FOUNDATIONS OF ALGEBRAIC GEOMETRY CLASS 16

CONTENTS

Last day: much more on quasicoherence. Quasicoherent sheaves form an abelian category. Finite type (= "locally finitely generated") and coherent sheaves. Support of a sheaf. Rank of a finite type sheaf at a point. Closed subschemes at a point.

Today: effective Cartier divisors; quasicoherent sheaves on projective A-schemes corresponding to graded modules, line bundles on projective A-schemes, O(n), generated by global sections, Serre's theorem, the adjoint functors $\sim$ and Γ_*.

1. YET MORE ON CLOSED SUBSCHEMES

Here are few more notions about closed subschemes.

In analogy with closed subsets of a topological space, we can define finite unions and arbitrary intersections of closed subschemes. On affine open set $\operatorname{Spec} R$, if for each i in an index set, I_i corresponds to a closed subscheme, the scheme-theoretic intersection of the closed subschemes corresponds to the ideal generated by the I_i (here the index set may be infinite), and the scheme-theoretic union corresponds to the intersection of by all I_i (here the index should be finite).

Exercise: Describe the scheme-theoretic intersection of $(y - x^2)$ and y in $\mathbb{A}^2$. Describe the scheme-theoretic union. Draw a picture.

Exercise: Prove some fact of your choosing showing that closed subschemes behave similarly to closed subsets. For example, if X, Y, and Z are closed subschemes of W, show that $(X \cap Z) \cup (Y \cap Z) = (X \cup Y) \cap Z$.

Date: Wednesday, November 16, 2005. Minor update January 31, 2007.

1.1. From closed subschemes to effective Cartier divisors. There is a special name for a closed subscheme locally cut out by one equation that is not a zero-divisor. More precisely, it is a closed subscheme such that there exists an affine cover such that on each one it is cut out by a single equation, not a zero-divisor. (This does not mean that on *any* affine it is cut out by a single equation — this notion doesn't satisfy the "gluability" hypothesis of the Affine Communication Theorem. If $I \subset R$ is generated by a non-zero divisor, then $I_f \subset R_f$ is too. But "not conversely". I might give an example later.) We call this an *effective Cartier divisor*. (This admittedly unwieldy terminology! But there is a reason for it.) By Krull, it is pure codimension 1.

Remark: if $I = (u) = (v)$, and u is not a zero-divisor, then u and v differ by a unit in R. Proof: $u \in (v)$ implies $u = av$. Similarly $v = bu$. Thus $u = abu$, from which $u(1-ab) = 0$. As u is not a zero-divisor, $1 = ab$, so a and b are units.

Reason we care: effective Cartier divisors give invertible sheaves. If $\mathcal{I}$ is an effective Cartier divisor on X, then $\mathcal{I}$ is an invertible sheaf. Reason: locally, sections are multiples of a single generator u, and there are no "relations".

The invertible sheaf corresponding to an effective Cartier divisor is for various reasons defined to be the dual of the ideal sheaf. This line bundle has a canonical section: Dualizing $0 \to \mathcal{I} \to \mathcal{O}$ gives us $\mathcal{O} \to \mathcal{I}^*$. **Exercise.** This section vanishes along our actual effective Cartier divisor.

1.2. Exercise. Describe the invertible sheaf in terms of transition functions. More precisely, on any affine open set where the effective Cartier divisor is cut out by a single equation, the invertible sheaf is trivial. Determine the transition functions between two such overlapping affine open sets. Verify that there is indeed a canonical section of this invertible sheaf, by describing it.

To describe the tensor product of such invertible sheaves: if $I = (u)$ (locally) and $J = (v)$, then the tensor product corresponds to (uv).

We get a monoid of effective Cartier divisors, with unit $\mathcal{I} = \mathcal{O}$. Notation: D is an effective Cartier divisor. $\mathcal{O}(D)$ is the corresponding line bundle. $\mathcal{O}(-D)$ is the ideal sheaf.

$$0 \to \mathcal{O}(-D) \to \mathcal{O} \to \mathcal{O}_D \to 0.$$

D is associated to the closed subscheme.

Hence we can get a bunch of invertible sheaves, by taking differences of these two. Surprising fact: we almost get them all! (True for all nonsingular schemes. This is true for all projective schemes. It is very hard to describe an invertible sheaf that is not describable in such a way.)

2. Quasicoherent sheaves on projective A-schemes

We now describe quasicoherent sheaves on a projective A-scheme. Recall that a projective A-scheme is produced from the data of $\mathbb{Z}^{\geq 0}$-graded ring S_*, with $S_0 = A$, and S_+ finitely generated as an A-module. The resulting scheme is denoted $\operatorname{Proj} S_*$.

