The Spectrum of a Commutative Ring

Mumford, *Red Book*, Ch. II, §1

Jacob Monzel

University of Pennsylvania

July 31, 2026

Spec as a Set

We will always let R be a commutative unital ring.

Definition 1

Let R be a ring. We define the set

$$\mathrm{Spec}(R) := \{\mathfrak{p} \subset R : \mathfrak{p} \ \text{ is a prime ideal}\}$$

We write $[\mathfrak{p}]$ for the element of $\mathrm{Spec}(R)$ given by the prime ideal $\mathfrak{p}$.

Definition 2

Let $M \subset R$. We define the subset of $\mathrm{Spec}(R)$:

$$V(M) := \{\mathfrak{p} \in \mathrm{Spec}(R) : \mathfrak{p} \supset M\}$$

What can we say about $V(M)$ vs $V(MR)$? (A prime containing M contains the ideal M generates, so $V(M) = V(MR)$: nothing is lost by only considering ideals.)

Why Primes and Not Just Maximal Ideals?

For varieties (Ch. I) the points were maximal ideals. Why does Spec take *all* primes?

Remark 3 (Functoriality)

If $\varphi : R \to S$ is a ring homomorphism and $\mathfrak{q} \subset S$ is prime, then $\varphi^{-1}(\mathfrak{q})$ is prime — so every ring map induces $\mathrm{Spec}(S) \to \mathrm{Spec}(R)$, $[\mathfrak{q}] \mapsto [\varphi^{-1}(\mathfrak{q})]$. But the preimage of a *maximal* ideal need not be maximal: for $\mathbb{Z} \hookrightarrow \mathbb{Q}$, the preimage of (0) is (0). So MaxSpec is not functorial, and we could not even define morphisms with it.

Varieties dodge this: for finite-type algebras over a field, preimages of maximal ideals ARE maximal (Zariski's lemma / Nullstellensatz), so Ch. I got away with closed points only.

(Foreshadowing §6: taking all primes buys us morphisms, at the cost of a strange point set — the functor of points is the systematic repair.)

The Zariski Topology

Lemma 4

The sets of the form $V(J)$ for some ideal $J \subset R$ satisfy the following properties:

1. $V(\sum J_k) = \bigcap V(J_k)$
2. $V(IJ) = V(I \cap J) = V(I) \cup V(J)$
3. $V(R) = \emptyset$ *and* $V(0) = \mathrm{Spec}(R)$.

▸ *proof*

This tells us that the set $\mathrm{Spec}(R)$ plus the collection of closed sets $V(J)$ form a topological space called the Zariski topology.

Lemma 5

For $f \in R$, let $D(f) = \mathrm{Spec}(R) \setminus V(f)$. Then the sets $D(f)$ are a basis of open sets for the Zariski topology.

▸ *proof*

The Zariski Topology

Remark 6

$\mathrm{Spec}(R)$ need not be a T_1 space. This is because the closure of a point $[\mathfrak{p}]$ is exactly $V(\mathfrak{p}) = \{[\mathfrak{q}] : \mathfrak{q} \supset \mathfrak{p}\}$ (from $\mathfrak{p} \in V(I) \iff I \subset \mathfrak{p} \iff V(\mathfrak{p}) \subset V(I)$). So a point is closed if and only if it is a maximal ideal. If R is a domain, (0) is prime with closure $V(0) = \mathrm{Spec}(R)$: a generic point.

Definition 7

Suppose Z is an irreducible closed subset of $\mathrm{Spec}(R)$. Then a point $z \in Z$ is a generic point of Z if Z is the closure of z, i.e. every *nonempty* open subset of Z contains z.

Proposition 8 (Mumford, Prop. 1)

If $x \in \mathrm{Spec}(R)$, then the closure of $\{x\}$ is irreducible and x is a generic point of this set. Conversely, every irreducible closed subset $Z \subset \mathrm{Spec}(R)$ equals $V(\mathfrak{p})$ for some prime ideal $\mathfrak{p} \subset R$, and $[\mathfrak{p}]$ is its unique generic point.

proof

The Galois Connection

Packaging what we know: writing $I(Z) = \bigcap_{[\mathfrak{p}] \in Z} \mathfrak{p}$ for closed $Z \subset \mathrm{Spec}(R)$, the maps $A \mapsto V(A)$ and $Z \mapsto I(Z)$ give an *order-reversing bijection*

$$\{\text{radical ideals of } R\} \longleftrightarrow \{\text{closed subsets of } \mathrm{Spec}(R)\},$$

since $V(A) = V(\sqrt{A})$ and $I(V(A)) = \sqrt{A}$ (a radical ideal is the intersection of the primes containing it). Special cases:

$$\begin{array}{ccc} \text{radical ideals} & \longleftrightarrow & \text{closed sets} \\ \text{prime ideals} & \longleftrightarrow & \text{irreducible closed sets} \\ \text{maximal ideals} & \longleftrightarrow & \text{closed points} \end{array}$$

The middle row is exactly the Proposition; this is the statement to reach for at a board.

