Math 203B: Algebraic Geometry
UCSD, winter 2016, Kiran S. Kedlaya
Affine schemes

In this lecture, we start with a ring R and construct from it an *affine scheme* in a manner analogous to the construction of an affine variety. Basic commutative algebra will be assumed; see Atiyah-Macdonald if you need a refresher.

See also: Hartshorne II.1, II.2.

1 The Zariski prime spectrum

We start with the underlying topological space. Let $\mathrm{Spec}(R)$ be the set of prime ideals of R, i.e., the set of ideals $\mathfrak{p} \subseteq R$ such that $R/\mathfrak{p}$ is an integral domain. By convention, the zero ring is not an integral domain, so the unit ideal is not prime.

For each ideal I of R, define the set

$$V(I) = \{\mathfrak{p} \in \mathrm{Spec}(R) : I \subseteq \mathfrak{p}\}.$$

The sets $V(I)$ satisfy the usual rules for closed sets in a topological space:

- The empty set is $V(R)$, while the set $\mathrm{Spec}(R)$ is $V(0)$.
- Given any number of sets $V(I_j)$, their intersection is $V(I)$ where I is the ideal of R generated by $\cup_j I_j$.
- Given two sets $V(I_1), V(I_2)$, their union is $V(I_1 \cap I_2)$. (It is obvious that $V(I_1) \cup V(I_2) \subseteq V(I_1 \cap I_2)$. Conversely, if $\mathfrak{p}$ is an ideal not containing I_1 or I_2, we can choose $f_1 \in I_1, f_2 \in I_2$ not in $\mathfrak{p}$, and then $f_1 f_2 \in I_1 \cap I_2$ is not in $\mathfrak{p}$ either.)

The resulting topology is called the *Zariski topology* on $\mathrm{Spec}(R)$.

For any ring homomorphism $f : R \to S$ and any $\mathfrak{p} \in \mathrm{Spec}(S)$, the induced map $R/f^{-1}(\mathfrak{p}) \to S/\mathfrak{p}$ is injective, and moreover $1 \notin f^{-1}(\mathfrak{p})$ because $f(1) = 1 \notin \mathfrak{p}$. So $R/f^{-1}(\mathfrak{p})$ is an integral domain, i.e., $f^{-1}(\mathfrak{p})$ is prime, and we have an induced map $f^* : \mathrm{Spec}(S) \to \mathrm{Spec}(R)$. Since $f^{-1}(V(I)) = V(f^{-1}(I))$, f^* is continuous.

2 Distinguished open subsets

For $f \in R$, define the *distinguished open subset*

$$D(f) = \mathrm{Spec}(R) - V((f))$$

to be the set of all prime ideals not containing f. Every open set is a union of these:

$$\mathrm{Spec}(R) - V(I) = \bigcap_{f \in I} D(f).$$

Since $D(f) \cap D(g) = D(fg)$, such open subsets form a basis of the Zariski topology.

What makes this kind of open subset special? Let R_f be the localization of R at the multiplicative set generated by f, which can also be described as $R[f^{-1}] = R[T]/(Tf - 1)$; then the map $R \to R_f$ induces a map $\operatorname{Spec}(R_f) \to \operatorname{Spec}(R)$ which in turn induces a homeomorphism $\operatorname{Spec}(R_f) \cong D(f)$. That is, $D(f)$ is itself naturally thought of as the prime spectrum of R_f; this perspective will be crucial in leading us to the construction of the structure sheaf.

3 A look ahead: properties of the structure sheaf

I like to call this the *first fundamental theorem of schemes.*

Theorem 1 *There exists a unique (up to unique isomorphism) sheaf of rings $\mathcal{O}$ on $\operatorname{Spec}(R)$ such that there exist isomorphisms*

$$\mathcal{O}(D(f)) \cong R_f \qquad (f \in R)$$

compatible with restriction: for $f, g \in R$, the diagram

$$\begin{array}{ccc} \mathcal{O}(D(f)) & \longrightarrow & \mathcal{O}(D(fg)) \\ \downarrow & & \downarrow \\ R_f & \longrightarrow & R_g \end{array}$$

commutes.

This description is not itself a definition of a sheaf, as that requires specifying the value of $\mathcal{O}(U)$ for *every* open set U, not just the distinguished ones. What we are using here to get away with only referring to distinguished opens is the local nature of sheaves.

4 A key corollary

The following statement is an immediate corollary of the theorem, but we will actually prove it first and then use it to deduce the theorem.

Corollary 2 *Let $f_1, \ldots, f_n \in R$ be a finite sequence of elements which generate the unit ideal (equivalently, $\operatorname{Spec}(R) = D(f_1) \cup \cdots \cup D(f_n)$). Then*

$$R \cong \ker \left(\prod_{i=1}^{n} R_{f_i} \to \prod_{i,j=1}^{n} R_{f_i f_j}, \quad (s_i)_i \mapsto (s_i - s_j)_{i,j} \right)$$

via the diagonal map.

