.2.E.

The cotangent exact sequence and the conormal exact sequence for schemes now directly follow.

22.2.25. Theorem. —

(a) (cotangent exact sequence) Suppose $X \xrightarrow{\pi} Y \xrightarrow{\rho} Z$ are morphisms of schemes. Then there is an exact sequence of quasicoherent sheaves on X

$$\pi^* \Omega_{Y/Z} \longrightarrow \Omega_{X/Z} \longrightarrow \Omega_{X/Y} \longrightarrow 0,$$

globalizing Theorem 22.2.9.

(b) (conormal exact sequence) Suppose $\rho: Y \rightarrow Z$ is a morphism of schemes, and $i: X \hookrightarrow Y$ is a closed embedding, with conormal sheaf $\mathcal{N}_{X/Y}^\vee$. Then there is an exact sequence of sheaves on X :

$$\mathcal{N}_{X/Y}^\vee \xrightarrow{\delta} i^* \Omega_{Y/Z} \longrightarrow \Omega_{X/Z} \longrightarrow 0,$$

globalizing Theorem 22.2.12.

22.2.R. EXERCISE. Prove Theorem 22.2.25. (What needs to be checked?)

You should expect these exact sequences to be left-exact as well in the presence of appropriate smoothness (see Remark 22.2.10 and §22.2.13). The following is an important special case for the cotangent exact sequence. (Theorem 22.3.7 will give a similar statement for the conormal exact sequence, for smooth varieties.)

22.2.S. EXERCISE (LEFT-EXACTNESS OF THE COTANGENT EXACT SEQUENCE WHEN π AND ρ ARE SMOOTH). Show that the cotangent exact sequence is exact on the left if π and ρ are smooth. Hint: the $(m+n) \times (m+n)$ matrix you used in Exercise 13.6.D is “block upper triangular”.

22.2.26. Pulling back relative differentials. Not surprisingly, the sheaf of relative differentials pull back, and behave well under base change.

22.2.27. Theorem (pullback of differentials). —

(a) If

$$\begin{array}{ccc} X' & \xrightarrow{\mu} & X \\ \downarrow & & \downarrow \\ Y' & \longrightarrow & Y \end{array}$$

is a commutative diagram of schemes, there is a natural homomorphism $\mu^* \Omega_{X/Y} \rightarrow \Omega_{X'/Y'}$ of quasicoherent sheaves on X' . An important special case is $Y = Y'$.

(b) (Ω behaves well under base change) If furthermore the above diagram is a Cartesian square (so $X' \cong X \times_Y Y'$) then $\mu^* \Omega_{X/Y} \rightarrow \Omega_{X'/Y'}$ is an isomorphism.

22.2.T. EXERCISE. Derive Theorem 22.2.27 from Exercise 22.2.K. (Why does the construction of Exercise 22.2.K(a) “glue well”?)

As a particular case of Theorem 22.2.27(b), the fiber of the sheaf of relative differentials is indeed the sheaf of differentials of the fiber. Thus the sheaf of differentials notion indeed “glues together” the differentials on each fiber.

22.2.U. EXERCISE. Suppose $\alpha: X \rightarrow Z$ and $\beta: Y \rightarrow Z$ are two morphisms. Describe an isomorphism $\Omega_{X \times_Z Y/Z} \xrightarrow{\sim} \alpha^* \Omega_{X/Z} \oplus \beta^* \Omega_{Y/Z}$.

22.3 Smoothness of varieties revisited

Suppose k is a field. Since §13.2.5, we have used an awkward definition of k -smoothness, and we now make our definition of k -smoothness more robust.

22.3.1. Redefinition. A k -scheme X is **k -smooth of dimension n** or **smooth of dimension n over k** if it is locally of finite type, of pure dimension n , and $\Omega_{X/k}$ is locally free of rank n . The dimension n is often omitted, and one might (possibly) want to call something smooth if it is the (scheme-theoretic) disjoint union of things smooth of various dimensions.

22.3.A. EXERCISE. Verify that this definition is equivalent to the one given in Definition 13.2.6.

22.3.B. ★★ EXERCISE (FOR THOSE WITH BACKGROUND IN COMPLEX GEOMETRY). Suppose X is a complex algebraic variety. Show that the analytification X^{an} of X (defined in Exercise 7.3.16) is smooth (in the differential-geometric sense) if and only if X is smooth (in the algebro-geometric sense, over \mathbb{C}). In this case, show that complex dimension of the complex manifold X^{an} (half the real dimension) is $\dim X$. Hint: the Jacobian criterion applies in both settings.

22.3.2. Important observation. As a consequence of our better definition of smoothness, we see that it can be checked on *any* affine cover by using the Jacobian criterion on each affine open set in the cover, as hinted in §13.2.7.

22.3.C. EXERCISE (SMOOTHNESS IS INSENSITIVE TO EXTENSION OF BASE FIELD). Prove the converse to Exercise 13.2.G.

22.3.D. IMPORTANT EXERCISE. Suppose k is perfect, and X is a finite type k -scheme. Show that X is smooth if and only if X is regular at all closed points. In particular, Theorem 13.2.10(a) holds. Hint: let $X_{\bar{k}} := X \times_k \bar{k}$. Explain why X is smooth if and only if $X_{\bar{k}}$ is smooth if and only if $X_{\bar{k}}$ is regular at all closed points if and only if X is regular at all closed points. See Exercise 13.2.O for the last step.

22.3.3. Remark: A short definition of smoothness of varieties. Exercise 22.3.D gives another pithy characterization of smoothness of varieties: *a finite type k -scheme is smooth if it is geometrically regular.*

22.3.4. Generic smoothness.

We can now verify something you may already have intuited. In positive characteristic, this is a hard theorem, in that it uses a result from commutative algebra that we have not proved.

22.3.5. Theorem (generic smoothness of varieties). — *If X is an irreducible variety over a perfect field k of dimension n , there is a dense open subset U of X such that U is smooth of dimension n .*

Hence, by Fact 13.8.2, U is regular. Theorem 26.3.1 will generalize this to smooth *morphisms*, at the expense of restricting to characteristic 0.

Proof. The $n = 0$ case is immediate, so we assume $n > 0$.

We will show that the rank at the generic point is n . Then by upper semicontinuity of the rank of a coherent sheaf (Exercise 14.3.J), it must be n in an open neighborhood of the generic point, and we are done.

We thus have to check that if K is the fraction field of a dimension n integral finite type k -scheme, i.e., (by Theorem 12.2.1) if K/k is a transcendence degree

n extension, then $\Omega_{K/k}$ is an n -dimensional vector space. But every extension of transcendence degree $n > 1$ is *separably generated* (see Exercise 22.2.M(b)): we can find n algebraically independent elements of K over k , say x_1, \dots, x_n , such that $K/k(x_1, \dots, x_n)$ is separable. (In characteristic 0, this is automatic from transcendence theory, see Exercise 12.2.A, as all finite extensions are separable. But it also holds for perfect fields in positive characteristic, see [E, Cor. A1.7] or [Mat2, Thm. 26.2].) Then $\Omega_{K/k}$ is generated by dx_1, \dots, dx_n (by Exercise 22.2.M(b)). \square

22.3.6. Left-exactness of the conormal exact sequence for embeddings of smooth varieties.

As described in §22.2.13, we expect the conormal exact sequence to be exact on the left in appropriately “smooth” situations.

22.3.7. Theorem (conormal exact sequence for smooth varieties). — Suppose $i: X \hookrightarrow Y$ is a closed embedding of smooth varieties over a field k , with conormal sheaf $\mathcal{N}_{X/Y}^\vee$. Then the conormal exact sequence (Theorem 22.2.25(b)) is exact on the left:

$$0 \longrightarrow \mathcal{N}_{X/Y}^\vee \xrightarrow{\delta} i^* \Omega_{Y/k} \longrightarrow \Omega_{X/k} \longrightarrow 0$$

is exact.

By dualizing, i.e., applying $\mathcal{H}om(\cdot, \mathcal{O}_X)$, we obtain the **normal exact sequence**

$$0 \rightarrow \mathcal{T}_{X/k} \rightarrow \mathcal{T}_{Y/k}|_X \rightarrow \mathcal{N}_{X/Y} \rightarrow 0$$

which is geometrically more intuitive (see Figure 22.3 and the discussion after the proof of Theorem 22.2.12).

22.3.8. Remark. The conormal sequence is exact on the left in even more general circumstances. Essentially all that is required is appropriate smoothness of $i \circ g: X \rightarrow Z$, see [Stacks, tag 06B7], and [Stacks, tag 06BB] for related facts.

Proof. We use the fact that smooth k -varieties are regular at their closed points, which we have only proved for perfect fields, but which we shall later prove fully (see §13.2.9).

Let \mathcal{I} be the ideal sheaf of $i: X \hookrightarrow Y$. By Exercise 13.2.L(b), i is a regular embedding in the open neighborhood of all closed points, and hence everywhere. Thus by Proposition 22.2.16(b), $\mathcal{I}/\mathcal{I}^2$ is locally free of rank r . By Exercise 6.5.F, the associated points of $\ker \delta$ are a subset of the associated points of $\mathcal{I}/\mathcal{I}^2$, so to show that $\ker \delta = 0$, it suffices to check this at the associated points of $\mathcal{I}/\mathcal{I}^2$, which are precisely the associated points of X (as $\mathcal{I}/\mathcal{I}^2$ is locally free).

As stated above, X is regular at closed points, hence (by Theorem 13.2.13) reduced (the nonreduced locus is closed, §9.4.10, and hence if nonempty would necessarily contain a closed point). Thus we need only check at the generic points of (the irreducible components of) X (as X has no embedded points, see §6.5.2 or Exercise 6.5.R). But this involves checking the left-exactness of a right-exact sequence of vector spaces, and as the dimension of the left (the codimension of X in Y by Proposition 22.2.16(b)) is precisely the difference of the other two, we are done. \square

22.3.9. Important Corollary. — Suppose Y is a smooth k -variety, and $q \in Y$ is a closed point whose residue field $\kappa(q)$ is separable over k . (This is automatic if $\text{char } k = 0$ or if k is a finite field.) Then the conormal exact sequence for $q \hookrightarrow Y$ yields an isomorphism of the Zariski cotangent space of Y at q with the fiber of $\Omega_{Y/k}$ at q .

Proof. Apply Theorem 22.3.7 with $X = q$. Note that q is indeed a smooth k -variety (using separability of $\kappa(q)/k$), and that $\Omega_{q/k} = 0$. \square

22.3.10. Remark. This result generalizes Proposition 22.2.18 to separable closed points. As described after the statement of Proposition 22.2.18, this has a number of consequences. For example, it extends the Jacobian description of the Zariski tangent space to separable closed points.

22.3.E. EXERCISE (CF. §13.2.16). Suppose p is a closed point of a k -variety X , with residue field $\kappa(p)$ that is separable over k . Define $\pi: X_{\bar{k}} := X \times_k \bar{k} \rightarrow X$ by base change from $\text{Spec } \bar{k} \rightarrow \text{Spec } k$. Suppose $q \in \pi^{-1}(p)$. Show that $X_{\bar{k}}$ is regular at q if and only if X is regular at p .

22.3.11. Proving Theorem 13.8.3, that the localization of regular local rings are regular in the case of varieties over perfect fields. Theorems 22.3.5 and 22.3.7 allow us to prove Theorem 13.8.3 (an important case of Fact 13.8.2).

22.3.12. Proof of Theorem 13.8.3. Suppose Y is a variety over a perfect field k that is regular at its closed points (so Y is smooth, by Exercise 22.3.D), and let η be a point of Y . We will show that Y is regular at η . Let $X = \bar{\eta}$. By Theorem 22.3.5, X contains a dense (=nonempty) open subset of smooth points. By shrinking Y by discarding the points of X outside that open subset, we may assume X is smooth.

Then Theorem 22.3.7 (exactness of the conormal exact sequence for smooth varieties) implies that $\mathcal{I}/\mathcal{I}^2$ is a locally free sheaf of rank $\text{codim}_{X/Y} = \dim \mathcal{O}_{Y,\eta}$.

22.3.F. EXERCISE. Let \mathfrak{m} be the maximal ideal of $\mathcal{O}_{Y,\eta}$. Identify the stalk of $\mathcal{I}/\mathcal{I}^2$ at the generic point η of X with $\mathfrak{m}/\mathfrak{m}^2$. Conclude the proof of Theorem 13.8.3. \square

22.3.G. TRICKY EXERCISE (LOCALIZATION OF REGULAR LOCAL RINGS OF VARIETIES ARE REGULAR, PROMISED JUST AFTER THEOREM 13.8.3). Suppose (A, \mathfrak{m}) is a regular local ring that is the localization of a finitely generated k -algebra, where k is perfect. Show that the localization of A at a prime is also a regular local ring.

22.4 Examples

The examples below are organized by topic, not by difficulty.

22.4.1. The geometric genus of a curve. A regular irreducible projective curve C (over a field k) has **geometric genus** $h^0(C, \Omega_{C/k})$. (This will be generalized to higher dimension in §22.5.3.) This is always finite, as $\Omega_{C/k}$ is coherent (Exercise 22.2.P), and coherent sheaves on projective k -schemes have finite-dimensional spaces of sections (Theorem 19.1.3(i)). Sadly, this isn't really a new invariant. We

will see in Exercise 22.5.C that this agrees with our earlier definition of genus, i.e., $h^0(C, \Omega_{C/k}) = h^1(C, \mathcal{O}_C)$.

Now let's compute some differentials!

22.4.2. The projective line. As an important first example, consider \mathbb{P}_k^1 , with the usual projective coordinates x_0 and x_1 . As usual, the first patch corresponds to $x_0 \neq 0$, and is of the form $\text{Spec } k[x_{1/0}]$ where $x_{1/0} = x_1/x_0$. The second patch corresponds to $x_1 \neq 0$, and is of the form $\text{Spec } k[x_{0/1}]$ where $x_{0/1} = x_0/x_1$.

Both patches are isomorphic to \mathbb{A}_k^1 , and $\Omega_{\mathbb{A}_k^1} \cong \mathcal{O}_{\mathbb{A}_k^1}$. (More precisely, $\Omega_{k[x]/k} = k[x] dx$.) Thus $\Omega_{\mathbb{P}_k^1}$ is an invertible sheaf (a line bundle). The invertible sheaves on \mathbb{P}_k^1 are of the form $\mathcal{O}(m)$. So which invertible sheaf is $\Omega_{\mathbb{P}_k^1}$?

Let's take a section, $dx_{1/0}$ on the first patch. It has no zeros or poles there, so let's check what happens on the other patch. As $x_{1/0} = 1/x_{0/1}$, we have $dx_{1/0} = -(1/x_{0/1}^2) dx_{0/1}$. Thus this section has a double pole where $x_{0/1} = 0$. Hence $\Omega_{\mathbb{P}_k^1/k} \cong \mathcal{O}(-2)$. As a consequence, we have shown that the geometric genus of \mathbb{P}_k^1 is 0.

Note that the above argument works equally well if k were replaced by \mathbb{Z} : our theory of Weil divisors and line bundles of Chapter 15 applies ($\mathbb{P}_{\mathbb{Z}}^1$ is factorial), so the previous argument essentially without change shows that $\Omega_{\mathbb{P}_{\mathbb{Z}}^1/\mathbb{Z}} \cong \mathcal{O}(-2)$. And because Ω behaves well with respect to base change (Theorem 22.2.27(b)), and any scheme maps to $\text{Spec } \mathbb{Z}$, this implies that $\Omega_{\mathbb{P}_B^1/B} \cong \mathcal{O}_{\mathbb{P}_B^1}(-2)$ for *any* base scheme B .

(Also, as suggested by §19.5.2, this shows that $\Omega_{\mathbb{P}_k^1/k}$ is the dualizing sheaf for \mathbb{P}_k^1 ; see also Example 19.5.4. But given that we haven't yet proved Serre duality, this isn't so meaningful.)

Side Remark: the fact that the degree of the tangent bundle is 2 is related to the "Hairy Ball Theorem" (the dimension 2 case of [Hat, Thm. 2.28]).

22.4.3. Hyperelliptic curves. Throughout this discussion of hyperelliptic curves, we suppose that $k = \bar{k}$ and $\text{char } k \neq 2$, so we may apply the discussion of §20.5. Consider a double cover $\pi: C \rightarrow \mathbb{P}_k^1$ by a regular projective curve C , branched over $2g + 2$ distinct points. We will use the explicit coordinate description of hyperelliptic curves of (20.5.2.1). By Theorem 20.5.1, C has genus g .

22.4.A. EXERCISE: DIFFERENTIALS ON HYPERELLIPTIC CURVES. What is the degree of the invertible sheaf $\Omega_{C/k}$? (Hint: let x be a coordinate on one of the coordinate patches of \mathbb{P}_k^1 . Consider π^*dx on C , and count poles and zeros. Use the explicit coordinates of §20.5. You should find that π^*dx has $2g + 2$ zeros and 4 poles (counted with multiplicity), for a total of $2g - 2$.) Doing this exercise will set you up well for the Riemann-Hurwitz formula, in §22.7.

22.4.B. EXERCISE. Show that $h^0(C, \Omega_{C/k}) = g$ (and hence that the geometric genus of C is g) as follows.

- Show that $\frac{dx}{y}$ is a (regular) differential on $\text{Spec } k[x, y]/(y^2 - f(x))$ (i.e., an element of $\Omega_{(k[x, y]/(y^2 - f(x)))/k}$).
- Show that for $0 \leq i < g$, $x^i(dx)/y$ extends to a global differential ω_i on C (i.e., with no poles).

(c) Show that the ω_i ($0 \leq i < g$) are linearly independent differentials. (Hint: Show that the valuation of ω_i at $x = 0$ is i . If $\omega := \sum_{j=0}^{g-1} a_j \omega_j$ is a nontrivial linear combination, with $a_j \in k$, and $a_i \neq 0$, show that the valuation of ω at $x = 0$ is i , and hence $\omega \neq 0$.)

★ (d) Show that the ω_i form a basis for the differentials. (Hint: consider the order of zeros of the ω_i at $\pi^{-1}(\infty)$.)

22.4.C. ★ EXERCISE (TOWARD SERRE DUALITY). (You may later see this as an example of Serre duality in action.)

(a) Show that $h^1(C, \Omega_{C/k}) = 1$. Interpret a generator of $H^1(C, \Omega_{C/k})$ as $x^{-1} dx$. (In particular, the pullback map $H^1(\mathbb{P}^1, \Omega_{\mathbb{P}^1/k}) \rightarrow H^1(C, \Omega_{C/k})$ is an isomorphism.)

(b) Describe a natural perfect pairing

$$H^0(C, \Omega_{C/k}) \times H^1(C, \mathcal{O}_C) \rightarrow H^1(C, \Omega_{C/k}).$$

In terms of our explicit coordinates, you might interpret it as follows. Recall from the proof of the hyperelliptic Riemann-Hurwitz formula (Theorem 20.5.1) that $H^1(C, \mathcal{O}_C)$ can be interpreted as

$$\left\langle \frac{y}{x}, \frac{y}{x^2}, \dots, \frac{y}{x^g} \right\rangle.$$

Then the pairing

$$\left\langle \frac{dx}{y}, \frac{x dx}{y}, \dots, x^{g-1} \frac{dx}{y} \right\rangle \times \left\langle \frac{y}{x}, \frac{y}{x^2}, \dots, \frac{y}{x^g} \right\rangle \rightarrow \langle x^{-1} dx \rangle$$

is basically “multiply and read off the $x^{-1} dx$ term”. Or in fancier terms: “multiply and take the residue”. (You may want to compare this to Example 19.5.4.)

22.4.4. Discrete valuation rings. The following exercise is used in the proof of the Riemann-Hurwitz formula, §22.7.

22.4.D. EXERCISE. Suppose that the discrete valuation ring (A, \mathfrak{m}, k) is a localization of a finitely generated k -algebra. Let t be a uniformizer of A . Show that the differentials are free of rank one and generated by dt , i.e., $\Omega_{A/k} = A dt$, as follows. (It is also possible to show this using the ideas from §22.5.)

(a) Show that $\Omega_{A/k}$ is a finitely generated A -module.

(b) Show that $\langle dt \rangle = \Omega_{A/k}$, i.e., that $\times dt: A \rightarrow \Omega_{A/k}$ is a surjection, as follows. Let π be the projection $A \rightarrow A/\mathfrak{m} = k$, so for $a \in A$, $a - \pi(a) \in \mathfrak{m}$. Define $\sigma(a) = (a - \pi(a))/t$. Show that $\Omega_{A/k} = \langle dt \rangle + \mathfrak{m}\Omega_{A/k}$, using the fact that for every $a \in A$,

$$da = \sigma(a) dt + t d\sigma(a).$$

Apply Nakayama’s Lemma version 3 (Exercise 8.2.G) to $\langle dt \rangle \subset \Omega_{A/k}$. (This argument, with essentially no change, can be used to show that if (A, \mathfrak{m}, k) is a localization of a finitely generated algebra over k , and t_1, \dots, t_n generate \mathfrak{m} , then dt_1, \dots, dt_n generate $\Omega_{A/k}$.)

(c) By part (b), $\Omega_{A/k}$ is a principal A -module. Show that $\times dt$ is an injection as follows. By the classification of finite generated modules over discrete valuation rings (Remark 13.5.14), it suffices to show that $t^m dt \neq 0$ for all m . The surjection $A \rightarrow A/(t^N)$ induces a map $\Omega_{A/k} \rightarrow \Omega_{(A/(t^N))/k}$, so it suffices to show that

$t^m dt$ is nonzero in $\Omega_{(A/(t^N))/k}$. Show that $A/(t^N) \cong k[t]/(t^N)$. The usual differentiation rule for polynomials gives a map $k[t]/(t^N) \rightarrow k[t]/(t^{N-1})$ which is a derivation of $k[t]/(t^N)$ over k , and $t^m dt$ will not map to 0 so long as N is sufficiently large. Put the pieces together and complete the proof. (An extension of these ideas can show that if (A, \mathfrak{m}, k) is a localization of a finitely generated algebra over k that is a regular local ring, then $\Omega_{A/k}$ is free of rank $\dim A$.)

22.4.5. Projective space and the Euler exact sequence. We next examine the differentials of projective space \mathbb{P}_k^n , or more generally \mathbb{P}_A^n where A is an arbitrary ring. As projective space is covered by affine open sets of the form \mathbb{A}^n , on which the differentials form a rank n locally free sheaf, $\Omega_{\mathbb{P}_A^n/A}$ is also a rank n locally free sheaf.

22.4.6. Theorem (the Euler exact sequence). — *The sheaf of differentials $\Omega_{\mathbb{P}_A^n/A}$ satisfies the following exact sequence*

$$0 \longrightarrow \Omega_{\mathbb{P}_A^n/A} \longrightarrow \mathcal{O}_{\mathbb{P}_A^n}(-1)^{\oplus(n+1)} \xrightarrow{\Phi} \mathcal{O}_{\mathbb{P}_A^n} \longrightarrow 0.$$

This is handy, because you can get a hold of $\Omega_{\mathbb{P}_A^n/A}$ in a concrete way. See Exercise 22.5.Q for an application. By dualizing this exact sequence, we have an exact sequence $0 \rightarrow \mathcal{O}_{\mathbb{P}_A^n} \rightarrow \mathcal{O}_{\mathbb{P}_A^n}(1)^{\oplus(n+1)} \rightarrow \mathcal{T}_{\mathbb{P}_A^n/A} \rightarrow 0$.

22.4.7. ★ Proof of Theorem 22.4.6. (What is really going on in this proof is that we consider those differentials on $\mathbb{A}_A^{n+1} \setminus \{0\}$ that are pullbacks of differentials on \mathbb{P}_A^n . For a different explanation, in terms of the Koszul complex, see [E, §17.5].)

We first describe a map $\phi: \mathcal{O}(-1)^{\oplus(n+1)} \rightarrow \mathcal{O}$, and later identify the kernel with $\Omega_{\mathbb{P}_A^n/A}$. The map is given by

$$\phi: (s_0, s_1, \dots, s_n) \mapsto x_0 s_0 + x_1 s_1 + \dots + x_n s_n.$$

(Do you see why?) You should think of this as a “degree 1” map, as each x_i has degree 1.

22.4.8. Remark. The dual $\phi^\vee: \mathcal{O} \rightarrow \mathcal{O}(1)^{\oplus(n+1)}$ of ϕ gives a map $\mathbb{P}^n \rightarrow \mathbb{P}^n$ via Important Theorem 17.4.1 on maps to projective space. This map is the identity, and this is one way of describing ϕ in a “natural” (coordinate-free) manner.

22.4.E. EASY EXERCISE. Show that ϕ is surjective, by checking on the open set $D(x_i)$. (There is a one-line solution.)

Now we must identify the kernel of ϕ with the differentials, and we can do this on each $D(x_i)$, so long as we do it in a way that works simultaneously for each open set. So we consider the open set U_0 , where $x_0 \neq 0$, and we have the usual coordinates $x_{j/0} = x_j/x_0$ ($1 \leq j \leq n$). Given a differential

$$f_1(x_{1/0}, \dots, x_{n/0}) dx_{1/0} + \dots + f_n(x_{1/0}, \dots, x_{n/0}) dx_{n/0}$$

we must produce $n + 1$ sections of $\mathcal{O}(-1)$. As motivation, we just look at the first term, and pretend that the projective coordinates are actual coordinates.

$$\begin{aligned} f_1 dx_{1/0} &= f_1 d(x_1/x_0) \\ &= f_1 \frac{x_0 dx_1 - x_1 dx_0}{x_0^2} \\ &= -\frac{x_1}{x_0^2} f_1 dx_0 + \frac{f_1}{x_0} dx_1 \end{aligned}$$

Note that x_0 times the “coefficient of dx_0 ” plus x_1 times the “coefficient of dx_1 ” is 0, and also both coefficients are of homogeneous degree -1 . Motivated by this, we take:

$$(22.4.8.1) \quad f_1 dx_{1/0} + \cdots + f_n dx_{n/0} \mapsto \left(-\frac{x_1}{x_0^2} f_1 - \cdots - \frac{x_n}{x_0^2} f_n, \frac{f_1}{x_0}, \frac{f_2}{x_0}, \dots, \frac{f_n}{x_0} \right)$$

Note that over U_0 , this indeed gives an injection of $\Omega_{\mathbb{P}^n/\Lambda}$ to $\mathcal{O}(-1)^{\oplus(n+1)}$ that surjects onto the kernel of $\mathcal{O}(-1)^{\oplus(n+1)} \rightarrow \mathcal{O}_X$ (if (g_0, \dots, g_n) is in the kernel, take $f_i = x_0 g_i$ for $i > 0$).

Let’s make sure this construction, applied to two different coordinate patches (say U_0 and U_1) gives the same answer. (This verification is best ignored on a first reading.) Note that

$$\begin{aligned} f_1 dx_{1/0} + f_2 dx_{2/0} + \cdots &= f_1 d\frac{1}{x_{0/1}} + f_2 d\frac{x_{2/1}}{x_{0/1}} + \cdots \\ &= -\frac{f_1}{x_{0/1}^2} dx_{0/1} + \frac{f_2}{x_{0/1}} dx_{2/1} - \frac{f_2 x_{2/1}}{x_{0/1}^2} dx_{0/1} + \cdots \\ &= -\frac{f_1 + f_2 x_{2/1} + \cdots}{x_{0/1}^2} dx_{0/1} + \frac{f_2 x_{2/1}}{x_{0/1}} dx_{2/1} + \cdots \end{aligned}$$

Under this map, the $dx_{2/1}$ term goes to the second factor (where the factors are indexed 0 through n) in $\mathcal{O}(-1)^{\oplus(n+1)}$, and yields f_2/x_0 as desired (and similarly for $dx_{j/1}$ for $j > 2$). Also, the $dx_{0/1}$ term goes to the “zero” factor, and yields

$$\left(\sum_{i=1}^n f_i (x_i/x_1) / (x_0/x_1)^2 \right) / x_1 = f_i x_i / x_0^2$$

as desired. Finally, the “first” factor must be correct because the sum over i of x_i times the i th factor is 0.

22.4.F. EXERCISE. Finish the proof of Theorem 22.4.6, by verifying that this map $\Omega_{\mathbb{P}^n/\Lambda} \rightarrow \mathcal{O}_{\mathbb{P}^n}(-1)^{\oplus(n+1)}$ identifies $\Omega_{\mathbb{P}^n/\Lambda}$ with $\ker \phi$. □

22.4.G. EXERCISE (PROMISED IN §19.5.7). Show that $h^1(\mathbb{P}_k^2, \Omega_{\mathbb{P}_k^2/k}) > 0$. Show that $\Omega_{\mathbb{P}_k^2/k}$ is not the direct sum of line bundles, and can’t even be *filtered* into line bundles. Show that Theorem 19.5.6 (that all finite rank vector bundles on \mathbb{P}^1 split into line bundles) cannot be extended to rank 2 vector bundles on \mathbb{P}_k^2 .

22.4.9. Generalizations of the Euler exact sequence. Generalizations of the Euler exact sequence are quite useful. We won’t use them later, so no proofs will be given.

First, the argument generalizes readily if $\text{Spec } A$ is replaced by an arbitrary base scheme. The Euler exact sequence generalizes further in a number of ways. As a first step, suppose \mathcal{V} is a rank $n + 1$ locally free sheaf (or vector bundle) on a scheme X . Suppose $\pi : \mathbb{P}\mathcal{V} \rightarrow X$ is the corresponding projective bundle. Then $\Omega_{\mathbb{P}\mathcal{V}/X}$ sits in an Euler exact sequence:

$$0 \rightarrow \Omega_{\mathbb{P}\mathcal{V}/X} \rightarrow \mathcal{O}_{\mathbb{P}\mathcal{V}}(-1) \otimes \pi^* \mathcal{V}^\vee \rightarrow \mathcal{O}_{\mathbb{P}\mathcal{V}} \rightarrow 0$$

If $\pi: \mathbb{P}\mathcal{V} \rightarrow X$, the map $\mathcal{O}(-1) \otimes \mathcal{V}^\vee \rightarrow \mathcal{O}_X$ is induced by $\mathcal{V}^\vee \otimes \pi_* \mathcal{O}(1) \cong (\mathcal{V}^\vee \otimes \mathcal{V}) \otimes \mathcal{O}_X \rightarrow \mathcal{O}_X$, where $\mathcal{V}^\vee \otimes \mathcal{V} \rightarrow \mathcal{O}_X$ is the trace map (§14.3.1). (You may wish to compare this to Remark 22.4.8.)

It is not obvious that this is useful, but we have already implicitly seen it in the case of \mathbb{P}^1 -bundles over curves, in Exercise 21.2.J, where the normal bundle to a section was identified in this way.

22.4.10. ★★ Generalization to the Grassmannian. For another generalization, fix a base field k , and let $G(m, n+1)$ be the space of sub-vector spaces of dimension m in an $(n+1)$ -dimensional vector space V (the Grassmannian, §17.7). Over $G(m, n+1)$ we have a short exact sequence of locally free sheaves

$$0 \rightarrow \mathcal{S} \rightarrow \mathcal{O}_{G(m, n+1)} \otimes V^\vee \rightarrow \mathcal{Q} \rightarrow 0$$

where $\mathcal{O}_{G(m, n+1)} \otimes V^\vee$ is the “trivial bundle whose fibers are V^\vee ” (do you understand what that means?), and \mathcal{S} is the “universal vector bundle” or “tautological vector (sub)bundle”. Then there is a canonical isomorphism

$$(22.4.10.1) \quad \Omega_{G(m, n+1)/k} \xrightarrow{\sim} \mathcal{H}om(\mathcal{Q}, \mathcal{S}).$$

22.4.H. EXERCISE. Recall that in the case of projective space, i.e., $m = 1$, $\mathcal{S} = \mathcal{O}(-1)$ (Exercise 18.1.H). Verify (22.4.10.1) in this case using the Euler exact sequence (Theorem 22.4.6).

22.4.I. EXERCISE. Prove (22.4.10.1), and explain how it generalizes Exercise 21.2.I. (The hint to Exercise 21.2.I may help.)

This Grassmannian fact generalizes further to Grassmannian bundles, and to flag varieties, and to flag bundles.

22.5 Understanding smooth varieties using their cotangent bundles

In this section, we construct birational invariants of varieties over algebraically closed fields (such as the geometric genus), motivate the notion of an unramified morphism, show that varieties are “smooth almost everywhere”, and get a first glimpse of Hodge theory.

22.5.1. The geometric genus, and other birational invariants from i-forms.

Suppose X is a projective scheme over k . Then $h^0(X, \Omega_{X/k})$ is an invariant of X . In fact it, and related invariants, are *birational invariants* if X is smooth, as shown in the following Exercise 22.5.A. We first define the **sheaf of (relative) i-forms**

$\Omega_{X/Y}^i := \wedge^i \Omega_{X/Y}$. Sections of $\Omega_{X/Y}^i$ (over some open set) are called **(relative) i-forms** (over that open set).

22.5.2. Joke. Old Macdonald had a form; $e_i \wedge e_i = 0$.

22.5.A. EXERCISE ($h^0(X, \Omega_{X/k}^i)$ ARE BIRATIONAL INVARIANTS). Suppose X and X' are birational irreducible smooth projective k -varieties. Show (for each i) that $H^0(X, \Omega_{X/k}^i) \cong H^0(X', \Omega_{X'/k}^i)$. Hint: fix a birational map $\phi: X \dashrightarrow X'$. By Exercise 17.5.B, the complement of the domain of definition U of ϕ is codimension at least 2. By pulling back i -forms from X' to U , we get a map $\phi^*: H^0(X', \Omega_{X'/k}^i) \rightarrow H^0(U, \Omega_{X/k}^i)$. Use Algebraic Hartogs's Lemma 12.3.15 and the fact that Ω^i is locally free to show the map extends to a map $\phi^*: H^0(X', \Omega_{X'/k}^i) \rightarrow H^0(X, \Omega_{X/k}^i)$. If $\psi: X' \dashrightarrow X$ is the inverse rational map, we similarly get a map $\psi^*: H^0(X, \Omega_{X/k}^i) \rightarrow H^0(X', \Omega_{X'/k}^i)$. Show that ϕ^* and ψ^* are inverses by showing that each composition is the identity on a dense open subset of X or X' .

22.5.3. The canonical bundle \mathcal{K}_X and the geometric genus $p_g(X)$. If X is a dimension n smooth k -variety, the invertible sheaf (or line bundle) $\det \Omega_{X/k} = \Omega_{X/k}^n$ (the sheaf of “algebraic volume forms”) has particular importance, and is called the **canonical (invertible) sheaf**, or the **canonical (line) bundle**. It is denoted \mathcal{K}_X (or $\mathcal{K}_{X/k}$). As mentioned in §19.5.2, if X is projective, then \mathcal{K}_X is the dualizing sheaf ω_X appearing in the statement of Serre duality, something we will establish in §31.4 (see Desideratum 31.1.1).

22.5.B. EXERCISE (THE ADJUNCTION FORMULA FOR \mathcal{K}_X). Suppose X is a smooth variety, and Z is a smooth subvariety of X . Show that

$$\mathcal{K}_Z \cong \mathcal{K}_X|_Z \otimes \det \mathcal{N}_{Z/X}.$$

(Hint: apply Exercise 14.2.W to Theorem 22.3.7.) In particular, by Exercise 22.2.H, if Z is codimension 1, then

$$\mathcal{K}_Z \cong (\mathcal{K}_X \otimes \mathcal{O}_X(Z))|_Z,$$

which is more commonly and compactly written as $\mathcal{K}_X(Z)|_Z$. The adjunction formula is often used inductively, for complete intersections; see Exercise 22.5.H for an example. (And see Exercise 21.2.B(a) and §31.4 for other versions of the adjunction formula.)

22.5.4. Definition. If X is a projective (or even proper) smooth k -variety, the birational invariant $h^0(X, \mathcal{K}_X) = h^0(X, \Omega_{X/k}^n)$ has particular importance. It is called the **geometric genus**, and is denoted $p_g(X)$. We saw this in the case of curves in §22.4.1. If X is an irreducible variety that is *not* smooth or projective, the phrase geometric genus refers to $h^0(X', \mathcal{K}_{X'})$ for some smooth projective X' *birational* to X . (By Exercise 22.5.A, this is independent of the choice of X' .) For example, if X is an irreducible reduced projective curve over k , the geometric genus is the geometric genus of the normalization of X . (But in higher dimension, it is not clear if there exists such an X' . It is a nontrivial fact that this is true in characteristic 0, and it is not yet known in positive characteristic; see Remark 23.4.6.)

It is a miracle that for a complex curve the geometric genus is the same as the topological genus and the arithmetic genus. We will connect the geometric genus

to the topological genus in our discussion of the Riemann-Hurwitz formula soon (Exercise 22.7.I). We connect the geometric genus to the arithmetic genus in the following exercise.

22.5.C. EASY EXERCISE. Assuming Miracle 19.5.2 (that the canonical bundle is Serre-dualizing), show that the geometric genus of an irreducible smooth projective curve over $k = \bar{k}$ equals its arithmetic genus.

22.5.D. EXERCISE. Suppose Z is a regular degree d surface in $\mathbb{P}_{\bar{k}}^3$. Compute the geometric genus $p_g(Z)$ of Z . Show that no regular quartic surface in $\mathbb{P}_{\bar{k}}^3$ is rational (i.e., birational to $\mathbb{P}_{\bar{k}}^2$, Definition 7.5.4). (Such quartic surfaces are examples of *K3 surfaces*, see Exercise 22.5.I.)

22.5.5. Important classes of varieties: Fano, Calabi-Yau, general type.

Suppose X is a smooth projective k -variety. Then X is said to be **Fano** if \mathcal{K}_X^\vee is ample, and **Calabi-Yau** if $\mathcal{K}_X \cong \mathcal{O}_X$. (Caution: there are other definitions of Calabi-Yau.)

22.5.E. EXERCISE. For all $j \geq 0$, the j th **plurigenus** of a smooth projective k -variety is $h^0(X, \mathcal{K}_X^{\otimes j})$. Show that the j th plurigenus is a birational invariant.

(By contrast, $h^0(X, \mathcal{K}_X^{\otimes j})$ is not a birational invariant if $j < 0$. For example, you can show that $h^0(X, \mathcal{K}_X^\vee)$ differs for $X = \mathbb{P}_k^2$ and $X = \text{Bl}_p \mathbb{P}_k^2$.)

The **Kodaira dimension** of X , denoted $\kappa(X)$, tracks the rate of growth of the plurigenera. It is the smallest k such that $h^0(X, \mathcal{K}_X^{\otimes j})/j^k$ is bounded (as j varies through the positive integers), except that if all the plurigenera $h^0(X, \mathcal{K}_X^{\otimes j})$ are 0, we say $\kappa(X) = -1$. It is a nontrivial fact that the Kodaira dimension always exists, and that it is the maximum of the dimensions of images of the j th “pluricanonical rational maps”. The latter fact implies that the Kodaira dimension is an integer between -1 and $\dim X$ inclusive. Exercise 22.5.E shows that the Kodaira dimension is a birational invariant. If $\kappa(X) = \dim X$, we say that X is of **general type**. (For more information on the Kodaira dimension, see [II, §10.5].)

22.5.F. EXERCISE. Show that if \mathcal{K}_X is ample then X is of general type.

22.5.G. EXERCISE. Show that a smooth geometrically irreducible projective curve (over a field k) is Fano (resp. Calabi-Yau, general type) if its genus is 0 (resp. 1, greater than 1).

The “trichotomy” of Exercise 22.5.G is morally the reason that curves behave differently depending on which of the three classes they lie in (see §20.6.1). In some sense, important aspects of this trichotomy extend to higher dimension, which is part of the reason for making these definitions. We now explore this trichotomy in the case of complete intersections.

22.5.H. EXERCISE. Suppose X is a smooth complete intersection in \mathbb{P}_k^N of hypersurfaces of degree d_1, \dots, d_n , where k is algebraically closed.

(a) Show that $\mathcal{K}_X \cong \mathcal{O}(-N - 1 + d_1 + \dots + d_n)|_X$.

(b) Show that $-N - 1 + d_1 + \dots + d_n$ is negative (resp. zero, positive) then X is Fano (resp. Calabi-Yau, general type).

(c) Find all possible values of N and d_1, \dots, d_n where X is Calabi-Yau of dimension at most 3. Notice how small this list is.

22.5.I. EXERCISE. A **K3 surface** over a field k is a proper smooth geometrically connected Calabi-Yau surface X over k such that $H^1(X, \mathcal{O}_X) = 0$. (Weil has written that K3 surfaces were named in honor of Kummer, Kähler, and Kodaira, and the mountain K2, [BPV, p. 288].) Prove that the dimension 2 (smooth) Calabi-Yau complete intersections of Exercise 22.5.H(c) are all K3 surfaces.

22.5.6. Tantalizing side remark. If you compare the degrees of the hypersurfaces cutting out complete intersection K3 surfaces (in Exercise 22.5.H(c)), with the degrees of hypersurfaces cutting out complete intersection canonical curves (see §20.8.2 and Exercise 20.8.E), you will notice a remarkable coincidence. Of course this is not a coincidence at all.

22.5.J. EXERCISE (NOT EVERYTHING FITS INTO THIS TRICHOTOMY). Suppose C is a smooth projective irreducible complex curve of genus greater than 1. Show that $C \times_{\mathbb{C}} \mathbb{P}_{\mathbb{C}}^1$ is neither Fano, nor Calabi-Yau, nor general type. Hint: Exercise 22.2.U.

22.5.7. ★ Kodaira vanishing.

The Kodaira vanishing theorem is an important result that is an important tool in a number of different areas of algebraic geometry. We state it for the sake of culture, but do not prove it (and hence will not use it).

22.5.8. The Kodaira Vanishing Theorem. — Suppose k is a field of characteristic 0, and X is a smooth projective k -variety. Then for any ample invertible sheaf \mathcal{L} , $H^i(X, \mathcal{K}_X \otimes \mathcal{L}) = 0$ for $i > 0$.

22.5.9. The restriction on characteristic is necessary; Raynaud gave an example where Kodaira vanishing fails in positive characteristic, [Ra]. The original proof is by proving the complex-analytic version, and then transferring it into a complex algebraic statement using Serre's GAGA Theorem [Se3]. The general characteristic 0 case can be reduced to \mathbb{C} — any reduction of this sort is often called (somewhat vaguely) an application of the *Lefschetz principle*. (For a precise formulation of the Lefschetz principle, with a short proof which applies in any characteristic see [Ek]. See [FR] for a generalization. See [MO90551] for more context.) Raynaud later gave a dramatic algebraic proof of Kodaira vanishing *using* positive characteristic (!), see [Del] or [II].

22.5.K. EXERCISE. Prove the Kodaira Vanishing Theorem in the case $\dim X = 1$, and for \mathbb{P}_k^n . (Neither case requires the characteristic 0 hypothesis.)

22.5.10. ★ A first glimpse of Hodge theory.

The invariant $h^j(X, \Omega_{X/k}^i)$ is called the **Hodge number** $h^{i,j}(X)$. By Exercise 22.5.A, $h^{i,0}$ are birational invariants. We will soon see (in Exercise 22.5.O) that this isn't true for all $h^{i,j}$.

22.5.L. EXERCISE. Suppose X is a smooth projective variety. Assuming Miracle 19.5.2 (that the canonical bundle is Serre-dualizing), show that Hodge numbers satisfy the symmetry $h^{p,q} = h^{n-p, n-q}$. (Exercise 14.2.V will be useful.)

22.5.M. EXERCISE (THE HODGE NUMBERS OF PROJECTIVE SPACE). Show that $h^{p,q}(\mathbb{P}_k^n) = 1$ if $0 \leq p = q \leq n$ and $h^{p,q}(\mathbb{P}_k^n) = 0$ otherwise. Hint: use the Euler exact sequence (Theorem 22.4.6) and apply Exercise 14.2.U.

22.5.11. Remark: The Hodge diamond. Over $k = \mathbb{C}$, further miracles occur. If X is an irreducible smooth projective complex variety, then it turns out that there is a direct sum decomposition

$$(22.5.11.1) \quad H^m(X, \mathbb{C}) = \bigoplus_{i+j=m} H^j(X, \Omega_{X/\mathbb{C}}^i),$$

from which $h^m(X, \mathbb{C}) = \sum_{i+j=m} h^{i,j}$, so the Hodge numbers (purely algebraic objects) yield the Betti numbers (a priori topological information). Moreover, complex conjugation interchanges $H^j(X, \Omega_{X/\mathbb{C}}^i)$ with $H^i(X, \Omega_{X/\mathbb{C}}^j)$, from which

$$(22.5.11.2) \quad h^{i,j} = h^{j,i}.$$

(Aside: This additional symmetry holds in characteristic 0 in general, but can fail in positive characteristic, see for example [Ig, p. 966], [Mu1, §II], [Se4, Prop. 16].) This is the beginning of the vast and rich subject of Hodge theory (see [GH1, §0.6] for more, or [Vo] for much more).

If we write the Hodge numbers in a diamond, with $h^{i,j}$ the i th entry in the $(i+j)$ th row, then the diamond has the two symmetries coming from Serre duality and complex conjugation. For example, the Hodge diamond of an irreducible smooth projective complex surface will be of the following form:

$$\begin{array}{ccccc} & & 1 & & \\ & q & & q & \\ p_g & & h^{1,1} & & p_g \\ & q & & q & \\ & & 1 & & \end{array}$$

where p_g is the geometric genus of the surface, and $q = h^{0,1} = h^{1,0} = h^{2,1} = h^{1,2}$ is called the **irregularity** of the surface. As another example, by Exercise 22.5.M, the Hodge diamond of \mathbb{P}^n is all 0 except for 1's down the vertical axis of symmetry.

You won't need the unproved statements (22.5.11.1) or (22.5.11.2) to solve the following problems.

22.5.N. EXERCISE. Assuming Miracle 19.5.2 (that the canonical bundle is Serre-dualizing), show that the Hodge diamond of a smooth projective geometrically irreducible genus g curve over a field k is the following.

$$\begin{array}{ccccc} & & 1 & & \\ & g & & g & \\ & & 1 & & \end{array}$$

22.5.O. EXERCISE. Show that the Hodge diamond of $\mathbb{P}_k^1 \times \mathbb{P}_k^1$ is the following.

$$\begin{array}{ccccc} & & 1 & & \\ & 0 & & 0 & \\ 0 & & 2 & & 0 \\ & 0 & & 0 & \\ & & 1 & & \end{array}$$

By comparing your answer to the Hodge diamond of \mathbb{P}_k^2 (Exercise 22.5.M), show that $h^{1,1}$ is not a birational invariant.

Notice that in both cases, $h^{1,1}$ is the Picard number ρ (defined in §19.4.9). In general, $\rho \leq h^{1,1}$ (see [GH1, §3.5, p. 456-7] in characteristic 0).

22.5.12. ★ Aside: Infinitesimal deformations and automorphisms.

It is beyond the scope of this book to make this precise, but if X is a variety, $H^0(X, \mathcal{T}_X)$ parametrizes infinitesimal automorphisms of X , and $H^1(X, \mathcal{T}_X)$ parametrizes infinitesimal deformations. As an example if $X = \mathbb{P}^1$ (over a field), $\mathcal{T}_{\mathbb{P}^1} \cong \mathcal{O}(2)$ (§22.4.2), so $h^0(\mathbb{P}^1, \mathcal{T}_{\mathbb{P}^1}) = 3$, which is precisely the dimension of the automorphism group of \mathbb{P}^1 (Exercise 17.4.B).

22.5.P. EXERCISE. Compute $h^0(\mathbb{P}_k^n, \mathcal{T}_{\mathbb{P}_k^n})$ using the Euler exact sequence (Theorem 22.4.6). Compare this to the dimension of the automorphism group of \mathbb{P}_k^n (Exercise 17.4.B).

22.5.Q. EXERCISE. Show that $H^1(\mathbb{P}_A^n, \mathcal{T}_{\mathbb{P}_A^n}) = 0$. (Thus projective space can't deform, and is “rigid”.)

22.5.R. EXERCISE. Assuming Miracle 19.5.2 (that the canonical bundle is Serre-dualizing), compute $h^i(C, \mathcal{T}_C)$ for a genus g smooth projective geometrically irreducible curve over k , for $i = 0$ and 1 . You should notice that $h^1(C, \mathcal{T}_C)$ for genus $0, 1$, and $g > 1$ is $0, 1$, and $3g - 3$ respectively; after doing this, re-read §20.8.3.

22.6 Unramified morphisms

Suppose $\pi: X \rightarrow Y$ is a morphism of schemes. The support of the quasicoherent sheaf $\Omega_\pi = \Omega_{X/Y}$ is called the **ramification locus**, and the image of its support, $\pi(\text{Supp } \Omega_{X/Y})$, is called the **branch locus**. If $\Omega_\pi = 0$, we say that π is **formally unramified**, and if π is also furthermore locally of finite type, we say π is **unramified**. (Caution: there is some lack of consensus in the definition of “unramified”; “locally of finite type” is sometimes replaced by “locally of finite presentation”, which was the definition originally used in [Gr-EGA].)

22.6.A. EASY EXERCISE (EXAMPLES OF UNRAMIFIED MORPHISMS).

- (a) Show that locally closed embeddings are unramified.
- (b) Show that if S is a multiplicative subset of the ring B , then $\text{Spec } S^{-1}B \rightarrow \text{Spec } B$ is formally unramified. (Thus if η is the generic point of an integral scheme Y , then $\text{Spec } \mathcal{O}_{Y,\eta} \rightarrow Y$ is formally unramified.)
- (c) Show that finite separable field extensions (or more correctly, the corresponding maps of schemes) are unramified.

22.6.B. EXERCISE (PRACTICE WITH THE CONCEPT).

- (a) Show that the normalization of the node in Exercise 10.7.E (see Figure 8.4) is unramified.
- (b) Show that the normalization of the cusp in Exercise 10.7.F (see Figure 10.4) is *not* unramified.

22.6.C. EASY EXERCISE (UNRAMIFIED MORPHISMS ARE A “REASONABLE” CLASS OF MORPHISMS). Show that the class of unramified morphisms is “reasonable” in the sense of §8.1.

22.6.D. EXERCISE (CHARACTERIZATIONS OF UNRAMIFIED MORPHISMS BY THEIR FIBERS). Suppose $\pi: X \rightarrow Y$ is locally of finite type.

(a) Show that π is unramified if and only if for each $q \in Y$, $\pi^{-1}(q)$ is the (scheme-theoretic) disjoint union of schemes of the form $\text{Spec } K$, where K is a finite separable extension of $\kappa(q)$.

(b) Show that π is unramified if and only if for each *geometric point* \bar{q} , $\pi^{-1}(\bar{q}) := \bar{q} \times_Y X$ is the (scheme-theoretic) disjoint union of copies of \bar{q} .

22.6.E. UNIMPORTANT EXERCISE (FOR NUMBER THEORISTS). Suppose $\phi: (A, \mathfrak{m}) \rightarrow (B, \mathfrak{n})$ is a local homomorphism of local rings. In algebraic number theory, such a ring morphism is said to be *unramified* if $B/\phi(\mathfrak{m})B$ is a finite separable extension of A/\mathfrak{m} . Show that if ϕ is finite type, this agrees with our definition.

22.6.F. EXERCISE (FOR USE IN EXERCISE 22.6.G(B) AND EXERCISE 26.2.G).

(a) Suppose $\pi: X \rightarrow Y$ is a locally finite type morphism of locally Noetherian schemes. Show that π is unramified if and only if $\delta_\pi: X \rightarrow X \times_Y X$ is an open embedding. Hint: Show the following. If $\phi: X \rightarrow Z$ is a closed embedding of Noetherian schemes, and the ideal sheaf \mathcal{I} of ϕ satisfies $\mathcal{I} = \mathcal{I}^2$, then ϕ is also an open embedding. For that, show that if (A, \mathfrak{m}) is a Noetherian local ring, and I is a proper ideal of A satisfying $I = I^2$, then $I = 0$. For that in turn, use Nakayama version 2 (Lemma 8.2.9). Also use the fact that $\text{Supp } \mathcal{I}$ is closed (using Exercise 14.3.D), so its complement is open.

★ (b) Adapt your proof of (a) to drop the locally Noetherian hypothesis. Hint: show that if π is locally of finite type, then δ_π is locally finitely presented.

22.6.G. EASY EXERCISE. Suppose $\pi: X \rightarrow Y$ and $\rho: Y \rightarrow Z$ are locally of finite type. Let $\tau = \rho \circ \pi$:

$$\begin{array}{ccc} X & \xrightarrow{\pi} & Y \\ & \searrow \tau & \swarrow \rho \\ & Z & \end{array}$$

(a) Show that if τ is unramified, then so is π .

(b) Show that if ρ is unramified, and τ is smooth of relative dimension n (e.g., étale if $n = 0$), then π is smooth of relative dimension n .

(Does this agree with your geometric intuition?)

22.6.H. EXERCISE. Suppose $\pi: X \rightarrow Y$ is locally of finite type. Show that the locus in X where π is unramified is open.

22.6.I. EXERCISE (THE MEANING OF UNRAMIFIEDNESS IN “VERY GEOMETRIC” SETTINGS). Show that a morphism $\pi: X \rightarrow Y$ of finite type schemes over an algebraically closed field is unramified if and only if π is injective on tangent vectors at closed points. (This was mentioned in the remark immediately after Theorem 20.1.1.)

22.6.1. Arithmetic side remark: the different and discriminant. If B is the ring of integers in a number field (§10.7.1), the **different ideal** of B is the annihilator of $\Omega_{B/\mathbb{Z}}$. It measures how “ramified” $\text{Spec } B \rightarrow \text{Spec } \mathbb{Z}$ is, and is a ring-theoretic version of the ramification locus. The **discriminant ideal** can be interpreted as the ideal of \mathbb{Z} corresponding to effective divisor on $\text{Spec } \mathbb{Z}$ that is the “push forward”

(not defined here, but defined as you might expect) of the divisor corresponding to the different. It is a ring-theoretic version of the branch locus. If B/A is an extension of rings of integers of number fields, the **relative different ideal** (of B) and **relative discriminant ideal** (of A) are defined similarly. (We won't use these ideas.)

22.7 The Riemann-Hurwitz Formula

The Riemann-Hurwitz formula generalizes our calculation of the genus g of a double cover of \mathbb{P}^1 branched at $2g + 2$ points, Theorem 20.5.1, to higher degree covers, and to higher genus target curves.

22.7.1. Definition. A finite morphism between integral schemes $X \rightarrow Y$ is said to be **separable** if it is dominant, and the induced extension of function fields $K(X)/K(Y)$ is a separable extension. Similarly, a generically finite morphism is **generically separable** if it is dominant, and the induced extension of function fields is a separable extension. Note that finite morphisms of integral schemes are automatically separable in characteristic 0.

22.7.2. Proposition. — *If $\pi: X \rightarrow Y$ is a generically separable morphism of irreducible smooth varieties of the same dimension n , then the cotangent exact sequence (Theorem 22.2.25) is exact on the left as well:*

$$(22.7.2.1) \quad 0 \longrightarrow \pi^* \Omega_{Y/k} \xrightarrow{\phi} \Omega_{X/k} \longrightarrow \Omega_{X/Y} \longrightarrow 0.$$

This is an example of left-exactness of the cotangent exact sequence in the presence of appropriate “smoothness”, see Remark 22.2.10.

Proof. We must check that ϕ is injective. Now $\Omega_{Y/k}$ is a rank n locally free sheaf on Y , so $\pi^* \Omega_{Y/k}$ is a rank n locally free sheaf on X . A locally free sheaf on an integral scheme (such as $\pi^* \Omega_{Y/k}$) is torsion-free (any section over any open set is nonzero at the generic point, see §6.1.4), so if a subsheaf of it (such as $\ker \phi$) is nonzero, it is nonzero at the generic point. Thus to show the injectivity of ϕ , we need only check that ϕ is an inclusion at the generic point. We thus tensor with \mathcal{O}_η where η is the generic point of X . Tensoring with \mathcal{O}_η is an exact functor (localization is exact, Exercise 1.6.F), and $\mathcal{O}_\eta \otimes \Omega_{X/Y} = 0$ (as $K(X)/K(Y)$ is a separable extension by hypothesis, and Ω for separable field extensions is 0 by Exercise 22.2.F(a)). Also, $\mathcal{O}_\eta \otimes \pi^* \Omega_{Y/k}$ and $\mathcal{O}_\eta \otimes \Omega_{X/k}$ are both n -dimensional \mathcal{O}_η -vector spaces (they are the stalks of rank n locally free sheaves at the generic point). Thus by considering

$$\mathcal{O}_\eta \otimes \pi^* \Omega_{Y/k} \rightarrow \mathcal{O}_\eta \otimes \Omega_{X/k} \rightarrow \mathcal{O}_\eta \otimes \Omega_{X/Y} \rightarrow 0$$

(which is $\mathcal{O}_\eta^{\oplus n} \rightarrow \mathcal{O}_\eta^{\oplus n} \rightarrow 0 \rightarrow 0$) we see that $\mathcal{O}_\eta \otimes \pi^* \Omega_{Y/k} \rightarrow \mathcal{O}_\eta \otimes \Omega_{X/k}$ is injective, and thus that $\pi^* \Omega_{Y/k} \rightarrow \Omega_{X/k}$ is injective. \square

People not confined to characteristic 0 should note what goes wrong for non-separable morphisms. For example, suppose k is a field of characteristic p , and consider the map $\pi: \mathbb{A}_k^1 = \operatorname{Spec} k[t] \rightarrow \mathbb{A}_k^1 = \operatorname{Spec} k[u]$ given by $u = t^p$. Then

Ω_π is the trivial invertible sheaf generated by dt . As another (similar but different) example, if $K = k(x)$ and $K' = k(x^p)$, then the inclusion $K' \hookrightarrow K$ induces $\pi: \operatorname{Spec} K[t] \rightarrow \operatorname{Spec} K'[t]$. Once again, Ω_π is an invertible sheaf, generated by dx (which in this case is pulled back from $\Omega_{K/K'}$ on $\operatorname{Spec} K$). In both of these cases, we have maps from one affine line to another, and there are vertical tangent vectors.

22.7.A. EXERCISE. If X and Y are smooth varieties of dimension n , and $\pi: X \rightarrow Y$ is generically separable, show that the ramification locus is pure codimension 1, and has a natural interpretation as an effective divisor, as follows. Interpret ϕ as an $n \times n$ Jacobian matrix (13.1.7.1) in appropriate local coordinates, and hence interpret the locus where ϕ is not an isomorphism as (locally) the vanishing scheme of the determinant of an $n \times n$ matrix. Hence the branch locus is also pure codimension 1. Hence we use the terms **ramification divisor** and **branch divisor**.

Before getting down to our case of interest, dimension 1, we begin with something (literally) small but fun. Suppose $\pi: X \rightarrow Y$ is a surjective k -morphism from a smooth k -scheme that contracts a subset of codimension greater than 1. More precisely, suppose π is an isomorphism over an open subset of Y , from an open subset U of X whose complement has codimension greater than 1. Then by Exercise 22.7.A, Y cannot be smooth. (*Small resolutions*, to be defined in Exercise 23.4.N, are examples of such π . In particular, you can find an example there.)

Suppose now that X and Y are dimension 1. Then the ramification locus is a finite set (**ramification points**) of X , and the branch locus is a finite set (**branch points**) of Y . (Figure 22.4 shows a morphism with two ramification points and one branch point.) Now assume that $k = \bar{k}$. We examine $\Omega_{X/Y}$ near a point $p \in X$.




FIGURE 22.4. An example where the branch divisor appears with multiplicity 2 (see Exercise 22.7.C)

As motivation for what we will see, we note that in complex geometry, nonconstant maps from (complex) curves to curves may be written in appropriate local coordinates as $x \mapsto x^m = y$, from which we see that dy pulls back to $mx^{m-1}dx$, so $\Omega_{X/Y}$ locally looks like functions times dx modulo multiples of $mx^{m-1}dx$.

Consider now our map $\pi: X \rightarrow Y$, and fix $p \in X$, and $q = \pi(p)$. Because the construction of Ω behaves well under base change (Theorem 22.2.27(b)), we may replace Y with Spec of the local ring $\mathcal{O}_{Y,q}$ at q , i.e., we may assume $Y = \operatorname{Spec} B$,

where B is a discrete valuation ring (as Y is a regular curve), with residue field $k = \bar{k}$ corresponding to q . Then as π is finite, X is affine too. Similarly, as the construction of Ω behaves well with respect to localization (Exercise 22.2.L), we may replace X by $\text{Spec } \mathcal{O}_{X,p}$, and thus assume $X = \text{Spec } A$, where A is a discrete valuation ring, and π corresponds to $B \rightarrow A$, inducing an isomorphism of residue fields (with k).

Suppose their uniformizers are s and t respectively, with $t \mapsto us^n$ where u is an invertible element of A .

$$\begin{array}{ccc} X & A & us^n \\ \downarrow & \uparrow & \uparrow \\ Y & B & t \end{array}$$

Recall that the differentials of a discrete valuation ring over k are generated by d of the uniformizer (Exercise 22.4.D). Then

$$dt = d(us^n) = uns^{n-1} ds + s^n du.$$

This differential on $\text{Spec } A$ vanishes to order at least $n-1$, and precisely $n-1$ if n is not divisible by the characteristic. If the order is $n-1$, we call it **tame** ramification, while if the order is greater, it is called **wild** ramification. We call this order the **ramification order** at this point of X .

22.7.B. EXERCISE. Show that the degree of $\Omega_{X/Y}$ at $p \in X$ is precisely the ramification order of π at p . (The degree of a coherent sheaf on a curve was defined in §19.4.5. To do this exercise, you will have to explain why a coherent sheaf \mathcal{F} supported on a finite set of points has a “degree” at each of these points, which sum to the total degree of \mathcal{F} .)

22.7.C. EXERCISE: INTERPRETING THE RAMIFICATION DIVISOR IN TERMS OF NUMBER OF PREIMAGES. For this exercise, we take the base field k to be algebraically closed. Suppose all the ramification above $q \in Y$ is tame (which is always true in characteristic 0). Show that the degree of the branch divisor at q is $\deg \pi - |\pi^{-1}(q)|$. Thus the multiplicity of the branch divisor counts the extent to which the number of preimages is less than the degree (see Figure 22.4).

22.7.3. Theorem (the Riemann-Hurwitz formula). — Suppose $\pi: X \rightarrow Y$ is a finite separable morphism of projective regular curves, of pure degree n . Let R be the ramification divisor. Then

$$(22.7.3.1) \quad 2g(X) - 2 = n(2g(Y) - 2) + \deg R.$$

22.7.D. EXERCISE. Prove the Riemann-Hurwitz formula. Hint: Apply the fact that degree is additive in exact sequences (Exercise 19.4.J) to (22.7.2.1). Recall that degrees of line bundles pull back well under finite morphisms of integral projective curves, Exercise 19.4.G. A torsion coherent sheaf on a reduced curve (such as Ω_π) is supported in dimension 0 (Exercise 14.3.G(b)), so $\chi(\Omega_\pi) = h^0(\Omega_\pi)$. Show that the degree of R as a divisor is the same as its degree in the sense of h^0 .

Here are some applications of the Riemann-Hurwitz formula.

22.7.4. Example. The degree of R is always even: any cover of a curve must be branched over an even number of points (counted with appropriate multiplicity).

22.7.E. EASY EXERCISE. Show that there is no nonconstant map from a smooth projective irreducible genus 2 curve to a smooth projective irreducible genus 3 curve. (Hint: $\deg R \geq 0$.)

22.7.5. Definition. If the only finite étale cover of a connected scheme X is an isomorphism, we say that X is **simply connected in the étale topology**. Do you see why “simply connected” is a reasonable name for the concept? We won’t use this terminology beyond the following brief discussion.

22.7.6. Example (“ \mathbb{P}_k^1 is simply connected”). If $k = \bar{k}$, the only connected unbranched finite separable cover of \mathbb{P}_k^1 is the isomorphism, for the following reason. Suppose X is connected and $X \rightarrow \mathbb{P}_k^1$ is unramified. Then X is a curve, and regular by Exercise 26.2.E(a). Applying the Riemann-Hurwitz theorem, using that the ramification divisor is 0, we have $2 - 2g_X = 2d$ with $d \geq 1$ and $g_X \geq 0$, from which $d = 1$ and $g_X = 0$.

22.7.F. EXERCISE (“IN CHARACTERISTIC 0, \mathbb{A}_k^1 IS SIMPLY CONNECTED”). Show that if $k = \bar{k}$ has characteristic 0, the only connected unbranched cover of \mathbb{A}_k^1 is itself. (Aside: in characteristic p , this needn’t hold; $\text{Spec } k[x, y]/(y^p - y - x) \rightarrow \text{Spec } k[x]$ is such a map, as you can show yourself. Once the theory of the algebraic fundamental group is developed, this translates to: “ \mathbb{A}^1 is not simply connected in characteristic p .” This cover is an example of an *Artin-Schreier cover*. Fun fact: the group \mathbb{Z}/p acts on this cover via the map $y \mapsto y + 1$. This is an example of a *Galois cover*; you can check that the extension of function fields is Galois.)

22.7.G. UNIMPORTANT EXERCISE. Extend Example 22.7.6 and Exercise 22.7.F, by removing the $k = \bar{k}$ hypothesis, and changing “connected” to “geometrically connected”.

22.7.7. Example: Lüroth’s Theorem. Continuing the notation of Theorem 22.7.3, suppose $g(X) = 0$. Then from the Riemann-Hurwitz formula (22.7.3.1), $g(Y) = 0$. (Otherwise, if $g(Y)$ were at least 1, then the right side of the Riemann-Hurwitz formula would be non-negative, and thus couldn’t be -2 , which is the left side.) Informally: the only maps from a genus 0 curve to a curve of positive genus are the constant maps. This has a nonobvious algebraic consequence, by our identification of covers of curves with field extensions (Theorem 18.4.3): all subfields of $k(x)$ containing k are of the form $k(y)$ where $y = f(x)$ for some $f \in k(x)$.

$$\begin{array}{ccc} k(x) & & \mathbb{P}^1 \\ \uparrow & & \downarrow \\ K(C) & & C \end{array} = \mathbb{P}^1$$

(It turns out that the hypotheses $\text{char } k = 0$ and $k = \bar{k}$ are not necessary.) This is Lüroth’s Theorem.

22.7.H. EXERCISE. Use Lüroth’s Theorem to give new geometric solutions to Exercises 7.5.J and 7.5.L. (These arguments will be less ad hoc, and more suitable for generalization, than the algebraic solutions suggested in the hints to those exercises.)

22.7.I. ★ EXERCISE (GEOMETRIC GENUS EQUALS TOPOLOGICAL GENUS). This exercise is intended for those with some complex background, who know that the Riemann-Hurwitz formula holds in the complex analytic category. Suppose C is an irreducible regular projective complex curve. Show that there is an algebraic nonconstant map $\pi: C \rightarrow \mathbb{P}_{\mathbb{C}}^1$. Describe the corresponding map of Riemann surfaces. Use the Riemann-Hurwitz formula to show that the algebraic notion of genus (as computed using the branched cover π) agrees with the topological notion of genus (using the same branched cover). (Recall that assuming Miracle 19.5.2 — that the canonical bundle is Serre-dualizing — we know that the geometric genus equals the arithmetic genus, Exercise 22.5.C.)

22.7.J. UNIMPORTANT EXERCISE. Suppose $\pi: X \rightarrow Y$ is a dominant morphism of regular curves, and R is the ramification divisor of π . Show that $\Omega_X(-R) \cong \pi^*\Omega_Y$. (This exercise is geometrically pleasant, but we won’t use it.) Hint: This says that we can interpret the invertible sheaf $\pi^*\Omega_Y$ over an open set U of X as precisely those differentials on U vanishing along the ramification divisor.

22.7.8. Informal example: The degree of the discriminant of degree d polynomials in one variable. You may be aware that there is a degree $2d - 2$ polynomial in the coefficients a_d, \dots, a_0 of the degree d polynomial

$$f(x) = a_dx^d + a_{d-1}x^{d-1} + \cdots + a_0 = 0$$

that vanishes precisely when $f(x)$ has a multiple root. For example, when $d = 2$, the discriminant is $a_1^2 - 4a_0a_2$. We can “compute” this degree $2d - 2$ using the Riemann-Hurwitz formula as follows. (You should try to make sense of the following informal and imprecise discussion.) We work over an algebraically closed field k of characteristic 0 for the sake of simplicity. If we take two general degree d polynomials, $g(x)$ and $h(x)$, the degree of the discriminant “should be” the number of $\lambda \in k$ for which $g(x) - \lambda h(x)$ has a double root. Consider the morphism $\mathbb{P}^1 \rightarrow \mathbb{P}^1$ given by $x \mapsto [g(x), h(x)]$. (Here we use “affine coordinates” on the source \mathbb{P}^1 : by x we mean $[x, 1]$.) Then this morphism has degree d . A ramification point a mapping to the branch point $[\lambda, 1]$ in the target \mathbb{P}^1 corresponds to a being a double root of $g(x) - \lambda h(x)$. Thus the number of branch points should be the desired degree of the discriminant. By the Riemann-Hurwitz formula there are $2d - 2$ branch points (admittedly, with multiplicity). It is possible to turn this into a proof, and it is interesting to do so.

22.7.9. Bounds on automorphism groups of curves.

It is a nontrivial fact that irreducible smooth projective curves of genus $g \geq 2$ have finite automorphism groups. (See [FK, §5.1] for a proof over \mathbb{C} . See [Schm] for the first proof in arbitrary characteristic, although better approaches are now available.) Granting this fact, we can show that in characteristic 0, the automorphism group has order at most $84(g - 1)$ (Hurwitz’s Automorphisms Theorem), as follows.

Suppose C is an irreducible smooth projective curve over an algebraically closed field $k = \bar{k}$ of characteristic 0, of genus $g \geq 2$. Suppose that G is a *finite* group of automorphisms of C . We now show that $|G| \leq 84(g-1)$. (The case where k is not algebraically closed is quickly dispatched by base-changing to \bar{k} .)

22.7.K. EXERCISE.

(a) Let C' be the smooth projective curve corresponding to the field extension $K(C)^G$ of k (via Theorem 18.4.3). ($K(C)^G$ means the G -invariants of $K(C)$.) Describe a morphism $\pi: C \rightarrow C'$ of degree $|G|$, as well as a faithful G -action on C that commutes with π .

(b) Show that above each branch point of π , the preimages are all ramified to the same order (as G acts transitively on them). Suppose there are n branch points and the i th one has ramification r_i (each $|G|/r_i$ times).

(c) Use the Riemann-Hurwitz formula to show that

$$(2g-2) = |G| \left(2g(C') - 2 + \sum_{i=1}^n \frac{r_i-1}{r_i} \right)$$

To maximize $|G|$, we wish to minimize

$$(22.7.9.1) \quad 2g(C') - 2 + \sum_{i=1}^n \frac{r_i-1}{r_i}$$

subject to (22.7.9.1) being positive. Note that $1/42$ is possible: take $g(C') = 0$, $n = 3$, and $(r_1, r_2, r_3) = (2, 3, 7)$.

22.7.L. EXERCISE. Show that you can't do better than $1/42$ by considering the following cases separately:

- (a) $g(C') > 1$,
- (b) $g(C') = 1$,
- (c) $g(C') = 0$ and $n \geq 5$,
- (d) $g(C') = 0$ and $n = 4$, and
- (e) $g(C') = 0$ and $n = 3$.

22.7.M. EXERCISE. Use the fact that (22.7.9.1) is at least $1/42$ to prove the result.

22.7.10. Remark. In positive characteristic, there can be many more automorphisms. For example, [Stic], in characteristic p , if p^n is a prime power that is not 2, then the completion of the affine curve $y^{p^n} = x + x^{p^n+1}$ has genus $g = p^n(p^n-1)/2$ and automorphism group of order $p^{3n}(p^{3n}+1)(p^{2n}-1)$.

CHAPTER 23

★ Blowing up

We next discuss an important construction in algebraic geometry, the blow-up of a scheme along a closed subscheme (cut out by a finite type ideal sheaf). The theory could mostly be developed immediately after Chapter 18, but the interpretation in terms of the conormal cone/bundle/sheaf of many classical examples makes it natural to discuss blowing up after differentials.

We won't use blowing up much in later chapters, so feel free to skip this topic for now. But it is an important tool. For example, one can use it to resolve singularities, and more generally, indeterminacy of rational maps. In particular, blow-ups can be used to relate birational varieties to each other.

We will start with a motivational example that will give you a picture of the construction in a particularly important (and the historically earliest) case, in §23.1. We will then see a formal definition, in terms of a universal property, §23.2. The definition won't immediately have a clear connection to the motivational example. We will deduce some consequences of the definition (assuming that the blow-up actually exists). We then prove that the blow-up exists, by describing it quite explicitly, in §23.3. As a consequence, we will find that the blow-up morphism is projective, and we will deduce more consequences from this. In §23.4, we will do a number of explicit computations, to see various sorts of applications, and to see that many things can be computed by hand.

23.1 Motivating example: Blowing up the origin in the plane

We will generalize the following notion, which will correspond to “blowing up” the origin of \mathbb{A}_k^2 (Exercise 10.3.F). Our discussion will be informal. Consider the subset of $\mathbb{A}^2 \times \mathbb{P}^1$ corresponding to the following. We interpret \mathbb{P}^1 as parametrizing the lines through the origin. Consider the subvariety

$$\mathrm{Bl}_{(0,0)} \mathbb{A}^2 := \{(p \in \mathbb{A}^2, [\ell] \in \mathbb{P}^1) : p \in \ell\},$$

which is the data of a point p in the plane, and a line ℓ containing both p and the origin. Algebraically: let x and y be coordinates on \mathbb{A}^2 , and X and Y be projective coordinates on \mathbb{P}^1 (“corresponding” to x and y); we will consider the subset $\mathrm{Bl}_{(0,0)} \mathbb{A}^2$ of $\mathbb{A}^2 \times \mathbb{P}^1$ corresponding to $xY - yX = 0$. We have the useful diagram

$$\begin{array}{ccccc} \mathrm{Bl}_{(0,0)} \mathbb{A}^2 & \hookrightarrow & \mathbb{A}^2 \times \mathbb{P}^1 & \twoheadrightarrow & \mathbb{P}^1 \\ & \searrow \beta & \downarrow & & \\ & & \mathbb{A}^2 & & \end{array}$$

You can verify that it is smooth over k (Definition 13.2.6 or 22.3.1) directly (you can now make the paragraph after Exercise 10.3.F precise), but here is an informal argument, using the projection $\mathrm{Bl}_{(0,0)} \mathbb{A}^2 \rightarrow \mathbb{P}^1$. The projective line \mathbb{P}^1 is smooth, and for each point $[\ell]$ in \mathbb{P}^1 , we have a smooth choice of points on the line ℓ . Thus we are verifying smoothness by way of a fibration over \mathbb{P}^1 .

We next consider the projection to \mathbb{A}^2 , $\beta: \mathrm{Bl}_{(0,0)} \mathbb{A}^2 \rightarrow \mathbb{A}^2$. This is an isomorphism away from the origin. Loosely speaking, if p is not the origin, there is precisely one line containing p and the origin. On the other hand, if p is the origin, then there is a full \mathbb{P}^1 of lines containing p and the origin. Thus the preimage of $(0,0)$ is a curve, and hence a divisor (an effective Cartier divisor, as the blown-up surface is regular). This is called the *exceptional divisor* of the blow-up.

If we have some curve $C \subset \mathbb{A}^2$ singular at the origin, it can be potentially partially desingularized, using the blow-up, by taking the closure of $C \setminus \{(0,0)\}$ in $\mathrm{Bl}_{(0,0)} \mathbb{A}^2$. (A **desingularization** or a **resolution of singularities** of a variety X is a proper birational morphism $\tilde{X} \rightarrow X$ from a regular scheme.) For example, consider the curve $y^2 = x^3 + x^2$, which is regular except for a node at the origin. We can take the preimage of the curve minus the origin, and take the closure of this locus in the blow-up, and we will obtain a regular curve; the two branches of the node downstairs are separated upstairs. (You can check this in Exercise 23.4.B once we have defined things properly. The result will be called the *proper transform* (or *strict transform*) of the curve.) We are interested in desingularizations for many reasons. Because we understand regular curves quite well, we could hope to understand other curves through their desingularizations. This philosophy holds true in higher dimension as well.

More generally, we can blow up \mathbb{A}^n at the origin (or more informally, “blow up the origin”), getting a subvariety of $\mathbb{A}^n \times \mathbb{P}^{n-1}$. Algebraically, If x_1, \dots, x_n are coordinates on \mathbb{A}^n , and X_1, \dots, X_n are projective coordinates on \mathbb{P}^{n-1} , then the blow-up $\mathrm{Bl}_{\bar{0}} \mathbb{A}^n$ is given by the equations $x_i X_j - x_j X_i = 0$. Once again, this is smooth: \mathbb{P}^{n-1} is smooth, and for each point $[\ell] \in \mathbb{P}^{n-1}$, we have a smooth choice of $p \in \ell$.

We can extend this further, by blowing up \mathbb{A}^{n+m} along a coordinate m -plane \mathbb{A}^m by adding m more variables x_{n+1}, \dots, x_{n+m} to the previous example; we get a subset of $\mathbb{A}^{n+m} \times \mathbb{P}^{m-1}$.

Because in complex geometry, submanifolds of manifolds locally “look like” coordinate m -planes in n -space, you might imagine that we could extend this to blowing up a regular subvariety of a regular variety. In the course of making this precise, we will accidentally generalize this notion greatly, defining the blow-up of any finite type quasicoherent sheaf of ideals in a scheme. In general, blowing up may not have such an intuitive description as in the case of blowing up something regular inside something regular — it can do great violence to the scheme — but even then, it is very useful.

Our description will depend only on the closed subscheme being blown up, and not on coordinates. That remedies a defect that was already present in the first example, of blowing up the plane at the origin. It is not obvious that if we picked different coordinates for the plane (preserving the origin as a closed subscheme) that we wouldn’t have two different resulting blow-ups.

As is often the case, there are two ways of understanding the notion of blowing up, and each is useful in different circumstances. The first is by universal property,

which lets you show some things without any work. The second is an explicit construction, which lets you get your hands dirty and compute things (and implies for example that the blow-up morphism is projective).

The motivating example here may seem like a very special case, but if you understand the blow-up of the origin in n -space well enough, you will understand blowing up in general.

23.2 Blowing up, by universal property

We now define the blow-up by a universal property. The disadvantage of starting here is that this definition won't obviously be the same as (or even related to) the examples of §23.1.

Suppose $X \hookrightarrow Y$ is a closed subscheme corresponding to a finite type quasicoherent sheaf of ideals. (If Y is locally Noetherian, the “finite type” hypothesis is automatic, so Noetherian readers can ignore it.)

The blow-up of $X \hookrightarrow Y$ is a Cartesian diagram

$$(23.2.0.1) \quad \begin{array}{ccc} E_X Y & \hookrightarrow & \mathrm{Bl}_X Y \\ \downarrow & & \downarrow \beta \\ X & \hookrightarrow & Y \end{array}$$

such that $E_X Y$ (the scheme-theoretic preimage of X by β) is an effective Cartier divisor (defined in §9.5.1) on $\mathrm{Bl}_X Y$, such that any other such Cartesian diagram

$$(23.2.0.2) \quad \begin{array}{ccc} D & \hookrightarrow & W \\ \downarrow & & \downarrow \\ X & \hookrightarrow & Y, \end{array}$$


where D is an effective Cartier divisor on W , factors uniquely through it:

$$\begin{array}{ccc} D & \hookrightarrow & W \\ \downarrow & & \downarrow \\ E_X Y & \hookrightarrow & \mathrm{Bl}_X Y \\ \downarrow & & \downarrow \\ X & \hookrightarrow & Y. \end{array}$$

We call $\mathrm{Bl}_X Y$ the **blow-up** (of Y along X , or of Y with **center** X). (Other somewhat archaic terms for this are *monoidal transformation*, *σ -process*, and *quadratic transformation*, and *dilation*.) We call $E_X Y$ the **exceptional divisor** of the blow-up. (Bl and β stand for “blow-up”, and E stands for “exceptional”.)

By a typical universal property argument, if the blow-up exists, it is unique up to unique isomorphism. (We can even recast this more explicitly in the language of Yoneda’s Lemma: consider the category of diagrams of the form (23.2.0.2), where

morphisms are diagrams of the form



Then the blow-up is a final object in this category, if one exists.)

If $Z \hookrightarrow Y$ is any closed subscheme of Y , then the scheme-theoretic preimage $\beta^{-1}Z$ is called the **total transform** of Z . We will soon see that β is an isomorphism away from X (Observation 23.2.2). $\overline{\beta^{-1}(Z - X)}$ is called the **proper transform** or **strict transform** of Z . (We will use the first terminology. We will also define it in a more general situation.) We will soon see (in the Blow-up Closure Lemma 23.2.6) that the proper transform is naturally isomorphic to $\text{Bl}_{Z \cap X} Z$, where $Z \cap X$ is the scheme-theoretic intersection.