The Zariski Topology

Proposition 9 (Mumford, Prop. 2)

Let $\{f_\alpha \mid \alpha \in S\}$ *be a set of elements of* R. *Then*

$$\mathrm{Spec}(R) = \bigcup_{\alpha \in S} D(f_\alpha)$$

if and only if $1 \in (\ldots, f_\alpha, \ldots)$, *the ideal generated by the* f_α*'s.*

proof

Note $1 \in (\ldots, f_\alpha, \ldots)$ iff finitely many f_{α_i} and $g_i \in R$ satisfy $1 = \sum_{i=1}^n g_i f_{\alpha_i}$: the algebraic analog of a partition of unity

Corollary 10 (Mumford, p. 68)

Spec(R) *is quasicompact.*

▸ proof

Remark 11

The same circle of ideas shows every *distinguished* open $D(f)$ is quasicompact.

▸ proof

However, open sets of Spec(R) in general need not be quasicompact: take $R = k[x_1, x_2, x_3, \ldots]$ and

$$U = \bigcup_{i \geq 1} D(x_i) = \mathrm{Spec}(R) \setminus V((x_1, x_2, \ldots)).$$

Any finite subfamily $D(x_{i_1}) \cup \cdots \cup D(x_{i_n})$ misses the point $[(x_{i_1}, \ldots, x_{i_n})] \in U$.

Aside: Which Spaces Are Spectra? (Hochster, 1969)

Everything proved so far assembles into an answer to a question Mumford could not have answered in these 1966–67 lectures. Call a space *spectral* if it is quasicompact, T_0, *sober* (every irreducible closed set has a unique generic point), and has a basis of quasicompact opens stable under finite intersection.

We have shown $\operatorname{Spec}(R)$ is spectral: quasicompact (the Corollary), T_0 (distinct primes have distinct closures), sober (the Proposition), and the $D(f)$ are a quasicompact basis with $D(f) \cap D(g) = D(fg)$.

Theorem 12 (Hochster, 1969)

A topological space is homeomorphic to $\operatorname{Spec}(R)$ *for some commutative ring R if and only if it is spectral.*

The surprise is sufficiency: the obvious necessary conditions already characterize spectra. Looking ahead: Zariski–Riemann spaces of fields and Huber's adic spaces are spectral; Berkovich spaces are the deliberate exception, trading spectrality for Hausdorffness.

Quick Recall of Localization UMP

Let R be a commutative ring and $S \subseteq R$ a multiplicative subset. The localization map $\ell\colon R \to S^{-1}R$ satisfies the following universal property:

Universal Property

For any ring homomorphism $\varphi\colon R \to T$ such that $\varphi(s) \in T^\times$ for all $s \in S$, there exists a unique homomorphism $\widetilde{\varphi}\colon S^{-1}R \to T$ making the diagram commute:

$$\begin{tikzcd} R \arrow[r, "\ell"] \arrow[dr, "\varphi"'] & S^{-1}R \arrow[d, dashed, "\exists!\,\widetilde{\varphi}"] \\ & T \end{tikzcd}$$

i.e. $\widetilde{\varphi} \circ \ell = \varphi$.

In other words, ℓ is initial among ring maps out of R inverting S.

The Structure Sheaf

We now want to endow the space $\mathrm{Spec}(R)$ with a sheaf of rings $\mathcal{O}_{\mathrm{Spec}(R)}$ called the structure sheaf. We will write $X = \mathrm{Spec}(R)$ for simplicity. To this end, we need some properties of distinguished open sets:

a) $D(f) \cap D(g) = D(fg)$ for all $f, g \in R$. (just take the complement of $V(fg) = V(f) \cup V(g)$)
b) $D(f) \supset D(g)$ if and only if $g \in \sqrt{(f)}$. In fact, since $\sqrt{(f)} = \bigcap\{\mathfrak{p} : f \in \mathfrak{p}\}$,

$$g \notin \sqrt{(f)} \iff \exists \mathfrak{p},\ f \in \mathfrak{p},\ g \notin \mathfrak{p} \iff \exists \mathfrak{p},\ [\mathfrak{p}] \notin D(f),\ [\mathfrak{p}] \in D(g).$$

We want to associate the localization R_f at the multiplicative system $(f, f^2, f^3, \ldots)$ to the open set $D(f)$. If $D(g) \subset D(f)$, then by (b) we have $g^n = hf$ for some $h \in R$ and $n \geq 1$, and there is a canonical map

$$R_f \longrightarrow R_g, \qquad \frac{a}{f^m} \longmapsto \frac{a \cdot h^m}{g^{nm}}.$$

(One needs to check this is independent of the choice of h and n AND well-defined AND even a ring homomorphism ... or one can use the UMP)

The Structure Sheaf

In particular, if $D(g) = D(f)$, the canonical maps $R_f \to R_g$ and $R_g \to R_f$ are mutually inverse, so we may identify R_f and R_g: the ring R_f really only depends on the open set $D(f)$. Whenever $D(k) \subset D(g) \subset D(f)$ the canonical maps form a commutative triangle:

$$\begin{tikzcd} R_f \arrow[rr] \arrow[dr] & & R_k \\ & R_g \arrow[ur] & \end{tikzcd}$$

