FOUNDATIONS OF ALGEBRAIC GEOMETRY CLASS 11

CONTENTS

Last day: finite type A-scheme, locally of finite type A-scheme, projective schemes over A or k.

Today: Smoothness=regularity=nonsingularity, Zariski tangent space and related notions, Nakayama's Lemma.

Warning: I've changed problem B6 to make it more general (reposted on web). The proof is the same as the original problem, but I'll use it in this generality.

1. Projective k-schemes and projective A-schemes

Last day, I defined $\operatorname{Proj} S_*$ where: S_* is a graded ring (with grading $\mathbb{Z}^{\geq 0}$). Last day I said: Suppose S_0 is an A-algebra. I've changed my mind: I'd like to take $S_0 = A$. $S_+ := \oplus_{i>0} S_i$ is the *irrelevant ideal*; suppose that it is finitely generated over S.

Set: The points of $\operatorname{Proj} S_*$ are defined to be the homogeneous prime ideals, except for any ideal containing the irrelevant ideal.

Topology: The closed subsets are of the form $V(I)$, where I is a homogeneous ideal. Particularly important open sets will the *distinguished open sets* $D(f) = \operatorname{Proj} S_* - V(f)$, where $f \in S_+$ is homogeneous. They form a base.

Structure sheaf: $\mathcal{O}_{\operatorname{Proj} S_*}(D(f)) := (S_f)_0$, where $(S_f)_0$ means the 0-graded piece of the graded ring (S_f). This is a sheaf. One method:

$$(D(f), \mathcal{O}_{\operatorname{Proj} S_*}|_{D(f)}) \cong \operatorname{Spec}(S_f)_0.$$

1.1. If S_* is generated by S_1 (as an S_0-algebra — we say S_* is *generated in degree* 1), say by $n+1$ elements $x_0, \ldots, x_n$, then $\operatorname{Proj} S_*$ “sits in $\mathbb{P}^n_A$” as follows. (X “in” Y currently means that the topological space of X is a subspace of the topological space of Y.) Consider A^{n+1} as a free module with generators $t_0, \ldots, t_n$ associated to $x_0, \ldots, x_n$.

Date: November 2, 2005. Minor updates January 31, 2007.

$k[\mathrm{Sym}^* A^{n+1}] = k[t_0, t_1, \dots, t_n] \twoheadrightarrow S_*$ implies $S = k[t_0, t_1, \dots t_n]/I$, where I is a homogeneous ideal. Example: $S_* = k[x, y, z]/(x^2 + y^2 - z^2)$ sits naturally in $\mathbb{P}^2$.

1.2. Easy exercise (silly example). $\mathbb{P}^0_A = \mathrm{Proj}\, A[T] \cong \mathrm{Spec}\, A$. Thus "$\mathrm{Spec}\, A$ is a projective A-scheme".

Here are some useful facts.

A **quasiprojective** A**-scheme** is an open subscheme of a projective A-scheme. The "A" is omitted if it is clear from the context; often A is some field.)

1.3. Exercise. Show that all projective A-schemes are quasicompact. (Translation: show that any projective A-scheme is covered by a finite number of affine open sets.) Show that $\mathrm{Proj}\, S_*$ is finite type over $A = S_0$. If S_0 is a Noetherian ring, show that $\mathrm{Proj}\, S_*$ is a Noetherian scheme, and hence that $\mathrm{Proj}\, S_*$ has a finite number of irreducible components. Show that any quasiprojective scheme is locally of finite type over A. If A is Noetherian, show that any quasiprojective A-scheme is quasicompact, and hence of finite type over A. Show this need not be true if A is not Noetherian.

I'm now going to ask a somewhat rhetorical question. It's going to sound complicated because of all the complicated words in it. But all the complicated words just mean simple concepts.

Question (open for now): are there any quasicompact finite type k-schemes that are not quasiprojective? (Translation: if we're gluing together a finite number of schemes each sitting in some $\mathbb{A}^n$, can we ever get something not quasiprojective?) The difficulty of answering this question shows that this is a good notion! We will see before long that the affine line with the doubled origin is not projective, but we'll call that kind of bad behavior "non-separated", and then the question will still stand: is every separated quasicompact finite type k-scheme quasiprojective?

