Math 8865: Introduction to Algebraic Geometry

Brian Lehmann

May 27, 2021

Contents

References

I have drawn from/plagarized many sources while writing these course notes; my approach has been heavily influenced by [Vak17] and [Gat20]. I used [Har95] as a reference for Chapter III and Chapter VI, [Vak17] for Chapters IV and V, [EO20] for Section 19, [Gat20] for Sections 24 and 35, and [GW10] for Section 33. I have taken exercises from all of these references and also Fulton, Hartshorne.

Part I

Schemes over a field

Chapter 1

Affine schemes

Our construction of schemes is inspired by the theory of manifolds. A manifold M is a topological space equipped with an open covering by charts which are isomorphic to open sets in a vector space. We can obtain different types of manifolds by allowing different types of "structural functions" on M: continuous, smooth, holomorphic, etc.

In this chapter we will construct affine schemes, the "charts" in the theory of algebraic geometry. In our setting the "structural functions" will be polynomial functions. There are a couple distinctive features of the algebro-geometric approach:

(1) Typically one constructs a manifold by first identifying the underlying set and then imposing a set of functions upon it. In algebraic geometry we view functions as the foundational objects: we use the ring of functions to construct a set and a topology. Sections 1-7 are dedicated to this construction.

(2) Any open subset of a manifold is also a manifold in an obvious way. In contrast, an open subset of an affine scheme need not be an affine scheme. In Sections 8-12 we analyze how to put open subsets on an equal footing with affine schemes.

Algebraic sets

Let's briefly review the classical theory of algebraic sets. Fix a field $\mathbb{K}$ and consider the vector space $\mathbb{K}^n$. If we let $x_1, \ldots, x_n$ denote the coordinate functions on $\mathbb{K}^n$, then the polynomial ring $\mathbb{K}[x_1, \ldots, x_n]$ is the set of algebraic functions on $\mathbb{K}^n$. An algebraic set in $\mathbb{K}^n$ is a subset of $\mathbb{K}^n$ defined by polynomial equations.

Unfortunately this notion has some serious flaws. The main issue is that an algebraic set can fail to encapsulate the richness of the equations we used to define it. For example, suppose we are working in $\mathbb{R}^2$. The locus defined by the equation $x^2 + y^2 = 0$ is simply a point, while the equation $x^2 + y^2 + 1 = 0$ defines the empty set. In particular there is no way to recover the defining equation from the set. In this chapter we will to correct these inadequacies.

In classical algebraic geometry, a morphism between two algebraic sets X, Y is defined as follows. First one identifies vector spaces $\mathbb{K}^n$ and $\mathbb{K}^m$ containing X and Y respectively. A polynomial function from $\mathbb{K}^n$ to $\mathbb{K}^m$ is any map given in coordinates by

$$(f_1(x_1, \ldots, x_n), \ldots, f_m(x_1, \ldots, x_n))$$

where each f_i is a polynomial function on $\mathbb{K}^n$. A morphism $f : X \to Y$ is the restriction of any polynomial function $\mathbb{K}^n \to \mathbb{K}^m$ such that the image of X is contained in Y. Just as we need to replace algebraic sets with a different construction, we will also need a new notion of a morphism. Our notion will have the advantage of being "intrinsic", i.e. independent of a choice of embedding.

Primer on $\mathbb{K}$-algebras

Throughout the text $\mathbb{K}$ will denote a field. A $\mathbb{K}$-algebra is a unital commutative ring R such that there is a non-zero homomorphism from $\mathbb{K}$ into the center of R. We also formally include the 0-ring in our set of $\mathbb{K}$-algebras. A $\mathbb{K}$-algebra homomorphism is a unital ring homomorphism that is also a morphism of $\mathbb{K}$-vector spaces. Note that any $\mathbb{K}$-algebra R admits a unique $\mathbb{K}$-algebra homomorphism $\mathbb{K} \to R$ which (by necessity) sends the unit in $\mathbb{K}$ to the unit in R.

A finitely generated $\mathbb{K}$-algebra is the same thing as a quotient of a polynomial ring $\mathbb{K}[x_1, \ldots, x_n]$ by an ideal. Every finitely generated $\mathbb{K}$-algebra R will be a Noetherian ring. This implies that every ideal $I \subset R$ admits a finite set of associated primes. (In particular, each ideal I admits a finite set of minimal primes lying over I.)

An important property of finitely generated $\mathbb{K}$-algebras that is not shared by every Noetherian ring is that they are Jacobson rings: for any ideal I the radical $\sqrt{I}$ is the intersection of the maximal ideals containing I.

A finitely generated $\mathbb{K}$-algebra R will be an Artinian ring when any of the following equivalent criteria are satisfied:

(1) R has only finitely many maximal ideals.

(2) Every prime ideal in R is a maximal ideal.

(3) R is a product of finitely many local rings R_i (each of which is a finitely generated $\mathbb{K}$-algebra).

1.1 Affine space

As discussed in the introduction, the classical notion of an "algebraic set" should be replaced by a different construction. Our first task will be to revisit the vector space $\mathbb{K}^n$ from this new perspective. In accordance with this philosophy that "function rings are the basic objects," we start by reinterpreting points in the vector space $\mathbb{K}^n$ from the perspective of functions.

Let $\vec{a} = (a_1, \ldots, a_n) \in \mathbb{K}^n$ be a point. Given any function

$$f \in \mathbb{K}[x_1, \ldots, x_n]$$

we can evaluate f at $\vec{a}$ by taking the image of f under the map $\mathbb{K}[x_1, \ldots, x_n] \to \mathbb{K}$ that sends $x_i \mapsto a_i$. In other words, evaluation at $\vec{a}$ corresponds to quotienting $\mathbb{K}[x_1, \ldots, x_n]$ by the maximal ideal $(x_1 - a_1, \ldots, x_n - a_n)$ consisting of polynomials which vanish at $\vec{a}$.

Conversely, suppose that $\mathfrak{m}$ is a maximal ideal of $\mathbb{K}[x_1, \ldots, x_n]$ such that the quotient $\mathbb{K}[x_1, \ldots, x_n]/\mathfrak{m}$ is isomorphic to $\mathbb{K}$. Then we claim that $\mathfrak{m}$ is the ideal of functions which vanish at some point of $\mathbb{K}^n$. Indeed, if we let $a_i \in \mathbb{K}$ be the image of x_i under this quotient map then $x_i - a_i \in \mathfrak{m}$ for $i = 1, \ldots, n$. Since $\mathfrak{m}$ is a maximal ideal, it must coincide with $(x_1 - a_1, \ldots, x_n - a_n)$. Altogether we see that there is a bijection

$$\text{Points in } \mathbb{K}^n \qquad \leftrightarrow \qquad \begin{array}{c} \text{Maximal ideals } \mathfrak{m} \subset \mathbb{K}[x_1, \ldots, x_n] \\ \text{such that } \mathbb{K}[x_1, \ldots, x_n]/\mathfrak{m} \cong \mathbb{K} \end{array}$$

From the ring-theoretic viewpoint, there is no special reason to single out the maximal ideals in $\mathbb{K}[x_1, \ldots, x_n]$ whose quotient is $\mathbb{K}$. Our first construction will enrich the bijection above to include all maximal ideals.

Definition 1.1.1. Let $\mathbb{K}$ be a field. We define affine space $\mathbb{A}^n_{\mathbb{K}}$ to be the set of maximal ideals in $\mathbb{K}[x_1, \ldots, x_n]$. (We will often suppress the field $\mathbb{K}$ in the notation when it is understood.)

The points of $\mathbb{A}^n_{\mathbb{K}}$ come in two types: the "traditional" points $\mathfrak{m}$ such that $R/\mathfrak{m} \cong \mathbb{K}$, and the "non-traditional points" $\mathfrak{m}$ such that $R/\mathfrak{m} \not\cong \mathbb{K}$. The traditional points are in bijection with the points of $\mathbb{K}^n$ and we will frequently pass back and forth between the algebraic notation $(x_1 - a_1, \ldots, x_n - a_n)$ and the geometric notation $(a_1, \ldots, a_n)$ for these points. We can derive some intuition for non-traditional points using Hilbert's Nullstellensatz.

Theorem 1.1.2 (Hilbert's Nullstellensatz)**.** *Let R be a finitely generated $\mathbb{K}$-algebra and let $\mathfrak{m}$ be a maximal ideal in R. Then the quotient $R/\mathfrak{m}$ is a finite field extension of $\mathbb{K}$.*

Suppose we are given a point $\mathfrak{m} \in \mathbb{A}^n_{\mathbb{K}}$. Then we can "evaluate" any function $f \in \mathbb{K}[x_1, \ldots, x_n]$ at $\mathfrak{m}$ by taking the image of f in the quotient $\mathbb{K}[x_1, \ldots, x_n]/\mathfrak{m}$. While evaluation at $\mathfrak{m}$ doesn't take values in $\mathbb{K}$, perhaps we can agree that taking values in a finite extension is good enough!

The field $\mathbb{K}[x_1, \ldots, x_n]/\mathfrak{m}$ is called the residue field of $\mathfrak{m}$. For example, the points of $\mathbb{A}^n_{\mathbb{K}}$ which have residue field $\mathbb{K}$ are what we have been calling "traditional points."

Example 1.1.3. In this example we describe the points of $\mathbb{A}^1_{\mathbb{K}}$ for various fields $\mathbb{K}$. Recall that for any field $\mathbb{K}$ the ring $\mathbb{K}[x]$ is a PID. Thus there is a bijection between the maximal ideals of $\mathbb{K}[x]$ and the monic irreducible polynomials in $\mathbb{K}[x]$.

(1) The only irreducible polynomials in $\mathbb{C}[x]$ are linear. Thus every point of $\mathbb{A}^1_{\mathbb{C}}$ is a traditional point; $\mathbb{A}^1_{\mathbb{C}}$ can be identified with $\mathbb{C}$ via the correspondence $(x-a) \leftrightarrow a$.

(2) The irreducible polynomials in $\mathbb{R}[x]$ are either linear or quadratic with negative discriminant. Thus $\mathbb{A}^1_{\mathbb{R}}$ has two types of points: traditional points of the form $(x-a)$ and non-traditional points of the form (x^2+bx+c) where $b^2-4c<0$.

(3) The number of monic irreducible polynomials in $\mathbb{F}_q[x]$ of degree n can be counted using the inclusion-exclusion principle: there are $\frac{1}{n}\sum_{d|n}\mu(n/d)q^d$ of them. This is the number of points of $\mathbb{A}^1_{\mathbb{F}_q}$ which have residue field $\mathbb{F}_{q^n}$.

(4) The set of monic irreducible polynomials in $\mathbb{Q}[x]$ is incredibly complicated; it encodes the entire richness of Galois theory! Suppose that $\mathbb{L}$ is a Galois extension of $\mathbb{Q}$. Then each point of $\mathbb{A}^2_{\mathbb{Q}}$ with residue field $\mathbb{L}$ defines a surjection $\mathbb{Q}[x] \to \mathbb{L}$. Furthermore two surjections $f, g : \mathbb{Q}[x] \to \mathbb{L}$ have the same kernel if and only if the elements $f(x)$ and $g(x)$ in $\mathbb{L}$ have the same minimal polynomial, or equivalently, lie in the same $\mathrm{Gal}(\mathbb{L}/\mathbb{Q})$-orbit. Thus for any Galois extension the points of $\mathbb{A}^1_{\mathbb{Q}}$ with residue field $\mathbb{L}$ are in bijection with Galois orbits of primitive elements in $\mathbb{L}$.

1.1.1 Vanishing loci

As mentioned in the introduction, a classical algebraic set is a subset of $\mathbb{K}^n$ defined by polynomial equations. It is traditional to put all the terms of the equation on one side so that each of our equations has the form $f = 0$. The resulting subset of $\mathbb{K}^n$ is known as the vanishing locus of the set of polynomials $\{f\}$.

In our setting, the analogous construction is as follows:

Definition 1.1.4. Let $J \subset \mathbb{K}[x_1, \ldots, x_n]$ be a set of polynomials. The vanishing locus of J is

$$V(J) := \{\mathfrak{m} \in \mathbb{A}^n_{\mathbb{K}} \,|\, \mathfrak{m} \supset J\}.$$

Exercise 1.1.5. Show that if J is any set of polynomials and I is the ideal generated by J then $V(J) = V(I)$. Thus we may (and henceforth will) assume that the set J is an ideal.

The following exercise describes why this construction deserves the appellation "vanishing locus".

Exercise 1.1.6. Let $I \subset \mathbb{K}[x_1, \ldots, x_n]$ be an ideal. Construct a bijection between $V(I)$ and the set of points $\mathfrak{m} \in \mathbb{A}^n_{\mathbb{K}}$ such that every function in I evaluates to 0 at $\mathfrak{m}$.

It is important to note that different ideals can have the same vanishing set. More precisely:

Proposition 1.1.7. *Let I and J be two ideals in $\mathbb{K}[x_1, \ldots, x_n]$. Then $V(I) = V(J)$ if and only if $\sqrt{I} = \sqrt{J}$.*

Proof. In a Jacobson ring the radical of an ideal I is the same as the intersection of the maximal ideals which contain I. Thus the information encoded by $V(I)$ is the same as the information encoded by $\sqrt{I}$. □

This is the first instance where we see the value of working with non-traditional points. If we did not include non-traditional points, then the forward implication of Proposition 1.1.7 would fail.

Exercise 1.1.8. Give an example of a non-algebraically closed field $\mathbb{K}$ and two ideals $I, J \subset \mathbb{K}[x_1, \ldots, x_n]$ such that $V(I) \cap \mathbb{K}^n = V(J) \cap \mathbb{K}^n$ but $\sqrt{I} \neq \sqrt{J}$.

A key property of V is that it is "inclusion reversing": large ideals in $\mathbb{K}[x_1, \ldots, x_n]$ are identified with small subsets of $\mathbb{A}^n_{\mathbb{K}}$ and vice versa.

Exercise 1.1.9.
(1) Prove that if I and J are ideals in $\mathbb{K}[x_1, \ldots, x_n]$ satisfying $I \subset J$ then $V(I) \supset V(J)$.
(2) Prove that if I and J are ideals in $\mathbb{K}[x_1, \ldots, x_n]$ such that $V(I) \subset V(J)$ then $\sqrt{I} \supset \sqrt{J}$.

1.1.2 Base change

When $\mathbb{K}$ is not algebraically closed the non-traditional points of $\mathbb{A}^n_{\mathbb{K}}$ can look quite complicated. How do mathematicians think of these points in practice? Suppose that $\mathbb{K} \subset \mathbb{L}$ is a Galois field extension. Then $\mathrm{Gal}(\mathbb{L}/\mathbb{K})$ admits an action on $\mathbb{L}[x_1, \ldots, x_n]$ by acting on the coefficients. Every point of $\mathbb{A}^n_{\mathbb{K}}$ corresponds to a $\mathrm{Gal}(\mathbb{L}/\mathbb{K})$-orbit of points of $\mathbb{A}^n_{\mathbb{L}}$ via the prescription

$$\mathfrak{m} \in \mathbb{A}^n_{\mathbb{K}} \leftrightarrow V(\mathfrak{m}) \in \mathbb{A}^n_{\mathbb{L}}$$

where we think of $\mathfrak{m}$ as an ideal in $\mathbb{L}[x_1, \ldots, x_n]$ via the inclusion $\mathbb{K}[x_1, \ldots, x_n] \subset \mathbb{L}[x_1, \ldots, x_n]$. In particular points on $\mathbb{A}^n_{\mathbb{K}}$ are the same as $\mathrm{Gal}(\overline{\mathbb{K}}/\mathbb{K})$-orbits of traditional points on $\mathbb{A}^n_{\overline{\mathbb{K}}}$.

Example 1.1.10. Let's try to understand the points of $\mathbb{A}^2_{\mathbb{R}}$. Using base change, we can think of points of $\mathbb{A}^2_{\mathbb{R}}$ as orbits of points of $\mathbb{A}^2_{\mathbb{C}}$ under the conjugation action.

Concretely, this means the following. Suppose that $\mathfrak{m} \in \mathbb{A}^2_{\mathbb{R}}$ is a non-traditional point. Consider the corresponding conjugate pair of points in $\mathbb{A}^2_{\mathbb{C}}$. The line between them will be invariant under conjugation, giving us a linear equation ℓ with real coefficients that vanishes along both points. Then the maximal ideal $\mathfrak{m} \in \mathbb{A}^2_{\mathbb{R}}$ will have the form (ℓ, q) where q is a quadratic whose restriction to the line $V(\ell)$ has negative discriminant.

For example, suppose that $\mathfrak{m} \in \mathbb{A}^2_{\mathbb{R}}$ corresponds to the Galois orbit of the point $(x - i, y - i) \in \mathbb{A}^2_{\mathbb{C}}$. How can we identify $\mathfrak{m}$ explicitly? In algebraic terms, we are looking for the preimage of $(x - i, y - i)$ under the inclusion $\mathbb{R}[x, y] \to \mathbb{C}[x, y]$. Equivalently, since the conjugate point is $(x + i, y + i)$ we would like to find $\mathbb{R}$-generators for the intersection $(x - i, y - i) \cap (x + i, y + i) \subset \mathbb{C}[x, y]$. Since the two maximal ideals are coprime, their intersection is the same as their product:

$$(x^2 + 1, xy + ix - iy + 1, xy - ix + iy + 1, y^2 + 1).$$

Taking the difference of the two middle terms we find the $\mathbb{R}$-generators $(x - y, x^2 + 1)$. Note that $(x - y)$ is the equation of the line containing the pair of points (i, i) and $(-i, -i)$ in $\mathbb{C}^2$.

1.1.3 Exercises

Exercise 1.1.11. Let $\mathfrak{m} \in \mathbb{A}^2_{\mathbb{Q}}$ be a point whose residue field $\mathbb{L}$ is Galois over $\mathbb{Q}$ of degree n. Assume that every point in the corresponding Galois orbit of points in $\mathbb{A}^2_{\mathbb{L}}$ lies on the same line. Show that $\mathfrak{m}$ is generated by one linear equation and one irreducible degree n equation. (In particular this applies to every $\mathfrak{m} \in \mathbb{A}^2$ whose residue field has degree 2 over $\mathbb{Q}$.)

Exercise 1.1.12. Let ξ_7 denote a primitive 7th root of unity and let $\sigma = \xi_7 + \xi_7^{-1}$. The minimal polynomial of σ is $x^3 + x^2 - 2x - 1$. $\mathbb{Q}(\sigma)$ is a degree 3 Galois extension of $\mathbb{Q}$ and the Galois action sends $\sigma \mapsto (\sigma^2 - 2) \mapsto (-\sigma^2 - \sigma + 1)$. Compute the point $\mathfrak{m} \in \mathbb{A}^2_{\mathbb{Q}}$ corresponding to:

(1) The Galois orbit of $(x - \sigma, y - \sigma)$ in $\mathbb{A}^2_{\mathbb{Q}(\sigma)}$.

(2) The Galois orbit of $(x - \sigma, y - \sigma^2 + 2)$ in $\mathbb{A}^2_{\mathbb{Q}(\sigma)}$.

Exercise 1.1.13. How many points of $\mathbb{A}^2_{\mathbb{F}_q}$ have residue field $\mathbb{F}_{q^n}$?

Exercise 1.1.14. Suppose that $\mathbb{K}$ is an infinite field. Prove that if $X \subset \mathbb{A}^n_{\mathbb{K}}$ is an closed subset that contains every traditional point then $X = \mathbb{A}^n_{\mathbb{K}}$. In contrast, for the finite field $\mathbb{F}_q$ identify a proper closed subset of $\mathbb{A}^n_{\mathbb{F}_q}$ that contains every traditional point.

Exercise 1.1.15. Which ideal defines the union of the three coordinate axes in $\mathbb{A}^3$?

Exercise 1.1.16. The best way to visualize a closed subset of $\mathbb{A}^2$ or $\mathbb{A}^3$ defined by polynomials with integer coefficients is usually to sketch the points in $\mathbb{R}^2$ or $\mathbb{R}^3$ points where these polynomials vanish. (Of course, such sketches must be taken with a grain of salt.) Quickly sketch the following curves:

(1) $y^2 = x^3 - x$ (an elliptic curve).

(2) $y^2 = x^3 + x^2$ (a nodal cubic).

(3) $y^2 = x^3$ (a cuspidal cubic).

Which points of the curve look different than the others? How is this reflected in the algebra of the defining equation?

1.2 Spectrum of a ring

Suppose R is a finitely generated $\mathbb{K}$-algebra. We would like to construct a geometric set associated to R. Our discussion from last lecture suggests that we should be looking at the maximal ideals $\mathfrak{m} \subset R$ such that $R/\mathfrak{m} \cong \mathbb{K}$.

Definition 1.2.1. Let R be a finitely generated $\mathbb{K}$-algebra. The max-spectrum of R is

$$\mathrm{mSpec}(R) = \{\mathfrak{m} \subset R \,|\, \mathfrak{m} \text{ is a maximal ideal}\}.$$

Just as with $\mathbb{A}^n_{\mathbb{K}}$, max-spectrums have both traditional points $\mathfrak{m}$ such that $R/\mathfrak{m} \cong \mathbb{K}$ and non-traditional points $\mathfrak{m}$ such that $R/\mathfrak{m} \ncong \mathbb{K}$.

If we realize R via a surjection $\mathbb{K}[x_1, \ldots, x_n] \to R$ with kernel I, then as a set $\mathrm{mSpec}(R)$ can be identified with the vanishing locus $V(I) \subset \mathbb{A}^n_{\mathbb{K}}$. Thus, set-theoretically max-spectrums are no different from the vanishing loci we saw in the previous chapter. However, it will be useful to have an "embedding-free" way of working with these sets.

Remark 1.2.2. You may have noticed that it somewhat artificial to work with the maximal ideals instead of the set of all prime ideals. In this semester we will only work with Jacobson rings, so we don't lose any information by focusing only on maximal ideals. But in the future when we want to expand our theory to include non-Jacobson rings it will be absolutely essential to work with all prime ideals. The (non-max) spectrum of a ring $\mathrm{Spec}(R)$ is defined to be the set of prime ideals in R.

1.2.1 Zariski topology

We will construct a topology on $\mathrm{mSpec}(R)$ using vanishing loci. The definition is the same as for affine space:

Definition 1.2.3. Let R be a finitely generated $\mathbb{K}$-algebra. Given any subset $J \subset R$, we define the vanishing locus $V(J) \subset \mathrm{mSpec}(R)$ as the set of points $\mathfrak{m} \in \mathrm{mSpec}(R)$ such that $\mathfrak{m} \supset J$. (Note that $V(J)$ is the same as vanishing locus of the ideal $\langle J \rangle$ generated by J, so we will usually just work with ideals.)

Vanishing loci in $\mathrm{mSpec}(R)$ will continue to satisfy the properties discussed in Section 1.1.1 (see Exercise 1.2.17). In this section it will be more important to understand how set-theoretic operations on vanishing loci correspond to algebraic operations on ideals.

Proposition 1.2.4. *Let R be a finitely generated $\mathbb{K}$-algebra.*

(1) Given two ideals I, J we have

$$V(I) \cup V(J) = V(I \cap J) = V(IJ).$$

(2) Given any collection of ideals $\{I_\alpha\}_{\alpha\in\mathcal{A}}$ we have

$$\bigcap_{\alpha\in\mathcal{A}} V(I_\alpha) = V\left(\sum_{\alpha\in\mathcal{A}} I_\alpha\right).$$

Proof. (1) It is clear that $V(I)\cup V(J)\subset V(I\cap J)\subset V(IJ)$ so we only need to show that $V(IJ)\subset V(I)\cup V(J)$. Suppose that $\mathfrak{m}$ is a maximal ideal containing IJ. Let $\{f_i\}$ be a set of generators for I and let $\{g_j\}$ be a set of generators for J. Then IJ is generated by the products $\{f_ig_j\}$. Suppose that $\mathfrak{m}\not\supset I$. Then there is some fixed generator f_i not contained in $\mathfrak{m}$. Since for every j the product f_ig_j is in $\mathfrak{m}$, by primality we see that $\mathfrak{m}$ contains all the generators of J. We conclude that for any $\mathfrak{m}\in V(IJ)$ either $I\subset\mathfrak{m}$ or $J\subset\mathfrak{m}$.

(2) A maximal ideal $\mathfrak{m}$ will contain $\sum I_\alpha$ if and only if it contains each individual I_α. The statement follows. $\square$

Note furthermore that $V(0)=\mathrm{mSpec}(R)$ and $V(R)=\emptyset$. We conclude:

Theorem 1.2.5. *Let R be a finitely generated $\mathbb{K}$-algebra. Sets of the form $V(I)$ form the closed sets in a topology on* $\mathrm{mSpec}(R)$*.*

Definition 1.2.6. The topology on $\mathrm{mSpec}(R)$ defined above is called the Zariski topology. We will denote the Zariski topology by $\mathbf{Zar}_{\mathrm{mSpec}(R)}$. Unless otherwise specified, we will *always* assume it is the underlying topology for our space.

1.2.2 Examples

Let's work out what the Zariski topology looks like in a few easy examples.

Example 1.2.7. Since $\mathbb{K}[x]$ is a PID, every closed subset of $\mathbb{A}^1_{\mathbb{K}}$ is the vanishing locus of a single polynomial f. Note that the vanishing locus of a polynomial f is the finite set of points generated by the monic irreducible factors of f. Conversely, any finite subset of $\mathbb{A}^1_{\mathbb{K}}$ is the vanishing locus for the polynomial obtained by multiplying the generators of the corresponding ideals. Thus the Zariski topology on $\mathbb{A}^1$ is the cofinite topology – the non-trivial closed sets are exactly the finite subsets of $\mathbb{A}^1$.

Example 1.2.8. There are two basic types of closed subset of $\mathbb{A}^2_{\mathbb{K}}$: points (which are the vanishing loci of maximal ideals) and hypersurfaces (which are defined to be closed subsets of the form $V(f)$ for a single polynomial f). We will later see that the non-trivial closed subsets of $\mathbb{A}^2$ are finite unions of sets of this type.

Example 1.2.9. Consider the ring $R=\mathbb{K}[x,y]/(y-x^2)$. According to our earlier discussion we can think of $\mathrm{mSpec}(R)$ as the locus in $\mathbb{A}^2_{\mathbb{K}}$ defined by the equation $y=x^2$. Don't forget that $\mathrm{mSpec}(R)$ can have non-traditional points – for example, if $\mathbb{K}=\mathbb{R}$ then in addition to the traditional points $(x-a,y-a^2)$ we will have non-traditional points of the form $(x^2+bx+c,y+bx+c)$.

In fact, there will be a bijection between points in mSpec(R) and points of $\mathbb{A}^1_{\mathbb{K}}$ since the rings R and $\mathbb{K}[x]$ are isomorphic. Nevertheless it is worth practicing how to identify the points of mSpec(R) explicitly using the identification as a parabola.

1.2.3 Functions

We will think of the ring R as the space of functions on mSpec(R). As discussed earlier, we can "evaluate" a function $f \in R$ along any point $\mathfrak{m}$ by taking the image of f under the quotient map $R \to R/\mathfrak{m}$. However, we emphasize that a function is *not* determined by its evaluations! For example, let $R = \mathbb{K}[x]/(x^2)$. Then mSpec(R) has a unique point. The function $x \in R$ evaluates to 0 at this point, but it is not the zero function in R. We will discuss this issue in more depth in Section 1.4.

We now define the first fundamental geometric object of the course.

Definition 1.2.10. An affine $\mathbb{K}$-scheme consists of a finitely generated $\mathbb{K}$-algebra R equipped with the set mSpec(R) and the Zariski topology:

$$(\text{set, topology, functions}) = (\mathrm{mSpec}(R), \mathbf{Zar}_{\mathrm{mSpec}(R)}, R).$$

We call $\mathbb{K}$ the ground field; we will often omit it from the notation.

There is of course some redundancy in our definition of an affine scheme – one can recover the set and the topology from the space of functions. Nevertheless it is helpful conceptually to think of the functions as a "separate" feature of an affine scheme.

Remark 1.2.11. As discussed in Remark 1.2.2 the actual definition of a scheme involves all the prime ideals in R, not just the maximal ideals. Nevertheless we will continue to use this mild abuse of notation throughout the semester.

1.2.4 Exercises

Exercise 1.2.12. Let $Y \subset \mathrm{mSpec}(R)$ be any subset. Prove that the closure of Y is the vanishing locus of the ideal I of functions which vanish at every point of Y.

Exercise 1.2.13. In contrast to traditional geometry, every open set in the Zariski topology is quite "large." For example, suppose that R is a finitely generated $\mathbb{K}$-algebra which is a domain.

(1) Show that any two non-empty open subsets of mSpec(R) will have non-empty intersection.

(2) Show that every non-empty open subset of mSpec(R) is dense.

Exercise 1.2.14. Consider $\mathrm{mSpec}(\mathbb{R}[x,y]/(x^2+y^2+1))$. Prove that there is a bijection between the points of this affine scheme and the set of equivalence classes of non-constant linear functions ℓ where we set $\ell_1 \sim \ell_2$ if the two equations are the same up to rescaling. (Hint: by comparing to Example 1.1.10 first show that a point $\mathfrak{m}$ will contain a non-constant linear function ℓ. Prove that the quotient $\mathbb{R}[x,y]/(x^2+y^2+1,\ell)$ is a field.)

Exercise 1.2.15. Let R be a finitely generated $\mathbb{K}$-algebra. Show that $\mathrm{mSpec}(R)$ is a finite set if and only if R is an Artinian ring. (Depending on how much you are willing to assume from the theory of Artinian rings this exercise may be trivial.)

Exercise 1.2.16. Fix positive integers m and n. We can identify the traditional points of $\mathbb{A}^{mn}$ with the $(m \times n)$-matrices with values in $\mathbb{K}$. Prove that for any integer r there is a closed subset Z_r whose traditional points are the matrices with rank $\leq r$.

Exercise 1.2.17. Let R be a finitely generated $\mathbb{K}$-algebra with ideals I, J. Prove the following statements:

(1) $V(I) = V(J)$ if and only if $\sqrt{I} = \sqrt{J}$.

(2) If $I \subset J$ then $V(I) \supset V(J)$.

(3) If $V(I) \subset V(J)$ then $\sqrt{I} \supset \sqrt{J}$.

1.3 Properties of the Zariski topology

In this chapter we begin exploring some of the oddities of the Zariski topology. You may have noticed in the previous chapter that the Zariski topology is quite different from the topologies we usually use in geometry. While it takes some time to adjust to the new situation, we will soon see that it comes with many benefits.

In a typical geometric situation, the two most important properties of a topology are the Hausdorff property and compactness. However, these properties are uninteresting for the Zariski topology:

Exercise 1.3.1. Prove that the Zariski topology on $\mathrm{mSpec}(R)$ is Hausdorff if and only if $\mathrm{mSpec}(R)$ is a finite set. (Hint: if R is not Artinian then it contains a prime ideal $\mathfrak{p}$ which is not maximal. Show that it is impossible to separate two maximal ideals containing $\mathfrak{p}$.)

Exercise 1.3.2. Prove that the Zariski topology on $\mathrm{mSpec}(R)$ is always compact.

(Algebraic geometers refer to this fact by saying that the Zariski topology on $\mathrm{mSpec}(R)$ is "quasi-compact." It is admittedly slightly perverse to introduce this entirely new notation which means exactly the same thing as the old notation. Nevertheless, the change in notation helps emphasize that the compactness of the Zariski topology has nothing to do with "compact geometric spaces.")

Later on we will discuss the algebro-geometric analogues of these two important properties. For now we will focus on some new features of the Zariski topology.

1.3.1 Noetherian property

Recall that every finitely generated $\mathbb{K}$-algebra is Noetherian. This translates into a topological property.

Definition 1.3.3. We say that a topological space X is Noetherian if it satisfies the DCC for closed sets: any decreasing chain of closed subsets

$$Z_1 \supset Z_2 \supset Z_3 \supset \dots$$

eventually stabilizes, i.e. there is some index N such that $Z_i = Z_{i+1}$ for every $i \geq N$.

Exercise 1.3.4. Show that if R is a finitely generated $\mathbb{K}$-algebra then $\mathrm{mSpec}(R)$ is a Noetherian topological space.

The Noetherian property is frequently useful for induction arguments. Suppose you want to prove that some property P holds for affine schemes. Assume that you can prove:

(1) If P holds for every proper closed subset of $\mathrm{mSpec}(R)$ then P holds for $\mathrm{mSpec}(R)$.

(2) (Base case:) If R is an Artinian ring then P holds for $\mathrm{mSpec}(R)$.

Then you can conclude that P holds for all affine schemes. The point is that the DCC property guarantees that the inductive argument will stop after a finite sequence of steps.

1.3.2 Irreducible sets

The following definition identifies a vital feature of the Zariski topology.

Definition 1.3.5. We say that a closed subset X of a topological space is irreducible if it satisfies the following property. Suppose that $C_1, C_2 \subset X$ are closed subsets such that $C_1 \cup C_2 = X$. Then either $C_1 = X$ or $C_2 = X$.

For most topologies you will have worked with before this notion is not useful. (See Exercise 1.3.18.) However, it is the most fundamental topological property when working with the Zariski topology. Its importance is clarified by the following result.

Lemma 1.3.6. *Let R be a finitely generated $\mathbb{K}$-algebra. A closed subset X of* $\mathrm{mSpec}(R)$ *is irreducible if and only if there is a prime ideal $\mathfrak{p}$ with $V(\mathfrak{p}) = X$.*

We emphasize that there will usually be other (non-prime) ideals whose vanishing locus is also X. However, the lemma shows that the unique radical ideal whose vanishing locus is X will be a prime ideal.

Proof. Since X is closed, it is the vanishing locus of some ideal I and we may suppose that I is radical.

First suppose that I is prime. Suppose that C_1, C_2 are closed subsets of $\mathrm{mSpec}(R)$ satisfying $C_1, C_2 \subset X$ and $X = C_1 \cup C_2$. Let J_1, J_2 be ideals defining C_1, C_2. By Proposition 1.2.4 we have $I \supset J_1 J_2$. Since I is prime we conclude that either $J_1 \subset I$ or $J_2 \subset I$. This means that either $X \subset C_1$ or $X \subset C_2$, concluding the proof.

Conversely suppose that I is not prime. Recall that any radical ideal is the intersection of the prime ideals containing it. Since R is Noetherian there is a finite list $\mathfrak{p}_1, \ldots, \mathfrak{p}_r$ of such primes. Let $C_i = V(\mathfrak{p}_i)$. Then $C_1 \cup \ldots \cup C_r = X$ but X is not contained in any C_i. $\square$

Example 1.3.7. We verify that the irreducible closed subsets of $\mathbb{A}^2$ are points and hypersurfaces of the form $V(f)$ where $f \in \mathbb{K}[x, y]$ is an irreducible polynomial. This will follow from:

Claim 1.3.8. Let $f, g \in \mathbb{K}[x, y]$ be polynomials with no common factors. Then $V(f, g)$ consists of a finite set of points.

Proof of claim: Since f, g have no common factors in $\mathbb{K}[x, y]$, they also have no common factors in $\mathbb{K}(x)[y]$. This latter ring is a PID, so there are elements $u, v \in \mathbb{K}(x)$ such that $uf + vg = 1$. Clearing denominators, we find a polynomial $h_x \in \mathbb{K}[x]$ such that there is an equality $u'f + v'g = h_x$ in $\mathbb{K}[x, y]$. In particular, we see that $V(f, g) \subset V(h_x)$. Arguing symmetrically, we find a polynomial $h_y \in \mathbb{K}[y]$ such that $V(f, g) \subset V(h_y)$. Finally, we claim that $V(h_x) \cap V(h_y)$ is a finite set. Indeed, the quotient $\mathbb{K}[x, y]/(h_x, h_y) = \mathbb{K}[x]/(h_x) \times \mathbb{K}[y]/(h_y)$ is a product of Artinian rings and thus has only finitely many maximal ideals. $\square$

Suppose now that $Z \subset \mathbb{A}^2$ is any irreducible closed set. If Z is the vanishing locus of a single polynomial f, then f must be irreducible. Otherwise, the claim shows that Z must be a finite set, hence a single point.

The key theorem governing the Zariski topology is:

Theorem 1.3.9. *Let R be a finitely generated $\mathbb{K}$-algebra and let X be a closed subset of* $\mathrm{mSpec}(R)$*. Then:*

(1) X is the union of a finite set $\{X_i\}_{i=1}^r$ of irreducible closed subsets X_i.

(2) The decomposition $X = \cup_{i=1}^r X_i$ of X into irreducible closed subsets is unique (in the sense that if we assume that for all indices $i \neq j$ we have $X_i \not\subset X_j$ then the decomposition is unique up to reordering).

We call the X_i the irreducible components of X.

This is a geometric application of the algebraic theory of primary decompositions – the irreducible components X_i will be in bijection with the minimal prime ideals lying over the ideal I defining X. In the exercises you will have the opportunity to give a purely topological proof.

Proof. Let I be a radical ideal whose vanishing locus is X. From the theory of primary decompositions we know that there is a finite set of prime ideals $\{\mathfrak{p}_i\}$ that are minimal over I (that is, which contain I and are minimal amongst all such primes with respect to inclusion). Since a radical ideal is the intersection of all the primes that contain it, we see that $\cap_i \mathfrak{p}_i = I$. Set $X_i = V(\mathfrak{p}_i)$. We know that X_i is irreducible by Lemma 1.3.6 and we have already shown that $X = \cup_i X_i$. Finally, the uniqueness of the X_i follows from the uniqueness of the associated primes in the theory of primary decomposition. □

Example 1.3.10. A prototypical example of Theorem 1.3.9 is the vanishing locus of the ideal (xy) in $\mathbb{A}^2$. The two minimal primes over (xy) are the ideals (x) and (y). Correspondingly the reducible set $V(xy)$ is the union of two irreducible components $V(x)$ and $V(y)$. This algebra is reflecting the fact that $V(xy)$ is the union of the two coordinate axes.

1.3.3 Exercises

Exercise 1.3.11. Find the irreducible components of $V(x - yz, xz - y^2)$ in $\mathbb{A}^3_{\mathbb{C}}$.

Exercise 1.3.12. Find the irreducible components of $V(x^2 + y^2 - 1, x^2 - z^2 - 1)$ in $\mathbb{A}^3_{\mathbb{C}}$.

Exercise 1.3.13. The notion of connectedness is still a well-behaved notion for affine schemes.

Let R be a finitely generated $\mathbb{K}$-algebra. Prove that the following are equivalent:

(1) $\mathrm{mSpec}(R)$ is disconnected.

(2) There exist non-zero idempotents $e_1, e_2 \in R$ such that $e_1 e_2 = 0$ and $e_1 + e_2 = 1$.

(3) R is isomorphic to a direct product of $\mathbb{K}$-algebras.

We can think of the idempotents e_1, e_2 as indicator functions for unions of connected components of X.

Exercise 1.3.14. Prove that every connected component of $\mathrm{mSpec}(R)$ is a finite union of irreducible components of $\mathrm{mSpec}(R)$.

Exercise 1.3.15. Let X be a Noetherian topological space. Suppose that Z is an irreducible closed subset of X. Prove that there is some irreducible component X_i of X such that $Z \subset X_i$. (Is there a unique such X_i?)

Exercise 1.3.16. Some of the properties of the Zariski topology we discussed in this section will hold more generally for any Noetherian topological space.

(1) Verify that every Noetherian topological space is compact.

(2) Verify the analogue of Theorem 1.3.9: every closed subset can be written as a union of irreducible closed subsets in an essentially unique way. (Hint: let $\mathcal{W}$ denote the set of closed sets which cannot be written as a union of irreducible sets. Show that if $\mathcal{W}$ is non-empty then it has a minimal element. Use this minimal element to derive a contradiction.)

Exercise 1.3.17. Prove that if X is an irreducible topological space then every non-empty open set in X is dense.

Exercise 1.3.18. Let X be a Hausdorff topological space. Prove that X is irreducible if and only if X is a single point.

Exercise 1.3.19. Let $X \subset \mathbb{C}^2$ be a subset defined by the vanishing locus of a set of polynomials. Prove that X is compact in the Euclidean topology if and only if it is a finite set.

(The analogous statement is true in higher dimensions but is a bit harder to prove. As a consequence, if we hope to formulate a good algebro-geometric analogue of the notion of "compactness", we should expect that this notion fails for every affine scheme that is not a finite set. Thus we will delay introducing this analogue until we have more examples of schemes.)

1.4 Zero divisors

We have declared that R is the "space of functions" on $\mathrm{mSpec}(R)$. In this section we discuss the zero divisors in the function ring R.

1.4.1 Nilpotents

As discussed earlier, for any point $\mathfrak{m} \in \mathrm{mSpec}(R)$ we can evaluate a function $f \in R$ at $\mathfrak{m}$ by taking the image under the quotient map $R \to R/\mathfrak{m}$. As we saw there, functions are *not* determined by how they evaluate at points. In this subsection we give a brief account of the functions which are "invisible" with respect to the topology.

Recall that the nilradical $\mathrm{Nil}(R)$ of a ring R is the radical of the 0 ideal. Since every finitely generated $\mathbb{K}$-algebra is a Jacobson ring, the nilradical coincides with the intersection of all maximal ideals in R. This shows that:

Proposition 1.4.1. *Let R be a finitely generated $\mathbb{K}$-algebra and let $f \in R$. Then $V(f) = \mathrm{mSpec}(R)$ if and only if $f \in \mathrm{Nil}(R)$.*

In particular, two functions $f, g \in R$ have the same evaluation at every point of $\mathrm{mSpec}(R)$ if and only if their difference is nilpotent.

Since nilpotent functions are topologically trivial, it may be tempting to expect that they do not play an important role in algebraic geometry. This is definitely not the case! Even when working over an algebraically closed field we can miss out on essential information if we insist on working merely with affine schemes which carry no nilpotent functions.

Example 1.4.2. Consider the set of ideals $(x^2 - a) \in \mathbb{R}[x]$ as we vary $a \in \mathbb{R}$. For $a > 0$, we have $(x^2 - a) = (x - \sqrt{a}) \cap (x + \sqrt{a})$ so the vanishing locus consists of two points, each with residue field $\mathbb{R}$. When $a = 0$, we have the ideal (x^2) whose vanishing locus consists of a single point. Finally, if $a < 0$, the ideal $(x^2 - a)$ defines a single point with residue field $\mathbb{C}$.

Note that in all these situations the quotient $\mathbb{R}[x]/(x^2 - a)$ has dimension 2 as a $\mathbb{R}$-vector space. This reflects the fact that the ideals defining these schemes fit into a nice family (even though the quotients look very different). In particular, the nilpotent structure of $\mathrm{mSpec}(\mathbb{R}[x]/(x^2))$ is "remembering" the fact that we constructed this scheme as a limit of pairs of disjoint points.

While this example is relatively innocuous, the extra information recorded using nilpotents will be vital in future constructions. (See Exercise 1.6.14 for a similar phenomenon in the setting of fibers of morphisms.)

The affine schemes which carry no nilpotent functions have a special name.

Definition 1.4.3. Let R be a finitely generated $\mathbb{K}$-algebra. We say that $\mathrm{mSpec}(R)$ is reduced if $\mathrm{Nil}(R) = 0$.

1.4.2 Varieties

The following is one of the key definitions in the course.

Definition 1.4.4. Let R be a finitely generated $\mathbb{K}$-algebra. We say that $\mathrm{mSpec}(R)$ is an affine variety if R is a domain.

Recall that we have a containment of ideals

$$\begin{tikzcd}[column sep=small]
\mathfrak{p}_1 \arrow[drr, dash] & \mathfrak{p}_2 \arrow[dr, dash] & \cdots & \mathfrak{p}_r \arrow[dl, dash] \\
 & & \mathrm{Nil}(R) = \cap_i \mathfrak{p}_i \arrow[d, dash] & \\
 & & 0 &
\end{tikzcd}$$

where the primes $\mathfrak{p}_i$ are the finite set of minimal primes in R. Thus R will be a domain if and only if $\mathrm{Nil}(R) = 0$ and there is a unique minimal prime in R. In other words:

Corollary 1.4.5. *Let R be a finitely generated $\mathbb{K}$-algebra. Then $\mathrm{mSpec}(R)$ is an affine variety if and only if it is both irreducible and reduced.*

1.4.3 Support

Recall that for an element $f \in R$ the annihilator $\mathrm{Ann}(f)$ is the ideal of functions g such that $fg = 0$.

Definition 1.4.6. Let R be a finitely generated $\mathbb{K}$-algebra. For any $f \in R$ the support of f is the closed subset $\mathrm{Supp}(f) \subset \mathrm{mSpec}(R)$ which is the vanishing locus of $\mathrm{Ann}(f)$.

Loosely speaking, the support of f describes the closed subset where f "lives" – the function f will vanish identically on the complement of $\mathrm{Supp}(f)$. Note that if f is not a zero-divisor then $\mathrm{Supp}(f) = \mathrm{mSpec}(R)$, so the support is only an interesting construction for zero divisors.

Remark 1.4.7. Be careful not to confuse the support $\mathrm{Supp}(f)$ and the vanishing locus $V(f)$. They are very different constructions! The relationship between the two is discussed in Exercise 1.4.13 and Remark 1.11.12.

It turns out that the support is controlled by the associated primes for the zero ideal.

Lemma 1.4.8. *Let R be a finitely generated $\mathbb{K}$-algebra and let $\mathfrak{p}_1, \ldots, \mathfrak{p}_s$ be the associated primes for the zero ideal. For any zero divisor $f \in R$, there is some subset $S \subset \{1, \ldots, s\}$ such that* $\mathrm{Supp}(f) = \cup_{i \in S} V(\mathfrak{p}_i)$.

Conversely, for any associated prime $\mathfrak{p}_i$ there is some zero divisor f such that $\mathrm{Supp}(f) = V(\mathfrak{p}_i)$.

In the proof we will need to use the colon ideal construction. Given an ideal I and an element $f \in R$, we define $(I : f) = \{g \in R | gf \in I\}$. For example $\mathrm{Ann}(f) = (0 : f)$.

Proof. Recall that an associated prime for an ideal I is a prime ideal which annihilates some element in R/I. In particular the last statement is immediate from the definition.

To see the first statement, the theory of primary decompositions shows there are $\mathfrak{p}_i$-primary ideals $\mathfrak{q}_i$ such that $0 = \cap_i \mathfrak{q}_i$. For any $f \in R$ we have

$$\mathrm{Ann}(f) = (0 : f) = (\cap_i \mathfrak{q}_i : f) = \cap_i (\mathfrak{q}_i : f).$$

Thus

$$\sqrt{\mathrm{Ann}(f)} = \cap_i \sqrt{(\mathfrak{q}_i : f)}.$$

Since $\mathfrak{q}_i$ is primary, if $fg \in \mathfrak{q}_i$ then either $f \in \mathfrak{q}_i$ or $g \in \sqrt{\mathfrak{q}_i}$. Thus $(\mathfrak{q}_i : f)$ is either R (if $f \in \mathfrak{q}_i$) or $\sqrt{\mathfrak{q}_i}$ (if $f \notin \mathfrak{q}_i$). We conclude

$$\sqrt{\mathrm{Ann}(f)} = \cap_{f \notin \mathfrak{q}_i} \sqrt{\mathfrak{q}_i} = \cap_{f \notin \mathfrak{q}_i} \mathfrak{p}_i$$

proving the statement. □

In Theorem 1.3.9 we showed that the minimal primes correspond to the irreducible components of $\mathrm{mSpec}(R)$. The embedded (i.e. non-minimal) associated primes will represent closed sets where there is "extra nilpotent structure."

Example 1.4.9. Let $R = \mathbb{K}[x, y]/(xy)$. The associated primes of the zero ideal are (x) and (y). The support of the function x is the y-axis, and the support of the function y is the x-axis.

Example 1.4.10. Let $R = \mathbb{K}[x, y]/(xy, x^2)$. The associated primes of the zero ideal are (x, y) and (x). The minimal prime (x) defines the unique irreducible component of $\mathrm{mSpec}(R)$; it is the support of the function y. The embedded prime (x, y) is the support of the nilpotent function x.

As mentioned above, the complement of $\mathrm{Supp}(f)$ is the largest open set $U \subset \mathrm{mSpec}(R)$ such that the restriction $f|_U$ is identically zero. (Note that being identically zero is stronger than evaluating to zero at every point.) Unfortunately we do not yet have the tools and language to explain this claim – we will revisit it in Lemma 1.11.11 after developing the theory of open sets in affine schemes.

1.4.4 Exercises

Exercise 1.4.11. Let $Y \subset \mathbb{A}^3$ be the vanishing locus of the ideal (yz, xz, y^3, x^2y). Compute a primary decomposition of the ideal. Use this decomposition to identify the irreducible components of Y and all possible supports of functions on Y.

Exercise 1.4.12. Let $\mathrm{mSpec}(R)$ be an affine $\mathbb{K}$-scheme. For any extension $\mathbb{L}/\mathbb{K}$ we define the base change of $\mathrm{mSpec}(R)$ by $\mathbb{L}$ to be the affine $\mathbb{L}$-scheme $\mathrm{mSpec}(R \otimes_{\mathbb{K}} \mathbb{L})$.

(1) Find an example of affine $\mathbb{K}$-scheme $\mathrm{mSpec}(R)$ which is irreducible and a finite extension $\mathbb{L}/\mathbb{K}$ such that the base change of $\mathrm{mSpec}(R)$ by $\mathbb{L}$ is reducible.

(2) Find an example of affine $\mathbb{K}$-scheme $\mathrm{mSpec}(R)$ which is reduced and a finite extension $\mathbb{L}/\mathbb{K}$ such that the base change of $\mathrm{mSpec}(R)$ by $\mathbb{L}$ is not reduced. (Hint: you will need to use a non-perfect field $\mathbb{K}$.)

On the other hand, prove that if $\mathrm{mSpec}(R \otimes_{\mathbb{K}} \mathbb{L})$ is irreducible/reduced then so is $\mathrm{mSpec}(R)$

Exercise 1.4.13. Let R be a finitely generated $\mathbb{K}$-algebra and let $f \in R$. Show that the complement of $\mathrm{Supp}(f)$ is contained in $V(f)$.

Exercise 1.4.14. Let R be a finitely generated $\mathbb{K}$-algebra and let $f \in R$. Prove that $\mathfrak{m} \in \mathrm{Supp}(R)$ if and only if f is in the kernel of the localization map $R \to R_{\mathfrak{m}}$.

Exercise 1.4.15. For any element $a \in \mathbb{K}$ consider the product I_a of the ideals (x, y) and $(x - a, z)$. The vanishing locus of I_a is two skew lines for $a \neq 0$ and two intersecting lines when $a = 0$. Show that I_a is a radical ideal for $a \neq 0$ but that I_0 is not. What is the geometric interpretation of $\sqrt{I_0}/I_0$?

Exercise 1.4.16. Here is another famous example similar in spirit to Example 1.4.2. For any element $a \in \mathbb{K}$ consider the ideal $I \subset \mathbb{K}[x, y, z]$ defined by $(a^2(x+1) - z^2, ax(x+1) - yz, xz - ay, y^2 - x^2(x+1))$.

(1) Show that for any value of a the vanishing locus X_a is an irreducible subset of $\mathbb{A}^3$.

(2) Prove that when $a \neq 0$ then X_a is reduced but that X_0 is not reduced.

Geometrically, the X_a form a family of curves in $\mathbb{A}^3$. When $a \neq 0$ this curve is smooth and spans all of $\mathbb{A}^3$. (We will later recognize these curves as twisted cubics.) However the limit X_0 is a nodal cubic contained in the plane $z = 0$ along with some extra nilpotent structure. The nilpotent structure on X_0 is recording the fact that we obtained it as a "limit" of curves lying outside of the plane.

1.5 Morphisms

In a typical geometry one defines a morphism $f : X \to Y$ to be a function which is continuous/differentiable/holomorphic (depending upon our choice of "structural functions" in our category). One then obtains a pullback map from the structural functions on Y to the structural functions on X by precomposing with f. In fact, one can use this condition to *define* morphisms: a morphism f is any set-theoretic function such that composition with f preserves "structural functions."

In this section we define morphisms of affine $\mathbb{K}$-schemes. As usual, our definition will directly refer to ring of polynomial functions. We will then obtain a continuous set-theoretic map ex post facto.

Definition 1.5.1. Let R and S be affine $\mathbb{K}$-schemes. A $\mathbb{K}$-morphism from $\mathrm{mSpec}(R)$ to $\mathrm{mSpec}(S)$ is a $\mathbb{K}$-algebra homomorphism $f^\sharp : S \to R$.

When the ground field $\mathbb{K}$ is understood we will frequently drop it from the notation.

Warning 1.5.2. It is also possible to define morphisms of affine schemes without reference to a base field. We will not attempt to carry this out since we will only work in the setting of a fixed ground field $\mathbb{K}$ (with the exception of the base change operation; see Exercise 1.5.16).

Our first task is to verify that a morphism naturally defines a continuous topological map.

Proposition 1.5.3. *Let R and S be affine $\mathbb{K}$-schemes. Let $f^\sharp : S \to R$ be a $\mathbb{K}$-algebra homomorphism. Then the function $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ defined by $f(\mathfrak{m}) = (f^\sharp)^{-1}\mathfrak{m}$ is a continuous topological map.*

Remember, $f^\sharp$ is supposed to define the "pullback of functions." This is compatible with our definition of f: the image of $\mathfrak{m}$ should be the vanishing locus of the set of functions whose pullbacks vanish at $\mathfrak{m}$.

Proof. Suppose that $\mathfrak{m}$ is a maximal ideal in R. We first show that $\mathfrak{n} := (f^\sharp)^{-1}(\mathfrak{m})$ is a maximal ideal in S. Since $\mathfrak{m}$ is prime, $\mathfrak{n}$ is also prime. Consider the diagram

$$\begin{tikzcd} S \arrow[r, "f^\sharp"] \arrow[d] & R \arrow[d] \\ S/\mathfrak{n} \arrow[r, "g"] & R/\mathfrak{m} \end{tikzcd}$$

The Nullstellensatz tells us that $R/\mathfrak{m}$ is a finite extension of $\mathbb{K}$. Since g is an injection, $S/\mathfrak{n}$ is a finite $\mathbb{K}$-module, or equivalently, an Artinian $\mathbb{K}$-algebra. Finally, an Artinian $\mathbb{K}$-algebra which is a domain is also a field. We conclude that $\mathfrak{n}$ is a maximal ideal. This defines a set-theoretic map $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$.

We also must show that f is continuous. Let Z be a closed subset of $\mathrm{mSpec}(S)$. Then Z has the form $Z = V(I)$ for some ideal I in S. The preimage of Z will consist of all maximal ideals $\mathfrak{m}$ of R such that $(f^\sharp)^{-1}\mathfrak{m}$ contains I. This is the same as the set of maximal ideals which contain $f^\sharp(I)$. Thus $f^{-1}(Z) = V(\langle f^\sharp(I)\rangle)$ is a closed set. □

Going forward we will denote a morphism from $\mathrm{mSpec}(R)$ to $\mathrm{mSpec}(S)$ by $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$, with the understanding that the "real" data is the ring map $f^\sharp$ implicit in the notation.

It is *absolutely crucial* to feel comfortable translating back and forth between the "algebraic" morphism $f^\sharp$ and the "geometric" morphism f. The following exercise is designed to help you think through this correspondence.

Exercise 1.5.4. Set $\mathbb{A}^n = \mathbb{K}[x_1, \ldots, x_n]$ and $\mathbb{A}^m = \mathbb{K}[y_1, \ldots, y_m]$. We will write traditional points using coordinates – that is, we will write $(a_1, \ldots, a_n) \in \mathbb{K}^n$ for the point $(x_1 - a_1, \ldots, x_n - a_n) \in \mathbb{A}^n$ and similarly will write $(b_1, \ldots, b_m) \in \mathbb{K}^m$.

Suppose we define a function $f : \mathbb{K}^n \to \mathbb{K}^m$ via the prescription

$$f(a_1, \ldots, a_n) = (f_1(a_1, \ldots, a_n), \ldots, f_m(a_1, \ldots, a_n))$$

where each f_j is a polynomial on $\mathbb{K}^n$. Show that the morphism $f : \mathbb{A}^n \to \mathbb{A}^m$ defined by the equation

$$f^\sharp(y_j) = f_j(a_1, \ldots, a_m)$$

will induce the set map f on traditional points.

Conversely, given a ring morphism $f^\sharp : \mathbb{K}[y_1, \ldots, y_m] \to \mathbb{K}[x_1, \ldots, x_n]$ describe the map on traditional points using coordinates.

Of course the correspondence of Exercise 1.5.4 is not limited to affine space; there are many situations where a morphism is defined by its action on traditional points (but see Exercise 1.5.5). In such situations it is common to describe a morphism f via its action on traditional points instead of describing the ring map $f^\sharp$.

Exercise 1.5.5. Show that the Frobenius morphism $f : \mathbb{A}^1_{\mathbb{F}_q} \to \mathbb{A}^1_{\mathbb{F}_q}$ that sends $x \mapsto x^q$ will restrict to the identity map on the set of traditional points but is not the identity as a map of schemes.

1.5.1 Examples

Let's work through a few important examples of morphisms of affine schemes.

Example 1.5.6. For any ideal $I \subset R$, consider the quotient map $f^\sharp : R \to R/I$. This defines a morphism $f : \mathrm{mSpec}(R/I) \to \mathrm{mSpec}(R)$. Note that the image of f is $V(I)$ and in fact f is a homeomorphism from $\mathrm{mSpec}(R/I)$ onto its image.

This example shows that any closed subset $V(I) \subset \mathrm{mSpec}(R)$ can naturally be given the structure of an affine scheme by identifying it with $\mathrm{mSpec}(R/I)$. In fact, a closed set admits many different scheme structures corresponding to the different choices of ideal with the same vanishing locus. By "a closed subscheme of $\mathrm{mSpec}(R)$" we will mean a closed subset equipped with the scheme structure induced by an ideal I.

Example 1.5.7. Suppose we fix a subset $I \subset \{1, \ldots, n\}$ of size k. Consider the inclusion $\mathbb{K}[x_i]_{i\in I} \to \mathbb{K}[x_i]_{i=1}^n$ which sends $x_i \mapsto x_i$. The corresponding map of schemes is the projection $\mathbb{A}^n \to \mathbb{A}^k$ onto the coordinates corresponding to I.

Example 1.5.8. Let R be a finitely generated $\mathbb{K}$-algebra and fix an element $g \in R$. We denote by R_g the localization of R along the set of non-negative powers of g – this is still a finitely generated $\mathbb{K}$-algebra. The localization map $f^\sharp : R \to R_g$ defines a morphism $f : \mathrm{mSpec}(R_g) \to \mathrm{mSpec}(R)$. Recall that $f^\sharp$ induces a bijection between the maximal ideals of R_g and the maximal ideals of R not containing g. In other words, the image of f is the open set in $\mathrm{mSpec}(R)$ which is the complement of $V(g)$ and f is a homeomorphism onto its image.

It is helpful to see the geometry of this map more explicitly. We will focus on a specific example: the localization of $\mathbb{K}[x]$ along the element x. We can identify $\mathbb{K}[x]_x \cong \mathbb{K}[x,y]/(xy-1)$. Note that $\mathrm{mSpec}(\mathbb{K}[x,y]/(xy-1))$ is just the hyperbola $xy = 1$ in $\mathbb{A}^2$. Then the map f arising from localization is the same as the restriction to the hyperbola of the projection map $\mathbb{A}^2 \to \mathbb{A}^1$. The image of this map is the complement of the origin in $\mathbb{A}^1$.

Example 1.5.9. Let $\mathbb{L}$ be a finite extension of the ground field $\mathbb{K}$. A $\mathbb{L}$-point of $\mathrm{mSpec}(R)$ is a morphism $f : \mathrm{mSpec}(\mathbb{L}) \to \mathrm{mSpec}(R)$, or equivalently, a (necessarily surjective) $\mathbb{K}$-algebra homomorphism $R \to \mathbb{L}$.

There is a bijection between $\mathbb{K}$-points of $\mathrm{mSpec}(R)$ and the points with residue field $\mathbb{K}$, and we will use these two terms interchangeably from now on. However, it is important to note that the residue field of the image of a $\mathbb{L}$-point need not be $\mathbb{L}$. Rather, the residue field of the image will be a field $\mathbb{F}$ that satisfies $\mathbb{K} \subset \mathbb{F} \subset \mathbb{L}$.

1.5.2 Exercises

Exercise 1.5.10. Even though a morphism $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ is a continuous set-theoretic function, most continuous set-theoretic functions will fail to be morphisms. For example, show that *every* bijective set-theoretic function $f : \mathbb{A}^1 \to \mathbb{A}^1$ is continuous in the Zariski topology.

Exercise 1.5.11. Suppose $\mathbb{F}_q$ is a finite field and let $\mathbb{L}/\mathbb{F}_q$ be a finite extension. It is a little easier to count $\mathbb{L}$-points than it is points with residue field $\mathbb{L}$. Show that $\mathbb{A}^n_{\mathbb{F}_q}$ has $|\mathbb{L}|^n$ different $\mathbb{L}$-points. (Compare against Example 1.1.3 and Exercise 1.1.13.)

Exercise 1.5.12. Show that the ring map $f^\sharp : \mathbb{K}[x,y]/(y^2-x^3) \to \mathbb{K}[t]$ sending $x \mapsto t^2, y \mapsto t^3$ induces a bijective homeomorphism of affine $\mathbb{K}$-schemes that is not an isomorphism. What is the geometric interpretation of this map?

Exercise 1.5.13. Let $\mathrm{mSpec}(R)$ be an affine $\mathbb{K}$-scheme. The reduced scheme underlying $\mathrm{mSpec}(R)$ is the closed subscheme $\mathrm{mSpec}(R)_{red} := \mathrm{mSpec}(R/\operatorname{Nil}(R))$.

Show that $\mathrm{mSpec}(R)_{red}$ satisfies the following universal property: if $\mathrm{mSpec}(S)$ is any reduced affine $\mathbb{K}$-scheme, then any morphism $f : \mathrm{mSpec}(S) \to \mathrm{mSpec}(R)$ factors through $\mathrm{mSpec}(R)_{red}$.

Example 1.5.14. Suppose that p is a prime number and $q = p^r$. Let $\mathrm{mSpec}(R)$ be an affine $\mathbb{F}_q$-scheme. The function $f^\sharp : R \to R$ defined by $f^\sharp(g) = g^q$ is a $\mathbb{F}_q$-algebra homomorphism, defining a morphism $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(R)$ called the Frobenius morphism. Prove that f is a homeomorphism but need not be an isomorphism.

Exercise 1.5.15.
(1) Suppose that $f : X \to Y$ is a continuous function between topological spaces. Show that if X is irreducible then $f(X)$ is also irreducible.
(2) Suppose that $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ is a morphism. Show that the set-theoretic image of any irreducible component of $\mathrm{mSpec}(R)$ is contained in an irreducible component of $\mathrm{mSpec}(S)$. (See Exercise 1.3.15.)

Exercise 1.5.16. Let $\mathbb{K}$ be a field and let $\mathbb{L}/\mathbb{K}$ be an extension. Suppose that $\mathrm{mSpec}(R)$ is an affine $\mathbb{K}$-scheme. As in Exercise 1.4.12 we define the base change of $\mathrm{mSpec}(R)$ by $\mathbb{L}$ to be the affine $\mathbb{L}$-scheme $\mathrm{mSpec}(R \otimes_{\mathbb{K}} \mathbb{L})$.

Prove that the assignment $\mathfrak{m} \mapsto \mathfrak{m} \cap R$ yields a well-defined continuous function $\mathrm{mSpec}(R \otimes_{\mathbb{K}} \mathbb{L}) \to \mathrm{mSpec}(R)$. Show that the preimage of every $\mathbb{K}$-point in $\mathrm{mSpec}(R)$ consists of a single $\mathbb{L}$-point in $\mathrm{mSpec}(R \otimes_{\mathbb{K}} \mathbb{L})$.

(When $\mathbb{L}/\mathbb{K}$ is a finite extension then this function fits into our framework of morphisms of affine $\mathbb{K}$-schemes via the inclusion $\mathbb{K} \subset R \otimes_{\mathbb{K}} \mathbb{L}$. But when $\mathbb{L}/\mathbb{K}$ is an arbitrary extension then this construction can leave the framework of finitely generated $\mathbb{K}$-algebras.)

Exercise 1.5.17. A conic in $\mathbb{A}^2$ is the vanishing locus of a single quadratic equation. The classification of conics up to isomorphism depends upon the ground field $\mathbb{K}$.

Suppose that the ground field is $\mathbb{C}$. Show that after a linear change of coordinates every conic will be isomorphic to one of the following schemes:

(1) Parabola: $y - x^2 = 0$.
(2) Irreducible non-parabola: $x^2 + y^2 = 1$.
(3) Intersecting lines: $xy = 0$.
(4) Non-intersecting lines: $x(x+1) = 0$.

(5) Double line: $x^2 = 0$.

Furthermore prove that no two of the conics above are isomorphic, so we have really classified all possible types. How does the list change if we work over $\mathbb{R}$ instead?

We will see in Section 3.1 that if we "compactify" these conics inside of projective space $\mathbb{P}^2$ then the classification becomes much simpler!

1.6 Images and preimages

In this section we continue our discussion of morphisms by analyzing images and preimages.

1.6.1 Preimages

Suppose that $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ is a morphism of affine schemes. The set of points $\mathfrak{m} \in \mathrm{mSpec}(R)$ which map to a fixed point $\mathfrak{n} \in \mathrm{mSpec}(S)$ are exactly the maximal ideals in R which contain $f^\sharp(\mathfrak{n})$. Note that this is the same as the set $V(f^\sharp(\mathfrak{n}))$, and thus it is quite natural to give this set the scheme structure coming from the ideal $\langle f^\sharp(\mathfrak{n})\rangle$. More generally:

Definition 1.6.1. Let $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ be a morphism of affine schemes defined by $f^\sharp : S \to R$. Let Z be the closed subscheme of $\mathrm{mSpec}(S)$ defined by the vanishing of an ideal I. We define the preimage $f^{-1}(Z)$ to be the vanishing locus in $\mathrm{mSpec}(R)$ of $f^\sharp(I)$ equipped with the scheme structure $\mathrm{mSpec}(R/\langle f^\sharp(I)\rangle)$.

Example 1.6.2. Recall that the $\mathbb{K}$-algebra homomorphism $f^\sharp : \mathbb{K}[x] \to \mathbb{K}[x, y]$ sending $x \mapsto x$ defines the projection $f : \mathbb{A}^2_{\mathbb{K}} \to \mathbb{A}^1_{\mathbb{K}}$ onto the x-axis. Suppose we fix a point $\mathfrak{m} \in \mathbb{A}^1_{\mathbb{K}}$ corresponding to the monic irreducible polynomial g. Then $f^{-1}(\mathfrak{m})$ is the vanishing locus of $f^\sharp(g)$. If the residue field of $\mathfrak{m}$ is $\mathbb{L}$, then the preimage is defined by the ring $\mathbb{K}[x, y]/(f^\sharp(g)) \cong \mathbb{L}[y]$. We conclude that the fiber over $\mathfrak{m}$ is isomorphic to $\mathbb{A}^1_{\mathbb{L}}$.

Note that a morphism $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ will be set-theoretically injective precisely when every maximal ideal $\mathfrak{m} \subset S$ satisfies the property that $\langle f^\sharp(\mathfrak{m})\rangle$ is also a maximal ideal. We have already seen several examples of such maps in Example 1.5.6, Example 1.5.8, and Exercise 1.5.12. Characterizing the injective morphisms of affine schemes is actually somewhat delicate; see Section 4.2 for a related discussion.

1.6.2 Images

We next discuss the image of a morphism $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$. Unfortunately now the story is not nearly so nice. The set-theoretic image of a morphism can fail to be either open or closed, and in particular may not admit the structure of an affine scheme.

Example 1.6.3. Consider the function $f : \mathbb{A}^2 \to \mathbb{A}^2$ defined by the ring map $f^\sharp : \mathbb{K}[x, y] \to \mathbb{K}[u, v]$ sending $x \mapsto u, y \mapsto uv$. We claim that the set-theoretic image of f is

$$\big(\mathbb{A}^2 \backslash V(x)\big) \cup \{(x, y)\}.$$

In particular the image of f is neither open nor closed. The claim is easy to see on the level of traditional points – the set theoretic map is just $f(a, b) = (a, ab)$ – but for the sake of completeness we will demonstrate the claim using algebra.

Let's analyze the image of a point $\mathfrak{m} \subset \mathbb{K}[u, v]$. First suppose that $\mathfrak{m} \in V(u)$, or in other words, that $\mathfrak{m}$ contains u. Then $(f^\sharp)^{-1}(\mathfrak{m})$ contains both x and y. Since we know the preimage is a maximal ideal the only option is that $(f^\sharp)^{-1}(\mathfrak{m}) = (x, y)$. In other words, the f-image of any point in $V(u)$ will be the origin in $\mathbb{A}^2$. Now suppose that $\mathfrak{m} \notin V(u)$, or in other words, that $\mathfrak{m}$ does not contain u. Then we see that $(f^\sharp)^{-1}(\mathfrak{m})$ also cannot contain x, that is, $(f^\sharp)^{-1}(\mathfrak{m}) \notin V(x)$. Altogether we see that the only point in $V(x) \cap f(\mathbb{A}^2)$ is the origin (and the preimage of the origin is $V(u)$).

Finally we claim that f induces a bijection between the points of $\mathbb{A}^2 \backslash V(u)$ and the points of $\mathbb{A}^2 \backslash V(x)$. By Example 1.5.8 these points are in bijective correspondence with $\mathrm{mSpec}(\mathbb{K}[u, v]_u)$ and $\mathrm{mSpec}(\mathbb{K}[x, y]_x)$ respectively. But $f^\sharp$ induces an isomorphism of these localized rings, implying that f does indeed induce a bijection between the two sets of maximal ideals.

Algebraic geometers are not willing to leave the realm of schemes when discussing images of morphisms. Thus we usually use the following convention:

Definition 1.6.4. Let $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ be a morphism of affine schemes. The scheme-theoretic image of f is the smallest closed subscheme $\mathrm{mSpec}(S/I)$ of $\mathrm{mSpec}(S)$ such that f factors through $\mathrm{mSpec}(S/I)$.

We must be very careful to distinguish the scheme-theoretic image from the set-theoretic image.

Warning 1.6.5. There are instances where we follow a different convention for images. For example, as discussed in Example 1.5.8 localization at an element induces a morphism whose set-theoretic image is an open set and in this case it makes sense to treat this open set as the "image." This issue will reoccur several times throughout the notes.

The following lemma explains how to construct the scheme-theoretic image for a morphism of affine schemes.

Lemma 1.6.6. *Consider a morphism* $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ *defined by* $f^\sharp : S \to R$. *The scheme-theoretic image of* f *is the closed subscheme* $\mathrm{mSpec}(S/\ker(f^\sharp))$ *of* $\mathrm{mSpec}(S)$.

In other words, the scheme-theoretic image is the geometric construction corresponding to the algebraic operation of pullback of ideals.

Proof. Note that $f^\sharp$ factors as $S \to S/\ker(f^\sharp) \to R$, and that $S/\ker(f^\sharp)$ is universal amongst all factorings of f through quotients of S. This yields the desired statement. □

If $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ is surjective then the discussion above shows that $f^\sharp$ must be injective. However, injectivity of $f^\sharp$ is not sufficient for f to be surjective; it only guarantees that the image of $f^\sharp$ is dense (see e.g. Example 1.5.8). This property is useful enough to earn its own definition.

Definition 1.6.7. A morphism $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ is dominant if the set-theoretic image $f(\mathrm{mSpec}(R))$ is dense in $\mathrm{mSpec}(S)$.

1.6.3 Constructible sets

We have seen that the set-theoretic image of a morphism of affine schemes need not be open or closed. Nevertheless, there are some topological constraints on the set-theoretic images of morphisms. The following definition is the key.

Definition 1.6.8. Let X be a topological space. A constructible subset of X is a finite union of locally closed subsets. (A locally closed subset is the intersection of an open subset and a closed subset.)

Note that the set of constructible subsets is closed under the operations of taking finite unions and taking complements. In fact, if we start with the set of all closed sets and repeatedly apply the finite union and complement operations the result will be the set of all constructible subsets.

Exercise 1.6.9. Let X be a topological space and let $S \subset X$ be a constructible subset. Show that S contains a dense open subset of $\overline{S}$.

Theorem 1.6.10 (Chevalley's theorem)**.** *Let $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ be a morphism of affine schemes. The set-theoretic image of any constructible set Z in* $\mathrm{mSpec}(R)$ *will be a constructible set in* $\mathrm{mSpec}(S)$*.*

Proof. The proof is by Noetherian induction on $\mathrm{mSpec}(S)$. The base case – when S is an Artinian ring – is clear. Thus we may assume that the statement holds for every proper closed subscheme of $\mathrm{mSpec}(S)$.

It suffices to prove the statement when Z is a locally closed subset. By replacing $\mathrm{mSpec}(R)$ by $\overline{Z}$ (equipped with its reduced structure) we may suppose that Z is an open set in $\mathrm{mSpec}(R)$. We may also suppose that $\mathrm{mSpec}(R)$ is irreducible. If f fails to be dominant then we can conclude the desired statement by the induction assumption applied to $\overline{f(\mathrm{mSpec}(R))}$. So we may assume f is dominant. We then appeal to the most important special case of the theorem:

Lemma 1.6.11. *Let $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ be a dominant morphism of irreducible affine schemes. Let U be any open subset of* $\mathrm{mSpec}(R)$*. Then the set-theoretic image $f(U)$ contains an open subset of* $\mathrm{mSpec}(S)$*.*

Proof. This result takes a bit of work. We will prove it using dimension theory, so we postpone the proof to Exercise 4.4.12. □

Let $V \subset \mathrm{mSpec}(S)$ be the open subset obtained by applying the previous lemma to Z. To show that $f(Z)$ is constructible, it suffices to show that $f(Z) \cap (\mathrm{mSpec}(S) \backslash V)$ is constructible. This follows from the induction assumption applied to the closed subscheme $\mathrm{mSpec}(S) \backslash V$. □

1.6.4 Exercises

Exercise 1.6.12. Let $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ be a morphism of affine schemes.

(1) Show that if $f^\sharp$ is surjective then f is injective. (What is an example where the converse fails?)

(2) Show that if $f^\sharp$ is injective then f is dominant. (What is an example where the converse fails?)

(3) Show that if the nilradical of S is 0 then $f^\sharp$ is injective if and only if f is dominant.

Exercise 1.6.13. Let $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ be a morphism of affine $\mathbb{K}$-schemes and let $\mathbb{L}$ be a finite extension of $\mathbb{K}$. Prove or disprove the following statements:

(1) If $\mathfrak{m}$ is a point of $\mathrm{mSpec}(R)$ with residue field $\mathbb{L}$, then $f(\mathfrak{m})$ is a point of $\mathrm{mSpec}(S)$ with residue field $\mathbb{L}$.

(2) If $\mathfrak{m}$ is an point of $\mathrm{mSpec}(S)$ with residue field $\mathbb{L}$, then some point in $f^{-1}(\mathfrak{m})$ is a point of $\mathrm{mSpec}(R)$ with residue field $\mathbb{L}$.

Exercise 1.6.14. Consider the parabola X defined by the equation $x^2 - y$ in $\mathbb{A}^2_{\mathbb{K}}$. By projecting onto the y-coordinate we obtain a map $f : X \to \mathbb{A}^1_{\mathbb{K}}$. Show that the fiber of f over a $\mathbb{K}$-point will have one of the following three types:

- A disjoint union of two $\mathbb{K}$-points.
- A single point with residue field $\mathbb{L}$ where $[\mathbb{L} : \mathbb{K}] = 2$.
- A single point with residue field $\mathbb{K}$ whose ring of functions R satisfies $\dim_{\mathbb{K}}(R) = 2$.

How can you determine which of the three possibilities happens at a given point? (Be careful if the characteristic is 2.)

Prove more generally that for any point of $\mathbb{A}^1_{\mathbb{K}}$ with residue field $\mathbb{L}$ the ring of functions of the fiber satisfies $\dim_{\mathbb{L}}(R) = 2$.

Exercise 1.6.15. Set $X = \mathrm{mSpec}(\mathbb{K}[x, y, z]/(xy, xz, yz))$. Describe all the fibers of the map $f : X \to \mathbb{A}^1$ defined by $f^\sharp : \mathbb{K}[t] \to \mathbb{K}[x, y, z]/(xy, xz, yz)$ sending $t \mapsto x + y + z$. What is the geometric interpretation of this map?

Exercise 1.6.16. Set $X = \mathrm{mSpec}(\mathbb{K}[x, y]/(y^2 - x^3 - x^2))$. Describe all the fibers of the map $f : \mathbb{A}^1 \to X$ defined by $f^\sharp : \mathbb{K}[x, y]/(y^2 - x^3 - x^2) \to \mathbb{K}[t]$ sending $x \mapsto t^2 - 1$, $y \mapsto t(t^2 - 1)$.

In this example the fibers are not "continuous" as they were in Exercise 1.6.14. What is the geometric interpretation of this map?

1.7 Category of affine $\mathbb{K}$-schemes

The goal of this section is to study the category of affine $\mathbb{K}$-schemes.

Definition 1.7.1. The category **AffSch**/$\mathbb{K}$ has as objects all affine $\mathbb{K}$-schemes and as morphisms the $\mathbb{K}$-morphisms between affine $\mathbb{K}$-schemes.

From now on we will use "categorical" notation instead of the "algebraic" notation we have used thus far:

- Affine $\mathbb{K}$-schemes will be denoted by letters such as X, Y, Z.
- A $\mathbb{K}$-morphism of affine $\mathbb{K}$-schemes will be written in the form $f : X \to Y$ (which implicitly includes the data of a $\mathbb{K}$-algebra homomorphism $f^\sharp$ from the ring defining Y to the ring defining X).
- If X is an affine $\mathbb{K}$-scheme and $\mathbb{L}/\mathbb{K}$ is an extension of fields, we denote by $X_{\mathbb{L}}$ the affine $\mathbb{L}$-scheme obtained by base change to $\mathbb{L}$ (i.e. the scheme obtained by tensoring the $\mathbb{K}$-algebra defining X by $\mathbb{L}$).

The following observation is an immediate consequence of the definitions.

Theorem 1.7.2. *There is a contravariant equivalence of categories*

$$\mathbf{AffSch}/\mathbb{K} \qquad \leftrightarrow \qquad \left\{ \begin{array}{c} \textit{finitely generated } \mathbb{K}\textit{-algebras} \\ \textit{equipped with } \mathbb{K}\textit{-algebra homomorphisms} \end{array} \right\}.$$

Remark 1.7.3. The category **AffVar**/$\mathbb{K}$ of affine $\mathbb{K}$-varieties is the full subcategory of **AffSch**/$\mathbb{K}$ whose objects are affine varieties. This category admits a contravariant equivalence with the category of finitely generated $\mathbb{K}$-algebra domains equipped with $\mathbb{K}$-algebra homomorphisms.

We can use the equivalence of Theorem 1.7.2 to discuss some categorical constructions for affine schemes. In brief, finite limits exist in **AffSch**/$\mathbb{K}$ but finite colimits need not exist.

1.7.1 Special objects

Just as every $\mathbb{K}$-algebra R admits a unique $\mathbb{K}$-algebra homomorphism $\mathbb{K} \to R$, every affine $\mathbb{K}$-scheme X admits a unique morphism $f : X \to \mathrm{mSpec}(\mathbb{K})$. Thus:

Lemma 1.7.4. $\mathrm{mSpec}(\mathbb{K})$ *is a terminal object in the category of affine* $\mathbb{K}$*-schemes.*

On the other hand, recall that by our convention the 0 ring is a $\mathbb{K}$-algebra. This is a terminal object in the category of $\mathbb{K}$-algebras, thus:

Lemma 1.7.5. $\mathrm{mSpec}(0)$ *is an initial object in the category of affine* $\mathbb{K}$*-schemes.*

1.7.2 Products

Let's recall the categorical definition of a product. Given two objects X, Y, the product of X and Y (if it exists) is an object $X \times Y$ equipped with morphisms $\pi_1 : X \times Y \to X$ and $\pi_2 : X \times Y \to Y$ (called projection morphisms) which satisfy the following universal property: if W is an object equipped with morphisms $\psi_1 : W \to X$ and $\psi_2 : W \to Y$, there is a unique morphism $f \times g : Z \to X \times Y$ making the following diagram commute:

$$\begin{tikzcd}
W \arrow[drr, bend left, "\psi_1"] \arrow[ddr, bend right, "\psi_2"'] \arrow[dr, dotted, "\exists!"] & & \\
& X \times Y \arrow[r, "\pi_1"] \arrow[d, "\pi_2"] & X \\
& Y &
\end{tikzcd}$$

Under our equivalence of categories, the product of $\mathrm{mSpec}(R)$ and $\mathrm{mSpec}(S)$ should be associated with the coproduct of R and S in the category of finitely generated $\mathbb{K}$-algebras. In other words, the product of affine schemes corresponds to the ring operation $\otimes_{\mathbb{K}}$.

Definition 1.7.6. Let $X = \mathrm{mSpec}(R)$ and $Y = \mathrm{mSpec}(S)$ be two affine $\mathbb{K}$-schemes. The product of X and Y is

$$X \times Y := \mathrm{mSpec}(R \otimes_{\mathbb{K}} S).$$

When we want to emphasize the ground field, we will instead write $X \times_{\mathrm{mSpec}(\mathbb{K})} Y$.

Example 1.7.7. We have $\mathrm{mSpec}(R) \times \mathbb{A}^n \cong \mathrm{mSpec}(R[x_1, \ldots, x_n])$.

It is interesting to compare how this construction compares with the product of the underlying sets or topological spaces.

Exercise 1.7.8. Using $\mathrm{mSpec}(\mathbb{C}) \times_{\mathrm{mSpec}(\mathbb{R})} \mathrm{mSpec}(\mathbb{C})$ as an example, show that the set underlying $X \times Y$ need *not* be the product of the sets underlying X and Y.

It turns out that the product and the set-theoretic product agree for traditional points – the traditional points of $X \times Y$ are in fact the product of the traditional points on X and the traditional points on Y. (In fact for any finite field extension $\mathbb{L}/\mathbb{K}$ the $\mathbb{L}$-points of the product $X \times Y$ are the same as the product of the $\mathbb{L}$-points on X and the $\mathbb{L}$-points on Y – this follows immediately from the universal property of the product.) However:

Exercise 1.7.9. Using $\mathbb{A}^n_{\mathbb{C}} \times \mathbb{A}^m_{\mathbb{C}} = \mathbb{A}^{n+m}_{\mathbb{C}}$, show that even when the set underlying $X \times Y$ is the product of the sets underlying X and Y the Zariski topology on $X \times Y$ need *not* be the product of the Zariski topologies on X and Y.

More generally, given morphisms $f : X \to Z$ and $g : Y \to Z$ we can construct the pullback $X \times_Z Y$ which will be universal for diagrams of the shape

$$\begin{tikzcd}
W \arrow[dr, dotted, "\exists!"] \arrow[drr, bend left, "\psi_1"] \arrow[ddr, bend right, "\psi_2"'] & & \\
& X \times_Z Y \arrow[r, "\pi_1"] \arrow[d, "\pi_2"] & X \arrow[d, "f"] \\
& Y \arrow[r, "g"'] & Z
\end{tikzcd}$$

Definition 1.7.10. Suppose $X = \mathrm{mSpec}(R)$, $Y = \mathrm{mSpec}(S)$, $Z = \mathrm{mSpec}(T)$ and we have morphisms $f : X \to Z$ and $g : Y \to Z$. Then the product of X and Y over Z is

$$X \times_Z Y = \mathrm{mSpec}(R \otimes_T S)$$

where R and S are given the structure of T-modules using $f^\sharp$ and $g^\sharp$ respectively.

Remark 1.7.11. Since $\mathrm{mSpec}(\mathbb{K})$ is a terminal object we have $X \times Y \cong X \times_{\mathrm{mSpec}(\mathbb{K})} Y$, justifying our earlier terminology.

In most circumstances it is best to visualize $X \times_Z Y$ via the fibers of the map $\pi_2 : X \times_Z Y \to Y$. Suppose that $y \in Y$ has residue field $\mathbb{L}$. Then

$$\pi_1^{-1}(Y) \cong F_{\mathbb{L}}$$

where F is the fiber of $f : X \to Z$ over $g(y)$. This is a consequence of the following computation for a maximal ideal $\mathfrak{m} \subset S$:

$$(R \otimes_T S)/\langle \mathfrak{m} \rangle \cong R/\langle g^{\sharp,-1}(\mathfrak{m})\rangle \otimes_{\mathbb{K}} (S/\mathfrak{m})$$

Thus, in a loose sense we can construct $X \times_Z Y$ by taking the fibers of f and "pulling them back" to form a fibration over Y. The following exercise gives our first illustration of this principle (see Exercise 1.7.22 for another application).

Exercise 1.7.12. Let $f : X \to Y$ be a morphism of affine schemes. Let $Z \subset Y$ be a closed subscheme. Show that the preimage $f^{-1}(Z)$ is isomorphic to $Z \times_Y X$.

The existence of products and pullbacks implies that equalizers also exist. Given two morphisms $f, g : X \to Y$, their equalizer is the same as the relative product $X \times_{Y \times Y} Y$ where the map $X \to Y \times Y$ is induced by (f, g) and the map $Y \to Y \times Y$ is induced by the identity map on both components. Altogether we conclude:

Theorem 1.7.13. *Finite limits exist in the category of affine $\mathbb{K}$-schemes.*

1.7.3 Colimits

The coproduct of two affine $\mathbb{K}$-schemes is simply the disjoint union. (Check!) However, in general pushouts need not exist. Equivalently, pullbacks do not exist in the category of finitely generated $\mathbb{K}$-algebras.

Example 1.7.14 ([Mon17] Example 1.3)**.** We first construct two finitely generated subrings of $\mathbb{K}[x, y]$ whose intersection is not finitely generated. Set $R = \mathbb{K}[x^2, x^3, y]$ and $R_2 = \mathbb{K}[x^2, y - x]$. Then the intersection $S = R_1 \cap R_2$ is $\mathbb{K}[x^{2a}(y - x)^b]_{(a,b)\in\mathcal{S}}$ where

$$\mathcal{S} = \{(a, b) \in \mathbb{Z}^2_{\geq 0} \mid \text{either } b = 0 \text{ or } a \geq 1\}.$$

Note that $\mathcal{S}$ is not finitely generated over $\mathbb{K}$.

We claim that the two inclusion maps $i_1 : R_1 \to \mathbb{K}[x, y]$ and $i_2 : R_2 \to \mathbb{K}[x, y]$ do not admit a pullback diagram. In the category of all $\mathbb{K}$-algebras, the pullback is given by S equipped with the inclusion maps. If i_1 and i_2 admitted a pullback S' in the category of finitely generated $\mathbb{K}$-algebras, then by the universal property any map from a finitely generated $\mathbb{K}$-algebra T to S would have to factor through S'. This is clearly impossible.

Remark 1.7.15. All finite limits exist in the category of arbitrary rings: for example, the equalizer is simply the subring where two ring homomorphisms coincide. Dually, all finite colimits exist in the category of arbitrary affine schemes. However, it turns out that in the category of arbitrary schemes finite colimits need not exist. This highlights the fact that the pullback is the most useful and important construction for schemes.

1.7.4 Exercises

Exercise 1.7.16. Suppose $X = \mathrm{mSpec}(R)$, $Y = \mathrm{mSpec}(S)$, $Z = \mathrm{mSpec}(T)$ and we have morphisms $f : X \to Z$ and $g : Y \to Z$. Prove that $X \times_Z Y$ is isomorphic to the closed subscheme of $X \times Y$ defined by the ideal generated by the image of T in $R \otimes S$.

Exercise 1.7.17. The product does not interact well with the notions of irreducibility or reducedness.

(1) Show that $\mathrm{mSpec}(\mathbb{C})$ is an irreducible $\mathbb{R}$-scheme but that $\mathrm{mSpec}(\mathbb{C}) \times_{\mathbb{R}} \mathrm{mSpec}(\mathbb{C})$ is not irreducible.

(2) Set $\mathbb{K} = \mathbb{F}_p(u)$. Show that $\mathrm{mSpec}(\mathbb{K}[t]/(t^p - u))$ is a reduced $\mathbb{K}$-scheme but that $\mathrm{mSpec}(\mathbb{K}[t]/(t^p - u)) \times_{\mathbb{K}} \mathrm{mSpec}(\mathbb{K}[t]/(t^p - u))$ is not reduced.

Exercise 1.7.18. Let X be an affine scheme. The diagonal $\Delta \subset X \times X$ is the scheme-theoretic image of the map $X \to X \times X$ induced by the identity map on both components.

Suppose that $X = \mathrm{mSpec}(R)$. First describe the $\mathbb{K}$-algebra homomorphism that yields the diagonal embedding $X \to X \times X$. Then find generators for the ideal of Δ in $R \otimes_{\mathbb{K}} R$.

Exercise 1.7.19. Let $f, g : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ be two morphisms. Show that the equalizer of f, g is the map $\mathrm{mSpec}(R/I) \to \mathrm{mSpec}(R)$ induced by quotienting R by the ideal generated by $\{f^\sharp(s) - g^\sharp(s)\}_{s \in S}$.

Exercise 1.7.20. Let $\mathbb{L}/\mathbb{K}$ be a field extension. Show that base change defines a functor from the category of affine $\mathbb{K}$-schemes to the category of affine $\mathbb{L}$-schemes. (Is this functor full? Is it faithful?)

Exercise 1.7.21. Suppose given a morphism $f : X \to Y$ of affine $\mathbb{K}$-schemes. The identity map $id : X \to X$ and the map $f : X \to Y$ together induce a morphism $id \times f : X \to X \times Y$. The graph Γ of f is defined to be the scheme-theoretic image of $id \times f$.

(1) Suppose that $X = \mathrm{mSpec}(R)$ and $Y = \mathrm{mSpec}(S)$. Write down generators for the ideal of Γ.

(2) Show that Γ is isomorphic to X.

(3) Show that Γ is the preimage of the diagonal $\Delta \subset Y \times Y$ under the map $(f \circ \pi_1) \times \pi_2 : X \times Y \to Y \times Y$.

Exercise 1.7.22. Consider the morphism $f : \mathbb{A}^2 \to \mathbb{A}^2$ defined by the map $f^\sharp : \mathbb{K}[x, y] \to \mathbb{K}[u, v]$ which sends $x \mapsto u, y \mapsto uv$. In Example 1.6.3 we saw that f contracts the u-axis to a point and defines a bijection on the complement U of the u-axis.

Consider the relative product

$$\begin{tikzcd} X \arrow[r, "\pi_1"] \arrow[d, "\pi_2"] & \mathbb{A}^2 \arrow[d, "f"] \\ \mathbb{A}^2 \arrow[r, "f"'] & \mathbb{A}^2 \end{tikzcd}$$

Explicitly compute the ring defining X and the homomorphism defining π_2. Use these computations to prove the following statements:

(1) Show that the restriction of π_2 to the preimage of U is bijective.

(2) Show that the restriction of π_2 to the preimage of the u-axis is a projection map $\mathbb{A}^2 \to \mathbb{A}^1$.

Geometrically, X looks like the union of two planes: one maps isomorphically onto $\mathbb{A}^2$ under π_2 and the other is mapped onto the u-axis via a coordinate projection. Think carefully about how this coheres with the description of the fibers of π_2 given in the text.