Moreover, if $[\mathfrak{p}] \in D(f)$, then $f \notin \mathfrak{p}$, so $R \setminus \mathfrak{p}$ contains $\{f, f^2, \dots\}$ and there is an obvious map $R_f \to R_{\mathfrak{p}}$. One can check that

$$R_{\mathfrak{p}} = \varinjlim_{f \in R \setminus \mathfrak{p}} R_f = \varinjlim_{[\mathfrak{p}] \in D(f)} R_f .$$

where we take the direct limit over the "obvious" direct system. (Again one needs to prove all of this is actually well-defined and makes sense ... but this is easy using the UMP)

Recall: Zero Is a Local Property

Lemma 13 (Local–global principle for vanishing)

Let M be an R-module and $m \in M$. Then $m = 0$ if and only if the image of m in $M_{\mathfrak{p}}$ is 0 for every prime $\mathfrak{p} \subset R$.

▸ proof

The closed set $V(\mathrm{Ann}(m))$ appearing in the proof is called the *support* of m. The same argument, applied to R_f as a module over itself, gives the relative version: an element $g \in R_f$ is 0 iff its image in $R_{\mathfrak{p}}$ is 0 for every $[\mathfrak{p}] \in D(f)$.

The Structure Sheaf: Local Vanishing

Lemma 14 (Local vanishing; Mumford, Lemma 1)

Suppose $D(f) = \bigcup_{\alpha \in S} D(f_\alpha)$. If $g \in R_f$ has image 0 in every R_{f_α}, then $g = 0$. proof

Idea in one line: covers let every $R_f \to R_{\mathfrak{p}}$ factor through some R_{f_α}, and then the local–global principle (Lemma 13) finishes.

In sheaf language: $R_{\mathfrak{p}}$ will be the stalk $\mathcal{O}_{X,[\mathfrak{p}]}$, and this lemma is exactly separatedness of $\mathcal{O}_X$ on the basis of distinguished opens ... i.e. the identity sheaf axiom.

The Structure Sheaf: Gluing

Lemma 15 (Gluing; Mumford, Lemma 2)

Suppose $D(f) = \bigcup_{\alpha \in S} D(f_\alpha)$*, and let* $g_\alpha \in R_{f_\alpha}$ *have equal images in* $R_{f_\alpha f_\beta}$ *for all* α, β*. Then some* $g \in R_f$ *maps to* g_α *in* R_{f_α} *for all* α*.*

proof

Idea in one line: quasicompactness reduces to a *finite* subcover, and then the algebraic partition of unity $1 = \sum h_i f_i^N$ lets you write down the glued element explicitly.

In sheaf language: this is the gluing sheaf axiom on the basis of distinguished opens.

What the Two Lemmas Really Say

Fix a cover $D(f) = \bigcup_\alpha D(f_\alpha)$. The above lemmas together say that the sequence

$$0 \longrightarrow R_f \longrightarrow \prod_\alpha R_{f_\alpha} \longrightarrow \prod_{\alpha,\beta} R_{f_\alpha f_\beta}$$

is exact, where the first map is the product of the restrictions and the second sends $(g_\alpha) \mapsto (g_\alpha - g_\beta)$ (computed in $R_{f_\alpha f_\beta}$):

- Local vanishing (Lemma 14) = exactness at R_f;
- Gluing (Lemma 15) = exactness at $\prod_\alpha R_{f_\alpha}$.

The Structure Sheaf: Definition

The lemmas say that $D(f) \mapsto R_f$ is "as close to a sheaf as a basis allows," and there is exactly one way to extend it to all opens. For each open $U \subset X$, define $\Gamma(U, \mathcal{O}_X)$ to be the set of elements

$$\{s_\mathfrak{p}\} \in \prod_{[\mathfrak{p}] \in U} R_\mathfrak{p}$$

for which there exists a covering of U by distinguished opens $D(f_\alpha)$ and elements $s_\alpha \in R_{f_\alpha}$ such that $s_\mathfrak{p}$ equals the image of s_α in $R_\mathfrak{p}$ whenever $[\mathfrak{p}] \in D(f_\alpha)$. This is a ring, and for $V \subset U$ the coordinate projection $\prod_{[\mathfrak{p}] \in U} R_\mathfrak{p} \to \prod_{[\mathfrak{p}] \in V} R_\mathfrak{p}$ takes $\Gamma(U, \mathcal{O}_X)$ into $\Gamma(V, \mathcal{O}_X)$; taking this as restriction we get a presheaf $\mathcal{O}_X$. Then:

1. $\mathcal{O}_X$ is a sheaf (this is trivially true)
2. $\Gamma(D(f), \mathcal{O}_X) = R_f$.
3. The stalk of $\mathcal{O}_X$ at $[\mathfrak{p}]$ is $R_\mathfrak{p}$.

proofs of 2, 3

Maps Between Stalks

Since X has non-closed points, we get maps *between stalks*. If $\mathfrak{p}_1 \subset \mathfrak{p}_2$ and $x_i = [\mathfrak{p}_i]$, then $x_2 \in \overline{\{x_1\}}$, so every neighbourhood of x_2 contains x_1, giving

$$\mathcal{O}_{x_2} = \varinjlim_{x_2 \in U} \Gamma(U, \mathcal{O}_X) \longrightarrow \varinjlim_{x_1 \in V} \Gamma(V, \mathcal{O}_X) = \mathcal{O}_{x_1},$$

which is just the natural localization map $R_{\mathfrak{p}_2} \to R_{\mathfrak{p}_1}$.