1.4. Exercise. Show that $\mathbb{P}^n_k$ is normal. More generally, show that $\mathbb{P}^n_R$ is normal if R is a Unique Factorization Domain.

I said earlier that the *affine cone* is $\mathrm{Spec}\, S_*$. (We'll soon see that we'll have a map from cone minus origin to Proj.) The *projective cone* of $\mathrm{Proj}\, S_*$ is $\mathrm{Proj}\, S_*[T]$. We have an intuitive picture of both.

1.5. Exercise (better version of exercise from last day). Show that the projective cone of $\mathrm{Proj}\, S_*$ has an open subscheme $D(T)$ that is the affine cone, and that its complement $V(T)$ can be identified with $\mathrm{Proj}\, S_*$ (as a topological space). More precisely, setting $T = 0$ "cuts out" a scheme isomorphic to $\mathrm{Proj}\, S_*$ — see if you can make that precise.

A lot of what we did for affine schemes generalizes quite easily, as you'll see in these exercises.

1.6. Exercise. Show that the irreducible subsets of dimension $n-1$ of $\mathbb{P}^n_k$ correspond to homogeneous irreducible polynomials modulo multiplication by non-zero scalars.

1.7. Exercise.

(a) Suppose I is any homogeneous ideal, and f is a homogeneous element. Suppose f vanishes on $V(I)$. Show that $f^n \in I$ for some n. (Hint: mimic the proof in the affine case.)
(b) If $Z \subset \operatorname{Proj} S_*$, define $I(\cdot)$. Show that it is a homogeneous ideal. For any two subsets, show that $I(Z_1 \cup Z_2) = I(Z_1) \cap I(Z_2)$.
(c) For any homogeneous ideal I with $V(I) \neq \emptyset$, show that $I(V(I)) = \sqrt{I}$. [They may need the next exercise for this.]
(d) For any subset $Z \subset \operatorname{Proj} S_*$, show that $V(I(Z)) = \overline{Z}$.

1.8. Exercise. Show that the following are equivalent. (a) $V(I) = \emptyset$ (b) for any f_i (i in some index set) generating I, $\cup D(f_i) = \operatorname{Proj} S_*$ (c) $\sqrt{I} \supset S_+$.

Now let's go back to some interesting geometry. Here is a useful construction. Define $S_{n*} := \oplus_i S_{ni}$. (We could rescale our degree, so "old degree" n is "new degree" 1.)

1.9. Exercise. Show that $\operatorname{Proj} S_{n*}$ is isomorphic to $\operatorname{Proj} S_*$.

1.10. Exercise. Suppose S_* is generated over S_0 by $f_1, \ldots, f_n$. Suppose $d = \operatorname{lcm}(\deg f_1, \ldots, \deg f_n)$. Show that S_{d*} is generated in "new" degree 1 (= "old" degree d). (Hint: I like to show this by induction on the size of the set $\{\deg f_1, \ldots, \deg f_n\}$.) This is handy, because we can stick every Proj in some projective space via the construction of 1.1.

1.11. Exercise. If S_* is a Noetherian domain over k, and $\operatorname{Proj} S_*$ is non-empty show that $\dim \operatorname{Spec} S_* = \dim \operatorname{Proj} S_* + 1$. (Hint: throw out the origin. Look at a distinguished $D(f)$ where $\deg f = 1$. Use the fact mentioned in Exercise 2.3 of Class 9. By the previous exercise, you can assume that S_* is generated in degree 1 over $S_0 = A$.)

Example: Suppose $S_* = k[x, y]$, so $\operatorname{Proj} S_* = \mathbb{P}^1_k$. Then $S_{2*} = k[x^2, xy, y^2] \subset k[x, y]$. What is this subring? Answer: let $u = x^2, v = xy, w = y^2$. I claim that $S_{2*} = k[u, v, w]/(uw - v^2)$.