We will soon show that the blow-up always exists, and describe it explicitly. We first make a series of observations, *assuming that the blow up exists*.

23.2.1. Observation. If X is the empty set, then $\text{Bl}_X Y = Y$. More generally, if X is an effective Cartier divisor, then the blow-up is an isomorphism. (Reason: $\text{id}_Y: Y \rightarrow Y$ satisfies the universal property.)

23.2.A. EXERCISE. If U is an open subset of Y , then $\text{Bl}_{U \cap X} U \cong \beta^{-1}(U)$, where $\beta: \text{Bl}_X Y \rightarrow Y$ is the blow-up.

Thus “we can compute the blow-up locally.”

23.2.B. EXERCISE. Show that if Y_α is an open cover of Y (as α runs over some index set), and the blow-up of Y_α along $X \cap Y_\alpha$ exists, then the blow-up of Y along X exists.

23.2.2. Observation. Combining Observation 23.2.1 and Exercise 23.2.A, we see that the blow-up is an isomorphism away from the locus you are blowing up:

$$\beta|_{\text{Bl}_X Y - E_X Y}: \text{Bl}_X Y - E_X Y \rightarrow Y - X$$

is an isomorphism.

23.2.3. Observation. If $X = Y$, then the blow-up is the empty set: the only map $W \rightarrow Y$ such that the pullback of X is an effective Cartier divisor is $\emptyset \hookrightarrow Y$. In this case we have “blown Y out of existence”!

23.2.C. EXERCISE (BLOW-UP PRESERVES IRREDUCIBILITY AND REDUCEDNESS). Show that if Y is irreducible, and X doesn’t contain the generic point of Y , then $\text{Bl}_X Y$ is irreducible. Show that if Y is reduced, then $\text{Bl}_X Y$ is reduced.

23.2.4. Existence in a first nontrivial case: blowing up a locally principal closed subscheme.

We next see why $\text{Bl}_X Y$ exists if $X \hookrightarrow Y$ is locally cut out by one equation. As the question is local on Y (Exercise 23.2.B), we reduce to the affine case $\text{Spec } A/(t) \hookrightarrow$

$\text{Spec } A$. (A good example to think through is $A = k[x, y]/(xy)$ and $t = x$.) Let

$$I = \ker(A \rightarrow A_t) = \{a \in A : t^n a = 0 \text{ for some } n > 0\},$$

and let $\phi: A \rightarrow A/I$ be the projection.

23.2.D. EXERCISE. Show that $\phi(t)$ is not a zerodivisor in A/I .

23.2.E. EXERCISE. Show that $\beta: \text{Spec } A/I \rightarrow \text{Spec } A$ is the blow up of $\text{Spec } A$ along $\text{Spec } A/t$. In other words, show that

$$\begin{array}{ccc} \text{Spec } A/(t, I) & \longrightarrow & \text{Spec } A/I \\ \downarrow & & \downarrow \beta \\ \text{Spec } A/t & \longrightarrow & \text{Spec } A \end{array}$$

is a “blow up diagram” (23.2.0.1). Hint: In checking the universal property reduce to the case where W (in (23.2.0.2)) is affine. Then solve the resulting problem about rings. Depending on how you proceed, you might find Exercise 11.4.G, about the uniqueness of extension of maps over effective Cartier divisors, helpful.

23.2.F. EXERCISE. Show that $\text{Spec } A/I$ is the scheme-theoretic closure of $D(t)$ in $\text{Spec } A$.

Thus you might geometrically interpret $\text{Spec } A/I \rightarrow \text{Spec } A$ as “shaving off any fuzz supported in $V(t)$ ”. In the Noetherian case, this can be interpreted as removing those associated points lying in $V(t)$. This is intended to be vague, and you should think about how to make it precise only if you want to.

23.2.5. The Blow-up Closure Lemma.

Suppose we have a fibered diagram

$$\begin{array}{ccc} W & \xrightarrow{\text{cl. emb.}} & Z \\ \downarrow & & \downarrow \\ X & \xrightarrow{\text{cl. emb.}} & Y \end{array}$$

where the bottom closed embedding corresponds to a finite type ideal sheaf (and hence the upper closed embedding does too). The first time you read this, it may be helpful to consider only the special case where $Z \rightarrow Y$ is a closed embedding.

Then take the fibered product of this square by the blow-up $\beta: \text{Bl}_X Y \rightarrow Y$, to obtain the Cartesian diagram

$$\begin{array}{ccc} W \times_Y E_X Y & \longrightarrow & Z \times_Y \text{Bl}_X Y \\ \downarrow & & \downarrow \\ E_X Y & \xrightarrow{\text{Cartier}} & \text{Bl}_X Y. \end{array}$$

(Why is the top row the base change of the bottom as claimed?) The bottom closed embedding is locally cut out by one equation, and thus the same is true of the top closed embedding as well. However, the local equation on $Z \times_Y \text{Bl}_X Y$ need not be a non-zerodivisor, and thus the top closed embedding is not necessarily an effective Cartier divisor.

Let \bar{Z} be the scheme-theoretic closure of

$$(Z \times_Y \text{Bl}_X Y) \setminus (W \times_Y \text{Bl}_X Y)$$

in $Z \times_Y \text{Bl}_X Y$. (As $W \times_Y \text{Bl}_X Y$ is locally principal, we are in precisely the situation of §23.2.4, so the scheme-theoretic closure is not mysterious.) Note that in the special case where $Z \rightarrow Y$ is a closed embedding, \bar{Z} is the proper transform, as defined in §23.2. For this reason, it is reasonable to call \bar{Z} the *proper transform* of Z even if Z isn't a closed embedding. Similarly, it is reasonable to call $Z \times_Y \text{Bl}_X Y$ the *total transform* of Z even if Z isn't a closed embedding.

Define $E_{\bar{Z}} \hookrightarrow \bar{Z}$ as the pullback of $E_X Y$ to \bar{Z} , i.e., by the fibered diagram

$$\begin{array}{ccc} E_{\bar{Z}} & \xrightarrow{\quad} & \bar{Z} \\ \downarrow \text{cl. emb.} & & \downarrow \text{cl. emb.} \\ W \times_Y E_X Y & \xrightarrow{\text{loc. prin.}} & Z \times_Y \text{Bl}_X Y \\ \downarrow & & \downarrow \\ E_X Y & \xrightarrow{\text{Cartier}} & \text{Bl}_X Y \end{array} \quad \begin{array}{l} \text{proper transform} \\ \\ \text{total transform} \end{array}$$

Note that $E_{\bar{Z}}$ is an effective Cartier divisor on \bar{Z} . (It is locally cut out by one equation, pulled back from a local equation of $E_X Y$ on $\text{Bl}_X Y$. Can you see why this is locally not a zerodivisor?) It can be helpful to note that the top square of the diagram above is a blow-up square, by Exercises 23.2.E and 23.2.F (and the fact that blow-ups can be computed affine-locally).

23.2.6. Blow-up Closure Lemma. — $(\text{Bl}_W Z, E_W Z)$ is canonically isomorphic to $(\bar{Z}, E_{\bar{Z}})$. More precisely: if the blow-up $\text{Bl}_X Y$ exists, then $(\bar{Z}, E_{\bar{Z}})$ is the blow-up of Z along W .

This will be very useful. We make a few initial comments. The first three apply to the special case where $Z \rightarrow Y$ is a closed embedding, and the fourth comment basically tells us we shouldn't have concentrated on this special case.

(1) First, note that if $Z \rightarrow Y$ is a closed embedding, then the Blow-Up Closure Lemma states that the proper transform (as defined in §23.2) is the blow-up of Z along the scheme-theoretic intersection $W = X \cap Z$.


(2) In particular, the Blow-Up Closure Lemma lets you actually compute blow-ups, and we will do lots of examples soon. For example, suppose C is a plane curve, singular at a point p , and we want to blow up C at p . Then we could instead blow up the plane at p (which we have already described how to do, even if we haven't yet proved that it satisfies the universal property of blowing up), and then take the scheme-theoretic closure of $C \setminus \{p\}$ in the blow-up.

(3) More generally, if W is some nasty subscheme of Z that we wanted to blow-up, and Z were a finite type k -scheme, then the same trick would work. We could work locally (Exercise 23.2.A), so we may assume that Z is affine. If W is cut out by r equations $f_1, \dots, f_r \in \Gamma(Z, \mathcal{O}_Z)$, then complete the f 's to a generating set f_1, \dots, f_n of $\Gamma(Z, \mathcal{O}_Z)$. This gives a closed embedding of Z into \mathbb{A}^n such that W is the scheme-theoretic intersection of Z with a coordinate linear space \mathbb{A}^{n-r} .

23.2.7. (4) Most generally still, this reduces the existence of the blow-up to a specific special case. (If you prefer to work over a fixed field k , feel free to replace \mathbb{Z} by

k in this discussion.) Suppose that for each n , $\text{Bl}_{(x_1, \dots, x_n)} \text{Spec } \mathbb{Z}[x_1, \dots, x_n]$ exists. Then I claim that the blow-up always exists. Here's why. We may assume that Y is affine, say $\text{Spec } B$, and $X = \text{Spec } B/(f_1, \dots, f_n)$. Then we have a morphism $Y \rightarrow \mathbb{A}_{\mathbb{Z}}^n$ given by $x_i \mapsto f_i$, such that X is the scheme-theoretic preimage of the origin. Hence by the blow-up closure lemma, $\text{Bl}_X Y$ exists.

23.2.G. ★ TRICKY EXERCISE. Prove the Blow-up Closure Lemma 23.2.6. Hint: obviously, construct maps in both directions, using the universal property. Constructing the following diagram may or may not help.



Hooked arrows indicate closed embeddings; and when morphisms are furthermore locally principal or even effective Cartier, they are so indicated. Exercise 11.4.G, on the uniqueness of extension of maps over effective Cartier divisors, may or may not help as well. Note that if $Z \rightarrow Y$ is actually a closed embedding, then so is $Z \times_Y \text{Bl}_X Y \rightarrow \text{Bl}_X Y$ and hence also $\bar{Z} \rightarrow \text{Bl}_X Y$.

23.3 The blow-up exists, and is projective

23.3.1. It is now time to show that the blow-up always exists. We will see two arguments, which are enlightening in different ways. Both will imply that the blow-up morphism is projective, and hence quasicompact, proper, finite type, and separated. In particular, if $Y \rightarrow Z$ is quasicompact (resp. proper, finite type, separated), so is $\text{Bl}_X Y \rightarrow Z$. (And if $Y \rightarrow Z$ is projective, and Z is quasicompact, then $\text{Bl}_X Y \rightarrow Z$ is projective. See the solution to Exercise 18.3.B for the reason for this annoying extra hypothesis.) The blow-up of a k -variety is a k -variety (using the fact that reducedness is preserved, Exercise 23.2.C), and the blow-up of an irreducible k -variety is an irreducible k -variety (using the fact that irreducibility is preserved, also Exercise 23.2.C).

Approach 1. As explained in §23.2.7, it suffices to show that $\mathrm{Bl}_{V(x_1, \dots, x_n)} \mathrm{Spec} \mathbb{Z}[x_1, \dots, x_n]$ exists. But we know what it is supposed to be: the locus in $\mathrm{Spec} \mathbb{Z}[x_1, \dots, x_n] \times \mathrm{Proj} \mathbb{Z}[X_1, \dots, X_n]$ cut out by the equations $x_i X_j - x_j X_i = 0$. We will show this by the end of the section.

Approach 2. We can describe the blow-up all at once as a *Proj*.

23.3.2. Theorem (Proj description of the blow-up). — Suppose $X \hookrightarrow Y$ is a closed subscheme cut out by a finite type quasicoherent sheaf of ideals $\mathcal{I} \hookrightarrow \mathcal{O}_Y$. Then

$$\mathrm{Proj}_Y (\mathcal{O}_Y \oplus \mathcal{I} \oplus \mathcal{I}^2 \oplus \mathcal{I}^3 \oplus \cdots) \rightarrow Y$$

satisfies the universal property of blowing up.

(We made sense of products of ideal sheaves, and hence \mathcal{I}^n , in Exercise 15.3.E.)

We will prove Theorem 23.3.2 soon (§23.3.3), after seeing what it tells us. Because \mathcal{I} is finite type, the graded sheaf of algebras has degree 1 piece that is finite type. The graded sheaf of algebras is also clearly generated in degree 1. Thus the sheaf of algebras satisfy Hypotheses 18.2.1 (“finite generation in degree 1”).

But first, we should make sure that the preimage of X is indeed an effective Cartier divisor. We can work affine-locally (Exercise 23.2.A), so we may assume that $Y = \mathrm{Spec} B$, and X is cut out by the finitely generated ideal I . Then

$$\mathrm{Bl}_X Y = \mathrm{Proj}_B (B \oplus I \oplus I^2 \oplus \cdots).$$

(You may recall that the ring $B \oplus I \oplus \cdots$ is called the *Rees algebra* of the ideal I in B , §13.9.1.) We are slightly abusing notation by using the notation $\mathrm{Bl}_X Y$, as we haven’t yet shown that this satisfies the universal property.

The preimage of X isn’t just any effective Cartier divisor; it corresponds to the invertible sheaf $\mathcal{O}(1)$ on this *Proj*. Indeed, $\mathcal{O}(1)$ corresponds to taking our graded ring, chopping off the bottom piece, and sliding all the graded pieces to the left by 1 (§16.2); it is the invertible sheaf corresponding to the graded module

$$I \oplus I^2 \oplus I^3 \oplus \cdots$$

(where that first summand I has grading 0). But this can be interpreted as the scheme-theoretic preimage of X , which corresponds to the ideal I of B :

$$I (B \oplus I \oplus I^2 \oplus \cdots) \hookrightarrow B \oplus I \oplus I^2 \oplus \cdots.$$

Thus the scheme-theoretic preimage of $X \hookrightarrow Y$ to $\mathrm{Proj}(\mathcal{O}_Y \oplus \mathcal{I} \oplus \mathcal{I}^2 \oplus \cdots)$, the invertible sheaf corresponding to $\mathcal{I} \oplus \mathcal{I}^2 \oplus \mathcal{I}^3 \oplus \cdots$, is an effective Cartier divisor in class $\mathcal{O}(1)$. Once we have verified that this construction is indeed the blow-up, this divisor will be our exceptional divisor $E_X Y$.

Moreover, we see that the exceptional divisor can be described beautifully as a *Proj* over X :

$$(23.3.2.1) \quad E_X Y = \mathrm{Proj}_X (\mathcal{O}_Y / \mathcal{I} \oplus \mathcal{I} / \mathcal{I}^2 \oplus \mathcal{I}^2 / \mathcal{I}^3 \oplus \cdots).$$

We will later see that in good circumstances (if X is a regular embedding in Y), this is a projectivization of a vector bundle (the “projectivized normal bundle”), see Exercise 23.3.D(a).

23.3.3. Proof of the universal property, Theorem 23.3.2. Let’s prove that this *Proj* construction satisfies the universal property. Then Approach 1 will also follow, as a special case of Approach 2.

23.3.4. Aside: why Approach 1? Before we begin, you may be wondering why we bothered with Approach 1. One reason is that you may find it more comfortable to work with this one nice ring, and the picture may be geometrically clearer to you (in the same way that thinking about the Blow-up Closure Lemma 23.2.6 in the case where $Z \rightarrow Y$ is a closed embedding is more intuitive). Another reason is that, as you will find in the exercises, you will see some facts more easily in this explicit example, and you can then pull them back to more general examples. Perhaps most important, Approach 1 lets you actually compute blow-ups by working affine-locally: if f_1, \dots, f_n are elements of a ring A , cutting a subscheme $X = \text{Spec } A/(f_1, \dots, f_n)$ of $Y = \text{Spec } A$, then $\text{Bl}_X Y$ can be interpreted as a closed subscheme of \mathbb{P}_A^{n-1} , by pulling back from $\text{Bl}_{V(x_1, \dots, x_n)} \text{Spec } \mathbb{Z}[x_1, \dots, x_n]$, and taking the closure of the locus “above X ” as dictated by the Blow-up Closure Lemma 23.2.6.

Proof. Reduce to the case of affine target $\text{Spec } R$ with ideal $I \subset R$. Reduce to the case of affine source, with principal effective Cartier divisor t . (A principal effective Cartier divisor is locally cut out by a single non-zerodivisor.) Thus we have reduced to the case $\text{Spec } S \rightarrow \text{Spec } R$, corresponding to $f: R \rightarrow S$. Say $(x_1, \dots, x_n) = I$, with $(f(x_1), \dots, f(x_n)) = (t)$. We will describe *one* map $\text{Spec } S \rightarrow \text{Proj } R[I]$ that will extend the map on the open set $\text{Spec } S_t \rightarrow \text{Spec } R$. It is then unique, by Exercise 11.4.G. We map $R[I]$ to S as follows: the degree one part is $f: R \rightarrow S$, and $f(X_i)$ (where X_i corresponds to x_i , except it is in degree 1) goes to $f(x_i)/t$. Hence an element X of degree d goes to $X/(t^d)$. On the open set $D_+(X_1)$, we get the map $R[X_2/X_1, \dots, X_n/X_1]/(x_2 - X_2/X_1 x_1, \dots, x_i X_j - x_j X_i, \dots) \rightarrow S$ (where there may be many relations) which agrees with f away from $V(t)$. Thus this map does extend over $V(I)$. \square

Here are some applications and observations arising from this construction of the blow-up. First, we can verify that our initial motivational examples are indeed blow-ups. For example, blowing up \mathbb{A}^2 (with coordinates x and y) at the origin yields: $B = k[x, y]$, $I = (x, y)$, and $\text{Proj}(B \oplus I \oplus I^2 \oplus \dots) = \text{Proj } B[X, Y]$ where the elements of B have degree 0, and X and Y are degree 1 and “correspond to” x and y respectively.

23.3.5. Normal bundles to exceptional divisors. The normal bundle to an effective Cartier divisor D is the (space associated to the) invertible sheaf $\mathcal{O}(D)|_D$, the invertible sheaf corresponding to the D on the total space, then restricted to D (Exercise 22.2.H). Thus in the case of the blow-up of a point in the plane, the exceptional divisor has normal bundle $\mathcal{O}(-1)$. (As an aside: Castelnuovo’s Criterion, Theorem 30.7.1, states that conversely given a smooth surface containing $E \cong \mathbb{P}^1$ with normal bundle $\mathcal{O}(-1)$, E can be blown-down to a point on another smooth surface.) In the case of the blow-up of a regular subvariety of a regular variety, the blow up turns out to be regular (see Theorem 23.3.10), the exceptional divisor is a projective bundle over X , and the normal bundle to the exceptional divisor restricts to $\mathcal{O}(-1)$ (Exercise 23.3.D).

23.3.A. HARDER BUT ENLIGHTENING EXERCISE. If $i: X \hookrightarrow \mathbb{P}^n$ is a projective scheme, identify the exceptional divisor of the blow up of the affine cone over X (§9.3.11) at the origin with X itself, and that its normal bundle (§23.3.5) is isomorphic to $\mathcal{O}_X(-1) := i^* \mathcal{O}_{\mathbb{P}^n}(-1)$. (In the case $X = \mathbb{P}^1$, we recover the blow-up of

the plane at a point. In particular, we recover the important fact that the normal bundle to the exceptional divisor is $\mathcal{O}(-1)$.)

23.3.6. The normal cone. Motivated by (23.3.2.1), as well as Exercise 23.3.D below, we make the following definition. If X is a closed subscheme of Y cut out by \mathcal{I} , then the **normal cone** $N_X Y$ of X in Y is defined as

$$(23.3.6.1) \quad N_X Y := \operatorname{Spec}_X (\mathcal{O}_Y / \mathcal{I} \oplus \mathcal{I} / \mathcal{I}^2 \oplus \mathcal{I}^2 / \mathcal{I}^3 \oplus \cdots).$$

This can profitably be thought of as an algebro-geometric version of an infinitesimal “tubular neighborhood”. But some cautions are in order. If Y is smooth, $N_X Y$ may not be smooth. (You can work out the example of $Y = \mathbb{A}_k^2$ and $X = V(xy)$.) And even if X and Y are smooth, then although $N_X Y$ is smooth (as we will see shortly, Exercise 23.3.D), it doesn’t “embed” in any way in Y (see Remark 23.3.9).

If X is a closed point p , then the normal cone is called the **tangent cone** to Y at p . The **projectivized tangent cone** is the exceptional divisor $E_X Y$ (the *Proj* of the same graded sheaf of algebras). Following §9.3.12, the tangent cone and the projectivized tangent cone can be put together in the projective completion of the tangent cone, which contains the tangent cone as an open subset, and the projectivized tangent cone as a complementary effective Cartier divisor.

In Exercise 23.3.D, we will see that at a regular point of Y , the tangent cone may be identified with the tangent space, and the normal cone may often be identified with the total space of the normal bundle.

23.3.B. EXERCISE. Suppose $Y = \operatorname{Spec} k[x, y] / (y^2 - x^2 - x^3)$ (the nodal cubic, see the bottom of Figure 8.4). Assume (to avoid distraction) that $\operatorname{char} k \neq 2$. Show that the tangent cone to Y at the origin is isomorphic to $\operatorname{Spec} k[x, y] / (y^2 - x^2)$. Thus, informally, the tangent cone “looks like” the original variety “infinitely magnified”.

23.3.C. EXERCISE. Suppose S_\bullet is a finitely generated graded algebra over a field k . Exercise 23.3.A gives an isomorphism of $\operatorname{Proj} S_\bullet$ with the exceptional divisor to the blow-up of $\operatorname{Spec} S_\bullet$ at the origin. Show that the tangent cone to $\operatorname{Spec} S_\bullet$ at the origin is isomorphic to $\operatorname{Spec} S_\bullet$ itself. (Your geometric intuition should lead you to find these facts believable.)

23.3.7. Blowing up regular embeddings.

The case of blow-ups of regular embeddings $X \subset Y$ is particularly pleasant. For example, the exceptional divisor is a projective bundle over X . The central reason is the following result.

23.3.8. Theorem. — *If $I \subset A$ is generated by a regular sequence a_1, \dots, a_d , then the natural map $\operatorname{Sym}_A^n(I/I^2) \rightarrow I^n/I^{n+1}$ is an isomorphism.*

Hence if a closed embedding $i: X \hookrightarrow Y$ is a regular embedding with ideal sheaf $\mathcal{I} \subset \mathcal{O}_Y$, then the natural map $\operatorname{Sym}^n(\mathcal{I}/\mathcal{I}^2) \rightarrow \mathcal{I}^n/\mathcal{I}^{n+1}$ is an isomorphism. Furthermore, in combination with Proposition 22.2.16, we see that $\mathcal{I}^n/\mathcal{I}^{n+1}$ is a locally free sheaf.

Before starting the proof of Theorem 23.3.8 in §23.3.11, we show its utility.

23.3.D. EXERCISE (ASSUMING THEOREM 23.3.8).

(a) Suppose $X \rightarrow Y$ is a regular embedding with ideal sheaf \mathcal{I} , identify the total space (§18.1.4) of the normal sheaf (the “normal bundle”) with the normal cone

$N_X Y$ (23.3.6.1), and show that the exceptional divisor $E_X Y$ is a projective bundle (the “projectivized normal bundle”) over X .

(b) Show that the normal bundle to $E_X Y$ in $\text{Bl}_X Y$ is $\mathcal{O}(-1)$ (for the projective bundle over X).

(c) Assume further that X is a reduced closed point p . Show that p is a regular point of Y . Identify the total space of the tangent space to p with the tangent cone to Y at p .

23.3.9. Remark. We can now make sense of a comment made in §23.3.6, that even if X and Y are smooth, then although $N_X Y$ is smooth, it needn’t admit an open embedding in Y . To do this, start with a smooth complex quartic surface Y containing a line X . (Most smooth quartic surfaces don’t contain a line, by Exercise 12.2.L, so you will have to construct one by hand.) Then $N_X Y$ is a line bundle over X , and thus rational (i.e., birational to $\mathbb{A}_{\mathbb{C}}^2$, Definition 7.5.4). But $N_X Y$ cannot admit an open embedding into Y , as otherwise Y would be rational, contradicting Exercise 22.5.D.

23.3.10. Theorem. — Suppose $X \hookrightarrow Y$ is a closed embedding of smooth varieties over k . Then $\text{Bl}_X Y$ is also smooth.

Proof. (We use the fact that smooth varieties are regular, the Smoothness-Regularity Comparison Theorem 13.2.10(b), whose proof we still have to complete.)

We need only check smoothness of $\text{Bl}_X Y$ at the points of $E_X Y$ (by Observation 23.2.2). By Exercise 13.2.L(b), $X \hookrightarrow Y$ is a regular embedding. Then by Exercise 23.3.D(a), $E_X Y$ is a projective bundle over X , and thus smooth, and hence regular at its closed points. But $E_X Y$ is an effective Cartier divisor on $\text{Bl}_X Y$. By the slicing criterion for regularity (Exercise 13.2.C), it follows that $\text{Bl}_X Y$ is regular at the closed points of $E_X Y$, hence smooth at all points of $E_X Y$. \square

23.3.11. * Proving Theorem 23.3.8.

The proof of Theorem 23.3.8 may reasonably be skipped on a first reading. We prove Theorem 23.3.8 following [F2, A.6.1], which in turn follows [Dav]. The proof will be completed in §23.3.13. To begin, let α be the map of graded rings

$$\alpha: (A/I)[X_1, \dots, X_d] \rightarrow \bigoplus_{n=0}^{\infty} I^n / I^{n+1}$$

which takes X_i to the image of a_i in I/I^2 . Clearly α is surjective.

23.3.E. EXERCISE. Show that Theorem 23.3.8 would follow from the statement that α is an isomorphism.

Because a_1 is a non-zerodivisor, we can interpret $A[a_2/a_1, \dots, a_d/a_1]$ as a subring of the total fraction ring (defined in §6.5.33). In particular, as A is a subring of its total fraction ring, the map $A \rightarrow A[a_2/a_1, \dots, a_d/a_1]$ is an injection. Define

$$\beta: A[T_2, \dots, T_d] \rightarrow A[a_2/a_1, \dots, a_d/a_1]$$

by $T_i \mapsto a_i/a_1$. Clearly, the map β is surjective, and $L_i := a_1 T_i - a_i$ lies in $\ker(\beta)$.

23.3.12. Lemma. — The kernel of β is (L_2, \dots, L_d) .

Proof. We prove the result by induction on d . We consider first the base case $d = 2$. Suppose $F[T_2] \in \ker \beta$, so $F(a_2/a_1) = 0$. Then applying the algorithm for the

Remainder Theorem, dividing $a_1^{\deg F} F(T_2)$ by $a_1 T_2 - a_2 = L_2$,

$$(23.3.12.1) \quad a_1^{\deg F} F(T_2) = G(T_2)(a_1 T_2 - a_2) + R$$

where $G(T_2) \in A[T_2]$, and $R \in A$ is the remainder. Substituting $T_2 = a_2/a_1$ (and using the fact that $A \rightarrow A[a_2/a_1]$ is injective), we have that $R = 0$. Then $(a_1 T_2 - a_2)G(T_2) \equiv 0 \pmod{(a_1^{\deg F})}$. Using the fact that a_2 is a non-zerodivisor modulo $(a_1^{\deg F})$ (as $a_1^{\deg F}, a_2$ is a regular sequence by Exercise 9.5.E), a short induction shows that the coefficients of $G(T_2)$ must all be divisible by $a_1^{\deg F}$. Thus $F(T_2)$ is divisible by $a_1 T_2 - a_2 = L_2$, so the case $d = 2$ is proved.

We now consider the general case $d > 2$, assuming the result for all smaller d . Let $A' = A[a_2/a_1]$. Then $a_1, a_3, a_4, \dots, a_d$ is a regular sequence in A' . Reason: a_1 is a non-zerodivisor in A' . $A'/(a_1) = (A[T_2]/(a_1 T_2 - a_2))/(a_1) = A[T_2]/(a_1 T_2 - a_2, a_1) = A[T_2]/(a_1, a_2)$. Then a_3 is a non-zerodivisor in this ring because it is a non-zerodivisor in $A/(a_1, a_2)$, a_4 is a non-zerodivisor in $A[T_2]/(a_1, a_2, a_3)$ because it is a non-zerodivisor in $A/(a_1, a_2, a_3)$, and so forth. Condition (ii) of the Definition 9.5.4 of regular sequence also holds, as

$$A'/(a_1, a_3, a_4, \dots, a_d) = A[T_2]/(a_1, a_2, \dots, a_d) \neq 0.$$

Consider the composition

$$A[T_2, \dots, T_d] \rightarrow A'[T_3, \dots, T_d] \rightarrow A'[a_3/a_1, \dots, a_d/a_1] = A[a_2/a_1, \dots, a_d/a_1].$$

By the case $d = 2$, the kernel of the first map is L_2 . By the inductive hypothesis, the kernel of the second map is (L_3, \dots, L_d) . The result follows. \square

23.3.13. Proof of Theorem 23.3.8. By Exercise 23.3.E, it suffices to prove that the surjection α is an isomorphism. Suppose $F \in \ker(\alpha)$; we wish to show that $F = 0$. We may assume that F is homogeneous, say of degree n . Consider the map $\alpha': A[X_1, \dots, X_d] \rightarrow \bigoplus_{n=0}^{\infty} I^n/I^{n+1}$ lifting α . Lift F to $A[X_1, \dots, X_d]$, so $F \in \ker(\alpha')$. We wish to show that $F \in IA[X_1, \dots, X_d]$. Suppose $F(a_1, \dots, a_d) = x \in I^{n+1}$. Then we can write x as $F'(a_1, \dots, a_d)$, where F' is a homogeneous polynomial of the same degree as F , with coefficients in I . Then by replacing F by $F - F'$, we are reduced to the following problem: suppose $F \in A[X_1, \dots, X_d]$ is homogeneous of degree n , and $F(a_1, \dots, a_d) = 0$, we wish to show that $F \in IA[X_1, \dots, X_n]$. But if $F(a_1, \dots, a_d) = 0$, then $F(1, a_2/a_1, \dots, a_d/a_1) = 0$ in $A[a_2/a_1, \dots, a_d/a_1]$. But Lemma 23.3.12 identifies the kernel of β , so

$$F(1, T_2, T_3, \dots, T_d) \in (a_1 T_2 - a_2, a_1 T_3 - a_3, \dots, a_1 T_d - a_d).$$

Thus the coefficients of F are in $(a_1, \dots, a_d) = I$ as desired. \square

23.4 Examples and computations

In this section we will work through a number of explicit examples, to get a sense of how blow-ups behave, how they are useful, and how one can work with them explicitly. **Throughout we work over a field k , and we assume throughout that $\text{char } k = 0$ to avoid distraction.** The examples and exercises are loosely arranged by topic, but not in order of importance.

23.4.1. Example: Blowing up the plane along the origin.

Let's first blow up the plane \mathbb{A}^2 along the origin, and see that the result agrees with our discussion in §23.1. Let x and y be the coordinates on \mathbb{A}^2 . The blow-up is $\text{Proj } k[x, y, X, Y]$, where $xY - yX = 0$. (Here x and y have degree 0 and X and Y have degree 1.) This is naturally a closed subscheme of $\mathbb{A}^2 \times \mathbb{P}^1$, cut out (in terms of the projective coordinates X and Y on \mathbb{P}^1) by $xY - yX = 0$. We consider the two usual patches on \mathbb{P}^1 : $[X, Y] = [s, 1]$ and $[1, t]$. The first patch yields $\text{Spec } k[x, y, s]/(sy - x)$, and the second gives $\text{Spec } k[x, y, t]/(y - xt)$. Notice that both are smooth: the first is $\text{Spec } k[y, s] \cong \mathbb{A}^2$, and the second is $\text{Spec } k[x, t] \cong \mathbb{A}^2$.

We now describe the exceptional divisor. We first consider the first (s) patch. The ideal is generated by (x, y) , which in our ys -coordinates is $(ys, y) = (y)$, which is indeed principal. Thus on this patch the exceptional divisor is generated by y . Similarly, in the second patch, the exceptional divisor is cut out by x . (This can be a little confusing, but there is no contradiction!) This explicit description will be useful in working through some of the examples below.

23.4.A. EXERCISE. Let p be a k -valued point of \mathbb{P}^2 . Exhibit an isomorphism between $\text{Bl}_p \mathbb{P}^2$ and the Hirzebruch surface $\mathbb{F}_1 = \mathbb{P}_{\mathbb{P}^1}(\mathcal{O}_{\mathbb{P}^1} \oplus \mathcal{O}_{\mathbb{P}^1}(1))$ (Example 18.2.4). (The map $\text{Bl}_p \mathbb{P}^2 \rightarrow \mathbb{P}^1$ informally corresponds to taking a point to the line connecting it to the origin. Do not be afraid: You can do this by explicitly working with coordinates.)

23.4.2. Resolving singularities.

23.4.3. The proper transform of a nodal curve (Figure 23.1). (You may wish to flip to Figure 8.4 while thinking through this exercise.) Consider next the curve $y^2 = x^3 + x^2$ inside the plane \mathbb{A}^2 (the nodal cubic, see Exercise 23.3.B). Let's blow up the origin, and compute the total and proper transform of the curve. (By the Blow-up Closure Lemma 23.2.6, the latter is the blow-up of the nodal curve at the origin.) In the first patch, we get $y^2 - s^2y^2 - s^3y^3 = 0$. This factors: we get the exceptional divisor y with multiplicity two, and the curve $1 - s^2 - s^3y = 0$. You can easily check that the proper transform is regular. Also, notice that the proper transform \tilde{C} meets the exceptional divisor at two points, $s = \pm 1$. This corresponds to the two tangent directions at the origin (as $s = x/y$).

23.4.B. EXERCISE (FIGURE 23.1). Describe both the total and proper transform of the curve C given by $y = x^2 - x$ in $\text{Bl}_{(0,0)} \mathbb{A}^2$. Show that the proper transform of C is isomorphic to C . Interpret the intersection of the proper transform of C with the exceptional divisor E as the slope of C at the origin.

23.4.C. EXERCISE: BLOWING UP A CUSPIDAL PLANE CURVE (CF. EXERCISE 10.7.F). Describe the proper transform of the cuspidal curve C given by $y^2 = x^3$ in the blown-up plane $\text{Bl}_{(0,0)} \mathbb{A}^2$. Show that it is regular. Show that the proper transform of C meets the exceptional divisor E at one point, and is tangent to E there.

The previous two exercises are the first in an important sequence of singularities, which we now discuss.




FIGURE 23.1. Resolving curve singularities (§23.4.3, Exercise 23.4.B, and Exercise 23.4.C)

23.4.D. EXERCISE: RESOLVING A_n CURVE SINGULARITIES. Resolve the singularity $y^2 = x^{n+1}$ in \mathbb{A}^2 , by first blowing up its singular point, then considering its proper transform and deciding what to do next.

23.4.4. Toward a definition of A_n curve singularities. You will notice that your solution to Exercise 23.4.D depends only on the “power series expansion” of the singularity at the origin, and not on the precise equation. For example, if you compare §23.4.3 with the $n = 1$ case of Exercise 23.4.D, you will see that they are “basically the same”. We will make this precise in Definition 30.3.C.

23.4.E. EXERCISE (WARM-UP TO EXERCISE 23.4.F). Blow up the cone point $z^2 = x^2 + y^2$ (Figure 3.4) at the origin. Show that the resulting surface is regular. Show that the exceptional divisor is isomorphic to \mathbb{P}^1 . (Remark: You can check that the normal bundle to this \mathbb{P}^1 is not $\mathcal{O}(-1)$, as is the case when you blow up a point on a smooth surface, see §23.3.5; it is $\mathcal{O}(-2)$.)

23.4.F. EXERCISE (RESOLVING A_n SURFACE SINGULARITIES). Resolve the singularity $z^2 = y^2 + x^{n+1}$ in \mathbb{A}^3 by first blowing up its singular point, then considering its proper transform, and deciding what to do next. (A k -surface singularity analytically isomorphic to this is called an A_n **surface singularity**. For example, the cone shown in Figure 3.4 is an A_1 surface singularity. We make this precise in Exercise 30.3.C.) This exercise is a bit time consuming, but is rewarding in that it shows that you can really resolve singularities by hand.

23.4.5. Remark: ADE-surface singularities and Dynkin diagrams (see Figure 23.2). A k -singularity analytically isomorphic to $z^2 = x^2y + y^{n-1}$ (resp. $z^2 = x^3 + y^4$, $z^2 = x^3 + xy^3$, $z^2 = x^3 + y^5$) is called a D_n surface singularity (resp. E_6 , E_7 , E_8 surface singularity). We will make this precise in Exercise 30.3.C, and you will then be able to guess the definition of the corresponding curve singularity. If you

(minimally) desingularize each of these surfaces by sequentially blowing up singular points as in Exercise 23.4.F, and look at the arrangement of exceptional divisors (the various exceptional divisors and how they meet), you will discover the corresponding Dynkin diagram. More precisely, if you create a graph, where the vertices correspond to exceptional divisors, and two vertices are joined by an edge if the two divisors meet, you will find the underlying graph of the corresponding Dynkin diagram. This is the start of several very beautiful stories; see Remark 27.4.5 for a first glimpse of one of them.




FIGURE 23.2. The exceptional divisors for resolutions of some ADE surface singularities, and their corresponding dual graphs (see Remark 23.4.5)

23.4.6. Remark: Resolution of singularities. Hironaka's Theorem on resolution of singularities implies that this idea of trying to resolve singularities by blowing up singular loci in general can succeed in characteristic 0. More precisely, if X is a variety over a field of characteristic 0, then X can be resolved by a sequence of blow-ups, where the n th blow-up is along a regular subvariety that lies in the singular locus of the variety produced after the $(n - 1)$ st stage (see [Hir], and [Ko2]). (The case of dimension 1 will be shown in §30.5.4, and the case of dimension 2 will be discussed in §30.7.4.) It is not known if an analogous statement is true in positive characteristic (except in dimension at most 2, by Lipman, see [Ar3]), but de Jong's Alteration Theorem [dJ] gives a result which is good enough for most applications. Rather than producing a birational proper map $\tilde{X} \rightarrow X$ from something regular, it produces a proper map from something regular that is generically finite (and the corresponding extension of function fields is separable).

Here are some other exercises related to resolution of singularities.

23.4.G. EXERCISE. Blowing up a nonreduced subscheme of a regular scheme can give you something singular, as shown in this example. Describe the blow up of the closed subscheme $V(y, x^2)$ in $\text{Spec } k[x, y] = \mathbb{A}^2$. Show that you get an A_1 surface singularity.

23.4.H. EXERCISE. Desingularize the tacnode (see Exercise 10.7.G and Definition 30.3.3) $y^2 = x^4$, not in two steps (as in Exercise 23.4.D), but in a single step by blowing up (y, x^2) .

23.4.I. EXERCISE (RESOLVING A SINGULARITY BY AN UNEXPECTED BLOW-UP). Suppose Y is the cone $x^2 + y^2 = z^2$, and X is the line cut out by $x = 0, y = z$ on Y . Show that $\text{Bl}_X Y$ is regular. (In this case we are blowing up a codimension 1 locus that

is not an effective Cartier divisor, see Problem 13.1.3. But it *is* an effective Cartier divisor away from the cone point, so you should expect your answer to be an isomorphism away from the cone point.)

23.4.7. Multiplicity of a function at a regular point. In order to pose Exercise 23.4.J, we introduce a useful concept. If f is a function on a locally Noetherian scheme X , its **multiplicity at a regular point** p is the largest m such that f lies in the m th power of the maximal ideal in the local ring $\mathcal{O}_{X,p}$. For example, if $f \neq 0$, $V(f)$ is singular at p if and only if $m > 1$. (Do you see why?)

23.4.J. EXERCISE. Show that the multiplicity of the exceptional divisor in the total transform of a subscheme Z of \mathbb{A}^n when you blow up the origin is the smallest multiplicity (at the origin) of a defining equation of Z . (For example, in the case of the nodal and cuspidal curves above, Example 23.4.3 and Exercise 23.4.C respectively, the exceptional divisor appears with multiplicity 2.)

23.4.8. Resolving rational maps.

23.4.K. EXERCISE (UNDERSTANDING THE BIRATIONAL MAP $\mathbb{P}^2 \dashrightarrow \mathbb{P}^1 \times \mathbb{P}^1$ VIA BLOW-UPS). Let p and q be two distinct k -points of \mathbb{P}^2 , and let r be a k -point of $\mathbb{P}^1 \times \mathbb{P}^1$. Describe an isomorphism $\mathrm{Bl}_{\{p,q\}} \mathbb{P}^2 \xrightarrow{\sim} \mathrm{Bl}_r \mathbb{P}^1 \times \mathbb{P}^1$. (Possible hint: Consider lines ℓ through p and m through q ; the choice of such a pair corresponds to a point of $\mathbb{P}^1 \times \mathbb{P}^1$. A point s of \mathbb{P}^2 not on line \overline{pq} yields a pair of lines $(\overline{ps}, \overline{qs})$ of $\mathbb{P}^1 \times \mathbb{P}^1$. Conversely, a choice of lines (ℓ, m) such that neither ℓ and m is line \overline{pq} yields a point $s = \ell \cap m \in \mathbb{P}^2$. This describes a birational map $\mathbb{P}^2 \dashrightarrow \mathbb{P}^1 \times \mathbb{P}^1$. Exercise 23.4.A is related.)

Exercise 23.4.K is an example of the general phenomenon explored in the next two exercises.

23.4.L. HARDER BUT USEFUL EXERCISE (BLOW-UPS RESOLVE BASE LOCI OF RATIONAL MAPS TO PROJECTIVE SPACE). Suppose we have a scheme Y , an invertible sheaf \mathcal{L} , and a number of sections s_0, \dots, s_n of \mathcal{L} (a *linear series*, Definition 16.3.6). Then away from the closed subscheme X cut out by $s_0 = \dots = s_n = 0$ (the base locus of the linear series), these sections give a morphism to \mathbb{P}^n . Show that this morphism extends uniquely to a morphism $\mathrm{Bl}_X Y \rightarrow \mathbb{P}^n$, where this morphism corresponds to the invertible sheaf $(\beta^* \mathcal{L})(-E_X Y)$, where $\beta: \mathrm{Bl}_X Y \rightarrow Y$ is the blow-up morphism. In other words, “blowing up the scheme-theoretic base locus resolves this rational map”. Hint: it suffices to consider an affine open subset of Y where \mathcal{L} is trivial. Uniqueness might use Exercise 11.4.G.

23.4.9. Remarks. (i) Exercise 23.4.L immediately implies that blow-ups can be used to resolve rational maps to projective schemes $Y \dashrightarrow Z \hookrightarrow \mathbb{P}^n$.

(ii) The following interpretation is enlightening. The linear series on Y pulls back to a linear series on $\mathrm{Bl}_X Y$, and the base locus of the linear series on Y pulls back to the base locus on $\mathrm{Bl}_X Y$. The base locus on $\mathrm{Bl}_X Y$ is $E_X Y$, an effective Cartier divisor. Because $E_X Y$ is not just locally principal, but also locally a non-zerodivisor, it can be “divided out” from the $\beta^* s_i$ (yielding a section of $(\beta^* \mathcal{L})(-E_X Y)$, thereby removing the base locus, and leaving a base-point-free linear series. (In a sense

that can be made precise through the universal property, this is the smallest “modification” of Y that can remove the base locus.) If X is already Cartier (as for example happens with any nontrivial linear system if Y is a regular pure-dimensional curve), then we can remove a base locus by just “dividing out X ”.

(iii) You may wish to revisit Exercise 20.7.D, and interpret it in terms of Exercise 23.4.L.

23.4.10. Examples. (i) The rational map $\mathbb{P}^n \dashrightarrow \mathbb{P}^{n-1}$ given by $[x_0, \dots, x_n] \dashrightarrow [x_1, \dots, x_n]$, defined away from $p = [1, 0, \dots, 0]$, is resolved by blowing up p . Then by the Blow-up Closure Lemma 23.2.6, if Y is any locally closed subscheme of \mathbb{P}^n , we can project to \mathbb{P}^{n-1} once we blow up p in Y , and the invertible sheaf giving the map to \mathbb{P}^{n-1} is (somewhat informally speaking) $\beta^*(\mathcal{O}_{\mathbb{P}^n}(1)) \otimes \mathcal{O}(-E_p Y)$.

(ii) Consider two general homogeneous cubic functions C_1 and C_2 in three variables, yielding two cubic curves in \mathbb{P}^2 . They are smooth (by Bertini’s Theorem 13.4.2), and meet in 9 points p_1, \dots, p_9 (using our standing assumption that we work over an algebraically closed field). Then $[C_1, C_2]$ gives a rational map $\mathbb{P}^2 \dashrightarrow \mathbb{P}^1$. To resolve the rational map, we blow up p_1, \dots, p_9 . The result is (generically) an *elliptic fibration* $\text{Bl}_{\{p_1, \dots, p_9\}} \mathbb{P}^2 \rightarrow \mathbb{P}^1$. (This is by no means a complete argument.)

(iii) Fix six general points p_1, \dots, p_6 in \mathbb{P}^2 . There is a four-dimensional vector space of cubics vanishing at these points, and they vanish scheme-theoretically precisely at these points. This yields a rational map $\mathbb{P}^2 \dashrightarrow \mathbb{P}^3$, which is resolved by blowing up the six points. The resulting morphism turns out to be a closed embedding, and the image in \mathbb{P}^3 is a (smooth) cubic surface. This is the famous fact that the blow up of the plane at six general points may be represented as a (smooth) cubic in \mathbb{P}^3 . (Again, this argument is not intended to be complete.) See 27.4.3 for a more precise and more complete discussion.

In reasonable circumstances, Exercise 23.4.L has an interpretation in terms of graphs of rational maps.

23.4.M. EXERCISE. Suppose s_0, \dots, s_n are sections of an invertible sheaf \mathcal{L} on an integral scheme X , not all 0. By Remark 17.4.3, these data give a rational map $\phi: X \dashrightarrow \mathbb{P}^n$. Give an isomorphism between the graph of ϕ (§11.4.4) and $\text{Bl}_{V(s_0, \dots, s_n)} X$. (Your argument will not require working over a field k ; it will work in general.)

You may enjoy exploring the previous idea by working out how the Cremona transformation $\mathbb{P}^2 \dashrightarrow \mathbb{P}^2$ (Exercise 7.5.I) can be interpreted in terms of the graph of the rational map $[x, y, z] \dashrightarrow [1/x, 1/y, 1/z]$.

23.4.N. ★ EXERCISE. Resolve the rational map

$$\text{Spec } k[w, x, y, z]/(wz - xy) \dashrightarrow \mathbb{P}^1$$

from the cone over the quadric surface to the projective line. Let X be the resulting variety, and $\pi: X \rightarrow \text{Spec } k[w, x, y, z]/(wz - xy)$ the projection to the cone over the quadric surface. Show that π is an isomorphism away from the cone point, and that the preimage of the cone point is isomorphic to \mathbb{P}^1 (and thus has codimension 2, and therefore is different from the resolution obtained by simply blowing up the

cone point). Possible hint: if Q is the quadric in \mathbb{P}^3 cut out by $wz - xy = 0$, then factor the rational map as $\text{Spec } k[w, x, y, z]/(wz - xy) \setminus \{0\} \rightarrow Q$ (cf. Exercise 9.3.N), followed by the isomorphism $Q \xrightarrow{\sim} \mathbb{P}^1 \times \mathbb{P}^1$ (Example 10.6.2), followed by projection onto one of the factors.

This is an example of a small resolution. (A **small resolution** $X \rightarrow Y$ is a resolution where the locus of points of Y where the fiber has dimension r is of codimension greater than $2r$. We will not use this notion again in any essential way.) Notice that this resolution of the morphism involves blowing up the base locus $w = x = 0$, which is a cone over one of the lines on the quadric surface $wz = xy$. We are blowing up an effective Weil divisor, which is necessarily not Cartier as the blow-up is not an isomorphism. In Exercise 13.1.E, we saw that (w, x) was not principal, while here we see that (w, x) is not even locally principal. Essentially by Exercise 15.2.R, $V(w, x)$ cannot even be the support of a locally principal divisor.

23.4.11. Remark: Non-isomorphic small resolutions. If you instead resolved the map $[w, y]$, you would obtain a similar looking small resolution

$$\pi': X' \rightarrow \text{Spec } k[w, x, y, z]/(wz - xy)$$

(it is an isomorphism away from the origin, and the fiber over the origin is \mathbb{P}^1). But it is different! More precisely, there is no morphism $X \rightarrow X'$ making the following the diagram commute.

$$\begin{array}{ccc} X & \xrightarrow{\quad \quad} & X' \\ & \searrow \pi & \swarrow \pi' \\ & \text{Spec } k[w, x, y, z]/(wz - xy) & \end{array}$$

The birational map $X \dashrightarrow X'$ is called the **Atiyah flop**, [At1].

23.4.12. Factorization of birational maps. We end our discussion of resolution of rational maps by noting that just as Hironaka's Theorem states that one may resolve all singularities of varieties in characteristic 0 by a sequence of blow-ups along smooth subvarieties (§23.4.6), the *Weak Factorization Theorem* (first proved by Włodarczyk) states that any two birational varieties X and Y in characteristic 0 may be related by blow-ups and “blow-downs” along smooth subvarieties. More precisely, there are varieties $X_0, \dots, X_n, X_{01}, \dots, X_{(n-1)n}$, with $X_0 = X$ and $X_n = Y$, with morphisms $X_{i(i+1)} \rightarrow X_i$ and $X_{i(i+1)} \rightarrow X_{i+1}$ ($0 \leq i < n$) which are blow-ups of smooth subvarieties.

23.4.13. Blow-ups and line bundles.

23.4.O. EXERCISE (GENERALIZING EXERCISE 21.2.D). Suppose X is a regular projective surface over k , and p is a k -valued point. Let $\beta: \text{Bl}_p X \rightarrow X$ be the blow-up morphism, and let $E = E_p X$ be the exceptional divisor. Consider the exact sequence

$$\mathbb{Z} \xrightarrow{\gamma: 1 \mapsto [E]} \text{Pic}(\text{Bl}_p X) \xrightarrow{\alpha} \text{Pic}(\text{Bl}_p X \setminus E) \longrightarrow 0$$

from (15.2.8.1). (Technically, (15.2.8.1) has Pic replaced by Cl . But because regular local rings are unique factorization domains by §13.8.5, $\text{Pic} = \text{Cl}$ in this case by

Proposition 15.2.10.) Note that $\text{Bl}_p X \setminus E = X \setminus p$. Show that $\text{Pic}(X \setminus p) = \text{Pic } X$. Show that $\beta^*: \text{Pic } X \rightarrow \text{Pic } \text{Bl}_p X$ gives a section to α . Use §23.3.5 to show that $\mathcal{O}_X(E)|_E \cong \mathcal{O}_E(-1)$ (so in the language of Chapter 21, E is a (-1) -curve, Definition 21.2.7), and from that show that γ is an injection. Conclude that $\text{Pic } \text{Bl}_p X \cong \text{Pic } X \oplus \mathbb{Z}$. Describe how to find the intersection matrix on $N_{\mathbb{Q}}^1(\text{Bl}_p X)$ from that of $N_{\mathbb{Q}}^1(X)$.

23.4.P. EXERCISE. Suppose D is an effective Cartier divisor (a curve) on X . Let $\text{mult}_p D$ be the multiplicity of D at p (Exercise 23.4.J), and let D^{pr} be the proper transform of D . Show that $\pi^* D = D^{pr} + (\text{mult}_p D)E$ as effective Cartier divisors. More precisely, show that the product of the local equation for D^{pr} and the $(\text{mult}_p D)$ th power of the local equation for E is the local equation for $\pi^* D$, and hence that (i) $\pi^* D$ is an effective Cartier divisor, and (ii) $\pi^* \mathcal{O}_X(D) \cong \mathcal{O}_{\text{Bl}_p X}(D^{pr}) \otimes \mathcal{O}_{\text{Bl}_p X}(E)^{\otimes (\text{mult}_p D)}$. (A special case is the equation $\ell = e + m$ in Hint 21.2.6.)

23.4.14. Change of the canonical line bundle under blow-ups.

As motivation for how the canonical line bundle changes under blowing up, consider $\pi: \text{Bl}_{(0,0)} \mathbb{A}^2 \rightarrow \mathbb{A}^2$. Let $X = \text{Bl}_{(0,0)} \mathbb{A}^2$ and $Y = \mathbb{A}^2$ for convenience. We use Exercise 22.7.A to relate $\pi^* \mathcal{K}_Y$ to \mathcal{K}_X .

We pick a generator for \mathcal{K}_Y near $(0,0)$: $dx \wedge dy$. (This is in fact a generator for \mathcal{K}_Y everywhere on \mathbb{A}^2 , but for the sake of generalization, we point out that all that matters is that is a generator at $(0,0)$, and hence *near* $(0,0)$ by geometric Nakayama, Exercise 14.3.E.) When we pull it back to X , we can interpret it as a section of \mathcal{K}_X , which will generate \mathcal{K}_X away from the exceptional divisor E , but may contain E with some multiplicity μ . Recall that X can be interpreted as the data of a point in \mathbb{A}^2 as well as the choice of a line through the origin. We consider the open subset U where the line is not vertical, and thus can be written as $y = mx$. Here we have natural coordinates: $U = \text{Spec } k[x, y, m]/(y - mx)$, which we can interpret as $\text{Spec } k[x, m]$. The exceptional divisor E meets U , at $x = 0$ (in the coordinates on U), so we can calculate μ on this open set. Pulling back $dx \wedge dy$ to U , we get

$$dx \wedge dy = dx \wedge d(xm) = m(dx \wedge dx) + x(dx \wedge dm) = x(dx \wedge dm)$$

as $dx \wedge dx = 0$. Thus $\pi^*(dx \wedge dy)$ vanishes to order 1 along E .

23.4.Q. EXERCISE (CF. UNIMPORTANT EXERCISE 22.7.J). Explain how this determines a canonical isomorphism $\mathcal{K}_X \xrightarrow{\sim} (\pi^* \mathcal{K}_Y)(E)$.

23.4.R. EXERCISE. Repeat the above calculation in dimension n . Show that the exceptional divisor appears with multiplicity $(n - 1)$.

23.4.S. * EXERCISE. Suppose k is perfect.

- (a) Suppose Y is a surface over k , and p is a regular k -valued point, and let $\beta: X \rightarrow Y$ be the blow-up of Y at p . Show that $\mathcal{K}_X \cong (\beta^* \mathcal{K}_Y)(E)$. Hint: to find a generator of \mathcal{K}_X near p , choose generators \bar{x} and \bar{y} of $\mathfrak{m}/\mathfrak{m}^2$ (where \mathfrak{m} is the maximal ideal of $\mathcal{O}_{Y,p}$), and lift them to elements of $\mathcal{O}_{X,p}$. Why does $dx \wedge dy$ generate \mathcal{K}_X at p ?
- (b) Repeat part (a) in arbitrary dimension (following Exercise 23.4.R).
- (c) Suppose Z is a smooth m -dimensional (closed) subvariety of a smooth n -dimensional variety Y , and let $\beta: X \rightarrow Y$ be the blow-up of Y along Z . Show that $\mathcal{K}_X \cong (\beta^* \mathcal{K}_Y)((n - m - 1)E)$. (Recall from Theorem 23.3.10 that $X = \text{Bl}_Z Y$ is smooth.)

23.4.15. ★ Dimensional cohomology vanishing for quasiprojective schemes (promised in §19.2.7).

Using the theory of blowing up, Theorem 19.2.6 (dimensional cohomology vanishing for quasicoherent sheaves on projective k -schemes) can be extended to quasiprojective k -schemes. Suppose X is a quasiprojective k -variety of dimension n . We show that X may be covered by $n + 1$ affine open subsets. As X is quasiprojective, there is some projective variety Y with an open embedding $X \hookrightarrow Y$. By replacing Y with the closure of X in Y , we may assume that $\dim Y = n$. Put any subscheme structure Z on the complement of X in Y (for example the reduced subscheme structure, §9.4.9). Let $Y' = \text{Bl}_Z Y$. Then Y' is a projective variety (§23.3.1), which can be covered by $n + 1$ affine open subsets. The complement of X in Y' is an effective Cartier divisor $(E_Z Y)$, so the restriction to X of each of these affine open subsets of Y is also affine, by Exercise 8.3.E.

23.4.16. Remarks. (i) You might then hope that *any* dimension n variety can be covered by $n + 1$ affine open subsets. This is not true. For each integer m , there is a threefold that requires at least m affine open sets to cover it, see [RotV, Ex. 4.9]. By the discussion above, this example is necessarily not quasiprojective. (ii) Here is a fact useful in invariant theory, which can be proved in the same way. Suppose p_1, \dots, p_n are closed points on a quasiprojective k -variety X . Then there is an affine open subset of X containing all of them.

23.4.T. EXERCISE (DIMENSIONAL COHOMOLOGY VANISHING FOR QUASIPROJECTIVE VARIETIES). Suppose X is a quasiprojective k -scheme of dimension d . Show that for any quasicoherent sheaf \mathcal{F} on X , $H^i(X, \mathcal{F}) = 0$ for $i > d$.

23.4.17. ★★ Deformation to the normal cone (Figure 23.3).

The following construction is key to the modern understanding of intersection theory in algebraic geometry, as developed by Fulton and MacPherson, [F2].

23.4.U. EXERCISE: DEFORMATION TO THE NORMAL CONE. Suppose Y is a k -variety, and $X \hookrightarrow Y$ is a closed subscheme.

(a) Show that the exceptional divisor of $\beta: \text{Bl}_{X \times 0}(Y \times \mathbb{A}^1) \rightarrow Y \times \mathbb{A}^1$ is isomorphic to the projective completion of the normal cone to X in Y .

(b) Let $\pi: \text{Bl}_{X \times 0}(Y \times \mathbb{A}^1) \rightarrow \mathbb{A}^1$ be the composition of β with the projection to \mathbb{A}^1 . Show that $\pi^{-1}(0)$ is the scheme-theoretic union of $\text{Bl}_X Y$ with the projective completion of the normal cone to X in Y , and the intersection of these two subschemes may be identified with $E_X Y$, which is a closed subscheme of $\text{Bl}_X Y$ in the usual way (as the exceptional divisor of the blow-up $\text{Bl}_X Y \rightarrow Y$), and a closed subscheme of the projective completion of the normal cone as described in Exercise 9.3.O.

The map

$$\text{Bl}_{X \times 0}(Y \times \mathbb{A}^1) \setminus \text{Bl}_X Y \longrightarrow \mathbb{A}^1$$

is called the **deformation to the normal cone** (short for *deformation of Y to the normal cone of X in Y*). Notice that the fiber above every k -point away from $0 \in \mathbb{A}^1$ is canonically isomorphic to Y , and the fiber over 0 is the normal cone. Because this family is “nice” (more precisely, *flat*, the topic of Chapter 25), we can prove things about general Y (near X) by way of this degeneration. (We will see in §25.4.10

that the deformation to the normal cone is a *flat* morphism, which is useful in intersection theory.)




FIGURE 23.3. Deformation to the normal cone (§23.4.17)

Part VI

More cohomological tools

CHAPTER 24

Derived functors

Ça me semble extrêmement plaisant de ficher comme ça beaucoup de choses, pas drôles quand on les prend séparément, sous le grand chapeau des foncteurs dérivés.

I find it very agreeable to stick all sorts of things, which are not much fun when taken individually, together under the heading of derived functors.

— A. Grothendieck, letter to J.-P. Serre, February 18, 1955 [GrS, p. 6]

In this chapter, we discuss derived functors, introduced by Grothendieck in his celebrated “Tôhoku article” [Gr1], and their applications to sheaves. For quasicohherent sheaves on quasicompact separated schemes, derived functor cohomology will agree with Čech cohomology (§24.5). Čech cohomology will suffice for most of our purposes, and is quite down to earth and computable, but derived functor cohomology is worth seeing. First, it will apply much more generally in algebraic geometry (e.g., étale cohomology) and elsewhere, although this is beyond the scope of this book. Second, it will easily provide us with some useful notions, such as the Ext functors and the Leray spectral sequence. But derived functors can be intimidating the first time you see them, so feel free to just skim the main results, and to return to them later.

24.1 The Tor functors

We begin with a warm-up: the case of Tor. This is a hands-on example, but if you understand it well, you will understand derived functors in general. Tor will be useful to prove facts about flatness, which we will discuss in §25.3. Tor is short for “torsion” (see Remark 25.3.1).

If you have never seen this notion before, you may want to just remember its properties. But I will prove everything anyway — it is surprisingly easy.

The idea behind Tor is as follows. Whenever we see a right-exact functor, we always hope that it is the end of a long exact sequence. Informally, given a short exact sequence

$$(24.1.0.1) \quad 0 \longrightarrow N' \longrightarrow N \longrightarrow N'' \longrightarrow 0,$$

we hope $M \otimes_A N' \rightarrow M \otimes_A N \rightarrow M \otimes_A N'' \rightarrow 0$ will extend to a long exact sequence
(24.1.0.2)

$$\begin{aligned} \cdots &\longrightarrow \mathrm{Tor}_i^A(M, N') \longrightarrow \mathrm{Tor}_i^A(M, N) \longrightarrow \mathrm{Tor}_i^A(M, N'') \longrightarrow \cdots \\ &\longrightarrow \mathrm{Tor}_1^A(M, N') \longrightarrow \mathrm{Tor}_1^A(M, N) \longrightarrow \mathrm{Tor}_1^A(M, N'') \\ &\longrightarrow M \otimes_A N' \longrightarrow M \otimes_A N \longrightarrow M \otimes_A N'' \longrightarrow 0. \end{aligned}$$

More precisely, we are hoping for *covariant functors* $\mathrm{Tor}_i^A(M, \cdot)$ from A -modules to A -modules (functoriality giving 2/3 of the morphisms in (24.1.0.2)), with $\mathrm{Tor}_0^A(M, N) \equiv M \otimes_A N$, and natural connecting homomorphisms $\delta: \mathrm{Tor}_{i+1}^A(M, N'') \rightarrow \mathrm{Tor}_i^A(M, N')$ for every short exact sequence (24.1.0.1) giving the long exact sequence (24.1.0.2). (“Natural” means: given a morphism of short exact sequences, the natural square you would write down involving the δ -morphism must commute.)

It turns out to be not too hard to make this work, and this will also motivate derived functors. Let’s now define $\mathrm{Tor}_i^A(M, N)$.

Take any resolution \mathcal{R} of N by free modules:

$$\cdots \longrightarrow A^{\oplus n_2} \longrightarrow A^{\oplus n_1} \longrightarrow A^{\oplus n_0} \longrightarrow N \longrightarrow 0.$$

More precisely, build this resolution from right to left. Start by choosing generators of N as an A -module, giving us $A^{\oplus n_0} \rightarrow N \rightarrow 0$. Then choose generators of the kernel, and so on. Note that we are not requiring the n_i to be finite (although we could, if N is a finitely generated module and A is Noetherian). Truncate the resolution, by stripping off the last term N (replacing $\rightarrow N \rightarrow 0$ with $\rightarrow 0$). Then tensor with M (which does not preserve exactness). Note that $M \otimes (A^{\oplus n_i}) = M^{\oplus n_i}$, as tensoring with M commutes with arbitrary direct sums (Exercise 1.3.M). Let $\mathrm{Tor}_i^A(M, N)_{\mathcal{R}}$ be the homology of this complex at the i th stage ($i \geq 0$). The subscript \mathcal{R} reminds us that our construction depends on the resolution, although we will soon see that it is independent of \mathcal{R} .

We make some quick observations.

- $\mathrm{Tor}_0^A(M, N)_{\mathcal{R}} \cong M \otimes_A N$, canonically. Reason: as tensoring is right-exact, and $A^{\oplus n_1} \rightarrow A^{\oplus n_0} \rightarrow N \rightarrow 0$ is exact, we have that $M^{\oplus n_1} \rightarrow M^{\oplus n_0} \rightarrow M \otimes_A N \rightarrow 0$ is exact, and hence that the homology of the truncated complex $M^{\oplus n_1} \rightarrow M^{\oplus n_0} \rightarrow 0$ is $M \otimes_A N$.

- If $M \otimes \cdot$ is exact (i.e., M is *flat*, §1.6.11), then $\mathrm{Tor}_i^A(M, N)_{\mathcal{R}} = 0$ for all $i > 0$. (This characterizes flatness, see Exercise 24.1.D.)

Now given two modules N and N' and resolutions \mathcal{R} and \mathcal{R}' of N and N' , we can “lift” any morphism $N \rightarrow N'$ to a morphism of the two resolutions:

$$\begin{array}{ccccccc} \cdots & \longrightarrow & A^{\oplus n_i} & \longrightarrow & \cdots & \longrightarrow & A^{\oplus n_1} & \longrightarrow & A^{\oplus n_0} & \longrightarrow & N & \longrightarrow & 0 \\ & & \downarrow & & & & \downarrow & & \downarrow & & \downarrow & & \\ \cdots & \longrightarrow & A^{\oplus n'_i} & \longrightarrow & \cdots & \longrightarrow & A^{\oplus n'_1} & \longrightarrow & A^{\oplus n'_0} & \longrightarrow & N' & \longrightarrow & 0 \end{array}$$

We do this inductively on i . Here we use the freeness of $A^{\oplus n_i}$: if a_1, \dots, a_{n_i} are generators of $A^{\oplus n_i}$, to lift the map $b: A^{\oplus n_i} \rightarrow A^{\oplus n'_{i-1}}$ (the composition of the differential $A^{\oplus n_i} \rightarrow A^{\oplus n_{i-1}}$ with the previously constructed $A^{\oplus n_{i-1}} \rightarrow A^{\oplus n'_{i-1}}$) to $c: A^{\oplus n_i} \rightarrow A^{\oplus n'_i}$, we arbitrarily lift $b(a_j)$ from $A^{\oplus n'_{i-1}}$ to $A^{\oplus n'_i}$, and declare this to be $c(a_j)$. (Warning for people who care about such things: we are using the Axiom of Choice here.)

Denote the choice of lifts by $\mathcal{R} \rightarrow \mathcal{R}'$. Now truncate both complexes (remove column $N \rightarrow N'$) and tensor with M . Maps of complexes induce maps of homology (Exercise 1.6.D), so we have described maps (a priori depending on $\mathcal{R} \rightarrow \mathcal{R}'$)

$$\mathrm{Tor}_i^A(M, N)_{\mathcal{R}} \rightarrow \mathrm{Tor}_i^A(M, N')_{\mathcal{R}'}.$$

We say two maps of complexes $f: C_{\bullet} \rightarrow C'_{\bullet}$ and $g: C_{\bullet} \rightarrow C'_{\bullet}$ are **homotopic** if there is a sequence of maps $w: C_i \rightarrow C'_{i+1}$ such that $f - g = dw + wd$.

24.1.A. EXERCISE. Show that two homotopic maps give the same map on homology.

24.1.B. CRUCIAL EXERCISE. Show that any two lifts $\mathcal{R} \rightarrow \mathcal{R}'$ are homotopic.

We now pull these observations together. (Be sure to digest these completely!)

(1) We get a map of A -modules $\mathrm{Tor}_i^A(M, N)_{\mathcal{R}} \rightarrow \mathrm{Tor}_i^A(M, N')_{\mathcal{R}'}$, independent of the lift $\mathcal{R} \rightarrow \mathcal{R}'$.

(2) Hence for any two resolutions \mathcal{R} and \mathcal{R}' of an A -module N , we get a *canonical* isomorphism $\mathrm{Tor}_i^A(M, N)_{\mathcal{R}} \xrightarrow{\sim} \mathrm{Tor}_i^A(M, N)_{\mathcal{R}'}$. Here's why. Choose lifts $\mathcal{R} \rightarrow \mathcal{R}'$ and $\mathcal{R}' \rightarrow \mathcal{R}$. The composition $\mathcal{R} \rightarrow \mathcal{R}' \rightarrow \mathcal{R}$ is homotopic to the identity (as it is a lift of the identity map $N \rightarrow N$). Thus if $f_{\mathcal{R} \rightarrow \mathcal{R}'}: \mathrm{Tor}_i^A(M, N)_{\mathcal{R}} \rightarrow \mathrm{Tor}_i^A(M, N)_{\mathcal{R}'}$ is the map induced by $\mathcal{R} \rightarrow \mathcal{R}'$, and similarly $f_{\mathcal{R}' \rightarrow \mathcal{R}}$ is the map induced by $\mathcal{R}' \rightarrow \mathcal{R}$, then $f_{\mathcal{R}' \rightarrow \mathcal{R}} \circ f_{\mathcal{R} \rightarrow \mathcal{R}'}$ is the identity, and similarly $f_{\mathcal{R} \rightarrow \mathcal{R}'} \circ f_{\mathcal{R}' \rightarrow \mathcal{R}}$ is the identity.

(3) Hence $\mathrm{Tor}_i^A(M, \cdot)$ doesn't depend on the choice of resolution. It is a covariant functor $\mathrm{Mod}_A \rightarrow \mathrm{Mod}_A$.

24.1.1. Remark. Note that if N is a free module, then $\mathrm{Tor}_i^A(M, N) = 0$ for all M and all $i > 0$, as N has the trivial resolution $0 \rightarrow N \rightarrow N \rightarrow 0$ (it is "its own resolution").

Finally, we get long exact sequences:

24.1.2. Proposition. — For any short exact sequence (24.1.0.1) we get a long exact sequence of Tor 's (24.1.0.2).

Proof. Given a short exact sequence (24.1.0.1), choose resolutions of N' and N'' . Then use these to get a resolution for N as follows.

$$(24.1.2.1) \quad \begin{array}{ccccccc} & & 0 & & 0 & & 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ \cdots & \longrightarrow & A^{\oplus n'_1} & \longrightarrow & A^{\oplus n'_0} & \longrightarrow & N' \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ \cdots & \longrightarrow & A^{\oplus (n'_1 + n''_1)} & \longrightarrow & A^{\oplus (n'_0 + n''_0)} & \longrightarrow & N \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ \cdots & \longrightarrow & A^{\oplus n''_1} & \longrightarrow & A^{\oplus n''_0} & \longrightarrow & N'' \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ & & 0 & & 0 & & 0 \end{array}$$

The map $A^{\oplus (n'_{i+1} + n''_{i+1})} \rightarrow A^{\oplus (n'_i + n''_i)}$ is the composition $A^{\oplus n'_{i+1}} \rightarrow A^{\oplus n'_i} \hookrightarrow A^{\oplus (n'_i + n''_i)}$ along with a lift of $A^{\oplus n'_{i+1}} \rightarrow A^{\oplus n'_i}$ to $A^{\oplus (n'_i + n''_i)}$ ensuring that the middle row is a *complex*.

24.1.C. EXERCISE. Verify that it is possible to choose such a lift of $A^{\oplus n'_{i+1}} \rightarrow A^{\oplus n'_i}$ to $A^{\oplus (n'_i + n''_i)}$.

Hence the middle row of (24.1.2.1) is *exact* (not just a complex), using the long exact sequence in cohomology (Theorem 1.6.6), and the fact that the top and bottom rows are exact. Thus the middle row is a resolution, and (24.1.2.1) is a short exact sequence of resolutions. (This is sometimes called the *horseshoe construction*, as the filling in of the middle row looks like filling in the middle of a horseshoe.) It may be helpful to notice that the columns other than the “ N -column” are all “direct sum exact sequences”, and the horizontal maps in the middle row are “block upper triangular”.

Then truncate (removing the right column $0 \rightarrow N' \rightarrow N \rightarrow N'' \rightarrow 0$), tensor with M (obtaining a short exact sequence of complexes) and take cohomology, yielding the desired long exact sequence. \square

24.1.D. EXERCISE. Show that the following are equivalent conditions on an A -module M .

- (i) M is flat.
- (ii) $\mathrm{Tor}_i^A(M, N) = 0$ for all $i > 0$ and all A -modules N .
- (iii) $\mathrm{Tor}_1^A(M, N) = 0$ for all A -modules N .

24.1.3. Caution. Given that free modules are immediately seen to be flat, you might think that Exercise 24.1.D implies Remark 24.1.1. This would follow if we knew that $\mathrm{Tor}_i^A(M, N) \cong \mathrm{Tor}_i^A(N, M)$, which is clear for $i = 0$ (as \otimes is symmetric), but we won’t know this about Tor_i when $i > 0$ until Exercise 24.3.A.

24.1.E. EXERCISE. Show that the connecting homomorphism δ constructed above is independent of all of choices (of resolutions, etc.). Try to do this with as little

annoyance as possible. (Possible hint: given two sets of choices used to build (24.1.2.1), build a map — a three-dimensional diagram — from one version of (24.1.2.1) to the other version.)


24.1.F. UNIMPORTANT EXERCISE. Show that $\mathrm{Tor}_i^A(M, \cdot)$ is an *additive* functor (Definition 1.6.1). (We won't use this later, so feel free to skip it.)

We have thus established the foundations of Tor.

24.2 Derived functors in general

24.2.1. Projective resolutions. We used very little about free modules in the above construction of Tor — in fact we used only that free modules are **projective**, i.e., those modules P such that for any surjection $M \twoheadrightarrow N$, it is possible to lift any morphism $P \rightarrow N$ to $P \rightarrow M$:

(24.2.1.1)



(As noted in §24.1, this uses the axiom of choice.) Equivalently, $\mathrm{Hom}(P, \cdot)$ is an exact functor (recall that $\mathrm{Hom}(Q, \cdot)$ is always left-exact for any Q). More generally, the same idea yields the definition of a **projective object in any abelian category**. Hence by following through our entire argument with projective modules replacing free modules throughout, (i) we can compute $\mathrm{Tor}_i^A(M, N)$ by taking any projective resolution of N , and (ii) $\mathrm{Tor}_i^A(M, N) = 0$ for any projective A -module N .

24.2.A. EXERCISE. Show that an object P is projective if and only if every short exact sequence $0 \rightarrow Q \rightarrow R \rightarrow P \rightarrow 0$ splits. Hence show that an A -module P is projective if and only if P is a direct summand of a free module.

24.2.B. EXERCISE. Show that projective modules are flat. Hint: Exercise 24.2.A. Be careful if you want to use Exercise 24.1.D; see Caution 24.1.3. (In fact, finitely generated projective modules over local rings are even free, see Remark 25.4.6.)

24.2.2. Definition: Derived functors.

The above description was low-tech, but immediately generalizes drastically. All we are using is that $M \otimes_A \cdot$ is a right-exact functor, and that for any A -module N , we can find a surjection $P \twoheadrightarrow N$ from a projective module. In general, if F is *any* right-exact covariant functor from the category of A -modules to any abelian category, this construction will define a sequence of functors $L_i F$ such that $L_0 F = F$ and the $L_i F$'s give a long exact sequence. We can make this more general still. We say that an abelian category **has enough projectives** if for any object N there is a surjection onto it from a projective object. Then if F is any right-exact covariant functor from an abelian category with enough projectives to any abelian category, then we can define the **left derived functors** to F , denoted $L_i F$ ($i \geq 0$). You should

reread §24.1 and see that throughout we only use the fact we have a projective resolution (repeatedly lifting maps as in (24.2.1.1)), as well as the fact that F sends products to products (a consequence of additivity of the functor, see Remark 1.6.2) to show that F applied to (24.1.2.1) preserves the exactness of the columns.

24.2.C. EXERCISE. The notion of an **injective object** in an abelian category is dual to the notion of a projective object.

- (a) State precisely the definition of an injective object.
- (b) Define derived functors for (i) covariant left-exact functors (these are called **right derived functors**), (ii) contravariant left-exact functors (also called **right derived functors**), and (iii) contravariant right-exact functors (these are called **left derived functors**), making explicit the necessary assumptions of the category having enough injectives or projectives.

24.2.3. Notation. Recall from Definition 24.2.2 that if F is a right-exact functor, its (left) derived functors are denoted $L_i F$ ($i \geq 0$, with $L_0 F = F$). Similarly, if F is a left-exact functor, its (right-) derived functors are denoted $R^i F$. The i is a superscript, to indicate that the long exact sequence is “ascending in i ”.

24.2.4. The Ext functors.

24.2.D. EASY EXERCISE (AND DEFINITION): Ext FUNCTORS FOR A -MODULES, FIRST VERSION. As $\text{Hom}(\cdot, N)$ is a contravariant left-exact functor in Mod_A , which has enough projectives, define $\text{Ext}_A^i(M, N)$ as the i th right derived functor of $\text{Hom}(\cdot, N)$, applied to M . State the corresponding long exact sequence for Ext-modules.

24.2.E. EASY EXERCISE (AND DEFINITION): Ext FUNCTORS FOR A -MODULES, SECOND VERSION. The category Mod_A has enough injectives (see §24.2.5). As $\text{Hom}(M, \cdot)$ is a covariant left-exact functor in Mod_A , define $\text{Ext}_A^i(M, N)$ as the i th right derived functor of $\text{Hom}(M, \cdot)$, applied to N . State the corresponding long exact sequence for Ext-modules.

We seem to have a problem with the previous two exercises: we have defined $\text{Ext}^i(M, N)$ twice, and we have two different long exact sequences! Fortunately, these two definitions of Ext agree (see Exercise 24.3.B), and two long exact sequences for Ext are better than one.

24.2.F. EASY EXERCISE. Use the definition of Ext in Exercise 24.2.D to show that if A is a Noetherian ring, and M and N are finitely generated A -modules, then $\text{Ext}_A^i(M, N)$ is a finitely generated A -module.

Ext-functors (for sheaves) will play a key role in the proof of Serre duality, see §31.2.

24.2.5. ★ The category of A -modules has enough injectives. We will need the fact that Mod_A has enough injectives, but the details of the proof won’t come up again, so feel free to skip this discussion.

24.2.G. EXERCISE. Suppose Q is an A -module, such that for every ideal $I \subset A$, every homomorphism $I \rightarrow Q$ extends to $A \rightarrow Q$. Show that Q is an injective A -module. Hint: suppose $N \subset M$ is an inclusion of A -modules, and we are given $\beta: N \rightarrow Q$. We wish to show that β extends to $M \rightarrow Q$. Use the Axiom of Choice to show that among those A -modules N' with $N \subset N' \subset M$, such that β extends to N' , there is a maximal one. If this N' is not M , give an extension of β to $N' + Am$, where $m \in M \setminus N'$, obtaining a contradiction.

24.2.H. EASY EXERCISE (USING THE AXIOM OF CHOICE, IN THE GUISE OF ZORN'S LEMMA). Show that a \mathbb{Z} -module (i.e., abelian group) Q is injective if and only if it is **divisible** (i.e., for every $q \in Q$ and $n \in \mathbb{Z}^{\neq 0}$, there is $q' \in Q$ with $nq' = q$). Hence show that any quotient of an injective \mathbb{Z} -module is also injective.

24.2.I. EXERCISE. Show that the category of \mathbb{Z} -modules $\text{Mod}_{\mathbb{Z}} = \text{Ab}$ has enough injectives. (Hint: if M is a \mathbb{Z} -module, then write it as the quotient of a free \mathbb{Z} -module F by some K . Show that M is contained in the divisible group $(F \otimes_{\mathbb{Z}} Q)/K$.)

24.2.J. EXERCISE. Suppose Q is an injective \mathbb{Z} -module, and A is a ring. Show that $\text{Hom}_{\mathbb{Z}}(A, Q)$ is an injective A -module. Hint: First describe the A -module structure on $\text{Hom}_{\mathbb{Z}}(A, Q)$. You will only use the fact that \mathbb{Z} is a ring, and that A is an algebra over that ring.

24.2.K. EXERCISE. Show that Mod_A has enough injectives. Hint: suppose M is an A -module. By Exercise 24.2.I, we can find an inclusion of \mathbb{Z} -modules $M \hookrightarrow Q$ where Q is an injective \mathbb{Z} -module. Describe a sequence of inclusions of A -modules

$$M \hookrightarrow \text{Hom}_{\mathbb{Z}}(A, M) \hookrightarrow \text{Hom}_{\mathbb{Z}}(A, Q).$$

(The A -module structure on $\text{Hom}_{\mathbb{Z}}(A, M)$ is via the A -action on the left argument A , not via the A -action on the right argument M .) The right term is injective by the previous Exercise 24.2.J.

24.2.6. ★ Universal δ -functors.

(This discussion is best skipped on a first reading; you should move directly to §24.3. We won't use the notion of δ -functors until §31.1.10.) We now describe a more general variant of derived functors, as you may use them in the discussion of Serre duality in Chapter 31. The advantage of the notion of universal δ -functor is that we can apply it even in cases where \mathcal{A} does not have enough injectives.