(Compare: on a variety, specialization $x_1 \rightsquigarrow x_2$ means "x_2 lies on the subvariety whose generic point is x_1," and the map restricts germs of functions from the bigger locus to the smaller one.)

Examples

Example 16

Let K be a field. Then $\mathrm{Spec}(K)$ is just a point and the structure sheaf is trivial. More generally, if R has DCC then R is the direct sum of its primary subrings $R = \bigoplus_{i=1}^{n} R_i$, and $\mathrm{Spec}(R \oplus S) = \mathrm{Spec}(R) \amalg \mathrm{Spec}(S)$; each R_i has the single prime $\mathrm{Nil}(R_i)$, so $\mathrm{Spec}(R)$ is n discrete points with R_i sitting as the stalk on the i-th point.

Example 17

Let K be a field. Then $\mathrm{Spec}(K[X])$ is called the affine line over K. Since $K[X]$ is a PID we know the only points are (0) and (f) where f is an irreducible polynomial. What does $\mathrm{Spec}(\mathbb{C}[X])$ look like? What about $\mathrm{Spec}(\mathbb{R}[X])$? What about $\mathrm{Spec}(\mathbb{F}_q[X])$?

Example 18

What does $\mathrm{Spec}(\mathbb{Z})$ look like? What are the points? The stalks? What is the function field, i.e. the stalk at the generic point? What are the residue fields?

Aside on Residue Fields

Question 19

Let $x \in \operatorname{Spec}(R)$ *correspond to the prime* $\mathfrak{p} \subset R$. *What is the residue field* $\kappa(x)$ *in general?*
Hint: localization is exact.

Aside on Residue Fields

Question 19

Let $x \in \mathrm{Spec}(R)$ correspond to the prime $\mathfrak{p} \subset R$. What is the residue field $\kappa(x)$ in general? Hint: localization is exact.

By definition $\kappa(x) = R_\mathfrak{p}/\mathfrak{p}R_\mathfrak{p}$. Localizing the exact sequence

$$0 \longrightarrow \mathfrak{p} \longrightarrow R \longrightarrow R/\mathfrak{p} \longrightarrow 0$$

at $\mathfrak{p}$ preserves exactness, so $R_\mathfrak{p}/\mathfrak{p}R_\mathfrak{p} \cong (R/\mathfrak{p})_\mathfrak{p}$. But $\mathfrak{p}$ maps to (0) in the domain $R/\mathfrak{p}$, so localizing inverts every nonzero element:

$$\kappa(x) \cong \mathrm{Frac}(R/\mathfrak{p}).$$

Sanity checks:

- $\mathfrak{p}$ maximal: $R/\mathfrak{p}$ is already a field, so $\kappa(x) = R/\mathfrak{p}$.
- R a domain, $\mathfrak{p} = (0)$ (the generic point): $\kappa(x) = \mathrm{Frac}(R)$, the function field.

Examples

Example 20

What can you say about the spectrum of Dedekind domains? What about DVRs?

Example 21

Consider $\mathrm{Spec}(\mathbb{C}[x, y])$. What are the closed points? What are the other non-closed points?

Example 22

Let K be a field. What can you say about $\mathrm{Spec}(\prod_{i=1}^{\infty} K)$? What are the naive prime ideals? Are we missing any? What are the stalks? Hint: consider ultraproducts and note that such a ring is VNR.

Example: The Flower

Example 23 (Mumford's Example F)

Let $X = \mathrm{Spec}(k[X, Y])$ and let

$$\mathcal{O} = \left\{ \frac{f(X,Y)}{g(X,Y)} : f, g \in k[X, Y],\ g(0,0) \neq 0 \right\},$$

the stalk of $\mathcal{O}_X$ at the origin. Its primes are: the maximal ideal (X, Y); the principal primes (f), f irreducible with $f(0,0) = 0$; and (0). So $\mathrm{Spec}(\mathcal{O})$ has only *one* closed point.

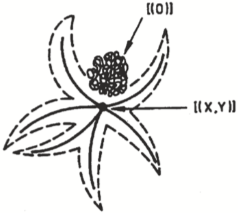

Think of $\mathrm{Spec}(\mathcal{O})$ as what survives of the plane arbitrarily close to the origin: the closed point itself; for each irreducible curve through the origin, its generic point (the "germ" of the curve – the petals of the flower); and the big generic point $[(0)]$. Curves missing the origin, and all other closed points of the plane, are gone.