Let $X = \operatorname{Proj} S_*$. Suppose M_* is a graded S_* module, *graded by* $\mathbb{Z}$. We define the quasicoherent sheaf $\widetilde{M_*}$ as follows. (I will avoid calling it $\tilde{M}$, as this might cause confusion with the affine case.) On the distinguished open $D(f)$, we let

$$\widetilde{M_*}(D(f)) \cong \widetilde{(M_f)_0}.$$

(More correctly: we define a sheaf $\widetilde{M_*(f)}$ on $D(f)$ for each f. We give identifications of the restriction of $\widetilde{M_*(f)}$ and $\widetilde{M_*(g)}$ to $D(fg)$. Then by an earlier problem set problem telling how to glue sheaves, these sheaves glue together to a single sheaf on $\widetilde{M_*}$ on X. We then discard the temporary notation $\widetilde{M_*(f)}$.)

This is clearly quasicoherent, because it is quasicoherent on each $D(f)$. If M_* is finitely generated over S_*, then so $\tilde{M}_*$ is a finite type sheaf.

I will now give some straightforward facts.

If M_* and M'_* agree in high enough degrees, then $\widetilde{M_*} \cong \widetilde{M'_*}$. Thus the map from graded S_*-modules to quasicoherent sheaves on $\operatorname{Proj} S_*$ is not a bijection.

Given a map of graded modules $\phi : M_* \to N_*$, we we get an induced map of sheaves $\widetilde{M_*} \to \widetilde{N_*}$. Explicitly, over $D(f)$, the map $M_* \to N_*$ induces $M_*[1/f] \to N_*[1/f]$ which induces $\phi_f : (M_*[1/f])_0 \to (N_*[1/f])_0$; and this behaves well with respect to restriction to smaller distinguished open sets, i.e. the following diagram commutes.

$$\begin{tikzcd}
(M_*[1/f])_0 \arrow[r, "\phi_f"] \arrow[d] & (N_*[1/f])_0 \arrow[d] \\
(M_*[1/(fg)])_0 \arrow[r, "\phi_{fg}"] & (N_*[1/(fg)])_0.
\end{tikzcd}$$

In fact $\sim$ is a functor from the category of graded S_*-modules to the category of quasicoherent sheaves on $\operatorname{Proj} S_*$. This isn't quite an isomorphism, but it is close. The relationship is akin to that between presheaves and sheaves, and the sheafification functor, as we will see before long.

2.1. Exercise. Show that $\widetilde{M_*} \otimes \widetilde{N_*} \cong \widetilde{M_* \otimes_{S_*} N_*}$. (Hint: describe the isomorphism of sections over each $D(f)$, and show that this isomorphism behaves well with respect to smaller distinguished opens.)

If $I_* \subset S_*$ is a graded ideal, we get a closed subscheme. Example: Suppose $S_* = k[w, x, y, z]$, so $\operatorname{Proj} S_* \cong \mathbb{P}^3$. Suppose $I_* = (wx - yz, x^2 - wy, y^2 - xz)$. Then we get the exact sequence of graded S_*-modules

$$0 \to I_* \to S_* \to S_*/I_* \to 0.$$

Which closed subscheme of $\mathbb{P}^3$ do we get? The twisted cubic!

2.2. Exercise. Show that if I_* is a graded ideal of S_*, show that we get a closed immersion $\operatorname{Proj} S_*/I_* \hookrightarrow \operatorname{Proj} S_*$.

3. Invertible sheaves (line bundles) on projective A-schemes

We now come to one of the most fundamental concepts in projective geometry.

First, I want to mention something that I should have mentioned some time ago.

3.1. 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 Exercise 2.2.

Suppose that S_* is generated in degree 1. By the previous exercise, this is not a huge assumption. Suppose M_* is a graded S_*-module. Define the graded module $M(n)_*$ so that $M(n)_m := M_{n+m}$. Thus the quasicoherent sheaf $\widetilde{M(n)_*}$ is given by

$$\Gamma(D(f), \widetilde{M(n)_*}) = \widetilde{(M_f)_n}$$

where here the subscript means we take the nth graded piece. (These subscripts are admittedly confusing!)

As an incredibly important special case, define $\mathcal{O}_{\operatorname{Proj} S_*}(n) := \widetilde{S(n)_*}$. When the space is implicit, it can be omitted from the notation: $\mathcal{O}(n)$ (pronounced "oh of n").

3.2. Important exercise. If S_* is generated in degree 1, show that $\mathcal{O}_{\operatorname{Proj} S_*}(n)$ is an invertible sheaf.