Abstracting key properties of derived functors, we define the data of a (cohomological) δ -functor, from an abelian category \mathcal{A} to another abelian category \mathcal{B} . A **δ -functor** is a collection of additive functors $T^i: \mathcal{A} \rightarrow \mathcal{B}$ (where T^i is taken to be 0 if $i < 0$), along with morphisms $\delta^i: T^i(A'') \rightarrow T^{i+1}(A')$ for each short exact sequence

$$(24.2.6.1) \quad 0 \rightarrow A' \rightarrow A \rightarrow A'' \rightarrow 0$$

in \mathcal{A} , satisfying two properties:

- (i) (*short exact sequences yield long exact sequences*) For each short exact sequence (24.2.6.1), the sequence

$$\cdots \longrightarrow T^{i-1}(A'') \xrightarrow{\delta^{i-1}} T^i(A') \longrightarrow T^i(A) \longrightarrow T^i(A'') \xrightarrow{\delta^i} T^{i+1}(A') \longrightarrow \cdots$$

(where the unlabeled maps come from the covariance of the T^i) is exact.

In particular, T^0 is left-exact.

(ii) (*functoriality of (i)*) For each morphism of short exact sequences in \mathcal{A}

$$\begin{array}{ccccccc} 0 & \longrightarrow & a' & \longrightarrow & a & \longrightarrow & a'' \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & A' & \longrightarrow & A & \longrightarrow & A'' \longrightarrow 0 \end{array}$$

(where the squares commute), the δ^i 's give a commutative diagram

$$\begin{array}{ccc} T^i(a'') & \xrightarrow{\delta^i} & T^{i+1}(a') \\ \downarrow & & \downarrow \\ T^i(A'') & \xrightarrow{\delta^i} & T^{i+1}(A') \end{array}$$

(where the vertical arrows come from the covariance of T^i and T^{i+1}).

Derived functor cohomology is clearly an example of a δ -functor; Čech cohomology of sheaves on quasicompact separated schemes is another. (You can make these statements precise if you wish.)

24.2.L. EXERCISE. Figure out the right definition of **morphism of δ -functors** $\mathcal{A} \rightarrow \mathcal{B}$. (It should then be clear that the δ -functors from \mathcal{A} to \mathcal{B} form a category.)

24.2.7. Definition. A (cohomological) δ -functor (T^i, δ^i) is **universal** if for any other δ -functor (T'^i, δ'^i) , and any natural transformation of functors $\alpha: T^0 \rightarrow T'^0$, there is a unique morphism of δ -functors $(T^i, \delta^i) \rightarrow (T'^i, \delta'^i)$ extending α . By universal property nonsense (and Exercise 24.2.L), given any covariant left-exact functor $F: \mathcal{A} \rightarrow \mathcal{B}$, there is at most one universal δ -functor (T^i, δ^i) extending F (i.e., with a natural isomorphism $T^0 \xrightarrow{\sim} F$). The key fact about universal δ -functors is the following.

24.2.8. Theorem. — Suppose (T^i, δ^i) is a covariant δ -functor from \mathcal{A} to \mathcal{B} , and for all $A \in \mathcal{A}$, there exists a monomorphism $A \rightarrow J$ with $T^i J = 0$ for all $i > 0$. Then (T^i, δ^i) is universal.

24.2.M. ★★ EXERCISE. Prove Theorem 24.2.8. Partial hint: motivated by Corollary 24.2.10 below, follow our discussion of derived functors. Better hint (because this exercise is hard): follow the hints in [Weib, Exercise 2.4.5], or follow the proof of [Lan, Ch. XX, Thm. 7.1].

24.2.9. Remark. An additive functor $F: \mathcal{A} \rightarrow \mathcal{B}$ is said to be **effaceable** if for every $A \in \mathcal{A}$, there is a monomorphism $A \rightarrow J$ with $F(J) = 0$. The hypotheses of Theorem 24.2.8 can be weakened to require only that T^i is effaceable for each $i > 0$, and you are welcome to prove that instead. (Indeed, [Weib, Exercise 2.4.5], [Lan, Ch. XX, Thm. 7.1], and the original source [Gr1, II.2.2.1] deal with this case.) We give the statement of Theorem 24.2.8 for simplicity, as we will only use this version.

24.2.10. Corollary. — If \mathcal{A} has enough injectives, and F is a left-exact covariant functor $\mathcal{A} \rightarrow \mathcal{B}$, then the $R^i F$ (with the δ^i that accompany them) form a universal δ -functor.

Proof. Each element of \mathcal{A} admits a monomorphism into an injective element; this is just the definition of “enough injectives” (Exercise 24.2.C). Higher derived functors of an injective elements I are always 0: just compute the higher derived functor by taking the injective resolution of I “by itself”. \square

24.3 Derived functors and spectral sequences

A number of useful facts can be easily proved using spectral sequences. By doing these exercises, you will lose any fear of spectral sequence arguments in similar situations, as you will realize they are all the same.

Before you read this section, you should read §1.7 on spectral sequences.

24.3.1. Symmetry of Tor.

24.3.A. EXERCISE (SYMMETRY OF Tor). Show that there is an isomorphism $\mathrm{Tor}_i^A(M, N) \xrightarrow{\sim} \mathrm{Tor}_i^A(N, M)$. (Hint: take a free resolution of M and a free resolution of N . Take their “product” to somehow produce a double complex. Use both orientations of the obvious spectral sequence and see what you get.)

On a related note:

24.3.B. EXERCISE. Show that the two definitions of $\mathrm{Ext}^i(M, N)$ given in Exercises 24.2.D and 24.2.E agree.

24.3.2. Derived functors can be computed using acyclic resolutions. Suppose $F: \mathcal{A} \rightarrow \mathcal{B}$ is a right-exact additive functor of abelian categories, and that \mathcal{A} has enough projectives. (In other words, the hypotheses ensure the existence of left derived functors of F . Analogous facts will hold with the other types of derived functors, Exercise 24.2.C(b).) We say that $A \in \mathcal{A}$ is **F-acyclic** (or just **acyclic** if the F is clear from context) if $L_i F A = 0$ for $i > 0$.

The following exercise is a good opportunity to learn a useful trick (Hint 24.3.3).

24.3.C. EXERCISE. Show that you can also compute the derived functors of an objects B of \mathcal{A} using **acyclic resolutions** (not just projective resolutions), i.e., by taking a resolution

$$\cdots \longrightarrow A_2 \longrightarrow A_1 \longrightarrow A_0 \longrightarrow B \longrightarrow 0$$

by F -acyclic objects A_i , truncating, applying F , and taking homology. Hence $\mathrm{Tor}_i(M, N)$ can be computed with a flat resolution of M or N .

24.3.3. *Hint for Exercise 24.3.C (and a useful trick: building a “projective resolution of a complex”).* Show that you can construct a double complex

(24.3.3.1)

$$\begin{array}{ccccccc}
 & \vdots & & \vdots & & \vdots & & \vdots \\
 & \downarrow & & \downarrow & & \downarrow & & \downarrow \\
 \cdots & \longrightarrow & P_{2,1} & \longrightarrow & P_{1,1} & \longrightarrow & P_{0,1} & \longrightarrow & P_1 & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
 \cdots & \longrightarrow & P_{2,0} & \longrightarrow & P_{1,0} & \longrightarrow & P_{0,0} & \longrightarrow & P_0 & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
 \cdots & \longrightarrow & A_2 & \longrightarrow & A_1 & \longrightarrow & A_0 & \longrightarrow & B & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
 & & 0 & & 0 & & 0 & & 0 & &
 \end{array}$$

where the rows and columns are exact and the P_i 's are projective. Do this by constructing the P_i 's inductively from the bottom right. Remove the bottom row, and the right-most nonzero column, and then apply F , to obtain a new double complex. Use a spectral sequence argument to show that (i) the double complex has homology equal to $L_i F B$, and (ii) the homology of the double complex agrees with the construction given in the statement of the exercise. If this is too confusing, read more about the Cartan-Eilenberg resolution below.

24.3.4. The Grothendieck composition-of-functors spectral sequence.

Suppose \mathcal{A} , \mathcal{B} , and \mathcal{C} are abelian categories, $F: \mathcal{A} \rightarrow \mathcal{B}$ and $G: \mathcal{B} \rightarrow \mathcal{C}$ are a left-exact additive covariant functors, and \mathcal{A} and \mathcal{B} have enough injectives. Thus right derived functors of F , G , and $G \circ F$ exist. A reasonable question is: how are they related?

24.3.5. Theorem (Grothendieck composition-of-functors spectral sequence). —

Suppose $F: \mathcal{A} \rightarrow \mathcal{B}$ and $G: \mathcal{B} \rightarrow \mathcal{C}$ are left-exact additive covariant functors, and \mathcal{A} and \mathcal{B} have enough injectives. Suppose further that F sends injective elements of \mathcal{A} to G -acyclic elements of \mathcal{B} . Then for each $X \in \mathcal{A}$, there is a spectral sequence with $\rightarrow E_2^{p,q} = R^q G(R^p F(X))$ converging to $R^(G \circ F)(X)$.*

We will soon see the Leray spectral sequence as an application (Theorem 24.4.5).

There is more one might want to extract from the proof of Theorem 24.3.5. For example, although E_0 page of the spectral sequence will depend on some choices (of injective resolutions), the E_2 page will be independent of choice. For our applications, we won't need this refinement.

We will have to work to establish Theorem 24.3.5, so the proof is possibly best skipped on a first reading.

24.3.6. ★ Proving Theorem 24.3.5.

Before we give the proof (in §24.3.8), we begin with some preliminaries to motivate it. In order to discuss derived functors applied to X , we choose an injective

resolution of X :

$$0 \longrightarrow X \longrightarrow I^0 \longrightarrow I^1 \longrightarrow \cdots .$$

To compute the derived functors $R^p F(X)$, we apply F to the injective resolution I^\bullet :

$$0 \longrightarrow F(I^0) \longrightarrow F(I^1) \longrightarrow F(I^2) \longrightarrow \cdots .$$

Note that $F(I^p)$ is G -acyclic, by hypothesis of Theorem 24.3.5. If we were to follow our nose, we might take simultaneous injective resolutions $I^{\bullet,\bullet}$ of the terms in the above complex $F(I^\bullet)$ (the “dual” of Hint 24.3.3 — note that only the columns are required to be exact), and apply G , and consider the resulting double complex:

(24.3.6.1)

$$\begin{array}{ccccccc}
 & & \vdots & & \vdots & & \vdots \\
 & & \uparrow & & \uparrow & & \uparrow \\
 0 & \longrightarrow & G(I^{0,2}) & \longrightarrow & G(I^{1,2}) & \longrightarrow & G(I^{2,2}) \longrightarrow \cdots \\
 & & \uparrow & & \uparrow & & \uparrow \\
 0 & \longrightarrow & G(I^{0,1}) & \longrightarrow & G(I^{1,1}) & \longrightarrow & G(I^{2,1}) \longrightarrow \cdots \\
 & & \uparrow & & \uparrow & & \uparrow \\
 0 & \longrightarrow & G(I^{0,0}) & \longrightarrow & G(I^{1,0}) & \longrightarrow & G(I^{2,0}) \longrightarrow \cdots \\
 & & \uparrow & & \uparrow & & \uparrow \\
 & & 0 & & 0 & & 0
 \end{array}$$

24.3.D. EXERCISE. Consider the spectral sequence with upward orientation, starting with (24.3.6.1) as page E_0 . Show that $E_2^{p,q}$ is $R^p(G \circ F)(X)$ if $q = 0$, and 0 otherwise.

We now see half of the terms in the conclusion of Theorem 24.3.5; we are halfway there. To complete the proof, we would want to consider another spectral sequence, with rightward orientation, but we need to know more about (24.3.6.1); we will build it more carefully.

24.3.7. Cartan-Eilenberg resolutions.

Suppose $\cdots \rightarrow C^{p-1} \rightarrow C^p \rightarrow C^{p+1} \rightarrow \cdots$ is a complex in an abelian category \mathcal{B} . We will build an injective resolution of C^\bullet

(24.3.7.1)

$$\begin{array}{ccccccc}
 & & \vdots & & \vdots & & \vdots \\
 & & \uparrow & & \uparrow & & \uparrow \\
 0 & \longrightarrow & I^{0,2} & \longrightarrow & I^{1,2} & \longrightarrow & I^{2,2} \longrightarrow \cdots \\
 & & \uparrow & & \uparrow & & \uparrow \\
 0 & \longrightarrow & I^{0,1} & \longrightarrow & I^{1,1} & \longrightarrow & I^{2,1} \longrightarrow \cdots \\
 & & \uparrow & & \uparrow & & \uparrow \\
 0 & \longrightarrow & I^{0,0} & \longrightarrow & I^{1,0} & \longrightarrow & I^{2,0} \longrightarrow \cdots \\
 & & \uparrow & & \uparrow & & \uparrow \\
 0 & \longrightarrow & C^0 & \longrightarrow & C^1 & \longrightarrow & C^2 \longrightarrow \cdots \\
 & & \uparrow & & \uparrow & & \uparrow \\
 & & 0 & & 0 & & 0
 \end{array}$$

satisfying some further properties.

We first define some notation for functions on a complex.

- Let $Z^p(K^\bullet)$ be the kernel of the p th differential of a complex K^\bullet .
- Let $B^{p+1}(K^\bullet)$ be the image of the p th differential of a complex K^\bullet . (The superscript is chosen so that $B^{p+1}(K^\bullet) \subset K^{p+1}$.)
- As usual, let $H^p(K^\bullet)$ be the cohomology at the p th step of a complex K^\bullet .

For each p , we have complexes

(24.3.7.2)

$$\begin{array}{ccccc}
 & \vdots & & \vdots & & \vdots \\
 & \uparrow & & \uparrow & & \uparrow \\
 Z^p(I^{\bullet,1}) & & B^p(I^{\bullet,1}) & & H^p(I^{\bullet,1}) \\
 & \uparrow & & \uparrow & & \uparrow \\
 Z^p(I^{\bullet,0}) & & B^p(I^{\bullet,0}) & & H^p(I^{\bullet,0}) \\
 & \uparrow & & \uparrow & & \uparrow \\
 Z^p(C^\bullet) & & B^p(C^\bullet) & & H^p(C^\bullet) \\
 & \uparrow & & \uparrow & & \uparrow \\
 & 0 & & 0 & & 0
 \end{array}$$

We will construct (24.3.7.1) so that the three complexes (24.3.7.2) are all injective resolutions (of their first nonzero terms). We begin by choosing injective resolutions $B^{p,*}$ of $B^p(C^\bullet)$ and $H^{p,*}$ of $H^p(C^\bullet)$; these will eventually be the last two lines of (24.3.7.2).

24.3.E. EXERCISE. Describe an injective resolution $Z^{p,*}$ of $Z^p(C^\bullet)$ (the first line of (24.3.7.2)) making the following diagram a short exact sequence of complexes.

(24.3.7.3)

$$\begin{array}{ccccccc}
 & \vdots & & \vdots & & \vdots & \\
 & \uparrow & & \uparrow & & \uparrow & \\
 0 & \longrightarrow & B^{p,1} & \longrightarrow & Z^{p,1} & \longrightarrow & H^{p,1} \longrightarrow 0 \\
 & & \uparrow & & \uparrow & & \uparrow \\
 0 & \longrightarrow & B^{p,0} & \longrightarrow & Z^{p,0} & \longrightarrow & H^{p,0} \longrightarrow 0 \\
 & & \uparrow & & \uparrow & & \uparrow \\
 0 & \longrightarrow & B^p(C^\bullet) & \longrightarrow & Z^p(C^\bullet) & \longrightarrow & H^p(C^\bullet) \longrightarrow 0 \\
 & & \uparrow & & \uparrow & & \uparrow \\
 & & 0 & & 0 & & 0
 \end{array}$$

Hint: the “dual” problem was solved in (24.1.2.1), by a “horseshoe construction”.

24.3.F. EXERCISE. Describe an injective resolution $I^{p,*}$ of C^p making the following diagram a short exact sequence of complexes.

(24.3.7.4)

$$\begin{array}{ccccccc}
 & \vdots & & \vdots & & \vdots & \\
 & \uparrow & & \uparrow & & \uparrow & \\
 0 & \longrightarrow & Z^{p,1} & \longrightarrow & I^{p,1} & \longrightarrow & B^{p+1,1} \longrightarrow 0 \\
 & & \uparrow & & \uparrow & & \uparrow \\
 0 & \longrightarrow & Z^{p,0} & \longrightarrow & I^{p,0} & \longrightarrow & B^{p+1,0} \longrightarrow 0 \\
 & & \uparrow & & \uparrow & & \uparrow \\
 0 & \longrightarrow & Z^p(C^\bullet) & \longrightarrow & C^p & \longrightarrow & B^{p+1}(C^\bullet) \longrightarrow 0 \\
 & & \uparrow & & \uparrow & & \uparrow \\
 & & 0 & & 0 & & 0
 \end{array}$$

(The hint for the previous problem applies again. We remark that the bottom nonzero rows of (24.3.7.3) and (24.3.7.4) appeared in (1.6.5.3).)

24.3.G. EXERCISE/DEFINITION. Build an injective resolution (24.3.7.1) of C^\bullet such that $Z^{p,*} = Z^p(I^{\bullet,*})$, $B^{p,*} = B^p(I^{\bullet,*})$, $H^{p,*} = H^p(I^{\bullet,*})$, so the three complexes (24.3.7.2) are injective resolutions. This is called a **Cartan-Eilenberg resolution** of C^\bullet .

24.3.8. Proof of the Grothendieck spectral sequence, Theorem 24.3.5. We pick up where we left off before our digression on Cartan-Eilenberg resolutions. Choose an injective resolution I^\bullet of X . Apply the functor F , then take a Cartan-Eilenberg resolution $I^{\bullet,*}$ of FI^\bullet , and then apply G , to obtain (24.3.6.1).

Exercise 24.3.D describes what happens when we take (24.3.6.1) as E_0 in a spectral sequence with upward orientation. So we now consider the rightward orientation.

From our construction of the Cartan-Eilenberg resolution, we have injective resolutions (24.3.7.2), and short exact sequences

$$(24.3.8.1) \quad 0 \longrightarrow B^p(I^\bullet, q) \longrightarrow Z^p(I^\bullet, q) \longrightarrow H^p(I^\bullet, q) \longrightarrow 0$$

$$(24.3.8.2) \quad 0 \longrightarrow Z^p(I^\bullet, q) \longrightarrow I^{p, q} \longrightarrow B^{p+1}(I^\bullet, q) \longrightarrow 0$$

of *injective* objects (from the columns of (24.3.7.3) and (24.3.7.4)). This means that both are *split* exact sequences (the central term can be expressed as a direct sum of the outer two terms), so upon application of G , both exact sequences remain exact.

Applying the left-exact functor G to

$$0 \longrightarrow Z^p(I^\bullet, q) \longrightarrow I^{p, q} \longrightarrow I^{p+1, q},$$

we find that $GZ^p(I^\bullet, q) = \ker(GI^{p, q} \rightarrow GI^{p+1, q})$. But this kernel is the *definition* of $Z^p(GI^\bullet, q)$, so we have an induced isomorphism $GZ^p(I^\bullet, q) = Z^p(GI^\bullet, q)$ (“ G and Z^p commute”). From the exactness of (24.3.8.2) upon application of G , we see that $GB^{p+1}(I^\bullet, q) = B^{p+1}(GI^\bullet, q)$ (both are $\text{coker}(GZ^p(I^\bullet, q) \rightarrow GI^{p, q})$). From the exactness of (24.3.8.1) upon application of G , we see that $GH^p(I^\bullet, q) = H^p(GI^\bullet, q)$ (both are $\text{coker}(GB^p(I^\bullet, q) \rightarrow GZ^p(I^\bullet, q))$) — so “ G and H^p commute”.

We return to considering the rightward-oriented spectral sequence with (24.3.6.1) as E_0 . Taking cohomology in the rightward direction, we find $E_1^{p, q} = H^p(GI^\bullet, q) = GH^p(I^\bullet, q)$ (as G and H^p commute). Now $H^p(I^\bullet, q)$ is an injective resolution of $(R^p F)(X)$ (the last resolution of (24.3.7.2)). Thus when we compute E_2 by using the vertical arrows, we find $E_2^{p, q} = R^q G(R^p F(X))$.

You should now verify yourself that this (combined with Exercise 24.3.D) concludes the proof of Theorem 24.3.5. \square

24.4 Derived functor cohomology of \mathcal{O} -modules

We wish to apply the machinery of derived functors to define cohomology of quasicoherent sheaves on a scheme X . Rather than working in the category $QCoh_X$, for a number of reasons it is simpler to work in the larger category $Mod_{\mathcal{O}_X}$ (but see Unimportant Remark 24.5.7).

24.4.1. Theorem. — *Suppose (X, \mathcal{O}_X) is a ringed space. Then the category of \mathcal{O}_X -modules $Mod_{\mathcal{O}_X}$ has enough injectives.*

As a side benefit (of use to others more than us), taking $\mathcal{O}_X = \mathbb{Z}$, we see that the category of sheaves of abelian groups on a fixed topological space have enough injectives.

24.4.2. Proof. We prove Theorem 24.4.1 in a series of exercises, following Gode-ment, [GrS, p. 27-28]. Suppose \mathcal{F} is an \mathcal{O}_X -module. We will exhibit an injection $\mathcal{F} \hookrightarrow \mathcal{Q}'$ into an injective \mathcal{O}_X -module. For each $p \in X$, choose an inclusion $\mathcal{F}_p \hookrightarrow Q_p$ into an injective $\mathcal{O}_{X,p}$ -module (possible as the category of $\mathcal{O}_{X,p}$ -modules has enough injectives, Exercise 24.2.K).

24.4.A. EXERCISE. Show that the skyscraper sheaf $\mathcal{Q}_p := i_{p,*}Q_p$, with module Q_p at point $p \in X$, is an injective \mathcal{O}_X -module. (You can cleverly do this by abstract nonsense, using Exercise 24.5.B, but it is just as quick to do it by hand.)

24.4.B. EASY EXERCISE. Show that the direct product (possibly infinite) of injective objects in an abelian category is also injective.

By the previous two exercises, $\mathcal{Q}' := \prod_{p \in X} \mathcal{Q}_p$ is an injective \mathcal{O}_X -module.

24.4.C. EASY EXERCISE. By considering stalks, show that the natural map $\mathcal{F} \rightarrow \mathcal{Q}'$ is an injection.

This completes the proof of Theorem 24.4.1. \square

We can now make a number of definitions.

24.4.3. Definition. If (X, \mathcal{O}_X) is a ringed space, and \mathcal{F} is an \mathcal{O}_X -module, define $H^i(X, \mathcal{F})$ as $R^i\Gamma(X, \mathcal{F})$. If furthermore $\pi: (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ is a map of ringed spaces, we have **higher pushforwards** (or **derived pushforwards**) $R^i\pi_*: \text{Mod}_{\mathcal{O}_X} \rightarrow \text{Mod}_{\mathcal{O}_Y}$.

We have defined these notions earlier in special cases, for quasicoherent sheaves on quasicompact separated schemes (for H^i), or for quasicompact separated morphisms of schemes (for $R^i\pi_*$), in Chapter 19. We will soon (§24.5) show that these older definitions agree with Definition 24.4.3. Thus the derived functor definition applies much more generally than our Čech definition. But it is worthwhile to note that almost everything we use will come out of the Čech definition. A notable exception is the Leray spectral sequence, which we now discuss.

24.4.4. The Leray spectral sequence.

24.4.5. Theorem (Leray spectral sequence). — Suppose $\pi: (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ is a morphism of ringed spaces. For any \mathcal{O}_X -module \mathcal{F} , there is a spectral sequence with E_2 term given by $H^q(Y, R^p\pi_*\mathcal{F})$ abutting to $H^{p+q}(X, \mathcal{F})$.

This is an immediate consequence of the Grothendieck composition-of-functors spectral sequence (Theorem 24.3.5) once we prove that the pushforward of an injective \mathcal{O} -module is an acyclic \mathcal{O} -module. We do this now.

24.4.6. Definition. We make an intermediate definition that is independently important. A sheaf \mathcal{F} on a topological space is **flasque** (also sometimes called *flabby*) if all restriction maps are surjective, i.e., if $\text{res}_{U \subset V}: \mathcal{F}(V) \rightarrow \mathcal{F}(U)$ is surjective for all $U \subset V$.

24.4.D. EXERCISE. Suppose (X, \mathcal{O}_X) is a ringed space.

(a) Show that if

$$(24.4.6.1) \quad 0 \longrightarrow \mathcal{F}' \longrightarrow \mathcal{F} \longrightarrow \mathcal{F}'' \longrightarrow 0$$

is an exact sequence of \mathcal{O}_X -modules, and \mathcal{F}' is flasque, then (24.4.6.1) is exact on sections over any open set U , i.e. $0 \longrightarrow \mathcal{F}'(U) \longrightarrow \mathcal{F}(U) \longrightarrow \mathcal{F}''(U) \longrightarrow 0$ is exact.

(b) Given an exact sequence (24.4.6.1), if \mathcal{F}' is flasque, show that \mathcal{F} is flasque if and only if \mathcal{F}'' is flasque.

24.4.E. EASY EXERCISE (PUSHFORWARD OF FLASQUES ARE FLASQUE).

(a) Suppose $\pi: X \rightarrow Y$ is a continuous map of topological spaces, and \mathcal{F} is a flasque sheaf of sets on X . Show that $\pi_*\mathcal{F}$ is a flasque sheaf on Y .

(b) Suppose $\pi: (X, \mathcal{O}_X) \rightarrow (Y, \mathcal{O}_Y)$ is a morphism of ringed spaces, and \mathcal{F} is a flasque \mathcal{O}_X -module. Show that $\pi_*\mathcal{F}$ is a flasque \mathcal{O}_Y -module.

24.4.7. Extension by zero, an occasional left adjoint to the inverse image functor. In order to prove that injective sheaves are flasque (Exercise 24.4.H), we introduce the “extension by zero” functor. In addition to always being a left adjoint, π^{-1} can sometimes be a right adjoint, when π is an inclusion of an open subset. Suppose $i: U \hookrightarrow Y$ is an inclusion of an open set into Y . Define the **extension of i by zero** $i_!: \text{Mod}_{\mathcal{O}_U} \rightarrow \text{Mod}_{\mathcal{O}_Y}$ as follows. Suppose \mathcal{F} is an \mathcal{O}_U -module. For open $W \subset Y$, define $(i_!^{\text{pre}}\mathcal{F})(W) = \mathcal{F}(W)$ if $W \subset U$, and 0 otherwise (with the obvious restriction maps). This is clearly a presheaf \mathcal{O}_Y -module. Define $i_!$ as $(i_!^{\text{pre}})^{\text{sh}}$. Note that $i_!\mathcal{F}$ is an \mathcal{O}_Y -module, and that this defines a functor. (The symbol “!” is read as “shriek”, presumably because people shriek when they see it. Thus “ $i_!$ ” is read as “ i -lower-shriek”.)

24.4.F. EXERCISE. Show that $i_!^{\text{pre}}\mathcal{F}$ satisfies the identity axiom, but need not be a sheaf. (We won’t need this, but it may give some insight into why this is called “extension by zero”. Possible source for an example: continuous functions on \mathbb{R} .)

24.4.G. EXERCISE.

(a) For $q \in Y$, show that $(i_!\mathcal{F})_q = \mathcal{F}_q$ if $q \in U$, and 0 otherwise.

(b) Show that $i_!$ is an exact functor.

(c) If \mathcal{G} is an \mathcal{O}_Y -module, describe an inclusion $i_!i^{-1}\mathcal{G} \hookrightarrow \mathcal{G}$. (Interesting remark we won’t need: Let Z be the complement of U , and $j: Z \rightarrow Y$ the natural inclusion. Then there is a short exact sequence

$$0 \rightarrow i_!i^{-1}\mathcal{G} \rightarrow \mathcal{G} \rightarrow j_*j^{-1}\mathcal{G} \rightarrow 0.$$

This is best checked by describing the maps, then checking exactness at stalks.)

(d) Show that $(i_!, i^{-1})$ is an adjoint pair, so there is a natural bijection

$$\text{Hom}_{\mathcal{O}_Y}(i_!\mathcal{F}, \mathcal{G}) \leftrightarrow \text{Hom}_{\mathcal{O}_U}(\mathcal{F}, \mathcal{G}|_U)$$

for any \mathcal{O}_U -module \mathcal{F} and \mathcal{O}_Y -module \mathcal{G} . (In particular, the sections of \mathcal{G} over U can be identified with $\text{Hom}_{\mathcal{O}_Y}(i_!\mathcal{O}_U, \mathcal{G})$.)

We are ready to use the tool of the extension by zero functor.

24.4.H. EXERCISE (INJECTIVE SHEAVES ARE FLASQUE). Suppose (X, \mathcal{O}_X) is a ringed space, and \mathcal{Q} is an injective \mathcal{O}_X -module. Show that \mathcal{Q} is flasque. Hint: If $U \subset V \subset X$, then describe an injection of \mathcal{O}_X -modules $0 \rightarrow (i_U)_!\mathcal{O}_U \rightarrow (i_V)_!\mathcal{O}_V$, where $i_U: U \hookrightarrow X$ and $i_V: V \hookrightarrow X$ are the obvious open embeddings. Apply the exact contravariant functor $\text{Hom}(\cdot, \mathcal{Q})$.

24.4.I. EXERCISE (FLASQUE IMPLIES Γ -ACYCLIC). Suppose \mathcal{F} is a flasque \mathcal{O}_X -module. Show that \mathcal{F} is Γ -acyclic (that $H^i(X, \mathcal{F}) = 0$ for $i > 0$, §24.3.2) as follows. As $\text{Mod}_{\mathcal{O}_X}$ has enough injectives, choose an inclusion of \mathcal{F} into some injective \mathcal{I} , and call its cokernel \mathcal{G} :

$$0 \rightarrow \mathcal{F} \rightarrow \mathcal{I} \rightarrow \mathcal{G} \rightarrow 0.$$

Then \mathcal{I} is flasque by Exercise 24.4.H, so \mathcal{G} is flasque by Exercise 24.4.D(b). Take the long exact sequence in (derived functor) cohomology, and show that $H^1(X, \mathcal{F}) = 0$. Your argument works for *any* flasque sheaf \mathcal{F} , so $H^1(X, \mathcal{G}) = 0$ as well. Show that $H^2(X, \mathcal{F}) = 0$. Turn this into an induction.

Thus if $\pi: X \rightarrow Y$ is a morphism of ringed spaces, and \mathcal{I} is an injective \mathcal{O}_X -module, then \mathcal{I} is flasque (Exercise 24.4.H), so $\pi_*\mathcal{I}$ is flasque (Exercise 24.4.E(b)), so $\pi_*\mathcal{I}$ is acyclic for the functor Γ (Exercise 24.4.I), so this completes the proof of the Leray spectral sequence (Theorem 24.4.5). \square

24.4.J. EXERCISE. Extend the Leray spectral sequence (Theorem 24.4.5) to deal with a composition of higher pushforwards for

$$(X, \mathcal{O}_X) \xrightarrow{\pi} (Y, \mathcal{O}_Y) \xrightarrow{\rho} (Z, \mathcal{O}_Z).$$

24.4.8. ★★ The category of \mathcal{O}_X -modules on a ringed space need not have enough projectives. In contrast to Theorem 24.4.1, the category of \mathcal{O}_X -modules on a ringed space need not have enough projectives. For example, let X be \mathbb{P}_k^1 with the Zariski-topology (in fact we will need very little about X — only that it is not an *Alexandrov space*), but take \mathcal{O}_X to be the constant sheaf \mathbb{Z} . We will see that $\text{Mod}_{\mathcal{O}_X}$ — i.e., the category of sheaves of abelian groups on X — does not have enough projectives. If $\text{Mod}_{\mathcal{O}_X}$ had enough projectives, then there would be a surjection $\psi: P \rightarrow \mathbb{Z}$ from a projective sheaf. Fix a closed point $q \in X$. We will show that the map on stalks $\psi_q: P_q \rightarrow \mathbb{Z}_q$ is the zero map, contradicting the surjectivity of ψ . For each open subset U of X , denote by \mathbb{Z}_U the extension by zero from U to X of the constant sheaf with values in \mathbb{Z} (§24.4.7 — $\mathbb{Z}_U(V) = \mathbb{Z}$ if $V \subset U$, and $\mathbb{Z}_U(V) = 0$ otherwise). For each open neighborhood V of q , let W be a strictly smaller open neighborhood. Consider the surjection $\mathbb{Z}_{X-q} \oplus \mathbb{Z}_W \rightarrow \mathbb{Z}$. By projectivity of P , the surjection ψ lifts to $P \rightarrow \mathbb{Z}_{X-q} \oplus \mathbb{Z}_W$. The map $P(V) \rightarrow \mathbb{Z}(V)$ factors through $\mathbb{Z}_{X-q}(V) \oplus \mathbb{Z}_W(V) = 0$, and hence must be the zero map. Thus the map $\psi_q: P_q \rightarrow \mathbb{Z}_q$ map is zero as well (Do you see why?) as desired.

24.5 Čech cohomology and derived functor cohomology agree

We next prove that Čech cohomology and derived functor cohomology agree, where the former is defined.

24.5.1. Theorem. — Suppose X is a quasicompact separated scheme, and \mathcal{F} is a quasicoherent sheaf. Then the Čech cohomology of \mathcal{F} agrees with the derived functor cohomology of \mathcal{F} .

This statement is not as precise as it should be. We would want to know that this isomorphism is functorial in \mathcal{F} , and that it respects long exact sequences (so the connecting homomorphism defined for Čech cohomology agrees with that for derived functor cohomology). There is also an important extension to higher push-forwards. We leave these issues for the end of this section, §24.5.5

In case you are curious: so long as it is defined appropriately, Čech cohomology agrees with derived functor cohomology in a wide variety of circumstances outside of scheme theory (if the underlying topological space is paracompact), but not always (see [Gr1, §3.8] for a counterexample). Remark 24.5.6 is also related.

The central idea in the proof (albeit with a twist) is a spectral sequence argument in the same style as those of §24.3, and uses two “cohomology-vanishing” ingredients, one for each orientation of the spectral sequence.

(A) If (X, \mathcal{O}_X) is a ringed space, \mathcal{Q} is an injective \mathcal{O}_X -module, and $X = \cup_i U_i$ is a finite open cover, then \mathcal{Q} has no i th Čech cohomology with respect to this cover for $i > 0$.

(B) If X is an affine scheme, and \mathcal{F} is a quasicoherent sheaf on X , then $(R^i \Gamma)(\mathcal{F}) = 0$ for $i > 0$.

Translation: (A) says that building blocks of derived functor cohomology have no Čech cohomology, and (B) says that building blocks of Čech cohomology have no derived functor cohomology.

24.5.A. PRELIMINARY EXERCISE. Suppose (X, \mathcal{O}_X) is a ringed space, \mathcal{Q} is an injective \mathcal{O} -module, and $i: U \hookrightarrow X$ is an open subset. Show that $\mathcal{Q}|_U$ is injective on U . Hint: use the fact that i^{-1} has an *exact* left adjoint $i_!$ (extension by zero), see §24.4.7, and the following diagrams.

$$\begin{array}{ccc}
 0 & \longrightarrow & \mathcal{A} \longrightarrow \mathcal{B} \\
 & & \downarrow \quad \swarrow ? \\
 & & \mathcal{Q}|_U
 \end{array}
 \qquad
 \begin{array}{ccc}
 0 & \longrightarrow & i_! \mathcal{A} \longrightarrow i_! \mathcal{B} \\
 & & \downarrow \quad \swarrow ? \\
 & & \mathcal{Q}
 \end{array}$$

In the course of Exercise 24.5.A, you will have proved the following fact, which we shall use again in Exercise 31.3.C. (You can also use it to solve Exercise 24.4.A.)

24.5.B. EXERCISE (RIGHT ADJOINTS TO EXACT FUNCTORS PRESERVE INJECTIVES). Show that if (F, G) is an adjoint pair of additive functors between abelian categories, and F is exact, then G sends injective elements to injective elements.

24.5.2. Proof of Theorem 24.5.1, assuming (A) and (B). As with the facts proved in §24.3, we take the only approach that is reasonable: we choose an injective resolution $0 \rightarrow \mathcal{F} \rightarrow \mathcal{Q}_\bullet$ of \mathcal{F} and a Čech cover of X , mix these two types of information in a double complex, and toss it into our spectral sequence machine (§1.7). More precisely, choose a finite affine open cover $X = \cup_i U_i$ and an injective resolution

$$0 \rightarrow \mathcal{F} \rightarrow \mathcal{Q}_0 \rightarrow \mathcal{Q}_1 \rightarrow \cdots$$

Consider the double complex

$$\begin{array}{ccccccc}
 (24.5.2.1) & \vdots & & \vdots & & \vdots & & \vdots \\
 & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 & 0 & \longrightarrow & \oplus_i \mathcal{Q}_2(\mathcal{U}_i) & \longrightarrow & \oplus_{i,j} \mathcal{Q}_2(\mathcal{U}_{ij}) & \longrightarrow & \oplus_{i,j,k} \mathcal{Q}_2(\mathcal{U}_{ijk}) \longrightarrow \cdots \\
 & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 & 0 & \longrightarrow & \oplus_i \mathcal{Q}_1(\mathcal{U}_i) & \longrightarrow & \oplus_{i,j} \mathcal{Q}_1(\mathcal{U}_{ij}) & \longrightarrow & \oplus_{i,j,k} \mathcal{Q}_1(\mathcal{U}_{ijk}) \longrightarrow \cdots \\
 & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 & 0 & \longrightarrow & \oplus_i \mathcal{Q}_0(\mathcal{U}_i) & \longrightarrow & \oplus_{i,j} \mathcal{Q}_0(\mathcal{U}_{ij}) & \longrightarrow & \oplus_{i,j,k} \mathcal{Q}_0(\mathcal{U}_{ijk}) \longrightarrow \cdots \\
 & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 & 0 & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & 0 \longrightarrow \cdots
 \end{array}$$

We take this as the E_0 term in a spectral sequence. First, we use the rightward filtration. As higher Čech cohomology of injective \mathcal{O} -modules is 0 (assumption **(A)**), we get 0's everywhere except in "column 0", where we get $\mathcal{Q}_i(X)$ in row i :

$$\begin{array}{ccccccc}
 & \vdots & & \vdots & & \vdots & & \vdots \\
 & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 & 0 & & \mathcal{Q}_2(X) & & 0 & & 0 & \cdots \\
 & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 & 0 & & \mathcal{Q}_1(X) & & 0 & & 0 & \cdots \\
 & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 & 0 & & \mathcal{Q}_0(X) & & 0 & & 0 & \cdots \\
 & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 & 0 & & 0 & & 0 & & 0
 \end{array}$$

Then we take cohomology in the vertical direction, and we get derived functor cohomology of \mathcal{F} on X on the E_2 page:

$$\begin{array}{ccccccc}
 & & \vdots & & \vdots & & \vdots & & \vdots \\
 & & & & & & & & \\
 & & 0 & & R^2\Gamma(X, \mathcal{F}) & & 0 & & 0 & \cdots \\
 & \nearrow & & \nearrow & & \nearrow & & \nearrow & & \\
 & 0 & & R^1\Gamma(X, \mathcal{F}) & & 0 & & 0 & & \cdots \\
 & \nearrow & & & \nearrow & & \nearrow & & \nearrow & \\
 & 0 & & \Gamma(X, \mathcal{F}) & & 0 & & 0 & & \cdots \\
 & \nearrow & & & \nearrow & & \nearrow & & \nearrow & \\
 & & 0 & & 0 & & 0 & & &
 \end{array}$$

We then start over on the E_0 page, and this time use the filtration corresponding to choosing the upward arrow first. By Proposition 24.5.A, $\mathcal{Q}_i|_{U_j}$ is injective on U_j , so we are computing the derived functor cohomology of \mathcal{F} on U_j . Then the higher derived functor cohomology of \mathcal{F} on U_j is 0 (assumption (B)), so all entries are 0 except possibly on row 0. Thus the E_1 term is:

(24.5.2.2)

$$0 \longrightarrow 0 \longrightarrow 0 \longrightarrow 0 \longrightarrow \cdots$$

$$0 \longrightarrow 0 \longrightarrow 0 \longrightarrow 0 \longrightarrow \cdots$$

$$0 \longrightarrow \oplus_i \Gamma(U_i, \mathcal{F}) \longrightarrow \oplus_{i,j} \Gamma(U_{ij}, \mathcal{F}) \longrightarrow \oplus_{i,j,k} \Gamma(U_{ijk}, \mathcal{F}) \longrightarrow \cdots$$

$$0 \longrightarrow 0 \longrightarrow 0 \longrightarrow 0 \longrightarrow \cdots$$

Row 0 is precisely the Čech complex of \mathcal{F} , so the spectral sequence converges at the E_2 term, yielding the Čech cohomology. Since one orientation yields derived functor cohomology and one yields Čech cohomology, we are done. \square

So it remains to show (A) and (B).

24.5.3. Ingredient (A): injectives have no Čech cohomology.

24.5.C. EXERCISE. Suppose $X = \cup_j U_j$ is a finite cover of X by open sets, and \mathcal{F} is a flasque sheaf (Definition 24.4.6) on X . Show that the Čech complex for \mathcal{F} with respect to $\cup_{j=1}^n U_j$ has no cohomology in positive degree, i.e., that it is exact except in degree 0 (where it has cohomology $\mathcal{F}(X)$, by the sheaf axioms). Hint: use induction on n . Consider the short exact sequence of complexes (19.2.4.2) (see

also (19.2.3.1)). The corresponding long exact sequence will immediately give the desired result for $i > 1$, and flasqueness will be used for $i = 1$.

Thus flasque sheaves have no Čech cohomology, so injective \mathcal{O} -modules in particular (Exercise 24.4.H) have none. This is all we need for our algebro-geometric applications, but to show you how general this machinery is, we give an entertaining application.

24.5.D. UNIMPORTANT EXERCISE (PERVERSE PROOF OF INCLUSION-EXCLUSION THROUGH COHOMOLOGY OF SHEAVES). The inclusion-exclusion principle is (equivalent to) the following: suppose that X is a finite set, and U_i ($1 \leq i \leq n$) are finite sets covering X . As usual, define $U_I = \cap_{i \in I} U_i$ for $I \subset \{1, \dots, n\}$. Then

$$|X| = \sum |U_i| - \sum_{|I|=2} |U_I| + \sum_{|I|=3} |U_I| - \sum_{|I|=4} |U_I| + \cdots.$$

Prove this by endowing X with the discrete topology, showing that the constant sheaf \mathbb{Q} is flasque, considering the Čech complex computing $H^i(X, \mathbb{Q})$ using the cover $\{U_i\}$, and using Exercise 1.6.B.

24.5.4. ★ Ingredient (B): quasicoherent sheaves on affine schemes have no derived functor cohomology.

We show the following statement by induction on k . Suppose X is an affine scheme, and \mathcal{F} is a quasicoherent sheaf on X . Then $R^i \Gamma(X, \mathcal{F}) = 0$ for $0 < i \leq k$. The result is vacuously true for $k = 0$; so suppose we know the result for all $0 < k' < k$ (with X replaced by *any* affine scheme). Suppose $\alpha \in R^k \Gamma(X, \mathcal{F})$. We wish to show that $\alpha = 0$. Choose an injective resolution by \mathcal{O}_X -modules

$$\begin{array}{c} \vdots \\ \uparrow d_1 \\ \mathcal{Q}_1 \\ \uparrow d_0 \\ \mathcal{Q}_0 \\ \uparrow \\ \mathcal{F} \\ \uparrow \\ 0 \end{array}$$

Then α has a representative α' in $\mathcal{Q}_k(X) = \Gamma(X, \mathcal{Q}_k)$, such that $d\alpha' = 0$. Because the injective resolution is exact, α' is locally a boundary. In other words, in the open neighborhood of any point $p \in X$, there is an open set V_p such that $\alpha|_{V_p} = d\alpha'$ for some $\alpha' \in \mathcal{Q}_{k-1}(V_p)$. (Be sure you see why this is true! Recall that taking cokernel of a map of sheaves requires sheafification, see Proposition 2.6.1.) By shrinking V_p if necessary, we can assume V_p is affine. By the quasicoompactness of X , we can choose a finite number of the V_p 's that cover X . Rename these as U_i , so we have an affine cover X . Consider the Čech cover of X with respect to *this*

affine cover (*not* the affine cover you might have thought we would use — that of X by itself — but instead an affine cover tailored for our particular α). Consider the double complex (24.5.2.1), as the E_0 term in a spectral sequence.

First consider the rightward orientation. As in the argument in §24.5.2, the spectral sequence converges at E_2 , where we get 0 everywhere, except that the derived functor cohomology appears in the 0th column.

Next, start over again, choosing the upward filtration. On the E_1 page, row 0 is the Čech complex, as in (24.5.2.2). All the rows between 1 and $k-1$ are 0 by our inductive hypothesis, but we don't yet know anything about the higher rows. Because we are interested in the k th derived functor, we focus on the k th antidiagonal ($E_{\bullet}^{p,k-p}$). The only possibly nonzero terms in this antidiagonal are $E_1^{k,0}$ and $E_1^{0,k}$. We look first at the term on the bottom row $E_1^{k,0} = \prod_{|I|=k} \Gamma(U_I, \mathcal{F})$, which is part of the Čech complex:

$$\cdots \longrightarrow \prod_{|I|=k-1} \Gamma(U_I, \mathcal{F}) \longrightarrow \prod_{|I|=k} \Gamma(U_I, \mathcal{F}) \longrightarrow \prod_{|I|=k+1} \Gamma(U_I, \mathcal{F}) \longrightarrow \cdots$$

But we have already verified that the Čech cohomology of a quasicoherent sheaf on an affine scheme vanishes — this is the one spot where we use the quasicoherence of \mathcal{F} . Thus this term vanishes by the E_2 page (i.e., $E_i^{k,0} = 0$ for $i \geq 2$).

So the only term of interest in the k th antidiagonal of E_1 is $E_1^{0,k}$, which is the homology of

$$(24.5.4.1) \quad \prod_i \mathcal{Q}_{k-1}(U_i) \longrightarrow \prod_i \mathcal{Q}_k(U_i) \longrightarrow \prod_i \mathcal{Q}_{k+1}(U_i),$$

which is $\prod_i R^k \Gamma(U_i, \mathcal{F})$ (using Preliminary Exercise 24.5.A which stated that the $\mathcal{Q}_j|_{U_i}$ are injective on U_i , so they can be used to compute $R^k(\Gamma(U_i, \mathcal{F}))$). So $E_2^{0,k}$ is the homology of

$$0 \longrightarrow \prod_i R^k \Gamma(U_i, \mathcal{F}) \longrightarrow \prod_{i,j} R^k \Gamma(U_{ij}, \mathcal{F})$$

and thereafter all differentials to and from the $E_{\bullet}^{0,k}$ terms will be 0, as the sources and targets of those arrows will be 0. Consider now our lift of α' of our original class $\alpha \in R^k \Gamma(X, \mathcal{F})$ to $\prod_i R^k \Gamma(U_i, \mathcal{F})$. Its image in the homology of (24.5.4.1) is zero — this was how we chose our cover U_i to begin with! Thus $\alpha = 0$ as desired, completing our proof. \square

24.5.E. ★★ EXERCISE. The proof is not quite complete. We have a class $\alpha \in R^k \Gamma(X, \mathcal{F})$, and we have interpreted $R^k \Gamma(X, \mathcal{F})$ as

$$\ker \left(\prod_i R^k \Gamma(U_i, \mathcal{F}) \longrightarrow \prod_{i,j} R^k \Gamma(U_{ij}, \mathcal{F}) \right).$$

We have two maps $R^k \Gamma(X, \mathcal{F}) \rightarrow R^k \Gamma(U_i, \mathcal{F})$, one coming from the natural restriction (under which we can see that the image of α is zero), and one coming from the actual spectral sequence machinery. Verify that they are the same map. (Possible hint: with the filtration used, the $E_{\infty}^{0,k}$ term is indeed the quotient of the homology of the double complex, so the map goes the right way.)

24.5.5. ★ Tying up loose ends.

24.5.F. IMPORTANT EXERCISE. State and prove the generalization of Theorem 24.5.1 to higher pushforwards $R^i\pi_*$, where $\pi: X \rightarrow Y$ is a quasicompact separated morphism of schemes.

24.5.G. EXERCISE. Show that the isomorphism of Theorem 24.5.1 is functorial in \mathcal{F} , i.e., given a morphism $\mathcal{F} \rightarrow \mathcal{G}$, the diagram

$$\begin{array}{ccc} H^i(X, \mathcal{F}) & \longleftrightarrow & R^i\Gamma(X, \mathcal{F}) \\ \downarrow & & \downarrow \\ H^i(X, \mathcal{G}) & \longleftrightarrow & R^i\Gamma(X, \mathcal{G}) \end{array}$$

commutes, where the horizontal arrows are the isomorphisms of Theorem 24.5.1, and the vertical arrows come from functoriality of H^i and $R^i\Gamma$. (Hint: “spectral sequences are functorial in E_0 ”, which can be easily seen from the construction, although we haven’t said it explicitly. See §1.7.7, which remarks on the functoriality of spectral sequences.)

24.5.H. EXERCISE. Show that the isomorphisms of Theorem 24.5.1 induce isomorphisms of long exact sequences.

24.5.6. Remark. If you wish, you can use the above argument to prove the following theorem of Leray. Suppose we have a sheaf of abelian groups \mathcal{F} on a topological space X , and some covering $\{U_i\}$ of X such that the (derived functor) cohomology of \mathcal{F} in positive degree vanishes on every finite intersection of the U_i . Then the cohomology of \mathcal{F} can be calculated by the Čech cohomology of the cover $\{U_i\}$; there is no need to pass to the colimit of all covers, as is the case for Čech cohomology in general.

24.5.7. Unimportant Remark: Working in $QCoh_X$ rather than $Mod_{\mathcal{O}_X}$. In our definition of derived functors of quasicoherent sheaves on X , we could have tried to work in the category of quasicoherent sheaves $QCoh_X$ itself, rather than in the larger category $Mod_{\mathcal{O}_X}$. There are several reasons why this would require more effort. It is not hard to show that $QCoh_X$ has enough injectives if X is Noetherian (see for example [Ha1, Exer. III.3.6(a)]). Because we don’t have “extension by zero” (§24.4.7) in $QCoh$, the proofs that injective quasicoherent sheaves on an open set U restrict to injective quasicoherent sheaves on smaller open subsets V (the analog of Exercise 24.5.A) and that injective quasicoherent sheaves are flasque (the analog of Exercise 24.4.H) are harder. You can use this to show that H^i and $R^i\pi_*$ computed in $QCoh$ are the same as those computed in $Mod_{\mathcal{O}}$ (once you make the statements precise). It is true that injective elements of $QCoh_X$ (X Noetherian) are injective in $Mod_{\mathcal{O}_X}$, but this requires work (see [Mur, Prop. 68]).

It is true that $QCoh_X$ has enough injectives for *any* scheme X , but this is much harder, see [EE]. And as is clear from the previous paragraph, “enough injectives” is only the beginning of what we want.

24.5.8. Unimportant Remark. Theorem 24.5.1 implies that if $\pi: X \rightarrow Y$ is quasicompact and separated, then $R^i\pi_*$ sends $QCoh_X$ to $QCoh_Y$ (by showing an isomorphism with Čech cohomology). If X and Y are Noetherian, the hypothesis “separated” can be relaxed to “quasiseparated”, using the ideas of Unimportant Remark 24.5.7. But

because morphisms from a Noetherian scheme are automatically quasiseparated (Easy Exercise 8.3.B(b)), we can then remove the redundant “quasiseparated” assumption. This sounds exciting, but it is not nearly as useful as the separated case, because without Čech cohomology, it is hard to compute anything.

CHAPTER 25

Flatness

The concept of flatness is a riddle that comes out of algebra, but which technically is the answer to many prayers.

— D. Mumford [Mu7, III.10]

It is a riddle, wrapped in a mystery, inside an enigma; but perhaps there is a key.

— W. Churchill [Ch]

25.1 Introduction

We come next to the important concept of flatness (first introduced in §17.3.8). We could have discussed flatness at length as soon as we had discussed quasicoherent sheaves and morphisms. But it is an unexpected idea, and the algebra and geometry are not obviously connected, so we have left it for relatively late. “Exact” seems more descriptive terminology than “flat” (so we would have “exact modules” and “exact morphisms”), but this has not caught on.

Serre has stated that he introduced flatness purely for reasons of algebra in his landmark “GAGA” paper [Se3], and that it was Grothendieck who recognized its geometric significance.

A flat morphism $\pi: X \rightarrow Y$ is the right notion of a “nice”, or “nicely varying” family over Y . For example, if π is a projective flat family over a connected Noetherian base (translation: $\pi: X \rightarrow Y$ is a projective flat morphism, with Y connected and Noetherian), we will see that various numerical invariants of fibers are constant, including the dimension (§25.5.6), and numbers interpretable in terms of an Euler characteristic (see §25.7):

- (a) the Hilbert polynomial (Corollary 25.7.2),
- (b) the degree (in projective space) (Exercise 25.7.B(a)),
- (b') (as a special case of (b)) if π is finite, the degree of π (recovering and extending the fact that the degree of a projective map between regular curves is constant, §18.4.4, see Exercise 25.4.G and §25.4.12),
- (c) the arithmetic genus (Exercise 25.7.B(b)),
- (d) the degree of a line bundle if the fiber is a curve (Corollary 25.7.3), and
- (e) intersections of divisors and line bundles (Exercise 25.7.D).

One might think that the right hypothesis might be smoothness, or more generally some sort of equisingularity, but we only need something weaker. And this is a good thing: branched covers are not fibrations in any traditional sense, yet they still behave well — the double cover $\mathbb{A}^1 \rightarrow \mathbb{A}^1$ given by $y \mapsto x^2$ has constant degree 2 (§10.3.4, revisited in §18.4.9). Another key example is that of a family of smooth curves degenerating to a nodal curve (Figure 25.1) — the topology of the

(underlying analytic) curve changes, but the arithmetic genus remains constant. One can prove things about regular curves by first proving them about a nodal degeneration, and then showing that the result behaves well in flat families. Degeneration techniques such as this are ubiquitous in algebraic geometry.




FIGURE 25.1. A flat family of smooth curves degenerating to a nodal curve: $y^2 = x^3 + x^2 + tx$.

Given the cohomological nature of the constancy of Euler characteristic result, you should not be surprised that the hypothesis needed (flatness) is cohomological in nature — it can be characterized by vanishing of Tor (Exercise 24.1.D), which we use to great effect in §25.3.

But flatness is important for other reasons too. As a start: as this is the right notion of a “nice family”, it allows us to correctly define the notion of moduli space. For example, the *Hilbert scheme* of \mathbb{P}^n “parametrizes closed subschemes of \mathbb{P}^n ”. Maps from a scheme B to the Hilbert scheme correspond to (finitely presented) closed subschemes of \mathbb{P}^n_B flat over B . By universal property nonsense, this defines the Hilbert scheme up to unique isomorphism (although we of course must show that it exists, which takes some effort — see Theorem 28.3.1 for more). The moduli space of smooth projective curves is defined by the universal property that maps to the moduli space correspond to projective flat (finitely presented) families whose geometric fibers are smooth curves. (Sadly, this moduli space does not exist...) On a related note, flatness is central in deformation theory: it is key to understanding how schemes (and other geometric objects, such as vector bundles) can deform (cf. §22.5.12). Finally, the notion of Galois descent generalizes to (faithfully) “flat descent”, which allows us to “glue” in more exotic Grothendieck topologies in the same way we do in the Zariski topology (or more classical topologies); but this is beyond the scope of our current discussion.

Happy families are all alike; every unhappy family is unhappy in its own way. (See, for example, Figure 25.2, and Exercises 25.2.G(d), 25.2.K, and 25.4.H.)

— L. Tolstoy, *Anna Karenina*, [To, I. 1]

25.1.1. Structure of the chapter.

Flatness has many aspects of different flavors, and it is easy to lose sight of the forest for the trees. Because the algebra of flatness seems so unrelated to the geometry, it can be nonintuitive. We will necessarily begin with algebraic foundations, but you should focus on the following points: methods of showing things are flat (both general criteria and explicit examples), and classification of flat modules over particular kinds of rings. You should try every exercise dealing with explicit examples such as these.

Here is an outline of the chapter, to help focus your attention.

- In §25.2, we discuss some of the easier facts, which are algebraic in nature.
- §25.3, §25.4, and §25.6 give ideal-theoretic criteria for flatness. §25.3 and §25.4 should be read together. The first uses Tor to understand flatness, and the second uses these insights to develop ideal-theoretic criteria for flatness. §25.6, on local criteria for flatness, is harder.
- §25.5 is relatively free-standing, and could be read immediately after §25.2. It deals with topological aspects of flatness, such as the fact that flat morphisms are open in good situations.
- In §25.7, we explain the fact that “the Euler characteristic of quasicoherent sheaves is constant in flat families” (with appropriate hypotheses), and its many happy consequences. This section is surprisingly easy given its utility.

You should focus on what flatness implies and how to “picture” it, but also on explicit criteria for flatness in different situations, such as for integral domains (Observation 25.2.2), principal ideal domains (Exercise 25.4.B), discrete valuation rings (Exercise 25.4.C), the dual numbers (Exercise 25.4.D), and local rings (Theorem 25.4.5).

25.2 Easier facts

Many facts about flatness are easy or immediate, although a number are tricky. As always, I will try to make clear which is which, to help you remember the easy facts and the key ideas of proofs of the harder facts. We will pick the low-hanging fruit first.

We recall the definition of a *flat A-module* (§1.6.11). If $M \in \text{Mod}_A$, $M \otimes_A \cdot$ is always right-exact (Exercise 1.3.H). We say that M is a **flat A-module** (or *flat over A* or *A-flat*) if $M \otimes_A \cdot$ is an exact functor. We say that a *ring morphism* $B \rightarrow A$ is **flat** if A is flat as a B -module. (In particular, the algebra structure of A is irrelevant.)

25.2.1. Two key examples.

(i) Free A -modules (even of infinite rank) are clearly flat. More generally, projective modules are flat (Exercise 24.2.B).

(ii) Localizations are flat: Suppose S is a multiplicative subset of B . Then $B \rightarrow S^{-1}B$ is a flat ring morphism because localization is exact (Exercise 1.6.F(a)).

25.2.A. EASY EXERCISE: FIRST EXAMPLES.

(a) (*trick question*) Classify flat modules over a field k .

(b) Show that $A[x_1, \dots, x_n]$ is a flat A -module.

(c) Show that the ring morphism $\mathbb{Q}[x] \rightarrow \mathbb{Q}[y]$, with $x \mapsto y^2$, is flat. (This will help us understand Example 10.3.4 better, see §25.4.12.)

We make some quick but important observations.

25.2.2. Important Observation. If x is a non-zerodivisor of A , and M is a flat A -module, then

$$M \xrightarrow{\times x} M$$

is injective. (Reason: apply the exact functor $M \otimes_A \cdot$ to the exact sequence $0 \longrightarrow A \xrightarrow{\times x} A \rightarrow 0$.) In particular, *flat modules are torsion-free*. (Torsion-freeness was defined in §6.1.4.) This observation gives an easy way of recognizing when a module is *not* flat. We will use it many times.

25.2.B. EXERCISE: ANOTHER EXAMPLE. Show that a finitely generated module over a principal ideal domain is flat if and only if it is torsion-free if and only if it is free. Hint: Remark 13.5.14 classifies finitely generated modules over a principal ideal domain. (Exercise 25.4.B sheds more light on flatness over a principal ideal domain. Proposition 14.3.2 is also related.)

25.2.C. EXERCISE (FLATNESS IS PRESERVED BY CHANGE OF BASE RING). Show that if M is a flat B -module, $B \rightarrow A$ is a homomorphism, then $M \otimes_B A$ is a flat A -module.

25.2.D. EXERCISE (FLATNESS IS PRESERVED BY COMPOSITION). Show that if A is a flat B -algebra, and M is A -flat, then M is also B -flat.

25.2.3. Proposition (flatness is a stalk/prime-local property). — *An A -module M is flat if and only if M_p is a flat A_p -module for all prime ideals p .*

Proof. Suppose first that M is a flat A -module. Given any exact sequence of A_p -modules

$$(25.2.3.1) \quad 0 \longrightarrow N' \longrightarrow N \longrightarrow N'' \longrightarrow 0,$$

$$0 \longrightarrow M \otimes_A N' \longrightarrow M \otimes_A N \longrightarrow M \otimes_A N'' \longrightarrow 0$$

is exact too. But $M \otimes_A N$ is canonically isomorphic to $M_p \otimes_{A_p} N$ (Do you see why?), so M_p is a flat A_p -module.

Suppose next that M_p is a flat A_p -module for all p . Given any short exact sequence (25.2.3.1), tensoring with M yields

$$(25.2.3.2) \quad 0 \longrightarrow K \longrightarrow M \otimes_A N' \longrightarrow M \otimes_A N \longrightarrow M \otimes_A N'' \longrightarrow 0$$

(using right-exactness of \otimes , Exercise 1.3.H) where K is the kernel of $M \otimes_A N' \rightarrow M \otimes_A N$. We wish to show that $K = 0$. It suffices to show that $K_p = 0$ for every prime $p \subset A$ (see the comment after Exercise ??). Given any p , localizing (25.2.3.1) at p and tensoring with the exact A_p -module M_p yields

$$(25.2.3.3) \quad 0 \longrightarrow M_p \otimes_{A_p} N'_p \longrightarrow M_p \otimes_{A_p} N_p \longrightarrow M_p \otimes_{A_p} N''_p \longrightarrow 0.$$

But localizing (25.2.3.2) at p and using the isomorphisms $M_p \otimes_{A_p} N_p \xrightarrow{\sim} (M \otimes_A N)_p$, we obtain the exact sequence

$$0 \longrightarrow K_p \longrightarrow M_p \otimes_{A_p} N'_p \longrightarrow M_p \otimes_{A_p} N_p \longrightarrow M_p \otimes_{A_p} N''_p \longrightarrow 0,$$

which is the same as the exact sequence (25.2.3.3) except for the K_p . Hence $K_p = 0$ as desired. \square

25.2.4. Flatness for schemes.

Motivated by Proposition 25.2.3, the extension of the notion of flatness to schemes is straightforward.

25.2.5. Definition: flat quasicoherent sheaves. We say that a quasicoherent sheaf \mathcal{F} on a scheme X is **flat at** $p \in X$ if \mathcal{F}_p is a flat $\mathcal{O}_{X,p}$ -module. We say that a quasicoherent sheaf \mathcal{F} on a scheme X is **flat (over X)** if it is flat at all $p \in X$. In light of Proposition 25.2.3, we can check this notion on an affine open cover of X .

25.2.6. Definition: flat morphism. Similarly, we say that a morphism of schemes $\pi: X \rightarrow Y$ is **flat at** $p \in X$ if $\mathcal{O}_{X,p}$ is a flat $\mathcal{O}_{Y,\pi(p)}$ -module. We say that a morphism of schemes $\pi: X \rightarrow Y$ is **flat** if it is flat at all $p \in X$. We can check flatness locally on the source and target.

We can combine these two definitions into a single fancy definition.

25.2.7. Definition: flat quasicoherent sheaf over a base. Suppose $\pi: X \rightarrow Y$ is a morphism of schemes, and \mathcal{F} is a quasicoherent sheaf on X . We say that \mathcal{F} is **flat (over Y) at** $p \in X$ if \mathcal{F}_p is a flat $\mathcal{O}_{Y,\pi(p)}$ -module. We say that \mathcal{F} is **flat (over Y)** if it is flat at all $p \in X$.

Definitions 25.2.5 and 25.2.6 correspond to the cases $X = Y$ and $\mathcal{F} = \mathcal{O}_X$ respectively. (Definition 25.2.7 applies without change to the category of ringed spaces, but we won't use this.)

25.2.E. EASY EXERCISE (REALITY CHECK). Show that open embeddings are flat.

Our results about flatness over rings above carry over easily to schemes.

25.2.F. EXERCISE. Show that a map of rings $B \rightarrow A$ is flat if and only if the corresponding morphism of schemes $\text{Spec } A \rightarrow \text{Spec } B$ is flat. More generally, if $B \rightarrow A$ is a map of rings, and M is an A -module, show that M is B -flat if and only if \widetilde{M} is flat over $\text{Spec } B$. (Be careful: this requires more than merely invoking Proposition 25.2.3.)

Thus if $\pi: X \rightarrow Y$ is an affine morphism, and \mathcal{F} is a quasicoherent sheaf on X , then \mathcal{F} is flat over Y if and only if $\pi_*\mathcal{F}$ is flat over Y .

25.2.G. EASY EXERCISE (EXAMPLES AND REALITY CHECKS).

(a) If X is a scheme, and $p \in X$, show that the natural morphism $\text{Spec } \mathcal{O}_{X,p} \rightarrow X$ is flat. (Hint: localization is flat, §25.2.1.)

(b) Show that $\mathbb{A}_A^n \rightarrow \text{Spec } A$ is flat.

(c) If \mathcal{F} is a locally free sheaf on a scheme X , show that $\mathbb{P}\mathcal{F} \rightarrow X$ (Definition 18.2.3) is flat.

(d) Show that $\text{Spec } k \rightarrow \text{Spec } k[t]/(t^2)$ is not flat. (Draw a picture to try to see what is not “nice” about this morphism. Some more insight about flatness of the dual numbers will be given in the criterion of Exercise 25.4.D.)

25.2.H. EXERCISE (FLATNESS IS PRESERVED BY COMPOSITION). Suppose $\pi: X \rightarrow Y$ and \mathcal{F} is a quasicoherent sheaf on X , flat over Y . Suppose also that $\psi: Y \rightarrow Z$ is a flat morphism. Show that \mathcal{F} is flat over Z .

25.2.I. EXERCISE (FLATNESS IS PRESERVED BY BASE CHANGE). Suppose $\pi: X \rightarrow Y$ is a morphism, and \mathcal{F} is a quasicoherent sheaf on X , flat over Y . If $\rho: Y' \rightarrow Y$ is any morphism, and $\rho': X \times_Y Y' \rightarrow X$ is the induced morphism, show that $(\rho')^*\mathcal{F}$

is flat over Y' .

$$\begin{array}{ccc}
 (\rho')^* \mathcal{F} & & \mathcal{F} \\
 \downarrow & & \downarrow \\
 X \times_Y Y' & \xrightarrow{\rho'} & X \\
 \downarrow & & \downarrow \pi \\
 Y' & \xrightarrow{\rho} & Y
 \end{array}$$

In particular, using Exercise 25.2.A(a), if X and Y' are k -schemes, and \mathcal{F} is *any* quasicoherent sheaf on X , then $(\rho')^* \mathcal{F}$ is flat over Y' . For example, $X \times_k Y'$ is *always* flat over Y' — “products over a field are flat over their factors”. (Feel free to immediately generalize this further; for example, \mathcal{F} can be a quasicoherent sheaf on a scheme Z over X , flat over Y .)

25.2.8. Flat morphisms behave reasonably. Since flatness is clearly local on the target from the definition, the previous two exercises show that the class of flat morphisms is “reasonable” in the sense of §8.1.

The following exercise is very useful for visualizing flatness and non-flatness (see for example Figure 25.2).

25.2.J. EXERCISE (FLAT MAPS SEND ASSOCIATED POINTS TO ASSOCIATED POINTS). Suppose $\pi: X \rightarrow Y$ is a flat morphism of locally Noetherian schemes. Show that any associated point of X must map to an associated point of Y . (Feel free to immediately generalize this to a coherent sheaf \mathcal{F} on X , flat over Y , without π itself needing to be flat.) Hint: suppose $\pi^\sharp: (B, \mathfrak{n}) \rightarrow (A, \mathfrak{m})$ is a local morphism of Noetherian local rings (i.e., $\pi^\sharp(\mathfrak{n}) \subset \mathfrak{m}$, §7.3.1). Suppose \mathfrak{n} is not an associated prime of B . Show that there is an element $f \in \mathfrak{n}$ not in any associated prime of B (perhaps using prime avoidance, Proposition 12.2.13), and hence is a non-zerodivisor. Show that $\pi^\sharp f \in \mathfrak{m}$ is a non-zerodivisor of A using Observation 25.2.2, and thus show that \mathfrak{m} is not an associated prime of A .

25.2.K. EXERCISE. Use Exercise 25.2.J to show that the following morphisms are not flat (see Figure 25.2):

- (a) $\text{Spec } k[x, y]/(xy) \rightarrow \text{Spec } k[x]$,
- (b) $\text{Spec } k[x, y]/(y^2, xy) \rightarrow \text{Spec } k[x]$,
- (c) $\text{Bl}_{(0,0)} \mathbb{A}_k^2 \rightarrow \mathbb{A}_k^2$.

Hint for (c): first pull back to a line through the origin to obtain something akin to (a). (This foreshadows the statement and proof of Proposition 25.5.7, which says that for flat morphisms “there is no jumping of fiber dimension”.)

25.2.L. EXERCISE. Show that $\pi: X \rightarrow Y$ is flat if and only if $\pi^*: \text{QCoh}_Y \rightarrow \text{QCoh}_X$ is an exact functor. (Translation: the earlier definition of flat morphism in Remark 17.3.8 agrees with Definition 25.2.6. But Definition 25.2.6 is “better”!)




FIGURE 25.2. Morphisms that are not flat (Exercise 25.2.K)

25.2.9. Theorem (cohomology commutes with flat base change). — Suppose

$$\begin{array}{ccc} X' & \xrightarrow{\rho'} & X \\ \pi' \downarrow & & \downarrow \pi \\ Y' & \xrightarrow{\rho} & Y \end{array}$$

is a Cartesian diagram, and π (and thus π') is quasicompact and separated (so higher pushforwards of quasicoherent sheaves by π and π' exist, as described in §19.8). Suppose also that ρ is flat, and \mathcal{F} is a quasicoherent sheaf on X . Then the natural morphisms (Exercise 19.8.B(a)) $\rho^*(R^i\pi_*\mathcal{F}) \rightarrow R^i\pi'_*(\rho'^*\mathcal{F})$ are isomorphisms.

25.2.M. EXERCISE. Prove Theorem 25.2.9. Hint: Exercise 19.8.B(b) is the special case where ρ is affine. Extend it to the quasicompact and separated case using the same idea as the solution to Exercise 17.2.C (which used Exercise 6.2.G). Your proof of the case $i = 0$ will only need a quasiseparated hypothesis in place of the separated hypothesis.

A useful special case is where Y' is the generic point of a reduced component of Y . In other words, in light of Exercise 25.2.G(a), the stalk of the higher pushforward of \mathcal{F} at the generic point is the cohomology of \mathcal{F} on the fiber over the generic point. This is a first example of something important: understanding cohomology of (quasicoherent sheaves on) fibers in terms of higher pushforwards. (We would certainly hope that higher pushforwards would tell us something about higher cohomology of fibers, but this is certainly not a priori clear!) In comparison to this result, which shows that cohomology of *any* quasicoherent sheaf commutes with *flat* base change, §25.7 and Chapter 28 deal with when and how cohomology of a *flat* quasicoherent sheaf commutes with *any* base change.

25.2.10. Pulling back closed subschemes (and ideal sheaves) by flat morphisms.

Closed subschemes pull back particularly well under flat morphisms, and this can be helpful to keep in mind. As pointed out in Remarks 17.3.8 and 17.3.9, in the case of flat morphisms, pullback of ideal sheaves as *quasicoherent sheaves* agrees with pullback in terms of the pullback of the corresponding closed subschemes.

In other words, closed subscheme exact sequences pull back (remain exact) under flat pullbacks. This is in fact not just a necessary condition for flatness; it is also sufficient, which can be shown using the ideal-theoretic criterion for flatness (Theorem 25.4.1). There is an analogous fact about pulling ideal sheaves of *flat* subschemes by *arbitrary* pullbacks, see §25.3.2.

25.2.N. EXERCISE.

- (a) Suppose D is an effective Cartier divisor on Y and $\pi: X \rightarrow Y$ is a flat morphism. Show that the pullback of D to X (by π) is also an effective Cartier divisor.
 (b) Use part (a) to show that under a flat morphism, regular embeddings pull back to regular embeddings.

25.2.O. UNIMPORTANT EXERCISE.

- (a) Suppose $\pi: X \rightarrow Y$ is a morphism, and $Z \hookrightarrow Y$ is a closed embedding cut out by an ideal sheaf $\mathcal{I} \subset \mathcal{O}_Y$. Show that $(\pi^* \mathcal{I})^n = \pi^*(\mathcal{I}^n)$. (Products of quasicoherent ideal sheaves were defined in Exercise 15.3.E.)
 (b) Suppose further that π is flat, $Y = \mathbb{A}_k^n$, and Z is the origin. Let $\mathcal{J} = \pi^* \mathcal{I}$ be the quasicoherent sheaf of algebras on X cutting out the pullback W of Z . Prove that the graded sheaf of algebras $\bigoplus_{n \geq 0} \mathcal{J}^n / \mathcal{J}^{n+1}$ (do you understand the multiplication?) is isomorphic to $\mathcal{O}_W[x_1, \dots, x_n]$ (interpreted as a graded sheaf of algebras). (Hint: first show that $\mathcal{J}^n / \mathcal{J}^{n+1} \cong \text{Sym}^n(\mathcal{J} / \mathcal{J}^2)$.)

25.2.P. UNIMPORTANT EXERCISE.

- (a) Show that blowing up commutes with flat base change. More precisely, if $\pi: X \rightarrow Y$ is any flat morphism, and $Z \hookrightarrow Y$ is any closed embedding, give a canonical isomorphism $(\text{Bl}_Z Y) \times_Y X \xrightarrow{\sim} \text{Bl}_{Z \times_Y X} X$. (You can proceed by universal property, using Exercise 25.2.N(a), or by using the Proj construction of the blow up and Exercise 25.2.O.)
 (b) Give an example to show that blowing up does not commute with base change in general.

25.3 Flatness through Tor

We defined the Tor (bi-)functor in §24.1: $\text{Tor}_i^A(M, N)$ is obtained by taking a free resolution of N , removing the N , tensoring it with M , and taking homology. Exercise 24.1.D characterized flatness in terms of Tor: M is A -flat if $\text{Tor}_1^A(M, N) = 0$ for all N . In this section, we reap the easier benefits of this characterization, recalling key properties of Tor when needed. In §25.4, we work harder to extract more from Tor.

It is sometimes possible to compute Tor from its definition, as shown in the following exercise that we will use repeatedly.

25.3.A. EXERCISE. Define $(M : x)$ as $\{m \in M : xm = 0\} \subset M$ — it consists of the elements of M annihilated by x . If x is not a zerodivisor, show that

$$\text{Tor}_i^A(M, A/(x)) = \begin{cases} M/xM & \text{if } i = 0; \\ (M : x) & \text{if } i = 1; \\ 0 & \text{if } i > 1. \end{cases}$$

Hint: use the resolution

$$0 \longrightarrow A \xrightarrow{\times x} A \longrightarrow A/(x) \longrightarrow 0$$

of $A/(x)$.

25.3.1. Remark. As a corollary of Exercise 25.3.A, we see again that flat modules over an integral domain are torsion-free (and more generally, Observation 25.2.2). Also, Exercise 25.3.A gives the reason for the notation Tor — it is short for *torsion*.

25.3.B. EXERCISE. If B is A -flat, use the FHHF Theorem (Exercise 1.6.H(c)) to give an isomorphism $B \otimes_A \text{Tor}_i^A(M, N) \xleftarrow{\sim} \text{Tor}_i^B(B \otimes M, B \otimes N)$.

Recall that the Tor functor is symmetric in its entries (there is an isomorphism $\text{Tor}_i^A(M, N) \xleftarrow{\sim} \text{Tor}_i^A(N, M)$, Exercise 24.3.A). This gives us a quick but very useful result.

25.3.C. EASY EXERCISE. If $0 \rightarrow N' \rightarrow N \rightarrow N'' \rightarrow 0$ is an exact sequence of A -modules, and N'' is flat (e.g., free), show that $0 \rightarrow M \otimes_A N' \rightarrow M \otimes_A N \rightarrow M \otimes_A N'' \rightarrow 0$ is exact for *any* A -module M .

We would have cared about this result long before seeing Tor , so it gives some motivation for learning about Tor . (Unimportant side question: Can you prove this without Tor , using a diagram chase?)

25.3.D. EXERCISE (IMPORTANT CONSEQUENCE OF EXERCISE 25.3.C). Suppose $0 \rightarrow \mathcal{F}' \rightarrow \mathcal{F} \rightarrow \mathcal{F}'' \rightarrow 0$ is a short exact sequence of quasicoherent sheaves on a scheme Y , and \mathcal{F}'' is flat (e.g., locally free). Show that if $\pi: X \rightarrow Y$ is any morphism of schemes, the pulled back sequence $0 \rightarrow \pi^*\mathcal{F}' \rightarrow \pi^*\mathcal{F} \rightarrow \pi^*\mathcal{F}'' \rightarrow 0$ remains exact.

25.3.E. EXERCISE (CF. EXERCISE 14.2.R FOR THE ANALOGOUS FACTS ABOUT VECTOR BUNDLES). Suppose $0 \rightarrow M' \rightarrow M \rightarrow M'' \rightarrow 0$ is an exact sequence of A -modules.

- (a) If M and M'' are both flat, show that M' is too. (Hint: Recall the long exact sequence for Tor , Proposition 24.1.2. Also, use that N is flat if and only if $\text{Tor}_i(N, N') = 0$ for all $i > 0$ and all N' , Exercise 24.1.D.)
- (b) If M' and M'' are both flat, show that M is too. (Same hint.)
- (c) If M' and M are both flat, show that M'' need not be flat.

25.3.F. EXERCISE. If $\cdots \longrightarrow M_{-2} \longrightarrow M_{-1} \longrightarrow M_0 \longrightarrow 0$ is an exact sequence of flat A -modules, show that it remains exact upon tensoring with any other A -module. (One possible approach: break the exact sequence into short exact sequences.)

25.3.G. EASY EXERCISE. If $0 \rightarrow M_0 \rightarrow M_1 \rightarrow \cdots \rightarrow M_n \rightarrow 0$ is an exact sequence, and M_i is flat for $i > 0$, show that M_0 is flat too. (Hint: as always, break the exact sequence into short exact sequences.)

We will use the Exercises 25.3.F and 25.3.G later in this chapter.

25.3.2. Pulling back quasicoherent ideal sheaves of flat closed subschemes by arbitrary morphisms (promised in §25.2.10). Suppose

$$\begin{array}{ccc} W & \xrightarrow{\alpha} & X \\ \downarrow & & \downarrow \\ Y & \xrightarrow{\beta} & Z \end{array}$$

is a fibered product, and $V \hookrightarrow X$ is a closed subscheme. Then $Y \times_Z V$ is a closed subscheme of W (§10.2.2). There are two possible senses in which $\mathcal{I}_{V/X}$ can be “pulled back” to W : as a quasicoherent sheaf $\alpha^* \mathcal{I}_{V/X}$, and as the ideal of the “pulled back” closed subscheme $\mathcal{I}_{Y \times_Z V/W}$. As pointed out in Remark 17.3.9, these are not necessarily the same, but they *are* the same if β is flat. We now give another important case in which they are the same.

25.3.H. EXERCISE. If V is flat over Z (with no hypotheses on β), show that $\alpha^* \mathcal{I}_{V/X} \cong \mathcal{I}_{Y \times_Z V/W}$. Hint: Exercise 25.3.D.

25.4 Ideal-theoretic criteria for flatness

The following theorem will allow us to classify flat modules over a number of rings. It is a refined version of Exercise 24.1.D, that M is a flat A -module if and only if $\mathrm{Tor}_1^A(M, N) = 0$ for all A -modules N .

25.4.1. Theorem (ideal-theoretic criterion for flatness). — *The A -module M is flat if and only if $\mathrm{Tor}_1^A(M, A/I) = 0$ for every ideal I .*

25.4.2. Remarks. Before getting to the proof, we make some side remarks that may give some insight into how to think about flatness. Theorem 25.4.1 is profitably stated without the theory of Tor . It is equivalent to the statement that M is flat if and only if for all ideals $I \subset A$, $I \otimes_A M \rightarrow M$ is an injection, and you can reinterpret the proof in this guise. Perhaps better, M is flat if and only if $I \otimes_A M \rightarrow IM$ is an isomorphism for every ideal I .

Flatness is often informally described as “continuously varying fibers”, and this can be made more precise as follows. An A -module M is flat if and only if it restricts nicely to closed subschemes of $\mathrm{Spec} A$. More precisely, what we lose in this restriction, the submodule IM of elements which “vanish on Z ”, is easy to understand: it consists of formal linear combinations of elements $i \otimes m$, with no surprise relations among them — i.e., the tensor product $I \otimes_A M$. This is the content of the following exercise, in which you may use Theorem 25.4.1.