Example: The Flower

Now the interesting part: throw away the closed point *AND* the Y-axis, i.e. pass to the distinguished open set

$$D(X) = \{X \neq 0\} \subset \mathrm{Spec}(\mathcal{O}).$$

This open set contains *none* of the closed points of $\mathrm{Spec}(\mathcal{O})$: its points are $[(0)]$ and the germs $[(f)]$ with $f \neq X$. A nonempty open subset of a scheme can miss *every* closed point of the ambient scheme! (Of course $D(X)$ has closed points in its *own* topology, the germs $[(f)]$, as it must since it is itself affine, i.e. rings always have maximal ideals)

On the other hand, let $K = k(X)$. Then this scheme is $\mathrm{Spec}\big(K[Y]_S\big)$ for a certain multiplicative system $S \subset K[Y]$: it is part of the affine line over K.

Moral: localization can trade "a punctured infinitesimal neighbourhood in a surface over k" for "a chunk of a line over the bigger field K." What looks like a point of a curve from one side is a whole curve (its generic point) from the other, i.e. dimension is only meaningful relative to the base.

Example

Consider Spec($\mathbb{Z}[X]$). This is an arithmetic surface. The prime ideals are (0), (f) for f either a prime number or a $\mathbb{Q}$-irreducible polynomial with content 1, OR (p, f) where p is a prime and f a monic integral polynomial irreducible modulo p.

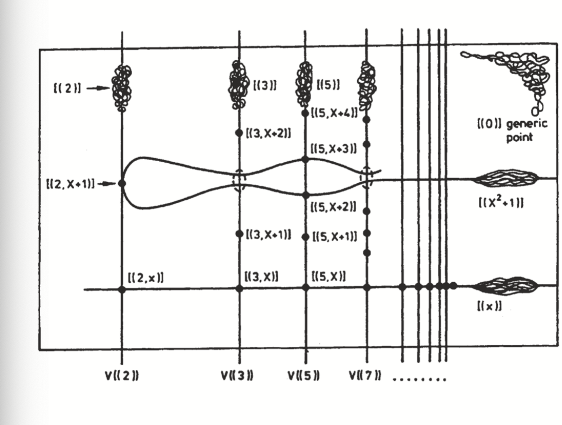

Figure:

Why is Spec($\mathbb{Z}[X]$) a "Surface"?

Each closed fiber $V((p))$ is a copy of $\mathbb{A}^1_{\mathbb{Z}/p\mathbb{Z}} = \operatorname{Spec}(\mathbb{F}_p[X])$; the picture pastes all of them together, with the "horizontal" curves $V((f))$ running across the fibers. The whole set-up is called a surface for two reasons:

1. All maximal chains of irreducible proper closed subsets have length 2, just as in $\mathbb{A}^2_k$: point $\subsetneq$ curve, e.g.

$$\{[(2, X+1)]\} \subsetneq V((X+1)) \qquad \text{and} \qquad \{[(2, X+1)]\} \subsetneq V((2)).$$

2. If $\mathcal{O}$ is the local ring at a closed point $x = [(p, f)]$, then $\mathcal{O}$ has Krull dimension 2: its maximal ideal $\mathfrak{m}$ is generated by p and f, and there is *no single element* $g \in \mathcal{O}$ with $\mathfrak{m} = \sqrt{(g)}$. The point is a genuine codimension-2 intersection, cut out by one "arithmetic direction" (p) and one "geometric direction" (f).

So Spec($\mathbb{Z}[X]$) fibers over the "curve" Spec($\mathbb{Z}$) the way a surface fibers over a curve.

Exercise (Mumford, p. 75)

Question 24

On $\mathrm{Spec}(\mathbb{Z}[X])$*:*

1. *What is* $V((p)) \cap V((f))$*, for* p *a prime and* f *a* $\mathbb{Q}$*-irreducible polynomial of content 1?*
2. *What is* $V((f)) \cap V((g))$*, for* f *and* g *distinct* $\mathbb{Q}$*-irreducible polynomials of content 1?*

Hints: intersections of closed sets are V of the sum: $V((p)) \cap V((f)) = V((p, f))$. Which primes contain (p, f)?

Use the picture: where does the horizontal curve $V((f))$ cross the vertical fiber $V((p))$? And where can two distinct horizontal curves possibly meet?

Exercise: Answers

1. $V((p)) \cap V((f)) = V((p, f))$, and primes containing (p, f) correspond to primes of $\mathbb{Z}[X]/(p) = \mathbb{F}_p[X]$ containing $\bar{f} = f \bmod p$. Content 1 forces $\bar{f} \neq 0$, so these are the ideals $(\bar{h})$ for $\bar{h}$ an irreducible factor of $\bar{f}$: we get the *finite* set of closed points

$$\{[(p, h)] : h \text{ monic irreducible dividing } f \bmod p\},$$

at most $\deg f$ of them (empty exactly when $\bar{f}$ is a nonzero constant, e.g. $f = pX + 1$). The curve meets each fiber in the roots of $f \bmod p$.