3.3. Essential exercise. Calculate $\dim_k \Gamma(\mathbb{P}^m_k, \mathcal{O}_{\mathbb{P}^m_k}(n))$.

I'll get you started on this. As always, consider the "usual" affine cover. Consider the $n = 1$ case. Say $S_* = k[x_0, \ldots, x_m]$. Suppose we have a global section of $\mathcal{O}(1)$. On $D(x_0)$, the sections are of the form $f(x_0, \ldots, x_n)/x_0^{\deg f - 1}$. On $D(x_1)$, the sections are of the form $g(x_0, \ldots, x_n)/x_1^{\deg g - 1}$. They are supposed to agree on the overlap, so

$$x_0^{\deg f - 1} g(x_0, \ldots, x_n) = x_1^{\deg g - 1} f(x_0, \ldots, x_n).$$

How is this possible? Well, we must have that $g = x_1^{\deg g - 1} \times$ some linear factor, and $f = x_0^{\deg f - 1} \times$ the same linear factor. Thus on $D(x_0)$, this section must be some linear form. On $D(x_1)$, this section must be the same linear form. By the same argument, on each

$D(x_i)$, the section must be the same linear form. Hence (with some argumentation), the global sections of $\mathcal{O}(1)$ correspond to the linear forms in $x_0, \ldots, x_m$, of which there are $m+1$.

Thus $x+y+2z$ is a section of $\mathcal{O}(1)$ on $\mathbb{P}^2$. It isn't a function, but I can say when this section vanishes — precisely where $x+y+2z=0$.

3.4. Exercise. Show that $\mathcal{F}(n) \cong \mathcal{F} \otimes \mathcal{O}(n)$.

3.5. Exercise. Show that $\mathcal{O}(m+n) \cong \mathcal{O}(m) \otimes \mathcal{O}(n)$.

3.6. Exercise. Show that if $m \neq n$, then $\mathcal{O}_{\mathbb{P}^l_k}(m)$ is not isomorphic to $\mathcal{O}_{\mathbb{P}^l_k}(n)$ if $l > 0$. (Hence we have described a countable number of invertible sheaves (line bundles) that are non-isomorphic. We will see later that these are *all* the line bundles on projective space $\mathbb{P}^n_k$.)

4. Generation by global sections, and Serre's Theorem

(I discussed this in class 15, but should have discussed them here. Hence they are in the class 16 notes.)

Suppose $\mathcal{F}$ is a sheaf on X. We say that $\mathcal{F}$ is *generated by global sections at a point* p if we can find $\phi : \mathcal{O}^{\oplus I} \to \mathcal{F}$ that is surjective at the stalk of p: $\phi_p : \mathcal{O}_p^{\oplus I} \to \mathcal{F}_p$ is surjective. (Some what more precisely, the stalk of $\mathcal{F}$ at p is generated by global sections of $\mathcal{F}$. The global sections in question are the images of the $|I|$ factors of $\mathcal{O}_p^{\oplus I}$.) We say that $\mathcal{F}$ is *generated by global sections* if it is generated by global sections at all p, or equivalently, if we can we can find $\mathcal{O}^{\oplus I} \to \mathcal{F}$ that is surjective. (By our earlier result that we can check surjectivity at stalks, this is the same as saying that it is surjective at all stalks.) If I can be taken to be finite, we say that $\mathcal{F}$ is generated by a finite number of global sections. We'll see soon why we care.

4.1. *Exercise.* If quasicoherent sheaves $\mathcal{L}$ and $\mathcal{M}$ are generated by global sections at a point p, then so is $\mathcal{L} \otimes \mathcal{M}$. (This exercise is less important, but is good practice for the concept.)

4.2. *Easy exercise.* An invertible sheaf $\mathcal{L}$ on X is generated by global sections if and only if for any point $x \in X$, there is a section of $\mathcal{L}$ not vanishing at x. (Hint: Nakayama.)

4.3. *Lemma. — Suppose $\mathcal{F}$ is a finite type sheaf on X. Then the set of points where $\mathcal{F}$ is generated by global sections is an open set.*

Proof. Suppose $\mathcal{F}$ is generated by global sections at a point p. Then it is generated by a finite number of global sections, say m. This gives a morphism $\phi : \mathcal{O}^{\oplus m} \to \mathcal{F}$, hence

$\operatorname{im}\phi \hookrightarrow \mathcal{F}$. The support of the (finite type) cokernel sheaf is a closed subset not containing p. $\square$

(Back to class 16!)