25.4.A. ★ EXERCISE (THE EQUATIONAL CRITERION FOR FLATNESS). Show that an A -module M is flat if and only if for every relation $\sum a_i m_i = 0$ with $a_i \in A$ and $m_i \in M$, there exist $m'_j \in M$ and $a_{ij} \in A$ such that $\sum_j a_{ij} m'_j = m_i$ for all i and $\sum_i a_i a_{ij} = 0$ in A for all j . (Translation: whenever elements of M satisfy an A -linear relation, this is “because” of linear equations holding in A .)

25.4.3. Unimportant remark. In the statement of Theorem 25.4.1, it suffices to check only finitely generated ideals. This is essentially the content of the following statement, which you can prove if you wish: Show that an A -module M is flat if and only if for all *finitely generated ideals* I , the natural map $I \otimes_A M \rightarrow M$ is an injection. Hint: if there is a counterexample for an ideal J that is not finitely generated, use it to find another counterexample for an ideal I that *is* finitely generated.

25.4.4. Proof of the ideal-theoretic criterion for flatness, Theorem 25.4.1. By Exercise 24.1.D, we need only show that $\mathrm{Tor}_1^A(M, A/I) = 0$ for all I implies $\mathrm{Tor}_1^A(M, N) = 0$ for all A -modules N , and hence that M is flat.

We first prove that $\mathrm{Tor}_1^A(M, N) = 0$ for all *finitely generated* modules N , by induction on the number n of generators a_1, \dots, a_n of N . The base case (if $n = 1$, so $N \cong A/\mathrm{Ann}(a_1)$ — our first use of the annihilator ideal in a long time) is our assumption. If $n > 1$, then $Aa_n \cong A/\mathrm{Ann}(a_n)$ is a submodule of N , and the quotient Q is generated by the images of a_1, \dots, a_{n-1} , so the result follows by considering the Tor_1 portion of the Tor long exact sequence for

$$0 \longrightarrow A/\mathrm{Ann}(a_n) \longrightarrow N \longrightarrow Q \longrightarrow 0.$$

We deal with the case of general N by abstract nonsense. Notice that N is the union of its finitely generated submodules $\{N_\alpha\}$. In fancy language, this union is a filtered colimit — any two finitely generated submodules are contained in a finitely generated submodule (specifically, the submodule they generate). Filtered colimits of modules commute with cohomology (Exercise 1.6.L), so $\mathrm{Tor}_1(M, N)$ is the colimit over α of $\mathrm{Tor}_1(M, N_\alpha) = 0$, and is thus 0. \square

We now use Theorem 25.4.1 to get explicit characterizations of flat modules over three (types of) rings: principal ideal domains, dual numbers, and some local rings.

Recall Observation 25.2.2, that flatness implies torsion-free. The converse is true for principal ideal domains:

25.4.B. EXERCISE (FLAT = TORSION-FREE FOR A PID). Show that a module over a principal ideal domain is flat if and only if it is torsion-free.

25.4.C. EXERCISE (FLATNESS OVER A DVR). Suppose M is a module over a discrete valuation ring A with uniformizer t . Show that M is flat if and only if t is not a zerodivisor on M , i.e., (using the notation defined in Exercise 25.3.A) $(M : t) = 0$. (See Exercise 25.2.B for the case of finitely generated modules.) This yields a simple and very important geometric interpretation of flatness over a regular curve, which we discuss in §25.4.9.

25.4.D. EXERCISE (FLATNESS OVER THE DUAL NUMBERS). Show that M is flat over $k[t]/(t^2)$ if and only if the “multiplication by t ” map $M/tM \rightarrow tM$ is an isomorphism. (This fact is important in deformation theory and elsewhere.) Hint: $k[t]/(t^2)$ has only three ideals.

25.4.5. Important Theorem (flat = free = projective for finitely presented modules over local rings). — Suppose (A, \mathfrak{m}) is a local ring (not necessarily Noetherian), and M is a finitely presented A -module. Then M is flat if and only if it is free if and only if it is projective.

25.4.6. Remarks. Warning: modules over local rings can be flat without being free: \mathbb{Q} is a flat $\mathbb{Z}_{(p)}$ -algebra ($\mathbb{Z}_{(p)}$ is the localization of \mathbb{Z} at p , not the p -adics), as all localizations are flat (§25.2.1), but it is not free (Do you see why?).

Also, non-Noetherian people may be pleased to know that with a little work, “finitely presented” can be weakened to “finitely generated”: use [Mat2, Thm. 7.10] in the proof below, where finite presentation comes up.

Proof. For any ring, free modules are projective (§24.2.1), and projective modules are flat (Exercise 24.2.B), so we need only show that flat modules are free for a local ring.

(At this point, you should see Nakayama coming from a mile away.) Now $M/\mathfrak{m}M$ is a finite-dimensional vector space over the field A/\mathfrak{m} . Choose a basis of $M/\mathfrak{m}M$, and lift it to elements $m_1, \dots, m_n \in M$. Consider $A^{\oplus n} \rightarrow M$ given by $e_i \mapsto m_i$. We will show this is an isomorphism. It is surjective by Nakayama’s Lemma (see Exercise 8.2.H): the image is all of M modulo the maximal ideal, hence is everything. As M is finitely presented, by Lemma 6.4.4 (“finitely presented implies always finitely presented”), the kernel K is finitely generated. Tensor $0 \rightarrow K \rightarrow A^{\oplus n} \rightarrow M \rightarrow 0$ with A/\mathfrak{m} . As M is flat, the result is still exact (Exercise 25.3.C):

$$0 \longrightarrow K/\mathfrak{m}K \longrightarrow (A/\mathfrak{m})^{\oplus n} \longrightarrow M/\mathfrak{m}M \longrightarrow 0.$$

But $(A/\mathfrak{m})^{\oplus n} \rightarrow M/\mathfrak{m}M$ is an isomorphism by construction, so $K/\mathfrak{m}K = 0$. As K is finitely generated, $K = 0$ by Nakayama’s Lemma 8.2.9. \square

Here is an immediate and useful corollary — really just a geometric interpretation.

25.4.7. Corollary (flat = locally free for finitely presented sheaves). — *A finitely presented sheaf \mathcal{F} on X is flat (over X) if and only if it is locally free.*

Proof. Local freeness of a finitely presented sheaf can be checked at the stalks, Exercise 14.3.F. \square

25.4.E. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a finite flat morphism of locally Noetherian schemes, and \mathcal{F} is a finite rank locally free sheaf on X .

- Show that $\pi_*\mathcal{F}$ is a finite rank locally free sheaf on Y .
- If Y is irreducible with generic point η , the degree of π above η is n , and \mathcal{F} is locally free of rank r , show that $\pi_*\mathcal{F}$ is locally free of rank nr . (This may be trickier than it looks.)

25.4.8. Remark. In particular, taking $\mathcal{F} = \mathcal{O}_X$, we see that in the locally Noetherian setting, *finite flat morphisms have locally constant degree*.

25.4.F. ★ EXERCISE (INTERESTING VARIANT OF THEOREM 25.4.5, BUT UNIMPORTANT FOR US). Suppose A is a ring (not necessarily local), and M is a finitely presented A -module. Show that M is flat if and only if it is projective. Hint: show that M is projective if and only if $M_{\mathfrak{m}}$ is free for every maximal ideal \mathfrak{m} . The harder direction of this implication uses the fact that $\mathrm{Hom}_{A_{\mathfrak{m}}}(M_{\mathfrak{m}}, N_{\mathfrak{m}}) = \mathrm{Hom}_A(M, N)_{\mathfrak{m}}$, which follows from Exercise 1.6.G. (Remark: There exist finitely *generated* flat modules that are not projective. They are necessarily not finitely presented. Example

without proof: let $A = \prod_{i=1}^{\infty} \mathbb{F}_2$, interpreted as functions $\mathbb{Z}^{\geq 0} \rightarrow \mathbb{Z}/2$, and let M be the module of functions modulo those of proper support, i.e., those vanishing at almost all points of $\mathbb{Z}^{\geq 0}$.)

25.4.G. EXERCISE. Prove the following useful criterion for flatness: Suppose $\pi: X \rightarrow Y$ is a finite morphism, and Y is reduced and locally Noetherian. Then π is flat if and only if $\pi_* \mathcal{O}_X$ is locally free, if and only if the rank of $\pi_* \mathcal{O}_X$ is locally constant ($\dim_{\kappa(q)}(\pi_* \mathcal{O}_X)_q \otimes \kappa(q)$ is a locally constant function of $q \in Y$). Partial hint: Exercise 14.3.K.

25.4.H. EXERCISE. Show that the normalization of the node (see Figure 8.4) is not flat. Hint: use Exercise 25.4.G. (This exercise can be strengthened to show that nontrivial normalizations are *never* flat.)

25.4.I. EXERCISE. In $\mathbb{A}_k^4 = \text{Spec } k[w, x, y, z]$, let X be the union of the wx -plane with the yz -plane:

$$(25.4.8.1) \quad X = \text{Spec } k[w, x, y, z]/(wy, wz, xy, xz).$$

The projection $\mathbb{A}_k^4 \rightarrow \mathbb{A}_k^2$ given by $k[a, b] \rightarrow k[w, x, y, z]$ with $a \mapsto w - y$, $b \mapsto x - z$ restricts to a morphism $X \rightarrow \mathbb{A}_k^2$. Show that this morphism is not flat.

25.4.9. Flat families over regular curves. Exercise 25.4.C gives an elegant geometric criterion for when morphisms to regular curves are flat.

25.4.J. EXERCISE (CRITERION FOR FLATNESS OVER A REGULAR CURVE). Suppose $\pi: X \rightarrow Y$ is a morphism from a locally Noetherian scheme to a regular (locally Noetherian) curve. (The local Noetherian hypothesis on X is so we can discuss its associated points.) Show that π is flat if and only if all associated points of X map to a generic point of Y . (This is a partial converse to Exercise 25.2.J, that flat maps always send associated points to associated points. As with Exercise 25.2.J, feel free to immediately generalize your argument to a coherent sheaf \mathcal{F} on X .)

25.4.10. For example, a nonconstant map from an integral (locally Noetherian) scheme to a regular curve must be flat. (As another example, the deformation to the normal cone, discussed in the double-starred section §23.4.17 and depicted in Figure 23.3, is flat.) Exercise 25.4.H (and the comment after it) shows that the regular condition is necessary. The example of two planes meeting at a point in Exercise 25.4.I shows that the dimension 1 condition is necessary.

25.4.11. ★ Unimportant remark: A valuative criterion for flatness. Exercise 25.4.J shows that flatness over a regular curves is geometrically intuitive (and is “visualizable”). It gives a criterion for flatness in general: suppose $\pi: X \rightarrow Y$ is finitely presented morphism. If π is flat, then for every morphism $Y' \rightarrow Y$ where Y' is the Spec of a discrete valuation ring, $\pi': X \times_Y Y' \rightarrow Y'$ is flat, so no associated points of $X \times_Y Y'$ map to the closed point of Y' . If Y is reduced and locally Noetherian, then this is a sufficient condition; this can reasonably be called a *valuative criterion for flatness*. (Reducedness is necessary: consider Exercise 25.2.G(d).) This gives an excellent way to visualize flatness, which you should try to put into words (perhaps after learning about flat limits below). See [Gr-EGA, IV₃.11.8.1] for a proof (and an extension without Noetherian hypotheses).

25.4.12. Revisiting the degree of a projective morphism from a curve to a regular curve. As hinted after the statement of Proposition 18.4.5, we can now better understand why nonconstant projective morphisms from a curve to a regular curve have a well-defined degree, which can be determined by taking the preimage of any point (§18.4.4). (Example 10.3.4 was particularly enlightening.) This is because such maps are flat by Exercise 25.4.J, and then the degree is constant by Remark 25.4.8 (see also Exercise 25.4.G). Also, Exercise 25.4.G yields a new proof of Proposition 18.4.5.

25.4.13. Flat limits. Here is an important consequence of Exercise 25.4.J, which we can informally state as: we can take flat limits over one-parameter families. More precisely: suppose A is a discrete valuation ring, and let 0 be the closed point of $\text{Spec } A$ and η the generic point. Suppose X is a locally Noetherian scheme over A , and Y is a closed subscheme of $X|_{\eta}$. Let Y' be the scheme-theoretic closure of Y in X . Then Y' is flat over A . Similarly, suppose Z is a one-dimensional Noetherian scheme, 0 is a regular point of Z , and $\pi: X \rightarrow Z$ is a morphism from a locally Noetherian scheme to Z . Assume that $\pi^{-1}(Z - \{0\})$ is flat over $Z - \{0\}$. If Y is a closed subscheme of $\pi^{-1}(Z - \{0\})$, and Y' is the scheme-theoretic closure of Y in X , then Y' is flat over Z . In both cases, the closure $Y'|_0$ is often called the **flat limit** of Y . (Feel free to weaken the Noetherian hypotheses on X .)

25.4.K. EXERCISE. Suppose (with the language of the previous paragraph) that A is a discrete valuation ring, X is a locally Noetherian A -scheme, and Y is a closed subscheme of the generic fiber $X|_{\eta}$. Show that there is only one closed subscheme Y' of X such that $Y'|_{\eta} = Y$, and Y' is flat over A .

25.4.L. HARDER EXERCISE (AN EXPLICIT FLAT LIMIT). Let $X = \mathbb{A}^3 \times \mathbb{A}^1 \rightarrow Z = \mathbb{A}^1$ over a field k , where the coordinates on \mathbb{A}^3 are x, y , and z , and the coordinates on \mathbb{A}^1 are t . Define Y away from $t = 0$ as the union of the two lines $y = z = 0$ (the x -axis) and $x = z - t = 0$ (the y -axis translated by t). Find the flat limit at $t = 0$. (Hints: (i) it is *not* the union of the two axes, although it includes this union. The flat limit is nonreduced at the node, and the “fuzz” points out of the plane they are contained in. (ii) $(y, z)(x, z) \neq (xy, z)$. (iii) Once you have a candidate flat limit, be sure to check that it *is* the flat limit. (iv) If you get stuck, read Example 25.4.14 below.)

Consider a projective version of the previous example, where two lines in \mathbb{P}^3 degenerate to meet. The limit consists of two lines meeting at a node, with some nonreduced structure at the node. Before the two lines come together, their space of global sections is two-dimensional. When they come together, it is not immediately obvious that their flat limit also has two-dimensional space of global sections as well. The reduced version (the union of the two lines meeting at a point) has a one-dimensional space of global sections, but the effect of the nonreduced structure on the space of global sections may not be immediately clear. However, we will see that “cohomology groups can only jump up in flat limits”, as a consequence (indeed the main moral) of the Semicontinuity Theorem 28.1.1.

25.4.14. ** Example of variation of cohomology groups in flat families. We can use a variant of Exercise 25.4.L to see an example of a cohomology group actually jumping. We work over an algebraically closed field to avoid distractions. Before

we get down to explicit algebra, here is the general idea. Consider a twisted cubic C in \mathbb{P}^3 . A projection pr_p from a random point $p \in \mathbb{P}^3$ will take C to a nodal plane cubic. Picture this projection “dynamically”, by choosing coordinates so p is at $[1, 0, 0, 0]$, and considering the map $\phi_t: [w, x, y, z] \mapsto [tw, x, y, z]$; ϕ_1 is the identity on \mathbb{P}^3 , ϕ_t is an automorphism of \mathbb{P}^3 for $t \neq 0$, and ϕ_0 is the projection. The limit of $\phi_t(C)$ as $t \rightarrow 0$ will be a nodal cubic, with nonreduced structure at the node “analytically the same” as what we saw when two lines came together (Exercise 25.4.L). (The phrase “analytically the same” can be made precise once we define completions in §30.)

Let’s now see this in practice. Rather than working directly with the twisted cubic, we use another example where we saw a similar picture. Consider the nodal (affine) plane cubic $y^2 = x^3 + x^2$. Its normalization (see Figure 8.4, Example (3) of §8.3.5, Exercise 10.7.E, ...) was obtained by adding an extra variable m corresponding to y/x (which can be interpreted as blowing up the origin, see §23.4.3). We use the variable m rather than t (used in §8.3.5) in order to reserve t for the parameter for the flat family.

We picture the nodal cubic C as lying in the xy -plane in $\mathbb{A}^3 = \text{Spec } k[x, y, m]$, and the normalization \tilde{C} projecting to it, with $m = y/x$. What are the equations for \tilde{C} ? Clearly, they include the equations $y^2 = x^3 + x^2$ and $y = mx$, but these are not enough — the m -axis (i.e., $x = y = 0$) is also in $V(y^2 - x^3 - x^2, y - mx)$. A little thought (and the algebra we have seen earlier in this example) will make clear that we have a third equation $m^2 = (x + 1)$, which along with $y = mx$ implies $y^2 = x^2 + x^3$. Now we have enough equations: $k[x, y, m]/(m^2 - (x + 1), y - mx)$ is an integral domain, as it is clearly isomorphic to $k[m]$. Indeed, you should recognize this as the algebra appearing in Exercise 10.7.E.

Next, we want to formalize our intuition of the dynamic projection to the xy -plane of $\tilde{C} \subset \mathbb{A}^3$. We picture it as follows. Given a point (x, y, m) at time 1, at time t we want it to be at (x, y, mt) . At time $t = 1$, we “start with” \tilde{C} , and at time $t = 0$ we have (set-theoretically) C . Thus at time $t \neq 0$, the curve \tilde{C} is sent to the curve cut out by equations

$$k[x, y, m]/(m^2 - t^2(x + 1), ty - mx).$$

The family over $\text{Spec } k[t, t^{-1}]$ is thus

$$k[x, y, m, t, t^{-1}]/(m^2 - t^2(x + 1), ty - mx).$$

Notice that we have inverted t because we are so far dealing only with nonzero t . For $t \neq 0$, this is certainly a “nice” family, and so surely flat. Let’s make sure this is true.

25.4.M. EXERCISE. Check this, as painlessly as possible! Hint: by a clever change of coordinates, show that the family is constant “over $\text{Spec } k[t, t^{-1}]$ ”, and hence pulled back (in some way you must figure out) via $\text{Spec } k[t, t^{-1}] \rightarrow \text{Spec } k$ from

$$\text{Spec } k[X, Y, M]/(M^2 - (X + 1), Y - MX) \longrightarrow \text{Spec } k,$$

which is flat by Trick Question 25.2.A(a).

We now figure out the flat limit of this family over $t = 0$, in $\text{Spec } k[x, y, m, t] \rightarrow \mathbb{A}^1 = \text{Spec } k[t]$. We first hope that our flat family is given by the equations we have

already written down:

$$\operatorname{Spec} k[x, y, m, t]/(m^2 - t^2(x+1), ty - mx).$$

But this is *not* flat over $\mathbb{A}^1 = \operatorname{Spec} k[t]$, as the fiber dimension jumps (§25.5.6): substituting $t = 0$ into the equations (obtaining the fiber over $0 \in \mathbb{A}^1$), we find $\operatorname{Spec} k[x, y, m]/(m^2, mx)$. This is set-theoretically the xy -plane ($m = 0$), which of course has dimension 2. Notice for later reference that this “false limit” is scheme-theoretically the xy -plane, *with some nonreduced structure along the y -axis*. (This may remind you of Figure 4.4.)

So we are missing at least one equation. One clue as to what equation is missing: the equation $y^2 = x^3 + x^2$ clearly holds for $t \neq 0$, and does *not* hold for our naive attempt at a limit scheme $m^2 = mx = 0$. So we put this equation back in, and have a second hope for describing the flat family over \mathbb{A}^1 :

$$\operatorname{Spec} k[x, y, m, t]/(m^2 - t^2(x+1), ty - mx, y^2 - x^3 - x^2) \longrightarrow \operatorname{Spec} k[t].$$

Let $A = k[x, y, m, t]/(m^2 - t^2(x+1), ty - mx, y^2 - x^3 - x^2)$ for convenience. The morphism $\operatorname{Spec} A \rightarrow \mathbb{A}^1$ is flat at $t = 0$. How can we show it? We could hope to show that A is an integral domain, and thus invoke Exercise 25.4.J. Instead we use Exercise 25.4.B, and show that t is not a zerodivisor on A . We do this by giving a “normal form” for elements of A .

25.4.N. EXERCISE. Show that each element of A can be written uniquely as a polynomial in x, y, m , and t such that no monomial in it is divisible by m^2, mx , or y^2 . Then show that t is not a zerodivisor on A , and conclude that $\operatorname{Spec} A \rightarrow \mathbb{A}^1$ is indeed flat.

Thus the flat limit when $t = 0$ is given by

$$\operatorname{Spec} k[x, y, m]/(m^2, mx, y^2 - x^3 - x^2).$$

25.4.O. EXERCISE. Show that the flat limit is nonreduced, and the “nonreducedness has length 1 and supported at the origin”. More precisely, if $X = \operatorname{Spec} A/(t)$, show that $\mathcal{I}_{X^{\text{red}}}$ is a skyscraper sheaf, with value k , supported at the origin. Sketch this flat limit X .

25.4.15. Note that we have a nonzero global function on X , given by m , which is supported at the origin (i.e., 0 away from the origin).

We now use this example to get a projective example with interesting behavior. We take the projective completion of this example, to get a family of cubic curves in \mathbb{P}^3 degenerating to a nodal cubic C with a nonreduced point.

25.4.P. EXERCISE. Do this: describe this family (in $\mathbb{P}^3 \times \mathbb{A}^1$) precisely.

Take the long exact sequence corresponding to

$$0 \longrightarrow \mathcal{I}_{C^{\text{red}}} \longrightarrow \mathcal{O}_C \longrightarrow \mathcal{O}_{C^{\text{red}}} \longrightarrow 0,$$

to get

$$\begin{aligned}
 0 &\longrightarrow H^0(C, \mathcal{I}_{C^{\text{red}}}) \xrightarrow{\alpha} H^0(C, \mathcal{O}_C) \longrightarrow H^0(C, \mathcal{O}_{C^{\text{red}}}) \\
 &\longrightarrow H^1(C, \mathcal{I}_{C^{\text{red}}}) \longrightarrow H^1(C, \mathcal{O}_C) \longrightarrow H^1(C, \mathcal{O}_{C^{\text{red}}}) \\
 &\longrightarrow H^2(C, \mathcal{I}_{C^{\text{red}}})
 \end{aligned}$$

We have $H^1(C, \mathcal{I}_{C^{\text{red}}}) = H^2(C, \mathcal{I}_{C^{\text{red}}}) = 0$ as $\mathcal{I}_{C^{\text{red}}}$ is supported in dimension 0 (by dimensional cohomology vanishing, Theorem 19.2.6). Also, $H^i(C^{\text{red}}, \mathcal{O}_{C^{\text{red}}}) = H^i(C, \mathcal{O}_{C^{\text{red}}})$ (property (v) of cohomology, see §19.1). The (reduced) nodal cubic C^{red} has $h^0(\mathcal{O}) = 1$ (§11.5.7) and $h^1(\mathcal{O}) = 1$ (cubic plane curves have genus 1, (19.6.6.1)). Also, $h^0(C, \mathcal{I}_{C^{\text{red}}}) = 1$ as observed above. Finally, α is not 0, as there exists a nonzero function on C vanishing on C^{red} (§25.4.15 — convince yourself that this function extends from the affine patch $\text{Spec } A$ to the projective completion).

Using the long exact sequence, we conclude $h^0(C, \mathcal{O}_C) = 2$ and $h^1(C, \mathcal{O}_C) = 1$. Thus in this example we see that $(h^0(\mathcal{O}), h^1(\mathcal{O})) = (1, 0)$ for the general member of the family (twisted cubics are isomorphic to \mathbb{P}^1), and the special member (the flat limit) has $(h^0(\mathcal{O}), h^1(\mathcal{O})) = (2, 1)$. Notice that both cohomology groups have jumped, yet the Euler characteristic has remained the same. The first behavior, as stated after Exercise 25.4.L, is an example of the Semicontinuity Theorem 28.1.1. The second, constancy of Euler characteristics in flat families, is what we turn to next. (It is no coincidence that the example had a singular limit, see §28.1.3.)

25.5 Topological aspects of flatness

We now discuss some topological aspects and consequences of flatness, that boil down to the Going-Down Theorem for flat morphisms (§25.5.3), which in turn comes from faithful flatness. Because dimension in algebraic geometry is a topological notion, we will show that dimensions of fibers behave well in flat families (§25.5.6).

25.5.1. Faithful flatness. The notion of faithful flatness is handy for many reasons, and we describe only a few. A B -module M is **faithfully flat** if for all complexes of B -modules

$$(25.5.1.1) \quad N' \rightarrow N \rightarrow N'',$$

(25.5.1.1) is exact if and only if $(25.5.1.1) \otimes_B M$ is exact. A B -algebra A is **faithfully flat** if it is faithfully flat as a B -module.

25.5.A. EXERCISE. Show that a flat B -module M is faithfully flat if and only if for all B -modules N , $M \otimes_B N = 0$ implies that $N = 0$.

25.5.B. EXERCISE. Suppose M is a flat B -module. Show that the following are equivalent.

- (a) M is faithfully flat;
- (b) for all prime ideals $\mathfrak{p} \subset B$, $M \otimes_B \kappa(\mathfrak{p})$ is nonzero;
- (c) for all maximal ideals $\mathfrak{m} \subset B$, $M \otimes_B \kappa(\mathfrak{m}) = M/\mathfrak{m}M$ is nonzero.

Suppose $\pi: X \rightarrow Y$ is a morphism of schemes. We say that π is **faithfully flat** if it is flat and surjective. (Unlike flatness, faithful flatness is not particularly useful for quasicoherent sheaves, so we do not define faithfully flat quasicoherent sheaves over a base.)

25.5.C. EXERCISE (CF. EXERCISE 25.5.B(B)). Suppose $B \rightarrow A$ is a ring morphism. Show that A is faithfully flat over B if and only if $\text{Spec } A \rightarrow \text{Spec } B$ is faithfully flat.

25.5.2. Faithfully flat morphisms behave “reasonably”. Faithfully flat morphisms behave “reasonably” in the sense of §8.1 (preserved by composition and base change, and local on the target) because both flatness and surjectivity have this property (§25.2.8 and Exercise 10.4.F respectively).

25.5.D. EXERCISE. Suppose $\pi: \text{Spec } A \rightarrow \text{Spec } B$ is flat.

- (a) Show that π is faithfully flat if and only if every *closed* point $q \in \text{Spec } B$ is in the image of π . (Hint: Exercise 25.5.B(c).)
- (b) Hence show that *every* flat (local) morphism of local rings (Definition 7.3.1) is faithfully flat. (Morphisms of local rings are assumed to be local, i.e., the maximal ideal pulls back to the maximal ideal.)

25.5.3. Going-Down for flat morphisms. A consequence of Exercise 25.5.D is the following useful result, whose statement makes no mention of faithful flatness. (The statement is not coincidentally reminiscent of the Going-Down Theorem for finite extensions of integrally closed domains, Theorem 12.2.12.)

25.5.E. EXERCISE (GOING-DOWN THEOREM FOR FLAT MORPHISMS).

- (a) Suppose that $B \rightarrow A$ is a flat morphism of rings, corresponding to a map $\pi: \text{Spec } A \rightarrow \text{Spec } B$. Suppose $q \subset q'$ are prime ideals of B , and \mathfrak{p}' is a prime ideal of A with $\pi([\mathfrak{p}']) = [q']$. Show that there exists a prime $\mathfrak{p} \subset \mathfrak{p}'$ of A with $\pi([\mathfrak{p}]) = [q]$. Hint: show that $B_{q'} \rightarrow A_{\mathfrak{p}'}$ is a flat local ring homomorphism, and hence faithfully flat by the Exercise 25.5.D(b).
- (b) Part (a) gives a geometric consequence of flatness. Draw a picture illustrating this.
- (c) Recall the Going-Up Theorem, described in §8.2.4. State the Going-Down Theorem for flat morphisms in a way parallel to Exercise 8.2.F, and prove it.

25.5.F. EXERCISE. Suppose $\pi: X \rightarrow Y$ is an integral (e.g., finite) flat morphism, and Y has pure dimension n . Show that X has pure dimension n . (This generalizes Exercise 12.1.G(a).) Hint: π satisfies both Going-Up (see Exercise 8.2.F) and Going-Down.

25.5.G. IMPORTANT EXERCISE: FLAT MORPHISMS ARE OPEN (IN REASONABLE SITUATIONS). Suppose $\pi: X \rightarrow Y$ is locally of finite type and flat, and Y (and hence X) is locally Noetherian. Show that π is an open map (i.e., sends open sets to open sets). Hint: reduce to showing that $\pi(X)$ is open for all such π . Reduce to the case where X and Y are both affine. Use Chevalley’s Theorem 8.4.2 to show that $\pi(X)$ is

constructible. Use the Going-Down Theorem for flat morphisms, Exercise 25.5.E, to show that $\pi(X)$ is closed under generization. Conclude using Exercise 8.4.C.

25.5.4. Remark (proof of Proposition 10.5.6). Suppose A and B are finite type k -algebras. Then by Exercise 25.5.G, $\text{Spec } A \times_{\text{Spec } k} \text{Spec } B \rightarrow \text{Spec } B$ is an open map.

25.5.5. Follow-ups to Exercise 25.5.G.

- (i) Of course, not all open morphisms are flat: witness $\text{Spec } k[t]/(t) \rightarrow \text{Spec } k[t]/(t^2)$.
- (ii) Also, in quite reasonable circumstances, flat morphisms are *not* open: witness $\text{Spec } k(t) \rightarrow \text{Spec } k[t]$ (flat by Example 25.2.1(ii)).
- (iii) On the other hand, you can weaken the hypotheses of “locally of finite type” and “locally Noetherian” to just “locally finitely presented” [Gr-EGA, IV₂.2.4.6] — as with the similar generalization in Exercise 10.3.I of Chevalley’s Theorem 8.4.2, use the fact that any such morphism is “locally” pulled back from a Noetherian situation. We won’t use this, and hence omit the details.

25.5.6. Dimensions of fibers are well-behaved for flat morphisms.

25.5.7. Proposition. — Suppose $\pi: X \rightarrow Y$ is a flat morphism of locally Noetherian schemes, with $p \in X$ and $q \in Y$ such that $\pi(p) = q$. Then

$$\text{codim}_X p = \text{codim}_Y q + \text{codim}_{\pi^{-1}(q)} p$$

(see Figure 12.3).

Informal translation: the dimension of the fibers is the difference of the dimensions of X and Y (at least locally). Compare this to Exercise 12.4.A, which stated that without the flatness hypothesis, we would only have inequality (\leq).

25.5.H. EXERCISE. Prove Proposition 25.5.7 as follows. As just mentioned, Exercise 12.4.A gives one inequality, so show the other. Given a chain of irreducible closed subsets in Y containing \bar{q} , and a chain of irreducible closed subsets in $\pi^{-1}(q) \subset X$ containing \bar{p} , construct a chain of irreducible closed subsets in X containing \bar{p} , using the Going-Down Theorem for flat morphisms (Exercise 25.5.E).

As a consequence of Proposition 25.5.7, if $\pi: X \rightarrow Y$ is a flat map of irreducible varieties, then the fibers of π all have pure dimension $\dim X - \dim Y$. (Warning: $\text{Spec } k[t]/(t) \rightarrow \text{Spec } k[t]/(t^2)$ does not exhibit dimensional jumping of fibers, is open, and sends associated points to associated points, cf. Exercise 25.2.J, but is not flat. If you prefer a reduced example, the normalization of the cuspidal plane cubic $\text{Spec } k[t] \rightarrow \text{Spec } k[x, y]/(y^2 - x^3)$, shown in Figure 10.4, also has these properties.) This leads us to the following useful definition.

25.5.8. Definition. If $\pi: X \rightarrow Y$ is a flat morphism that is locally of finite type, and all fibers of π have pure dimension n , we say that π is **flat of relative dimension n** . (In particular, when one says a morphism is flat of relative dimension n , the locally finite type hypotheses are implied. Remark 25.5.10 motivates this hypothesis.)

25.5.I. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a flat morphism of finite type k -schemes, and Y is equidimensional (so “codimension is the difference of dimensions”, cf. Theorem 12.2.9). Show that the following are equivalent.

- (i) The scheme X has pure dimension $\dim Y + n$.
- (ii) The morphism π is flat of relative dimension n .

25.5.J. EXERCISE. Suppose $\pi: X \rightarrow Y$ and $\rho: Y \rightarrow Z$ are morphisms of locally Noetherian schemes, flat of relative dimension m and n respectively (hence locally of finite type). Show that $\rho \circ \pi$ is flat of relative dimension $m+n$. Hint: Exercise 25.5.I.

25.5.K. EXERCISE. Show that the notion of a morphism being “flat of relative dimension n ” is preserved by arbitrary base change. Hint to show that fiber dimension is preserved: Exercise 12.2.K.

25.5.9. Remark. For practice, you can now give a second solution to Exercise 12.2.F, that dimension is additive for products of varieties.

25.5.10. Remark. The reason for the “locally finite type” assumption in the definition of “flat of relative dimension n ” is that we want this to be a “reasonable” class of morphisms. In particular, we want our notion of “flatness of relative dimension n ” to be preserved by base change. Consider the fibered diagram

$$\begin{array}{ccc} \operatorname{Spec} k(x) \otimes_k k(y) & \longrightarrow & \operatorname{Spec} k(y) \\ \pi' \downarrow & & \downarrow \pi \\ \operatorname{Spec} k(x) & \longrightarrow & \operatorname{Spec} k. \end{array}$$

Both π and π' are trivially flat, because they are morphisms to spectra of fields (Exercise 25.2.A(a)). But the dimension of the fiber of π is 0, while (as described in Remark 12.2.16) the dimension of the fiber of π' is 1.

25.5.11. Generic flatness. “Generic flatness” (which would be better called “general flatness”), stated in Exercise 25.5.R, will be a consequence of the following fact. For the purposes of this discussion only, we say a B -algebra A satisfies (\dagger) if for each finitely generated A -module M , there exists a nonzero $f \in B$ such that M_f is a free B_f -module.

25.5.12. Grothendieck’s Generic Freeness Lemma. — Suppose B is a Noetherian integral domain. Then every finitely generated B -algebra satisfies (\dagger) .

Proof. We prove the Generic Freeness Lemma 25.5.12 in a series of exercises.

25.5.L. EXERCISE. Show that B itself satisfies (\dagger) .

25.5.M. EXERCISE. Reduce the proof of Lemma 25.5.12 to the following statement: if A is a finitely generated B -algebra satisfying (\dagger) , then $A[T]$ does too. (Hint: induct on the number of generators of A as a B -algebra.)

We now prove this statement. Suppose A satisfies (\dagger) , and let M be a finitely generated $A[T]$ -module, generated by the finite set S . Let $M_0 = 0$, and let M_1 be the sub- A -module of M generated by S . For $n > 0$, inductively define

$$M_{n+1} = M_n + TM_n,$$

a sub- A -module of M . Note that M is the increasing union of the A -modules M_n .

25.5.N. EXERCISE. Show that multiplication by T induces a surjection

$$\psi_n: M_n/M_{n-1} \longrightarrow M_{n+1}/M_n.$$

25.5.O. EXERCISE. Show that for $n \gg 0$, ψ_n is an isomorphism. Hint: use the ascending chain condition on M_1 .

25.5.P. EXERCISE. Show that there is a nonzero $f \in B$ such that $(M_{i+1}/M_i)_f$ is free as a B_f -module, for all i . Hint: as i varies, M_{i+1}/M_i passes through only finitely many isomorphism classes.

25.5.Q. EXERCISE (NOT REQUIRING NOETHERIAN HYPOTHESES). Suppose M is a B -module that is an increasing union of submodules M_i , with $M_0 = 0$, and that every M_{i+1}/M_i is free. Show that M is free. Hint: first construct compatible isomorphisms $\phi_n: \bigoplus_{i=0}^{n-1} M_{i+1}/M_i \xrightarrow{\sim} M_n$ by induction on n . Then show that the colimit $\phi := \operatorname{colim} \phi_n: \bigoplus_{i=0}^{\infty} M_{i+1}/M_i \rightarrow M$ is an isomorphism. Side Remark: More generally, your argument will show that if the M_{i+1}/M_i are all projective modules (to be defined in §24.2.1), then M is (non-naturally) isomorphic to their direct sum.

□

25.5.R. EASY EXERCISE (GENERIC FLATNESS). Suppose $\pi: X \rightarrow Y$ is a finite type morphism to a Noetherian integral scheme, and \mathcal{F} is a coherent sheaf on X . Show that there is a dense open subset $U \subset Y$ over which \mathcal{F} is flat. Translation: \mathcal{F} is flat at all points of X mapping to U . (An important special case is if $\mathcal{F} = \mathcal{O}_X$, which shows that there is a dense open subset U over which π is flat.) Hint: Grothendieck's Generic Freeness Lemma 25.5.12.

25.5.13. Interpretation of the degree of a generically finite morphism. We use this to interpret the degree of a rational map of varieties in terms of counting preimages, fulfilling a promise made in §12.2.2. Suppose $\pi: X \dashrightarrow Y$ is a generically finite rational map of k -varieties of degree d . For simplicity we assume that X and Y are irreducible. (But feel free to relax this.) Replace X by an open subset on which π is a morphism. By Proposition 12.4.3 (“generically finite implies generally finite”), there is a dense open subset V of Y over which π is finite. By “generic flatness” (Exercise 25.5.R), there is a dense open subset V' of V over which π is *finite and flat*. Then by Remark 25.4.8 (“finite flat morphisms have locally constant degree”), over this open subset V' , π has “locally constant degree”, which is necessarily d .

25.5.14. Exercise 25.5.R can be improved:

25.5.15. Theorem (generic flatness, improved version, [Stacks, tag 052B]). — If $\pi: X \rightarrow Y$ is a finite type morphism of schemes, Y is reduced, and \mathcal{F} is a finite type quasi-coherent sheaf on X , then there is an open dense subscheme $U \subset Y$ such that $\mathcal{F}|_{\pi^{-1}(U)}$ is of finite presentation over $\pi^{-1}(U)$, and flat over U .

We won't use this result, so we omit the proof.

25.5.16. Flatness is an open condition. Generic flatness can be used to show that in reasonable circumstances, the locus where a quasicoherent sheaf is flat over a base is an open subset. More precisely:

25.5.17. Theorem (flatness is an open condition). — Suppose $\pi: X \rightarrow Y$ is a locally finite type morphism of locally Noetherian schemes, and \mathcal{F} is a coherent sheaf on X .

- (a) The locus of points of X at which \mathcal{F} is Y -flat is an open subset of X .
- (b) If π is closed (e.g., proper), then the locus of points of Y over which \mathcal{F} is flat (i.e. those $q \in Y$ for which \mathcal{F} is flat at all $p \in \pi^{-1}(q)$) is an open subset of Y .

Part (b) follows immediately from part (a). Part (a) reduces to a nontrivial statement in commutative algebra, see for example [Mat2, Thm. 24.3] or [Gr-EGA, IV₃.11.1.1]. As is often the case, Noetherian hypotheses can be dropped in exchange for local finite presentation hypotheses on the morphism π , see [Gr-EGA, IV₃.11.3.1] or [Stacks, tag 00RC].

25.6 Local criteria for flatness

(This is the hardest section on ideal-theoretic criteria for flatness, and could profitably be postponed to a second reading.)

In the case of a Noetherian local ring, there is a greatly improved version of the ideal-theoretic criterion of Theorem 25.4.1: we need check only *one* ideal — the maximal ideal. The price we pay for the simplicity of this “local criterion for flatness” is that it is harder to prove.

25.6.1. Theorem (local criterion for flatness). — Suppose (A, \mathfrak{m}) is a Noetherian local ring, and M is a finitely generated A -module. Then M is flat if and only if $\mathrm{Tor}_1^A(M, A/\mathfrak{m}) = 0$.

This is a miracle: flatness over all of $\mathrm{Spec} A$ is determined by what happens over the closed point. (Caution: the finite generation hypothesis is necessary. Let $A = k[x, y]_{(x, y)}$ and $M = k(x)$, with y acting as 0. Then M is not flat by Observation 25.2.2, but it turns out that it satisfies the local criterion otherwise.)

Theorem 25.6.1 is an immediate consequence of the following more general statement.

25.6.2. Theorem (local criterion for flatness, more general version). — Suppose $(B, \mathfrak{n}) \rightarrow (A, \mathfrak{m})$ is a local morphism of Noetherian local rings (§7.3.1), and that M is a finitely generated A -module. Then M is B -flat if and only if $\mathrm{Tor}_1^B(M, B/\mathfrak{n}) = 0$.

25.6.3. ★ Proof of Theorem 25.6.2. A sign of the difficulty of this result is that the Artin-Rees Lemma 13.9.3 (or a consequence thereof) is used twice — once for the local ring (A, \mathfrak{m}) (in the guise of the Krull Intersection Theorem, Exercise 13.9.A), and once for the local ring (B, \mathfrak{n}) .

Recall that a B -module N has finite length if there is a finite sequence $0 = N_0 \subset N_1 \subset \cdots \subset N_n = N$ with $N_m/N_{m-1} \cong B/\mathfrak{n}$ for $1 \leq m \leq n$.

25.6.A. EASY PRELIMINARY EXERCISE. With the same hypotheses as Theorem 25.6.2, suppose that $\text{Tor}_1^B(M, B/\mathfrak{n}) = 0$. Show that $\text{Tor}_1^B(M, N) = 0$ for all B -modules N of finite length, by induction on the length of N .

By Exercise 24.1.D, if M is B -flat, then $\text{Tor}_1^B(M, B/\mathfrak{n}) = 0$, so it remains to assume that $\text{Tor}_1^B(M, B/\mathfrak{n}) = 0$ and show that M is B -flat.

By the ideal theoretic criterion for flatness (Theorem 25.4.1, see §25.4.2), we wish to show that $\phi: I \otimes_B M \rightarrow M$ is an injection for all ideals I of B , i.e., that $\ker \phi = 0$.

Note that $I \otimes_B M$ inherits an A -module structure (as M is an A -module). It is furthermore a *finitely generated* A -module (Do you see why?), so by the Krull Intersection Theorem (Exercise 13.9.A), $\cap_t \mathfrak{m}^t(I \otimes_B M) = 0$. Thus it suffices to show that $\ker \phi \subset \mathfrak{m}^t(I \otimes_B M)$ for all t .

As $\mathfrak{n} \subset \mathfrak{m}$ (or more correctly, the image of \mathfrak{n} is contained in \mathfrak{m}), it suffices to show that $\ker \phi \subset \mathfrak{n}^t(I \otimes_B M)$ for all t . Notice that $\mathfrak{n}^t(I \otimes_B M)$ is (the image in $I \otimes_B M$ of) $(\mathfrak{n}^t I) \otimes_B M$.

25.6.B. EXERCISE. Show that for each s , $\mathfrak{n}^t \cap I \subset \mathfrak{n}^s I$ for $t \gg 0$, so it suffices to show that $\ker \phi \subset (\mathfrak{n}^t \cap I) \otimes_B M$ for all t . Hint: Use the Artin-Rees Lemma 13.9.3, taking A there to be B here, $M_t = \mathfrak{n}^t$, $I = \mathfrak{n}$, and $L = I$.

Consider the short exact sequence

$$0 \longrightarrow \mathfrak{n}^t \cap I \longrightarrow I \longrightarrow I/(\mathfrak{n}^t \cap I) \longrightarrow 0.$$

Applying $(\cdot) \otimes_B M$, and using the fact that $I/(\mathfrak{n}^t \cap I)$ is finite length, we have that

$$0 \longrightarrow (\mathfrak{n}^t \cap I) \otimes_B M \longrightarrow I \otimes_B M \longrightarrow (I/(\mathfrak{n}^t \cap I)) \otimes_B M \longrightarrow 0$$

is exact using Exercise 25.6.A. Our goal is thus to show that $\ker \phi$ maps to 0 in

$$(25.6.3.1) \quad (I/(\mathfrak{n}^t \cap I)) \otimes_B M = ((I + \mathfrak{n}^t)/\mathfrak{n}^t) \otimes_B M.$$

Applying $(\cdot) \otimes_B M$ to the short exact sequence

$$(25.6.3.2) \quad 0 \longrightarrow (I + \mathfrak{n}^t)/\mathfrak{n}^t \longrightarrow B/\mathfrak{n}^t \longrightarrow B/(I + \mathfrak{n}^t) \longrightarrow 0,$$

and using Exercise 25.6.A (as $B/(I + \mathfrak{n}^t)$ is finite length), the top row of the diagram (25.6.3.3)

$$(25.6.3.3) \quad \begin{array}{ccccccc} 0 & \longrightarrow & ((I + \mathfrak{n}^t)/\mathfrak{n}^t) \otimes_B M & \xrightarrow{\alpha} & (B/\mathfrak{n}^t) \otimes_B M & \longrightarrow & (B/(I + \mathfrak{n}^t)) \otimes_B M \longrightarrow 0 \\ & & \uparrow & & \uparrow & & \\ & & I \otimes_B M & \xrightarrow{\phi} & B \otimes_B M & & \end{array}$$

is exact, and the square clearly commutes. But then any element of $I \otimes_B M$ mapping to 0 in $B \otimes_B M = M$ must map to 0 (under the right vertical arrow) in $(B/\mathfrak{n}^t) \otimes_B M$, and hence must have mapped to 0 in $((I + \mathfrak{n}^t)/\mathfrak{n}^t) \otimes_B M$ by the injectivity of α , as desired. \square

This argument basically shows that flatness is an “infinitesimal” property, depending only on the completion of the scheme at the point in question. This is made precise as follows.

Suppose $(B, \mathfrak{n}) \rightarrow (A, \mathfrak{m})$ is a (local) homomorphism of Noetherian local rings, and M is an A -module. If M is flat over B , then for each $t \in \mathbb{Z}^{\geq 0}$, $M/(\mathfrak{n}^t M)$ is

flat over B/n^t (flatness is preserved by base change, 25.2.I). (You should of course restate this in your mind in the language of schemes and quasicoherent sheaves.) The *infinitesimal criterion for flatness* states that this necessary criterion for flatness is actually sufficient.

25.6.C. ★ EXERCISE (THE INFINITESIMAL CRITERION FOR FLATNESS). Suppose $(B, n) \rightarrow (A, m)$ is a (local) homomorphism of local rings, and M is a finitely generated A -module. Suppose further that for each $t \in \mathbb{Z}^{\geq 0}$, $M/(n^t M)$ is flat over B/n^t . Show that M is flat over B . (In combination with Exercise 25.6.A, this gives another proof of the local criterion of flatness, Theorem 25.6.2.) Hint: follow the proof of Theorem 25.6.2. Given the hypothesis, then for each t , we wish to show that $\ker \phi$ maps to 0 in (25.6.3.1). We wish to apply $(\cdot) \otimes_B M$ to (25.6.3.2) and obtain the (exactness of the) top row of (25.6.3.3). To do this, show that applying $(\cdot) \otimes_B M$ to (25.6.3.2) is the same as applying $(\cdot) \otimes_{B/n^t} (M/n^t M)$. Then proceed as in the rest of the proof of Theorem 25.6.2.

Exercise 25.6.C implies the useful fact that if A is a Noetherian local ring, then its completion (at its maximal ideal) is flat over A .

25.6.4. The slicing criterion for flatness.

A useful variant of the local criterion is the following. Suppose t is a non-zerodivisor of B in n (geometrically: an effective Cartier divisor on the target passing through the closed point). If M is flat over B , then t is not a zerodivisor of M (Observation 25.2.2). Also, M/tM is a flat B/tB -module (flatness commutes with base change, Exercise 25.2.I). The next result says that this is a characterization of flatness, at least when M is finitely generated, or somewhat more generally.

25.6.5. Theorem (slicing criterion for flatness). — Suppose $(B, n) \rightarrow (A, m)$ is a local morphism of Noetherian local rings, M is a finitely generated A -module, and $t \in n$ is a non-zerodivisor on B . Then M is B -flat if and only if

- (i) t is not a zerodivisor on M , and
- (ii) M/tM is flat over $B/(t)$.

(For slicing criteria for other properties, see Exercise 13.2.C and Theorem 29.2.3.)

Proof. Assume that t is not a zerodivisor on M , and M/tM is flat over $B/(t)$. We will show that M is B -flat. (As stated at the start of §25.6.4, the other implication is a consequence of what we have already shown.)

By the local criterion for flatness, Theorem 25.6.2, we know

$$\mathrm{Tor}_1^{B/(t)}(M/tM, (B/(t))/n) = 0,$$

and we wish to show that $\mathrm{Tor}_1^B(M, B/n) = 0$. (Note that $(B/(t))/n = B/n$.) The result then follows from the following lemma. \square

25.6.6. Lemma. — Suppose M is a B -module, and $t \in B$ is not a zerodivisor on M . Then for any $B/(t)$ -module N , we have

$$(25.6.6.1) \quad \mathrm{Tor}_i^B(M, N) = \mathrm{Tor}_i^{B/(t)}(M/tM, N).$$

Proof. We calculate the left side of (25.6.6.1) by taking a free resolution of M :

$$(25.6.6.2) \quad \cdots \longrightarrow F_2 \longrightarrow F_1 \longrightarrow F_0 \longrightarrow M \longrightarrow 0.$$

By Exercise 25.3.A, $\mathrm{Tor}_i^B(M, B/(t)) = 0$ for $i > 0$ (here we use that t is not a zero-divisor on M , to show that $\mathrm{Tor}_1^B(M, B/(t)) = 0$). But this Tor module is computed by tensoring the free resolution (25.6.6.2) of M with $B/(t)$. Thus the complex

$$(25.6.6.3) \quad \cdots \longrightarrow F_2/tF_2 \longrightarrow F_1/tF_1 \longrightarrow F_0/tF_0 \longrightarrow M/tM \longrightarrow 0$$

is exact (exactness except at the last term comes from the vanishing of Tor_i). This is a free resolution of M/tM over the ring $B/(t)$! The left side of (25.6.6.1) is obtained by tensoring (25.6.6.2) by N and truncating and taking homology, and the right side is obtained by tensoring (25.6.6.3) by N and truncating and taking homology. As $(\cdot) \otimes_B N = (\cdot \otimes_B (B/t)) \otimes_{B/t} N$, we have established (25.6.6.1) as desired. \square

25.6.D. EXERCISE. Give a second (admittedly less direct) proof of the criterion for flatness over a discrete valuation ring of Exercise 25.4.J, using the slicing criterion for flatness (Theorem 25.6.5).

25.6.E. EXERCISE. Use the slicing criterion to give a second solution to Exercise 25.4.I, on two planes in \mathbb{P}^4 meeting at a point.

The following exercise gives a sort of slicing criterion for flatness in the *source*.

25.6.F. EXERCISE. Suppose A is a B -algebra, A and B are Noetherian, M is a finitely generated A -module, and $f \in A$ has the property that for all maximal ideals $\mathfrak{n} \subset B$, multiplication by f is injective on $M/\mathfrak{n}M$. Show that if M is B -flat, then M/fM is also B -flat. (Hint: Use the local criterion for flatness, Theorem 25.6.2. Notice that

$$0 \longrightarrow M \xrightarrow{\times f} M \longrightarrow M/fM \longrightarrow 0$$

is a flat resolution of M/fM .)

Exercise 25.6.F has an immediate geometric interpretation: “Suppose $\pi: X \rightarrow Y$ is a morphism of Noetherian schemes, \mathcal{F} is a coherent sheaf on X , and $Z \hookrightarrow X$ is a locally principal subscheme ...” In the special case where $\mathcal{F} = \mathcal{O}_X$, this leads to the notion of a **relative effective Cartier divisor**: a locally principal subscheme of X that is an effective Cartier divisor on all the fibers of π . Exercise 25.6.F implies that if π is flat, then any relative effective Cartier divisor is also flat. (See Exercise 28.3.F for an important application.)

25.6.7. Remark: *Local slicing criterion for flatness in the source, without Noetherian assumptions.*

The Noetherian hypotheses in Exercise 25.6.F can be removed. Suppose $(B, \mathfrak{n}) \rightarrow (A, \mathfrak{m})$ is a flat (local) homomorphism of local rings (not necessarily Noetherian). Suppose further that A is the localization of a finitely presented B -algebra. (This means that A is *essentially finitely presented*, or *essentially of finite presentation*, over B , but we won’t need this language.) If $f \in A$ is a non-zero-divisor in $B/\mathfrak{m}B$, then A/f is flat over B . See [Stacks, tag 046Z] for a proof.

25.6.8. ★★ Fibral flatness. We conclude by mentioning a criterion for flatness that is useful enough to be worth recognizing, but not so useful as to merit proof here.

25.6.G. EXERCISE. Suppose we have a commuting diagram

$$(25.6.8.1) \quad \begin{array}{ccc} X & \xrightarrow{\pi} & Y \\ & \searrow \tau & \swarrow \rho \\ & Z & \end{array}$$

and a quasicoherent sheaf \mathcal{F} on X , and points $p \in X$, $q = \pi(p) \in Y$, $r = \tau(p) \in Z$. Suppose ρ is flat at q , and \mathcal{F} is π -flat at p . Show that \mathcal{F} is τ -flat at p , and $\mathcal{F}|_r$ is π_r -flat at p . Here $\mathcal{F}|_r$ is the restriction of \mathcal{F} to the fiber above r , and $\pi_r: \tau^{-1}(r) \rightarrow \rho^{-1}(r)$ is the restriction of π above r .

The Fibrational Flatness Theorem states that in good circumstances the converse is true.

25.6.9. The Fibrational Flatness Theorem ([Gr-EGA, IV₃.11.3.10], see also [Stacks, tag 039A]). — Suppose we have a commuting diagram (25.6.8.1) and a finitely presented sheaf \mathcal{F} on X , and points $p \in X$, $q = \pi(p) \in Y$, $r = \tau(p) \in Z$, with $\mathcal{F}_p \neq 0$. Suppose either X and Y are locally Noetherian, or ρ and τ are locally of finite presentation. Then the following are equivalent.

- (a) \mathcal{F} is τ -flat at p , and $\mathcal{F}|_r$ is π_r -flat at p .
- (b) ρ is flat at q , and \mathcal{F} is π -flat at p .

This is a useful way of showing that a \mathcal{F} is π -flat. The architecture of the argument is as follows. First reduce to the case where X , Y , and Z are affine. Cleverly reduce to the Noetherian case (see [Gr-EGA, IV.11.2.7]), then solve the resulting nontrivial problem in commutative algebra (see [Gr-EGA, IV.11.3.10.1]).

25.7 Flatness implies constant Euler characteristic

We come to an important consequence of flatness promised in §25.1. We will see that this result implies many answers and examples to questions that we would have asked before we even knew about flatness.

25.7.1. Important Theorem (Euler characteristic is constant in flat families). — Suppose $\pi: X \rightarrow Y$ is a projective morphism of locally Noetherian schemes, and \mathcal{F} is a coherent sheaf on X , flat over Y . Then $\chi(X_q, \mathcal{F}|_{X_q}) = \sum_{i \geq 0} (-1)^i h^i(X_q, \mathcal{F}|_{X_q})$ is a locally constant function of $q \in Y$ (where $X_q = \pi^{-1}(q)$).

This is the first sign that “cohomology behaves well in flat families.” (We will soon see a second: the Semicontinuity Theorem 28.1.1. A different proof of Theorem 25.7.1, giving an extension to the proper case, will be given in §28.2.5.) The Noetherian hypotheses are used to ensure that $\pi_* \mathcal{F}(m)$ is a coherent sheaf.

Theorem 25.7.1 gives a necessary condition for flatness. Converses (yielding a sufficient condition) are given in Exercise 25.7.A.

Proof. We make three quick reductions. (i) The question is local on the target Y , so we may reduce to case Y is affine, say $Y = \text{Spec } B$, so π factors through a closed embedding $X \hookrightarrow \mathbb{P}_B^n$ for some n . (ii) We may reduce to the case $X = \mathbb{P}_B^n$, by considering \mathcal{F} as a sheaf on \mathbb{P}_B^n . (iii) We may reduce to showing that for $m \gg 0$,

$h^0(X_q, \mathcal{F}(m)|_{X_q})$ is a locally constant function of $q \in Y$ (by Serre vanishing for $m \gg 0$, Theorem 19.1.3(ii), h^0 agrees with the Euler characteristic).

Twist by $\mathcal{O}(m)$ for $m \gg 0$, so that all the higher cohomology of \mathcal{F} vanishes. Now consider the Čech complex \mathcal{C}^\bullet for $\mathcal{F}(m)$. Note that all the terms in the Čech complex \mathcal{C}^\bullet are flat, because \mathcal{F} is flat. (Do you see why?) As all higher cohomology groups (higher pushforwards) vanish, \mathcal{C}^\bullet is exact except at the first term, where the cohomology is $\Gamma(\pi_* \mathcal{F}(m))$. We add the module $\Gamma(\pi_* \mathcal{F}(m))$ to the front of the complex, so it is once again exact:

(25.7.1.1)

$$0 \longrightarrow \Gamma(\pi_* \mathcal{F}(m)) \longrightarrow \mathcal{C}^1 \longrightarrow \mathcal{C}^2 \longrightarrow \dots \longrightarrow \mathcal{C}^{n+1} \longrightarrow 0.$$

(We have done this trick of tacking on a module before, for example in (19.2.4.1).) Thus by Exercise 25.3.G, as we have an exact sequence in which all but the first terms are flat, the first term is flat as well. Thus $\pi_* \mathcal{F}(m)$ is a flat coherent sheaf on Y , and hence locally free (Corollary 25.4.7), and thus has locally constant rank.

Suppose $q \in Y$. We wish to show that the Hilbert function $h_{\mathcal{F}|_{X_q}}(m)$ is a locally constant function of q . To compute $h_{\mathcal{F}|_{X_q}}(m)$, we tensor the Čech resolution with $\kappa(q)$ and take cohomology. Now the extended Čech resolution (with $\Gamma(\pi_* \mathcal{F}(m))$ tacked on the front), (25.7.1.1), is an exact sequence of flat modules, and hence remains exact upon tensoring with $\kappa(q)$ (Exercise 25.3.F). Thus

$$\Gamma(\pi_* \mathcal{F}(m)) \otimes \kappa(q) \cong \Gamma(\pi_* \mathcal{F}(m)|_q),$$

so the Hilbert function $h_{\mathcal{F}|_{X_q}}(m)$ is the rank at q of a locally free sheaf, which is a locally constant function of $q \in Y$. \square

Before we get to the interesting consequences of Theorem 25.7.1, we mention some converses.

25.7.A. UNIMPORTANT EXERCISE (CONVERSES TO THEOREM 25.7.1). (We won't use this exercise for anything.)

(a) Suppose A is a ring, and S_\bullet is a finitely generated A -algebra that is flat over A . Show that $\text{Proj } S_\bullet$ is flat over A .

(b) Suppose $\pi: X \rightarrow Y$ is a projective morphism of locally Noetherian schemes (which as always includes the data of an invertible sheaf $\mathcal{O}_X(1)$ on X), such that $\pi_* \mathcal{O}_X(m)$ is locally free for all $m \geq m_0$ for some m_0 . Show that π is flat. Hint: describe X as

$$\text{Proj} \left(\mathcal{O}_Y \bigoplus (\oplus_{m \geq m_0} \pi_* \mathcal{O}_X(m)) \right).$$

(c) More generally, suppose $\pi: X \rightarrow Y$ is a projective morphism of locally Noetherian schemes, and \mathcal{F} is a coherent sheaf on X , such that $\pi_* \mathcal{F}(m)$ is locally free for all $m \geq m_0$ for some m_0 . Show that \mathcal{F} is flat over Y .

(d) Suppose $\pi: X \rightarrow Y$ is a projective morphism of locally Noetherian schemes, and \mathcal{F} is a coherent sheaf on X , such that $\sum (-1)^i h^i(X_q, \mathcal{F}|_q)$ is a locally constant function of $q \in Y$. If Y is reduced, show that \mathcal{F} must be flat over Y . (Hint: Exercise 14.3.K shows that constant rank implies local freeness in particularly nice circumstances.)

We now give some ridiculously useful consequences of Theorem 25.7.1.

25.7.2. Corollary. — Assume the same hypotheses and notation as in Theorem 25.7.1. Then the Hilbert polynomial of $\mathcal{F}|_{X_q}$ is locally constant as a function of $q \in Y$.

25.7.B. CRUCIAL EXERCISE. Suppose $X \rightarrow Y$ is a projective flat morphism, where Y is connected. Show that the following functions of $q \in Y$ are constant: (a) the degree of the fiber, (b) the dimension of the fiber, (c) the arithmetic genus of the fiber.

25.7.C. EXERCISE. Use §25.4.9 and Exercise 25.7.B(a) to give another solution to Exercise 18.4.D(a) (“the degree of a finite map from a curve to a regular curve is constant”).

Another consequence of Corollary 25.7.2 is something remarkably useful.

25.7.3. Corollary. — An invertible sheaf on a flat projective family of curves has locally constant degree on the fibers.

(Recall that the definition of the degree of a line bundle on a projective curve requires no hypotheses on the curve such as regularity, see (19.4.6.1).)

Proof. An invertible sheaf \mathcal{L} on a flat family of curves is always flat (as locally it is isomorphic to the structure sheaf). Hence $\chi(X_q, \mathcal{L}_q)$ is a constant function of q . By the definition of degree given in (19.4.6.1), $\deg(X_q, \mathcal{L}_q) = \chi(X_q, \mathcal{L}_q) - \chi(X_q, \mathcal{O}_{X_q})$. The result follows from the local constancy of $\chi(X_q, \mathcal{L}_q)$ and $\chi(X_q, \mathcal{O}_{X_q})$ (Theorem 25.7.1). \square

The following exercise is a serious generalization of Corollary 25.7.3.

25.7.D. ★ EXERCISE FOR THOSE WHO HAVE READ STARRED CHAPTER 21: INTERSECTION NUMBERS ARE LOCALLY CONSTANT IN FLAT FAMILIES. Suppose $\pi: X \rightarrow Y$ is a flat projective morphism to a connected scheme; $\mathcal{L}_1, \dots, \mathcal{L}_n$ are line bundles on X ; and \mathcal{F} is a coherent sheaf on X , flat over Y , such that the support of \mathcal{F} when restricted to any fiber of π has dimension at most n . If q is any point of Y , define (the temporary notation) $(\mathcal{L}_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F})_q$ to be the intersection on the fiber X_q of $\mathcal{L}_1, \dots, \mathcal{L}_n$ with $\mathcal{F}|_{X_q}$ (Definition 21.1.1). Show that $(\mathcal{L}_1 \cdot \mathcal{L}_2 \cdots \mathcal{L}_n \cdot \mathcal{F})_q$ is independent of q .

Corollary 25.7.3 motivates the following discussion.

25.7.4. Definition. Suppose \mathcal{L}_1 and \mathcal{L}_2 are line bundles on a k -variety X . We say that \mathcal{L}_1 and \mathcal{L}_2 are **algebraically equivalent** if there exists a connected (but not necessarily irreducible) k -variety Y with two k -valued points q_1 and q_2 , and a line bundle \mathcal{L} on $X \times Y$ such that the restriction of \mathcal{L} to the fibers over q_1 and q_2 are isomorphic to \mathcal{L}_1 and \mathcal{L}_2 respectively.

25.7.E. EXERCISE. Show that “algebraic equivalence” really is an equivalence relation. Show that the line bundles algebraically equivalent to \mathcal{O} form a subgroup of $\text{Pic } X$. This subgroup is denoted $\text{Pic}^0 X$.

25.7.5. The Néron-Severi group is the group of line bundles $\text{Pic } X$ modulo algebraic equivalence, $\text{Pic } X / \text{Pic}^0 X$. (This definition was promised in §19.4.9.) By Corollary 25.7.3, algebraic equivalence implies numerical equivalence (§19.4.9):

$\text{Pic}^0 X \subset \text{Pic}^\tau X$. (Side Remark: a line bundle on a proper k -scheme X is numerically trivial if and only if there exists an integer $m \neq 0$ with $L^{\otimes m}$ algebraically trivial. Thus $\text{Pic}^\tau X / \text{Pic}^0 X$ is torsion. See [SGA6, XIII, Thm. 4.6] for a proof, or [Laz, Cor. 1.4.38] for the projective case.)

25.7.6. ★ Hironaka's example of a proper nonprojective smooth threefold.

In §17.4.11, we produced a proper nonprojective variety, but it was singular. We can use Corollary 25.7.3 to give a *smooth* example, due to Hironaka.

Inside \mathbb{P}_k^3 , fix two conics C_1 and C_2 , which meet in two (k -valued) points, p_1 and p_2 . We construct a proper map $\pi: X \rightarrow \mathbb{P}_k^3$ as follows. Away from p_i , we blow up C_i and then blow up the proper transform of C_{3-i} (see Figure 25.3). This is well-defined, as away from p_1 and p_2 , C_1 and C_2 are disjoint, blowing up one and then the other is the same as blowing up their union, and thus the order doesn't matter.




FIGURE 25.3. Hironaka's example of a proper nonprojective smooth threefold

Note that π is proper, as it is proper away from p_1 , and proper away from p_2 , and the notion of properness is local on the target (Proposition 11.5.4(b)). As \mathbb{P}_k^3 is proper (over k , Theorem 8.4.7), and compositions of proper morphisms are proper (Proposition 11.5.4(c)), X is proper.

25.7.F. EXERCISE. Show that X is smooth. (Hint: Theorem 23.3.10.) Let E_i be the preimage of $C_i \setminus \{p_1, p_2\}$. Show that $\pi|_{E_i}: E_i \rightarrow C_i \setminus \{p_1, p_2\}$ is a \mathbb{P}^1 -bundle (and flat).

25.7.G. EXERCISE. Let \bar{E}_i be the closure of E_i in X . Show that $\bar{E}_i \rightarrow C_i$ is flat. (Hint: Exercise 25.4.J.)

25.7.H. EXERCISE. Show that $\pi^{-1}(p_i)$ is the union of two \mathbb{P}^1 's, say Y_i and Z_i , meeting at a point, such that $Y_i, Y_{3-i}, Z_{3-i} \in \bar{E}_i$ but $Z_i \notin \bar{E}_i$.

25.7.I. EXERCISE. Show that X is not projective as follows. Suppose otherwise \mathcal{L} is a very ample line bundle on X , so \mathcal{L} has positive degree on every curve (including the Y_i and Z_i). Using flatness of $\bar{E}_i \rightarrow C_i$, and constancy of degree in flat families (Exercise 25.7.D), show that $\deg_{Y_i} \mathcal{L} = \deg_{Y_{3-i}} \mathcal{L} + \deg_{Z_{3-i}} \mathcal{L}$. Obtain a contradiction. (This argument will remind you of the argument of §17.4.11.)

25.7.7. *The notion of “projective morphism” is not local on the target.* Note that $\pi: X \rightarrow \mathbb{P}^3$ is not projective, as otherwise X would be projective (as the composition of projective morphisms is projective if the final target is quasicompact, Exercise 18.3.B). But away from each p_i , π is projective (as it is a composition of blow-ups, which are projective by construction, and the final target is quasicompact, so Exercise 18.3.B applies). Thus the notion of “projective morphism” is not local on the target.

Smooth and étale morphisms, and flatness

We have defined smooth and étale morphisms earlier (Definition 13.6.2), but we are now in a position to understand them much better. We will see that the notion of unramified morphism (§22.6) is a natural companion to them.

Our three algebro-geometric definitions won't be so obviously a natural triplet, but we will discuss the definitions given in [Gr-EGA] (§26.2.6), and in this context the three types of morphisms look very similar. (We briefly mention other approaches and definitions in §26.2.7.)

26.1 Some motivation

The three classes of morphisms we will discuss in this chapter are the analogs of the following types of maps of manifolds, in differential geometry.

- *Submersions* are maps inducing surjections of tangent spaces everywhere. They are useful in the notion of a fibration. (Perhaps a more relevant notion from differential geometry, allowing singularities, is: “locally on the source a smooth fibration”.)
- *Isomorphisms locally on the source* (or *local isomorphisms*) are maps inducing isomorphisms of tangent spaces.
- *Immersions* are maps inducing injections of tangent spaces.

(Recall our warning from §9.2: “immersion” is often used in algebraic geometry with a different meaning.)

[picture to be made later]

FIGURE 26.1. Sketches of notions from differential geometry: (a) maps “locally on the source a smooth fibration”, (b) local isomorphisms, and (c) immersions. (In algebraic geometry: (a) smooth morphisms, (b) étale morphisms, and (c) unramified morphisms.)

In order to better understand smooth and étale morphisms, we temporarily forget our earlier definitions, and consider some examples of things we want to be analogs of “local isomorphism” (or “locally on the source an isomorphism”) and “locally on the source a smooth fibration”, and see if they help us make good definitions.

26.1.1. “Local isomorphisms” (étale morphisms). Consider the parabola $x = y^2$ projecting to the x -axis, over the complex numbers. (This example has come

up repeatedly, in one form or another, including in Exercise 13.6.E.) We might reasonably want this to be a local isomorphism (on the source) away from the origin. We might also want the notion of local isomorphism space to be an open condition: the locus where a morphism is a local isomorphism should be open on the source. This is true for the differential geometric definition. But then this morphism should be a local isomorphism over the generic point, and here we get a non-trivial residue field extension $(\mathbb{C}(y)/\mathbb{C}(y^2))$, not an isomorphism. Thus we are forced to consider (the Spec's of) certain finite extensions of fields to be "isomorphisms locally on the source" — in this case, even "covering spaces". (We will see in Exercise 26.2.C that we want precisely the finite separable extensions; you could have shown this earlier.)

Note also in this example there are no (nonempty) Zariski-open subsets $U \subset \operatorname{Spec} \mathbb{C}[x] \setminus \{0\}$ and $V \subset \operatorname{Spec} \mathbb{C}[x, y]/(x - y^2) \setminus \{(0, 0)\}$ where the map sends U into V isomorphically (Exercise 13.6.E), so this is not a local isomorphism in a way you may have seen before. This leads to the notion of the étale topology, which is not even a topology in the usual sense, but a "Grothendieck topology" (§6.2.9). The étale topology is beyond the scope of this book.

26.1.2. Submersions (smooth morphisms).

26.1.3. Fibers are smooth varieties. As a first approximation of the algebro-geometric version of submersion, we will want the fibers to be smooth varieties (over the residue field). So the very first thing we need is to generalize the notion of "variety" over a base. It is reasonable to do this by having a locally finite type hypothesis. For somewhat subtle reasons, we will require the stronger condition that the morphism be locally of finite presentation. (If you really care, you can see where it comes up in our discussion. But of course there is no difference for Noetherian readers.)

The fibers are not just varieties; they should be *smooth* varieties (of dimension n , say). From the case of smoothness over a field, or from our intuition of what smooth varieties should look like, we expect the sheaf of differentials on the fibers to be locally free of rank n , or even better, that the sheaf of relative differentials be locally free of rank n .

26.1.4. Flatness. At this point, our first approximation of "smooth morphism" is some version of "locally finitely presented, and fibers are smooth varieties". But that isn't quite enough. For example, a horrible map from a scheme X to a curve Y that maps a different regular variety to each point of Y (X is the infinite disjoint union of these) should not be considered a smooth fibration in any reasonable sense. Also, we might not want to consider $\operatorname{Spec} k \rightarrow \operatorname{Spec} k[\epsilon]/(\epsilon^2)$ to be a submersion; for example, this isn't surjective on tangent spaces, and more generally the picture "doesn't look like a fibration".

Both problems are failures of $\pi: X \rightarrow Y$ to be a nice, "continuous" family. Whenever we are looking for some vague notion of "niceness" we know that "flatness" will be in the definition.

For comparison, note that "unramified" has no flatness hypothesis, and indeed we didn't expect it, as we would want the inclusion (closed embedding) of the origin into \mathbb{A}^1 to be unramified. But then weird things may be unramified.




FIGURE 26.2. We don't want this to be a "smooth morphism"

For example, if $X = \coprod_{z \in \mathbb{C}} \operatorname{Spec} \mathbb{C}$, then the morphism $X \rightarrow \mathbb{A}_{\mathbb{C}}^1$ sending the point corresponding to z to the point $z \in \mathbb{A}_{\mathbb{C}}^1$ is unramified. Such is life.

26.1.5. Desired Alternate Definitions. We might hope that a morphism $\pi: X \rightarrow Y$ is **smooth of relative dimension n** if and only if

- (i) π is locally of finite presentation,
- (ii) π is flat of relative dimension n , and
- (iii) $\Omega_{X/Y}$ is locally free of rank n .

We might similarly hope that a morphism $\pi: X \rightarrow Y$ is **étale** if and only if

- (i) π is locally of finite presentation,
- (ii) π is flat, and
- (iii) $\Omega_{X/Y} = 0$.

We will shortly show (in Theorem 26.2.2 and Exercise 26.2.B respectively) that these desired definitions are equivalent to Definition 13.6.2.

26.2 Different characterizations of smooth and étale morphisms

The main result in this section is a description of equivalent characterizations of smooth morphisms, Theorem 26.2.2. We will state the theorem, and then give consequences, and then finally give a proof.

26.2.1. But first, we discuss a preliminary fact. It requires something from §29.2.7: if k is a field, and $f_1, \dots, f_i \in k[x_1, \dots, x_N]$ cut out a dimension $N - i$ subscheme X of \mathbb{A}_k^N , and $f_{i+1} \in k[x_1, \dots, x_N]$ does not vanish on any irreducible component of X , then f_{i+1} is not a zerodivisor on X . For now, take this fact for granted; you shouldn't read all of Chapter 29 in order to do this exercise. There are other ways around this issue, but it is not worth the effort.

26.2.A. EXERCISE. Suppose $\pi: X \rightarrow Y$ is smooth of relative dimension n . Either (i) assume Y (and hence X) is locally Noetherian, or (ii) accept Remark 25.6.7 (the slicing criterion for flatness in the source without Noetherian hypotheses, which had a

reference rather than a proof). Show that π is flat. Hint: suppose π is the morphism $\text{Spec } B[x_1, \dots, x_{n+r}]/(f_1, \dots, f_r) \rightarrow \text{Spec } B$. Consider $B[x_1, \dots, x_{n+r}]/(f_1, \dots, f_i)$ as i runs from 0 to r , and use induction on i and the slicing criterion for flatness in the source ((i) Exercise 25.6.F or (ii) Remark 25.6.7).

We now come to the central result of this section.

26.2.2. Theorem. — Suppose $\pi: X \rightarrow Y$ is a morphism of schemes. Then the following are equivalent.

- (i) The morphism π is smooth of relative dimension n (Definition 13.6.2).
- (ii) The morphism π is locally finitely presented and flat of relative dimension n ; and Ω_π is locally free of rank n (Desired Alternate Definition 26.1.5).
- (iii) The morphism π is locally finitely presented and flat, and the fibers are smooth k -schemes of pure dimension n .
- (iv) The morphism π is locally finitely presented and flat, and the geometric fibers are smooth k -schemes of pure dimension n .

Before proving Theorem 26.2.2, we motivate it, by giving a number of exercises and results assuming it.

26.2.B. EXERCISE. Show that a morphism π is étale (Definition 13.6.2) if and only if π is locally of finite presentation, flat, and $\Omega_\pi = 0$ (Desired Alternate Definition 26.1.5).

26.2.C. EASY EXERCISE. Suppose $\pi: X \rightarrow Y$ is a morphism. Show that the following are equivalent.

- (a) π is étale.
- (b) π is smooth and unramified.
- (c) π is locally finitely presented, flat, and unramified.
- (d) π is locally finitely presented, flat, and for each $q \in Y$, $\pi^{-1}(q)$ is the disjoint union of schemes of the form $\text{Spec } K$, where K is a finite separable extension of $\kappa(q)$.
- (e) π is locally finitely presented, flat, and for each geometric point \bar{q} of Y , $\pi^{-1}(\bar{q})$ is the (scheme-theoretic) disjoint union of copies of \bar{q} .

26.2.D. EXERCISE. If $\pi: X \rightarrow Y$ is étale, show that any preimage $p \in X$ of any regular point $q \in Y$ whose local ring has dimension n is also a regular point whose local ring has dimension n . Hint: Prove the result by induction on $\dim \mathcal{O}_{Y,q}$. “Slice” by an element of $\mathfrak{m}_{Y,q} \setminus \mathfrak{m}_{Y,q}^2$. Use the slicing criterion for regularity (Exercise 13.2.C).

26.2.3. Proof of the Smoothness-Regularity Comparison Theorem 13.2.10(b) (every smooth k -scheme is regular). By Observation 13.6.4, any dimension n smooth k -scheme X can locally be expressed as an étale cover of \mathbb{A}_k^n , which is regular by Exercise 13.3.O. Then by Exercise 26.2.D, X is regular. \square

26.2.E. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a morphism of equidimensional k -varieties, and Y is smooth. Use the cotangent exact sequence (Theorem 22.2.25) to show the following.

- (a) Suppose that $\dim X = \dim Y = n$, and π is unramified. Show that X is smooth.

★ (b) Suppose $\dim X = m > \dim Y = n$, and the (scheme-theoretic) fibers of π over closed points are smooth of dimension $m - n$. Show that X is smooth.

26.2.F. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a morphism (over k) of smooth equidimensional k -varieties, where $\dim X = m$ and $\dim Y = n$. (These two parts are clearly essentially the same geometric situation.)

(a) Show that if the fibers of π are all smooth of dimension $m - n$, then π is smooth of relative dimension $m - n$.

(b) Show that if $\pi^* \Omega_Y \rightarrow \Omega_X$ is injective (“ T_π is surjective”, or “the cotangent exact sequence is exact on the left”), then π is smooth of relative dimension $m - n$.

Hint: for each point $q \in Y$, $\mathcal{O}_{Y,q}$ is a regular local ring. Work by induction on $\dim \mathcal{O}_{Y,q}$, and “slice” by an element of $\mathfrak{m}_{Y,q} \setminus \mathfrak{m}_{Y,q}^2$.

26.2.G. EXERCISE. Suppose $\pi: X \rightarrow Y$ is locally finitely presented, $\rho: Y \rightarrow Z$ is étale, and $\tau = \rho \circ \pi$. Show that π is smooth of dimension n (e.g., étale, taking $n = 0$) if and only if τ is. (We showed earlier that π is unramified if and only if τ is. Do you see why and where?) Hint: use the Cancellation Theorem 11.2.1 for flat morphisms and Exercise 22.6.F. For the Ω part of the problem, use the cotangent exact sequence, Theorem 22.2.25. (If you only solved Exercise 22.6.F(a), then you will still be able to use it to prove this in the Noetherian case.)

26.2.4. ★★ Proof of Theorem 26.2.2.

(i) *implies* (ii). As noted in §13.6.2, local finite presentation is clear. Flatness follows from Exercise 26.2.A. Finally, the sheaf of relative differentials is locally free of rank n by Exercise 22.2.Q.

(ii) *implies* (iii). Suppose π satisfies (ii). We wish to show that the fibers are smooth. The hypotheses of (ii) are preserved by base change, so we may assume that $Y = \text{Spec } k$. Then X is indeed smooth (see the equivalent Definition 22.3.1).

(iii) *implies* (iv) by Exercise 13.2.G.

We now come to the main part of the argument: (iv) *implies* (i). (Part of this argument will be reminiscent of the proof of Theorem 13.6.5, in which we also had to prove smoothness.)

Fix a closed point $p \in X$. Let $q = \pi(p) \in Y$, say $q = \text{Spec } k$. We will show that there is an open neighborhood U_i of p and V_i of q of the form stated in Definition 13.6.2. It suffices to deal with the case that p is a closed point in $\pi^{-1}(q)$, because closed points are dense in finite type schemes (Exercise 5.3.F). (We cannot and do not assume that q is a closed point of Y .)

Because π is finitely presented, there are affine open neighborhoods U of p and V of q , with $\pi(U) \subset V$, such that the morphism $\pi|_U: U \rightarrow V$ can be written as

$$\text{Spec } B[x_1, \dots, x_N]/I \longrightarrow \text{Spec } B,$$

where I is finitely generated.

Fix a geometric point $\bar{q} = \text{Spec } \bar{k}$ of Y mapping to q . Fix a point \bar{p} of the geometric fiber over \bar{q} , which maps to $p \in X$. (Why is there such a \bar{p} ?) We then

have a diagram

$$\begin{array}{ccccc}
 U_{\bar{k}} & \xrightarrow{\quad} & U_k & \xrightarrow{\quad} & U \\
 \parallel & & \parallel & & \parallel \\
 \operatorname{Spec} \bar{k}[x_1, \dots, x_N]/I_{\bar{k}} & \longrightarrow & \operatorname{Spec} k[x_1, \dots, x_N]/I_k & \hookrightarrow & \operatorname{Spec} B[x_1, \dots, x_N]/I \\
 \downarrow & & \downarrow & & \downarrow \pi|_U \\
 \operatorname{Spec} \bar{k} & \longrightarrow & \operatorname{Spec} k & \longrightarrow & \operatorname{Spec} B \\
 \parallel & & \parallel & & \parallel \\
 \bar{q} & \longrightarrow & q & \longrightarrow & V
 \end{array}$$

where $U_{\bar{k}}$ (resp. U_k) are the fibers of U over \bar{q} (resp. q), and $I_{\bar{k}}$ (resp. I_k) is the ideal cutting out $U_{\bar{k}}$ (resp. U_k) in $\mathbb{A}_{\bar{k}}^N$ (resp. \mathbb{A}_k^N).