2. $V((f)) \cap V((g)) = V((f, g))$. In $\mathbb{Q}[X]$ we have $\gcd(f, g) = 1$, so $af + bg = 1$; clearing denominators gives $uf + vg = N$ for some nonzero integer N (one can take $N = \mathrm{Res}(f, g)$). Any prime containing f and g therefore contains N, hence lies over one of the finitely many primes $p \mid N$; over such p, part 1 gives the points $[(p, h)]$ with h a *common* irreducible factor of $\bar{f}$ and $\bar{g}$. Conclusion: two distinct horizontal curves meet in finitely many closed points, exactly over the primes where f and g acquire a common root mod p – detected by the resultant.

Examples: Nilpotents

Example 25

Elements $a \in R$ are "functions" on $\mathrm{Spec}(R)$: the value $a(x)$ at $x = [\mathfrak{p}]$ is the image of a in the residue field $\Bbbk(x) = R_{\mathfrak{p}}/\mathfrak{p}R_{\mathfrak{p}} = \mathrm{Frac}(R/\mathfrak{p})$. But these values lie in different fields at different points, and the function does not determine the element:

$$a(x) = 0 \text{ for all } x \in \mathrm{Spec}(R) \iff a \in \mathfrak{p} \text{ for all primes } \mathfrak{p} \iff a \text{ is nilpotent.}$$

So unlike for varieties, $\mathcal{O}_X$ does not embed in a sheaf of k-valued functions when R has nilpotents.

Example 26 (Dual numbers)

Let $Z = \mathrm{Spec}\big(k[X]/(X^2)\big)$: a single point supporting the nonzero function X whose value at that point is 0. Embedding Z into $\mathbb{A}^1$ at the origin, its sheaf is $\mathcal{O}_{0,\mathbb{A}^1}$ modulo functions vanishing to *second* order: from $f|_Z$ we recover both $f(0)$ and $f'(0)$, i.e. nilpotents record infinitesimal information in some sense.

Examples: Subsets $V(A)$ in Affine Space

Example 27

Most generally, $\mathrm{Spec}\big(k[X_1, \ldots, X_n]/A\big)$ can be interpreted as the subset $V(A) \subset \mathbb{A}^n_k$, but with a sheaf of rings "fattened" so as to include the first few terms of the Taylor expansion of polynomials f in directions *normal* to $V(A)$. E.g. for $\mathrm{Spec}\big(k[X, Y]/(g^2)\big)$, g irreducible: this is the curve $g = 0$ "with multiplicity 2," and from the image of f in $k[X, Y]/(g^2)$ one can reconstruct all first partials of f, including the one normal to $g = 0$. If R is any ring with nilpotents, picture $\mathrm{Spec}(R)$ as carrying extra *normal material*, in the direction of which one can take partial derivatives, but which is not tangent to any dimension actually present in the space $\mathrm{Spec}(R)$.

Distinguished Open Subschemes

Proposition 28 (Mumford, Prop. 3)

Let R be a ring and let $f \in R$. Then the topological space $D(f)$ together with the restriction of the sheaf of rings $\mathcal{O}_{\mathrm{Spec}(R)}$ to $D(f)$ is isomorphic to $\mathrm{Spec}(R_f)$ with its sheaf of rings $\mathcal{O}_{\mathrm{Spec}(R_f)}$.

▸ *proof*

Idea in one line: primes of R_f are primes of R avoiding f (a homeomorphism $D(f) \leftrightarrow \mathrm{Spec}(R_f)$), and on basic opens the sections match via $R_{fg} \cong (R_f)_g$.

Consequence: distinguished opens of affine schemes are again affine — the workhorse fact behind every "reduce to the affine/basic-open case" argument in the sequel.

Appendix: Full Proofs

Every result above has its complete proof here; use the ◂ back buttons (or the navigation arrows, bottom right) to jump back to the statement.

Proof: The Zariski Topology Axioms & the Basis

Proof of Lemma 4.

(1) A prime contains $\sum J_k$ iff it contains every J_k. (2) From $IJ \subset I \cap J \subset I, J$ we get $V(I) \cup V(J) \subset V(I \cap J) \subset V(IJ)$. Conversely, if $\mathfrak{p} \supset IJ$ but $\mathfrak{p} \not\supset I$, pick $a \in I \setminus \mathfrak{p}$; for every $b \in J$ we have $ab \in IJ \subset \mathfrak{p}$, and $\mathfrak{p}$ prime with $a \notin \mathfrak{p}$ forces $b \in \mathfrak{p}$ — so $\mathfrak{p} \supset J$. (3) No prime contains 1; every prime contains 0. $\square$

Proof of Lemma 5.