Exercise 1.7.23. The following exercise shows that the ring of functions on an affine scheme $\mathrm{mSpec}(R)$ can be identified with the set of morphisms $f : \mathrm{mSpec}(R) \to \mathbb{A}^1$. Let R be a finitely generated $\mathbb{K}$-algebra.

(1) Show the set of morphisms $f : \mathrm{mSpec}(R) \to \mathbb{A}^1$ has the structure of a $\mathbb{K}$-algebra, where:

- We add two morphisms $f, g : \mathrm{mSpec}(R) \to \mathbb{A}^1$ by taking the product map $f \times g : \mathrm{mSpec}(R) \to \mathbb{A}^2$ and composing it with the map $a : \mathbb{A}^2 \to \mathbb{A}^1$ defined by $a^\sharp(t) = x + y$.
- We multiply two morphisms $f, g : \mathrm{mSpec}(R) \to \mathbb{A}^1$ by taking the product map $f \times g : \mathrm{mSpec}(R) \to \mathbb{A}^2$ and composing it with the map $m : \mathbb{A}^2 \to \mathbb{A}^1$ defined by $m^\sharp(t) = xy$.
- The element $a \in \mathbb{K}$ corresponds to the morphism $\mathrm{mSpec}(R) \to \mathrm{mSpec}(\mathbb{K}) \to \mathbb{A}^1$ where the last map is the inclusion of the point $(x - a)$.

(2) Show that the space of functions R on $\mathrm{mSpec}(R)$ is isomorphic to the $\mathbb{K}$-algebra of functions $f : \mathrm{mSpec}(R) \to \mathbb{A}^1$.

1.8 Sheaves

In Example 1.5.6 we showed that every closed subset of an affine scheme naturally admits the structure of an affine scheme. In contrast, it turns out that open subsets may or may not admit the structure of an affine scheme. (This is perhaps surprising – after all, it's easy to see that open subsets of manifolds are still manifolds.)

Our next goal is to put open subsets of affine schemes on an equal footing with affine schemes. The key idea is to associate a ring of functions $\mathcal{O}_X(U)$ to any open subset U of an affine scheme X. However, before we construct the rings of functions $\mathcal{O}_X(U)$, we will introduce a new concept: a sheaf of abelian groups. A sheaf is an object that encodes the function rings for all open subsets U simultaneously. While the study of abstract sheaves is unfortunately a bit of a distraction, the payoff in convenience will more than compensate for the time spent.

Warning 1.8.1. Although sheaves play an important role in algebraic geometry, in Part I we will only ever see one type of sheaf (the structure sheaf on a scheme). For this reason we defer our discussion of a general theory of sheaves to Part II. For now we will only prove exactly what we need for applications to the structure sheaf.

1.8.1 Sheaves

Suppose that X is a topological manifold. For any open subset $U \subset X$, let $\mathcal{C}(U)$ denote the space of continuous functions on U. There are two important ways in which these functions interact as we vary our open set:

(1) Restriction: given an inclusion of open sets $V \subset U$, any function on U also induces a function on V.

(2) Gluing: given an open cover $\{V_i\}$ of U and functions f_i on V_i which agree on the common overlaps, we obtain a (unique) function f on U by gluing the f_i.

The definition of a sheaf formalizes these two important properties to give us a general language for discussing "functions":

Definition 1.8.2. Let X be a topological space. A sheaf $\mathcal{F}$ of abelian groups on X consists of the following data:

(1) for every open subset U, an abelian group $\mathcal{F}(U)$ whose elements are known as "sections of U", and

(2) for every inclusion of non-empty open subsets $V \subset U$, a homomorphism $\rho_{U,V}$: $\mathcal{F}(U) \to \mathcal{F}(V)$ known as a "restriction map"

satisfying the following conditions:

(1) (Normalization) $\mathcal{F}(\emptyset) = 0$.

(2) (Compatibility) The assignment $U \mapsto \mathcal{F}(U)$ and $(V \subset U) \mapsto \rho_{U,V}$ defines a contravariant functor from the category of open subsets of X (with morphisms = inclusions) to the category of abelian groups. In other words, $\rho_{U,U} = id$ and if $W \subset V \subset U$ then $\rho_{U,V} \circ \rho_{V,W} = \rho_{U,W}$.

(3) (Identity) Suppose that $\{V_i\}$ is an open cover of U. Suppose that $f_1, f_2 \in \mathcal{F}(U)$ satisfy $\rho_{U,V_i}(f_1) = \rho_{U,V_i}(f_2)$ for every i. Then $f_1 = f_2$.

(4) (Gluing) Suppose that $\{V_i\}$ is an open cover of U. Suppose that for every i we have an element $f_i \in \mathcal{F}(V_i)$. Furthermore suppose that for every pair of indices i, j we have $\rho_{V_i,V_i \cap V_j}(f_i) = \rho_{V_j,V_i \cap V_j}(f_j)$. Then there exists an element $f \in \mathcal{F}(U)$ satisfying $\rho_{U,V_i}(f) = f_i$ for every i.

Note that the Identity and Gluing axioms together show that given an open cover $\{V_i\}_{i \in I}$ of U and a "compatible" set of elements $f_i \in \mathcal{F}(V_i)$ there exists a *unique* gluing $f \in \mathcal{F}(U)$. In other words, there is an exact sequence

$$0 \to \mathcal{F}(U) \to \prod_i \mathcal{F}(V_i) \to \prod_{i,j} \mathcal{F}(V_i \cap V_j)$$

where each morphism is the product of restriction maps. (This is condition is often used to replace (1),(3),(4) above.)

We will say that a sheaf of abelian groups $\mathcal{F}$ is a sheaf of finitely generated $\mathbb{K}$-algebras if for every non-empty open subset $\mathcal{F}(U)$ is a finitely generated $\mathbb{K}$-algebra and if for every inclusion of non-empty open subsets the restriction map is a $\mathbb{K}$-algebra homomorphism.

Remark 1.8.3. All our sheaves will be sheaves of abelian groups, so we will henceforth drop the "abelian groups" from our notation.

Example 1.8.4. Let X be a topological (resp. differentiable, holomorphic) manifold. For each open set $U \subset X$ let $\mathcal{C}(U)$ denote the set of continuous (resp. differentiable, holomorphic) functions on U. Then $\mathcal{C}$ is a sheaf.

Example 1.8.5. Let X be a topological (resp. differentiable, holomorphic) manifold and let $\pi : \mathcal{E} \to X$ be a vector bundle over X. For each open set $U \subset X$ let $\mathcal{F}(U)$ denote the set of continuous (resp. differentiable, holomorphic) sections of π over U. Then $\mathcal{F}$ is a sheaf.

Example 1.8.6. Let X be a topological space. For every non-empty open $U \subset X$ assign $\mathcal{F}(U) = \mathbb{Z}$ and let $\rho_{U,V}$ be the identity map. This is *not* a sheaf on X. The issue is that it fails the gluing axiom: for example, if U, V are disjoint open subsets then the gluing axiom should imply that $\mathcal{F}(U \cup V) = \mathcal{F}(U) \times \mathcal{F}(V)$ but this fails in our example.

If we want to turn this assignment into a sheaf, we should instead let $\mathcal{F}(U)$ denote the product of copies of $\mathbb{Z}$ indexed by the connected components of U. (The restriction map $\rho_{U,V}$ is the identity on the factors corresponding to an inclusion of connected components and the zero map otherwise.)

1.8.2 Maps of sheaves

Suppose that X is a topological space and that $\mathcal{F}, \mathcal{G}$ are sheaves on X. By definition a morphism of sheaves $\phi : \mathcal{F} \to \mathcal{G}$ assigns to each open set U a homomorphism $\phi(U) : \mathcal{F}(U) \to \mathcal{G}(U)$ in such a way that ϕ is compatible with restriction: for any open $V \subset U$

$$\begin{tikzcd} \mathcal{F}(U) \arrow[r, "\phi(U)"] \arrow[d, "\rho_{U,V}"] & \mathcal{G}(U) \arrow[d, "\rho_{U,V}"] \\ \mathcal{F}(V) \arrow[r, "\phi(V)"] & \mathcal{G}(V). \end{tikzcd}$$

is a commuting diagram. An isomorphism of sheaves on X is a morphism with an inverse.

Suppose now we are given a continuous morphism of topological spaces $f : X \to Y$, and sheaves $\mathcal{O}_X$ on X and $\mathcal{O}_Y$ on Y. We would like to construct a map of sheaves from $\mathcal{O}_Y$ to $\mathcal{O}_X$ representing the pullback of functions. In accordance with Warning 1.8.1, we do not explain a general theory but instead specify what such a map should mean in this special case.

Definition 1.8.7. In the setting above, a pullback map $f^\sharp$ from $\mathcal{O}_Y$ to $\mathcal{O}_X$ assigns to every open set $V \subset Y$ a homomorphism $f^\sharp(V) : \mathcal{O}_Y(V) \to \mathcal{O}_X(f^{-1}(V))$ in such a way that $f^\sharp$ is compatible with restriction: for open subsets $V_1, V_2 \subset Y$ we require that

$$\begin{tikzcd} \mathcal{O}_Y(V_1) \arrow[r, "f^\sharp(V_1)"] \arrow[d, "\rho_{V_1,V_2}"] & \mathcal{O}_X(f^{-1}(V_1)) \arrow[d, "\rho_{f^{-1}(V_1),f^{-1}(V_2)}"] \\ \mathcal{O}_Y(V_2) \arrow[r, "f^\sharp(V_2)"] & \mathcal{O}_X(f^{-1}(V_2)). \end{tikzcd}$$

is a commuting diagram.

Note that given two continuous maps $f : X \to Y$ and $g : Y \to Z$ and pullback maps of sheaves $f^\sharp, g^\sharp$, we can compose the pullback maps in a natural way.

1.8.3 Stalks

The following definition identifies one of the most powerful tools for working with sheaves.

Definition 1.8.8. Let X be a topological space and let $\mathcal{F}$ be a sheaf on X. For any point $x \in X$ the stalk $\mathcal{F}_x$ is defined to be the direct limit

$$\mathcal{F}_x = \varinjlim_{U \ni x} \mathcal{F}(U).$$

In other words, consider the set of pairs (U, f) where U is an open neighborhood of x and $f \in \mathcal{F}(U)$. Say that two pairs (U, f) and (V, g) are equivalent if there is some open set $W \subset U \cap V$ that contains x such that the restrictions of f and g to $\mathcal{F}(W)$ coincide. Then $\mathcal{F}_x$ is the set of equivalence classes of pairs (U, f). We call these equivalence classes germs of sections.

Since the direct limit of abelian groups *receives* a map from each group, for any open neighborhood U of x there is a canonical restriction map $\rho_{U,x} : \mathcal{F}(U) \to \mathcal{F}_x$.

Conceptually speaking, the stalk $\mathcal{F}_x$ records information about all open neighborhoods of x at once without the need to specify a particular neighborhood. One reason why this construction is so useful is that sections of a sheaf can be "determined locally."

Exercise 1.8.9. Let X be a topological space equipped with a sheaf $\mathcal{F}$. Prove that for any open set U the product of the restriction maps

$$\rho : \mathcal{F}(U) \to \prod_{x \in U} \mathcal{F}_x$$

is injective.

Construction 1.8.10. Let X be a topological space. Suppose $\phi : \mathcal{F} \to \mathcal{G}$ is a morphism of sheaves on X. Then for every $x \in X$ the map ϕ determines a morphism on stalks $\phi_x : \mathcal{F}_x \to \mathcal{G}_x$ in the following way. For every open neighborhood U of x consider the composition

$$\mathcal{F}(U) \xrightarrow{\phi(U)} \mathcal{G}(U) \xrightarrow{\rho_{U,x}} \mathcal{G}_x.$$

This collection of homomorphisms determines a homomorphism $\mathcal{F}_x \to \mathcal{G}_x$ using the universal property of the direct limit.

Similarly, a continuous map $f : X \to Y$ and a pullback map of sheaves $f^\sharp : \mathcal{O}_Y \to \mathcal{O}_X$ yield morphisms $\mathcal{O}_{Y,f(x)} \to \mathcal{O}_{X,x}$ for every point $x \in X$.

1.8.4 Exercises

Exercise 1.8.11. Let X be a topological space with a sheaf $\mathcal{F}$. Suppose that $U \subset X$ is an open subset. Show that we can define a sheaf $\mathcal{F}|_U$ on U by restricting the functor $\mathcal{F}$ to the open sets of X contained in U. This sheaf is called the restriction of $\mathcal{F}$ to U.

Exercise 1.8.12. Let X and Y be topological spaces carrying sheaves $\mathcal{O}_X$ and $\mathcal{O}_Y$ respectively. Suppose that $f : X \to Y$ is a continuous map and that $f^\sharp : \mathcal{O}_Y \to \mathcal{O}_X$ is a pullback map of sheaves. Explain how for any open set $U \subset X$ and restricted map $f|_U$ the pullback map of sheaves $f^\sharp$ induces a pullback map of sheaves $f^\sharp|_U : \mathcal{O}_Y \to \mathcal{O}_X|_U$. We will call this the restriction of the pullback map to U.

(In particular, when we apply this construction to the identity map $f : Y \to Y$ we obtain natural pullback maps $\mathcal{O}_Y \to \mathcal{O}_Y|_V$ for open sets $V \subset Y$.)

Exercise 1.8.13. Let X be a topological space and let $\mathcal{F}$ be a sheaf on X. Suppose that U is an open subset and $f \in \mathcal{F}(U)$. The support of f is defined to be

$$\mathrm{Supp}(f) = \{x \in U \mid \rho_{U,x}(f) \neq 0\}.$$

Prove that $\mathrm{Supp}(f)$ is a closed subset of U. (Hint: what does it mean for an equivalence class to be the zero element in $\mathcal{F}_x$?)

Exercise 1.8.14. Prove carefully the claim in the text that pullback morphisms of sheaves can be composed in a natural way.

1.9 Constructing sheaves locally

In this section we discuss a method for constructing sheaves; we will later apply this method to construct the sheaf of functions on a scheme.

1.9.1 Constructing a sheaf from a base

Let X be a topological space. Suppose that $\mathcal{B} = \{V_i\}$ is a base for the topology of X – that is, each V_i is open and every open set in X is a union of elements in our base. We might hope that we can construct a sheaf $\mathcal{F}$ on X by first determining its values on the V_i and then applying the gluing axiom to deduce its value on every open set.

Definition 1.9.1. Let X be a topological space and let $\mathcal{B} = \{V_i\}$ be a base for the topology. A $\mathcal{B}$-sheaf $\widetilde{\mathcal{F}}$ assigns to every open set $V_i \in \mathcal{B}$ an abelian group $\widetilde{\mathcal{F}}(V_i)$ and to each inclusion $V_i \subset V_j$ of open sets in $\mathcal{B}$ a restriction map $\widetilde{\rho}_{V_j,V_i}$ such that the following properties hold:

(1) $\widetilde{\mathcal{F}}(\emptyset) = 0$.

(2) The assignments $\widetilde{\mathcal{F}}, \widetilde{\rho}$ define a contravariant functor from the category of open subsets of X contained in $\mathcal{B}$ (with morphisms = inclusions) to the category of abelian groups.

(3) For any open set $V_i \in \mathcal{B}$ and any open cover of V_i by elements in $\mathcal{B}$ the identity and gluing axioms hold.

Theorem 1.9.2. *Let X be a topological space and let $\mathcal{B} = \{V_i\}_{i\in I}$ be a base for the topology. Suppose that $\widetilde{\mathcal{F}}, \widetilde{\rho}$ are a $\mathcal{B}$-sheaf on X. Then there is a sheaf $\mathcal{F}$ on X such that for every $i \in I$ we have an isomorphism $\phi_i : \mathcal{F}(V_i) \to \widetilde{\mathcal{F}}(V_i)$ and for every $V_j \subset V_i$ we have $\phi_j \circ \rho_{V_i,V_j} = \widetilde{\rho}_{V_i,V_j} \circ \phi_i$. Furthermore $\mathcal{F}$ is uniquely determined up to isomorphism by these properties.*

Proof. Let $U \subset X$ be an open subset. Define $I_U \subset I$ to be the subset of indices such that $V_i \subset U$. We define $\mathcal{F}(U)$ as a subset of the product $\prod_{i\in I_U} \widetilde{\mathcal{F}}(V_i)$:

$$\mathcal{F}(U) := \left\{ (f_i \in \widetilde{\mathcal{F}}(V_i))_{i\in I_U} \;\middle|\; \widetilde{\rho}_{V_{i_1},V_{i_1}\cap V_{i_2}}(f_{i_1}) = \widetilde{\rho}_{V_{i_2},V_{i_1}\cap V_{i_2}}(f_{i_2}) \;\; \forall i_1, i_2 \in I_U \right\}. \tag{1.9.1}$$

For open sets $U_2 \subset U_1$ we define the restriction map ρ_{U_1,U_2} as follows. Note that if $U_2 \subset U_1$ then $R_{U_2} \subset R_{U_1}$. Then we define ρ_{U_1,U_2} as the restriction of the forgetful map $\prod_{i\in I_{U_1}} \widetilde{\mathcal{F}}(V_i) \to \prod_{i\in I_{U_2}} \widetilde{\mathcal{F}}(V_i)$ to the subset $\mathcal{F}(U)$. (It is clear that the image of ρ_{U_1,U_2} is contained in $\mathcal{F}(U_2)$.)

We next prove that $\mathcal{F}$ is a sheaf. The only axiom that takes some work to verify carefully is the gluing axiom. Suppose that $\{W_j\}_{j\in J}$ is an open cover of U and that we are given elements $f_j \in \mathcal{F}(W_j)$ such that for any $j_1, j_2 \in J$ the restriction of f_{j_1} and f_{j_2} to $\mathcal{F}(W_{j_1} \cap W_{j_2})$ coincide. This compatibility implies that if $i' \in \cup_{j\in J} I_{W_j}$ then there is a unique element $f_{i'} \in \widetilde{\mathcal{F}}(V_{i'})$ obtained by restricting f_j from any W_j containing $V_{i'}$.

We would like to glue to obtain an element $f \in \mathcal{F}(U)$. Fix any $i \in I_U$. Consider the subset of $\cup_{j\in J} I_{W_j}$ consisting of those i' such that $i' \in V_i$. By applying the sheaf axioms for $\widetilde{\mathcal{F}}$, we see that the various $f_{i'} \in \widetilde{\mathcal{F}}(V_{i'})$ glue to yield a unique element $f_i \in \widetilde{\mathcal{F}}(V_i)$. As we vary i, we obtain a well-defined element $f \in \mathcal{F}(U)$ which has the property that its restriction to each W_j is f_j.

We next show that $\mathcal{F}(V_i)$ is isomorphic to $\widetilde{\mathcal{F}}(V_i)$. In fact, the isomorphism is simply the forgetful map; note that the compatibility condition of Equation (1.9.1) implies that the forgetful map is an isomorphism. Furthermore, it is clear that under these isomorphisms ρ_{V_i,V_j} is identified with $\widetilde{\rho}_{V_i,V_j}$.

Finally, we must show that $\mathcal{F}$ is uniquely determined by this data. We leave this as an exercise in the gluing axiom. □

Exercise 1.9.3. Prove the uniqueness claim in Theorem 1.9.2.

While Theorem 1.9.2 feels quite technical, in practice it is not so hard to apply: given the data on a base for the topology, for any open set U one constructs the sections in $\mathcal{F}(U)$ by choosing an open cover from our base and then identifying which systems of sections are compatible under restriction.

Exercise 1.9.4. Let X be a topological space and let $\mathcal{B}$ be a base for the topology. Suppose that $\widetilde{\mathcal{F}}$ and $\widetilde{\mathcal{G}}$ are two $\mathcal{B}$-sheaves. A morphism $\widetilde{\phi}$ of $\mathcal{B}$-sheaves assigns to each open set $V_i \in \mathcal{B}$ a homomorphism $\widetilde{\phi}_{V_i} : \widetilde{\mathcal{F}}(V_i) \to \widetilde{\mathcal{G}}(V_i)$ in such a way that the various $\widetilde{\phi}_{V_i}$ commute with restriction.

Prove that a morphism of $\mathcal{B}$-sheaves $\widetilde{\phi} : \widetilde{\mathcal{F}} \to \widetilde{\mathcal{G}}$ induces a morphism of the corresponding sheaves $\phi : \mathcal{F} \to \mathcal{G}$ such that for every $V_i \in \mathcal{B}$ we have $\widetilde{\phi}_{V_i} = \phi_{V_i}$.

We will sometimes apply Theorem 1.9.2 to "glue" sheaves on open subsets of X.

Corollary 1.9.5. *Let X be a topological space equipped with an open cover $\{U_i\}$. Suppose that for each index i we have a sheaf $\mathcal{F}_i$ on U_i. Suppose furthermore that for every pair of indices i, j we have an isomorphism*

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

and that ϕ_{ii} is the identity map, $\phi_{ij} = \phi_{ji}^{-1}$ and $\phi_{jk} \circ \phi_{ij} = \phi_{ik}$ (as isomorphisms of sheaves on $U_i \cap U_j \cap U_k$). Then there is a sheaf $\mathcal{F}$ on X (unique up to isomorphism) such that $\mathcal{F}|_{U_i}$ is isomorphic to $\mathcal{F}_i$.

The conditions on the ϕ_{ij} are known as the "cocycle condition."

Proof. Let $\mathcal{B}$ denote the basis of X consisting of open sets that are contained in some open set U_i in our cover. We define a $\mathcal{B}$-sheaf $\widetilde{\mathcal{F}}$ as follows. For any open subset V contained in some U_i, we define $\widehat{\mathcal{F}}(V)$ as a subset of $\prod_{U_i \supset V} \mathcal{F}_i(V)$ via

$$\widehat{\mathcal{F}}(V) = \{\, (f_i \in \mathcal{F}_i(V)) \;|\; \phi_{ij}(V)(f_i) = f_j \;\; \forall i, j \,\}\,.$$

The restriction maps are defined coordinatewise. Our compatibility conditions guarantee that $\widehat{\mathcal{F}}(V)$ is isomorphic to $\mathcal{F}_i(V)$ for any i such that $U_i \supset V$ and that the restriction maps in $\mathcal{F}_i$ are compatible with the restriction maps in $\widehat{\mathcal{F}}$. Thus Theorem 1.9.2 yields the desired sheaf $\mathcal{F}$. □

1.9.2 Constructing a morphism of sheaves on an open cover

Using the gluing property, we might hope that a morphism of sheaves (or even better, a pullback morphism of sheaves) can be described locally. Let's first recall what happens for continuous maps of topological spaces.

Exercise 1.9.6. Let X and Y be topological spaces. Suppose that we have an open cover $\{U_i\}$ of X and for each index i we have a continuous function $f_i : U_i \to Y$. Suppose furthermore that every pair of indices i, j we have $f_i|_{U_i \cap U_j} = f_j|_{U_i \cap U_j}$. Then there is a unique continuous function $f : X \to Y$ such that $f|_{U_i} = f_i$.

When we enrich our spaces with sheaves, the statement is very similar. In keeping with our philosophy, we will not aim for the most general statement but for the statement we will need in the immediate future. We will use the "restriction of pullback maps" defined in Exercise 1.8.12.

Proposition 1.9.7. *Let X and Y be topological spaces and let $\mathcal{O}_X$ and $\mathcal{O}_Y$ be sheaves of abelian groups on X and Y respectively. Suppose that we have an open cover $\{U_i\}_{i \in I}$ of X, continuous maps $f_i : U_i \to Y$, and pullback maps $f_i^\sharp : \mathcal{O}_Y \to \mathcal{O}_X|_{U_i}$. For every pair of indices $i, j \in I$ let $U_{ij} = U_i \cap U_j$. Suppose furthermore that for all $i, j \in I$ we have $f_i|_{U_{ij}} = f_j|_{U_{ij}}$ as functions and that the two restricted pullback maps of sheaves $\mathcal{O}_Y \to \mathcal{O}_X|_{U_{ij}}$ also coincide.*

Then there is a morphism $f : X \to Y$ and a pullback map of sheaves $f^\sharp : \mathcal{O}_Y \to \mathcal{O}_X$ such that for every $i \in I$ we have $f|_{U_i} = f_i$ and the restriction of $f^\sharp$ to U_i is $f_i^\sharp$. Furthermore f and $f^\sharp$ are uniquely determined.

Proof. As in Exercise 1.9.6 we get a continuous map $f : X \to Y$. Suppose that V is an open subset of Y; we must describe a map $f^\sharp(V) : \mathcal{O}_Y(V) \to \mathcal{O}_X(f^{-1}(V))$. To do this, we identify $\mathcal{O}_X(f^{-1}(V))$ as the subset of $\prod_{i \in I} \mathcal{O}_X(f^{-1}(V) \cap U_i)$ satisfying compatibility:

$$\mathcal{O}_X(f^{-1}(V)) \cong \big(g_i \in \mathcal{O}_X(f^{-1}(V) \cap U_i) \bigm| g_i|_{U_{ij}} = g_j|_{U_{ij}} \forall i, j \in I\big)$$

We then set $f^\sharp(V) = \prod_i f_i^\sharp(V) : \mathcal{O}_Y(V) \to \prod_{i \in I} \mathcal{O}_X(f^{-1}(V) \cap U_i)$. The image of $f^\sharp$ actually lies in $\mathcal{O}_X(f^{-1}(V))$ due to the assumption the the restricted pullback maps coincide on overlaps.

We then need to verify that $f^\sharp$ commutes with restriction maps. But this follows from the combination of the gluing axiom for $\mathcal{O}_X$ and the fact that the $f_i^\sharp$ commute with restriction maps.

Finally we must prove the uniqueness of f and $f^\sharp$. Again we leave this as an exercise in the gluing axiom. □

1.9.3 Exercises

Exercise 1.9.8. Let X be a topological space with a sheaf $\mathcal{F}$. Suppose that $\{V_i\}$ is a base for the topology on X. Prove that

$$\mathcal{F}_x \cong \varinjlim_{V_i \ni x} \mathcal{F}(V_i)$$

where the direct limit is taken over the open sets in our base which contain x.

Exercise 1.9.9. Let X be a topological space and $\phi_1, \phi_2 : \mathcal{F} \to \mathcal{G}$ be two morphisms of sheaves on X. Suppose that $\{V_i\}$ is a base for the topology on X and that $\phi_1(V_i) = \phi_2(V_i)$ for every i. Prove that $\phi_1 = \phi_2$ as morphisms of sheaves.

Exercise 1.9.10. Let X be a topological space and $\phi : \mathcal{F} \to \mathcal{G}$ a morphism of sheaves on X. Suppose that $\{V_i\}$ is a base for the topology on X. Prove that ϕ is an isomorphism if and only if ϕ_{V_i} is an isomorphism for every i.

1.10 Sheaf of functions: varieties

We now return to our goal of constructing a sheaf of functions on an affine variety. (We will postpone the discussion of affine schemes to the next section.)

1.10.1 Functions on open subsets

Let $X = \mathrm{mSpec}(R)$ be an affine variety. In addition to the polynomial functions on X defined by R, one can consider the rational functions on X defined by the quotients of polynomials in R.

Definition 1.10.1. Let $X = \mathrm{mSpec}(R)$ be an affine variety. The function field $\mathbb{K}(X)$ is defined to be the fraction field $\mathrm{Frac}(R)$.

Note that we may not be able to evaluate a rational function $f/g \in \mathrm{Frac}(R)$ on all of X – we will run into trouble at points where g evaluates to 0. However, we can evaluate f/g on the open subset of X where g does not vanish. It is important to remember that when R is not a UFD there may be several different ways of writing the same element of $\mathrm{Frac}(R)$ as a quotient of elements of R (see Example 1.10.3) – thus we may be able to "unexpectedly" evaluate the fraction f/g at a point in $V(g)$ by changing the representative.

Definition 1.10.2. Let $X = \mathrm{mSpec}(R)$ be an affine variety. For any open subset $U \subset X$, the ring of functions on U is denoted by $\mathcal{O}_X(U)$ and is defined by

$$\mathcal{O}_X(U) := \left\{ r \in \mathbb{K}(X) \,\middle|\, \forall \mathfrak{m} \in U, \exists f, g \in R \text{ s.t. } r = \frac{f}{g} \text{ and } \mathfrak{m} \notin V(g) \right\}.$$

We say that an element $r \in \mathbb{K}(X)$ is regular along U if it is contained in $\mathcal{O}_X(U)$.

In other words, $\mathcal{O}_X(U)$ consists of the elements $r \in \mathbb{K}(X)$ which are well-defined on all of U. Note that we allow the representation of r as a quotient of polynomials to change as we vary the point $\mathfrak{m}$; we only require that r have *some* representation as a fraction f/g such that g does not vanish at $\mathfrak{m}$.

Example 1.10.3. Let $X = \mathrm{mSpec}(\mathbb{K}[x_1, x_2, x_3, x_4]/(x_1x_4 - x_2x_3))$ and let U be the complement of $V(x_2, x_4)$. In $\mathbb{K}(X)$ we have the element $\frac{x_1}{x_2} = \frac{x_3}{x_4}$. It is clear that this element is contained in $\mathcal{O}(U)$ since there is no point of U where both x_2 and x_4 vanish. However, this element does not admit a representation as a fraction such that the denominator is well-defined on all of U simultaneously.

Note that the ring of functions *increases* as the open subset *decreases*: smaller open subsets admit more functions. From this perspective, it is natural to think of $\mathbb{K}(X)$ as an object encoding the rings of functions of "all open sets at once", and this intuition will be a useful guide in the future.

Theorem 1.10.4. *Let X be an affine variety. For any open sets $V \subset U$ define the restriction function $\rho_{U,V}$ to be the inclusion $\mathcal{O}_X(U) \hookrightarrow \mathcal{O}_X(V)$ as subsets of $\mathbb{K}(X)$. Then the assignment $U \mapsto \mathcal{O}_X(U)$ and the restriction maps $\rho_{U,V}$ define a sheaf on X.*

We call $\mathcal{O}_X$ the structure sheaf on X.

Proof. The only axiom that needs to be verified carefully is the gluing axiom. Suppose that $\{V_i\}$ is an open cover of a set U. Suppose furthermore that we have elements $f_i \in \mathcal{O}_X(V_i)$ such that $\rho_{V_i,V_i \cap V_j}(f_i) = \rho_{V_j,V_i \cap V_j}(f_j)$. This means that the functions f_i and f_j represent the same element $r \in \mathbb{K}(X)$. Since r admits local expressions showing that is regular on each V_i and the V_i cover U we see that r admits local expressions which shows that it is regular on all of U. The corresponding $f \in \mathcal{O}_X(U)$ represents the gluing of the f_i. $\square$

1.10.2 Open affine subsets

We next introduce a particular type of open set that plays a key role in the construction of the sheaf of functions.

Definition 1.10.5. Let mSpec(R) be an affine scheme. For any function $f \in R$ we let D_f denote the complement of $V(f)$. Such an open set is called a distinguished open affine in mSpec(R).

One of the key properties of distinguished open affines is that they form a base for the Zariski topology.

Proposition 1.10.6. *Let* mSpec(R) *be an affine scheme. The distinguished open affines form a base for the Zariski topology on* mSpec(R)*.*

Proof. First, we must show that the intersection of any two distinguished open affines is a distinguished open affine. In fact, for any $f, g \in R$ we have $D_f \cap D_g = D_{fg}$.

Second, we must show that any open set $U \subset$ mSpec(R) is a union of distinguished open affines. The complement of U is closed, hence equal to $V(I)$ for some ideal $I \subset R$. If we choose a finite generating set $\{f_i\}_{i=1}^r$ for I then

$$\bigcap_{i=1}^{r} V(f_i) = V(I).$$

In other words, $\{D_{f_i}\}_{i=1}^r$ is an open cover of U. $\square$

The ring of functions for a distinguished open affine has a simple description. For any element $f \in R$ we will denote by R_f the localization of R along all non-negative powers of f.

Theorem 1.10.7. *Let $X =$ mSpec(R) be an affine variety and let $f \in R$. Then $\mathcal{O}_X(D_f) = R_f$.*

In particular, applying this theorem with $f = 1$ we see that $\mathcal{O}_X(X) = R$.

Proof. Note that both $\mathcal{O}_X(D_f)$ and R_f are subsets of $\mathrm{Frac}(R)$. Since every element of R_f can be written as g/f^n for some positive integer n it is clear that $R_f \subset \mathcal{O}_X(D_f)$.

Conversely, suppose $r \in \mathcal{O}_X(D_f)$. This means that for any point $\mathfrak{m} \in D_f$ we can find a representation $r = a/g$ such that h does not vanish at $\mathfrak{m}$. In fact the representation a/g will be well-defined on an open subset of D_f, so we may identify an open cover $\{U_i\}$ of D_f and representations $r = a_i/g_i$ such that for every index i we have $V(g_i) \subset X \backslash U_i$. Since distinguished open affines form a basis for the topology, we may replace each U_i by a union of distinguished open affines (while using the same representation of r). After relabeling we may suppose that every open set in our open cover is D_{h_i} for some i and r is represented by the element a_i/g_i in D_{h_i}. Since g_i does not vanish on D_{h_i}, we have $V(h_i) \supset V(g_i)$. This implies that $\sqrt{(h_i)} \subset \sqrt{(g_i)}$ so in particular there is a positive integer k_i such that $h_i^{k_i} \in (g_i)$. Writing $h_i^{k_i} = b_i g_i$, we will rewrite our fraction as $a_i b_i / h_i^{k_i}$.

We will need to understand how the fractions $a_i b_i / h_i^{k_i}$ are related as we change i. Since they each represent the same element in $\mathbb{K}(X)$, by definition the cross-products should yield the same element of R: for every $i, j \in I$ we have

$$a_i b_i h_j^{k_j} = a_j b_j h_i^{k_i}. \tag{1.10.1}$$

Proposition 1.10.6 shows that there is a finite subcover $\{D_{h_i}\}_{i=1}^n$ of D_f. Then

$$V(f) \subset \bigcap_{i=1}^n V(h_i) = V(\{h_i\}_{i=1}^n)$$

so just as before we see there is an integer m such that $f^m \in (\{h_i^{k_i}\})$. We write $f^m = \sum_{j=1}^n c_j h_j^{k_j}$. Define $g = \sum_{j=1}^n a_j b_j c_j$. We claim that the fraction $g/f^m \in R_f$ represents r. Indeed, for any index i we check that this fraction agrees with $a_i b_i / h_i^{k_i}$ using the cross product:

$$a_i b_i f^m = \sum_{j=1}^n a_i b_i c_j h_j^{k_j} = \sum_{j=1}^n a_j b_j c_j h_i^{k_i} = g h_i^{k_i}$$

where we use Equation (1.10.1) for the middle equality. Thus $\mathcal{O}_X(D_f) \subset R_f$. $\square$

This computation is more important than it appears. Recall from Example 1.5.8 that the localization map $R \to R_f$ defines a homeomorphism from $\mathrm{mSpec}(R_f)$ onto its image $D_f \subset \mathrm{mSpec}(R)$. The key property of D_f is that these two computations match up exactly. The following definition formalizes this property.

Definition 1.10.8. Let $X = \mathrm{mSpec}(R)$ be an affine variety and let $U \subset X$ be an open subset. We say that U is an open affine subset of X if the inclusion $R \to \mathcal{O}_X(U)$ defines a map $\mathrm{mSpec}(\mathcal{O}_X(U)) \to X$ which has U as its image.

For an open affine subset U the topology and the ring of functions are in harmony. Thus it makes sense to say that U is isomorphic to the affine variety $\mathrm{mSpec}(\mathcal{O}_X(U))$. However not every open subset of X will be an open affine subset.

Example 1.10.9. Let $U \subset \mathbb{A}^2$ denote the complement of the origin. Since $\mathbb{K}[x, y]$ is a UFD, every element of $\mathbb{K}(\mathbb{A}^2)$ will admit a unique representation as a fraction f/g where f, g are relatively prime. However, there is no polynomial in $\mathbb{K}[x, y]$ which vanishes only at the origin – indeed, since any radical ideal is the intersection of the maximal ideals containing it, any irreducible polynomial g contained in (x, y) must be contained in at least one other maximal ideal. Combining these two facts we deduce that $\mathcal{O}_X(U) = \mathbb{K}[x, y]$ even though $U \subsetneq \mathrm{mSpec}(\mathbb{K}[x, y])$.

1.10.3 Exercises

Exercise 1.10.10. Often the most convenient way to compute the ring of functions on an open set U is to combine Theorem 1.10.7 and the gluing axiom. For example, let U be the complement of the origin in $\mathbb{A}^2$ and let $U_1 = D_x$, $U_2 = D_y$. Note that $U = U_1 \cup U_2$. Use Theorem 1.10.7 and the gluing axiom to compute $\mathcal{O}_X(U)$ in a different way than Example 1.10.9.

Exercise 1.10.11. Example 1.10.9 is just the first instance of a general pattern. For example, show that if $I \subset \mathbb{K}[x_1, \ldots, x_n]$ is not contained in any principal ideal and U is the complement of $V(I)$ in $\mathbb{A}^n$ then the inclusion map $\mathbb{K}[x_1, \ldots, x_n] \to \mathcal{O}_X(U)$ is an isomorphism. (See Theorem 5.6.7 for a general statement.)

Exercise 1.10.12. Let $\mathfrak{K}$ be a finitely generated field extension of $\mathbb{K}$. Show that there is an affine $\mathbb{K}$-variety whose function field is $\mathfrak{K}$.

Exercise 1.10.13. Let $X = \mathrm{mSpec}(R)$ be an affine variety. Suppose that $\mathfrak{p} \subset R$ is a prime ideal and let $Z = V(\mathfrak{p})$. Suppose that $U \subset X$ is an open set such that $U \cap Z \neq \emptyset$. Prove that $\mathcal{O}_X(U) \subset R_{\mathfrak{p}}$ as subsets of $\mathbb{K}(X)$.

1.11 Sheaf of functions: schemes

In this section we construct the sheaf of functions on an affine scheme $X = \mathrm{mSpec}(R)$. Just as with affine varieties, we would like the ring of functions of an open set U to reflect the structure of localizations of R. Unfortunately when working with arbitrary affine schemes there is no field which contains all localizations of R. Thus we will take a different approach.

We will construct the sheaf of functions on U in the following steps. First, for any distinguished open affine we will declare by fiat that $\mathcal{O}_X(D_f) \cong R_f$. Since the distinguished open affines form a base for the topology, this information will be enough to construct a sheaf of functions via Theorem 1.9.2.

1.11.1 Constructing the structure sheaf

Recall from Theorem 1.9.2 the procedure to construct a sheaf from a topological base $\{V_i\}$. First, we must decide what the sections and restriction maps are for open sets in this base. Second, we must verify that our choices satisfy conditions (1-4) of Definition 1.8.2 with respect to open covers consisting of open sets in our base. Third, Theorem 1.9.2 guarantees that there is a unique sheaf on X that is compatible with this data.

Before continuing, we need a couple exercises concerning the geometry of distinguished open affines.

Exercise 1.11.1. Fix a set of elements $\{g_i\}$ in R. Show that $\{D_{g_i}\}$ is an open cover of $\mathrm{mSpec}(R)$ if and only if the g_i generate R.

Exercise 1.11.2. Let $\mathrm{mSpec}(R)$ be an affine scheme and let $f, g \in R$. Show that $D_f \subset D_g$ if and only if the image of g under the localization map $R \to R_f$ is invertible.

Exercise 1.11.3. Let $\mathrm{mSpec}(R)$ be an affine scheme and let $f \in R$. Show that the set of $g \in R$ such that $D_f \supset D_g$ is a multiplicative set.

For any distinguished open affine D_f we define $\mathcal{O}_X(D_f)$ to be the localization of R along all elements $g \in R$ such that $V(g) \subset V(f)$. Exercise 1.11.2 proves that this localized ring is isomorphic to R_f. However, our definition has the advantage that it is defined purely topologically – if we choose different elements $f, g \in R$ with $D_f = D_g$ then $\mathcal{O}_X(D_f) = \mathcal{O}_X(D_g)$ (whereas R_f and R_g are only canonically isomorphic).

For any inclusion $D_f \subset D_g$ we define the restriction map ρ_{D_g,D_f} to be the canonical map obtain from the universal property of localization.

We then need to verify conditions (1-4) of Definition 1.8.2 for this base. The key step is the following proposition.

Proposition 1.11.4. *Let R be a finitely generated $\mathbb{K}$-algebra. Fix a finite set of elements $\{g_i\}_{i=1}^r$ which generate R. Then there is an exact sequence of R-modules*

$$0 \to R \to \prod_i R_{g_i} \to \prod_{i,j} R_{g_i g_j}$$

where the first homomorphism is the product of the localization maps $\rho_i : R \to R_{g_i}$ and the second homomorphism sends the tuple (r_i) to the tuple $(\rho_{i,j} r_i - \rho_{j,i} r_j)$ where $\rho_{i,j} : R_{g_i} \to R_{g_i g_j}$ are the localization maps.

The proof is essentially the same as the proof of Proposition 1.11.4 (and for good reason – the two statements are saying essentially the same thing).

Proof. It is clear that the image of the leftmost map $\prod_i \rho_i$ is contained in the set of compatible elements

$$(f_i \in R_{g_i} | \rho_{i,j}(f_i) = \rho_{j,i}(f_j) \forall i, j)$$

and we must show this map is an isomorphism.

First we show injectivity. Suppose that $f \in R$ is mapped to 0. In other words, for every index i there is some positive integer k_i such that $fg^{k_i} = 0$. Set $N = \sup_i k_i$. Since $R = (g_1, \ldots, g_r)$, we also have $R = (g_1^N, \ldots, g_r^N)$. We deduce that $f = 0$.

Next we show surjectivity. Suppose we have a compatible set of elements f_i. Write $f_i = a_i / g_i^{k_i}$. By assumption we have $a_i / g_i^{k_i} = a_j / g_j^{k_j}$ as elements in $R_{g_i g_j}$. Thus for any pair of indices $i \neq j$ there is some non-negative integer t_{ij} such that

$$a_i g_j^{k_j + t_{ij}} g_i^{t_{ij}} = a_j g_i^{k_i + t_{ij}} g_j^{t_{ij}}.$$

We define $M = \sup_i k_i + \sup_{j \neq i} t_{ij}$. Then we can rewrite

$$f_i = \frac{a_i}{g_i^{k_i}} = \frac{a_i g_i^{M-k_i}}{g_i^M}.$$

For notational convenience we will write $b_i = a_i g_i^{M-k_i}$. The advantage of this change is that for $i \neq j$ we have the simpler relation

$$b_i g_j^M = a_i g_j^{k_j} (g_i^{M-k_i} g_j^{M-k_j}) = a_j g_i^{k_i} (g_i^{M-k_i} g_j^{M-k_j}) = b_j g_i^M.$$

Since $R = (g_1, \ldots, g_r)$, we also have $R = (g_1^M, \ldots, g_r^M)$. In particular we have an equality $1 = \sum_{i=1}^r c_i g_i^M$ for appropriate choices of c_i. Define $f = \sum_{i=1}^r c_i b_i$. We claim that $f \in R$ maps to $f_j \in R_{g_j}$ under the localization map. Indeed, we have

$$f g_j^M = \sum_{i=1}^r c_i b_i g_j^M = \sum_{i=1}^r c_i b_j g_i^M = b_j.$$

□

Exercise 1.11.5. Extend the result of Proposition 1.11.4 by proving that the same statement holds for arbitrary sets of elements $\{g_i\}_{i \in I}$ which generate R.

Corollary 1.11.6. *Let $X = \mathrm{mSpec}(R)$ be an affine scheme. Our definitions of $\mathcal{O}_X(D_f)$ and ρ_{D_g,D_f} satisfy conditions (1-4) of Definition 1.8.2 with respect to the base of open affines.*

Proof. Note that $\emptyset = D_0$ so that $\mathcal{O}_X(\emptyset) = 0$, verifying (1). Condition (2) is clear from the construction. We will show (3) and (4) simultaneously. Let D_f be any distinguished open affine and let $\{D_{g_i}\}_{i \in I}$ be an open cover. By Exercise 1.11.1 the g_i generate R_f. Applying Exercise 1.11.5 to R_f and $\{g_i\}_{i \in I}$ and using the canonical isomorphisms $R_h \cong \mathcal{O}_X(D_h)$ we obtain both (3) and (4). $\square$

We are now in a position to apply Theorem 1.9.2.

Definition 1.11.7. Let $X = \mathrm{mSpec}(R)$ be an affine scheme. The structure sheaf $\mathcal{O}_X$ is the sheaf obtained by applying Theorem 1.9.2 in the above setting.

In other words, if $U \subset X$ is any open subset which is covered by distinguished open affines D_{f_i}, we define $\mathcal{O}_X(U)$ via the exact sequence

$$0 \to \mathcal{O}_X(U) \to \prod_i R_{f_i} \to \prod_{i,j} R_{f_{i,j}}.$$

Exercise 1.11.8. Verify that if X is an affine variety then the structure sheaf as constructed in Definition 1.11.7 is the same as the structure sheaf as constructed in Theorem 1.10.4.

1.11.2 Stalks

We next compute the stalks of the structure sheaf.

Proposition 1.11.9. *Let $X = \mathrm{mSpec}(R)$ be an affine scheme. Then for any point $x \in X$ we have $\mathcal{O}_{X,x} = R_\mathfrak{m}$ where $\mathfrak{m}$ is the maximal ideal corresponding to x.*

Proof. By Exercise 1.9.8 we can compute stalks using the base of distinguished open affines. Thus we are reduced to the computation

$$R_\mathfrak{m} \cong \varinjlim_{f \notin \mathfrak{m}} R_f.$$

These two rings are isomorphic (as rings, and hence also as $\mathbb{K}$-algebras) because they satisfy the same universal property. That is, if $g : R \to S$ is a ring homomorphism such that the image of $R \backslash \mathfrak{m}$ consists of units in S, then g admits a unique factorization through both of these rings – through $R_\mathfrak{m}$ due to the universal property of localization, and through R_f due to the universal properties of localizations and direct limits. $\square$

The main feature of $\mathcal{O}_{X,x}$ is its unique maximal ideal which we will denote by $\mathfrak{m}_x$. (We will try to avoid the unfortunate but logically consistent notation $\mathfrak{m}_\mathfrak{m}$.) We know that $\mathfrak{m}_x$ is generated by a finite set of functions $f_1, \ldots, f_r$ all of which vanish at x. Loosely speaking, we can think of the f_i as giving "local coordinates" near x. Indeed, for any sufficiently small open affine neighborhood U of x the f_i will be regular functions on U whose common vanishing locus is only the point x.

Remark 1.11.10. This analogy should be taken with a grain of salt – it works well in some ways but not in others. For example, a smooth n-dimensional manifold is locally isomorphic to $\mathbb{R}^n$, but the local rings $\mathcal{O}_{X,x}$ usually will look quite different from the local rings of $\mathbb{A}^n$. (See Section 5.2 for one aspect in which the analogy does work well.)

It turns out that the closest analogy with the Euclidean situation is if we look not at the localization $\mathcal{O}_{X,x}$ but the completion $\widehat{\mathcal{O}_X(X)}_\mathfrak{m}$. In other words, if we use formal power series instead of polynomials we can come closer to capturing the behavior of analytic functions on small open neighborhoods. We will not pursue this direction in these notes.

We can now finally resolve our previous discussion on the support of a function. Given an open set U and an element $f \in \mathcal{O}_X(U)$ we define support of f to be the set of points $x \in U$ such that the image of f under $\rho_{U,x} : \mathcal{O}_X(U) \to \mathcal{O}_{X,x}$ is non-zero. (According to Exercise 1.4.14 this agrees with our old definition when U is affine.)

Lemma 1.11.11. *Let X be an affine scheme. Let $U \subset X$ be an open set and let $f \in \mathcal{O}_X(U)$. Then the complement of* $\operatorname{Supp}(f)$ *in U is the largest open subset V of U such that $\rho_{U,V}(f)$ is identically zero.*

Proof. This is an immediate consequence of Exercise 1.8.9. □

Remark 1.11.12. It is important to be clear on the difference between the support $\operatorname{Supp}(f)$ and the vanishing locus $V(f)$ (as defined in Exercise 1.11.13). Geometrically, the vanishing locus determines when f *evaluates* to 0. The complement of the support is the largest open set where f is *identically* zero. (Recall that being identically zero is stronger than evaluating to zero!) Algebraically, the vanishing locus determines when f is sent to 0 by *quotienting*. The support determines when f is sent to 0 by *localizing*.

1.11.3 Exercises

Exercise 1.11.13. Suppose that X is an affine scheme and that U is an open subset. Let $I \subset \mathcal{O}_X(U)$ be an ideal. The vanishing locus $V(I)$ is defined to be set of points $x \in U$ such that the stalk map $\rho_{U,x} : \mathcal{O}_X(U) \to \mathcal{O}_{X,x}$ satisfies the property that $\rho_{U,x}(I) \subset \mathfrak{m}_{X,x}$.

(1) Show that $V(I)$ is a closed subset of U.

(2) Show that if U is an open affine set then $V(I)$ coincides with Definition 1.2.3.

Exercise 1.11.14. Suppose that $f : X \to Y$ is a morphism of affine schemes. Show that the localizations of $f^\sharp$ induce a pullback morphism of sheaves $f^\sharp : \mathcal{O}_Y \to \mathcal{O}_X$.

Exercise 1.11.15. Let $X = \mathrm{mSpec}(R)$ be an affine scheme. Given any subset $Z \subset X$ we can define

$$\mathcal{O}_{X,Z} = \varinjlim_{U \supset Z} \mathcal{O}(U).$$

Prove that if $\mathfrak{p}$ is a prime ideal in R then $\mathcal{O}_{X,V(\mathfrak{p})} = R_{\mathfrak{p}}$.

1.12 Quasiaffine schemes

We are finally prepared to construct a category of "quasiaffine $\mathbb{K}$-schemes" which puts open subsets of affine schemes on equal footing. Example 1.10.9 demonstrated that an open subset U of an affine scheme X need not itself be an affine scheme. For a non-affine open subset, the ring of functions $\mathcal{O}(U)$ is not really sufficient information – instead, we should keep track of the entire sheaf of functions.

Definition 1.12.1. A quasiaffine $\mathbb{K}$-scheme consists of an open set $U \subset X$ of an affine $\mathbb{K}$-scheme X equipped with the following data:

$$(\text{set, topology, sheaf of functions}) = (U, \mathbf{Zar}|_U, \mathcal{O}|_U).$$

Exercise 1.12.10 shows that given any quasiaffine scheme U the map $f : U \to \mathrm{mSpec}(\mathcal{O}(U))$ is an injection realizing U as an open subset of an affine scheme. Thus one can make quasiaffineness an "intrinsic" property by requiring that this map f realizes U as an open subset of an affine scheme and that the sheaf of functions on U is the restriction from $\mathrm{mSpec}(\mathcal{O}(U))$.

Warning 1.12.2. Suppose X is a quasiaffine $\mathbb{K}$-scheme. Although $\mathcal{O}_X(X)$ is a $\mathbb{K}$-algebra, it need not be finitely generated. We will see a closely related example in a slightly different setting in Example 2.4.6.

1.12.1 Morphisms of quasiaffine schemes

Suppose we have two quasiaffine schemes X and Y. Then at the very least a morphism $f : X \to Y$ should consist of the following data:

(1) A continuous set-theoretic map $f : X \to Y$.

(2) A pullback morphism of sheaves $f^\sharp : \mathcal{O}_Y \to \mathcal{O}_X$.

However, this data is *not* sufficient to define a morphism. The issue is that we have unmoored the algebra from the topology. (For example, there is nothing to prevent us from letting f be any continuous map and setting the morphisms $f^\sharp$ to be identically zero.) We need to add a condition to ensure that the set-theoretic map is induced by the map of functions. Our condition will use the vanishing locus as defined in Exercise 1.11.13:

(*) Let V be an open subset of Y and let $I \subset \mathcal{O}(V)$ be an ideal. Then we have $f^{-1}(V(I)) = V(\langle f^\sharp(I)\rangle)$ inside of $f^{-1}(V)$.

Definition 1.12.3. A morphism of quasiaffine schemes $f : X \to Y$ consists of a set theoretic map f and a pullback map of sheaves $f^\sharp : \mathcal{O}_Y \to \mathcal{O}_X$ satisfying condition (*).

These definitions yield a category $\mathbf{QAffSch}/\mathbb{K}$ of quasiaffine schemes.

Exercise 1.12.4. Explain carefully how to compose two morphisms of quasiaffine schemes.

Exercise 1.12.5. Show that condition (*) is equivalent to:

(*') For any point $x \in X$ the induced map of stalks $f^\sharp_x : \mathcal{O}_{Y,f(x)} \to \mathcal{O}_{X,x}$ is a local homomorphism of local rings. (That is, the image of the maximal ideal in $\mathcal{O}_{Y,f(x)}$ is the maximal ideal in $\mathcal{O}_{X,x}$.)

We have finally achieved our goal of putting closed and open subsets on the same footing as affine schemes. The following definitions explain how to give a closed set or open set the structure of a quasiaffine scheme.

Definition 1.12.6. A closed embedding is a morphism $f : Z \to X$ such that

(1) f takes Z homeomorphically onto a closed subset of X, and

(2) for any open affine subset $V \subset X$ the map $f^\sharp(V) : \mathcal{O}(V) \to \mathcal{O}(f^{-1}(V))$ is surjective.

Definition 1.12.7. An open embedding is a morphism $f : Y \to X$ such that

(1) f takes Y homeomorphically onto an open subset of X, and

(2) for any open subset $V \subset X$ contained in $f(Y)$ the map $f^\sharp(V) : \mathcal{O}(V) \to \mathcal{O}(f^{-1}(V))$ is an isomorphism.

1.12.2 Revisiting affine $\mathbb{K}$-schemes

We currently have two competing definitions of a morphism of affine schemes (Definition 1.5.1 and Definition 1.12.3). The following result resolves this tension.

Proposition 1.12.8. *Let $X = \mathrm{mSpec}(R)$ and $Y = \mathrm{mSpec}(S)$ be affine schemes. Then there is a bijection between morphisms $f : X \to Y$ (in the sense of Definition 1.12.3) and $\mathbb{K}$-algebra homomorphisms $f^\sharp : S \to R$.*

Proof. First suppose given a $\mathbb{K}$-algebra homomorphism $f^\sharp : S \to R$. Proposition 1.5.3 constructs a continuous set-theoretic map $f : X \to Y$ from $f^\sharp$. Exercise 1.11.14 shows that $f^\sharp$ naturally induces a morphism of sheaves. Finally, the argument of Proposition 1.5.3 shows that the set-theoretic map f is compatible with the behavior of vanishing loci induced by $f^\sharp$. Together these define a morphism $f : X \to Y$ in the sense of Definition 1.5.1.

Conversely, suppose given a morphism $f : X \to Y$ as in Definition 1.5.1. By definition f includes the data of a $\mathbb{K}$-algebra map $f^\sharp(Y) : \mathcal{O}(Y) \to \mathcal{O}(X)$.

We need to verify that these are inverse constructions; there is only one non-trivial direction. Suppose that $f : X \to Y$ is a morphism and let $f^\sharp : S \to R$ be the induced map as above. Let $f' : X \to Y$ be the morphism constructed from $f^\sharp$ and let $f'^\sharp$ denote the induced map of sheaves. We must show that f' coincides with f. By construction $f^\sharp$ and $f'^\sharp$ agree as maps $S = \mathcal{O}_Y(Y) \to \mathcal{O}_X(X) = R$ (but we have not yet shown that the two sheaf maps coincide for other choices of open sets).

We start by discussing the topology of f'. Let $\mathfrak{m}$ be a maximal ideal of S. Then by applying condition (*) to the open sets X and Y we have

$$f^{-1}(V(\mathfrak{m})) = V(f^\sharp(\mathfrak{m})) = V(f'^\sharp(\mathfrak{m})) = f'^{-1}(V(\mathfrak{m})).$$

This shows that f and f' agree set-theoretically.

We then must show that $f'^\sharp$ and $f^\sharp$ are the same pullback maps of sheaves. Suppose that D_g is a distinguished open affine in X and D_h is a distinguished open affine in Y such that $D_g \subset f^{-1}(D_h)$. Applying condition (*) to the open sets X and Y we see that $f^\sharp(h) \in \sqrt{g}$. By analogous reasoning $f'^\sharp(h) \in \sqrt{g}$. Consider the commuting diagram

$$\begin{tikzcd} S \arrow[r, "f^\sharp = f'^\sharp"] \arrow[d, "\psi_h"] & R \arrow[d, "\psi_g"'] \\ S_h \arrow[r] & R_g \end{tikzcd}$$

where ψ_h, ψ_g are the localization maps. By the universal property of localization there is a unique homomorphism along the bottom that makes this diagram commute. Thus we see that $f'^\sharp(D_h) = f^\sharp(D_h) : \mathcal{O}(D_h) \to \mathcal{O}(D_g)$.

For arbitrary open sets, the equality of $f^\sharp$ and $f'^\sharp$ follows from the case of distinguished open affines using the uniqueness in Proposition 1.9.7. $\square$

In particular this implies:

Corollary 1.12.9. *The category* **AffSch**$/\mathbb{K}$ *is a full subcategory of* **QAffSch**$/\mathbb{K}$*.*

We will now retcon our definition of an affine scheme to include the sheaf of functions.

1.12.3 Exercises

Exercise 1.12.10. Suppose that U is a quasiaffine scheme, so in particular U is an open subset of an affine scheme X. Show that the restriction map $\rho : \mathcal{O}(X) \to \mathcal{O}(U)$ defines an open embedding $\mathrm{mSpec}(\mathcal{O}(U)) \to X$ whose image contains U.

Exercise 1.12.11. Consider the affine scheme X defined as $V(xy - zw) \subset \mathbb{A}^4_{\mathbb{K}}$. Let $U \subset X$ denote the complement of $V(y, z)$. Prove that U is isomorphic to the complement of a line in $\mathbb{A}^3$. Conclude that U is not an affine scheme. (This computation is related to Example 1.10.3 showing that there is an element of $\mathbb{K}(X)$ that is defined globally on U but has no globally defined fraction.)

Chapter 2

Projective schemes

Let's quickly review the construction of projective space from a classical viewpoint.

Let $\mathbb{K}$ be a field. Projective space $\mathbb{P}^n$ over $\mathbb{K}$ is the parameter space for the set of 1-dimensional subspaces of $\mathbb{K}^{n+1}$. We can construct this space as the quotient of the set of non-zero vectors in $\mathbb{K}^{n+1}$ by the equivalence relation which identifies any two proportional vectors:

$$\mathbb{P}^n := \frac{\mathbb{K}^{n+1}\backslash\{0\}}{\vec{a} \sim \lambda\vec{a}}.$$

Traditionally one identifies a point of $\mathbb{P}^n$ using an $(n+1)$-tuple $(a_0 : a_1 : \ldots : a_n)$ which represents the coordinates for a point in the corresponding line in $\mathbb{K}^{n+1}$. We use colons instead of commas to remind us that any rescaling of this vector represents the same point of $\mathbb{P}^n$:

$$(a_0 : a_1 : \ldots : a_n) = (\lambda a_0 : \lambda a_1 : \ldots : \lambda a_n) \qquad \forall\lambda \in \mathbb{K}\backslash\{0\}$$

When $\mathbb{K} = \mathbb{C}$ or $\mathbb{R}$ we can equip $\mathbb{P}^n$ with the quotient of the Euclidean topology on $\mathbb{K}^{n+1}$ and with this choice projective space is compact. In algebraic geometry we usually equip $\mathbb{P}^n$ with the quotient of the Zariski topology on $\mathbb{K}^{n+1}$. Concretely, this means that the closed sets in $\mathbb{P}^n$ are the subsets cut out by homogeneous polynomial equations in $\mathbb{K}[x_0, \ldots, x_n]$. (Although it does not make sense to evaluate a homogeneous polynomial at a point of projective space, it does make sense to ask whether a homogeneous polynomial vanishes at a point of projective space since this condition is scaling-invariant.)

Let D_i denote the open subset of points $\vec{a} \in \mathbb{P}^n$ where the ith coordinate does not vanish. (Note that this condition does not depend upon which representative of $\vec{a}$ we pick.) After rescaling the vector $\vec{a}$, we may ensure that the ith coordinate is equal to 1:

$$D_i = \{(a_0, \ldots, a_{i-1}, 1, a_{i+1}, \ldots, a_n) \in \mathbb{P}^n\}$$

This description gives D_i a bijective correspondence with the set $\mathbb{K}^n$. When $\mathbb{K} = \mathbb{C}$ or $\mathbb{R}$ we use these open sets as charts to give $\mathbb{P}^n$ the structure of a manifold. The complement of D_i is isomorphic to projective space of dimension one less:

$$\mathbb{P}^n \backslash D_i = \{(a_0, \ldots, a_{i-1}, 0, a_{i+1}, \ldots, a_n) \in \mathbb{P}^n\}.$$

Thus in geometric situations one can think of $\mathbb{P}^n$ as the compactification of $\mathbb{K}^n$ obtained by "adding a $\mathbb{P}^{n-1}$ at infinity." By induction we have

$$\mathbb{P}^n = \mathbb{K}^n \cup \mathbb{K}^{n-1} \cup \ldots \cup \mathbb{K} \cup \{pt\}.$$

Note that the only polynomial functions on $\mathbb{K}^{n+1}$ which descend to $\mathbb{P}^n$ are the constant functions. (This is analogous to Liouville's theorem in complex geometry which guarantees that every holomorphic function on $\mathbb{P}^n_{\mathbb{C}}$ is constant.) Nevertheless, $\mathbb{P}^n$ admits many polynomial maps to other projective spaces. Suppose that $f_0, \ldots, f_m \in \mathbb{K}[x_0, \ldots, x_n]$ are homogeneous polynomials of the same degree d. We can then define a map $f : \mathbb{P}^n \to \mathbb{P}^m$ via the prescription

$$f(\vec{a}) = (f_0(\vec{a}) : \ldots : f_m(\vec{a}))$$

If we rescale $\vec{a}$ by λ then $f(\vec{a})$ is rescaled by λ^d. This guarantees that f is a well-defined map for any $\vec{a}$ *except* possibly for the points in the closed set $f_0(\vec{a}) = \ldots = f_m(\vec{a}) = 0$. In other words, homogeneous polynomial functions naturally define maps on open subsets of $\mathbb{P}^n$.

Primer on graded rings

The only graded rings R we will consider will have $\mathbb{Z}$-gradings. Our standard example is the polynomial ring $\mathbb{K}[x_0, \ldots, x_n]$ given the grading by degree. We denote the dth graded piece of R by R_d. Each graded piece of R is naturally an R_0-module.

An element $f \in R$ is said to be homogeneous of degree d if it is contained in R_d. More generally, given any element $f \in R$ the dth homogeneous component of f is the image of f under the projection $R = \oplus_d R_d \to R_d$. If S is a multiplicatively closed subset of R consisting only of homogeneous elements, then R_S is a graded ring under the assignment $\deg(1/f) = -\deg(f)$.

An ideal $I \subset R$ is homogeneous if it satisfies any of the following equivalent conditions:

(1) I is generated by homogeneous elements.

(2) If $f \in I$ then every homogeneous component of f is contained in I.

(3) $I = \oplus_{d \geq 0}(I \cap R_d)$.

If I is a homogeneous ideal then R/I inherits a grading from R. If I, J are homogeneous ideals, then so are $I + J$, $I \cap J$, IJ, and $\sqrt{I}$. A homogeneous ideal is prime if and only if for every pair of *homogeneous* elements $f, g \in R$ satisfying $fg \in I$ either $f \in I$ or $g \in I$.

A finitely generated graded $\mathbb{K}$-algebra R is a finitely generated $\mathbb{K}$-algebra equipped with a $\mathbb{Z}_{\geq 0}$-grading such that the canonical copy of $\mathbb{K}$ in R has degree 0. (We also allow the 0 ring.) The main example is the polynomial ring $\mathbb{K}[x_0, \ldots, x_n]$ graded by degree. Note that the quotient of a polynomial ring by a homogeneous ideal will be a finitely generated graded $\mathbb{K}$-algebra, but not every finitely generated $\mathbb{K}$-algebra has this form. (There are two potential obstructions – any quotient of a polynomial ring will satisfy $R_0 \cong \mathbb{K}$ and will be generated as a $\mathbb{K}$-algebra by its degree 1 elements.)

A graded homomorphism of graded rings R, S is a homomorphism $f : R \to S$ such that every homogeneous element in R maps to a homogeneous element in S and $f(R_0) \subset S_0$. In other words, there will be some integer m such that $f(R_d) \subset S_{md}$ for every d. In particular pullback under f will preserve the homogeneity of ideals.

2.1 Projective space: quotient

The first task facing us is to describe projective space as a topological space. Just as for affine space, we would like to add "non-traditional" points to obtain a tighter link between algebra and geometry. There are two approaches: we can either construct $\mathbb{P}^n$ as a quotient of $\mathbb{A}^{n+1}\backslash 0$ or we can construct $\mathbb{P}^n$ by gluing together affine spaces $\mathbb{A}^n$. In this section we discuss the first approach.

2.1.1 Points

Our strategy for the construction of $\mathbb{P}^n$ is to use the "scaling invariant" geometry of $\mathbb{A}^{n+1}$. Note that the vanishing locus in $\mathbb{A}^{n+1}$ defined by a homogeneous ideal will be "invariant under rescaling" in the sense that it will be a union of lines through the origin. Conversely, any "scaling invariant" subset will be defined by a homogeneous ideal. Thus we will construct projective space using the homogeneous ideals in $\mathbb{K}[x_0, \ldots, x_n]$.

There is a unique homogeneous ideal in $\mathbb{K}[x_0, \ldots, x_n]$ which is also a maximal ideal, namely, the origin $(x_0, \ldots, x_n)$. The lines through the origin are the "next smallest" closed subsets which are scaling invariant, and these should give us points in projective space.

Definition 2.1.1. We say that a homogeneous ideal $\mathfrak{m} \subset \mathbb{K}[x_0, \ldots, x_n]$ is almost maximal if it is prime, it is properly contained in $(x_0, \ldots, x_n)$, and there is no prime homogeneous ideal I satisfying $\mathfrak{m} \subsetneq I \subsetneq (x_0, \ldots, x_n)$.

The points of projective space $\mathbb{P}^n_{\mathbb{K}}$ are the almost maximal homogeneous ideals $\mathfrak{m} \subset \mathbb{K}[x_0, \ldots, x_n]$.

If we have n linearly independent homogeneous linear functions $\ell_1, \ldots, \ell_n$ then the ideal $(\ell_1, \ldots, \ell_n)$ is an example of an almost maximal homogeneous ideal in $\mathbb{K}[x_0, \ldots, x_n]$. These are the "traditional points" of $\mathbb{P}^n$ which define 1-dimensional subspaces of $\mathbb{K}^{n+1}$. However, there can also be "non-traditional points" where some of the generators have degree larger than 1.

Exercise 2.1.2. Let $\mathbb{K}$ be a field. Show that every almost maximal homogeneous ideal in $\mathbb{K}[x, y]$ will be generated by a single irreducible homogeneous polynomial. (See Example 1.3.7.)

Explain why the assignment

$$\sum_{i=0}^{d} c_i x^i y^{d-i} \leftrightarrow \sum_{i=0}^{d} c_i t^i$$

describes a bijection between the set of irreducible homogeneous polynomials in $\mathbb{K}[x, y]$ *except for the polynomial* y and the set of irreducible polynomials in $\mathbb{K}[t]$. This bijection gives us an explicit description of the points on $\mathbb{P}^1_{\mathbb{K}}$ as being equal to $\mathbb{A}^1_{\mathbb{K}} \cup \{(y)\}$.

2.1.2 Topology

Definition 2.1.3. Let $I \subset \mathbb{K}[x_0, \ldots, x_n]$ be a homogeneous ideal. The vanishing locus $V_+(I) \subset \mathbb{P}^n$ is the set of almost maximal homogeneous ideals which contain I.

Here the "+" in V_+ emphasizes that the construction is reflecting the behavior of homogeneous elements of degree > 0; the degree 0 homogeneous elements don't behave in the same way as other homogeneous elements. It also reminds us that homogeneous functions aren't really functions on $\mathbb{P}^n$ – we save the notation V for honest functions (see Exercise 1.11.13).

The vanishing locus behaves just like you would expect with respect to radicals (see Exercise 2.1.16) and with respect to ideal operations.

Exercise 2.1.4. Verify that $\mathbb{P}^n$ admits a topology whose closed sets are the vanishing loci of homogeneous ideals. (Note that the maximal ideal $(x_0, \ldots, x_n)$ will define the empty set. For this reason the maximal ideal is sometimes called the "irrelevant ideal.")

The topology on $\mathbb{P}^n$ whose closed sets are the vanishing loci of homogeneous ideals is known as the Zariski topology on $\mathbb{P}^n$.

Example 2.1.5. Suppose $f \in \mathbb{K}[x_0, \ldots, x_n]$ is a homogeneous equation of degree d. The vanishing locus $V_+(f)$ is called a hypersurface of degree d.

Example 2.1.6. Suppose that $\ell_1, \ldots, \ell_{n-k}$ are homogeneous linear equations that are linearly independent (as elements of the $\mathbb{K}$-vector space $\mathbb{K}[x_0, \ldots, x_n]_1$). The vanishing locus of these equations is called a k-plane in $\mathbb{P}^n$. In the special case when we have a single linear equation, the vanishing locus is called a hyperplane – equivalently, a hyperplane is a degree 1 hypersurface.

On the level of traditional points, a k-plane is just the image of a $(k+1)$-dimensional subspace of $\mathbb{K}^{n+1}$ under the quotient map to projective space.

2.1.3 Functions

Now that we have constructed $\mathbb{P}^n$ as a topological space, we would like to construct a sheaf of algebraic functions on $\mathbb{P}^n$. In contrast to the situation for affine schemes, homogeneous functions on $\mathbb{K}[x_0, \ldots, x_n]$ do not define functions on $\mathbb{P}^n$ in any natural way. Indeed, if we rescale a traditional point in $\mathbb{P}^n$ by a constant λ then the value of a degree d homogeneous polynomial will be rescaled by λ^d. (It turns out that degree d homogeneous functions define sections of a line bundle on $\mathbb{P}^n$.)

However, if we take the quotient of two degree d polynomials, then we *do* get a well-defined function on projective space (on the open set where the denominator does not vanish).

Definition 2.1.7. Let $\mathbb{F}$ denote the field obtained by localizing $\mathbb{K}[x_0, \ldots, x_n]$ along all (non-zero) homogeneous elements. Then the function field $\mathbb{K}(\mathbb{P}^n)$ is the subfield of $\mathbb{F}$ consisting of fractions whose numerators and denominators have the same degree. More explicitly,

$$\mathbb{K}(\mathbb{P}^n) := \mathbb{K}\left(\frac{x_1}{x_0}, \ldots, \frac{x_n}{x_0}\right).$$

Definition 2.1.7 should be compared with the well-known theorem that every meromorphic function on complex projective space is a rational function whose numerators and denominators have the same degree.

Definition 2.1.8. Let $U \subset \mathbb{P}^n$ be an open subset. We define the function ring $\mathcal{O}_{\mathbb{P}^n}(U)$ to be

$$\mathcal{O}_{\mathbb{P}^n}(U) = \left\{ \frac{f}{g} \in \mathbb{K}(\mathbb{P}^n) \,\middle|\, V_+(g) \subset \mathbb{P}^n \backslash U \right\}$$

Given an inclusion of open subsets $V \subset U$, we define the restriction map $\rho_{U,V} : \mathcal{O}_{\mathbb{P}^n}(U) \to \mathcal{O}_{\mathbb{P}^n}(V)$ as the inclusion as subsets of $\mathbb{K}(\mathbb{P}^n)$.

Exercise 2.1.9. Prove that $\mathcal{O}_{\mathbb{P}^n}$ is a sheaf of $\mathbb{K}$-algebras.

Remark 2.1.10. Since $\mathbb{K}(\mathbb{P}^n)$ is the fraction field of the UFD $\mathbb{K}[\frac{x_1}{x_0}, \ldots, \frac{x_n}{x_0}]$ we do not need to define the sheaf of functions using local expressions like we did in Definition 1.10.2.

Example 2.1.11. According to Definition 2.1.8 we have $\mathcal{O}_{\mathbb{P}^n}(\mathbb{P}^n) = \mathbb{K}$. In other words, the only global functions on projective space are the constant functions. (This should be compared with Liouville's theorem showing that the only holomorphic functions on a compact space are the constant ones.)

2.1.4 Distinguished open affines

The next definition identifies the most important open subsets of projective space.

Definition 2.1.12. Fix a homogeneous element $f \in \mathbb{K}[x_0, \ldots, x_n]$ of degree ≥ 1. The distinguished open affine set corresponding to f is the open set $D_{+,f} := \mathbb{P}^n \backslash V_+(f)$.

Notation 2.1.13. Set $S = \mathbb{K}[x_0, \ldots, x_n]$ and consider the localization S_f at a homogeneous element f. If we assign $\deg(\frac{1}{f}) = -\deg(f)$ then S_f is a $\mathbb{Z}$-graded ring. As usual we let $(S_f)_d$ denote the degree d part.

Proposition 2.1.14. *For any homogeneous element $f \in \mathbb{K}[x_0, \ldots, x_n]$ of degree ≥ 1 we have*

$$\mathcal{O}_{\mathbb{P}^n}(D_{+,f}) = (\mathbb{K}[x_0, \ldots, x_n]_f)_0.$$

Proof. The containment $\supset$ is clear, and we just need to show the reverse containment. Since $\mathbb{K}[x_0, \ldots, x_n]$ is a UFD, each element in $\mathbb{K}(\mathbb{P}^n)$ admits a unique expression $\frac{g}{h}$ in lowest terms. Then $\mathcal{O}_{\mathbb{P}^n}(D_{+,f})$ will consist of those fractions such that h is a factor of some power of f. If we write $f^r = ah$, then our fraction can also be written as $\frac{ag}{f^r}$. This shows the containment $\subset$. $\square$

2.1.5 Exercises

Exercise 2.1.15. Show that $\mathbb{P}^n$ is an irreducible Noetherian topological space.

Exercise 2.1.16. Let I, J be homogeneous ideals properly contained in $\mathbb{K}[x_0, \ldots, x_n]$. Verify that the vanishing locus satisfies the following properties.

(1) $V_+(I) = V_+(\sqrt{I})$.

(2) $V_+(I) = V_+(J)$ if and only if $\sqrt{I} = \sqrt{J}$.

(3) If $I \subset J$ then $V_+(I) \supset V_+(J)$.

(4) If $V_+(I) \subset V_+(J)$ then $\sqrt{I} \supset \sqrt{J}$.

Exercise 2.1.17. Prove that a closed subset $X \subset \mathbb{P}^n$ is irreducible if and only if there is a homogeneous prime ideal $\mathfrak{p} \subset \mathbb{K}[x_0, \ldots, x_n]$ such that $X = V(\mathfrak{p})$.

Exercise 2.1.18. Suppose that $I \subset \mathbb{K}[x_0, \ldots, x_n]$ is a homogeneous ideal that is not contained in any homogeneous principal ideal. Let U be the complement of $V_+(I)$. Prove that $\mathcal{O}_{\mathbb{P}^n}(U) = \mathbb{K}$.

Exercise 2.1.19. For any point $x \in \mathbb{P}^n$, verify that the stalk $\mathcal{O}_{\mathbb{P}^n,x}$ is isomorphic to the degree 0 part of the localization of $\mathbb{K}[x_0, \ldots, x_n]$ along all *homogeneous* elements not contained in $\mathfrak{m}$.

Exercise 2.1.20. Prove that the distinguished open affines on $\mathbb{P}^n$ form a base for the Zariski topology.

2.2 Projective space: scheme

While the construction of $\mathbb{P}^n$ given in Definition 2.1.1 is very succinct, it is unfortunately a bit difficult to identify the points of $\mathbb{P}^n$ directly using this description. An alternative approach is to construct $\mathbb{P}^n$ by gluing together copies of affine space $\mathbb{A}^n$.

In this section we study the distinguished open affines $D_{+,f}$ in $\mathbb{P}^n$ in more detail. We will show that the sets $D_{+,f}$ equipped with the sheaves $\mathcal{O}_{\mathbb{P}^n}|_{D_{+,f}}$ are examples of the affine schemes we studied in the previous chapter. The most important example are the "affine charts" D_{+,x_i}, and for simplicity we will focus on this special case.

2.2.1 Distinguished charts

Recall that $D_{+,x_i} = \mathbb{P}^n \backslash V_+(x_i)$ is the locus where x_i does not vanish. As discussed in the introduction to the chapter, the traditional points in D_{+,x_i} can naturally be identified with the vector space $\mathbb{K}^n$. This identification is achieved by rescaling by the ith coordinate

$$(a_0 : \ldots : a_n) \leftrightarrow \left(\frac{a_0}{a_i} : \ldots : \frac{a_{i-1}}{a_i} : 1 : \frac{a_{i+1}}{a_i} : \ldots : \frac{a_n}{a_i}\right)$$

In other words, the linear coordinates on D_{+,x_i} are the restriction of the functions $\frac{a_j}{a_i}$ on $\mathbb{P}^n$.

We would like to "upgrade" this identification to show that D_{+,x_i} is actually isomorphic to $\mathbb{A}^n$ (as a topological space equipped with a sheaf). While we now need to argue using algebra in the place of traditional points, the construction is essentially the same.

The first step is to identify the ring of functions on D_{+,x_i} with the polynomial ring of functions on affine space. According to Proposition 2.1.14, we can identify $\mathcal{O}_{\mathbb{P}^n}(D_{+,x_i})$ with the degree 0 subring of $\mathbb{K}[x_0, \ldots, x_n]_{x_i}$.

Lemma 2.2.1. *Set $S = \mathbb{K}[x_0, \ldots, x_n]$. Then $(S_{x_i})_0$ is isomorphic to*

$$\mathbb{K}[y_0, \ldots, y_{i-1}, y_{i+1}, \ldots, y_n]$$

under the identification $\frac{x_j}{x_i} \leftrightarrow y_j$. More generally, each $(S_{x_i})_d$ is a free $(S_{x_i})_0$-module generated by x_i^d.

Not only does the lemma show that $\mathcal{O}_{\mathbb{P}^n}(D_{+,x_i})$ is isomorphic to a polynomial ring, it also shows that the identification is the same as the geometric identification described earlier – we simply invert the ith coordinate.

Exercise 2.2.2. Verify the previous lemma carefully.

The next step is to show that this identification of functions is compatible with the topology and sheaf structure of $\mathbb{P}^n$. In other words, we must show that the graded algebra of the ring $\mathbb{K}[x_0, \ldots, x_n]$ coincides with the (non-graded) algebra of the ring $\mathcal{O}_{\mathbb{P}^n}(D_{+,x_i})$ along the open set D_{+,x_i}.

Proposition 2.2.3. *Set $S = \mathbb{K}[x_0, \ldots, x_n]$. The function $f : D_{+,x_i} \to \mathbb{A}^n$ defined via the rule*

$$\mathfrak{m} \mapsto \mathfrak{m}_{x_i} \cap (S_{x_i})_0$$

is a homeomorphism. Furthermore, f identifies the sheaf $\mathcal{O}_{\mathbb{P}^n}|_{D_{+,x_i}}$ with the sheaf $\mathcal{O}_{\mathbb{A}^n}$.

Proof. We first check that f is a well-defined and bijective function. Recall that pullback under $S \to S_{x_i}$ yields a bijection between the prime ideals of S_{x_i} and the prime ideals of S which do not contain x_i. This bijection preserves the homogeneity of ideals. Thus the map $\mathfrak{m} \leftrightarrow \mathfrak{m}_{x_i}$ gives a bijection between the points of D_{+,x_i} and the homogeneous ideals in S_{x_i} which are maximal amongst all homogeneous ideals. (Note that the irrelevant ideal in S contains x_i and thus becomes identified with the entire ring S_{x_i}, so it is not associated with a "larger" homogeneous ideal.)

We claim that the homogeneous ideals in S_{x_i} which are maximal amongst all homogeneous ideals are exactly the sets of the form

$$\ldots \mathfrak{n}x_i^{-2} \oplus \mathfrak{n}x_i^{-1} \oplus \mathfrak{n} \oplus \mathfrak{n}x_i \oplus \mathfrak{n}x_i^2 \oplus \ldots \tag{2.2.1}$$

for some maximal ideal $\mathfrak{n}$ in $(S_{x_i})_0$. Any such set is an ideal, and an ideal of this form is clearly a homogeneous ideal. It only remains to show that these are the maximal elements in the set of homogeneous ideals. For any homogeneous ideal $I \subsetneq S_{x_i}$ the set $I \cap (S_{x_i})_0$ must be a proper ideal of $(S_{x_i})_0$. Let $\mathfrak{n}$ be any maximal ideal containing it. If $I \cap (S_{x_i})_d$ failed to be contained in $\mathfrak{n}(S_{x_i})_d$ then by multiplying by x_i^{-d} we would obtain an element of $I \cap (S_{x_i})_0$ not contained in $\mathfrak{n}$, an impossibility. We conclude that every homogeneous ideal is contained in an ideal described by Equation (2.2.1). On the other hand no two different sets described by Equation (2.2.1) will contain each other. Altogether we conclude that f is a well-defined and bijective function.

We next show that f is a homeomorphism. Let $I \subset S$ be a homogeneous ideal; its localization I_{x_i} is still homogeneous. The set of homogeneous ideals in S_{x_i} which are maximal amongst all homogeneous ideals and which contain I_{x_i} is in bijection with $D_{+,x_i} \cap V_+(I)$. By the argument in the previous paragraph, an ideal described by Equation (2.2.1) will contain I_{x_i} if and only if

$$\mathfrak{n} \supset I_{x_i} \cap (S_{x_i})_0.$$

This shows that f maps $V_+(I) \cap D_{+,x_i}$ to $V(I_{x_i} \cap (S_{x_i})_0)$, finishing the proof that f is a homeomorphism.

It only remains to check that the sheaves of functions are the same. The easiest way to make the comparison is to use the base coming from distinguished open affines. Suppose that $D_{+,g}$ is a distinguished open affine contained in $D_{+,f}$. This implies that f becomes a unit when we localize by g, giving us a canonical map $\mathbb{K}[x_0, \ldots, x_n]_f \to \mathbb{K}[x_0, \ldots, x_n]_g$ and hence also a canonical map on the degree 0 parts

$$(\mathbb{K}[x_0, \ldots, x_n]_f)_0 \to (\mathbb{K}[x_0, \ldots, x_n]_g)_0$$

This latter map agrees with localization along the element $g^{\deg f}/f^{\deg g}$. In this way we see that $\mathcal{O}_{\mathbb{P}^n}(D_{+,g})$ can be obtained from $\mathcal{O}_{\mathbb{P}^n}(D_{+,f})$ by localizing along a function which vanishes along its complement. This is the same description as the ring of functions for a distinguished open affine inside of an affine scheme. We conclude that these two sheaves are isomorphic along each distinguished open affine and thus, by Exercise 1.9.10, isomorphic everywhere. □

Since the topology and the sheaf of functions on D_{+,x_i} match up, it is fair to say (by analogy with Definition 1.10.8) that D_{+,x_i} is an open affine subset of $\mathbb{P}^n$. It turns out that the analogous statement is true for any distinguished open affine in $\mathbb{P}^n$, justifying the terminology.

Proposition 2.2.4. *For every homogeneous $f \in \mathbb{K}[x_0, \dots, x_n]$ the distinguished open affine $D_{+,f}$ is homeomorphic to* $\mathrm{mSpec}(\mathcal{O}_{\mathbb{P}^n}(D_{+,f}))$ *and the homeomorphism identifies the sheaves of functions on the two spaces.*

The map is the same as in Proposition 2.2.3: a point $\mathfrak{m} \in D_{+,f}$ is identified with $\mathfrak{m}_f \cap (\mathbb{K}[x_0, \dots, x_n]_f)_0$. We will not give the proof here, deferring it to the next section (Proposition 2.3.8). Proposition 2.2.4 has an important consequence:

Corollary 2.2.5. *The complement of a hypersurface in $\mathbb{P}^n$ is an affine variety.*

2.2.2 Working with affine charts

It is important to be able to pass explicitly back and forth projective space and affine charts. Let's review how this correspondence works according to Proposition 2.2.3.

- Let $I \subset \mathbb{K}[x_0, \dots, x_n]$ be a homogeneous ideal. Define the ideal $J \subset \mathbb{K}[\frac{x_0}{x_i}, \dots, \frac{x_n}{x_i}]$ by dividing every degree d generator of I by x_i^d. Then $V_+(I) \cap D_i$ is the vanishing locus of the ideal J.

Example 2.2.6. Consider the line $ax + by + cz = 0$ in $\mathbb{P}^2$. The intersection of this line with $D_{+,x}$ is defined by $a + b\frac{y}{x} + c\frac{z}{x} = 0$ in $\mathbb{K}[\frac{y}{x}, \frac{z}{x}]$.

More generally, a closed subset $X \subset \mathbb{P}^n$ will be linear if and only if the intersection of X with every chart D_{+,x_i} is an affine linear subset of $\mathbb{A}^n$.

- Let $J \subset \mathbb{K}[\frac{x_0}{x_i}, \dots, \frac{x_n}{x_i}]$ be an ideal. We define the homogenization of J to be the ideal I constructed by multiplying every generator of J by the smallest power of x_i that yields a polynomial. The closure of $V(J) \subset D_i$ in $\mathbb{P}^n$ is the vanishing locus of I.

The points added by taking the closure will be given by the intersection of $V_+(x_i)$ with $V_+(I)$.

Example 2.2.7. Consider the line $ax + by + c = 0$ in $\mathbb{A}^2$. The closure in $\mathbb{P}^2$ is the line $ax + by + cz = 0$. By taking a closure we have added the single point $(a : b : 0)$ recording the "asymptotic limit" of the line. Note that parallel lines yield the same "limit point" at infinity.

2.2.3 Exercises

Exercise 2.2.8. For each of the following closed subsets of $\mathbb{A}^2$, take the closure in $\mathbb{P}^2$ and identify explicitly which points are added when we take the closure.

(1) $x^2 + y^2 = 1$ in $\mathbb{A}^2$.

(2) $x^2 = y$ in $\mathbb{A}^2$.

(3) $y^2 = f(x)$ in $\mathbb{A}^2$.

Exercise 2.2.9. Let D_{+,x_i}, D_{+,x_j} be two different affine charts in $\mathbb{P}^n$. We have isomorphisms $\phi_i : \mathbb{A}^n \to D_{+,x_i}$ and $\phi_j : \mathbb{A}^n \to D_{+,x_j}$. Let $U_{ij} \subset \mathbb{A}^n$ be the preimage of $D_{+,x_i} \cap D_{+,x_j}$ under ϕ_i and define U_{ji} analogously. What are U_{ij}, U_{ji} explicitly? What is the map $\phi_j^{-1} \circ \phi_i : U_{ij} \to U_{ji}$?

2.3 mProj construction

We next describe how to construct a geometric object from an arbitrary finitely generated graded $\mathbb{K}$-algebra. Our construction will mimic the scheme-theoretic construction of $\mathbb{P}^n$ from the ring $\mathbb{K}[x_0, \dots, x_n]$. Note that we do not insist that our graded ring R be a quotient of a polynomial ring, nor do we insist that the degree 0 part be isomorphic to $\mathbb{K}$.

Definition 2.3.1. Let R be a finitely generated graded $\mathbb{K}$-algebra. (Remember this means that the canonical copy of $\mathbb{K}$ is contained in R_0.) A homogenous ideal $\mathfrak{m} \in R$ is said to be almost maximal if it is prime and it is a maximal element of the set of all prime homogeneous ideals which do not contain $R_{>0}$. We define $\mathrm{mProj}(R)$ to be the set of almost maximal homogeneous ideals in R.

Given any homogeneous ideal I, we define $V_+(I)$ to be the set of almost maximal homogeneous ideals containing I. The Zariski topology is the topology whose closed sets are the vanishing loci $V_+(I)$.

Exercise 2.3.2. Verify carefully that the proposed definition for the Zariski topology actually is a topology.

Exercise 2.3.3. Let $X \subset \mathbb{P}^n$ be a closed subset defined by the vanishing locus of a homogenous ideal $I \subset \mathbb{K}[x_0, \dots, x_n]$. Prove that X is homeomorphic to $\mathrm{mProj}(\mathbb{K}[x_0, \dots, x_n]/I)$.

Exercise 2.3.4. Suppose that $R = R_0[x]$ where R_0 is a finitely generated $\mathbb{K}$-algebra and x lies in degree 1. Show that $\mathrm{mProj}(R) \cong \mathrm{mSpec}(R_0)$.

2.3.1 Functions on mProjs

There are several ways to define the sheaf of functions on a mProj. Just like for affine schemes, we can no longer appeal to the existence of a "function field" and so we will construct the sheaf of functions via the base of distinguished open affines via a gluing argument.

Definition 2.3.5. Let R be a finitely generated graded $\mathbb{K}$-algebra. For any homogeneous $f \in R$ with degree > 0 we define the distinguished open affine $D_{+,f}$ in $\mathrm{mProj}(R)$ to be the complement of $V_+(f)$.

Exercise 2.3.6. Show that distinguished open affines form a base of the topology on $\mathrm{mProj}(R)$.

The following lemma is the key tool for working with distinguished open affines.

Lemma 2.3.7. *Let S be a $\mathbb{Z}_{\geq 0}$-graded ring and let f be a homogeneous element of S of degree > 0. The assignment*

$$\begin{array}{ccc} \mathfrak{p} \subset S & \xrightarrow{\textit{localization}} & \mathfrak{p}_f \subset S_f \\ & & \Big\downarrow \textit{intersection} \\ & & \mathfrak{p}_f \cap (S_f)_0 \subset (S_f)_0. \end{array}$$

defines a bijection between the homogeneous prime ideals in S not containing f and the prime ideals of $(S_f)_0$.

Proof. Standard properties of localization yield a bijection between the prime homogeneous ideals in S not containing f and the prime homogeneous ideals in S_f. It only remains to show that intersection defines a bijection between the prime homogeneous ideals in S_f and the prime ideals of $(S_f)_0$. Using the injection $(S_f)_0 \to S_f$ it is clear that the preimage of any homogeneous prime ideal is prime. Conversely, given a prime ideal $\mathfrak{p} \subset (S_f)_0$, consider the subset

$$\dots Q_{-2} \oplus Q_{-1} \oplus Q_0 \oplus Q_1 \oplus Q_2 \oplus \dots$$

where Q_i is the set of elements $h \in (S_f)_i$ such that $h^{\deg f}/f^i \in \mathfrak{p}$. By primality $Q_0 = \mathfrak{p}$. It is clear that if $h \in Q_i$ and $h' \in (S_f)_j$ then $hh' \in Q_{i+j}$. Furthermore, if $h_1, h_2 \in Q_i$ then

$$\frac{(h_1 + h_2)^{2\deg(f)}}{f^{2i}} = \frac{\sum_{i=0}^{2\deg(f)} h_1^i h_2^{2\deg(f)-i}}{f^{2i}} \in \mathfrak{p}$$

since every term of the sum is divisible by either $h_1^{\deg f}$ or $h_2^{\deg f}$. We deduce that $(h_1+h_2)^2 \in Q_{2i}$, and thus by primality of $\mathfrak{p}$ we have $h_1+h_2 \in Q_i$. Altogether this shows that the Q_i are the components of a homogeneous ideal $\mathfrak{q} \subset S_f$. Since primality of a homogeneous ideal can be detected using only homogeneous elements, the primality of $\mathfrak{q}$ follows directly from the primality of $\mathfrak{p}$. Thus we have verified that intersection does indeed define a bijection between the prime homogeneous ideals in S_f and the prime ideals of $(S_f)_0$. $\square$

Proposition 2.3.8. *Let $D_{+,f}$ be a distinguished open affine in* $\mathrm{mProj}(R)$*. Then $D_{+,f}$ is homeomorphic to* $\mathrm{mSpec}((R_f)_0)$ *via the correspondence*

$$\begin{array}{ccc} \mathfrak{m} \subset R & \xleftrightarrow{\textit{localization}} & \mathfrak{m}_f \subset R_f \\ & & \Big\updownarrow \textit{intersection} \\ & & \mathfrak{m}_f \cap (R_f)_0 \subset (R_f)_0. \end{array}$$

Exercise 2.3.9. Prove Proposition 2.3.8 by mimicking the proof of Proposition 2.2.4.