We make an important observation, using §25.3.2 (for a flat closed subscheme, the pullback of the defining ideal as a quasicoherent sheaf is the same as the ideal of the pulled back closed subscheme). Because $B[x_1, \dots, x_N]/I$ is flat over B , we have that $I_{\bar{k}} = I \otimes_B \bar{k}$ and $I_k = I \otimes_B k$.

By our hypothesis (iv), the geometric fiber $U_{\bar{k}} \subset \mathbb{A}_{\bar{k}}^N$ is smooth, and hence regular at its closed points by Exercise 13.2.D. (Here we use the fact that \bar{k} is algebraically closed.)

Let $r = N - n$ be the codimension of U_k in \mathbb{A}_k^N (and the codimension of $U_{\bar{k}}$ in $\mathbb{A}_{\bar{k}}^N$). By the Jacobian criterion for smoothness over a field (§22.3.2), the cotangent space for $U_{\bar{k}}$ at \bar{p} is cut out in $\Omega_{\mathbb{A}_{\bar{k}}^N|\bar{k}}|_{\bar{p}}$ by r linear equations (over \bar{k}). By the conormal exact sequence for $U_{\bar{k}} \subset \mathbb{A}_{\bar{k}}^N$ (Theorem 22.2.25(b)), these equations are spanned by the elements of $I_{\bar{k}}$, which in turn are spanned by elements of I (as $I_{\bar{k}} = I \otimes_B \bar{k}$, so elements of $I_{\bar{k}}$ are finite sums of pure tensors $i \otimes a$ with $i \in I$ and $a \in \bar{k}$). Thus we can choose f_1, \dots, f_r in I so that the images of f_1, \dots, f_r in $\bar{k}[x_1, \dots, x_N]$ cut out the Zariski tangent space of $U_{\bar{k}}$ at \bar{p} (as the Zariski tangent space at a closed point of a variety over an algebraically closed field is computed by the Jacobian matrix, Exercise 13.2.D).

26.2.H. EXERCISE. Show that f_1, \dots, f_r generate the ideal $(I_{\bar{k}})_{\bar{p}}$. Hint: look at your argument for Exercise 13.2.L(a).

Rearrange the x_i 's so that the Jacobian matrix of the f_i with respect to the *first* r of the x_i is invertible at \bar{p} , and hence at p , and hence in an open neighborhood of p . Then (an open neighborhood of p in) $\operatorname{Spec} B[x_1, \dots, x_N]/(f_1, \dots, f_r)$ is of the form we are looking for. We will show that U is the same as $\operatorname{Spec} B[x_1, \dots, x_N]/(f_1, \dots, f_r)$ (i.e., $I = (f_1, \dots, f_r)$) near p , thereby completing the proof.

For notational compactness (at the cost of having subscripts later on with confusingly different meanings), define

$$\begin{aligned}
 A &:= B[x_1, \dots, x_N]/(f_1, \dots, f_r), \\
 A_k &:= A \otimes_B k = k[x_1, \dots, x_N]/(f_1, \dots, f_r), \\
 A_{\bar{k}} &:= A \otimes_B \bar{k} = \bar{k}[x_1, \dots, x_N]/(f_1, \dots, f_r).
 \end{aligned}$$

We have the diagram

$$\begin{array}{ccc} A/J & \xleftarrow{\quad} & A \\ & \nwarrow \quad \nearrow & \\ & B & \end{array}$$

where J denotes the image of I in A . We wish to show that “ J is 0 near p ”.

Let \mathfrak{m} be the maximal ideal of A corresponding to p , and let $\bar{\mathfrak{m}}$ be the maximal ideal of $A_{\bar{k}}$ corresponding to \bar{p} .

Let $J_{\bar{k}}$ be the image of $I_{\bar{k}}$ in $A_{\bar{k}}$. We use §25.3.2 a second time in this proof. Because $A/J = B[x_1, \dots, x_N]/I$ is flat over B , $J_{\bar{k}}$ can be interpreted both as the pullback of J as a quasicoherent sheaf, or as the ideal sheaf of the “pulled back closed subscheme corresponding to J ”. More precisely,

$$(26.2.4.1) \quad J \otimes_B \bar{k} = J_{\bar{k}}.$$

Similarly,

$$(26.2.4.2) \quad J \otimes_B k = J_k.$$

Now Exercise 26.2.H means that

$$J \cdot (A_{\bar{k}})_{\bar{\mathfrak{m}}} = (0),$$

so taking the quotient by the maximal ideal $\bar{\mathfrak{m}}$ of $(A_{\bar{k}})_{\bar{\mathfrak{m}}}$ yields

$$(J \cdot A_{\bar{k}}) \otimes_{A_{\bar{k}}} (A_{\bar{k}}/\bar{\mathfrak{m}}) = (0),$$

which by (26.2.4.1) can be rewritten as

$$(J \otimes_B \bar{k}) \otimes_{A \otimes_B \bar{k}} (A_{\bar{k}}/\bar{\mathfrak{m}}) = (0),$$

which in turn can be readily rewritten as

$$(J \otimes_B k) \otimes_{A \otimes_B k} (A_{\bar{k}}/\bar{\mathfrak{m}}) = (0)$$

and then as

$$((J \otimes_B k) \otimes_{A \otimes_B k} (A_k/\mathfrak{m})) \otimes_{A_k/\mathfrak{m}} (A_{\bar{k}}/\bar{\mathfrak{m}}) = (0)$$

But $A_{\bar{k}}/\bar{\mathfrak{m}} = \bar{k}$ is a field extension of $A_k/\mathfrak{m} = k$, and a k -vector space V is 0 if (and only if) $V \otimes_k \bar{k} = 0$, so we have shown that

$$((J \otimes_B k) \otimes_{A \otimes_B k} (A_k/\mathfrak{m})) = (0).$$

Using (26.2.4.2) (essentially reversing our steps, with \bar{k} replaced by k), we have

$$(J \cdot A_k) \otimes_{A_k} (A_k/\mathfrak{m}) = 0,$$

which means that $J \otimes_A (A/\mathfrak{m}) = 0$.

26.2.I. EXERCISE. Show that $J_{\mathfrak{m}} = 0$, and from this conclude that there is an element $f \in A \setminus \mathfrak{m}$ such that $fJ = 0$. (Hint: Nakayama, and finite generation of J .)

Then $J = 0$ in the open neighborhood $D(f) \subset \text{Spec } A$ of p , implying that (i) holds. \square

This proof can be made notably easier by working in characteristic 0, where we can take (iii) rather than (iv) as our starting point to prove (i), using the fact that smooth schemes over perfect fields k are regular at closed points (see Exercise 22.3.D).

26.2.5. **Formally unramified, smooth, and étale.**

[Gr-EGA] takes a different starting point for the definition of unramified, smooth, and étale. The definitions there make clear that these three definitions form a family.

The cost of these definitions are that they are perhaps less immediately motivated by geometry, and it is harder to show some basic properties. The benefit is that it is possible to show more (for example, left-exactness of the cotangent and conormal exact sequences, and good interpretations in terms of completions of local rings). But we simply introduce these ideas here, and do not explore them. See [BLR, §2.2] for an excellent discussion. (You should largely ignore what follows, unless you find later in life that you really care.)

26.2.6. Definition. We say that $\pi: X \rightarrow Y$ is **formally smooth** (resp. **formally étale**, **formally unramified**) if for all affine schemes Z , and every closed subscheme $Z_0 \subset Z$ defined by a nilpotent ideal, and every morphism $Z \rightarrow Y$, the canonical map $\mathrm{Hom}_Y(Z, X) \rightarrow \mathrm{Hom}_Y(Z_0, X)$ is surjective (resp. bijective, injective). This is summarized in the following diagram, which is reminiscent of the valuative criteria for separatedness and properness.

$$\begin{array}{ccc}
 Z_0 & \longrightarrow & X \\
 \downarrow \text{nilpotent ideal} & \nearrow ? & \downarrow \pi \\
 Z & \longrightarrow & Y
 \end{array}$$

(You can check that this is the same as the definition we would get by replacing “nilpotent” by “square-zero”. This is sometimes an easier formulation to work with.)

Then [Gr-EGA] defines smooth as morphisms that are formally smooth and locally of finite presentation, and similarly for unramified and étale.

One can show that [Gr-EGA]’s definitions of formally unramified, and smooth agree with the definitions we give. For “formally unramified” (where our definition given in §22.6 is not obviously the same as the definition of [Gr-EGA] given here), see [Gr-EGA, IV₄.17.2.1] or [Stacks, tag 00UO]. For “smooth”, see [Gr-EGA, IV₄.17.5.1] or [Stacks, tag 00TN]. (Our characterization of étale as smooth of relative dimension 0 then agrees with [Gr-EGA]. Our definition of unramified as formally unramified plus locally of finite type disagrees with [Gr-EGA], as mentioned in §22.6.)

26.2.7. Other starting points. Unlike many other definitions in algebraic geometry, there are a number of quite different ways of defining smooth, étale, and unramified morphisms, and it is nontrivial to relate them. We have just described the approach of [Gr-EGA]. Another common approach is the characterization of smooth morphisms as locally finitely presented, flat, and with regular geometric fibers (see Theorem 26.2.2). Yet another definition is via a naive version of the cotangent complex; this is the approach taken by [Stacks, tag 00T2], and is less frightening than it sounds. Finally, the different characterizations of Exercise 26.2.C give a number of alternate initial definitions of étaleness.

26.3 Generic smoothness and the Kleiman-Bertini Theorem

We will now discuss a number of important results that fall under the rubric of “generic smoothness”, an idea we first met in §22.3.4. All will require working over a field of characteristic 0 in an essential way.

26.3.1. Theorem (generic smoothness on the source). — *Let k be a field of characteristic 0, and let $\pi: X \rightarrow Y$ be a dominant morphism of integral k -varieties. Then there is a nonempty (=dense) open set $U \subset X$ such that $\pi|_U$ is smooth of dimension $\dim X - \dim Y$.*

The key idea already appeared in Theorem 22.3.5, when we showed that every variety has an open subset that is regular.

Proof. Define $n = \dim X - \dim Y$ (the “relative dimension”). Now $K(X)/K(Y)$ is a finitely generated field extension of transcendence degree n (from transcendence theory, see for example Exercise 12.2.A), so $\Omega_{X/Y}$ has rank n at the generic point, by Exercise 22.2.M(b). (Here we use the hypothesis $\text{char } k = 0$, as we are using the fact that $K(X)/K(Y)$ is separably generated by *any* transcendence basis.) By upper semicontinuity of fiber rank of a coherent sheaf (Exercise 14.3.J), it is rank n for every point in a dense open set. On a reduced scheme, constant rank implies locally free of that rank (Exercise 14.3.K), so $\Omega_{X/Y}$ is locally free of rank n on that set. Also, by generic flatness (Exercise 25.5.R), it is flat on a dense open set. Let U be the intersection of these two open sets. \square

26.3.2. Example. An examination of the proof yields an example showing that this result fails in positive characteristic: consider the purely inseparable extension $\mathbb{F}_p(t)/\mathbb{F}_p(t^p)$. The same problem can arise even over an *algebraically closed* field of characteristic p : consider $\mathbb{A}_k^1 = \text{Spec } k[t] \rightarrow \text{Spec } k[u] = \mathbb{A}_k^1$, given by $u \mapsto t^p$.

If furthermore X is smooth, the situation is even better.

26.3.3. Theorem (generic smoothness on the target). — *Suppose $\pi: X \rightarrow Y$ is a morphism of k -varieties, where $\text{char } k = 0$, and X is smooth (over k). Then there is a dense open subset U of Y such that $\pi|_{\pi^{-1}(U)}$ is a smooth morphism.*

Note that $\pi^{-1}(U)$ may be empty! Indeed, if π is not dominant, we will have to take such a U .

To prove Theorem 26.3.3, we use a neat trick. Suppose $\pi: X \rightarrow Y$ is a morphism of schemes that are finite type over k , where $\text{char } k = 0$. Define

$$X_r = \left\{ p \in X : \text{rank} \left(\pi^* \Omega_{Y/k} \rightarrow \Omega_{X/k} \right) \Big|_p \leq r \right\}.$$

26.3.A. EXERCISE. Show that X_r is a closed subset of X . Hint: If X and Y are both smooth, show that the rank condition implies that X_r is cut out by “determinantal equations”, or use an appropriate form of uppersemicontinuity.

By Chevalley’s Theorem 8.4.2, $\pi(X_r)$ is constructible, so we can make sense of $\dim \pi(X_r)$.

26.3.4. Lemma. — *If $\pi: X \rightarrow Y$ is a morphism of schemes that are finite type over k , where $\text{char } k = 0$, then $\dim \pi(X_r) \leq r$.*

Proof. In this proof, we make repeated use of the identification of Zariski cotangent spaces at closed points with the fibers of cotangent sheaves, using Corollary 22.3.9. We can replace X by an irreducible component of X_r , and Y by the closure of that component's image in Y (with reduced subscheme structure, see §9.4.9). The resulting map will have all of X contained in X_r , so we may as well assume $X = X_r$. This boils down to the following linear algebra observation: if a linear map $\rho: V_1 \rightarrow V_2$ has rank at most r , and V'_i is a subspace of V_i , with ρ sending V'_1 to V'_2 , then the restriction of ρ to V'_1 has rank at most that of ρ itself.) Thus we have a dominant morphism $\pi: X \rightarrow Y$, and we wish to show that $\dim Y \leq r$. Using generic smoothness on the source (Theorem 26.3.1) for $Y \rightarrow \text{Spec } k$, we can shrink Y further so as to assume that it is smooth. By generic smoothness on the source for $\pi: X \rightarrow Y$, there is a nonempty open subset $U \subset X$ such that $\pi: U \rightarrow Y$ is smooth. But then for any point $p \in U$, the cotangent map $\pi^* \Omega_Y|_{\pi(p)} \rightarrow \Omega_X|_p$ is injective (Exercise 22.2.S), and has rank at most r . Taking p to be a closed point, we have $\dim Y = \dim_{\pi(p)} Y \leq \dim \Omega_Y|_{\pi(p)} \leq r$. \square

There is not much left to do to prove the theorem.

26.3.5. Proof of Theorem 26.3.3. Reduce to the case Y smooth over k (by restricting to a smaller open set, using generic smoothness of Y , Theorem 26.3.1). Say $n = \dim Y$. Now $\dim \pi(X_{n-1}) \leq n - 1$ by Lemma 26.3.4, so remove $\pi(X_{n-1})$ from Y as well. Then the rank of Ω_π is at least n for each closed point of X . But as Y is regular of dimension n , we have that $\Omega_\pi = (\pi^* \Omega_{Y/k} \rightarrow \Omega_{X/k})$ is injective for every closed point of X , and hence $\pi^* \Omega_{Y/k} \rightarrow \Omega_{X/k}$ is an injective map of sheaves (Do you see why?). Thus π is smooth by Exercise 26.2.F(b). \square

26.3.6. ★★ The Kleiman-Bertini Theorem. The same idea of bounding the dimension of some “bad locus” can be used to prove the Kleiman-Bertini Theorem 26.3.7 (due to Kleiman), which is useful in (for example) enumerative geometry. Throughout this discussion $k = \bar{k}$, although the definitions and results can be generalized. Suppose G is a group variety (over $k = \bar{k}$), and we have a G -action on a variety X . We say that the **action is transitive** if it is transitive on closed points. A better definition (that you can show is equivalent) is that the action morphism $G \times X \rightarrow X$ restricted to a fiber above any closed point of X is surjective. A variety X with a transitive G -action is said to be a **homogeneous space** for G . For example, G acts on itself by left-translation, and via this action G is a homogeneous space for itself.

26.3.B. EASY EXERCISE. Suppose G is a group variety over an algebraically closed field k of characteristic 0. Show that every homogeneous space X for G is smooth. (In particular, taking $X = G$, we see that G is smooth.) Hint: X has a dense open set U that is smooth by Theorem 26.3.1, and G acts transitively on the closed points of X , so we can cover X with translates of U .

26.3.7. The Kleiman-Bertini Theorem, [K12, Thm. 2]. — *Suppose X is homogeneous space for a group variety G (over a field $k = \bar{k}$ of characteristic 0). Suppose $\alpha: Y \rightarrow X$ and $\beta: Z \rightarrow X$ are morphisms from smooth k -varieties Y and Z .*

(a) Then there is a nonempty open subset $V \subset G$ such that for every $\sigma \in V(k)$, $Y \times_X Z$ defined by

$$\begin{array}{ccc} Y \times_X Z & \longrightarrow & Z \\ \downarrow & & \downarrow \beta \\ Y & \xrightarrow{\sigma \circ \alpha} & X \end{array}$$

(Y is “translated by σ ”) is smooth of dimension $\dim Y + \dim Z - \dim X$ (but possibly empty).

(b) Furthermore, there is a nonempty open subset $V \subset G$ such that

$$(26.3.7.1) \quad (G \times_k Y) \times_X Z \rightarrow G$$

is a smooth morphism of relative dimension $\dim Y + \dim Z - \dim X$ over V .

The first time you hear this, you should think of the special case where $Y \rightarrow X$ and $Z \rightarrow X$ are locally closed embeddings (Y and Z are smooth subvarieties of X). In this case, the Kleiman-Bertini theorem says that the second subvariety will meet a “general translate” of the first “transversely”.

Proof. It is more pleasant to describe this proof “backwards”, by considering how we would prove it ourselves. We will use generic smoothness twice.

Clearly (b) implies (a), so we prove (b).

In order to show that the morphism (26.3.7.1) is smooth over a nonempty open set $V \subset G$, it would suffice to apply generic smoothness on the target (Theorem 26.3.3) to (26.3.7.1). Hence it suffices to show that $(G \times_k Y) \times_X Z$ is a smooth k -variety. Now Z is smooth over k , so it suffices to show that $(G \times_k Y) \times_X Z \rightarrow Z$ is a smooth morphism (as the composition of two smooth morphisms is smooth, Exercise 13.6.D). But this is obtained by base change from $G \times_k Y \rightarrow X$, so it suffices to show that this latter morphism is smooth (as smoothness is preserved by base change).

Now $G \times_k Y \rightarrow X$ is a G -equivariant morphism. (By “ **G -equivariant**”, we mean that the G -action on both sides respects the morphism.) By generic smoothness of the target (Theorem 26.3.3), this is smooth over a dense open subset X . But then by transitivity of the G -action on X , this morphism is smooth everywhere. \square

26.3.C. EXERCISE (POOR MAN’S KLEIMAN-BERTINI). Prove Theorem 26.3.7(a) without the hypotheses on k (on algebraic closure or characteristic), and without the smoothness in the conclusion. Hint: This is a question about dimensions of fibers of morphisms, so you could have solved this after reading §12.4.

26.3.D. EXERCISE (IMPROVED CHARACTERISTIC 0 BERTINI). Suppose Z is a smooth k -variety, where $\text{char } k = 0$ and $k = \bar{k}$. Let V be a finite-dimensional base-point-free linear series on Z , i.e., a finite vector space of sections of some invertible sheaf \mathcal{L} on Z . Show that a general element of V , considered as a closed subscheme of Z , is regular. (More explicitly: each element $s \in V$ gives a closed subscheme of Z . Then for a general s , considered as a point of $\mathbb{P}V$, the corresponding closed subscheme is smooth over k .) Hint: figure out what this has to do with the Kleiman-Bertini Theorem 26.3.7. Let $n = \dim V$, $G = \text{GL}(V)$, $X = \mathbb{P}[V^\vee]$, take Z in Kleiman-Bertini to be the Z of the problem, and let Y be the “universal hyperplane” over $\mathbb{P}[V^\vee]$ (the incidence variety $I \subset \mathbb{P}V \times \mathbb{P}[V^\vee]$ of Definition 13.4.1).

26.3.E. EASY EXERCISE. Interpret Bertini's Theorem 13.4.2 over a characteristic 0 field as a corollary of Exercise 26.3.D.

In characteristic 0, Exercise 26.3.D is a good improvement on Bertini's Theorem. For example, we don't need \mathcal{L} to be very ample, or X to be projective. But unlike Bertini's Theorem, Exercise 26.3.D fails in positive characteristic, as shown by the one-dimensional linear series $\{pQ : Q \in \mathbb{P}^1\}$. This is essentially Example 26.3.2. (Do you see why this does not contradict Bertini's Theorem 13.4.2?)

CHAPTER 27

Twenty-seven lines

This topic is Chapter 27 for purely numerical reasons. As a result, it invokes one fact that we will only meet in a later chapter: Castelnuovo’s Criterion (Theorem 30.7.1) is used in the proof of Proposition 27.4.1. This is a feature rather than a bug: you needn’t read the details of the later chapters in order to enjoy this one. In fact, you should probably read this many chapters earlier, taking key facts as “black boxes”, to appreciate the value of seemingly technical definitions and theorems before you learn about them in detail. This topic is a fitting conclusion to an energetic introduction to algebraic geometry, even if you do not earlier cover everything needed to justify all details.

27.1 Introduction




FIGURE 27.1. Twenty-seven lines on a cubic surface

Wake an algebraic geometer in the dead of night, whispering: “27”. Chances are, he will respond: “lines on a cubic surface”.

— R. Donagi and R. Smith, [DS] (on page 27, of course)

Since the middle of the nineteenth century, geometers have been entranced by the fact that there are 27 lines on every smooth cubic surface, and by the remarkable structure of the configuration of the lines. Their discovery by Cayley and Salmon in 1849 has been called the beginning of modern algebraic geometry, [Do, p. 55].

The reason so many people are bewitched by this fact is because it requires some magic, and this magic connects to many other things, including fundamental ideas we have discussed, other beautiful classical constructions (such as Pascal’s Mystical Hexagon Theorem 20.4.4, the fact that most smooth quartic plane curves have 28 bitangents, exceptional Lie groups, ...), and many themes in modern algebraic geometry (deformation theory, intersection theory, enumerative geometry, arithmetic and diophantine questions, ...). It will be particularly pleasant for us, as it takes advantage of many of the things we have learned.

You are now ready to be initiated into the secret fellowship of the twenty-seven lines.

27.1.1. Theorem. — *Every smooth cubic surface in $\mathbb{P}_{\bar{k}}^3$ contains exactly 27 lines.*

Theorem 27.1.1 is closely related to the following.

27.1.2. Theorem. — *Every smooth cubic surface over \bar{k} is isomorphic to \mathbb{P}^2 blown up at 6 points.*

There are many reasons why people consider these facts magical. First, there is the fact that there are *always* 27 lines. Unlike most questions in enumerative geometry, there are no weasel words such as “a general cubic surface” or “most cubic surfaces” or “counted correctly” — as in, “every monic degree d polynomial has d roots — counted correctly”. And somehow (and we will see how) it is precisely the smoothness of the surface that makes it work.

Second, there is the magic that you *always* get the blow-up of the plane at six points (§27.4).

Third, there is the magical incidence structure of the 27 lines, which relates to E_6 in Lie theory. The Weyl group of E_6 is the symmetry group of the incidence structure (see Remark 27.3.5). In a natural way, the 27 lines form a basis of the 27-dimensional fundamental representation of E_6 .

27.1.3. Structure of this chapter. Throughout this chapter, X will be a smooth cubic surface over an algebraically closed field \bar{k} . In §27.2, we establish some preliminary facts. In §27.3, we prove Theorem 27.1.1. In §27.4, we prove Theorem 27.1.2. We remark here that the only input that §27.4 needs from §27.3 is Exercise 27.3.J. This can be done directly by hand (see for example [Rei, §7], [Be, Thm. IV.13], or [Sh, p. 246-7]), and Theorem 27.1.2 readily implies Theorem 27.1.1, using Exercise 27.4.E. We would thus have another, shorter, proof of Theorem 27.1.1. The reason for giving the argument of §27.3 is that it is natural given what we have done so far, it gives you some glimpse of some ideas used more broadly in the subject (the key idea is that a map from one moduli space to another is finite and flat), and it may help you further appreciate and digest the tools we have developed.

27.2 Preliminary facts

By Essential Exercise 15.1.C, there is a 20-dimensional vector space of cubic forms in four variables, so the cubic surfaces in \mathbb{P}^3 are parametrized by \mathbb{P}^{19} . Let $\Delta \subset \mathbb{P}^{19}$ be the locus corresponding to singular cubic hypersurfaces. Bertini’s

Theorem 13.4.2 shows that the locus of smooth cubics $\mathbb{P}^{19} \setminus \Delta$ is dense and open in the space \mathbb{P}^{19} of all cubics.

27.2.A. EXERCISE. Show that any smooth cubic surface X is “anticanonically embedded” — it is embedded by the complete anticanonical linear series \mathcal{K}_X^\vee . Hint: the adjunction formula, Exercise 22.5.B.

27.2.B. EXERCISE. Suppose $X \subset \mathbb{P}_k^3$ is a smooth cubic surface. Suppose C is a curve on X . Show that C is a line if and only if C is a (-1) -curve (Definition 21.2.7). Hint: the adjunction formula again, perhaps in the guise of Exercise 21.2.B(a); and also Exercise 19.6.L.

It will be useful to find a *single* cubic surface over a *single* field with exactly 27 lines:

27.2.C. HANDS-ON EXERCISE. Show that the **Fermat cubic surface**

$$(27.2.0.1) \quad x_0^3 + x_1^3 + x_2^3 + x_3^3 = 0$$

in $\mathbb{P}_\mathbb{C}^3$ has precisely 27 lines, each of the form

$$x_0 + \omega x_i = x_j + \omega' x_k = 0,$$

where $\{1, 2, 3\} = \{i, j, k\}$, $j < k$, and ω and ω' are cube roots of 1 (possibly the same). This will require some brute force. Possible hint: up to a permutation of coordinates, show that every line in \mathbb{P}^3 can be written $x_0 = ax_2 + bx_3$, $x_1 = cx_2 + dx_3$. Show that this line is on (27.2.0.1) if and only if

$$(27.2.0.2) \quad a^3 + c^3 + 1 = b^3 + d^3 + 1 = a^2b + c^2d = ab^2 + cd^2 = 0$$


Show that if a, b, c , and d are all nonzero, then (27.2.0.2) has no solutions.

27.2.1. ★ The singular cubic surfaces give an irreducible divisor Δ in the parameter space \mathbb{P}^{19} of all cubic surfaces. This discussion is starred not because it is hard, but because it is not needed in the rest of the chapter. Nonetheless, it is a pretty application of what we have learned, and it foreshadows key parts of the argument in §27.3.

Bertini’s Theorem 13.4.2 shows that the locus of smooth cubics in the \mathbb{P}^{19} of all cubics is dense and open, but the proof suggests more: that the complement Δ of this locus is of pure codimension 1. We now show that this is the case, and even that this “discriminant hypersurface” Δ is irreducible.

27.2.D. EXERCISE. (Hint for both: recall the solution to Exercise 12.2.L.)

(a) Define the incidence correspondence $Y \subset \mathbb{P}^{19} \times \mathbb{P}^3$ corresponding to the data of a cubic surface X , along with a *singular* point $p \in X$. (This is part of the exercise! We need Y as a scheme, not just as a set.) Let μ be the projection $Y \rightarrow \mathbb{P}^{19}$.



(b) Show that Y is an irreducible smooth variety of dimension 18, by describing it as a \mathbb{P}^{15} -bundle over \mathbb{P}^3 .

27.2.E. EXERCISE. Show that there exists a cubic surface with a single singular point. Feel free to assume k is your favorite field; the main point is to familiarize yourself with geometric ideas, not the peculiarities of positive characteristic fields. (Hint if $\text{char } k \neq 2$: Exercise 13.3.C.)

27.2.F. EXERCISE. Show that $\mu(Y) = \Delta$ is a closed irreducible subset of codimension 1. Hint for the codimension 1 statement: Exercise 12.4.A or Theorem 12.4.1, in combination with Exercise 27.2.E.


Your argument will generalize with essentially no change to deal with degree d hypersurfaces in \mathbb{P}^n .

27.3 Every smooth cubic surface (over \bar{k}) contains 27 lines

We are now ready to prove Theorem 27.1.1. Until Exercise 27.3.K, to avoid distraction, we assume $\text{char } \bar{k} = 0$. However, the following argument carries through without change if $\text{char } \bar{k} \neq 3$. The one required check — and the reason for the restriction on the characteristic — is that Exercise 27.2.C works with \mathbb{C} replaced by any such \bar{k} .

27.3.A. EXERCISE. (Hint for both: recall the solution to Exercise 12.2.L.)

(a) Define the incidence correspondence $Z \subset \mathbb{P}^{19} \times \mathbb{G}(1, 3)$ corresponding to the data of a line ℓ in \mathbb{P}^3 contained in a cubic surface X . (As in Exercise 27.2.D, this is part of the exercise.) Let π be the projection $Z \rightarrow \mathbb{P}^{19}$.



(b) Show that Z is an irreducible smooth variety of dimension 19, by describing it as a \mathbb{P}^{15} -bundle over the Grassmannian $\mathbb{G}(1, 3)$.

27.3.B. EXERCISE. Use the fact that there exists a cubic surface with a finite number of lines (Exercise 27.2.C), and the behavior of dimensions of fibers of morphisms (Exercise 12.4.A or Theorem 12.4.1) to show the following.

- (a) Every cubic surface contains a line, i.e., π is surjective. (Hint: show that π is projective.)
- (b) “Most cubic surfaces have a finite number of lines”: there is a dense open subset $U \subset \mathbb{P}^{19}$ such that the cubic surfaces parametrized by closed points of U have a positive finite number of lines. (Hint: upper semicontinuity of fiber dimension, Theorem 12.4.2.)

The following fact is the key result in the proof of Theorem 27.1.1, and one of the main miracles of the 27 lines, that ensures that the lines stay distinct on a smooth surface. It states, informally, that two lines can’t come together without damaging the surface. This is really a result in deformation theory: we are explicitly showing that a line in a smooth cubic surface has no first-order deformations.

27.3.1. Theorem. — *If ℓ is a line in a regular cubic surface X , then $\{\ell \subset X\}$ is a reduced isolated point of the fiber of $\pi: Z \rightarrow \mathbb{P}^{19}$.*

Before proving Theorem 27.3.1, we use it to prove Theorem 27.1.1.

27.3.2. *Proof of Theorem 27.1.1, assuming Theorem 27.3.1.* Now π is a projective morphism, and by Theorem 27.3.1, over $\mathbb{P}^{19} \setminus \Delta$, π has relative dimension 0, and hence has finite fibers. Hence by Theorem 19.1.7 (projective + finite fibers = finite), π is finite over $\mathbb{P}^{19} \setminus \Delta$.

Furthermore, Z and \mathbb{P}^{19} are both smooth of dimension 19 (over \mathbb{Z}), and the fibers of $\pi: Z \rightarrow \mathbb{P}^{19}$ are (geometrically) reduced points above $\mathbb{P}^{19} \setminus \Delta$, again by Theorem 27.3.1. Then by Exercise 26.2.F, π is étale over $\mathbb{P}^{19} \setminus \Delta$ (and in particular flat).

Thus, over $\mathbb{P}^{19} \setminus \Delta$, π is a finite flat morphism, and so the fibers of π (again, away from Δ) always have the same number of points, “counted correctly” (Remark 25.4.8). But by Theorem 27.3.1, above each closed point of $\mathbb{P}^{19} \setminus \Delta$, each point of the fiber of π counts with multiplicity one (here using $k = \bar{k}$). Finally, by Exercise 27.2.C, the Fermat cubic surface gives an example of one regular cubic surface with precisely 27 lines, so (as $\mathbb{P}^{19} \setminus \Delta$ is connected) we are done. \square

We have actually shown that away from Δ , $Z \rightarrow \mathbb{P}^{19}$ is a finite étale morphism of degree 27.

27.3.3. * *Proof of Theorem 27.3.1.* Choose projective coordinates so that the line ℓ is given, in a distinguished affine subset (with coordinates named x, y, z), by the z -axis. (We use affine coordinates to help visualize what we are doing, although this argument is better done in projective coordinates. On a second reading, you should translate this to a fully projective argument.)

27.3.C. EXERCISE. Consider the lines of the form $(x, y, z) = (a, b, 0) + t(a', b', 1)$ (where $(a, b, a', b') \in \mathbb{A}^4$ is fixed, and t varies in \mathbb{A}^1). Show that a, b, a', b' can be interpreted as the “usual” coordinates on one of the standard open subsets of the Grassmannian (see §7.7), with $[\ell]$ as the origin.

Having set up local coordinates on the moduli space, we can get down to business. Suppose $f(x, y, z) = 0$ is the (affine version) of the equation for the cubic surface X . Because X contains the z -axis ℓ , $f(x, y, z) \in (x, y)$. More generally, the line

$$(27.3.3.1) \quad (x, y, z) = (a, b, 0) + t(a', b', 1)$$

lies in X precisely when $f(a + ta', b + tb', t)$ is 0 as a cubic polynomial in t . This is equivalent to four equations in a, a', b , and b' , corresponding to the coefficients of t^3, t^2, t , and 1. This is better than just a set-theoretic statement:

27.3.D. EXERCISE. Verify that these four equations are local equations for the scheme-theoretic fiber (or scheme-theoretic preimage) $\pi^{-1}([\ell])$.

Now we come to the crux of the argument, where we use the regularity of X (along ℓ). We have a precise question in algebra. We are given a cubic surface X given by $f(x, y, z) = 0$, containing ℓ , and we know that X is regular (including “at ∞ ”, i.e., in \mathbb{P}^3). To show that $[\ell] = V(a, a', b, b')$ is a reduced isolated point in the fiber of π , we work in the ring $\bar{k}[a, a', b, b']/(a, a', b, b')^2$, i.e., we impose the




FIGURE 27.2. Parameters for the space of “lines near ℓ ”, in terms of where they meet the $z = 0$ plane and the $z = 1$ plane

equations

$$(27.3.3.2) \quad a^2 = aa' = \cdots = (b')^2 = 0,$$

and try to show that $a = a' = b = b' = 0$. (It is essential that you understand why we are setting $(a, a', b, b')^2 = 0$. You can also interpret this argument in terms of the derivatives of the functions involved — which after all can be interpreted as forgetting higher-order information and remembering only linear terms in the relevant variables, cf. Exercise 13.1.G. See [Mu6, §8D] for a description of this calculation in terms of derivatives.)

Suppose $f(x, y, z) = c_{x^3}x^3 + c_{x^2y}x^2y + \cdots + c_11 = 0$, where $c_{x^3}, c_{x^2y}, \dots \in \bar{k}$. Because $\ell \subset X$, i.e., $f \in (x, y)$, we have $c_1 = c_z = c_{z^2} = c_{z^3} = 0$. We now substitute (27.3.3.1) into f , and then apply (27.3.3.2). Only the coefficients of f of monomials involving precisely one x or y survive:

$$\begin{aligned} & c_x(a + a't) + c_{xz}(a + a't)t + c_{xz^2}(a + a't)t^2 \\ & + c_y(b + b't) + c_{yz}(b + b't)t + c_{yz^2}(b + b't)t^2 \\ & = (a + a't)(c_x + c_{xz}t + c_{xz^2}t^2) + (b + b't)(c_y + c_{yz}t + c_{yz^2}t^2) \end{aligned}$$

is required to be 0 as a polynomial in t . (Recall that c_x, \dots, c_{yz^2} are fixed elements of \bar{k} .) Let $C_x(t) = c_x + c_{xz}t + c_{xz^2}t^2$ and $C_y(t) = c_y + c_{yz}t + c_{yz^2}t^2$ for convenience.

Now X is regular at $(0, 0, 0)$ precisely when c_x and c_y are not both 0 (as $c_z = 0$). More generally, X is regular at $(0, 0, t_0)$ precisely if $c_x + c_{xz}t_0 + c_{xz^2}t_0^2 = C_x(t_0)$ and $c_y + c_{yz}t_0 + c_{yz^2}t_0^2 = C_y(t_0)$ are not both zero. You should be able to quickly check that X is regular at the point of ℓ “at ∞ ” precisely if c_{xz^2} and c_{yz^2} are not both zero. We summarize this as follows: X is regular at every point of ℓ precisely if the two quadratics $C_x(t)$ and $C_y(t)$ have no common roots, including “at ∞ ”.

We now use this to force $a = a' = b = b' = 0$ using $(a + a't)C_x(t) + (b + b't)C_y(t) \equiv 0$.

We deal first with the special case where C_x and C_y have two distinct roots, both finite (i.e., c_{xz^2} and c_{yz^2} are nonzero). If t_0 and t_1 are the roots of $C_x(t)$, then substituting t_0 and t_1 into $(a + a't)C_x(t) + (b + b't)C_y(t)$, we obtain $b + b't_0 = 0$, and $b + b't_1 = 0$, from which $b = b' = 0$. Similarly, $a = a' = 0$.

27.3.E. EXERCISE. Deal with the remaining cases to conclude the proof of Theorem 27.3.1. (It is possible to do this quite cleverly. For example, you may be able to re-choose coordinates to ensure that C_x and C_y have finite roots.)

□

27.3.4. The configuration of lines.

By the “configuration of lines” on a cubic surface, we mean the data of which pairs of the 27 lines intersect. We can readily work this out in the special case of the Fermat cubic surface (Exercise 27.2.C). (It can be more enlightening to use the description of X as a blow-up of \mathbb{P}^2 , see Exercise 27.4.E.) We now show that the configuration is the “same” (interpreted appropriately) for all smooth cubic surfaces.

27.3.F. EXERCISE. Construct a degree $27!$ finite étale map $W \rightarrow \mathbb{P}^{19} \setminus \Delta$, that parametrizes a cubic surface along with an *ordered list* of 27 distinct lines. Hint: let W' be the 27th fibered power of Z over $\mathbb{P}^{19} \setminus \Delta$, interpreted as parametrizing a cubic surface with an ordered list of 27 lines, not necessarily distinct. Let W be the subset corresponding to where the lines are distinct, and show that W is open and closed in W' , and thus a union of connected components of W' . An alternative way of identifying W within W' , in a way that might generalize to Exercise 27.3.H: make sense of the fact for each i, j , “ $\ell_i \cdot \ell_j$ ” is a locally constant function on W' , so the locus where these functions are positive for all $i \neq j$ is both open and closed; show that this is the W we seek.

We now make sense of the statement that the configuration of lines on the Fermat cubic surface (call it X_0) is the “same” as the configuration on some other smooth cubic surface (call it X_1). Lift the point $[X_0]$ to a point $w_0 \in W$. Let W'' be the connected component of W containing w_0 .

27.3.G. EXERCISE. Show that the morphism $W'' \rightarrow \mathbb{P}^{19} \setminus \Delta$ is finite and étale (and surjective).

Choose a point $w_1 \in W''$ mapping to $[X_1]$. Because W parametrizes a “labeling” or ordering of the 27 lines on a surface, we now have chosen an identification of the lines on X_0 with those of X_1 . Let the lines be ℓ_1, \dots, ℓ_{27} on X_0 , and let the corresponding lines on X_1 be m_1, \dots, m_{27} .

27.3.H. EXERCISE (USING STARRED EXERCISE 25.7.D). Show that $\ell_i \cdot \ell_j = m_i \cdot m_j$ for all i and j .

27.3.I. EXERCISE. Show that for each smooth cubic surface $X \subset \mathbb{P}_{\mathbb{K}}^3$, each line on X meets exactly 10 other lines $\ell_1, \ell'_1, \dots, \ell_5, \ell'_5$ on X , where ℓ_i and ℓ'_i meet for each i , and no other pair of the lines meet.

27.3.J. EXERCISE. Show that every smooth cubic surface contains two disjoint lines ℓ and ℓ' , such that there are precisely five other lines ℓ_1, \dots, ℓ_5 meeting both ℓ and ℓ' . (In fact, for any two disjoint lines ℓ and ℓ' in a smooth cubic surface, there are precise five other lines meeting them; but this is strictly harder than what the problem asks.)

27.3.5. Remark (cf. Remark 27.4.5): The Weyl group $W(E_6)$. The symmetry group of the configuration of lines — i.e., the subgroup of the permutations of the 27 lines preserving the intersection data — magically turns out to be the Weyl group of E_6 , a group of order 51840. (You know enough to at least verify that the size of the group is 51840, using the Fermat surface of Exercise 27.2.C, but this takes some work.) It is no coincidence that the degree of W'' over $\mathbb{P}^{19} \setminus \Delta$ is 51840, and the Galois group of the Galois closure of $K(Z)/K(\mathbb{P}^{19} \setminus \Delta)$ is isomorphic to $W(E_6)$ (see [H1, III.3]).

27.3.K. ★★ EXERCISE. Prove Theorem 27.3.1 in arbitrary characteristic. Begin by figuring out the right statement of Exercise 27.3.A over \mathbb{Z} , and proving it. Then follow the argument given in this section, making changes when necessary.

27.3.6. ★ Fano varieties of lines, and Hilbert schemes. In Exercises 27.3.A and 27.3.K, you constructed a moduli space of lines contained in a X , as a scheme. Your argument can be generalized to any $X \subset \mathbb{P}^N$. This construction is called the *Fano variety of lines* of X (not to be confused with the notion of a *Fano variety*, §22.5.5), and is an example of a Hilbert scheme.

27.4 Every smooth cubic surface (over \bar{k}) is a blown up plane

We now prove Theorem 27.1.2. As stated in §27.1.3, this section is remarkably independent from the previous one; all we will need is Exercise 27.3.J, and it is possible to prove this in other ways.

Suppose X is a smooth cubic surface (over \bar{k}). Suppose ℓ is a line on X , and choose coordinates on the ambient \mathbb{P}^3 so that ℓ is cut out by x_0 and x_1 . Projection from ℓ gives a rational map $\mathbb{P}^3 \dashrightarrow \mathbb{P}^1$ (given by $[x_0, x_1, x_2, x_3] \mapsto [x_0, x_1]$), which extends to a morphism on X . The reason is that this rational map is resolved by blowing up the closed subscheme $V(x_0, x_1)$ (Exercise 23.4.L). But (x_0, x_1) cuts out the effective Cartier divisor ℓ on X , and blowing up an effective Cartier divisor does not change X (Observation 23.2.1).

Now choose two disjoint lines ℓ and ℓ' as in Exercise 27.3.J, and consider the morphism $\rho: X \rightarrow \mathbb{P}^1 \times \mathbb{P}^1$, where the map to the first \mathbb{P}^1 is projection from ℓ , and the map to the second \mathbb{P}^1 is the projection from ℓ' . The first \mathbb{P}^1 can then be identified with ℓ' , and the second with ℓ .

27.4.A. EXERCISE. Show that the morphism ρ is birational. Hint: given a general point of $(p, q) \in \ell' \times \ell$, we obtain a point of X as follows: the line pq in \mathbb{P}^3 meets the cubic X at three points by Bézout's Theorem 9.3.D: p , q , and some third point $x \in X$; send (p, q) to x . (This idea appeared earlier in the development of the group law on the cubic curve, see Proposition 20.10.3.) Given a general point $x \in X$, we obtain a point $(p, q) \in \ell' \times \ell$ by projecting from ℓ' and ℓ .

In particular, we have shown for the first time that every smooth cubic surface over \bar{k} is rational.

27.4.B. EXERCISE: ρ CONTRACTS PRECISELY ℓ_1, \dots, ℓ_5 . Show that ρ is an isomorphism away from the ℓ_i mentioned in Exercise 27.3.J, and that each $\rho(\ell_i)$ is a point $p_i \in \ell' \times \ell$.

27.4.1. Proposition. — *The morphism $\rho: X \rightarrow \mathbb{P}^1 \times \mathbb{P}^1$ is the blow-up of $\mathbb{P}^1 \times \mathbb{P}^1$ at the five p_i .*

Proof. By Castelnuovo's Criterion (Theorem 30.7.1), as the lines ℓ_i are (-1) -curves (Exercise 27.2.B), they can be contracted. More precisely, there is a morphism $\beta: X \rightarrow X'$ that is the blow-up of X' at five closed points p'_1, \dots, p'_5 , such that ℓ_i is the exceptional divisor at p'_i . We wish to show that X' is $\mathbb{P}^1 \times \mathbb{P}^1$.

The morphism $\rho: X \rightarrow \mathbb{P}^1 \times \mathbb{P}^1$ yields a morphism $\rho': X' \setminus \{p'_1, \dots, p'_5\} \rightarrow \mathbb{P}^1 \times \mathbb{P}^1$. We now show that ρ' extends over p'_i for each i , sending p'_i to p_i . Choose an open neighborhood of $p_i \in \mathbb{P}^1 \times \mathbb{P}^1$ isomorphic to \mathbb{A}^2 , with coordinates x and y . Then both x and y pull back to functions on a punctured open neighborhoods of p'_i (i.e., there is some open neighborhood U of p'_i such that x and y are functions on $U \setminus \{p'_i\}$). By Algebraic Hartogs's Lemma 12.3.15, they extend over p'_i , and this extension is unique as $\mathbb{P}^1 \times \mathbb{P}^1$ is separated — use the Reduced-to-Separated Theorem 11.4.2 if you really need to. Thus ρ' extends over p'_i . (Do you see why $\rho'(p'_i) = p_i$?)

27.4.C. EXERCISE. Show that the birational morphism $\rho': X' \rightarrow \mathbb{P}^1 \times \mathbb{P}^1$ is invertible. Hint: you can use Zariski's Main Theorem (in the guise of Exercise 30.6.D), but you needn't use something so powerful. Instead, note that the birational map ρ'^{-1} is a morphism away from p_1, \dots, p_5 . Use essentially the same argument as in the last paragraph to extend ρ'^{-1} over each p_i . □

27.4.2. Conclusion of the proof of Theorem 27.1.2. As a consequence we see that X is the blow-up of $\mathbb{P}^1 \times \mathbb{P}^1$ at 5 points. Because the blow-up of $\mathbb{P}^1 \times \mathbb{P}^1$ at one point is isomorphic to the blow-up of \mathbb{P}^2 at two points (Exercise 23.4.K), Theorem 27.1.2 then follows. □

27.4.3. Reversing the process.

(This is a more precise version of 23.4.10(iii).) The process can be reversed: we can blow-up \mathbb{P}^2 at six points, and embed it in \mathbb{P}^3 . We first explain why we can't blow up \mathbb{P}^2 at just any six points and hope to embed the result in \mathbb{P}^3 . Because the cubic surface is embedded anticanonically (Exercise 27.2.A), we see that any curve C in X must satisfy $(\mathcal{K}_X \cdot C) < 0$.

27.4.D. EXERCISE. Suppose \mathbb{P}^2 is sequentially blown up at p_1, \dots, p_6 , resulting in a smooth surface X .

- (a) If p_i lies on the exceptional divisor of the blow-up at p_j ($i > j$), then show that there is a curve $C \subset X$ isomorphic to \mathbb{P}^1 , with $(\mathcal{K}_X \cdot C) \geq 0$.
- (b) If the p_i are distinct points on \mathbb{P}^2 , and three of them are collinear, show that there is a curve $C \subset X$ isomorphic to \mathbb{P}^1 , with $(\mathcal{K}_X \cdot C) \geq 0$.

(c) If the six p_i are distinct points on a smooth conic, show that there is a curve $C \subset X$ isomorphic to \mathbb{P}^1 , with $(\mathcal{K}_X \cdot C) \geq 0$.

Thus the only chance we have of obtaining a smooth cubic surface by blowing up six points on \mathbb{P}^2 is by blowing up six distinct points, no three on a line and not all on a conic.

27.4.4. Proposition. — *The anticanonical map of \mathbb{P}^2 blown up at six distinct points, no three on a line and not all on a conic, gives a closed embedding into \mathbb{P}^3 , as a cubic surface.*

Because we won't use this, we only describe the main steps of the proof: first count sections of the *anticanonical bundle* $\mathcal{K}_{\mathbb{P}^2}^\vee \cong \mathcal{O}_{\mathbb{P}^2}(3)$ (there is a 4-dimensional vector space of cubics on \mathbb{P}^2 vanishing at the six points, and these correspond to sections of the anticanonical bundle of the blow-up via Exercise 23.4.S(a)). Then show that these sections separate points and tangent vectors of X , thus showing that the anticanonical linear series gives a closed embedding, Theorem 20.1.1. Judicious use of the Cremona transformation (Exercise 7.5.I) can reduce the amount of tedious case-checking in this step.

27.4.E. EXERCISE. Suppose X is the blow-up of $\mathbb{P}_{\mathbb{K}}^2$ at six distinct points p_1, \dots, p_6 , no three on a line and not all on a conic. Verify that the only (-1) -curves on X are the six exceptional divisors, the proper transforms of the 15 lines $p_i p_j$, and the proper transforms of the six conics through five of the six points, for a total of 27.

27.4.F. EXERCISE. Solve Exercises 27.3.I and 27.3.J again, this time using the description of X as a blow-up of \mathbb{P}^2 .

27.4.5. Remark. If you blow up $4 \leq n \leq 8$ points on \mathbb{P}^2 , with no three on a line and no six on a conic, then the symmetry group of the configuration of lines is a Weyl group, as shown in the following table.

n	4	5	6	7	8
	$W(A_4)$	$W(D_5)$	$W(E_6)$	$W(E_7)$	$W(E_8)$

(If you know about Dynkin diagrams, you may see the pattern, and may be able to interpret what happens for $n = 3$ and $n = 9$.) This generalizes part of Remark 27.3.5, and the rest of it can similarly be generalized.

Cohomology and base change theorems

28.1 Statements and applications

Higher pushforwards are easy to define, but it is hard to get a geometric sense of what they are, or how they behave. For example, given a morphism $\pi: X \rightarrow Y$, and a quasicohherent sheaf \mathcal{F} on X , you might reasonably hope that the fibers of $R^i\pi_*\mathcal{F}$ are the cohomologies of \mathcal{F} along the fibers. More precisely, given $\psi: q \rightarrow Y$ corresponding to the inclusion of a point (better: $\psi: \text{Spec } \kappa(q) \rightarrow Y$), yielding the fibered diagram

$$(28.1.0.1) \quad \begin{array}{ccc} X_q & \xrightarrow{\psi'} & X \\ \pi' \downarrow & & \downarrow \pi \\ q & \xrightarrow{\psi} & Y, \end{array}$$

one might hope that the morphism

$$\phi_q^p: \psi^*(R^p\pi_*\mathcal{F}) \longrightarrow H^p(X_q, \mathcal{F}|_{X_q})$$

(given in Exercise 19.8.C) is an isomorphism. (Note: $\mathcal{F}|_{X_q}$ and $(\psi')^*\mathcal{F}$ are symbols for the same thing. The first is often preferred, but we sometimes use the second because we will consider more general ψ and ψ' .) We could then picture $R^i\pi_*\mathcal{F}$ as somehow fitting together the cohomology groups of fibers into a coherent sheaf. (Warning: this is too much to hope for, see Exercise 28.1.A.)

It would also be wonderful if $h^p(X_q, (\psi')^*\mathcal{F})$ was constant, and ϕ_q^p put them together into a nice locally free sheaf (vector bundle) $R^p\pi_*\mathcal{F}$.

The base change $\psi: q \rightarrow Y$ should not be special. As long as we are dreaming, we may as well hope that in good circumstances, given a Cartesian diagram (19.8.4.1)

$$(28.1.0.2) \quad \begin{array}{ccc} W & \xrightarrow{\psi'} & X \\ \pi' \downarrow & & \downarrow \pi \\ Z & \xrightarrow{\psi} & Y, \end{array}$$

the natural morphism

$$(28.1.0.3) \quad \phi_Z^p: \psi^*(R^p\pi_*\mathcal{F}) \longrightarrow R^p\pi'_*(\psi')^*\mathcal{F}$$

of sheaves on Z (Exercise 19.8.B(a)) is an isomorphism. (In some cases, we can already address this question. For example, cohomology commutes with flat base change, Theorem 25.2.9, so the result holds if ψ is flat.)

We formalize our dreams into three nice properties that we might wish in this situation. We will see that they are closely related. Suppose \mathcal{F} is a coherent sheaf on X , $\pi: X \rightarrow Y$ is proper, Y (hence X) is Noetherian, and \mathcal{F} is flat over Y .

- (a) Given a Cartesian square (28.1.0.1), is $\phi_q^p: R^p\pi_*\mathcal{F} \otimes \kappa(q) \rightarrow H^p(X_q, \mathcal{F}|_{X_q})$ an isomorphism?
- (b) Given a Cartesian square (28.1.0.2), is $\phi_Z^p: \psi^*(R^p\pi_*\mathcal{F}) \rightarrow R^p\pi'_*(\psi')^*\mathcal{F}$ an isomorphism?
- (c) Is $R^p\pi_*\mathcal{F}$ locally free?

We turn first to property (a). The dimension of the left side $R^p\pi_*\mathcal{F} \otimes \kappa(q)$ is an upper semicontinuous function of $q \in Y$ by upper semicontinuity of rank of finite type quasicoherent sheaves (Exercise 14.3.J). The Semicontinuity Theorem states that the dimension of the right is also upper semicontinuous. More formally:

28.1.1. Semicontinuity Theorem. — *Suppose $\pi: X \rightarrow Y$ is a proper morphism of Noetherian schemes, and \mathcal{F} is a coherent sheaf on X flat over Y . Then for each $p \geq 0$, the function $Y \rightarrow \mathbb{Z}$ given by $q \mapsto \dim_{\kappa(q)} H^p(X_q, \mathcal{F}|_{X_q})$ is an upper semicontinuous function of $q \in Y$.*

Translation: ranks of cohomology groups are upper semicontinuous in proper flat families. (A proof will be given in §28.2.4.)

28.1.2. Example. You may already have seen an example of cohomology groups jumping, in §25.4.14. Here is a simpler example, albeit not of the structure sheaf. Let (E, p_0) be an elliptic curve over a field k , and consider the projection $\pi: E \times E \rightarrow E$ to the second factor. Let \mathcal{L} be the invertible sheaf (line bundle) on $E \times E$ corresponding to the divisor that is the diagonal, minus the section of $p_0 \times E$ of π (where $p_0 \in E$). Then $\mathcal{L}|_{p_0}$ (i.e., $\mathcal{L}|_{E \times p_0}$) is trivial, but $\mathcal{L}|_p$ is non-trivial for any $p \neq p_0$ (as we showed in our study of genus 1 curves, in §20.9). Thus $h^0(E, \mathcal{L}|_p)$ is 0 in general, but jumps to 1 for $p = p_0$.

28.1.A. EXERCISE. Show that $\pi_*\mathcal{L} = 0$. Thus we cannot picture $\pi_*\mathcal{L}$ as “gluing together” h^0 of the fibers; in this example, cohomology does not commute with “base change” or “taking fibers”.

28.1.3. Side Remark. In characteristic 0, the cohomology of \mathcal{O} doesn’t jump in smooth families. Over \mathbb{C} , this is because Betti numbers are constant in connected families, and (22.5.11.1) (from Hodge theory) expresses the Betti constants h_{Betti}^k as sums (over $i + j = k$) of upper semicontinuous functions $h^j(\Omega^i)$, so the Hodge numbers $h^j(\Omega^i)$ must in fact be constant. The general characteristic 0 case can be reduced to \mathbb{C} by an application of the Lefschetz principle (which also arose in §22.5.9). But ranks of cohomology groups of \mathcal{O} for smooth families of varieties *can* jump in positive characteristic (see for example [MO70920]). Also, the example of §25.4.14 shows that the “smoothness” hypothesis cannot be removed.

28.1.4. Grauert’s Theorem. If $R^p\pi_*\mathcal{F}$ is locally free (property (c)) and ϕ_q^p is an isomorphism (property (a)), then $h^p(X_q, \mathcal{F}|_{X_q})$ is clearly locally constant. The following is a partial converse.

28.1.5. Grauert's Theorem. — *If $\pi: X \rightarrow Y$ is proper, Y is reduced and locally Noetherian, \mathcal{F} is a coherent sheaf on X flat over Y , and $h^p(X_q, \mathcal{F}|_{X_q})$ is a locally constant function of $q \in Y$, then $R^p\pi_*\mathcal{F}$ is locally free, and ϕ_Z^p is an isomorphism for all $\psi: Z \rightarrow Y$.*

In other words, if cohomology groups of fibers have locally constant dimension (over a reduced base), then they can be fit together to form a vector bundle, and the fiber of the pushforward is identified with the cohomology of the fiber. Our dreams at the start of this chapter have come true!

(See §28.2.12 to remove Noetherian hypotheses.)

We further note that if Y is integral, π is proper, and \mathcal{F} is a coherent sheaf on X flat over Y , then by the Semicontinuity Theorem 28.1.1 there is a dense open subset of Y on which $R^p\pi_*\mathcal{F}$ is locally free (and on which the fiber of the p th higher pushforward is the p th cohomology of the fiber).

The following statement is even more magical than Grauert's Theorem 28.1.5.

28.1.6. Cohomology and Base Change Theorem. — *Suppose π is proper, Y is locally Noetherian, \mathcal{F} is coherent over X and flat over Y , and ϕ_q^p is surjective. Then the following hold.*

- (i) *There is an open neighborhood U of q such that for any $\psi: Z \rightarrow U$, ϕ_Z^p is an isomorphism. In particular, ϕ_q^p is an isomorphism.*
- (ii) *Furthermore,*
 - (a) ϕ_q^{p-1} *is surjective (hence an isomorphism by (i)) if and only if*
 - (b) $R^p\pi_*\mathcal{F}$ *is locally free in some open neighborhood of q (or equivalently, $(R^p\pi_*\mathcal{F})_q$ is a free $\mathcal{O}_{Y,q}$ -module, Exercise 14.3.F).**(These then imply that $h^p(X_r, \mathcal{F}|_{X_r})$ is constant for r in an open neighborhood of q .)*

(Proofs of Theorems 28.1.5 and 28.1.6 will be given in §28.2. Note in (ii) that if $p = 0$, ϕ_q^{p-1} is automatically surjective, as ϕ_q^{-1} is the zero map. See §28.2.12 to remove Noetherian hypotheses.)

This is amazing: the hypothesis that ϕ_q^p is surjective involves what happens only over points q of X , with *reduced* structure, yet it has implications over the (possibly nonreduced) scheme as a whole! This might remind you of the local criterion for flatness (Theorem 25.6.2), which indeed is the key technical ingredient of the proof.

Here are some consequences.

28.1.B. EXERCISE. Use Theorem 28.1.6 to give a second solution to Exercise 25.4.E. (This is a big weapon to bring to bear on this problem, but it is still enlightening; your original solution to Exercise 25.4.E may foreshadow the proof of the Cohomology and Base Change Theorem 28.1.6.)

28.1.C. EXERCISE. Suppose $\pi: X \rightarrow Y$ is proper, Y is locally Noetherian, and \mathcal{F} is a coherent sheaf on X , flat over Y . Suppose further that $H^p(X_q, \mathcal{F}|_{X_q}) = 0$ for some $q \in Y$. Show that there is an open neighborhood U of q such that $(R^p\pi_*\mathcal{F})|_U = 0$.

28.1.D. EXERCISE. Suppose $\pi: X \rightarrow Y$ is proper, Y is locally Noetherian, and \mathcal{F} is a coherent sheaf on X , flat over Y . Suppose further that $H^p(X_q, \mathcal{F}|_{X_q}) = 0$ for all

$q \in Y$. Show that the $(p-1)$ st cohomology commutes with arbitrary base change: ϕ_Z^{p-1} is an isomorphism for all $\psi: Z \rightarrow Y$.

28.1.E. EXERCISE. Suppose π is proper, Y is locally Noetherian, and \mathcal{F} is a coherent sheaf on X flat over Y . Suppose further that $R^p\pi_*\mathcal{F} = 0$ for $p \geq p_0$. Show that $H^p(X_q, \mathcal{F}|_{X_q}) = 0$ for all $q \in Y$, $p \geq p_0$.

28.1.F. EXERCISE. Suppose π is proper, Y is locally Noetherian, and \mathcal{F} is a coherent sheaf on X , flat over Y . Suppose further that Y is reduced. Show that there exists a dense open subset U of Y such that ϕ_Z^p is an isomorphism for all $\psi: Z \rightarrow U$ and all p . (Hint: find suitable open neighborhoods of the generic points of Y . See Exercise 25.2.M and the paragraph following it.)

28.1.7. An important class of morphisms: Proper, \mathcal{O} -connected morphisms $\pi: X \rightarrow Y$ of locally Noetherian schemes.

If a morphism $\pi: X \rightarrow Y$ satisfies the property that the natural map $\mathcal{O}_Y \rightarrow \pi_*\mathcal{O}_X$ is an isomorphism, we say that π is **\mathcal{O} -connected**. Clearly the notion of \mathcal{O} -connectedness is local on the target, and preserved by composition.

28.1.G. EASY EXERCISE. Show that proper \mathcal{O} -connected morphisms of locally Noetherian schemes are surjective.

28.1.8. We will soon meet Zariski's Connectedness Lemma 30.5.1, which shows that proper, \mathcal{O} -connected morphisms of locally Noetherian schemes have connected fibers. In some sense, this class of morphisms is really the right class of morphisms capturing what we might want by "connected fibers"; this is the motivation for the terminology. The following result gives some evidence for this point of view, in the flat context.

28.1.H. IMPORTANT EXERCISE. Suppose $\pi: X \rightarrow Y$ is a proper *flat* morphism of locally Noetherian schemes, whose fibers satisfy $h^0(X_q, \mathcal{O}_{X_q}) = 1$. (Important remark: this is satisfied if π has geometrically connected and geometrically reduced fibers, by §11.5.7.) Show that π is \mathcal{O} -connected. Hint: consider

$$\mathcal{O}_Y \otimes \kappa(q) \longrightarrow (\pi_*\mathcal{O}_X) \otimes \kappa(q) \xrightarrow{\phi_q^0} H^0(X_q, \mathcal{O}_{X_q}) \cong \kappa(q).$$

The composition is surjective, hence ϕ_q^0 is surjective, hence it is an isomorphism by the Cohomology and Base Change Theorem 28.1.6(i). Then by the Cohomology and Base Change Theorem 28.1.6(ii), $\pi_*\mathcal{O}_X$ is locally free, thus of rank 1. Perhaps use Nakayama's Lemma to show that a map of invertible sheaves $\mathcal{O}_Y \rightarrow \pi_*\mathcal{O}_X$ that is an isomorphism on fibers over points (with reduced structure) is necessarily an isomorphism of sheaves.

28.1.9. * Unimportant remark. This class of proper, \mathcal{O} -connected morphisms is not preserved by arbitrary base change, and thus is not "reasonable" in the sense of §8.1. But you can show that they are preserved by *flat* base change, using the fact that cohomology commutes with flat base change, Theorem 25.2.9. Furthermore, the conditions of Exercise 28.1.H behave well under base change, and Noetherian hypotheses can be removed from the Cohomology and Base Change Theorem 28.1.6 (at the expense of finitely presented hypotheses, see §28.2.12), so the

class of morphisms $\pi: X \rightarrow Y$ that are proper, finitely presented, and flat, with geometrically connected and geometrically reduced fibers, is “reasonable” (and useful).

28.1.10. We next address the following question. Suppose $\pi: X \rightarrow Y$ is a morphism of schemes. Given an invertible sheaf \mathcal{L} on X , we ask when it is the pullback of an invertible sheaf \mathcal{M} on Y . For this to be true, we certainly need that \mathcal{L} is trivial on the fibers. We will see that if π is a proper \mathcal{O} -connected morphism of locally Noetherian schemes, then this often suffices. Given \mathcal{L} , we recover \mathcal{M} as $\pi_*\mathcal{L}$; the fibers of \mathcal{M} are one-dimensional, and glue together to form a line bundle. We now begin to make this precise.

28.1.I. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a proper, \mathcal{O} -connected morphism of locally Noetherian schemes. Show that if \mathcal{M} is any invertible sheaf on Y , then the natural morphism $\mathcal{M} \rightarrow \pi_*\pi^*\mathcal{M}$ is an isomorphism. In particular, we can recover \mathcal{M} from $\pi^*\mathcal{M}$ by applying the pushforward π_* .

28.1.11. Proposition. — Suppose $\pi: X \rightarrow Y$ is a **flat**, proper morphism of locally Noetherian schemes with geometrically connected and geometrically reduced fibers (hence \mathcal{O} -connected, by Exercise 28.1.H). Suppose also that Y is reduced, and \mathcal{L} is an invertible sheaf on X that is trivial on the fibers of π (i.e., $\mathcal{L}|_{X_q}$ is a trivial invertible sheaf on X_q for all $q \in Y$). Then $\pi_*\mathcal{L}$ is an invertible sheaf on Y (call it \mathcal{M}), and the natural map $\pi^*\mathcal{M} \rightarrow \mathcal{L}$ is an isomorphism.

Proof. By Grauert’s Theorem 28.1.5, $\pi_*\mathcal{L}$ is locally free of rank 1 (again, call it \mathcal{M}), and $\mathcal{M} \otimes_{\mathcal{O}_Y} \kappa(q) \rightarrow H^0(X_q, \mathcal{L}|_{X_q})$ is an isomorphism. We have a natural map of invertible sheaves $\pi^*\mathcal{M} = \pi^*\pi_*\mathcal{L} \rightarrow \mathcal{L}$. To show that it is an isomorphism, we need only show that it is surjective. (Do you see why? If A is a ring, and $\phi: A \rightarrow A$ is a surjection of A -modules, why is ϕ an isomorphism?) For this, it suffices to show that it is surjective on the fibers of π . (Do you see why? Hint: if the cokernel of the map is not 0, then it is not 0 above some point of Y .) But this follows from the first line of the proof (using for example that $\mathcal{M} \cong \mathcal{O}$ in a neighborhood of q). \square

Proposition 28.1.11 has some pleasant consequences. For example, if you have two invertible sheaves \mathcal{A} and \mathcal{B} on X that are isomorphic on every fiber of π , then they differ by a pullback of an invertible sheaf on Y : just apply Proposition 28.1.11 to $\mathcal{A} \otimes \mathcal{B}^\vee$.

28.1.12. Projective bundles.

28.1.J. EXERCISE. Let X be a locally Noetherian scheme, and let $\text{pr}_1: X \times \mathbb{P}^n \rightarrow X$ be the projection onto the first factor. Suppose \mathcal{L} is a line bundle on $X \times \mathbb{P}^n$, whose degree on every fiber of pr_1 is zero. Use the Cohomology and Base Change Theorem 28.1.6 to show that $(\text{pr}_1)_*\mathcal{L}$ is an invertible sheaf on X . Show that the natural map $\text{pr}_1^*((\text{pr}_1)_*\mathcal{L}) \rightarrow \mathcal{L}$ is an isomorphism, as at the end of proof of Proposition 28.1.11.

Your argument will apply just as well to the situation where $\text{pr}_1: X \times \mathbb{P}^n \rightarrow X$ is replaced by a \mathbb{P}^n -bundle over X , $\text{pr}_1: Z \rightarrow X$; or by $\text{pr}_1: Z \rightarrow X$ which is a proper smooth morphism whose geometric fibers are integral curves of genus 0.

Furthermore, the locally Noetherian hypotheses can be removed, see §28.2.12.

28.1.K. EXERCISE. Suppose X is a connected Noetherian scheme. Show that $\mathrm{Pic}(X \times \mathbb{P}^n) \cong \mathrm{Pic} X \times \mathbb{Z}$. Hint: the map $\mathrm{Pic} X \times \mathrm{Pic} \mathbb{P}^n \rightarrow \mathrm{Pic}(X \times \mathbb{P}^n)$ is given by $(\mathcal{L}, \mathcal{O}(m)) \mapsto \mathrm{pr}_1^* \mathcal{L} \otimes \mathrm{pr}_2^* \mathcal{O}(m)$, where $\mathrm{pr}_1: X \times \mathbb{P}^n \rightarrow X$ and $\mathrm{pr}_2: X \times \mathbb{P}^n \rightarrow \mathbb{P}^n$ are the projections from $X \times \mathbb{P}^n$ to its factors. (The notation \boxtimes is often used for this construction, see §17.4.9.)

A very similar argument will show that if Z is a \mathbb{P}^n -bundle over X , then $\mathrm{Pic} Z \cong \mathrm{Pic} X \times \mathbb{Z}$. You will undoubtedly also be able to figure out the right statement if X is not connected.

28.1.13. Remark. As mentioned in §20.10.1, the Picard group of a scheme often “wants to be a scheme”. You may be able to make this precise in the case of $\mathrm{Pic} \mathbb{P}_{\mathbb{Z}}^n$. In this case, the *scheme* $\mathrm{Pic} \mathbb{P}_{\mathbb{Z}}^n$ is “ \mathbb{Z} copies of $\mathrm{Spec} \mathbb{Z}$ ”, with the “obvious” group scheme structure. Can you figure out what functor it represents? Can you show that it represents this functor? This will require extending Exercise 28.1.K out of the Noetherian setting, using §28.2.12.

28.1.L. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a projective flat morphism over a Noetherian integral scheme, all of whose geometric fibers are isomorphic to \mathbb{P}^n (over the appropriate field). Show that π is a projective bundle if and only if there is an invertible sheaf \mathcal{L} on X that restricts to $\mathcal{O}(1)$ on all the geometric fibers. (One direction is clear: if it is a projective bundle, then it has a $\mathcal{O}(1)$ which comes from the projectivization, see Exercise 18.2.D. In the other direction, the candidate vector bundle is $\pi_* \mathcal{L}$. Show that it is indeed a locally free sheaf of the desired rank. Show that its projectivization is indeed $\pi: X \rightarrow Y$.)

Caution: the map $\pi: \mathrm{Proj} \mathbb{R}[x, y, z]/(x^2 + y^2 + z^2) \rightarrow \mathrm{Spec} \mathbb{R}$ shows that not every projective flat morphism over a Noetherian integral scheme, all of whose geometric fibers are isomorphic to \mathbb{P}^n , is necessarily a \mathbb{P}^n -bundle. However, *Tsen’s Theorem* implies that if the target is a *smooth curve over an algebraically closed field*, then the morphism is a \mathbb{P}^n -bundle (see [GS, Thm. 6.2.8]). Example 19.4.7 shows that “curve” cannot be replaced by “5-fold” in this statement — the “universal smooth plane conic” is not a \mathbb{P}^1 -bundle over the parameter space $U \subset \mathbb{P}^5$ of smooth plane conics. If you wish, you can extend Example 19.4.7 to show that “curve” cannot even be replaced by “surface”. (Just replace the \mathbb{P}^5 of all conics with a generally chosen \mathbb{P}^2 of conics — but then figure out what goes wrong if you try to replace it with a generally chosen \mathbb{P}^1 of conics.)

28.1.M. EXERCISE. Suppose $\pi: X \rightarrow Y$ is the projectivization of a vector bundle \mathcal{F} over a locally Noetherian scheme (i.e., $X \cong \mathrm{Proj} \mathrm{Sym}^\bullet \mathcal{F}$). Recall from §18.2.3 that for any invertible sheaf \mathcal{L} on Y , $X \cong \mathrm{Proj} \mathrm{Sym}^\bullet(\mathcal{F} \otimes \mathcal{L})$. Show that these are the only ways in which it is the projectivization of a vector bundle. (Hint: recover \mathcal{F} by pushing forward $\mathcal{O}(1)$.)

28.1.14. The Hodge bundle.

28.1.N. EXERCISE (THE HODGE BUNDLE). Suppose $\pi: X \rightarrow Y$ is a flat proper morphism of locally Noetherian schemes, and the fibers of π are regular irreducible

curves of genus g . (By Theorem 26.2.2 it is a smooth morphism of relative dimension 1, and $\Omega_{X/Y}$ is a line bundle.) Show that $\pi_*\Omega_{X/Y}$ is a locally free sheaf on Y of rank g , and that the construction of π commutes with base change: given a Cartesian square

$$(28.1.14.1) \quad \begin{array}{ccc} X' & \xrightarrow{\psi'} & X \\ \pi' \downarrow & & \downarrow \pi \\ Y' & \xrightarrow{\psi} & Y, \end{array}$$

there is a canonical isomorphism

$$(\pi'_*\Omega_{X'/Y'}) \xrightarrow{\sim} \psi^*(\pi_*\Omega_{X/Y}).$$

(The locally free sheaf $\pi_*\Omega_{X/Y}$ is called the **Hodge bundle**.) Hint: use the Cohomology and Base Change Theorem 28.1.6 twice, once with $p = 2$, and once with $p = 1$.

28.2 ★ Proofs of cohomology and base change theorems

The key to proving the Semicontinuity Theorem 28.1.1, Grauert's Theorem 28.1.5, and the Cohomology and Base Change Theorem 28.1.6 is the following wonderful idea of Mumford (see [Mu3, p. 47 Lem. 1]). It turns questions of pushforwards (and how they behave under arbitrary base change) into something computable with vector bundles (hence questions of linear algebra). After stating it, we will interpret it.

28.2.1. Key Theorem. — Suppose $\pi: X \rightarrow \operatorname{Spec} B$ is a proper morphism, and \mathcal{F} is a coherent sheaf on X , flat over $\operatorname{Spec} B$, where B is Noetherian. Then there is a complex

$$(28.2.1.1) \quad \cdots \xrightarrow{\delta^{-2}} K^{-1} \xrightarrow{\delta^{-1}} K^0 \xrightarrow{\delta^0} K^1 \xrightarrow{\delta^1} \cdots \longrightarrow K^n \xrightarrow{\delta^n} 0$$

of finitely generated free B -modules and an isomorphism of functors

$$(28.2.1.2) \quad H^p(X \times_B A, \mathcal{F} \otimes_B A) \xrightarrow{\sim} H^p(K^\bullet \otimes_B A)$$

for all p , for all ring maps $B \rightarrow A$. (Here A needn't be Noetherian.)

Because (28.2.1.1) is a complex of free B -modules, all of the information is contained in the maps, which are matrices with entries in B . This will turn questions about cohomology (and base change) into questions about linear algebra. For example, semicontinuity will turn into the fact that ranks of matrices (with functions as entries) drop on closed subsets (§12.4.4(ii)).

Although the complex (28.2.1.1) is infinite, by (28.2.1.2) it has no cohomology in negative degree, even after any ring extension $B \rightarrow A$ (as the left side of (28.2.1.2) is 0 for $p < 0$).

The idea behind the proof is as follows: take the Čech complex, produce a complex of finite rank free modules mapping to it “with the same cohomology” (a *quasiisomorphic complex*, §19.2.3). We first construct the complex so that (28.2.1.2) holds for $B = A$ in the next lemma, and then show the same complex works for general A , in Lemma 28.2.3 immediately thereafter.

28.2.2. Lemma. — Let B be a Noetherian ring. Suppose C^\bullet is a complex of B -modules such that $H^i(C^\bullet)$ are finitely generated B -modules, and such that $C^p = 0$ for $p > n$. Then there exists a complex K^\bullet of finite rank free B -modules such that $K^p = 0$ for $p > n$, and a homomorphism of complexes $\alpha: K^\bullet \rightarrow C^\bullet$ such that α induces isomorphisms $H^i(K^\bullet) \xrightarrow{\sim} H^i(C^\bullet)$ for all i .

Proof. We build this complex inductively. (This may remind you of Hint 24.3.3.) Assume we have defined $(K^p, \alpha^p, \delta^p)$ for $p \geq m+1$ (as in (28.2.2.1)) such that the squares commute, and the top row is a complex, and α^q defines an isomorphism of cohomology $H^q(K^\bullet) \rightarrow H^q(C^\bullet)$ for $q \geq m+2$ and a surjection $\ker(\delta^{m+1}) \rightarrow H^{m+1}(C^\bullet)$, and the K^p are finite rank free B -modules. (Our base case is $m = p$: take $K^n = 0$ for $n > p$.)

$$(28.2.2.1) \quad \begin{array}{ccccccc} & & & & K^{m+1} & \xrightarrow{\delta^{m+1}} & K^{m+2} \xrightarrow{\delta^{m+2}} \dots \\ & & & & \downarrow \alpha^{m+1} & & \downarrow \alpha^{m+2} \\ \dots & \longrightarrow & C^{m-1} & \xrightarrow{\epsilon^{m-1}} & C^m & \xrightarrow{\epsilon^m} & C^{m+1} \xrightarrow{\epsilon^{m+1}} C^{m+2} \longrightarrow \dots \end{array}$$

We construct $(K^m, \delta^m, \alpha^m)$. Choose generators of $H^m(C^\bullet)$, say c_1, \dots, c_M . Let

$$D^{m+1} := \ker \left(\ker(\delta^{m+1}) \xrightarrow{\alpha^{m+1}} H^{m+1}(C^\bullet) \right).$$

Choose generators of D^{m+1} , say d_1, \dots, d_N . (This is where we use the Noetherian hypotheses — to ensure this kernel D^{m+1} is finitely generated.) Let $K^m = B^{\oplus(M+N)}$. Define $\delta^m: K^m \rightarrow K^{m+1}$ by sending the last N generators to d_1, \dots, d_N , and the first M generators to 0. Define α^m by sending the first M generators of $B^{\oplus(M+N)}$ to (lifts of) c_1, \dots, c_M , and sending the last N generators to arbitrarily chosen lifts of the $\alpha^{m+1}(d_i)$ (as the $\alpha^{m+1}(d_i)$ are 0 in $H^{m+1}(C^\bullet)$, and thus lie in the image of ϵ^m), so the square (with upper left corner K^m) commutes. Then by construction, we have completed our inductive step:

$$\begin{array}{ccccccc} & & & & K^m & \xrightarrow{\delta^m} & K^{m+1} \xrightarrow{\delta^{m+1}} K^{m+2} \longrightarrow \dots \\ & & & & \downarrow \alpha^m & & \downarrow \alpha^{m+1} \\ \dots & \longrightarrow & C^{m-1} & \xrightarrow{\epsilon^{m-1}} & C^m & \xrightarrow{\epsilon^m} & C^{m+1} \xrightarrow{\epsilon^{m+1}} C^{m+2} \longrightarrow \dots \end{array}$$

□

28.2.3. Lemma. — Suppose $\alpha: K^\bullet \rightarrow C^\bullet$ is a morphism of complexes of **flat** B -modules, bounded on the right (i.e., $K^n = C^n = 0$ for $n \gg 0$), inducing isomorphisms of cohomology (a quasiisomorphism, §19.2.3). Then “this quasiisomorphism commutes with arbitrary change of base ring”: for every B -algebra A , the maps $H^p(K^\bullet \otimes_B A) \rightarrow H^p(C^\bullet \otimes_B A)$ are isomorphisms.

Proof. The mapping cone M^\bullet of $\alpha: K^\bullet \rightarrow C^\bullet$ is exact by Exercise 1.7.E. Then $M^\bullet \otimes_B A$ is still exact, by Exercise 25.3.F. But $M^\bullet \otimes_B A$ is the mapping cone of

$$\alpha \otimes_B A: K^\bullet \otimes_B A \rightarrow C^\bullet \otimes_B A,$$

so by Exercise 1.7.E, $\alpha \otimes_B A$ induces an isomorphism of cohomology (i.e., is a quasiisomorphism) too. \square

Proof of Key Theorem 28.2.1. Choose a finite affine covering of X . Take the Čech complex C^\bullet for \mathcal{F} with respect to this cover. Recall that Grothendieck's Coherence Theorem 19.9.1 (which had Noetherian hypotheses) showed that the cohomology of \mathcal{F} is coherent. (Theorem 19.9.1 required serious work. If you need Theorem 28.2.1 only in the projective case, the analogous statement with projective hypotheses, Theorem 19.8.1(d), was much easier.) Apply Lemma 28.2.2 to get the nicer variant K^\bullet of the same complex C^\bullet . By Lemma 28.2.3, if we tensor with A and take cohomology, we get the same answer whether we use K^\bullet or C^\bullet . \square

We now use Theorem 28.2.1 to prove some of the fundamental results stated earlier: the Semicontinuity Theorem 28.1.1, Grauert's Theorem 28.1.5, and the Cohomology and Base Change Theorem 28.1.6. In the course of proving Semicontinuity, we will give a new proof of Theorem 25.7.1, that Euler characteristics are locally constant in flat families (that applies more generally in proper situations).

28.2.4. Proof of the Semicontinuity Theorem 28.1.1. The result is local on Y , so we may assume Y is affine. Let K^\bullet be a complex as in Key Theorem 28.2.1.