The $D(f)$ are open. Any open set is $\mathrm{Spec}(R) \setminus V(J)$ for an ideal J, and

$$[\mathfrak{p}] \notin V(J) \iff J \not\subset \mathfrak{p} \iff \exists f \in J,\ f \notin \mathfrak{p} \iff [\mathfrak{p}] \in D(f) \text{ for some } f \in J,$$

so $\mathrm{Spec}(R) \setminus V(J) = \bigcup_{f \in J} D(f)$. $\square$

Proof: Irreducible Closed Sets I

Proof (first statement). Let Z be the closure of $\{x\}$. If $Z = W_1 \cup W_2$ with each W_i closed, then $x \in W_1$ or $x \in W_2$; in either case W_i is a closed set containing x and contained in its closure, so $W_i = Z$. Hence Z is irreducible with generic point x.

Proof (converse), part 1. Suppose $V(A)$ is irreducible. Since $V(A) = V(\sqrt{A})$ we may assume $A = \sqrt{A}$ (and $A \neq R$, since irreducible sets are nonempty). *Claim: A is prime.* If not, there exist $f, g \in R$ with $fg \in A$ but $f, g \notin A$. Let $B = A + (f)$ and $C = A + (g)$. Then $A = B \cap C$: if $h = \alpha f + a_1 = \beta g + a_2 \in B \cap C$ ($a_i \in A$; $\alpha, \beta \in R$), then

$$h^2 = \alpha\beta \cdot fg + a_1(\beta g + a_2) + a_2(\alpha f) \in A,$$

so $h \in A$ (as $A = \sqrt{A}$). Therefore $V(A) = V(B) \cup V(C)$.

Proof: Irreducible Closed Sets II

Proof (converse), part 2. Now, since $A = \sqrt{A}$, A is the intersection of the primes containing it; in particular some prime $\mathfrak{p} \supset A$ has $f \notin \mathfrak{p}$. Then $[\mathfrak{p}] \in V(A) \setminus V(B)$, so $V(A) \not\subset V(B)$, and similarly $V(A) \not\subset V(C)$, contradicting irreducibility. Hence A is prime, and $V(A) = \overline{\{[A]\}}$ as noted earlier, so $[A]$ is a generic point of $V(A)$.

Uniqueness. If $[\mathfrak{p}']$ were another generic point, then $[\mathfrak{p}'] \in V(A)$ gives $A \subset \mathfrak{p}'$, while $[A]$ in the closure of $[\mathfrak{p}']$ gives $A \supset \mathfrak{p}'$; so $A = \mathfrak{p}'$. $\square$

Remark 29

Notice that the real content of this result is that we have a bijection between irreducible closed subsets of $\mathrm{Spec}(R)$ and points of $\mathrm{Spec}(R)$ via $\mathfrak{p} \mapsto V(\mathfrak{p})$.

back

Proof: Covers and Partitions of Unity

Proof of Proposition 9.

The $D(f_\alpha)$ cover $\mathrm{Spec}(R)$ iff every prime $\mathfrak{p}$ fails to contain some f_α, i.e. iff no prime ideal contains $(\dots, f_\alpha, \dots)$, i.e. iff $1 \in (\dots, f_\alpha, \dots)$. $\square$

back

Proof: Quasicompactness

Proof of the Corollary.

It suffices to check every covering by basic open sets has a finite subcover. If $\mathrm{Spec}(R) = \bigcup_\alpha D(f_\alpha)$, then $1 \in (\ldots, f_\alpha, \ldots)$, so already $1 \in (f_{\alpha_1}, \ldots, f_{\alpha_n})$ for some finite subset, and $\mathrm{Spec}(R) = D(f_{\alpha_1}) \cup \cdots \cup D(f_{\alpha_n})$. □

Proof that $D(f)$ is quasicompact.

If $D(f) \subset \bigcup_\alpha D(g_\alpha)$, then $V((\ldots, g_\alpha, \ldots)) \subset V(f)$, so $f \in \sqrt{(\ldots, g_\alpha, \ldots)}$, i.e. $f^n = \sum_{i=1}^k h_i g_{\alpha_i}$ for some n; then $D(f) \subset D(g_{\alpha_1}) \cup \cdots \cup D(g_{\alpha_k})$. □

back

Proof: Zero Is a Local Property

Proof of Lemma 13.

The image of m in $M_{\mathfrak{p}}$ vanishes iff $cm = 0$ for some $c \notin \mathfrak{p}$, i.e. iff $\mathrm{Ann}(m) \not\subset \mathfrak{p}$. Hence

$$\{\, \mathfrak{p} : m \neq 0 \text{ in } M_{\mathfrak{p}} \,\} = V(\mathrm{Ann}(m)).$$

If $m \neq 0$, then $\mathrm{Ann}(m)$ is a proper ideal, so it is contained in some prime, and this set is nonempty. $\square$

(The same argument, applied to R_f as a module over itself, gives the relative version stated on the main slide.)

back

Proof: Local Vanishing

Proof of Lemma 14.