1.12. Exercise. Prove this.

We have a graded ring with three generators; thus we think of it as sitting "in" $\mathbb{P}^2$. This is $\mathbb{P}^1$ as a conic in $\mathbb{P}^2$.

1.13. *Side remark: diagonalizing quadrics.* Suppose k is an algebraically closed field of characteristic not 2. Then any quadratic form in n variables can be "diagonalized" by changing coordinates to be a sum of squares (e.g. $uw - v^2 = ((u+v)/2)^2 + (i(u-v)/2)^2 + (iv)^2$), and the number of such squares is invariant of the change of coordinates. (Reason:

write the quadratic form on $x_1, \dots, x_n$ 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 — here you are using characteristic $\neq 2$. Then diagonalize M — here you are using algebraic closure.) Thus 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 0 (not really 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 squares is isomorphic to $\mathbb{P}^1$.

We now soup up this example.

1.14. Exercise. Show that $\operatorname{Proj} S_{3*}$ is the *twisted cubic* "in" $\mathbb{P}^3$.

1.15. Exercise. Show that $\operatorname{Proj} S_{d*}$ 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" $\mathbb{P}^d$.

More generally, if $S_* = k[x_0, \dots, x_n]$, then $\operatorname{Proj} S_{d*} \subset \mathbb{P}^{N-1}$ (where N is the number of degree d polynomials in $x_0, \dots, x_n$) is called the d*-uple embedding* or d*-uple Veronese embedding*. **Exercise.** Show that $N = \binom{n+d}{d}$.

1.16. Exercise. Find the equations cutting out the *Veronese surface* $\operatorname{Proj} S_{2*}$ where $S_* = k[x_0, x_1, x_2]$, which sits naturally in $\mathbb{P}^5$.

1.17. *Example.* If we put a non-standard weighting on the variables of $k[x_1, \dots, x_n]$ — say we give x_i degree d_i — then $\operatorname{Proj} k[x_1, \dots, x_n]$ is called *weighted projective space* $\mathbb{P}(d_1, d_2, \dots, d_n)$.

1.18. Exercise. Show that $\mathbb{P}(m, n)$ is isomorphic to $\mathbb{P}^1$. Show that $\mathbb{P}(1,1,2) \cong \operatorname{Proj} k[u, v, w, z]/(uw - v^2)$. Hint: do this by looking at the even-graded parts of $k[x_0, x_1, x_2]$, cf. Exercise 1.9. (Picture: this is a projective cone over a conic curve.)

1.19. *Important exercise for later.* (a) *(Hypersurfaces meet everything of dimension at least* 1 *in projective space — unlike in affine space.)* Suppose X is a closed subset of $\mathbb{P}^n_k$ of dimension at least 1, and H a nonempty hypersurface in $\mathbb{P}^n_k$. Show that H meets X. (Hint: consider the affine cone, and note that the cone over H contains the origin. Use Krull's Principal Ideal Theorem.)

(b) (Definition: Subsets in $\mathbb{P}^n$ cut out by linear equations are called *linear subspaces*. Dimension $1, 2$ linear subspaces are called *lines* and *planes* respectively.) Suppose $X \hookrightarrow \mathbb{P}^n_k$ is a closed subset of dimension r. Show that any codimension r linear space meets X. (Hint: Refine your argument in (a).)
(c) Show that there is a codimension $r+1$ complete intersection (codimension $r+1$ set that is the intersection of $r+1$ hypersurfaces) missing X. (The key step: show that there is a hypersurface of sufficiently high degree that doesn't contain every generic point of X.) 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.)

1.20. Exercise. Describe all the lines on the quadric surface $wx - yz = 0$ in $\mathbb{P}^3_k$. (Hint: they come in two "families", called the *rulings* of the quadric surface.)

Hence by Remark 1.13, 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.