Then for $q \in Y$,

$$\begin{aligned} \dim_{\kappa(q)} H^p(X_q, \mathcal{F}|_{X_q}) &= \dim_{\kappa(q)} \ker(\delta^p \otimes_B \kappa(q)) - \dim_{\kappa(q)} \operatorname{im}(\delta^{p-1} \otimes_B \kappa(q)) \\ &= \dim_{\kappa(q)} (K^p \otimes_B \kappa(q)) - \dim_{\kappa(q)} \operatorname{im}(\delta^p \otimes_B \kappa(q)) \\ &\quad - \dim_{\kappa(q)} \operatorname{im}(\delta^{p-1} \otimes_B \kappa(q)) \end{aligned} \tag{28.2.4.1}$$

Now $\dim_{\kappa(q)} \operatorname{im}(\delta^p \otimes_B \kappa(q))$ is a lower semicontinuous function on Y . (Reason: the locus where the dimension is less than some number N is obtained by setting all $N \times N$ minors of the matrix $K^p \rightarrow K^{p+1}$ to 0; cf. §12.4.4(ii)). The same is true for $\dim_{\kappa(q)} \operatorname{im}(\delta^{p-1} \otimes_B \kappa(q))$. The result follows. \square

28.2.5. A new proof (and extension to the proper case) of Theorem 25.7.1 that Euler characteristics of flat sheaves are locally constant.

If K^\bullet were finite “on the left” as well — if $K^p = 0$ for $p \ll 0$ — then we would have a short proof of Theorem 25.7.1. By taking alternating sums (over p) of (28.2.4.1), we would have that

$$\chi(X_q, \mathcal{F}|_{X_q}) = \sum (-1)^p h^p(X_q, \mathcal{F}|_{X_q}) = \sum (-1)^p \operatorname{rank} K^p,$$

which is locally constant. The only problem is that the sums are infinite. We patch this problem by truncating the complex K^\bullet below where there is cohomology. Define J^\bullet by $J^p = K^p$ for $p \geq 0$, $J^p = 0$ for $p < -1$, and $J^{-1} := \ker(K^{-1} \rightarrow K^0)$. Then J^\bullet is a complex in the obvious way, and the map of complexes $K^\bullet \rightarrow C^\bullet$ clearly

factors through J^\bullet :

$$\begin{array}{ccccccccccc}
 \cdots & \longrightarrow & J^{-3} & \longrightarrow & J^{-2} & \longrightarrow & J^{-1} & \longrightarrow & J^0 & \longrightarrow & J^1 & \longrightarrow & J^2 & \longrightarrow & \cdots \\
 & & \parallel & & \parallel & & \parallel & & \parallel & & \parallel & & \parallel & & \\
 \cdots & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & \ker(K^{-1} \rightarrow K^0) & \longrightarrow & K^0 & \longrightarrow & K^1 & \longrightarrow & K^2 & \longrightarrow & \cdots \\
 & & \downarrow \sim & & \downarrow \sim & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
 \cdots & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & C^0 & \longrightarrow & C^1 & \longrightarrow & C^2 & \longrightarrow & \cdots
 \end{array}$$

Clearly $J^\bullet \rightarrow C^\bullet$ induces an isomorphism on cohomology (recall both have 0 cohomology for $p < 0$).

Now J^{-1} (the kernel of a map of coherent modules) is coherent. Consider the mapping cone M^\bullet of $\beta: J^\bullet \rightarrow C^\bullet$:

$$0 \rightarrow J^{-1} \rightarrow C^{-1} \oplus J^0 \rightarrow C^0 \oplus J^1 \rightarrow \cdots \rightarrow C^{n-1} \oplus J^n \rightarrow C^n \rightarrow 0.$$

From Exercise 1.7.E, as $J^\bullet \rightarrow C^\bullet$ induces an isomorphism on cohomology, the mapping cone has no cohomology — it is exact. All terms in it are flat except possibly J^{-1} (the C^p are flat by assumption, and J^i is free for $i \neq -1$). Hence J^{-1} is flat too, by Exercise 25.3.G. But flat coherent sheaves are locally free (Corollary 25.4.7). Then Theorem 25.7.1 follows from

$$\chi(X_q, \mathcal{F}|_{X_q}) = \sum (-1)^p h^p(X_q, \mathcal{F}|_{X_q}) = \sum (-1)^p \text{rank } J^p.$$

□

28.2.6. Proof of Grauert's Theorem 28.1.5 and the Cohomology and Base Change Theorem 28.1.6 (following Eric Larson).

(** Experts: You might see in the proof that what makes p th cohomology commute with base change is that the complex of Key Theorem 28.2.1 can be “broken” into two complexes at the p th step. You might even want to interpret this in terms of the Čech complex as an object of the derived category of B -modules.)

Thanks to Theorem 28.2.1, Theorems 28.1.5 and 28.1.6 are now statements about a complex of free modules over a Noetherian ring.

28.2.7. Definition. Suppose $\phi: \mathcal{E} \rightarrow \mathcal{F}$ is a morphism of finite rank locally free sheaves on a scheme X . (More precisely: \mathcal{E} and \mathcal{F} are finite rank locally free sheaves, and $\phi: \mathcal{E} \rightarrow \mathcal{F}$ is a morphism of quasicoherent sheaves.) We say ϕ is **strongly of constant rank a** if for every point $p \in X$, there are integers b and c , and there is an open neighborhood U of p with a commutative diagram

$$\begin{array}{ccccc}
 \mathcal{E}|_U & \xrightarrow{\quad \phi \quad} & \mathcal{F}|_U & & \\
 \uparrow \sim & & \uparrow \sim & & \\
 \mathcal{O}_U^{\oplus(a+b)} & \longrightarrow & \mathcal{O}_U^{\oplus a} & \longrightarrow & \mathcal{O}_U^{\oplus(a+c)}
 \end{array}$$

where the map $\mathcal{O}_U^{\oplus(a+b)} \rightarrow \mathcal{O}_U^{\oplus a}$ is projection to the first a summands, and the map $\mathcal{O}_U^{\oplus a} \rightarrow \mathcal{O}_U^{\oplus(a+c)}$ is inclusion as the first a summands. We say that ϕ is **strongly of constant rank** if near every point $p \in X$, ϕ is strongly of constant rank a for some a .

28.2.8. Important but straightforward observations. (i) The notion “strongly of constant rank α ” commutes with any base change $Y \rightarrow X$. (ii) The quasicoherent sheaves $\ker \phi$, $\operatorname{im} \phi$, and $\operatorname{coker} \phi$ are locally free (of finite rank b , α , and c respectively), and their construction commutes with any base change $Y \rightarrow X$.

$$\begin{array}{ccccccccc}
 0 & \longrightarrow & \ker \phi & \longrightarrow & \mathcal{E} & \xrightarrow{\phi} & \mathcal{F} & \longrightarrow & \operatorname{coker} \phi & \longrightarrow & 0 \\
 & & \uparrow \sim & & \uparrow \sim & & \uparrow \sim & & \uparrow \sim & & \\
 0 & \longrightarrow & \mathcal{O}_U^{\oplus b} & \longrightarrow & \mathcal{O}_U^{\oplus (a+b)} & \longrightarrow & \mathcal{O}_U^{\oplus (a+c)} & \longrightarrow & \mathcal{O}_U^{\oplus c} & \longrightarrow & 0
 \end{array}$$

(iii) If $\phi: \mathcal{E} \rightarrow \mathcal{F}$ is strongly of constant rank α , then at any point $p \in X$, the rank of $\phi|_p: \mathcal{E}_p \rightarrow \mathcal{F}_p$ is α .

In the next exercises we give three criteria for when a morphism ϕ of finite rank locally free sheaves is strongly of constant rank.

28.2.A. EXERCISE. Suppose $\phi: \mathcal{E} \rightarrow \mathcal{F}$ is a morphism of finite rank locally free sheaves. Show that ϕ is strongly of constant rank α if and only if $\operatorname{coker} \phi$ is locally free of rank α .

28.2.B. EXERCISE. Suppose X is reduced, and $\phi: \mathcal{E} \rightarrow \mathcal{F}$ is a morphism of finite rank locally free sheaves. Show that ϕ is strongly of constant rank α if and only if $\phi|_p: \mathcal{E}_p \rightarrow \mathcal{F}_p$ is rank α for all points $p \in X$. (Hint: use the previous exercise, and Exercise 14.3.K.)

28.2.C. EXERCISE. Suppose $\phi: \mathcal{E} \rightarrow \mathcal{F}$ is a morphism of finite rank locally free sheaves, and $p \in X$. Show that the natural map $(\ker \phi)|_p \rightarrow \ker(\phi|_p)$ is surjective if and only if ϕ is strongly of constant rank in some neighborhood of p .

28.2.9. Proof of Grauert’s Theorem 28.1.5. By hypothesis, $h^p(X_q, \mathcal{F}|_{X_q})$ is a locally constant function of $q \in Y$. From (28.2.4.1), $h^p(X_q, \mathcal{F}|_{X_q}) = \operatorname{rank} K^p - \operatorname{rank} \operatorname{im}(\delta^p|_q) - \operatorname{rank} \operatorname{im}(\delta^{p-1}|_q)$. But $\operatorname{rank} K^p$ is constant, and $\operatorname{rank} \operatorname{im}(\delta^p|_q)$ and $\operatorname{rank} \operatorname{im}(\delta^{p-1}|_q)$ are lower semicontinuous, so in fact $\operatorname{rank} \operatorname{im}(\delta^p|_q)$ and $\operatorname{rank} \operatorname{im}(\delta^{p-1}|_q)$ must be locally constant. By Exercise 28.2.B, both δ^{p-1} and δ^p are strongly of constant rank. Then by Observation 28.2.8(ii) $\operatorname{coker} \delta^{p-1}$ and $\operatorname{im} \delta^p$ are both locally free (of finite rank). In the short exact sequence

$$(28.2.9.1) \quad 0 \rightarrow H^p(K^\bullet) \rightarrow \operatorname{coker} \delta^{p-1} \rightarrow \operatorname{im} \delta^p \rightarrow 0$$

(Exercise 1.6.5.4, the “dual” definition of cohomology), both $\operatorname{coker} \delta^{p-1}$ and $\operatorname{im} \delta^p$ correspond to finite rank locally free sheaves. Thus $H^p(K^\bullet)$ does as well, by Exercise 14.2.R(a).

28.2.D. EXERCISE. Show (perhaps using (28.2.9.1)) that the construction of $H^p(K^\bullet)$ commutes with any base change, thereby completing the proof of Grauert’s Theorem 28.1.5. \square

In order to prove the Cohomology and Base Change Theorem 28.1.6, we need a preliminary result.

28.2.10. Lemma. — Suppose

$$\begin{array}{ccccc} K^{p-1} & \xrightarrow{\delta_K^{p-1}} & K^p & \xrightarrow{\delta_K^p} & K^{p+1} \\ \downarrow & & \downarrow & & \downarrow \\ J^{p-1} & \xrightarrow{\delta_J^{p-1}} & J^p & \xrightarrow{\delta_J^p} & J^{p+1} \end{array}$$

is a map of complexes, with the left vertical arrow surjective. Then $H^p(K^\bullet) \rightarrow H^p(J^\bullet)$ is surjective if and only if $\ker \delta_K^p \rightarrow \ker \delta_J^p$ is surjective.

Proof. The map $\operatorname{im} \delta_K^{p-1} \rightarrow \operatorname{im} \delta_J^{p-1}$ is surjective: any element α of $\operatorname{im} \delta_J^{p-1}$ lifts to J^{p-1} , then can lift to K^{p-1} , which then can map to K^p , which maps to α . Then apply the Snake Lemma 1.7.5 to

$$\begin{array}{ccccccc} 0 & \longrightarrow & \operatorname{im} \delta_K^{p-1} & \longrightarrow & \ker \delta_K^p & \longrightarrow & H^p(K^\bullet) \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & \operatorname{im} \delta_J^{p-1} & \longrightarrow & \ker \delta_J^p & \longrightarrow & H^p(J^\bullet) \longrightarrow 0. \end{array}$$

□

28.2.11. Proof of the Cohomology and Base Change Theorem 28.1.6. We focus on the complex near the p th step, and its restriction to the point $q \in X$:

$$\begin{array}{ccccc} K^{p-1} & \longrightarrow & K^p & \xrightarrow{\delta^p} & K^{p+1} & \text{free } B\text{-modules} \\ \downarrow & & \downarrow & & \downarrow & \\ K^{p-1} \otimes_B \kappa(q) & \longrightarrow & K^p \otimes_B \kappa(q) & \xrightarrow{\delta^p \otimes \kappa(q)} & K^{p+1} \otimes_B \kappa(q) \end{array}$$

By hypothesis, $\phi_q^p : (R^p \pi_* \mathcal{F})|_q \rightarrow H^p(X_q, \mathcal{F}|_{X_q})$ is surjective. By Lemma 28.2.10, this is equivalent to $(\ker \delta^p) \otimes \kappa(q) \rightarrow \ker(\delta^p \otimes \kappa(q))$ being surjective. By Exercise 28.2.C, this is equivalent to δ^p being strongly of constant rank near q , which implies (by Observation 28.2.8(ii)) that $\ker \delta^p$ is finite rank locally free, and the construction of $\ker \delta^p$ commutes with any base change in a neighborhood of q .

Now $H^p(K^\bullet) = \operatorname{coker}(K^{p-1} \rightarrow \ker \delta^p)$, i.e., we have

$$K^{p-1} \longrightarrow \ker \delta^p \longrightarrow H^p(K^\bullet) \longrightarrow 0.$$

Thus $H^p(K^\bullet)$ commutes with any base change (as tensor product is right exact). This completes the proof of part (i) of the Theorem.

For part (ii), consider again the map $K^{p-1} \rightarrow \ker \delta^p$ of finite rank locally free sheaves, whose cokernel is $H^p(K^\bullet)$. Now $H^p(K^\bullet)$ is locally free if and only if $K^{p-1} \rightarrow \ker \delta^p$ is strongly of constant rank, if and only if (since δ^p is strongly of constant rank) δ^{p-1} is strongly of constant rank, if and only if $H^{p-1}(K^\bullet) \rightarrow H^{p-1}(K^\bullet|_q)$ is surjective. This completes the proof of part (ii). □

28.2.12. ★ Removing Noetherian conditions.

It can be helpful to have versions of the theorems of §28.1 without Noetherian conditions; important examples come from moduli theory, and will be discussed

in the next section. Noetherian conditions can often be exchanged for finite presentation conditions. We begin with an extension of Exercise 10.3.H.

28.2.E. EXERCISE. Suppose $\pi: X \rightarrow \operatorname{Spec} B$ is a finitely presented morphism, and \mathcal{F} is a finitely presented sheaf on X . Show that there exists a base change diagram of the form

$$(28.2.12.1) \quad \begin{array}{ccc} \mathcal{F} & & \mathcal{F}' \\ \downarrow & & \downarrow \\ X & \xrightarrow{\sigma} & X' \\ \downarrow \pi & & \downarrow \pi' \\ \operatorname{Spec} B & \xrightarrow{\rho} & \operatorname{Spec} \mathbb{Z}[x_1, \dots, x_N]/I \end{array}$$

where N is some integer, $I \subset \mathbb{Z}[x_1, \dots, x_N]$, and π' is finitely presented (= finite type as the target is Noetherian, see §8.3.13), and a finitely presented (= coherent) quasicoherent sheaf \mathcal{F}' on X' with $\mathcal{F} \cong \sigma^* \mathcal{F}'$.

28.2.13. Properties of π' . (The ideal I appears in the statement of Exercise 28.2.E not because it is needed there, but to make the statement of this remark correct.) If π is proper, then diagram (28.2.12.1) can be constructed so that π' is also proper (using [Gr-EGA, IV₃.8.10.5]). Furthermore, if \mathcal{F} is flat over $\operatorname{Spec} B$, then (28.2.12.1) can be constructed so that \mathcal{F}' is flat over $\operatorname{Spec} \mathbb{Z}[x_1, \dots, x_N]/I$ (using [Gr-EGA, IV₃.11.2.6]). This requires significantly more work.

28.2.F. EXERCISE. Assuming the results stated in §28.2.13, prove the following results, with the “locally Noetherian” hypotheses removed, and “finite presentation” hypotheses added:

- (a) the constancy of Euler characteristic in flat families (Theorem 25.7.1, extended to the proper case as in §28.2.5);
- (b) the Semicontinuity Theorem 28.1.1;
- (c) Grauert’s Theorem 28.1.5 (you will have to show that $\mathbb{Z}[x_1, \dots, x_N]/I$ in (28.2.12.1) can be taken to be reduced); and
- (d) the Cohomology and Base Change Theorem 28.1.6.

28.2.14. Necessity of finite presentation conditions. The finite presentation conditions are necessary. There is a projective flat morphism to a connected target where the fiber dimension jumps. There is a finite flat morphism where the degree of the fiber is not locally constant. There is a projective flat morphism to a connected target where the fibers are curves, and the arithmetic genus is not constant. See [Stacks, tag 05LB] for the first example; the other two use the same idea.

28.3 Applying cohomology and base change to moduli problems

The theory of moduli relies on ideas of cohomology and base change. We explore this by examining two special cases of one of the primordial moduli spaces, the Hilbert scheme: the Grassmannian, and the fact that degree d hypersurfaces in

projective space are “parametrized” by another projective space (corresponding to degree d polynomials, see Remark 4.5.3).

As suggested in §25.1, the Hilbert functor $\mathrm{Hilb}_Y \mathbb{P}^n$ of \mathbb{P}_Y^n parametrizes finitely presented closed subschemes of \mathbb{P}_Y^n , where Y is an arbitrary scheme. More precisely, it is a contravariant functor sending the Y -scheme X to the set of finitely presented closed subschemes of $X \times_Y \mathbb{P}_Y^n = \mathbb{P}_X^n$ flat over X (and sending morphisms $X_1 \rightarrow X_2$ to pullbacks of flat families). An early achievement of Grothendieck was the construction of the Hilbert scheme, which can then be cleverly used to construct many other moduli spaces.

28.3.1. Theorem (Grothendieck). — $\mathrm{Hilb}_{\mathbb{Z}} \mathbb{P}^n$ is representable by a scheme locally of finite type.

(Grothendieck’s original argument is in [Gr5]. A readable construction is given in [Mu2], and in [FGIKNV, Ch. 5].)

28.3.A. EASY EXERCISE. Assuming Theorem 28.3.1, show that $\mathrm{Hilb}_Y \mathbb{P}^n$ is representable, by showing that it is represented by $\mathrm{Hilb}_{\mathbb{Z}} \mathbb{P}^n \times_{\mathbb{Z}} Y$. Thus the general case follows from the “universal” case of $Y = \mathbb{Z}$.

28.3.B. EXERCISE. Assuming Theorem 28.3.1, show that $\mathrm{Hilb}_{\mathbb{Z}} \mathbb{P}^n$ is the disjoint union of schemes $\mathrm{Hilb}_{\mathbb{Z}}^{p(m)} \mathbb{P}^n$, each one corresponding to finitely presented closed subschemes of $\mathbb{P}_{\mathbb{Z}}^n$ whose fibers have fixed Hilbert polynomial $p(m)$. Hint: Corollary 25.7.2.

28.3.2. Theorem (Grothendieck). — Each $\mathrm{Hilb}_{\mathbb{Z}}^{p(m)} \mathbb{P}^n$ is projective over \mathbb{Z} .

In order to get some feeling for the Hilbert scheme, we discuss two important examples, without relying on Theorem 28.3.1.

28.3.3. The Grassmannian.

We have defined the Grassmannian $G(k, n)$ twice before, in §7.7 and §17.7. The second time involved showing the representability of a (contravariant) functor (from *Sheaves* to *Sets*), of rank k locally free quotient sheaves of a rank n free sheaf.

We now consider a parameter space for a more geometric problem. The space will again be $G(k, n)$, but because we won’t immediately know this, we invent some temporary notation. Let $G'(k, n)$ be the contravariant functor (from *Schemes* to *Sets*) which assigns to a scheme B the set of *finitely presented* closed subschemes of \mathbb{P}_B^{n-1} , flat over B , whose fiber over any point $b \in B$ is a (linearly embedded) $\mathbb{P}_{\kappa(b)}^{k-1}$ in $\mathbb{P}_{\kappa(b)}^{n-1}$:

(28.3.3.1)

$$\begin{array}{ccc} X & \xrightarrow{\text{cl. subscheme}} & \mathbb{P}_B^{n-1} \\ \text{flat, f. pr.} \downarrow & \swarrow \pi & \\ B & & \end{array}$$

(This describes the map to *Sets*; you should think through how pullback makes this into a contravariant functor.)

28.3.4. Theorem. — The functor $G'(k, n)$ is represented by $G(k, n)$.

Translation: there is a natural bijection between diagrams of the form (28.3.3.1) (where the fibers are \mathbb{P}^{k-1} 's) and diagrams of the form (17.7.0.1) (the diagrams that $G(k, n)$ parametrizes, or represents).

One direction is notably easier. Suppose we are given a diagram of the form (17.7.0.1) over a scheme B ,

$$(28.3.4.1) \quad \mathcal{O}_B^{\oplus n} \twoheadrightarrow \mathcal{Q},$$

where \mathcal{Q} is locally free of rank k . Applying Proj_B to the Sym^\bullet construction on both $\mathcal{O}_B^{\oplus n}$ and \mathcal{Q} , we obtain a closed embedding

$$(28.3.4.2) \quad \begin{array}{ccc} \text{Proj}_B(\text{Sym}^\bullet \mathcal{Q}) & \hookrightarrow & \text{Proj}_B(\text{Sym}^\bullet \mathcal{O}_B^{\oplus n}) \\ & \searrow & \swarrow \\ & B & \end{array} \quad = \mathbb{P}^{n-1} \times B$$

(as, for example, in Exercise 18.2.H).

The fibers are linearly embedded \mathbb{P}^{k-1} 's (as base change, in this case to a point of B , commutes with the Proj construction, Exercise 18.2.E). Note that $\text{Proj}(\text{Sym}^\bullet \mathcal{Q})$ is flat and finitely presented over B , as it is a projective bundle. We have constructed a diagram of the form (28.3.3.1).

We now need to reverse this. The trick is to produce (28.3.4.1) from our geometric situation (28.3.3.1), and this is where cohomology and base change will be used.

Given a diagram of the form (28.3.3.1) (where the fibers are \mathbb{P}^{k-1} 's), consider the closed subscheme exact sequence for X :

$$0 \longrightarrow \mathcal{I}_X \longrightarrow \mathcal{O}_{\mathbb{P}_B^{n-1}} \longrightarrow \mathcal{O}_X \longrightarrow 0.$$

Tensor this with $\mathcal{O}_{\mathbb{P}_B^{n-1}}(1)$:

$$(28.3.4.3) \quad 0 \longrightarrow \mathcal{I}_X(1) \longrightarrow \mathcal{O}_{\mathbb{P}_B^{n-1}}(1) \longrightarrow \mathcal{O}_X(1) \longrightarrow 0.$$

Note that $\mathcal{O}_X(1)$ restricted to each fiber of π is $\mathcal{O}(1)$ on \mathbb{P}^{k-1} (over the residue field), for which all higher cohomology vanishes (§19.3).

28.3.C. EXERCISE. Show that $R^i \pi_* \mathcal{O}_X(1) = 0$ for $i > 0$, and $\pi_* \mathcal{O}_X(1)$ is locally free of rank k . Hint: use the Cohomology and Base Change Theorem 28.1.6. Either use the non-Noetherian discussion of §28.2.12 (which we haven't proved), or else just assume B is locally Noetherian.

28.3.D. EXERCISE. Show that the long exact sequence obtained by applying π_* to (28.3.4.3) is just a short exact sequence of locally free sheaves

$$0 \longrightarrow \pi_* \mathcal{I}_X(1) \longrightarrow \pi_* \mathcal{O}_{\mathbb{P}_B^{n-1}}(1) \longrightarrow \pi_* \mathcal{O}_X(1) \longrightarrow 0.$$

of ranks $n - k$, n , and k respectively, where the middle term is canonically identified with $\mathcal{O}_B^{\oplus n}$.

The surjection $\mathcal{O}_B^{\oplus n} \twoheadrightarrow \pi_* \mathcal{O}_X(1)$ is precisely a diagram of the sort we wished to construct, (17.7.0.1).

28.3.E. EXERCISE. Close the loop, by using these two “inverse” constructions to show that $G(k, n)$ represents the functor $G'(k, n)$.

28.3.5. Hypersurfaces.

Ages ago (in Remark 4.5.3), we informally said that hypersurfaces of degree d in \mathbb{P}^n are parametrized by a $\mathbb{P}^{\binom{n+d}{d}-1}$. We now make this precise. We work over a base \mathbb{Z} for suitable generality. You are welcome to replace \mathbb{Z} by a field of your choice, but by the same argument as in Easy Exercise 28.3.A, all other cases are obtained from this one by base change.

Define the contravariant functor $H_{d,n} : \text{Sch} \rightarrow \text{Sets}$ from schemes to sets as follows. To a scheme B , we associated the set of all closed subschemes $X \hookrightarrow \mathbb{P}_B^n$, flat and finitely presented over B , all of whose fibers are degree d hypersurfaces in \mathbb{P}^n (over the appropriate residue field). To a morphism $B_1 \rightarrow B_2$, we obtain a map $H_{d,n}(B_2) \rightarrow H_{d,n}(B_1)$ by pullback.

28.3.6. Proposition. — *The functor $H_{d,n}$ is represented by $\mathbb{P}^{\binom{n+d}{d}-1}$.*

As with the case of the Grassmannian, one direction is easy, and the other requires cohomology and base change.

28.3.F. EASY EXERCISE. Over $\mathbb{P}^{\binom{n+d}{d}-1}$, described a closed subscheme $\mathcal{X} \hookrightarrow \mathbb{P}^n \times \mathbb{P}^{\binom{n+d}{d}-1}$ that will be the universal hypersurface. Show that \mathcal{X} is flat and finitely presented over $\mathbb{P}^{\binom{n+d}{d}-1}$. (For flatness, you can use the local criterion of flatness on the source, Exercise 25.6.F, but it is possible to deal with it easily by working by hand.)

Thus given any morphism $B \rightarrow \mathbb{P}^{\binom{n+d}{d}-1}$, by pullback, we have a degree d hypersurface X over B (an element of $H_{d,n}(B)$).

Our goal is to reverse this process: from a degree d hypersurface $\pi: X \rightarrow B$ over B (an element of $H_{d,n}(B)$), we want to describe a morphism $B \rightarrow \mathbb{P}^{\binom{n+d}{d}-1}$.

Consider the closed subscheme exact sequence for $X \hookrightarrow \mathbb{P}_B^n$, twisted by $\mathcal{O}_{\mathbb{P}_B^n}(d)$:

$$(28.3.6.1) \quad 0 \longrightarrow \mathcal{I}_X(d) \longrightarrow \mathcal{O}_{\mathbb{P}_B^n}(d) \longrightarrow \mathcal{O}_X(d) \longrightarrow 0.$$

28.3.G. EXERCISE (CF. EXERCISE 28.3.C). Show that the higher pushforwards (by π) of each term of (28.3.6.1) is 0, and that the long exact sequence of pushforwards of (28.3.6.1) is

$$0 \longrightarrow \pi_* \mathcal{I}_X(d) \longrightarrow \pi_* \mathcal{O}_{\mathbb{P}_B^n}(d) \longrightarrow \pi_* \mathcal{O}_X(d) \longrightarrow 0.$$

where the middle term is free of rank $\binom{n+d}{d}$ (whose summands can be identified with degree d monomials in the projective variables x_1, \dots, x_n (see Exercise 9.3.J), and the left term $\pi_* \mathcal{I}_X(d)$ is locally free of rank 1 (basically, a line bundle).

(It is helpful to interpret the middle term $\mathcal{O}_B^{\oplus \binom{n+d}{d}}$ as parametrizing homogeneous degree d polynomials in $n+1$ variables, and the rank 1 subsheaf of $\pi_* \mathcal{I}_X(d)$ as “the equation of X ”. This will motivate what comes next.)

Taking the dual of the injection $\pi_* \mathcal{I}_X(d) \hookrightarrow \mathcal{O}_B^{\oplus \binom{n+d}{d}}$, we have a surjection

$$\mathcal{O}_B^{\oplus \binom{n+d}{d}} \twoheadrightarrow \mathcal{L}$$

from a free sheaf onto an invertible sheaf $\mathcal{L} = (\pi_* \mathcal{I}_X(d))^\vee$, which (by the universal property of projective space) yields a morphism $B \rightarrow \mathbb{P}^{\binom{n+d}{d}-1}$.

28.3.H. EXERCISE. Close the loop: show that these two constructions are inverses, thereby proving Proposition [28.3.6](#).

28.3.7. Remark. The proof of the representability of the Hilbert scheme shares a number of features of our arguments about the Grassmannian and the parameter space of hypersurfaces.

Depth and Cohen-Macaulayness

We now introduce the notion of depth. Depth is an algebraic rather than geometric concept, so we concentrate on developing some geometric sense of what it means. Most important is the geometric bound on depth by dimension of (closures of) associated points (Theorem 29.1.2). A central tool to understanding depth is the Koszul complex, but we avoid this approach, as we can prove what we need directly.

When the depth of a local ring equals its dimension, the ring is said to be Cohen-Macaulay, and this is an important way in which schemes can be “nice”. For example, regular local rings are Cohen-Macaulay (§29.2.5), as are regular embeddings in smooth varieties (Proposition 29.2.6). Cohen-Macaulayness will be key to the proof of Serre duality in Chapter 31, through the Miracle Flatness Theorem 29.2.11.

Another application of depth is Serre’s $R_1 + S_2$ criterion for normality (Theorem 29.3.2), which we will use to prove that regular schemes are normal (§29.3.5) without having to show that regular local rings are factorial (the Auslander-Buchsbaum Theorem, Fact 13.8.5), and to prove that regular embeddings in smooth schemes are normal if they are regular in codimension 1 (§29.3.3).

29.1 Depth

Recall the theory of regular sequences from §9.5.3.

29.1.1. Definition. Suppose (A, \mathfrak{m}) is a Noetherian local ring, and M is a finitely generated A -module. The **depth of M** (denoted $\text{depth } M$) is the length of the longest M -regular sequence with elements in \mathfrak{m} . (More generally, if R is a Noetherian ring, $I \subset R$ is an ideal, and M is a finitely generated R -module, then the **I -depth of M** , denoted $\text{depth}_I M$, is the length of the longest M -regular sequence with elements in I . We won’t need this.)

29.1.A. IMPORTANT EXERCISE. Suppose M is a nonzero finitely generated module over a Noetherian local ring (A, \mathfrak{m}) . Show that $\text{depth } M = 0$ if and only if every element of \mathfrak{m} is a zerodivisor of M if and only if \mathfrak{m} is an associated prime of M .

29.1.B. EXERCISE. Suppose M is a finitely generated module over a Noetherian local ring A . Show that $\text{depth } M \leq \dim \text{Supp } M$. In particular,

$$(29.1.1.1) \quad \text{depth } A \leq \dim A.$$

(We will improve this result in Theorem 29.1.6.)

At this point, it is hard to determine the depth of an A -module M . You can start trying to build an M -regular sequence by successively choosing x_1, x_2, \dots , but how do you know you have made the right choices to find the *longest* one? The happy answer is that you can't go wrong; this is the content of the next result. (We then describe how to find the depth of a module M in practice, in §29.1.4.)

29.1.2. Theorem. — Suppose M is a finitely generated module over a Noetherian local ring (A, \mathfrak{m}) . Then all maximal M -regular sequences contained in \mathfrak{m} (M -regular sequences x_1, \dots, x_n in \mathfrak{m} that cannot be extended to x_1, \dots, x_{n+1}) have the same length. Thus the depth of M is the length of any maximal M -regular sequence.

We prove Theorem 29.1.2 by giving a cohomological criterion, Theorem 29.1.3, for a regular sequence to be maximal. (You will then prove Theorem 29.1.2 in Exercise 29.1.D.) We will also use this criterion to give a better bound on the depth in Theorem 29.1.6. Theorem 29.1.2 is the key technical result of this Chapter. An important moral is that depth should be understood as a “cohomological” property.

29.1.3. Theorem (cohomological criterion for existence of regular sequences). — Suppose (A, \mathfrak{m}) is a Noetherian local ring, and M is a finitely generated A -module. The following are equivalent.

- (i) For every finitely generated A -module N with $\text{Supp } N = \{\mathfrak{m}\}$ (i.e., N has finite length), $\text{Ext}_A^i(N, M) = 0$ for all $i < n$.
- (ii) $\text{Ext}_A^i(A/\mathfrak{m}, M) = 0$ for all $i < n$.
- (iii) There exists an M -regular sequence in \mathfrak{m} of length n .

This result can be extended in various ways, see for example [Mat1, Thm. 28].

Proof. Clearly (i) implies (ii).

29.1.C. EXERCISE. Prove that (ii) implies (i). Hint: apply Exercise 6.5.P to N , so that it admits a filtration with each subquotient isomorphic to A/\mathfrak{m} . Use induction on the length of the filtration, and the long exact sequence for $\text{Ext}_A^i(\cdot, M)$.

Proof that (iii) implies (ii). The case $n = 0$ is vacuous. We inductively prove the result for all n . Suppose (iii) is satisfied, where $n \geq 1$, and assume that we know (ii) for “all smaller n ”. Choose a regular sequence x_1, \dots, x_n of length n . Then x_1 is a non-zerodivisor on M , so we have an exact sequence

$$(29.1.3.1) \quad 0 \longrightarrow M \xrightarrow{\times x_1} M \longrightarrow M/x_1 M \longrightarrow 0.$$

Then $M/x_1 M$ has a regular sequence x_2, \dots, x_{n-1} of length $n-1$, so by the inductive hypothesis, $\text{Ext}_A^i(A/\mathfrak{m}, M/x_1 M) = 0$ for $i < n-1$. Taking the Ext long exact sequence for $\text{Ext}_A^i(A/\mathfrak{m}, \cdot)$ for (29.1.3.1), we find that

$$\text{Ext}_A^i(A/\mathfrak{m}, M) \xrightarrow{\times x_1} \text{Ext}_A^i(A/\mathfrak{m}, M)$$

is an injection for $i < n$. Now $\text{Ext}_A^i(A/\mathfrak{m}, M)$ can be computed by taking an injective resolution of M , and applying $\text{Hom}_A(A/\mathfrak{m}, \cdot)$. Hence as x_1 lies in \mathfrak{m} (and thus annihilates A/\mathfrak{m}), multiplication by x_1 is the zero map. Thus (ii) holds for n as well.

Proof that (ii) implies (iii). The case $n = 0$ is vacuous.

We deal next with the case $n = 1$, by showing the contrapositive. Assume that there are no non-zero-divisors in \mathfrak{m} on M , so by Exercise 29.1.A, \mathfrak{m} is an associated prime of M . Thus (§6.5.7) we have an injection $A/\mathfrak{m} \hookrightarrow M$, yielding $\text{Hom}_A(A/\mathfrak{m}, M) \neq 0$ as desired.

We now inductively prove the result for all $n > 1$. Suppose (ii) is satisfied, where $n \geq 2$, and assume that we know (iii) for “all smaller n ”. Then by the case $n = 1$, there exists a non-zero-divisor x_1 on M , so we have a short exact sequence (29.1.3.1). A portion of the Ext long exact sequence for $\text{Ext}_A^i(A/\mathfrak{m}, \cdot)$ for (29.1.3.1) is

$$\text{Ext}_A^i(A/\mathfrak{m}, M) \longrightarrow \text{Ext}_A^i(A/\mathfrak{m}, M/x_1M) \longrightarrow \text{Ext}_A^{i+1}(A/\mathfrak{m}, M).$$

By assumption, both $\text{Ext}_A^i(A/\mathfrak{m}, M)$ and $\text{Ext}_A^{i+1}(A/\mathfrak{m}, M)$ are 0 for $i < n - 1$, so $\text{Ext}_A^i(A/\mathfrak{m}, M/x_1M) = 0$ for $i < n - 1$, so by the inductive hypothesis, we have an (M/x_1M) -regular sequence x_2, \dots, x_n of length $n - 1$ in \mathfrak{m} . Adding x_1 to the front of this sequence, we are done. \square

29.1.D. EXERCISE. Prove Theorem 29.1.2. Hint: by Theorem 29.1.3 (notably, the equivalence of (ii) and (iii)), you have control of how long an M -regular sequence in \mathfrak{m} can be. Use the criterion, and the long exact sequence used in the proof of Theorem 29.1.3, to show that any M -regular sequence in \mathfrak{m} can be extended to this length.

29.1.E. EXERCISE. Suppose M is a finitely generated module over a Noetherian local ring (A, \mathfrak{m}) . If x is a non-zero-divisor in \mathfrak{m} , show that $\text{depth}(M/xM) = \text{depth } M - 1$. Hint: Theorem 29.1.2.

29.1.4. Finding the depth of a module. We can now compute the depth of M by successively finding non-zero-divisors, as follows. Is there a non-zero-divisor x on M in \mathfrak{m} ?

- (a) If not, then $\text{depth } M = 0$ (Exercise 29.1.A).
- (b) If so, then choose any such x , and (using the previous exercise) repeat the process with M/xM .

The process must terminate by Exercise 29.1.B.

29.1.F. IMPORTANT EXERCISE. Suppose (A, \mathfrak{m}) is a dimension d regular local ring. Show that $\text{depth } A = d$. Hint: Exercise 13.2.B.

29.1.G. EXERCISE (CF. EXERCISE 25.4.I). Suppose $X = \text{Spec } R$, where

$$R = k[w, x, y, z]/(wy, wz, xy, xz),$$

the union of two co-ordinate two-planes in \mathbb{A}_k^4 meeting at the origin. Show that the depth of the local ring of X at the origin is 1. Hint: Show that $w - y$ is not a zero-divisor, and that $R/(w - y)$ has an embedded point at the origin.

29.1.5. Depth is bounded by the dimension of associated prime ideals.

Theorem 29.1.3 can be used to give an important improvement of the bound (29.1.1.1) on depth by the dimension:

29.1.6. Theorem. — *The depth of a module M is at most the smallest $\dim A/\mathfrak{p}$ as \mathfrak{p} runs over the associated prime ideals of M .*

(The example of two planes meeting at a point in Exercise 29.1.G shows that this bound is not sharp.) The key step in the proof of Theorem 29.1.6 is the following result of Ischebeck.

29.1.7. Lemma. — Suppose (A, \mathfrak{m}) is a Noetherian local ring, and M and N are nonzero finitely generated A -modules. Then $\text{Ext}_A^i(N, M) = 0$ for $i < \text{depth } M - \dim \text{Supp } N$.

Proof. Consider the following statements.

(*general*_r) Lemma 29.1.7 holds for $\dim \text{Supp } N \leq r$.

(*prime*_r) Lemma 29.1.7 holds for $\dim \text{Supp } N \leq r$ and $N = A/\mathfrak{p}$ for some prime \mathfrak{p} .

Note that (*general*₀) (and hence (*prime*₀)) is true, as in this case $\text{Supp } N = \{[\mathfrak{m}]\}$, and the result follows from Theorem 29.1.3 (from “(iii) implies (i)”).

29.1.H. EXERCISE. Show that (*prime*_r) implies (*general*_r). Hint: For a short exact sequence $0 \rightarrow M' \rightarrow M \rightarrow M'' \rightarrow 0$, $\text{Supp } M = \text{Supp } M' \cup \text{Supp } M''$ (Exercise 6.5.B).

We conclude the proof by showing that (*general*_{r-1}) implies (*prime*_r) for $r \geq 1$. Fix a prime $\mathfrak{p} \subset A$ with $\dim A/\mathfrak{p} = r$. Since $\dim A/\mathfrak{p} > 0$, we can choose $x \in \mathfrak{m} \setminus \mathfrak{p}$. Consider the exact sequence

$$(29.1.7.1) \quad 0 \longrightarrow A/\mathfrak{p} \xrightarrow{\times x} A/\mathfrak{p} \longrightarrow A/(\mathfrak{p} + (x)) \longrightarrow 0,$$

noting that $\dim \text{Supp } A/(\mathfrak{p} + (x)) \leq r-1$ (Do you see why?). Then the Ext long exact sequence obtained by applying $\text{Hom}_A(\cdot, M)$ to (29.1.7.1), along with the vanishing of $\text{Ext}_A^i(A/(\mathfrak{p} + (x)), M)$ for $i < \text{depth } M - r + 1$ (by (*general*_{r-1})), implies that

$$\text{Ext}_A^i(A/\mathfrak{p}, M) \xrightarrow{\times x} \text{Ext}_A^i(A/\mathfrak{p}, M)$$

is an isomorphism for $i < \text{depth } M - r$. But $\text{Ext}_A^i(A/\mathfrak{p}, M)$ is a finitely generated A -module (Exercise 24.2.F), so by Nakayama’s Lemma 8.2.8, $\text{Ext}_A^i(A/\mathfrak{p}, M) = 0$ for $i < \text{depth } M - r$. \square

29.1.I. EASY EXERCISE. Prove Theorem 29.1.6. Hint: if $\mathfrak{p} \in \text{Ass}(M)$, then $\text{Hom}_A(A/\mathfrak{p}, M) \neq 0$ (§6.5.7).

29.2 Cohen-Macaulay rings and schemes

29.2.1. Definition. A Noetherian local ring (A, \mathfrak{m}) is **Cohen-Macaulay** (or often **CM** for short) if $\text{depth } A = \dim A$, i.e., if equality holds in (29.1.1.1). (One may define **Cohen-Macaulay module** similarly, but we won’t need this concept.) A locally Noetherian scheme is **Cohen-Macaulay** if all of its local rings are Cohen-Macaulay. We won’t use this language, but a ring A is a **Cohen-Macaulay ring** if all of its localizations $A_{\mathfrak{p}}$ at all primes \mathfrak{p} are Cohen-Macaulay, or equivalently, if $\text{Spec } A$ is Cohen-Macaulay. (Fact 29.2.14 below, that Cohen-Macaulayness of a local ring is preserved by localization, ensures that there is no contradiction between this definition, and the definition of a Cohen-Macaulay local ring.)

29.2.A. EXERCISE. Show that every locally Noetherian scheme of dimension 0 is Cohen-Macaulay. Show that a locally Noetherian scheme of dimension 1 is Cohen-Macaulay if and only if it has no embedded points.

29.2.2. (Counter)example. Let X be the example of Exercise 29.1.G — two planes meeting at a point. By Exercise 29.1.G, X is not Cohen-Macaulay.

29.2.B. EXERCISE. Suppose A is Cohen-Macaulay Noetherian local ring. Use Theorem 29.1.6 to show that $\text{Spec } A$ is equidimensional, and has no embedded points. (It is not true that Noetherian local rings of pure dimension having no embedded prime ideals are Cohen-Macaulay, see Example 29.2.2.)

29.2.3. Theorem (slicing criterion for Cohen-Macaulayness). — Suppose (A, \mathfrak{m}) is a Noetherian local ring, and $x \in \mathfrak{m}$ is a non-zerodivisor. Then (A, \mathfrak{m}) is Cohen-Macaulay if and only if $(A/x, \mathfrak{m})$ is Cohen-Macaulay.

Compare this to the slicing criteria for regularity and flatness (Exercise 13.2.C and Theorem 25.6.5 respectively).

29.2.4. Geometric interpretation of the slicing criterion. Suppose X is a locally Noetherian scheme, and D is an effective Cartier divisor. If X is Cohen-Macaulay, then so is D . If D is Cohen-Macaulay, then X is Cohen-Macaulay at the points of D .

29.2.C. EXERCISE. Prove Theorem 29.2.3, using Theorem 29.1.2, the fact that maximal regular sequences (in \mathfrak{m}) all have the same length.

29.2.D. EXERCISE. Show that if (A, \mathfrak{m}) is Cohen-Macaulay, then a set of elements $x_1, \dots, x_r \in \mathfrak{m}$ is a regular sequence (for A) if and only if $\dim A/(x_1, \dots, x_r) = \dim A - r$.

29.2.5. By Exercise 29.1.F, regular local rings (Noetherian by definition) are Cohen-Macaulay. In particular, as smooth schemes over a field k are regular (Theorem 13.2.10(b)), we see that smooth k -schemes are Cohen-Macaulay. (In combination with Exercise 29.2.B, this explains why effective Cartier divisors on smooth varieties have no embedded points, which justifies a comment made in Aside 21.2.2.) Combining this with Exercise 29.2.D or §29.2.4, we have the following.

29.2.6. Proposition. — Regular embeddings in smooth k -schemes are Cohen-Macaulay.

29.2.7. As a consequence of Proposition 29.2.6 and the fact that Cohen-Macaulay schemes have no embedded points (Exercise 29.2.B), we see that regular embeddings in smooth k -schemes (in \mathbb{A}_k^n or \mathbb{P}_k^n for example) have no embedded points, generalizing Exercise 6.5.E (the hypersurface in \mathbb{A}_k^n case). This is not clear without the theory of Cohen-Macaulayness! (This fact was used in Exercise 26.2.A, see §26.2.1.)

29.2.8. Alternate definition of Cohen-Macaulayness. The slicing criterion (Theorem 29.2.3) gives an enlightening alternative inductive definition of Cohen-Macaulayness in terms of effective Cartier divisors, in the spirit of the method of §29.1.4 for computing depth. Suppose as before that (A, \mathfrak{m}) is a Noetherian local ring.

- (i) If $\dim A = 0$, then A is Cohen-Macaulay (Exercise 29.2.A).

- (iia) If $\dim A > 0$, and every element of \mathfrak{m} is a zerodivisor, then A is *not* Cohen-Macaulay (by Exercise 29.1.A).
- (iib) Otherwise, choose *any* non-zerodivisor x in \mathfrak{m} . Then A is Cohen-Macaulay if and only if $A/(x)$ (necessarily of dimension $\dim A - 1$ by Krull's Principal Ideal Theorem 12.3.3) is Cohen-Macaulay.

The following example could have been stated (but not proved) before we knew any algebraic geometry at all. (We work over \mathbb{C} rather than over an arbitrary field simply to ensure that the statement requires as little background as possible.)

29.2.E. FUN EXERCISE (MAX NOETHER'S AF + BG THEOREM). Suppose $f, g \in \mathbb{C}[x_0, x_1, x_2]$ are two homogeneous polynomials, cutting out two curves in $\mathbb{P}_{\mathbb{C}}^2$ that meet “transversely”, i.e., at a finite number of reduced points. Suppose $h \in \mathbb{C}[x_0, x_1, x_2]$ is a homogeneous polynomial vanishing at these points. Show that $h \in (f, g)$. Hint: show that the intersection of the affine cones $V(f)$ and $V(g)$ in \mathbb{A}^3 has no embedded points. (This problem is quite nontrivial to do without the theory developed in this chapter! As a sign that this is subtle: you can easily construct *three* quadratics $e, f, g \in \mathbb{C}[x_0, x_1, x_2]$ cutting out precisely the two points $[1, 0, 0]$ and $[0, 1, 0] \in \mathbb{P}_{\mathbb{C}}^2$, yet the line $z = 0$ is not in the ideal (e, f, g) for degree reasons.)

29.2.9. Miracle flatness.

We conclude with a remarkably simple and useful criterion for flatness, which we shall use in the proof of Serre duality. The main content is the following algebraic result.

29.2.10. Miracle Flatness Theorem (algebraic version). — Suppose $\phi: (B, \mathfrak{n}) \rightarrow (A, \mathfrak{m})$ is a (local) homomorphism of Noetherian local rings, such that A is Cohen-Macaulay, and B is regular, and $A/\mathfrak{n}A = A \otimes_B (B/\mathfrak{n})$ (the ring corresponding to the fiber) has pure dimension $\dim A - \dim B$. Then ϕ is flat.

Proof. We prove Theorem 29.2.10 by induction on $\dim B$. If $\dim B = 0$, then B is a field (Exercise 13.2.A), so the result is immediate, as everything is flat over a field (Exercise 25.2.A(a)). Assume next that $\dim B > 0$, and we have proved the result for all “ B of smaller dimension”. Choose $x \in \mathfrak{n} - \mathfrak{n}^2$, so B/x is a regular local ring of dimension $\dim B - 1$. Then

$$\begin{aligned} \dim A/xA &\leq \dim B/(x) + \dim A/\mathfrak{n}A && \text{(Key Exercise 12.4.A)} \\ &= \dim B - 1 + \dim A/\mathfrak{n}A && \text{(Krull's Principal Ideal Theorem 12.3.3)} \\ &= \dim A - 1 && \text{(by hypothesis of Theorem 29.2.10).} \end{aligned}$$

By Krull's Principal Ideal Theorem 12.3.3, $\dim A/xA \geq \dim A - 1$, so we have $\dim A/xA = \dim B/(x) + \dim A/\mathfrak{n}A = \dim A - 1$. By Exercise 29.2.D, A/xA is a Cohen-Macaulay ring, and x is a non-zerodivisor on A . The inductive hypothesis then applies to $(B/(x), \mathfrak{n}) \rightarrow (A/xA, \mathfrak{m})$, so A/xA is flat over $B/(x)$. Then by the local slicing criterion for flatness (Theorem 25.6.5), $B \rightarrow A$ is flat, as desired. \square

29.2.11. Miracle Flatness Theorem. — Suppose $\pi: X \rightarrow Y$ is a morphism of equidimensional finite type k -schemes, where X is Cohen-Macaulay, Y is regular, and the fibers of π have dimension $\dim X - \dim Y$. Then π is flat.

29.2.F. EXERCISE. Prove the Miracle Flatness Theorem 29.2.11. (Do not forget that schemes usually have non-closed points!)

The geometric situation in the Miracle Flatness Theorem 29.2.11 is part of the following pretty package.

29.2.G. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a map of locally finite type k -schemes, where both X and Y are equidimensional, and Y is regular. Show that if any two of the following hold, then the third does as well:

- (i) π is flat of relative dimension $\dim X - \dim Y$.
- (ii) X is Cohen-Macaulay.
- (iii) Every fiber X_y is Cohen-Macaulay of pure dimension $\dim X - \dim Y$.

Hint: if $\phi: B \rightarrow A$ is a flat ring map, then ϕ sends non-zerodivisors to non-zerodivisors (Observation 25.2.2).

The statement of Exercise 29.2.G can be improved, at the expense of killing the symmetry. In the implication (ii) and (iii) imply (i), the Cohen-Macaulay hypotheses in (iii) are not needed.

29.2.12. Example. As an example of Exercise 29.2.G in action, we consider the example of two planes meeting a point, continuing the notation of Exercise 29.1.G. Consider the morphism $\pi: X \rightarrow Y := \mathbb{A}^2$ given by $(w - y, x - z)$. Then Y is regular (by the Smoothness-Regularity Comparison Theorem 13.2.10(b)). But X is not Cohen-Macaulay (Exercise 29.1.G) and π is not flat (Exercise 25.4.I), so by Exercise 29.2.G, you can use either of these to prove the other.

29.2.13. Fancy properties of Cohen-Macaulayness.

We mention a few additional properties of Cohen-Macaulayness without proof. They are worth seeing, but we will not use them.

If X is a locally Noetherian scheme, the locus of Cohen-Macaulay points is open, see [Mat2, Thm 24.5] (and also [Stacks, tag 00RD]). In particular:

29.2.14. Fact (cf. Fact 13.8.5). — *Any localization of a Cohen-Macaulay Noetherian local ring at a prime ideal is also Cohen-Macaulay.*

(See [Stacks, tag 00NB], [E, Prop. 18.8], [Mat2, Thm. 17.3(iii)], or [Mat1, Thm. 30] for a direct proof.) In particular, for quasicompact schemes, Cohen-Macaulayness may be checked at closed points. The fact that Cohen-Macaulayness is preserved by localization can be used to quickly show that Cohen-Macaulay local rings are catenary (see [E, Cor. 18.10], [Mat1, Thm. 31(ii)], or [Mat2, Thm. 17.9]). A ring A is Cohen-Macaulay if and only if $A[x]$ is Cohen-Macaulay ([E, Prop 18.9], [Mat2, Thm 17.7], [Mat1, Thm. 33]), if and only if $A[[x]]$ is [Mat2, p. 137]. A local ring is Cohen-Macaulay if and only if its completion is Cohen-Macaulay ([E, Prop. 18.8], [Mat2, Thm 17.5]).

29.3 ★ Serre's $R_1 + S_2$ criterion for normality

The notion of depth yields a useful criterion for normality, due to Serre.

29.3.1. Definition. Suppose A is a Noetherian ring, and $k \in \mathbb{Z}^{\geq 0}$. We say A has **property R_k** (A is **regular in codimension $\leq k$** , or more sloppily, A is **regular in codimension k**) if for every prime $\mathfrak{p} \subset A$ of codimension at most k , $A_{\mathfrak{p}}$ is regular. (In light of Fact 13.8.2, a Noetherian local ring is regular and only if it has property R_k for all k .)

We say A has **property S_k** if for every prime $\mathfrak{p} \subset A$, the local ring $A_{\mathfrak{p}}$ has depth at least $\min(k, \dim A_{\mathfrak{p}})$ — “the local rings are Cohen-Macaulay up until codimension k , and have depth at least k thereafter”. (In light of Fact 29.2.14, a Noetherian local ring is Cohen-Macaulay if and only if it has property S_k for all k .)

29.3.A. EASY EXERCISE. Note that a Noetherian ring A trivially always has property S_0 . Show that a Noetherian ring A has property R_0 if $\text{Spec } A$ is “generically reduced”: it is reduced at the generic point of each of its irreducible components. Show that a Noetherian ring A has property S_1 if and only if $\text{Spec } A$ has no embedded points. (Possible hint: Exercise 29.1.A.)

29.3.B. EXERCISE. Show that a Noetherian ring A has no nilpotents ($\text{Spec } A$ is reduced) if and only if it has properties R_0 and S_1 . (Hint: show that a Noetherian scheme is reduced if and only if each irreducible component is generically reduced, and it has no embedded points.)

Incrementing the subscripts in Exercise 29.3.B yields Serre’s criterion.

29.3.2. Theorem (Serre’s criterion for normality). — *A Noetherian ring A is normal if and only if it has properties R_1 and S_2 .*

(Recall from §5.4.1 that a Noetherian ring A is normal if its localizations $A_{\mathfrak{p}}$ at all primes \mathfrak{p} are integrally closed domains.) Thus failure of normality can have two possible causes: it can be failure of R_1 (something we already knew, from the equivalence of (a) and (g) in Theorem 13.5.8), or it can be the more subtle failure of S_2 . Examples of varieties satisfying R_1 but not S_2 are given in Example 29.3.4 (two planes meeting at a point) and Exercise 29.3.D (the pinched plane).

29.3.3. Applications. As usual, before giving a proof, we give some applications to motivate the result. First, it implies that Cohen-Macaulay schemes are normal if and only if they are regular in codimension 1. Thus to check normality of hypersurfaces (or more generally, regular embeddings) in \mathbb{P}_k^n (or more generally, in any smooth variety), it suffices to check that their singular locus has codimension greater than one. (You should think through the details of why these statements are true.) In particular, this gives a new (more complicated) proof of Exercise 5.4.I(b) and (c).

29.3.C. EXERCISE (PRACTICE WITH THE CONCEPT). Show that two-dimensional normal varieties are Cohen-Macaulay.

29.3.4. Example: two planes meeting at a point. The variety X of Exercise 29.1.G (two planes meeting at a point) is not normal (why?), but it is regular away from the origin. This implies that X does not have property S_2 (and hence is not Cohen-Macaulay), without the algebraic manipulations of Exercise 29.1.G.

We already knew that this example was not Cohen-Macaulay, but the same idea can show that the pinched plane $\text{Spec } k[x^3, x^2, xy, y]$ (appearing in Exercise 13.5.I) is not Cohen-Macaulay. Because of the “extrinsic” description of the ring, it is difficult to do this in another way.

29.3.D. EXERCISE: THE PINCHED PLANE IS NOT COHEN-MACAULAY. Let A be the subring $k[x^3, x^2, xy, y]$ of $k[x, y]$. Show that A is not Cohen-Macaulay. Hint: Exercise 13.5.I showed that A is not integrally closed.

29.3.5. Regular local rings are integrally closed. Serre’s criterion can be used to show that regular local rings are integrally closed without going through the hard Fact 13.8.5 (the Auslander-Buchsbaum Theorem) that they are unique factorization domains. Regular local rings are Cohen-Macaulay (§29.2.5), and (by Exercise 29.3.B) regular in codimension 0 as they are integral domains (Theorem 13.2.13), so we need only show that regular local rings are regular in codimension 1. We can invoke a different hard Fact 13.8.2 that localizations of regular local rings are again regular, but at least we have shown this for localizations of finitely generated algebras over a perfect field (see Theorem 13.8.3 and Exercise 22.3.G).

29.3.6. Caution. As is made clear by the following exercise, the condition S_2 is a condition on *all* prime ideals, not just those of codimension at most 2.

29.3.E. EXERCISE. Give an example of a variety satisfying R_1 , and Cohen-Macaulay at all points of codimension at most 2, which is not normal.

29.3.7. Proof of Theorem 29.3.2. The proof of Serre’s criterion will take us until the end of this section.

29.3.8. Normal implies $R_1 + S_2$.

Suppose first that A is normal. It is reduced, and thus satisfies (R_0) and S_1 by Exercise 29.3.B. It satisfies R_1 from the equivalence of (a) and (g) in Theorem 13.5.8. All that is left is to verify property S_2 . We thus must show that if (A, \mathfrak{m}) is an integrally closed Noetherian local ring (implicitly, integral domain) of dimension greater than 1, then A has depth at least 2.

Choose any nonzero $x \in \mathfrak{m}$. (Don’t forget that normality implies that the local ring A is an integral domain.) We wish to show that $\text{depth } A/(x) > 0$. Assume otherwise that $\text{depth } A/(x) = 0$, so (by Exercise 29.1.A) \mathfrak{m} is an associated prime of $A/(x)$. Thus there is some nonzero $y \in A/(x)$ (i.e., $y \notin (x)$) with $y\mathfrak{m} = 0$ in $A/(x)$, i.e., *multiplying any element $z \in \mathfrak{m}$ by y gives you a multiple of x* . Let $w = y/x \in K(A)$. We will show that $w \in A$, and thus that $y \in (x)$, yielding a contradiction.

For each prime ideal $\mathfrak{p} \subset A$ of codimension 1, choose any $z \in \mathfrak{m} \setminus \mathfrak{p}$. Then $zw = zy/x \in (y/x)\mathfrak{m} \subset A$ by the italicized statement in the previous paragraph, so $w \in A_{\mathfrak{p}}$. Thus $w \in A_{\mathfrak{p}}$ for all codimension 1 prime ideals $\mathfrak{p} \subset A$, so by Algebraic Hartogs’s Lemma 12.3.15 (using the hypothesis that A is integrally closed), $w \in A$ as desired.

29.3.9. $R_1 + S_2$ implies normal: the integral domain case. Normality is a stalk-local condition by definition (§5.4.1), so it suffices to show:

(†) if (A, \mathfrak{m}) is a dimension d Noetherian local ring satisfying R_1 and S_2 then it is an integrally closed integral domain.

(As A satisfies R_0 and S_1 , A is reduced by Exercise 29.3.B. Also, d is finite by Exercise 12.3.H.)

The case in which A is an integral domain is notably easier, and may help motivate the statement of the criterion, so we deal with this case first, and leave the general case as something to read when you have too much time on your hands. Note that this case suffices for many of the consequences we discuss, notably the normality of regular schemes (§29.3.5).

Suppose now that A is an integral domain. Suppose $x \in K(A)$ is integral over A . We must prove that $x \in A$. Write $x = f/g$, where $f, g \in A$. We wish to show that $f \in (g)$, or equivalently that f restricts to 0 on the closed subscheme $W := V(g)$ of $\text{Spec } A$.

29.3.F. EXERCISE. Show that $W := V(g)$ has no embedded points. Translation: $A/(g)$ satisfies S_1 .

Let η_1, \dots, η_m be the generic points of the irreducible components of W , corresponding to prime ideals $\mathfrak{p}_1, \dots, \mathfrak{p}_m$ of A . By Exercise 29.3.F, η_1, \dots, η_m are all the associated points of W , so the natural map

$$(29.3.9.1) \quad A/(g) \rightarrow \prod_{i=1}^m \mathcal{O}_{W, \eta_i} = \prod_{i=1}^m (A/(g))_{\mathfrak{p}_i}$$

is an injection, by Exercise 6.5.I.

The ring $A_{\mathfrak{p}_i}$ is regular in codimension 1, so by parts (a) and (g) of Theorem 13.5.8, $A_{\mathfrak{p}_i}$ is integrally closed. Now x is integral over A and thus over $A_{\mathfrak{p}_i}$, so by the integral closedness of $A_{\mathfrak{p}_i}$, $x \in A_{\mathfrak{p}_i}$. Hence f is 0 in each $A_{\mathfrak{p}_i}/(g) = (A/(g))_{\mathfrak{p}_i}$. From the injectivity of (29.3.9.1), we have that f is 0 on W as desired.

29.3.10. $\star\star R_1 + S_2$ implies normal: the general case. We prove (†) by induction on d .

29.3.G. EXERCISE. Prove the desired result if $d = 0$ and $d = 1$.

Our next step is to show that A is integrally closed in its total fraction ring, the localization of A at the multiplicative subset of non-zero-divisors (defined in §6.5.33). Suppose we have $f, g \in A$, with g a non-zero-divisor, such that f/g satisfies the monic equation

$$(29.3.10.1) \quad (f/g)^n + \sum_{i=1}^n a_i (f/g)^{n-i} = 0$$

with $a_i \in A$. We prove $f \in (g)$ by induction on d . The cases $d = 0$ and 1 follow from Exercise 29.3.G. Assume now that $d > 1$, and that the result is known for “smaller d ”.

As $\text{depth } A \geq 2$, $\text{depth } A/(g) \geq 1$, so there is a non-zero-divisor t on $A/(g)$ with $t \in \mathfrak{m}$. Consider the localization of (29.3.10.1) at all the prime ideals corresponding to points of $D(t) \subset \text{Spec } A$ (those prime ideals not containing t), each of which has codimension $< d$. By the inductive hypothesis, $f \in (g)$ in each of these localizations, so f is zero in $(A/(g))_t$. But as t is a non-zero-divisor on $A/(g)$, the map $A/(g) \rightarrow (A/(g))_t$ is an injection (Exercise 1.3.C), so f is zero in $A/(g)$, as desired. We have thus shown that A is integrally closed in its total fraction ring.

29.3.11. Lemma. — Suppose R is a reduced ring with finitely many minimal prime ideals $\mathfrak{p}_1, \dots, \mathfrak{p}_n$. If R is integrally closed in its total fraction ring, then R is a finite product of integral domains.

29.3.12. Proof of Lemma 29.3.11. The \mathfrak{p}_i are the associated prime ideals of R , by Exercise 6.5.R, so the natural map $\phi: R \rightarrow \prod_{i=1}^n R_{\mathfrak{p}_i}$ is an inclusion (by Exercise 6.5.I).

29.3.H. EXERCISE. Show that the map ϕ identifies the total fraction ring of R with $\prod_{i=1}^n R_{\mathfrak{p}_i}$.

Suppose e_i is the i th idempotent of the total fraction ring $\prod R_{\mathfrak{p}_i}$ (i.e., $(0, \dots, 0, 1, 0, \dots, 0)$, where the 1 is in the i th spot). Since R is integrally closed in its total fraction ring, it contains each of the e_i .

29.3.I. EXERCISE. Conclude the proof of Lemma 29.3.11, by describing $R = \prod_{i=1}^n Re_i$, and showing that each Re_i is an integrally closed integral domain. Possible hint: Remark 3.6.3. \square

We now return to our proof of Serre's criterion. By Lemma 29.3.11, we now know that A is a product of integral domains. But $\text{Spec } A$ is connected (A is a local ring!), so A is an integral domain. \square

CHAPTER 30

Power series and the Theorem on Formal Functions

30.1 Introduction

Power series are a central tool in analytic geometry. Their analog in algebraic geometry, completion, is similarly useful. We will only touch on some aspects of the subject.

In §30.2, we deal with some algebraic preliminaries. In §30.3, we use completions to (finally) give a definition of various singularities, such as nodes. (We won't use these definitions in what follows.) In §30.4, we state the main technical result of the chapter, the Theorem on Formal Functions 30.4.2. The subsequent three sections give applications. In §30.5, we prove Zariski's Connectedness Lemma 30.5.1 and the Stein Factorization Theorem 30.5.3. In §30.6, we prove a commonly used version of (Grothendieck's version of) Zariski's Main Theorem 30.6.1; we rely on §30.5. In §30.7, we prove Castelnuovo's Criterion for contracting (-1) -curves, which we used in Chapter 27. The proof of Castelnuovo's Criterion also uses §30.5. Finally, in §30.8, we prove the Theorem on Formal Functions 30.4.2.

There are deliberately many small sections in this chapter, so you can see that they tend not to be as hard at they look, with the exception of the proof of the Theorem on Formal Functions itself, and possibly Theorem 30.2.6 relating completion to exactness and flatness.

30.2 Algebraic preliminaries

Suppose A is a ring, and $J_1 \supset J_2 \supset \cdots$ is a decreasing sequence of ideals. Then we may take the limit $\lim A/J_n$. This ring is often denoted \hat{A} , and is called the **completion** of A (for this sequence of ideals); the sequence of ideals is left implicit. The most important case of completion is if $J_n = I^n$, where I is an ideal of A . The limit $\lim A/I^n$ is called the **I -adic completion of A** , or the **completion of A along I or at I** .

30.2.1. Example. We define $k[[x_1, \dots, x_n]]$ as the completion of $k[x_1, \dots, x_n]$ at the maximal ideal (x_1, \dots, x_n) . This is the ring of **formal power series** in n variables over k .

The p -adic numbers (Example 1.4.4) are another example.

30.2.A. EXERCISE. Suppose that $J'_1 \supset J'_2 \supset \cdots$ is a decreasing series of ideals that is **cofinal** with J_n . In other words, for every J_n , there is some J'_N with $J_n \supset J'_N$, and for every J'_n , there is some J_N with $J'_n \supset J_N$. Show that there is a canonical isomorphism $\lim A/J'_n \xrightarrow{\sim} \lim A/J_n$. Thus what matters is less the specific sequence of ideals than the “cofinal” equivalence class.

30.2.2. Preliminary remarks. We have an obvious morphism $A \rightarrow \hat{A}$.

If we put the discrete topology on A/J_n , then \hat{A} naturally has the structure of a **topological ring** (a ring that is a topological space, where all the ring operations are continuous, or equivalently, a ring object in the category of topological spaces). We then can have the notion of a **topological module** M over a topological ring A' — a module over the underlying ring A' , with a topology, such that the action of A' on M is continuous.

30.2.B. EXERCISE. Suppose \mathfrak{m} is a maximal ideal of a ring A . Show that the completion of A at \mathfrak{m} is canonically isomorphic to the completion of $A_{\mathfrak{m}}$ at \mathfrak{m} .

If (A, \mathfrak{m}) is a Noetherian local ring, then the natural map $A \rightarrow \hat{A}$ is an injection: anything in the kernel must lie in $\bigcap \mathfrak{m}^i$, which is 0 by the Krull Intersection Theorem (Exercise 13.9.A(b)). Thus “no information is lost by completing”, just as analytic functions are (locally) determined by their power series expansion.

(In the case of completion at an ideal I , we say that A is **I-adically separated** if $A \rightarrow \hat{A}$ is an injection, and **complete with respect to I** if $A \rightarrow \hat{A}$ is an isomorphism, although we won’t need these phrases. For example, the Krull Intersection Theorem implies that if I is a proper ideal of a Noetherian integral domain or a Noetherian local ring, then A is I -adically separated.)

30.2.C. EXERCISE. Suppose that (A, \mathfrak{m}) is Noetherian local ring containing its residue field k (i.e., it is a k -algebra). Let x_1, \dots, x_n be elements of A whose images are a basis for $\mathfrak{m}/\mathfrak{m}^2$. Show that the map of k -algebras

$$(30.2.2.1) \quad k[[t_1, \dots, t_n]] \rightarrow \hat{A}$$

defined by $t_i \mapsto x_i$ is a surjection. (First explain why there *is* such a map!) As usual, for local rings, the completion is assumed to be at the maximal ideal.

Exercise 30.2.C is a special case of the *Cohen Structure Theorem*. (See [E, §7.4] for more on this topic.)

30.2.D. EXERCISE. Let X be a locally Noetherian scheme over k , let $p \in X$ be a rational (k -valued) point. Suppose $p \in X$ is regular of codimension d (i.e., $\dim \mathcal{O}_{X,p} = d$). Describe an isomorphism $\hat{\mathcal{O}}_{X,p} \xrightarrow{\sim} k[[x_1, \dots, x_d]]$ as topological rings. (Hint: As in Exercise 30.2.C, choose d elements of \mathfrak{m} that restrict to a basis of $\mathfrak{m}/\mathfrak{m}^2$; these will be your x_1, \dots, x_d . Show that the map (30.2.2.1) has no kernel. It may help to identify $\mathfrak{m}^n/\mathfrak{m}^{n+1}$ with $\text{Sym}^n(\mathfrak{m}/\mathfrak{m}^2)$ using Theorem 23.3.8.)

The converse also holds: if (A, \mathfrak{m}, k) is a Noetherian local ring that is a k -algebra, and $\hat{A} \cong k[[x_1, \dots, x_d]]$, then A is a regular local ring of dimension d ; see [AtM, Prop. 11.24] for the key step.

30.2.3. Remark. Suppose p is a smooth rational (k -valued) point of a k -variety of dimension d , and $f \in \mathcal{O}_{X,p}$ (f is a “local function”). By way of the isomorphism of

Exercise 30.2.D, we can interpret f as an element of $k[[x_1, \dots, x_d]]$. This should be interpreted as the “power series expansion of f at p ”, in the “local coordinates x_1, \dots, x_d ”.

30.2.4. Aside: Some more geometric motivation.

Exercise 30.2.D may give some motivation for completion: “in the completion, regular schemes look like affine space”. This is often stated in the suggestive language of “formally locally, regular schemes are isomorphic to affine space”; this will be made somewhat more precise in §30.4.1.

We now give a little more geometric motivation for completion, that we will not use later on. Recall from §13.6.1 that étale morphisms are designed to look like “local isomorphisms” in differential geometry. But Exercise 13.6.E showed that this metaphor badly fails in one important way. More precisely, let $Y = \operatorname{Spec} k[t]$, and $X = \operatorname{Spec} k[u, 1/u]$, and let $p \in X$ be given by $u = 1$, and $q \in Y$ be given by $t = 1$. Suppose $\operatorname{char} k \neq 2$. The morphism $\pi: X \rightarrow Y$ induced by $t \mapsto u^2$ is étale, and $\pi(p) = q$. But there is no open neighborhood of p that π maps isomorphically onto an open neighborhood of q . However, the following Exercise shows that π induces an isomorphism of completions (an isomorphism of “formal neighborhoods”).

30.2.E. EXERCISE. Continuing the notation of the previous paragraph, show that π induces an isomorphism of completions $\hat{\mathcal{O}}_{Y,q} \rightarrow \hat{\mathcal{O}}_{X,p}$.

30.2.F. EXERCISE. Suppose $Y = \operatorname{Spec} A$, and $X = \operatorname{Spec} A[t]/(f(t))$, and $\pi: X \rightarrow Y$ is the morphism induced by $A \rightarrow A[t]/(f(t))$. Suppose $p \in X$ is a closed point, $q \in Y$ is a closed point, $\pi(p) = q$, and π is étale at p and induces an isomorphism of residue fields. Show that π induces an isomorphism of completions $\hat{\mathcal{O}}_{Y,q} \rightarrow \hat{\mathcal{O}}_{X,p}$. (If you are not familiar with Hensel’s Lemma, you will rediscover its central idea in the course of solving this exercise.)

With a little more care, you can show more generally that if $\pi: X \rightarrow Y$ is an étale morphism, $\pi(p) = q$, and π induces an isomorphism of residue fields at p , then π induces an isomorphism of completions $\hat{\mathcal{O}}_{Y,q} \rightarrow \hat{\mathcal{O}}_{X,p}$. (You may even wish to think about how to remove the hypothesis of isomorphism of residue fields.)

These ideas are close to the definition of “formal étaleness” discussed in §26.2.6.

You can interpret these results as a statement that “the implicit function theorem works formally locally”, even though it doesn’t work Zariski-locally. (The étale topology is somewhere between these two; it is partially designed so that the implicit function theorem in this sense always holds, essentially by fiat. But this is not the place to discuss the étale topology.)

30.2.5. Completion and exactness.

We conclude this section with an interesting and useful statement.

30.2.6. Theorem. — Suppose A is a Noetherian ring, and $I \subset A$ is an ideal. For any A -module M , let $\hat{M} = \varprojlim M/I^j M$ be the completion of M with respect to I .

(a) The completion \hat{A} (of A with respect to I) is flat over A .

(b) If M is finitely generated, then the natural map $\hat{A} \otimes_A M \rightarrow \hat{M}$ is an isomorphism.

(c) If $0 \rightarrow M \rightarrow N \rightarrow P \rightarrow 0$ is a short exact sequence of finitely generated A -modules, then $0 \rightarrow \hat{M} \rightarrow \hat{N} \rightarrow \hat{P} \rightarrow 0$ is exact. (Thus completion preserves exact sequences of finitely generated modules, by Exercise 1.6.E.)

We will use (a) in §30.3, and (b) in §30.8.

30.2.7. Remark. Before proving Theorem 30.2.6, we make some remarks. Parts (a) and (b) together clearly imply part (c), but we will use (c) to prove (a) and (b). Also, note a delicate distinction (which helps me remember the statement): if $0 \rightarrow M \rightarrow N \rightarrow P \rightarrow 0$ is an exact sequence of A -modules, *not necessarily finitely generated*, then

$$(30.2.7.1) \quad 0 \rightarrow \hat{A} \otimes_A M \rightarrow \hat{A} \otimes_A N \rightarrow \hat{A} \otimes_A P \rightarrow 0$$

is *always* exact, but

$$(30.2.7.2) \quad 0 \rightarrow \hat{M} \rightarrow \hat{N} \rightarrow \hat{P} \rightarrow 0$$

need *not* be exact — and when it *is* exact, it is often because the modules are finitely generated, and thus (30.2.7.2) is really (30.2.7.1).

Caution: completion is not always exact. Consider the exact sequence of $k[t]$ -modules

$$0 \longrightarrow \bigoplus_{n=1}^{\infty} k[t] \xrightarrow{\times(t, t^2, t^3, \dots)} \bigoplus_{n=1}^{\infty} k[t] \longrightarrow \bigoplus_{n=1}^{\infty} k[t]/(t^n) \longrightarrow 0.$$

After completion with respect to the ideal $((t, t, t, \dots))$, the sequence is no longer exact in the middle: (t^2, t^3, t^4, \dots) maps to 0, but is not in the image of the completion of the previous term.

★ *Proof.* The key step is to prove (c), which we do through a series of exercises. Suppose that $0 \rightarrow M \rightarrow N \rightarrow P \rightarrow 0$ is a short exact sequence of finitely generated A -modules.

30.2.G. EXERCISE. Show that $\hat{N} \rightarrow \hat{P}$ is surjective. Hint: consider an element of \hat{P} as a sequence $(p_j \in P/I^j P)_{j \geq 0}$, where the image of p_{j+1} is p_j , cf. Exercise 1.4.C. Build a preimage $(n_j \in N/I^j N)_{j \geq 0}$ by induction on j .

We now wish to identify $\ker(\hat{N} \rightarrow \hat{P})$ with \hat{M} .

30.2.H. EXERCISE. Show that for each $j \geq 0$,

$$(30.2.7.3) \quad 0 \rightarrow M/(M \cap I^j N) \rightarrow N/I^j N \rightarrow P/I^j P \rightarrow 0$$

is exact. (Possible hint: show that $0 \rightarrow M \cap I^j N \rightarrow M \rightarrow N/I^j N \rightarrow P/I^j P \rightarrow 0$ is exact.)

The short exact sequences (30.2.7.3) form an inverse system as j varies. Its limit is left-exact (because limits always are), but it is also right-exact by Exercise 1.6.M, as the “transition maps on the left” $M/(M \cap I^{j+1} N) \rightarrow M/(M \cap I^j N)$ are clearly surjective. Thus

$$(30.2.7.4) \quad 0 \rightarrow \lim M/(M \cap I^j N) \rightarrow \hat{N} \rightarrow \hat{P} \rightarrow 0$$

is exact. To complete the proof of (c), it suffices (by Exercise 30.2.A) to show that the sequence of submodules $I^j M$ is cofinal with the sequence $M \cap I^j N$, so that the term $\lim M/(M \cap I^j N)$ on the left of (30.2.7.4) is naturally identified with \hat{M} .

30.2.I. EXERCISE. Prove this. Hint: clearly $I^j M \subset M \cap I^j N$. By Corollary 13.9.4 to the Artin-Rees Lemma 13.9.3, for some integer s , $M \cap I^{j+s} N = I^j(M \cap I^s N)$ for all $j \geq 0$, and clearly $I^j(M \cap I^s N) \subset I^j M$.

This completes the proof of part (c) of Theorem 30.2.6.

For part (b), present M as

$$(30.2.7.5) \quad A^{\oplus m} \xrightarrow{\alpha} A^{\oplus n} \longrightarrow M \longrightarrow 0$$

where α is an $m \times n$ matrix with coefficients in A . Completion is exact in this case by part (c), and clearly commutes with finite direct sums, so

$$\hat{A}^{\oplus m} \longrightarrow \hat{A}^{\oplus n} \longrightarrow \hat{M} \longrightarrow 0$$

is exact. Tensor product is right-exact, and commutes with direct sums (Exercise 1.3.M), so

$$\hat{A}^{\oplus m} \longrightarrow \hat{A}^{\oplus n} \longrightarrow \hat{A} \otimes_A M \longrightarrow 0$$

is exact as well. Notice that the maps from $\hat{A}^{\oplus m}$ to $\hat{A}^{\oplus n}$ in both right-exact sequences are the same; they are both α . Thus their cokernels are identified, and (b) follows.

Finally, to prove (a), we need to extend the ideal-theoretic criterion for flatness (Theorem 25.4.1) slightly. Recall (§25.4.2) that it is equivalent to the fact that an A -module M is flat if and only if for all ideals I , the natural map $I \otimes_A M \rightarrow M$ is an injection.

30.2.J. EXERCISE (STRONGER FORM OF THE IDEAL-THEORETIC CRITERION FOR FLATNESS). Show that an A -module M is flat if and only if for all *finitely generated ideals* I , the natural map $I \otimes_A M \rightarrow M$ is an injection. (Hint: if there is a counterexample for an ideal J that is not finitely generated, use it to find another counterexample for an ideal I that *is* finitely generated.)

By this criterion, to prove (a) it suffices to prove that the multiplication map $I \otimes_A \hat{A} \rightarrow \hat{A}$ is an injection for all finitely generated ideals I . But by part (b), this is the same showing that $\hat{I} \rightarrow \hat{A}$ is an injection; and this follows from part (c). \square

30.3 Defining types of singularities

Singularities are best defined in terms of completions. As an important first example, we finally define “node”.

30.3.1. Definition. Suppose X is a dimension 1 variety over \bar{k} , and $p \in X$ is a closed point. We say that X has a **node** at p if the completion of $\mathcal{O}_{X,p}$ at $m_{X,p}$ is isomorphic (as topological rings) to $\bar{k}[[x, y]]/(xy)$.

30.3.A. EXERCISE. Suppose $k = \bar{k}$ and $\text{char } k \neq 2$. Show that the curve in \mathbb{A}_k^2 cut out by $(y^2 - x^2 - x^3)$, (which we have been studying repeatedly since Figure 8.4) has a node at the origin.

30.3.B. EXERCISE. Suppose $k = \bar{k}$ and $\text{char } k \neq 2$, and we have $f(x, y) \in k[x, y]$. Show that $\text{Spec } k[x, y]/(f(x, y))$ has a node at the origin if and only if f has no terms of degree 0 or 1, and the degree 2 terms are not a perfect square. (This generalizes Exercise 30.3.A.)

The definition of node outside the case of varieties over algebraically closed fields is more problematic, and we give some possible ways forward. For varieties over a non-algebraically closed field k , one can always base change to the closure \bar{k} . As an alternative approach, if p is a k -valued point of a variety over k (not necessarily algebraically closed), then we could take the same Definition 30.3.1; this might reasonably be called a *split node*, because the branches (or more precisely, the tangent directions) are distinguished. Those singularities that are not split nodes, but which become nodes after base change to \bar{k} (such as the origin in $\text{Spec } \mathbb{R}[x, y]/(x^2 + y^2)$) might reasonably be called *non-split nodes*.

30.3.2. Other singularities. We may define other singularities similarly. To avoid complications, we do so over an algebraically closed field.