This homeomorphism allows us to make the following definition.

Definition 2.3.10. Let R be a finitely generated graded $\mathbb{K}$-algebra. For every distinguished open affine $D_{+,f}$ in R we let $\widetilde{\mathcal{O}}(D_{+,f})$ denote be the degree 0 part of the localization of R along all homogeneous elements h such that $D_{+,f} \cap V_+(h) = \emptyset$. We define $\widetilde{\rho}_{D_{+,f},D_{+,g}}$ using the universal property of localization.

We then define the structure sheaf $\mathcal{O}_{\mathrm{mProj}(R)}$ by applying Theorem 1.9.2 using the base of distinguished open affines equipped with $\widetilde{O}$, $\widetilde{\rho}$ as constructed above.

Exercise 2.3.11. Verify carefully that the construction above gives a sheaf on $\mathrm{mProj}(R)$. (Hint: show that if $D_{+,f} \subset D_{+,g}$ then there is a commutative diagram

$$\begin{tikzcd} D_{+,f} \arrow[r, "\cong"] \arrow[d] & \mathrm{mSpec}((R_f)_0) \arrow[d] \\ D_{+,g} \arrow[r, "\cong"] & \mathrm{mSpec}((R_g)_0). \end{tikzcd}$$

where the map on the right is induced by localizing $(R_g)_0$ along $f^{\deg g}/g^{\deg f}$. Then verify that the desired gluing property matches the sheaf property for affine schemes.)

Exercise 2.3.12. Let R be a finitely generated graded $\mathbb{K}$-algebra. Since an almost maximal homogeneous ideal $\mathfrak{m}$ is prime, the set S of homogeneous elements not contained in $\mathfrak{m}$ is multiplicatively closed. Show that the stalk of $\mathcal{O}_{\mathrm{mProj}(R)}$ at a point $\mathfrak{m}$ is $(R_S)_0$.

Warning 2.3.13. Let R be a finitely generated graded $\mathbb{K}$-algebra. In general it is not possible to recover R from the topology and sheaf of functions on $\mathrm{mProj}(R)$: there can be many different graded rings which have the same mProj. (Note the contrast with the mSpec construction for affine schemes.)

The key issue is that the mProj construction comes with an extra piece of data: it actually returns both a scheme X and (if R is generated in degree 1) a line bundle $\mathcal{L}$ on X. (In the example $\mathrm{mProj}(\mathbb{K}[x_1, \ldots, x_n])$ this line bundle $\mathcal{L}$ is the dual of the tautological line bundle on $\mathbb{P}^n$.) Even when we fix both the scheme X and the line bundle $\mathcal{L}$, it is still not true that we can recover the graded ring we started with; however see Remark 2.7.9.

2.3.2 Structural maps

Suppose that R is a finitely generated graded $\mathbb{K}$-algebra. The degree 0 piece R_0 plays an important role in understanding the geometry of $\mathrm{mProj}(R)$.

Consider the distinguished open affine $D_{+,f} \subset \mathrm{mProj}(R)$. Note that the localization map yields a natural inclusion $R_0 \hookrightarrow \mathcal{O}_{\mathrm{mProj}(R)}(D_{+,f})$. Furthermore these maps commute with further localizations, in the sense that if $D_{+,g} \subset D_{+,f}$ we have commutative diagrams

$$\begin{tikzcd} R_0 \arrow[r] \arrow[dr] & \mathcal{O}_{\mathrm{mProj}(R)}(D_{+,f}) \arrow[d, "restriction"] \\ & \mathcal{O}_{\mathrm{mProj}(R)}(D_{+,g}). \end{tikzcd}$$

By the gluing property we conclude that for every open subset U there is an inclusion $R_0 \hookrightarrow \mathcal{O}_{\mathrm{mProj}(R)}(U)$. In particular:

- R_0 injects into the space of global sections on $\mathrm{mProj}(R)$.

R_0 need not be isomorphic to $\mathcal{O}_{\mathrm{mProj}(R)}(\mathrm{mProj}(R))$. However, it turns out that the global sections $\mathcal{O}_{\mathrm{mProj}(R)}(\mathrm{mProj}(R))$ will always be a module-finite extension of R_0. For now we will only use a weaker property:

Theorem 2.3.14. *Let R be a finitely generated graded $\mathbb{K}$-algebra. Then $\mathcal{O}_{\mathrm{mProj}(R)}(\mathrm{mProj}(R))$ is a finitely generated $\mathbb{K}$-algebra.*

Unfortunately I do not know of any easy proof of this statement; we will only see a proof (much) later in the course.

Here is the geometric interpretation of R_0. For any distinguished open affine $D_{+,f}$ the inclusion $R_0 \hookrightarrow \mathcal{O}_{\mathrm{mProj}(R)}(D_{+,f})$ determines a dominant morphism $D_{+,f} \to \mathrm{mSpec}(R_0)$ of affine schemes. Furthermore, given any two distinguished open affines $D_{+,f}, D_{+,g}$ the two induced functions on the intersection are the same. Thus we can conclude:

- There is a continuous function $p : \mathrm{mProj}(R) \to \mathrm{mSpec}(R_0)$ with dense image.

After we define morphisms of quasiprojective schemes in the next section, we will recognize that p is actually a morphism in our category. In fact, it turns out that the function p is surjective; see Exercise 2.11.9.

Conceptually it is best to think of the map $\mathrm{mProj}(R) \to \mathrm{mSpec}(R_0)$ as part of the data of the mProj construction. In other words, mProj is a “relative construction” over $\mathrm{mSpec}(R_0)$. As we will see later on, it is only when we use this perspective that we can identify the “natural” geometric properties of mProj.

2.3.3 Exercises

Exercise 2.3.15. Consider the distinguished open affine $D_{+,x^2+y^2+z^2}$ in $\mathbb{P}^2$. Identify explicitly the finitely generated $\mathbb{K}$-algebra R such that $D_{+,x^2+y^2+z^2} \cong \mathrm{mSpec}(R)$ by writing R as a quotient of a polynomial ring.

Exercise 2.3.16. Let R be a finitely generated graded $\mathbb{K}$-algebra. Show that a closed subset $X \subset \mathrm{mProj}(R)$ is irreducible if and only if there is a homogeneous prime ideal I such that $X = V(I)$. Conclude that every closed subset in $\mathrm{mProj}(R)$ is a finite union of the vanishing loci of homogeneous prime ideals.

Exercise 2.3.17. Let R be a finitely generated graded $\mathbb{K}$-algebra. Prove that $\mathrm{mProj}(R)$ is a Noetherian topological space.

Exercise 2.3.18. Compute $\mathrm{mProj}(\mathbb{K}[x,y]/(xy))$. (Check to make sure that your algebraic computation matches your geometric intuition concerning $V_+(xy)$ in $\mathbb{P}^1$.)

Exercise 2.3.19. Compute $\mathrm{mProj}(\mathbb{K}[x,y])$ where x lies in degree 1 and y lies in degree 2. (Choose a covering of this Proj by affine charts and describe how these charts glue. Do you recognize the resulting variety?)

Exercise 2.3.20. Consider $X = \mathrm{mProj}(\mathbb{K}[x_0,x_1,x_2])$ where x_0, x_1 are in degree 1 and x_2 is in degree 2. Prove that X is isomorphic to the quadric cone in $\mathbb{P}^3_{\mathbb{K}}$ defined by the equation $y^2 - xz$.

Exercise 2.3.21. Let $U \subset \mathbb{P}^2_{\mathbb{K}}$ be the complement of a $\mathbb{K}$-point. Prove that $\mathcal{O}_{\mathbb{P}^2}(U) \cong \mathbb{K}$.

2.4 Quasiprojective schemes and projective schemes

A projective scheme X is a closed set of projective space equipped with a scheme structure. Just as with affine schemes, we will formulate projective schemes in an "embedding free" way. Actually, our construction will yield a more general class of schemes than just projective schemes.

2.4.1 Quasiprojective schemes

The following definition identifies the most general class of geometric object that we will work with.

Definition 2.4.1. A quasiprojective $\mathbb{K}$-scheme X is an open subset of some $\mathrm{mProj}(R)$ where R is a finitely generated graded $\mathbb{K}$-algebra. It comes equipped with the following data:

$$(\text{set, topology, sheaf of functions}) = (X, \mathbf{Zar}_{\mathrm{mProj}(R)}|_X, \mathcal{O}_{\mathrm{mProj}(R)}|_X).$$

In particular, each $\mathrm{mProj}(R)$ is an example of a quasiprojective scheme. (Despite the notation, $\mathrm{mProj}(R)$ need not be a projective scheme; see Section 2.4.2.)

Exercise 2.4.2. Verify that every quasiaffine $\mathbb{K}$-scheme is also a quasiprojective $\mathbb{K}$-scheme.

We say that an open subset $U \subset X$ is an open affine if U is homeomorphic to $\mathrm{mSpec}(\mathcal{O}_X(U))$ and this homeomorphism identifies $\mathcal{O}_X|_U$ with the sheaf of functions on the mSpec. Note that every distinguished open affine is an open affine in this sense so that such opens form a base for the Zariski topology on X.

Definition 2.4.3. Suppose that X is a quasiprojective scheme. For any open subset $U \subset X$ and any element $g \in \mathcal{O}_X(U)$ we define $V(g) = \{x \in U | \rho_{U,x}(g) \in \mathfrak{m}_x\}$.

Using Definition 2.4.3, we define morphisms in the same way that we did for quasiaffine schemes:

Definition 2.4.4. A morphism $f : X \to Y$ of quasiprojective schemes consists of the following data:

- A continuous set-theoretic map $f : X \to Y$.
- A pullback map of sheaves $f^\sharp : \mathcal{O}_Y \to \mathcal{O}_X$.

which satisfy the following compatibility criterion:

(*) For any open subset $V \subset Y$ and for any ideal $I \subset \mathcal{O}(V)$ we have $f^{-1}(V(I)) = V(\langle f^\sharp(I) \rangle)$ inside $f^{-1}(V)$.

Remark 2.4.5. We could also replace (*) by condition (*') as in Exercise 1.12.5.

In this way we get a category **QProSch**/$\mathbb{K}$. Note that **AffSch**/$\mathbb{K}$ is a full subcategory.

Example 2.4.6. Suppose X is a quasiprojective $\mathbb{K}$-scheme. In contrast to affine schemes and the mProj construction, it is possible that $\mathcal{O}_X(X)$ is not finitely generated over $\mathbb{K}$. For example, suppose that W is the union of two planes P_1, P_2 in $\mathbb{P}^3$ which meet along a line ℓ. Let ℓ' be a line in P_1 that is different from ℓ (so ℓ and ℓ' meet along a point p) and set $X = W \backslash \ell'$.

We can think of $\mathcal{O}_X(X)$ as the combination of all functions on $P_1 \backslash \ell' \cong \mathbb{A}^2$ and all functions on $P_2 \backslash p$ whose restrictions to $\ell \backslash p$ agree. (Verify this carefully!) By Exercise 2.3.21, the only functions on $P_2 \backslash p$ are the constant functions. Thus $\mathcal{O}_X(X)$ consists of all functions on $\mathbb{A}^2$ which restrict to a constant function on a line in $\mathbb{A}^2$. Without loss of generality we can identify this line as $V(y)$, in which case $\mathcal{O}_X(X)$ is the subring of $\mathbb{K}[x, y]$ generated as a $\mathbb{K}$-algebra by 1 and by monomials of the form $x^m y^{1+n}$ with $m, n \geq 0$.

2.4.2 Projective schemes

The most important case of the Proj construction is when the degree 0 piece of R is exactly $\mathbb{K}$ (and not some larger ring).

Definition 2.4.7. A quasiprojective scheme X is projective if there is a finitely generated graded $\mathbb{K}$-algebra R satisfying $R_0 \cong \mathbb{K}$ such that X is isomorphic to $\mathrm{mProj}(R)$.

More generally, we will see later that for any finitely generated graded $\mathbb{K}$-algebra R the structural morphism $\mathrm{mProj}(R) \to \mathrm{mSpec}(R_0)$ has projective fibers (Exercise 2.10.16). In other words, we can think of $\mathrm{mProj}(R)$ as being "projective over $\mathrm{mSpec}(R_0)$" – this motivates the notation mProj.

2.4.3 Morphisms to affine schemes

A morphism of affine schemes $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ is determined by a $\mathbb{K}$-algebra homomorphism $f^\sharp : S \to R$. The following important result gives a related statement for all quasiprojective schemes.

Theorem 2.4.8. *Let X be a quasiprojective scheme and $\mathrm{mSpec}(S)$ an affine scheme. There is a bijection between morphisms $f : X \to \mathrm{mSpec}(S)$ and $\mathbb{K}$-algebra homomorphisms $f^\sharp : S \to \mathcal{O}_X(X)$.*

Proof. Fix an open cover $\{U_i\}$ of X by open affines. Consider the commutative diagram

$$\begin{tikzcd}
\mathrm{Hom}(X, \mathrm{mSpec}(S)) \arrow[r] \arrow[d, "\rho_1"] & \mathrm{Hom}(S, \mathcal{O}_X(X)) \arrow[d, "\rho_2"] \\
\prod_i \mathrm{Hom}(U_i, \mathrm{mSpec}(S)) \arrow[r] & \prod_i \mathrm{Hom}(S, \mathcal{O}_X(U_i)).
\end{tikzcd}$$

Since a continuous map is determined by its restrictions to an open cover, the map ρ_1 is injective. It is clear that the image of ρ_1 is contained in the set

$$\{\, (f_i : U_i \to \mathrm{mSpec}(S)) \,|\, f_i|_{U_i \cap U_j} = f_j|_{U_i \cap U_j} \}$$

By Exercise 2.4.16 ρ_1 is a bijection onto this set.

Consider now the map ρ_2. By the gluing property for sheaves we know that $\mathcal{O}_X(X)$ is the same as

$$\{\, (g_i \in \mathcal{O}_X(U_i)) \,|\, g_i|_{U_i \cap U_j} = g_j|_{U_i \cap U_j} \}$$

Thus, ρ_2 takes $\mathrm{Hom}(S, \mathcal{O}_X(X))$ bijectively to the subset of $\prod_i \mathrm{Hom}(S, \mathcal{O}_X(U_i))$ defined by

$$\{\, (f_i^\sharp : S \to \mathcal{O}_X(U_i)) \,|\, f_i^\sharp \circ \rho_{ij} = f_j^\sharp \circ \rho_{ji} \}.$$

By Proposition 1.12.8 the map along the bottom of the diagram is a bijection. Furthermore, it identifies the image of ρ_1 with the image of ρ_2. Altogether we see that the map along the top of the diagram is a bijection. □

2.4.4 Open and closed embeddings

Next we discuss how to give open and closed sets of a quasiprojective scheme the structure of a quasiprojective scheme.

Definition 2.4.9. A morphism $f : X \to Y$ of quasiprojective schemes is an open embedding if f defines an isomorphism from X to an open subset of Y and the structure sheaf of X is isomorphic to the restriction of the structure sheaf to this open subset.

We will call X an open subscheme of Y if additionally f is an inclusion map.

The terminology is slightly confusing, but not in an interesting way: an open embedding $f : X \to Y$ defines an isomorphism between X and an open subscheme of Y. Thus for many purposes the terms open embedding/subscheme can be used interchangeably.

Exercise 2.4.10. Prove that a composition of open embeddings is an open embedding.

It is somewhat harder to work with closed subsets. Although we give define closed embedding/subscheme here, we will not have the tools to work with them until Section 2.8.

Definition 2.4.11. A morphism $f : X \to Y$ of quasiprojective schemes is a closed embedding if f is a homeomorphism onto a closed subset of Y and for every open affine $V \subset Y$ the map $f^\sharp(V) : \mathcal{O}_Y(V) \to \mathcal{O}_X(f^{-1}(V))$ is surjective.

We call X a closed subscheme of Y if additionally f is an inclusion map.

The relationship between closed embeddings and closed subschemes is the same as the relationship between open embeddings and open subschemes: a closed embedding $f : X \to Y$ defines an isomorphism between X and a closed subscheme of Y. However, this relationship is no longer immediate from the definitions and we defer the proof to Section 2.8.

Warning 2.4.12. If $f : X \to Y$ is a closed embedding we cannot conclude that $f^\sharp(U) : \mathcal{O}_Y(U) \to \mathcal{O}_X(f^{-1}(U))$ is surjective for every open set $U \subset Y$. For example, consider the closed embedding $f : X \to \mathbb{P}^1$ where X is the disjoint union of two $\mathbb{K}$-points. Then $\mathcal{O}_{\mathbb{P}^1}(\mathbb{P}^1) \cong \mathbb{K}$ but $\mathcal{O}_X(f^{-1}(\mathbb{P}^1)) = \mathcal{O}_X(X) \cong \mathbb{K}^2$. Surjectivity does not behave compatibly with gluing!

While our definition of projective and quasiprojective schemes is very explicit, it has the disadvantage of obscuring the geometric significance. The following theorem clarifies the geometric meaning of projectiveness and quasiprojectiveness:

Theorem 2.4.13. *A quasiprojective scheme is projective if and only if admits a closed embedding to some projective space.*

Every quasiprojective scheme admits an open embedding to a projective scheme.

The proof will be given in Theorem 2.8.12 and in Theorem 2.10.15.

Remark 2.4.14. Note the definition is exactly analogous to the affine case: an affine scheme is a closed subscheme of affine space and a quasiaffine scheme is an open subscheme of an affine scheme.

2.4.5 Exercises

Exercise 2.4.15. Set $X = \mathrm{mProj}(\mathbb{K}[x, y, z])$ where x, y lie in degree 1 and z lies in degree 0. Consider the structural morphism $p : X \to \mathbb{A}^1$. Prove that every fiber of π is isomorphic to $\mathbb{P}^1_{\mathbb{L}}$ for some finite extension $\mathbb{L}/\mathbb{K}$. (We will later identify X as the product of $\mathbb{P}^1$ and $\mathbb{A}^1$.)

Exercise 2.4.16. Let X, Y be quasiprojective schemes. Suppose that we have an open cover $\{U_i\}_{i\in I}$ of X and morphisms of quasiprojective schemes $f_i : U_i \to Y$. Suppose furthermore that for all $i, j \in I$ we have $f_i|_{U_i\cap U_j} = f_j|_{U_i\cap U_j}$ (as functions equipped with pullback maps of sheaves).

Prove that there is a unique morphism $f : X \to Y$ such that for every $i \in I$ we have $f|_{U_i} = f_i$. (Hint: apply Proposition 1.9.7. The main point is to verify the extra condition (*).)

Exercise 2.4.17. Construct the quotient morphism $\mathbb{A}^{n+1}\backslash\{0\} \to \mathbb{P}^n$ as a morphism of quasiprojective schemes. (Hint: one way is to define this morphism on affine charts and show that these morphisms are compatible on the overlaps as in Exercise 2.4.16.)

Exercise 2.4.18. Use Chevalley's Theorem for affine schemes to prove that for any morphism $f : X \to Y$ of quasiprojective schemes the image of a constructible subset of X is a constructible subset of Y.

Exercise 2.4.19. Let X be a quasiprojective scheme. Let $U \subset X$ be an open subset and fix $f \in \mathcal{O}_X(U)$. Let U_f denote the complement of the vanishing locus of f in U. Prove that $\mathcal{O}_X(U_f) = \mathcal{O}_X(U)_f$. (We have already proved this when U is an open affine; the goal is to extend this property to arbitrary open sets.)

2.5 Properties of quasiprojective schemes

The goal of this section is to extend several definitions and constructions from the setting of affine schemes to the setting of quasiprojective schemes.

2.5.1 Local properties

Suppose that P is a property we have defined for affine schemes. We would like to say that a quasiprojective scheme X satisfies P if X admits an open cover of distinguished open affines, each of which satisfies property P. In order for this to be a well-behaved notion, we need to verify that property P is independent of the choice of affine cover.

It turns out that these "well-behaved" properties P are precisely those which are compatible with localization. (Since so many properties of rings are compatible with localization, we can obtain many properties of schemes in this way!)

Definition 2.5.1. Let P be a property of rings. We say that P is a local property if for any finitely generated $\mathbb{K}$-algebra R the following conditions are equivalent:

(1) R satisfies P.

(2) R_f satisfies P for every $f \in R$.

(3) There is a finite set of elements f_i which generate R such that R_{f_i} satisfies P for every i.

The following lemma is crucial for passing information from one open affine to another.

Lemma 2.5.2 (Nike's lemma)**.** *Let X be a quasiprojective scheme. Suppose that U and V are open affines in X. Then $U \cap V$ admits a cover by open sets which are simultaneously distinguished open affines in both U and V.*

Proof. Fix any point $x \in U \cap V$. Choose an $f \in \mathcal{O}_X(U)$ such that the corresponding distinguished open affine W is an open neighborhood of x contained in $U \cap V$. Choose a $g \in \mathcal{O}_X(V)$ such that the corresponding distinguished open affine D is an open neighborhood of x contained in W. Let g' denote the image of g under the restriction map $\mathcal{O}_X(V) \to \mathcal{O}_X(W)$.

Note that D is the distinguished open affine in W corresponding to g'; indeed, according to Exercise 1.11.13 the vanishing locus can be detected by restricting to stalks and by construction the restrictions of g' and g to every stalk in W agree. We can write $g' = g''/f^n$ for some $g'' \in R$ and some integer n. Then D is the distinguished open affine in U which is the complement of $V(fg'')$. $\square$

As a consequence, we obtain:

Theorem 2.5.3 (Affine Communication Lemma)**.** *Let P be a local property of finitely generated $\mathbb{K}$-algebras. Let X be a quasiprojective scheme. Then every open affine in X satisfies P if and only if X admits an open cover by open affines which satisfy P.*

Exercise 2.5.4. Prove the Affine Communication Lemma.

There are some properties of rings which satisfy a different chain of equivalences.

Definition 2.5.5. Let P be a property of rings. We say that P is a max-local property if for any finitely generated $\mathbb{K}$-algebra R the following conditions are equivalent:

(1) R satisfies P.

(2) $R_{\mathfrak{m}}$ satisfies P for every maximal ideal $\mathfrak{m} \in R$.

Proposition 2.5.6. *Let X be a quasiprojective scheme. Suppose that P is a max-local property. Then the following are equivalent:*

(1) There is an open cover of X by open affines $\{U_i\}$ such that each $\mathcal{O}_X(U_i)$ satisfies P.

(2) Every stalk $\mathcal{O}_{X,x}$ satisfies P.

(3) For every open affine $\{U\}$ we have that $\mathcal{O}_X(U)$ satisfies P.

In this case we say that X satisfies property P.

Proof. (1) $\implies$ (2): Fix a point $x \in X$. Since $\{U_i\}$ is an open cover, $x \in U_j$ for some j. By Proposition 1.11.9 the stalk $\mathcal{O}_{X,x}$ is a localization of $\mathcal{O}_X(U_j)$ along a maximal ideal. Thus $\mathcal{O}_{X,x}$ satisfies P.

(2) $\implies$ (3): Fix an open affine U. For every $x \in U$ Proposition 1.11.9 shows that the stalk $\mathcal{O}_{X,x}$ is the localization of $\mathcal{O}_X(U)$ along the corresponding maximal ideal. Thus P holds for $\mathcal{O}_X(U)$.

(3) $\implies$ (1): Immediate. □

2.5.2 Irreducible and reduced

Exercise 2.5.7. Prove that every quasiprojective scheme is a Noetherian topological space.

In particular, by Exercise 1.3.16 every quasiprojective scheme can be written as a finite union of irreducible components. When we are working with $\mathrm{mProj}(R)$, the irreducible components will be determined by the minimal homogeneous prime ideals. We say that a quasiprojective scheme is irreducible if it has a unique irreducible component.

Reducedness takes a little more work. The key is:

Exercise 2.5.8. Let R be a finitely generated $\mathbb{K}$-algebra. The property $\mathrm{Nil}(R) = 0$ is a stalk-local property.

As discussed above, this gives us a very clean way to work with the reduced property for quasiprojective schemes: we define reducedness using local rings.

Definition 2.5.9. Let X be a quasiprojective scheme and let $x \in X$. We say that X is reduced at x if the nilradical of $\mathcal{O}_{X,x}$ is 0.

Exercise 2.5.10.
(1) Suppose that $X = \mathrm{mSpec}(R)$ is an affine scheme. Let $\{\mathfrak{p}_i\}$ be the set of associated primes for the zero ideal such that $R_{\mathfrak{p}_i}$ is not reduced. Prove that $x \in X$ is non-reduced if and only if it lies in $V(\mathfrak{p}_i)$ for some i.

(2) Suppose that X is an arbitrary quasiprojective scheme. Prove that the set of nonreduced points in X is closed.

We say that a quasiprojective scheme X is reduced if it is reduced at every local point. By Proposition 2.5.6, this is equivalent to saying that X admits an open covering by open affines which are reduced.

2.5.3 Varieties

Definition 2.5.11. A quasiprojective variety is a quasiprojective scheme that is both irreducible and reduced.

Exercise 2.5.12. Let X be a quasiprojective scheme. Show that X is a variety if and only if for every open subset $U \subset X$ we have that $\mathcal{O}_X(U)$ is a domain. (Warning: this is not equivalent to saying that $\mathcal{O}_{X,x}$ is a domain for every point x. Can you think of a counterexample?)

Exercise 2.5.13. Show that if R is a finitely generated graded $\mathbb{K}$-algebra which is a domain then $\mathrm{mProj}(R)$ is a quasiprojective variety.

For a quasiprojective variety X every local ring $\mathcal{O}_{X,x}$ is a domain. Note that if U is an open affine in X, then for any two points in U the fraction fields of the local rings will be isomorphic. Using a covering by open affines, we see that the fraction fields of *all* the local rings $\mathcal{O}_{X,x}$ will be isomorphic, and in particular every open affine $U \subset X$ will have isomorphic function fields.

Definition 2.5.14. Let X be a quasiprojective variety. The function field $\mathbb{K}(X)$ is defined to be the function field of any open affine $U \subset X$.

Just as we saw earlier for projective space in Definition 2.1.7, if $X \cong \mathrm{mProj}(R)$ for a domain R then the function field will be the degree 0 elements in the localization of R along all non-zero homogeneous elements.

Exercise 2.5.15. Prove that if X is a quasiprojective variety then we can use Definition 1.10.2 to define the structure sheaf on X.

2.5.4 Exercises

Exercise 2.5.16. Decide whether the following projective schemes are varieties. If not, are they reducible or non-reduced (or both)?

(1) $\mathrm{mProj}(\mathbb{K}[x, y, z]/(xy))$.

(2) $\mathrm{mProj}(\mathbb{K}[x, y, z]/(x^3 - yz^2, x^2 - yz))$.

(3) $\mathrm{mProj}(\mathbb{K}[w, x, y, z]/(xy - zw))$.

(4) $\mathrm{mProj}(\mathbb{K}[w, x, y, z]/(xy - zw, x^2 - wy))$.

It may be helpful to sketch a picture of these vanishing loci in affine charts.

Exercise 2.5.17. Let $U \subset \mathbb{P}^2$ be the complement of a point. Show that U is not a quasiaffine scheme. (U is also not a projective scheme because it is not proper; see Section 2.11.)

2.6 Rational maps

Suppose that X is a quasiprojective variety. Many of the properties of X can be detected on any non-empty open subset:

- Every non-empty open subset of X is dense.
- For any non-empty open affine subset $U \subset X$ we have $\mathbb{K}(U) = \mathbb{K}(X)$.

Based on these properties, it is reasonable to expect that much of the interesting information about X can be detected on a non-empty open subset. Here is another instance of this principle:

Lemma 2.6.1. *Let X, Y be quasiprojective varieties and let $f_1, f_2 : X \to Y$ be two morphisms. Suppose that there is a non-empty open subset $U \subset X$ such that $f_1|_U = f_2|_U$. Then $f_1 = f_2$.*

Proof. Note that U will intersect every open subset of X. It suffices to prove the theorem if we replace Y by a dense open set and X by its preimage. Furthermore, by Exercise 2.4.16 we can check the equality of f_1 and f_2 on any open cover of X. Altogether it suffices to prove the case when $X = \mathrm{mSpec}(R)$ and $Y = \mathrm{mSpec}(S)$ are affine varieties.

After shrinking U we may suppose that it is a distinguished open affine D_g of X. The morphisms $f_1^\sharp, f_2^\sharp : S \to R$ become identical if we compose with the localization map $R \to R_g$. Since R is a domain, this implies that $f_1^\sharp = f_2^\sharp$. □

This lemma indicates that all the interesting information about a morphism of quasiprojective varieties can be detected on an open subset. Our next definition codifies this principle.

Definition 2.6.2. Consider a tuple (X, Y, U, f) where X and Y are quasiprojective varieties, U is a non-empty open subset of X, and f is a morphism of quasiprojective schemes $f : U \to Y$. We say that two such tuples (X, Y, U, f) and (X, Y, U', f') are equivalent if $f|_{U \cap U'} = f'|_{U \cap U'}$.

A rational map $f : X \dashrightarrow Y$ is an equivalence class of tuples (X, Y, U, f).

Exercise 2.6.3. Check carefully that the relation described in Definition 2.6.2 is actually an equivalence relation.

This definition is particularly useful for projective varieties since (as we will see in Section 2.7) the natural maps between projective varieties are often only rational maps and not morphisms.

Remark 2.6.4. One can also define rational maps for arbitrary quasiprojective schemes. However, a morphism is no longer uniquely determined by its behavior on an arbitrary open subset. To recover this uniqueness property, we must ensure that the scheme-theoretic image of inclusion of the open set is the entire scheme; see Construction 2.8.11.

Note that the ambiguity about the locus U where f is defined is baked into the notation of a rational map. It turns out that every rational map admits a largest open set U where f is defined. This set is known as the "locus of definition" of f.

Exercise 2.6.5. Verify that any rational map is represented by a tuple (X, Y, U, f) such that for any other tuple (X, Y, U', f') in the same equivalence class we have $U \supset U'$. (Hint: see Exercise 2.4.16.)

Example 2.6.6. Recall that $\mathbb{P}^1$ parametrizes the space of lines in $\mathbb{A}^2$ through the origin. By sending any non-zero point in $\mathbb{A}^2$ to the unique line containing it, we get a rational map $\phi : \mathbb{A}^2 \dashrightarrow \mathbb{P}^1$ defined away from the origin.

This rational map is known as projection away from the origin. In terms of $\mathbb{K}$-points, this is exactly the map $(a, b) \mapsto (a : b)$. (In other words, the coordinates on $\mathbb{P}^1$ record the slope of the line connecting the point to the origin.)

We can define this function in a rigorous way using charts. Let's use coordinates x_0, x_1 on $\mathbb{A}^2$ (note the unusual indexing!) and y_0, y_1 on $\mathbb{P}^1$. Then the map $D_{x_0} \to D_{+,y_0}$ is defined by sending $\frac{y_0}{y_1} \mapsto \frac{x_0}{x_1}$, and the map $D_{x_1} \to D_{+,y_1}$ sends $\frac{y_1}{y_0} \mapsto \frac{x_1}{x_0}$.

2.6.1 Rational maps and function fields

We would like to build a category whose morphisms are rational maps. However, in general it is not possible to compose rational maps – the image of the first map might be contained in the locus where the second map is undefined. Instead we must restrict our attention to a particular type of rational map:

Definition 2.6.7. We say that a rational map $f : X \dashrightarrow Y$ of quasiprojective varieties is dominant if it is represented by a morphism $f : U \to Y$ such that the set-theoretic image of U is dense in Y.

Exercise 2.6.8. Show that if $f : X \dashrightarrow Y$ is dominant then *any* representative $f : U \to Y$ will have the property that the set-theoretic image contains an open subset of Y.

Dominant rational maps have the key advantage that they can be composed.

Lemma 2.6.9. *Suppose that $f : X \dashrightarrow Y$ and $g : Y \dashrightarrow Z$ are rational maps of quasiprojective varieties. Then there is a representative (Y, Z, V, g) of g and a representative (X, Y, U, f) of f such that $f^{-1}(V) = U$.*

Thus we can define the composition $g \circ f : X \dashrightarrow Z$ using $(X, Z, U, g \circ f)$.

Proof. Follows immediately from Exercise 2.6.8. □

Given a dominant rational map $f : X \dashrightarrow Y$, one obtains a map $f^\sharp : \mathbb{K}(Y) \to \mathbb{K}(X)$ representing pullback of rational functions. To define this map carefully, we choose open affines $\operatorname{Spec}(R) \subset X$, $\operatorname{Spec}(S) \subset Y$ such that f yields a dominant morphism $f : \operatorname{Spec}(R) \to \operatorname{Spec}(S)$. Since f is dominant, the ring map $f^\sharp : S \to R$ is injective. After localizing we obtain a map $f^\sharp : \operatorname{Frac}(S) \to \operatorname{Frac}(R)$. One can check without too much trouble that this map $\mathbb{K}(Y) \to \mathbb{K}(X)$ is independent of our choices.

Theorem 2.6.10. *There is a contravariant equivalence of categories*

$$\left\{ \begin{array}{c} \textit{quasiprojective } \mathbb{K}\textit{-varieties} \\ \textit{equipped with dominant} \\ \textit{rational maps} \end{array} \right\} \quad \leftrightarrow \quad \left\{ \begin{array}{c} \textit{finitely generated field extensions} \\ \textit{of } \mathbb{K} \textit{ equipped with} \\ \mathbb{K}\textit{-algebra homomorphisms} \end{array} \right\}$$

defined by the functor that sends $X \mapsto \mathbb{K}(X)$ *and* $f : X \dashrightarrow Y$ *to* $f^\sharp : \mathbb{K}(Y) \to \mathbb{K}(X)$.

We could equally well use affine $\mathbb{K}$-varieties in place of quasiprojective $\mathbb{K}$-varieties in our category on the left.

Proof. To check that this functor defines an equivalence of categories, we must ensure that it is full, faithful, and essentially surjective.

Essential surjectivity was done in Exercise 1.10.12: suppose we write $\mathfrak{K} = \mathbb{K}(g_1, \ldots, g_r)$. The subring $R := \mathbb{K}[g_1, \ldots, g_r]$ is a domain so that $\operatorname{mSpec}(R)$ is a variety, and the function field of $\operatorname{mSpec}(R)$ is isomorphic to $\mathfrak{K}$.

We next check faithfulness. Suppose that $f_1, f_2 : X \dashrightarrow Y$ are two rational maps. There are open affines $\operatorname{mSpec}(R) \subset X$ and $\operatorname{mSpec}(S) \subset Y$ such that f_1, f_2 define morphisms $\operatorname{mSpec}(R) \to \operatorname{mSpec}(S)$. If the localized maps $f_1^\sharp, f_2^\sharp : \operatorname{Frac}(S) \to \operatorname{Frac}(R)$ are the same, then by restricting these maps to $S \subset \operatorname{Frac}(S)$ we see that the ring maps $f_1^\sharp, f_2^\sharp$ are also the same.

Finally, we check full. Suppose X, Y are quasiprojective varieties and we are given a map $f^\sharp : \mathbb{K}(Y) \to \mathbb{K}(X)$. Choose open affines $\operatorname{mSpec}(R) \subset X$, $\operatorname{mSpec}(S) \subset Y$, so we obtain a map $f^\sharp : \operatorname{Frac}(S) \to \operatorname{Frac}(R)$. Suppose that $S = \mathbb{K}[g_1, \ldots, g_r]$. We can write $f^\sharp(g_i) = \frac{a_i}{b_i}$ for some $a_i, b_i \in R$. Thus the image of S in $\operatorname{Frac}(R)$ is contained in the localization $R_{b_1 b_2 \ldots b_r}$. In other words, $f^\sharp$ defines a morphism $f : \operatorname{mSpec}(D_{b_1 b_2 \ldots b_r}) \to \operatorname{mSpec}(S)$, and hence a rational map $f : X \dashrightarrow Y$. It's clear that f induces our original map $f^\sharp$ on the level of function fields. $\square$

2.6.2 Birational maps

We conclude by studying the isomorphisms in the categories of Theorem 2.6.10.

Theorem 2.6.11. *Let $f : X \dashrightarrow Y$ be a dominant rational map of quasiaffine $\mathbb{K}$-varieties. The following are equivalent:*

(1) f induces an isomorphism of function fields $f : \mathbb{K}(X) \cong \mathbb{K}(Y)$.

(2) *There is a dominant rational map $g : Y \dashrightarrow X$ such that the compositions $f \circ g$ and $g \circ f$ are both the identity map on their locus of definition.*

(3) *There are open subsets $U \subset X$ and $V \subset Y$ such that $f|_U : U \to V$ is an isomorphism.*

In this situation we say that f is a birational map. More generally, we say that X and Y are birationally equivalent if we want to assert that there is some birational map between them.

Proof. By Theorem 2.6.10 (1) is equivalent to saying that f is an isomorphism in the category of affine $\mathbb{K}$-varieties equipped with dominant rational maps. In other words, (1) is equivalent to the existence of an inverse map $g : Y \dashrightarrow X$ such that the compositions $f \circ g$ and $g \circ f$ are in the equivalence class of the identity map. By Lemma 2.6.1 this is equivalent to (2).

The equivalence of (2) and (3) is clear. □

Warning 2.6.12. The open subsets U in Theorem 2.6.11.(3) need not coincide with the locus of definition of f. Example 2.6.13 shows that the locus of definition can be strictly larger than the set U in Theorem 2.6.11.(3).

Example 2.6.13. Consider the morphism $f : \mathbb{A}^2 \to \mathbb{A}^2$ defined by the ring map $f^\sharp : \mathbb{K}[x, y] \to \mathbb{K}[u, v]$ sending $x \mapsto u, y \mapsto uv$. This function was considered in Example 1.6.3. We showed that f is not an isomorphism but that we do obtain an isomorphism after passing to the open subsets D_x and D_u. In particular we see that f is a birational map.

2.6.3 Exercises

Exercise 2.6.14. Show that the cuspidal curve $\mathrm{mSpec}(\mathbb{K}[x, y]/(y^2 - x^3))$ is birational to $\mathbb{A}^1_{\mathbb{K}}$, but the two spaces are not isomorphic. Write down an open set and an isomorphism realizing the birational equivalence. Identify the induced isomorphism of function fields explicitly.

Exercise 2.6.15. Show that the morphism $f : \mathbb{A}^1_{\mathbb{K}} \to \mathrm{mSpec}(\mathbb{K}[x, y]/(y^2 - x^3 - x^2))$ is birational to $\mathbb{A}^1_{\mathbb{K}}$, but the two spaces are not isomorphic. Write down an open set and an isomorphism realizing the birational equivalence. Identify the induced isomorphism of function fields explicitly.

Exercise 2.6.16. Consider the Cremona transformation $f : \mathbb{P}^2 \dashrightarrow \mathbb{P}^2$ defined by the map of homogeneous coordinate rings $\mathbb{K}[x, y, z] \to \mathbb{K}[u, v, w]$ sending $x \mapsto vw, y \mapsto uw, z \mapsto uv$. Show that f is a birational map. Identify the largest open locus where f is defined. Identify the largest open locus where f is an isomorphism.

Exercise 2.6.17. A quasiprojective variety X is said to be rational if it is birationally equivalent to a projective space $\mathbb{P}^n$ for some n. Show that the hypersurface $V_+(wz - xy) \subset \mathbb{P}^3$ is birational to $\mathbb{P}^2$.

2.7 Graded homomorphisms

Let R, S be finitely generated graded $\mathbb{K}$-algebras. Given a graded $\mathbb{K}$-algebra homomorphism $f^\dagger : S \to R$ we might hope to construct a morphism $f : \mathrm{mProj}(R) \to \mathrm{mProj}(S)$. This is too optimistic – as discussed in the beginning of the chapter, such a morphism can only be defined on the complement of the locus where the defining polynomials vanish.

The following propositions show that a graded $\mathbb{K}$-algebra homomorphism $f^\dagger : S \to R$ defines a morphism from an open subset of $\mathrm{mProj}(R)$ to $\mathrm{mProj}(S)$. In particular, graded homomorphisms give us a rich source of rational maps.

Proposition 2.7.1. *Let $f^\dagger : S \to R$ be a graded homomorphism of finitely generated graded $\mathbb{K}$-algebras. Let $U \subset \mathrm{mProj}(R)$ be the complement of $V_+(f^\dagger(S_{>0}))$. Then $f^\dagger$ defines a morphism $f : U \to \mathrm{mProj}(S)$.*

Proof. By Exercise 2.4.16, we can construct f by defining its restriction to the sets in a base of U consisting of open affines (and ensuring these maps are compatible with restriction). Fix any homogeneous element $g \in S$. Then by localizing $f^\dagger$ and passing to degree 0 parts we obtain a map $(S_g)_0 \to (R_{f^\dagger(g)})_0$. This induces a morphism $f : \mathrm{mSpec}(D_{+,f^\dagger(g)}) \to \mathrm{mSpec}(D_{+,g})$.

To ensure these maps glue, we must check they are compatible under restriction. Note that if $V_+(g) \subset V_+(h)$ then g defines an invertible element in S_h. Thus we have a canonical map $(S_g)_0 \to (S_h)_0$ obtained by localizing $(S_g)_0$ along $h^{\deg g}/g^{\deg h}$. This shows that the map of affine varieties induced by the ring map $(S_h)_0 \to (R_{f^\dagger(h)})_0$ agrees with the restriction of the map of affine varieties $\mathrm{mSpec}(D_{+,f^\dagger(g)}) \to \mathrm{mSpec}(D_{+,g})$ to the open set $D_{+,h}$. This finishes the proof. $\square$

Example 2.7.2. Fix some subset $I \subset \{0, \dots, n\}$ of size $k+1$. Let L be the $(n-k-1)$ dimensional hyperplane defined as the common vanishing locus of the x_i. The inclusion $\mathbb{K}[x_I]_{i\in I} \to \mathbb{K}[x_0, \dots, x_n]$ induces a morphism $\pi : \mathbb{P}^n \backslash L \to \mathbb{P}^k$. On each coordinate chart D_{+,x_i} for $i \in I$ the restriction of π is simply a coordinate projection $\mathbb{A}^n \to \mathbb{A}^k$. The closure of each fiber of π will be a hyperplane of dimension $n-k$ which is defined by k independent homogeneous linear equations in the variables x_I.

Let's discuss the geometry of this important construction. Starting from the fixed $(n-k-1)$-plane L, we have assigned to any point $x \in \mathbb{P}^n$ the $n-k$ plane spanned by x and L. In particular, we are thinking of the points $\mathbb{P}^k$ as the parameter space for the $(n-k)$-planes containing L. The corresponding morphism $\mathbb{P}^n \backslash L \to \mathbb{P}^k$ sends a point of $\mathbb{P}^n$ to the unique $(n-k)$-plane which contains it. Any map of this type is called projection away from the $(n-k-1)$-plane L.

Warning 2.7.3. Suppose that $f^\dagger : S \to R$ is graded homomorphism of finitely generated graded $\mathbb{K}$-algebras. While Theorem 2.7.1 guarantees that we obtain a morphism f that is defined away from the common vanishing locus of the defining functions, it is possible that f can actually be defined on a larger open subset in $\mathrm{mProj}(R)$.

Consider for example the map $f^\dagger : \mathbb{C}[s,t] \to \mathbb{C}[x,y,z]/(x^2+y^2-z^2)$ sending $s \mapsto x, t \mapsto z-y$. Geometrically, this is the composition of the inclusion $X := V_+(x^2+y^2-z^2) \subset \mathbb{P}^2$ with projection away from the point $x = y - z = 0$. Theorem 2.7.1 yields a map $f : X \backslash \{(x, z-y)\} \to \mathbb{P}^1$. We claim that we can extend f to be a morphism on all of X by sending $(x, z-y) \mapsto (t)$.

We prove this carefully using affine charts. X is covered by the two affine charts $X \cap D_{+,z-y}$ and $X \cap D_{+,z+y}$ in $\mathbb{P}^2$. $\mathbb{P}^1$ admits the two affine charts $D_{+,s}$ and $D_{+,t}$. By localizing $f^\dagger$ we obtain the morphism $X \cap D_{+,z-y} \to D_{+,t}$ defined by $\frac{s}{t} \mapsto \frac{x}{z-y}$.

On the other charts we define $X \cap D_{+,z+y} \to D_{+,s}$ via $\frac{t}{s} \mapsto \frac{x}{z+y}$. The key point is that this morphism agrees with our original map on the locus $X \cap D_{+,(z+y)} \cap D_{+,(z-y)}$ where both maps are defined. This follows from the computation

$$\frac{z-y}{x} = \frac{z^2-y^2}{x(z+y)} = \frac{x}{z+y}.$$

By gluing the maps on these two charts we extend f to all of X.

Warning 2.7.4. Just as many different graded rings can define isomorphic mProjs, many different graded homomorphisms can define the same map of schemes. Even worse, if we fix finitely generated graded $\mathbb{K}$-algebras R and S there may be a morphism $f : U \to \mathrm{mProj}(S)$ from an open subset $U \subset \mathrm{mProj}(R)$ which is not induced by any graded homomorphism involving S and R.

(One example is to take an elliptic curve $E \subset \mathbb{P}^2_{\mathbb{C}}$ defined by a cubic equation and define $f : E \to E$ to be translation by a non-torsion point. Unfortunately we don't yet have the tools to analyze this example, but as you might guess it arises from some kind of incompatibility of the line bundles $\mathcal{L}$ arising from the mProj construction with the geometry of the map f.)

2.7.1 Veronese subrings

As mentioned in the introduction to the chapter it is not true that every finitely generated graded $\mathbb{K}$-algebra R with $R_0 \cong \mathbb{K}$ is a quotient of a polynomial ring. Thus it is not immediately clear whether every projective scheme is homeomorphic to a closed subset of projective space. In this subsection we will address this issue using the following special example of a graded homomorphism.

Definition 2.7.5. Given a $\mathbb{Z}$-graded ring R, we define the degree d Veronese subring $R^{(d)}$ to be the $\mathbb{Z}$-graded ring such that $R^{(d)}_k = R_{dk}$ (equipped with the natural addition and multiplication structures).

Exercise 2.7.6. Prove that if R is a finitely generated graded $\mathbb{K}$-algebra then $R^{(d)}$ is as well for any positive integer d.

Proposition 2.7.7. *Let R be a finitely generated graded $\mathbb{K}$-algebra and fix a positive integer d. Then* $\mathrm{mProj}(R^{(d)})$ *and* $\mathrm{mProj}(R)$ *are isomorphic.*

We will study the geometry of this map in Section 3.2.

Proof. Consider the natural inclusion $i : R^{(d)} \to R$. The i-pullback of a homogeneous ideal is a homogeneous ideal. We claim that if $\mathfrak{p}$ is a homogeneous prime ideal in R and I is any homogeneous ideal in R then $\mathfrak{p} \supset I$ if and only if $\mathfrak{p} \cap R^{(d)} \supset I \cap R^{(d)}$. The forward implication is obvious. To see the reverse implication, suppose that $\mathfrak{p} \cap R^{(d)} \supset I \cap R^{(d)}$ and let f be any element in I. Then $f^d \in I \cap R^{(d)} \subset \mathfrak{p} \cap R^{(d)}$ so that $f \in \mathfrak{p}$.

Conversely, suppose that $\mathfrak{q}$ is a homogeneous prime ideal in $R^{(d)}$. We claim that $\sqrt{\mathfrak{q}}$ is a prime ideal in R. It suffices to check that if a product of homogeneous elements $f, g \in R$ lies in $\sqrt{\mathfrak{q}}$ then either f or g lie in $\sqrt{\mathfrak{q}}$. Choose some positive integer j such that $f^{dj}g^{dj} \in \mathfrak{q}$. Due to the primality of $\mathfrak{q}$ in $R^{(d)}$ we see that either f^{dj} or g^{dj} lies in $\mathfrak{q}$. This proves the desired statement.

The previous two paragraphs show that pullback by i induces a bijection between homogeneous prime ideals in R and in $R^{(d)}$ and that this bijection preserves the inclusion relation. In particular we obtain a bijection between the almost maximal homogeneous ideals of R and the almost maximal homogeneous ideals of R. Furthermore, since the topology of R and $R^{(d)}$ is controlled by homogeneous prime ideals (see Exercise 2.3.16) we see that the induced set-theoretic map $\mathrm{mProj}(R) \to \mathrm{mProj}(R^{(d)})$ is a homeomorphism.

Finally, we must show an equality of sheaves of functions. Fix $f \in R$. We claim that $(R_f)_0$ is isomorphic to $(R^{(d)}_{f^d})_0$. The inclusion morphism

$$\begin{aligned} i : (R^{(d)}_{f^d})_0 &\to (R_f)_0 \\ a/f^{dr} &\mapsto a/f^{dr} \end{aligned}$$

is injective. It is also surjective: any element $b/f^a \in (R_f)_0$ with $a = qd + t$ is equivalent to the element $bf^{d-t}/f^{d(q+1)}$ which is in the image of i.

The computation shows that for every distinguished open affine $D_{+,f} \subset \mathrm{mProj}(R)$ and for the corresponding distinguished open affine $D_{+,f^d} \subset \mathrm{mProj}(R^{(d)})$ we have that $\mathcal{O}_{\mathrm{mProj}(R)}(D_{+,f})$ and $\mathcal{O}_{\mathrm{mProj}(R^{(d)})}(D_{+,f^d})$ are isomorphic. Using the gluing property of sheaves this implies that that $\mathcal{O}_{\mathrm{mProj}(R)}$ and $\mathcal{O}_{\mathrm{mProj}(R^{(d)})}$ are the same. $\square$

Exercise 2.7.8. Verify carefully the last sentence in the proof of Proposition 2.7.7.

Remark 2.7.9. Let R, R' be finitely generated graded $\mathbb{K}$-algebras that are generated in degree 1. Suppose that $\mathrm{mProj}(R)$ and $\mathrm{mProj}(R')$ are isomorphic and that under this isomorphism the line bundles $\mathcal{L}, \mathcal{L}'$ from Warning 2.3.13 are identified. Then there is some positive integer d such that $R^{(d)}$ and $R'^{(d)}$ are isomorphic. In other words, even though one cannot recover R from $\mathrm{mProj}(R)$ equipped with the line bundle $\mathcal{L}$, one can recover its "Veronese equivalence class".

The following exercise is a key step in showing that projective schemes are indeed closed subschemes of projective space.

Exercise 2.7.10. Let R be a finitely generated graded $\mathbb{K}$-algebra. Prove that there exists a sufficiently large d so that $R^{(d)}$ is generated as a $\mathbb{K}$-algebra by its degree 1 piece. Deduce that this $R^{(d)}$ is a quotient of a polynomial ring by a homogeneous ideal. (Hint: suppose that the degrees of the generators $g_1, \ldots, g_s$ of R are $d_1, \ldots, d_s$. Show that $d = sd_1 \ldots d_s$ works.)

Deduce that if R is a finitely generated graded $\mathbb{K}$-algebra such that $R_0 \cong \mathbb{K}$ then $\mathrm{mProj}(R)$ is isomorphic to the Proj of a quotient of a polynomial ring.

2.7.2 Exercises

Exercise 2.7.11. Prove that any element of $\mathrm{PGL}_{n+1}(\mathbb{K})$ defines an automorphism of $\mathbb{P}^n$ via multiplication. (In Theorem 6.4.8 we will show that these are the only automorphisms of $\mathbb{P}^n$, at least if $\mathbb{K}$ is algebraically closed.)

Prove that the action of $\mathrm{PGL}_{n+1}(\mathbb{K})$ on the $\mathbb{K}$-points of $\mathbb{P}^n$ is 2-transitive and faithful.

Exercise 2.7.12. Let $R = \mathbb{K}[x, y, z]/(xz, yz, z^2)$. Show that the two maps $f^\dagger, g^\dagger : R \to R$ defined by

$$\begin{array}{lll} f^\dagger(x) = x & f^\dagger(y) = y & f^\dagger(z) = z \\ g^\dagger(x) = x & g^\dagger(y) = y & g^\dagger(z) = 0 \end{array}$$

define the same morphism $\mathrm{mProj}(R) \to \mathrm{mProj}(R)$. (This is simply reflecting the map that $f^\dagger$ and $g^\dagger$ are identified after passing to the 2-Veronese subring of R.)

Exercise 2.7.13. Consider the graded ring $\mathbb{K}[s, t]$ where s has degree $a \geq 1$, t has degree $b \geq 1$ and $\gcd(a, b) = 1$. Prove that $\mathrm{mProj}(\mathbb{K}[s, t])$ is isomorphic to $\mathbb{P}^1$. (Hint: show that the Veronese subring $\mathbb{K}[s, t]^{(ab)}$ is isomorphic to $\mathbb{K}[x, y]$ where we assign both x and y the degree ab.)

Exercise 2.7.14. Just as different graded rings can define the same mProj, different ideals can define the same closed subset of a mProj. This exercise explores this ambiguity.

Let R be a finitely generated graded $\mathbb{K}$-algebra which is generated by R_1. Let $I \subset R$ be a homogeneous ideal. The saturation I_{sat} is defined to be

$$I_{sat} = \{g \in R \,|\, \exists j \in \mathbb{Z}_{\geq 0} \text{ s.t. } gR_j \subset I\}.$$

(1) Show that I_{sat} is a homogeneous ideal and that the degree m pieces satisfy $I_{sat,m} = I_m$ for m sufficiently large. Show that $V(I) = V(I_{sat})$ as closed subsets of $\mathrm{mProj}(R)$.

(2) Show that $\mathrm{mProj}(R/I)$ is isomorphic to $\mathrm{mProj}(R/I_{sat})$.

(3) Suppose that J is any homogeneous ideal such that $V(J) = V(I)$ and the induced homeomorphism between $\mathrm{mProj}(R/I)$, $\mathrm{mProj}(R/J)$ yields an isomorphism of sheaves of functions. Prove that $J_{sat} = I_{sat}$.

Remark 2.7.15. When R is not generated by R_1 the analogous statements are no longer true if one uses saturated ideals. Instead, let m denote the lcm of the generators of R. If one defines the "weak saturation"

$$I_{wsat} := \{g \in R \,|\, Rg \cap R_{dm} \subset I \ \forall d \gg 0\}$$

then one can develop an analogous theory.

2.8 Closed embeddings

In this section we systematically develop the theory of closed embeddings. Let's recall the definition:

> A morphism $f : X \to Y$ of quasiprojective schemes is a closed embedding if f is a homeomorphism onto a closed subset of Y and for every open affine $V \subset Y$ the map $f^\sharp(V) : \mathcal{O}_Y(V) \to \mathcal{O}_X(f^{-1}(V))$ is surjective.

Exercise 2.8.1. Suppose $X = \mathrm{mSpec}(R)$ is an affine scheme. Show that every closed subscheme of X is an affine scheme and is defined by the vanishing locus of an ideal in R.

Since the definition involves *every* open affine in Y, it can be challenging to check directly. The following result shows we only need to check surjectivity for a cover by open affines.

Proposition 2.8.2. *Let $f : X \to Y$ be a morphism of quasiprojective schemes. Suppose that f is a homeomorphism onto a closed subset of Y and there exists an open cover of Y by open affines V such that $f^\sharp(V) : \mathcal{O}_Y(V) \to \mathcal{O}_X(f^{-1}(V))$ is surjective. Then f is a closed embedding.*

Proof. Suppose we fix an open affine U of Y. By Lemma 2.5.2, U is covered by distinguished open affines D_g each of which is also a distinguished open affines inside an open affine V in our cover. Since surjectivity of a ring homomorphism is preserved by localization, for each D_g the map $f^\sharp(D_g) : \mathcal{O}_Y(D_g) \to \mathcal{O}_X(f^{-1}(D_g))$ is surjective. Since surjectivity of a ring homomorphism can be checked locally, we deduce that $f^\sharp(U) : \mathcal{O}_Y(U) \to \mathcal{O}_X(f^{-1}(U))$ is surjective. □

Exercise 2.8.3. Let R be a finitely generated graded $\mathbb{K}$-algebra and let I be a homogeneous ideal. Show that the quotient map $R \to R/I$ induces (via Proposition 2.7.1) a closed embedding $\mathrm{mProj}(R/I) \to \mathrm{mProj}(R)$.

2.8.1 Ideal sheaves

When working with closed embeddings it is often more convenient to focus on the kernel of $f^\sharp$. The following definition describes the relevant construction.

Definition 2.8.4. Let X be a quasiprojective scheme. A quasicoherent ideal sheaf $\{I_U\}_{U\subset X}$ on X assigns to each open affine $U \subset X$ an ideal $I_U \subset \mathcal{O}_X(U)$ such that the following property holds:

♣ For any open affine U and any $f \in \mathcal{O}_X(U)$ we have $I_{D_f} = (I_U)_f$.

Remark 2.8.5. Using the gluing property for local sections of $\mathcal{O}_X$, a quasicoherent ideal sheaf (as defined above) allows us to construct a "subsheaf" $\mathcal{I} \subset \mathcal{O}_X$. We will not need this perspective.

Recall that for any quasiprojective scheme X the structure sheaf $\mathcal{O}_X$ is compatible with the structure induced by localization maps. Thus, if we have any morphism of quasiprojective schemes $f : X \to Y$, the kernel of $f^\sharp$ will also be compatible with localizations and will thus be a quasicoherent ideal sheaf. This motivates the condition (♣).

Our next goal is to prove a converse statement: any quasicoherent ideal sheaf must come from a closed embedding. We start by considering the mProj construction.

Theorem 2.8.6. *Let R be a finitely generated graded $\mathbb{K}$-algebra and let $X = \mathrm{mProj}(R)$. Suppose we fix a quasicoherent ideal sheaf $\{I_U\}_{U\subset X}$ on X. Then there is a closed subscheme $Z \subset X$ such that for every open affine $U \subset X$ the kernel of $\mathcal{O}_X(U) \to \mathcal{O}_Z(U \cap Z)$ is equal to I_U.*

The strategy is to recombine the various I_U to obtain a homogeneous ideal $J \subset R$. We will then verify that $Z = \mathrm{mProj}(R/J)$ has the desired properties.

Proof. Consider a distinguished open affine $U = D_{+,g}$ inside of $X = \mathrm{mProj}(R)$. The quasicoherent ideal sheaf gives an ideal $I_U \subset (R_g)_0$. We define the ideal $J_U \subset R$ by taking the preimage of the ideal generated by I_U under the localization map $R \to R_g$. We then define $J = \cap_U J_U$ as we vary over all distinguished open affines U in X.

Set $Z = \mathrm{mProj}(R/J)$. According to Exercise 2.8.3, the quotient $R \to R/J$ yields a closed embedding $Z \to X$. We claim that this closed embedding has the desired property: the kernel of $f^\sharp$ on open affines is determined by our given quasicoherent ideal sheaf.

We first check this for distinguished open affines $U = D_{+,g}$. Since $Z \cap U$ is defined by the ideal $(J_g)_0$, we must verify that $(J_g)_0 = I_U$. Since $J \subset J_U$ and by construction $(J_{U,g})_0 = I_U$, we have the containment $\subset$. Conversely, given any distinguished open affine $V = D_{+,g}$ of X consider the diagram of localizations

$$\begin{tikzcd} R \arrow[r] \arrow[d] & R_{g'} \arrow[d] \\ R_g \arrow[r] & R_{gg'}. \end{tikzcd}$$

By the condition (♣) the ideal generated by I_V in $R_{g'}$ extends to the ideal generated by $I_{U\cap V}$ in $R_{gg'}$. In other words, for any element h in the ideal generated by I_V in $R_{g'}$ there is some power k of g such that hg^k is contained in the contraction of the ideal generated by $I_{U\cap V}$ to $R_{g'}$. Taking preimages, we see that for any element h' of J_V there is some power k of g such that $h'g^k$ is contained in $J_{U\cap V}$. Thus J_V and $J_{U\cap V}$ have the same extension to R_g. Since the extension of $J_{U\cap V}$ to R_g contains the ideal generated by I_U, J_V also has this property. Since V was an arbitrary distinguished open affine, we conclude that J_g contains the ideal generated by I_U.

It only remains to verify the desired property for open affines $U \subset X$ which are not distinguished. Conceptually, this follows from the fact that both I_U and the structure sheaf for $Z \cap U$ are determined by gluing sections for the distinguished open affines contained in U. We leave the details to the motivated reader. □

One consequence of our argument is the following important corollary. It shows the converse to Exercise 2.8.3.

Corollary 2.8.7. *Let R be a finitely generated graded $\mathbb{K}$-algebra and let $X = \mathrm{mProj}(R)$. For any closed subscheme $Z \subset X$, there is a homogeneous ideal $I \subset R$ such that $Z \cong \mathrm{mProj}(R/I)$ and the inclusion map is induced by the quotient $R \to R/I$.*

Warning 2.8.8. While every closed *subscheme* of $\mathrm{mProj}(R)$ is defined by a surjection $R \mapsto R/I$, this is not true for every closed *embedding* of $\mathrm{mProj}(R)$. As discussed in Warning 2.7.4, an elliptic curve in $\mathbb{P}^2$ admits automorphisms which cannot be induced by any graded homomorphism of its homogeneous coordinate ring.

Finally, we extend Theorem 2.8.6 to arbitrary quasiprojective schemes.

Theorem 2.8.9. *Let X be a quasiprojective scheme. Suppose we fix a quasicoherent ideal sheaf $\{I_U\}_{U \subset X}$ on X. Then there is a closed subscheme $Z \subset X$ such that for every open affine $U \subset X$ the kernel of $\mathcal{O}_X(U) \to \mathcal{O}_Z(U \cap Z)$ is equal to I_U.*

Proof. Since X is a quasiprojective scheme, by definition it admits an open embedding $X \subset Y \cong \mathrm{mProj}(R)$ for some finitely generated graded $\mathbb{K}$-algebra R. Our first goal is to extend the quasicoherent ideal sheaf to Y.

Suppose V is an open affine in Y. If $V \cap X = \emptyset$, we define $J_V = \mathcal{O}_X(V)$. Next suppose that $V \cap X \neq \emptyset$. Choose an open affine W that is a distinguished open affine in V and is contained in X. The quasicoherent ideal sheaf on X yields an ideal $I_W \subset \mathcal{O}_X(W) = \mathcal{O}_Y(W)$. We let $J_{V,W}$ denote the preimage of I_W under the localization map $\mathcal{O}_Y(V) \to \mathcal{O}_Y(W)$. We then define $J_V = \bigcap_W I_{V,W}$ as W varies over all distinguished open affines of V contained in X.

We first verify that if $V \subset X$ is an open affine then $I_V = J_V$ (so this construction does not change the structure of the quasicoherent ideal sheaf for open affines in X). Suppose that $W \subset V$ is a distinguished open affine. Then (♣) guarantees that I_W is obtained by localizing I_V. Thus $J_{V,W} \supset I_V$ and so I_V is the unique minimal element amongst all the ideals whose intersection defines J_V.

We next verify that the various $\{J_V\}_{V \subset Y}$ yield a quasicoherent ideal sheaf on Y. Suppose given an open affine $V \subset Y$ and a distinguished open affine W in V. For any smaller open affine W' that is distinguished in W (and hence also in V) and is contained in X, both $J_{V,W'}$ and $J_{W,W'}$ are the pullbacks of the same ideal under the localization maps. Thus $J_{W,W'}$ is the localization of $J_{V,W'}$. By varying W' and taking intersections we obtain the compatibility of J_V and J_V with localization.

By Theorem 2.8.6 the quasicoherent ideal sheaf on Y defines a closed subscheme Z_Y of Y. We define the quasiprojective scheme $Z = Z_Y \cap X$. Using $J_V = I_V$ for open affines $V \subset X$ it is straightforward to verify that Z has the desired property. □

Exercise 2.8.10. Suppose that $f : X \to Y$ is a closed embedding of quasiprojective schemes. Let $\{I_U\}_{U \subset Y}$ be the quasicoherent ideal sheaf defined via the kernel of $f^\sharp$. By Theorem 2.8.9 this quasicoherent ideal sheaf defines a closed subscheme $Z \subset Y$. Prove that X is isomorphic to Z.

2.8.2 Scheme theoretic image

The argument used in the proof of Theorem 2.8.9 is a special case of a more general construction.

Construction 2.8.11. Let $f : X \to Y$ be a morphism of quasiprojective schemes. We define the scheme-theoretic image of f as follows. For every open affine $U \subset Y$, we have a map $f^\sharp(U) : \mathcal{O}_Y(U) \to \mathcal{O}_X(f^{-1}U)$. We define I_U to be the kernel of this map.

Suppose that $g \in \mathcal{O}_Y(U)$ and let $V = D_g$ be the corresponding distinguished open affine. Using Exercise 2.4.19, we see that $\mathcal{O}_X(f^{-1}V) = \mathcal{O}_X(f^{-1}U)_{f^\sharp(U)(g)}$. Thus the $\{I_U\}_{U \subset Y}$ form a quasicoherent ideal sheaf on Y. Then Theorem 2.8.9 yields a closed subscheme $Z \subset Y$. We call Z the scheme-theoretic image of f.

It is straightforward to check that f factors through the inclusion $Z \subset Y$ and that the set-theoretic image of f is dense in Z. In some sense Z is the "smallest" closed subscheme of Y which has this factoring property.

2.8.3 Projective schemes

We are now equipped to verify "half" of Theorem 2.4.13.

Theorem 2.8.12. *Let X be a quasiprojective scheme. Then X is projective if and only if it admits a closed embedding into $\mathbb{P}^n$ for some n.*

Proof. First suppose that X is projective. By definition X is isomorphic to mProj(R) for some finitely generated graded ring R with $R_0 \cong \mathbb{K}$. By Exercise 2.7.10 X is isomorphic to the mProj of a quotient of a polynomial ring by a homogeneous ideal. By Exercise 2.8.3 we see that X admits a closed embedding into projective space.

Conversely, suppose that X admits a closed embedding into projective space. Then Corollary 2.8.7 shows that X is the mProj of the quotient of a polynomial ring by a homogeneous ideal. Note that the 0th graded piece of a such a ring will be isomorphic to $\mathbb{K}$. $\square$

2.8.4 Exercises

Exercise 2.8.13. Prove that a composition of closed embeddings is a closed embedding.

Exercise 2.8.14. Suppose that $f : X \to Y$ is both an open embedding and a closed embedding. Prove that f is an isomorphism.

Exercise 2.8.15. Let X be a quasiprojective scheme. Show that the nilpotent functions on open affines in X form a quasicoherent ideal sheaf. Use this quasicoherent ideal sheaf to construct a closed embedding $i : X_{red} \to X$ such that i is a homeomorphism and X_{red} is reduced.

Prove that if $f : Y \to X$ is any morphism from a reduced quasiprojective scheme Y then f factors through i. In particular the scheme-theoretic image (in the sense of Construction 2.8.11) of a morphism from a reduced variety will always be reduced.

Exercise 2.8.16. Let X be a projective $\mathbb{K}$-scheme.

(1) Prove that for any finite set of points in X there is an open affine $U \subset X$ containing the entire set.

(2) Suppose that $Z \subset X$ is a closed subscheme. Prove that there is an open affine U in X such that $Z \cap U$ is dense in Z.

(Hint: embed X into a projective space $\mathbb{P}^n$ and find hypersurfaces in $\mathbb{P}^n$ which do not contain the components of the given closed subsets in X.)

2.9 Products

The most important construction in algebraic geometry is the (relative) product. As with affine schemes, by "product" we will always mean the categorial product defined via a universal property.

In this section we will construct the product of two quasiprojective schemes. Our strategy is to reduce the question to the affine case using open covers.

2.9.1 Products and open sets

We first need to know how taking products interacts with taking open sets. The following sequence of lemmas explains this relationship.

Lemma 2.9.1. *Let X, Y, Z be quasiprojective schemes with morphisms $f : X \to Z$, $g : Y \to Z$. Let $U \subset X$ be an open subset. Suppose that the product $X \times_Z Y$ exists in the category of quasiprojective schemes. Then the product $U \times_Z Y$ exists and is equal to the preimage of U under the projection map $\pi_1 : X \times_Z Y \to X$.*

Exercise 2.9.2. Prove Lemma 2.9.1. (The key point is that if X is a quasiprojective scheme, $i : U \to X$ is an open embedding, and $f, g : W \to U$ are two morphisms such that $i \circ f = i \circ g$, then $f = g$. Everything else is a diagram chase using the universal property of $X \times_Z Y$.)

Lemma 2.9.3. *Let X, Y, Z be quasiprojective schemes with morphisms $f : X \to Z$, $g : Y \to Z$. Let $\{U_i\}_{i \in I}$ be an open cover of X. Assume that each relative product $U_i \times_Z Y$ exists in the category of quasiprojective schemes.*

Suppose that P is a quasiprojective scheme admitting morphisms $\pi_1 : P \to X$, $\pi_2 : P \to Y$ that fit into a commutative diagram

$$\begin{tikzcd} P \arrow[r, "\pi_1"] \arrow[d, "\pi_2"] & X \arrow[d, "f"] \\ Y \arrow[r, "g"] & Z. \end{tikzcd}$$

Suppose also that for every i the open set $\pi_1^{-1}(U_i)$ is isomorphic to $U_i \times_Z Y$ when equipped with the restrictions of the maps π_1, π_2. Then the product $X \times_Z Y$ exists and is isomorphic to P (equipped with the maps π_1, π_2).

Proof. We must show that P satisfies the universal property of a product. Suppose that M is any quasiprojective scheme equipped with maps $h_1 : M \to X$ and $h_2 : M \to Y$ which make a commuting diagram with the maps to Z. For each open subset U_i, let $N_i = h_1^{-1}(U_i)$. By the universal property of products, we have a system of maps $\phi_i : N_i \to U_i \times_Z Y$.

To verify that the ϕ_i glue we must check that they agree on overlaps. Let $U_{ij} = U_i \cap U_j$. By applying Lemma 2.9.1, we see that $U_{ij} \times_Z Y$ is the same as $\pi_1^{-1}(U_{ij})$. Let $N_{ij} = N_i \cap N_j$. By applying the uniqueness of factorings for products, we see that

$$\phi_i|_{N_{ij}} = \phi_j|_{N_{ij}}.$$

By applying Exercise 2.4.16 we deduce that the ϕ_i glue together to give a morphism $\phi : M \to P$ making all the diagrams commute.

Finally, we must verify that ϕ is the unique such map. If there were two such maps ϕ, ψ then by applying uniqueness of products we see $\phi|_{N_i} = \psi|_{N_i}$. We deduce that $\phi = \psi$ everywhere. □

To construct the product $X \times_Z Y$, Lemma 2.9.1 and Lemma 2.9.3 allow us to reduce to open affine covers of X and Y. The final step is to allow us to also replace Z by an open affine cover. This is accomplished in the next lemma.

Lemma 2.9.4. *Let X, Y, Z be quasiprojective schemes with morphisms $f : X \to Z$, $g : Y \to Z$. Suppose that $W \subset Z$ is an open subset. Suppose $U \subset f^{-1}(W)$ and $V \subset g^{-1}(W)$ are open subsets. If $U \times_W V$ exists in the category of quasiprojective schemes, then $U \times_Z V$ does as well and is isomorphic to it.*

Exercise 2.9.5. Prove Lemma 2.9.4 by chasing arrows around.

2.9.2 Constructing the product

The basic idea is to construct products by passing to open covers consisting of open affine sets. We already know how to construct the relative product for affine schemes inside of **AffSch**/$\mathbb{K}$. In fact, the construction of the product inside of **QProSch**/$\mathbb{K}$ is exactly the same.

Exercise 2.9.6. Suppose given affine schemes $\mathrm{mSpec}(R), \mathrm{mSpec}(S), \mathrm{mSpec}(T)$ equipped with $\mathbb{K}$-algebra homomorphisms $f^\sharp : T \to R$ and $g^\sharp : T \to S$. Using Theorem 2.4.8, prove that $\mathrm{mSpec}(R \otimes_T S)$ satisfies the universal property needed to be the relative product $\mathrm{mSpec}(R) \times_{\mathrm{mSpec}(T)} \mathrm{mSpec}(S)$ in the category of quasiprojective schemes.

We next turn to the product of mProjs. We start by constructing the absolute product.

Construction 2.9.7. Suppose that $X = \mathrm{mProj}(R)$ and $Y = \mathrm{mProj}(S)$. Let Q denote the finitely generated graded $\mathbb{K}$-algebra

$$Q = \bigoplus_{d \geq 0} (R_d \otimes_{\mathbb{K}} S_d).$$

We set $P = \mathrm{mProj}(Q)$. We claim that P is the product of X and Y. We first need to construct morphisms $\pi_1 : P \to X$ and $\pi_2 : P \to Y$. Given any homogeneous $y \in Y$ of degree d, we define a graded homomorphism $\pi_{1,y}^{\dagger} : R^{(d)} \to Q$ by sending $r \mapsto r \otimes y^{\deg(r)/d}$. This defines a morphism from an open subset of P to $\mathrm{mProj}(R^{(d)}) \cong X$. As we vary y, these maps agree on the overlaps and thus define a morphism $\pi_1 : P \to X$. We construct π_2 in an analogous way.

To verify the universal property we will apply Lemma 2.9.3. Fix homogeneous elements $f \in R$ and $g \in S$. After replacing both f and g by suitable powers, we may assume that both have the same degree d in R and S respectively. Then the degree 0 part of the localization of Q at $f \times g$ is

$$\begin{aligned}(Q_{f\times g})_0 &\cong \bigoplus_{m\geq 0}(R_{md} \otimes_{\mathbb{K}} S_{md}) \cdot (f \otimes g)^{-m} \\ &\cong (R_f)_0 \otimes_{\mathbb{K}} (S_g)_0.\end{aligned}$$

Varying f and g, we see that P admits a cover by open affines which are products of the corresponding open affines in X and Y. Then by applying Lemma 2.9.3 first in one factor, then the other, we deduce that $P \cong X \times Y$.

In particular:

Corollary 2.9.8. *If X and Y are projective schemes then $X \times Y$ is a projective scheme.*

Proof. Write $X \cong \mathrm{mProj}(R)$, $Y \cong \mathrm{mProj}(S)$ with $R_0 \cong S_0 \cong \mathbb{K}$. Then $X \times Y$ is also projective since $R_0 \otimes S_0 \cong \mathbb{K}$. □

We will study the geometry of the product of projective schemes in more detail in Section 3.3.

Since a morphism $\mathrm{mProj}(R) \to \mathrm{mProj}(S)$ need not be induced by a graded homomorphism $S \to R$ (see Warning 2.7.4) we need to take a different approach for relative products.

Exercise 2.9.9. Let X and Y be projective schemes equipped with morphisms $f : X \to Z$ and $g : Y \to Z$ where Z is a quasiprojective scheme. Suppose that $W \subset Z$, $U \subset f^{-1}(W)$, $V \subset f^{-1}(W)$ are open affine subsets in their respective schemes. From Exercise 1.7.16 we know that $U \times_W V$ is a closed subscheme of $U \times V$. Show that as we vary U, V, W the various ideals defining the closed subsets $U \times_W V$ define a quasicoherent ideal sheaf on $X \times Y$.

By Theorem 2.8.9 this quasicoherent ideal sheaf corresponds to a closed subscheme Q of $X \times Y$. Use Lemma 2.9.3 and Lemma 2.9.4 to verify that Q is isomorphic to the relative product $X \times_Z Y$. (In particular this shows that $X \times_Z Y$ is a projective scheme.)

Finally, by applying Lemma 2.9.1 and Lemma 2.9.4 we deduce the existence of relative products for all quasiprojective schemes:

Theorem 2.9.10. *Let $f : X \to Z$ and $g : Y \to Z$ be morphisms of quasiprojective schemes. Then the product $X \times_Z Y$ exists in the category of quasiprojective schemes.*

This construction is also known as the pullback of f by $Y \to Z$, or the base change of f by $Y \to Z$. This language is often used when we are thinking of $X \to Z$ as a "family of schemes", so that $X \times_Z Y$ represents "pulling the fibers back over $Y \to Z$".

2.9.3 Exercises

Exercise 2.9.11. Suppose that $X = \mathrm{mSpec}(R)$ is an affine scheme. Show that $X \times \mathbb{P}^n$ is isomorphic to $\mathrm{mProj}(R[x_0, \ldots, x_n])$ with the grading given by the degree in the x_i variables.

More generally, suppose that $X = \mathrm{mSpec}(R)$ is an affine scheme and $Y = \mathrm{mProj}(S)$ is a projective scheme. Show that $X \times Y$ is isomorphic to $\mathrm{mProj}(R \otimes_{\mathbb{K}} S)$ where the grading is only in the S factor.

Exercise 2.9.12. Let X and Z be quasiprojective schemes.

(1) Suppose that $f : X \to Z$ is a closed embedding. Show that for any morphism $g : Y \to Z$ the projection map $X \times_Z Y \to Y$ is a closed embedding.

(2) Suppose that $f : X \to Z$ is an open embedding. Show that for any morphism $g : Y \to Z$ the projection map $X \times_Z Y \to Y$ is an open embedding.

We refer to these important properties by saying closed embeddings and open embeddings are "stable under base change."

Exercise 2.9.13. By arguing locally on affine charts, show that the map $(id, id) : X \to X \times X$ is a closed embedding. The image of this map is called the diagonal Δ.

Exercise 2.9.14. Let $f : X \to Y$ be a morphism of quasiprojective schemes. Show that $(f, id) : X \to X \times Y$ is a closed embedding. The image of this map is called the graph Γ of f.

2.10 Applications of products

In the previous section we constructed the relative product for quasiprojective schemes. In this section we continue our discussion of products by focusing on geometric applications.

2.10.1 Constructions

Certain constructions in algebraic geometry are best formulated using the relative product.

Definition 2.10.1. Let $f : X \to Y$ be a morphism of quasiprojective schemes. Suppose that $i : Z \to Y$ is any morphism. The fiber of f over i is $X \times_Y Z$.

Of course, we normally think of "fibers" when $i : Z \to Y$ is the inclusion of a point. However, it is useful to make this more general definition. Note that this definition coincides with Definition 1.6.1 when X, Y are affine schemes and $i : Z \to Y$ is a closed embedding.

Exercise 2.10.2. Show that the definition of the fiber is compatible with set-theoretic fibers. That is, given a morphism $f : X \to Y$ of quasiprojective schemes and a point $y \in Y$ the fiber of f over y is homeomorphic to the closed subset $f^{-1}(y)$ as a topological space.

The fiber product can also be used to define intersections. The advantage of using the fiber product is that it automatically gives us the "correct" scheme structure on the underlying topological space.

Definition 2.10.3. Let X be a quasiprojective scheme. Suppose that U, V are closed or open subschemes of X. We define the intersection of U and V to be $U \times_X V$.

Exercise 2.10.4. Prove that if U, V are closed or open subschemes of X then the underlying set of $U \times_X V$ is homeomorphic to $U \cap V$.

2.10.2 Base change and products

Suppose we have an extension of fields $\mathbb{L}/\mathbb{K}$. In Exercise 1.4.12 we discussed the base change operation for affine schemes: to any affine $\mathbb{K}$-scheme $X = \mathrm{mSpec}(R)$ we can associate the affine $\mathbb{L}$-scheme $X_{\mathbb{L}} := \mathrm{mSpec}(R \otimes_{\mathbb{K}} \mathbb{L})$. This operation naturally extends to the mProj construction, and hence also to arbitrary quasiprojective schemes.

When $\mathbb{L}/\mathbb{K}$ is a finite extension, the product gives us an alternative way of describing base change. Let X be a quasiprojective $\mathbb{K}$-scheme. Note that both X and $\mathrm{mSpec}(\mathbb{L})$ admit a canonical morphism to $\mathrm{mSpec}(\mathbb{K})$, allowing us to define the product $X \times_{\mathrm{mSpec}(\mathbb{K})} \mathrm{mSpec}(\mathbb{L})$.

Exercise 2.10.5. Let $\mathbb{L}/\mathbb{K}$ be a finite extension. Suppose that R is a finitely generated graded $\mathbb{K}$-algebra. Prove that $\mathrm{mProj}(R \otimes \mathbb{L})$ is isomorphic (as a quasiprojective $\mathbb{L}$-scheme) to the product $\mathrm{mProj}(R) \times_{\mathrm{mSpec}(\mathbb{K})} \mathrm{mSpec}(\mathbb{L})$. Deduce that the base change of any quasiprojective scheme can be defined using the product construction.

Remark 2.10.6. The only issue with non-finite extensions $\mathbb{L}/\mathbb{K}$ is that we leave the realm of finitely generated $\mathbb{K}$-algebras. In the setting of arbitrary schemes, it is always true that $X_{\mathbb{L}} = X \times_{\mathrm{mSpec}(\mathbb{K})} \mathrm{mSpec}(\mathbb{L})$.

2.10.3 Geometric properties and products

It is important to remember that the product does not interact well with the notions of irreducibility or reducedness (Exercise 1.7.17). However, there is a way to correct this deficiency.

Definition 2.10.7. Fix an algebraic closure $\overline{\mathbb{K}}$ of $\mathbb{K}$. We say that a quasiprojective $\mathbb{K}$-scheme X is:

- geometrically irreducible, if the base change $X_{\overline{\mathbb{K}}}$ is irreducible.
- geometrically reduced, if the base change $X_{\overline{\mathbb{K}}}$ is reduced.

Exercise 2.10.8. Prove that geometrically irreducible implies irreducible and geometrically reduced implies reduced, but that the converse implications are false. (Exercise 1.4.12 asks you to prove this for affine schemes.)

The following results (which we will not prove) are crucial for understanding these two concepts.

Theorem 2.10.9. *Let R be a finitely generated $\mathbb{K}$-algebra. Fix a separable closure $\mathbb{K}^{sep}$ of $\mathbb{K}$. If $R \otimes_{\mathbb{K}} \mathbb{K}^{sep}$ has a unique minimal prime, then for every field extension $\mathbb{L}/\mathbb{K}$ there is a unique minimal prime in $R \otimes_{\mathbb{K}} \mathbb{L}$.*

Theorem 2.10.10. *Let R be a finitely generated $\mathbb{K}$-algebra. Fix a perfect closure $\mathbb{K}^{per}$ of $\mathbb{K}$. If $R \otimes_{\mathbb{K}} \mathbb{K}^{per}$ has no nilpotents, then for every field extension $\mathbb{L}/\mathbb{K}$ there are no nilpotents in $R \otimes_{\mathbb{K}} \mathbb{L}$.*

Exercise 2.10.11. Using Theorem 2.10.9 and Theorem 2.10.10 show that geometric irreducibility can be checked on any separable closure and geometric reducedness can be checked on any perfect closure. In particular, a reduced scheme over a perfect field is geometrically reduced.

Theorem 2.10.12. *Let X and Y be quasiprojective $\mathbb{K}$-schemes.*

(1) If X and Y are geometrically irreducible, then $X \times Y$ is (geometrically) irreducible.

(2) If X and Y are geometrically reduced, then $X \times Y$ is (geometrically) reduced.

Proof. It suffices to prove the statement when $X = \mathrm{mSpec}(R)$ and $Y = \mathrm{mSpec}(S)$ are affine and our ground field $\mathbb{K}$ is algebraically closed.

(1) Suppose that $R \otimes_{\mathbb{K}} S$ had a non-nilpotent zero divisor. Since this property is preserved upon quotienting by a nilpotent ideal, the ring $R \otimes_{\mathbb{K}} S/\operatorname{Nil}(S)$ would also have a non-nilpotent zero divisor. But this latter ring injects into $R \otimes_{\mathbb{K}} \operatorname{Frac}(S/\operatorname{Nil}(S))$ which has no non-nilpotent zero divisors by Theorem 2.10.9.

(2) Since S is nilpotent free, it admits an injection into $\prod_i S/\mathfrak{p}_i$ where we let $\mathfrak{p}_i$ vary over the finitely many minimal primes in S. Thus $R \otimes_{\mathbb{K}} S$ admits an injection into $\prod_i R \otimes \operatorname{Frac}(S/\mathfrak{p}_i)$. Since the latter ring has no nilpotents by Theorem 2.10.10, $R \otimes_{\mathbb{K}} S$ also has no nilpotents. $\square$

Warning 2.10.13. A relative product of geometrically irreducible/reduced schemes need not be irreducible/reduced, see e.g. Exercise 1.7.22.

2.10.4 Quasiprojectives are open in projectives

We are finally equipped to prove the remaining "half" of Theorem 2.4.13. The first step is to show that we can switch the order of open and closed embeddings.

Proposition 2.10.14. *Suppose that $f : X \to Y$ is a closed embedding and $g : Y \to Z$ is an open embedding. Let W denote the scheme-theoretic image of the composition $g \circ f : X \to Z$. Then the induced map $X \to W$ is an open embedding.*

Proof. Let $p_1 : Y \times_Z W \to Y$ and $p_2 : Y \times_Z W \to W$ be the projection maps. By Exercise 2.9.12 we see that p_1 is a closed embedding and p_2 is an open embedding. In particular it suffices to show that X is isomorphic to $Y \times_Z W$.

Using the universal property of the relative product we have a morphism $h : X \to Y \times_Z W$. Since f and p_1 are closed embeddings and f factors through p_1, we see that h is injective and has closed image. Since p_2 is an open embedding and the image of X in W is dense, we see that h has dense image. Altogether we see that g is a homeomorphism of the underlying sets.

Since the map $f^\sharp(U) : \mathcal{O}_Y(U) \to \mathcal{O}_X(f^{-1}U)$ is surjective for every open affine in Y, we see that $Y \times_Z W$ is covered by open affines for which $h^\sharp$ is surjective. Since W is the scheme-theoretic image of $g \circ f$, for every open affine in W the map on sheaves induced by $X \to W$ is injective. Since this map factors through h, we see that $Y \times_Z W$ is covered by open affines for which $h^\sharp$ is injective. Since localization is exact, we can refine these two open covers to find a cover by open affines for which $h^\sharp$ is an isomorphism. This implies that $h^\sharp$ induces an isomorphism of sheaves. $\square$

Theorem 2.10.15. *Every quasiprojective scheme X admits an open embedding into a projective scheme.*

Proof. By definition X admits an open embedding into $\operatorname{mProj}(R)$ for some finitely generated graded ring R. By Exercise 2.7.10 and Exercise 2.8.3, $\operatorname{mProj}(R)$ admits a closed embedding into $\operatorname{mSpec}(R_0) \times \mathbb{P}^n$ for some positive integer n. Since $\operatorname{mSpec}(R_0)$ admits a closed embedding into some affine space $\mathbb{A}^m$, by Exercise 2.9.12 the scheme $\mathrm{mSpec}(R_0) \times \mathbb{P}^n$ admits a closed embedding into $\mathbb{A}^m \times \mathbb{P}^n$. Finally, $\mathbb{A}^m \times \mathbb{P}^n$ admits an open embedding into $\mathbb{P}^m \times \mathbb{P}^n$.

Altogether, we have found a morphism $f : X \to \mathbb{P}^m \times \mathbb{P}^n$ which is a composition of closed embeddings and open embeddings. Applying Proposition 2.10.14 repeatedly, we can factor f as an open embedding $f_1 : X \to Y$ followed by a closed embedding $f_2 : Y \to \mathbb{P}^m \times \mathbb{P}^n$. Since $\mathbb{P}^m \times \mathbb{P}^n$ is projective, Theorem 2.8.12 shows that Y is also projective. $\square$

2.10.5 Exercises

Exercise 2.10.16. Show that for a finitely generated graded $\mathbb{K}$-algebra R the fibers of the structural map $p : \mathrm{mProj}(R) \to \mathrm{mSpec}(R_0)$ are projective. (Hint: show that the fiber over a point $\mathfrak{m} \in \mathrm{mSpec}(R_0)$ will be $\mathrm{mProj}(R/\mathfrak{m})$.)

Exercise 2.10.17. Compute some fibers...

Exercise 2.10.18. Set $X = \mathrm{mProj}(\mathbb{K}[w, x, y, z]/(wz - yx))$. Consider the graded homomorphism $f^\dagger : \mathbb{K}[s, t, u] \to \mathbb{K}[w, x, y, z]/(wz - yx)$ sending $s \mapsto w$, $t \mapsto x - y$, $u \mapsto z$.

(1) Prove that $f^\dagger$ induces a morphism $f : X \to \mathbb{P}^2$.

(2) Show that the set-theoretic fibers of f are the same as the orbits of the involution on X defined by the coordinate map $x \leftrightarrow y$.

2.11 Properness

In this section we develop the algebraic geometer's version of "compactness". In order to motivate our construction, we briefly review the topological situation. Suppose that X and Y are topological spaces that are locally compact and Hausdorff. If X is compact and $f : X \to Y$ is a continuous map, then:

(1) f is closed – that is, the f-image of any closed set in X is a closed set in Y – and

(2) the f-preimage of any closed subset of Y is compact in X.

More generally, a morphism $f : X \to Y$ is said to be proper if it is a closed map and every fiber of f is compact. (Note that X is compact iff the map from X to a point is proper iff every $f : X \to Y$ is proper. Thus properness is a generalization of the notion of compactness that applies to morphisms instead of spaces.)

Our analogue is:

Definition 2.11.1. Let $f : X \to Y$ be a morphism of quasiaffine schemes. We say that f is proper if it is universally closed, i.e. for every morphism $g : Z \to Y$ the induced map $\widetilde{f} : X \times_Y Z \to Z$ is a closed topological map.

We say that X is proper if the map $X \to \mathrm{mSpec}(\mathbb{K})$ is proper.

Remark 2.11.2. It turns out that a map of locally compact second-countable Hausdorff topological spaces is proper if and only if it is universally closed, so that our definition really is a close analogue of the geometric notion.

Example 2.11.3. The affine line $\mathbb{A}^1$ is not proper. (Remember, we interpret this to mean that the map $\mathbb{A}^1 \to \mathrm{mSpec}(\mathbb{K})$ is not proper.)

To see that $\mathbb{A}^1 \to \mathrm{mSpec}(\mathbb{K})$ is not proper, we base-change over the map $\mathbb{A}^1 \to \mathrm{mSpec}(\mathbb{K})$ to obtain the projection map $\mathbb{A}^2 \to \mathbb{A}^1$. This map is very far from being closed; for example, the image of the closed set $V(xy - 1)$ in $\mathbb{A}^2$ is an open set in $\mathbb{A}^1$.

Our first example of a proper map is:

Lemma 2.11.4. *Suppose that $f : X \to Y$ is a closed embedding. Then f is proper.*

Proof. Let $g : Z \to Y$ be any morphism. Then Exercise 2.9.12 shows that the map $X \times_Y Z \to Z$ is also a closed embedding. Since any closed embedding is a closed map, we conclude that f is universally closed. □

Exercise 2.11.5. Suppose that $f : X \to Y$ and $g : Y \to Z$ are proper morphisms of quasiprojective $\mathbb{K}$-schemes. Show that $g \circ f$ is proper.

Exercise 2.11.6. Prove that properness is stable under base change (in the sense of Exercise 2.9.12).

The main result of this section is:

Theorem 2.11.7. *If X is a projective scheme then every map $f : X \to Y$ of quasiprojective schemes is proper. In particular X itself is proper.*

This should be no surprise – after all, every closed subset of $\mathbb{P}^n_{\mathbb{C}}$ is compact in the Euclidean topology. The argument is known as the Fundamental Theorem of Elimination Theory.