4.4. Important Exercise (an important theorem of Serre). Suppose S_0 is a Noetherian ring, and S_* is generated in degree 1. Let $\mathcal{F}$ be any finite type sheaf on $\operatorname{Proj} S_*$. Then for some integer n_0, for all $n \geq n_0$, $\mathcal{F}(n)$ can be generated by a finite number of global sections.

I'm going to sketch how you should tackle this exercise, after first telling you the reason we will care.

4.5. *Corollary.* — Thus any coherent sheaf $\mathcal{F}$ on $\operatorname{Proj} S_*$ can be presented as:

$$\oplus_{\text{finite}} \mathcal{O}(-n) \to \mathcal{F} \to 0.$$

We're going to use this a lot!

Proof. Suppose we have m global sections $s_1, \ldots, s_m$ of $\mathcal{F}(n)$ that generate $\mathcal{F}(n)$. This gives a map

$$\oplus_m \mathcal{O} \longrightarrow \mathcal{F}(n)$$

given by $(f_1, \ldots, f_m) \mapsto f_1 s_1 + \cdots + f_m s_m$ on any open set. Because these global sections generate $\mathcal{F}$, this is a surjection. Tensoring with $\mathcal{O}(-n)$ (which is exact, as tensoring with any locally free is exact) gives the desired result. $\square$

Here is now a hint/sketch for the Serre exercise 4.4.

We can assume that S_* is generated in degree 1; we can do this thanks to Exercise 3.1. Suppose $\deg f = 1$. Say $\mathcal{F}|_{D(f)} = \tilde{M}$, where M is a $(S_*[1/f])_0$-module, generated by $m_1, \ldots, m_n$. These elements generate all the stalks over all the points of $D(f)$. They are sections over this big distinguished open set. It would be wonderful if we knew that they had to be restrictions of global sections, i.e. that there was a global section m_i' that restricted to m_i on $D(f)$. If that were always true, then we would cover X with a finite number of each of these $D(f)$'s, and for each of them, we would take the finite number of generators of the corresponding module. Sadly this is not true.

However, we will see that $f^N m$ "extends", where m is any of the m_i's, and N is sufficiently large. We will see this by (easily) checking first that $f^N m$ extends over another distinguished open $D(g)$ (i.e. that there is a section of $\mathcal{F}(N)$ over $D(g)$ that restricts to $f^n m$ on $D(g) \cap D(f) = D(fg)$). But we're still not done, because we don't know that the extension over $D(g)$ and over some other $D(g')$ agree on the overlap $D(g) \cap D(g') = D(gg')$ — in fact, they need not agree! But after multiplying both extensions by $f^{N'}$ for large enough N', we will see that they agree on the overlap. By quasicompactness, we need to to extend over only a finite number of $D(g)$'s, and to make sure extensions agree over the finite number of pairs of such $D(g)$'s, so we will be done.

Great, let's make this work. Let's investigate this on $D(g) = \operatorname{Spec} A$, where the degree of g is also 1. Say $\mathcal{F}|D(g) \cong \tilde{N}$. Let $f' = f/g$ be "the function corresponding to f on $D(g)$". We have a section over $D(f')$ on the affine scheme $D(g)$, i.e. an element of $N_{f'}$, i.e. something of the form $n/(f')^N$ for some $n \in N$. So then if we multiply it by f'^N, we can certainly extend it! So if we multiply by a big enough power of f, m certainly extends over any $D(g)$.

As described earlier, the only problem is, we can't guarantee that the extensions over $D(g)$ and $D(g')$ agree on the overlap (and hence glue to a single extensions). Let's check on the intersection $D(g) \cap D(g') = D(gg')$. Say $m = n/(f')^N = n'/(f')^{N'}$ where we can take $N = N'$ (by increasing N or N' if necessary). We certainly may not have $n = n'$, but by the (concrete) definition of localization, after multiplying with enough f''s, they become the same.

In conclusion after multiplying with enough f's, our sections over $D(f)$ extend over each $D(g)$. After multiplying by even more, they will all agree on the overlaps of any two such distinguished affine. Exercise 4.4 is to make this precise.

5. Every quasicoherent sheaf on a projective A-scheme arises from a graded module

We have gotten lots of quasicoherent sheaves on $\operatorname{Proj} S_*$ from graded S_*-modules. We'll now see that we can get them all in this way.

We want to figure out how to "undo" the $\tilde{M}$ construction. When you do the essential exercise 3.3, you'll suspect that in good situations,

$$M_n \cong \Gamma(\operatorname{Proj} S_*, \tilde{M}(n)).$$

Motivated by this, we define

$$\Gamma_n(\mathcal{F}) := \Gamma(\operatorname{Proj} S_*, \mathcal{F}_n).$$

Then $\Gamma_*(\mathcal{F})$ is a graded S_*-module, and we can dream that $\Gamma_*(\mathcal{F})^{\sim} \cong \mathcal{F}$. We will see that this is indeed the case!