30.3.3. Definition. Suppose X is a variety over \bar{k} , and p is a closed point of X , where $\text{char } k \neq 2, 3$. We say that X has a **cusp** (resp. **tacnode**, **triple point**) at p if $\hat{\mathcal{O}}_{X,p}$ is isomorphic to the completion of the curve $\text{Spec } \bar{k}[x, y]/(y^2 - x^3)$ (resp. $\text{Spec } \bar{k}[x, y]/(y^2 - x^4)$, $\text{Spec } \bar{k}[x, y]/(y^3 - x^3)$). We say that X has an **ordinary multiple point of multiplicity m** , or **ordinary m -fold point**, if $\hat{\mathcal{O}}_{X,p}$ is isomorphic to the completion of the curve $\text{Spec } \bar{k}[x, y]/(f(x, y))$ where f is a homogeneous polynomial of degree m with no repeated roots. (You can see quickly that an ordinary 2-fold point — or *ordinary double point* — is precisely a node, and an ordinary 3-fold point is a triple point.)

30.3.C. TRIVIAL EXERCISE. (For this exercise, work over an algebraically closed field for simplicity.) Define A_n curve singularity (see §23.4.4). Define A_n surface singularity (see Exercise 23.4.F). Define D_n , E_6 , E_7 , and E_8 surface and curve singularities (see Remark 23.4.5).

30.3.4. Using this definition. We now give an example of how this definition can be used.

30.3.D. EXERCISE. Suppose X is a \bar{k} -variety with a node at a closed point p . Show that the blow up of X at p yields a morphism $\beta: \tilde{X} \rightarrow X$, where the exceptional divisor $\beta^{-1}(p)$ consists of two reduced smooth points. Hint: use the fact that completion is flat (Theorem 30.2.6(a)), and that blowing up commutes with flat base change (Exercise 25.2.P(a)), to turn this into a calculation on the “formal model” of the node, $\text{Spec } k[[x, y]]/(xy)$.

30.3.E. EXERCISE. Continuing the terminology of the previous exercise, describe an exact sequence

$$(30.3.4.1) \quad 0 \rightarrow \mathcal{O}_X \rightarrow \beta_* \mathcal{O}_{\tilde{X}} \rightarrow \mathcal{O}_p \rightarrow 0.$$

Hint: cohomology commutes with flat base change, Theorem 25.2.9 (or more simply, the affine case version Exercise 19.8.B(b)).

30.3.F. EXERCISE. We continue the terminology of the previous two exercises. If X is a pure-dimensional reduced projective curve, show that $p_a(X) = p_a(\tilde{X}) + 1$.

Thus “resolving a node” of a curve reduces the arithmetic genus of the curve by 1. If you wish, you can readily show that “resolving a cusp” reduces the genus by 1, and “resolving a tacnode” reduces the genus by 2. In general, for each type of

curve singularity, the contribution it makes to the genus — the difference between the genus of the curve and that of its normalization, when the curve is projective — is called the δ -invariant. Thus δ for a node or cusp is 1, and δ for a tacnode is 2.

30.3.G. EXERCISE. Show that for *any* singularity type (other than a smooth point), $\delta > 0$.

30.3.H. EXERCISE. Show that a reduced irreducible degree d plane curve can have at most $\binom{d-1}{2}$ singularities.

30.3.I. EXERCISE. Here is a trick for speedily working out the genus of a nodal (projective reduced) curve when it comes up in examples. (The description will be informal.) Assuming it is possible, draw a sketch on a piece of paper (i.e., the real plane) where nodes “look like nodes”, and a curve of genus g is drawn with g holes, as in Figure 30.1. (In Figure 30.1, we see three irreducible components, of geometric genus 0, 2, and 3 respectively. The first two components are smooth, while the third meets itself twice. There are also three additional nodes, joining each pair of the irreducible components.) Then if the curve is connected, count the number of holes visible in the picture. Prove that this recipe works. (Feel free to figure out what to do if the curve is not connected.)




FIGURE 30.1. This curve has genus 8 (count the “holes”!)

30.3.5. This can be quite useful. For example, the sketches in Figure 30.2 can remind you that a degree d plane curve has arithmetic genus $\binom{d-1}{2}$, the curve on $\mathbb{P}^1 \times \mathbb{P}^1$ coming from a section of $\mathcal{O}(a, b)$ (§17.4.9) has genus $(a-1)(b-1)$, and a curve in class $2C + 2F$ on the Hirzebruch surface \mathbb{F}_2 (using the language of Exercise 21.2.M) has arithmetic genus 3. In each case, we are taking a suitable curve in the class, taking advantage of the fact that arithmetic genus is constant in flat families (Crucial Exercise 25.7.B).

30.3.6. Other definitions. If $\bar{k} = \mathbb{C}$, then definitions of this sort agree with the analytic definitions (see [Ar1, §1]). For example, a complex algebraic curve singularity is a node if and only if it is *analytically* isomorphic to an open neighborhood




FIGURE 30.2. These curves have arithmetic genus 6, 6, and 3 respectively


of $xy = 0$ in \mathbb{C}^2 . There is also a notion of isomorphism “étale-locally”, which we do not define here. Once again, this leads to the same definition of these types of singularities (see [Ar2, §2]).

30.4 The Theorem on Formal Functions

Suppose $\pi: X \rightarrow Y$ is proper morphism of locally Noetherian schemes, and \mathcal{F} is a coherent sheaf on X , so $R^i\pi_*\mathcal{F}$ is a coherent sheaf (Grothendieck’s Coherence Theorem 19.9.1). Fix a point $q \in Y$. We already have a sense that there is an imperfect relationship between the fiber of $R^i\pi_*\mathcal{F}$ at q and the cohomology of \mathcal{F} “on the fiber X_q ”, and much of Chapter 28 was devoted to making this precise.

The Theorem on Formal Functions deals with this issue in a different way. Rather than comparing the fiber of $R^i\pi_*\mathcal{F}$ with the cohomology of \mathcal{F} restricted to

the fiber, it gives an isomorphism between information on “infinitesimal thickenings” of the fiber. Informally:



30.4.1. We now make this more precise. If Z is a closed subscheme of W cut out by ideal sheaf \mathcal{I} , we say that the closed subscheme cut out by \mathcal{I}^{n+1} is the **n th order formal (or infinitesimal) neighborhood (or thickening)** of Z (in W). The phrase “formal neighborhood” without mention of an “order” refers to the information contained in all of these neighborhoods at once, often in the form of a limit of the sort we will soon describe. (This also leads us to the notion of *formal schemes*, an important notion we will nonetheless not need.)

We turn now to our situation of interest. Rather than the fiber of $R^i \pi_* \mathcal{F}$ at q , we consider its completion: we take the stalk at q , and complete it at $\mathfrak{m}_{Y,q}$. Rather than the cohomology of \mathcal{F} along the fiber, we consider the cohomology of \mathcal{F} when restricted to the n th formal neighborhood X_n of the fiber, and take the limit. The Theorem of Formal Functions says that these are canonically identified.

Before we state the Theorem on Formal Functions, we need to be precise about what these limits are, and how they are defined. To concentrate on the essential, we do this in the special case where q is a *closed* point of Y , and leave the (mild) extension to the general case to you (Exercise 30.4.A).

We deal first with $R^i \pi_* \mathcal{F}$. It is helpful to note that by Exercise 30.2.B, we can compute the restriction to the n th formal neighborhood either by pulling back to $\text{Spec } \mathcal{O}_{Y,q}$ (“the germ of Y at q ”, i.e., tensoring with $\mathcal{O}_{Y,q}$; or in any affine open neighborhood of $q \in Y$. For convenience, we pick an affine open neighborhood $\text{Spec } B \subset Y$ of q . (Again, by Exercise 30.2.B, it won’t matter which affine open subset we take.) Define $X_B := \pi^{-1}(\text{Spec } B)$ for convenience. Let $\mathfrak{m} \subset B$ be the maximal ideal corresponding to q . Then $R^i \pi_* (\mathcal{F})_q$ is canonically the completion of the coherent B -module $H^i(X_B, \mathcal{F}|_{X_B})$ at \mathfrak{m} :

$$R^i \pi_* (\mathcal{F})_q^\wedge = \lim H^i(X_B, \mathcal{F}|_{X_B}) / \mathfrak{m}^n H^i(X_B, \mathcal{F}|_{X_B}).$$

We turn next to the cohomology of \mathcal{F} on thickenings of the fiber. We have closed embeddings $X_n \hookrightarrow X_{n+1}$, and thus maps of cohomology groups $H^i(X_{n+1}, \mathcal{F}|_{X_{n+1}}) \rightarrow H^i(X_n, \mathcal{F}|_{X_n})$. We have base change maps

$$(30.4.1.1) \quad H^i(X_B, \mathcal{F}|_{X_B})/\mathfrak{m}^n H^i(X_B, \mathcal{F}|_{X_B}) \rightarrow H^i(X_n, \mathcal{F}|_{X_n})$$

(see (28.1.0.3)) such that the square

$$\begin{array}{ccc} H^i(X_B, \mathcal{F}|_{X_B})/\mathfrak{m}^{n+1} H^i(X_B, \mathcal{F}|_{X_B}) & \longrightarrow & H^i(X_{n+1}, \mathcal{F}|_{X_{n+1}}) \\ \downarrow & & \downarrow \\ H^i(X_B, \mathcal{F}|_{X_B})/\mathfrak{m}^n H^i(X_B, \mathcal{F}|_{X_B}) & \longrightarrow & H^i(X_n, \mathcal{F}|_{X_n}) \end{array}$$

commutes. (Do you see why? Basically, this is again because they are base change maps.) Thus we have an induced map of limits:

$$(30.4.1.2) \quad R^i \pi_* (\mathcal{F})_{\mathfrak{q}}^{\wedge} \longrightarrow \lim_{n \rightarrow \infty} H^i(X_n, \mathcal{F}|_{X_n}).$$

30.4.A. EXERCISE. Extend the previous discussion to the case where \mathfrak{q} is not closed.

The Theorem on Formal Functions states that this is an isomorphism.

30.4.2. Theorem on Formal Functions. — Suppose $\pi: X \rightarrow Y$ is a proper morphism of locally Noetherian schemes, \mathcal{F} is a coherent sheaf on X , and $\mathfrak{q} \in Y$. Then (30.4.1.2) is an isomorphism for all $i \geq 0$.

Warning: the Theorem on Formal Functions does *not* imply anything about the maps “at finite level”, i.e., about (30.4.1.1) for “finite n ”.

The proof of Theorem 30.4.2 is subtle, and is postponed to the double-starred §30.8. We first give four important applications to keep you motivated and excited: Zariski’s Connectedness Lemma 30.5.1, Stein Factorization 30.5.3, Zariski’s Main Theorem 30.6.1, and Castelnuovo’s Criterion 30.7.1 for contracting (-1) -curves on surfaces.

30.5 Zariski’s Connectedness Lemma and Stein Factorization

We now state and prove Zariski’s Connectedness Lemma, which was mentioned in §28.1.8.

30.5.1. Zariski’s Connectedness Lemma. — If a proper morphism $\pi: X \rightarrow Y$ of locally Noetherian schemes is \mathcal{O} -connected (see §28.1.7), then $\pi^{-1}(\mathfrak{q})$ is connected for every $\mathfrak{q} \in Y$.

The proof requires the following result.

30.5.A. EASY EXERCISE (THE COMPLETION OF A LOCAL RING IS A LOCAL RING). Suppose (A, \mathfrak{m}) is a Noetherian local ring. Show that the completion of A along \mathfrak{m} , $\hat{A} := \lim A/\mathfrak{m}^n$, is a local ring, with maximal ideal $\mathfrak{m}\hat{A}$. Hint: show that any element of \hat{A} not in $\mathfrak{m}\hat{A}$ is invertible.

Proof. Assume otherwise that there is some $q \in Y$ such that $\pi^{-1}(q)$ is not connected, say $\pi^{-1}(q) = X_1 \sqcup X_2$ (where X_1 and X_2 are nonempty open subsets of the fiber, see Figure 30.3). Then the n th order formal neighborhood of $\pi^{-1}(q)$, having the same topological space, is also disconnected. We use the useful trick of idempotents (Remark 3.6.3). Let e_1 be the function 1 on X_1 and 0 on X_2 , and let e_2 be the function 1 on X_2 and 0 on X_1 . These functions make sense for any order formal neighborhood, and they have natural images in the inverse limit. Thus we get nonzero elements $e_1, e_2 \in \hat{\mathcal{O}}_{Y,q}$ with $e_1 + e_2 = 1$, $e_1 e_2 = 0$. Now, $\hat{\mathcal{O}}_{Y,q}$ is a local ring (Exercise 30.5.A). But $e_1 e_2 = 0$ implies that neither e_1 nor e_2 is invertible (or else the other one would be 0), so both e_1 and e_2 are in the maximal ideal, and hence $e_1 + e_2$ can't be 1. We thus have a contradiction. \square




FIGURE 30.3. “Formal neighborhoods” of X_1 and X_2 in the proof of Zariski’s Connectedness Lemma 30.5.1

30.5.2. Stein factorization. We next show the existence of a Stein factorization. We could have given this construction long before, but now Zariski’s Connectedness Lemma 30.5.1 will give us some impressive consequences.

30.5.3. Stein Factorization Theorem. — *Any proper morphism $\pi: X \rightarrow Y$ of locally Noetherian schemes can be factored into $\beta \circ \alpha$, where $\alpha: X \rightarrow Y'$ is proper and \mathcal{O} -connected (hence has connected fibers, by Zariski’s Connectedness Lemma 30.5.1), and $\beta: Y' \rightarrow Y$ is a finite morphism.*

(30.5.3.1)

$$\begin{array}{ccc} X & \xrightarrow[\alpha]{\mathcal{O}\text{-conn.}} & Y' \\ & \searrow \pi & \swarrow \beta \\ & Y & \end{array} \quad \text{finite}$$

We note that by the Cancellation Theorem 11.2.1 for projective morphisms, if π happens to be projective, then so is α (as δ_β , being a closed embedding, is projective).

Although it is not in the statement of the theorem, the proof produces a *specific* factorization, which is called *the Stein factorization* of π . The picture to have in mind is that the Stein factorization (roughly) contracts the continuous parts of the fibers, in some canonical way; see Figure 30.4.

As usual, the Noetherian hypotheses can be removed (see [Stacks, tag 03H2]).




FIGURE 30.4. Stein Factorization

Proof of the Stein Factorization Theorem 30.5.3. By Grothendieck's Coherence Theorem 19.9.1, $\pi_* \mathcal{O}_X$ is a coherent sheaf (of algebras) on Y . Define $Y' := \text{Spec } \pi_* \mathcal{O}_X$, so (as $\pi_* \mathcal{O}_X$ is finite type) the structure morphism $\beta: Y' \rightarrow Y$ is finite. We have a factorization

$$\begin{array}{ccc} X & \xrightarrow{\alpha} & Y' \\ & \searrow \pi & \swarrow \beta \\ & Y & \end{array} \quad \begin{array}{c} \text{finite} \end{array}$$

By the Cancellation Theorem 11.2.1 for proper morphisms, α is proper. Finally, we show that α is \mathcal{O} -connected. By the quasicoherence of both $\mathcal{O}_{Y'}$ and $\alpha_* \mathcal{O}_X$, it suffices to check on an affine cover of Y' . We choose as our affine cover of Y' the pullback of an affine cover of Y . If U is an affine open subset of Y , then on $\beta^{-1}(U)$, on the level of rings, $\mathcal{O}_{Y'} \rightarrow \alpha_* \mathcal{O}_X$ is precisely the isomorphism $\Gamma(U, \pi_* \mathcal{O}_X) \xrightarrow{\sim} \Gamma(U, \pi_* \mathcal{O}_X)$. \square

30.5.B. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a proper morphism of locally Noetherian schemes, that is an isomorphism over a dense open subset of Y . Suppose further that Y is normal. Show that π is \mathcal{O} -connected. (This applies, for example, to blow-ups of smooth varieties under smooth centers.) Hence, by Zariski's Connectedness Lemma 30.5.1, for every $q \in Y$, $\pi^{-1}(q)$ is connected. This is called **Zariski's Main Theorem** (the version for birational morphisms). Hint: Exercise 10.7.P.

30.5.C. EXERCISE. Suppose $\pi: X \rightarrow Y$ and $\pi': X \rightarrow Y'$ are two proper morphisms of locally Noetherian schemes, each contracting the same *connected* set of X (and an isomorphism elsewhere), and suppose Y and Y' are normal. Show that π and π' are the same (or more precisely, that there is an isomorphism $i: Y \xrightarrow{\sim} Y'$ such that $\pi' = i \circ \pi$). Informal translation: If $X \rightarrow Y$ has connected fibers, then Y is determined by just knowing the locus contracted on X . Hint: Identify Y and Y' as topological spaces, and then identify π and π' as maps of topological spaces. Use Stein factorization to recover the structure sheaf as $\pi_* \mathcal{O}_X$.

30.5.D. EXERCISE. Show that the construction of the Stein factorization of a morphism $\pi: X \rightarrow Y$ of locally Noetherian schemes commutes with any flat base change $Z \rightarrow Y$ of locally Noetherian schemes.

30.5.4. Resolution of singularities of curves.

If C is a reduced projective curve over a field k , then we can resolve its singularities by normalization: $\nu: \tilde{C} \rightarrow C$. (See Remark 23.4.6 for some discussion of resolution of singularities, and Theorem 13.5.8 for the reason why normal curves are regular.) But as we have seen from examples, it requires luck and insight to figure out how to take an integral closure, and I often have neither. We now take the (inspired) guesswork out of desingularization by explaining how to desingularize by blowing up. The algorithm is simple: we find a singular point, then blow it up, then look for more singular points to repeat the process. Clearly if there are no singular points to be found, then we are done; the problem is to show that this is guaranteed to terminate. We do this by making use of an integer invariant: the arithmetic genus. We must show (i) that if $p_a(C) = p_a(\tilde{C})$, then ν is an isomorphism — C is already nonsingular; (ii) if $C' \rightarrow C$ is a birational morphism, then $p_a(C') \leq p_a(C)$ (so $p_a(\tilde{C})$ will be a lower bound for the arithmetic genus throughout the process); and (iii) that blowing up a singular point decreases the arithmetic genus.

To set up these results, we consider the following situation. Suppose $\pi: C' \rightarrow C$ is a finite morphism (where C is as described in the previous paragraph, and C' is a reduced curve), that is an isomorphism away from a finite closed subset of C . Then the pullback map $\mathcal{O}_C \rightarrow \pi_* \mathcal{O}_{C'}$ has no kernel, as C is reduced, and the map is an isomorphism on the generic points of the components of C . The cokernel \mathcal{G} is supported on a finite set of closed points, and thus $H^i(C, \mathcal{G}) = 0$ for $i > 0$, and if $H^0(C, \mathcal{G}) = 0$, then $\mathcal{G} = 0$. From the exact sequence

$$0 \longrightarrow \mathcal{O}_C \longrightarrow \pi_* \mathcal{O}_{C'} \longrightarrow \mathcal{G} \longrightarrow 0,$$

and the fact that π is affine, $\chi(C', \mathcal{O}_{C'}) = \chi(C, \mathcal{O}_C) + \chi(C, \mathcal{G})$, from which $p_a(C) \geq p_a(C')$, with equality if and only if π is \mathcal{O} -connected (i.e., $\mathcal{G} = 0$).

30.5.E. EXERCISE. Show that if π is \mathcal{O} -connected, then π is an isomorphism. Hint: use Zariski's Connectedness Lemma 30.5.1 to show that the fibers are connected, and hence that π is bijective. Identify C and C' as topological spaces. Then use the isomorphism $\mathcal{O}_C \xrightarrow{\sim} \pi_* \mathcal{O}_{C'}$ to identify C and C' as ringed spaces.

To complete our strategy, it remains to show that if $\beta: C' \rightarrow C$ is the blow-up of C at a point p , and β is an isomorphism, then C is already nonsingular at that point. But if β is an isomorphism, then p is the exceptional divisor, and hence effective Cartier, and hence cut out locally by a singular equation. But then $\mathfrak{m}_{C,p} \subset \mathcal{O}_{C,p}$ is principal, so $\mathcal{O}_{C,p}$ is a regular local ring of dimension 1.

(You should verify that we have completed our strategy. Can you extend this argument to the case where C is not projective?)

Resolution of singularities for surfaces will be briefly discussed in §30.7.4.

30.6 Zariski's Main Theorem

Zariski's Main Theorem is misnamed in all possible ways (although not horribly so). It is not a single *Theorem*, because there are many results that go by this name, and they often seem quite unrelated. What they have in common is that

they are intellectual descendants of a particular result of Zariski. This result is not the *Main* Theorem of Zariski's prolific career; the name comes because it was the Main Theorem of a particular paper, [Z]. And finally, it is not *Zariski's* in any stronger sense than this; the modern versions are due to Grothendieck.

We have already seen one version of Zariski's Main Theorem: the “birational” form of Exercise 30.5.B. Our goal in this section is to prove a strong statement that is close to the optimal version of the one most often used.


Before we start, we note that the fiber of a proper morphism over a point $q \in Y$ is a proper (hence finite type) scheme over the residue field $\kappa(q)$, and we say that a point is *isolated in its fiber* if it forms a component of the fiber of dimension 0. (Here component means both connected and irreducible component.)

30.6.1. Zariski's Main Theorem (Grothendieck's form). — Suppose $\pi: X \rightarrow Y$ is a proper morphism of locally Noetherian schemes.

(a) The set of points of X that are isolated in their fiber forms an open subset $X_0 \subset X$.

(b) The morphism $\pi|_{X_0}: X_0 \rightarrow Y$ factors into an open embedding $X_0 \hookrightarrow Y'$ followed by a finite morphism $Y' \rightarrow Y$. Moreover, the morphism $\pi: X \rightarrow Y$ factors through Y' .


(30.6.1.1)



As with Stein factorization, the proof of the theorem yields a *specific* factorization.

If π is a morphism of *varieties*, we already know part (a), by upper semicontinuity of fiber dimension Theorem 12.4.2(a). For this reason, the proof in the case of varieties is easier, and (even if you are interested in the more general case) you are advised to read the proof restricting to this simpler case, to concentrate on the main ideas. But even in the case of varieties, we will need some of the ideas arising in the proof of (a), so you should read it until advised to skip ahead to (b).

Proof. (a) Take the Stein factorization (30.5.3.1) of π :



By the finiteness of β , a point of X is an isolated point in its fiber of π if and only if it is an isolated point in its fiber of α . Thus we may replace π by α . (Basically, we are reducing to the case where π is \mathcal{O} -connected, but we prefer to carefully call the morphism under consideration α so we can return to considering π in the proof of (b).)

Because fibers of proper \mathcal{O} -connected morphisms of locally Noetherian schemes are connected (Zariski's Connectedness Lemma 30.5.1), a point is isolated in its fiber of α if and only if it is its fiber.

If you are considering only the variety case, at this point you should jump to the proof of part (b).

To show that X_0 is open, we will identify it with the locus where $\Omega_\alpha = 0$ (the unramified locus of α — not of π !), which is open by Exercise 22.6.H.

30.6.A. EXERCISE. If $\Omega_\alpha|_p = 0$, show that p is isolated in its fiber of α (and hence of π).

We now show the other direction. Suppose p is isolated in its fiber, and let $r \in Y'$ be $\alpha(p)$. We will show that α induces an isomorphism $\alpha^\sharp: \mathcal{O}_{Y',r} \xrightarrow{\sim} \mathcal{O}_{X,p}$, so $\Omega_{\mathcal{O}_{X,p}/\mathcal{O}_{Y',r}} = 0$. As the construction of Ω behaves well with respect to localization on the source and target (Exercise 22.2.L), this implies that $(\Omega_\alpha)_p = 0$, so Ω_α is 0 at p .

The stalk $\mathcal{O}_{Y',r}$ is obtained by taking the colimit (of sections of \mathcal{O}) over all open subsets of Y' containing r . For any such open subset $V \subset Y'$, the condition $\alpha_* \mathcal{O}_X = \mathcal{O}_{Y'}$ gives an isomorphism $\Gamma(V, \mathcal{O}_{Y'}) \xrightarrow{\sim} \Gamma(\alpha^{-1}(V), \mathcal{O}_X)$ (compatible with inclusions). We will show that the system of open subsets $\alpha^{-1}(V)$ (as V varies through open neighborhoods of r), which each necessarily contain p , is cofinal with the system of open neighborhoods of p in X . To do this, we must show that for any open neighborhood U of p in X , there is an open neighborhood V of r in Y' such that $\alpha^{-1}(V) \subset U$. To do this, note that $X \setminus U$ is closed in X , so $\alpha(X \setminus U)$ is closed in Y' (as α is proper), so its complement $V := Y' \setminus \alpha(X \setminus U)$ is open. Simply note that $\alpha^{-1}(V) \subset U$ (implicitly using that α , being proper and \mathcal{O} -connected, is surjective, see Exercise 28.1.G), and we are done.

(b) In our proof of (a), we have established that α gives a bijection between X_0 and an open subset of Y' . Notice that this is furthermore a homeomorphism: any open subset of Y' pulls back to an open subset of X_0 by the continuity of α ; and if U is an open subset of X_0 , then $\alpha(X \setminus U)$ is a closed subset of Y' , so its complement (which is $\alpha(U)$, using that proper \mathcal{O} -connected morphisms are surjective, Exercise 28.1.G) is open.

Using this isomorphism of topological spaces, the condition of \mathcal{O} -connectedness of α (restricted to $\alpha(X_0)$) shows that α gives an isomorphism of ringed spaces (i.e., of schemes) $X_0 \xrightarrow{\sim} \alpha(X_0)$. We have found our desired factorization. \square

30.6.B. EXERCISE. Prove Zariski's Main Theorem with slightly different hypotheses: with Y affine, say $\text{Spec } A$, and X a *quasiprojective* A -scheme. (Hint: by the Definition 4.5.9 of quasiprojective A -scheme, we can find an open embedding $X \hookrightarrow X'$ into a projective A -scheme. Apply Theorem 30.6.1 to $X' \rightarrow \text{Spec } A$.)

For the next two applications (Exercises 30.6.D and 30.6.E), we will need an annoying hypothesis, which is foreshadowed by the previous exercise. We say that a morphism $\pi: X \rightarrow Y$ of locally Noetherian schemes satisfies (\dagger) if for all affines U_i in some open cover of Y , the restriction $\pi|_{\pi^{-1}(U_i)}: \pi^{-1}(U_i) \rightarrow U_i$ factors as an open embedding into a scheme proper over U_i . A morphism satisfying (\dagger) is necessarily separated and finite type. It turns out that this is sufficient: Nagata's Compactification Theorem states that *every* separated finite type morphism

of Noetherian schemes $\pi: X \rightarrow Y$ can be factored into a dense open embedding into a scheme proper over Y . (See [Lü] for a proof. The Noetherian hypotheses can be replaced by the condition that Y is quasicompact and quasiseparated, see [Co].)

30.6.C. EXERCISE. Show that if $\pi: X \rightarrow Y$ is a morphism of varieties over k , and X is an open subset of a proper variety Z , then π satisfies (\dagger) . (Hint: X is an open subscheme of its closure in $Z \times_k Y$.)

30.6.D. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a quasifinite birational morphism of integral locally Noetherian schemes, and that Y is normal. Suppose further π satisfies (\dagger) . Show that π is an open embedding. (Hence if furthermore π is a bijection, then π must be an isomorphism.)

30.6.E. EXERCISE. Suppose $\pi: X \rightarrow Y$ is a quasifinite morphism of Noetherian schemes satisfying (\dagger) . Show that π factors as an open embedding into a finite morphism.

$$\begin{array}{ccc} X & \xrightarrow{\text{open}} & Y' \\ & \searrow \pi & \swarrow \text{finite} \\ & Y & \end{array}$$

This shows that it is reasonable to think of quasifinite morphisms as “open subsets of finite morphisms”, as promised in §8.3.12.

As a final application of Zariski’s Main Theorem, we can finally prove a characterization of finite morphisms that we have mentioned a number of times.

30.6.2. Theorem. — Suppose $\pi: X \rightarrow Y$ is a morphism of locally Noetherian schemes. The following are equivalent.

- (a) π is finite.
- (b) π is affine and proper.
- (c) π is proper and quasifinite.

As usual, Noetherian hypotheses can be removed: [Gr-EGA, IV₃.8.11.1] shows that proper, quasifinite, locally finitely presented morphisms are finite.

Proof. Finite morphisms are affine by definition, and proper by Proposition 11.5.3, so (a) implies (b).

To show that (b) implies (c), we need only show that affine proper morphisms have finite fibers (as the “finite type” part of the definition of quasifiniteness is taken care of by properness). This was shown in Exercise 19.9.A.

Finally, we assume (c), that π is proper and quasifinite, and show (a), that π is finite. By Zariski’s Main Theorem 30.6.1, we have a factorization (30.6.1.1).

$$\begin{array}{ccc} X_0 & \xrightarrow[\gamma]{\text{open}} & Y' \\ & \searrow \pi & \swarrow \text{finite} \\ & Y & \end{array}$$

Now $X = X_0$. (Do you see why? Hint: Exercise 8.4.D followed by Exercise 8.3.G.) By the Cancellation Theorem 11.2.1 for proper morphisms (using that $Y' \rightarrow Y$ is finite hence separated), γ is proper.

30.6.F. EXERCISE. Show that proper quasicompact open embeddings are closed embeddings. Hint: use Corollary 9.4.5 to show that the image is open and closed in the target.

Applying this to our situation, we see that $X \rightarrow Y$ is the composition of two finite morphisms $\gamma: X \hookrightarrow Y'$ and $Y' \rightarrow Y$, and is thus finite itself. \square

30.6.3. **Other versions of Zariski's Main Theorem.**

The Noetherian conditions in Theorem 30.6.1 can be relaxed (see [Gr-EGA, IV₃.8.12.6] or [GW, Thm. 12.73]): if $\pi: X \rightarrow Y$ is a separated morphism of finite type, and Y is quasicompact and quasiseparated, then the set of points X_0 isolated in their fiber is open in X , and for every quasicompact open subset U of X_0 , there is a factorization

$$\begin{array}{ccc} U & \xrightarrow{\gamma} & Y' \\ & \searrow \pi|_U & \swarrow \beta \\ & Y & \end{array} \quad \begin{array}{c} \text{finite} \end{array}$$

such that β is finite, $\gamma|_U$ is a quasicompact open embedding, and $\gamma^{-1}(\gamma(U)) = U$.

There are also topological and power series forms of Zariski's Main Theorem, see [Mu7, §III.9].

30.7 Castelnuovo's Criterion for contracting (-1) -curves

We now prove a result used in Chapter 27. We showed in Exercise 23.4.O that if we blow up a regular surface at a (reduced) point, the exceptional divisor is a (-1) -curve (see Definition 21.2.7: it is isomorphic to \mathbb{P}^1 , and has normal bundle $\mathcal{O}_C(-1)$). Castelnuovo's Criterion is the converse: if we have a quasiprojective surface containing a (-1) -curve, that surface is obtained by blowing up another surface at a reduced regular point. We say that we can *blow down* the (-1) -curve.

30.7.1. Theorem (Castelnuovo's Criterion). — *Let $C \subset X$ be a (-1) -curve on a smooth projective surface over k . Then there exists a birational morphism $\pi: X \rightarrow Y$ such that Y is a smooth projective surface, $\pi(C)$ is a k -valued point, π is the blow-up of Y at that point, and C is the exceptional divisor of the blow-up.*

By Exercise 30.5.C, there is only one way to “blow down C ”: this contraction is unique.

Proof. The proof is in three steps. *Step 1.* We construct $\pi: X \rightarrow Y$ that contracts C to a point q (and is otherwise an isomorphism). At this point we will know that Y is a projective variety. *Step 2.* Then we show that Y is smooth; this is the hard step, and requires the Theorem on Formal Functions 30.4.2. *Step 3.* Finally, we recognize π as the blow-up $\text{Bl}_q Y \rightarrow Y$.

Step 1. We do this by finding an invertible sheaf \mathcal{L} whose complete linear series will do the job. We start with a very ample invertible sheaf \mathcal{H} on X such that $H^1(X, \mathcal{H}) = 0$. We can find such an invertible sheaf \mathcal{H} by choosing any very ample invertible sheaf, and then taking a suitably large multiple, invoking

Exercise 17.6.C (very ample \otimes very ample = very ample) and Serre vanishing (Theorem 19.1.3(ii)). Let $k = \mathcal{H} \cdot C$; $k > 0$ by the very ampleness of \mathcal{H} (Exercise 21.1.K).

30.7.A. EXERCISE. Show that $H^1(X, \mathcal{H}(iC)) = 0$ for $0 \leq i \leq k$. Hint: use induction on i , using

$$(30.7.1.1) \quad 0 \rightarrow \mathcal{H}((i-1)C) \rightarrow \mathcal{H}(iC) \rightarrow \mathcal{H}(iC)|_C \rightarrow 0.$$

Hint: Note that $\mathcal{H}(iC)|_C$ is a line bundle on $C \cong \mathbb{P}^1$; which $\mathcal{O}(n)$ is it?

Define $\mathcal{L} := \mathcal{H}(kC)$. We will use the sections of \mathcal{H} to obtain sections of \mathcal{L} , via the map $\phi: H^0(X, \mathcal{H}) \rightarrow H^0(X, \mathcal{H}(kC))$.

30.7.B. EXERCISE. Show that \mathcal{L} is base-point-free. Hint: To show that \mathcal{L} has no base points away from C , consider the image of ϕ . To show that \mathcal{L} has no base points on C , use that from (30.7.1.1) for $i = k$, we have that $H^0(X, \mathcal{L}) \rightarrow H^0(C, \mathcal{L}|_C)$ is surjective, and $\mathcal{L}|_C \cong \mathcal{O}_C$.

Thus the complete linear series $|\mathcal{L}|$ yields a morphism π' from X to some projective space \mathbb{P}^N .

30.7.C. EXERCISE. Show that π' precisely contracts C to a point q' . More explicitly, show that π' sends C to a point q' , and $\pi'|_{X \setminus C}: X \setminus C \rightarrow \mathbb{P}^N \setminus \{q'\}$ is a closed embedding. Hint: To show that π' gives a closed embedding away from C : consider the image of ϕ . To show that C is contracted by π' , use the fact that $\deg \mathcal{L}|_C = 0$.

Let Y' be the image of π' , so we have a morphism $X \rightarrow Y'$ (which we also call π'). Let $\nu: Y \rightarrow Y'$ be the normalization of Y' , so π' lifts to $\pi: X \rightarrow Y$ by the universal property of normalization (§10.7). (We are normalizing in order to use Exercise 30.5.B in Step 2.)

30.7.D. EASY EXERCISE. Show that $\nu^{-1}(q') \subset Y$ is a single point, which we call q ; and that $\pi(C) = q$. (Hint: the image of π is closed, and $\pi(C)$ is connected.)

30.7.E. EXERCISE. Let X be a smooth projective surface, containing a curve $C \subset X$ isomorphic to \mathbb{P}_k^1 . If $C \cdot C < 0$, show that there is a morphism $\pi: X \rightarrow X'$ contracting C to a point, and leaving the rest of X unchanged. (This will not be used, and is included to give practice with the argument of Step 1.)

Step 2. We now show that Y is a smooth surface. As Y is normal, we have that $\pi_* \mathcal{O}_X \cong \mathcal{O}_Y$ ($\pi: X \rightarrow Y$ is \mathcal{O} -connected), by Exercise 30.5.B. We will show that Y is smooth at q by showing that $\hat{\mathcal{O}}_{Y,q} \cong k[[x, y]]$ as topological rings, as then $\dim_k(\mathfrak{m}/\mathfrak{m}^2) = 2$. By the Theorem on Formal Functions 30.4.2, $\hat{\mathcal{O}}_{Y,q} \cong \lim H^0(C'_n, \mathcal{O}_{C'_n})$, where C'_n is the closed subscheme of X that is the scheme-theoretic preimage of the n th order neighborhood of q .

We are not precisely sure what the n th order neighborhood is, but that is fine; this inverse system is cofinal with $\mathcal{O}(-nC) = \mathcal{I}_{C/X}^n$ (Do you see why?), so we can take this inverse limit instead (by Exercise 30.2.A). It suffices to show that for all $n \geq 0$, $H^0(C_n, \mathcal{O}_{C_n}) \cong k[[x, y]]/(x, y)^n$ where C_n is defined as $0 \rightarrow \mathcal{O}(-nC) \rightarrow \mathcal{O}_X \rightarrow \mathcal{O}_{C_n} \rightarrow 0$.

We do this by induction on n . As a k -vector space, this is easy, using

$$0 \rightarrow \mathcal{I}^n / \mathcal{I}^{n+1} \rightarrow \mathcal{O}_{C_{n+1}} \rightarrow \mathcal{O}_{C_n} \rightarrow 0$$

(the closed subscheme exact sequence (9.1.2.1) for C_n in C_{n+1}), using the canonical isomorphism $H^0(\mathbb{P}^1, \mathcal{O}_{\mathbb{P}^1}(n)) \xrightarrow{\sim} \text{Sym}^n(H^0(\mathbb{P}^1, \mathcal{O}_{\mathbb{P}^1}(1)))$. The tricky thing is that we want this isomorphism as *rings*.

30.7.F. EXERCISE. Prove this. (This may remind you of how we found the ring of functions on the total space of an invertible sheaf in Exercise 20.11.E. This is no coincidence.)

Hence we have smoothness, completing Step 2.

Step 3. We now must recognize π as the blow-up of Y at q . As C is an effective Cartier divisor on X , by the universal property of blowing up means that there is a unique morphism α making the following diagram commute

$$\begin{array}{ccc} X & \xrightarrow{\alpha} & \text{Bl}_q Y \\ & \searrow \pi & \swarrow \beta \\ & Y & \end{array}$$

By the Cancellation Theorem 11.2.1 for projective morphisms, α is projective.

30.7.G. EXERCISE. Show that α has finite fibers.

As projective morphisms with finite fibers are finite (Theorem 19.1.7), α is a finite morphism. But then α is a birational finite morphism to a normal scheme, and thus an isomorphism by Exercise 10.7.P. \square

So far we have only used the Theorem on Formal Functions 30.4.2 for H^0 . The next exercise will give you a chance to see it in action for other cohomology groups.

30.7.H. EXERCISE.

(a) Suppose $\beta: X = \text{Bl}_q Y \rightarrow Y$ is the blow-up of a smooth surface over k at a smooth k -valued point. Show that $R^1\beta_*\mathcal{O}_X = R^2\beta_*\mathcal{O}_X = 0$.

(b) Show that the natural maps $\beta^*: H^1(Y, \mathcal{O}_Y) \rightarrow H^1(X, \mathcal{O}_X)$ and $\beta^*: H^2(Y, \mathcal{O}_Y) \rightarrow H^2(X, \mathcal{O}_X)$ are isomorphisms.

30.7.2. Elementary transformations.

Recall the definition of the Hirzebruch surfaces \mathbb{F}_n (over a field k) from Example 18.2.4. If $n > 0$, then on \mathbb{F}_n , we have a unique curve E of maximally negative self-intersection (Exercise 21.2.R). It has self-intersection $E \cdot E = -n$. It is a section of the projection $\pi: \mathbb{F}_n \rightarrow \mathbb{P}^1$. (If $n = 0$, so $\mathbb{F}_n \cong \mathbb{P}^1 \times \mathbb{P}^1$, E is usually defined as any “constant” section of $\pi: \mathbb{P}^1 \times \mathbb{P}^1 \rightarrow \mathbb{P}^1$ — it is not unique.)

Suppose q is a k -valued point of \mathbb{F}_n . Let F be the fiber of π containing q .

30.7.I. EXERCISE. Show that $F \cdot F = 0$.

Let $\beta: \text{Bl}_q \mathbb{F}_n \rightarrow \mathbb{F}_n$ be the blow-up at q , and let Z be the exceptional divisor (as the name “ E ” is taken). Let F' be the proper transform of F , and let E' be the proper transform of E .

30.7.J. EXERCISE. Show that $F' \cdot F' = -1$.

Thus by Castelnuovo’s Criterion, we can blow down F' , to obtain a new surface Y .

30.7.K. EXERCISE. Suppose $n > 0$. If $q \notin E$, show that $Y \cong \mathbb{F}_{n-1}$. If $q \in E$, show that $Y \cong \mathbb{F}_{n+1}$. Possible hint: If you knew that Y was a Hirzebruch surface \mathbb{F}_m , you could recognize it by the self-intersection of the unique curve of negative self-intersection if $m > 0$.

30.7.L. EXERCISE. Suppose $n = 0$. Show that $Y \cong \mathbb{F}_1$. Hint: Exercise 23.4.K.

This discussion can be generalized to \mathbb{P}^1 -bundles over more general curves. The following is just a first step in the story.

30.7.M. EXERCISE. Suppose C is a smooth curve over k , $\pi: X \rightarrow C$ is a \mathbb{P}^1 -bundle over C . If q is a k -valued point of X , then we can blow up q and blow down the proper transform of the fiber through q . Show that the result is another \mathbb{P}^1 -bundle X' over C . (This is called an **elementary transformation** of the ruled surface.)

30.7.N. EXERCISE. Suppose $X = \text{Proj}_C(\mathcal{L} \oplus \mathcal{M})$, and q lies on the section corresponding to

$$0 \rightarrow \mathcal{L} \rightarrow \mathcal{L} \oplus \mathcal{M} \rightarrow \mathcal{M} \rightarrow 0.$$

Show that $X' \cong \text{Proj}_C(\mathcal{L}(\pi(q)) \oplus \mathcal{M})$.

Can you speculate (or show) what happens if you are just told that $X \cong \text{Proj}_C \mathcal{V}$, where \mathcal{V} is a locally free sheaf of rank 2, and q lies on a section corresponding to (21.2.8.1)?

30.7.3. Minimal models of surfaces.

A surface over an algebraically closed field is called **minimal** if it has no (-1) -curves. If a surface X is not minimal, then we can choose a (-1) -curve, and blow it down. If the resulting surface is not minimal, we can again blow down a (-1) -curve, and so on. By the finiteness of the Picard number ρ (which we have admittedly not proved, see §19.4.9), as blowing down reduces ρ by 1 (by Exercise 23.4.O), this process must terminate. By this means we construct a *minimal model* of X — a minimal surface birational to X . If minimal models were unique, this would provide a means of classifying surfaces up to birationality: if X and X' were smooth projective surfaces, to see if they were birational, we would find their minimal models, and see if they were isomorphic.

Sadly, minimal models are not unique. The example of $\mathbb{P}^1 \times \mathbb{P}^1$ and \mathbb{P}^2 show that there can be more than one minimal model in a birational equivalence class. (Why are these two surfaces minimal?) Furthermore, the isomorphism between the blow-up of \mathbb{P}^2 at two points with the blow-up of $\mathbb{P}^1 \times \mathbb{P}^1$ at one point (Exercise 23.4.K) shows that a smooth surface can have more than one minimal model.

Nonetheless, the failure of uniqueness is well-understood: it was shown by Zariski that the only surfaces with more than one minimal model are those birational to $C \times \mathbb{P}^1$ for some curve C . This is part of the Enriques-Kodaira classification of surfaces. For more on this, see [Be], [Ba], or [BPV, §6].

30.7.4. Here are two more facts worth mentioning but not proving here. First, singular surfaces can be desingularized by normalizing and then repeatedly blowing up points; this explicit desingularization is analogous to the version for curves in §30.5.4, and is generalized by Hironaka's Theorem on resolution of singularities (Remark 23.4.6). Second, birational maps $X \dashrightarrow Y$ between smooth surfaces can

be factored into a finite number of blow-ups followed by a finite number of blow-downs (cf. §23.4.12, on weak factorization). As a consequence, by Exercise 30.7.H, birational projective surfaces have the same $H^i(\mathcal{O})$ for $i = 0, 1, 2$. (Using facts stated in §22.5.10, you can show that over \mathbb{C} , this implies that the entire Hodge diamond except for $h^{1,1}$ is a birational invariant of surfaces.)

30.8 ★★ Proof of the Theorem on Formal Functions 30.4.2

In this section, we will prove the following.

30.8.1. Theorem. — Suppose $X \rightarrow \operatorname{Spec} B$ is a proper morphism, with B Noetherian, and \mathcal{F} is a coherent sheaf on X . Suppose I is an ideal of B , and \hat{B} is the I -adic completion $\varprojlim B/I^n$ of B . Let $X_n := X \times_B (B/I^{n+1})$ (the n th order formal neighborhood of the fiber X_0) and let \mathcal{F}_n be the pullback of \mathcal{F} to X_n . Then for all i , the natural map

$$(30.8.1.1) \quad H^i(X, \mathcal{F}) \otimes_B \hat{B} = \widehat{H^i(X, \mathcal{F})} \longrightarrow \varprojlim H^i(X_n, \mathcal{F}_n)$$

is an isomorphism.

This proof is due to Serre, by way of B. Conrad.

30.8.A. EXERCISE. Prove that Theorem 30.8.1 implies Theorem 30.4.2. (This is essentially immediate if q is a closed point, and might require a little thought if q is not closed; cf. Exercise 30.4.A.)

On the right side of (30.8.1.1), we have morphisms $X_1 \hookrightarrow X_2 \hookrightarrow \cdots \hookrightarrow X$, from which we get restriction maps $H^i(X_1, \mathcal{F}_1) \leftarrow H^i(X_2, \mathcal{F}_2) \leftarrow \cdots \leftarrow H^i(X, \mathcal{F})$. But notice that the X_i are all supported on the same underlying set, and it is helpful to keep in mind that the \mathcal{F}_i 's are all sheaves on the same topological space.

We will most often apply this in the case where I is a maximal ideal. But we may as well prove the result in this generality. The result can be relativized (with B and I replaced by a scheme and a closed subscheme), but we won't bother giving a precise statement.

Our proof will only use properness through the fact that the pushforward of coherent sheaves under proper morphisms are coherent (Grothendieck's Coherence Theorem 19.9.1), which required hard work. If you haven't read that proof, you can retreat to the projective case with little loss.

We will show that both sides of (30.8.1.1) are \hat{B} -modules, and that the map of (30.8.1.1) is an isomorphism of *topological* \hat{B} -modules.

Write $I^n \mathcal{F}$ for the coherent sheaf on X defined as you might expect: if \mathcal{I} is the pullback of \tilde{I} from $\operatorname{Spec} B$ to X , $I^n \mathcal{F}$ is the image of $\mathcal{I}^n \otimes \mathcal{F} \rightarrow \mathcal{F}$. Then $H^i(X_n, \mathcal{F}_n)$ can be interpreted as $H^i(X, \mathcal{F}/I^n \mathcal{F})$ (by §19.1 property (v)), so the terms in the limit on the right side of (30.8.1.1) all live on the single space X . From now on we work on X .

The map in the statement of the theorem arises from the restriction $H^i(X, \mathcal{F}) \rightarrow H^i(X, \mathcal{F}/I^n \mathcal{F})$. This appears in the long exact sequence associated to

$$0 \rightarrow I^n \mathcal{F} \rightarrow \mathcal{F} \rightarrow \mathcal{F}/I^n \mathcal{F} \rightarrow 0,$$

and in particular in the portion

$$H^i(X, I^n \mathcal{F}) \xrightarrow{u_n} H^i(X, \mathcal{F}) \longrightarrow H^i(X, \mathcal{F}/I^n \mathcal{F}) \longrightarrow H^{i+1}(X, I^n \mathcal{F}) \xrightarrow{v_n} H^{i+1}(X, \mathcal{F}).$$

We shrink this further to:

(30.8.1.2)

$$0 \longrightarrow H^i(X, \mathcal{F})/\mathrm{im}(u_n) \longrightarrow H^i(X, \mathcal{F}/I^n \mathcal{F}) \longrightarrow \ker(v_n) \longrightarrow 0$$

and then apply \lim to both sides. Limits are left-exact (§1.6.12), so the result is left-exact. But the situation is even better than that. The transition maps of the left term for our exact sequences,

$$H^i(X, \mathcal{F})/\mathrm{im}(u_{n+1}) \rightarrow H^i(X, \mathcal{F})/\mathrm{im}(u_n),$$

are clearly surjective, so by Exercise 1.6.M we have an *exact* sequence

(30.8.1.3)

$$0 \longrightarrow \lim H^i(X, \mathcal{F})/\mathrm{im}(u_n) \xrightarrow{\mu} \lim H^i(X, \mathcal{F}/I^n \mathcal{F}) \longrightarrow \lim \ker(v_n) \longrightarrow 0.$$

Two key facts will imply our result. (i) We will show that $\{\mathrm{im}(u_n)\}$ is cofinal with the I -adic topology on $H^i(X, \mathcal{F})$. Thus by Exercise 30.2.A, we have a natural isomorphism of the left side of (30.8.1.3) with

$$\lim H^i(X, \mathcal{F})/I^n H^i(X, \mathcal{F}) \xleftarrow{\sim} H^i(X, \mathcal{F})^\wedge.$$

Because $H^i(X, \mathcal{F})$ is a finitely generated B -module (by Grothendieck's Coherence Theorem 19.9.1, or more simply Exercise 19.1.A in the projective case), this is isomorphic to $H^i(X, \mathcal{F}) \otimes_B \hat{B}$ by Theorem 30.2.6(b), the left side of (30.8.1.1). (ii) We will show that there is some d such that the transition map $\ker(v_{n+d}) \rightarrow \ker(v_n)$ is the zero map for $n \gg 0$. This implies that the right term $\lim \ker(v_n)$ in (30.8.1.3) is 0. (Do you see why?) As a result, the map μ in (30.8.1.3) becomes our desired isomorphism. We now begin these two tasks.

(i) We will show that there is some n_0 such that for $n \geq n_0$,

$$I^n H^i(X, \mathcal{F}) \subset \mathrm{im}(u_n) \subset I^{n-n_0} H^i(X, \mathcal{F}).$$

The first inclusion is straightforward: note that $H^i(X, \mathcal{F}) \rightarrow H^i(X, \mathcal{F}/I^n \mathcal{F})$ is B -linear with kernel $\mathrm{im}(u_n)$, and target annihilated by I^n .

We will make use of the ring $R_\bullet := B \oplus I \oplus I^2 \oplus \cdots$ (the Rees algebra, defined in §13.9.1). It is a graded ring generated in degree 1. Because I is a finitely generated ideal of B , R_\bullet is a finitely generated B -algebra, and thus Noetherian.

For each $a \in I^m$, we have a commutative diagram of quasicoherent sheaves

$$\begin{array}{ccc} I^n \mathcal{F} & \xrightarrow{\times a} & I^{n+m} \mathcal{F} \\ \downarrow & & \downarrow \\ \mathcal{F} & \xrightarrow{\times a} & \mathcal{F} \end{array}$$

which induces a commutative diagram of B -modules

$$\begin{array}{ccc} H^i(X, I^n \mathcal{F}) & \xrightarrow{\times a} & H^i(X, I^{n+m} \mathcal{F}) \\ \downarrow & & \downarrow \\ H^i(X, \mathcal{F}) & \xrightarrow{\times a} & H^i(X, \mathcal{F}). \end{array}$$

In this way, $\oplus_{n \geq 0} H^i(X, I^n \mathcal{F})$ and $\oplus_{n \geq 0} H^i(X, \mathcal{F})$ are graded R_\bullet -modules (where a is considered as an element of R_m). Then $\oplus_{n \geq 0} \text{im}(u_n)$ is the image of the map of R_\bullet -modules

$$(30.8.1.4) \quad \oplus H^i(X, I^n \mathcal{F}) \rightarrow \oplus H^i(X, \mathcal{F}),$$

and is thus itself an R_\bullet -module.

We will show that $\oplus_{n \geq 0} H^i(X, I^n \mathcal{F})$ is a *finitely generated* R_\bullet -module. This will imply that its image $\oplus_{n \geq 0} \text{im}(u_n)$ is also finitely generated, and thus by Proposition 13.9.2, for some n_0 $\text{im}(u_{n+1}) = I \text{im}(u_n)$ for $n \geq n_0$, which in turn implies (by induction) that for $n \geq n_0$, we have $\text{im}(u_n) \subset I^{n-n_0} H^i(X, \mathcal{F})$, which will complete task (i).

Now cohomology of quasicoherent sheaves commutes with direct sums (§19.1 (vii)), so $\oplus_{n \geq 0} H^i(X, I^n \mathcal{F}) = H^i(X, \oplus_{n \geq 0} I^n \mathcal{F})$.

30.8.B. EXERCISE. Consider the fibered diagram

$$\begin{array}{ccc} X \times_B R_\bullet & \xrightarrow{\sigma} & X \\ \rho \downarrow & & \downarrow \\ \text{Spec } R_\bullet & \longrightarrow & \text{Spec } B. \end{array}$$

Interpret $\oplus_n H^i(X, I^n \mathcal{F})$ as $R^i \rho_* \sigma^* \mathcal{F}$. (Be sure to check that the R_\bullet -action determined by the pushforward agrees with the R_\bullet -action on $\oplus_n H^i(X, I^n \mathcal{F})$ described above!)

But $\sigma^* \mathcal{F}$ is coherent (pullback takes finite type quasicoherent sheaves to finite type quasicoherent sheaves, and all schemes here are locally Noetherian), so $\rho_* \sigma^* \mathcal{F}$ is coherent (proper pushforward of coherent sheaves on Noetherian schemes are coherent, by Grothendieck's Coherence Theorem 19.9.1). We have now completed task (i).

(ii) For the same reason that $\oplus_{n \geq 0} H^i(X, I^n \mathcal{F})$ is a coherent R_\bullet -module, $\oplus_{n \geq 0} H^{i+1}(X, I^n \mathcal{F})$ is a coherent R_\bullet -module. (Just replace i with $i+1$ throughout the discussion in (i).)

Notice that $\oplus_{n \geq 0} \ker v_n$ is the kernel of the map of graded R_\bullet -modules

$$\oplus_{n \geq 0} H^{i+1}(X, I^n \mathcal{F}) \xrightarrow{\oplus_{n \geq 0} v_n} \oplus_{n \geq 0} H^{i+1}(X, \mathcal{F})$$

(which is just (30.8.1.4) with i replaced by $i+1$) and thus itself has the structure of a graded R_\bullet -module. It is a submodule of the finitely generated module $\oplus_{n \geq 0} H^{i+1}(X, I^n \mathcal{F})$, so $\oplus_{n \geq 0} \ker v_n$ is finitely generated as well. We invoke Proposition 13.9.2 a second time, to show that there is some d such that $\ker(v_{n+1}) = I \ker(v_n)$ for $n \geq d$, from which

$$(30.8.1.5) \quad \ker(v_{A+B}) = I^A \ker v_B \quad \text{for } A \geq 0 \text{ and } B \geq d.$$

For any $a \in I^d$, the map

$$\times a: I^n \mathcal{F} \rightarrow I^{n+d} \mathcal{F}$$

induces a map $H^{i+1}(X, I^n \mathcal{F}) \rightarrow H^{i+1}(X, I^{n+d} \mathcal{F})$, which sends $\ker(v_n)$ to $\ker(v_{n+d})$. (This can be interpreted as coming from the action of R_\bullet on the module $\oplus \ker(v_n)$,

taking a as an element of R_d .) As a runs through I^d , the B -linear spans of the images of

$$\ker(v_n) \longrightarrow H^{i+1}(X, I^n \mathcal{F}) \xrightarrow{\times a} H^{i+1}(X, I^{n+d} \mathcal{F})$$

for $n \geq d$ is $I^d \ker(v_n) = \ker(v_{n+d})$ by (30.8.1.5). The composition

$$H^{i+1}(X, I^n \mathcal{F}) \xrightarrow{\times a} H^{i+1}(X, I^{n+d} \mathcal{F}) \longrightarrow H^{i+1}(X, I^n \mathcal{F})$$

(the latter map coming from inclusion of sheaves) is exactly the map

$$\times a: H^{i+1}(X, I^n \mathcal{F}) \rightarrow H^{i+1}(X, I^n \mathcal{F}).$$

Thus the image of $\ker(v_{n+d}) \rightarrow \ker(v_n)$ is exactly the B -linear span of the images of the composition maps

$$\ker(v_n) \hookrightarrow H^{i+1}(X, I^n \mathcal{F}) \xrightarrow{\times a} H^{i+1}(X, I^n \mathcal{F})$$

as a runs through I^d .

It thus suffices to show that the multiplication by $a \in I^d$ on $H^{i+1}(X, I^n \mathcal{F})$ annihilates $\ker(v_n)$ for $n \geq d$. Because $\ker(v_n) = I^{n-d} \ker(v_d)$ for all $n \geq d$ ((30.8.1.5) again), it suffices to show that any element a of I^d annihilates $\ker(v_d)$. But $\ker(v_d)$ is the image of $H^i(X, \mathcal{F}/I^d \mathcal{F})$ (from the long exact sequence, see (30.8.1.2)), and thus is annihilated by any element of I^d , as desired. \square

CHAPTER 31

★ Proof of Serre duality

Réfléchissant un peu à ton théorème de dualité, je m'aperçois que sa formulation générale est à peu près évidente, et d'ailleurs je viens de vérifier qu'elle se trouve implicitement (dans le cas de l'espace projectif) dans ton théorème donnant les $T^q(M)$ par des Ext. (J'ai bien l'impression, salaud, que tes §3 et 4 du Chap. 3 peuvent se faire aussi sans aucun calcul).

Thinking a bit about your duality theorem, I notice that its general form is almost obvious, and in fact I just checked that (for a projective space) it is implicitly contained in your theorem giving the $T^q(M)$ in terms of Exts. (I have the impression, you bastard, that §3 and 4 in your Chap. 3 could be done without any computation).

— A. Grothendieck, letter to J.-P. Serre, December 15, 1955 [GrS, p. 19]

31.1 Introduction

Throughout this chapter, we work over a fixed “ground field” k .

We first met Serre duality in §19.5 (or even Theorem 19.1.2), and we have repeatedly seen how useful it is. We will prove Theorem 19.5.1 (Corollary 31.3.10, combined with Exercise 31.4.I, see Remark 31.4.9), as well as stronger versions, and we will be left with a desire to prove even more. We give three statements (*Serre duality for vector bundles*; *Serre duality for Hom*; and *Serre duality for Ext*), in two versions (*functorial* and *trace*). (These names are idiosyncratic and nonstandard.) We give several variants for a number of reasons. First, the easier statements will be easier to prove, and the hardest statements we won't be able to prove here. Second, they may help give you experience in knowing how to know what to hope for, and what to try to prove.

Throughout this chapter, X will be a projective k -scheme of pure dimension n . We will want a coherent sheaf ω (or with more precision, ω_X , or even better, $\omega_{X/k}$) on X , the **dualizing sheaf**, which will play a role in the statements of duality. For the best statements, we will want a **trace morphism**

$$(31.1.0.1) \quad \boxed{t: H^n(X, \omega_X) \rightarrow k.}$$

31.1.1. Desideratum: the determinant of the cotangent bundle is dualizing for smooth varieties. If X is smooth, we will want $\omega_X = \mathcal{K}_X$ in this case (recall $\mathcal{K}_X = \det \Omega_X$) — the miracle that the canonical bundle is Serre-dualizing (§19.5.2). In particular, ω_X is an invertible sheaf. This will be disturbingly harder to prove

than the basic duality statements we show; we will only get to it later (Exercise 31.4.I). But we will prove more, for example that ω_X is an invertible sheaf if X is a regular embedding in a smooth variety (Exercise 31.4.H).

31.1.2. Desideratum: Serre duality for vector bundles. The first version of duality, which (along with Desideratum 31.1.1) gives Theorem 19.5.1, is the following: if \mathcal{F} is locally free of finite rank, then we have a functorial isomorphism

$$(31.1.2.1) \quad H^i(X, \mathcal{F}) \cong H^{n-i}(X, \mathcal{F}^\vee \otimes \omega_X)^\vee.$$

More precisely, we want to construct a *particular* isomorphism (31.1.2.1), or equivalently, a particular perfect pairing $H^i(X, \mathcal{F}) \times H^{n-i}(X, \mathcal{F}^\vee \otimes \omega_X) \rightarrow k$. This isomorphism will be *functorial in \mathcal{F}* , i.e., it gives a natural isomorphism of covariant functors

$$(31.1.2.2) \quad H^i(X, \cdot) \xrightarrow{\sim} H^{n-i}(X, \cdot^\vee \otimes \omega_X)^\vee.$$

We call this **functorial Serre duality for vector bundles**.

Better still, there should be a cup product in cohomology, which can be used to construct a map $H^i(X, \mathcal{F}) \times H^{n-i}(X, \mathcal{F}^\vee \otimes \omega_X) \rightarrow H^n(X, \omega_X)$. This should be functorial in \mathcal{F} (in the sense that we get a natural transformation of functors $H^i(X, \cdot) \rightarrow H^{n-i}(X, \cdot^\vee \otimes \omega_X)^\vee \otimes H^n(X, \omega_X)$). Combined with the trace map (31.1.0.1), we get a map $H^i(X, \mathcal{F}) \times H^{n-i}(X, \mathcal{F}^\vee \otimes \omega_X) \rightarrow k$, which should yield (31.1.2.2). We call this the **trace version of Serre duality for vector bundles**. (This was hinted at after the statement of Theorem 19.5.1.)

In fact, the trace version of Serre duality for vector bundles (and hence the functorial version) is true when X is Cohen-Macaulay, and in particular when X is smooth. We will prove the functorial version (see Corollary 31.3.10). We will give an indication of how the trace version can be proved in Remark 31.3.15.

31.1.3. Desideratum: duality for more general X and \mathcal{F} . A weaker sort of duality will hold with weaker hypotheses. We will show that (without Cohen-Macaulay hypotheses) for any coherent sheaf \mathcal{F} on a pure n -dimensional projective k -scheme X , there is a *functorial isomorphism*

$$(31.1.3.1) \quad \mathrm{Hom}(\mathcal{F}, \omega_X) \xrightarrow{\sim} H^n(X, \mathcal{F})^\vee.$$

We call this **functorial Serre duality for Hom**.

In parallel with Serre duality for vector bundles, we have a natural candidate for the perfect pairing:

$$(31.1.3.2) \quad \mathrm{Hom}(\mathcal{F}, \omega_X) \times H^n(X, \mathcal{F}) \longrightarrow H^n(X, \omega_X) \xrightarrow{t} k,$$

where the trace map t is some linear functional on $H^n(X, \omega_X)$. If this composition is a perfect pairing, we say that X (with the additional data of (ω_X, t)) satisfies the **trace version of Serre duality for Hom**. Unlike the trace version of Serre duality for vector bundles, we already know what the “cup product” map $\mathrm{Hom}(\mathcal{F}, \omega_X) \times H^n(X, \mathcal{F}) \rightarrow H^n(X, \omega_X)$ is: an element $[\sigma: \mathcal{F} \rightarrow \omega_X]$ of $\mathrm{Hom}(\mathcal{F}, \omega_X)$ induces — by covariance of $H^n(X, \cdot)$, see §19.1 — a map $H^n(X, \mathcal{F}) \rightarrow H^n(X, \omega_X)$. The resulting pairing is clearly functorial in \mathcal{F} , so the trace version of Serre duality for Hom implies the functorial version. Unlike the case of Serre duality for vector bundles, the functorial version of Serre duality for Hom also implies the trace version:

31.1.A. EXERCISE. Show that the functorial version of Serre duality for Hom implies the trace version. In other words, the trace map is already implicit in the functorial isomorphism $\text{Hom}(\mathcal{F}, \omega_X) \xrightarrow{\sim} H^n(X, \mathcal{F})^\vee$. Hint: consider the commuting diagram (coming from functoriality)

$$\begin{array}{ccc} \text{Hom}(\mathcal{F}, \omega_X) & \longrightarrow & H^n(X, \mathcal{F})^\vee \\ \uparrow & & \uparrow \\ \text{Hom}(\omega_X, \omega_X) & \longrightarrow & H^n(X, \omega_X)^\vee. \end{array}$$

31.1.4. Definition. Suppose X is a projective k -scheme of pure dimension n . A coherent sheaf $\omega = \omega_X = \omega_{X/k}$ along with a map $t: H^n(X, \omega_X) \rightarrow k$ is called **dualizing** if the natural map (cf. (31.1.3.1))

$$(31.1.4.1) \quad \text{Hom}(\mathcal{F}, \omega_X) \times H^n(X, \mathcal{F}) \longrightarrow H^n(X, \omega_X) \xrightarrow{t} k$$

is a perfect pairing for all coherent sheaves \mathcal{F} on X . We call ω_X the **dualizing sheaf** and t the **trace map**. (The earlier discussion of ω_X and t was aspirational. This now is a definition.) If X has such (ω_X, t) , we say that X satisfies **Serre duality** (for Hom). The following proposition justifies the use of the word “the” (as opposed to “a”) in the phrase “the dualizing sheaf”.

31.1.5. Proposition. — *If a dualizing sheaf and trace (ω_X, t) exists for X , this data is unique up to unique isomorphism.*

Proof. Suppose we have two such (ω_X, t) and (ω'_X, t') . From the two morphisms

$$(31.1.5.1) \quad \text{Hom}(\mathcal{F}, \omega_X) \times H^n(X, \mathcal{F}) \longrightarrow H^n(X, \omega_X) \xrightarrow{t} k$$

$$\text{Hom}(\mathcal{F}, \omega'_X) \times H^n(X, \mathcal{F}) \longrightarrow H^n(X, \omega'_X) \xrightarrow{t'} k,$$

we get a natural bijection $\text{Hom}(\mathcal{F}, \omega_X) \cong \text{Hom}(\mathcal{F}, \omega'_X)$, which is functorial in \mathcal{F} . By the typical universal property argument (Exercise 1.3.Z), this induces a (unique) isomorphism $\omega_X \xrightarrow{\sim} \omega'_X$. From (31.1.5.1), under this isomorphism, the two trace maps t and t' must be the same too. \square

We will prove the functorial and trace versions of Serre duality for Hom in Corollary 31.3.12. The special case of projective space will be a key ingredient; we prove this case now.

31.1.6. Serre duality (for Hom) for projective space. Define $\omega_{\mathbb{P}_k^n}$ (or just ω for convenience) as $\mathcal{O}_{\mathbb{P}_k^n}(-n-1)$. Let t be any isomorphism(a) $H^n(\mathbb{P}_k^n, \omega_{\mathbb{P}_k^n}) \xrightarrow{\sim} k$ (Theorem 19.1.2(b)). As the notation suggests, $(\omega_{\mathbb{P}_k^n}, t)$ will be dualizing for projective space \mathbb{P}_k^n .

31.1.B. EXERCISE. Suppose $\mathcal{F} = \mathcal{O}_{\mathbb{P}_k^n}(m)$. Show that the natural map (31.1.4.1) is a perfect pairing. (Hint: do this by hand. See the discussion after Theorem 19.1.2.) Hence show that if \mathcal{F} is a direct sum of line bundles on \mathbb{P}_k^n , the natural map (31.1.4.1) is a perfect pairing.

31.1.7. Proposition. — *The functorial version (and hence the trace version by Exercise 31.1.A) of Serre duality for Hom holds for \mathbb{P}_k^n .*

Proof. We wish to show that the (functorial) natural map

$$\text{Hom}(\mathcal{F}, \omega_{\mathbb{P}_k^n}) \xrightarrow{\sim} H^n(\mathbb{P}_k^n, \mathcal{F})^\vee,$$

coming from the cup pairing and the choice of isomorphism t , is an isomorphism (cf. (31.1.3.1)). Fix a coherent sheaf \mathcal{F} on \mathbb{P}_k^n . By Theorem 16.3.1, we can present \mathcal{F} as

$$(31.1.7.1) \quad 0 \longrightarrow \mathcal{G} \longrightarrow \mathcal{E} \longrightarrow \mathcal{F} \longrightarrow 0$$

where \mathcal{E} is a finite direct sum of line bundles, and \mathcal{G} is coherent. Applying the left-exact functor $\text{Hom}(\cdot, \omega_{\mathbb{P}_k^n})$ to (31.1.7.1), we have the exact sequence

$$(31.1.7.2) \quad 0 \longrightarrow \text{Hom}(\mathcal{F}, \omega_{\mathbb{P}_k^n}) \longrightarrow \text{Hom}(\mathcal{E}, \omega_{\mathbb{P}_k^n}) \longrightarrow \text{Hom}(\mathcal{G}, \omega_{\mathbb{P}_k^n}).$$

Taking the long exact sequence in cohomology for (31.1.7.1) and dualizing, we have the exact sequence

$$(31.1.7.3) \quad 0 \longrightarrow H^n(\mathbb{P}_k^n, \mathcal{F})^\vee \longrightarrow H^n(\mathbb{P}_k^n, \mathcal{E})^\vee \longrightarrow H^n(\mathbb{P}_k^n, \mathcal{G})^\vee$$

The (functorial) pairing (31.1.3.2) gives a map from (31.1.7.2) to (31.1.7.3):

$$(31.1.7.4) \quad \begin{array}{ccccccc} 0 & \longrightarrow & 0 & \longrightarrow & H^n(\mathbb{P}_k^n, \mathcal{F})^\vee & \longrightarrow & H^n(\mathbb{P}_k^n, \mathcal{E})^\vee & \longrightarrow & H^n(\mathbb{P}_k^n, \mathcal{G})^\vee \\ \alpha \uparrow & & \beta \uparrow & & \gamma \uparrow & & \delta \uparrow & & \epsilon \uparrow \\ 0 & \longrightarrow & 0 & \longrightarrow & \text{Hom}(\mathcal{F}, \omega_{\mathbb{P}_k^n}) & \longrightarrow & \text{Hom}(\mathcal{E}, \omega_{\mathbb{P}_k^n}) & \longrightarrow & \text{Hom}(\mathcal{G}, \omega_{\mathbb{P}_k^n}) \end{array}$$

Maps α and β are obviously isomorphisms, and Exercise 31.1.B shows that δ is an isomorphism. Thus by the subtle version of the Five Lemma (Exercise 1.7.D, as β and δ are injective and α is surjective), γ is injective. This shows that the natural map $\text{Hom}(\mathcal{F}', \omega_{\mathbb{P}_k^n}) \rightarrow H^n(\mathbb{P}_k^n, \mathcal{F}')^\vee$ is injective for *all* coherent sheaves \mathcal{F}' , and in particular for $\mathcal{F}' = \mathcal{G}$. Thus ϵ is injective. Then by the dual of the subtle version of the Five Lemma (as β and δ are surjective, and ϵ is injective), γ is surjective. \square

31.1.8. Mathematical puzzle. Here is a puzzle to force you to confront a potentially confusing point. We will see that Desideratum 31.1.1 holds for \mathbb{P}^1 , so $\omega_{\mathbb{P}^1} \cong \Omega_{\mathbb{P}^1}$. What then is the trace map $t: H^1(\mathbb{P}^1, \Omega_{\mathbb{P}^1}) \rightarrow k$? The Čech complex for $H^1(\mathbb{P}^1, \Omega_{\mathbb{P}^1})$ (with the usual cover of \mathbb{P}^1) is given by

$$(31.1.8.1) \quad 0 \longrightarrow \Omega_{\mathbb{P}^1}(U_0) \times \Omega_{\mathbb{P}^1}(U_1) \xrightarrow{\alpha} \Omega_{\mathbb{P}^1}(U_0 \cap U_1) \longrightarrow 0.$$

If $U_0 = \text{Spec } k[x]$, and $U_0 \cap U_1 = \text{Spec } k[x, 1/x]$, then the differentials on $U_0 \cap U_1$ are those of the form $f(x) dx$ where $f(x)$ is a Laurent polynomial (for example: $(x^{-3} + x^{-1} + 3 + 17x^4) dx$). To compute $H^1(\mathbb{P}^1, \Omega_{\mathbb{P}^1})$, we need to find which such differentials on $U_0 \cap U_1$ are in the image of α in (31.1.8.1). Clearly any term of the form $x^i dx$ (for $i \geq 0$) extends to a differential on U_0 , and thus is in the image of α . A short calculation shows that any term of the form $x^i dx$ ($i < -1$) extends to a differential on U_1 . Thus the cokernel of α can be described as the one-dimensional k -vector space generated by $x^{-1} dx$. We have an obvious isomorphism to k : take

the coefficient of $x^{-1} dx$, which can be interpreted as “take the residue at 0”. But there is another choice, which is equally good: take the residue at ∞ — certainly there is no reason to privilege 0 over ∞ (or U_0 over U_1)! But these two residues are *not* the same — they add to 0 (as you can quickly calculate — you may also believe it because of the Residue Theorem in the theory of Riemann surfaces). So: which one is the trace?

31.1.9. Desideratum: a stronger version of duality, involving Ext.

Ta formule:

$$H^{n-p}(X, \mathcal{F})' = \text{Ext}_{\mathcal{O}}^p(X, F, \Omega^n)$$

m'excite beaucoup, car je suis bien convaincu que c'est la bonne façon d'énoncer le théorème de dualité à la fois dans le cas analytique et dans le cas algébrique (le plus important — pour moi!).

I find your formula

$$H^{n-p}(X, \mathcal{F})' = \text{Ext}_{\mathcal{O}}^p(X, F, \Omega^n)$$

very exciting, as I am quite convinced that it is the right way to state the duality theorem in both the analytic case and the algebraic case (the more important one — for me!).

— J.-P. Serre, letter to A. Grothendieck, December 22, 1955 [GrS, p. 21]

The vector bundle and Hom versions of Serre duality have a common extension. If we have an isomorphism of functors

$$(31.1.9.1) \quad \text{Ext}^i(\cdot, \omega_X) \xrightarrow{\sim} H^{n-i}(X, \cdot)^\vee,$$

we say that X satisfies **functorial Serre duality for Ext**. (The case $i = 0$ is functorial Serre duality for Hom, or by Exercise 31.1.A, the trace version.) Here Ext^i is the derived functor of Hom; see the start of §31.2 below.

In Exercise 31.2.I, we will find that the functorial version of Serre duality for Ext (resp. the trace version) implies the functorial version of Serre duality for vector bundles (resp. the trace version). Thus to prove Theorem 19.5.1, it suffices to prove functorial Serre duality for Ext, and Desideratum 31.1.1 (see Remark 31.4.9). We will prove the functorial version of Serre duality for Ext when X is Cohen-Macaulay in Corollary 31.3.14.

31.1.10. Functorial Serre duality for Ext holds for projective space.

We now prove that functorial Serre duality for Ext holds for projective space. We will use the machinery of universal δ -functors (§24.2.6), so you may wish to either quickly skim that section, or else ignore this discussion.

31.1.C. EXERCISE. Show that $(\text{Ext}_{\mathbb{P}_k^n}^i(\cdot, \omega_{\mathbb{P}_k^n}))$ is a (contravariant) universal δ -functor. Hint: Ext is not a derived functor in its first argument, so you can't use the “projective” version of Corollary 24.2.10. Instead, use Theorem 24.2.8, and the existence of a surjection $\mathcal{O}(m)^{\oplus N} \rightarrow \mathcal{F}$ for each \mathcal{F} , for some $m < 0$.

31.1.D. EXERCISE. Show that $(H^{n-i}(\mathbb{P}_k^n, \cdot)^\vee)$ is a universal δ -functor. (What are the δ -maps?) Hint: try the same idea as in the previous exercise.

Proposition 31.1.7 gives an isomorphism of functors $\mathrm{Ext}_{\mathbb{P}_k^n}^0(\cdot, \omega_{\mathbb{P}_k^n}) \xleftarrow{\sim} H^n(\mathbb{P}_k^n, \cdot)^\vee$, so by the Definition 24.2.7 of universal δ -functor, we have an isomorphism of δ -functors $(\mathrm{Ext}_{\mathbb{P}_k^n}^i(\cdot, \omega_{\mathbb{P}_k^n})) \xleftarrow{\sim} (H^{n-i}(\mathbb{P}_k^n, \cdot)^\vee)$, thereby proving functorial Serre duality for Ext for \mathbb{P}_k^n .

31.1.11. Trace version of Serre duality for Ext . As with the previous versions, the functoriality of functorial Serre duality for Ext should come from somewhere. We should expect a natural cup product $\mathrm{Ext}^i(\mathcal{F}, \omega_X) \times H^{n-i}(X, \mathcal{F}) \rightarrow H^n(X, \omega_X)$, which coupled with the trace map (31.1.0.1) should yield the isomorphism (31.1.9.1). We call this the **trace version of Serre duality for Ext** . We will not be able to prove the trace version, as we will not define this cup product. But we will give some indication of how it works in §31.2.3.

31.1.12. Necessity of Cohen-Macaulay hypotheses. We remark that the Cohen-Macaulay hypotheses are necessary everywhere they are stated. The following example applies in all cases. Let X be the union of two 2-planes in \mathbb{P}_k^4 , meeting at a point p . If there were a coherent dualizing sheaf ω_X on X , then for $d \gg 0$, we would have $h^1(X, \mathcal{O}_X(-d)) = h^1(X, \omega_X(d))$ by Serre duality, which must be 0 by Serre vanishing (Theorem 19.1.3(ii)). We show that this is not the case.

31.1.E. EXERCISE. Let H be a hyperplane in \mathbb{P}^4 not passing through p ; say H is hyperplane $x_0 = 0$. Let Z be the intersection of dH (the divisor $x_0^d = 0$) with X . Show (cheaply) that for $d \geq 0$, $h^0(Z, \mathcal{O}_Z) > 1$. Show that $h^1(X, \mathcal{O}_X) = 0$ (perhaps using an easy Čech cover, or perhaps by comparing the cohomology of X to that of its normalization, as we did in with curves in (30.3.4.1)). Use the exact sequence

$$0 \rightarrow \mathcal{O}_X(-Z) \rightarrow \mathcal{O}_X \rightarrow \mathcal{O}_Z \rightarrow 0$$

to show that $h^1(X, \mathcal{O}_X(-d)) > 0$.

31.2 Ext groups and Ext sheaves for \mathcal{O} -modules

Recall that for any ringed space X , the category $\mathrm{Mod}_{\mathcal{O}_X}$ has enough injectives (Theorem 24.4.1). Thus for any \mathcal{O}_X -module \mathcal{F} on X , we may define

$$\mathrm{Ext}_X^i(\mathcal{F}, \cdot): \mathrm{Mod}_{\mathcal{O}_X} \rightarrow \mathrm{Mod}_{\Gamma(X, \mathcal{O}_X)}$$

as the i th right derived functor of $\mathrm{Hom}_X(\mathcal{F}, \cdot)$, and we have a corresponding long exact sequence for $\mathrm{Ext}_X^i(\mathcal{F}, \cdot)$. We similarly define a sheaf version of this

$$\mathcal{E}xt_X^i(\mathcal{F}, \cdot): \mathrm{Mod}_{\mathcal{O}_X} \rightarrow \mathrm{Mod}_{\mathcal{O}_X}$$

as the i th right derived functor of $\mathcal{H}om_X(\mathcal{F}, \cdot)$. In both cases, the subscript X is often omitted when it is clear from the context, although this can be dangerous when more than one space is relevant to the discussion. (We saw Ext functors for A -modules in §24.2.4.)

Warning: it is not clear (and in fact not true, see §24.4.8) that $\mathrm{Mod}_{\mathcal{O}_X}$ has enough projectives, so we cannot define Ext^i as a derived functor in its left argument. Nonetheless, we will see that it behaves as though it *is* a derived functor — it is “computable by acyclicity”, and has a long exact sequence (Remark 31.2.1).

Another warning: with this definition, it is not clear that if \mathcal{F} and \mathcal{G} are quasi-coherent sheaves on a scheme, then the $\mathcal{E}xt^i(\mathcal{F}, \mathcal{G})$ are quasicoherent, and indeed the aside in Exercise 14.3.A(a) points out this need not be true even for $i = 0$. But Exercise 31.2.F will reassure you.

Exercise 24.5.A (an injective \mathcal{O}_X -module, when restricted to an open subset $U \subset X$, is injective on U) has a number of useful consequences.

31.2.A. EXERCISE. Suppose \mathcal{I} is an injective \mathcal{O}_X -module. Show that $\mathcal{H}om_{\mathcal{O}_X}(\cdot, \mathcal{I})$ is an exact contravariant functor. (A related fact: $\mathcal{H}om_{\mathcal{O}_X}(\cdot, \mathcal{I})$ is exact, by the definition of injectivity, Exercise 24.2.C(a).)

31.2.B. EXERCISE. Suppose X is a ringed space, \mathcal{F} and \mathcal{G} are \mathcal{O}_X -modules, and U is an open subset. Describe a canonical isomorphism $\mathcal{E}xt_X^i(\mathcal{F}, \mathcal{G})|_U \xleftarrow{\sim} \mathcal{E}xt_U^i(\mathcal{F}|_U, \mathcal{G}|_U)$. Hint: take an injective resolution of \mathcal{G} on X , and restrict it to U . Use Exercise 31.2.A to show that the result is an injective resolution of $\mathcal{G}|_U$.

31.2.C. EXERCISE. Suppose X is a ringed space, and \mathcal{G} is an \mathcal{O}_X -module.
(a) Show that

$$\mathcal{E}xt^i(\mathcal{O}_X, \mathcal{G}) = \begin{cases} \mathcal{G} & \text{if } i = 0, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

(b) Describe a canonical isomorphism $\mathcal{E}xt^i(\mathcal{O}_X, \mathcal{G}) \xleftarrow{\sim} H^i(X, \mathcal{G})$.