2. "SMOOTHNESS" = REGULARITY = NONSINGULARITY

The last property of schemes I want to discuss is something very important: when they are "smooth". For unfortunate historic reasons, *smooth* is a name given to certain morphisms of schemes, but I'll feel free to use this to use it also for schemes themselves. The more correct terms are *regular* and *nonsingular*. A point of a scheme that is not smooth=regular=nonsingular is, not surprisingly, *singular*.

The best way to describe this is by first defining the tangent space to a scheme at a point, what we'll call the *Zariski tangent space*. This will behave like the tangent space you know and love at smooth points, but will also make 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.

This definition is short but surprising. I'll have to convince you that it deserves to be called the tangent space. I've always found this tricky to explain, and that is because we want to show that it agrees with our intuition; but unfortunately, our intuition is crappier than we realize. So I'm just going to define it for you, and later try to convince you that it is reasonable.

Suppose A is a ring, and $\mathfrak{m}$ is a point. Translation: we have a point $[\mathfrak{m}]$ on a scheme $\operatorname{Spec} A$. Let $k = A/\mathfrak{m}$ be the residue field. Then $\mathfrak{m}/\mathfrak{m}^2$ is a vector space over the residue field $A/\mathfrak{m}$: it is an A-module, and $\mathfrak{m}$ acts like 0. This is defined to be the **Zariski cotangent space**. The dual is the **Zariski tangent space**. Elements of the Zariski cotangent space are called **cotangent vectors** or **differentials**; elements of the tangent space are called **tangent vectors**.

Note: This is intrinsic; it doesn't depend on any specific description of the ring itself (e.g. the choice of generators over a field k = choice of embedding in affine space). An interesting feature: in some sense, the cotangent space is more algebraically natural than the tangent space. 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.

I'm now going to give you a bunch of plausibility arguments that this is a reasonable definition.

First, I'll make a moral argument that this definition is plausible for the cotangent space of the origin of $\mathbb{A}^n$. Functions on $\mathbb{A}^n$ should restrict to a linear function on the tangent space. What 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.

Here is a second argument, for those of you who think of the tangent space as the space of derivations. (I didn't say this in class, because I didn't realize that many of you thought in this way until later.) More precisely, in differential geometry, the tangent space at a point is sometimes defined as the vector space of derivations at that point. A derivation is a function that takes in functions near the point that vanish at the point, and gives elements of the field k, and satisfies the Leibniz rule $(fg)' = f'g + g'f$. Translation: a derivation is a map $\mathfrak{m} \to k$. But $\mathfrak{m}^2 \to 0$, as if $f(p) = g(p) = 0$, then $(fg)'(p) = f'(p)g(p) + g'(p)f(p) = 0$. Thus we have a map $\mathfrak{m}/\mathfrak{m}^2 \to k$, i.e. an element of $(\mathfrak{m}/\mathfrak{m}^2)^*$. **Exercise (for those who think in this way).** Check that this is reversible, i.e. that any map $\mathfrak{m}/\mathfrak{m}^2 \to k$ gives a derivation — i.e., check the Leibniz rule.

2.1. Here is an old-fashioned example to help tie this down to earth. This is not currently intended to be precise. In $\mathbb{A}^3$, we have a curve cut out by $x + y + z^2 + xyz = 0$ and $x - 2y + z + x^2y^2z^3 = 0$. 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 analytic sense). I give questions like this in multivariable calculus. The students do a page of calculus to get the answer, because I can't tell them to just pick out the linear terms.

Another example: $x + y + z^2 = 0$ 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 a page of calculus which would boil down to picking off the linear terms. They would end up with the equations $x + y = 0$ and $x + y = 0$, which cuts 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 cotangent space is cut out by $x + y = 0$, and *define* this as a singular point. (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.

2.2. *Proposition. — Suppose* $(A, \mathfrak{m})$ *is a Noetherian local ring. Then* $\dim A \leq \dim_k \mathfrak{m}/\mathfrak{m}^2$.

We'll prove this on Friday.

If equality holds, we say that A is **regular** at $\mathfrak{p}$. If A is a local ring, then we say that A is a **regular local ring**. If A is regular at all of its primes, we say that A is a **regular ring**.