Proof. We start by making several reductions. By assumption X admits a closed embedding $i : X \to \mathbb{P}^n$. We will consider the map $f : X \to Y$ as the composition of the morphisms

$$\begin{tikzcd} X \arrow[r, "{(id,f)}"] \arrow[drr, "f"'] & X \times Y \arrow[r, "{(i,id)}"] & \mathbb{P}^n \times Y \arrow[d, "\pi_2"] \\ & & Y \end{tikzcd}$$

As shown in Exercise 2.9.14 the map $(id, f) : X \to X \times Y$ is a closed embedding. Lemma 2.11.4 shows that (id, f) is proper. By assumption the inclusion map $i : X \to \mathbb{P}^n$ is a closed embedding, thus Exercise 2.11.4 and Exercise 2.11.6 show that i is proper. Since properness is closed under composition by Exercise 2.11.5, it suffices to show that the projection map $\pi_2 : \mathbb{P}^n \times Y \to Y$ is proper.

Note that if $g : Z \to Y$ is any morphism then $\mathbb{P}^n \times_Y Z$ is the same as the projection map $\mathbb{P}^n \times Z \to Z$. Thus by varying Y it suffices to prove that $\mathbb{P}^n \times Y \to Y$ is closed for any quasiprojective scheme Y. Suppose we take an open cover of Y by open affines $\{U_i\}$. Then the projection map $\mathbb{P}^n \times Y \to Y$ will be a closed map if and only if each projection map $\mathbb{P}^n \times U_i \to U_i$ is a closed map. Thus it suffices to consider the case when Y is affine. By arguing on each component of Y separately we may suppose that Y is irreducible. Finally, since we only care about the topology of the map we may suppose that Y is reduced. Thus we have reduced to the case when $Y = \mathrm{mSpec}(R)$ is an affine variety so that R is a domain.

By Exercise 2.9.11 we know that $\mathbb{P}^n \times \mathrm{mSpec}(R)$ is isomorphic to $\mathrm{mProj}(R[x_0, \ldots, x_n])$. For convenience we denote the ring $R[x_0, \ldots, x_n]$ by S. Fix a homogeneous ideal $I = (h_1, \ldots, h_r)$. Then we have a map

$$\psi_m : S_{m-\deg(h_1)} \oplus \ldots \oplus S_{m-\deg(h_r)} \xrightarrow{(\cdot h_1, \ldots, \cdot h_r)} S_m.$$

whose image is the mth homogeneous part of I. Note that ψ_m is a map of finitely generated free R-modules and thus defined by a matrix M_m with entries in R. Given any maximal ideal $\mathfrak{m}$, the tensor product of ψ_m with $R/\mathfrak{m}$ will be surjective if and only if none of the $\dim_R(S_m)$-minors of M_m vanish at $\mathfrak{m}$.

We want to show that $f(V_+(I))$ is closed in $\mathrm{mSpec}(R)$ for any homogeneous ideal I in $R[x_0, \ldots, x_n]$. Note that for a maximal ideal $\mathfrak{m} \subset R$ we have

$$\begin{aligned} \mathfrak{m} \in f(V_+(I)) \quad &\Leftrightarrow \quad I/\mathfrak{m} \subset (R/\mathfrak{m})[x_0, \ldots, x_n] \text{ has non-empty vanishing locus} \\ &\Leftrightarrow \quad \sqrt{I/\mathfrak{m}} \neq (x_0, \ldots, x_n) \\ &\Leftrightarrow \quad (I/\mathfrak{m})_m \not\supset (R/\mathfrak{m})[x_0, \ldots, x_n]_m \text{ for any } m > 0 \end{aligned}$$

As discussed earlier, for any fixed m the locus of $\mathfrak{m}$ such that $\psi_m \otimes R/\mathfrak{m}$ fails to be surjective will be a closed subset of $\mathrm{mSpec}(R)$. The set of $\mathfrak{m}$ such that $\psi_m \otimes R/\mathfrak{m}$ fails to be surjective for all m will be an intersection of an infinite number of closed subsets of $\mathrm{mSpec}(R)$, hence closed. As shown above $f(V_+(I))$ is precisely this closed subset. □

Remark 2.11.8. It is natural to wonder whether the projective schemes are the only proper schemes. In the setting of quasiprojective $\mathbb{K}$-schemes the answer is yes. Indeed, Theorem 2.10.15 shows that every quasiprojective scheme admits an open embedding $i : U \to X$ to a projective scheme. If U is not projective then the image of i is not closed so that i is not proper.

However, in the more general setting of arbitrary schemes there are proper schemes which are not projective.

Exercise 2.11.9. Let R be a finitely generated graded $\mathbb{K}$-algebra. Prove that the canonical morphism $\mathrm{mProj}(R) \to \mathrm{mSpec}(R_0)$ is proper. This is the correct generalization of Theorem 2.11.7 for the mProj construction. (Hint: first show that $\mathrm{mProj}(R)$ admits a closed embedding into $\mathbb{P}^n \times \mathrm{mSpec}(R_0)$ for some n that forms a commutative diagram with the structural maps to $\mathrm{mSpec}(R_0)$.)

We give a couple applications of Theorem 2.11.7.

Corollary 2.11.10. *Let X be a projective scheme and let $f : X \to Y$ be a morphism of quasiprojective schemes. Then the scheme-theoretic image of f is a projective scheme.*

Recall that the scheme-theoretic image was constructed in Construction 2.8.11.

Proof. Since Y is quasiprojective it admits an open embedding $i : Y \to Z$ into a projective scheme Z. The scheme-theoretic image of $i \circ f$ is a closed subscheme of Z, hence a projective scheme. But by Theorem 2.11.7 $i \circ f$ is a closed map, so that every point of Z is in the set-theoretic image of $i \circ f$. In particular Z is contained in the set-theoretic image of f. Thus the scheme-theoretic image of f is also equal to Z. □

Corollary 2.11.11. *Let X be a projective $\mathbb{K}$-scheme. Then $\mathcal{O}_X(X)$ is a finite dimensional algebra over $\mathbb{K}$.*

Proof. The hard part of the argument is that $\mathcal{O}_X(X)$ is a finitely generated $\mathbb{K}$-algebra, which we will assume for now. Suppose for a contradiction that $\mathcal{O}_X(X)$ fails to be an Artinian ring. Thus there is a function $f \in \mathcal{O}_X(X)$ which is not a unit and not a zero divisor.

By Theorem 2.4.8 there is a function $f : X \to \mathbb{A}^1$ defined by the ring map $\mathbb{K}[t] \to \mathcal{O}_X(X)$ sending $t \mapsto f$. We claim that f is dominant. Indeed, if we choose any open affine $U \subset X$ such that f does not vanish on U then the map $f|_U$ is defined by $\mathbb{K}[t] \to \mathcal{O}_X(X) \xrightarrow{\rho_{X,U}} \mathcal{O}_X(U)$. Since this ring morphism is injective, the map $f|_U$ is dominant.

Since X is projective, the dominant map $f : X \to \mathbb{A}^1$ must be surjective. However, if we then compose f with the inclusion $\mathbb{A}^1 \hookrightarrow \mathbb{P}^1$ we get a morphism from X to a quasiprojective scheme whose image is not closed. By Theorem 2.11.7 this gives us a contradiction. $\square$

2.11.1 Exercises

Exercise 2.11.12. Suppose that X is a projective variety over an algebraically closed field $\mathbb{K}$. Prove that $\mathcal{O}_X(X) = \mathbb{K}$. (What is the correct analogue over an arbitrary ground field?)

Show that the converse fails: a scheme X can be reducible or non-reduced and still have $\mathcal{O}_X(X) = \mathbb{K}$. (Hint: consider either the union of two lines or the double line in $\mathbb{P}^2$.)

Exercise 2.11.13. Prove that the following are equivalent:

(1) X is a quasiprojective scheme with finitely many points.

(2) X is both projective and affine.

Exercise 2.11.14. Let $H \subset \mathbb{P}^n$ be a hypersurface. Suppose that $X \subset \mathbb{P}^n$ does not have only finitely many points. Show that $H \cap X \neq \emptyset$.

Chapter 3

First examples

In this chapter we will study some first examples of projective varieties. The most fundamental examples of algebraic varieties are all motivated by linear algebra. (These are also the examples which can be studied without using more advanced tools from scheme theory.) All the examples we will see in this chapter are still being actively studied by mathematicians today.

Suppose that we are given a closed subscheme $X \subset \mathbb{P}^n$. There are three different ways in which an ideal I could be associated to X:

(1) The ideal I could define X set-theoretically, i.e. $V_+(I) = X$ as closed subsets.

(2) The ideal I could define X scheme-theoretically, i.e. $\mathrm{mProj}(\mathbb{K}[x_0, \ldots, x_n]/I) \cong X$.

(3) The ideal I could be the saturated ideal defining X, i.e. every homogeneous function that vanishes along X is contained in I.

These three possibilities successively tell us more about X but are successively harder to prove rigorously.

3.0.1 Algebraic preliminaries

When we study the Grassmannian, we will need to know some properties of the exterior product. Suppose that V is an n-dimensional $\mathbb{K}$-vector space. Fix an integer k satisfying $0 < k < n$. We have two perfect pairings

$$\bigwedge^{n-k} V^\vee \times \bigwedge^{n-k} V \to \mathbb{K}$$
$$\bigwedge^{k} V \times \bigwedge^{n-k} V \to \bigwedge^{n} V$$

Choosing an isomorphism $\bigwedge^n V \to \mathbb{K}$ we can identify $\bigwedge^k V$ and $\bigwedge^{n-k} V^\vee$. This isomorphism is only natural up to scaling, but when we projectivize this ambiguity won't matter.

Suppose we choose a basis $v_1, \ldots, v_r$ of V with dual basis $v_1^\vee, \ldots, v_r^\vee$. For any subset $I \subset \{1, \ldots, r\}$ of size k our isomorphism identifies

$$v_{i_1} \wedge \ldots \wedge v_{i_k} \leftrightarrow \pm v_{j_1}^\vee \wedge \ldots \wedge v_{j_{n-k}}^\vee$$

where J is the complement of I. Given an element $\omega \in \bigwedge^k V$ we will denote the corresponding element in $\bigwedge^{n-k} V^\vee$ by ω^*.

Given any element $\omega \in \bigwedge^k V$ we will associate two subspaces D_ω, L_ω of V. The goal of this subsection is to define and analyze these two subspaces.

Definition 3.0.1. Given any $\omega \in \bigwedge^k V$, we define D_ω to be the set of elements $v \in V$ such that $v \wedge \rho = \omega$ for some $\rho \in \bigwedge^{k-1} V$.

The following result verifies that D_ω is actually a subspace.

Lemma 3.0.2. *Let V be a $\mathbb{K}$-vector space of dimension n and fix $0 < k < n$. Suppose that $\omega \in \bigwedge^k V$. Then D_ω is the kernel of the map*

$$\begin{aligned} \varphi_\omega : V &\to \bigwedge^{k+1} V \\ v &\mapsto v \wedge \omega \end{aligned}$$

Proof. Certainly D_ω is in the kernel of φ_ω. Conversely, if $v \in \ker(\varphi_\omega)$, then we can choose a basis $\{v_1, v_2, \ldots, v_n\}$ of V with $v = v_1$. If we write ω in this basis, we get an expression $\omega = \sum_I a_I v_I$ where I varies over strictly increasing k-tuples of integers in the set $\{1, \ldots, n\}$. By assumption

$$0 = v_1 \wedge \omega = \sum_I a_I (v_1 \wedge v_I).$$

It is clear that no canceling is possible for the indices I which do not contain 1. Thus the coefficients for such indices I must vanish, showing that every summand of ω must involve v_1 somewhere. This implies that $v \in D_\omega$. $\square$

Corollary 3.0.3. *Let V be a $\mathbb{K}$-vector space of dimension n and fix $0 < k < n$. Suppose that $\omega \in \bigwedge^k V$. Choose a basis $\{v_1, \ldots, v_n\}$ of V such that $v_1, \ldots, v_s$ span D_ω. Then $\omega = v_1 \wedge \ldots \wedge v_s \wedge \rho$ for some $\rho \in \bigwedge^{k-s} V$. Furthermore any other subspace of V satisfying this property will be contained in D_ω.*

Proof. Suppose we expand $\omega = \sum a_I v_I$ in our chosen basis. The proof of Lemma 3.0.2 shows that for $i = 1, \ldots, s$ the vector v_i must appear in every summand of ω which has a non-zero coefficient. The first statement follows. The last property is clear from the definition of D_ω. $\square$

We will define our next subspace using duality.

Definition 3.0.4. We define L_ω to be the annihilator of $D_{\omega^*} \subset V^\vee$.

By Lemma 3.0.2 L_ω is the image of the map $(\varphi_{\omega^*})^\vee : \bigwedge^{n-k+1} V \to V$ dual to the map $\varphi_{\omega^*} : V^\vee \to \bigwedge^{n-k+1} V^\vee$. The following result clarifies the significance of L_ω.

Lemma 3.0.5. *Let V be a $\mathbb{K}$-vector space of dimension n and fix $0 < k < n$. Suppose that $\omega \in \bigwedge^k V$. Then ω is in the image of $\bigwedge^k L_\omega \to \bigwedge^k V$. Furthermore any other subspace of V satisfying this property will contain L_ω.*

Proof. Fix a basis $\{v_1, v_2, \ldots, v_n\}$ of V such that $\{v_i^\vee\}_{i=1}^s$ is a basis for D_{ω^*}. By Corollary 3.0.3 we can write $\omega^* = v_1^\vee \wedge \ldots \wedge v_s^\vee \wedge v^\vee$ for some $v^\vee \in \bigwedge^{n-k-s} V^\vee$. Using the coordinate description of the identification $\omega \leftrightarrow \omega^*$, we see that no summand of ω can involve $v_1, v_2, \ldots, v_s$. In other words, ω is contained in the kth exterior product of the annihilator $\text{Span}\{v_{s+1}, \ldots, v_n\}$ of D_{ω^*}.

To see the final statement, suppose that L is any subspace of V satisfying the desired property. Choose a basis $\{v_1, \ldots, v_n\}$ of V so that L is the span of $\{v_i\}_{i=1}^s$. Using the coordinate description of the identification $\omega \leftrightarrow \omega^*$ we see that the annihilator of L is contained in D_{ω^*}. □

The duality between D_ω and L_ω can be emphasized in the following way. We say that $\rho \in \bigwedge^a V$ divides $\psi \in \bigwedge^b V$ if $a \leq b$ and there exists some $\xi \in \bigwedge^{b-a} V$ such that $\rho \wedge \xi = \psi$. Then

- D_ω is the largest subspace of V such that $\bigwedge^{\dim D_\omega} D_\omega$ divides ω.
- L_ω is the smallest subspace of V such that ω divides $\bigwedge^{\dim L_\omega} L_\omega$.

It is not a priori obvious that there must be a largest or a smallest such subspace, but this follows from the proofs above. Altogether we have:

Proposition 3.0.6. *Let V be a $\mathbb{K}$-vector space of dimension n and fix $0 < k < n$. Suppose that $\omega \in \bigwedge^k V$. Then the following are equivalent:*

(1) *ω is a pure wedge product, i.e. $\omega = v_1 \wedge \ldots \wedge v_k$ for a set of linearly independent vectors $\{v_i\}$ in V.*

(2) *D_ω has dimension k.*

(3) *L_ω has dimension k.*

3.1 Quadric hypersurfaces

Recall that a hypersurface in projective space is the vanishing locus of a single homogeneous equation $f \in \mathbb{K}[x_0, \ldots, x_n]$. We say that the hypersurface X has degree d if the defining equation f has degree d. Hypersurfaces of degree 1 are the (codimension 1) hyperplanes in $\mathbb{P}^n$ and each is isomorphic to $\mathbb{P}^{n-1}$.

In this section we will study the next simplest example. Throughout we will assume that $\mathrm{char}(\mathbb{K}) \neq 2$.

Definition 3.1.1. A quadric hypersurface $X \subset \mathbb{P}^n$ is a degree 2 hypersurface.

Example 3.1.2. The most familiar examples of quadric hypersurfaces are the conics $C \subset \mathbb{P}^2$. Note that the intersection of C with each affine chart D_{+,x_i} will be the vanishing locus in $\mathbb{A}^2$ of an equation of degree at most 2. Furthermore, there will be at least one chart such that $C \cap D_{+,x_i}$ is the vanishing locus of an equation of degree exactly 2.

The classification of conics in $\mathbb{A}^2_{\mathbb{C}}$ was discussed in Exercise 1.5.17. The classification becomes much simpler when we work in $\mathbb{P}^2$ instead.

Quadric hypersurfaces have a close relationship with symmetric bilinear forms. Suppose that

$$f = \sum_{i=0}^{n} a_i x_i^2 + \sum_{i<j} a_{ij} x_i x_j.$$

We associate to f the symmetric bilinear form Q on $\mathbb{K}^{n+1}$ defined by $Q(\vec{v}, \vec{w}) = \frac{1}{2}(f(\vec{v} + \vec{w}) - f(\vec{v}) - f(\vec{w}))$. If we define the symmetric matrix M which has diagonal entries a_i and off-diagonal entries $a_{ij}/2$, then $Q(\vec{v}, \vec{w}) = \vec{v}^t \cdot M \cdot \vec{w}$. Note that we can recover the hypersurface from M.

We next turn to the problem of classifying the isomorphism classes of quadric hypersurfaces. Note that the isomorphism type of a quadric hypersurface X is unchanged by a linear change of coordinates. Furthermore, the action of coordinate changes on quadric hypersurfaces is equivalent to the action of coordinate changes on the corresponding symmetric bilinear forms. (Check!) Since every symmetric matrix over a field $\mathbb{K}$ is diagonalizable, we see:

Lemma 3.1.3. *Every quadric hypersurface $X \subset \mathbb{P}^n_{\mathbb{K}}$ is isomorphic to the vanishing locus of an equation of the form*

$$f = \sum_{i=0}^{n} a_i x_i^2$$

for some $a_i \in \mathbb{K}$.

This isn't yet quite enough to determine the isomorphism types – for example, the behavior of $\mathbb{K}^\times/(\mathbb{K}^\times)^2$ will need to be accounted for. Depending on the ground field the classification can be a bit complicated.

Example 3.1.4. Suppose that $\mathbb{K}$ is an algebraically closed field (with $\mathrm{char}(\mathbb{K}) \neq 2$). Then any quadric hypersurface Q in $\mathbb{P}^n_{\mathbb{K}}$ is isomorphic to the vanishing locus of an equation of the form

$$f = \sum_{i=0}^{r} x_i^2$$

for some $0 \leq r \leq n$. In this setting the number $(r+1)$ is called the rank of the quadric hypersurface.

It turns out that two quadrics are isomorphic if and only if they have the same rank. The easiest way to prove this is to use the notion of "singular points" from later in the course: the singular locus of Q is the $(n-r-1)$-dimensional plane L contained in Q defined by the equations $x_0 = \ldots = x_r = 0$. For example, there are four isomorphism types of quadrics in $\mathbb{P}^3$ – the "smooth" quadric (rank 4), the quadric cone (rank 3), the union of two planes (rank 2), and the double plane (rank 1).

Here is a geometric description of Q when $0 < r < n$. Let T denote the r-dimensional plane $x_{r+1} = \ldots = x_n = 0$. Note that $T \cap Q$ is a full rank quadric hypersurface in T. Then Q is the cone over $T \cap Q$ with vertex L – in other words, Q is the union of the $(n-r)$-dimensional planes which are spanned by L and a point of $T \cap Q$.

Example 3.1.5. Every quadric hypersurface in $\mathbb{P}^n_{\mathbb{R}}$ can be described by an equation of the form

$$f = \sum_{i=0}^{r} x_i^2 - \sum_{j=r+1}^{s} x_j^2.$$

Example 3.1.6. Over an algebraically closed field (with $\mathrm{char}(\mathbb{K}) \neq 2$) there are three isomorphism classes of conic in $\mathbb{P}^2$ – the irreducible reduced conics, the unions of two lines, and the double lines – corresponding to the three equations $x_0^2 + x_1^2 + x_2^2$, $x_0^2 + x_1^2$, and x_0^2.

Over other fields the classification becomes more complicated. We will just focus on the irreducible reduced conics. It turns out that the $\mathbb{K}$-isomorphism classes of such conics are in bijection with the $\mathbb{K}$-isomorphism classes of generalized quaternion algebras (or equivalently, the 2-torsion elements in the Brauer group $\mathrm{Br}(\mathbb{K})$). Precisely, the generalized quaternion algebra $Q(a, b)$ corresponds to the conic $ax_0^2 + bx_1^2 = x_2^2$. For example, over $\mathbb{R}$ we have the trivial matrix algebra and the usual quaternions, and these correspond respectively to the conic $V_+(x_0^2 + x_1^2 - x_2^2)$ and the conic $V_+(x_0^2 + x_1^2 + x_2^2)$. We can see directly that these two conics are not isomorphic because the first admits many $\mathbb{R}$-points and the second does not admit any $\mathbb{R}$-points.

For a number field $\mathbb{K}$ the Brauer-Hasse-Noether theorem describes an exact sequence

$$0 \to \mathrm{Br}(\mathbb{K}) \to \bigoplus_{\text{places } v} \mathrm{Br}(\mathbb{K}_v) \to \mathbb{Q}/\mathbb{Z} \to 0.$$

Each local field at a finite place has Brauer group $\mathbb{Q}/\mathbb{Z}$ and the map on the right is the addition map. Thus there are infinitely many different isomorphism types of irreducible conics over a number field. In Exercise 3.1.12 and Exercise 3.1.13 we will see some examples of non-isomorphic conics over $\mathbb{Q}$.

It turns out that the classification of quadrics over $\mathbb{K}$ is the same as the classification of symmetric bilinear forms:

Theorem 3.1.7. *Two quadric hypersurfaces in $\mathbb{P}^n$ are isomorphic if and only if they can be identified under a linear change of variables.*

We do not yet have the tools to prove this result.

Remark 3.1.8. It turns out that the classification of quadric hypersurfaces up to birational equivalence is much trickier! Apparently this is still unsolved for general fields $\mathbb{K}$.

3.1.1 Moduli of quadric hypersurfaces

Note that a quadric hypersurface in $\mathbb{P}^n$ is determined by a choice of $\binom{n+2}{2}$ coefficients up to rescaling. In other words, the quadric hypersurfaces are parametrized by the traditional points of a $\mathbb{P}^{(n+2)(n+1)/2-1}$. Thus it makes sense to think of this projective space as the "moduli space" parametrizing quadric hypersurfaces. For clarity we will write M to denote this moduli space.

It is natural to ask for the meaning of the non-traditional points on M. Recall that a non-traditional point corresponds to a $\mathrm{Gal}(\overline{\mathbb{K}}/\mathbb{K})$-orbit of points on the base change $M_{\overline{\mathbb{K}}}$. Thus, the non-traditional points of $M_{\mathbb{K}}$ define the hypersurfaces defined over $\mathbb{K}$ which are unions of the Galois conjugates of a $\overline{\mathbb{K}}$-quadric.

Example 3.1.9. The moduli space of conics on $\mathbb{P}^2_{\mathbb{R}}$ is $M \cong \mathbb{P}^5_{\mathbb{R}}$. We will use coordinates $y_0, \ldots, y_5$ on M and will associate the traditional point $(a : b : c : d : e : f)$ with the equation $ax^2 + bxy + cy^2 + dxz + eyz + fz^2$.

Note that $x^2 + y^2 + iz^2 = 0$ defines a conic over $\mathbb{C}$ but not over $\mathbb{R}$; its Galois conjugate is $x^2 + y^2 - iz^2 = 0$. The union of these two hypersurfaces is given by the equation

$$(x^2 + y^2 + iz^2)(x^2 + y^2 - iz^2) = x^4 + 2x^2y^2 + y^4 + z^4.$$

This hypersurface is defined over $\mathbb{R}$ and corresponds to the non-traditional point $(y_1, y_3, y_4, y_2 - y_0, y_5^2 + y_0^2)$ of M.

3.1.2 Exercises

Exercise 3.1.10. Let $Q \subset \mathbb{A}^2_{\mathbb{Q}}$ be the conic defined by $x^2 + y^2 = 1$. Write down the equations for the map $Q \dashrightarrow \mathbb{P}^1$ obtained by projecting away from $x = (-1, 0)$. Show that this map induces a bijection on the $\mathbb{Q}$-points in $Q\backslash\{x\}$ and the $\mathbb{Q}$-points in $\mathbb{P}^1\backslash\{\infty\}$. Explain how the formula you wrote down describes all primitive Pythagorean triples.

Exercise 3.1.11. Suppose that $Q \subset \mathbb{P}^2_{\mathbb{K}}$ is an irreducible conic which has a $\mathbb{K}$-point x. Projection away from the point x in $\mathbb{P}^2$ defines a rational map $\phi : Q \dashrightarrow \mathbb{P}^1$.

(1) Let $U = Q \backslash \{x\}$. Show that ϕ is an isomorphism from U to its image, and in particular is a birational map.

(2) Prove that ϕ extends to an isomorphism on all of Q.

(We saw an example of this behavior in Warning 2.7.3.)

Exercise 3.1.12. Let Q_1 be the conic $x^2+y^2+z^2 = 0$ and let Q_2 be the conic $x^2+y^2 = 3z^2$ in $\mathbb{P}^2_{\mathbb{Q}}$.

(1) Show that neither Q_1 nor Q_2 have any $\mathbb{Q}$-points. (Thus we cannot use rational points to distinguish these two conics.)

(2) Prove that Q_2 has $\mathbb{Q}(\sqrt{3})$-points but that Q_1 does not. Deduce that Q_1 and Q_2 are not isomorphic.

Exercise 3.1.13. As discussed in Example 3.1.6, one can also show that two conics over a number field are non-isomorphic by base-changing to local fields. (See Section 2 of Serre's book "A course in arithmetic.")

Let Q_1 be the conic $2x^2 + 5y^2 = z^2$ and let Q_2 be the conic $x^2 + y^2 = 3z^2$ in $\mathbb{P}^2_{\mathbb{Q}}$. Show that the base changes of Q_1 and Q_2 to $\mathbb{Q}_5$ are not isomorphic by showing that Q_2 admits $\mathbb{Q}_5$-points but that Q_1 does not.

Deduce that Q_1 and Q_2 are not isomorphic over $\mathbb{Q}$.

Exercise 3.1.14. Suppose that $\mathbb{F}_q$ is a finite field of characteristic $\neq 2$.

(1) Show that for non-zero $a, b \in \mathbb{F}_q$ the equation $ax^2 + by^2 = 1$ always has a solution $(x, y) \in \mathbb{F}_q^2$ by using a counting argument.

(2) Apply Exercise 3.1.11 to deduce that every irreducible conic over $\mathbb{F}_q$ is isomorphic to $\mathbb{P}^1$.

Exercise 3.1.15. In this exercise and the next we study the lines on quadrics in $\mathbb{P}^3_{\mathbb{C}}$. (The rank 1 and rank 2 cases are both unions of planes, so the lines are easy to understand.)

Consider the quadric $Q \subset \mathbb{P}^3_{\mathbb{C}}$ defined by the equation $wz - xy = 0$. Check that this quadric has full rank.

For any $(s : t) \in \mathbb{P}^1$ consider the line in $\mathbb{P}^3$ defined by the equations $tw = sy$ and $tx = sz$. Prove that for any choice of s, t the corresponding line is contained in Q. In other words, in this way we obtain a one-dimensional family of lines parametrized by $\mathbb{P}^1$.

Similarly, for any $(a : b) \in \mathbb{P}^1$ consider the line given by $bw = ax$ and $by = az$. Prove that this gives a different family of lines on Q.

Exercise 3.1.16. Consider the quadratic cone Q defined by the equation $x^2 + y^2 = z^2$. Prove that the intersection of Q with a hyperplane P is either:

- a irreducible reduced conic in P if P does not contain the cone point $(1 : 0 : 0 : 0)$, or
- a union of lines or a double line if P does contain the cone point.

In particular, deduce that every line in Q contains the cone point.

3.2 Veronese embeddings

Suppose that V is a vector space and consider the map

$$\begin{aligned} V &\to \mathrm{Sym}^d(V) \\ v &\mapsto v^d \end{aligned}$$

This function is known as the dth Veronese map. It is natural to wonder what the image looks like: which elements of $\mathrm{Sym}^d(V)$ are dth powers? It is more convenient to work with the projectivized version $\mathbb{P}(V) \to \mathbb{P}(\mathrm{Sym}^d(V))$. In this section we will identify the image of this map as a projective variety and study its properties.

3.2.1 Rational normal curves

We start with the case when V is two-dimensional. In practice it is a little easier to fix coordinates: if we fix a basis e_1, e_2 of V then $\mathrm{Sym}^d(V)$ is spanned by $e_1^d, e_1^{d-1}e_2, \ldots, e_2^d$. Thus on traditional points the Veronese map can be described via

$$\begin{aligned} &f : \mathbb{P}^1 \to \mathbb{P}^d \\ &(s:t) \mapsto (s^d : s^{d-1}t : s^{d-2}t^2 : \ldots : st^{d-1} : t^d) \end{aligned}$$

The corresponding map of projective varieties is defined by

$$\mathbb{K}[y_0, \ldots, y_d] \to \mathbb{K}[x_0, x_1]$$

sending $y_i \mapsto x_0^i x_1^{d-i}$. In other words, the Veronese map is the geometric counterpart of passing to a degree d Veronese subring.

Note that we could equally well have used a different basis for $\mathrm{Sym}^d(V)$ without fundamentally changing the nature of this map. Thus we define:

Definition 3.2.1. A rational normal curve X of degree d in $\mathbb{P}^d$ is the image of the Veronese map described above *or* its translate under any linear change of coordinates.

For simplicity we will focus on the rational normal curve X which is the image of the original map.

Proposition 3.2.2. *The dth Veronese map $f : \mathbb{P}^1 \to \mathbb{P}^d$ is an isomorphism from $\mathbb{P}^1$ onto the vanishing locus X of the ideal*

$$I = \{y_i y_j - y_k y_l \,|\, i + j = k + l\}.$$

Furthermore I is the saturated ideal defining X.

Note that I is generated by the quadratic equations which "obviously" vanish along X. The main point is to show that there are no other functions which vanish on X.

On the level of traditional points the inverse rational map $X \dashrightarrow \mathbb{P}^1$ is given by choosing an index $i < d$ and sending $(a_0 : \ldots : a_d) \mapsto (a_i : a_{i+1})$. As we vary i, the resulting maps all lie in the same equivalence class of rational maps and are naturally defined on different open subsets of X. Together the domains cover X so that this rational map is actually a morphism.

Proof. The map on homogeneous coordinate rings is the composition of $\mathbb{K}[y_0, \ldots, y_d] \to \mathbb{K}[x_0, x_1]^{(d)}$ and the dth Veronese inclusion $\mathbb{K}[x_0, x_1]^{(d)} \to \mathbb{K}[x_0, x_1]$. By Proposition 2.7.7 the second map defines an isomorphism $\mathbb{P}^1 \to \mathrm{mProj}(\mathbb{K}[x_0, x_1]^{(d)})$. The first map is a surjection, so by Exercise 2.8.3 it defines a closed embedding $\mathrm{mProj}(\mathbb{K}[x_0, x_1]^{(d)}) \to \mathbb{P}^d$. Thus the map takes $\mathbb{P}^1$ isomorphically to its image.

We next verify that the vanishing locus of I defines the image of the Veronese map. It is clear that I vanishes along this subscheme. To check that it actually defines this subscheme we restrict our attention to the distinguished open affine D_{+,y_i}. The preimage of this open subset is D_{+,x_0x_1} and the map $f : D_{+,x_0x_1} \to D_{+,y_i}$ is defined by

$$\begin{aligned} f^\sharp : \mathbb{K}\left[\frac{y_0}{y_i}, \ldots, \frac{y_d}{y_i}\right] &\to \mathbb{K}\left[\frac{x_0}{x_1}, \frac{x_1}{x_0}\right] \\ \frac{y_j}{y_i} &\mapsto x_0^{j-i} x_1^{i-j} \end{aligned}$$

As we observed earlier the localization of I is contained in the kernel of this map. In fact using the relations determined by the generators of I we see that in the quotient $\mathbb{K}[\frac{y_0}{y_i}, \ldots, \frac{y_d}{y_i}]/I_{y_i}$ we have

$$\frac{y_{i-1}}{y_i} = \left(\frac{y_{i+1}}{y_i}\right)^{-1} \qquad \frac{y_{i+j}}{y_i} = \left(\frac{y_{i+1}}{y_i}\right)^{j} \qquad \frac{y_{i-k}}{y_i} = \left(\frac{y_{i-1}}{y_i}\right)^{k}.$$

Together these relations show that after quotienting by the localization of I the map $f^\sharp$ becomes an isomorphism, so the kernel of $f^\sharp$ is the localization of I. By varying the chart, we see that I defines the image of X as a closed subscheme of $\mathbb{P}^d$.

It is a little trickier to prove that I is saturated. It suffices to prove that any homogeneous polynomial that vanishes on X is contained in I. By induction on degree one can show that any homogeneous polynomial $P(y_0, \ldots, y_d)$ can be written as

$$P(y_0, \ldots, y_d) = R_0(y_0, y_d) + R_1(y_0, y_d)y_1 + \ldots + R_{d-1}(y_0, y_d)y_{d-1} + T$$

for some polynomials $R_0, \ldots, R_{d-1}$ and some polynomial $T \in I$. If P vanishes on X, then $P(s^d, s^{d-1}t, \ldots, t^d)$ is identically zero. In the expression above, $T(s^d, \ldots, t^d)$ is also identically zero. Furthermore, the exponents of s, t occurring in $R_0, R_1, \ldots, R_{d-1}$ will be different modulo d. In order for this function to vanish identically we must have that

$R_0 = \ldots = R_{d-1} = 0$ identically. We conclude that P lies in I. (In other words, I is the kernel of the map $f^\dagger$, thus a prime ideal, thus saturated.) $\square$

Remark 3.2.3. Proposition 3.2.2 describes the saturated ideal defining the image of the Veronese map, but there can be ideals with fewer generators whose vanishing locus is the same set. For example, when $d \geq 4$ some of the quadratic generators are redundant for defining the closed subscheme X.

Example 3.2.4. A rational normal curve in $\mathbb{P}^3$ is known as a twisted cubic. In the coordinates w, x, y, z the image X of the Veronese map is defined by the ideal

$$I = (wz - xy, x^2 - wy, y^2 - xz)$$

We will verify in Proposition 6.3.6 that there is no homogeneous ideal I with two generators such that $X = V_+(I)$. This is traditionally the first example demonstrating a general phenomenon: a "codimension k set" in $\mathbb{P}^n$ may require more than k equations to define it.

3.2.2 Veronese varieties

The picture in higher dimension is essentially the same. It will be convenient to write coordinates on the target $\mathbb{P}^{\binom{n+d}{d}-1}$ in a slightly unusual way: we let $\mathbb{K}[y_I]$ be the polynomial ring whose generators are indexed by the n-tuples $I = (i_0, \ldots, i_n)$ of non-negative integers satisfying $i_0 + \ldots + i_n = d$. Define the map $\mathbb{K}[y_I] \to \mathbb{K}[x_0, \ldots, x_n]$ sending $y_I \mapsto x^I$.

Proposition 3.2.5. *The dth Veronese map $f : \mathbb{P}^n \to \mathbb{P}^{\binom{n+d}{d}-1}$ defines an isomorphism from $\mathbb{P}^n$ to a closed subvariety X. The saturated ideal defining X has generators*

$$\{y_I y_J - y_K y_L \mid I + J = K + L \text{ as vectors}\}.$$

As in Proposition 3.2.2 the main point is to show that this ideal defines X and that it is saturated. The first step is similar to what we did for Proposition 3.2.2.

Exercise 3.2.6. Check that the ideal defines the closed subscheme X by passing to affine charts. (If you prefer, you may suppose $\mathbb{K}$ is algebraically closed and argue using $\mathbb{K}$-points.)

To prove that the ideal is saturated, one could use Gröbner basis techniques. Here is an alternative approach based on representation theory:

Exercise 3.2.7. Define the map $f^\dagger : \mathbb{K}[y_I] \to \mathbb{K}[x_0, \ldots, x_n]^{(d)}$ sending $y_I \mapsto x^I$. Let V denote the vector space of homogeneous linear functions in the x_i. The linear homogeneous functions in the y_I can be identified with $\mathrm{Sym}^d(V)$. Thus, on the degree n level $f^\dagger$ is defined as

$$\phi_n : \mathrm{Sym}^n(\mathrm{Sym}^d(V)) \to \mathrm{Sym}^{nd}(V).$$

We can write the kernel as a direct sum of Schur functors $S^\lambda(V)$; for example, we have

$$\ker(\phi_2) = \bigoplus_{i=1}^{d} S^{(2d-2i,2i)}(V).$$

Show that each irreducible summand of $\ker(\phi_n)$ is contained in $\ker(\phi_2)\otimes \mathrm{Sym}^{n-2}(\mathrm{Sym}^d(V))$. Use this to prove Proposition 3.2.5.

One of the main applications of the Veronese embedding is to "linearize" the space of homogeneous degree d polynomials on $\mathbb{P}^n$. For example, the image of a degree d hypersurface under the Veronese map is simply a hyperplane section of the Veronese variety. In particular, Exercise 3.2.12 uses "linearizing" to show that any projective scheme is isomorphic to a subscheme of some projective space defined by linear and quadratic equations. (This result is interesting but not particularly useful.)

3.2.3 Exercises

Exercise 3.2.8. Prove that any $(n+1)$ distinct $\mathbb{K}$-points of the rational normal curve in $\mathbb{P}^n$ are linearly independent.

Exercise 3.2.9. Show that the rational normal curve in $\mathbb{P}^n$ is isomorphic to the intersection of the image of the 2nd Veronese embedding of $\mathbb{P}^n \to \mathbb{P}^{\binom{n+2}{2}-1}$ with a linear subspace.

Exercise 3.2.10. Let $\mathbb{K}$ be an algebraically closed field of characteristic 0.

As we saw when discussing rational normal curves, we can identify the coordinates $y_0, \ldots, y_d$ on $\mathbb{P}^d$ with a basis for the dth symmetric powers of the coordinates x_0, x_1 on $\mathbb{P}^1$. Under this identification $\mathbb{P}^d$ obtains an action by the group $\mathrm{PGL}_2(\mathbb{K})$.

(1) Show that the action of PGL_2 on $\mathbb{P}^2$ has two orbits: the conic $C = V_+(xz - y^2)$ and $\mathbb{P}^2\backslash C$.

(2) Show that the action of PGL_2 on $\mathbb{P}^3$ has three orbits: a twisted cubic C, the complement $Y\backslash C$ in a degree 4 hypersurface Y containing C, and $\mathbb{P}^3\backslash Y$.

Exercise 3.2.11. Consider the 2-Veronese map $\mathbb{P}^2 \to \mathbb{P}^5$ defined by $f^\dagger : \mathbb{K}[y_0, \ldots, y_5] \to \mathbb{K}[x_0, x_1, x_2]$ sending

$$f^\dagger(y_0) = x_0^2 \qquad f^\dagger(y_1) = x_0x_1 \qquad f^\dagger(y_2) = x_0x_2$$
$$f^\dagger(y_3) = x_1^2 \qquad f^\dagger(y_4) = x_1x_2 \qquad f^\dagger(y_5) = x_2^2$$

Consider the rational map $\phi : \mathbb{P}^5 \dashrightarrow \mathbb{P}^4$ away from the point defined by the ideal $(y_0, y_3, y_5, y_1 - y_4, y_2 - y_4)$. Prove that the composition $\phi \circ f$ defines a morphism from $\mathbb{P}^2$ to $\mathbb{P}^4$ which is an isomorphism onto its image.

(A famous theorem of Steiner shows that the 2-Veronese surface in $\mathbb{P}^5$ is the only smooth surface S in $\mathbb{P}^5$ such that projection away from a point defines an isomorphism from S to its image in $\mathbb{P}^4$.)

Exercise 3.2.12. Let $I \subset \mathbb{K}[x_0, \dots, x_n]$ be a homogeneous ideal defining a projective scheme X. Show that for some sufficiently large d the ideal $J = \langle I_d \rangle$ will have the same saturation as I.

Prove that the image of $V_+(J)$ under the dth Veronese map $f : \mathbb{P}^n \to \mathbb{P}^{\binom{n+d}{d}-1}$ will be the intersection of the image of f with a linear subspace. Conclude that every projective scheme X admits an embedding into some projective space such that the homogeneous ideal of X is generated by linear and quadratic equations.

3.3 Segre varieties

Suppose that V, W are vector spaces and consider the map

$$V \oplus W \to V \otimes W$$
$$(v, w) \mapsto v \otimes w$$

The image will be the elements of $V \otimes W$ which are pure tensors. Again, it is more convenient to pass to the projectivized version $\mathbb{P}(V) \times \mathbb{P}(W) \to \mathbb{P}(V \otimes W)$. If we introduce coordinates, the map is defined as

$$(a_0 : \ldots : a_n) \times (b_0 : \ldots : b_m) \mapsto (a_0 b_0 : \ldots : a_i b_j : \ldots : a_n b_m)$$

We will show that this map defines a morphism $f : \mathbb{P}^n \times \mathbb{P}^m \to \mathbb{P}^{nm+n+m}$ giving the product the structure of a projective variety.

Just as with the Veronese map, it will be useful to index the coordinates of $\mathbb{P}^{nm+n+m}$ a non-traditional way. We define the ring $\mathbb{K}[z_{ij}]$ with coordinates indexed by $0 \leq i \leq n$ and $0 \leq j \leq m$. If we picture these coordinates in a $(n+1 \times m+1)$-matrix, then the image of the $\mathbb{K}$-points under the map f will consist precisely of those (non-zero) matrices which have rank 1. In other words, we should expect X to be the vanishing locus of the ideal I generated by the 2×2 minors:

$$I = \{z_{ij} z_{kl} - z_{il} z_{kj} \,|\, 0 \leq i, k \leq n \text{ and } 0 \leq j, l \leq m\}.$$

Definition 3.3.1. The Segre variety $\Sigma_{n,m}$ is defined to be $\mathrm{mProj}(\mathbb{K}[z_{ij}]/I)$.

We next equip $\Sigma_{n,m}$ with morphisms to the projective spaces $\mathbb{P}^n$ and $\mathbb{P}^m$. The rational map $\Sigma_{n,m} \to \mathbb{P}^n$ is given by fixing an index l and defining the graded homomorphism $\mathbb{K}[x_i] \to \mathbb{K}[z_{ij}]/I$ that sends $x_i \mapsto z_{il}$. As we vary l to get different rational maps, we see that these maps agree on the overlaps by using the defining equations of I. Together these rational maps define a morphism we call $p_1 : \Sigma_{n,m} \to \mathbb{P}^n$. Analogously we define a morphism $p_2 : \Sigma_{n,m} \to \mathbb{P}^m$ by fixing k and sending $y_j \mapsto z_{kj}$.

Proposition 3.3.2. *$\Sigma_{n,m}$ (equipped with p_1, p_2) is isomorphic to $\mathbb{P}^n \times \mathbb{P}^m$ (equipped with the projection maps).*

Proof. Let $R = \mathbb{K}[x_0, \ldots, x_n]$ and $S = \mathbb{K}[y_0, \ldots, y_m]$. Recall that $\mathbb{P}^n \times \mathbb{P}^m$ is the mProj of the graded ring

$$Q = \bigoplus_{d \geq 0} (R_d \otimes_{\mathbb{K}} S_d).$$

Consider the map of graded rings $f^\dagger : \mathbb{K}[z_{ij}] \to Q$ sending $z_{ij} \mapsto x_i \otimes y_j$. We claim that this induces an isomorphism $f : \mathbb{P}^n \times \mathbb{P}^m \to \Sigma_{n,m}$ that is compatible with the projection maps.

We will verify this on affine charts. Fix indices k, l and consider the chart $D_{+,z_{kl}}$ on $\Sigma_{n,m}$ with ring of functions $\mathbb{K}[z_{ij}/z_{kl}]/I$. Using the defining equations for I we see that for any $i \neq k, j \neq l$ we can write

$$\frac{z_{ij}}{z_{kl}} = \frac{z_{il}}{z_{kl}} \cdot \frac{z_{kj}}{z_{kl}}.$$

Thus the quotient ring is isomorphic to $\mathbb{K}[z_{kj}/z_{kl}] \otimes \mathbb{K}[z_{il}/z_{kl}]$.

The preimage of this affine chart in $\mathbb{P}^n \times \mathbb{P}^m$ is defined by localizing the element $x_i \otimes y_j$ in Q. We have already shown in Construction 2.9.7 that this open affine has ring of functions $(R_{x_i})_0 \otimes (S_{y_j})_0$. It is then clear that the two affine charts are both isomorphic to $\mathbb{A}^n \times \mathbb{A}^m$ and that the product structure on each is preserved by f. Furthermore this identification is compatible when restricted to the intersection of two charts. We conclude that f is an isomorphism. $\square$

Example 3.3.3. The Segre variety $\Sigma_{1,1} \cong \mathbb{P}^1 \times \mathbb{P}^1$ is a subvariety of $\mathbb{P}^3$. Switching from the coordinates $\{z_{ij}\}_{i,j=0}^1$ to the more traditional coordinates $\{x_k\}_{k=0}^3$ we see that it is defined by the equation $x_0x_3 - x_1x_2 = 0$. In other words, it is a quadric hypersurface. Exercise 3.1.15 explicitly identifies the two families of $\mathbb{P}^1$s inside of this quadric corresponding to the fibers of the projection maps.

If $\mathbb{K}$ is an algebraically closed field (with $\mathrm{char}(\mathbb{K}) \neq 2$) then Example 3.1.4 implies that every full rank quadric in $\mathbb{P}^3$ is isomorphic to $\mathbb{P}^1 \times \mathbb{P}^1$.

Example 3.3.4. The Segre variety $\Sigma_{2,1} \cong \mathbb{P}^2 \times \mathbb{P}^1$ is the subvariety of $\mathbb{P}^5$ defined by the equations

$$y_0y_3 - y_1y_2 = y_0y_5 - y_1y_4 = y_2y_5 - y_3y_4 = 0.$$

If we choose any pair of these equations then the vanishing locus will be the union of $\Sigma_{2,1}$ with a three-dimensional plane. For example, the first two equations define the union of $\Sigma_{2,1}$ with the plane $y_0 = y_1 = 0$. (Verify this carefully!)

It is also true that the ideal I is the saturated ideal defining $\Sigma_{n,m}$. You can show this using either of the following exercises.

Exercise 3.3.5. Consider the map $f^\dagger : \mathbb{K}[z_{ij}] \to Q$ sending $z_{ij} \mapsto x_iy_j$ as above. Prove that I is the kernel of this map as follows. Suppose that $f \in \mathbb{K}[z_{ij}]$ becomes identically zero upon setting $z_{ij} = x_iy_j$. By looking at some "highest degree" terms explain how you can repeatedly subtract the product of a monomial and a generator of I to eventually bring f down to 0.

Exercise 3.3.6. Let V and W denote the vector space of homogeneous linear functions in the x_i and y_j respectively. Then the homogeneous linear functions in the z_{ij} can be identified as $V \otimes W$. Thus, we can identify

$$\mathbb{K}[z_{ij}]_d = \mathrm{Sym}^d(V \otimes W) \cong \bigoplus_{\lambda \vdash d} S^\lambda(V) \otimes S^\lambda(W)$$

as λ varies over all partitions of d and S^λ denotes the corresponding Schur functor. Similarly, we can identify

$$Q_d = \mathrm{Sym}^d(V) \otimes \mathrm{Sym}^d(W).$$

The map on degree d components $f_d^\dagger : \mathbb{K}[z_{ij}]_d \to Q_d$ is just the projection map. For example, the kernel of $f_2^\dagger$ corresponds to the summand $\bigwedge^2(V) \otimes \bigwedge^2(W)$. Prove that for any $d > 2$ the kernel of $f_d^\dagger$ is a subrepresentation of $(\bigwedge^2(V) \otimes \mathrm{Sym}(V)^{\otimes d-2}) \otimes (\bigwedge^2(W) \otimes \mathrm{Sym}(W)^{\otimes d-2})$ and deduce that it lies in the ideal I.

3.3.1 Subschemes of products of projective space

We have now seen two ways of thinking about $\mathbb{P}^n \times \mathbb{P}^m$, but the most convenient way to work with this variety requires a third approach. We write $\{x_i\}_{i=0}^n$ for the projective coordinates on $\mathbb{P}^n$ and $\{y_j\}_{j=0}^m$ for the projective coordinates on $\mathbb{P}^m$.

Definition 3.3.7. We say that a polynomial f in the variables $\{x_i, y_j\}$ is bihomogeneous of degree (d, e) if it is homogeneous of degree e in the variables x_i and homogeneous of degree e in the variables y_j.

For example, the polynomial $x_0y_0^2 + x_1y_1^2 + x_2y_2^2$ is bihomogeneous of degree $(1, 2)$. When we defined $\mathbb{P}^n \times \mathbb{P}^m$ we only used polynomials which were bihomogeneous of equal degree (d, d). (In other words, under the Segre embedding a degree d equation on the ambient projective space will become a bihomogeneous equation of bidegree (d, d).) However, we can equally well make sense of closed subschemes defined by the vanishing locus of bihomogeneous polynomials of unequal degree. Indeed, a bihomogeneous equation of bidegree (d, e) with $d > e$ can be identified with the system of equations of bidegree (d, d) obtained by multiplying against every polynomial in the $\{y_j\}$ of degree $d - e$. It is often more convenient to work directly with bihomogeneous coordinates due to their increased flexibility.

Example 3.3.8. The twisted cubic X in $\mathbb{P}^3$ that is the image of the Veronese map lies on the Segre surface $\Sigma_{1,1}$ because the equation $wz - xy$ is one of the generators of its ideal. Let's see how to write X in bihomogeneous equations using coordinates s, t and u, v on the two $\mathbb{P}^1$ factors and identifying $w = su, x = sv, y = tu, z = tv$.

The equation $x^2 - wy$ is also in the ideal of X. When we restrict to Σ, it becomes the bidegree $(2, 2)$ polynomial $s^2v^2 - stu^2$. This factors into s and $sv^2 - tu^2$, reflecting the fact that $V_+(x^2 - wy) \cap \Sigma_{1,1}$ is the union of a line and X. Similarly, the other equation $y^2 - xz$ factors into t and $tu^2 - sv^2$. Altogether this shows that X is defined by the single bihomogeneous equation $tu^2 - sv^2$ in $\Sigma_{1,1}$.

(It is also defined by the two equations $s^2v^2 - stu^2$, $t^2u^2 - stv^2$ pulled back from $\mathbb{P}^3$. As discussed above, these equations have equal degree in each coordinate system. However it is much easier to work with a single equation.)

3.3.2 Exercises

Exercise 3.3.9. Consider the Segre variety $\Sigma_{2,1} \subset \mathbb{P}^5$.

(1) We know that $\Sigma_{2,1} \cong \mathbb{P}^2 \times \mathbb{P}^1$. Write down the equations describing the subvarieties which are the fibers over a $\mathbb{K}$-point for one of the projection maps. (Since there can be infinitely many fibers, the equations will depend upon some parameters; see Exercise 3.1.15 for the analogous construction for $\Sigma_{1,1}$.)

(2) Suppose that ℓ is a line in $\mathbb{P}^5$ that is contained in the Segre variety $\Sigma_{2,1}$. Show that ℓ must be contained in a fiber of one of the projection maps.

Exercise 3.3.10. Identify a linear subspace $L \subset \mathbb{P}^{n^2+2n}$ of dimension $\binom{n+2}{2} - 1$ such that the image of the diagonal $\Delta \subset \mathbb{P}^n \times \mathbb{P}^n$ under the Segre embedding in $\mathbb{P}^{n^2+2n}$ is isomorphic to the 2-Veronese variety of $\mathbb{P}^n$ lying inside of L.

Exercise 3.3.11. Fix a non-negative integer r. Consider the subvariety $S_r \subset \mathbb{P}^2 \times \mathbb{P}^1$ defined by the bihomogeneous equation

$$s^r x + t^r y = 0$$

where s, t are coordinates on $\mathbb{P}^1$ and x, y, z are coordinates on $\mathbb{P}^2$. S_r is known as the rth Hirzebruch surface.

(1) Show that the fiber of the projection map $p : S_r \to \mathbb{P}^1$ over a point of residue field $\mathbb{L}$ is isomorphic to $\mathbb{P}^1_{\mathbb{L}}$.

(2) Show that there is a section $\sigma : \mathbb{P}^1 \to S_r$ whose image is the point $x = y = 0$ in every fiber. Show that the image of this section is the only curve contracted by the projection map $S_r \to \mathbb{P}^2$.

(3) Show that S_r is covered by four charts $\{U_i\}_{i=1}^4$ each isomorphic to $\mathbb{A}^2$. If we write $\psi_i : \mathbb{A}^2 \to U_i$ for this isomorphism, identify the "overlap maps" $\psi_{ij} : \psi_i^{-1}(U_i \cap U_j) \to \psi_j^{-1}(U_i \cap U_j)$.

3.4 Blow-ups

The blow-up is one of the fundamental constructions in algebraic geometry. It is a basic tool for understanding birational maps and birational equivalence.

Let $X = \mathrm{mSpec}(R)$ be an affine scheme. Fix an ideal $I = (g_0, \ldots, g_r)$ in R. (Note that we do not insist that the g_i are a minimal set of generators of I.) We can define a rational map $\phi : X \dashrightarrow \mathbb{P}^r$ by sending $x \mapsto (g_0(x) : \ldots : g_r(x))$. More precisely, if use coordinates $y_0, \ldots, y_r$ on $\mathbb{P}^{r-1}$, then we can define the map on charts $D_{g_i} \to D_{+,y_i}$ by sending $\frac{y_j}{y_I} \mapsto \frac{g_j}{g_i}$, and one can verify that these maps glue to yield a morphism on the complement U of $V(I)$. In particular, we can define the graph $\Gamma_\phi \subset U \times \mathbb{P}^r$.

Definition 3.4.1. The blow-up $\mathrm{Bl}_I(X)$ of X along I is the closure of Γ_ϕ inside of $X \times \mathbb{P}^r$ (or more accurately, the scheme-theoretic image of the inclusion $\Gamma_\phi \to X \times \mathbb{P}^r$).

For convenience we will focus on the case when X is an affine variety. In this case $\mathrm{Bl}_I(X)$ will also be a variety since the closure of an irreducible subset is irreducible and the scheme-theoretic image of a reduced scheme is reduced.

Note that the two projection maps induce a diagram

$$\begin{tikzcd} & \mathrm{Bl}_I(X) \arrow[dl, "\pi_1"'] \arrow[dr, "\pi_2"] & \\ X \arrow[rr, dashed, "\phi"] & & \mathbb{P}^r \end{tikzcd}$$

such that $\pi_2 = \phi \circ \pi_1$ as rational maps. By construction, the map $\mathrm{Bl}_I(X) \to X$ will be an isomorphism over U, and in particular will be birational. The locus in $\mathrm{Bl}_I(X)$ where π_1 is not an isomorphism is known as the "exceptional locus" and is often denoted by E.

The key perspective is that the blow-up turns the rational map ϕ into a morphism at the cost of replacing X by a birationally equivalent variety. (This construction is more general than it might appear – any rational map can be written in the form ϕ above.) There are other ways of thinking about blow-ups which we will see later on.

It is not difficult to identify the blow-up using explicit equations:

Theorem 3.4.2. *The blow-up of $X = \mathrm{mSpec}(R)$ along $I = (g_0, \ldots, g_r)$ is the closed subscheme of $X \times \mathbb{P}^r$ defined by the homogenous ideal in $R[y_0, \ldots, y_r]$ defined by the equations*

$$g_i y_j = g_j y_i$$

for $i, j \in \{0, \ldots, k\}$ with $i \neq j$.

Here we are implicitly using the identification $X \times \mathbb{P}^r \cong \mathrm{mProj}(R[y_0, \ldots, y_r])$ of Exercise 2.9.11.

Proof. As discussed earlier, the rational map ϕ restricts to a morphism $\phi_i : D_{g_i} \to D_{+,y_i}$ for any i. The graph of ϕ_i is the affine subscheme of $\mathrm{mSpec}(R_{g_i}[\frac{y_0}{y_i}, \ldots, \frac{y_r}{y_i}])$ defined by the equations

$$\frac{y_j}{y_i} = \frac{g_j}{g_i}$$

as we vary j. It is clear that these affine charts also describe the vanishing locus of the homogeneous ideal in the statement. $\square$

3.4.1 Examples

Example 3.4.3. Suppose that $I = (g)$ is a principal ideal. Since $\mathbb{P}^0 \cong \mathrm{mSpec}(\mathbb{K})$ the map $\phi : X \dashrightarrow \mathbb{P}^0$ is just the structure map and $X \times \mathbb{P}^0 \cong X$. By definition, $\mathrm{Bl}_I(X)$ will be the closure of the complement of $V(g)$ in X. If g is not a zero-divisor then $\mathrm{Bl}_I(X) \cong X$, but if g is a zero divisor then $\mathrm{Bl}_I(X)$ will be the closure of the complement of the components of X where g vanishes.

Example 3.4.4. In this example we blow-up the origin in $\mathbb{A}^2$. We will write x_0, x_1 for the coordinates on $\mathbb{A}^2$. (Note the unusual indexing!) The ideal for the origin is $I = (x_0, x_1)$, and the map $\phi : \mathbb{A}^2 \dashrightarrow \mathbb{P}^1$ will be projection away from the origin discussed in Example 2.6.6. We let U denote the complement of the origin where ϕ is defined.

The blow-up is defined by the single equation $x_0 y_1 = y_0 x_1$. Over the complement of the origin, the map $\mathrm{Bl}_I(\mathbb{A}^2) \to \mathbb{A}^2$ will be an isomorphism. On the other hand, the preimage of the origin will be the entire $\mathbb{P}^1$ fiber. We can check these claims very explicitly using the two affine charts $\mathbb{A}^2 \times D_{+,y_i}$ of $\mathbb{A}^2 \times \mathbb{P}^1$. On the chart where y_0 does not vanish, $\mathrm{Bl}_I(\mathbb{A}^2)$ is defined by the equation

$$x_0 \cdot \frac{y_1}{y_0} = x_1 \qquad \text{in } \mathbb{K}\left[x_0, x_1, \frac{y_1}{y_0}\right].$$

Abstractly this chart is isomorphic to $\mathbb{A}^2$ (in the coordinates x_0 and y_1/y_0). Under this identification with $\mathbb{A}^2$, the projection map from this chart of X to the D_{+,y_0} factor is just a coordinate projection $\mathbb{A}^2 \to \mathbb{A}^1$. The projection map from this chart of X to $\mathbb{A}^2$ is the birational map discussed in Example 1.6.3. As discussed there, the map will contract the locus $V(x_0)$ down to a point and will be an isomorphism on the complement of this vanishing locus. In particular, we see that on this chart the exceptional locus E is the vanishing locus of x_0. The situation for the other chart is completely symmetric.

Let's discuss the geometry of this construction. The fibers of ϕ are the lines through the origin in $\mathbb{A}^2$ intersected with U. When we blow-up, we "separate" these lines by replacing the origin by the $\mathbb{P}^1$ representing tangent directions at the origin. The resulting surface will now admit a map to $\mathbb{P}^1$ with fibers isomorphic to $\mathbb{A}^1$. In summary:

- the fiber of $\mathrm{Bl}_I(\mathbb{A}^2) \to \mathbb{P}^1$ over a $\mathbb{K}$-point q is the line ℓ corresponding to q (as a subvariety of $\mathbb{A}^2 \times \{q\}$),

- the fiber of $\mathrm{Bl}_I(\mathbb{A}^2) \to \mathbb{A}^2$ over 0 is $\mathbb{P}^1$ (representing the family of lines through 0), and the map $\mathrm{Bl}_I(\mathbb{A}^2) \to \mathbb{A}^2$ an isomorphism over the complement of 0.

Remark 3.4.5. One should think of the blow-up of the origin in $\mathbb{A}^2$ as "separating the tangent directions" at the origin. We will revisit this idea in Section 5.5.

Example 3.4.6. Let X denote the blow-up of the line L defined by $x_0 = x_1 = 0$ inside of $\mathbb{A}^3$. If Y denotes the blow-up of $\mathbb{A}^2$ at the origin, then it is easy to see that $X \cong Y \times \mathbb{A}^1$ (where the $\mathbb{A}^1$ factor corresponds to the x_2-coordinate). Here the exceptional locus E is isomorphic to $\mathbb{P}^1 \times L$ and the birational map $X \to \mathbb{A}^3$ restricts to the second projection map on E.

In this case, one should think of the blow-up as "separating the directions in the normal bundle of L". For example, suppose that ℓ_1, ℓ_2 are two lines in $\mathbb{A}^3$ that intersect at a point in L. Then the strict transforms of ℓ_1 and ℓ_2 will intersect in X precisely when their tangent directions map to the same quotient in the normal bundle of L at the point.

3.4.2 Compatibility of blow-ups

The blow-up satisfies several important compatibility properties which makes it easier to think about and work with. The most fundamental property is that the blow-up $\mathrm{Bl}_I(X)$ does not depend upon the choice of generators of I. (Note that this property is implicit in our notation.)

Proposition 3.4.7. *Let $X = \mathrm{mSpec}(R)$ be an affine scheme and let I be an ideal. Suppose we choose different sets of generators $(g_0, \dots, g_r)$ and $(g'_0, \dots, g'_s)$ for I. Let Y and Y' be the blow-up with respect to these two sets of generators equipped with the projection maps $p : Y \to X$ and $p' : Y' \to X$. Then there is an isomorphism $\psi : Y \cong Y'$ such that $p = p' \circ \psi$.*

Proof. Write $g'_k = \sum_{i=0}^r h_{i,k} g_i$. Let $\{y_i\}_{i=0}^r$ be coordinates on $\mathbb{P}^r$ and let $\{z_k\}_{j=0}^s$ be coordinates on $\mathbb{P}^s$. Then we define a rational map $X \times \mathbb{P}^r \dashrightarrow X \times \mathbb{P}^s$ using the graded homomorphism $f^\dagger : R[z_k] \mapsto R[y_i]$ sending $z_k \mapsto \sum_{i=0}^r h_{i,k} y_i$.

We first claim that this rational map restricts to a morphism on Y. Indeed, the rational map $X \times \mathbb{P}^r \dashrightarrow X \times \mathbb{P}^s$ is defined away from the locus defined by the equations $\{\sum_{i=0}^r h_{i,k} y_i\}_{k=0}^s$. On the other hand Y is defined by the equations $g_i y_j - g_j y_i$. The intersection is given by taking the sum J of the ideals. The relation

$$g_j \sum_{i=0}^r h_{i,k} y_i - (g_i y_j - g_j y_i) \sum_{i=0}^r h_{i,k} = \sum_{i=0}^r h_{i,k} g_i y_j = g'_k y_j$$

shows that the ideal $I(y_0, \dots, y_s)$ is contained in J. Since $V_+(I(y_0, \dots, y_s)) = \emptyset$, we also have $V_+(J) = \emptyset$ so that the rational map defines a morphism on Y.

We next need to check that the image of Y is Y'. Indeed, in the quotient ring of $R[y_i]$ by the homogeneous ideal of Y we have the relation

$$\begin{aligned} g'_k \sum_{i=0}^{r} h_{i,l} y_i - g'_l \sum_{i=0}^{r} h_{i,k} y_i &= \sum_{j=0}^{r}\sum_{i=0}^{r} h_{j,k} h_{i,l} g_j y_i - \sum_{j=0}^{r}\sum_{i=0}^{r} h_{j,l} h_{i,k} g_j y_i \\ &= 0 \end{aligned}$$

since $g_i y_j = g_j y_i$ in this quotient ring. This shows that the homogeneous ideal for Y' is in the kernel of $f^\dagger$.

Finally, we can construct a map $Y' \to Y$ in the opposite way using the analogous construction from a relation $g_i = \sum_{k=0}^{s} h'_{k,i} g'_k$. It is clear that the composition of the two maps on the homogeneous coordinate rings of Y and Y' is the identity, showing that the induced maps are isomorphisms. The compatibility of the projection maps to $X = \mathrm{mSpec}(R)$ follows from the fact that the maps $R[z_k] \to R[y_i]$ and $R[y_i] \to R[z_k]$ are isomorphisms on the 0-graded piece R. $\square$

The blow-up operation also satisfies two "geometric" compatibilities. First, blow-ups are compatible with taking closed sets. To make sense of this, we need the following definition:

Definition 3.4.8. Let $X = \mathrm{mSpec}(R)$ be an affine scheme and let $I = (g_0, \ldots, g_r)$ be an ideal in R. Let $Z = V(J)$ be a closed subscheme of X. Suppose that no associated prime for J contains I. The strict transform of Z in $\mathrm{Bl}_I(X)$ is the closure of $Z \cap (X \backslash V(I))$, or more precisely, the scheme-theoretic image of $Z \cap (X \backslash V(I))$ in $\mathrm{Bl}_I(X)$ under the inclusion map.

Exercise 3.4.9. Let $X = \mathrm{mSpec}(R)$ be an affine scheme and let $I = (g_0, \ldots, g_r)$ be an ideal in R. Consider a closed subscheme $Z = V(J)$ in X. Then the blow-up of Z along the quotient ideal $\overline{I}$ in R/J (with the same generators) is the same as the strict transform of Z in $\mathrm{Bl}_I(X)$.

This compatibility is quite useful for computations. For example, if we would like to blow-up a plane curve along an ideal I, we can instead blow-up $\mathbb{A}^2$ along the corresponding ideal and take the strict transform of the curve (see Exercise 3.4.11).

The second geometric compatibility is with taking open sets.

Exercise 3.4.10. Let $X = \mathrm{mSpec}(R)$ be an affine scheme and let $I = (g_0, \ldots, g_r)$ be an ideal in R. Consider a distinguished open affine D_f in X. Then the blow-up of D_f along the localized ideal I_f (with the same generators) is the same as

$$\pi_1^{-1}(D_f) \cap \mathrm{Bl}_I(X) \subset X \times \mathbb{P}^r.$$

This exercise indicates that one can blow-up a quasiprojective scheme along a sheaf of ideals by doing the blow-up on affine charts and gluing (see Exercise 3.4.13).

3.4.3 Exercises

Exercise 3.4.11. Compute the blow-ups of the following curves in $\mathbb{A}^2$ along the ideal (x, y). What is the preimage of the origin? If E denotes the exceptional locus for $\mathrm{Bl}_0(\mathbb{A}^2)$, how does the strict transform of the curve in $\mathrm{Bl}_0(\mathbb{A}^2)$ intersect E?

(1) The curve $x = 0$.

(2) The curve $xy = 0$.

(3) The curve $y^2 = x^3 + x^2$.

(4) The curve $y^2 = x^3$.

(5) The curve $y^2 = x^4$.

Exercise 3.4.12. Let $\mathbb{A}^n$ have coordinates $x_0, \ldots, x_{n-1}$ and let $I = (x_0, \ldots, x_k)$ for some $k \leq n-1$. Describe the fibers of the maps $\mathrm{Bl}_I(\mathbb{A}^n) \to \mathbb{A}^n$ and $\mathrm{Bl}_i(\mathbb{A}^n) \to \mathbb{P}^k$.

Exercise 3.4.13. Let X be a projective scheme and let $Y \subset X$ be a closed subscheme. For every open affine $U \subset X$, the intersection $U \cap Y$ is a closed subscheme of U which identifies an ideal $I_U \subset \mathcal{O}_X(U)$. Show that as we vary U the various blow-ups $\mathrm{Bl}_I(U)$ can be glued together to define a projective variety $\mathrm{Bl}_Y(X)$.

For example, show that the blow-up of $\mathbb{P}^n$ along a linear space L is the closure of the graph of the projection $\mathbb{P}^n \dashrightarrow \mathbb{P}^k$ away from L. If L is the coordinate plane $x_0 = \ldots = x_k = 0$, show that the blow-up is defined in $\mathbb{P}^n \times \mathbb{P}^k$ by the bihomogeneous equations $x_i y_j = x_j y_i$ for $0 \leq i, j \leq k$.

However, in contrast to the blow-up of $\mathbb{A}^n$ along a linear space the exceptional locus E is no longer a product of two varieties!

Exercise 3.4.14. Let $X, Y \subset \mathbb{A}^n$ be closed subvarieties defined by ideals I, J. Recall that the intersection is defined by $I + J$. Assume that $V(I + J)$ is strictly contained in X and strictly contained in Y. Prove that the strict transforms of X and Y in the blow-up of $I + J$ are disjoint.

Exercise 3.4.15. Let X be an affine variety and consider the blow-up $Bl_I(X)$ along an ideal I. Let Z denote the preimage of $V(I)$. Prove that $Bl_I(X)$ is covered by open affine subsets U such that $Z \cap U$ is the vanishing locus of a single equation. (One sometimes says that the blow-up "principalizes" the ideal I.)

3.5 Grassmannians: projective structure

The Grassmannian $G(k,n)$ is the parameter space for k-dimensional subspaces of an n-dimensional vector space. Similarly, $\mathbb{G}(k,n)$ denotes the parameter space for k-dimensional planes in $\mathbb{P}^n$. Note that $\mathbb{G}(k,n) = G(k+1,n+1)$. In this section we give the Grassmannian $G(k,n)$ the structure of a projective variety whose $\mathbb{K}$-points represent the corresponding k-dimensional subspaces.

Example 3.5.1. We have $G(1,n) = \mathbb{P}^{n-1}$. The Grassmannian $G(n-1,n)$ can be identified with the dual projective space $(\mathbb{P}^{n-1})^\vee$.

We will construct the Grassmannian as a closed subscheme of $\mathbb{P}^{\binom{n}{k}-1}$. The embedding $G(k,n) \hookrightarrow \mathbb{P}^{\binom{n}{k}-1}$ is known as the Plücker embedding and can be described set theoretically as follows. Suppose $W \in G(k,n)$ is a k-dimensional subspace. Fix any basis $v_1, \ldots, v_k$ of W and consider the map

$$\begin{aligned} \phi : G(k,n) &\to \mathbb{P}\left(\bigwedge^k \mathbb{K}^n\right) \\ W &\mapsto v_1 \wedge \ldots \wedge v_k \end{aligned}$$

Note that $\phi(W)$ does not depend upon the choice of basis. The $\mathbb{K}$-points in the image of the Plücker embedding will be the $\mathbb{K}$-points in $\mathbb{P}\left(\bigwedge^k \mathbb{K}^n\right)$ which are the projectivizations of decomposable vectors (i.e. which can be written as a pure wedge product).

3.5.1 Defining a closed subset

To identify the ideal that defines the Plücker embedding, we need to think more carefully about what it means for an element of the exterior algebra to be a pure wedge product. For clarity we will write V for the vector space $\mathbb{K}^n$. Fix an element $\omega \in \bigwedge^k \mathbb{K}^n$ and consider the linear map

$$\begin{aligned} \varphi_\omega : V &\to \bigwedge^{k+1} V \\ v &\mapsto v \wedge \omega \end{aligned}$$

Then ω will be a pure wedge product if and only if the rank of this map is $\leq n-k$. This is a consequence of Lemma 3.0.2 and Proposition 3.0.6.

The map $\mathbb{P}(\bigwedge^k V) \to \mathbb{P}(\mathrm{Hom}(V, \bigwedge^{k+1} V))$ which sends ω to the map φ_ω is a linear function. By Exercise 1.2.16 the subset of $\mathbb{P}(\mathrm{Hom}(V, \bigwedge^{k+1} V))$ which corresponds to maps of rank $\leq n-k$ is a closed subset defined by the vanishing of the $(n-k+1)\times(n-k+1)$ minors of a matrix of linear functions in the homogeneous coordinate ring. By pulling these functions back under the linear map defined above, we obtain an ideal whose vanishing locus is the Grassmannian. This shows how to define the Grassmannian as a projective scheme.

3.5.2 Plücker relations

Unfortunately, the ideal we have constructed is very far from the ideal I of *all* homogeneous functions which vanish on the Grassmannian. The saturated ideal I defining the Grassmannian is generated by certain quadratic functions known as Plücker relations. The construction is similar, but a little more complicated.

Note that we have two perfect pairings

$$\bigwedge^{n-k} V^\vee \times \bigwedge^{n-k} V \to \mathbb{K}$$
$$\bigwedge^{k} V \times \bigwedge^{n-k} V \to \bigwedge^{n} V$$

Choosing an isomorphism $\bigwedge^n V \to \mathbb{K}$, we can identify $\bigwedge^k V$ and $\bigwedge^{n-k} V^\vee$. This isomorphism is only natural up to scaling, but when we projectivize this ambiguity won't matter. Suppose we choose a basis $v_1, \ldots, v_r$ of V with dual basis $v_1^\vee, \ldots, v_r^\vee$. For any subset $I \subset \{1, \ldots, r\}$ of size k we have

$$v_{i_1} \wedge \ldots \wedge v_{i_k} \leftrightarrow \pm v_{j_1}^\vee \wedge \ldots \wedge v_{j_{n-k}}^\vee$$

where J is the complement of I.

Fix any $\omega \in \bigwedge^k V$ and let ω^* denote the corresponding element in $\bigwedge^{n-k} V^\vee$. Consider the composition

$$\psi : \bigwedge^{n-k+1} V \xrightarrow{(\wedge\omega^*)^\vee} V \xrightarrow{\wedge\omega} \bigwedge^{k+1} V$$

We show that ψ is the zero map if and only if ω is a pure wedge power. By Lemma 3.0.5 the image of the first map is the smallest subspace W such that ω is in the image of $\bigwedge^k W \to \bigwedge^k V$. If $\omega = v_1 \wedge \ldots \wedge v_k$ is a pure wedge product then the image W of the first map is the span of the v_i and the vanishing of the map follows. If ω is not a pure wedge product then there is some element in W which does not have vanishing wedge power with ω and so ψ does not vanish identically. Note that the two functions whose composition is ψ have entries which are linear in the coordinates on $\mathbb{P}(\bigwedge^k V)$; thus, the vanishing of the composed function ψ will be described by a quadratic equation in the coordinates of $\mathbb{P}(\bigwedge^k V)$.

By choosing a basis of V we can impose coordinates $\mathbb{K}[x_I]$ on $\mathbb{P}(\bigwedge^k V)$ where $I = i_1 < i_2 < \ldots < i_k$ is a size k subset of $\{1, \ldots, n\}$ written in increasing order. It is useful to allow indices which are ordered k-tuples of different elements in $\{1, 2, \ldots, n\}$ that are not necessarily increasing: by $x_{i_1 i_2 \ldots i_k}$ we will mean

$$x_{i_1 \ldots i_k} := (-1)^{\mathrm{sgn}(\sigma)} x_I$$

where I is the corresponding increasing ordered subset of $\{1, \ldots, n\}$ and σ is the permutation rearranging the i_j into increasing order. One can write the Plücker relations concretely in this basis:

Definition 3.5.2. Fix a pair of positive integers (k,n) with $k<n$. Fix an ordered subset $i_1<\ldots<i_{k-1}$ and an ordered subset $j_1<\ldots<j_{k+1}$ of $\{1,\ldots,n\}$. For each such choice, we obtain a corresponding Plücker relation

$$\sum_{\ell=1}^{k+1}(-1)^{\ell}x_{i_1,\ldots,i_{k-1},j_\ell}x_{j_1,\ldots,\widehat{j_\ell},\ldots,j_{k+1}}=0$$

Here we interpret a coordinate as 0 if it has a repeated index.

We define the Grassmanian $G(k,n)$ to be the vanishing locus of the ideal I generated by all Plücker relations for (k,n). We will show in the next section that I defines a variety. It is also true that I is the saturated ideal that defines the Grassmannian, but we will not prove this.

The $\mathbb{K}$-points of $G(k,n)$ will parametrize k-dimensional subsets of $\mathbb{K}^n$. The non-traditional points will parametrize Galois orbits of planes defined over extensions of $\mathbb{K}$; see Example 3.1.9 for a similar phenomenon in a different setting.

Example 3.5.3. The Plücker embedding realizes the Grassmannian $G(2,4)$ as a subvariety of $\mathbb{P}^5$. Using the coordinates x_I indexed by pairs of elements in $\{1,2,3,4\}$, all the Plücker relations yield a single non-trivial equation

$$x_{12}x_{34}-x_{13}x_{24}+x_{14}x_{23}=0.$$

In other words, $G(2,4)$ is a quadric hypersurface in $\mathbb{P}^5$.

Example 3.5.4. The Plücker embedding realizes the Grassmannian $G(2,5)$ as a subvariety of $\mathbb{P}^9$. Using the coordinates x_I indexed by pairs of elements in $\{1,2,3,4,5\}$, the Plücker relations are

$$\begin{aligned}
x_{12}x_{34}-x_{13}x_{24}+x_{14}x_{23}&=0\\
x_{12}x_{35}-x_{13}x_{25}+x_{15}x_{23}&=0\\
x_{12}x_{45}-x_{14}x_{25}+x_{15}x_{24}&=0\\
x_{13}x_{45}-x_{14}x_{35}+x_{15}x_{34}&=0\\
x_{23}x_{45}-x_{24}x_{35}+x_{25}x_{34}&=0
\end{aligned}$$

3.5.3 Exercises

Exercise 3.5.5. Show that the Plücker relations for $G(2,n)$ span an $\binom{n}{4}$ dimensional subspace of the vector space of homogeneous quadrics on $\mathbb{P}^{\binom{n}{2}-1}$.

Exercise 3.5.6. Suppose V has basis $e_1,\ldots,e_n$. As above we let $\mathbb{K}[x_I]$ denote the homogeneous coordinate ring on $\mathbb{P}(\bigwedge^k V)$ as I varies over all subsets of $\{1,\ldots,n\}$ of size k. Explain why the $\mathbb{K}$-points of the intersection $V_+(x_I)\cap G(k,n)$ represent the k-planes in V which fail to intersect the subspace $\mathrm{Span}\{e_i\}_{i\in\{1,2,\ldots,n\}\setminus I}$ transversally.

Exercise 3.5.7. Consider the Plücker embedding $G(k,n) \hookrightarrow \mathbb{P}(\bigwedge^{\binom{n}{k}} V)$. Fix two k-planes L_1, L_2. Show that the line between the corresponding points of $\mathbb{P}(\bigwedge^{\binom{n}{k}} V)$ is contained in $G(k,n)$ if and only if $L_1 \cap L_2$ has dimension $k-1$. In this case, the line ℓ parametrizes the family of k-planes which contain this $(k-1)$-dimensional plane.

Exercise 3.5.8. Let $\mathbb{K}$ be an algebraically closed field. Consider the Plücker embedding $\mathbb{G}(1,3) \hookrightarrow \mathbb{P}^5$.

(1) Fix a point $p \in \mathbb{P}^3$. Let $\Sigma_p \subset \mathbb{G}(1,3)$ denote the set of lines in $\mathbb{P}^3$ which contain p. Prove that the Plücker embedding sends Σ_p to a 2-plane in $\mathbb{P}^5$.

(2) Fix a hyperplane $H \subset \mathbb{P}^3$. Let $\Sigma_H \subset \mathbb{G}(1,3)$ denote the set of lines in $\mathbb{P}^3$ which are contained in H. Prove that the Plücker embedding sends Σ_H to a 2-plane in $\mathbb{P}^5$.

(3) Prove that every 2-plane in $\mathbb{P}^5$ that is contained in $\mathbb{G}(1,3)$ has one of the two forms above.

3.6 Grassmannians: chart structure

We continue our study of the Grassmannian $G(k,n)$ building upon our definition of the Plücker embedding $G(k,n) \hookrightarrow \mathbb{P}(\bigwedge^k \mathbb{K}^n)$ in the previous section.

3.6.1 Affine charts set-theoretically

Our next goal is to describe a covering of $G(k,n)$ by affine varieties. We first explain the construction set-theoretically.

Fix a basis $e_1, \dots, e_n$ of $\mathbb{K}^n$. Let $U \subset G(k,n)$ be the set of planes which meet the subspace $\mathrm{Span}(e_{k+1}, \dots, e_n)$ transversally. Equivalently, U is the set of planes W such that if we write a basis for W as the columns of an $n \times k$ matrix M then the top $k \times k$ minor will not vanish. Any such W admits a *unique* basis such that the matrix M has the form

$$M = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ 0 & 0 & \dots & 1 \\ a_{1,1} & a_{1,2} & \dots & a_{1,k} \\ a_{2,1} & a_{2,2} & \dots & a_{2,k} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n-k,1} & a_{n-k,2} & \dots & a_{n-k,k} \end{bmatrix}.$$

By assigning to W the coordinates $(a_{1,1}, \dots, a_{n-k,k})$ we obtain a bijection $U \leftrightarrow \mathbb{K}^{k(n-k)}$. As we choose different subsets of k vectors in our basis $\{e_i\}_{i=1}^n$ then the corresponding open sets form an open cover of $G(k,n)$.

Let's analyze how this description interacts with the Plücker embedding. Recall that $\mathbb{P}(\bigwedge^k \mathbb{K}^n)$ has homogeneous coordinates given by x_I as I varies over all subsets of $\{1, \dots, n\}$ of size k. The subset U is the sublocus of $G(k,n)$ where the homogeneous coordinate $x_{12\dots k}$ does not vanish. Thus each quotient $x_I/x_{12\dots k}$ yields a well-defined function on U. Up to a sign change this function is the $k \times k$ minor of the matrix M corresponding to the choice of rows defined by I. Note that each coordinate $a_{i,j}$ is defined by one of these minors: take the first k rows of M and replace the jth row by the ith row of M. Explicitly,

$$a_{i,j} = (-1)^{k-j} \frac{x_{12\dots\widehat{j}\dots k(k+i)}}{x_{12\dots k}} = \frac{x_{12\dots(k+i)\dots k}}{x_{12\dots k}}$$

where we are using our usual sign convention when the indices are not in increasing order. In particular, this means that the affine coordinates on U are (up to sign) defined by the restriction of the functions $x_I/x_{12\dots k}$ from $\mathbb{P}(\bigwedge^k \mathbb{K}^n)$ where $I \cap \{1,2,\dots,k\}$ has size $k-1$.

Remark 3.6.1. Note that by taking minors of the matrix M we obtain some equations in the homogeneous coordinate ring which vanish along $G(k,n)$. For example, since $a_{1,1}a_{2,2} - a_{1,2}a_{2,1}$ is a 2×2 minor, this quadratic function of certain affine coordinates can be identified with a different affine coordinate. Clearing denominators, this minor yields a quadratic relation $x_{(k+1)23\ldots k}x_{1(k+2)3\ldots k} - x_{(k+2)23\ldots k}x_{1(k+1)3\ldots k} = x_{(k+1)(k+2)34\ldots k}x_{12\ldots k}$.

3.6.2 Affine charts algebraically

We next repeat this computation algebraically. Consider the open affine $D_{+,x_I} \subset \mathbb{P}(\bigwedge^k \mathbb{K}^n)$.

Claim 3.6.2. The intersection $G(k,n) \cap D_{+,x_I}$ is isomorphic to $\mathbb{A}^{k(n-k)}$.

Proof. The coordinate ring of D_{+,x_I} is the set of polynomials in the fractions x_J/x_I for subsets $J \neq I$. For each subset $J \subset \{1,\ldots,n\}$, we let $s(J,I)$ denote the size of $J \cap I$.

Suppose that J is a subset of size k satisfying $s(I,J) \leq k-2$. Fix an ordered subset $i_1 < \ldots < i_{k-1}$ of $k-1$ elements of I and an ordered subset $j_1 < \ldots < j_{k+1}$ that is the union of J with the missing element of I. Corresponding to these choices we have the Plücker relation

$$\sum_{\ell=1}^{k+1}(-1)^{\ell} x_{i_1,\ldots,i_{k-1},j_\ell} x_{j_1,\ldots,\widehat{j_\ell},\ldots,j_{k+1}} = 0.$$

This equation has $x_I x_J$ as one summand; all the other summands have the form $x_K x_L$ where $s(K,I) > s(J,I)$ and $s(L,I) > s(J,I)$. After localizing at x_I, we can write x_J/x_I in terms of variables whose indices have larger intersection with I. Substituting repeatedly, we can write $\frac{x_J}{x_I}$ as a polynomial expression in variables x_K/x_I with $s(K,I) = k-1$. Precisely, these equations have the form

$$\frac{x_J}{x_I} = \sum_{\sigma: J\setminus(J\cap I) \to I\setminus(J\cap I)} (-1)^{\mathrm{sgn}(\sigma)} \left(\prod_{j \in J\setminus J\cap I} x_{\{j\}\cup(I\setminus\sigma(j))} \right)$$

as we let σ vary over all bijections between $J\setminus(J\cap I) \to I\setminus(J\cap I)$ (where $\mathrm{sgn}(\sigma)$ is an appropriately chosen sign). Although the term $x_I x_J$ can appear in many different Plücker relations, all the possibilities will yield the same expression for x_J/x_I.

The Plücker relations which do not involve x_I become identically zero after localizing and substituting in the equations above. We will not verify this carefully.

Let $\mathfrak{q} \subset \mathbb{K}[x_J/x_I]$ denote the ideal obtained by restricting the Plücker relations to the localization of x_I. Let R denote the polynomial ring in the variables x_K/x_I such that $s(K,I) = k-1$. The argument above shows that R is isomorphic to $\mathbb{K}[x_J/x_I]/\mathfrak{q}$. Since are exactly $k(n-k)$ subsets of $\{1,2,\ldots,n\}$ satisfying $s(K,I) = k-1$ we have that $\mathrm{mSpec}(R) \cong \mathbb{A}^{k(n-k)}$. □

3.6.3 The universal plane

The "universal plane" over $\mathbb{G}(k,n)$ is the following scheme.

Proposition 3.6.3. *There is a closed subvariety $\mathcal{U} \subset \mathbb{G}(k,n) \times \mathbb{P}^n$ such that the fiber of the projection map $\mathcal{U} \to \mathbb{G}(k,n)$ over a $\mathbb{K}$-point $x \in \mathbb{G}(k,n)$ is the k-plane parametrized by x.*

In fact, the fiber of $\mathcal{U} \to \mathbb{G}(k,n)$ over any point will be a subscheme which is a union of planes defined over a finite field extension.

Proof. We denote the homogeneous coordinate ring of $\mathbb{P}^n$ by $\mathbb{K}[y_0, \ldots, y_n]$. Let $\{D_{+,x_I}\}$ denote the set of affine charts on $\mathbb{G}(k,n)$ discussed in Claim 3.6.2. For each D_{+,x_I} we will construct a closed subvariety of $D_{+,x_I} \times \mathbb{P}^n$ which satisfies the desired property over D_{+,x_I}. We will then glue these subschemes to obtain $\mathcal{U}$.

Any subset $K \subset \{1, \ldots, n\}$ with $s(K,I) = k-1$ can be identified by removing an element i of I and adding an element j from the complement of I. As we vary $i \in I$ and $j \in \{1, \ldots, n\} \backslash I$, we write $b_{j,i}$ for the restriction of the function $x_{\{j\} \cup I \backslash \{i\}}/x_I$ on D_{+,x_I}. (Note that the index and sign conventions are different than for the a-variables used before.) Claim 3.6.2 shows that $D_{x_I} \cong \mathbb{A}^{k(n-k)}$ with coordinate ring $\mathbb{K}[b_{j,i}]_{i \in I, j \in \{1,\ldots,n\} \backslash I}$.

By Exercise 2.9.11 the vanishing locus of any set of polynomials in $\mathbb{K}[b_{j,i}, y_0, \ldots, y_n]$ that is homogeneous in the y variables will define a closed subset of $D_{x_I} \times \mathbb{P}^n$. For every $j \in \{1, 2, \ldots, n\} \backslash I$ consider the equation

$$f_j := y_j + \sum_{i \in I} (-1)^{\mathrm{sgn}(\sigma_{i,j})} b_{j,i} y_i$$

where $\sigma_{i,j}$ is the permutation which takes the ordered k-tuple obtained by replacing i by j and rearranges it to increasing order. The system of $(n-k)$ equations $\{f_j\}_{j \in \{1,\ldots,n\} \backslash I}$ will cut out a scheme U_I over D_{+,x_I}. It is clear that U_I is a variety – in fact, it is isomorphic to $\mathbb{A}^{k(n-k)} \times \mathbb{P}^k$. Fix a $\mathbb{K}$-point $x \in D_{+,x_I}$. It is clear that when we evaluate the coordinates $b_{j,i}$ at x the linear equations f_j define the corresponding k-plane in $\mathbb{P}^n$. Thus the fiber of the map $U_I \to D_{x_I}$ over x is indeed the plane parametrized by x.

We claim that as we vary I, the U_I glue together to give a closed subvariety of $\mathbb{G}(k,n) \times \mathbb{P}^n$. Indeed, for any subset $K \subset \{1, \ldots, n\}$ of size $k+1$ consider the equation

$$\widetilde{f}_K := \sum_{j \in K} (-1)^{\mathrm{sgn}(\sigma_j)} x_{K \backslash \{j\}} y_j.$$

where σ_j is the permutation which rearranges the ordered $(k+1)$-tuple K by moving j to the end. If I is a subset of size k such that $K \backslash I$ is a single element j, then the restriction of $\widetilde{f}_K$ to D_{+,x_I} is the equation f_j. (The $\widetilde{f}_K$ such that $I \not\subset K$ will vanish identically on D_{+,x_I}.) Thus we can use the $\widetilde{f}_K$ to define $\mathcal{U}$. $\square$

3.6.4 Exercises

Exercise 3.6.4. Let $\mathbb{K}$ be an algebraically closed field. Fix two closed subsets $X, Y \subset \mathbb{P}^n$. Let $Z = X \cap Y$. Show that the set of lines through a point in $X \backslash Z$ and a point in $Y \backslash Z$

is a quasiprojective variety $V \subset \mathbb{G}(1,n)$. The join $\mathrm{Join}(X,Y)$ is the union of all the lines parametrized by the closure $\overline{V}$. Show that $\mathrm{Join}(X,Y)$ is a closed subset of $\mathbb{P}^n$.

Exercise 3.6.5. Let $\mathbb{K}$ be an algebraically closed field. Let $X \subset \mathbb{P}^n$ be a closed subset. Show that the set of lines $V \subset \mathbb{G}(1,n)$ which connect two distinct points in X is a quasiprojective variety V. The secant variety $\mathrm{Sec}(X)$ is the union of all the lines parametrized by the closure $\overline{V}$. Show that $\mathrm{Sec}(X)$ is a closed subset of $\mathbb{P}^n$.

Exercise 3.6.6. Let $\mathbb{K}$ be an algebraically closed field. Let $C \subset \mathbb{P}^3$ be the conic defined by the equations $x_3 = x_0x_2 - x_1^2 = 0$. Find the equations of the subvariety $X \subset \mathbb{G}(1,3)$ that parametrizes the lines which meet C.

Repeat this exercise for the standard twisted cubic in $\mathbb{P}^3$.

Chapter 4

Dimension

Dimension is a surprisingly subtle concept. For example, in order to verify that a (topological) manifold has a well-defined dimension one must prove "invariance of domain": if $U \subset \mathbb{R}^m$ and $V \subset \mathbb{R}^n$ are open and there is a homeomorphism $f : U \to V$ then $n = m$. This was proved in 1912 by Brouwer using tools from algebraic topology.

In this chapter we will construct a theory of dimension for quasiprojective $\mathbb{K}$-schemes. There are several possible definitions one could use; we will show that they all coincide for $\mathbb{K}$-schemes. (For general schemes these definitions could give different values, highlighting the difficulties with this notion!)