Since the purported isomorphism is a morphism of R-modules, we may check that it is indeed an isomorphism by doing so locally; that is, it is enough to check that for each prime ideal $\mathfrak{p}$,

$$R_{\mathfrak{p}} \cong \ker\left(\prod_{i=1}^n R_{f_i} \to \prod_{i,j=1}^n R_{f_i f_j}, \quad (s_i)_i \mapsto (s_i - s_j)_{i,j}\right)_{\mathfrak{p}}.$$

(Reminder: for M an R-module and S a multiplicative subset of R, M_S is the set of formal quotients m/s with $s \in S$ modulo the relations $m/s = (ms')/(ss')$.) Since localization preserves kernels, I can move that localization inside the parentheses, to rewrite the claim as

$$R_{\mathfrak{p}} \cong \ker\left(\prod_{i=1}^n (R_{f_i})_{\mathfrak{p}} \to \prod_{i,j=1}^n (R_{f_i f_j})_{\mathfrak{p}}, \quad (s_i)_i \mapsto (s_i - s_j)_{i,j}\right).$$

Here $(R_{f_i})_{\mathfrak{p}}$ can be reinterpreted as $(R_{\mathfrak{p}})_{f_i}$, because both of them equal R_S where S is the multiplicative subset of R generated by f_i and the complement of $\mathfrak{p}$. The claim thus becomes

$$R_{\mathfrak{p}} \cong \ker\left(\prod_{i=1}^n (R_{\mathfrak{p}})_{f_i} \to \prod_{i,j=1}^n (R_{\mathfrak{p}})_{f_i f_j}, \quad (s_i)_i \mapsto (s_i - s_j)_{i,j}\right).$$

Here now is the key point: there exists at least one index i for which $\mathfrak{p} \in D(f_i)$, and for any such index we have

$$(R_{\mathfrak{p}})_{f_i} \cong R_{\mathfrak{p}}, \qquad (R_{\mathfrak{p}})_{f_i f_j} \cong (R_{\mathfrak{p}})_{f_j}$$

because f_i becomes a unit in $R_{\mathfrak{p}}$. Thus we may project $\prod_{i=1}^n (R_{\mathfrak{p}})_{f_i}$ onto one factor to obtain a map from the right side to the left, and this is easily seen to be inverse to the map the other way.

5 Localization and stalks

Before continuing, let us clarify the relationship between algebraic localization of a ring at a prime ideal and the formation of stalks of a sheaf. For this purpose, let us temporarily assume the properties of the structure sheaf. Then the stalk $\mathcal{O}_{\mathfrak{p}}$ is by definition the direct limit of $\mathcal{O}(U)$ as U runs through all open subsets of $\operatorname{Spec}(R)$ containing $\mathfrak{p}$. Since the distinguished open sets form a neighborhood basis, this is the same as taking the direct limit of R_f as f runs through all elements of R not contained in $\mathfrak{p}$. This is patently equal to the localization of R at the whole multiplicative set $R - \mathfrak{p}$, which by definition is $R_{\mathfrak{p}}$.

Thanks to one of the homework problems, this identification implies that the diagonal map

$$R \to \prod_{\mathfrak{p} \in \operatorname{Spec}(R)} R_{\mathfrak{p}}$$

is injective; that is, we can view elements of R as functions from $\operatorname{Spec}(R)$ to the disjoint union $\sqcup_{\mathfrak{p} \in \operatorname{Spec}(R)} R_{\mathfrak{p}}$ without losing any information.

6 Construction of the structure sheaf

We now turn this picture around and use it to establish the existence and uniqueness of the sheaf $\mathcal{O}$. The uniqueness is already clear now: up to unique isomorphism, we must take $\mathcal{O}(U)$ to be the set of functions $s : U \to \sqcup_{\mathfrak{p}\in\mathrm{Spec}(R)} R_{\mathfrak{p}}$ such that $s(\mathfrak{p}) \in R_{\mathfrak{p}}$ for all $\mathfrak{p} \in U$ which arise locally from ring elements; that is, U can be covered by some distinguished open subsets $D(f)$ for each of which we can find a ring element r_f such that $s_{D(f)} = r_f$ (that is, for $\mathfrak{p} \in D(f)$, $s(\mathfrak{p})$ is the image of r_f in $R_{\mathfrak{p}}$.)

The construction of this presheaf makes it clear that it is actually a sheaf. What is less obvious is that $\mathcal{O}(D(f)) = R_f$. More precisely, there is a natural map $R_f \to \mathcal{O}(D(f))$ which we want to be an isomorphism.

To see that this map is injective, it will be enough to check directly the claim from the previous section, that

$$R_f \to \prod_{\mathfrak{p}\in D(f)} R_{\mathfrak{p}}$$

is injective. For ease of notation, we may just do the case $f = 1$, and then apply it with R replaced by R_f. Suppose $r \in R$ maps to zero in every $R_{\mathfrak{p}}$. The set

$$\mathrm{Ann}(r) = \{r' \in R : rr' = 0\}$$

is an ideal of R, but by hypothesis it cannot be contained in any prime ideal; it is thus the trivial ideal, so $r = 0$.

To see that it is surjective, again we need only treat the case $f = 1$. Let $s \in \mathcal{O}(\mathrm{Spec}(R))$ be a section. By hypothesis, we can cover $\mathrm{Spec}(R)$ with open subsets $D(f_i)$ on each of which can represent s using an element $r_i \in R_{f_i}$. By the injectivity argument we just made, the elements r_i, r_j must have the same image in R_{ij}. Moreover, the elements f_i generate the unit ideal, so I can write 1 as a linear combination using only finitely many of them (that is, $\mathrm{Spec}(R)$ is a *quasicompact* topological space: it is not typically Hausdorff but every open covering has a finite subcovering). So now I have a finite list of elements $f_1, \dots, f_n$ and a set of elements $r_i \in R_{f_i}$ such that r_i, r_j have the same image in $R_{f_i f_j}$; by the earlier corollary, these come from a single element $r \in R$. This completes the proof of the theorem.