A scheme X is **regular** or **nonsingular** or **smooth** at a point p if the local ring $\mathcal{O}_{X,p}$ is regular. It is **singular** at the point otherwise. A scheme is **regular** or **nonsingular** or **smooth** if it is regular at all points. It is **singular** otherwise (i.e. if it is singular at *at least one* point.

In order to prove Proposition 2.2, we're going to pull out another algebraic weapon: Nakayama's lemma. This was done in Math 210, so I didn't discuss it in class. You should read this short exposition. If you have never seen Nakayama before, you'll see a complete proof. If you want a refresher, here it is. And even if you are a Nakayama expert, please take a look, because there are several related facts that go by the name of Nakayama's Lemma, and we should make sure we're talking about the same one(s). Also, this will remind you that the proof wasn't frightening and didn't require months of previous results.

2.3. *Nakayama's Lemma version 1. — Suppose R is a ring, I an ideal of R, and M is a finitely-generated R-module. Suppose $M = IM$. Then there exists an $a \in R$ with $a \equiv 1 \pmod I$ with $aM = 0$.*

Proof. Say M is generated by $m_1, \ldots, m_n$. Then as $M = IM$, we have $m_i = \sum_j a_{ij} m_j$ for some $a_{ij} \in I$. Thus

$$(\mathrm{Id}_n - A) \begin{pmatrix} m_1 \\ \vdots \\ m_n \end{pmatrix} = 0 \qquad (1)$$

where Id_n is the $n \times n$ identity matrix in R, and $A = (a_{ij})$. We can't quite invert this matrix, but we almost can. Recall that any $n \times n$ matrix M has an adjoint $\mathrm{adj}(M)$ such that $\mathrm{adj}(M)M = \det(M)\mathrm{Id}_n$. The coefficients of $\mathrm{adj}(M)$ are polynomials in the coefficients of M. (You've likely seen this in the form a formula for M^{-1} when there is an inverse.) Multiplying both sides of (1) on the left by $\mathrm{adj}(\mathrm{Id}_n - A)$, we obtain

$$\det(\mathrm{Id}_n - A) \begin{pmatrix} m_1 \\ \vdots \\ m_n \end{pmatrix} = 0.$$

But when you expand out $\det(\mathrm{Id}_n - A)$, you get something that is $1 \pmod I$. □

Here is why you care: Suppose I is contained in all maximal ideals of R. (The intersection of all the maximal ideals is called the *Jacobson radical*, but I won't use this phrase. Recall that the nilradical was the intersection of the *prime ideals* of R.) Then I claim that any $a \equiv 1 \pmod I$ is invertible. For otherwise $(a) \neq R$, so the ideal (a) is contained in some maximal ideal $\mathfrak{m}$ — but $a \equiv 1 \pmod{\mathfrak{m}}$, contradiction. Then as a is invertible, we have the following.

2.4. *Nakayama's Lemma version 2. — Suppose* R *is a ring,* I *an ideal of* R *contained in all maximal ideals, and* M *is a finitely-generated* R*-module. (Most interesting case:* R *is a local ring, and* I *is the maximal ideal.) Suppose* $M = IM$*. Then* $M = 0$*.*

2.5. Important exercise (Nakayama's lemma version 3). Suppose R is a ring, and I is an ideal of R contained in all maximal ideals. Suppose M is a *finitely generated* R-module, and $N \subset M$ is a submodule. If $N/IN \hookrightarrow M/IM$ an isomorphism, then $M = N$.

2.6. Important exercise (Nakayama's lemma version 4). Suppose $(R, \mathfrak{m})$ is a local ring. Suppose M is a finitely-generated R-module, and $f_1, \dots, f_n \in M$, with (the images of) $f_1, \dots, f_n$ generating $M/\mathfrak{m}M$. Then $f_1, \dots, f_n$ generate M. (In particular, taking $M = \mathfrak{m}$, if we have generators of $\mathfrak{m}/\mathfrak{m}^2$, they also generate $\mathfrak{m}$.)

E-mail address: vakil@math.stanford.edu