Primer on finite ring homomorphisms

A ring homomorphism $f^\sharp : S \to R$ is said to be finite if it gives R the structure of a finitely generated S-module. (NB: this is much more restrictive than being finitely generated as an S-algebra!) A ring homomorphism $f^\sharp$ is finite if and only if it is an integral homomorphism of rings and it realizes R as a finitely generated S-algebra. In particular, if S and R are finitely generated $\mathbb{K}$-algebras then $f^\sharp$ is finite iff it is integral.

We will need to know that finiteness can be detected locally.

Proposition 4.0.1. *Let $f^\sharp : B \to A$ be a homomorphism of rings. Suppose that $\{g_j\}_{j=1}^r$ is a finite set of elements in B which generate the unit ideal. Then A is a finite B-module if and only if $A_{f^\sharp(g_j)}$ is a finite B_{g_j}-module for every j.*

Finite ring homomorphisms induce a close correspondence between prime ideals.

Theorem 4.0.2 (Lying Over)**.** *Suppose $g^\sharp : B \to A$ is a finite injective ring homomorphism. For any prime ideal $\mathfrak{p} \subset B$, there exists a prime ideal $\mathfrak{q} \subset A$ such that $(g^\sharp)^{-1}(\mathfrak{q}) = \mathfrak{p}$.*

Applying Lying Over repeatedly, we obtain:

Theorem 4.0.3 (Going Up). *Let $g^\sharp : B \to A$ be a finite injective ring homomorphism. Suppose that*

$$\mathfrak{q}_0 \subsetneq \mathfrak{q}_1 \subsetneq \cdots \subsetneq \mathfrak{q}_m$$

is a chain of primes in B, and that

$$\mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \cdots \subsetneq \mathfrak{p}_n$$

is a chain of primes in A with $n \leq m$ such that $(g^\sharp)^{-1}(\mathfrak{p}_i) = \mathfrak{q}_i$ for $i = 0, 1, \ldots, n$. Then the chain of $\mathfrak{p}_i$s can be continued via a chain

$$\mathfrak{p}_n \subsetneq \mathfrak{p}_{n+1} \subsetneq \cdots \subsetneq \mathfrak{p}_m$$

such that $(g^\sharp)^{-1}(\mathfrak{p}_i) = \mathfrak{q}_i$ for $i = 0, \ldots, m$.

When A is integrally closed, there is also a Going Down theorem which allows us to extend chains of prime ideals “downwards” instead of ”upwards”.

Theorem 4.0.4 (Going Down). *Let $g^\sharp : B \to A$ be a finite injective ring homomorphism such that B is an integrally closed domain and A is a domain. Given any inclusion of prime ideals $\mathfrak{q} \subset \mathfrak{q}'$ of B and a prime $\mathfrak{p}'$ of A such that $\mathfrak{p}' \cap B = \mathfrak{q}'$ there exists a prime $\mathfrak{p}$ of A contained in $\mathfrak{p}'$ such that $\mathfrak{p} \cap B = \mathfrak{q}$.*

4.1 Finite maps

In this section we start building a loose analogy between the following two constructions:

(1) In algebra, injective finite $\mathbb{K}$-algebra homomorphisms $f^\sharp : S \to R$.

(2) In topology, covering maps with ramification (where we allow sheets of a covering to come together along a Zariski closed subset).

Covering maps (and, if correctly defined, covering maps with ramification) are proper with finite fibers. Our algebro-geometric analogues will also satisfy these two key geometric properties.

4.1.1 Affine morphisms

We start with an auxiliary notion that is important in its own right.

Definition 4.1.1. Let $f : X \to Y$ be a morphism of quasiprojective schemes. We say that f is affine if the preimage of every open affine subset in Y is an open affine subset in X.

Note that a morphism of projective schemes will have projective fibers, and thus (by Exercise 2.11.13) will almost never be affine. On the other hand, we will soon see that morphisms of affine schemes are always affine.

Since Definition 4.1.1 involves every open affine subset of Y it is difficult to check directly. The following key lemma shows that it suffices to check what happens for a single cover of Y by open affines.

Lemma 4.1.2. *Let $f : X \to Y$ be a morphism of quasiprojective schemes. Suppose that Y admits an open cover by open affines $\{V_i\}$ such that the preimage of each V_i is an open affine in X. Then the preimage of* every *open affine V in Y will be an open affine in X.*

For example, this implies that a morphism of affine schemes is always affine.

Proof. By Lemma 2.5.2 there is an open cover $\{W_j\}$ of V by open affines such that each W_j is simultaneously a distinguished open affine in V and a distinguished open affine for some V_i in our open cover. We let $g_j \in \mathcal{O}_Y(V)$ be an element whose localization yields W_j and let $h_j \in \mathcal{O}_Y(V_i)$ be an element whose localization yields W_j. Note that the preimage of W_j is an open affine D_j of X. Indeed, it is the distinguished open affine of $f^{-1}(V_i)$ corresponding to the element $f^\sharp(V_i)(h_j)$.

Let U be the preimage of V in X. Set $R = \mathcal{O}_X(U)$. Let $\widetilde{g}_j$ be the image of g_j under the pullback $f^\sharp(V) : \mathcal{O}_Y(V) \to \mathcal{O}_X(U)$. Note that $D_j := f^{-1}(W_j)$ is the complement of the vanishing locus of $\widetilde{g}_j$. Furthermore, by construction each D_j is an open affine in X. Finally, by Exercise 2.4.19 we know that $\mathcal{O}_X(D_j) \cong R_{\widetilde{g}_i}$.

We must verify that $U \cong \mathrm{mSpec}(R)$. By Theorem 2.4.8 there is a morphism $h : U \to \mathrm{mSpec}(R)$. Furthermore, again by Theorem 2.4.8 the map $f|_U : U \to \mathrm{mSpec}(S)$ is the composition of h with the morphism $\mathrm{mSpec}(R) \to \mathrm{mSpec}(S)$ induced by $f^\sharp$. This means that the restriction $h|_{D_j}$ must be the isomorphism $D_j \cong \mathrm{mSpec}(R_j)$ for every j. Since h is an isomorphism along each open set in this open cover, h is itself an isomorphism. □

4.1.2 Finite morphisms

Finite morphisms are defined by adding one additional property to the definition of an affine morphism.

Definition 4.1.3. We say that a morphism of quasiprojective schemes $f : X \to Y$ is finite if for every open affine V in Y the preimage U in X is an open affine such that $f^\sharp(V) : \mathcal{O}_Y(V) \to \mathcal{O}_X(U)$ is a finite $\mathbb{K}$-algebra morphism.

Again this definition is hard to check directly, so we should develop a criterion that is easier to verify.

Lemma 4.1.4. *Let $f : X \to Y$ be a morphism of quasiprojective schemes. The following conditions are equivalent:*

(1) There is an open covering of Y by open affines $\{V_i\}$ such that for every i the preimage $f^{-1}(V_i)$ is an open affine in X and the map $f^\sharp(V_i) : \mathcal{O}_Y(V_i) \to \mathcal{O}_X(f^{-1}(V_i))$ is a finite $\mathbb{K}$-algebra homomorphism.

(2) f is a finite morphism.

Proof. Lemma 4.1.2 shows that a morphism satisfying (1) will be an affine morphism. It only remains to show that for every open affine V in Y with preimage U we have that $\mathcal{O}_X(U)$ is a finite $\mathcal{O}_Y(V)$-module. Tracing through the proof of Lemma 4.1.2, we can assume there is a set $\{g_j\} \subset \mathcal{O}_Y(V)$ which generates the unit ideal and that $\mathcal{O}_X(U)_{g_j}$ is a finite $\mathcal{O}_Y(V)_{g_j}$-module for every j. We conclude by Proposition 4.0.1.

(2) $\implies$ (1) is immediate. □

Example 4.1.5. Let $f : \mathbb{A}^1 \to \mathbb{A}^1$ be the morphism induced by the map $f^\sharp : \mathbb{K}[x] \to \mathbb{K}[t]$ sending $x \mapsto P(t)$ where P is a degree n polynomial. Then $\mathbb{K}[t]$ is generated as a $\mathbb{K}[x]$-module by $1, t, t^2, \ldots, t^{n-1}$. Thus f is a finite morphism.

Example 4.1.6. The open embedding $\mathbb{A}^1 \backslash \{0\} \to \mathbb{A}^1$ is not finite.

We have already seen one general source of examples of finite maps.

Exercise 4.1.7. Prove that every closed embedding is a finite map.

Warning 4.1.8. The correct analogy between finite maps and covering maps with ramification is when we require the finite morphism to be dominant. We do not impose this as a hypothesis, so "smaller" maps like closed embeddings will also be finite.

4.1.3 Basic properties of finite maps

We next translate several facts about integral ring homomorphisms to the setting of schemes.

Exercise 4.1.9. Prove that a composition of finite morphisms is finite.

Lemma 4.1.10. *Suppose $f : X \to Y$ is a finite morphism. Suppose that $g : Z \to Y$ is an morphism. Then the induced $\widetilde{f} : X \times_Y Z \to Z$ is a finite morphism. (In other words, finite morphisms are stable under base change.)*

Proof. Let $\{V_k\}$ denote an open cover of Y by open affines. For each V_k let $\{U_k\}$ denote the preimage; since f is finite U_k is an open affine in X. Finally, let $\{W_{jk}\}$ be an open cover of $g^{-1}(V_k)$ by open affines. Then the various $U_k \times_{V_k} W_{jk}$ form an open cover of $X \times_Z Y$ by open affines. Thus it suffices to prove the statement when X, Y, Z are affine.

Translating to the algebraic situation, we need to prove the following:

Claim 4.1.11. Suppose that $f^\sharp : S \to R$ is a finite $\mathbb{K}$-algebra homomorphism. Let $g^\sharp : S \to T$ be a $\mathbb{K}$-algebra homomorphism. Then the induced map $T \to R \otimes_S T$ is finite.

To see this claim, for some positive integer n we have a surjective homomorphism $S^{\oplus n} \to R$ of S-modules. Since tensor product is right exact, we obtain a surjection $T^{\oplus n} \to R \otimes_S T$ of T-modules, proving the claim. $\square$

Lemma 4.1.12. *A finite morphism $f : X \to Y$ is topologically closed.*

This is the geometric interpretation of the "Lying Over" theorem.

Proof. One can check if a set is closed by intersecting against all sets in an open cover. Thus we may suppose that Y is affine. Since f is finite this implies that X is also affine, so f is induced by a $\mathbb{K}$-algebra homomorphism $f^\sharp : S \to R$.

Suppose that $Z \subset X$ is a closed subset which is the vanishing locus of an ideal I. Set $J = (f^\sharp)^{-1}I$. Consider the map $\overline{f}^\sharp : S/J \hookrightarrow R/I$. Applying Lying Over (Theorem 4.0.2) with $B = S/J$, $A = R/I$, and $\mathfrak{m}$ any maximal ideal of S/J, we find a prime ideal $\mathfrak{q} \in R/I$ such that $(\overline{f}^\sharp)^{-1}(\mathfrak{q}) = \mathfrak{m}$. In fact, this $\mathfrak{q}$ must be a maximal ideal of R/I. Indeed, we have an injection $S/\mathfrak{m} \hookrightarrow R/\mathfrak{q}$ which is a finite map. Note that $B/\mathfrak{m}$ is a field and $A/\mathfrak{q}$ is a domain. Since the map $\psi_x : A/\mathfrak{q} \to A/\mathfrak{q}$ induced by multiplication by a non-zero element $x \in A/\mathfrak{q}$ is injective and $A/\mathfrak{q}$ is a finite-dimensional $B/\mathfrak{m}$-space, in fact ψ_x must be bijective. This shows that $A/\mathfrak{q}$ is a field and so $\mathfrak{q}$ is a maximal ideal.

Geometrically, this implies that for any point $\mathfrak{m} \in V(J)$ there is a point $\mathfrak{q} \in V(I)$ such that $f(\mathfrak{q}) = \mathfrak{m}$. In other words, the image of Z coincides with the closed set $V(J)$ in Y. This shows that f is a closed map. $\square$

Example 4.1.13. The ring morphism $f^\sharp : \mathbb{K}[x] \to \mathbb{K}[x]_x$ is not finite. Thus the open embedding $D_x \to \mathbb{A}^1$ is not finite. We can also see this by noticing that Lemma 4.1.12 fails.

By combining Lemma 4.1.10 and Lemma 4.1.12 we deduce:

Theorem 4.1.14. *Finite morphisms are proper.*

4.1.4 Exercises

Exercise 4.1.15. Let $X \subset \mathbb{P}^{n+1}$ be a hypersurface and suppose that $p \in \mathbb{P}^{n+1}$ is a $\mathbb{K}$-point not contained in X. Show that projection away from p defines a finite morphism $X \to \mathbb{P}^n$.

4.2 Fibers of finite maps

In this section we continue the analogy between dominant finite morphisms of schemes and covering maps with ramification in topology. We will focus on the fibers of finite morphisms. The following easy lemma describes one of the defining properties of a finite morphism.

Lemma 4.2.1. *Let $f : X \to Y$ be a finite morphism. Then every fiber of f is a finite scheme.*

Proof. Fix a point $y \in Y$ and let $F := X \times_Y y$ be the fiber over y. By Lemma 4.1.10, the induced map $f : F \to y$ is a finite map. Since $y = \mathrm{mSpec}(S)$ for a finite $\mathbb{K}$-module S, we see that $F = \mathrm{mSpec}(R)$ for a finite $\mathbb{K}$-module R. In other words, R is an Artinian ring. We conclude that F is a finite set. □

It is important to note that the converse of Lemma 4.2.1 is not true: for example, most open embeddings are not finite. (See Example 4.1.13.) However, if we add in the properness condition then we do get equivalent definitions.

Theorem 4.2.2. *Let $f : X \to Y$ be a morphism of quasiprojective schemes that is proper and has finite fibers. Then f is a finite morphism.*

Unfortunately, the proof requires tools that we have not developed so far. In fact, even more is true: the following important theorem describes the structure of morphisms with finite fibers. (Such morphisms are called "quasifinite.")

Theorem 4.2.3 (Zariski's Main Theorem)**.** *Let $f : X \to Y$ be a morphism of quasiprojective schemes such that every fiber of f is finite. Then f is the composition of an open embedding $i : X \to Z$ followed by a finite morphism $g : Z \to Y$.*

Exercise 4.2.4. Deduce Theorem 4.2.2 from Theorem 4.2.3.

4.2.1 Degree

Since a finite morphism has finite fibers, it is interesting to "count" the number of points in a fiber.

Definition 4.2.5. Let $f : X \to Y$ be a finite morphism of quasiprojective schemes. For any point $y \in Y$, the fiber $f^{-1}(y)$ is an affine scheme defined by an Artinian $\kappa(y)$-algebra R. We define the degree of f over y to be

$$\deg_y(f) := \dim_{\kappa(y)}(R).$$

(Note that if $f^{-1}(y)$ is empty then our convention is that $R = 0$ and $\deg_y(f) = 0$.)

It turns out that $\deg_y$ has nice topological properties.

Theorem 4.2.6. *Let $f : X \to Y$ be a finite morphism of quasiprojective schemes. Then $\deg_y(f)$ is an upper semicontinuous function on Y.*

Proof. We must show that for any positive integer d the set $Z_d := \{y \in Y | \deg_y(f) \leq d\}$ is an open subset. When $d = 0$, we must show that the image of f is a closed subset of Y which follows from the properness of finite morphisms.

When $d > 0$, we can check whether a subset of Y is open by intersecting against every element in an open cover. In this way we reduce to the case where $Y = \mathrm{mSpec}(S)$ (and hence also $X = \mathrm{mSpec}(R)$) is affine.

Fix a maximal ideal $\mathfrak{m} \subset S$ representing a point $y \in Z_d$. Then we have

$$\dim_{\kappa(y)} R/\langle f^\sharp(\mathfrak{m})\rangle \leq d.$$

If we consider $R \otimes S_\mathfrak{m}$ as a module over the local ring $S_\mathfrak{m}$ and apply Nakayama's lemma, we see that there is a set of d elements $g_1, \ldots, g_d$ which generate $R \otimes S_\mathfrak{m}$. By definition each g_i is a function in some sufficiently small open neighborhood of y. Using the fact that distinguished open affines form a base for the topology, we can find some $h \in S$ such that $g_1, \ldots, g_d \in S_h$. Thus $R \otimes S_h$ is generated by $g_1, \ldots, g_d$ as an S_h-module. We deduce that for any maximal ideal $\mathfrak{n} \subset S_h$ we have

$$\dim_{S/\mathfrak{m}} R/\langle f^\sharp(\mathfrak{n})\rangle \leq d.$$

In other words, every point y in the open set D_h satisfies $\deg_y(f) \leq d$, proving the desired statement. □

In particular this implies:

Corollary 4.2.7. *Let $f : X \to Y$ be a finite morphism of quasiprojective schemes. Suppose that Y is irreducible. Then there is a non-empty open subset $U \subset Y$ such that $\deg_y(f)$ is constant for $y \in U$.*

Exercise 4.2.8. Show that $\deg_y(f)$ can actually increase on closed sets by considering the following two examples of finite morphisms.

(1) The morphism $f : \mathbb{A}^1 \to \mathrm{mSpec}(\mathbb{K}[x, y]/(y^2 - x^3 - x^2))$ defined by $x \mapsto t^2 - 1, y \mapsto t(t^2 - 1)$.

(2) The morphism $f : \mathbb{A}^1 \to \mathrm{mSpec}(\mathbb{K}[x, y]/(y^2 - x^3))$ defined by $x \mapsto t^2, y \mapsto t^3$.

Corollary 4.2.7 shows that for a dominant finite morphism $f : X \to Y$ of quasiprojective varieties there is a non-empty open subset $U \subset Y$ such that $\deg_y(f)$ is constant for $y \in U$. Our next goal is to reinterpret this number in terms of the global geometry of X and Y.

Definition 4.2.9. Let $f : X \to Y$ be a dominant finite morphism of quasiprojective varieties. This induces an inclusion $f^\sharp : \mathbb{K}(Y) \hookrightarrow \mathbb{K}(X)$. We define the degree $\deg(f)$ to be $[\mathbb{K}(X) : \mathbb{K}(Y)]$.

The degree of a dominant finite map is an important invariant! It is used frequently when studying the relationship between varieties which have the same dimension. For example, note that a finite dominant morphism f will be birational if and only if it has degree 1.

Theorem 4.2.10. *Let $f : X \to Y$ be a dominant finite morphism of quasiprojective varieties. There is a non-empty open subset $U \subset Y$ such that for every $y \in U$ we have*

$$\deg_y(f) = \deg(f).$$

Exercise 4.2.8 shows that U need not equal all of Y.

Proof. By replacing Y by an open affine subset, it suffices to consider the case when $Y = \mathrm{mSpec}(S)$ (and hence also $X = \mathrm{mSpec}(R)$) is an affine variety. We will appeal to the following important algebraic theorem.

Theorem 4.2.11 (Grothendieck's Generic Freeness)**.** *Suppose that B is a finitely generated $\mathbb{K}$-algebra (or more generally, a finitely generated algebra over a Noetherian domain). For any finitely generated B-module M, there is an element $g \in B$ such that M_g is a free B_f-module.*

In our setting, this theorem implies that there is an element $g \in S$ such that R_g is a free S_g-module. Denote the rank of this module by r. Then for any maximal ideal $\mathfrak{m} \subset S_g$, we have that $R_g/f_g^\sharp(\mathfrak{m})$ is a free $S_g/\mathfrak{m}$-module of rank r. In other words, for any $y \in D(g)$ we have that $\deg_y(f) = r$.

Finally, we must show that $r = \deg(f)$. It is clear that $R_g \otimes \mathbb{K}(Y)$ has dimension r over $\mathbb{K}(Y)$. But $R_g \otimes \mathbb{K}(Y)$ is a finite extension of $\mathbb{K}(Y)$, and hence must coincide with the field $\mathbb{K}(X)$. □

4.2.2 Criterion for closed embedding

We'll close this section with an interesting criterion for a morphism to be a closed embedding. Certainly a closed embedding must be proper and be set-theoretically injective. Note however that these two conditions are not sufficient (see Exercise 4.2.8.(2)). If we replace "set-theoretically injective" by a condition guaranteeing that the fibers are single points scheme-theoretically, then we do obtain an equivalent criterion.

Theorem 4.2.12. *Let $f : X \to Y$ be a finite morphism of quasiprojective schemes such that $\deg_y(f) \leq 1$ for every point $y \in Y$. Then f is a closed embedding.*

According to Theorem 4.2.2 we may equivalently assume that f is proper and that if the fiber over y is non-empty then it is a single reduced $\kappa(y)$-point.

Proof. It suffices to consider the case when $Y = \mathrm{mSpec}(S)$ (and hence also $X = \mathrm{mSpec}(R)$) is an affine scheme. We want to show that the ring map $f^\sharp : S \to R$ is surjective. Let M denote the cokernel of this B-module map. For any maximal ideal $\mathfrak{m} \subset S$, we have an induced exact sequence

$$S_{\mathfrak{m}} \to R_{\mathfrak{m}} \to M_{\mathfrak{m}} \to 0.$$

If we tensor by $S/\mathfrak{m}$, then the leftmost map becomes $S/\mathfrak{m} \to R/\langle f^\sharp(\mathfrak{m})\rangle$ which by assumption is an isomorphism whenever the rightmost term is not zero. By Nakayama's lemma, we deduce that $M_{\mathfrak{m}} = 0$ for every maximal ideal $\mathfrak{m}$. Since we can test whether a B-module vanishes after localizing at maximal ideals, we see that $M = 0$ and that $f^\sharp$ is surjective. $\square$

4.3 Dimension

In this section we define the dimension of a quasiprojective scheme. We will give several definitions and verify that they all coincide. This is not just pedantry – different properties of the dimension can be most easily observed using different definitions.

Our first definition makes sense for any Noetherian topological space.

Definition 4.3.1. Let X be a Noetherian topological space. The Krull dimension of X is defined to be the supremum over all integers r such that there is a strictly descending chain

$$Z_0 \supsetneq Z_1 \supsetneq \ldots \supsetneq Z_r$$

where each Z_i is an irreducible closed subset of X. We denote the Krull dimension by $\mathrm{krdim}(X)$.

By the DCC condition for closed subsets any descending chain in a Noetherian topological space will eventually stop. However, this does *not* guarantee that the dimension is finite! It is true, but not obvious, that every quasiprojective $\mathbb{K}$-scheme has finite dimension.

Remark 4.3.2. Suppose that $X = \mathrm{mSpec}(R)$ is an affine scheme. Then the Krull dimension is the same as the largest integer r such that there is a chain

$$\mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \ldots \subsetneq \mathfrak{p}_r$$

such that each $\mathfrak{p}_i$ is a prime ideal in R.

Example 4.3.3. We have $\mathrm{krdim}(\mathbb{A}^1) = 1$ corresponding to the chain $0 \subset (x)$ of prime ideals in $\mathbb{K}[x]$. In Example 1.3.7 we checked that $\mathrm{krdim}(\mathbb{A}^2) = 2$.

The following easy exercises clarify the meaning of the Krull dimension:

Exercise 4.3.4. Let X be a quasiprojective scheme. Prove that as we vary X_i over all the irreducible components of X we have $\mathrm{krdim}(X) = \sup_{X_i} \mathrm{krdim}(X_i)$.

Exercise 4.3.5. Let X be a quasiprojective scheme. Prove that $\mathrm{krdim}(X) = \mathrm{krdim}(X_{red})$.

Thus the Krull dimension is determined by its value for quasiprojective varieties. From now on, it will be most convenient to focus on defining the dimension for quasiprojective varieties (and to use Exercise 4.3.4 and Exercise 4.3.5 to extend it to quasiprojective schemes).

Exercise 4.3.6. Let X be a quasiprojective variety of finite dimension of dimension r. Prove that X contains an open affine subset U such that $\mathrm{krdim}(U) = \mathrm{krdim}(X)$.

Prove that if X has dimension ∞ then for any $r > 0$ there is an open affine subset $U \subset X$ with $\mathrm{krdim}(U) = r$.

Even though the Krull dimension is very general and is easy to define, it is quite hard to work with. For example, it is not obvious that $\mathrm{krdim}(\mathbb{A}^n) = n$. Luckily, for $\mathbb{K}$-schemes we have access to alternative definitions that are easier to work with.

4.3.1 Transcendence dimension

Loosely speaking, we expect that the dimension of $\mathbb{A}^n$ should be n since its function ring requires n generators. The following definition makes this intuition precise:

Definition 4.3.7. Let X be a quasiprojective variety. We define the transcendence dimension of X to be

$$\operatorname{trdim}(X) = \operatorname{tr.deg.}(\mathbb{K}(X)/\mathbb{K}).$$

Exercise 4.3.8. Let X be a quasiprojective variety. Show that for any open subset $U \subset X$ we have $\operatorname{trdim}(U) = \operatorname{trdim}(X)$.

A helpful advantage of the transcendence dimension is that it is often easy to compute; for example, we have $\operatorname{trdim}(\mathbb{P}^n) = \operatorname{trdim}(\mathbb{A}^n) = n$.

Example 4.3.9. Let f be an irreducible element of $\mathbb{K}[x_1, \dots, x_n]$. Without loss of generality we may suppose that x_1 divides some term in f. Let X denote the vanishing locus of f. Then $\mathbb{K}(X)$ is a finite algebraic extension of $\mathbb{K}(x_2, \dots, x_n)$. Thus we see that $\operatorname{trdim}(X) = n - 1$.

4.3.2 Dominance dimension

Definition 4.3.10. Let X be a quasiprojective variety. We define the dominance dimension of X to be the supremum of all integers n such that there exists a dominant rational map $f : X \dashrightarrow \mathbb{P}^n$. We denote this quantity by $\operatorname{domdim}(X)$.

Exercise 4.3.11. Let X be a quasiprojective variety. Show that for any open subset U we have $\operatorname{domdim}(U) = \operatorname{domdim}(X)$.

4.3.3 Comparison of dimensions

Before proving our main theorem, we need the following result.

Proposition 4.3.12. *Let $f : X \to Y$ be a dominant finite morphism of affine varieties. Then* $\operatorname{krdim}(X) = \operatorname{krdim}(Y)$.

The main ingredient to this theorem is the "Going Up" theorem for finite extensions, allowing us to relate prime ideals in the two rings of functions.

Proof. Write $X = \operatorname{mSpec}(R)$ and $Y = \operatorname{mSpec}(S)$. Note that both R and S are domains and by Exercise 1.6.12 $f^\sharp$ is an inclusion. Choose a chain of prime ideals in Y. By Lying Over (Theorem 4.0.2) we can lift the minimal prime in the chain to X; by Going Up (Theorem 4.0.3) we can lift the entire chain to a chain in X. In this way we see that $\operatorname{krdim}(X) \geq \operatorname{krdim}(Y)$.

Conversely, suppose given a chain of prime ideals

$$(0) \subsetneq \mathfrak{p}_1 \subsetneq \ldots \subsetneq \mathfrak{p}_n$$

in R. Set $\mathfrak{q}_i = (f^\sharp)^{-1}(\mathfrak{p}_i)$. Then we have

$$(0) \subset \mathfrak{q}_1 \subset \ldots \subset \mathfrak{q}_n$$

and we would like to show that every containment is strict. Suppose for a contradiction that $\mathfrak{q}_i = \mathfrak{q}_{i+1}$. Consider the inclusion of quotients $\overline{f}^\sharp : S/\mathfrak{q}_i \hookrightarrow R/\mathfrak{p}_i$. This is still a finite ring homomorphism. We then localize to get an inclusion $\widetilde{f}^\sharp : \mathrm{Frac}(S/\mathfrak{q}_i) \hookrightarrow (R/\mathfrak{p}_i)_{\overline{f}^\sharp(S/\mathfrak{q}_i)}$. This is again a finite ring homomorphism. Furthermore, since by assumption $\mathfrak{q}_i = \mathfrak{q}_{i+1}$ we see that the image of the prime $\mathfrak{p}_{i+1}$ in this ring will be a non-zero prime ideal.

Since $\mathrm{Frac}(S/\mathfrak{q}_i)$ is a field, we see that $(R/\mathfrak{p}_i)_{\overline{f}^\sharp(S/\mathfrak{q}_i)}$ is an Artinian ring. In particular, since this ring is a domain it should be a field. However, this contradicts the existence of the non-zero prime arising from $\mathfrak{p}_{i+1}$. We conclude that each $\mathfrak{q}_i \subsetneq \mathfrak{q}_{i+1}$ and thus that $\mathrm{krdim}(Y) \geq \mathrm{krdim}(X)$. □

Theorem 4.3.13. *Let X be a quasiprojective variety. Then*

$$\mathrm{domdim}(X) = \mathrm{trdim}(X) = \mathrm{krdim}(X).$$

Henceforth we will denote this common quantity by $\dim(X)$. Note that this result implies that the Krull dimension of any quasiprojective $\mathbb{K}$-scheme will be finite.

Proof. By applying Exercise 4.3.6, Exercise 4.3.8, and Exercise 4.3.11 we reduce to the case when $X = \mathrm{mSpec}(R)$ is an affine variety.

We first show $\mathrm{domdim}(X) = \mathrm{trdim}(X)$. Recall that a dominant map $f : X \dashrightarrow \mathbb{P}^n$ induces an inclusion of function fields $\mathbb{K}(\mathbb{P}^n) \subset \mathbb{K}(X)$. In this way we see that $\mathrm{domdim}(X) \leq \mathrm{trdim}(X)$. To see the converse equality, we will need the celebrated Noether Normalization Theorem.

Theorem 4.3.14 (Noether Normalization)**.** *Let R be a finitely generated $\mathbb{K}$-algebra domain. Suppose that the transcendence degree of $\mathrm{Frac}(R)$ over $\mathbb{K}$ is n. Then there exists algebraically independent elements $x_1, \ldots, x_n \in R$ such that R is a finite integral extension of $\mathbb{K}[x_1, \ldots, x_n]$.*

By Noether Normalization, if $\mathrm{trdim}(X) = n$ then $\mathrm{mSpec}(R)$ admits a finite morphism to $\mathbb{A}^n$. We conclude that $\mathrm{domdim}(X) \geq \mathrm{trdim}(X)$, so that the two quantities are equal.

We finish the proof by showing $\mathrm{trdim}(X) = \mathrm{krdim}(X)$. The proof is by induction on $\mathrm{trdim}(X)$. First suppose that $\mathrm{trdim}(X) = 0$. This implies that R is an Artinian ring. Thus every prime ideal in R is a maximal ideal, showing that $\mathrm{krdim}(X) = 0$ as well.

Now we prove the induction step. Suppose that $\mathrm{trdim}(X) = n$. Using Noether normalization and our argument in the paragraph above, this condition implies that X admits a finite morphism $f : X \to \mathbb{A}^n$. By Proposition 4.3.12 we have $\mathrm{krdim}(X) = \mathrm{krdim}(\mathbb{A}^n)$. Thus it suffices to show that $\mathrm{krdim}(\mathbb{A}^n) = n$.

Consider a maximal chain of prime ideals in $\mathbb{K}[x_1, \ldots, x_n]$:

$$(0) \subsetneq \mathfrak{p}_1 \subsetneq \ldots \subsetneq \mathfrak{p}_r$$

Since $\mathbb{K}[x_1, \ldots, x_n]$ is a UFD, $\mathfrak{p}_1$ must be a principal ideal generated by an irreducible element f. (If it were not, we could find a smaller non-zero prime ideal contained in $\mathfrak{p}_1$.) Set $Y = V(f)$. By our induction assumption $\mathrm{krdim}(Y) = \mathrm{trdim}(Y)$. In Example 4.3.9 we showed that Y has transcendence dimension $n - 1$, showing that the chain of ideals $\mathfrak{p}_1 \subsetneq \ldots \subsetneq \mathfrak{p}_r$ has n elements. We conclude that $\mathrm{krdim}(\mathbb{A}^n) = n$ as desired. □

Notation 4.3.15. A quasiprojective scheme is said to be equidimensional if every component has the same dimension. An equidimensional quasiprojective scheme of dimension 1 is called a curve, of dimension 2 is called a surface, and of dimension $n \geq 3$ is called an n-fold.

4.3.4 Exercises

Exercise 4.3.16. Let X be a quasiprojective scheme. Let $U \subset X$ be an open subset. Prove that $\dim(U) \leq \dim(X)$. Prove that if U intersects the irreducible component of X of maximal dimension then $\dim(U) = \dim(X)$.

Exercise 4.3.17. Let X be a quasiprojective scheme. Suppose that $f : Y \to X$ is a closed embedding. Prove that $\dim(Y) \leq \dim(X)$.

Exercise 4.3.18. (1) Suppose that $f : X \to Y$ is a finite morphism of quasiprojective schemes. Prove that $\dim(X) \leq \dim(Y)$.

(2) Suppose that $f : X \to Y$ is a dominant finite morphism of quasiprojective varieties. Prove that $\dim(X) = \dim(Y)$.

Exercise 4.3.19. Let X be a quasiprojective variety. Suppose that $f : X \dashrightarrow Y$ is a dominant rational map. Prove that $\dim(X) \geq \dim(Y)$.

Exercise 4.3.20. Let X be a quasiprojective scheme. Prove that $\dim(X) = 0$ if and only X is a finite set.

Exercise 4.3.21. Let X and Y be quasiprojective varieties. Prove that $\dim(X \times Y) = \dim(X) + \dim(Y)$.

More generally, if we have morphisms $f : X \to Z$ and $g : Y \to Z$, prove that

$$\dim(X \times_Z Y) \leq \dim(X) + \dim(Y).$$

Find examples which exhibit the various dimensions allowed by this bound.

Exercise 4.3.22. Prove that $\dim(G(k,n)) = k(n-k)$.

Exercise 4.3.23. Let X be a quasiprojective $\mathbb{K}$-scheme and let $\mathbb{L}/\mathbb{K}$ be a field extension. Prove that $\dim(X)$ is the same as $\dim(X_{\mathbb{L}})$.

Exercise 4.3.24. Let X be a quasiprojective scheme. For any $x \in X$, we define the local dimension $\operatorname{locdim}_x(X)$ to be the Krull dimension of the ring $\mathcal{O}_{X,x}$.

Prove that the local dimension of x is the same as the maximal dimension of any irreducible component of X that contains x.

4.4 Properties of dimension

In Section 4.3 we saw several basic properties of the dimension. In this section we will discuss two difficult properties.

4.4.1 Krull's Prinicipal Ideal Theorem

Krull's Principal Ideal Theorem is the following statement:

Theorem 4.4.1 (Krull's Principal Ideal Theorem). *Let A be a Noetherian ring and let $f \in A$ be a non-unit. Let $\mathfrak{p}$ be a minimal element in the set of prime ideals containing (f). Then either*

(1) $\mathfrak{p}$ is a minimal prime in A, or

(2) the maximal length chain of primes descending from $\mathfrak{p}$ has the form $\mathfrak{p}_0 \subsetneq \mathfrak{p}$.

If f is a non-zero divisor, then we are guaranteed to be in case (2). If $\mathrm{Nil}(A) = 0$ and f is a zero divisor, then we are guaranteed to be in case (1).

Example 4.4.2. When A is not reduced then zero-divisors can land in case (2). Consider for example the zero-divisor y in the ring $\mathbb{K}[x, y]/(x^2, xy)$. The prime ideal (x, y) is a minimal prime containing y and we have a chain of prime ideals $(x) \subset (x, y)$.

Suppose that X is a quasiprojective scheme and that $f \in \mathcal{O}_X(X)$ is a non-unit. Let $\mathfrak{p}$ be a minimal prime containing f. Krull's PIT shows that we cannot "sandwich" any irreducible subsets in the middle of $V(\mathfrak{p}) \subsetneq X$. We might optimistically hope that if we extend a chain of irreducible subsets downward from $V(\mathfrak{p})$ then (together with $V(\mathfrak{p})$) we obtain a maximal length chain in X. This is exactly the content of the following two results.

Theorem 4.4.3 (Geometric Krull's PIT). *Let X be an irreducible quasiprojective variety of dimension n. Suppose that $f \in \mathcal{O}_X(X)$ is a non-unit and set $Z = V(f)$. Then* $\dim(Z) = \dim(X) - 1$.

Proof. Since f is not a unit, there must be an open affine $U \subset X$ such that the restriction of f to U is also not a unit. (If the restriction of f to every open affine were a unit, we could take multiplicative inverses along each open set and glue to obtain a multiplicative inverse on X.) Thus $V(f)$ is a proper closed subset of X. By Exercise 2.8.16 we can find an open affine in X whose intersection with $V(f)$ is dense. Replacing X by this open affine, we may assume that $X = \mathrm{mSpec}(R)$ is an affine variety.

Let $\mathfrak{p}$ be a prime ideal minimal amongst all primes containing f and let W be the vanishing locus of $\mathfrak{p}$. Since f cannot be a zero-divisor, Krull's PIT shows that $\mathfrak{p}$ has height 1. In other words, there is no irreducible closed subset Y satisfying $W \subsetneq Y \subsetneq X$.

Applying Noether normalization to R we obtain a finite map $f : X \to \mathbb{A}^n$. The restriction of f to W is also finite, so by Exercise 4.3.18 we have $\dim(W) = \dim(f(W))$. Note that $f(W)$ is irreducible (since W is) and closed (since f is finite) and thus contained in an irreducible hypersurface H in $\mathbb{A}^n$. We claim that $f(W) \subsetneq H$. Indeed, if we had strict containments $f(W) \subsetneq H \subsetneq \mathbb{A}^n$ then by applying the Going Down theorem (Theorem 4.0.4) to the ideals of H and $f(Z)$ and the extension $\mathbb{K}[x_1, \ldots, x_n] \subset R$ we would find that $f^{-1}(H)$ is an irreducible subset that is properly contained in X and that contains Z but is not equal to it, a contradiction.

Since any hypersurface in $\mathbb{A}^n$ has dimension $n-1$ (Example 4.3.9), we see that $\dim(W) = \dim(f(W)) = n-1$. Since $W \subset Z$ we have $\dim(Z) \geq \dim(X) - 1$. On the other hand, any chain of prime ideals in $R/(f)$ will yield a chain of prime ideals in R whose length is larger by one, showing that $\dim(Z) + 1 \leq \dim(X)$. This proves the equality. □

Remark 4.4.4. The converse of Theorem 4.4.3 is false: a codimension 1 subset need not be defined as the vanishing locus of a single equation (see Example 3.2.4 and Exercise 4.4.14).

The corresponding statement for quasiprojective schemes is:

Corollary 4.4.5 (Geometric Krull's PIT)**.** *Let X be a quasiprojective scheme and let $f \in \mathcal{O}_X(X)$. Set $Z = V(f)$. If Z_0 is a component of Z and X_0 is a component of X which contains Z_0 then Z_0 has codimension 0 or 1 in X_0.*

Proof. We know that Z_0 will be a component of the restriction of f to X_0. Thus we may suppose that X is irreducible. If f is nilpotent, then $V(f) = X$ will have codimension 0. If f is not nilpotent, then its restriction to X_{red} is non-zero and we conclude by Theorem 4.4.3. □

Exercise 4.4.6. Let X be a quasiprojective variety. Suppose that $f : Y \to X$ is a closed embedding. Prove that if $\dim(Y) = \dim(X)$ then f is an isomorphism.

Krull's Principal Ideal Theorem is surprisingly useful. Here is a couple applications.

Proposition 4.4.7. *Let X be an affine scheme. Let $x \in X$ and let n be the largest dimension of any component of X that contains x. Then there are functions $f_1, \ldots, f_n$ such that x is an irreducible component of $V(f_1, \ldots, f_n)$.*

Proof. The proof is by induction on n. If $n = 0$ then the statement is clear. If $n \geq 1$, the maximal ideal $\mathfrak{m}$ defining x will not be a minimal prime in $\mathcal{O}_X(X)$. By the Prime Avoidance Lemma $\mathfrak{m}$ is not contained in the union of the minimal primes of $\mathcal{O}_X(X)$. Thus $\mathfrak{m}$ contains a non-zero divisor f_n. For every irreducible component X_i of X that contains x, the intersection $V(f_n) \cap X_i$ will have dimension one less than $\dim(X_i)$ by Theorem 4.4.5. We apply the induction hypothesis to $V(f_n)$ to obtain functions $\overline{f}_1, \ldots, \overline{f}_{n-1}$ in $\mathcal{O}_X(X)/\mathfrak{m}$. Lifting these functions to $\mathcal{O}_X(X)$ we obtain the result. □

Suppose that the finitely generated $\mathbb{K}$-algebra R is a domain. Then Geometric Krull's PIT shows that for any $f \in R$ the vanishing locus $V(f)$ has codimension 1 in $\mathrm{mSpec}(R)$. As mentioned above the converse does not hold in general, but the following proposition shows that the converse does hold when R is a UFD.

Proposition 4.4.8. *Let R be a finitely generated $\mathbb{K}$-algebra that is a UFD. Suppose that $Z \subset \mathrm{mSpec}(R)$ is a subvariety of codimension 1. Then $Z = V(f)$ for some $f \in R$.*

In fact, this condition characterizes the UFDs: if R is a domain which is not a UFD, then there will be some codimension 1 prime which is not principal.

Proof. Choose some $g \in R$ such that $Z \subset V(g)$. Write $g = \prod g_i$ as a decomposition into irreducibles. Then $Z \subset V(g_1) \cup V(g_2) \cup \ldots \cup V(g_r)$. Since Z is irreducible, it is contained in some $V(g_i)$. Note that $V(g_i)$ is also an irreducible subvariety and by Krull's PIT it has codimension 1 in $\mathrm{mSpec}(R)$. Applying Exercise 4.4.6 we see that $Z = V(g_i)$. □

4.4.2 Fiber dimension

The following result is another surprisingly useful theorem governing the behavior of dimension.

Theorem 4.4.9. *Let $f : X \to Y$ be a morphism of quasiprojective varieties.*

(1) For every non-empty fiber F we have $\dim(X) \leq \dim(Y) + \dim(F)$.

(2) Suppose that f is dominant. Then there is an open set $V \subset Y$ such that every fiber F over a point in V is non-empty and satisfies $\dim(X) = \dim(Y) + \dim(F)$.

Remark 4.4.10. If a fiber F has many components, then the dimension inequality in (1) holds true for each component of F. We can deduce this by replacing X by an open subset which removes the other components of F.

Remark 4.4.11. The "correct" statement is upper semicontinuity of fiber dimension, which is a little bit stronger. See Exercise 4.4.19.

Let's first give an informal proof of (2). Suppose that $X = \mathrm{mSpec}(R)$ has dimension n and $Y = \mathrm{mSpec}(S)$ has dimension m. If we knew that the map $f^\sharp : S \to R$ were the composition of the inclusion $S \to S[t_1, \ldots, t_{n-m}]$ with a finite inclusion $S[t_1, \ldots, t_{n-m}] \to R$, then the fibers of $f : X \to Y$ would have the same dimension as the fibers of $Y \times \mathbb{A}^{n-m} \to Y$. Unfortunately such a factorization is not always possible – although it looks reminiscent of Noether Normalization, we are not in a setting to apply this result. However, we *can* apply Noether Normalization to the $\mathrm{Frac}(S)$-algebra $R \otimes_{\mathbb{K}} \mathrm{Frac}(S)$ to get a factorization of the inclusion $\mathrm{Frac}(S) \to R \otimes_{\mathbb{K}} \mathrm{Frac}(S)$. As we've seen before, we can "spread out" this result on $\mathrm{Frac}(S)$ to obtain a factorization over an open subset of Y.

Proof. After replacing X and Y by suitably chosen open affines, we may suppose that X and Y are affine schemes. We write $X = \mathrm{mSpec}(R)$ and $Y = \mathrm{mSpec}(S)$. We set $n = \dim(X)$ and $m = \dim(Y)$.

(1) Suppose F is the fiber over $y \in Y$. Apply Proposition 4.4.7 to find m functions $f_1, \ldots, f_m$ on Y such that y is an irreducible component of $V(g_1, \ldots, g_m)$. After replacing Y by an open neighborhood of y and X by the preimage of this open neighborhood, we may suppose that $y = V(g_1, \ldots, g_m)$. Then the fiber over Y is defined by the equations $\{f^\sharp g_i\}$. By Krull's PIT we see that $\dim(F) + m \geq n$.

(2) Since f is dominant Exercise 1.6.12 shows that $f^\sharp$ is injective. As discussed above, we consider the injection $f^\sharp : \mathrm{Frac}(S) \to R \otimes_{\mathbb{K}} \mathrm{Frac}(S)$. Note that the rightmost term still has fraction field $\mathrm{Frac}(R)$ and that the transcendence degree of $\mathrm{Frac}(R)/\mathrm{Frac}(S)$ is $n - m$. Noether Normalization (Theorem 4.3.14) shows that for some algebraically independent elements $t_1, \ldots, t_{n-m}$ we have a factorization

$$\begin{tikzcd}
\mathrm{Frac}(S) \arrow[rr, "f^\sharp"] \arrow[dr] & & R \otimes_{\mathbb{K}} \mathrm{Frac}(S) \\
& \mathrm{Frac}(S)[t_1, \ldots, t_{n-m}] \arrow[ur] &
\end{tikzcd}$$

Since each $t_i \in R \otimes_{\mathbb{K}} \mathrm{Frac}(S)$, by clearing denominators we can find a single element $s \in S$ such that $st_i \in R$ for every i. Consider the diagram

$$\begin{tikzcd}
S_s \arrow[r] \arrow[d] & S_s[t_1, \ldots, t_{n-m}] \arrow[r] \arrow[d] & R_s \arrow[d] \\
\mathrm{Frac}(S) \arrow[r] & \mathrm{Frac}(S)[t_1, \ldots, t_{n-m}] \arrow[r] & R \otimes_{\mathbb{K}} \mathrm{Frac}(S)
\end{tikzcd}$$

Unfortunately the top right horizontal map may not be a finite inclusion. However, suppose we let $u_1, \ldots, u_k$ be a finite set of generators for R_s as a $S_s[t_1, \ldots, t_{n-m}]$-algebra. By comparing to the bottom row of the diagram, we see that each u_i satisfies a monic equation with coefficients in $\mathrm{Frac}(S)[t_1, \ldots, t_{n-m}]$. The coefficients may not be in $S_s[t_1, \ldots, t_{n-m}]$, but there is an element $\widetilde{s}$ so that all coefficients are in $S_{s\widetilde{s}}[t_1, \ldots, t_n]$.

Let $V = D_{ss'}$ in Y and set $U = f^{-1}(V)$. The computation above shows that $f|_U$ factors as a dominant finite morphism $U \to V \times \mathbb{A}^{n-m}$ followed by the projection $V \times \mathbb{A}^{n-m} \to V$. Let F be any fiber of $U \to V$. Since finiteness of a morphism is preserved by base change, we obtain a surjective finite map $F \to \mathbb{A}^{n-m}$. Since finite maps are closed, there must be an irreducible component F_i of F which admits a finite surjective map $F_i \to \mathbb{A}^{n-m}$. By Exercise 4.3.18 $\dim(F_i) \leq n - m$, proving the statement. $\square$

Exercise 4.4.12. Use Theorem 4.4.9 to finish the proof of Chevalley's Theorem (Theorem 1.6.10).

4.4.3 Exercises

Exercise 4.4.13. Let $f : X \to Y$ be a dominant morphism of quasiprojective varieties. Suppose that $\dim(X) = \dim(Y)$. Show that there is some non-empty open subset $U \subset Y$ such that the map $f : f^{-1}(U) \to U$ is a finite map. Such maps are called generically finite. (Hint: use the proof of Theorem 4.4.9.)

More generally, if we have a dominant morphism of quasiprojective schemes $f : X \to Y$ such that $\dim(X) = \dim(Y)$, again show that there is some non-empty open subset $U \subset Y$ such that the map $f : f^{-1}(U) \to U$ is a finite map.

Exercise 4.4.14. Consider the affine variety $X = \mathrm{mSpec}(\mathbb{K}[w, x, y, z]/(wz - yx))$. Show that (w, x) is a prime ideal that defines a codimension 1 subvariety $Z \subset X$. Show that there is no principal ideal that defines Z. (Hint: one way is to use the fact that the ring defining X is a graded ring. A more general approach will be presented in Exercise 5.1.16.)

Exercise 4.4.15. Let X, Y, and Z be quasiprojective varieties equipped with *dominant* morphisms $f : X \to Z$ and $g : Y \to Z$. Prove that

$$\dim(X) + \dim(Y) - \dim(Z) \leq \dim(X \times_Z Y) \leq \dim(X) + \dim(Y).$$

Find examples which exhibit the various dimensions allowed by this bound.

Exercise 4.4.16. Let $X \subset \mathbb{P}^n$ be a closed subscheme of dimension ≥ 1. Prove that if H is a hyperplane that does not contain any component of X then the intersection $X \cap H$ has dimension one less than X. (Also, if X has dimension 0 then $X \cap H$ will be empty.)

Exercise 4.4.17. Let X be an affine variety and consider the blow-up $\mathrm{Bl}_I(X)$ along an ideal I. Let Z denote the preimage of $V(I)$ (as defined in Exercise 3.4.15). Prove that Z has codimension 1 in $\mathrm{Bl}_I(X)$.

Exercise 4.4.18. Let $f : X \to Y$ be a proper morphism of quasiprojective schemes. Suppose that Y is irreducible and that every fiber F of f is irreducible of the same dimension. Prove that X is irreducible. (What is a counterexample when f fails to be proper?)

Exercise 4.4.19. This exercise shows the upper semicontinuity of fiber dimension.

Suppose $f : X \to Y$ is a morphism of quasiprojective schemes. For any point $x \in X$ we define $\mu(x)$ to be the local dimension of x in the fiber F of f containing x. (The local dimension was defined in Exercise 4.3.24.)

(1) Prove that for any point x we have $\mathrm{locdim}_x(X) \leq \dim(Y) + \mu(x)$. (Hint: apply Theorem 4.4.9.(1) and Exercise 4.3.24.)

(2) Prove that the map $x \mapsto \mu(x)$ is upper semicontinuous. (Hint: apply Theorem 4.4.9.(2) and argue by Noetherian induction.)

4.5 Application: Moduli spaces and incidence correspondences

A moduli space in algebraic geometry is a scheme M whose points parametrize some other kind of object. We have seen two examples already: the moduli space $\mathbb{P}^{(n+2)(n+1)/2-1}$ of quadric hypersurfaces and the Grassmannian $\mathbb{G}(k,n)$ parametrizing k-dimensional planes in $\mathbb{P}^n$. As we saw in these two cases, the $\mathbb{K}$-points of the moduli space parameterize the corresponding objects which are defined over $\mathbb{K}$, while the $\mathbb{L}$-points parametrize the “orbits” of objects defined over $\mathbb{L}$.

Here is another example:

Example 4.5.1. Consider the hypersurfaces $H \subset \mathbb{P}^n$ which are defined by a single equation f of degree d. By identifying the coefficients of f as coordinates on a projective space, we see that the degree d hypersurfaces are parametrized by the $\mathbb{K}$-points on $\mathbb{P}^{\binom{n+d}{d}-1}$. We call this projective space the moduli space of degree d hypersurfaces.

Another important property of moduli spaces is the existence of a family over M: is there a scheme $\mathcal{U}$ with a morphism $u : \mathcal{U} \to M$ such that the fiber of u over a point $m \in M$ is exactly the object parametrized by M? We have already seen this construction for the Grassmannian. Fortunately, it is also easy to construct a family over the moduli space of hypersurfaces.

Example 4.5.2. Fix a projective space $\mathbb{P}^n$ and consider the moduli space $M = \mathbb{P}^{\binom{n+d}{d}-1}$ of degree d hypersurfaces. We let $\mathcal{I}$ denote the set of ordered $(n+1)$-tuples of non-negative integers which add up to d. Each degree d polynomial on $\mathbb{P}^n$ can be written the form $\sum_{I\in\mathcal{I}} a_I x^I$. Note that M is equipped with the homogeneous coordinate ring $\mathbb{K}[y_I]_{I\in\mathcal{I}}$ where the value of y_I represents the coefficient of the monomial x^I.

The family over the moduli space is the hypersurface $\mathcal{H}$ on $M \times \mathbb{P}^n$ defined in bihomogeneous coordinates (as in Section 3.3.1) by the equation $\sum_{I\in\mathcal{I}} y_I x^I$. It comes equipped with a projection map $u : \mathcal{H} \to M$.

Broadly speaking, a key goal in moduli theory is to understand the sublocus of the moduli space M parametrizing objects with special properties. In particular it is interesting to know when “most” objects parametrized by M have a special property.

Definition 4.5.3. Let M be a moduli space. We say that a property P “holds for a general object parametrized by M” if the Zariski closure of the set of objects which fail P is a proper closed subset of M.

4.5.1 Incidence correspondences

Incidence correspondences are one of the best ways to construct and study special subloci of a moduli space M. Suppose given two moduli spaces M_1, M_2. An incidence correspondence is a subset $Z \subset M_1 \times M_2$ consisting of the set of pairs (X, Y) of parametrized objects such that X, Y interact in some special way.

Often given an incidence correspondence Z we will be in a situation where one map $Z \to M_2$ is easy to understand and the other map $Z \to M_1$ is the one we would like to study. Below we will see a an example of incidence correspondences at work.

4.5.2 Lines in hypersurfaces

In this subsection we work over an algebraically closed field $\mathbb{K}$.

Proposition 4.5.4. *Fix a projective space $\mathbb{P}^n$ and a degree d. Suppose that $d > 2n - 3$. Then the general degree d hypersurface does not contain any lines.*

Let M denote the moduli space of degree d hypersurfaces and let $\mathbb{G}(1, n)$ denote the moduli space of lines in $\mathbb{P}^n$. We would like to define the incidence correspondence $Z \subset M \times \mathbb{G}(1, n)$ which is the set of pairs (H, ℓ) such that $\ell \subset H$.

Claim 4.5.5. There is a closed subset $Z \subset M \times \mathbb{G}(1, n)$ whose $\mathbb{K}$-points represent pairs (H, ℓ) such that $\ell \subset H$.

Proof. Consider the family of hyperplanes $\mathcal{H} \subset M \times \mathbb{P}^n$ and the family of lines $\mathcal{U} \subset \mathbb{G}(1, n) \times \mathbb{P}^n$. The preimages $p_{13}^{-1}(\mathcal{H})$ and $p_{23}^{-1}(\mathcal{U})$ will be closed subsets of $M \times \mathbb{G}(1, n) \times \mathbb{P}^n$. Let $\widehat{Z}$ denote their intersection, a closed set. Note that $\widehat{Z}$ denotes the triples (H, ℓ, p) such that $p \in H$ and $p \in \ell$.

Consider the projection map $p_{12} : M \times \mathbb{G}(1, n) \times \mathbb{P}^n \to M \times \mathbb{G}(1, n)$. The fiber of p_{12} over a point (H, ℓ) will represent the intersection of H and ℓ. By Exercise 4.4.19 there will be a closed subset $\widetilde{Z} \subset \widehat{Z}$ consisting of points where the local fiber dimension of p_{12} is ≥ 1. The set-theoretic image of $\widetilde{Z}$ in $M \times \mathbb{G}(1, n)$ will be the locus Z of pairs (H, ℓ) such that $\ell \subset H$. Since p_{12} is proper, the image Z of $\widetilde{Z}$ is closed. $\square$

Proof of Proposition 4.5.4: Our strategy is to show that the dimension of Z is less than the dimension of M. To prove this, we need to compute the dimension of the fibers of the map $Z \to \mathbb{G}(1, n)$. Let's consider the preimage of a $\mathbb{K}$-point. Since the PGL_{n+1}-action on the space of lines is transitive, this dimension is independent of the choice of line. We may as well choose the line ℓ whose equation is $(x_2, \ldots, x_n)$. A degree d hypersurface f will contain the line ℓ if and only if $f(x_0, x_1, 0, \ldots, 0)$ is identically 0. In other words, we see that f cannot involve any monomials which only use the variables x_0, x_1. Degree d monomials in x_0, x_1 form a $(d + 1)$-dimensional vector space, showing that every fiber of $Z \to \mathbb{G}(1, n)$ over a $\mathbb{K}$-point is a projective space of dimension $\dim(M) - (d + 1)$.

By Theorem 4.4.9 we deduce that

$$\dim(Z) \leq \dim(\mathbb{G}(1, n)) + \dim(F) = \dim(M) + 2(n - 1) - (d + 1).$$

When $d > 2n - 3$ then $\dim(Z) < \dim(M)$ so the image of $Z \to M$ will be contained in a proper closed subset of M. $\square$

Remark 4.5.6. Although it does not prove it, the argument above suggests that a general hypersurface of degree $d < 2n-3$ will contain a $(2n-3-d)$-dimensional family of lines. This turns out to be the case. For example, when $d = 1$ then any degree d hypersurface H is isomorphic to $\mathbb{P}^{n-1}$ and thus contains a $2(n-2)$-dimensional family of lines. When $d = 2$ we saw an example in Exercise 3.1.15.

4.5.3 Exercises

Exercise 4.5.7. Let $\mathbb{K}$ be an algebraically closed field. Let $X \subset \mathbb{P}^n$ be a closed subscheme of dimension ≥ 1. Prove that a general hyperplane H satisfies $\dim(X \cap H) = \dim(X) - 1$. (Hint: you will need Exercise 2.11.14.)

Exercise 4.5.8. Let $\mathbb{K}$ be an algebraically closed field. Prove that there is some positive integer $d = d(k,n)$ such that a general hypersurface of degree $\geq d$ in $\mathbb{P}^n$ will not contain any k-planes. (What is the function $d(k,n)$?)

Exercise 4.5.9. There are exactly 27 lines in the Fermat cubic hypersurface $x_0^3+x_1^3+x_2^3+x_3^3$ in $\mathbb{P}^3_{\mathbb{C}}$. Can you find them all? (It turns out that every smooth cubic hypersurface over an algebraically closed field will contain exactly 27 lines.)

Exercise 4.5.10. Fix a projective space $\mathbb{P}^n$ and a degree d. Suppose that $d \leq 2n-3$. Suppose that you can find a degree d hypersurface which has a $(2n-3-d)$-dimensional family of lines. Prove that a general degree d hypersurface admits a $(2n-3-d)$-dimensional family of lines.

Exercise 4.5.11. Let $X \subset \mathbb{P}^n$ be a closed subscheme. If we fix a $\mathbb{K}$-point $p \in \mathbb{P}^n$ that is not contained in X then projection away from p defines a morphism $\phi : X \to \mathbb{P}^{n-1}$. Prove that for a general point p the morphism ϕ has finite fibers.

(Remember, this means that there is a non-empty open subset $U \subset \mathbb{P}^n$ such that any $\mathbb{K}$-point contained in U has this property. We do not insist that U actually contain any $\mathbb{K}$-points when our ground field is finite.)

Chapter 5

Smoothness

Suppose that M is a smooth manifold. There are several approaches one can take to defining the tangent space at a point $x \in M$:

(1) Chart structure: one can first define the tangent space for points on $\mathbb{R}^n$ and then "transform" these spaces to M using the chart structure. One must verify that the definition does not depend on the choice of chart.

(2) Jets of curves: consider the set of curves $\sigma : J \to M$ which are smooth at x. We can define an equivalence relation on such curves by setting $\sigma_1 \sim \sigma_2$ if for every smooth function $f : M \to \mathbb{R}$ defined on a neighborhood of x the derivatives of $f \circ \sigma_1$ and $f \circ \sigma_2$ coincide. We can define the tangent space to be the set of equivalence classes of such σ.

(3) Derivations: a tangent vector at x allows us to take directional derivatives of functions near x. Alternatively, we can use this feature as a way to define tangent vectors.

Let $\mathcal{C}_x^\infty$ denote the set of germs of smooth real-valued functions near x. A derivation is a linear map $T : \mathcal{C}_x^\infty \to \mathbb{R}$ satisfying the product rule $T(fg) = f(x)T(g)+g(x)T(f)$. Then the vector space of derivations is the tangent space at x.

It will be helpful to modify this definition slightly. Let $(\mathcal{C}_x^\infty)_0$ denote the germs of all functions which vanish at x. Note that any derivation is determined by its values on the subset $(\mathcal{C}_x^\infty)_0 \subset \mathcal{C}_x^\infty$; indeed, for any function $f \in \mathcal{C}_x^\infty$ the behavior of a derivation on f is determined by its behavior for $f - f(x) \in (\mathcal{C}_x^\infty)_0$ where $f(x)$ denotes the constant function. Thus we can define the tangent space as the space of derivations of $(\mathcal{C}_x^\infty)_0$. The advantage of this perspective is that the product rule now becomes $T(fg) = 0$.

In algebraic geometry there is no easy analogue of the first definition – while every scheme admits a covering by affine charts, most schemes do not admit coverings by charts isomorphic to $\mathbb{A}^n$. However, the second and third definitions both admit analogues in our setting.

Recall that the "foundational" objects for schemes are not topological spaces but rings of functions. Correspondingly, the cotangent space is the natural construction for schemes; we then define tangent spaces by taking a dual.

5.1 Zariski tangent space

Our first definition of a "tangent space" in algebraic geometry is based on derivations. We will give a second definition – which is slightly different – in Section 5.3.

Definition 5.1.1. Let X be a quasiprojective scheme. Fix $x \in X$ and let $\mathfrak{m}_x$ denote the maximal ideal of the stalk $\mathcal{O}_{X,x}$. The Zariski cotangent space at x is defined to be

$$T_{X,x}^{\vee} := \mathfrak{m}_x/\mathfrak{m}_x^2$$

considered as a finite-dimensional vector space over the residue field $\kappa(x)$.

The Zariski tangent space $T_{X,x}$ is the dual of this $\kappa(x)$-vector space.

Note that there is no need to pass all the way to the local ring $\mathcal{O}_{X,x}$ to compute the Zariski cotangent space: if we choose any open affine mSpec(R) containing x and let $\mathfrak{m}$ be the maximal ideal associated to x, the Zariski cotangent space is simply $\mathfrak{m}/\mathfrak{m}^2$ (since localization and quotients commute).

Remark 5.1.2. The Zariski tangent space models the derivation approach to the tangent space of a manifold. In our setting $\mathcal{O}_{X,x}$ is analogous to the space of germs $\mathcal{C}_x^{\infty}$ and the maximal ideal $\mathfrak{m}_x$ is analogous to $(\mathcal{C}_x^{\infty})_0$. Based on our discussion in the introduction to the chapter, a derivation should be a map $\mathfrak{m}_x \to \mathbb{L}_x$ that sends $\mathfrak{m}_x^2$ to zero. In other words, a derivation should be an element of $(\mathfrak{m}_x/\mathfrak{m}_x^2)^{\vee}$.

Example 5.1.3. Suppose that $X \subset \mathbb{A}^2$ is the vanishing locus of a polynomial f which contains the origin $\mathfrak{m} = (x, y)$. We can write

$$f = ax + by + f'$$

where $f' \in \mathfrak{m}^2$. Using the usual notion of derivatives we see that the tangent line of X at the origin should have equation $ax + by = 0$. Let's analyze this from the perspective of the Zariski tangent space.

We can write the cotangent space $(x, y)/(x, y)^2$ of the origin as $\mathbb{K}x \oplus \mathbb{K}y$ and the dual space as $\mathbb{K}x^{\vee} \oplus \mathbb{K}y^{\vee}$ (although it would be more traditional to use the notation $\mathbb{K}dx \oplus \mathbb{K}dy$ for the cotangent space and $\mathbb{K}\frac{d}{dx} \oplus \mathbb{K}\frac{d}{dy}$ for the tangent space). With this notation:

(1) The cotangent space of X at the origin is the quotient of $\mathbb{K}x \oplus \mathbb{K}y$ by the subspace spanned by $a\,x + b\,y$.

(2) The tangent space of X at the origin is the subspace of $\mathbb{K}x^{\vee} \oplus \mathbb{K}y^{\vee}$ defined by the equation $ax^{\vee} + by^{\vee} = 0$.

We expect that X should be "smooth" at the origin precisely when f has non-zero linear part. We can reinterpret this condition using the Zariski cotangent space: X should be "smooth" at the origin when $\dim_{\mathbb{K}} T_{X,0} = 1$ and should fail to be "smooth" when $\dim_{\mathbb{K}} T_{X,0} = 2$.

Example 5.1.4. More generally, the Zariski tangent space at the origin in $\mathbb{A}^n$ is isomorphic to $\mathbb{K}^n$ under the identification

$$(a_1, \ldots, a_n) \leftrightarrow \left(g \mapsto \sum_{i=1}^{n} a_i \frac{\partial g}{\partial x_i}(0) \right).$$

Note that the role of the derivative $\frac{\partial}{\partial x_i}$ is just to pick off the x_i-coefficient of the linear term of g. If we use a different point x with residue field $\mathbb{K}$, we evaluate the partial derivatives at x instead.

Example 5.1.5. Consider the point $\mathfrak{m} = (x-y, x^2+1) \in \mathbb{A}^2_{\mathbb{R}}$. As discussed in Example 1.1.10 this point represents the Galois orbit $(x+i, y+i) \cup (x-i, y-i)$ of complex-valued points. Thus we can expect the Zariski cotangent space at $\mathfrak{m}$ to be the "union" of the cotangent spaces at these two points. Computing, we find

$$\begin{aligned} T^\vee_{\mathbb{A}^2,\mathfrak{m}} &= \frac{(x-y, x^2+1)}{(x^2-2y^2+y^2, x^3-x^2y+x-y, x^4+2x^2+1)} \\ &= \mathbb{R}(x-y) \oplus \mathbb{R}(x^2+1) \oplus \mathbb{R}(x^2-xy) \oplus \mathbb{R}(x^3+x) \end{aligned}$$

as an $\mathbb{R}$-vector space. Recall however that we should be thinking of the Zariski cotangent space as a vector space over the residue field $\mathbb{L} = \mathbb{R}[x,y]/(x-y, x^2+1)$. Identifying $\mathbb{L} = \mathbb{R} \oplus \mathbb{R}x$, we should instead write

$$T^\vee_{\mathbb{A}^2,\mathfrak{m}} = \mathbb{L}(x-y) \oplus \mathbb{L}(x^2+1)$$

5.1.1 Computing the Zariski cotangent space

The following results give a some general methods for calculating the Zariski cotangent space. Vakil calls the following result "Krull's PIT for the Zariski tangent space."

Exercise 5.1.6. Let $X = \mathrm{mSpec}(R)$ and let $\mathfrak{m}$ be a maximal ideal in R. Suppose that $f \in R$ vanishes at $\mathfrak{m}$ and set $Y = V(f)$. Show that $T_{Y,\mathfrak{m}}$ is the subspace of $T_{X,\mathfrak{m}}$ defined by the equation $f (\mathrm{mod}\, \mathfrak{m}^2) = 0$.

More generally, the Zariski tangent space of a $\mathbb{K}$-point in an affine $\mathbb{K}$-scheme can be computed using the Jacobian. Suppose that $X \subset \mathbb{A}^n$ is an affine scheme defined by the ideal $I = (f_1, \ldots, f_r)$. We define the Jacobian matrix

$$\mathrm{Jac}_{f_1,\ldots,f_r} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_r}{\partial x_1} & \cdots & \frac{\partial f_r}{\partial x_n} \end{bmatrix}.$$

Here we are thinking of the entries in Jac as elements of $\mathbb{K}[x_1, \ldots, x_n]$.

Fix a $\mathbb{K}$-point x in X. The inclusion $X \subset \mathbb{A}^n$ induces a surjective map of $\mathbb{K}$-vector spaces

$$\mathfrak{m}_{\mathbb{A}^n,x}/\mathfrak{m}^2_{\mathbb{A}^n,x} \to \mathfrak{m}_{X,x}/\mathfrak{m}^2_{X,x}.$$

The kernel of this map is spanned by the linear terms of the f_i. Dually, we see that:

Proposition 5.1.7. *Suppose that $X \subset \mathbb{A}^n$ is an affine scheme defined by the ideal $I = (f_1, \ldots, f_r)$. For any $x \in X$ with residue field $\mathbb{K}$ the Zariski tangent space $T_{X,x}$ is the kernel of* $\mathrm{Jac}_{f_1,\ldots,f_r}(x)$.

Warning 5.1.8. This computation only works for points with residue field $\mathbb{K}$. This might strike you as a little strange, and you would be right. In Section 5.3 we will give a different definition of "tangent space" for which the Jacobian construction works all of the time.

Example 5.1.9. Let $\mathbb{K} = \mathbb{F}_p(u)$ and let $X = \mathrm{mSpec}(\mathbb{K}[x, y]/(y^2 - x^p + u))$. We will show that the Zariski tangent space of the point $\mathfrak{m} = (y, x^p - u)$ is not computed by the Jacobian.

Note that the residue field of $\mathfrak{m}$ is $\mathbb{L} = \mathbb{F}_p(u^{1/p})$. As an $\mathbb{L}$-vector space the quotient $\mathfrak{m}/\mathfrak{m}^2$ is just $\mathbb{L}y$, hence one-dimensional. On the other hand, the Jacobian matrix is:

$$\mathrm{Jac}_f(x) = \left[\begin{array}{cc} 0 & 2y \end{array}\right].$$

When evaluated at the point $\mathfrak{m}$, the coordinate y vanishes and thus the Jacobian has rank 0 and its kernel has dimension 2.

5.1.2 Morphisms

Suppose given a morphism $f : X \to Y$ of quasiprojective schemes taking the point $x \in X$ to the point $y \in Y$. Recall that the residue field of y and the residue field of x may not be the same – we always have an inclusion $\mathcal{O}_{Y,y}/\mathfrak{m}_y \hookrightarrow \mathcal{O}_{X,x}/\mathfrak{m}_x$ but the map may not be an isomorphism. Using the pullback $f^\sharp$ we obtain a map

$$f^* : \mathfrak{m}_y/\mathfrak{m}_y^2 \to \mathfrak{m}_x/\mathfrak{m}_x^2$$

of $\mathcal{O}_{Y,y}/\mathfrak{m}_y$-vector spaces. However, if we are thinking of these sets as Zariski cotangent spaces then the field of definition need not be the same.

Note that $f^* : T^\vee_{Y,y} \to T^\vee_{X,x}$ only induces a map of Zariski tangent spaces $f_* : T_{X,x} \to T_{Y,y}$ when x and y have the same residue field. This is an indication that the Zariski tangent space is a bit unnatural. The alternative notion introduced in Section 5.3 has better behavior with respect to morphisms.

5.1.3 Exercises

Example 5.1.10. Suppose that p is a prime number and $q = p^r$. Let $\mathrm{mSpec}(R)$ be an affine $\mathbb{F}_q$-scheme. Recall from Example 1.5.14 that the Frobenius map $f : \mathrm{mSpec}(R) \to \mathrm{mSpec}(R)$ is induced by the ring homomorphism $f^\sharp(g) = g^q$. Prove that for every $x \in X$ the induced map of Zariski tangent spaces is the zero map.

Exercise 5.1.11. In this exercise we compute a "basis free" description of the tangent space of a $\mathbb{K}$-point in $\mathbb{P}^n$.

Let $x \in \mathbb{P}^n$ be a point with residue field $\mathbb{K}$. Consider the quotient morphism $f : \mathbb{A}^{n+1} \backslash \{0\} \to \mathbb{P}^n$ and let $\widetilde{x} \in \mathbb{A}^{n+1}$ be any $\mathbb{K}$-point whose image is x. Show that if we identify $T_{\mathbb{A}^{n+1},\widetilde{x}} \cong \mathbb{K}^{n+1}$ as in Example 5.1.4 then the kernel of the map $f_* : \mathbb{K}^{n+1} \to T_{\mathbb{P}^n,x}$ is the line $\ell \subset \mathbb{K}^{n+1}$ corresponding to x.

More generally, suppose $X = V_+(f_1, \ldots, f_r)$ in $\mathbb{P}^n$ and $x \in X$ is a $\mathbb{K}$-point. Let $\widetilde{X} = V(f_1, \ldots, f_r)$ be the cone over X in $\mathbb{A}^{n+1} \backslash \{0\}$ and let $\widetilde{x} \in \mathbb{A}^{n+1}$ be any $\mathbb{K}$-point whose image is x. Show that $T_{X,x}$ is the quotient of $T_{\widetilde{X},\widetilde{x}}$ by the line corresponding to x.

Exercise 5.1.12. Suppose that Y, Z are closed subschemes of the quasiprojective scheme X. Suppose $x \in Y \cap Z$.

(1) Prove that $T_{Y \cap Z,x} = T_{Y,x} \cap T_{Z,x}$ as subsets of $T_{X,x}$.

(2) Show that $T_{Y \cup Z,x} \supset \mathrm{Span}(T_{Y,x}, T_{Z,x})$. Give an example where the containment is strict.

Exercise 5.1.13. Let X be a quasiprojective $\mathbb{K}$-scheme. Fix a point $x \in X$ with residue field $\mathbb{K}$. Let $\mathcal{T}$ denote the set of homomorphisms $\mathrm{mSpec}(\mathbb{K}[t]/(t^2)) \to X$ whose set-theoretic image is x. For each homomorphism g, let $\overline{g}$ denote the map $\mathfrak{m}_x/\mathfrak{m}_x^2 \to \mathbb{K}$ arising from the map on stalks induced by g. Prove that this rule defines a bijection between $\mathcal{T}$ and the Zariski tangent space at x.

(This exercise illustrates how the Zariski tangent space can be thought of as maps from "infinitesimal curves" to $x \in X$. Note that if $X = \mathrm{mSpec}(R)$ is affine, then $\mathcal{T}$ is just the set of quotients $R \to \mathbb{K}[t]/(t^2)$ which contain $\mathfrak{m}$ in the kernel. Such a map must have the form $r \mapsto r(x) + a_r t$ where $r(x)$ denotes evaluation and the function a_r is essentially equivalent to the choice of a vector in the Zariski tangent space.)

Exercise 5.1.14. Let X be a quasiprojective $\mathbb{K}$-scheme and let $x \in X$ be a point with residue field $\mathbb{K}$. Let D (for "dual numbers") denote $\mathrm{mSpec}(\mathbb{K}[t]/(t^2))$. As in Exercise 5.1.13 let $\mathcal{T}$ denote the set of homomorphisms $D \to X$ whose set-theoretic image is x. Show that we can give $\mathcal{T}$ the structure of a $\mathbb{K}$-vector space using the following prescriptions:

(1) To add two morphisms $f, g : D \to X$, we take the composition

$$D \xrightarrow{\Delta} D \times D \xrightarrow{f \times g} X.$$

(Explicitly, when $X = \mathrm{mSpec}(R)$ this rule combines the two maps $r \mapsto r(x) + a_r t$ and $r \mapsto r(x) + b_r t$ to the map $r \mapsto r(x) + (a_r + b_r)t$, justifying calling this map addition.)

(2) To rescale a morphism $f : D \to X$ by a constant $a \in \mathbb{K}$, we precompose by the map $m_a : D \to D$ which sends $t \mapsto at$. (Explicitly, when $X = \mathrm{mSpec}(R)$ this rule changes the map $r \mapsto r(x) + b_r t$ to the map $r \mapsto r(x) + ab_r t$, justifying calling this map rescaling.)

Show that with this prescription the bijection of Exercise 5.1.13 is an isomorphism of $\mathbb{K}$-vector spaces.

Exercise 5.1.15. Let X, Y be quasiprojective $\mathbb{K}$-schemes. Let $w \in X \times Y$ be a $\mathbb{K}$-point and let x and y denote the two projections of w to X and Y respectively. Show that

$$T_w(X \times Y) \cong T_x X \oplus T_y Y.$$

(Hint: use Exercise 5.1.13.)

Exercise 5.1.16. Here is an interesting application of the Zariski tangent space. Consider the ideal (x, z) in the variety $X = \mathrm{mSpec}(\mathbb{K}[x, y, z]/(xy - z^2))$.

(1) Check that (x, z) defines a codimension 1 subvariety of X.

(2) Prove that (x, z) is not a principal ideal as follows. Note that $Z = V(x, z)$ contains the origin. If Z were principal, then by Exercise 5.1.6 $T_{Z,0}$ would have codimension 1 in $T_{X,0}$. Show that this is not the case.

5.2 Regularity

In this section we give our first definition of "smoothness" using the Zariski cotangent space.

5.2.1 Dimension and the Zariski cotangent space

The first step is to relate the dimension of a quasiprojective scheme X to the dimension of its Zariski cotangent space. Recall from Exercise 4.3.24 that the local dimension of a quasiprojective scheme X at a point x is the Krull dimension of $\mathcal{O}_{X,x}$, or equivalently, the largest dimension of any irreducible component of X containing x.

Theorem 5.2.1. *Let X be a quasiprojective scheme. For any $x \in X$ we have*

$$\dim_{\mathcal{O}_{X,x}/\mathfrak{m}_x} T^{\vee}_{X,x} \geq \operatorname{locdim}_x(X).$$

Proof. Suppose that $\overline{f}_1, \ldots, \overline{f}_r$ is a $\mathcal{O}_{X,x}/\mathfrak{m}_x$-basis for $\mathfrak{m}_x/\mathfrak{m}_x^2$. Let $f_1, \ldots, f_r$ be any lifts of these elements to $\mathfrak{m}_x$. By Nakayama's Lemma, we have $\mathfrak{m}_x = (f_1, \ldots, f_r)$. We have then reduced to the following lemma. □

Lemma 5.2.2. *Let R be Noetherian local ring with maximal ideal $\mathfrak{m}$. Suppose that $f_1, \ldots, f_r \in R$ are chosen so that $\mathfrak{m}$ is a minimal prime over $(f_1, \ldots, f_r)$. Then the Krull dimension of R is at most r.*

In particular this shows that Noetherian local rings always have finite Krull dimension (in contrast to arbitrary Noetherian rings). If we think of the generators of $\mathfrak{m}_x$ as "local coordinates" near a point x, this lemma is saying that the number of local coordinates gives an upper bound on the dimension of X near x. This lemma is a variant of Krull's Principal Ideal Theorem.

Proof. The proof is by induction on r. If $r = 1$ then we conclude by Krull's Prinicipal Ideal Theorem.

Suppose that $r > 1$. Let $\mathfrak{p}$ be any prime ideal such that there is no prime ideal $\mathfrak{q}$ satisfying $\mathfrak{p} \subsetneq \mathfrak{q} \subsetneq \mathfrak{m}$. Note that $\mathfrak{p}$ cannot contain every generator f_i; without loss of generality $f_1 \notin \mathfrak{p}$. Note that $\mathfrak{m}$ is the unique prime ideal that can contain $(f_1, \mathfrak{p})$. In other words, $\sqrt{(f_1, \mathfrak{p})} = \mathfrak{m}$, so there is some positive integer N such that $f_i^N \in (f_1, \mathfrak{p})$ for every i. We write $f_i^N = g_i + a_i f_1$ where $g_i \in \mathfrak{p}$, $a_i \in R$. Note that

$$V(f_1, g_2, \ldots, g_r) = V(f_1, f_2^N, \ldots, f_r^N) = V(\mathfrak{m}).$$

We claim that $\mathfrak{p}$ is minimal amongst all the prime ideals which contain $(g_2, \ldots, g_r)$. Note that in the ring $R/(g_2, \ldots, g_r)$ the quotient of $\mathfrak{m}$ is principal, generated by f_1. By Krull's PIT the ideal $\mathfrak{m}$ is codimension at most 1. Thus the quotient of $\mathfrak{p}$ must be a minimal prime ideal in $R/(g_2, \ldots, g_r)$.

By the induction assumption, the Krull dimension of $R_{\mathfrak{p}}$ is at most $r-1$. Thus the longest chain of primes descending from $\mathfrak{p}$ has length at most r, showing that the longest chain in R is $\leq r+1$. $\square$

5.2.2 Regularity

A local ring is said to be regular if its Krull dimension is the same as the minimal number of generators for its maximal ideal. (In other words, a local ring is regular if it achieves the equality in Lemma 5.2.2.) The following definition is the geometric version.

Definition 5.2.3. Let X be a quasiprojective scheme. We say that a point $x \in X$ is regular if $\mathcal{O}_{X,x}$ is a regular local ring, or equivalently, if

$$\dim T_{X,x}^{\vee} = \operatorname{locdim}_x(X).$$

We say that X is regular if it is regular at every point.

There are certain types of points that "look" nonregular, for example, non-reduced points or points that lie at the intersection of two components. The following lemmas confirm this intuition.

Lemma 5.2.4. *Let R be a regular local ring with maximal ideal $\mathfrak{m}$ with Krull dimension n. For any $f \in \mathfrak{m} \backslash \mathfrak{m}^2$ the quotient $R/(f)$ is a regular local ring with Krull dimension $n-1$.*

Conversely, if $f \in \mathfrak{m} \backslash \mathfrak{m}^2$ satisfies that $R/(f)$ is a regular local ring with Krull dimension $n-1$ then R is also regular.

Proof. By Exercise 5.1.6, the Zariski cotangent space of $R/(f)$ has dimension one less than the Zariski cotangent space of R. By Krull's Principal Ideal Theorem, the Krull dimension of $R/(f)$ is at most one less than the Krull dimension of R. The first statement then follows from Theorem 5.2.1.

Conversely, if $R/(f)$ is regular of dimension $n-1$ then in particular the Zariski cotangent space of $R/(f)$ has dimension $n-1$. By Exercise 5.1.6 we see that R has Zariski cotangent space of dimension n, and thus R is regular. $\square$

Proposition 5.2.5. *A regular local ring R is a domain.*

Proof. The proof is by induction on the Krull dimension n of R. In the base case R is Artinian and the claim is clear.

In general, suppose that $\mathfrak{p}$ is the smallest prime in a chain of prime ideals in R of maximal length. We would like to prove that $\mathfrak{p} = (0)$. Note that $R/\mathfrak{p}$ is still a local ring with Krull dimension n. Since the number of generators of the maximal ideal can only decrease, Theorem 5.2.1 shows that $R/\mathfrak{p}$ is a regular local ring.

Choose any element $f \in \mathfrak{m} \backslash \mathfrak{m}^2$. By Lemma 5.2.4 $R/(f)$ and $R/(\mathfrak{p} + (f))$ are both regular local rings of Krull dimension $n-1$. By induction these are both domains. But a surjective ring homomorphism from one domain to another which preserves the Krull dimension must be an isomorphism. We conclude that $(f) = \mathfrak{p} + (f)$.

This implies that $\mathfrak{p} \subset (f)$. In particular, every element $r \in \mathfrak{p}$ can be written as $r = af$ for some $a \in R$. Since $f \notin \mathfrak{p}$ (since $R/(\mathfrak{p} + (f))$ is not the same as $R/\mathfrak{p}$), this implies that $r = af$ for $a \in \mathfrak{p}$. In other words, $\mathfrak{p} \subset f\mathfrak{p}$.

It is clear that $f\mathfrak{p} \subset \mathfrak{p}$, so $\mathfrak{p} = f\mathfrak{p}$. By Nakayama's Lemma this implies that $\mathfrak{p} = 0$. □

Corollary 5.2.6. *Let X be a quasiprojective scheme. Suppose that $x \in X$ is a regular point. Then:*

- *x lies on a unique irreducible component of X, and*
- *x is a reduced point.*

Exercise 5.2.7. Prove the previous corollary.

5.2.3 Affine space

We will now compute one example: we will show that $\mathbb{A}^n$ is regular. In fact, our argument will prove something more general.

Proposition 5.2.8. *Let S be a regular local ring with maximal ideal $\mathfrak{n}$ and residue field $\mathbb{K}$. Set $R = S[x]$ and let $\mathfrak{p}$ be any prime ideal in R such that $\mathfrak{n}R \subset \mathfrak{p}$. Then $R_\mathfrak{p}$ is also a regular local ring.*

Proof. Set $d = \dim(S)$. Since S is a regular local ring, the ideal $\mathfrak{n}$ is generated by d elements $f_1, f_2, \ldots, f_d$.

There are two possibilities to consider. First suppose that $\mathfrak{n}R = \mathfrak{p}$. Since the elements $f_1, \ldots, f_d$ generate $\mathfrak{p}$ we see that it requires at most d generators. However, $\dim(R_\mathfrak{p})$ is at least d, since a chain of prime ideals of length d in S

$$0 = \mathfrak{q}_0 \subset \mathfrak{q}_1 \subset \ldots \subset \mathfrak{q}_d = \mathfrak{n}$$

induces the chain $\mathfrak{q}_i R_\mathfrak{p}$ in $R_\mathfrak{p}$. By Theorem 5.2.1 we conclude that $R_\mathfrak{p}$ is regular.

Next suppose that $\mathfrak{n}R \subsetneq \mathfrak{p}$. Since $R/\mathfrak{n}R \cong \mathbb{K}[x]$ we see that the image of $\mathfrak{p}$ in this ring is a maximal ideal generated by a polynomial $g(x)$. Thus $\mathfrak{p}$ is a maximal ideal in R and if $\overline{g}(x)$ denotes any lift of $g(x)$ then $f_1, \ldots, f_d, \overline{g}(x)$ is a generating set for $\mathfrak{p}$ of length $d + 1$. However, $\dim(R_\mathfrak{p})$ is at least $d + 1$, since a chain of prime ideals of length d in S

$$0 = \mathfrak{q}_0 \subset \mathfrak{q}_1 \subset \ldots \subset \mathfrak{q}_d = \mathfrak{n}$$

induces the chain $\mathfrak{q}_i R_\mathfrak{p}$ in $R_\mathfrak{p}$ with an additional $\mathfrak{p}$ at the end. By Theorem 5.2.1 we conclude that $R_\mathfrak{p}$ is regular. □

Exercise 5.2.9. Using Proposition 5.2.8, show that if X is a $\mathbb{K}$-variety with a regular point x then every point of $X \times \mathbb{A}^1$ that maps to x is regular. In particular, show that $\mathbb{A}^n$ is regular.

5.2.4 Exercises

Exercise 5.2.10. Let $X \subset \mathbb{A}^n$ be a closed subscheme defined by the ideal $(f_1, \ldots, f_r)$. Let $x \in X$ be a $\mathbb{K}$-point. Show that x is a regular point of X if and only if the rank of the affine Jacobian $(r \times n)$-matrix $\mathrm{Jac}_{f_1,\ldots,f_r}(x)$ is equal to $n - \dim_x(X)$.

Exercise 5.2.11. Let $X \subset \mathbb{P}^n$ be a closed subscheme defined by a homogeneous ideal $(f_1, \ldots, f_r)$. Let $x \in X$ be a $\mathbb{K}$-point. Show that x is a regular point of X if and only if the rank of the projective Jacobian $(r \times (n+1))$-matrix $\mathrm{Jac}_{f_1,\ldots,f_r}(x)$ is equal to $n - \dim_x(X)$.

Note that the partial derivatives of the projective Jacobian matrix are not well-defined functions on the points of X – however, the rank of the matrix is well-defined. This condition is called the Projective Jacobian Criterion.

5.3 Relative tangent spaces and smoothness

Let X be a quasiprojective $\mathbb{K}$-scheme and fix a finite extension $\mathbb{L}/\mathbb{K}$. Earlier we saw that the set of points in X with residue field $\mathbb{L}$ might be a bit complicated (see Example 1.1.3.(3) and Exercise 1.1.13). However, if we work with the space of $\mathbb{L}$-points – that is, the maps $\mathrm{mSpec}(\mathbb{L}) \to X$ – the structure is a lot nicer and more intuitive (see Exercise 1.5.11 and the discussion after Exercise 1.7.8).

Today we will give our second construction of a tangent space, known as the "relative tangent space." The relationship between this construction and the Zariski tangent space is closely analogous to the relationship between $\mathbb{L}$-points and points with a fixed residue field $\mathbb{L}$. In particular, the relative tangent space is more "functorial" than the Zariski tangent space and thus exhibits many desirable features that the Zariski tangent space lacks.

5.3.1 Relative tangent space

Definition 5.3.1. Let X be a quasiprojective $\mathbb{K}$-scheme and let $\mathbb{L}$ be a finite extension of $\mathbb{K}$. Consider a $\mathbb{L}$-valued point $\sigma : \mathrm{mSpec}(\mathbb{L}) \to X$. The relative tangent space of σ over $\mathbb{K}$, denoted by $T_{X/\mathbb{K},\sigma}$ is the set of morphisms $\mathrm{mSpec}(\mathbb{L}[t]/(t^2)) \to X$ such that the composition with the inclusion $\mathrm{mSpec}(\mathbb{L}) \to \mathrm{mSpec}(\mathbb{L}[t]/(t^2))$ is σ.

Note that the image of σ need not have residue field $\mathbb{L}$. In the case when the image of $\sigma : \mathrm{mSpec}(\mathbb{L}) \to X$ is a point x with residue field $\mathbb{L}$, we will denote the relative tangent space by $T_{X/\mathbb{K},x}$.

Exercise 5.3.2. Prove that we can equip $T_{X/\mathbb{K},\sigma}$ with the structure of a $\mathbb{L}$-vector space using the prescription of Exercise 5.1.14.

When σ is a $\mathbb{K}$-point the Zariski tangent space and the relative tangent space are isomorphic; this follows from Exercise 5.1.13. However, when the residue field is larger than $\mathbb{K}$ it is possible for the two spaces to differ (see Example 5.3.11). It turns out that the relative tangent space is much better behaved than the Zariski tangent space. For example, given any morphism $f : X \to Y$ and any $\mathbb{L}$-point σ we obtain a pushforward map $f_* : T_{X/\mathbb{K},\sigma} \to T_{Y/\mathbb{K},f\circ\sigma}$ by composing with f. One of the most useful properties of the relative tangent space is that it is invariant under base change.

Proposition 5.3.3. *Let X be a quasiprojective $\mathbb{K}$-scheme, let $\mathbb{L}$ be a finite extension of $\mathbb{K}$, and let $\sigma : \mathrm{mSpec}(\mathbb{L}) \to X$ be a $\mathbb{L}$-point. Suppose that $\mathbb{F}$ is a field satisfying $\mathbb{K} \subset \mathbb{F} \subset \mathbb{L}$. Then $T_{X/\mathbb{K},\sigma}$ is isomorphic to $T_{X_{\mathbb{F}}/\mathbb{F},\sigma_{\mathbb{F}}}$.*

Here $\sigma_{\mathbb{F}}$ denotes the point $\sigma \times i : \mathrm{mSpec}(\mathbb{L}) \to X \times \mathrm{mSpec}(\mathbb{F})$ where $i^\sharp : \mathbb{F} \to \mathbb{L}$ is the given inclusion.

Proof. Note that there exists a unique $\mathbb{K}$-morphism $\mathrm{mSpec}(\mathbb{L}[t]/(t^2)) \to \mathrm{mSpec}(\mathbb{F})$. Thus, by the universal property of the product there is a bijection between $T_{X/\mathbb{K},\sigma}$ and the set of

$\mathbb{K}$-morphisms $\mathrm{mSpec}(\mathbb{L}[t]/(t^2)) \to X \times_{\mathrm{mSpec}(\mathbb{K})} \mathrm{mSpec}(\mathbb{F})$ which yield $\sigma_{\mathbb{F}}$ upon composition with $\mathrm{mSpec}(\mathbb{L}) \to \mathrm{mSpec}(\mathbb{L}[t]/(t^2))$. But every such $\mathbb{K}$-morphism is a $\mathbb{F}$-morphism (and conversely any $\mathbb{F}$-morphism is also a $\mathbb{K}$-morphism). Thus such morphisms are in bijection with $T_{X_{\mathbb{F}}/\mathbb{F},\sigma_{\mathbb{F}}}$. □

By combining this with the fact that the relative tangent space and Zariski tangent space agree for points over the ground field, we see that the relative tangent space is the same as the Zariski tangent space after a base change.