Let $[\mathfrak{p}] \in D(f)$. Since the $D(f_\alpha)$ cover $D(f)$, we have $[\mathfrak{p}] \in D(f_\alpha)$ for some α, i.e. $f_\alpha \notin \mathfrak{p}$. Then f_α is invertible in $R_\mathfrak{p}$, so by the universal property of localization the map $R_f \to R_\mathfrak{p}$ factors as

$$R_f \longrightarrow R_{f_\alpha} \longrightarrow R_\mathfrak{p}.$$

Since $g \mapsto 0$ in R_{f_α}, also $g \mapsto 0$ in $R_\mathfrak{p}$. As $[\mathfrak{p}] \in D(f)$ was arbitrary, the local–global principle (Lemma 13, relative form) gives $g = 0$. $\square$

back

Proof: Gluing, Reduction to Finite Covers

We treat the case $f = 1$; the general case is identical.

Step 1: X is quasicompact. A cover by the $D(f_\alpha)$ means $1 \in (f_\alpha : \alpha \in S)$, so $1 = \sum_{i=1}^k h_i f_i$ is a *finite* sum, and $D(f_1), \ldots, D(f_k)$ already cover X.

Step 2: gluing over the subcover suffices. Suppose $g \in R$ restricts to g_i for $i = 1, \ldots, k$. For any other α, the elements g and g_α of R_{f_α} agree in each $R_{f_\alpha f_i}$, and $D(f_\alpha) = \bigcup_i D(f_\alpha f_i)$; by local vanishing (Lemma 14) applied to $g - g_\alpha$, they are equal.

Proof: Gluing, Partition-of-Unity Computation

Step 3: glue over $X = D(f_1) \cup \cdots \cup D(f_k)$. Write $g_i = b_i / f_i^n$ with one n for all i (possible: finitely many). Agreement in $R_{f_i f_j}$ means

$$(f_i f_j)^m \left[b_i f_j^n - b_j f_i^n \right] = 0$$

for a common m. Set $N = n + m$ and $b_i' = b_i f_i^{N-n}$; then $b_i' f_j^N = b_j' f_i^N$ *exactly*.

Since $D(f_i) = D(f_i^N)$, we still have $1 = \sum h_i f_i^N$. Set $g = \sum_i h_i b_i'$. Then for each j,

$$f_j^N \cdot g = \sum_i f_j^N b_i' h_i = \sum_i f_i^N b_j' h_i = b_j' \sum_i h_i f_i^N = b_j',$$

so $g \mapsto b_j' / f_j^N = g_j$ in R_{f_j}. $\square$

back

Proof: Sections and Stalks of $\mathcal{O}_X$

Proof of 2. The map $R_f \to \prod_{[\mathfrak{p}]\in D(f)} R_\mathfrak{p}$ is injective by the relative local–global principle (Lemma 13) and lands in $\Gamma(D(f), \mathcal{O}_X)$ by definition (take the cover by $D(f)$ itself). For surjectivity: a section of $\Gamma(D(f), \mathcal{O}_X)$ gives a cover $D(f) = \bigcup D(f_\alpha)$ and $s_\alpha \in R_{f_\alpha}$ agreeing *stalkwise* on overlaps; applying Lemma 13 to $s_\alpha - s_\beta$ on $D(f_\alpha f_\beta)$, they agree in $R_{f_\alpha f_\beta}$, so Lemma 15 (gluing) produces the desired element of R_f. □

Proof of 3. Compute the stalk as a direct limit over the distinguished opens containing $[\mathfrak{p}]$:

$$(\mathcal{O}_X)_{[\mathfrak{p}]} = \varinjlim_{[\mathfrak{p}]\in D(f)} \Gamma(D(f), \mathcal{O}_X) = \varinjlim_{f\notin\mathfrak{p}} R_f = R_\mathfrak{p}. \qquad \square$$

back

Proof: Distinguished Open Subschemes

Proof.

Let $i : R \to R_f$, $r \mapsto r/1$. If $\mathfrak{p}$ is a prime of R with $f \notin \mathfrak{p}$, then $i(\mathfrak{p}) \cdot R_f$ is a prime of R_f; if $\mathfrak{q}$ is a prime of R_f, then $i^{-1}(\mathfrak{q})$ is a prime of R not containing f. These mutually inverse maps give a bijection $D(f) \leftrightarrow \operatorname{Spec}(R_f)$, which one checks is a homeomorphism matching, for each $g \in R$, the basic open sets $D(fg) \subset D(f)$ and $D(g/1) \subset \operatorname{Spec}(R_f)$. The sections of the two structure sheaves on these are $\mathcal{O}_{\operatorname{Spec}(R)}(D(fg)) = R_{fg}$ and $\mathcal{O}_{\operatorname{Spec}(R_f)}(D(g/1)) = (R_f)_g$, and the canonical map $R_{fg} \to (R_f)_g$ is an isomorphism, since inverting fg inverts both f and g and vice versa. These identifications on a basis are compatible with restriction, hence glue to an isomorphism $\mathcal{O}_{\operatorname{Spec}(R)}|_{D(f)} \cong \mathcal{O}_{\operatorname{Spec}(R_f)}$ over the homeomorphism of spaces. $\square$

back