5.1. Exercise. Show that Γ_* gives a functor from the category of quasicoherent sheaves on $\operatorname{Proj} S_*$ to the category of graded S_*-modules. (In other words, show that if $\mathcal{F} \to \mathcal{G}$ is a morphism of quasicoherent sheaves on $\operatorname{Proj} S_*$, describe the natural map $\Gamma_*(\mathcal{F}) \to \Gamma_*(\mathcal{G})$, and show that such natural maps respect the identity and composition.)

Note that $\sim$ and Γ_* cannot be inverses, as $\sim$ can turn two different graded modules into the same quasicoherent sheaf.

Our initial goal is to show that there is a natural isomorphism $\widetilde{\Gamma_*(\mathcal{F})} \to \mathcal{F}$, and that there is a natural map $M_* \to \Gamma_*(\widetilde{M_*})$. We will show something better: that $\sim$ and Γ_* are adjoint.

We start by describing the natural map $M_* \to \Gamma_* \widetilde{M_*}$. We describe it in degree n. Given an element m_n, we seek an element of $\Gamma(\operatorname{Proj} S_*, \widetilde{M_*}(n)) = \Gamma(\operatorname{Proj} S_*, \widetilde{M_{(n+*)}})$. By shifting the grading of M_* by n, we can assume $n = 0$. For each $D(f)$, we certainly have an element of $(M[1/f])_0$ (namely m), and they agree on overlaps, so the map is clear.

5.2. Exercise. Show that this canonical map need not be injective, nor need it be surjective. (Hint: $S_* = k[x]$, $M_* = k[x]/x^2$ or $M_* = \{$ polynomials with no constant terms $\}$.)

The natural map $\widetilde{\Gamma_* \mathcal{F}} \to \mathcal{F}$ is more subtle (although it will have the advantage of being an isomorphism).

5.3. Exercise. Describe the natural map $\widetilde{\Gamma_* \mathcal{F}} \to \mathcal{F}$ as follows. First describe it 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 $f^N(f^{n'} m - f^n m') = 0$. Show that your map behaves well on overlaps $D(f) \cap D(g) = D(fg)$.

5.4. Longer Exercise. Show that the natural map $\widetilde{\Gamma_* \mathcal{F}} \to \mathcal{F}$ is an isomorphism, by showing that it is an isomorphism over $D(f)$ for any f. Do this by first showing that it is surjective. This will require following some of the steps of the proof of Serre's theorem (Exercise 4.4). Then show that it is injective.

5.5. *Corollary. — Every quasicoherent sheaf arises from this tilde construction. Each closed subscheme of* $\operatorname{Proj} S_*$ *arises from a graded ideal* $I_* \subset S_*$.

In particular, let $x_0, \dots, x_n$ be generators of S_1 as an A-module. Then we have a surjection of graded rings

$$A[t_0, \dots, t_n] \to S_*$$

where $t_i \mapsto x_i$. Then this describes $\operatorname{Proj} S_*$ as a closed subscheme of $\mathbb{P}^n_A$.

5.6. Exercise (Γ_* and $\sim$ are adjoint functors). Prove part of the statement that Γ_* and $\sim$ are adjoint functors, by describing a natural bijection $\operatorname{Hom}(M_*, \Gamma_*(\mathcal{F})) \cong \operatorname{Hom}(\widetilde{M_*}, \mathcal{F})$. For the map from left to right, start with a morphism $M_* \to \Gamma_*(\mathcal{F})$. Apply $\sim$, and postcompose with the isomorphism $\widetilde{\Gamma_* \mathcal{F}} \to \mathcal{F}$, to obtain

$$\widetilde{M_*} \to \widetilde{\Gamma_* \mathcal{F}} \to \mathcal{F}.$$

Do something similar to get from right to left. Show that "both compositions are the identity in the appropriate category". (Is there a clever way to do that?)

5.7. Saturated S_*-modules. We end with a remark: different graded S_*-modules give the same quasicoherent sheaf on $\operatorname{Proj} S_*$, but the results of this section show that there is a "best" graded module for each quasicoherent sheaf, and there is a map form each graded module to its "best" version, $M_* \to \Gamma_*(\widetilde{M_*})$. A module for which this is an isomorphism (a "best" module) is called *saturated*. I don't think we'll use this notation, but other people do.

This "saturation" map $M_* \to \Gamma_*(\widetilde{M_*})$ should be seen as analogous to the sheafification map, taking presheaves to sheaves.

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