Proposition 5.3.4. *Let X be a quasiprojective $\mathbb{K}$-scheme, let $\mathbb{L}$ be a finite extension of $\mathbb{K}$, and let $\sigma : \mathrm{mSpec}(\mathbb{L}) \to X$ be a $\mathbb{L}$-point. Then $T_{X/\mathbb{K},\sigma}$ is isomorphic to the Zariski tangent space of $X_{\mathbb{L}}$ at the $\mathbb{L}$-point $\sigma_{\mathbb{L}}$.*

As a corollary, we can use the Jacobian to describe the relative tangent space.

Corollary 5.3.5. *Let X be an affine $\mathbb{K}$-scheme, let $\mathbb{L}$ be a finite extension of $\mathbb{K}$, and let $\sigma : \mathrm{mSpec}(\mathbb{L}) \to X$ be a $\mathbb{L}$-point with image x. Suppose that X is a closed subscheme of $\mathbb{A}^n$ defined by the ideal $(f_1, \ldots, f_r)$. Then $T_{X/\mathbb{K},\sigma}$ is the tensor product of $\mathbb{L}$ with the $\kappa(x)$-vector space which is the kernel of the matrix*

$$\mathrm{Jac}_{f_1,\ldots,f_r}(x) = \begin{bmatrix} \frac{\partial f_1}{\partial x_1}(x) & \cdots & \frac{\partial f_1}{\partial x_n}(x) \\ \vdots & \ddots & \vdots \\ \frac{\partial f_r}{\partial x_1}(x) & \cdots & \frac{\partial f_r}{\partial x_n}(x) \end{bmatrix}.$$

This yields the following important property of the relative tangent space.

Proposition 5.3.6. *Let X be a quasiprojective $\mathbb{K}$-scheme. Then the dimension of the relative tangent space is upper semicontinuous: for any positive integer d, the subset*

$$Z_d = \{x \in X \mid \dim_{\kappa(x)}(T_{X/\mathbb{K},x}) \geq d\}$$

is a closed set in X.

Proof. It suffices to prove this result when $X = \mathrm{mSpec}(R)$ is affine. Choose an embedding $X \hookrightarrow \mathbb{A}^n$ and write $(f_1, \ldots, f_r)$ for its defining ideal. Then Z_d is the vanishing locus of the ideal generated by the $(n-d+1)$-minors of $\mathrm{Jac}_{f_1,\ldots,f_r}$. □

By Proposition 5.3.4 and Theorem 5.2.1 the relative tangent space of $\sigma : \mathbb{L} \to X$ always has $\mathbb{L}$-dimension at least as large as the local dimension $\mathrm{locdim}_{\sigma(\mathrm{mSpec}(\mathbb{L}))}(X)$.

5.3.2 Smoothness

Definition 5.3.7. Let X be a quasiprojective $\mathbb{K}$-scheme. Let $\mathbb{L}/\mathbb{K}$ be a finite extension. We say that a $\mathbb{L}$-point σ of X is smooth if the $\mathbb{L}$-dimension of $T_{X/\mathbb{K},\sigma}$ is the local dimension $\operatorname{locdim}_{\sigma(\mathrm{mSpec}(\mathbb{L}))}(X)$. If σ is not smooth, we call it a singular $\mathbb{L}$-point.

When the image x of $\sigma : \mathrm{mSpec}(\mathbb{L}) \to X$ has residue field $\mathbb{L}$, we simply say that x is a smooth or a singular point. The smooth locus of X is the union of the smooth points in X, and the singular locus is its complement. We say that X is smooth if every point is smooth.

A direct consequence of Proposition 5.3.6 is that the singular locus is closed.

Lemma 5.3.8. *Let X be a quasiprojective $\mathbb{K}$-scheme. The smooth locus of X is a (possibly empty) open subset of X.*

Remark 5.3.9. In the setting of affine $\mathbb{K}$-schemes, it is also true that the nonregular locus is closed. However the proof takes some work – the statement does not hold for arbitrary rings.

We have already seen that regularity and smoothness coincide for $\mathbb{K}$-points. The following theorem (which we do not prove) identifies other situations where the two coincide.

Theorem 5.3.10. *Let X be a quasiprojective scheme.*

- *If $x \in X$ is a smooth point, then it is a regular point.*
- *Suppose that $x \in X$ is a point with perfect residue field. Then x is regular if and only if it is a smooth point.*

In particular, if x is a smooth point then the local ring $\mathcal{O}_{X,x}$ is a domain. Arguing as in Corollary 5.2.6, we see that any smooth point x is reduced and is contained in a unique irreducible component of X. The following example shows that regular and smooth can be different for points x whose residue field is non-perfect and is different from our ground field.

Example 5.3.11. Let $\mathbb{K} = \mathbb{F}_p(u)$ and let $X = \mathrm{mSpec}(\mathbb{K}[x, y]/(y^2 - x^p + u))$. By Krull's PIT this affine scheme has dimension 1. We will show that the point $\mathfrak{m} = (y, x^p - u)$ is regular but not smooth. Note that its residue field is $\mathbb{L} = \mathbb{F}_p(u^{1/p})$. To see this point is regular, note that as an $\mathbb{L}$-vector space the quotient $\mathfrak{m}/\mathfrak{m}^2$ is just $\mathbb{L}y$, hence one-dimensional. To this point is not smooth, we compute the Jacobian:

$$\mathrm{Jac}_f(x) = \begin{bmatrix} 0 & 2y \end{bmatrix}.$$

When evaluated at the point $\mathfrak{m}$, the coordinate y vanishes and thus the Jacobian has rank 0 and its kernel has dimension 2. This shows that $\mathfrak{m}$ is not a smooth point.

5.3.3 Exercises

Exercise 5.3.12. Find all the singular points of the following plane curves in $\mathbb{A}^2_{\mathbb{C}}$.

(1) $y^3 - y^2 + x^3 - x^2 + 3xy^2 + 3x^2y + 2xy$.

(2) $x^4 + y^4 - x^2y^2$.

(3) $x^3 + y^3 - 3x^2 - 3y^2 + 3xy + 1$.

Exercise 5.3.13. Let $\mathbb{K}$ be a field of characteristic 0. For any positive integers d and n show that the Fermat hypersurface of degree d defined by the equation $x_0^d + \ldots + x_n^d = 0$ is smooth in $\mathbb{P}^n$. (What can go wrong in characteristic p?)

Exercise 5.3.14. Let $\mathbb{K}$ be an algebraically closed field of characteristic $\neq 2$. We consider the geometry of the quadric $X = V_+(\sum_{i=0}^r x_i^2)$ in $\mathbb{P}^n$.

(1) Show that the singular points of X are given by the linear space $L \subset X$ defined by the ideal $(x_0, \ldots, x_r)$.

(2) Let $\phi : \mathrm{Bl}_L \mathbb{P}^n \to \mathbb{P}^n$ be the blow-up of $\mathbb{L}$. Show that the strict transform of X in $\mathrm{Bl}_L \mathbb{P}^n$ is smooth.

If we let $\widetilde{X}$ denote the strict transform of X, then the map $\phi : \widetilde{X} \to X$ is a birational morphism such that $\widetilde{X}$ is smooth and ϕ is an isomorphism over the smooth locus of X. Such a map is called a resolution of singularities.

Exercise 5.3.15. Let X be a quasiprojective $\mathbb{K}$-scheme, and consider a chain of finite extensions $\mathbb{K} \subset \mathbb{F} \subset \mathbb{L}$. Let $\sigma : \mathrm{mSpec}(\mathbb{F}) \to X$ be a $\mathbb{F}$-point and let $\widetilde{\sigma} : \mathrm{mSpec}(\mathbb{L}) \to X$ denote the composition with the natural map $\mathrm{mSpec}(\mathbb{L}) \to \mathrm{mSpec}(\mathbb{F})$. Show that $T_{X/\mathbb{K},\sigma} \otimes \mathbb{L}$ is isomorphic to $T_{X/\mathbb{K},\widetilde{\sigma}}$.

Exercise 5.3.16. Let $\mathbb{L}/\mathbb{K}$ be a finite field extension and let $\mathbb{F}/\mathbb{L}$ be an arbitrary field extension. Let $\sigma : \mathrm{mSpec}(\mathbb{L}) \to X$ be an $\mathbb{L}$-point and let $\sigma' : \mathrm{mSpec}(\mathbb{F}) \to X_{\mathbb{F}}$ be the base change of σ to $\mathbb{F}$. Prove that σ is a smooth point if and only if σ' is a smooth point.

Exercise 5.3.17. Let $\mathbb{K}$ be an algebraically closed field. Suppose that $X \subset \mathbb{P}^n$ is a closed subscheme. Fix a point $p \in X$. Suppose that ℓ is a line through p that is not contained in X. We say that ℓ is a tangent line to X at p if $T_{\ell,p}$ is contained in $T_{X,p}$ (as subspaces of $T_{\mathbb{P}^n,p}$).

Prove that ℓ is a tangent line if and only if the intersection $X \cap \ell$ is non-reduced at p.

5.4 Geometric properties

In this section we address two key properties of smoothness.

5.4.1 Generic smoothness

Theorem 5.4.1. *Suppose $\mathbb{K}$ is a perfect field. Any quasiprojective $\mathbb{K}$-variety X of dimension n is birational to a hypersurface $V(f) \subset \mathbb{A}^{n+1}$ defined by an irreducible polynomial f.*

Here we will give an algebraic argument.

Proof. Let $\mathfrak{K}$ be the function field of X. Since $\mathbb{K}$ is perfect, $\mathfrak{K}/\mathbb{K}$ is finitely separably generated. By Noether Normalization, there are algebraically independent elements $t_1, \dots, t_n$ such that $\mathfrak{K}$ is a finite separable extension of $\mathbb{K}(t_1, \dots, t_n)$. By the primitive element theorem, there is an element θ such that $\mathfrak{K} = \mathbb{K}(t_1, \dots, t_n, \theta)$. We know that θ satisfies an irreducible polynomial equation whose coefficients are rational functions in the t_i. Clear denominators to get a polynomial $f \in \mathbb{K}[t_1, \dots, t_n, \theta]$. This polynomial defines a hypersurface in $\mathbb{A}^{n+1}$ with the same function field as X. □

Theorem 5.4.2. *Let X be a geometrically integral quasiprojective $\mathbb{K}$-variety of dimension n. There is a non-empty open subset $U \subset X$ such that every point $x \in U$ is a smooth point.*

The statement can fail if X is not geometrically integral; consider for example the $\mathbb{F}_p(t)$-variety $\mathrm{mSpec}(\mathbb{F}_p(t)[x]/(x^p - t))$.

Proof. By Exercise 5.3.16 smoothness can be detected after base-changing to an algebraically closed field, so we may suppose that $\mathbb{K}$ is algebraically closed. Then by Theorem 5.4.1 it suffices to prove the statement when $X = V(f)$ is a hypersurface in $\mathbb{A}^{n+1}$ defined by an irreducible polynomial f. The set of singular points is defined by the equations $f = 0$ and $\frac{\partial f}{\partial x_i} = 0$ for $i = 1, \dots, n+1$. If $X_{sing} = X$, then each $\frac{\partial f}{\partial x_i}$ is contained in the ideal (f). Since taking derivatives drops degree, this can only happen if all the partial derivatives of f are identically 0.

When $\mathbb{K}$ has characteristic 0 this is impossible. When $\mathbb{K}$ has characteristic p, we see that every term of f has exponents which are all divisible by p. Since all the coefficients are also pth powers (as $\mathbb{K}$ is algebraically closed), we see that $f = g^p$ for some polynomial g. But this contradicts the irreducibility of f. □

5.4.2 Bertini's Theorem

Suppose $X \subset \mathbb{P}^n$ is a projective variety. A hyperplane section of X is the intersection $X \cap H$ for some $\mathbb{K}$-hyperplane H. An important principle in algebraic geometry is that when H is general (in the sense of Definition 4.5.3) the intersection $X \cap H$ will inherit many of the nice properties of X. Our first example of this principle is Bertini's Theorem.

Theorem 5.4.3. *Suppose that $X \subset \mathbb{P}^n$ is a smooth projective variety. Consider the parameter space $\mathbb{P}^{n\vee}$ for hyperplanes in $\mathbb{P}^n$. There is some non-empty open subset $V \subset \mathbb{P}^{n\vee}$ such that every hyperplane H parametrized by a $\mathbb{K}$-point in V the intersection $X \cap H$ is smooth.*

Proof. By Exercise 5.3.16 smoothness can be detected after base changing to the algebraic closure of $\mathbb{K}$. Thus we may assume that our ground field is algebraically closed.

Let $U \subset (\mathbb{P}^n)^\vee$ denote the open subset parametrizing hyperplanes which do not contain X. Consider the incidence correspondence $I \subset U \times X$ consisting of pairs (H, x) such that x is a singular point of $X \cap H$. Our first goal is to show that I is a closed subset of $U \times X$.

Fix a point $x \in X$. We can choose an open affine $\mathbb{A}^n \subset \mathbb{P}^n$ which has x as the origin and let $\mathbb{K}[y_1, \ldots, y_n]$ denote its coordinate ring and let $I = (f_1, \ldots, f_r)$ denote the ideal of $X \cap \mathbb{A}^n$. If we fix a hyperplane containing x then $\mathbb{A}^n \cap H$ is the vanishing locus of a homogeneous linear function ℓ. Krull's Principal Ideal Theorem shows that the local dimension of x in $X \cap H$ is one less than the dimension of X. By the Jacobian criterion $X \cap H$ will fail to be smooth at x precisely when the matrix

$$\begin{bmatrix} \frac{\partial f_1}{\partial x_1}(x) & \ldots & \frac{\partial f_1}{\partial x_n}(x) \\ \vdots & \ddots & \vdots \\ \frac{\partial f_r}{\partial x_1}(x) & \ldots & \frac{\partial f_r}{\partial x_n}(x) \\ \frac{\partial \ell}{\partial x_1}(x) & \ldots & \frac{\partial \ell}{\partial x_n}(x) \end{bmatrix}.$$

has rank $\dim(X)$. This will happen when the last row is in the span of the first r rows. Note that the last row consists of the coefficients of ℓ; thus, as we vary x the vanishing of the corresponding $(\dim(X)+1)$-minors of this matrix will be described by equations in the homogeneous coordinates on $U \times X$.

We next show that the projection map $I \to U$ is not dominant. It suffices to prove that $\dim(I) < \dim(U)$. The argument above shows that the hyperplanes such that $X \cap H$ is singular at x must contain a fixed $\dim(X)$-subplane in the projectivized tangent space of $\mathbb{P}^n$ at x. In other words, in the parameter space $\mathbb{P}^{n-1}$ of hyperplanes H that contain x, the subset whose intersection with X is singular at x will lie in a $(n-1-\dim(X))$-dimensional subplane. By Theorem 4.4.9 we see that I has dimension $\leq (n-1-\dim(X)) + \dim(X) = n-1$, finishing the proof. $\square$

Remark 5.4.4. If $\mathbb{K}$ is an infinite field then any open subset of $\mathbb{P}^{n\vee}$ will contain a $\mathbb{K}$-point. Thus there will always be a hyperplane for which Bertini Theorem's applies.

If $\mathbb{K}$ is a finite field this is no longer true. [[Kat99] Question 10] gives the example of the hypersurface $\sum_{i=0}^n x_i y_i^q - x_i^q y_i = 0$ in $\mathbb{P}^{2n+1}_{\mathbb{F}_q}$. The intersection of this hypersurface with every $\mathbb{F}_q$-hyperplane is singular.

5.4.3 Exercises

Exercise 5.4.5. Suppose that $\mathbb{K}$ is an infinite field. Use Bertini's theorem to show that a general degree d hypersurface in $\mathbb{P}^n$ is smooth. More generally, prove that if we choose general homogeneous polynomials $f_1, \ldots, f_r \in \mathbb{K}[x_0, \ldots, x_n]$ with $r \leq n$ then the subscheme $X = V_+(f_1, \ldots, f_r)$ is smooth.

Exercise 5.4.6. Suppose that $\mathbb{K}$ is algebraically closed. Suppose that $X \subset \mathbb{P}^n$ is a projective variety of dimension d. Show that the intersection of X with d general hyperplanes will be a finite set of reduced points.

(More precisely, let $M \cong \mathbb{P}^n$ denote the parameter space of hyperplanes on $\mathbb{P}^n$. We can think of $M^{\times d}$ as the parameter space of sets of d hyperplanes. Show that there is an open subset $U \subset M^{\times d}$ such that for every point $p \in U$ the intersection of M with the d hyperplanes corresponding to p will be a finite set of reduced points.)

5.5 Tangent cones and blow-ups

In this section we will define the tangent cone construction, a modification of the notion of a tangent space. For the sake of simplicity we will focus only on $\mathbb{K}$-points. We will need the following definition:

Definition 5.5.1. Let $f \in \mathbb{K}[x_1, \dots, x_n]$. Let d be the smallest degree of any term of f. We define f_{min} to be the sum of all the terms of f which have degree d.

Let $I \subset \mathbb{K}[x_1, \dots, x_n]$ be an ideal (not necessarily homogeneous). The initial ideal I_{min} is the ideal generated by f_{min} as we vary f over all elements in I.

It is important to note that I_{min} need not be the same as the ideal generated by the minimal degree parts of the generators of I. For example, if $I = (x + y^2, xy)$ then $I_{min} = (x, y^3)$.

Definition 5.5.2. Let $X \subset \mathbb{A}^n$ be an affine scheme and let $x \in X$ be a $\mathbb{K}$-point. Without loss of generality we may suppose that x is the origin in $\mathbb{A}^n$. Let $I \subset \mathbb{K}[x_1, \dots, x_n]$ be the ideal of X. The tangent cone of x in X is the vanishing locus of the ideal consisting of the lowest degree homogeneous part of I.

Exercise 5.5.3. Sketch the following curves in $\mathbb{A}^2$ and their tangent cones at the origin. How do the tangent cones compare to the graph of the variety?

(1) $y^2 = x^3 + x^2$.

(2) $y^2 = x^3$.

(3) $(x^2 + y^2)^2 + 3x^2y - y^3 = 0$.

As illustrated by the previous exercise, the tangent cone is the "limit" of the tangent lines of all nearby points. Note that it is possible for this limit to have many components of different dimensions.

Exercise 5.5.4. Let $X, Y \subset \mathbb{A}^n$ be closed subschemes. Show that the tangent cone of $X \cup Y$ at a point x is the union of the tangent cones of X and Y at x.

Remark 5.5.5. We can also think of the tangent cone as an ideal in $\mathcal{O}_{X,x}$ as follows. Consider the finitely generated graded $\mathbb{K}$-algebra

$$S = \mathcal{O}_{X,x}/\mathfrak{m}_x \oplus \mathfrak{m}_x/\mathfrak{m}_x^2 \oplus \mathfrak{m}_x^2/\mathfrak{m}_x^3 \oplus \dots$$

Note that we have a surjection $\mathrm{Sym}(\mathfrak{m}/\mathfrak{m}^2) \to S$. This yields a closed subscheme $\mathrm{mSpec}(S) \subset \mathrm{mSpec}(\mathrm{Sym}(\mathfrak{m}/\mathfrak{m}^2)) \cong T_{X,x}$ which is the (intrinsic) tangent cone of X. By composing with the natural map $T_{X,x} \to T_{\mathbb{A}^n,x}$ we obtain the tangent cone as defined earlier.

With this definition the tangent cone is *intrinsic* to X, i.e. does not depend on a choice of embedding. Thus the definition naturally extends to all quasiprojective varieties. We will not pursue this line of reasoning here, being content to work with embedded varieties.

5.5.1 Tangent cones and blowing-up a point

Tangent cones provide new insight into the blow-up construction.

Theorem 5.5.6. *Let $X \subset \mathbb{A}^n$ be a closed subscheme. Let $x \in X$ be a $\mathbb{K}$-point. Let $\phi : \mathrm{Bl}_x \mathbb{A}^n \to \mathbb{A}^n$ be the blow-up and let $\phi_X : \widetilde{X} \to X$ be the strict transform of X. The fiber of ϕ_X over x is isomorphic to the* mProj *of the ideal defining the tangent cone as a subscheme of the exceptional divisor $E \cong \mathbb{P}^{n-1}$ of ϕ.*

Recall that ϕ_X can equally well be thought of as the blow-up of the ideal defining the point x. Thus one can use Theorem 5.5.6 to show that the tangent cone is intrinsic to X, independent of the choice of an embedding.

Proof. Without loss of generality we may assume x is the origin. Recall that the blow-up of $\mathbb{A}^n$ at the origin (with coordinates $x_0, \dots, x_{n-1}$) is defined by the ideal

$$(x_i y_j - x_j y_i)_{i,j=0}^{n-1}$$

inside of $\mathbb{A}^n \times \mathbb{P}^{n-1}$. Suppose that $X = V(I)$. We can compute the blow-up $\widetilde{X}$ by taking the preimage of X and then removing the "extra component" represented by the exceptional divisor.

Consider the open subset $U_0 \subset \mathrm{Bl}_x \mathbb{A}^n$ where $y_0 \neq 0$. This open set is isomorphic to n-dimensional affine space in the coordinates $x_0, \frac{y_1}{y_0}, \dots, \frac{y_{n-1}}{y_0}$. On this chart, the preimage of X is defined by $f\left(x_0, x_0\frac{y_1}{y_0}, \dots, x_0\frac{y_1}{y_0}\right)$ as we vary $f \in I$. Since the exceptional divisor is cut out by the equation $x_0 = 0$, the strict transform of X is defined by the set of equations

$$\frac{f\left(x_0, x_0\frac{y_1}{y_0}, \dots, x_0\frac{y_1}{y_0}\right)}{x_0^{d_{min}}}$$

where d_{min} is the minimal degree of any term in f.

Consider now the intersection of the strict transform of X with the exceptional divisor E in this open subset. The intersection is given by setting $x_0 = 0$. For each function f, the only remaining terms will be those whose degree is d_{min} – more precisely, they will be the terms of minimal degree which do not involve x_0 at all. Varying the charts, we see that $\widetilde{X} \cap E$ is generated by the projectivization of I_{min}. □

Theorem 5.5.6 illustrates how one should think about the blow-up at a $\mathbb{K}$-point x: the operation replaces x with the projectivization of its tangent cone. In particular, when we blow-up $\mathbb{A}^n$ at the origin the exceptional divisor E can be identified with the $\mathbb{P}^{n-1}$ parametrizing tangent directions at 0. When we take the strict transform $\widetilde{X}$ of $X \subset \mathbb{A}^n$, the intersection of $\widetilde{X}$ with the exceptional divisor E is the set of "limit" tangent lines as we approach 0 from directions in X.

5.5.2 Tangent cones and general blow-ups

More generally, suppose that $X = \mathrm{mSpec}(R)$ is an affine scheme and I is an ideal in R. Then a fiber of the blow-up map $\phi_X : Bl_I(X) \to X$ over a point in $V(I)$ will represent the directions in the tangent cone of X that are not tangent to $V(I)$. More precisely, the construction replaces $V(I)$ by the projectivization of its normal sheaf – since we haven't discussed vector bundles on schemes, we won't attempt to make this precise.

Example 5.5.7. We use coordinates x_0, x_1, x_2 on $\mathbb{A}^3$. Suppose that $L \subset \mathbb{A}^3$ is the linear subspace defined by the equations $x_0 = x_1 = 0$. By taking a preimage in the blow-up, we see that the exceptional divisor of $\phi : \mathrm{Bl}_L \mathbb{A}^3 \to \mathbb{A}^3$ is isomorphic to $\mathrm{mProj}(\mathbb{K}[y_0, y_1, x_2])$ where we only give positive weighting to the first two variables. In particular, the fiber over the origin is given by setting $x_2 = 0$.

Let $X \subset \mathbb{A}^3$ be a closed subscheme containing the origin. We let $\widetilde{X}$ denote the strict transform of X under $\phi : \mathrm{Bl}_L \mathbb{A}^3 \to \mathbb{A}^3$. Arguing as in Theorem 5.5.6, we see that on the chart $y_0 \neq 0$ the intersection with $\widetilde{X}$ is defined by varying the functions

$$\frac{f(x_0, x_0\frac{y_1}{y_0}, x_2)}{x_0^{d_{min}}}$$

where d_{min} is the minimal degree of any term of f in the x_0, x_1 variables. A similar description applies to the chart $y_1 \neq 0$. Thus we see that the intersection of the strict transform of $\widetilde{X}$ with the exceptional divisor E is defined by the projectivization of the initial ideal where the weight is purely in the first two variables.

Next consider the fiber of $\phi|_{\widetilde{X}}$ over the origin. On the chart $y_0 \neq 0$ this fiber is given by setting $x_0 = x_2 = 0$. The only remaining functions will be the terms of minimal degree which only involve x_1. Arguing similarly on the other chart, we see that this fiber of $\phi|_{\widetilde{X}}$ is defined by quotienting the initial ideal by the variable x_2. In other words, it is the projectivization of the quotient of the tangent cone of X by the tangent direction of L.

5.5.3 Exercises

Exercise 5.5.8. An irreducible plane curve is said to have a node at a point if its tangent cone is the union of two different lines. It is said to have a cusp at a point if its tangent cone is a double line.

Prove that an irreducible cubic curve in $\mathbb{P}^2$ has at most one singular point and this singular point is either a node or a cusp.

Exercise 5.5.9. For each of the following plane curves in $\mathbb{P}^2_{\mathbb{C}}$, find the singular points and the tangent cones at each singular point.

(1) $xy^4 + yz^4 + zx^4$.

(2) $x^2y^3 + x^2z^3 + y^2z^3$.

(3) $y^2z - x(x-z)(x-\lambda z)$ for some $\lambda \in \mathbb{C}$.

(4) $x^n + y^n + z^n$.

5.6 Normality

We next turn to a weaker variant of regularity that is motivated by algebra. Suppose that $X = \mathrm{mSpec}(R)$ is an affine variety. We are interested in the situation when R is integrally closed in $\mathrm{Frac}(R)$. Since this condition is a stalk-local property when R is a domain, we can define it for quasiprojective varieties as well.

Lemma 5.6.1. *Let X be a quasiprojective variety. The following conditions are equivalent:*

(1) There is an open cover $\{U_i\}$ of X by open affines such that every ring $\mathcal{O}_X(U_i)$ is integrally closed in $\mathbb{K}(X)$.

(2) For every open affine $U \subset X$ we have that $\mathcal{O}_X(U)$ is integrally closed in $\mathbb{K}(X)$.

(3) For every point $x \in X$ the stalk $\mathcal{O}_{X,x}$ is integrally closed in $\mathbb{K}(X)$.

Proof. This follows from the fact that the integral closure condition for domains can be checked by localizing at maximal ideals combined with Proposition 2.5.6. □

Definition 5.6.2. Let X be a quasiprojective variety. If X satisfies the equivalent conditions of Lemma 5.6.1 we say that X is a normal variety.

We will only work with normality for varieties (although one can define it for arbitrary schemes by requiring that all local rings be integrally closed).

Example 5.6.3. Every regular quasiprojective variety is normal. In Proposition 5.2.5 we showed that a regular local ring is a domain, and it only remains to show that a regular local ring is integrally closed in its fraction field. This is a consequence of the Auslander-Buchsbaum Theorem which shows that a regular local ring will be a UFD (and thus integrally closed).

Example 5.6.4. Not every normal quasiprojective variety is regular. For example, set $R = \mathbb{K}[x, y, z]/(xy - z^2)$ and consider the quadric cone $X = \mathrm{mSpec}(R)$. Then X is not regular at the origin.

However, X is normal if $\mathrm{char}(\mathbb{K}) \neq 2$. To see this, note that $\mathbb{K}(X) = \mathbb{K}(x, y)[z]/(z^2 - xy)$ is a degree 2 extension of $\mathbb{K}(x, y)$. Suppose that $\alpha = g + hz$ is an element of $\mathbb{K}(X)$ where $g, h \in \mathbb{K}(x, y)$. The minimal polynomial of this element is

$$t^2 - 2gt + (g^2 - xyh^2)$$

If α is integral over R, then $g \in R$ and $g^2 - xyh^2 \in R$ so that $h \in R$ as well. This means that g, h are contained in $\mathbb{K}[x, y]$. Thus $\alpha \in R$.

5.6.1 Properties of normal varieties

Normal varieties can be characterized using two important geometric properties. Although normal varieties can be singular, these properties show that the singularities of X are "mild." The first property shows that the singular locus of a normal variety will have codimension ≥ 2. The statement relies on the following fact about regularity: the locus of nonregular points in a quasiprojective variety is a proper closed subset. (We proved this for smoothness in Theorem 5.4.2 but we will not prove the analogous statement for regularity.)

Theorem 5.6.5 (Property R1). *Let X be a normal quasiprojective variety. Let X_{nonreg} denote the closed locus of nonregular points in X. Then $\dim(X_{nonreg}) \leq \dim(X) - 2$.*

In particular, every normal variety of dimension 1 is regular.

Proof. It suffices to prove the theorem when $X = \mathrm{mSpec}(R)$ is an affine variety of dimension n.

Suppose for a contradiction that X_{nonreg} has codimension 1. Let Z be a codimension 1 irreducible component of X_{nonreg} equipped with its reduced structure. Thus Z is the vanishing locus of a height 1 prime ideal $\mathfrak{p}$. Consider the local ring $R_{\mathfrak{p}}$: this is a Noetherian local ring of dimension 1. Since integral closure is preserved by localization, $R_{\mathfrak{p}}$ is a DVR.

Let f be a generator of the maximal ideal in $R_{\mathfrak{p}}$. Then there is some $g \in R \backslash \mathfrak{p}$ such that $f \in \mathcal{O}_X(D_g)$. Note that $D_g \cap Z$ will be the vanishing locus of the localized ideal $\mathfrak{p}_g$. This ideal contains the ideal (f), and the cokernel of the inclusion $(f) \subset \mathfrak{p}_g$ vanishes after tensoring by $R_{\mathfrak{p}}$. Since the cokernel is finitely generated, there is a single function $h \in R_g$ such that tensoring by R_{gh} will also kill the cokernel of $(f) \subset \mathfrak{p}_g$. In other words, the intersection $D_{gh} \cap Z$ will be the vanishing locus of the single function $f \in \mathcal{O}_X(D_{gh})$. For simplicity of notation, we let $\widetilde{X} = D_{gh}$ and $\widetilde{Z} = D_{gh} \cap Z$.

Since $\widetilde{Z}$ is irreducible and reduced it has a regular point z. By Exercise 5.1.6 $T_{\widetilde{Z},z}$ is the vanishing locus of a single linear equation in $T_{\widetilde{X},z}$. Thus

$$\dim(X) - 1 = \dim(Z) = \dim(T_{\widetilde{Z},z}) = \dim(T_{\widetilde{X},z}) - 1.$$

This shows that z is also a regular point in X, yielding a contradiction. $\square$

Remark 5.6.6. More generally, we say a quasiprojective scheme X is Rk if the nonregular locus of X has codimension $\geq k + 1$.

The second key result addresses the behavior of functions. We have already seen that in some situations we can remove a codimension 2 subset without affecting the ring of functions (Exercise 1.10.11 and Example 1.10.9). The following result gives us a systematic way of thinking about this property.

Theorem 5.6.7 (Property S2)**.** *Let X be a normal quasiprojective variety. Let Z be a closed subset of X of codimension ≥ 2. Then the natural injection $\mathcal{O}_X(X) \hookrightarrow \mathcal{O}_X(X\backslash Z)$ is an isomorphism.*

Equivalently, any rational function in $\mathbb{K}(X)$ will fail to be defined along a codimension 1 subset. The complex analytic analogue of this result is Hartog's theorem showing that a holomorphic function defined away from a codimension 2 subset can be extended to the entire manifold.

Proof. We first consider the case when X is an affine variety. We will use the following theorem from commutative algebra:

Theorem 5.6.8. *Let R be a Noetherian domain that is integrally closed. Then*

$$R = \bigcap_{\textit{height } 1 \textit{ primes } \mathfrak{p}} R_{\mathfrak{p}}$$

as subsets of $\mathrm{Frac}(R)$*.*

Thus it suffices to show that $\mathcal{O}_X(X\backslash Z) \subset R_{\mathfrak{p}}$ for every height 1 prime $\mathfrak{p}$. Suppose $r \in \mathcal{O}_X(X\backslash Z)$. Fix a height 1 prime $\mathfrak{p}$ so that $V(\mathfrak{p})$ has codimension 1 in X. Since Z has codimension ≥ 2 we have $V(\mathfrak{p}) \cap X\backslash Z \neq \emptyset$. Exercise 1.10.13 shows that $r \in R_{\mathfrak{p}}$. (In other words, since $V(\mathfrak{p})$ meets the locus where r is defined, the denominator of r must be contained in $R\backslash\mathfrak{p}$.)

When X is an arbitrary quasiprojective variety, the statement follows from the affine case using the gluing property of the sheaf of functions. □

Remark 5.6.9. More generally, we say that a ring R satisfies Sk if the localization of R along any prime ideal $\mathfrak{p}$ satisfies

$$\mathrm{depth}(R_{\mathfrak{p}}) \geq \min\{k, \dim(R_{\mathfrak{p}})\}$$

and this definition naturally extends to schemes. It is true, but not obvious, that the S2 property as we have defined it matches up with the S2 property defined using depth.

It turns out that the R1 and S2 properties characterize normality for quasiprojective varieties:

Theorem 5.6.10 (Serre's Criterion)**.** *Let X be a quasiprojective variety. Then X is normal if and only if X satisfies R1 and S2.*

We will not prove this result.

Example 5.6.11. Using the "depth" definition of the S2 property, one can see that an irreducible reduced hypersurface in $\mathbb{P}^n$ will always be S2. Thus a hypersurface will be normal if and only if it is regular in codimension 1.

5.6.2 Normalization

Definition 5.6.12. Let $X = \mathrm{mSpec}(R)$ be an affine variety. The normalization $\nu : X^\nu \to X$ of X is the morphism of affine varieties corresponding to the inclusion $\nu^\sharp : R \to R_{int}$ to the integral closure of R.

Example 5.6.13. Set $R = \mathbb{K}[x,y]/(y^2 - x^3)$. Then R is not integrally closed: the element $\frac{y}{x} \in \mathrm{Frac}(R)$ lies in the integral closure but not in R. There is an isomorphism $\mathbb{K}[t] \to R_{int}$ via the map $t \mapsto \frac{y}{x}$. Thus the normalization of $\mathrm{mSpec}(R)$ is $\mathbb{A}^1$ equipped with the map $\mathbb{A}^1 \to \mathrm{mSpec}(R)$ defined by $x \mapsto t^2 - 1, y \mapsto t^3 - t$.

Using the compatibility of the integral closure operation with localization, we can extend the definition to arbitrary quasiprojective varieties as follows.

Construction 5.6.14. Note that taking integral closures is compatible with localization: for any multiplicative subset $S \subset R$ the integral closure of R_S is the localization of R_{int} localized along S. In particular, given an open affine subset U of an affine variety X, the normalization of U will be the preimage of U in the normalization of X.

Thus we can define the normalization of any quasiprojective variety X as follows. First suppose X is a projective variety. If we take a closed embedding from X into projective space, then the ideal of homogeneous functions which vanish on X is a homogeneous prime ideal $\mathfrak{p}$. In particular this means that $X = \mathrm{mProj}(R)$ where R is the domain defined by quotienting by $\mathfrak{p}$. Then we can define the normalization of X by taking the mProj of an integral closure of R in the field obtained by localizing along all non-zero homogeneous elements. In general, a quasiprojective variety X is an open subset of some projective variety Y and we can define the normalization of X by taking its preimage in the normalization of Y.

The normalization map $\nu : X^\nu \to X$ is the "minimal" map from a normal variety to X (see Exercise 5.6.17). The following result shows that ν has excellent geometric properties.

Theorem 5.6.15. *Let X be a quasiprojective variety. The normalization map $\nu : X^\nu \to X$ is finite and birational.*

Proof. By taking an open cover by open affines it suffices to prove the result when $X = \mathrm{mSpec}(R)$ is an affine variety. Since the inclusion $\nu^\sharp : R \to R_{int}$ is an integral homomorphism of finitely generated $\mathbb{K}$-algebras, it is finite. Choose a finite set of elements $f_i/g_i \in \mathrm{Frac}(R)$ which give a finite set of generators of R_{int} as an R-module. Let s be the product of the g_i. Then the localization of $\nu^\sharp$ along s is an isomorphism. Thus ν is an isomorphism along the open subset D_s of $\mathrm{mSpec}(R)$. □

Remark 5.6.16. In fact, the normalization map ν will be an isomorphism over the smooth locus of X (which is a non-empty open subset by Theorem 5.4.2).

5.6.3 Exercises

Exercise 5.6.17. Suppose that $f : X \to Y$ is a dominant morphism from a normal quasiprojective variety X to a quasiprojective variety Y. Show that f admits a unique factorization through the normalization of Y.

Exercise 5.6.18. Suppose that $f : X \to Y$ is a finite birational morphism. Show there is a unique map $g : Y^\nu \to X$ from the normalization Y^ν such that the composition $f \circ g$ is the normalization map.

Exercise 5.6.19. This exercise gives an example of a variety which is regular in codimension 1 but not normal.

Consider the subring $R = \mathbb{K}[x^3, x^2, xy, y^2]$ of $\mathbb{K}[x, y]$. Show that the inclusion map of rings defines the normalization of mSpec(R). (mSpec(R) is known as the "pinched plane" – can you see why?)

Prove explicitly that mSpec(R) is not S_2 by finding an open subset $U \subset \text{mSpec}(R)$ whose complement has codimension 2 such that $\mathcal{O}_{\text{mSpec}(R)}(\text{mSpec}(R)) \subsetneq \mathcal{O}_{\text{mSpec}(R)}(U)$.

Exercise 5.6.20. Let X be a projective curve (i.e. a projective variety of dimension 1). Prove that X is normal if and only if X is regular. (Hint: show that both conditions are equivalent to requiring that every stalk $\mathcal{O}_{X,x}$ be a DVR.)

Chapter 6

Subvarieties of projective space

This chapter is devoted to the systematic study of closed subschemes of projective space. To every closed subscheme X we associate two constructions:

(1) The Hilbert polynomial, recording the dimensions of the graded pieces of the homogeneous coordinate ring of X.

(2) The degree, a positive integer reflecting the number of intersections points of X with a general linear space of complementary dimension.

These two constructions are closely related. In fact, the degree is determined by the leading coefficient of the Hilbert polynomial. This chapter is devoted to the study of these two constructions. Along the way, we will prove several fundamental properties of projective space which you should be sure to internalize (Proposition 6.2.1, Exercise 6.2.15, Theorem 6.4.8).

The main motivation for studying the degree comes from homology theory. Let's briefly review the homology of projective space. We saw earlier that $\mathbb{P}^n_{\mathbb{C}}$ can be written as a disjoint union

$$\mathbb{A}^n_{\mathbb{C}} \cup \mathbb{A}^{n-1}_{\mathbb{C}} \cup \ldots \mathbb{A}^1_{\mathbb{C}} \cup \{pt\}.$$

In fact this yields a cellular decomposition of $\mathbb{P}^n_{\mathbb{C}}$. Since the cells only occur in even dimension, all the boundary maps in the chain complex computing cellular homology are zero. This shows that:

$$H_i(\mathbb{P}^n_{\mathbb{C}}, \mathbb{Z}) \cong \begin{cases} \mathbb{Z} & \text{if } 0 \leq i \leq 2n \text{ and } i \text{ is even} \\ 0 & \text{otherwise} \end{cases}$$

By the universal coefficient theorem the cohomology groups have the same form.

Since $\mathbb{P}^n_{\mathbb{C}}$ is a complex Kähler manifold, the cohomology carries many structures: the cup product, Poincaré duality, the Hodge diamond, Hard Lefschetz, etc. However, since the cohomology groups of $\mathbb{P}^n_{\mathbb{C}}$ are so simple most of these constructions are uninteresting. The only interesting structure is the cup product in cohomology. As a ring, we have

$$H^*(\mathbb{P}^n_{\mathbb{C}}, \mathbb{Z}) \cong \mathbb{Z}[x]/(x^{n+1})$$

where x is identified with an additive generator for $H^2(\mathbb{P}^n_{\mathbb{C}}, \mathbb{Z})$. Concretely, if $\sigma \in H^2(\mathbb{P}^n_{\mathbb{C}}, \mathbb{Z})$ denotes the first Chern class of the dual of the tautological bundle then every cohomology class in degree $2k$ is a multiple of σ^k. (Note that in algebraic geometry the *dual* of the tautological bundle is naturally identified with the positive generator 1 and the tautological bundle is identified with -1.)

If we dualize the cup product, we get a "cap product" structure on the homology group. We will write L_k for the positive generator of $H_{2k}(\mathbb{P}^n_{\mathbb{C}}, \mathbb{Z})$ obtained by pushing forward the fundamental class of a k-plane in $\mathbb{P}^n_{\mathbb{C}}$. Then the cap product has the following description.

$$\begin{aligned} H_{2k}(\mathbb{P}^n_{\mathbb{C}}, \mathbb{Z}) \times H_{2l}(\mathbb{P}^n_{\mathbb{C}}, \mathbb{Z}) &\xrightarrow{\cap} H_{2k+2l-2n}(\mathbb{P}^n_{\mathbb{C}}, \mathbb{Z}) \\ (\alpha = aL_k, \beta = bL_l) &\mapsto \alpha \cap \beta = abL_{k+l-n} \end{aligned}$$

Suppose that α, β are two homology classes corresponding to oriented submanifolds L, M of $\mathbb{P}^n_{\mathbb{C}}$. If L and M intersect "transversally", then $\alpha \cap \beta$ is represented by the oriented submanifold $L \cap M$. In other words, the cap product on homology corresponds to the geometric operation of taking intersections (in sufficiently nice situations).

The degree is the "algebraic analogue" of the homology class: if $X \subset \mathbb{P}^n_{\mathbb{C}}$ is a smooth subvariety of dimension k, then X carries a fundamental class and its pushforward to $H_{2k}(\mathbb{P}^n_{\mathbb{C}}, \mathbb{Z}) \cong \mathbb{Z}$ is the integer $\deg(X)$. This follows from the defining property of the degree: $\deg(X)$ is computed by intersecting against a plane, which is a generator of the homology in complementary dimension.

Carrying this analogue further, we would like to construct a "cap product" in our setting by taking intersections of algebraic subschemes. Bezout's Theorem shows that (in good situations) the degree of $X \cap Y$ is the product of $\deg(X)$ and $\deg(Y)$ so that our "algebraic cap product" exactly matches the corresponding construction in homology.

The last section is devoted to Hilbert schemes. It turns out that whenever we have a "nice" (i.e. flat) family of closed subschemes of $\mathbb{P}^n$ the Hilbert polynomial stays constant in the family. (The converse is also true: one can test the flatness of a family by checking whether the Hilbert polynomial is constant.) Thus if we want to construct a moduli space of closed subschemes of $\mathbb{P}^n$, the best option is to look at all closed subschemes with a fixed Hilbert polynomial. Although we do not have the tools to give a rigorous construction of the Hilbert scheme, we will take an informal look at a few examples.

Throughout this chapter we assume that our ground field $\mathbb{K}$ is algebraically closed. This is mainly for convenience; the results (when correctly formulated) are true in more generality.

6.1 Hilbert polynomials

Notation 6.1.1. Let S be a graded ring and let M be a graded S-module. For any integer b we let $M(b)$ denote the graded S-module which is abstractly isomorphic to M but has the grading shifted by b, that is, $M(b)_r = M_{r+b}$.

The Hilbert polynomial is a construction from commutative algebra that measures the "size" of a finitely generated graded $\mathbb{K}$-algebra (or more generally, a graded module).

Definition 6.1.2. Let S be a finitely generated graded $\mathbb{K}$-algebra and suppose that M is a finitely generated graded S-module. The Hilbert function of M is the function $\chi_M : \mathbb{N} \to \mathbb{N}$ which assigns to each positive number r the number $\chi_M(r) := \dim_{\mathbb{K}}(M_r)$.

Example 6.1.3. If $S = \mathbb{K}[x_0, \ldots, x_n]$ is the polynomial ring in $n + 1$ variables then $\chi_S(r) = \binom{n+r}{r}$.

Exercise 6.1.4. Set $R = \mathbb{K}[x_0, \ldots, x_n]$. Consider a homogeneous function $f \in R$ of degree d and let $S = R/(f)$. Using the exact sequence of graded R-modules

$$0 \to R(-1) \xrightarrow{\cdot f} R \to S \to 0$$

show that $\chi_S(r) = \binom{n+r}{r} - \binom{n+r-d}{r-d}$ (where as usual we interpret a binomial coefficient as 0 if one of the inputs is negative).

Since different graded rings can have the same mProj, the Hilbert function is *not* an invariant of the projective scheme mProj(S). Thus we will primarily be interested in the following situation.

Definition 6.1.5. Let X be a closed subscheme of $\mathbb{P}^n$ defined by a saturated ideal I. We define the Hilbert function of X (as a subscheme of $\mathbb{P}^n$) to be the Hilbert function of the quotient ring $\mathbb{K}[x_0, \ldots, x_n]/I$. We will denote this function by χ_X.

6.1.1 Computing Hilbert functions

There are several ways to compute Hilbert functions. First, we can work directly from the definition. For any degree r we have an exact sequence

$$0 \to I_r \to \mathbb{K}[x_0, \ldots, x_n]_r \to S_r \to 0.$$

The kernel I_r is the space of homogeneous degree r polynomials which vanish along X. Thus S_r represents the space of homogeneous degree r polynomials "restricted" to X. (Remember, these polynomials are not actually functions on $\mathbb{P}^n$ and their restrictions are not functions on X. Rather, they are "homogeneous functions" on both spaces.)

When X is a particularly nice variety we can compute the dimension of S_r directly from this definition.

Example 6.1.6. Let X be the rational normal curve in $\mathbb{P}^n$. Recall that under the identification $X \cong \mathbb{P}^1$ the homogeneous degree r functions on $\mathbb{P}^n$ restrict to the homogeneous degree nr functions on X. Thus we have

$$\chi_X(r) = \dim \mathbb{K}[x, y]_r^{(n)} = nr + 1.$$

Example 6.1.7. Let X be the union of three distinct points in $\mathbb{P}^2$. We first show that restriction yields a map $\mathbb{K}[x, y, z]_r \to \mathbb{K}^3$ in the following sense. Let's choose a linear function ℓ which does not vanish on any of the three points. In particular, X is contained in the affine chart $D_{+,\ell}$. We then obtain a linear map

$$\begin{aligned} \psi_r : \mathbb{K}[x, y, z]_r &\to \mathcal{O}_X(X) \cong \mathbb{K}^3 \\ f &\mapsto f/\ell^r \end{aligned}$$

The Hilbert function $\chi_X(r)$ will be the dimension of the image of this map.

First, note that $\chi_X(1)$ depends upon whether or not the three points are collinear. If the three points are collinear, then the map ψ_1 has a one-dimensional kernel so that $\chi_X(1) = 3 - 1 = 2$. Otherwise we have $\chi_X(1) = 3 - 0 = 3$.

For $r \geq 2$, we claim that $\chi_X(r) = 3$. Indeed, for each point there is a linear function which vanishes at that point but not the others. By multiplying and adding these together in various combinations, we obtain polynomials of higher degrees which yield a surjection onto $\mathbb{K}^3$.

In general, the best way to compute the Hilbert polynomial of the vanishing locus of $I \subset \mathbb{K}[x_0, \ldots, x_n]$ is often to compute a free graded resolution of $\mathbb{K}[x_0, \ldots, x_n]/I$. The following example illustrates this technique.

Example 6.1.8. Suppose that $X \subset \mathbb{P}^n$ is the intersection of two hypersurfaces $V_+(f)$ and $V_+(g)$ which share no common components. Set $R = \mathbb{K}[x_0, \ldots, x_n]$ and $I = (f, g)$. Let $d = \deg(f)$ and $e = \deg(g)$. We have an exact sequence of graded modules

$$0 \to R(-d-e) \xrightarrow{\begin{bmatrix} \cdot(-e) \\ \cdot f \end{bmatrix}} R(-d) \oplus R(-e) \xrightarrow{[\cdot f, \cdot g]} R \to R/I \to 0.$$

Show that for $r \geq d + e$ we have $\chi_X(r) = \binom{n+r}{r} - \binom{n+r-d}{r} - \binom{n+r-e}{r} + \binom{n+r-d-e}{r}$. (What happens for smaller values of r?)

6.1.2 Hilbert polynomials

Generalizing Example 6.1.7, we see that the Hilbert function of a zero-dimensional scheme is eventually constant.

Lemma 6.1.9. *Let $X \subset \mathbb{P}^n$ be a closed subscheme with $\dim(X) = 0$. Then for d sufficiently large $\chi_X(d)$ is equal to $\dim_{\mathbb{K}}(\mathcal{O}_X(X))$.*

Proof. Let S be the quotient of $\mathbb{K}[x_0, \dots, x_n]$ by the saturated ideal I defining X. Up to a coordinate change, we may suppose that X is contained in the affine chart D_{+,x_0}. In particular, X is an affine scheme which is the vanishing locus of some ideal $J \subset \mathbb{K}[\frac{x_1}{x_0}, \dots, \frac{x_n}{x_0}]$. Since $\mathcal{O}_X(X)$ is Artinian, there are finitely many elements in the quotient $\mathbb{K}[\frac{x_1}{x_0}, \dots, \frac{x_n}{x_0}]/J$ which generate this ring as a $\mathbb{K}$-module. In particular, by clearing denominators we see that that for d sufficiently large we have $\mathcal{O}_X(X) \cong S_d$ as $\mathbb{K}$-vector spaces. $\square$

The most important property of the Hilbert function is that it is eventually polynomial. Furthermore, the various coefficients of the polynomial have interesting geometric interpretations.

Theorem 6.1.10. *Let $X \subset \mathbb{P}^n$ be a closed subscheme. There is a polynomial $P(d)$ such that $\chi_X(d) = P(d)$ for all d sufficiently large. Furthermore, the degree of P is the dimension of X.*

The polynomial P is known as the Hilbert polynomial of X. You may have seen a variant of this result before in a commutative algebra course. We will give a "geometric" proof.

Proof. We prove the theorem by induction on $\dim(X)$. The case when $\dim(X) = 0$ is handled by Lemma 6.1.9.

Suppose now that $\dim(X) > 0$. Define $S = \mathbb{K}[x_0, \dots, x_n]/I$. We claim that for a general hyperplane H its equation ℓ is not a zero divisor in S. Indeed, the zero divisors in S are the union of the associated (necessarily homogeneous) primes for the zero ideal in S. For each such prime consider the corresponding vanishing locus $Z_i \subset \mathbb{P}^n$. The hyperplanes whose equations are zero divisors in S are exactly the hyperplanes which contain some Z_i. By Exercise 4.4.16, such hyperplanes form a proper closed subset of the moduli space $(\mathbb{P}^n)^\vee$ parametrizing hyperplanes.

Fix a general hyperplane defined by an equation ℓ. Consider the exact sequence of abelian groups

$$0 \to S_{r-1} \xrightarrow{\cdot\ell} S_r \to (S/(\ell))_r \to 0. \tag{6.1.1}$$

The ideal (I, ℓ) defines $X \cap H$. Exercise 4.4.16 shows that $X \cap H$ has dimension $\dim(X) - 1$. Furthermore, the ideal (I, ℓ) will agree with its saturation in sufficiently high degree. Thus by induction on dimension we know that $\chi_{X\cap H}$ is eventually equal to a polynomial of degree $\dim(X) - 1$. For r sufficiently large we have $\chi_X(r) - \chi_X(r-1) = P_{X\cap H}(r)$. Using the properties of successive difference functions, we deduce that χ_X is eventually polynomial of degree $\dim(X)$. $\square$

Exercise 6.1.11. Prove carefully the claim above about difference functions implicit in the proof above: if $\chi_X(r) - \chi_X(r-1)$ is eventually polynomial, then $\chi_X(r)$ is eventually polynomial and its degree is one larger than the degree of the difference equation.

6.1.3 Exercises

Exercise 6.1.12. Prove that the Hilbert function of four points in $\mathbb{P}^2$ is $2, 3, 4, 4, \ldots$ if the points are collinear or $3, 4, 4, 4, \ldots$ if the points are not.

Exercise 6.1.13. Prove that the Hilbert polynomial for the d-uple Veronese embedding of $\mathbb{P}^n$ is

$$P(r) = \binom{n + r \cdot d}{d}.$$

Exercise 6.1.14. Compute the Hilbert polynomial for $\mathbb{P}^n \times \mathbb{P}^m$ under its Segre embedding.

Exercise 6.1.15. Let $X \subset \mathbb{P}^3$ be a conic curve and let $Y \subset \mathbb{P}^3$ be the union of two skew lines (i.e. two lines which span $\mathbb{P}^3$). Note that X and Y both have dimension 1 and degree 2. Compute the Hilbert polynomials of X and Y and show that they are different.

Exercise 6.1.16. Let $X \subset \mathbb{P}^2$ be the union of three lines defined by $V_+(xy(x - y))$ and let $Y \subset \mathbb{P}^3$ be the union of the three coordinate axes through the point $(0 : 0 : 0 : 1)$. Compute the Hilbert polynomials of X and Y and show that they are different.

6.2 Degree

Suppose that X is a closed subscheme of $\mathbb{P}^n$. As discussed in the introduction, the degree of X is a positive integer which represents the "homology class" of X. In this section we will study this important invariant.

6.2.1 Intersections in projective space

We first start by developing the theory of intersections in projective space. The following proposition is fundamental for understanding projective space. We have already seen the special case when Y is a hypersurface in Exercise 2.11.14.

Proposition 6.2.1. *Let X and Y be closed subschemes of $\mathbb{P}^n$ of dimensions a, b respectively. Suppose that $a + b \geq n$. Then $X \cap Y \neq \emptyset$ and $\dim(X \cap Y) \geq a + b - n$.*

The proof uses a clever trick to put us in a situation where we can use Krull's PIT.

Proof. Let $C(X), C(Y) \subset \mathbb{A}^{n+1}$ be the cones over X and Y (that is, the loci defined by the same ideals in $\mathbb{K}[x_0, \ldots, x_n]$ but considered as subsets of $\mathbb{A}^{n+1}$ instead of $\mathbb{P}^n$). We have $\dim(C(X)) = a + 1$ and $\dim(C(Y)) = b + 1$.

Consider the diagonal $i : \Delta \subset \mathbb{A}^{n+1} \times \mathbb{A}^{n+1}$. Note that i takes $C(X) \cap C(Y)$ isomorphically to $\Delta \cap (C(X) \times C(Y))$. Exercise 1.7.18 showed that Δ is defined by the equations $\{1 \otimes x_i - x_i \otimes 1\}_{i=0}^{n+1}$. By applying Krull's PIT $(n + 1)$ times, we see that the dimension of $\Delta \cap (C(X) \times C(Y))$ is at least $(a + b + 2) - (n + 1)$. Thus $C(X) \cap C(Y)$ has dimension $\geq (a + b - n) + 1$ and in particular has dimension ≥ 1. Taking the image in $\mathbb{P}^n$ drops the dimension by 1, yielding the desired statement. □

In particular, this shows that a closed subscheme X of $\mathbb{P}^n$ of dimension k will intersect every plane of dimension $n - k$. Our next result shows that if we drop the dimension by 1, then a general plane will not intersect X.

Proposition 6.2.2. *Let X be a closed subscheme of $\mathbb{P}^n$ of dimension k. A general plane of dimension $n - k - 1$ will not intersect X.*

Exercise 6.2.3. Prove Proposition 6.2.2. (Hint: use Exercise 4.5.7.)

6.2.2 Degree

Definition 6.2.4. Let X be a closed subscheme of $\mathbb{P}^n$ and let $P(x)$ denote its Hilbert polynomial. Suppose that $P(x)$ has degree r. Then the degree of X is $r!$ times the coefficient of x^r in $P(x)$.

Example 6.2.5. Let X be a hypersurface in $\mathbb{P}^n$ defined by an equation of degree d. Exercise 6.1.4 shows that the degree of X is also d. (Thus there is no conflict between our two competing definitions of the degree of a hypersurface.)

Example 6.2.6. Let X be a 0-dimensional subscheme of $\mathbb{P}^n$. Then Lemma 6.1.9 shows that

$$\deg(X) = \dim_{\mathbb{K}} \mathcal{O}_X(X).$$

Again, this definition is compatible with the usage of degree in other contexts (see Section 4.2.1).

Our next task is to understand the geometric significance of the degree.

Theorem 6.2.7. *Let X be a closed subscheme of $\mathbb{P}^n$. Let H be a general hyperplane. Then $\deg(X) = \deg(X \cap H)$.*

Proof. We return to the setting of the proof of Theorem 6.1.10. If we let S denote the quotient of $\mathbb{K}[x_0, \ldots, x_n]$ by the saturated ideal I defining X, then Equation (6.1.1) gives an exact sequence

$$0 \to S_{r-1} \xrightarrow{\cdot \ell} S_r \to (S/(\ell))_r \to 0.$$

For r sufficiently large, by taking dimensions we see that

$$P_X(r) - P_X(r-1) = P_{X \cap H}(r).$$

Using standard facts about difference of polynomials, we see that $(\dim X)!$ times the leading coefficient of P_X is the same as $(\dim X - 1)!$ times the leading coefficient of $P_{X\cap H}$. $\square$

Exercise 6.2.8. By mimicking the proof of Theorem 6.1.10 prove the following statement: if X is a closed subscheme of $\mathbb{P}^n$ and H is a general degree d hypersurface then $\deg(X \cap H) = \deg(X) \cdot \deg(H)$.

Repeating Theorem 6.2.7 inductively, we obtain the following. Suppose that X is a closed subscheme of $\mathbb{P}^n$ of dimension k. Let Z denote the 0-dimensional subscheme obtained by intersecting X against a general $(n-k)$-plane. Then the degree of X is the same as the $\mathbb{K}$-dimension of the Artinian ring $\mathcal{O}_Z(Z)$.

When X is a smooth subvariety, the Bertini Theorem shows that the intersection of X against a general $(n-k)$-plane L will be smooth, and thus also reduced. In this situation the degree of X is simply the number of points in $X \cap L$.

Remark 6.2.9. This discussion also explains why the degree matches up with the "homology class": we can compute the integer representing a fixed homology class by taking a cap product against a generator of the cohomology group in complementary dimension.

6.2.3 Degree and components

We next relate the degree of a closed subscheme X to the degrees of its irreducible components.

Definition 6.2.10. Let X be a quasiprojective scheme and let X_0 be an irreducible component of X equipped with the reduced structure. Let $U = \mathrm{mSpec}(R)$ be any open affine in X that is contained in X_0. Since X_0 is irreducible the localization of R at its unique minimal prime is an Artinian ring S over the field $\mathbb{K}(X_0)$.

We define the multiplicity of X_0 in X, denoted by $\mathrm{mult}(X_0, X)$, to be the dimension of S over $\mathbb{K}(X_0)$. (When X is irreducible, we will use the shorthand $\mathrm{mult}(X)$ instead of the more precise $\mathrm{mult}(X, X)$.)

The multiplicity of X_0 in X measures the failure of X to be reduced at a general point of X_0. It does not observe any "extra" non-reducedness along closed subschemes properly contained in X_0.

Exercise 6.2.11. Show that in the setting of Definition 6.2.10 $\mathrm{mult}(X)$ is the same as the degree of the tangent cone of X at a general point $x \in X$ (when considered as a projective subscheme of the projectivized tangent space of $\mathbb{P}^n$ at x). Deduce that the multiplicity of X along X_0 is independent of the choice of open affine $U \subset X_0$.

Theorem 6.2.12. *Let X be a closed subscheme of $\mathbb{P}^n$ of dimension k. Let $X_1, \dots, X_m$ denote the irreducible components of X which have dimension k. Then*

$$\deg(X) = \sum_{i=1}^{m} \mathrm{mult}(X_i, X) \deg(X_i).$$

In particular, the degree only reflects the top-dimensional components of X and the "generic" non-reducedness of the top-dimensional components. Conceptually, it is fair to say that the degree is really an invariant for closed subvarieties, since the computation for arbitrary closed subschemes reduces to this case.

Proof. We will use the following result from commutative algebra:

Theorem 6.2.13. *Let S be a Noetherian graded ring and let M be a finitely generated graded S-module. Then there exists a filtration*

$$0 = M^0 \subset M^1 \subset \ldots \subset M^{r-1} \subset M^r = M$$

such that for each i there is a homogeneous prime $\mathfrak{p}_i \subset S$ such that the quotient M^i/M^{i+1} is isomorphic to a shift $S/\mathfrak{p}_i(l_i)$ for some $l_i \in \mathbb{Z}$. Furthermore, for each minimal homogeneous prime $\mathfrak{p}$ of S the number of times this prime occurs in the set of quotients is equal to the length of $M_\mathfrak{p}$ over $S_\mathfrak{p}$.

In our situation we let S be the quotient of $\mathbb{K}[x_0, \dots, x_n]$ by the saturated ideal defining X and apply the theorem to S considered as a module over itself. We deduce that the Hilbert polynomial $P_X(r)$ is a sum of Hilbert polynomials $P_{Z_i}(r - l_i)$ for irreducible closed subschemes Z_i of X (where the presence of l_i accounts for the shift in degrees). Note that

(1) By Theorem 6.1.10 the only irreducible closed subschemes which can contribute to the leading coefficient of P_X will be the top-dimensional components of X.

(2) The contribution of such a component to the leading coefficient will not be affected by the "shift" in degrees.

(3) The number of times each top-dimensional component X_i of X occurs is the same as the length of $S_\mathfrak{p}$ over itself.

The last step in the proof is given by the following exercise. □

Exercise 6.2.14. Finish the proof of Theorem 6.2.12 by verifying that for a minimal prime $\mathfrak{p}$ corresponding to a top dimensional component X_i of X the length of $S_\mathfrak{p}$ over itself is the same as $\mathrm{mult}(X_i, X)$. (Hint: show that the length of $S_\mathfrak{p}$ is the same as the length of the local ring obtained by homogeneous localization along $\mathfrak{p}$.)

6.2.4 Exercises

Exercise 6.2.15. Prove that there is no non-constant morphism $f : \mathbb{P}^n \to X$ to a quasiprojective scheme X with $\dim(X) < n$. (Hint: if there were such a morphism, one could use the fibers of f to find disjoint subvarieties of $\mathbb{P}^n$.)

Exercise 6.2.16. Show that the degree of the dth Veronese embedding of $\mathbb{P}^m$ in $\mathbb{P}^{\binom{m+d}{d}-1}$ is d^m.

Exercise 6.2.17. Compute the degree of the Segre embedding of $\mathbb{P}^n \times \mathbb{P}^m$.

Exercise 6.2.18. Let X be a closed subvariety of $\mathbb{P}^n$. Suppose that p is a $\mathbb{K}$-point in $\mathbb{P}^n \backslash X$ so that projection away from p defines a morphism $\phi : X \to \mathbb{P}^{n-1}$. Suppose furthermore that ϕ is birational into its image. (One can show that if $\dim(X) \leq n-2$ then projection away from a general point in $\mathbb{P}^n$ will satisfy this property.) Prove that $\deg(X) = \deg(\phi(X))$.

6.3 Bezout's Theorem

As discussed in the introduction, the degree of a subscheme represents its "homology class" and we would like to develop an analogue of the cap product in our setting. We expect this product to be modeled on taking intersections. Bezout's theorem tells us when intersections of varieties are indeed compatible with this desired "algebraic cap product" obtained by multiplying degrees.

6.3.1 Complete intersections

The easiest version of Bezout's theorem addresses a very special type of subscheme of $\mathbb{P}^n$.

Definition 6.3.1. A closed subscheme $X \subset \mathbb{P}^n$ is said to be a complete intersection if its saturated homogeneous ideal is generated by a regular sequence of elements in $\mathbb{P}^n$.

Theorem 6.3.2. *Suppose that X is a complete intersection in $\mathbb{P}^n$ defined by the $f_1, \ldots, f_k$. Then*

$$\deg(X) = \deg(f_1) \cdot \deg(f_2) \cdot \ldots \cdot \deg(f_k).$$

Proof. The proof is by induction on k. For any $j \leq k$ we let X_j denote the complete intersection in $\mathbb{P}^n$ defined by the equations $f_1, \ldots, f_j$. We also let S^j denote the quotient of $\mathbb{K}[x_0, \ldots, x_n]$ by the saturated ideal defining X_j. We have exact sequences

$$0 \to S^j(-\deg(f_{j+1})) \to S^j \to S^{j+1} \to 0.$$

In sufficiently high degrees, the dimension of the graded pieces of these rings are measured by the Hilbert polynomials. Applying the theory of difference equations we obtain the desired statement. □

The geometric significance of complete intersections is clarified by the unmixedness theorem.

Proposition 6.3.3. *Suppose that X is a closed subscheme of $\mathbb{P}^n$ of codimension k. Then X is a complete intersection if and only if its saturated homogeneous ideal is generated by k equations.*

Note that by Krull's PIT k is the minimal possible number of generators for the ideal of X.

Proof. We only need to prove the reverse implication. Suppose that $f_1, \ldots, f_k$ generate the ideal for X. We will need a result originally proved by Macaulay:

Theorem 6.3.4. *Let R be a polynomial ring. Suppose that I is an ideal whose height is equal to the number of its generators. Then every associated prime of the 0 ideal in the quotient R/I is a minimal prime.*

We prove by induction that the sequence consisting of the first r polynomials $f_1, \ldots, f_r$ is regular. The base case is trivial. For the induction step, suppose that $(f_1, \ldots, f_r)$ is regular. Krull's PIT shows that every component of $V_+(f_1, \ldots, f_r)$ must have codimension r in $\mathbb{P}^n$ – otherwise it would be impossible to get an $(n-k)$-dimensional scheme by dropping dimension $k - r$ more times. Theorem 6.3.4 shows that f_{r+1} will be a zero divisor in $R/(f_1, \ldots, f_r)$ if and only if it vanishes along some component of $V_+(f_1, \ldots, f_r)$. But in this case it is again impossible to get a $(n-k)$-dimensional scheme by intersecting all the hypersurfaces. We conclude that the sequence remains regular when adding in f_{r+1}. □

By combining this result with Theorem 6.3.2, we obtain a version of Bezout's Theorem for complete intersections.

Corollary 6.3.5. *Suppose that $X, Y \subset \mathbb{P}^n$ are complete intersections of dimensions a and b respectively. Suppose that every component of $X \cap Y$ has dimension $a + b - n$. Then*

$$\deg(X) \cdot \deg(Y) = \deg(X \cap Y)$$

Theorem 6.3.2 can also be used to show that many closed subvarieties of $\mathbb{P}^n$ are not complete intersections. For example:

Proposition 6.3.6. *Let $C \subset \mathbb{P}^3$ be a twisted cubic. Then there is no homogeneous ideal $I \subset \mathbb{K}[w, x, y, z]$ with two generators such that $V_+(I) = C$.*

Proof. Suppose there were an ideal I with two generators f, g such that $C = V_+(I)$. In particular C is a complete intersection, so that

$$\deg(f) \cdot \deg(g) = \deg(C) = 3.$$

In particular, either f or g must be a linear equation. But C is not contained in any hyperplane in $\mathbb{P}^3$, a contradiction. □

6.3.2 Bezout's Theorem

Unfortunately the case of complete intersections is rather special. We really would like a version of Bezout's Theorem that holds under less restrictive hypotheses. At the very least we will need to require that $X \cap Y$ has the expected dimension.

Definition 6.3.7. Let X and Y be closed subvarieties of $\mathbb{P}^n$ which have dimensions a and b respectively. We say that X and Y meet dimensionally transversally if $a + b \geq n$ and every component of $X \cap Y$ has dimension $a + b - n$.

Unfortunately this condition is not restrictive enough; we will need to impose extra conditions for Bezout's Theorem to hold. We will not prove the following general statement:

Theorem 6.3.8. *Let X and Y be closed subschemes of $\mathbb{P}^n$. Suppose that X and Y meet dimensionally transversally. Suppose also that for every component $Z_i \subset X \cap Y$ the stalks $\mathcal{O}_{X,z}$ and $\mathcal{O}_{Y,z}$ are Cohen-Macaulay rings for general points $z \in Z_i$. Then*

$$\deg(X) \cdot \deg(Y) = \deg(X \cap Y).$$

Recall that a local ring is said to be Cohen-Macaulay if the maximal length of a regular sequence is the same as its dimension. Thus the Cohen-Macaulay condition should be viewed as a "smoothness hypothesis" that is somewhat weaker than regularity. Examples of Cohen-Macaulay schemes (i.e. schemes such that every local ring is Cohen-Macaulay) include smooth varieties, complete intersections in $\mathbb{P}^n$, and all reduced irreducible curves.

Theorem 6.3.8 fails if we drop the Cohen-Macaulay hypothesis.

Example 6.3.9. The easiest example of a non-Cohen-Macaulay scheme is the union X of two planes in $\mathbb{P}^4$ which meet transversally at a single point p. For example, L could be the vanishing locus of the ideal (wy, wz, xy, xz) in $\mathbb{K}[v, w, x, y, z]$. The stalk of $\mathcal{O}_X$ at the point p will fail to be Cohen-Macaulay.

Suppose we choose a plane L such that $X \cap L$ has dimension 0. If L is a general plane then $X \cap L$ has degree 2 as expected. However, if L meets X at p, then an easy computation shows that $X \cap L$ has degree 3.

Remark 6.3.10. It is natural to try to fix Bezout's Theorem by adding in a "correction term" to the degrees of the components of $X \cap Y$. This correction is known Serre's formula. Suppose that X and Y are locally defined in an affine chart $\mathbb{A}^n$ by ideals I, J in the coordinate ring . Assume that X and Y meet dimensionally transversally. For any component Z of $X \cap Y$ we define the intersection multiplicity

$$I_Z(X, Y) = \sum_{i=0}^{n} \mathrm{length}_{\mathcal{O}_{X\cap Y, Z}} \mathrm{Tor}_i^A(A/\widetilde{I}, A/\widetilde{J})$$

where A is the local ring $\mathcal{O}_{\mathbb{A}^n, Z}$ and $\widetilde{I}, \widetilde{J}$ are the ideals defining X, Y in this local ring. Note that the 0th Tor term is exactly the multiplicity of Z in $X \cap Y$ (in the sense of Definition 6.2.10); when X and Y are Cohen-Macaulay the higher Tor terms will vanish. Then Bezout's Theorem states that if X, Y meet dimensionally transversally we have

$$\deg(X) \cdot \deg(Y) = \sum I_Z(X, Y) \deg(Z)$$

as Z varies over all components of $X \cap Y$.

It is sometimes useful to know when the intersection of two varieties has a component of multiplicity one. The following proposition describes a situation known as "transversal intersection"

Proposition 6.3.11. *Let X and Y be closed subvarieties of $\mathbb{P}^n$. Suppose that X and Y meet dimensionally transversally. Let Z be a component of $X \cap Y$ such that X and Y are Cohen-Macaulay at general points of Z. Then* $\operatorname{mult}_Z(X \cap Y) = 1$ *if and only if for a general point $z \in Z$ we have that X and Y are smooth at z and their tangent spaces $T_{X,z}$, $T_{Y,z}$ meet transversally.*

Remark 6.3.12. We can drop the Cohen-Macaulay assumption in Proposition 6.3.11 if we use the "corrected" intersection multiplicity in place of $\operatorname{mult}_Z(X \cap Y)$.

6.3.3 Exercises

Exercise 6.3.13. Let X be the dth Veronese embedding of $\mathbb{P}^n$. Recall that under the identification $X \cong \mathbb{P}^n$ a homogeneous linear function on the ambient projective space restricts to define a homogeneous degree d equation on X. Apply Bezout's Theorem on $\mathbb{P}^n$ to compute the degree of X.

Exercise 6.3.14. Let $\Sigma_{n,m}$ denote the Segre embedding of $\mathbb{P}^n \times \mathbb{P}^m$.

(1) Prove that if we take a union $L_1 \cup L_2$ where L_1 is the preimage of a hyperplane under the first projection and L_2 is the preimage of a hyperplane under the second embedding then there is a hyperplane $H \subset \mathbb{P}^{nm+n+m}$ such that $H \cap \Sigma_{n,m} = L_1 \cup L_2$.

(2) Use this construction to show that $\deg(\Sigma_{n,m}) = \binom{n+m}{n}$.

(3) Use this construction to show that the vanishing locus of a bihomogeneous polynomial of bidegree (a, b) in $\Sigma_{n,m}$ will have degree $a \cdot \binom{n+m-1}{m} + b \cdot \binom{n+m-1}{n}$.

Exercise 6.3.15. Show that the Segre variety corresponding to $\mathbb{P}^1 \times \mathbb{P}^2$ is not a complete intersection.

6.4 Low degree subvarieties of projective space

In this section we turn to the problem of classifying closed subschemes of $\mathbb{P}^n$ of low degree. According to Theorem 6.2.12, the degree of X is determined by the degrees of its top-dimensional components and their multiplicities. In particular, the degree cannot detect non-reduced structure along proper closed subsets. Thus in our classification scheme we should focus on the case when X is a closed subvariety.

6.4.1 Curves

Our classification of low degree curves starts with the following lemma:

Lemma 6.4.1. *Let C be a closed subvariety of $\mathbb{P}^n$ of dimension 1. Suppose that $\deg(C) < n$. Then C is contained in a hyperplane.*

Proof. Note that any n points of $\mathbb{P}^n$ are contained in some hyperplane. In particular, if we choose n distinct points on C we can find a hyperplane H containing these points.

Suppose that $C \not\subset H$. Since C is reduced, Krull's PIT guarantees that the equation ℓ defining H is not a zero divisor in the homogeneous coordinate ring defining C. Repeating the argument of Theorem 6.2.7 we see that $\deg(C) = \deg(C \cap H)$. But the left hand side is $< n$ and by construction the right hand side is n, yielding a contradiction. We conclude that $C \subset H$. □

Exercise 6.4.2. More generally, show that if C is closed subvariety of $\mathbb{P}^n$ of dimension 1 and degree $n-k$ then X is contained in a linear subspace of codimension k.

By Lemma 6.4.1 the lowest possible degree of an irreducible curve C not contained in a hyperplane is n. We have already seen one example of this phenomenon: rational normal curves. Conversely, it turns out that rational normal curves provide the only examples.

Proposition 6.4.3. *Let $C \subset \mathbb{P}^n$ be a reduced irreducible curve of degree n. Suppose that C is not contained in any hyperplane. Then C is a rational normal curve.*

Proof. We first show that C is smooth. Suppose C has a singular point p. Choose a hyperplane H that contains p and $n-1$ other general points of C. Then $H \cap C$ has length $n+1$, contradicting Bezout's Theorem. We conclude that C is smooth.

We next show that C is isomorphic to $\mathbb{P}^1$. Fix $n-1$ points $\{p_i\}_{i=1}^{n-1}$ on the curve C and consider the $\mathbb{P}^1$ parametrizing hyperplanes H through the p_i. Bezout's Theorem shows that for every H the intersection $H \cap C$ is a dimension 0 scheme with length n. Thus, so long as H is not tangent to C at one of the p_i, the intersection of H with C will define one additional reduced point. Let L denote the $(n-2)$-plane spanned by the p_i. Then projection away from L defines an isomorphism from $C \backslash \{p_i\}$ to an open subset of $\mathbb{P}^1$. In fact this rational map extends to an isomorphism of C with $\mathbb{P}^1$.

As we intersect C with any hyperplane H, the ϕ-image of $H \cap C$ will be a length n subscheme of $\mathbb{P}^1$. This defines a $\mathbb{K}$-linear map from the space of linear functions on $\mathbb{P}^n$ to the space of degree n functions on $\mathbb{P}^1$. Since this map is injective, it is an isomorphism, showing that C is a rational normal curve. $\square$

Example 6.4.4. Let's classify all the reduced irreducible curves C in $\mathbb{P}^n$ which have low degree.

Degree 1: By Exercise 6.4.2 we see that C must be a line.

Degree 2: Let C be a reduced irreducible curve of degree 2. Then Exercise 6.4.2 and Proposition 6.4.3 show that C is a conic in a 2-plane in $\mathbb{P}^n$.

Degree 3: Let $C \subset \mathbb{P}^n$ be a reduced irreducible curve of degree 3. Then Exercise 6.4.2 and Proposition 6.4.3 show that either:

- C is a twisted cubic in some 3-plane in $\mathbb{P}^n$, or
- C is a degree three curve in a 2-plane in $\mathbb{P}^n$ (and in particular, is the intersection of a plane with a cubic hypersurface).

Degree 4: Let $C \subset \mathbb{P}^n$ be a reduced irreducible curve of degree 4. Then Exercise 6.4.2 and Proposition 6.4.3 show that one of the following holds.

- C is a rational normal curve cubic in some 4-plane in $\mathbb{P}^n$,
- C is a degree four curve in a 3-plane in $\mathbb{P}^n$. Although we do not have the tools to show it, it turns out that there are two ways to obtain such curves: one can take an intersection of two quadric hypersurfaces in $\mathbb{P}^3$, or one can choose a smooth quadric $Q \cong \mathbb{P}^1 \times \mathbb{P}^1$ and look at the vanishing locus of a bihomogeneous polynomial of bidegree $(1,3)$ in Q.
- C is a degree four curve in a 2-plane in $\mathbb{P}^n$ (and thus the intersection of a plane with a degree 4 hypersurface).

6.4.2 Higher dimensions

Exercise 6.4.5. Show that if $X \subset \mathbb{P}^n$ is a closed subvariety of dimension k and degree $< n-k+1$ then X is contained in a hyperplane.

Exercise 6.4.6. Suppose that $X \subset \mathbb{P}^n$ is a closed subvariety of degree 1. Show that X must be a plane.

The following statement generalizes Proposition 6.4.3. We will present it without proof.

Theorem 6.4.7. *Let $X \subset \mathbb{P}^n$ be a closed subvariety of dimension k. If $\deg(X) = n-k+1$ then X is either:*

(1) a quadric hypersurface,

(2) a cone over the Veronese surface in $\mathbb{P}^5$,

(3) a rational normal scroll.

6.4.3 Applications

We now give a couple applications of Bezout's theorem. First, we prove a result promised in Exercise 2.7.11.

Theorem 6.4.8. *Let $\mathbb{K}$ be an algebraically closed field. Then every automorphism of $\mathbb{P}^n$ is given by multiplication by an element of* $\mathrm{PGL}_{n+1}(\mathbb{K})$.

Using a similar approach one can show that the statement is valid for any field.

Proof. Suppose that $\phi : \mathbb{P}^n \to \mathbb{P}^n$ is an automorphism. Let H be any hyperplane in $\mathbb{P}^n$ and let ℓ be a line which meets H transversally at one point. Then $\phi(\ell)$ and $\phi(H)$ also meet transversally at one point. By Bezout's Theorem we must have $\deg(\phi(\ell)) = \deg(\phi(H)) = 1$. By Exercise 6.4.6 $\phi(H)$ is still a hyperplane.

Restricting ϕ to the complements of H and $\phi(H)$, we obtain a morphism $\phi' : \mathbb{A}^n \to \mathbb{A}^n$. By applying the same argument to other hyperplanes and intersecting with this open affine, we see that any linear function on the target must pullback to a power of a linear function on the domain. But the only way that such a map ϕ' can be an isomorphism is if the pullback of a linear function is a linear function. By extending this map to all of $\mathbb{P}^n$ using homogeneous coordinates we obtain the desired statement. □

Bezout's theorem also allows us to prove a classical statement about singular points of plane curves.

Corollary 6.4.9. *Let C be a reduced plane curve of degree d. Then C has at most $\frac{1}{2}d(d-1)$ singular points.*

This bound is achieved by a union of d lines. If C is irreducible, one can prove a stronger bound of $\frac{1}{2}(d-1)(d-2)$.

Proof. Let f be the equation defining C. Since C is reduced f has a derivative which does not vanish identically along any component of C. Without loss of generality we may assume this derivative is $\partial f/\partial x$. This is a homogeneous function of degree $(d-1)$, and thus its vanishing locus is a curve $D \subset \mathbb{P}^2$. By construction C and D do not share any components, so Bezout's theorem shows that $\deg(C \cap D) = d(d-1)$. Note that C and D will intersect at every singular point of C and that the intersection multiplicity at such a point is at least 2. Thus twice the number of singular points is at most $d(d-1)$, proving the statement. □

6.5 Hilbert schemes

Thus far we have seen two examples of moduli spaces of subvarieties of $\mathbb{P}^n$:

(1) The moduli space $\mathbb{P}^{\binom{n+d}{n}-1}$ of degree d hyperplanes.

(2) The moduli space $\mathbb{G}(k,n)$ of k-planes.

These are both special examples of a more general construction known as a Hilbert scheme. The Hilbert scheme is the prototypical example of a moduli space in algebraic geometry and is instrumental in the construction of many other types of moduli space.

Suppose we would like to construct a parameter space for subvarieties of $\mathbb{P}^n$. If we hope to obtain a finite dimensional moduli space, we should start by fixing some topological invariants of our subvarieties: the dimension and the degree. It is natural to ask whether there are any other invariants which are always preserved by deformations. It turns out that for "good" (i.e. flat) families of closed subschemes the precise list of invariants is given by the coefficients of the Hilbert polynomial.

The Hilbert scheme $M_{P,n}$ parametrizes all closed subschemes of $\mathbb{P}^n$ with Hilbert polynomial P. This moduli space has two key properties. First, it is projective: if we take a one-dimensional family of closed subschemes with Hilbert polynomial P, we can define a "limit" closed subscheme which still has Hilbert polynomial P. This compactness property shows that there are no "missing" points in our family.

The second is that Hilbert schemes are "fine moduli spaces". Consider the functor which assigns to a scheme Y the set of closed subschemes $\mathcal{C} \subset Y \times \mathbb{P}^n$ such that every fiber of $\mathcal{C} \to Y$ over a closed point has Hilbert polynomial P. This functor is representable: there is a bijection between such constructions over Y and morphisms $Y \to M_{P,n}$. In particular, this implies the existence of a universal family $\mathcal{U}$ over $M_{P,n}$ such that any family over Y is constructed by choosing a morphism $Y \to M_{P,n}$ and taking the construction $\mathcal{U} \times_{M_{P,n}} Y$.

6.5.1 Construction of the Hilbert scheme

Theorem 6.5.1. *Fix a polynomial $P(r)$. There is a projective scheme $M_{P,n}$ whose points parametrize the closed subschemes of $\mathbb{P}^n$ which have Hilbert polynomial equal to $P(r)$.*

Sketch of proof: The first step in the proof is to show the following theorem:

Theorem 6.5.2. *There is some fixed positive integer m (depending only on $P(r)$ and n) such that the following holds. Suppose that X is any closed subscheme whose Hilbert polynomial is $P(r)$ and let I be the saturated ideal defining it. Then:*

(1) The Hilbert function $\chi(r)$ is equal to the Hilbert polynomial $P(r)$ for every $r \geq m$.

(2) I is generated in degree $\leq m$. In particular, I is the saturation of the ideal generated by I_m.

The first statement shows that $I_m \subset \mathbb{K}[x_0, \ldots, x_n]_m$ is a subspace of codimension $P(m)$. Let G be the Grassmannian parametrizing subspaces of $\mathbb{K}[x_0, \ldots, x_n]_m$ with this codimension. Then the second condition shows that the map which associates to any closed subscheme X the corresponding subspace in $\mathbb{G}$ is injective.

In fact, the image of this map is a closed subscheme of G. To see that it is an algebraic subset, we need to write down conditions on G which guarantee that the ideal generated by a given codimension $P(m)$ subspace defines a closed subscheme with the correct Hilbert polynomial. We need to ensure that each $\chi(r)$ for $r > m$ has the correct value. Since the map $\mathbb{K}[x_0, \ldots, x_n]_{m-r} \times I_m \to I_r$ is surjective, for each r we obtain a system of equations expressing the fact that this linear map must have rank no more than $\dim(I_r)$. Combining all these equations we obtain a closed subset of G. □

6.5.2 Examples

Although we do not yet have the tools to give a rigorous construction of the Hilbert scheme, we can still study several special examples. Suppose we fix a polynomial P_r. The first step toward understanding the Hilbert scheme $M_{P,n}$ is to identify which types of closed subschemes of $\mathbb{P}^n$ have Hilbert polynomial P. We can then hope to leverage this information to obtain a better understanding of $\mathbb{P}^n$.

Example 6.5.3. Suppose that $P(r)$ is a non-negative constant a. The closed subschemes of $\mathbb{P}^n$ with Hilbert polynomial a will be 0-dimensional subschemes of degree a. For example, $M_{1,n}$ is just $\mathbb{P}^n$ itself.

A point in $M_{2,n}$ will either parametrize two distinct points or a single point with length 2 (i.e. a point "with a tangent direction"). $M_{2,n}$ can be constructed by blowing up the diagonal in $\mathbb{P}^n \times \mathbb{P}^n$ (which has the effect of separating tangent directions for the points contained in the diagonal) and then quotienting by the involution which swaps the two factors.

Surprisingly, $M_{a,n}$ can be very complicated even when n and a are small. While there will always be at least one component which generically parametrizes a distinct points in $\mathbb{P}^n$, there can also be other components (and the singularities and non-reduced structure can vary wildly). For example, the Hilbert scheme of 8 points in $\mathbb{P}^4$ is reducible – there is a component which only parametrizes non-reduced schemes.

Example 6.5.4. Suppose that $P(r) = r + 1$. Then $P(r)$ is the Hilbert polynomial of a line. We claim that any closed subscheme $X \subset \mathbb{P}^n$ with Hilbert polynomial $P(r)$ must be a line.

To see this, note that X has dimension 1 and degree 1. Suppose we choose a one-dimensional component $X_0 \subset X$ equipped with its reduced structure. **** shows that X_0 must be a line. Then we have an exact sequence

$$0 \to I_X \to I_{X_0} \to K \to 0$$

where K is the cokernel of the inclusion map.

We know that for X the Hilbert function eventually agrees with the Hilbert polynomial $r+1$. But X_0 already has Hilbert polynomial $r+1$. This implies that the cokernel K satisfies $K_m = 0$ for m sufficiently large. Since both ideals are saturated, this means that $I_X = I_{X_0}$.

The Hilbert scheme $M_{r+1,n}$ is of course the Grassmannian $\mathbb{G}(1,n)$.

Example 6.5.5. Suppose that $P(r) = r+2$. Again a closed subscheme X with this Hilbert polynomial will have dimension 1 and degree 1. However, in contrast to Example 6.5.4 it is now possible for X to be reducible or non-reduced. Indeed, by running through the argument again we see that either X will be either:

(1) the union of a line and a distinct point, or

(2) a line with a single non-reduced point such that I/I^2 has $\mathbb{K}$-dimension 1.

A general point of $M_{r+2,n}$ will parametrize the first type of closed subscheme, and there will be a proper sublocus of $M_{r+2,n}$ parametrizing the second. Note that $M_{r+2,n}$ is birational to $\mathbb{P}^n \times \mathbb{G}(1,n)$ but is not isomorphic to it since if we fix a point p on a line ℓ there are many different non-reduced structures we can impose at p.

Example 6.5.6. Suppose that $P(r) = 2r+1$. There are three types of curves in $\mathbb{P}^n$ of degree 2: conics, pairs of lines, and double lines. Accordingly, there are three different types of closed subschemes of $\mathbb{P}^n$ with Hilbert polynomial P.

(1) The Hilbert polynomial of a conic is $2r+1$. Thus if X contains a conic as a closed subscheme then X must itself be a conic.

(2) A pair of disjoint lines in $\mathbb{P}^n$ will have Hilbert polynomial $2r+2$ or $2r+1$ depending on whether or not the lines intersect. Thus it is impossible for X to contain a pair of skew lines. If X contains a pair of intersecting lines, then X is equal to these two lines.

(3) Suppose that X is supported on a single line. Since X has degree 2, we see that X must be a non-reduced scheme with generic multiplicity 2. (Note that while affine double lines are all isomorphic, projective double lines need not be.)

    Let I_ℓ denote the ideal of the line supporting X. Since X defines a double line, we see $I_\ell \supsetneq I_X$. Since I_X contains at least $n-1$ functions which vanish everywhere along ℓ, the smallest the degrees of such functions can be is $(n-2)$ linear functions and 1 quadratic function. However, such an ideal already defines a subscheme with Hilbert polynomial $2r+1$ (and any further reduced structure would only increase the Hilbert polynomial). We conclude that X is a planar double line, that is, a double line in a plane $\mathbb{P}^2$.

Note that any subscheme of type (1), (2), or (3) is (scheme-theoretically) contained in a unique 2-plane. Thus, we get a morphism $\rho : M_{2r+1,n} \to \mathbb{G}(2, n)$. Each fiber of ρ parametrizes the set of degree 2 curves in a fixed $\mathbb{P}^2$. Thus the fibers of ρ are isomorphic to $\mathbb{P}^5$ and $M_{2r+1,n}$ has dimension $5 + 3(n - 3)$.

6.5.3 Exercises

Exercise 6.5.7. Show that if $P(r) = \binom{k+r}{k}$ is the Hilbert polynomial of a k-dimensional plane in $\mathbb{P}^n$ then the Hilbert scheme is the Grassmannian $\mathbb{G}(k, n)$.

Exercise 6.5.8. Show that the Hilbert scheme parametrizing closed subschemes of $\mathbb{P}^n$ with Hilbert polynomial
$$P(r) = \binom{n+r}{n} - \binom{n+r-d}{n}$$
is the moduli space of hypersurfaces $\mathbb{P}^{\binom{n+d}{d}-1}$ discussed in Example 4.5.1.

Part II

Sheaves and schemes

Chapter 7

Sheaves

Suppose that X is a topological manifold. For any open subset $U \subset X$, let $\mathcal{C}(U)$ denote the space of continuous functions on U. There are two important ways in which these functions interact as we vary our open set:

(1) Restriction: given an inclusion of open sets $V \subset U$, any function on U also induces a function on V.

(2) Gluing: given an open cover $\{V_i\}$ of U and functions f_i on V_i which agree on the common overlaps, we obtain a unique function f on U by gluing the f_i.

The definition of a sheaf formalizes these two important properties to give us a general language for discussing "abstract functions" on a topological space. A sheaf $\mathcal{F}$ assigns to every open subset $U \subset X$ an abelian group $\mathcal{F}(U)$. The key axiom that a sheaf must satisfy is gluing: given "local objects" that agree on overlaps, there is a unique way to glue them to get a "global object". Thus sheaves are often used to compare local vs. global geometry.

Example 7.0.1. Let U be an open set of $\mathbb{C}$ and let $g : U \to \mathbb{C}$ be a holomorphic function. Can we write $g = e^h$ for some holomorphic function h? The key observation is that this is problem can always be solved locally: for any point $x \in U$ we can take a logarithm of g on a sufficiently small neighborhood of x. Thus the answer will depend only on the global geometry of U.

It turns out that our ability to find global logarithms can naturally be expressed in the language of sheaves. The exponential map defines a function $\mathcal{O}_{hol} \to \mathcal{O}_{hol}^{\times}$ taking the sheaf of holomorphic functions to the sheaf of invertible holomorphic functions. This map is surjective locally, but not globally – this is a common feature for morphisms of sheaves. The failure of the surjectivity of this map can be controlled using sheaf cohomology theory; as a result, one can see that the obstruction to the surjectivity of the exponential map comes from the singular homology $H^1(U, \mathbb{Z})$.

In addition to sheaves of continuous (or differentiable or etc.) functions, there are several other common geometric constructions of sheaves. The most important example comes from vector bundles.

Example 7.0.2. Let X be a topological space. Recall that a rank r vector bundle on X is a continuous surjection $\pi : \mathcal{V} \to X$ and for each point $x \in X$ an identification of the structure of the vector space $\mathbb{R}^k$ on the fiber $\pi^{-1}x$ such that the following compatibility is satisfied: for every point $x \in X$ there is an open neighborhood U of x and a homeomorphism

$$\psi_U : U \times \mathbb{R}^r \to \pi^{-1}U$$

such that $(\pi \circ \psi_U)(x, v) = x$ and if we fix $y \in U$ the map $v \mapsto \psi_U(y, v)$ is a linear isomorphism between $\mathbb{R}^k$ and $\pi^{-1}y$.