31.2.D. EXERCISE. Use Exercise 31.2.C(a) to show that if \mathcal{E} is a locally free sheaf on X , then $\mathcal{E}xt^i(\mathcal{E}, \mathcal{G}) = 0$ for $i > 0$.

In the category of modules over a ring, we like projectives more than injectives, because free modules are easy to work with. It would be wonderful if locally free sheaves on schemes were always projective, but sadly this is not true. For if \mathcal{O}_X were projective, then $\mathcal{E}xt^i(\mathcal{O}_X, \cdot)$ would be zero (as a functor) for $i > 0$, but by Exercise 31.2.C(b), it is the functor $H^i(X, \cdot)$, which is often nonzero. Nonetheless, we can still compute $\mathcal{E}xt$ using a locally free resolution, as shown by the following exercise.

31.2.E. IMPORTANT EXERCISE. Suppose X is a ringed space, and

$$(31.2.0.1) \quad \cdots \longrightarrow \mathcal{E}_1 \longrightarrow \mathcal{E}_0 \longrightarrow \mathcal{F} \longrightarrow 0$$

is a resolution of an \mathcal{O}_X -module \mathcal{F} by locally free sheaves. (Of course we are most interested in the case where X is a scheme, and \mathcal{F} is quasicoherent, or even coherent.) Let \mathcal{E}_\bullet denote the truncation of (31.2.0.1), where \mathcal{F} is removed. Describe an isomorphism $\mathcal{E}xt^i(\mathcal{F}, \mathcal{G}) \xleftarrow{\sim} H^i(\mathcal{H}om(\mathcal{E}_\bullet, \mathcal{G}))$. In other words, $\mathcal{E}xt^\bullet(\mathcal{F}, \mathcal{G})$ can be computed by taking a locally free resolution of \mathcal{F} , truncating, applying $\mathcal{H}om(\cdot, \mathcal{G})$, and taking homology. Hint: choose an injective resolution

$$0 \longrightarrow \mathcal{G} \longrightarrow \mathcal{I}_0 \longrightarrow \mathcal{I}_1 \longrightarrow \cdots$$

and consider the spectral sequence whose E_0 term is

$$\begin{array}{ccccc}
 \vdots & & \vdots & & \\
 \uparrow & & \uparrow & & \\
 \mathrm{Hom}(\mathcal{E}_0, \mathcal{I}_1) & \longrightarrow & \mathrm{Hom}(\mathcal{E}_1, \mathcal{I}_1) & \longrightarrow & \cdots \\
 \uparrow & & \uparrow & & \\
 \mathrm{Hom}(\mathcal{E}_0, \mathcal{I}_0) & \longrightarrow & \mathrm{Hom}(\mathcal{E}_1, \mathcal{I}_0) & \longrightarrow & \cdots
 \end{array}$$

This result is important: to compute $\mathcal{E}xt$, we can compute it using finite rank locally free resolutions. You can work affine by affine (by Exercise 31.2.B), and on each affine you can use a free resolution of the left argument. As another consequence of Exercise 31.2.E:

31.2.F. EXERCISE. Suppose \mathcal{F} and \mathcal{G} are coherent sheaves on a quasiprojective k -scheme X . Show that $\mathcal{E}xt_X^i(\mathcal{F}, \mathcal{G})$ is a coherent sheaf as well. (Your argument will work on any scheme for which there always exist resolutions by finite rank locally free sheaves.)

31.2.1. Remark. The statement “ $\mathcal{E}xt^i(\mathcal{F}, \mathcal{G})$ behaves like a derived functor in the first argument” is true in a number of ways. We can compute it using a resolution of \mathcal{F} by locally frees, which are acyclic for $\mathcal{E}xt^i(\cdot, \mathcal{G})$. And we even have a corresponding long exact sequence, as shown in the next exercise.

31.2.G. EXERCISE. Suppose $0 \rightarrow \mathcal{F}' \rightarrow \mathcal{F} \rightarrow \mathcal{F}'' \rightarrow 0$ is an exact sequence of \mathcal{O}_X -modules on a ringed space X . For any \mathcal{O}_X -module \mathcal{G} , describe a long exact sequence

$$\begin{aligned}
 0 \longrightarrow \mathrm{Hom}(\mathcal{F}'', \mathcal{G}) \longrightarrow \mathrm{Hom}(\mathcal{F}, \mathcal{G}) \longrightarrow \mathrm{Hom}(\mathcal{F}', \mathcal{G}) \\
 \longrightarrow \mathrm{Ext}^1(\mathcal{F}'', \mathcal{G}) \longrightarrow \mathrm{Ext}^1(\mathcal{F}, \mathcal{G}) \longrightarrow \mathrm{Ext}^1(\mathcal{F}', \mathcal{G}) \longrightarrow \cdots
 \end{aligned}$$

Hint: take an injective resolution $0 \rightarrow \mathcal{G} \rightarrow \mathcal{I}^0 \rightarrow \cdots$. Use the fact that if \mathcal{I} is injective, then $\mathrm{Hom}(\cdot, \mathcal{I})$ is exact (the definition of injectivity, Exercise 24.2.C(a)). Hence get a short exact sequence of complexes

$$0 \rightarrow \mathrm{Hom}(\mathcal{F}'', \mathcal{I}^\bullet) \rightarrow \mathrm{Hom}(\mathcal{F}, \mathcal{I}^\bullet) \rightarrow \mathrm{Hom}(\mathcal{F}', \mathcal{I}^\bullet) \rightarrow 0$$

and take the long exact sequence in cohomology.

Here are two useful exercises.

31.2.H. EXERCISE. Suppose X is a ringed space, \mathcal{F} and \mathcal{G} are \mathcal{O}_X -modules, and \mathcal{E} is a locally free sheaf on X . Describe isomorphisms

$$\begin{aligned}
 \mathcal{E}xt^i(\mathcal{F} \otimes \mathcal{E}^\vee, \mathcal{G}) &\xrightarrow{\sim} \mathcal{E}xt^i(\mathcal{F}, \mathcal{G} \otimes \mathcal{E}) \xrightarrow{\sim} \mathcal{E}xt^i(\mathcal{F}, \mathcal{G}) \otimes \mathcal{E} \\
 \text{and} \quad \mathrm{Ext}^i(\mathcal{F} \otimes \mathcal{E}^\vee, \mathcal{G}) &\cong \mathrm{Ext}^i(\mathcal{F}, \mathcal{G} \otimes \mathcal{E}).
 \end{aligned}$$

Hint: show that if \mathcal{I} is injective then $\mathcal{I} \otimes \mathcal{E}$ is injective.

31.2.I. EXERCISE. If \mathcal{F} is a locally free sheaf on a scheme X , and ω_X is any coherent sheaf on X , describe an isomorphism $\mathrm{Ext}^i(\mathcal{F}, \omega_X) \xrightarrow{\sim} H^i(X, \mathcal{F}^\vee \otimes \omega_X)$ (functorial in \mathcal{F}). Show that functorial Serre duality for Ext implies functorial Serre duality for vector bundles. (You may wish to ponder the trace versions as well.) As a consequence, by §31.1.10, functorial Serre duality for vector bundles holds for projective space. Hint: Exercises 31.2.H and 31.2.C(b).

31.2.2. The local-to-global spectral sequence for Ext .

“Sheaf” Ext and “global” Ext are related by a spectral sequence. This is a straightforward application of the Grothendieck composition-of-functors spectral sequence, once we show that $\mathrm{Hom}(\mathcal{F}, \mathcal{I})$ is acyclic for the functor Γ .

31.2.J. EXERCISE. Suppose \mathcal{I} is an injective \mathcal{O}_X -module. Show that $\mathrm{Hom}(\mathcal{F}, \mathcal{I})$ is *flasque* (and thus Γ -acyclic by Exercise 24.4.I). Hint: suppose $j: U \hookrightarrow V$ is an inclusion of open subsets. We wish to show that $\mathrm{Hom}(\mathcal{F}, \mathcal{I})(V) \rightarrow \mathrm{Hom}(\mathcal{F}, \mathcal{I})(U)$ is surjective. Note that $\mathcal{I}|_V$ is injective on V (Exercise 24.5.A). Apply the exact functor $\mathrm{Hom}_V(\cdot, \mathcal{I}|_V)$ to the inclusion $j_!(\mathcal{F}|_U) \hookrightarrow \mathcal{F}|_V$ of sheaves on V (§24.4.7).

31.2.K. EXERCISE (THE LOCAL-TO-GLOBAL SPECTRAL SEQUENCE FOR Ext). Suppose X is a ringed space, and \mathcal{F} and \mathcal{G} are \mathcal{O}_X -modules. Describe a spectral sequence with E_2 -term $H^j(X, \mathrm{Ext}^i(\mathcal{F}, \mathcal{G}))$ abutting to $\mathrm{Ext}^{i+j}(\mathcal{F}, \mathcal{G})$. (Hint: use the Grothendieck composition-of-functors spectral sequence, Theorem 24.3.5. Recall that $\mathrm{Hom}(\mathcal{F}, \cdot) = \Gamma(\mathrm{Hom}(\mathcal{F}, \cdot))$, Exercise 2.3.C.)

31.2.3. ★ Composing Ext ’s (and H^i ’s): the Yoneda cup product.

It is useful and reassuring to know that Ext ’s can be composed, in a reasonable sense. We won’t need this, and so just outline the ideas, so you can recognize them in the future should you need them. For more detail, see [Gr3, §2] or [Cart].

If \mathcal{C} is an abelian category, and A_\bullet and B_\bullet are complexes in \mathcal{C} , then define $\mathrm{Hom}_\bullet(A_\bullet, B_\bullet)$ as the integer-graded group of *graded homomorphisms*: the elements of $\mathrm{Hom}_n(A_\bullet, B_\bullet)$ are the maps from the complex A_\bullet to B_\bullet shifted “to the right by n ”, without the requirement that the squares commute. Define $\delta: \mathrm{Hom}_\bullet(A_\bullet, B_\bullet)$ by

$$\delta(u) = du + (-1)^{n+1}ud$$

for each $u \in \mathrm{Hom}_n(A_\bullet, B_\bullet)$ (where d sloppily denotes the differential in both A_\bullet and B_\bullet). Then $\delta^2 = 0$, turning $\mathrm{Hom}_\bullet(A_\bullet, B_\bullet)$ into a complex. Let $H^\bullet(A_\bullet, B_\bullet)$ be the cohomology of this complex. If C_\bullet is another complex in \mathcal{C} , then composition of maps of complexes yields a map $\mathrm{Hom}_\bullet(A_\bullet, B_\bullet) \times \mathrm{Hom}_\bullet(B_\bullet, C_\bullet) \rightarrow \mathrm{Hom}_\bullet(A_\bullet, C_\bullet)$ which induces a map on cohomology:

$$(31.2.3.1) \quad H^\bullet(A_\bullet, B_\bullet) \times H^\bullet(B_\bullet, C_\bullet) \rightarrow H^\bullet(A_\bullet, C_\bullet)$$

which can be readily checked to be (well-defined and) associative. In particular, $H^\bullet(A_\bullet, A_\bullet)$ has the structure of a graded associative *non-commutative* ring (with unit), and $H^\bullet(A_\bullet, B_\bullet)$ (resp. $H^\bullet(B_\bullet, A_\bullet)$) has a natural graded left-module (resp. right-module) structure over this ring. The cohomology groups $H^\bullet(A_\bullet, B_\bullet)$ are functorial in both A_\bullet and B_\bullet . A short exact sequence of complexes $0 \rightarrow A'_\bullet \rightarrow A_\bullet \rightarrow A''_\bullet \rightarrow 0$ induces long exact sequences

$$\cdots \longrightarrow H^i(A'_\bullet, B_\bullet) \longrightarrow H^i(A_\bullet, B_\bullet) \longrightarrow H^i(A''_\bullet, B_\bullet) \longrightarrow H^{i+1}(A'_\bullet, B_\bullet) \longrightarrow \cdots$$

and

$$\cdots \longrightarrow H^i(B_\bullet, A''_\bullet) \longrightarrow H^i(B_\bullet, A_\bullet) \longrightarrow H^i(B_\bullet, A'_\bullet) \longrightarrow \cdots$$

Suppose now that \mathcal{C} has enough injectives. Suppose $A, B \in \mathcal{C}$, and let I^\bullet_A be any injective resolution of A (more precisely: take an injective resolution of A , and remove the “leading” A), and similarly for I^\bullet_B . Then it is a reasonable exercise to describe canonical isomorphisms

$$H^\bullet(I^\bullet_A, I^\bullet_B) \xleftarrow{\sim} H^\bullet(A, I^\bullet_B) \xleftarrow{\sim} \text{Ext}^\bullet(A, B)$$

where in the middle term, the “ A ” is interpreted as a complex that is zero, except the 0th piece is A .

Then the map (31.2.3.1) induces a (graded) map

$$(31.2.3.2) \quad \text{Ext}^\bullet(A, B) \times \text{Ext}^\bullet(B, C) \rightarrow \text{Ext}^\bullet(A, C)$$

extending the natural map $\text{Hom}(A, B) \times \text{Hom}(B, C) \rightarrow \text{Hom}(A, C)$. (Of course, one must show that the map (31.2.3.2) is independent of choice of injective resolutions of B and C .)

In particular, in the category of \mathcal{O} -modules on a ringed space X , we have (using Exercise 31.2.C(b)) a natural map

$$H^i(X, \mathcal{F}) \times \text{Ext}^j(\mathcal{F}, \mathcal{G}) \rightarrow H^{i+j}(X, \mathcal{G}).$$

This is the source of the natural map in the trace version of Serre duality for Ext , discussed at the end of §31.1.

31.3 Serre duality for projective k -schemes

We now prove various versions of Serre duality for projective k -schemes, by leveraging what we know about Serre duality for projective space.

31.3.1. The weird functor $\pi_{\text{sh}}^!$.

The key construction is a *right adjoint* $\pi_{\text{sh}}^!$ to the pushforward π_* , when π is an affine morphism. (Caution: the notation $\pi_{\text{sh}}^!$ is nonstandard, and is introduced only for the purposes of the arguments we will give.) This is surprising, as we usually think of π_* as a *right-adjoint* (to the pullback π^*), not a left adjoint. This can be seen as very roughly analogous to the surprising occasional *left-adjoint* to the pullback: extension by zero (§24.4.7).

31.3.2. We begin with the ring-theoretic version. Suppose $B \rightarrow A$ is a ring morphism, M is an A -module, and N is a B -module. Note that $\text{Hom}_B(A, N)$ naturally has the structure of an A -module. Also, M naturally carries the structure of a B -module; when we wish to emphasize its structure as a B -module, we sometimes call it M_B (see Exercise 1.5.E).

Consider the map

$$(31.3.2.1) \quad \text{Hom}_A(M, \text{Hom}_B(A, N)) \rightarrow \text{Hom}_B(M_B, N)$$

defined as follows. Given $m \in M$, and an element ϕ of $\text{Hom}_A(M, \text{Hom}_B(A, N))$, send m to $\phi_m(1)$.

31.3.A. EXERCISE.

- (a) Show that (31.3.2.1) is a homomorphism of B -modules.
 (b) Show that (31.3.2.1) is a bijection. Thus $(M \mapsto M_B, N \mapsto \text{Hom}_B(A, N))$ is an adjoint pair $\text{Mod}_A \leftrightarrow \text{Mod}_B$.
 (c) Show that this bijection (31.3.2.1) behaves well with respect to localization at an element of B .

Exercise 31.3.A(c) implies that this naturally “sheafifies” to a construction for an affine morphism $\pi: X \rightarrow Y$.

31.3.B. EXERCISE. Suppose $\pi: X \rightarrow Y$ is an affine morphism.

- (a) Explain how the map

$$\pi_{\text{sh}}^!: \text{QCoh}_Y \longrightarrow \text{QCoh}_X$$

$$\mathcal{G} \longmapsto \pi^{-1} \text{Hom}_Y(\pi_* \mathcal{O}_X, \mathcal{G})$$

(where $\text{Hom}_Y(\pi_* \mathcal{O}_X, \mathcal{G})$ is interpreted as a quasicoherent sheaf on X via Exercise 18.1.E, and in this guise is denoted $\pi^{-1} \text{Hom}_Y(\pi_* \mathcal{O}_X, \mathcal{G})$) globalizes the construction of Exercise 31.3.A, yielding a covariant functor $\text{QCoh}_Y \rightarrow \text{QCoh}_X$.

- (b) Describe a natural isomorphism of quasicoherent sheaves on Y

$$\pi_* \pi_{\text{sh}}^! \mathcal{G} \xrightarrow{\sim} \text{Hom}_Y(\pi_* \mathcal{O}_X, \mathcal{G}).$$

- (c) Show that $(\pi_*, \pi_{\text{sh}}^!)$ is an adjoint pair between QCoh_X and QCoh_Y .

Caution: we have defined $\pi_{\text{sh}}^!$ only for categories of quasicoherent sheaves. If π is a finite morphism, and Y (and hence X) is locally Noetherian (the case that will be relevant for us), then $\pi_{\text{sh}}^!$ is a covariant functor from the category of *coherent* sheaves on Y to *coherent* sheaves on X . We may show this affine-locally, using the notation of §31.3.2. As A and N are both coherent B -modules, $\text{Hom}_B(A, N)$ is a coherent B -module (cf. Exercise 14.3.A(b)), hence a finitely generated B -module, and hence a finitely generated A -module, hence a coherent A -module.

Thus if π is a finite morphism of locally Noetherian schemes, $(\pi_*, \pi_{\text{sh}}^!)$ is an adjoint pair between Coh_X and Coh_Y .

31.3.3. If $\mathcal{F} \in \text{QCoh}_X$ and $\mathcal{G} \in \text{QCoh}_Y$, then there is a natural isomorphism

$$(31.3.3.1) \quad \pi_* \text{Hom}_X(\mathcal{F}, \pi_{\text{sh}}^! \mathcal{G}) \xrightarrow{\sim} \text{Hom}_Y(\pi_* \mathcal{F}, \mathcal{G}),$$

which affine-locally is the isomorphism (31.3.2.1) described in Exercise 31.3.A.

If \mathcal{G} is an \mathcal{O}_X -module (not necessarily quasicoherent), it is in general *not* clear how to make sense of this construction to define an \mathcal{O}_Y -module $\pi_{\text{sh}}^! \mathcal{G}$. (Try it and see!) However, in the special case where π is a closed embedding, we *can* make sense of $\pi_{\text{sh}}^! \mathcal{G}$, as discussed in the next exercise.

31.3.C. EXERCISE (USED IN §31.4). Suppose $\pi: X \rightarrow Y$ is a *closed embedding* of schemes and \mathcal{G} is an \mathcal{O}_Y -module.

- (a) Explain why $\pi_{\text{sh}}^! (\mathcal{G}) := \text{Hom}_Y(\pi_* \mathcal{O}_X, \mathcal{G})$ naturally has the structure of an \mathcal{O}_X -module. Hint: if \mathcal{I} is the ideal sheaf of X , explain how $\pi_{\text{sh}}^! (\mathcal{G})$ (over some open subset $U \subset Y$) is annihilated by “functions vanishing on X ” (elements of $\mathcal{I}(U)$).

Hence we have defined a map $\pi_{\text{sh}}^! : \text{Mod}_{\mathcal{O}_Y} \rightarrow \text{Mod}_{\mathcal{O}_X}$ extending the map of Exercise 31.3.B(a).

(b) Show that $(\pi_*, \pi_{\text{sh}}^!)$ is an adjoint pair between $\text{Mod}_{\mathcal{O}_X}$ and $\text{Mod}_{\mathcal{O}_Y}$.

(c) Show that $\pi_{\text{sh}}^!$ sends injective \mathcal{O}_Y -modules to injective \mathcal{O}_X -modules. (Hint: π_* is exact; use Exercise 24.5.B.)

31.3.4. Remark. If π is finite and flat, which is the case most of interest to us, $\pi_{\text{sh}}^!$ agrees (in the only possible sense of the word) with $\pi^!$, one of Grothendieck's "six operations", and indeed this motivates our notation ("a variant of $\pi^!$ for sheaves"). But $\pi^!$ naturally lives in the world of derived categories, so we will not discuss it here. (For more on $\pi^!$, see [KS1, Ch. III].)

We now apply the machinery of $\pi_{\text{sh}}^!$ to Serre duality.

31.3.5. Projective n -dimensional schemes are covers of \mathbb{P}^n .

31.3.6. Proposition. — Suppose X is a projective k -scheme of pure dimension n .

- (a) There exists a finite morphism $\pi: X \rightarrow \mathbb{P}^n$.
- (b) If furthermore X is Cohen-Macaulay, then π is flat.

Proof. Part (b) follows from part (a) by the Miracle Flatness Theorem 29.2.11, so we will prove part (a).

Choose a closed embedding $j: X \hookrightarrow \mathbb{P}^N$. For simplicity of exposition, first assume that k is an infinite field. By Exercise 12.3.B(d), there is a linear space L of codimension $n + 1$ (one less than complementary dimension) disjoint from X . Projection from L yields a morphism $\pi: X \rightarrow \mathbb{P}^n$. The morphism π is affine ($\mathbb{P}^N \setminus L \rightarrow \mathbb{P}^n$ is affine, and the closed embedding $X \hookrightarrow \mathbb{P}^N \setminus L$ is, like all closed embeddings, affine) and projective (by the Cancellation Theorem 11.2.1 for projective morphisms), so π is finite (projective affine morphisms of locally Noetherian schemes are finite, Corollary 19.1.6).

31.3.D. EXERCISE. Prove Proposition 31.3.6(a), without the assumption that k is infinite. Hint: using Exercise 12.3.B(c), show that there is some d such that there is an intersection of $n + 1$ degree d hypersurfaces missing X . Then apply the above argument to the d th Veronese embedding of \mathbb{P}^N (§9.3.6). □

31.3.7. Remark. More generally, the above argument shows that any projective k -scheme of dimension n (not necessarily equidimensional) can be expressed as a finite cover of \mathbb{P}^n . This might be seen as a projective version of Noether normalization.

31.3.8. Serre duality on X via $\pi_{\text{sh}}^!$.

Suppose $\pi: X \rightarrow Y$ is a finite morphism of projective k -schemes of pure dimension n , and we have a coherent sheaf ω_Y on Y . (We will soon apply this in the case where $Y = \mathbb{P}_k^n$, but we may as well avoid distraction and needless specificity.)

31.3.9. Proposition. — Suppose π is flat. If functorial Serre duality for Ext holds for Y , with dualizing sheaf ω_Y , then functorial Serre duality for vector bundles holds for X , with dualizing sheaf $\pi_{\text{sh}}^!(\omega_Y)$.

Note the mismatch of the hypotheses and the conclusion: we use functorial Serre duality for Ext in the hypothesis, and get functorial Serre duality only for vector bundles in the conclusion.

Proof. For each i , and each finite rank locally free sheaf \mathcal{F} on X , we have isomorphisms (functorial in \mathcal{F}):

$$\begin{aligned}
 H^{n-i}(X, \mathcal{F}^\vee \otimes \pi_{\text{sh}}^! \omega_Y) &\xrightarrow{\sim} H^{n-i}(Y, \pi_*(\mathcal{F}^\vee \otimes \pi_{\text{sh}}^! \omega_Y)) \quad (\text{affineness of } \pi, \S 19.1 \text{ (v)}) \\
 &\xrightarrow{\sim} H^{n-i}(Y, \pi_*(\text{Hom}_X(\mathcal{F}, \pi_{\text{sh}}^! \omega_Y))) \quad (\text{Exercise 14.3.B}) \\
 &\xrightarrow{\sim} H^{n-i}(Y, \text{Hom}_Y(\pi_* \mathcal{F}, \omega_Y)) \quad (\text{equ. (31.3.3.1)}) \\
 &\xrightarrow{\sim} H^{n-i}(Y, (\pi_* \mathcal{F})^\vee \otimes \omega_Y) \quad (\text{Exercises 25.4.E and 14.3.B}) \\
 &\xrightarrow{\sim} H^i(Y, \pi_* \mathcal{F})^\vee \quad (\text{functorial Serre duality for Ext on } Y) \\
 &\xrightarrow{\sim} H^i(X, \mathcal{F})^\vee \quad (\text{affineness of } \pi, \S 19.1 \text{ (v)})
 \end{aligned}$$

(The hypothesis of flatness is used to show that $\pi_* \mathcal{F}$ is locally free.) \square

31.3.10. Corollary. — *Functorial Serre duality for vector bundles holds for every Cohen-Macaulay equidimensional projective k -scheme.*

Proof. Combine Proposition 31.3.9 with Proposition 31.3.6(b) and §31.1.10. \square

31.3.11. Proposition. — *Suppose $\pi: X \rightarrow Y$ is a finite morphism of projective k -schemes. If the trace version of Serre duality for Hom holds for Y with dualizing sheaf (ω_Y, t_Y) , then functorial Serre duality for Hom holds for X with dualizing sheaf $\pi_{\text{sh}}^! \omega_Y$.*

Recall from Exercise 31.1.A that for Serre duality for Hom, the trace version is equivalent to the functorial version. We state the Proposition in this awkward way because we use t_Y , and conclude functoriality for Serre duality for Hom for X . But the trace map for X will come cheaply as the diagonal arrow in (31.3.11.1).

Note that we have no flatness hypotheses on π (unlike the corresponding propositions for Serre duality for vector bundles and, soon, for Ext).

Proof. The following diagram commutes, and is functorial in \mathcal{F} .
(31.3.11.1)

$$\begin{array}{ccc}
 \text{Hom}_X(\mathcal{F}, \pi_{\text{sh}}^! \omega_Y) \times H^n(X, \mathcal{F}) & \longrightarrow & H^n(X, \pi_{\text{sh}}^! \omega_Y) \\
 \downarrow & & \downarrow \\
 \text{Hom}_Y(\pi_* \mathcal{F}, \omega_Y) \times H^n(Y, \pi_* \mathcal{F}) & \longrightarrow & H^n(Y, \pi_* \pi_{\text{sh}}^! \omega_Y) \\
 & & \downarrow \\
 & & H^n(Y, \omega_Y) \xrightarrow{t_Y} k
 \end{array}$$

\nearrow^{t_X}

(The left vertical arrow is an isomorphism; it is an isomorphism on each factor.) The commutativity of the diagram relies on the adjointness of $(\pi_*, \pi_{\text{sh}}^!)$, and the affineness of π . \square

31.3.12. Corollary. — *The functorial and trace versions of Serre duality for Hom hold for all equidimensional projective k -schemes (with no Cohen-Macaulay hypotheses).*

Proof. Combine Proposition 31.3.11 with Proposition 31.3.6(a) and Proposition 31.1.7. \square

31.3.13. Proposition. — *If X is a Cohen-Macaulay projective k -scheme of pure dimension n , and $\pi: X \rightarrow \mathbb{P}^n$ is a finite flat morphism (as in Proposition 31.3.6(b)), then functorial Serre duality for Ext holds for X with dualizing sheaf $\omega_X := \pi_{sh}^!(\omega_{\mathbb{P}^n})$.*

Unlike Proposition 31.3.9 and 31.3.11, we did not state this for general $\pi: X \rightarrow Y$; we will use the fact that the target is \mathbb{P}^n .

Proof. The argument will parallel that of §31.1.10: we will show an isomorphism of δ -functors

$$(31.3.13.1) \quad \text{Ext}_X^i(\mathcal{F}, \pi_{sh}^!(\omega_{\mathbb{P}^n})) \xrightarrow{\sim} H^{n-i}(X, \mathcal{F})^\vee.$$

We already have an isomorphism for $i = 0$ (Corollary 31.3.12), so it suffices to show that both sides of (31.3.13.1) are *universal* δ -functors. We do this by verifying the criterion of Theorem 24.2.8. For any coherent sheaf \mathcal{F} on X we can find a surjection $\mathcal{O}(-m)^{\oplus N} \rightarrow \mathcal{F}$ for $m \gg 0$ (and some N), so it suffices to show that for $m \gg 0$,

$$\text{Ext}_X^i(\mathcal{O}(-m), \pi_{sh}^!(\omega_{\mathbb{P}^n})) = 0 \quad \text{and} \quad H^{n-i}(X, \mathcal{O}(-m))^\vee = 0$$

for $i > 0$ and $m \gg 0$. For the first, we have (for $i > 0$ and $m \gg 0$)

$$\begin{aligned} \text{Ext}_X^i(\mathcal{O}(-m), \pi_{sh}^!(\omega_{\mathbb{P}^n})) &\cong H^i(X, \pi_{sh}^!(\omega_{\mathbb{P}^n})(m)) \quad (\text{Exercise 31.2.H}) \\ &= 0 \quad (\text{Serre vanishing, Thm. 19.1.3(ii)}) \end{aligned}$$

For the second, we have

$$\begin{aligned} H^{n-i}(X, \mathcal{O}_X(-m)) &= H^{n-i}(X, \mathcal{O}_X \otimes \pi^* \mathcal{O}(-m)) \\ &\cong H^{n-i}(\mathbb{P}^n, (\pi_* \mathcal{O}_X) \otimes \mathcal{O}(-m)) \end{aligned}$$

by the projection formula (Exercise 19.8.E(b)). Now $\pi_* \mathcal{O}_X$ is locally free by Exercise 25.4.E, so by functorial Serre duality for vector bundles on \mathbb{P}^n , we have (for $m \gg 0$)

$$\begin{aligned} H^{n-i}(\mathbb{P}^n, (\pi_* \mathcal{O}_X) \otimes \mathcal{O}(-m)) &\cong H^i(\mathbb{P}^n, (\pi_* \mathcal{O}_X)^\vee \otimes \mathcal{O}(m) \otimes \omega_{\mathbb{P}^n}) \\ &\cong H^i(\mathbb{P}^n, ((\pi_* \mathcal{O}_X)^\vee \otimes \omega_{\mathbb{P}^n}) \otimes \mathcal{O}(m)) \\ &= 0 \quad (\text{Serre vanishing, Thm. 19.1.3(ii)}) \end{aligned}$$

\square

31.3.14. Corollary. — *Functorial Serre duality for Ext holds for all equidimensional Cohen-Macaulay projective k -schemes.*

Proof. Combine Proposition 31.3.13 with Proposition 31.3.6(b) and the projective space case of §31.1.10. \square

31.3.15. ★★ Remark: *the trace version of Serre duality for vector bundles.* Here is an outline of how the above proof can be extended to prove the trace version of Serre

duality for vector bundles on all equidimensional Cohen-Macaulay projective k -schemes. Continuing the notation of Proposition 31.3.13, the Yoneda cup product of §31.2.3 followed by the trace map gives a pairing

$$\mathrm{Ext}_X^i(\mathcal{F}, \omega_X) \times H^{n-i}(X, \mathcal{F}) \xrightarrow{\cup} H^n(X, \omega_X) \xrightarrow{t} k,$$

which is functorial in \mathcal{F} . This induces a functorial map (31.3.13.1). One shows that this particular map is a map of δ -functors. Then by the above argument, this particular map is an isomorphism.

31.4 The adjunction formula for the dualizing sheaf, and $\omega_X = \mathcal{K}_X$

The dualizing sheaf behaves well with respect to slicing by effective Cartier divisors, and this will be useful to obtaining Desideratum 31.1.1, i.e., the miracle that the canonical bundle is Serre-dualizing (§19.5.2). In order to show this, we give a description for the dualizing sheaf for a subvariety of a variety satisfying Serre duality for Ext.

31.4.1. Preliminary Observation. Suppose $\pi: X \rightarrow Y$ is a closed embedding. Then for any \mathcal{O}_Y -module \mathcal{F} , $\mathrm{Ext}_Y^i(\pi_*\mathcal{O}_X, \mathcal{F})$ naturally has the structure of an \mathcal{O}_X -module. (Reason: we compute this by taking an injective resolution of \mathcal{F} by \mathcal{O}_Y -modules, then truncating, then applying $\mathrm{Hom}_Y(\pi_*\mathcal{O}_X, \cdot)$. But for any \mathcal{O}_Y -module \mathcal{G} , $\mathrm{Hom}_Y(\pi_*\mathcal{O}_X, \mathcal{G})$ has the structure of an \mathcal{O}_X -module.) To emphasize its structure as an \mathcal{O}_X -module, we write it as $\pi^{-1}\mathrm{Ext}_Y^i(\pi_*\mathcal{O}_X, \mathcal{F})$.

31.4.A. EXERCISE. Show that if Y (and hence X) is locally Noetherian, and \mathcal{G} is a coherent sheaf on Y , then $\pi^{-1}\mathrm{Ext}_Y^i(\pi_*\mathcal{O}_X, \mathcal{G})$ is a coherent sheaf on X . (Hint: $\mathrm{Ext}_Y^i(\pi_*\mathcal{O}_X, \mathcal{G})$ is a coherent sheaf on Y , Exercise 31.2.F.)

31.4.2. The dualizing sheaf for a subvariety in terms of the dualizing sheaf of the ambient variety.

For the rest of this section, $\pi: X \rightarrow Y$ will be a closed embedding of equidimensional projective k -schemes of dimension n and N respectively, where Y satisfies functorial Serre duality for Ext. Let $r = N - n$ (the codimension of X in Y). By Corollary 31.3.12, X satisfies functorial Serre duality for Hom. We now identify ω_X in terms of ω_Y .

31.4.3. Theorem. — We have $\omega_X = \pi^{-1}\mathrm{Ext}_Y^r(\pi_*\mathcal{O}_X, \omega_Y)$.

Before we prove Theorem 31.4.3, we explain what happened to the “earlier” $\mathrm{Ext}_Y^i(\pi_*\mathcal{O}_X, \omega_Y)$ ’s.

31.4.4. Proposition. — Suppose that $\pi: X \hookrightarrow Y$ is a closed embedding of equidimensional projective k -schemes of dimension n and N respectively, and Y satisfies functorial Serre duality for Ext. Then for all $i < r := N - n$, $\mathrm{Ext}_Y^i(\pi_*\mathcal{O}_X, \omega_Y) = 0$.

Proof. As $\mathrm{Ext}_Y^i(\pi_*\mathcal{O}_X, \omega_Y)$ is coherent (Exercise 31.2.F), it suffices to show that $\mathrm{Ext}_Y^i(\pi_*\mathcal{O}_X, \omega_Y) \otimes \mathcal{O}(m)$ has no nonzero global sections for $m \gg 0$ (as for any

coherent sheaf \mathcal{G} on Y , $\mathcal{G}(m)$ is generated by global sections for $m \gg 0$ by Serre's Theorem A, Theorem 16.3.8). By Exercise 31.2.H,

$$H^j(Y, \mathcal{E}xt_Y^i(\pi_* \mathcal{O}_X, \omega_Y(m))) = H^j(Y, \mathcal{E}xt_Y^i(\pi_* \mathcal{O}_X, \omega_Y(m))).$$

If $j > 0$, then for $m \gg 0$, by Serre vanishing, $H^j(Y, \mathcal{E}xt_Y^i(\pi_* \mathcal{O}_X, \omega_Y(m))) = 0$. Thus by the local-to-global spectral sequence for Ext (Exercise 31.2.K),

$$H^0(Y, \mathcal{E}xt_Y^i(\pi_* \mathcal{O}_X, \omega_Y(m))) = \text{Ext}_Y^i(\pi_* \mathcal{O}_X, \omega_Y(m)).$$

By Exercise 31.2.H again, then functorial Serre duality for Ext on Y ,

$$\text{Ext}_Y^i(\pi_* \mathcal{O}_X, \omega_Y(m)) = \text{Ext}_Y^i(\pi_* \mathcal{O}_X(-m), \omega_Y) = H^{N-i}(Y, \pi_* \mathcal{O}_X(-m))$$

which is 0 if $N - i > n$, as the cohomology of a quasicoherent sheaf on a projective scheme vanishes in degree higher than the dimension of the sheaf's support (dimensional cohomology vanishing, Theorem 19.2.6). \square

The most difficult step in the proof of Theorem 31.4.3 is the following.

31.4.5. Lemma. — Suppose that $\pi: X \rightarrow Y$ is a closed embedding of codimension r of equidimensional projective k -schemes, and Y satisfies functorial Serre duality for Ext . We have an isomorphism, functorial in $\mathcal{F} \in \text{Coh}_X$:

$$(31.4.5.1) \quad \text{Hom}_X(\mathcal{F}, \pi^{-1} \mathcal{E}xt_Y^r(\pi_* \mathcal{O}_X, \omega_Y)) \xrightarrow{\sim} \text{Ext}_Y^r(\pi_* \mathcal{F}, \omega_Y).$$

Proof. Choose an injective resolution

$$0 \longrightarrow \omega_Y \longrightarrow \mathcal{I}^0 \longrightarrow \mathcal{I}^1 \longrightarrow \mathcal{I}^2 \longrightarrow \dots$$

of ω_Y . To compute the right side of (31.4.5.1), we drop the ω_Y from this resolution, and apply $\text{Hom}_Y(\pi_* \mathcal{F}, \cdot)$, to obtain

$$0 \longrightarrow \text{Hom}_Y(\pi_* \mathcal{F}, \mathcal{I}^0) \longrightarrow \text{Hom}_Y(\pi_* \mathcal{F}, \mathcal{I}^1) \longrightarrow \text{Hom}_Y(\pi_* \mathcal{F}, \mathcal{I}^2) \longrightarrow \dots,$$

or (by adjointness of π_* and $\pi_{\text{sh}}^!$ for \mathcal{O} -modules when π is a closed embedding, Exercise 31.3.C(b)):

$$0 \longrightarrow \text{Hom}_X(\mathcal{F}, \pi_{\text{sh}}^! \mathcal{I}^0) \longrightarrow \text{Hom}_X(\mathcal{F}, \pi_{\text{sh}}^! \mathcal{I}^1) \longrightarrow \text{Hom}_X(\mathcal{F}, \pi_{\text{sh}}^! \mathcal{I}^2) \longrightarrow \dots.$$

Motivated by this, consider the complex

$$(31.4.5.2) \quad 0 \longrightarrow \pi_{\text{sh}}^! \mathcal{I}^0 \longrightarrow \pi_{\text{sh}}^! \mathcal{I}^1 \longrightarrow \pi_{\text{sh}}^! \mathcal{I}^2 \longrightarrow \dots.$$

Note first that (31.4.5.2) is indeed a complex, and second that $\pi_{\text{sh}}^! \mathcal{I}^i$ are injective \mathcal{O}_X -modules (Exercise 31.3.C(c)).

31.4.B. EXERCISE. Show that the cohomology of (31.4.5.2) at the i th step is $\mathcal{E}xt_Y^i(\pi_* \mathcal{O}_X, \omega_Y)$.

Thus by Proposition 31.4.4, the complex (31.4.5.2) is exact before the r th step.

31.4.C. EXERCISE. Show that there exists a direct sum decomposition $\pi_{\text{sh}}^! \mathcal{I}^r = \mathcal{J} \oplus \mathcal{K}$, so that

$$0 \longrightarrow \pi_{\text{sh}}^! \mathcal{I}^0 \longrightarrow \pi_{\text{sh}}^! \mathcal{I}^1 \longrightarrow \dots \longrightarrow \pi_{\text{sh}}^! \mathcal{I}^{r-1} \longrightarrow \mathcal{J} \longrightarrow 0$$

is exact. Hint: work out the case $r = 1$ first. Another hint: Notice that if

$$0 \longrightarrow \mathcal{I}^0 \xrightarrow{\alpha} \mathcal{F}$$

is exact and \mathcal{I}^0 is injective, then there is a map $\beta: \mathcal{F} \rightarrow \mathcal{I}^0$ “splitting” α , allowing us to write \mathcal{F} as a direct sum $\mathcal{I}^0 \oplus \mathcal{I}^1$. If \mathcal{F} is furthermore injective, then \mathcal{I}^1 is injective too. This is the beginning of an induction.

31.4.D. EXERCISE. (We continue the notation of the previous exercise.) Identify $\ker(\mathcal{K} \rightarrow \pi_{\text{sh}}^! \mathcal{I}^{r+1})$ with $\pi^{-1} \mathcal{E}xt_Y^r(\pi_* \mathcal{O}_X, \omega_Y)$.

31.4.E. EXERCISE. Put together the pieces above to complete the proof of Lemma 31.4.5. \square

We are now ready to prove Theorem 31.4.3.

31.4.6. Proof of Theorem 31.4.3. Suppose \mathcal{F} is a coherent sheaf on X . We wish to find an isomorphism

$$\text{Hom}_X(\mathcal{F}, \pi^{-1} \mathcal{E}xt_Y^r(\pi_* \mathcal{O}_X, \omega_Y)) \xrightarrow{\sim} H^n(X, \mathcal{F})^\vee,$$

functorial in \mathcal{F} . By Lemma 31.4.5, we have a functorial isomorphism

$$\text{Hom}_X(\mathcal{F}, \pi^{-1} \mathcal{E}xt_Y^r(\pi_* \mathcal{O}_X, \omega_Y)) \xrightarrow{\sim} \text{Ext}_Y^r(\pi_* \mathcal{F}, \omega_Y).$$

By functorial Serre duality for Ext on Y , we have a functorial isomorphism

$$\text{Ext}_Y^r(\pi_* \mathcal{F}, \omega_Y) \xrightarrow{\sim} H^{N-r}(Y, \pi_* \mathcal{F})^\vee.$$

As \mathcal{F} is a coherent sheaf on X , and $N - r = n$, the right side is precisely $H^n(X, \mathcal{F})^\vee$. \square

31.4.7. Applying Theorem 31.4.3. We first apply Theorem 31.4.3 in the special case where X is an effective Cartier divisor on Y . We can compute the dualizing sheaf $\pi^{-1} \mathcal{E}xt_Y^1(\pi_* \mathcal{O}_X, \omega_Y)$ by computing $\mathcal{E}xt_Y^1(\pi_* \mathcal{O}_X, \omega_Y)$ using any locally free resolution (on Y) of $\pi_* \mathcal{O}_X$ (Exercise 31.2.E). But $\pi_* \mathcal{O}_X$ has a particularly simple resolution, the closed subscheme exact sequence (9.1.2.1) for X :

$$(31.4.7.1) \quad 0 \longrightarrow \mathcal{O}_Y(-X) \longrightarrow \mathcal{O}_Y \longrightarrow \pi_* \mathcal{O}_X \longrightarrow 0.$$

We compute $\mathcal{E}xt^\bullet(\pi_* \mathcal{O}_X, \omega_Y)$ by truncating (removing the $\pi_* \mathcal{O}_X$), and applying $\mathcal{H}om(\cdot, \omega_Y)$: $\mathcal{E}xt^\bullet(\pi_* \mathcal{O}_X, \omega_Y)$ is the cohomology of

$$0 \longrightarrow \mathcal{H}om_Y(\mathcal{O}_Y, \omega_Y) \longrightarrow \mathcal{H}om(\mathcal{O}_Y(-X), \omega_Y) \longrightarrow 0,$$

$$\text{i.e.,} \quad 0 \longrightarrow \omega_Y \longrightarrow \omega_Y \otimes \mathcal{O}_Y(X) \longrightarrow 0.$$

We immediately see that $\mathcal{E}xt^i(\pi_* \mathcal{O}_X, \omega_Y) = 0$ if $i \neq 0, 1$. Furthermore, $\mathcal{E}xt^0(\pi_* \mathcal{O}_X, \omega_Y) = 0$ by Proposition 31.4.4. (Unimportant aside: you can use this to show that ω_Y has no embedded points.)

We now consider $\omega_X = \text{coker}(\omega_Y \rightarrow \omega_Y \otimes \mathcal{O}_Y(X))$. Tensoring (31.4.7.1) with the invertible sheaf $\mathcal{O}_Y(X)$, and then tensoring with ω_Y , yields

$$\omega_Y \longrightarrow \omega_Y \otimes \mathcal{O}_Y(X) \longrightarrow \omega_Y \otimes \mathcal{O}_Y(X)|_X \longrightarrow 0$$

The right term is often (somewhat informally) written as $\omega_Y(X)|_X$. Thus

$$\omega_X = \text{coker}(\omega_Y \rightarrow \omega_Y \otimes \mathcal{O}_Y(X)) = \omega_Y(X)|_X,$$

and this identification is *canonical*.

We have shown the following.

31.4.8. Proposition (the adjunction formula). — Suppose that Y is a Cohen-Macaulay projective scheme of pure dimension n (which satisfies functorial Serre duality for Ext by Corollary 31.3.14), and X is an effective Cartier divisor on Y . (Hence X is Cohen-Macaulay by §29.2.4, and thus satisfies functorial Serre duality for Ext by Corollary 31.3.14.) Then $\omega_X = \omega_Y(X)|_X$. In particular, if ω_Y is an invertible sheaf on Y , then ω_X is an invertible sheaf on X .

As an immediate application, we have the following.

31.4.F. EXERCISE. Suppose X is a complete intersection in \mathbb{P}^n , of hypersurfaces of degrees d_1, \dots, d_r . (Note that X is Cohen-Macaulay by Proposition 29.2.6, and thus satisfies functorial Serre duality for Ext , by Corollary 31.3.14.) Show that $\omega_X \cong \mathcal{O}_X(-n - 1 + \sum d_i)$. If furthermore X is smooth, show that $\omega_X \cong \det \Omega_X$. (Hint for the last sentence: use the adjunction formula for \mathcal{K} , Exercise 22.5.B).

But we can say more.

31.4.G. EXERCISE. Suppose Y is a smooth equidimensional variety, and $\pi: X \hookrightarrow Y$ is a codimension r regular embedding with normal sheaf $\mathcal{N}_{X/Y}$ (which is locally free, by Proposition 22.2.16(b)). Suppose \mathcal{L} is an invertible sheaf on Y .

(a) Show that $\mathcal{E}xt_Y^i(\pi_* \mathcal{O}_X, \mathcal{L}) = 0$ if $i \neq r$.

(b) Describe a *canonical* isomorphism $\pi^{-1} \mathcal{E}xt_Y^r(\pi_* \mathcal{O}_X, \mathcal{L}) \xrightarrow{\sim} (\det \mathcal{N}_{X/Y}) \otimes_{\mathcal{O}_X} \mathcal{L}|_X$.

(Note that because Exercise 31.4.G has nothing explicitly to do with duality, we have no projectivity assumptions on Y ; it is a completely local question.) From Exercise 31.4.G we deduce the following.

31.4.H. IMPORTANT EXERCISE. Suppose X is a codimension r regular embedding in \mathbb{P}_k^n . (Then X is Cohen-Macaulay by Proposition 29.2.6, and thus satisfies functorial Serre duality for Ext by Corollary 31.3.14.) Show that

$$\omega_X = \pi^{-1} \mathcal{E}xt_{\mathbb{P}^n}^r(\pi_* \mathcal{O}_X, \omega_{\mathbb{P}^n}) \cong (\det \mathcal{N}_{X/\mathbb{P}^n}) \otimes \omega_{\mathbb{P}^n}|_X.$$

In particular, ω_X is an invertible sheaf.

31.4.I. IMPORTANT EXERCISE. Suppose X is a smooth pure codimension r subvariety of \mathbb{P}_k^n (and hence a regular embedding, by Exercise 13.2.L(b)). Show that $\omega_X \cong \mathcal{K}_X$. Hint: both sides satisfy adjunction (see Exercise 22.5.B for adjunction for Ω): they are isomorphic to $(\det \mathcal{N}_{X/Y}) \otimes \omega_{\mathbb{P}^n}|_X \cong (\det \mathcal{N}_{X/Y}) \otimes (\mathcal{K}_{\mathbb{P}^n})|_X$.

31.4.9. Remark. As long promised, the version of Serre duality given in Theorem 19.5.1 now follows by combining Corollary 31.3.10 with Exercise 31.4.I.

Aux yeux de ces amateurs d'inquiétude et de perfection, un ouvrage n'est jamais achevé, mot qui pour eux n'a aucun sens, mais abandonné; et cet abandon, qui le livre aux flammes ou au public (et qu'il soit l'effet de la lassitude ou de l'obligation de livrer), leur est une sorte d'accident, comparable à la rupture d'une réflexion . . .

In the eyes of those lovers of anxiety and perfection, a work is never finished — a word that for them has no meaning — but abandoned; and this abandonment, whether to

the flames or to the public (and which is the result of weariness or an obligation to deliver) is a kind of accident to them, like the breaking off of a reflection . . .

— P. Valéry, [Val, p. 1497]

Well, I've been talking to you for two years and now I'm going to quit. In some ways I would like to apologize, and other ways not. I hope – in fact, I know – that two or three dozen of you have been able to follow everything with great excitement, and have had a good time with it. But I also know that “the powers of instruction are of very little efficacy except in those happy circumstances in which they are practically superfluous.” So, for the two or three dozen who have understood everything, may I say I have done nothing but shown you the things. For the others, if I have made you hate the subject, I'm sorry. I never taught elementary physics before, and I apologize. I just hope that I haven't caused a serious trouble to you, and that you do not leave this exciting business. I hope that someone else can teach it to you in a way that doesn't give you indigestion, and that you will find someday that, after all, it isn't as horrible as it looks.

— R. Feynman, [FLS, Epilogue]

Bibliography

- [Al] P. Aluffi, *Algebra, Chapter 0*, Grad. Stud. in Math. **104**, Amer. Math. Soc., Providence, RI, 2009. [0.3.2](#)
- [Ar1] M. Artin, *On the solutions of analytic equations*, Invent. Math. **5** (1968), no. 4, 277–291. [30.3.6](#)
- [Ar2] ———, *Algebraic approximation of structures over complete local rings*, Publ. Math. IHES **36** (1969), 23–58. [30.3.6](#)
- [Ar3] ———, *Lipman’s proof of resolution of singularities for surfaces*, in *Arithmetic geometry* (Storrs, Conn., 1984), Springer-Verlag, New York, 1986, 267–287. [23.4.6](#)
- [Ar4] ———, *Algebra*, Prentice Hall, Inc., Englewood Cliffs, NJ, 1991. [5.4.N](#)
- [At1] M. F. Atiyah, *On analytic surfaces with double points*, Proc. Roy. Soc. London Ser. A **247** (1958), 237–244. [23.4.11](#)
- [At2] ———, *Mathematics in the 20th century*, Amer. Math. Monthly **108** (2001), no. 7, 654–666. [0.3](#), [8.2](#)
- [AtM] M. F. Atiyah and I. G. Macdonald, *Introduction to Commutative Algebra*, Addison-Wesley Publ. Co., Reading, MA, 1969. [0.3.2](#), [12.2.12](#), [12.3.16](#), [30.2.D](#)
- [Ba] L. Bădescu, *Algebraic Surfaces*, V. Mařek trans., Universitext, Springer-Verlag, New York, 2001. [11.5.6](#), [30.7.3](#)
- [BCDKT] M. Baker, B. Conrad, S. Dasgupta, K. Kedlaya, and J. Teitelbaum, *p-adic geometry*, D. Savitt and D. Thakur ed., Univ. Lect. Series **45**, Amer. Math. Soc., Providence, RI, 2008. [6.6](#)
- [BPV] W. Barth, C. Peters, and A. van de Ven, *Compact Complex Surfaces*, Ergeb. Math. Grenz. **4**, Springer-Verlag, Berlin, 1984. [22.5.I](#), [30.7.3](#)
- [Be] A. Beauville, *Complex Algebraic Surfaces*, 2nd ed., R. Barlow trans., Cambridge U.P., Cambridge, 1996. [27.1.3](#), [30.7.3](#)
- [BL] C. Birkenhake and H. Lange, *Complex Abelian Varieties*, 2nd ed., Grundle Math. Wiss. **302**, Springer-Verlag, Berlin, 2004. [11.5.11](#), [11.5.14](#)
- [BLR] S. Bosch, W. Lütkebohmert, and M. Raynaud, *Néron Models*, Ergeb. Math. Grenz. **21**, Springer-Verlag, Berlin, 1990. [26.2.5](#)
- [Bo] N. Bourbaki, *Éléments de mathématique*, Fasc. XXX: *Algèbre Commutative*, Actualités Scientifiques et Industrielles, no. 1308, Hermann, Paris, 1964. [5.4.G](#)
- [BGS] S. Boissière, O. Gabber, and O. Serman, *Sur le produit de variétés localement factorielles ou \mathbb{Q} -factorielles*, arxiv:1104.1861v1, preprint 2011. [5.4.3](#)
- [BP] R. Brown and T. Porter, *Analogy, concepts and methodology, in mathematics*, UWB Math Preprint, www.maths.bangor.ac.uk/research/ftp/rpam/06-08.pdf, May 26, 2006. [1.1](#)
- [CL] F. Call and G. Lyubeznik, *A simple proof of Grothendieck’s theorem on the parafactoriality of local rings*, in *Commutative algebra: syzygies, multiplicities, and birational algebra* (South Hadley, MA, 1992), Contemp. Math. **159**, Amer. Math. Soc., Providence, RI, 1994, 15–18. [13.8.6](#)
- [Cart] P. Cartier, *Les groupes $\text{Ext}^s(A, B)$* , in *Séminaire A. Grothendieck: Algèbre homologique*, t. 1, 1957, no. 3. [31.2.3](#)
- [Ch] W. Churchill, *The Russian Enigma*, BBC broadcast, October 1, 1939. [25](#)
- [Co] B. Conrad, *Deligne’s notes on Nagata compactifications*, J. Ramanujan Math. Soc. **22** (2007), no. 3, 205–257. [30.6.B](#)
- [Dan] V. I. Danilov, *Samuel’s conjecture*, Mat. Sb. (N.S.) **81 (123)** (1970), 132–144. [5.4.5](#)
- [Dav] E. D. Davis, *Ideals of the principal class, R-sequences and a certain monoidal transformation*, Pacific J. Math. **20** (1967), 197–205. [23.3.11](#)

- [DW] R. Dedekind and H. Weber, *Theory of algebraic functions of one variable*, J. Stillwell trans. History of Math. **39**, Amer. Math. Soc., Providence, RI; London Math. Soc. London, 2012. [19.5.6](#)
- [De] P. Deligne, *Courbes elliptiques: formulaire d'après Tate*, in *Modular Functions of One Variable IV (Proc. Internat. Summer School, Univ. Antwerp, Antwerp, 1972)*, 53–73, Lect. Notes in Math. **476**, Springer-Verlag, Berlin, 1975. [20.9.B](#)
- [DeI] P. Deligne and L. Illusie, *Relèvements modulo p^2 et décomposition du complexe de de Rham*, Invent. Math. **89** (1987), no. 2, 247–270. [22.5.9](#)
- [Di] J. Dieudonné, *On regular sequences*, Nagoya Math. J. **27** (1966), 355–356. [9.5.6](#)
- [Do] I. Dolgachev, *Luigi Cremona and cubic surfaces*, in *Luigi Cremona (1830–1903)*, 55–70, Incontr. Studio, **36**, Istituto Lombardo di Scienze e Lettere, Milan, 2005. [27.1](#)
- [DS] R. Donagi and R. C. Smith, *The structure of the Prym map*, Acta Math. **146** (1981), no. 1–2, 25–102. [27.1](#)
- [DZ] H. Duan and X. Zhao, *Schubert calculus and the integral cohomology of exceptional Lie groups*, arXiv:0711.2541v10, preprint 2007. [21.3.5](#)
- [DF] D. Dummit and R. Foote, *Abstract Algebra*, 3rd ed., Wiley and Sons, Inc., Hoboken, NJ, 2004. [8.2.1](#), [13.5.14](#), [20.9.B](#)
- [E] D. Eisenbud, *Commutative Algebra with a View to Algebraic Geometry*, Grad. Texts in Math. **150**, Springer-Verlag, New York, 1995. [0.3.2](#), [1.4.10](#), [5.4.G](#), [6.5.36](#), [9.5.5](#), [10.5.18](#), [10.7.3](#), [10.7.4](#), [12.2.12](#), [13.8.2](#), [13.8.5](#), [16.3.2](#), [22.2.10](#), [22.3.5](#), [22.4.7](#), [29.2.14](#), [30.2.C](#)
- [EGH] D. Eisenbud, M. Green, and J. Harris, *Cayley-Bacharach theorems and conjectures*, Bull. AMS **33** (1996), no. 3, 295–324. [20.4.B](#)
- [Ek] P. Eklof, *Lefschetz's principle and local functors*, Proc. Amer. Math. Soc. **37** (1973), no. 2, 1973. [22.5.9](#)
- [EE] E. Enochs and S. Estrada, *Relative homological algebra in the category of quasi-coherent sheaves*, Adv. Math. **194** (2005), 284–295. [24.5.7](#)
- [FGIKNV] B. Fantechi, L. Göttsche, L. Illusie, S. Kleiman, N. Nitsure and A. Vistoli, *Fundamental algebraic geometry: Grothendieck's FGA explained*, Mathematical Surveys and Monographs **123**, Amer. Math. Soc., Providence, RI, 2005. [20.10.1](#), [20.11.12](#), [21.1.4](#), [21.1.6](#), [21.2.16](#), [28.3.1](#)
- [FK] H. M. Farkas and I. Kra, *Riemann Surfaces*, Grad. Texts in Math. **71**, Springer-Verlag, New York, 1980. [22.7.9](#)
- [FLS] R. Feynman, R. Leighton, and M. Sands, *The Feynman Lectures on Physics, Volume III*, Addison Wesley, 1971. [31.4.9](#)
- [Fi] F. Scott Fitzgerald, *The Crack-Up: A desolately frank document from one for whom the salt of life has lost its savor*, Esquire Magazine, February 1936, Esquire Inc., Chicago, IL. [3.3](#)
- [FR] G. Frey and H.-G. Rück, *The strong Lefschetz principle in algebraic geometry*, manuscripta math. **55** (1986), no. 3–4, 385–401. [22.5.9](#)
- [F1] W. Fulton, *Young Tableaux: With Applications to Representation Theory and Geometry*, Cambridge U.P., Cambridge, 1997. [17.7.3](#)
- [F2] W. Fulton, *Intersection Theory*, 2nd ed., Ergeb. Math. Grenz. **2**, Springer-Verlag, New York, 1998. [21.1.1](#), [21.3.B](#), [21.3.B](#), [21.3.B](#), [23.3.11](#), [23.4.17](#)
- [Ge] S. Germain, *Œuvres Philosophiques de Sophie Germain* (nov. éd.), Librairie de Firmin-Didot et Cie, Paris, 1896, http://fr.wikisource.org/wiki/OEuvres_philosophiques_de_Sophie_Germain. [3.2.4](#)
- [GS] P. Gille and T. Szamuely, *Central Simple Algebras and Galois Cohomology*, Cambridge U.P., Cambridge, 2006. [28.1.L](#)
- [GW] U. Görtz and T. Wedhorn, *Algebraic Geometry I*, Vieweg + Teubner, Wiesbaden, 2010. (Appendices C, D, and E, collecting properties of schemes and morphisms and their relationships, are alone worth the price.) [0.1](#), [17.6.3](#), [19.5.7](#), [19.7.6](#), [19.9.7](#), [30.6.3](#)
- [GR] H. Grauert and R. Remmert, *Coherent Analytic Sheaves*, Grundle Math. Wiss. **265**, Springer-Verlag, Berlin, 1984. [6.4.6](#), [7.3.17](#)
- [GH1] P. Griffiths and J. Harris, *Principles of Algebraic Geometry*, Wiley-Interscience, New York, 1978. [9.3.10](#), [17.6.2](#), [19.4.9](#), [22.5.11](#), [22.5.O](#)
- [GH2] ———, *On the Noether-Lefschetz theorem and some remarks on codimension two cycles*, Math. Ann. **271** (1985), 31–51. [12.2.18](#)
- [Gr1] A. Grothendieck, *Sur quelques points d'algèbre homologique*, Tôhoku Math. J. (2) **9** (1957), 119–221. [1.1](#), [1.7.7](#), [24](#), [24.2.9](#), [24.5.1](#)

- [Gr2] ———, *Sur la classification des fibrés holomorphes sur la sphère de Riemann*, Amer. J. Math. **79** (1957), 121–138. [19.5.6](#)
- [Gr3] ———, *Théorèmes de dualité pour les faisceaux algébriques cohérents*, Séminaire N. Bourbaki **4**, exp. no. 149, Soc. Math. France, Paris, 1995, 169–193. [31.2.3](#)
- [Gr4] ———, *The cohomology theory of abstract algebraic varieties*, in *Proc. Internat. Congress Math. (Edinburgh, 1958)*, Cambridge U.P., New York, 1960. [0.2](#)
- [Gr5] ———, *Fondements de la géométrie algébrique*, in *Collected Bourbaki Talks*, Secrétariat Mathématique, Paris, 1962. [28.3.1](#)
- [Gr6] ———, *Récoltes et Semailles*, unpublished. ([document](#)), [0.0.3](#), II
- [Gr-EGA] ——— (w. J. Dieudonné), *Éléments de Géométrie Algébrique*, Publ. Math. IHES, Part I: **4** (1960); Part II: **8** (1961); Part III: **11** (1961), 17 (1963); Part IV: **20** (1964), **24** (1965), **28** (1966), **32** (1967). [0.1](#), [5.4.M](#), [8.3.18](#), [10.3.I](#), [13.7.E](#), [13.7.4](#), [13.7.6](#), [13.7.9](#), [13.8.5](#), [15.3.E](#), [17.6.2](#), [17.6.G](#), [18.3.3](#), [19.9.7](#), [22.6](#), [25.4.11](#), [25.5.5](#), [25.5.17](#), [25.6.9](#), [26](#), [26.2.5](#), [26.2.6](#), [26.2.7](#), [28.2.13](#), [30.6.2](#), [30.6.3](#)
- [Gr-EGA'] ——— (w. J. Dieudonné), *Éléments de Géométrie Algébrique I: Le langage des schémas*, Grund. Math. Wiss. **166** (2nd ed.), Springer-Verlag, Berlin-New York, 1971. [13.7.6](#)
- [GrS] A. Grothendieck and J.-P. Serre, *Grothendieck-Serre Correspondence*, P. Colmez and J.-P. Serre ed., C. Maclean trans., Amer. Math. Soc. and Soc. Math. de France, Paris, 2004. [1.6](#), [1.7](#), [13.7.10](#), [19.2.7](#), [24](#), [24.4.2](#), [31](#), [31.1.9](#)
- [H1] J. Harris, *Galois groups of enumerative problems*, Duke Math. J. **46** (1979), no. 4, 685–724. [27.3.5](#)
- [H2] ———, *Algebraic Geometry: A First Course*, Grad. Texts in Math. **133**, Springer-Verlag, New York, 1995. [13.4.5](#)
- [Ha1] R. Hartshorne, *Algebraic Geometry*, Grad. Texts in Math. **52**, Springer-Verlag, New York-Heidelberg, 1977. [13.7.6](#), [15.2.P](#), [18.3.2](#), [19.2.7](#), [19.5.7](#), [19.7.2](#), [24.5.7](#)
- [Ha2] ———, *Ample Subvarieties of Algebraic Varieties*, Lect. Notes in Math. **156**, Springer-Verlag, Berlin-New York, 1970. [7.3.16](#)
- [Hat] A. Hatcher, *Algebraic Topology*, Cambridge U.P., Cambridge, 2002. [22.4.2](#)
- [HM] M. Hazewinkel and C. Martin, *A short elementary proof of Grothendieck's theorem on algebraic vector bundles over the projective line*, J. Pure Appl. Algebra **25** (1982), no. 2, 207–211. [19.5.7](#)
- [He] R. C. Heitmann, *Examples of noncatenary rings*, Trans. Amer. Math. Soc. **247** (1979), 125–136. [12.2.10](#)
- [Hil] D. Hilbert, *Ueber die Theorie der algebraischen Formen*, Math. Ann. **36** (1890), 473–534. [3.6.17](#)
- [Hir] H. Hironaka, *Resolution of singularities of an algebraic variety over a field of characteristic zero. I, II*, Ann. of Math. (2) **79** (1964), 109–203; *ibid.* (2) **79** (1964), 205–326. [23.4.6](#)
- [Ho1] M. Hochster, *Prime Ideal Structure in Commutative Rings*, Ph.D. Thesis, Princeton University, 1967. [5.1.E](#)
- [Ho2] ———, *Prime ideal structure in commutative rings*, Trans. Amer. Math. Soc. **142** (1969), 43–60. [5.1.E](#)
- [Ig] J. Igusa, *On some problems in abstract algebraic geometry*, Proc. Nat. Acad. Sci. **41** (1955), 964–967. [22.5.11](#)
- [Ii] S. Iitaka, *Algebraic Geometry: An Introduction to Birational Geometry of Algebraic Varieties*, Grad. Texts in Math. **76**, Springer-Verlag, New York-Berlin, 1982. [22.5.E](#)
- [Ill] L. Illusie, *Frobenius and Hodge degeneration*, in *Introduction to Hodge Theory*, SMF/AMS Texts and Monographs **8**, 99–149, 2002. [22.5.9](#)
- [dJ] A. J. de Jong, *Smoothness, semi-stability and alterations*, IHES Publ. Math. **83** (1996), 51–93. [23.4.6](#)
- [KN] S. Kaji and M. Nakagawa, *The Chow rings of the algebraic groups of type E_6 , E_7 , and E_8* , arXiv:0709.3702v3, preprint 2007. [21.3.5](#)
- [Ka] I. Kaplansky, *Commutative Rings* (rev. ed.), U. of Chicago Press, Chicago, Ill.-London, 1974. [12.3.5](#)
- [KS1] M. Kashiwara and P. Schapira, *Sheaves on Manifolds*, Grund. Math. Wiss. **292**, Springer-Verlag, Berlin, 1994. [1.6.4](#), [1.6.5](#), [1.6.6](#), [31.3.4](#)
- [KS2] ———, *Categories and sheaves*, Grund. Math. Wiss. **332**, Springer-Verlag, Berlin, 2006. [0.3.1](#)

- [KPV] K. Kedlaya, B. Poonen, and R. Vakil, *The William Lowell Putnam Mathematical Competition 1985–2000: Problems, Solutions, and Commentary*, Math. Ass. of Amer., Washington, DC, 2002. [0.0.1](#)
- [Kl1] S. Kleiman, *Toward a numerical theory of ampleness*, Ann. Math. (2) **84** (1966), no. 3, 293–344. [19.4.9](#), [21.1.3](#), [21.4.2](#), [21.4.6](#), [21.4.7](#), [21.4.C](#)
- [Kl2] ———, *The transversality of a general translate*, Compositio Math. **28** (1974), no. 3, 287–297. [26.3.7](#)
- [Kl3] ———, *Cartier divisors versus invertible sheaves*, Comm. Alg. **28** (2000), no. 12, 5677–5678. [15.3.E](#)
- [Ko1] J. Kollár, *Rational Curves on Algebraic Varieties*, Ergeb. Math. Grenz. **32**, Springer-Verlag, Berlin, 1996. [12.3.21](#), [21.1.3](#), [21.1.J](#), [21.1.6](#), [21.3.3](#)
- [Ko2] ———, *Lectures on Resolutions of Singularities*, Annals of Math. Stud. **166**, Princeton U.P., Princeton, NJ, 2007. [23.4.6](#)
- [Lan] S. Lang, *Algebra* (rev. 3rd ed.), Grad. Texts in Math. **211**, Springer-Verlag, New York, 2002. [5.4.5](#), [12.2.15](#), [15.2.8](#), [21.2.U](#), [24.2.M](#), [24.2.9](#)
- [Laz] R. Lazarsfeld, *Positivity in Algebraic Geometry I: Classical setting: line bundles and linear series*, Ergeb. Math. Grenz. **48**, Springer-Verlag, Berlin, 2004. [25.7.5](#)
- [Lef] S. Lefschetz, *L'analysis situs et la géométrie algébrique*, Gauthier-Villars, Paris, 1924. [12.2.18](#)
- [Lew] R. Lewontin, *The Triple Helix: Gene, Organism, and Environment*, Harvard U.P., Cambridge MA, 2000. [4.2](#)
- [Lit] J. E. Littlewood, *Littlewood's Miscellany*, Cambridge U.P., Cambridge, 1986. [3.3](#)
- [Liu] Q. Liu, *Algebraic Geometry and Arithmetic Curves*, R. Ern  trans., Oxford Grad. Texts in Math. **6**, Oxford U.P., Oxford, 2002. [0.1](#), [5.1.E](#), [13.2.15](#), [19.2.9](#), [22.2.10](#)
- [L ] W. L tkebohmert, *On compactification of schemes*, Manuscripta Math. **80** (1), 95–111. [30.6.B](#)
- [Mac] S. Mac Lane, *Categories for the Working Mathematician*, 2nd ed., Grad. Texts in Math. **5**, Springer-Verlag, New York, 1998. [1.6.4](#)
- [Mat1] H. Matsumura, *Commutative Algebra*, 2nd ed., Math. Lecture Note Series **56**, Benjamin/Cummings Publ. Co., Inc., Reading, MA, 1980. [12.2.12](#), [29.1.3](#), [29.2.14](#)
- [Mat2] ———, *Commutative Ring Theory*, M. Reid trans., 2nd ed, Cambridge Stud. in Adv. Math. **8**, Cambridge U.P., Cambridge, 1989. [0.3.2](#), [5.4.5](#), [13.8.2](#), [13.8.5](#), [22.3.5](#), [25.4.6](#), [25.5.17](#), [29.2.13](#), [29.2.14](#)
- [MO33489] <http://mathoverflow.net/questions/33489/> [14.2.4](#)
- [MO68421] Answers by G. Elencwajg and M. Kim to <http://mathoverflow.net/questions/68421/> [7.3.16](#), [20.11.3](#)
- [MO70920] Answers by T. Ekedahl to <http://mathoverflow.net/questions/70920/> [28.1.3](#)
- [MO90551] Answer by M. Brandenburg to <http://mathoverflow.net/questions/90551/> [22.5.9](#)
- [MO129242] Answer by L. Moret-Bailly to <http://mathoverflow.net/questions/129242/> [9.5.H](#)
- [MO286390] Answers to <https://mathoverflow.net/questions/286390/> [20.7.B](#)
- [Mc] C. McLarty, *The Rising Sea: Grothendieck on simplicity and generality*, in J. J. Gray and K. H. Parshall eds., *Episodes in the History of Modern Algebra (1800–1950)*, Amer. Math. Soc., Providence, RI, 2007. [\(document\)](#), [0.0.3](#), [II](#), [3](#)
- [Mu1] D. Mumford, *Pathologies of modular algebraic surfaces*, Amer. J. Math. **83** (no. 2), Apr. 1961, 339–342. [22.5.11](#)
- [Mu2] ———, *Lectures on Curves on an Algebraic Surface*, Princeton U.P., Princeton NJ, 1966. [3.2.12](#), [28.3.1](#)
- [Mu3] ———, *Abelian Varieties*, TIFR Studies in Math. **5**, Oxford U.P., London, 1970. [21.2.H](#), [28.2](#)
- [Mu4] ———, *Curves and Their Jacobians*, Univ. of Michigan Press, Ann Arbor, MI, 1975. [\(document\)](#)
- [Mu5] ———, *Hilbert's fourteenth problem — the finite generation of subrings such as rings of invariants*, in *Mathematical Developments Arising from Hilbert Problems, Part 2*, F. Browder ed., Amer. Math. Soc., Providence RI, 1976. [20.11.13](#)
- [Mu6] ———, *Algebraic Geometry I: Complex Projective Varieties*, Classics in Math., Springer-Verlag, Berlin, 1995. [27.3.D](#)

- [Mu7] ———, *The Red Book of Varieties and Schemes* (2nd expanded ed.), Lect. Notes in Math. **1358**, Springer-Verlag, Berlin, 1999. (Caution: in order to keep readers on their toes, a number of typos were introduced between the first and second editions.) ([document](#)), [13.8.5](#), [25](#), [30.6.3](#)
- [Mur] D. Murfet, <http://therisingsea.org/notes/ModulesOverAScheme.pdf> [0.0.3](#), [24.5.7](#)
- [N] F. Nietzsche, *Twilight of the Idols, or, How to Philosophize with a Hammer*, 1889. [1](#)
- [OSS] C. Okonek, M. Schneider, and H. Spindler, *Vector Bundles on Complex Projective Spaces*, corrected reprint of the 1980 edition, Birkhäuser/Springer Basel AG, Basel, 2011. [19.5.7](#)
- [Pi] P. Picasso, *Statement to Marius de Zayas*, 1923, in *Picasso Speaks*, The Arts, New York, May 1923, pp. 315–26. [3.3](#)
- [Po1] H. Poincaré, *Analysis situs*, J. de l'École Polytechnique (2) **1** (1895), 1–123. [4.2.1](#)
- [Po2] H. Poincaré, *Papers on topology: Analysis situs and its five supplements*, J. Stillwell trans., Amer. Math. Soc., Providence, RI, 2010. [4.2.1](#)
- [P1] B. Poonen, *Varieties without extra automorphisms, I: Curves*, Math. Res. Lett. **7** (2000), no. 1, 67–76. [20.7.E](#)
- [P2] ———, *Rational Points on Varieties*, <http://www-math.mit.edu/~poonen/papers/Qpoints.pdf>, preprint. [11.5.10](#)
- [Ra] M. Raynaud, *Contre-exemple au vanishing theorem en caractéristique $p > 0$* , in C. P. Ramanujam — A Tribute, 273–278, TIFR Studies in Math. **8**, Springer, Berlin-New York, 1978. [22.5.9](#)
- [RS] M. Reed and B. Simon, *Methods of Modern Mathematical Physics I: Functional Analysis*, Academic Press, Inc., New York, 1980. [12.1](#)
- [Rei] M. Reid, *Undergraduate Algebraic Geometry*, London Math. Soc. Student Texts 12, Cambridge U.P., Cambridge, 1988. [27.1.3](#)
- [Rem] R. Remmert, *Local theory of complex spaces*, in *Several complex variables VII: Sheaf-theoretical methods in complex analysis*, Encl. of Math. Sci. **74**, Springer-Verlag, Berlin, 1994. [6.4.6](#)
- [Roh] F. Rohrer, *Quasicoherent sheaves on toric schemes*, Expo. Math. **32** (2014), 33–78. [16.1.2](#)
- [RotV] M. Roth and R. Vakil, *The affine stratification number and the moduli space of curves*, CRM Proceedings and Lecture Notes **38**, Univ. de Montréal, 2004, 213–227. [19.2.6](#), [23.4.16](#)
- [Sa] P. Salmon, *Sulla fattorialità delle algebre graduate e degli anelli locali*, Rend. Sem. Mat. Univ. Padova **41** (1968), 119–138. [5.4.5](#)
- [Schm] H. L. Schmid, *Über die Automorphismen eines algebraischen Funktionenkörpers von Primzahlcharakteristik*, J. Reine Angew. Math. **179** (1938), 5–14. [22.7.9](#)
- [Schw] K. Schwede, *Gluing schemes and a scheme without closed points*, in *Recent Progress in Arithmetic and Algebraic Geometry*, 157–172, Contemp. Math. **386**, Amer. Math. Soc., Providence, RI, 2005. [5.1.E](#)
- [Se1] J.-P. Serre, *Faisceaux algébriques cohérents*, Ann. of Math. (2) **61** (1955), 197–278. [6.4.6](#), [6.4.8](#)
- [Se2] ———, *Sur la dimension cohomologique des anneaux et des modules noethériens*, in Proc. Intern. Symp. on Alg. Number Theory, 176–189, Tokyo-Nikko, 1955. [9.5.3](#)
- [Se3] ———, *Géométrie algébrique et géométrie analytique*, Ann. de l'institut Fourier **6** (1956), 1–42. [7.3.16](#), [11.5](#), [22.5.9](#), [25.1](#)
- [Se4] ———, *Sur la topologie des variétés algébriques en caractéristique p* , in *Symposium Internacional de Topología Algebraica, International Symposium on Algebraic Topology*, 24–53, Univ. Nacional Autónoma de México and UNESCO, Mexico City, 1958. [22.5.11](#)
- [Se5] ———, *A Course in Arithmetic*, Grad. Texts in Math. **7**, Springer-Verlag, New York, 1973. [20.9.B](#)
- [SGA2] A. Grothendieck, *Séminaire de Géométrie Algébrique du Bois Marie 1962: Cohomologie locale des faisceaux cohérents et Théorèmes de Lefschetz locaux et globaux (SGA2)*, North-Holland Publ. Co., Amsterdam, 1968. [13.8.6](#)
- [SGA6] P. Berthelot, A. Grothendieck, and L. Illusie eds., *Séminaire de Géométrie Algébrique du Bois Marie 1966–67: Théorie des intersections et théorème de Riemann-Roch (SGA 6)*, Lect. Notes in Math. **225**, Springer-Verlag, Berlin-New York, 1971. [21.3.5](#), [25.7.5](#)
- [Sh] I. Shafarevich, *Basic Algebraic Geometry 1: Varieties in Projective Space*, 2nd ed., M. Reid trans., Springer-Verlag, Berlin, 1994. [27.1.3](#)
- [Si1] J. Silverman, *The Arithmetic of Elliptic Curves*, 2nd ed., Grad. Texts in Math. **206**, Springer-Verlag, New York-Berlin, 2009. [20.11.8](#)

- [Si2] ———, *Advanced Topics in the Arithmetic of Elliptic Curves*, Grad. Texts in Math. **151**, Springer-Verlag, New York-Berlin, 1994. [20.9.B](#)
- [Stacks] The Stacks Project Authors, *Stacks Project*, <http://stacks.math.columbia.edu>. (To look up tags: <http://stacks.math.columbia.edu/tags>. To search: <http://stacks.math.columbia.edu/search>.) [0.1](#), [1.6.13](#), [3.6.13](#), [5.3.C](#), [6.2.9](#), [7.5.2](#), [9.2.B](#), [9.5.7](#), [10.3.7](#), [10.3.G](#), [10.7.D](#), [12.2.10](#), [12.3.17](#), [13.7.E](#), [13.7.4](#), [13.7.9](#), [17.6.J](#), [18.1.G](#), [18.2](#), [18.2.A](#), [18.2.2](#), [18.2.6](#), [18.3](#), [18.3.H](#), [22.2.10](#), [22.3.8](#), [25.5.15](#), [25.5.17](#), [25.6.7](#), [25.6.9](#), [26.2.6](#), [26.2.7](#), [28.2.14](#), [29.2.13](#), [29.2.14](#), [30.5.3](#)
- [Stic] H. Stichtenoth, *Über die Automorphismengruppe eines algebraischen Funktionenkörpers von Primzahlcharakteristik. I. Eine Abschätzung der Ordnung der Automorphismengruppe*, Arch. Math. (Basel) **24** (1973), 527–544. [22.7.10](#)
- [Ta] A. Tarski, *A decision method for elementary algebra and geometry*, RAND Corp., Santa Monica, CA, 1948. [8.4.H](#)
- [TWi] R. Taylor and A. Wiles, *Ring-theoretic properties of certain Hecke algebras*, Ann. of Math. (2) **141** (3) (1995), 553–572. [20.6.1](#)
- [To] L. Tolstoy, *Anna Karenina*, <http://www.gutenberg.org/files/1399/1399-h/1399-h.htm>. [25.1](#)
- [Vak1] R. Vakil, *The Rising Sea: Foundations of Algebraic Geometry*, preprint. [1.6.H](#)
- [Vak2] ———, *Puzzling through exact sequences: A bedtime story with pictures*, <https://www.3blue1brown.com/content/blog/exact-sequence-picturebook/PuzzlingThroughExactSequences>. [1.7.7](#)
- [Val] P. Valéry, *Au sujet du Cimetière Marin*, in *Œuvres*, édition Jean Hytier, tome I, Gallimard, La Pléiade, Paris, 1957. [31.4.9](#)
- [vdW] B. L. van der Waerden, *Algebra* vol. II, J. R. Schulenberg trans., Springer-Verlag, New York, 1991. [10.5.18](#)
- [Vo] C. Voisin, *Hodge Theory and Complex Algebraic Geometry I and II*, L. Schneps trans., Cambridge Stud. in Adv. Math. **76** and **77**, Cambridge U.P., Cambridge, 2007. [22.5.11](#)
- [Weib] C. Weibel, *An Introduction to Homological Algebra*, Cambridge Stud. in Adv. Math. **38**, Cambridge U.P., Cambridge, 1994. [0.3.2](#), [1.6.4](#), [1.7](#), [24.2.M](#), [24.2.9](#)
- [Weil] A. Weil, *Scientific works — Collected papers Vol. II (1951–1964)*, Springer-Verlag, New York-Heidelberg, 1979. [20.10](#)
- [Wey] H. Weyl, *Philosophy of Mathematics and Natural Science*, Princeton U.P., Princeton, NJ, 2009. [4.4.F](#)
- [Wi] A. Wiles, *Modular elliptic curves and Fermat's last theorem*, Ann. of Math. (2) **141** (3) (1995), 443–551. [20.6.1](#)
- [Z] O. Zariski, *Foundations of a general theory of birational correspondences*, Trans. Amer. Math. Soc. **53** (1943), 490–542. [30.6](#)