Given any open set $U \subset X$, the space of sections of U is the set of continuous maps $\sigma : U \to \mathcal{V}$ such that $\pi \circ \sigma = id_U$. Since sections can be glued, we can define a sheaf of sections which associates to any open set U the set of sections of π over U.

The importance of this example is evident in the language we use for sheaves, which includes terms like "sections" and "fibers" (intended to generalize the corresponding notions in Example 7.0.2). We will eventually use this intuition to *define* vector bundles in the language of algebraic geometry.

7.0.1 Preliminaries

Let $\mathbf{C}$ and $\mathbf{D}$ be categories. An adjunction between $\mathbf{C}$ and $\mathbf{D}$ is a pair of functors $F : \mathbf{D} \to \mathbf{C}$ and $G : \mathbf{C} \to \mathbf{D}$ such that for all objects X in $\mathbf{C}$ and all objects Y in $\mathbf{D}$ we have a bijection

$$\mathrm{Hom}_{\mathbf{C}}(FY, X) \cong \mathrm{Hom}_{\mathbf{D}}(Y, GX).$$

We furthermore require that these bijections are natural in both entries: there is a natural isomorphism between the functors $\mathrm{Hom}_{\mathbf{D}}(F-, X)$ and $\mathrm{Hom}_{\mathbf{D}}(-, GX)$ from $\mathbf{D}$ to $\mathbf{Set}$, and similarly for the two functors $\mathrm{Hom}_{\mathbf{C}}(FY, -)$ and $\mathrm{Hom}_{\mathbf{C}}(Y, G-)$ from $\mathbf{C}$ to $\mathbf{Set}$.

Definition 7.0.3. In the situation above F is called a left adjoint to G and G is called a right adjoint to F.

The most important property of adjoint functors is limit preservation.

Theorem 7.0.4. *Right adjoint functors commute with limits. Left adjoint functors commute with colimits.*

For example, suppose that products exist in $\mathbf{C}$. Then for any two objects $X_1, X_2 \in \mathbf{C}$ the object $G(X_1 \times X_2)$ satisfies the right universal property to be identified with the product $G(X_1) \times G(X_2)$.

Suppose next that $\mathbf{C}$ and $\mathbf{D}$ are abelian categories. Then the kernel construction is an example of a categorical limit and the cokernel construction is an example of a categorical colimit. Thus:

Corollary 7.0.5. *Let F, G be an adjoint pair defining an adjunction of abelian categories. Then G is left exact and F is right exact.*

7.1 Presheaves and sheaves

We start by defining the notion of a presheaf.

Definition 7.1.1. Let X be a topological space. A presheaf $\mathcal{F}$ of abelian groups on X consists of the following data:

(1) for every open subset U, an abelian group $\mathcal{F}(U)$ whose elements are known as "sections of $\mathcal{F}$ on U", and

(2) for every inclusion of non-empty open subsets $V \subset U$, a homomorphism $\rho_{U,V} : \mathcal{F}(U) \to \mathcal{F}(V)$ known as a "restriction map"; we sometimes denote $\rho_{U,V}(s)$ using the notation $s|_V$

satisfying the following conditions:

(1) (Normalization) $\mathcal{F}(\emptyset) = 0$.

(2) (Compatibility) The assignment $U \mapsto \mathcal{F}(U)$ and $(V \subset U) \mapsto \rho_{U,V}$ defines a contravariant functor from the category of open subsets of X (with morphisms = inclusions) to the category of abelian groups. In other words, $\rho_{U,U} = id$ and if $W \subset V \subset U$ then $\rho_{U,V} \circ \rho_{V,W} = \rho_{U,W}$.

More generally, a presheaf with values in a category $\mathbf{C}$ is defined via a contravariant functor from the category of open subsets of X to $\mathbf{C}$. (In particular, the restriction maps must be morphisms in $\mathbf{C}$.) All the presheaves we work with will take values in an enrichment of the category of abelian groups $\mathbf{Ab}$, so we won't lose anything by working exclusively in this setting for now.

Presheaves do not have very much structure. As discussed in the introduction to the chapter, we will normally be interested in the situation when our sections admit some kind of gluing property. The notion of a sheaf was introduced in Section 1.8:

Definition 7.1.2. Suppose that X is a topological space and that $\mathcal{F}$ is a presheaf on X. If $\mathcal{F}$ satisfies the following additional two axioms we say that $\mathcal{F}$ is a sheaf:

(3) (Identity) Suppose that $\{V_i\}$ is an open cover of U. Suppose that $f_1, f_2 \in \mathcal{F}(U)$ satisfy $\rho_{U,V_i}(f_1) = \rho_{U,V_i}(f_2)$ for every i. Then $f_1 = f_2$.

(4) (Gluing) Suppose that $\{V_i\}$ is an open cover of U. Suppose that for every i we have an element $f_i \in \mathcal{F}(V_i)$. Furthermore suppose that for every pair of indices i, j we have $\rho_{V_i, V_i \cap V_j}(f_i) = \rho_{V_j, V_i \cap V_j}(f_j)$. Then there exists an element $f \in \mathcal{F}(U)$ satisfying $\rho_{U,V_i}(f) = f_i$ for every i.

Remark 7.1.3. A presheaf which just satisfies the identity axiom is called a "separated presheaf". We won't use this notation.

it is quite common to capture the Identity and Gluing axioms using the following exact sequence: suppose we fix an open subset U of X and an open cover $\{V_i\}_{i\in I}$ of U. Then the axioms show that

$$0 \to \mathcal{F}(U) \xrightarrow{\phi} \prod_{i\in I} \mathcal{F}(V_i) \xrightarrow{\psi} \prod_{(i,j)\in I^2} \mathcal{F}(V_i \cap V_j)$$

is an exact sequence. Here ϕ is the product of the restriction maps $\rho_{U,V_i} : \mathcal{F}(U) \to \mathcal{F}(V_i)$ and ψ sends a tuple (f_i) to $(f_i|_{U_i\cap U_j} - f_j|_{U_i\cap U_j})$. For the cover $\{V_i\}$ the identity axiom corresponds to the injectivity of ϕ and the gluing axiom corresponds to exactness at the middle place.

Example 7.1.4. Let X be a topological space. Fix an abelian group A. The constant presheaf $\mathcal{F}$ with value A associates to every open set U the abelian group A (and to every inclusion of open sets the identity map). This is not in general a sheaf. One issue is that it fails the gluing axiom: for example, if U, V are disjoint open subsets then the gluing axiom should imply that $\mathcal{F}(U \cup V) = \mathcal{F}(U) \times \mathcal{F}(V)$ but this fails in our example.

If we want to work with sheaves, we should instead consider the locally constant sheaf with value A: to every open set U we assign the set of locally constant functions $U \to A$. We denote the locally constant sheaf with value A on X by A_X.

7.1.1 Stalk

Just as for sheaves, the notion of a stalk is absolutely essential for working with presheaves.

Definition 7.1.5. Let X be a topological space and let $\mathcal{F}$ be a presheaf on X. For any point $x \in X$ the stalk $\mathcal{F}_x$ is defined to be the direct limit

$$\mathcal{F}_x = \varinjlim_{U \ni x} \mathcal{F}(U).$$

In other words, consider the set of pairs (U, f) where U is an open neighborhood of x and $f \in \mathcal{F}(U)$. Say that two pairs (U, f) and (V, g) are equivalent if there is some open set $W \subset U \cap V$ that contains x such that the restrictions of f and g to $\mathcal{F}(W)$ coincide. Then $\mathcal{F}_x$ is the set of equivalence classes of pairs (U, f). We call these equivalence classes germs of sections of $\mathcal{F}$.

Since the direct limit of abelian groups receives a map from each group, for any open neighborhood U of x there is a canonical restriction map $\rho_{U,x} : \mathcal{F}(U) \to \mathcal{F}_x$. When $\mathcal{F}$ is a sheaf of functions, this restriction map assigns to any function its germ at x.

Exercise 7.1.6. Let X be a topological space equipped with a sheaf $\mathcal{F}$. Prove that for any open set U the product of the restriction maps

$$\rho : \mathcal{F}(U) \to \prod_{x\in U} \mathcal{F}_x$$

is injective. (However, this need not be true when $\mathcal{F}$ is a presheaf.)

7.1.2 Examples

We have already extensively studied one class of sheaves: the structure sheaves on quasiprojective $\mathbb{K}$-schemes. We next introduce a couple other sheaves that play an important role in algebraic geometry.

Example 7.1.7. Let X be a topological space and let $x \in X$ be a point. Fix an abelian group A. The skyscraper sheaf at x with value A, denoted by $A(x)$, is defined via the assignment:

$$\mathcal{F}(U) = \begin{cases} A \text{ if } x \in U \\ 0 \text{ if } x \notin U \end{cases}$$

where the restriction maps are the identity map if both open sets contain x and 0 otherwise. If x is a closed point, then the stalk of $A(x)$ at x is equal to A and at any other point is equal to 0. (What happens if x is not a closed subset of X?)

The following example describes one of the most important examples of sheaves in algebraic geometry.

Example 7.1.8. Consider projective space $\mathbb{P}^n$. Define the field $\mathbb{F}$ to be the localization of the homogeneous coordinate ring $\mathbb{K}[x_0, \dots, x_n]$ along all non-zero homogeneous elements and let $\mathbb{F}_m \subset \mathbb{F}$ denote the set of degree m elements in $\mathbb{F}$. For any integer m we define the sheaf $\mathcal{O}_{\mathbb{P}^n}(m)$ via the rule

$$\mathcal{O}_{\mathbb{P}^n}(m)(U) := \left\{ \frac{f}{g} \in \mathbb{F}_m \,\middle|\, V_+(g) \cap U = \emptyset \right\}$$

where the restriction maps are the inclusions. Note that $\mathcal{O}_{\mathbb{P}^n}(0)$ is just the structure sheaf $\mathcal{O}_{\mathbb{P}^n}$ defined by degree 0 elements.

Let's analyze the sheaf $\mathcal{O}_{\mathbb{P}^n}(m)$ in more detail. First, note that the global sections are just the degree m homogeneous polynomials in the homogeneous coordinate ring:

$$\mathcal{O}_{\mathbb{P}^n}(m)(\mathbb{P}^n) = \mathbb{K}[x_0, \dots, x_n]_m.$$

In particular the global sections of the sheaves $\mathcal{O}_{\mathbb{P}^n}(m)$ allow us to recover the homogeneous coordinate ring in an "intrinsic" way from projective space. Next consider the affine charts $D_i = \mathbb{P}^n \backslash V_+(x_i)$. Recall that $\mathcal{O}_{\mathbb{P}^n}(D_i)$ is isomorphic to $\mathbb{K}[\frac{x_0}{x_i}, \dots, \frac{x_n}{x_i}]$. Then $\mathcal{O}_{\mathbb{P}^n}(m)(D_i)$ is the free $\mathcal{O}_{\mathbb{P}^n}(D_i)$-module generated by x_i^m.

Exercise 7.1.9. Show that for every point $x \in \mathbb{P}^n$ the stalk $\mathcal{O}_{\mathbb{P}^n}(m)_x$ is a free rank 1 module over the local ring $\mathcal{O}_{\mathbb{P}^n,x}$.

7.1.3 Exercises

Exercise 7.1.10. Let X be a topological space with a sheaf $\mathcal{F}$. Suppose that $U \subset X$ is an open subset. Show that we can define a sheaf $\mathcal{F}|_U$ on U by restricting the functor $\mathcal{F}$ to the open sets of X contained in U. This sheaf is called the restriction of $\mathcal{F}$ to U.

Exercise 7.1.11. Let X be a topological space and let $\mathcal{F}$ be a sheaf on X. Suppose that U is an open subset and $f \in \mathcal{F}(U)$. The support of f is defined to be

$$\mathrm{Supp}(f) = \{x \in U \mid \rho_{U,x}(f) \neq 0\}.$$

Prove that $\mathrm{Supp}(f)$ is a closed subset of U. (Hint: what does it mean for an equivalence class to be the zero element in $\mathcal{F}_x$?)

Exercise 7.1.12. Let X be a topological space and let $\mathcal{F}$ be a sheaf on X. The support of the sheaf $\mathcal{F}$ is defined to be the set of points x such that $\mathcal{F}_x$ is not zero. Show that the support of $\mathcal{F}$ need not be a closed subset of X. (Hint: try constructing a suitable sheaf on the two-pointed space $X = \{p, q\}$ where the closed subsets are $\emptyset, \{p\}, X$.)

Exercise 7.1.13. Let $\mathcal{F}$ and $\mathcal{G}$ be presheaves on X. The direct sum of $\mathcal{F}$ and $\mathcal{G}$, denoted by $\mathcal{F} \oplus \mathcal{G}$, is the presheaf which associates to every open set U the direct sum $\mathcal{F}(U) \oplus \mathcal{G}(U)$ and whose restriction maps act componentwise as $\rho_{\mathcal{F}}$ and $\rho_{\mathcal{G}}$.

Show that if $\mathcal{F}$ and $\mathcal{G}$ are sheaves then $\mathcal{F} \oplus \mathcal{G}$ is also a sheaf.

7.2 Morphisms and sheafification

In this section we will show that there is a canonical way of transforming a presheaf into a sheaf.

7.2.1 Morphisms of presheaves

The first step is to define the category of presheaves on X.

Definition 7.2.1. Suppose that X is a topological space and that $\mathcal{F}, \mathcal{G}$ are presheaves on X. A morphism of presheaves $\phi : \mathcal{F} \to \mathcal{G}$ assigns to each open set U a homomorphism $\phi(U) : \mathcal{F}(U) \to \mathcal{G}(U)$ in such a way that ϕ is compatible with restriction: for any open $V \subset U$ the diagram

$$\begin{tikzcd} \mathcal{F}(U) \arrow[r, "\phi(U)"] \arrow[d, "\rho_{\mathcal{F},U,V}"] & \mathcal{G}(U) \arrow[d, "\rho_{\mathcal{G},U,V}"] \\ \mathcal{F}(V) \arrow[r, "\phi(V)"] & \mathcal{G}(V). \end{tikzcd}$$

is a commuting diagram. A morphism of sheaves is defined in the same way.

The category of presheaves on X will be denoted by $\mathbf{PreSh}(X)$ and the category of sheaves on X will be denoted by $\mathbf{Sh}(X)$.

Note that a morphism of presheaves $\phi : \mathcal{F} \to \mathcal{G}$ induces morphisms of stalks in the following way. Fix a point $x \in X$. For every open neighborhood U of x consider the composition

$$\mathcal{F}(U) \xrightarrow{\phi(U)} \mathcal{G}(U) \xrightarrow{\rho_{U,x}} \mathcal{G}_x.$$

This collection of homomorphisms determines a homomorphism $\phi_x : \mathcal{F}_x \to \mathcal{G}_x$ using the universal property of the direct limit.

Example 7.2.2. Consider the sheaf $\mathcal{O}_{\mathbb{P}^1}(m)$ on $\mathbb{P}^1$. Fix a homogeneous polynomial $f \in \mathbb{K}[x, y]$ of degree d. For every open set $U \subset \mathbb{P}^1$ define the homomorphism

$$\begin{gathered} \phi(U) : \mathcal{O}_{\mathbb{P}^1}(m)(U) \to \mathcal{O}_{\mathbb{P}^1}(m+d)(U) \\ \frac{g}{h} \mapsto \frac{fg}{h} \end{gathered}$$

Since these $\phi(U)$ obviously commute with restriction, for any homogeneous degree d polynomial f we obtain a morphism of sheaves $\mathcal{O}_{\mathbb{P}^1}(m) \xrightarrow{\cdot f} \mathcal{O}_{\mathbb{P}^1}(m+d)$.

7.2.2 Sheafification

Suppose that we start with a presheaf $\mathcal{F}$ and would like to modify it to obtain a sheaf. When $\mathcal{F}$ fails the identity axiom, we will need to "decrease" the number of sections by identifying any sections whose restrictions agree. When $\mathcal{F}$ fails the gluing axiom, we will need to "increase" the number of sections to ensure that local sections can be glued.

We will describe a clean way to accomplish both goals at once.

Construction 7.2.3. Let X be a topological space and let $\mathcal{F}$ be a presheaf on X. The sheafification $\mathcal{F}^+$ of $\mathcal{F}$ is defined as follows. For any open set U, we define $\mathcal{F}^+(U)$ as a subset of $\prod_{x\in U}\mathcal{F}_x$ via

$$\mathcal{F}^+(U) = \left\{ (f_x \in \mathcal{F}_x)_{x\in U} \;\middle|\; \begin{array}{c} \text{for every } x \in U, \exists \text{ an open neighborhood} \\ V \ni x \text{ and a section } g \in \mathcal{F}(V) \text{ s.t.} \\ g|_x = f_x \text{ for every } x \in V \end{array} \right\}.$$

For $V \subset U$ the restriction map is induced by the projection maps $\prod_{x\in U}\mathcal{F}_x \to \prod_{x\in V}\mathcal{F}_x$.

It is clear that $\mathcal{F}^+$ is a presheaf. (Since the empty product of abelian groups is 0, the normalization axiom holds.) Note that every presheaf $\mathcal{F}$ admits a canonical presheaf map $sh : \mathcal{F} \to \mathcal{F}^+$ by sending a section $s \in \mathcal{F}(U)$ to the element $(s|_x \in \mathcal{F}_x)_{x\in U}$ of $\mathcal{F}^+(U)$.

Proposition 7.2.4. *Let X be a topological space and let $\mathcal{F}$ be a presheaf on X. The sheafification $\mathcal{F}^+$ is a sheaf.*

Proof. To verify the identity axiom, suppose we have an open set U and an open cover $\{V_i\}$ of U. If two sections $f, f' \in \mathcal{F}^+(U)$ have the same restriction to V_i then their components corresponding to points $x \in V_i$ coincide. Thus if f, f' have the same restriction to each V_i in our open cover then they must be the same section.

To verify the gluing axiom, suppose we have an open set U and an open cover $\{V_i\}$ of U. Suppose we have sections $f_i \in \mathcal{F}^+(V_i)$ whose restrictions to overlaps agree. This guarantees that the component of f_i associated to a point $x \in X$ depends only on x and not on the index i. By selecting these choices of germs at the various points $x \in X$ we obtain an element $f \in \prod_{x\in U}\mathcal{F}_x$. It is clear that f satisfies the necessary condition to lie in $\mathcal{O}_X(U)$. $\square$

A sheafification $\mathcal{F}^+$ comes equipped with a presheaf morphism $sh : \mathcal{F} \to \mathcal{F}^+$ which sends a section $s \in \mathcal{F}(U)$ to its stalks at points in U. We will call sh the sheafification map and consider it as part of the data of the sheafification.

Exercise 7.2.5. Show that sheafification defines a functor from the category of presheaves to the category of sheaves.

While the sheafification $\mathcal{F}^+$ can look quite different from $\mathcal{F}$, there is one important way in which the two presheaves are similar.

Exercise 7.2.6. Let X be a topological space and let $\mathcal{F}$ be a presheaf on X. Show that the sheafification map $sh : \mathcal{F} \to \mathcal{F}^+$ induces an isomorphism of stalks $sh : \mathcal{F}_x \xrightarrow{\cong} \mathcal{F}_x^+$ for every $x \in X$.

In fact, this property – the existence of a presheaf map from $\mathcal{F}$ inducing an isomorphism on stalks – characterizes the sheafification up to isomorphism by Proposition 7.2.8 and Proposition 7.3.14.

Example 7.2.7. Let X be a topological space. The sheafification of the constant presheaf $\mathcal{F}$ with value A is the locally constant sheaf with value A.

7.2.3 Universal property

Proposition 7.2.8. *Let X be a topological space and let $\mathcal{F}$ be a presheaf on X. For any sheaf $\mathcal{G}$ on X, every presheaf morphism $\phi : \mathcal{F} \to \mathcal{G}$ will factor uniquely through $sh : \mathcal{F} \to \mathcal{F}^+$.*

In particular this implies that the sheafification of a sheaf is isomorphic to itself, resolving a potential conflict in our notation.

Proof. The morphism ϕ induces stalk homomorphism $\phi_x : \mathcal{F}_x \to \mathcal{G}_x$. For any open set U, consider the induced map $\prod_{x \in U} \mathcal{F}_x \to \prod_{x \in U} \mathcal{G}_x$. Under this map, an element in $\mathcal{F}^+(U)$ will be sent to an element of

$$\left\{ (g_x \in \mathcal{G}_x)_{x \in U} \;\middle|\; \begin{array}{c} \text{for every } x \in U, \exists \text{ an open neighborhood} \\ V \ni x \text{ and a section } h \in \mathcal{G}(V) \text{ s.t.} \\ h|_x = g_x \text{ for every } x \in V \end{array} \right\}.$$

By combining the gluing axiom with Exercise 7.1.6, we see that this subset of $\prod_{x \in U} \mathcal{G}_x$ is isomorphic to $\mathcal{G}(U)$. In this way we obtain a morphism $\phi^+ : \mathcal{F}^+ \to \mathcal{G}$.

To prove that the factoring ϕ^+ is unique, note that by Exercise 7.2.6 the image of the composition

$$\mathcal{F}^+(U) \to \mathcal{G}(U) \to \prod_{x \in U} \mathcal{G}_x$$

is determined by the maps ϕ_x on stalks. Since the second map is injective by Exercise 7.1.6, we deduce that the factoring ϕ^+ is unique. $\square$

Using the universal property, we see that sheafification defines a functor from $\mathbf{PreSh}(X)$ to $\mathbf{Sh}(X)$. In fact, Proposition 7.2.8 says that the sheafification functor is a left adjoint to the forgetful functor $\mathbf{Sh}(X) \to \mathbf{PreSh}(X)$.

7.2.4 Exercises

Exercise 7.2.9. Suppose that $\phi : \mathcal{F} \to \mathcal{G}$ is a morphism of presheaves such that open set U the map $\phi(U) : \mathcal{F}(U) \to \mathcal{G}(U)$ is injective. Show that the induced map of sheafifications $\mathcal{F}^+ \to \mathcal{G}^+$ has the same property.

Exercise 7.2.10. For each of the following examples of a topological space X equipped with a presheaf $\mathcal{F}$, describe the sheafification $\mathcal{F}^+$.

(1) $X = \mathbb{R}^n$, $\mathcal{F}$ assigns to each open set U the set of bounded functions on U.

(2) $X = S^1$, $\mathcal{F}$ assigns to each open set U the set of continuous functions f on U which satisfy $f(x) = f(-x)$ for every pair of antipodal points $x, -x$ in U.

(Note that we can think of $\mathcal{F}$ as the presheaf of pullbacks of functions under the quotient map $S^1 \to \mathbb{RP}^1$.)

(3) $X = \mathbb{C}$, $\mathcal{F}$ assigns to each open set U the set of holomorphic functions on U which admit a square root.

7.3 Kernels, images, and cokernels

In this section we fix a topological space X. The goal of this section is transport constructions from the category **Ab** to the categories $\mathbf{PreSh}(X)$ and $\mathbf{Sh}(X)$. This is the first step toward giving these categories the structure of abelian categories.

Definition 7.3.1. Suppose that $\phi : \mathcal{F} \to \mathcal{G}$ is a morphism of presheaves of abelian groups. We can define the kernel presheaf (resp. image presheaf) of ϕ simply by assigning to every open subset U the kernel (resp. image) of $\phi(U)$ equipped with the restriction maps from $\mathcal{F}$ (resp. $\mathcal{G}$). We denote this construction by $\ker_{pre}(\phi)$ (resp. $\mathrm{im}_{pre}(\phi)$).

When $\mathcal{F}, \mathcal{G}$ are sheaves it is not obvious (and in fact not true) that these presheaf constructions will yield sheaves. We will correct this deficiency by systematically appealing to the sheafification construction.

The main theme of this section is the relationship between constructions for morphisms of sheaves and properties of the induced morphisms of stalks.

7.3.1 Kernels

Definition 7.3.2. Let $\phi : \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves of abelian groups. The kernel of ϕ, denoted $\ker(\phi)$, is the sheaf which associates to each open set U the kernel of $\phi(U)$ and to each inclusion of open subsets $V \subset U$ the restriction of $\rho_{F,U,V}$ to $\ker(\phi(U))$.

We must verify that this definition actually yields a sheaf. It is clear that the construction yields a presheaf. The identity axiom is an immediate consequence of the identity axiom for $\mathcal{F}$. To verify the gluing axiom, suppose that we are given an open set U, an open cover $\{V_i\}$, and local sections $f_i \in \ker(\phi)(V_i)$. By the gluing axiom in $\mathcal{F}$ there is a section $f \in \mathcal{F}(U)$ which restricts to the various f_i, and we just need to verify that f is in the kernel of ϕ. Note that

$$\phi(U)(f)|_{V_i} = \phi(V_i)(f|_{V_i}) = \phi(V_i)(f_i) = 0$$

for every i. By the identity axiom for the sheaf $\mathcal{G}$ we conclude that $\phi(U)(f) = 0$.

Remark 7.3.3. Note that $\ker(\phi)$ comes equipped with an inclusion morphism of sheaves $i : \ker(\phi) \to \mathcal{F}$.

Definition 7.3.4. We say that a morphism of sheaves $\phi : \mathcal{F} \to \mathcal{G}$ is injective if $\ker(\phi) = 0$. Equivalently, ϕ is injective if and only if $\phi(U)$ is injective for every open set U.

Exercise 7.3.5. Let $\phi : \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves. Prove that $\ker(\phi)_x \cong \ker(\phi_x)$ and that this isomorphism is the map on stalks induced by the inclusion map $i : \ker(\phi) \to \mathcal{F}$.

Deduce that ϕ is injective if and only if the induced stalk maps ϕ_x are injective for every $x \in X$.

7.3.2 Images

Definition 7.3.6. Let $\phi : \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves of abelian groups. The image of ϕ, denoted $\mathrm{im}(\phi)$, is the sheafification of the image presheaf $\mathrm{im}_{pre}(\phi)$.

It is important to note that the image presheaf need not be a sheaf. Even though it will always satisfy the identity axiom, it may fail to satisfy the gluing axiom. The issue is that if $\phi(f_i)$ are local sections of $\mathcal{G}$ which agree on overlaps, there is no reason to expect the corresponding local sections f_i of $\mathcal{F}$ to agree on overlaps. (Example 7.3.13 gives an explicit example.)

Remark 7.3.7. Since the only axiom that the image presheaf fails is the gluing axiom, we must "increase" the number of sections to obtain a sheaf. Indeed, it is easy to see that for any open set U there is an inclusion $\mathrm{im}(\phi(U)) \subset (\mathrm{im}\,\phi)(U)$.

Remark 7.3.8. By Exercise 7.2.9 $\mathrm{im}(\phi)$ comes equipped with an inclusion morphism of sheaves $i : \mathrm{im}(\phi) \to \mathcal{F}$.

Definition 7.3.9. We say that a morphism of sheaves $\phi : \mathcal{F} \to \mathcal{G}$ is surjective if $\mathrm{im}(\phi) = \mathcal{G}$.

Exercise 7.3.10. Let $\phi : \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves. Prove that $\mathrm{im}(\phi)_x \cong \mathrm{im}(\phi_x)$ and this isomorphism of stalks is induced by the inclusion map $i : \mathrm{im}(\phi) \to \mathcal{G}$.

Deduce that ϕ is surjective if and only if the induced stalk maps ϕ_x are surjective for every $x \in X$.

The following important exercise will be used many times in the future.

Exercise 7.3.11. Let X be a topological space and let $\phi : \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves. Prove that ϕ is surjective if and only if for every open set U and every section $g \in \mathcal{G}(U)$ there is an open cover $\{V_i\}$ of U and elements $f_i \in \mathcal{F}(V_i)$ such that $\phi(V_i)(f_i) = g_i$ for every i.

Warning 7.3.12. Suppose that a map of sheaves $\phi : \mathcal{F} \to \mathcal{G}$ is "locally surjective" in the sense that there is a base $\{U\}$ for the topology of X such that for every U the map $\phi(U)$ is surjective. Exercise 7.3.11 shows that ϕ is a surjective map of sheaves.

However, the converse is false: Exercise 7.3.21 shows that a surjective map of sheaves need not be "locally surjective." In fact, for a surjective morphism of sheaves there might not be any open set U such that $\phi(U)$ is surjective.

The following example shows that a surjective morphism of sheaves need not induce a surjection of sections on every open set U. (This also shows that the image presheaf need not be a sheaf.)

Example 7.3.13. Set $\mathbb{P}^1 = \mathrm{mProj}(\mathbb{K}[x, y])$. Consider the morphism $\phi : \mathcal{O}_{\mathbb{P}^1}(-1) \oplus \mathcal{O}_{\mathbb{P}^1}(-1) \to \mathcal{O}_{\mathbb{P}^1}$ obtained by multiplying by x on the first component and multiplying by y on the second component. Then ϕ is a surjective map. Indeed suppose we fix a distinguished open affine set $U \subset \mathbb{P}^1$ which is the complement to the vanishing locus of an irreducible homogeneous polynomial g. Then the induced map of sections on the open set U is

$$(\mathbb{K}[x,y]_g)_{-1} \oplus (\mathbb{K}[x,y]_g)_{-1} \to (\mathbb{K}[x,y]_g)_0.$$

This map is a surjection: if we write $g = xg_1 + yg_2$ for some polynomials g_1, g_2, an element h on the right will be the image of the element $(h\frac{g_1}{g}, h\frac{g_2}{g})$ on the left. Taking limits we see that ϕ is surjective on stalks, and thus surjective.

However, if we look at global sections the induced map

$$\begin{array}{ccc} \mathcal{O}_{\mathbb{P}^1}(-1)(\mathbb{P}^1) \oplus \mathcal{O}_{\mathbb{P}^1}(-1)(\mathbb{P}^1) & \longrightarrow & \mathcal{O}_{\mathbb{P}^1}(\mathbb{P}^1) \\ \| & & \| \\ 0 & & \mathbb{K} \end{array}$$

is the zero map.

7.3.3 Isomorphisms

There are many equivalent ways to describe an isomorphism of sheaves.

Proposition 7.3.14. *Let $\phi : \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves. The following are equivalent:*

(1) ϕ is an isomorphism.

(2) ϕ is injective and surjective.

(3) $\phi(U) : \mathcal{F}(U) \to \mathcal{G}(U)$ is an isomorphism for every open set U.

(4) For every $x \in X$ the induced map on stalks $\phi_x : \mathcal{F}_x \to \mathcal{G}_x$ is an isomorphism.

Proof. (1) $\Rightarrow$ (2): it is clear that $\ker(\phi)$ is contained in $\ker(\phi^{-1}\circ\phi)$, showing that $\ker(\phi) = 0$. Similarly, it is clear that the image of ϕ contains $\operatorname{im}(\phi \circ \phi^{-1})$, so that the image of ϕ must be all of $\mathcal{G}$.

(2) $\Rightarrow$ (3): the injectivity of ϕ implies that each $\phi(U)$ is injective. To see that $\phi(U)$ is surjective, choose some $g \in \mathcal{G}(U)$. By surjectivity of ϕ there is an open cover $\{V_i\}$ of U and elements $f_i \in \mathcal{F}(V_i)$ such that $\phi(V_i)(f_i) = g|_{V_i}$. Since ϕ is injective, the restrictions of f_i and f_j to $V_i \cap V_j$ agree. Thus the f_i glue to give an element $f \in \mathcal{F}(U)$. The image of this element is $g \in \mathcal{G}(U)$, showing surjectivity.

(3) $\Rightarrow$ (1): it is clear that for every inclusion of open sets $V \subset U$ we have $\rho_{U,V} \circ \phi(U)^{-1} = \phi(V)^{-1} \circ \rho_{U,V}$. Thus we can define an inverse morphism via the prescription $(\phi^{-1})(U) = \phi(U)^{-1}$.

(2) $\Leftrightarrow$ (4): apply Exercise 7.3.5 and Exercise 7.3.10. $\square$

Warning 7.3.15. Proposition 7.3.14.(4) does *not* claim that if $\mathcal{F}, \mathcal{G}$ are sheaves which satisfy $\mathcal{F}_x \cong \mathcal{G}_x$ for every $x \in X$ then $\mathcal{F}$ and $\mathcal{G}$ are isomorphic. We can only conclude the existence of an isomorphism when there is some global map ϕ which induces all these isomorphisms of stalks.

For example, consider the sheaves $\mathcal{O}_{\mathbb{P}^1}(d)$ on $\mathbb{P}^1$. Exercise 7.1.9 shows that for every point $x \in X$ and for every integer d the stalk $\mathcal{O}_{\mathbb{P}^1}(d)_x$ is a free rank 1 module over $\mathcal{O}_{\mathbb{P}^1,x}$. In particular for every integer d these stalks are isomorphic abelian groups. However, the sheaves $\mathcal{O}_{\mathbb{P}^1}(d)$ are definitely not isomorphic to each other for different values of d – for example, when $d \geq 0$ then the space of global sections of $\mathcal{O}_{\mathbb{P}^1}(d)$ is a $\mathbb{K}$-vector space of dimension $d+1$.

7.3.4 Cokernels

Definition 7.3.16. Let $\phi : \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves of abelian groups. The cokernel presheaf $\mathrm{cok}_{pre}(\phi)$ associates to each open set U the cokernel of $\phi(U)$ and to each inclusion of open subsets $V \subset U$ the quotient of the restriction map $\rho_{\mathcal{G},U,V}$. The cokernel of ϕ, denoted $\mathrm{cok}(\phi)$, is the sheafification of the cokernel presheaf.

Remark 7.3.17. Note that $\mathrm{cok}(\phi)$ comes equipped with a quotient morphism of sheaves $q : \mathcal{G} \to \mathrm{cok}(\phi)$ which is the composition of the quotient map to the cokernel presheaf and the sheafification map.

Exercise 7.3.18. Let $\phi : \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves. Prove that $\mathrm{cok}(\phi)_x \cong \mathrm{cok}(\phi_x)$ and that this isomorphism is the map on stalks induced by the quotient map $q : \mathcal{G} \to \mathrm{cok}(\phi)$.

The cokernel presheaf need not be a sheaf.

Exercise 7.3.19. Find a morphism of sheaves $\phi : \mathcal{F} \to \mathcal{G}$ such that the cokernel presheaf is not a sheaf. Can you find an example where the identity axiom fails? Where the gluing axiom fails?

7.3.5 Exercises

Exercise 7.3.20. Let X be a topological space. Suppose that $\phi_1, \phi_2 : \mathcal{F} \to \mathcal{G}$ are two morphisms of sheaves such that for every $x \in X$ the induced maps on stalks satisfy $\phi_{1,x} = \phi_{2,x}$. Prove that $\phi_1 = \phi_2$.

Exercise 7.3.21. Let $X = \mathbb{A}^1_{\mathbb{K}}$. Let $\mathcal{G}$ be the skyscraper sheaf at the origin with value $\mathbb{K}[x]_{(x)}$. Show that the localization maps induce a surjection morphism of sheaves $\phi : \mathcal{O}_X \to \mathcal{G}$. Show however that there is no open neighborhood U of the origin such that the map $\phi(U) : \mathcal{O}_X(U) \to \mathcal{G}(U)$ is surjective. (Can you leverage this idea to find a surjective morphism of sheaves such that there is no open set U with $\phi(U)$ surjective?)

Exercise 7.3.22. Consider the morphism of sheaves $\mathcal{O}_{\mathbb{P}^1}(-1) \to \mathcal{O}_{\mathbb{P}^1}$ defined by multiplication by x. Identify the image and cokernel of this map.

Exercise 7.3.23. Consider the morphism of sheaves $\mathcal{O}_{\mathbb{P}^1}(-d) \to \mathcal{O}_{\mathbb{P}^1}$ defined by multiplication by a degree d homogeneous polynomial f. Show that the cokernel is supported on $V_+(f)$.

7.4 Category of sheaves

Let X be a fixed topological space. In this section we show that the category $\mathbf{Sh}(X)$ is an abelian category and discuss the theory of exact sequences in this category.

7.4.1 Abelian category

The category of presheaves of abelian groups on X carries the structure of an abelian category in a natural way. As discussed in the previous section, one can transport constructions from $\mathbf{Ab}$ to $\mathbf{PreSh}(X)$ by performing them simultaneously for each open set U.

As we saw before, the corresponding constructions for sheaves are more difficult. However, the nice properties of the sheafification functor allow us to turn presheaf constructions into sheaf constructions without losing the "universality" of the construction. In particular, we will see that the category of sheaves of abelian groups on X is also an abelian category.

First we verify that $\mathbf{Sh}(X)$ is an additive category. Given any two sheaves $\mathcal{F}, \mathcal{G}$ we must endow the space $\mathrm{Hom}(\mathcal{F}, \mathcal{G})$ with the structure of an abelian group. Given two morphisms $\phi, \psi : \mathcal{F} \to \mathcal{G}$ we define

$$(\phi + \psi)(U) := \phi(U) + \psi(U)$$

for any open set U. It is clear that $\phi + \psi$ is still compatible with restriction so that this function is a well-defined sheaf morphism. It is also clear that composition of morphisms is bilinear. Finally, the zero object in $\mathbf{Sh}(X)$ is the 0 sheaf, and the biproduct is given by the direct sum operation $\oplus$ from Exercise 7.1.13.

Exercise 7.4.1. Verify that $\oplus$ (as defined in Exercise 7.1.13) is a biproduct in the category $\mathbf{Sh}(X)$.

7.4.2 Kernels and cokernels

We next check that our definitions of kernel and cokernel agree with the categorical definitions.

Lemma 7.4.2. *Let $\phi : \mathcal{F} \to \mathcal{G}$ denote a morphism of sheaves. Then the inclusion $i : \ker(\phi) \to \mathcal{F}$ is the equalizer of ϕ and the zero map.*

Proof. It is clear that $\phi \circ i = 0 \circ i$, since both are the zero map. Suppose that $\psi : \mathcal{H} \to \mathcal{F}$ also satisfies $\phi \circ \psi = 0 \circ \psi$. Then for every open set U the image of $\psi(U) : \mathcal{H}(U) \to \mathcal{F}(U)$ must lie in $\ker(\phi(U))$. In this way we see that ψ factors through $i : \ker(\phi) \to \mathcal{F}$ and it is clear this factorization is unique. □

Lemma 7.4.3. *Let $\phi : \mathcal{F} \to \mathcal{G}$ denote a morphism of sheaves. Then the quotient map $q : \mathcal{G} \to \mathrm{cok}(\phi)$ is the coequalizer of ϕ and the zero map.*

Proof. Let $\mathcal{Q}$ denote the presheaf cokernel and let $\widetilde{q} : \mathcal{G} \to \mathcal{Q}$ denote the presheaf cokernel map. Since the composition of ϕ with $\widetilde{q}$ is the zero map, the composition of ϕ with q is also the zero map.

Suppose that $\psi : \mathcal{G} \to \mathcal{H}$ satisfies $\psi \circ \phi = \psi \circ 0$. Then for every open set U the map $\psi(U) : \mathcal{G}(U) \to \mathcal{H}(U)$ admits a unique factorization through $\widetilde{q}(U) : \mathcal{G}(U) \to \mathcal{Q}(U)$. In this way we see that as a map of presheaves ψ admits a unique factorization through $\widetilde{q}$. Using the universal property of sheafification in Proposition 7.2.8, we see that as a map of sheaves ψ admits a unique factorization through q. □

7.4.3 Monomorphisms and epimorphisms

The next step is to identify the monomorphisms and epimorphisms in $\mathbf{Sh}(X)$.

Proposition 7.4.4. *Let $\phi : \mathcal{F} \to \mathcal{G}$ be a morphism of sheaves. Then:*

(1) ϕ is a monomorphism if and only if it is injective.

(2) ϕ is an epimorphism if and only if it is surjective.

Proof. (1) First suppose that ϕ is injective. Suppose that $\psi_1, \psi_2 : \mathcal{H} \to \mathcal{F}$ satisfy that $\phi \circ \psi_1 = \phi \circ \psi_2$. In particular, for every point $x \in X$ we have $\phi_x \circ \psi_{1,x} = \phi_x \circ \psi_{2,x}$. Since injections are monomorphisms in $\mathbf{Ab}$, we see that $\psi_{1,x} = \psi_{2,x}$ for every point x. We conclude that $\psi_1 = \psi_2$ by Exercise 7.3.20.

Conversely, suppose that ϕ is not injective. Then $\ker(\phi)$ is not the zero sheaf. Since the composition of ϕ with the zero map and the inclusion map from $\ker(\phi)$ to $\mathcal{F}$ are the same, we see that ϕ is not a monomorphism.

(2) First suppose that ϕ is surjective. Suppose that $\psi_1, \psi_2 : \mathcal{G} \to \mathcal{H}$ satisfy that $\psi_1 \circ \phi = \psi_2 \circ \phi$. In particular, for every point $x \in X$ we have $\psi_{1,x} \circ \phi_x = \psi_{2,x} \circ \phi_x$. Since surjections are epimorphisms in $\mathbf{Ab}$, we see that $\psi_{1,x} = \psi_{2,x}$ for every point x. We conclude that $\psi_1 = \psi_2$ by Exercise 7.3.20.

Conversely, suppose that ϕ is not surjective. By Exercise 7.3.10 and Exercise 7.3.18 we see that $\mathrm{cok}(\phi)$ is not the zero sheaf. Since the composition of ϕ with the zero map and the quotient map from $\mathcal{G}$ to $\mathrm{cok}(\phi)$ are the same, we see that ϕ is not an epimorphism. □

In particular, by combining Proposition 7.4.4 with Exercise 7.3.5 and Exercise 7.3.10 we see that the property of being a monomorphism or epimorphism can be checked on the level of stalks.

The final step is to check that every monomorphism is a kernel and that every epimorphism is a cokernel. This follows from what we have already shown. First, we claim that if $\phi : \mathcal{F} \to \mathcal{G}$ is a monomorphism then ϕ is the kernel of the quotient map $q : \mathcal{G} \to \mathrm{cok}(\phi)$. Indeed, this can be checked on the level of stalks, where it follows from Proposition 7.4.4, Exercise 7.3.5, and Exercise 7.3.18. The dual argument is similar. Thus we have shown:

Theorem 7.4.5. *The category* $\mathbf{Sh}(X)$ *is abelian.*

Before moving on, we point out that the image sheaf $\mathrm{im}(\phi)$ coincides with the notion of "image" in an abelian category – this follows from a stalk computation. Thus any morphism of sheaves $\phi : \mathcal{F} \to \mathcal{G}$ factors as an epimorphism $\mathcal{F} \to \mathrm{im}(\phi)$ composed with a monomorphism $\mathrm{im}(\phi) \to \mathcal{G}$.

7.4.4 Exact sequences

As with any abelian category, one of the fundamental notions in $\mathbf{Sh}(X)$ is that of an exact sequence. A complex of sheaves

$$\dots \xrightarrow{\phi_{i-2}} \mathcal{F}_{i-1} \xrightarrow{\phi_{i-1}} \mathcal{F}_i \xrightarrow{\phi_i} \mathcal{F}_{i+1} \xrightarrow{\phi_{i+1}} \dots$$

is exact if $\mathrm{im}(\phi_{i-1}) = \ker(\phi_i)$ for every i.

Exercise 7.4.6. Show that a complex of sheaves is exact if and only if for every point $x \in X$ the induced complex of stalks at x is exact.

Much of the course will be devoted to studying how various functors on $\mathbf{Sh}(X)$ interact with exact sequences. The following example will be particularly important for us.

Definition 7.4.7. Let X be a topological space. The global sections functor $\mathbf{Sh}(X) \to \mathbf{Ab}$ assigns to any sheaf $\mathcal{F}$ the abelian group $\mathcal{F}(X)$ and to any morphism of sheaves ϕ the homomorphism ϕ_X. This functor is denoted by $\Gamma(X, -) : \mathbf{Sh}(X) \to \mathbf{Ab}$.

One could of course use any open set in place of X, but the global sections functor is the most important one.

Lemma 7.4.8. *Let* X *be a topological space. The global sections functor is left-exact.*

Proof. Suppose that $0 \to \mathcal{F}' \xrightarrow{\phi} \mathcal{F} \xrightarrow{\psi} \mathcal{F}'' \to 0$ is a short exact sequence of sheaves. Since ϕ is injective, $\mathcal{F}'(X) \to \mathcal{F}(X)$ is injective.

It only remains to show exactness in the middle, i.e. that $\mathrm{im}(\phi(X)) = \ker(\psi(X))$. It is clear that $\mathrm{im}(\phi(X)) \subset \ker(\psi(X))$ since $\psi \circ \phi = 0$. Conversely, suppose that $f \in \mathcal{F}(X)$ lies in the kernel of $\psi(X)$. This means that f is also an element of $\ker(\phi)(X)$. Due to the exactness of the sequence, this implies that $f \in \mathrm{im}(\phi)(X)$. In other words, there is an open cover $\{V_i\}$ of X and elements $g_i \in \mathcal{F}'(V_i)$ such that $f|_{V_i} = \phi(V_i)(g_i)$. Since ϕ is injective, the restriction of g_i and g_j to $V_i \cap V_j$ agree. Applying the gluing axiom, we conclude that there is an element $g \in \mathcal{F}'(X)$ such that $\phi(X)(g) = f$. $\square$

The following example shows that the global sections functor is not exact.

Example 7.4.9. Consider the map $\phi : \mathcal{O}_{\mathbb{P}^1}(-2) \to \mathcal{O}_{\mathbb{P}^1}$ defined by multiplying by the equation xy. Then ϕ is an inclusion and its cokernel is the union $\mathbb{K}(p) \oplus \mathbb{K}(q)$ of the skyscraper sheaves with value $\mathbb{K}$ at the two points $p, q \in V_+(xy)$. If we take the exact sequence

$$0 \to \mathcal{O}_{\mathbb{P}^1}(-2) \to \mathcal{O}_{\mathbb{P}^1} \to \mathbb{K}(p) \oplus \mathbb{K}(q) \to 0$$

and apply the global sections functor we obtain

$$\begin{array}{ccccccc} 0 \longrightarrow & \mathcal{O}_{\mathbb{P}^1}(-2)(\mathbb{P}^1) & \longrightarrow & \mathcal{O}_{\mathbb{P}^1}(\mathbb{P}^1) & \longrightarrow & (\mathbb{K}(p) \oplus \mathbb{K}(q))(\mathbb{P}^1) \\ & \| & & \| & & \| \\ & 0 & & \mathbb{K} & & \mathbb{K}^2 \end{array}$$

In particular the rightmost map in this sequence is not surjective.

7.4.5 Exercises

Exercise 7.4.10. Let X be a topological space. Suppose that to each point $x \in X$ we assign a divisible abelian group Q_x. Define the sheaf $\mathcal{Q}$ by assigning to any open set U the product $\prod_{x \in U} Q_x$ and to any inclusion $V \subset U$ the corresponding projection map. Prove that $\mathcal{Q}$ is an injective object in $\mathbf{Sh}(X)$.

Exercise 7.4.11 ([Cla])**.** Let X be a manifold of dimension ≥ 1. Fix a point $x \in X$. For any open neighborhood $V \subset X$ define the sheaf $\mathbb{Z}_V$ via the prescription:

$$\mathbb{Z}_V(U) = \begin{cases} \mathbb{Z}^{\pi_0(U)} \text{ if } U \subset V \\ 0 \text{ if } U \not\subset V \end{cases}$$

with the obvious restriction maps.

(1) Show that for any open neighborhood V of x there is a surjection $\rho_V : \mathbb{Z}_V \to \mathbb{Z}(x)$ where $\mathbb{Z}(x)$ denotes the skyscraper sheaf at x with value $\mathbb{Z}$.

(2) Use the surjections ρ_V to show that there is no projective object in $\mathbf{Sh}(X)$.

7.5 Pushforward and inverse image

Suppose that $f : X \to Y$ is a continuous map between topological spaces. In this section we analyze how to pass sheaves of abelian groups back and forth between X and Y. The pushforward of a sheaf is easy to define, but can be tricky to work with. The pullback is more difficult to define, but it turns out to have better geometric behavior.

Definition 7.5.1. Let $\mathcal{F}$ be a presheaf of abelian groups on X. We define the presheaf $f_*\mathcal{F}$ which assigns to the open set $V \subset Y$ the group $\mathcal{F}(f^{-1}V)$ and assigns to an inclusion of open sets the restriction map for the preimages. When $\mathcal{F}$ is a sheaf then $f_*\mathcal{F}$ is a sheaf as well – the identity and gluing properties for $f_*\mathcal{F}$ follow immediately from the corresponding statements for $\mathcal{F}$.

To any morphism $\phi : \mathcal{F} \to \mathcal{G}$ on X we define the morphism $f_*\phi : f_*\mathcal{F} \to f_*\mathcal{G}$ on Y via the prescription $(f_*\phi)(V) := \phi(f^{-1}V)$. It is then clear that f_* defines a functor from $\mathbf{Sh}(X)$ to $\mathbf{Sh}(Y)$.

Example 7.5.2. Let X be a topological space. Suppose that $i : U \to X$ is the inclusion of an open subset. For any sheaf $\mathcal{F}$ on U, the pushforward sheaf $i_*\mathcal{F}$ assigns to any open set $V \subset X$ the abelian group $\mathcal{F}(U \cap V)$.

Example 7.5.3. Let X be a topological space with a sheaf $\mathcal{F}$. Suppose that $i : x \to X$ is the inclusion of a point. For any sheaf $\mathcal{F}$ on x, the pushforward sheaf $i_*\mathcal{F}$ is the skyscraper sheaf at x with value $\mathcal{F}(x)$.

The inverse image of a sheaf is harder to define but tends to be better behaved than the pushforward. (This is an avatar of the fact that the natural operation for vector bundles is the pullback operation.) We will reserve the term "pullback" for a later construction that is only defined for sheaves of modules. As demonstrated by Exercise 7.2.10.(2) we cannot expect the "naive pullback" of a sheaf to be a sheaf, so we should anticipate a sheafification when defining this operation.

Definition 7.5.4. Let $\mathcal{G}$ be a sheaf of abelian groups on Y. We define the inverse image presheaf $f^{-1}_{pre}\mathcal{G}$ which assigns to every open set U in X the abelian group

$$\varinjlim_{V \supset f(U)} \mathcal{G}(V)$$

and to every inclusion of open sets the homomorphism induced by the universal property of the direct limit applied to the various restriction maps for $\mathcal{G}$. The inverse image sheaf $f^{-1}\mathcal{G}$ is the sheafification of the inverse image presheaf.

To any morphism $\phi : \mathcal{F} \to \mathcal{G}$ on Y we define the morphism $f^{-1}\phi : f^{-1}\mathcal{F} \to f^{-1}\mathcal{G}$ on Y as follows. Fix an open set U on X. For every open set $V' \supset f(U)$ we have a composition

$$\mathcal{F}(V') \xrightarrow{\phi} \mathcal{G}(V') \to \varinjlim_{V \supset f(U)} \mathcal{G}(V).$$

By the universal property of direct limits, we obtain a morphism of inverse image presheaves $f_{pre}^{-1}\phi : f_{pre}^{-1}\mathcal{F} \to f_{pre}^{-1}\mathcal{G}$. By applying the sheafification functor we obtain the desired morphism $f^{-1}\phi$. With this definition f^{-1} becomes a functor from $\mathbf{Sh}(Y)$ to $\mathbf{Sh}(X)$.

Remark 7.5.5. While the inverse image presheaf always satisfies the identity property (check!), it can fail to be a sheaf even in very simple situations. For example, if the map $f : X \to Y$ contracts X to a point, the inverse image presheaf will be a constant presheaf (which is usually not a sheaf).

Example 7.5.6. Let X be a topological space with a sheaf $\mathcal{F}$. Suppose that $i : U \to X$ is the inclusion of an open subset. Then the inverse image sheaf $i^{-1}\mathcal{F}$ is isomorphic to $\mathcal{F}|_U$.

Example 7.5.7. Let X be a topological space with a sheaf $\mathcal{F}$. If we let $i : x \to X$ denote the inclusion of a point, then the inverse image sheaf $i^{-1}\mathcal{F}$ is isomorphic to the stalk $\mathcal{F}_x$.

7.5.1 Stalks

The inverse image construction behaves well with respect to stalks.

Lemma 7.5.8. *Let $f : X \to Y$ be a continuous map of topological spaces. Let $\mathcal{G}$ be a sheaf of abelian groups on Y. For any point $x \in X$ we have $(f^{-1}\mathcal{G})_x \cong \mathcal{G}_{f(x)}$.*

Proof. First we compute the stalk of the inverse image presheaf. The key is the realization that an open set $V \subset Y$ will contain the f-image of some open neighborhood $U \ni x$ if and only if V contains $f(x)$. Thus

$$\begin{aligned} f_{pre}^{-1}\mathcal{G}_x &= \varinjlim_{U \ni x} \left(\varinjlim_{V \supset f(U)} \mathcal{G}(V) \right) \\ &\cong \varinjlim_{V \ni f(x)} \mathcal{G}(V) \\ &\cong \mathcal{G}_{f(x)}. \end{aligned}$$

By Exercise 7.2.6 the stalk of the inverse image presheaf at x is isomorphic to the stalk of the inverse image sheaf. $\square$

The behavior of stalks under pushforward is more subtle. (This is another indication that the inverse image is more natural in geometric contexts.) Suppose that $f : X \to Y$ is a continuous map and that $\mathcal{F}$ is a sheaf of abelian groups on X. For a point $y \in Y$ we have

$$(f_*\mathcal{F})_y = \varinjlim_{V \ni y} \mathcal{F}(f^{-1}V)$$

If we let F denote the fiber over y, there is a natural morphism from $(f_*\mathcal{F})_y$ to $\lim_{U \supset F} \mathcal{F}(U)$ but this map need not be an isomorphism since there may be open neighborhoods of F which do not contain the preimage of any open neighborhood of y. Thus there is no way to define the stalk $(f_*\mathcal{F})_y$ using only the fiber F and the topology of X.

7.5.2 Adjointness

Theorem 7.5.9. *Let $f : X \to Y$ be a continuous map of topological spaces. Then f^{-1} and f_* form an adjoint pair: for any sheaves $\mathcal{F}$ on X and $\mathcal{G}$ on Y we have a bijection*

$$\mathrm{Hom}_{\mathbf{Sh}(X)}(f^{-1}\mathcal{G}, \mathcal{F}) \cong \mathrm{Hom}_{\mathbf{Sh}(Y)}(\mathcal{G}, f_*\mathcal{F})$$

natural in both entries.

This adjunction takes some work to set up carefully. We will just describe the bijection between Hom spaces, leaving the naturality as an exercise for the dedicated reader.

Proof. Suppose given a morphism $\phi_X : f^{-1}\mathcal{G} \to \mathcal{F}$. By precomposing with the sheafification map, we obtain a morphism of presheaves $\phi_{X,pre} : f^{-1}_{pre}\mathcal{G} \to \mathcal{F}$. In particular, for every open set $V \subset Y$ we obtain a map $f^{-1}_{pre}\mathcal{G}(f^{-1}V) \to \mathcal{F}(f^{-1}V)$. Furthermore these maps are compatible with the restriction maps for open sets of Y. Using the isomorphism $\mathcal{G}(V) \cong f^{-1}_{pre}\mathcal{G}(f^{-1}V)$ we obtain a morphism $\mathcal{G} \to f_*\mathcal{F}$.

Suppose given a morphism $\phi_Y : \mathcal{G} \to f_*\mathcal{F}$. In other words, for every open set $V \subset Y$ we have a map $\mathcal{G}(V) \to f_*\mathcal{F}(V)$. For any open subset U of X, we obtain a map

$$\lim_{V \supset f(U)} \mathcal{G}(V) \to \lim_{V \supset f(U)} f_*\mathcal{F}(V) = \lim_{f^{-1}V \supset U} \mathcal{F}(f^{-1}V) \to \mathcal{F}(U)$$

where the last map is defined by the restriction maps. It is clear that these maps are compatible with restriction for open sets in X. Thus we get a map of presheaves $f^{-1}_{pre}\mathcal{G} \to \mathcal{F}$. By the universal property of sheafifcation, we obtain a map of sheaves $f^{-1}\mathcal{G} \to \mathcal{F}$.

These two constructions are inverses. □

Note that f^{-1} is a left adjoint and f_* is a right adjoint in this adjoint pair of functors. In particular, they have good exactness properties.

Proposition 7.5.10. *Let $f : X \to Y$ be a continuous map of topological spaces. Then f_* is left-exact and f^{-1} is exact.*

Proof. Theorem 7.5.9 and Corollary 7.0.5 imply that f_* is left-exact and f^{-1} is right-exact. To see that f^{-1} is exact, we combine Lemma 7.5.8 (showing that $(f^{-1}\mathcal{G})_x \cong \mathcal{G}_{f(x)}$) and Exercise 7.4.6 (showing that exactness can be checked on the level of stalks.) □

Exercise 7.5.11. Prove that f_* is left-exact "by hand" without appealing to the general properties of adjoint functors.

7.5.3 Extension by zero

We saw previously that f^{-1} is always a left adjoint. In very special situations it can also be a right adjoint functor. Our next goal is to describe its paired functor.

Definition 7.5.12. Let X be a topological space and let $i : U \to X$ be the inclusion of an open subset. Given any sheaf of abelian groups $\mathcal{F}$ on U, we define the extension by zero presheaf $i_{!,pre}\mathcal{F}$ by setting

$$i_{!,pre}\mathcal{F}(V) = \begin{cases} \mathcal{F}(V) & \text{if } V \subset U \\ 0 & \text{if } V \not\subset U \end{cases}$$

We define the extension by zero sheaf $i_!\mathcal{F}$ by sheafifying the extension by zero presheaf.

The stalks of $i_!\mathcal{F}$ are identically zero for points outside of the open set U and are equal to the stalks of $\mathcal{F}$ for points in U. Note that $i_!$ is not the same as i_* for inclusions of open sets as in Example 7.5.2. (What is an example where the two constructions are different? Can you think of an example where they yield the same thing?)

Exercise 7.5.13. Let X be a topological space and let $i : U \to X$ be the inclusion of an open set. Prove that $i_!$ is a left adjoint and i^{-1} is a right adjoint which together form an adjoint pair.

7.5.4 Exercises

Exercise 7.5.14. Let $f : X \to Y$ and $g : Y \to Z$ be continuous maps of topological spaces. Prove that:

(1) $(g \circ f)_* = g_* \circ f_*$.

(2) $(g \circ f)^{-1} = f^{-1} \circ g^{-1}$. (Hint: since these constructions involve a sheafification, it is not easy to directly compare the values of the two constructions on open sets. Instead, you should construct a morphism between them and show that it induces an isomorphism of stalks.)

Exercise 7.5.15. Let X be a topological space. Let $i : U \to X$ be the inclusion of an open subset and let $j : Z \to X$ be the inclusion of the complement of U. Show that for any sheaf $\mathcal{F}$ on X we have an exact sequence

$$0 \to i_!i^{-1}\mathcal{F} \to \mathcal{F} \to j_*j^{-1}\mathcal{F} \to 0.$$

Calculate this sequence explicitly when $X = \mathbb{A}^n$, Z is the origin, and the sheaf is $\mathcal{O}_{\mathbb{A}^n}$.

Exercise 7.5.16. Finish the proof of Theorem 7.5.9 by proving naturality of the bijection.

7.6 Gluing sheaves

In this section we discuss how one can construct sheaves by gluing "local data" on a topological space. Suppose that X is a topological space and that $\mathcal{B}$ is a base for the topology on X. In Definition 1.9.1 we defined the notion of a $\mathcal{B}$-sheaf on X.

Definition 7.6.1. Let X be a topological space and let $\mathcal{B} = \{V_i\}$ be a base for the topology. A $\mathcal{B}$-sheaf $\widetilde{\mathcal{F}}$ assigns to every open set $V_i \in \mathcal{B}$ an abelian group $\widetilde{\mathcal{F}}(V_i)$ and to each inclusion $V_i \subset V_j$ of open sets in $\mathcal{B}$ a restriction map $\widetilde{\rho}_{V_j,V_i}$ such that the following properties hold:

(1) $\widetilde{\mathcal{F}}(\emptyset) = 0$.

(2) The assignments $\widetilde{\mathcal{F}}, \widetilde{\rho}$ define a contravariant functor from the category of open subsets of X contained in $\mathcal{B}$ (with morphisms = inclusions) to the category of abelian groups.

(3) For any open set $V_i \in \mathcal{B}$ and any open cover of V_i by elements in $\mathcal{B}$ the identity and gluing axioms hold.

In Section 1.9 we showed:

- Given a $\mathcal{B}$-sheaf $\widetilde{\mathcal{F}}$, there is a unique sheaf $\mathcal{F}$ on X such that $\mathcal{F}(U) = \widetilde{\mathcal{F}}(U)$ for every U in $\mathcal{B}$ (Theorem 1.9.2).

- Given a morphism of $\mathcal{B}$-sheaves $\widetilde{\phi} : \widetilde{\mathcal{F}} \to \widetilde{\mathcal{G}}$ we obtain an induced morphism of the corresponding sheaves $\phi : \mathcal{F} \to \mathcal{G}$ that agrees locally with $\widetilde{\phi}$ (Exercise 1.9.4).

Altogether we obtain:

Theorem 7.6.2. *Let X be a topological space equipped with a base $\mathcal{B}$ for the topology. There is an equivalence between the category of $\mathcal{B}$-sheaves and the category of sheaves.*

The categorical inverse of the "gluing functor" discussed above is the forgetful functor from the category of sheaves to the category of $\mathcal{B}$-sheaves. The most important special case of this construction is:

Corollary 7.6.3. *Let X be a topological space equipped with an open cover $\{U_i\}$. Suppose that for each index i we have a sheaf $\mathcal{F}_i$ on U_i. Suppose furthermore that for every pair of indices i, j we have an isomorphism*

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

and that ϕ_{ii} is the identity map, $\phi_{ij} = \phi_{ji}^{-1}$ and $\phi_{jk} \circ \phi_{ij} = \phi_{ik}$ (as isomorphisms of sheaves on $U_i \cap U_j \cap U_k$). Then there is a sheaf $\mathcal{F}$ on X (unique up to isomorphism) such that $\mathcal{F}|_{U_i}$ is isomorphic to $\mathcal{F}_i$.

We proved this in Corollary 1.9.5. The key point is that the different $\mathcal{F}_i$ give us many different ways of defining the group of sections associated to a given open set. The "cocycle condition" – that is, the condition on the maps ϕ_{ij} – is exactly what we need to canonically identify these various choices.

Just as we can glue sheaves on open sets, we can glue morphisms of sheaves on open sets, and one can again view this as a special case of Theorem 7.6.2.

Corollary 7.6.4. *Let X be a topological space equipped with two sheaves $\mathcal{F}, \mathcal{G}$. Let $\{U_i\}$ be an open cover of X. Suppose that for every open set U_i we are given a morphism $\phi_i : \mathcal{F}|_{U_i} \to \mathcal{G}|_{U_i}$ and that furthermore we have*

$$\phi_i|_{U_i \cap U_j} = \phi_j|_{U_i \cap U_j}.$$

Then there is a unique morphism of sheaves $\phi : \mathcal{F} \to \mathcal{G}$ such that $\phi|_{U_i} = \phi_i$ for every i.

The construction of sheaves through local data is quite common and it is important to feel comfortable with this perspective.

Example 7.6.5. Let's construct the sheaves $\mathcal{O}(d)$ on $\mathbb{P}^1$ via local data. Consider the open cover of $\mathbb{P}^1$ by the two coordinate charts isomorphic to $\mathbb{A}^1$. We set $U_0 = D_{+,x_0}$ (with ring of functions $\mathbb{K}[\frac{x_1}{x_0}]$) and $U_1 = D_{+,x_1}$ (with ring of functions $\mathbb{K}[\frac{x_0}{x_1}]$). Note that the intersection $U_0 \cap U_1$ is the affine scheme corresponding to the ring $\mathbb{K}[\frac{x_0}{x_1}, \frac{x_1}{x_0}]$.

We plan to glue together the structure sheaf $\mathcal{O}_{\mathbb{P}^1}|_{U_0}$ on U_0 and the structure sheaf $\mathcal{O}_{\mathbb{P}^1}|_{U_1}$ on U_1. According to Theorem 7.6.3, we only need to specify the data of an isomorphism

$$\phi_{01} : \mathcal{O}_{\mathbb{P}^1}|_{U_0 \cap U_1} \to \mathcal{O}_{\mathbb{P}^1}|_{U_0 \cap U_1}.$$

The easiest way to obtain an isomorphism is to multiply by a unit in the ring $\mathcal{O}_{\mathbb{P}^1}(U_0 \cap U_1)$. We define the morphism of sheaves ϕ_d which assigns to open set $V \subset U_0 \cap U_1$ the multiplication by $(\frac{x_0}{x_1})^d$. (These maps clearly commute with restriction and thus give a valid morphism.) Since $\frac{x_0}{x_1}$ is invertible on $U_0 \cap U_1$ the morphism ϕ_d is invertible. By Corollary 7.6.3 the isomorphism ϕ_d yields a corresponding sheaf $\mathcal{F}_d$.

Let's verify that $\mathcal{F}_d$ is isomorphic to $\mathcal{O}(d)$. For any open set $U \subset \mathbb{P}^1$, we have

$$\mathcal{F}_d(U) = \{\, (g_0, g_1) \in \mathcal{O}_{\mathbb{P}^1}(U \cap U_0) \times \mathcal{O}_{\mathbb{P}^1}(U \cap U_1) \mid \phi_d(g_0|_{U \cap U_0 \cap U_1}) = g_1|_{U \cap U_0 \cap U_1} \,\}.$$

Since ϕ_0 is the identity map, in this case the condition matches the usual gluing axiom for sheaves so that $\mathcal{F}_0 \cong \mathcal{O}_{\mathbb{P}^1}$. However, when ϕ_d is not the identity map then the gluing condition is different from the usual one. Precisely, the condition is that $x_0^d g_0 = x_1^d g_1$ when both sides are considered as degree d elements in $\mathbb{K}(x_0, x_1)$. Thus the map $\mathcal{O}(d)(U) \to \mathcal{F}_d(U)$ sending

$$g \mapsto \left(\frac{g}{x_0^d}, \frac{g}{x_1^d} \right)$$

defines the desired isomorphism.

Exercise 7.6.6. Consider the open cover of $\mathbb{P}^n$ by affine charts $U_i = D_{+,x_i}$. Show that $\mathcal{O}_{\mathbb{P}^n}(d)$ is isomorphic to the sheaf $\mathcal{F}_d$ defined by gluing the structure sheaves using the local homomorphisms ϕ_{ij} which is multiplication by $(\frac{x_i}{x_j})^d$.

7.6.1 Constructing a sheaf on stalks

There is an alternative perspective on sheaves coming from the "espace étalé" construction. We won't need this perspective in the future, but we include it here for the sake of completeness.

Let X be a topological space with a sheaf $\mathcal{F}$. Exercise 7.1.6 shows that a section $f \in \mathcal{F}(U)$ is determined by its restriction to the stalks $\mathcal{F}_x$ for stalks $x \in U$. In fact, one can recover $\mathcal{F}(U)$ from the combined information of the stalks and the "local sections" as follows.

Let $\mathfrak{F}$ denote the disjoint union of the stalks $\sqcup_{x\in X}\mathcal{F}_x$. We have a natural map $\pi : \mathfrak{F} \to X$. For any open set $U \subset X$ and any section $s \in \mathcal{F}(U)$ we obtain a subset $\{\rho_{U,x}(s)\}_{x\in U}$ of $\mathfrak{F}$. We give $\mathfrak{F}$ the topology which has these sets as a base.

Exercise 7.6.7. Given any open subset $U \subset X$, prove that $\mathcal{F}(U)$ can be identified with the space of continuous maps $\sigma : U \to \mathfrak{F}$ which are sections of π. Show that under this identification restriction maps $\rho_{U,V}$ can be defined as the restriction of continuous maps.

One advantage of the espace étalé perspective is that both the pullback functor and the sheafification functor admit easy and natural descriptions.

Chapter 8

Schemes

The theory of schemes is closely analogous to the theory of manifolds. However, our basic objects will be determined by algebra instead of geometry. To any ring R we can associate a geometric space known as the spectrum $\mathrm{Spec}(R)$. A scheme is defined to be a space that is locally isomorphic to the spectrum of a ring.

The theory of manifolds is nicest for spaces which have some additional properties: second countable, Hausdorff, compact, etc. Most of the chapter is dedicated to identifying analogous constructions in algebraic geometry.

We briefly remind the reader of our setting. All rings R are commutative unital rings; all homomorphisms of rings are unital (i.e. send 1 to 1). The 0 ring is a valid ring; it does not admit a homomorphism to any ring besides itself but it receives homomorphisms from every ring. We will accept the Axiom of Choice, so in particular every ideal is contained in a maximal ideal.

8.0.1 Algebraic preliminaries

We briefy review a few finiteness assumptions for rings.

Definition 8.0.1. A homomorphism of rings $\phi : S \to R$ is of finite type if it realizes R as a finitely generated S-algebra. In other words, ϕ has finite type if for some positive integer n there exists a surjection $S[x_1, \ldots, x_n] \to R$ that restricts to ϕ.

A homomorphism of rings $\phi : S \to R$ is of finite presentation if for some positive integer n there exists a surjection $S[x_1, \ldots, x_n] \to R$ whose kernel is generated by a finite set of polynomials.

Note that when S is a Noetherian ring the two definitions are equivalent. Let P denote the property "of finite type" or "of finite presentation". Then homomorphisms of type P have many nice properties. The most relevant properties for us are:

(1) Homomorphisms of type P are preserved under composition.

(2) Suppose that $\phi : S \to R$ is a homomorphism and that $\{r_i\}$ is a set of elements which generate the unit ideal in R. Then $\phi : S \to R$ satisfies P if and only if $\phi : S \to R_{r_i}$ satisfies P for every i.

(3) Homomorphisms of type P are preserved under tensor product: if $\phi : S \to R$ satisfies P and $S \to T$ is a ring homomorphism then $T \to R \otimes_S T$ satisfies P.

We will also use a much stronger version of finiteness:

Definition 8.0.2. A homomorphism of rings $\phi : S \to R$ is finite if it realizes R as a finitely generated S-module. In other words, there is some positive integer n such that there exists a surjection of S-modules $S^{\oplus n} \to R$.

A homomorphism of rings $\phi : S \to R$ is said to be integral if every element of R satisfies a monic equation with coefficients in S.

It turns out that a homomorphism is finite if and only if it is integral and has finite type. Let P denote the property "finite" or "integral". Then homomorphisms of type P satisfy a (slightly different) list of properties:

(1) Homomorphisms of type P are preserved under composition.

(2) Suppose that $\phi : S \to R$ is a homomorphism and that $\{s_j\}$ is a set of elements which generate the unit ideal in R. Then $\phi : S \to R$ satisfies P if and only if $\phi : S_{s_j} \to R_{\phi(s_j)}$ satisfies P for every j.

(3) Homomorphisms of type P are preserved under tensor product: if $\phi : S \to R$ satisfies P and $S \to T$ is a ring homomorphism then $T \to R \otimes_S T$ satisfies P.

Note in particular that the localization property of finite morphisms is weaker than the corresponding property for finite type morphisms.

8.1 Spectrum of a ring

The building blocks for schemes are spectrums of rings.

Definition 8.1.1. Let R be a ring. The spectrum of R, denoted $\mathrm{Spec}(R)$, is the set of prime ideals in R. For any ideal I in R, we define the vanishing locus $V(I) \subset \mathrm{Spec}(R)$ to be the set of prime ideals which contain I. These form the closed sets in a topology called the Zariski topology.

Exercise 8.1.2. Prove that the $V(I)$ satisfy the necessary properties to form the closed sets in a topology.

Exercise 8.1.3. If I and J are ideals in R prove that $V(I) \supset V(J)$ if and only if $\sqrt{I} \subset \sqrt{J}$.

It is important to note that not every point in $\mathrm{Spec}(R)$ is closed! Indeed, the closure of $\mathfrak{p} \in \mathrm{Spec}(R)$ is $V(\mathfrak{p})$. Thus the closed points of $\mathrm{Spec}(R)$ will be exactly the maximal ideals on R. The following definitions describe some more basic properties of points of $\mathrm{Spec}(R)$.

Definition 8.1.4. We say that a point $\mathfrak{p}$ in $\mathrm{Spec}(R)$ is a generic point if its closure is all of $\mathrm{Spec}(R)$. Equivalently, $\mathfrak{p}$ is a generic point if it is the unique minimal prime ideal in R.

Warning 8.1.5. Be careful not to confuse the notion of a "generic point" and a "general point". We say that a property is true for a "general point" of a scheme X if it holds for the points in some dense open subset U. Often properties of the generic point can be translated into properties of general points and vice versa, but it is important to keep your language straight.

Definition 8.1.6. Given a point $\mathfrak{p} \in \mathrm{Spec}(R)$ we define the residue field $\kappa(\mathfrak{p})$ to be $R/\mathfrak{p}$.

Loosely speaking, we can think of the residue field $\kappa(\mathfrak{p})$ as the "receiving space" for functions in R evaluated at the point $\mathfrak{p}$. From this perspective, the evaluation map is simply the quotient $R \to R/\mathfrak{p}$. Note that the functions which vanish at the point $\mathfrak{p}$ are exactly the same as the functions contained in $\mathfrak{p}$.

8.1.1 Morphisms

Suppose that $f^\sharp : S \to R$ is a homomorphism of rings. For every prime ideal $\mathfrak{p}$ of R the pullback $(f^\sharp)^{-1}$ is a prime ideal of S. In this way we obtain a function $f : \mathrm{Spec}(R) \to \mathrm{Spec}(S)$.

Exercise 8.1.7. Verify that if $f^\sharp : S \to R$ is a ring homomorphism then the pullback map $f : \mathrm{Spec}(R) \to \mathrm{Spec}(S)$ is continuous in the Zariski topology.

A morphism $f : \mathrm{Spec}(R) \to \mathrm{Spec}(S)$ between two spectra is defined to be a continuous map which is induced by a ring homomorphism $f^\sharp : S \to R$. Note that the fiber over a point $\mathfrak{q} \in \mathrm{Spec}(S)$ can be identified with the closed set $V(f^\sharp(\mathfrak{q}))$ of $\mathrm{Spec}(R)$.

Exercise 8.1.8. Let I be an ideal in a ring R. Show that the map $f : \mathrm{Spec}(R/I) \to \mathrm{Spec}(R)$ induced by the quotient $f^\sharp : R \to R/I$ induces a homeomorphism between $\mathrm{Spec}(R/I)$ and the closed subset $V(I) \subset \mathrm{Spec}(R)$.

8.1.2 Examples

The following examples help us develop our intuition for the points in $\mathrm{Spec}(R)$.

Example 8.1.9. There are two types of points in $\mathrm{Spec}(\mathbb{K}[x])$: the maximal ideals and the prime ideal (0). The maximal ideals are precisely the closed points of $\mathrm{Spec}(\mathbb{K}[x])$ and the zero ideal is the generic point of $\mathrm{Spec}(\mathbb{K}[x])$.

Example 8.1.10. More generally, suppose that R is a finitely generated $\mathbb{K}$-algebra. Recall that $\mathrm{mSpec}(R)$ denotes the set of maximal ideals of R. In other words, $\mathrm{mSpec}(R)$ is the subset of $\mathrm{Spec}(R)$ consisting of all the closed points. By assigning to any prime ideal $\mathfrak{p}$ the intersection $V(\mathfrak{p}) \cap \mathrm{mSpec}(R)$ we obtain a bijection

$$\{\text{points of } \mathrm{Spec}(R)\} \leftrightarrow \left\{ \begin{array}{c} \text{irreducible closed} \\ \text{subsets of } \mathrm{mSpec}(R) \end{array} \right\}.$$

The bijectivity follows from the fact that R is a Jacobson ring: each prime ideal is the intersection of the maximal ideals which contain it.

For some examples of $\mathbb{K}$-algebras we can list all the prime ideals explicitly, but in general the set of points of $\mathrm{Spec}(R)$ is very complicated. Thus we usually understand the non-closed points of $\mathrm{Spec}(R)$ via the corresponding irreducible closed subsets of $\mathrm{mSpec}(R)$.

We next turn to several fundamental examples which are not finitely generated $\mathbb{K}$-algebras.

Example 8.1.11. $\mathrm{Spec}(\mathbb{Z})$ contains one closed point for each prime $p > 0$ and one generic point corresponding to the prime ideal 0. Note the close similarity between $\mathrm{Spec}(\mathbb{Z})$ and $\mathrm{Spec}(\mathbb{K}[x])$ for a field $\mathbb{K}$. The main difference is that the residue fields for closed points of the latter are all finite extensions of $\mathbb{K}$ while the residue fields for closed points of the former all have different characteristics.

More generally, if $\mathcal{O}_K$ is the ring of integers in a number field then $\mathrm{Spec}(\mathcal{O}_K)$ will be a Dedekind domain so that every prime ideal besides 0 will be maximal. The inclusion $\mathbb{Z} \to \mathcal{O}_K$ defines a map $\mathrm{Spec}(\mathcal{O}_K) \to \mathrm{Spec}(\mathbb{Z})$ whose fibers are controlled by the splitting of primes in $\mathcal{O}_K$.

Example 8.1.12. The ring $\mathbb{Z}[x]$ has three types of prime ideals:

(1) Height 0: the zero ideal.

(2) Height 1: principal ideals generated by either a prime $p \in \mathbb{Z}$ or by a polynomial f that is irreducible (and whose coefficients are relatively prime).

(3) Height 2: maximal ideals of the form (p, f) where f is a polynomial that is irreducible mod p.

It may be helpful to see a geometric description of this ring. The inclusion $f^\sharp : \mathbb{Z} \to \mathbb{Z}[x]$ induces a morphism $f : \mathrm{Spec}(\mathbb{Z}[x]) \to \mathrm{Spec}(\mathbb{Z})$. The fiber over a closed point (p) is $\mathrm{Spec}(\mathbb{Z}[x] \otimes_{\mathbb{Z}} \mathbb{Z}/(p)) \cong \mathbb{A}^1_{\mathbb{F}_p}$. The fiber over the generic point 0 is $\mathrm{Spec}(\mathbb{Z}[x] \otimes_{\mathbb{Z}} \mathbb{Q}) \cong \mathbb{A}^1_{\mathbb{Q}}$. In this way it is fair to think of $\mathrm{Spec}(\mathbb{Z}[x])$ as "$\mathbb{A}^1$ over the ring $\mathbb{Z}$".

The previous example motivates the following definition.

Definition 8.1.13. Let R be any ring. We define $\mathbb{A}^n_R$ to be $\mathrm{Spec}(R[x_1, \ldots, x_n])$. Note that $\mathbb{A}^n_R$ comes equipped with a "structure morphism" $\mathbb{A}^n_R \to \mathrm{Spec}(R)$. The fiber of the structure map over a point $\mathfrak{p} \in \mathrm{Spec}(R)$ is isomorphic to $\mathbb{A}^n_{\kappa(\mathfrak{p})}$.

Warning 8.1.14. Suppose $\mathbb{K}$ is a field. We will frequently use the geometric notation $(a_1, \ldots, a_n)$ to denote the closed point in $\mathbb{A}^n_{\mathbb{K}}$ which is more precisely identified as the maximal ideal $(x_1 - a_1, \ldots, x_n - a_n)$. Unfortunately this notation has the potential to cause a mild confusion: the origin $(0, \ldots, 0)$ is a closed point of $\mathbb{A}^n_{\mathbb{K}}$ which is not the same as the generic point (0) of $\mathbb{A}^n_{\mathbb{K}}$.

Example 8.1.15. Let R be a finitely generated $\mathbb{K}$-algebra and let $\mathfrak{p}$ be a prime ideal. The local ring $R_{\mathfrak{p}}$ is usually not finitely generated over $\mathbb{K}$. There will be a bijection between points of $\mathrm{Spec}(R_{\mathfrak{p}})$ and points of $\mathrm{Spec}(R)$ which contain $\mathfrak{p}$ in their closure.

The localization map defines an injection $f : \mathrm{Spec}(R_{\mathfrak{p}}) \to \mathrm{Spec}(R)$ whose image is the intersection of all the open sets in $\mathrm{Spec}(R)$ containing $\mathfrak{p}$. (Note that the image does not actually contain any open subset of $\mathrm{Spec}(R)$.) As we discussed in Section 1.8.3, we should think of $R_{\mathfrak{p}}$ as capturing the behavior of "arbitrarily small" open neighborhoods of $\mathfrak{p}$ without actually identifying a specific example.

Example 8.1.16. Let Z be a DVR. Then $\mathrm{Spec}(Z)$ has two points: a closed point $\mathfrak{m}$ and a generic point 0. A key example is the local ring $\mathcal{O}_{C,p}$ of a curve C at a smooth point p. As discussed above, in this case we can think of the generic point 0 as identifying an "arbitrarily small" open neighborhood the closed point p in C.

More generally, recall that an integral domain R is called a valuation ring if for every element f in the fraction field we have either $f \in R$ or $f^{-1} \in R$. The key property of valuation rings is that their ideals are totally ordered by inclusion. In particular, every valuation ring has a unique maximal ideal $\mathfrak{m}$.

Loosely speaking, valuation rings play the role of "small open neighborhoods" of curves in topology. From this perspective, a morphism $f : \mathrm{Spec}(R) \to X$ represents a "small piece of a curve" in X extending off of the point $f(\mathfrak{m})$.

Example 8.1.17. Let R be a finitely generated $\mathbb{K}$-algebra and let $\mathfrak{m}$ be a maximal ideal. Let $\widehat{R}_{\mathfrak{m}}$ denote the completion of R with respect to $\mathfrak{m}$. Then $\mathrm{Spec}(\widehat{R}_{\mathfrak{m}})$ comes equipped with a canonical map $\mathrm{Spec}(\widehat{R}_{\mathfrak{m}}) \to \mathrm{Spec}(R)$. This is close to – but not quite the same as – the "formal neighborhood" of $\mathfrak{m}$, a ringed space (which is not a scheme!) defined by taking a completion of the structure sheaf. (In the current example, the "formal neighborhood" of $\mathfrak{m}$ is just a single topological point equipped with the sheaf of rings $\widehat{R}_{\mathfrak{m}}$.)

The formal neighborhood of a point is the algebro-geometric equivalent of looking at a limit of small open neighborhoods in the Euclidean topology. For example, the formal neighborhood of the origin in $\mathbb{A}^n$ is equipped with the ring $\mathbb{K}[[x_1, \ldots, x_n]]$. This ring is also the set of germs at the origin of holomorphic functions on $\mathbb{C}^n$. Although there are many more holomorphic functions on $\mathbb{C}^n$ then there are algebraic functions on $\mathbb{A}^n_{\mathbb{C}}$, when we pass to a formal local neighborhood the induced rings of functions in the two settings now coincide!

Sometimes geometric theorems which do not hold for Zariski open sets will have analogues for formal neighborhoods (e.g. the inverse function theorem). However we will not study formal neighborhoods in any depth.

8.1.3 Structure sheaf

We will define the structure sheaf on $\mathrm{Spec}(R)$ in a way that is exactly analogous to our definition for $\mathrm{mSpec}(R)$. The crucial property is that the behavior of functions on open subsets is determined by localization. Intuitively, we would like to define the functions on an open set U simply by inverting all functions f which do not vanish at any point of U. However, this is not quite correct; the right definition is:

Definition 8.1.18. Let R be a ring. For every open subset $U \in \mathrm{Spec}(R)$, define $\widetilde{\mathcal{O}}(U)$ to be the localization of R along all functions f such that $V(f) \cap U = \emptyset$. The restriction maps $\widetilde{\rho}_{U,V}$ are determined by the universal property of localization. This defines a presheaf of abelian groups on $\mathrm{Spec}(R)$.

The structure sheaf $\mathcal{O}_{\mathrm{Spec}(R)}$ is the sheafification of this presheaf $\widetilde{\mathcal{O}}$. Note that for every open set U the set $\mathcal{O}_{\mathrm{Spec}(R)}(U)$ naturally carries the structure of a ring.

A distinguished open affine subset of $\mathrm{Spec}(R)$ is an open set of the form $D_f := \mathrm{Spec}(R) \backslash V(f)$ for some element $f \in R$. It is easy to show that such subsets form a base of the topology of $\mathrm{Spec}(R)$. They are also the open subsets for which it is easy to compute the structure sheaf:

Proposition 8.1.19. *For any $f \in R$ we have $\mathcal{O}_{\mathrm{Spec}(R)}(D_f) \cong R_f$.*

Proof. We first claim that for the presheaf $\widetilde{O}$ we have $\widetilde{\mathcal{O}}(D_f) \cong R_f$. Indeed, if $g \in R$ is a function such that $V(g) \subset V(f)$ then g is already a unit in R_f so the map $R_f \to \widetilde{\mathcal{O}}(D_f)$ induced by the universal property of localization is an isomorphism.

An element of $\mathcal{O}_{\mathrm{Spec}(R)}(D_f)$ is obtained by fixing an open cover $\{V_i\}$ of D_f and choosing elements $g_i \in \widetilde{\mathcal{O}}(V_i)$ whose restrictions to intersections agree. By refining the open cover we may suppose that each V_i is itself a distinguished open affine in $\mathrm{Spec}(R)$. The desired statement then follows from the localization exact sequence of Proposition 1.11.4 applied to R_f. $\square$

We also have:

Proposition 8.1.20. *For any point $\mathfrak{p} \in \mathrm{Spec}(R)$ the stalk of the structure sheaf at $\mathfrak{p}$ is $\mathcal{O}_{\mathrm{Spec}(R),\mathfrak{p}} = R_{\mathfrak{p}}$.*

Proof. Since distinguished open affines are cofinal in the direct limit defining the stalk, it suffices to check that

$$R_{\mathfrak{p}} = \varinjlim_{f \notin \mathfrak{p}} R_f$$

which can be proved by comparing the universal properties of both sides. $\square$

Example 8.1.21. Suppose that R is a finitely generated $\mathbb{K}$-algebra. Then the structure sheaf on $\mathrm{mSpec}(R)$ is simply the restriction of the structure sheaf on $\mathrm{Spec}(R)$. In other words, one can recover the topology and structure sheaf on $\mathrm{Spec}(R)$ entirely from $\mathrm{mSpec}(R)$.

From now on we will replace work exclusively with Spec instead of mSpec. For example, we define $\mathbb{A}^n_{\mathbb{K}} = \mathrm{Spec}(\mathbb{K}[x_1, \ldots, x_n])$. We can be confident that these new constructions will behave in exactly the same way as the old ones; however, it will often be convenient theoretically to have access to the new non-closed points.

Finally, we discuss how morphisms of spectra interact with the structure sheaf. The key point is:

> A morphism $f : \mathrm{Spec}(R) \to \mathrm{Spec}(S)$ induces a map of sheaves $f^\sharp : \mathcal{O}_{\mathrm{Spec}(S)} \to f_*\mathcal{O}_{\mathrm{Spec}(R)}$.

Recall that a morphism $f : \mathrm{Spec}(R) \to \mathrm{Spec}(S)$ is simply a ring homomorphism $f^\sharp : S \to R$. Suppose we fix an open subset $U \subset \mathrm{Spec}(S)$. Note that a function $g \in S$ will satisfy $V(g) \cap U = \emptyset$ if and only if $V(f^\sharp(g)) \cap f^{-1}U = \emptyset$. Thus the universal property of localization induces a ring map $\widetilde{\mathcal{O}}_{\mathrm{Spec}(S)}(U) \to \widetilde{\mathcal{O}}_{\mathrm{Spec}(R)}(f^{-1}U)$. Varying U, we obtain a map of presheaves $\widetilde{\mathcal{O}}_{\mathrm{Spec}(S)} \to f_*\widetilde{\mathcal{O}}_{\mathrm{Spec}(R)}$. By Exercise 7.2.5 the universal property of sheafification yields a map $f^\sharp : \mathcal{O}_{\mathrm{Spec}(S)} \to f_*\mathcal{O}_{\mathrm{Spec}(R)}$.

Even though the sheaf map is completely determined by the original ring map via localization, we still consider the sheaf map to be "part of the data" of the morphism $f : \mathrm{Spec}(R) \to \mathrm{Spec}(S)$.

8.1.4 Exercises

Exercise 8.1.22. Consider the nodal cubic $X = V(y^2 - x^3 - x^2)$ in $\mathbb{A}^2_{\mathbb{K}}$. Show that every open subset of X is irreducible.

Consider now the completion of $\mathbb{K}[x, y]$ along the ideal (x, y). The equation $y^2 - x^3 - x^2$ defines a closed subset of the corresponding affine scheme. Show that that this closed set has two irreducible components corresponding to the two "Euclidean local" branches of the curve at the node. (Hint: what is the Taylor series for $\sqrt{x^2 + x^3}$?)

Exercise 8.1.23. Suppose that $f : \mathrm{Spec}(R) \to \mathrm{Spec}(S)$ is the morphism associated to the ring map $f^\sharp : S \to R$. Given a point $\mathfrak{p} \in S$, we define the scheme-theoretic fiber of f over $\mathfrak{p}$ to be $V((f(\mathfrak{p})))$. Prove that the set-theoretic fiber of f over $\mathfrak{p}$ is the same as the set underlying the scheme-theoretic fiber.

Exercise 8.1.24. Consider the ring $R = \mathbb{K}[w, x, y, z]/(wy, wz, xy, xz)$. Geometrically $\mathrm{Spec}(R)$ represents the union of two copies of $\mathbb{A}^2_{\mathbb{K}}$ meeting at the origin. Let $U \subset \mathrm{Spec}(R)$ be the complement of the origin.

Prove that $\mathcal{O}_{\mathrm{Spec}(R)}(U)$ is not the same as the localization of R along all the functions which vanish along the origin. (Hint: consider the function which is identically 1 on one component of U and identically 0 on the other.) This explains why we must include a sheafification when defining the structure sheaf.

Exercise 8.1.25. Consider the morphism $f : \mathrm{Spec}(\mathbb{Z}[i]) \to \mathrm{Spec}(\mathbb{Z})$ induced by the inclusion $\mathbb{Z} \hookrightarrow \mathbb{Z}[i]$. What is the fiber of f over a point of $\mathrm{Spec}(\mathbb{Z})$?

Exercise 8.1.26. Consider the equation $y^2 = x^3 + 1$. Discuss and contrast the vanishing locus of this equation:

(1) In $\mathbb{C}[x, y]$.

(2) In $\mathbb{Q}[x, y]$.

(3) In $\mathbb{F}_p[x, y]$ where $p = 2$, $p = 3$, or $p \geq 5$.

(4) In $\mathbb{Z}[x, y]$.

Hypothesize why it makes sense to say that the vanishing locus $y^2 = x^3 + 1$ is "an elliptic curve defined over $\mathbb{Z}[\frac{1}{6}]$."

8.2 Schemes

We will define a scheme to be a topological space X equipped with a sheaf of rings $\mathcal{O}_X$ such that X is "locally isomorphic" to the spectrum of a ring. Before we can make such a definition, we need to specify what "isomorphic" means for such spaces.

Definition 8.2.1. A ringed space is a topological space X equipped with a sheaf of rings $\mathcal{O}_X$. We say that X is a locally ringed space if for every $x \in X$ the stalk $\mathcal{O}_{X,x}$ is a local ring.

A morphism of locally ringed spaces consists of a continuous map $f : X \to Y$ and a morphism of sheaves $f^\sharp : \mathcal{O}_Y \to f_*\mathcal{O}_X$ such that for every $x \in X$ the induced morphism on stalks $f^\sharp_x : \mathcal{O}_{Y,f(x)} \to \mathcal{O}_{X,x}$ is a local homomorphism of local rings. (That is, the $f^\sharp$-image of the maximal ideal in $\mathcal{O}_{Y,f(x)}$ should be the maximal ideal in $\mathcal{O}_{X,x}$.)

Here the map of sheaves $f^\sharp$ represents the "pullback map" which associates to any function on Y the induced map on X obtained by composition with f.

Remark 8.2.2. It may feel more natural for the "pullback map" to be written as a map $f^{-1}\mathcal{O}_Y \to \mathcal{O}_X$ instead of a map $\mathcal{O}_Y \to f_*\mathcal{O}_X$. (Of course by adjunction the data represented by these two possible representations are completely equivalent.) One reason we prefer the latter is because for any open set $V \subset Y$ the ring $f_*\mathcal{O}_X(V)$ is naturally a $\mathcal{O}_Y(V)$-module, whereas the opposite statement is not true. Thus the map $\mathcal{O}_Y \to f_*\mathcal{O}_X$ fits more naturally into the theory of sheaves of modules which we will soon develop.

The "local homomorphism" requirement in Definition 8.2.1 is at first glance a bit obscure. In the setting of algebraic geometry, this requirement ensures that the topological map f and the sheaf map $f^\sharp$ are "compatible" in the sense that the pullback of the functions vanishing at $f(x)$ will be functions which vanish at x. This guarantees that the continuous function f is "induced locally" by the ring maps $f^\sharp$. (We discussed this issue at length in the setting of quasiaffine $\mathbb{K}$-schemes in Section 1.12.)

Definition 8.2.3. An affine scheme is a locally ringed space which is isomorphic (as a locally ringed space) to the spectrum of a ring. A scheme is a locally ringed space such that every point $x \in X$ admits an open neighborhood which is an affine scheme.

A morphism of schemes is a morphism of locally ringed spaces. We let **Sch** denote the category of schemes.

Right away we should settle a potential conflict of notation with the previous section.

Theorem 8.2.4. *Let $X = \mathrm{Spec}(R)$ and $Y = \mathrm{Spec}(S)$ be affine schemes. There is a bijection between the set of scheme morphisms $f : X \to Y$ and the set of ring homomorphisms $f^\sharp : S \to R$.*

This bijection assigns to any morphism f the map $f^\sharp : \mathcal{O}_Y(Y) \to \mathcal{O}_X(X)$ and to any ring map $f^\sharp$ the morphism $f : X \to Y$ described in the previous section. The proof is essentially the same as the proof of Proposition 1.12.8.

8.2.1 Gluing schemes

Our first task is to find a way to construct schemes that are not simply the spectrum of a ring. The following construction gives a very general and very useful method for creating new schemes.

Theorem 8.2.5. *Let $\{X_i\}_{i \in I}$ be a collection of schemes. Suppose that:*

- *for every pair of indices i, j we have an open subscheme $X_{ij} \subset X_i$ where $X_{ii} = X$.*
- *for every pair of indices i, j we have an isomorphism $f_{ij} : X_{ij} \to X_{ji}$ where $f_{ii} = id_X$.*

satisfying the cocycle condition

- *$f_{ij}(X_{ik} \cap X_{ij}) \subset X_{jk}$, and*
- *$f_{ik}|_{X_{ij} \cap X_{ik}} = f_{jk}|_{X_{ji} \cap X_{jk}} \circ f_{ij}|_{X_{ij} \cap X_{ik}}$.*

Then there is a scheme X (unique up to unique isomorphism) equipped with open sets U_i such that each U_i is isomorphic to X_i and $U_i \cap U_j$ is isomorphic to X_{ij}.

Note that the "cocycle condition" shows up again in this new setting.

Proof. We let X denote the quotient of $\sqcup X_i$ by the equivalence relation $x \sim f_{ij}(x)$ for every i, j. (The cocycle condition guarantees that this is an equivalence relation.) Note that for every index i there is a natural inclusion $X_i \hookrightarrow X$ which is a homeomorphism onto an open subset. We can then apply Corollary 7.6.3 to glue the structure sheaves $\mathcal{O}_{X_i}$ to obtain a sheaf of rings $\mathcal{O}$. Since the locally ringed space $(X, \mathcal{O})$ is locally isomorphic around any point to one of the X_i, it is a scheme. □

We can also naturally define morphisms of schemes via gluing; see Exercise 8.2.20. Here are some examples of the gluing construction in action.

Example 8.2.6. Suppose that A_1, A_2 are both isomorphic to $\mathbb{A}^1_{\mathbb{K}}$ and that U_1, U_2 are the complements of the origin in the two affine lines. The identity map on this open set induces an isomorphism $f_{12} : U_1 \to U_2$. By Theorem 8.2.5 we can glue the two affine lines along f_{12} to obtain a scheme X.

This scheme will look like a copy of the affine line with the origin "doubled". In other words, there are two origins $0_1, 0_2$ coming from the two affine lines. Any open neighborhood of 0_1 will intersect every open neighborhood of 0_2; nevertheless there are open sets containing 0_1 but not 0_2 and vice versa. Exercise 8.2.16 asks you to show that X is not an affine scheme.

Example 8.2.7. Suppose that A_1, A_2 are both isomorphic to $\mathbb{A}^1_{\mathbb{K}}$ and that U_1, U_2 are the complements of the origin in the two affine lines. Note that the complement of the origin has a non-trivial isomorphism defined by the ring map $x \mapsto x^{-1}$. Let $f_{12} : U_1 \to U_2$ denote this inversion map. According to Theorem 8.2.5 we can glue the two affine lines along f_{12} to obtain a scheme. The scheme obtained in this way is isomorphic to $\mathbb{P}^1_{\mathbb{K}}$; the inversion used in the gluing map f_{12} matches with the inversion identifying the coordinate ring $\mathbb{K}[\frac{y}{x}]$ on $D_{+,x}$ with the coordinate ring $\mathbb{K}[\frac{x}{y}]$ on $D_{+,y}$.

8.2.2 Proj construction

The Proj construction associates a scheme to a $\mathbb{Z}_{\geq 0}$-graded algebra R. The topology and functions of this scheme are determined by the structure of the graded ideals in R. Alternatively, one can view the Proj construction as a special case of Theorem 8.2.5 which glues affine schemes via the graded localizations of R.

Construction 8.2.8. Let R be a $\mathbb{Z}_{\geq 0}$-graded ring. We define $\mathrm{Proj}(R)$ to be the set of prime homogeneous ideals which do not contain $R_{>0}$. Given a homogeneous ideal I, we set $V_+(I)$ to be the set of prime homogeneous ideals which contain I. These form the closed sets in a topology on $\mathrm{Proj}(R)$.

A distinguished open affine in $\mathrm{Proj}(R)$ is an open set of the form $D_{+,f} = \mathrm{Proj}(R)\backslash V_+(f)$. One can show that $D_{+,f}$ is homeomorphic to $\mathrm{Spec}((R_f)_0)$ using the familiar construction

$$\begin{tikzcd} \mathfrak{p} \subset R \arrow[r, leftrightarrow, "\text{localization}"] & \mathfrak{p}_f \subset R_f \arrow[d, leftrightarrow, "\text{intersection}"] \\ & \mathfrak{p}_f \cap (R_f)_0 \subset (R_f)_0. \end{tikzcd}$$

The proof that this is a homeomorphism is identical to the special case discussed in Proposition 2.3.8

To define the structure sheaf $\mathcal{O}_{\mathrm{Proj}(R)}$, we glue together the structure sheaves of the various distinguished open affines $\mathrm{Spec}((R_f)_0)$ using Corollary 7.6.3. (Alternatively, we could directly glue the schemes $\mathrm{Spec}((R_f)_0)$ together along their common intersections using Theorem 8.2.5.)

The scheme $\mathrm{Proj}(R)$ always comes equipped with a map $\mathrm{Proj}(R) \to \mathrm{Spec}(R_0)$. This map can be defined by gluing the maps of distinguished open affines in $\mathrm{Proj}(R)$ corresponding to the maps of rings $R_0 \to (R_f)_0$ using Exercise 8.2.20.

Definition 8.2.9. For any ring R_0, we define $\mathbb{P}^n_{R_0}$ to be $\mathrm{Proj}(R_0[x_0, \ldots, x_n])$. This scheme is equipped with a structure map $\mathbb{P}^n_{R_0} \to \mathrm{Spec}(R_0)$; the fiber over a point $\mathfrak{p} \in \mathrm{Spec}(R_0)$ is $\mathbb{P}^n_{\kappa(\mathfrak{p})}$.

When R is a finitely generated graded $\mathbb{K}$-algebra this construction behaves exactly the same as the familiar mProj construction.

Example 8.2.10. Let R be a finitely generated $\mathbb{Z}_{\geq 0}$-graded $\mathbb{K}$-algebra. Then the points of $\mathrm{mProj}(R)$ are in bijection with the closed points of $\mathrm{Proj}(R)$. More generally, there is a bijection

$$\{\text{points of } \mathrm{Proj}(R)\} \leftrightarrow \left\{ \begin{array}{c} \text{irreducible closed} \\ \text{subsets of } \mathrm{mProj}(R) \end{array} \right\}.$$

As we discussed in Example 8.1.10, all the information about the topology of $\mathrm{Proj}(R)$ is captured by $\mathrm{mProj}(R)$. Furthermore, as with affine schemes the restriction of the structure sheaf on $\mathrm{Proj}(R)$ yields the structure sheaf on $\mathrm{mProj}(R)$. From now on we will replace our mProj construction by the Proj construction.

For rings that are not $\mathbb{K}$-algebras one must be a little careful – our intuition concerning vector spaces can lead us astray.

Example 8.2.11. Consider $\mathbb{P}^1_{\mathbb{Z}} = \mathrm{Proj}(\mathbb{Z}[x, y])$. By analogy with projective space over a field, we might expect the closed points to look like ordered pairs of integers $(a : b)$ up to equivalence by rescaling. It turns out that such pairs are in bijection with the morphisms $\mathrm{Spec}(\mathbb{Z}) \to \mathbb{P}^1_{\mathbb{Z}}$; however, the closed points of $\mathbb{P}^1_{\mathbb{Z}}$ will look quite different.

The best way to think about $\mathbb{P}^1_{\mathbb{Z}}$ is using the structure morphism $\mathbb{P}^1_{\mathbb{Z}} \to \mathrm{Spec}(\mathbb{Z})$. The fiber over a closed point (p) will be isomorphic to $\mathbb{P}^1_{\mathbb{F}_p}$ and the fiber over the generic point (0) will be $\mathbb{P}^1_{\mathbb{Q}}$. Note that any closed point of $\mathbb{P}^1_{\mathbb{Z}}$ will be contained in a fiber of the first type and thus such points can be identified by choosing a prime p and a closed point in $\mathbb{P}^1_{\mathbb{F}_p}$.

8.2.3 Relative schemes

Definition 8.2.12. Fix a scheme S. Then an S-scheme (or equivalently a scheme over S) is a scheme X equipped with a "structure morphism" $p_X : X \to S$. A morphism of S-schemes is a commuting diagram

$$\begin{tikzcd} X \arrow[rr, "f"] \arrow[dr, "p_X"'] & & Y \arrow[dl, "p_Y"] \\ & S & \end{tikzcd}$$

The category of S-schemes will be denoted by $\mathbf{Sch}/S$.

Many geometric constructions are best understood in a relative category rather than the category of all schemes. We have seen a couple examples of this principle in action already:

Example 8.2.13. Every quasiprojective $\mathbb{K}$-scheme comes equipped with a morphism $X \to \mathrm{Spec}(\mathbb{K})$. Furthermore the $\mathbb{K}$-morphisms of quasiprojective $\mathbb{K}$-schemes $f : X \to Y$ are exactly the same as the morphisms in the category of $\mathbb{K}$-schemes. (Indeed, requiring that

f locally be defined by a $\mathbb{K}$-algebra homomorphism is the same as insisting that f form a commuting diagram with the structural morphisms to $\operatorname{Spec}(\mathbb{K})$.)

Example 8.2.14. Let R be a $\mathbb{Z}_{\geq 0}$-graded ring. As discussed in Section 8.2.2 we have a structural morphism $\operatorname{Proj}(R) \to \operatorname{Spec}(R_0)$. The most natural setting for the Proj construction is the category $\mathbf{Sch}/\operatorname{Spec}(R_0)$.

8.2.4 Scheme valued points

Definition 8.2.15. Let X be a scheme. For any other scheme Z a Z-valued point of X is defined to be a morphism $Z \to X$. The set of all Z-points of X is denoted by $X(Z)$.

When X is an S-scheme, we usually insist that $Z \to X$ be a morphism in the category $\mathbf{Sch}/S$.

Scheme-valued points can be used to build up a systematic "categorical" perspective of X. In many ways the notion of scheme-theoretic points is more natural than the notion of topological points. For example, the set of $\operatorname{Spec}(\mathbb{K})$-valued points of X will usually exhibit better behavior than the set of points in X which have fixed residue field $\mathbb{K}$ (as we saw in Exercise 1.5.11).

8.2.5 Exercises

Exercise 8.2.16. Let X be the affine line with a doubled origin constructed in Example 8.2.6. Compute $\mathcal{O}_X(X)$ and conclude that X is not an affine scheme.

Exercise 8.2.17. Let X be a scheme. Recall that a point $\xi \in X$ is a generic point if its closure is all of X.

(1) Prove that if X has a generic point then it is topologically irreducible.

(2) Prove that if X has a generic point ξ then every open affine U in S contains ξ.

(3) Prove that if X has a generic point ξ then ξ is the only generic point.

Exercise 8.2.18. Let X be a scheme, $U \subset X$ an open set, and $f \in \mathcal{O}_X(U)$. We define the vanishing locus of f, denoted by $V(f)$, to be the set of points $x \in U$ such that $\rho_{U,x}(f) \in \mathfrak{m}_x$.

(1) Prove that $V(f)$ is a closed subset of U.

(2) Let U' denote the complement of $V(f)$ in U. Prove that $\rho_{U,U'}(f)$ has an inverse in $\mathcal{O}_X(U')$.

(We have seen both these properties before in the special case when U is affine.)

Exercise 8.2.19. Prove carefully the claim implicit in Example 8.2.11: if R is a PID then the R-points of the R-scheme $\mathbb{P}^1_R$ (or in other words, sections of the structure map $\mathbb{P}^1_R \to \operatorname{Spec}(R)$) are in bijection with

$$R^2 \backslash (0,0) \ / \ (a,b) \sim (ra, rb) \, \forall r \neq 0$$

You may prefer the equivalent description that the R-points of $\mathbb{P}^1_R$ are the pairs of coprime elements (r_1, r_2) up to rescaling by $R^\times$.

Exercise 8.2.20. Let X be a scheme. Suppose that $\{U_i\}$ is an open cover of X by schemes and for each i we have a morphism $f_i : U_i \to Y$. Suppose furthermore that $f_i|_{U_i \cap U_j} = f_j|_{U_i \cap U_j}$ for all i, j. Prove that there is a morphism $f : X \to Y$ such that $f|_{U_i} = f_i$ for all i.

8.3 First properties of schemes

8.3.1 Local properties

We first develop a systematic way of passing a ring-theoretic property P to the category of schemes. As in Section 2.5, the necessary ingredient is that P should be compatible with localization in some way. We discuss two types of "local properties" for ring.

We say that a ring property P is stalk-local when P holds for a ring R if and only if it holds for every localization $R_{\mathfrak{p}}$ at a prime ideal $\mathfrak{p}$. The following theorem (whose proof is immediate) shows that stalk-local properties can naturally be extended to arbitrary schemes.

Theorem 8.3.1. *Let P be a stalk-local property of rings. Then for any scheme X the following are equivalent:*

(1) X admits an open cover by affine open subsets U satisfying P.

(2) For every $x \in X$ the stalk $\mathcal{O}_{X,x}$ satisfies P.

(3) Every open affine U in X satisfies P.

In these cases we say that X satisfies P.

Some examples of stalk-local properties are:

- Reducedness, i.e. the property $\mathrm{Nil}(R) = 0$.
- Normality, i.e. the property that the localization of R along every prime ideal is an integrally closed domain.

  (The property of being an integrally closed domain is almost stalk-local – if we assume R is a domain, then we can check the integral closure property on stalks.)
- Regularity, i.e. the property that the localization of R along every prime ideal is regular.

We say that a scheme is reduced, normal, or regular if all of its local rings satisfy the corresponding property.

Just as with quasiprojective $\mathbb{K}$-schemes, the schemes which are irreducible and reduced play a special role in the theory. (We will reserve the terminology "variety" for schemes defined over a field.)

Definition 8.3.2. A scheme X is integral if it is reduced and topologically irreducible.

Exercise 8.3.3. Show that X is integral if and only if for every open affine $U \subset X$ the section ring $\mathcal{O}_X(U)$ is a domain.

We say that a ring property P is local if

- whenever R satisfies P then the localization R_f at an element $f \in R$ also satisfies P, and
- if $\{f_i\}$ generate the unit ideal in R and R_{f_i} satisfies P for every i then R satisfies P.

Note that every stalk-local property is also local, but the converse is not true. The property "R is Noetherian" is an important example of a local property that is not stalk-local. The key tool for working with local properties is:

Lemma 8.3.4 (Nike's lemma). *Let X be a scheme. Suppose that U and V are open affines in X. Then $U \cap V$ admits a cover by open sets which are simultaneously distinguished open affines in both U and V.*

Proof. The argument for Lemma 2.5.2 works equally well in this more general setting. □

Using Nike's Lemma, one easily deduces:

Theorem 8.3.5. *Let P be a local property of rings. Then for any scheme X the following are equivalent:*

(1) X admits an open cover by affine open subsets U satisfying P.

(2) Every open affine U in X satisfies P.

In the setting of Theorem 8.3.5 we will usually say that X is "locally P".

8.3.2 Morphisms to affine schemes

The following result is absolutely essential for understanding the category of schemes.

Theorem 8.3.6. *Let $X = \mathrm{Spec}(R)$ be an affine scheme. For any scheme Y, there is a bijection between morphisms $Y \to X$ and ring homomorphisms $R \to \mathcal{O}_Y(Y)$.*

Proof. Suppose we choose an open cover of Y consisting of open affines $\{V_i\}$. We have a diagram

$$\begin{tikzcd} \mathrm{Hom}(Y, X) \arrow[r] \arrow[d, "\rho_1"] & \mathrm{Hom}(R, \mathcal{O}_Y(Y)) \arrow[d, "\rho_2"] \\ \prod_i \mathrm{Hom}(V_i, X) \arrow[r] & \prod_i \mathrm{Hom}(R, \mathcal{O}_Y(V_i)). \end{tikzcd}$$

The two vertical maps are injective and by Theorem 8.2.4 the map along the bottom is a bijection. It only remains to check that the images of the two vertical maps are identified by this bijection. In other words, we must check that the gluing data for obtaining a morphism $Y \to X$ from a set of morphisms $V_i \to X$ is the same as the gluing data for constructing a morphism $f^\sharp(Y) : R \to \mathcal{O}_Y(Y)$ from morphisms $f^\sharp(V_i) : R \to \mathcal{O}_Y(V_i)$ using the gluing property for sheaves. This argument is explained in the proof of Theorem 2.4.8. □

8.3.3 Open and closed embeddings

It turns out that any open or a closed subset of a scheme can be given the structure of a scheme. In this section we begin discussing these important constructions.

Definition 8.3.7. A morphism $f : X \to Y$ is an open embedding if f takes X homeomorphically onto an open subset U of Y and the map $f^\sharp : \mathcal{O}_Y \to f_*\mathcal{O}_X$ induces an isomorphism between $\mathcal{O}_Y|_U$ and $\mathcal{O}_X$.

An open subscheme is an open embedding such that f is a set-theoretic inclusion.

It is clear that any open subset of a scheme can naturally be equipped with the structure of an open subscheme in a unique way.

We next turn to closed embeddings. The first step is the following proposition:

Proposition 8.3.8. *Let $f : X \to Y$ be a morphism of schemes. The following are equivalent:*

(1) f takes X homeomorphically onto a closed subset Z of Y and the map $f^\sharp : \mathcal{O}_Y \to f_*\mathcal{O}_X$ is surjective.

(2) For every open affine $V \subset Y$, the preimage $f^{-1}(V)$ is affine and the induced map $f^\sharp(V) : \mathcal{O}_Y(V) \to \mathcal{O}_X(f^{-1}V)$ is surjective.

(3) Y admits an open cover by open affines V satisfying that the induced map $f^\sharp(V) : \mathcal{O}_Y(V) \to \mathcal{O}_X(f^{-1}V)$ is surjective.

Although we could give the proof now, it will be much easier once we have developed more machinery. Thus we take this theorem for granted for now and postpone the proof until Lemma 9.3.11; the dedicated reader can check that we do not introduce any circular arguments in this way. (Note that the implications (2) $\implies$ (3) $\implies$ (1) are immediate, so the only missing step is (1) $\implies$ (2).)

Definition 8.3.9. A morphism $f : X \to Y$ is a closed embedding if it satisfies the three equivalent definitions of Proposition 8.3.8.

A closed subscheme is a closed embedding such that f is a set-theoretic inclusion.

Exercise 8.3.10. Prove that a closed subscheme of an affine scheme is affine.

Closed embeddings are significantly more subtle than open embeddings. For example, it is not immediately clear that a closed subset of a scheme will admit the structure of a closed subscheme. The following construction identifies the most useful tool for working with closed embeddings.

Definition 8.3.11. Let X be a scheme. A quasicoherent ideal sheaf $\mathcal{I}$ is a subsheaf of $\mathcal{O}_X$ such that:

(1) for every open affine $U \subset X$ we have that $\mathcal{I}(U)$ is an ideal in $\mathcal{O}_X(U)$,

(2) for every open affine U and every $f \in \mathcal{O}_X(U)$ we have the localization formula $\mathcal{I}(D_f) = \mathcal{I}(U)_f$.

It turns out the data of a closed subscheme $X \subset Y$ is equivalent to the data of a quasicoherent ideal sheaf on Y. The following exercise shows how to construct a closed subscheme from a quasicoherent ideal sheaf; we will postpone the construction of a quasicoherent ideal sheaf from a closed subscheme to Theorem 9.3.12.

Exercise 8.3.12. Suppose that $\mathcal{I}$ is a quasicoherent ideal sheaf on a scheme Y. Prove that there is a closed subscheme $f : X \to Y$ such that $\mathcal{I}$ is the kernel of $f^\sharp$. (Hint: for every open affine U in Y construct a closed subscheme $V(\mathcal{I}(U))$ of U. Show that these closed subschemes can be glued to yield a scheme X.)

8.3.4 Exercises

Exercise 8.3.13. Prove that open and closed embeddings are examples of monomorphisms in the category of schemes.

Exercise 8.3.14. Let X be a scheme. Given any affine open subset $U \subset X$, we can consider the ideal of nilpotents $N(U) \subset \mathcal{O}_X(U)$. Show that these ideals together form a quasicoherent ideal sheaf $\mathcal{N} \subset \mathcal{O}_X$.

The corresponding closed subscheme is known as the reduced subscheme $X_{red} \subset X$. Prove that the inclusion is a homeomorphism.

(Warning: it is not true that on every open subset U the ideal $\mathcal{N}(U)$ is the set of nilpotents in $\mathcal{O}_X(U)$. Consider for example $X = \coprod_{n\in\mathbb{N}} \mathrm{Spec}(\mathbb{K}[x]/(x^n))$.)

Exercise 8.3.15. Let $f : X \to Y$ be a morphism of schemes. Suppose that $\mathcal{I}$ is a quasicoherent ideal sheaf on Y. Give an example showing that $f^{-1}\mathcal{I}$ need not admit an injection into $\mathcal{O}_X$ (and thus cannot necessarily be considered as a quasicoherent ideal sheaf on X).

The inverse image ideal sheaf $f^{-1}\mathcal{I} \cdot \mathcal{O}_X$ is the image of $f^{-1}\mathcal{I}$ under the natural map $f^{-1}\mathcal{O}_Y \to \mathcal{O}_X$. (Conceptually, the inverse image ideal sheaf is obtained by "pulling back the functions in $\mathcal{I}$.) Prove that the inverse image ideal sheaf is a quasicoherent ideal sheaf.

Exercise 8.3.16. Let X be an integral scheme.

(1) Show that X is normal if and only if for every open affine U the ring $\mathcal{O}_X(U)$ is integrally closed in its fraction field.

(2) Construct a normalization map $\nu : X^\nu \to X$ from a normal scheme X^ν in the following way: for every open affine U in X let $\widetilde{U}$ be the scheme defined by the integral closure of $\mathcal{O}_X(U)$ in its fraction field. Show that as we vary U the affine schemes $\widetilde{U}$ can be glued together to give ν.

(3) Prove that the morphism $\nu : \widetilde{X} \to X$ has the following universal property: for every normal integral scheme Y and any dominant morphism $f : Y \to X$ there is a unique factoring of f through ν.

8.4 Category of schemes

In this section we work through some basic properties of the category of (relative) schemes. By far the most important property is the existence of relative products.

8.4.1 Initial and final objects

Let X be a scheme. Theorem 8.3.6 show that there is a bijection between morphisms $f : X \to \mathbb{Z}$ and ring maps $\mathbb{Z} \to \mathcal{O}_X(X)$. Since $\mathbb{Z}$ is an initial object in the category of rings, we conclude that $\mathrm{Spec}(\mathbb{Z})$ is a final object in the category **Sch**.

As in other geometric categories, the empty set $\mathrm{Spec}(0)$ is an initial object in **Sch**.

8.4.2 Products

It turns out that (relative) products of schemes exist in full generality. We achieve this construction via the following steps:

(1) We first construct products for morphisms of affine schemes. Given morphisms $f : \mathrm{Spec}(R) \to \mathrm{Spec}(T)$, $g : \mathrm{Spec}(S) \to \mathrm{Spec}(T)$ of affine schemes, Theorem 8.3.6 shows that $\mathrm{Spec}(R \otimes_T S)$ satisfies the right universal property to be identified with the relative product $\mathrm{Spec}(R) \times_{\mathrm{Spec}(T)} \mathrm{Spec}(S)$ in **Sch**.

(2) Let X, Y, Z be schemes equipped with morphisms $f : X \to Z$, $g : Y \to Z$. Let $\{U_i\}$ be an open cover of X. Suppose that we can construct each product $U_i \times_Z Y$ in the category of schemes. Then we can also construct $X \times_Z Y$ as follows. Using the universal property of products as in Lemma 2.9.1 we see that the preimage of $U_i \cap U_j$ in $U_i \times_Z Y$ is isomorphic to $(U_i \cap U_j) \times_Z Y$. Thus the preimage of $U_i \cap U_j$ in $U_i \times_Z Y$ is canonically isomorphic to its preimage in $U_j \times_Z Y$. Applying Theorem 8.2.5 we can glue the various $U_i \times_Z Y$ along these open subsets with these canonical identifications. The resulting scheme P is equipped with two natural maps $\pi_1 : P \to X$ and $\pi_2 : P \to Y$. Arguing as in Lemma 2.9.3 we see that P satisfies the universal property of $X \times_Z Y$.

(3) Combining the previous two steps, we can construct all relative products $X \times_{\mathrm{Spec}(T)} Y$ over an affine scheme. Arguing as in Lemma 2.9.4 we can leverage this case to construct arbitrary relative products $X \times_Z Y$.

Given any scheme Y we define $\mathbb{A}^n_Y := \mathbb{A}^n_{\mathbb{Z}} \times_{\mathrm{Spec}(\mathbb{Z})} Y$ and $\mathbb{P}^n_Y := \mathbb{P}^n_{\mathbb{Z}} \times_{\mathrm{Spec}(\mathbb{Z})} Y$. The following exercise verifies that there is no conflict with our earlier notation.

Exercise 8.4.1. Let A be any ring. Prove that $\mathbb{P}^n_A \cong \mathbb{P}^n_{\mathbb{Z}} \times_{\mathrm{Spec}(\mathbb{Z})} \mathrm{Spec}(A)$.

The commuting square obtained from the relative product is often referred to as "base change" or as a "Cartesian diagram". Many fundamental constructions for schemes are best formalized using the product construction.

Construction 8.4.2. Let $f : X \to Y$ be a morphism of schemes. For any morphism $i : Z \to X$, the fiber of f over Z is the base change $X \times_Y Z$. We are usually interested in the case when i is a closed or open embedding. In particular, given a point $y \in Y$ the fiber over y is often denoted by X_y.

Note an important consequence: fibers are "preserved" by base change. That is, suppose we have a morphism $f : X \to Y$ and a morphism $g : Z \to Y$. Then the fiber of $X \times_Y Z \to Z$ over a point z is isomorphic to the fiber of f over $g(z)$.

Construction 8.4.3. Let $i : Y \to X$ and $j : Z \to X$ be open or closed embeddings. The intersection of Y and Z is $Y \times_X Z$.

Construction 8.4.4. Let $f : X \to Z$ be a morphism of schemes. The diagonal of f is the morphism $\Delta_{X/Z} : X \to X \times_Z X$ defined by the identity morphisms to both factors. We also sometimes use $\Delta_{X/Z}$ to denote the set-theoretic image of the diagonal morphism.

Construction 8.4.5. Let $f : X \to Y$ be a morphism of S-schemes. The graph of f is the morphism $\Gamma_f : X \to X \times_S Y$ defined by the identity morphism to the first factor and f to the second. We also sometimes use Γ_f to denote the set-theoretic image of the graph morphism.

It is tempting to think that the diagonal and graph of a morphism will be closed embeddings, but this need not be true in general. We will return to this issue in the following sections.

8.4.3 Coproducts

Coproducts exist in the category of schemes; just like in other geometric categories, the coproduct of two schemes is simply their disjoint union. There is another type of relative coproduct that frequently occurs in geometric constructions: gluing along isomorphic closed subschemes. (Note that Theorem 8.2.5 guarantees that we can glue along isomorphic open subschemes in very general situations.)

Construction 8.4.6. Let X_1, X_2 be two schemes. Suppose that Z is a scheme admitting closed embeddings $f_1 : Z_1 \to X_1$ and $f_2 : Z_2 \to X_2$. We construct the gluing $X_1 \sqcup_Z X_2$.

We start by discussing the affine case. If X_1, X_2 are both affine, then Z is also affine (being isomorphic to a closed subscheme of an affine scheme). Suppose $X_i = \mathrm{Spec}(R_i)$ for $i = 1, 2$, that $Z = \mathrm{Spec}(S)$, and that f_i is defined by $f_i^\sharp : R_i \to S$. Then we define $Y := X_1 \sqcup_\phi X_2$ as the Spec of the subring

$$T := \{(r_1, r_2) \in R_1 \times R_2 \mid f_1^\sharp(r_1) = f_2^\sharp(r_2)\}.$$

Exercise 8.4.7. (1) Show that $\mathrm{Spec}(T)$ has a closed subset isomorphic to $\mathrm{Spec}(S)$ whose complement is isomorphic to $(X_1 \backslash Z) \cup (X_2 \backslash Z)$.

(2) Show that for any $t \in T$ the distinguished open affine D_t is isomorphic to the result of applying the gluing construction on the distinguished open affines $D_{\pi_1(t)} \subset X_1$ and $D_{\pi_2(t)} \subset X_2$ with respect to the closed subscheme isomorphic to $D_{f_1^\sharp(t)=f_2^\sharp(t)} \subset Z$.

We now describe the general construction. We first construct the topological space Y by gluing together X_1 and X_2 over the isomorphism ϕ; we let Z denote the (closed) locus of gluing. We next construct the structure sheaf on Y. For every point y, we choose an affine neighborhood in the obvious way: for points $y \in X_1 \backslash Z$ (resp. $X_2 \backslash Z$) we just use a sufficiently small affine in X_1 (resp. X_2) and for points $y \in Z$ we use the construction above. Exercise 8.4.7 verifies that the local structure sheaves glue together to give a global sheaf $\mathcal{O}_Y$ giving Y the structure of a scheme.

Warning 8.4.8. Despite the apparent similarity to Construction 8.4.6, it may not be possible to glue two isomorphic closed subschemes in a single scheme X. (Rather, such a construction yields an object in the larger category of algebraic spaces.)

Since arbitrary pullbacks exist in the category of rings, one might dually hope that arbitrary pushouts exist in the category of schemes. This turns out not to be the case. The argument below shows that you cannot obtain a scheme by gluing two copies of $\mathbb{A}^1$ along the complements of the generic point – the resulting object would not have "enough" closed points to be a scheme.

Example 8.4.9 ([Bra]). Let X_1, X_2 be two copies of $\mathbb{A}^1_{\mathbb{K}}$ and let $\eta \cong \mathrm{Spec}(\mathbb{K}(t))$ denote the generic point of $\mathbb{A}^1$. We claim that the two inclusion maps $i_1, i_2 : \eta \rightrightarrows X_1 \sqcup X_2$ do not admit a coequalizer in the category of schemes.

Suppose that there were a morphism $f : X_1 \sqcup X_2 \to Y$ that coequalized i_1, i_2. We first claim that no two closed points of $X_1 \sqcup X_2$ map to the same closed point in Y. To see this, it is enough to see that for any closed points $x_1 \in X_1, x_2 \in X_2$ there is some morphism of schemes $g : X_1 \sqcup X_2 \to Z$ such that $g(x_1) \neq g(x_2)$ but $g(\eta_1) = g(\eta_2)$. If x_1, x_2 represent different points of $\mathbb{A}^1$ one can take $g : \mathbb{A}^1 \sqcup \mathbb{A}^1 \to \mathbb{A}^1$ to be the identity on each component; if they represent the same point of $\mathbb{A}^1$ one can simply add a translation to one of the components.

On the other hand, if the coequalizer Y existed then we could find an open affine $V \subset Y$ which contains an open neighborhood of the f-image of the generic points on X_1 and X_2. Thus $f^{-1}(V)$ would contain an open subset $U_1 \subset X_1$ and an open subset $U_2 \subset X_2$. Let $W \subset U_1 \cap U_2$ be a distinguished open affine that lies in the intersection of these two open sets in $\mathbb{A}^1$. We denote the copies of W in our two affine lines by $W_1 \subset X_1$ and $W_2 \subset X_2$. Note that by assumption the two maps $W_1 \to V$ and $W_2 \to V$ define the same map on the generic point of the W_i. This implies that they are the same map: since V is affine both maps are determined by a ring homomorphism $\mathcal{O}_U(U) \to \mathcal{O}_W(W)$, and since the induced maps on fraction fields are the same the original maps must also coincide. In particular, for any closed point $w \in W$ the corresponding closed points $w_1 \in W_1$, $w_2 \in W_2$ are identified by f. This contradicts the previous paragraph.

8.4.4 Monomorphisms and epimorphisms

A good starting point is to recall how these two notions behave in the category of commutative rings. A monomorphism of rings is the same as an injective homomorphism. Epimorphisms of rings are a bit subtle; some examples include surjective homomorphisms and localization maps. Since scheme morphisms are dual to ring homomorphisms, we can loosely expect these notions to "flip" when we work with schemes.

Both monomorphisms and epimorphisms of schemes are a bit subtle. In general they cannot be characterized in the same way as their topological counterparts (see Exercise 8.4.12 and Exercise 8.4.13). However under the right hypotheses we can recover the geometric situation. Here we will give two statements, without proof and using concepts we have not yet defined, explaining when our geometric intuition is correct.

Theorem 8.4.10 (EGA IV, 17.2.6)**.** *Let $f : X \to Y$ be a morphism of schemes that is locally of finite type (see Definition 8.6.14). The following are equivalent:*

(1) f is a monomorphism.

(2) The fiber over every point $y \in Y$ is either empty or is isomorphic to $\mathrm{Spec}(\kappa(y))$.

Theorem 8.4.11 (Exercise 8.7.19)**.** *Let $f : X \to Y$ be a morphism of schemes that are reduced and separated (see Definition 8.7.1). The following are equivalent:*

(1) f is an epimorphism.

(2) The image of f is set-theoretically dense in Y.

8.4.5 Relative schemes

Fix a scheme S. The categorical properties of the category $\mathbf{Sch}/S$ are the same as those of $\mathbf{Sch}$. The identity map $id : S \to S$ is the final object in this category. Relative products exist in $\mathbf{Sch}/S$: one just takes the relative product in the larger category $\mathbf{Sch}$ and notes that there is a unique map from the product to S making all the natural maps to S commute.

8.4.6 Exercises

Exercise 8.4.12. Show that the normalization of a cuspidal cubic is an injective morphism that is not a monomorphism.

Exercise 8.4.13. Find an example of a set-theoretically surjective morphism of schemes which is not an epimorphism.

Exercise 8.4.14. Suppose that $\mathbb{L}/\mathbb{K}$ is a finite Galois extension of fields. Compute $\mathrm{Spec}(\mathbb{L}) \times_{\mathrm{Spec}(\mathbb{K})} \mathrm{Spec}(\mathbb{L})$. Then compute the product $\mathrm{Spec}(\mathbb{K}(x)) \times_{\mathrm{Spec}(\mathbb{K})} \mathrm{Spec}(\mathbb{K}(y))$.

Exercise 8.4.15. Set $X = Y = \mathbb{A}^1_{\mathbb{K}}$ and consider the morphism $f : X \to Y$ defined by $x \mapsto x^2$. Compute $X \times_Y X$. (Careful: what happens in characteristic 2?)

Exercise 8.4.16. Let X be a scheme over S. Suppose that $T \to S$ is any morphism and write $X_T := X \times_S T$. Show that the diagonal maps for X/S and for X_T/T are compatible in the sense that we have a pullback diagram

$$\begin{tikzcd} X_T \arrow[r, "\Delta_{X_T/T}"] \arrow[d] & X_T \times_T X_T \arrow[d] \\ X \arrow[r, "\Delta_{X/S}"] & X \times_S X \end{tikzcd}$$

Exercise 8.4.17. Suppose we have morphisms of schemes $f_1 : X_1 \to T$, $f_2 : X_2 \to T$, and $T \to S$. Show that the following diagram is a pullback diagram:

$$\begin{tikzcd} X_1 \times_T X_2 \arrow[r, "(id{,}id)"] \arrow[d, "f_1 \circ p_1 = f_2 \circ p_2"'] & X_1 \times_S X_2 \arrow[d, "(f_1{,}f_2)"] \\ T \arrow[r, "\Delta_{T/S}"] & T \times_S T \end{tikzcd}$$

Exercise 8.4.18. Let S be a scheme. A group scheme over S is an S-scheme $\rho : G \to S$ equipped with morphisms

$$\mu : G \times_S G \to G \qquad e : S \to G \qquad i : G \to G$$

such that the maps μ, e, i satisfy the usual axioms for groups: associativity of μ and the identity and inverse axioms. Note that each of these axioms should be interpreted as the equality of two maps:

$$\begin{array}{rl} \text{(associativity)} & \mu \circ (\mu \times id) = \mu \circ (id \times \mu) : G \times G \times G \to G \\ \text{(identity)} & \mu \circ (id \times e) = proj_1 : G \times_S S \to G \\ & \mu \circ (e \times id) = proj_2 : S \times_S G \to G \\ \text{(inverse)} & e \circ \rho = \mu \circ (id \times i) \circ \Delta_{G/S} : G \to G \\ & e \circ \rho = \mu \circ (i \times id) \circ \Delta_{G/S} : G \to G \end{array}$$

Show that the following are examples of group schemes by carefully defining the maps μ, e, i using explicit ring homomorphisms. (Note that for an affine group scheme the corresponding ring homomorphisms should define a "cogroup".)

(1) The closed points of $\mathbb{G}_a(\mathbb{C}) = \mathrm{Spec}(\mathbb{C}[x])$ are in bijection with $\mathbb{C}$. Show that $\mathbb{G}_a(\mathbb{C})$ carries a group scheme structure corresponding to the usual addition in $\mathbb{C}$.

(2) The closed points of $\mathbb{G}_m(\mathbb{C}) = \operatorname{Spec}(\mathbb{C}[x, x^{-1}])$ are in bijection with $\mathbb{C}^\times$. Show that $\mathbb{G}_m(\mathbb{C})$ carries a group scheme structure corresponding to the usual multiplication in $\mathbb{C}^\times$.

(3) The closed points of $\mu_n(\mathbb{C}) = \operatorname{Spec}(\mathbb{C}[x, x^{-1}])/(x^n - 1)$ are in bijection with the nth roots of unity in $\mathbb{C}$. Show that $\mu_n(\mathbb{C})$ carries a group scheme structure corresponding to the group structure on the roots of unity.

(4) Define analogous group schemes $\mathbb{G}_a(\mathbb{Z})$, $\mathbb{G}_m(\mathbb{Z})$, $\mu_n(\mathbb{Z})$ over the integers. What do these look like explicitly?

8.5 Noetherian schemes

The definition of a manifold includes the condition that the space be second countable. Similarly, in algebraic geometry we usually focus on schemes that are not "too large". In this section we discuss several finiteness conditions that ensure that schemes exhibit nice behavior. There are two types: "finiteness" conditions for the underlying rings and finiteness conditions for the gluing process.

8.5.1 Quasicompactness

The most important fundamental finiteness condition is the following:

Definition 8.5.1. We say that a scheme X is quasicompact if it is compact in the Zariski topology.

While it is a bit perverse to introduce new notation that means the same thing as a familiar definition, this change reminds of the important fact that quasicompactness has very little to do with the geometric notion of "compactness". In practice one often uses the equivalent definition provided by the following exercise.

Exercise 8.5.2. Show that a scheme X is quasicompact if and only if it admits a finite cover by open affines. (Hint: first show that every affine scheme is quasicompact; see Exercise 1.3.2.)

Example 8.5.3. Let $X = \mathrm{Spec}(\mathbb{K}[x_1, x_2, x_3, \ldots])$. Like all affine schemes, X is quasicompact. However, the open subscheme U which is the complement of the origin is not quasicompact. Indeed, the infinite cover $\cup_{i \in \mathbb{N}} D_{x_i}$ has no finite subcover.

8.5.2 Quasicompact quasiseparated

There is a somewhat technical topological condition that is even better than quasicompactness.

Definition 8.5.4. A scheme X is said to be quasicompact quasiseparated if it admits a finite open cover $\{U_i\}_{i=1}^k$ by open affines, and furthermore for each pair of indices the intersection $U_i \cap U_j$ is covered by a finite set of open affines.

This definition is vaguely analogous to the notion of "paracompact" in the classical topology. There is one particular type of argument (see Exercise 8.5.17, Theorem 9.3.8, Exercise 9.5.28, Exercise 9.3.20) which uses this assumption in an essential way. Due to the importance of this particular argument we cannot avoid using this definition. However, since we will only use this notion infrequently as a technical assumption, we will not systematically develop the theory of schemes of this type.

8.5.3 Noetherian schemes

Many of the basic features of quasiprojective $\mathbb{K}$-schemes – for example, the decomposition into irreducible components – relied on the Noetherian property of the underlying rings. The next definition identifies a property of schemes which allows us to develop a similar theory.

Definition 8.5.5. We say that a scheme X is locally Noetherian it satisfies the following equivalent conditions:

(1) X admits an open cover by sets isomorphic to spectra of Noetherian rings

(2) Every open affine in X is isomorphic to the spectrum of a Noetherian ring.

We say that X is Noetherian if it is quasicompact and locally Noetherian.

Note that the equivalence of the two conditions follows from the fact that the property of being a Noetherian ring is local. Noetherian schemes satisfy several important finiteness properties that are not shared by every quasicompact scheme.

Exercise 8.5.6. Show that if X is a Noetherian scheme then every open subset of X is quasicompact. Deduce that every open subset of a Noetherian scheme is quasicompact quasiseparated.

Exercise 8.5.7. Show that if X is a Noetherian scheme then the topological space underlying X is a Noetherian topological space (in the sense of Exercise 1.3.16).

As with any Noetherian topological space, Noetherian schemes admit decompositions into irreducible closed subsets. In particular we can define the dimension of a Noetherian scheme using chains of closed irreducible subsets.

Definition 8.5.8. The dimension of a Noetherian scheme is the maximal length r of a chain

$$X_0 \subsetneq X_1 \subsetneq \ldots \subsetneq X_r$$

of irreducible closed subsets of X.

However, when X is not a quasiprojective $\mathbb{K}$-scheme, you must be *very careful* with the notion of dimension.

Warning 8.5.9. Although the Noetherian condition guarantees that any particular chain will have finite length, it may be that the supremum over the lengths of all such chains is infinite. In other words, an affine scheme defined by a Noetherian ring can have infinite dimension. The most famous is "Nagata's example" obtained by localizing a polynomial ring in infinitely many variables along a carefully chosen multiplicative subset.

Despite this pathology, Noetherian induction works on every Noetherian scheme.

Warning 8.5.10. In general the vanishing locus of a single equation on a Noetherian scheme X need not have dimension $\geq \dim(X) - 1$. In other words, Geometric Krull's PIT (Theorem 4.4.3) does not hold for arbitrary Noetherian schemes.

For example, set $R = \mathbb{K}[x]_{(x)}[t]$. One can show that $\dim(\mathrm{Spec}(R)) = 2$ corresponding to the chain of prime ideals $0 \subset (t) \subset (x,t)$. However, the single equation $xt - 1$ (which is neither a unit nor a zero-divisor) defines a closed subscheme of codimension 1 and of dimension 0 since $R/(xt-1)$ is a field.

The way around these pathologies is to focus on the notion of codimension. Suppose that X is an irreducible Noetherian scheme. We can define the codimension of an irreducible closed subscheme $Y \subset X$ as the maximum length r of a chain of irreducible closed subsets $Y = Z_0 \subset Z_1 \subset \ldots \subset Z_r = X$. The correct analogues for Noetherian schemes of the results we proved in Section 4.4 (Krull's PIT, dimension of fibers) will involve the codimension.

For example, Krull's PIT shows that the codimension of the vanishing locus of a function in an irreducible Noetherian scheme will have codimension at most 1. However, the codimension of Y in X and the dimension of Y may not be "complementary" in the sense that they no longer need to add up to $\dim(X)$, and in such situations Geometric Krull's PIT can fail.

8.5.4 Associated points

Let R be a Noetherian ring. A prime ideal in R is said to be an associated prime if it annihilates some element in R. An associated prime is said to be isolated if it is a minimal prime in R; otherwise we call it an embedded prime. There will only be a finite set of associated primes for R and their intersection will be $\mathrm{Nil}(R)$.

Definition 8.5.11. Let X be a Noetherian scheme. We say that a point $x \in X$ is an associated point if the maximal ideal $\mathfrak{m}_x$ is an associated prime of the local ring $\mathcal{O}_{X,x}$, or equivalently, if $\mathfrak{m}_x$ consists entirely of zerodivisors.

If we localize a Noetherian ring R along a multiplicative set S, the associated primes of $S^{-1}R$ are the associated primes of R which do not intersect S. This means that a prime ideal $\mathfrak{p} \subset R$ is an associated prime if and only if $\mathfrak{p}$ is an associated point of $\mathrm{Spec}(R)$. In particular, a point $x \in X$ is an associated point if and only if it defines an associated prime in any affine open containing it.

Exercise 8.5.12. Show that a Noetherian scheme has only finitely many associated points.

We call an associated point $x \in X$ an isolated or embedded point depending on whether $\mathfrak{m}_x$ is isolated or embedded. The behavior of these points is determined by primary decomposition. In particular:

- There is a bijection between the irreducible components of X and the isolated points x. Each isolated point is the generic point of some irreducible component of X.

- The locus of points in X which are not reduced (i.e. such that $\mathrm{Nil}(\mathcal{O}_{X,x}) \neq 0$) is the closure of the non-reduced associated points. Every embedded point is non-reduced but isolated points may or may not be reduced.

- Given any open subset $U \subset X$ and any section $f \in \mathcal{O}_X(U)$, the support of f is the intersection of U with the closure of a finite set of associated points in U (see Lemma 1.4.8).

- Given any open affine $U \subset X$ and any section $f \in \mathcal{O}_X(U)$, f will be a zero divisor if and only if $V(f)$ contains an associated point of X.

Thus the associated points play an important role in controlling the geometry of a Noetherian scheme.

8.5.5 Chevalley's Theorem

Another important result that holds for Noetherian schemes is Chevalley's Theorem.

Theorem 8.5.13. *Let $f : X \to Y$ be a finite type morphism of Noetherian schemes. Then the set-theoretic image $f(X)$ is a constructible subset of Y.*

We have seen this result before in the context of quasiprojective $\mathbb{K}$-schemes. The proof in general is similar:

(1) Using elementary topological arguments, we reduce to showing that $f(X)$ is constructible when X, Y are integral and affine and $f(X)$ is dense in Y. In particular f induces an injection $f^\sharp : K(Y) \to K(X)$.

(2) Applying Noether normalization, we see that the injection $f^\sharp : K(Y) \to K(X)$ is the composition of a transcendental extension $K(Y) \to K(Y)(g_1, \ldots, g_r)$ followed by a finite algebraic extension.

(3) Spreading out, we see that there is an open set $V \subset Y$ such that over V the morphism f is the composition of a finite morphism and a projection $\mathbb{A}^n \times V \to V$. In particular, f is surjective over V.

(4) Appeal to Noetherian induction to deduce the constructibility the intersection of the image of f with $Y \backslash V$.

8.5.6 Exercises

Exercise 8.5.14. Define $X = \mathrm{Spec}(\mathbb{C}[x, y_1, y_2, y_3, \ldots]/(y_1^2, y_2^2, y_3^2, \ldots, (x-1)y_1, (x-2)y_2, (x-3)y_3, \ldots))$. Prove that the natural map $\mathbb{A}^1_{\mathbb{C}} \to X$ is a homeomorphism. Prove that the non-reduced locus of X is a countable union of points, and in particular, is not a closed subset of X. (This is in contrast to the situation for Noetherian schemes.)

Exercise 8.5.15. Consider the ring $R = \mathbb{K}[x_1, x_2, x_3, \ldots]/(x_1, x_2^2, x_3^3, \ldots)$. Show that $\mathrm{Spec}(R)$ is a Noetherian topological space (in fact it only has one point) even though R is not a Noetherian ring. Thus, a scheme whose underlying topological space is Noetherian need not be a Noetherian scheme.

Exercise 8.5.16. An open subscheme $i : U \to X$ is said to be scheme-theoretically dense if the induced map $i^\sharp : \mathcal{O}_X \to i_* \mathcal{O}_U$ is injective.

(1) Prove that if X is reduced then an open subscheme is scheme-theoretically dense if and only if U is set-theoretically dense.

(2) Prove that if X is a Noetherian scheme then an open subscheme is scheme-theoretically dense if and only if U contains every associated point of X.

Exercise 8.5.17. Let X be a scheme and let $f \in \mathcal{O}_X(X)$. Define X_f to be the complement of the vanishing locus $V(f)$ defined in Exercise 8.2.18. That exercise showed that the restriction of f to X_f is invertible. When X satisfies some finiteness hypotheses we can say more.

(1) Suppose that X is quasicompact. Suppose that $a \in \mathcal{O}_X(X)$ is an element whose restriction to X_f is 0. Prove that for some $n > 0$ we have $f^n a = 0$ in $\mathcal{O}_X(X)$.

(2) Suppose that X is quasicompact quasiseparated. Let $b \in \mathcal{O}_X(X_f)$. Show that for some $n > 0$ the element $f^n b$ is the restriction of an element of $\mathcal{O}_X(X)$.

(Hint: for each open affine U_i we can find an integer n_i such that $f^{n_i} b \in \mathcal{O}(U_i \cap X_f)$ extends to a function g_i on U_i. Apply (1) to each open set in the finite cover of $U_i \cap U_j$ by open affines to show that if we increase n further we can ensure that the $\{g_i\}$ agree on overlaps.)

(3) Suppose that X is quasicompact quasiseparated. Then $\mathcal{O}_X(X_f) = \mathcal{O}_X(X)_f$.

Exercise 8.5.18. Suppose that $f : X \to Y$ is a morphism of Noetherian schemes inducing a sheaf map $f^\sharp : \mathcal{O}_Y \to f_* \mathcal{O}_X$. Show that the kernel of $f^\sharp$ is a quasicoherent ideal sheaf. (Hint: one strategy is to appeal to Exercise 8.5.17.)

The closed subscheme of Y defined by this ideal sheaf is called the scheme-theoretic image of f.

Exercise 8.5.19. There is no natural way to define the scheme-theoretic image of a morphism $f : X \to Y$ without some finiteness properties. For example, let $Y = \mathbb{A}^1_{\mathbb{K}}$ and let $X = \sqcup_{n \in \mathbb{N}} \mathbb{K}[x]/(x^n)$. Let $f : X \to Y$ be the morphism whose restriction to each component of X is the inclusion map. Prove that the kernel of $f^\sharp$ does not define the same closed subset as the set-theoretic closure of $f(X)$.

8.6 Properties of morphisms

One of Grothendieck's insights is that many properties of schemes are best understood as properties of morphisms. (This principle is implicit in Definition 8.2.12.) If P is a property of schemes, we will often say that a morphism $f : X \to Y$ satisfies property P if for every open affine subset $U \subset Y$ the scheme $f^{-1}(U)$ satisfies P.

In this section we will set up a general theory of properties of morphisms. It turns out that there are three key properties which all "well-behaved" properties of morphisms will satisfy.

Definition 8.6.1. Suppose that P is a property of morphisms of schemes. We say that

(1) P is preserved under composition, if whenever $f : X \to Y$ satisfies P and $g : Y \to Z$ satisfies P then $g \circ f$ satisfies P.

(2) P is stable under base change, if whenever a morphism $f : X \to Z$ satisfies P and $g : Y \to Z$ is any morphism the induced map $X \times_Z Y \to Y$ satisfies P.

(3) P is local on the target, if the following two properties hold. First, if a morphism $f : X \to Y$ satisfies P then for any open $V \subset Y$ the morphism $f|_{f^{-1}V}$ satisfies P. Second, given a morphism $f : X \to Y$ and an open covering $\{V_i\}$ of Y, if $f|_{f^{-1}V_i}$ satisfies P for every i then f satisfies P.

Remark 8.6.2. Suppose the property P is stable under base change. This implies that if a morphism $f : X \to Y$ satisfies P then so does $f : f^{-1}U \to U$ for any open subset $U \subset Y$. Thus, if P is a property that is stable under base change then to verify "local on the target" we just need to verify the second half of Definition 8.6.1.(3). We will often use this shortcut without further mention.

The most basic examples of well-behaved classes of morphisms are open and closed embeddings.

Exercise 8.6.3. Prove that open embeddings satisfy the three properties of Definition 8.6.1. (Hint: to show stable under base change, verify that if $U \to Y$ is an open embedding and $f : X \to Y$ is any morphism then $U \times_Y X$ is isomorphic to $f^{-1}U$.)

Exercise 8.6.4. Prove that closed embeddings satisfy the three properties of Definition 8.6.1. (Hint: Proposition 8.3.8 shows that closed embeddings are local on the target; this property will be useful for proving the other two.)

8.6.1 Properties of morphisms arising from the diagonal

We will be particularly interested in properties of morphisms which model topological properties. Usually a direct translation of a topological property into algebraic geometry is not so interesting due to the oddities of the Zariski topology. Instead, the best analogue is obtained by systematically using diagonal morphisms. The following lemma gives a general framework for constructing "well-behaved" properties of morphisms using the diagonal.

Lemma 8.6.5. *Suppose that P is a property of morphisms that is preserved under composition, stable under base change, and local on the target. Say that a morphism $f : X \to Y$ satisfies property Q if the induced map $\Delta_{X/Y} : X \to X \times_Y X$ satisfies P. Then property Q is also (1) preserved under composition, (2) stable under base change, and (3) local on the target.*

Proof. (1) Suppose $X \to T$ and $T \to S$ are morphisms which satisfy Q. By applying Exercise 8.4.17 with $X_1 = X_2 = X$, we obtain a pullback diagram on the right side of the commuting diagram:

$$\begin{tikzcd} X \arrow[r, "\Delta_{X/T}"] & X \times_T X \arrow[r, "g"] \arrow[d] & X \times_S X \arrow[d] \\ & T \arrow[r, "\Delta_{T/S}"] & T \times_S T \end{tikzcd}$$

By assumption $\Delta_{X/S}$ satisfies P; since P is stable under base change, g also satisfies P. Since P is closed under composition, we conclude that $\Delta_{X/S} : X \to X \times_S X$ satisfies P.

(2) Suppose that $X \to S$ satisfies P and that $T \to S$ is any morphism. Set $X_T := X \times_S T$. By Exercise 8.4.16 we have a pullback diagram

$$\begin{tikzcd} X_T \arrow[r, "\Delta_{X_T/T}"] \arrow[d] & X_T \times_T X_T \arrow[d] \\ X \arrow[r, "\Delta_{X/S}"] & X \times_S X \end{tikzcd}$$

By assumption $\Delta_{X/S}$ satisfies P. Since P is stable under base change, $\Delta_{X_T/T}$ also satisfies P.

(3) Let $f : X \to Y$ be a morphism of schemes. As in Remark 8.6.2, it suffices to show that if P holds over the open sets in an open cover of Y then it holds for f. Let $\{V_i\}$ be an open cover of Y and suppose that $f : f^{-1}V_i \to V_i$ satisfies P for every i. Consider the map $g : X \times_Y X \to Y$. By applying Lemma 2.9.1 twice we see that the preimage of V_i under this map is isomorphic to $f^{-1}V_i \times_{V_i} f^{-1}V_i$. Since f satisfies P over each V_i and since by Exercise 8.4.16 we have

$$\left(f^{-1}V_i \times_{V_i} f^{-1}V_i\right) \cap \Delta_{X/Y} \cong \Delta_{f^{-1}V_i/V_i}$$

we see that the diagonal $X \to X \times_Y X$ satisfies P over the open set $g^{-1}V_i$. As we vary i the sets $g^{-1}V_i$ form an open cover of $X \times_Y X$. Since P is local on the target, we see that $X \to X \times_Y X$ satisfies P. $\square$

There is also a useful cancellation theorem for "well-behaved" properties of morphisms.

Proposition 8.6.6. *Let P be a property of morphisms of schemes that is closed under composition and stable under base change. Suppose we have a diagram*

$$\begin{tikzcd} X \arrow[rr, "f"] \arrow[dr, "h"'] & & Y \arrow[dl, "g"] \\ & S & \end{tikzcd}$$

such that $h : X \to S$ and $\Delta_{Y/S} : Y \to Y \times_S Y$ satisfy P. Then $f : X \to Y$ satisfies P.

Proof. Applying Exercise 8.4.17 to the maps $f : X \to Y$, $id : Y \to Y$, and $g : Y \to S$, we obtain a pullback diagram

$$\begin{tikzcd} X \arrow[r, "\Gamma_{X/Y}"] \arrow[d, "f"'] & X \times_S Y \arrow[d] \\ Y \arrow[r, "\Delta_{Y/S}"] & Y \times_S Y \end{tikzcd}$$

Since P is stable under base change, the graph morphism $\Gamma_{X/Y}$ satisfies P. Similarly, since $h : X \to S$ satisfies P the the projection morphism $p_2 : X \times_S Y \to Y$ obtained by base change satisfies P. Since P is closed under composition, the composed map $f = p_2 \circ \Gamma_{X/Y} : X \to Y$ satisfies P. $\square$

8.6.2 Geometric examples

We end this section by discussing four geometric examples of well-behaved morphisms. We will only outline the proofs; careful arguments can be found in [Sta15].

Definition 8.6.7. A morphism $f : X \to Y$ is quasicompact (resp. quasicompact quasiseparated) if the preimage of every open affine $V \subset Y$ is quasicompact (resp. quasicompact quasiseparated).

Exercise 8.6.8. Prove that a morphism of affine schemes is quasicompact quasiseparated. Thus (just as with the absolute versions) these two notions pertain to the finiteness of the gluing structure.

Proposition 8.6.9. *The quasicompact property and the quasicompact quasiseparated property for morphisms are preserved under composition, stable under base change, and local on the target.*

Proof. First we consider quasicompactness. Preserved under composition is easy, and stable under base change follows quickly from local on the target. To prove local on the target, we need to show that for an open affine $V \subset Y$ the preimage $f^{-1}V$ is quasicompact if and only if for any finite set of elements $\{f_i\}$ in $\mathcal{O}_Y(V)$ which generate the unit ideal the preimage of each $f^{-1}D_{f_i}$ is quasicompact. The reverse implication is clear, since a finite union of quasicompact sets is quasicompact. For the forward implication, suppose $f^{-1}V = \cup_{j=1}^n U_j$ is a finite union of open affines. If we set $g_{ij} = f^\sharp(V)(f_i|_{U_j})$ then $f^{-1}D_{f_i} = \cup_{j=1}^n D_{g_{ij}}$ is also a finite union of open affines, hence quasicompact.

It turns out that the quasiseparated condition is equivalent to requiring that the diagonal $\Delta_{X/Y} : X \to X \times_Y X$ is a quasicompact morphism (Exercise 8.6.21). We conclude the three desired properties of quasicompact quasiseparated morphisms from Lemma 8.6.5. □

Definition 8.6.10. We say that a morphism $f : X \to Y$ is affine if for every open affine subset $V \subset Y$ the preimage $f^{-1}V$ is also affine.

Proposition 8.6.11. *The affineness property for morphisms is preserved under composition, stable under base change, and local on the target. In fact, if $f : X \to Y$ is a morphism and Y admits an open cover by open affines $\{V_i\}$ such that $f^{-1}(V_i)$ is affine, then f is affine.*

Proof. Once the last statement is proved, the first statement follows without too much difficulty. We proved this statement for quasiprojective $\mathbb{K}$-schemes in Lemma 4.1.2, and the proof in this more general context is similar. There is one important subtlety: in Lemma 4.1.2 we used the formula $\mathcal{O}_X(U \backslash V(f)) = \mathcal{O}_X(U)_f$ when U is the preimage of an open affine in Y. For general schemes this formula only holds under certain finiteness hypotheses. Under our open cover assumption, Proposition 8.6.9 shows that f is quasicompact quasiseparated and thus Exercise 8.5.17 shows that this localization formula is valid whenever U is the preimage of an open affine $V \subset Y$. □

Definition 8.6.12. A morphism $f : X \to Y$ is finite if it is affine and furthermore for every open affine $V \subset Y$ we have that $f^\sharp : \mathcal{O}_Y(V) \to \mathcal{O}_X(f^{-1}V)$ realizes $\mathcal{O}_X(f^{-1}V)$ as a finitely generated $\mathcal{O}_Y(V)$-module.

Proposition 8.6.13. *Finite morphisms are closed under composition, stable under base change, and local on the target.*

Proof. Since affine morphisms satisfy these three properties, we just need to address the module structure, and thus the desired properties follow from the algebraic results discussed in Section 8.0.1. □

Definition 8.6.14. A morphism $f : X \to Y$ is locally of finite type if for every open affine $V = \mathrm{Spec}(S)$ of Y and every open affine $U = \mathrm{Spec}(R)$ contained in $f^{-1}(V)$ the ring extension $f^\sharp : S \to R$ realizes R as a finitely generated S-algebra. It has finite type if it is locally of finite type and quasicompact.

We define locally finitely presented and finitely presented in the analogous way.

Despite the superficial similarity between finite morphisms and morphisms of finite type, the two types of morphisms play a vastly different role. "Finite type" is often a standing assumption in geometric settings; for example, every morphism of quasiprojective $\mathbb{K}$-schemes will have finite type. In contrast finite morphisms have a number of special geometric and cohomological properties.

Remark 8.6.15. There is a slightly unsettling asymmetry between the definition of finite and the definition of finite type. For finite morphisms, we needed to assume the morphism was affine and we considered the map $\mathcal{O}_Y(V) \to \mathcal{O}_X(f^{-1}V)$. For finite type morphisms, we do not need an affineness assumption and we allowed ourselves to take any open affine U in $f^{-1}V$ and consider the map $\mathcal{O}_Y(V) \to \mathcal{O}_X(U)$

The reason for this asymmetry is that finite type morphisms satisfy a stronger localization property than finite morphisms. Thus we can get away with a weaker set of assumptions when we define them; see the following proof for more details.

Proposition 8.6.16. *Morphisms of finite type (or locally of finite type, or finitely presented, or locally finitely presented) are closed under composition, stable under base change, and local on the target.*

Proof. The key realization is that the localization result for ring extensions of this type in Section 8.0.1 is *stronger* than the localization result for finite homomorphisms. It implies that morphisms of this type are local on the domain – if we take an open cover $\{U_i\}$ of X and every map $U_i \to Y$ is locally of finite type, then $X \to Y$ is also locally of finite type. The three properties then follow easily from the algebraic results discussed in Section 8.0.1. □

8.6.3 Exercises

Exercise 8.6.17. Let $f : X \to \mathbb{P}^2$ denote the blow-up of a point. Show that f is not an affine morphism.

Exercise 8.6.18. Suppose that $f : X \to Z$ is a finite morphism. Prove that the fiber of f over every point of Z is finite.

Exercise 8.6.19. Prove that an open embedding is locally of finite type. Prove that a closed embedding has finite type.

Exercise 8.6.20. Suppose that $f : X \to Y$ is a morphism. Exercise 8.4.17 shows that if U, V are open subschemes of X then the following diagram is a pullback diagram:

$$\begin{tikzcd}
U \times_X V \arrow[r] \arrow[d] & U \times_Y V \arrow[d] \\
X \arrow[r, "\Delta_{X/Y}"] & X \times_Y X
\end{tikzcd}$$

Show that the downwards map on the right is an open embedding. Deduce that if U, V are open sets in X, then the intersection $U \cap V$ is isomorphic to $\Delta_{X/Y} \cap (U \times_Y V)$. (This result is most intuitive when f is the map to $\mathrm{Spec}(\mathbb{Z})$, where we see that $U \cap V \cong \Delta \cap (U \times V)$.)

Exercise 8.6.21. Let $f : X \to Y$ be a quasicompact morphism. Prove that f is quasicompact quasiseparated if and only if $\Delta_{X/Y} : X \to X \times_Y X$ is quasicompact. (Hint: use Exercise 8.6.20.)

More generally, an arbitrary morphism f is said to be quasiseparated if the diagonal morphism is quasicompact. We will not use this concept without the quasicompactness hypothesis.

8.7 Separatedness

Another basic assumption we impose on manifolds is the Hausdorff condition. In this section we analyze the analogue of this property in algebraic geometry. As discussed in Exercise 1.3.1, the Zariski topology on a scheme is almost never Hausdorff. In order to translate "Hausdorffness" into algebraic geometry, we will rely on an alternative formulation: a topological space X is Hausdorff if and only if the diagonal $\Delta \subset X \times X$ is a closed subset. This definition generalizes well to algebraic geometry.

Definition 8.7.1. Let $f : X \to Y$ be a morphism of schemes. We say that f is separated if the diagonal morphism $\delta : X \to X \times_Y X$ is a closed embedding.

If X is an S-scheme, we say that X is separated over S if the structural morphism $X \to S$ is separated. (When S is not specified, we assume $S = \mathrm{Spec}(\mathbb{Z})$ is the final object in the category of schemes.)

Remark 8.7.2. Note that the algebraic geometric definition of separatedness does not simply recover the topological notion of Hausdorfness because the product in **Sch** is not compatible with the topological product of the underlying topological spaces.

Remark 8.7.3. Note that if a morphism is quasicompact and separated then it is quasicompact quasiseparated.

Exercise 8.7.4. Show that every morphism of affine schemes is separated.

Use this property to show for any morphism $f : X \to Y$ the set-theoretic image of $\Delta_{X/Y}$ is a locally closed subset of $X \times_Y X$ (i.e. a closed subset of an open set).

Separatedness has all three of our standard properties for morphisms.

Theorem 8.7.5. *Separatedness is preserved by composition, stable under base change, and local on the target.*

Proof. This follows from Lemma 8.6.5 and Exercise 8.6.4. □

The next result shows that the separatedness condition is a purely topological condition on the diagonal.

Proposition 8.7.6. *Let $f : X \to Y$ be a morphism. Then f is separated if and only if the set-theoretic image of the diagonal map $\Delta_{X/Y}$ is a closed subset of $X \times_Y X$.*

Proof. We need to show that if the image of $\Delta_{X/Y}$ is closed then $\Delta_{X/Y}$ is a homeomorphism and $\Delta^\sharp_{X/Y}$ is surjective. In fact, since the composition of $\Delta_{X/Y}$ with the first projection map $\pi_1 : X \times_Y X \to X$ is the identity map $\Delta_{X/Y}$ is always a homeomorphism. To see that $\Delta^\sharp_{X/Y} : \mathcal{O}_{X \times_Y X} \to \Delta_{X/Y,*}\mathcal{O}_X$ is surjective, it suffices to prove that the map is surjective in an affine neighborhood of a point in the set-theoretic image of $\Delta_{X/Y}$. For any point

$x \in X$, we can choose an open affine neighborhood V of $f(x)$ and an open neighborhood $U \subset f^{-1}(V)$ of x. Then $U \times_V U$ is affine and so by Exercise 8.7.4 the restriction of $\Delta_{X/Y}$ to U is a closed embedding. This shows that $\Delta^\sharp_{X/Y}$ is surjective locally around the point $\Delta_{X/Y}(x)$, finishing the proof. □

8.7.1 Examples

We next discuss some basic examples of separated and non-separated morphisms. One important example is that closed embeddings and open embeddings are separated (see Exercise 8.7.16).

Example 8.7.7. The "line with a doubled origin" described in Example 8.2.6 is non-separated. Let's recall the construction. Let $U \subset \mathbb{A}^1_{\mathbb{C}}$ denote the complement of the origin. We can define a scheme X by taking two copies of $\mathbb{A}^1_{\mathbb{C}}$ and identifying the open set U in each copy using the identity isomorphism. This scheme X will look like $\mathbb{A}^1_{\mathbb{C}}$ with the origin "doubled". In other words, the points of X are $(\mathbb{A}^1\backslash\{0\}) \cup \{0_1, 0_2\}$.

Let's check that the structure map $X \to \operatorname{Spec}(\mathbb{K})$ is not separated. The diagonal in $X \times_{\operatorname{Spec}(\mathbb{K})} X$ will consist of all points of the form $\{(x,x) | x \in X\}$. However, the closure will also contain the points $(0_1, 0_2)$ and $(0_2, 0_1)$ which are not contained in the diagonal. (Check this claim carefully using affine charts!) Thus the diagonal is not closed.

Example 8.7.8. The diagonal map for projective space $\Delta_{\mathbb{P}^n} : \mathbb{P}^n \to \mathbb{P}^n \times \mathbb{P}^n$ will have closed image. In fact, according to Exercise 3.3.10 the image is the intersection of the Segre embedding of $\mathbb{P}^n \times \mathbb{P}^n$ with a linear space. (We will reprove this in Proposition 8.7.9.)

By leveraging the argument for projective space, we can show that:

Proposition 8.7.9. *Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is finitely generated as an S_0-algebra. Then* $\operatorname{Proj}(S)$ *is separated.*

Since open embeddings are separated, Proposition 8.7.9 in turn implies that every quasiprojective $\mathbb{K}$-scheme is separated.

Proof. Exercise 8.7.4 shows that $\operatorname{Spec}(S_0)$ is separated. Since separatedness is preserved by composition it suffices to show that the structural morphism $\operatorname{Proj}(S) \to \operatorname{Spec}(S_0)$ is separated. By Exercise 2.7.10 there is some positive integer n such that this morphism factors as

$$\operatorname{Proj}(S) \to \mathbb{P}^n_{S_0} \to \operatorname{Spec}(S_0)$$

where the first morphism is a closed embedding. The first map is separated by Exercise 8.7.16, thus it suffices to prove that the projection map $\mathbb{P}^n_{S_0} \to \operatorname{Spec}(S_0)$ is separated.

We only need to show that the diagonal is a closed subscheme of $\mathbb{P}^n_{S_0} \times_{\operatorname{Spec}(S_0)} \mathbb{P}^n_{S_0}$. We can cover $\mathbb{P}^n_{S_0} \times_{\operatorname{Spec}(S_0)} \mathbb{P}^n_{S_0}$ with affine charts of the form $U_{ij} := D_{+,x_i} \times_{\operatorname{Spec}(S_0)} D_{+,x_j}$.

The intersection of the diagonal with U_{ij} is isomorphic to $D_{+,x_i} \cap D_{+,x_j}$ and the map $\Delta \cap U_{ij} \to U_{ij}$ is given by the surjection of rings

$$\begin{aligned} S_0[\tfrac{x_0}{x_i}, \ldots, \tfrac{x_n}{x_i}] \otimes_{S_0} S_0[\tfrac{y_0}{y_j}, \ldots, \tfrac{y_n}{y_j}] &\to S_0[\tfrac{z_0}{z_i}, \ldots, \tfrac{z_n}{z_i}, \tfrac{z_i}{z_j}] \\ \tfrac{x_k}{x_i} &\mapsto \tfrac{z_k}{z_i} \\ \tfrac{y_k}{y_j} &\mapsto \tfrac{z_k}{z_i} \cdot \tfrac{z_i}{z_j} \end{aligned}$$

In particular, $\Delta \cap U_{ij}$ is a closed subscheme of U_{ij}. Since closed embeddings are local on the target, by varying i, j we see that $\mathbb{P}^n_{S_0}$ is separated. $\square$

8.7.2 Properties of separated schemes

We will next prove a couple useful properties of separated schemes.

Proposition 8.7.10. *Suppose that $f : X \to \operatorname{Spec}(R)$ is a separated morphism. Then for any open affines U, V in X the intersection $U \cap V$ is also an open affine.*

Proof. Exercise 8.6.20 shows that $U \cap V$ is isomorphic to

$$(U \times_{\operatorname{Spec}(R)} V) \cap \Delta_{X/\operatorname{Spec}(R)}.$$

Note that $U \times_{\operatorname{Spec}(R)} V$ is affine since all three terms are affine. Using the separatedness assumption we see that $(U \times_{\operatorname{Spec}(R)} V) \cap \Delta$ is a closed subscheme of $U \times_{\operatorname{Spec}(R)} V$ so that it is also affine. Thus $U \cap V$ is affine. $\square$

If $f, g : X \to Y$ are two continuous maps of topological spaces and Y is Hausdorff, then the equalizer $\{x \in X | f(x) = g(x)\}$ is closed in X. The following statement gives an analogue in the setting of algebraic geometry.

Proposition 8.7.11. *Suppose that $f, g : X \to Y$ are two morphisms of S-schemes such that Y is separated over S. Suppose there is an open subscheme $i : U \to X$ with scheme-theoretically dense image such that $f \circ i = g \circ i$. Then $f = g$.*

Recall from Exercise 8.5.16 that an open subscheme $i : U \to X$ is said to be scheme-theoretically dense if the map $i^\sharp : \mathcal{O}_X \to i_*\mathcal{O}_U$ is an injection.

Proof. It suffices to prove the statement when X is an affine scheme. The pair of maps $f, g : X \to Y$ yield a morphism $h : X \to Y \times_S Y$. Define E via the pullback diagram

$$\begin{tikzcd} E \arrow[r, "\pi_2"] \arrow[d, "\pi_1"'] & Y \arrow[d, "\Delta_{Y/S}"] \\ X \arrow[r, "h"] & Y \times_S Y \end{tikzcd}$$

Note that since Y is separated π_1 is a closed embedding.

Consider the map $i : U \to X$. The pair of maps $(i, f \circ i = g \circ i)$ induces a morphism $\psi : U \to E$ through which i factors. Since π_1 is a closed embedding, $\mathcal{O}_X \to \pi_{1*}\mathcal{O}_E$ is a surjection. Since the open embedding i factors through E, the map $E \to X$ is set theoretically bijective and $\mathcal{O}_X \to \pi_{1*}\mathcal{O}_E$ is also an injection, hence an isomorphism. We conclude that h factors through the diagonal $\Delta_{Y/S}$ showing that $f = g$. □

8.7.3 Valuative criterion

One way of describing the Hausdorff property is to require that all convergent sequences of points have a unique limit. The following criterion describes the analogous statement in algebraic geometry. Recall that a valuation ring R is an integral domain such that for every element f in the fraction field either f or f^{-1} is contained in R. As discussed in Example 8.1.16, such rings are loosely analogous to the role of small open sets of curves in topology.

Theorem 8.7.12 (Valuative criterion for separatedness)**.** *Let $f : X \to Y$ be a morphism of schemes such that X is locally Noetherian. Then f is separated if and only if for any valuation ring R with field of fractions K and any diagram*

$$\begin{tikzcd} \mathrm{Spec}(K) \arrow[r] \arrow[d] & X \arrow[d, "f"] \\ \mathrm{Spec}(R) \arrow[r] & Y \end{tikzcd}$$

there is at most one morphism $\mathrm{Spec}(R) \to X$ *making the diagram commute.*

Note that the forward implication follows from Proposition 8.7.11. The reverse implication is more difficult.

Consider the composition $\mathrm{Spec}(K) \to X \to Y$. The map $\mathrm{Spec}(R) \to Y$ represents a "small arc of a curve" through this K-point, and the valuative criterion requires that there is at most one "lift" of this arc to X. Thus the valuative criterion is analogous to the description of the Hausdorff property via uniqueness of limits. Although this result is important conceptually, we will not use this theorem and so refer to [Har77, Theorem II.4.3] for a proof.

Remark 8.7.13. If f is a finite type morphism of Noetherian schemes, then it suffices to check the valuative criterion for all DVRs.

8.7.4 $\mathbb{K}$-schemes

Definition 8.7.14. Let $\mathbb{K}$ be a field. A $\mathbb{K}$-scheme is a scheme X equipped with a morphism $X \to \mathrm{Spec}(\mathbb{K})$ that is finite type and separated.

Remark 8.7.15. Some authors call the schemes satisfying Definition 8.7.14 “varieties over $\mathbb{K}$”. We have chosen instead to reserve the term “variety” for those $\mathbb{K}$-schemes which are also irreducible and reduced. Some authors only use the term “variety” when the ground field $\mathbb{K}$ is algebraically closed. All these usages are common in the literature; be careful to keep track!

$\mathbb{K}$-schemes are the basic “geometric” objects over the field $\mathbb{K}$. Amongst all $\mathbb{K}$-schemes, we have two important subclasses:

- Projective $\mathbb{K}$-schemes: the $\mathbb{K}$-schemes which admit a closed embedding into projective space. (As shown in Theorem 2.8.12, these are also the $\mathbb{K}$-schemes isomorphic to $\mathrm{Proj}(R)$ where R is a finitely generated graded $\mathbb{K}$-algebra with $R_0 \cong \mathbb{K}$.)

- Quasiprojective $\mathbb{K}$-schemes: the $\mathbb{K}$-schemes which admit an open embedding into a projective $\mathbb{K}$-scheme.

The fact that these are $\mathbb{K}$-schemes follows from Proposition 8.7.9.

8.7.5 Exercises

Exercise 8.7.16. (1) Show that the diagonal and the graph of a morphism are monomorphisms in the category of schemes (even though they need not be closed or open embeddings).

(2) Show that if $f : X \to Y$ is a monomorphism then $\Delta_{X/Y} : X \to X \times_Y X$ is an isomorphism. Conclude that monomorphisms are separated. In particular, closed embeddings, open embeddings, diagonal morphisms, and graph morphisms are separated.

(3) Suppose $f : X \to Y$ and $g : Y \to Z$ are morphisms of schemes. Show that if $g \circ f$ is separated then f is separated.

Exercise 8.7.17. Suppose $f : X \to Y$ and $g : Y \to Z$ are morphisms of schemes such that g is separated. Use Proposition 8.6.6 to prove the following statements.

(1) If $g \circ f$ is a closed embedding then f is a closed embedding.

(2) If $g \circ f$ is an affine morphism then f is an affine morphism.

(3) If $g \circ f$ is a finite morphism then f is a finite morphism.

Consider the map $g : X \to \mathbb{A}^1$ where X denotes the line with the doubled origin and g is the function which “glues” the two origins in X. Use this morphism to show that all three of the statements can fail when g is not separated.

Exercise 8.7.18. Give an example of a non-separated scheme X and open affines U, V in X such that $U \cap V$ is not affine.

Exercise 8.7.19. Prove that a morphism $f : X \to Y$ of separated reduced schemes is an epimorphism if and only if the image of f is dense in Y. (Hint: for the forward implication, let Z denote the closure of the image of f. Consider the two maps from Y to the scheme $Y \coprod_Z Y$ constructed in Construction 8.4.6.)

8.8 Properness

Amongst all manifolds, the compact manifolds have the nicest geometric properties. Since most of the schemes we work with are (quasi)compact in the Zariski topology, we are led to look for a different way of formulating an analogous construction for schemes.

Definition 8.8.1. A morphism $f : X \to Y$ is universally closed if for every morphism $Z \to Y$ the induced map $X \times_Y Z \to Z$ is topologically closed.

As we discussed in Section 2.11, the universally closed condition is closely analogous the notion of "properness" in topology, and can thus be used to recapture our intuition concerning compactness.

Definition 8.8.2. A morphism of schemes $f : X \to Y$ is proper if it is separated, finite type, and universally closed.

If X is an S-scheme, we say that X is proper over S if the structural morphism $X \to S$ is separated. (When S is not specified, we assume $S = \mathrm{Spec}(\mathbb{Z})$ is the final object in the category of schemes.)

Note that separatedness is included as an assumption, so that proper schemes are analogous to compact Hausdorff manifolds. Since all three conditions in Definition 8.8.2 are well-behaved properties of morphisms, we obtain:

Theorem 8.8.3. *Proper morphisms are preserved under composition, stable under base change, and local on the target.*

8.8.1 Examples

We next discuss several basic examples of proper maps.

Exercise 8.8.4. Prove that closed embeddings are proper.

The following result shows that projective $\mathbb{K}$-schemes are proper.

Theorem 8.8.5. *Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is finitely generated as an S_0-algebra. Then the structure map $\mathrm{Proj}(S) \to \mathrm{Spec}(S_0)$ is proper.*

Proof. Applying Exercise 8.8.4 and using Theorem 8.8.3 we reduce to showing that $\mathbb{P}^n_{R_0} \to \mathrm{Spec}(R_0)$ is proper. This follows from the Fundamental Theorem of Elimination Theory as described in Theorem 2.11.7 and Exercise 2.11.9. □

Proposition 8.8.6. *Finite morphisms are proper.*

Proof. Finite morphisms are affine, hence separated by Exercise 8.7.4. It is also clear that a finite morphism has finite type. Finally, one can show that finite morphisms are topologically closed using the Going Up theorem; see Lemma 4.1.12. Since finite morphisms are also stable under base change, this means they are in fact universally closed. □

8.8.2 Proper versus projective

Theorem 8.8.5 shows that projective $\mathbb{K}$-schemes are proper. However, there are examples of proper $\mathbb{K}$-schemes which are not projective. (Unfortunately, we do not yet have the tools to give an example.) The following important result shows that proper $\mathbb{K}$-schemes are not too far from projective $\mathbb{K}$-schemes: any proper $\mathbb{K}$-variety is birationally equivalent to a projective $\mathbb{K}$-variety. More precisely:

Theorem 8.8.7 (Chow's Lemma)**.** *Let S be a Noetherian scheme. Suppose that $f : X \to S$ is proper. Then there exists an S-scheme X' such that:*

(1) X' admits a proper surjective map $\phi : X' \to X$ that is an isomorphism over a dense open subset of X.

(2) For some positive integer n the structure map $X' \to S$ factors through a closed embedding $X' \hookrightarrow \mathbb{P}^n_S$.

It turns out that every $\mathbb{K}$-scheme admits an open embedding into a proper $\mathbb{K}$-scheme. (In other words, the relationship between quasiprojective $\mathbb{K}$-schemes and projective $\mathbb{K}$-schemes is the same as the relationship between arbitrary $\mathbb{K}$-schemes and proper $\mathbb{K}$-schemes.) This is a consequence of:

Theorem 8.8.8 (Nagata's Compactification Theorem)**.** *Let S be a Noetherian scheme. Suppose that $f : X \to S$ is separated and finite type. Then f factors as an open embedding $X \hookrightarrow P$ composed with a proper map $P \to S$.*

8.8.3 Global sections

One of the most important properties of proper varieties is a finiteness result for the ring of global sections.

Theorem 8.8.9. *Let X be a scheme equipped with a proper morphism $f : X \to \operatorname{Spec}(R)$. Then $\mathcal{O}_X(X)$ is integral over R.*

The argument is conceptually the same as the argument for Corollary 2.11.11: to any global function on X we associate a morphism $X \to \mathbb{A}^1_R$ and use the inclusion $\mathbb{A}^1_R \subset \mathbb{P}^1_R$ to prove the result.

Proof. Suppose $f \in \mathcal{O}_X(X)$. If f is nilpotent, then certainly it is integral over R. Otherwise, we claim there is an open subset $U \subset X$ such that $f|_U$ is invertible. Indeed, since f has finite type over an affine scheme we see that X is quasicompact. If f vanishes along all of X, then the restriction of f to every open affine is nilpotent, and by choosing a finite open cover of X by open affines we conclude that there is some global N such that $f^N = 0$. Thus any non-nilpotent function in $\mathcal{O}_X(X)$ has $V(f) \subsetneq X$ and by Exercise 8.2.18 the restriction of f to the complement of $V(f)$ is invertible.

By Theorem 8.3.6 we get a $\mathrm{Spec}(R)$-morphism $g : U \to \mathbb{A}^1_R$ corresponding to the ring map $g^\sharp : R[t] \to \mathcal{O}_X(U)$ sending $t \mapsto f^{-1}$. By composing the graph $\Gamma : U \to U \times_{\mathrm{Spec}(R)} \mathbb{A}^1_R$ with the open embedding $U \times_{\mathrm{Spec}(R)} \mathbb{A}^1_R \to X \times_{\mathrm{Spec}(R)} \mathbb{A}^1_R$ we obtain a map $h : U \to X \times_{\mathrm{Spec}(R)} \mathbb{A}^1_R$. We claim that the set-theoretic image of this map is a closed subset G. In fact, since f^{-1} is identified with t along G we see that G is the preimage of the hyperbola $V(xy-1)$ under the map $(f, id) : X \times_{\mathrm{Spec}(R)} \mathbb{A}^1_R \to \mathbb{A}^1_R \times_{\mathrm{Spec}(R)} \mathbb{A}^1_R$.

Since $X \to \mathrm{Spec}(R)$ is proper, the map $\pi : X \times_{\mathrm{Spec}(R)} \mathbb{A}^1_R \to \mathbb{A}^1_R$ is topologically closed. Thus the image $\pi(G)$ is a closed subset of $\mathbb{A}^1_R$ and is defined by some ideal I. Since we defined G by sending t to an invertible function, $\pi(G)$ also must be contained in the complement of $V(t)$. This implies that $I + (t) = R[t]$. In particular, we can write

$$1 = r(t) + ts(t)$$

for some polynomials $r(t) \in I$ and $s(t) \in R[t]$. Since $r(t)$ vanishes on the image of $f^{-1} : U \to \mathbb{A}^1$, by pulling back this equation via f^{-1} we see that

$$1 = 0 + f^{-1}s(f^{-1}).$$

After multiplying by an appropriate power of f to clear denominators we see that f is integral over R. □

8.8.4 Valuative criterion

Just as with separatedness, one can define properness via a valuative criterion. In a Hausdorff manifold, the compactness property corresponds to the existence of limits of sequences. Thus, in the valuative criterion of properness we add an existence condition to the uniqueness we require for separatedness.

Theorem 8.8.10 (Valuative criterion for properness)**.** *Let $f : X \to Y$ be a morphism of schemes of finite type such that X is locally Noetherian. Then f is proper if and only if for any valuation ring R with field of fractions K and any diagram*

$$\begin{tikzcd} \mathrm{Spec}(K) \arrow[r] \arrow[d] & X \arrow[d, "f"] \\ \mathrm{Spec}(R) \arrow[r] & Y \end{tikzcd}$$

there is exactly one morphism $\mathrm{Spec}(R) \to X$ making the diagram commute.

We refer to [Har77, Theorem II.4.7] for a proof.

Remark 8.8.11. If f is a finite type morphism of Noetherian schemes, then it suffices to check the valuative criterion for all DVRs.

The most important and commonly used consequence of the valuative criterion is the following. By a curve over a field $\mathbb{K}$, we will mean a 1-dimensional integral $\mathbb{K}$-scheme.

Theorem 8.8.12. *Let X be a proper $\mathbb{K}$-scheme. Suppose that C is a regular curve over $\mathbb{K}$. Then any rational map $\phi : C \dashrightarrow X$ extends to a morphism $f : C \to X$.*

Of course the restriction to $\mathbb{K}$-schemes is not necessary; any situation where we have a regular scheme C that is locally defined by a Dedekind ring should work just as well. However, the regularity hypothesis is essential (see Exercise 8.8.17).

Proof. Suppose that $p \in C$ is a point where f is not defined. Since C is regular, the stalk $\mathcal{O}_{C,p}$ is a DVR. Let Z denote $\operatorname{Spec}(\mathcal{O}_{C,p})$ and let ξ denote $\operatorname{Spec}(K(C))$. By the valuative criterion the map $f : \xi \to X$ extends to a morphism $\overline{f} : Z \to X$.

Let $U \subset X$ be an open affine that contains the $\overline{f}$-image of the unique closed point in Z and let $V \subset C$ be an open affine neighborhood of p such that $V \backslash p$ is contained in $f^{-1}(U)$. We have two ring homomorphisms:

$$f^{\sharp}(U) : \mathcal{O}_X(U) \to \mathcal{O}_C(V \backslash p)$$
$$\overline{f}^{\sharp}(U) : \mathcal{O}_X(U) \to \mathcal{O}_{C,p}$$

and by construction these maps are the same map when we post-compose by the map to the function field of C. This implies that the image of the map $\overline{f}^{\sharp}(U) : \mathcal{O}_X(U) \to \mathcal{O}_{C,p}$ actually lies in $\mathcal{O}_C(V)$. In this way we obtain a morphism $\overline{f} : V \to U$ which agrees with f on $V \backslash p$. By gluing f and $\overline{f}$ we obtain a rational map from C to X that is well-defined at p and is in the same equivalence class of rational maps as f. $\square$

We have the following interesting consequence:

Corollary 8.8.13. *Let X be a regular curve over a field $\mathbb{K}$. Then X is proper over $\mathbb{K}$ if and only if it is a projective $\mathbb{K}$-scheme.*

Here the regularity hypothesis is not essential; it turns out that every proper curve is projective.

Proof. We first show that any regular curve is birational to a regular projective curve. Let C be a regular curve and let $U \subset C$ be an open affine. Then there is a closed embedding $U \to \mathbb{A}^n_{\mathbb{K}}$. By composing with $\mathbb{A}^n_{\mathbb{K}} \to \mathbb{P}^n_{\mathbb{K}}$ and taking the scheme theoretic image, we see that U admits an open embedding into a projective curve Y. The normalization map $\nu : Z \to Y$ will be an isomorphism over U. The curve Z will be regular (since an integrally closed DVR is a regular local ring) and will be projective by Construction 5.6.14, thus it is a regular projective curve birational to C.

Suppose now that X is a proper curve. We have already shown that there is a birational morphism $\phi : X \dashrightarrow Z$ to a regular projective curve. By Theorem 8.8.12 ϕ and ϕ^{-1} both extend to morphisms g and g' respectively. Since the compositions $g \circ g'$ and $g' \circ g$ agree with the identity map on an open subset, by Proposition 8.7.11 g and g' are inverses. This shows that X is isomorphic to Z. $\square$

Remark 8.8.14. In fact, the argument of Corollary 8.8.13 shows that every regular $\mathbb{K}$-curve is birationally equivalent to a unique regular projective curve.

8.8.5 Exercises

Exercise 8.8.15. Suppose $f : X \to Y$ and $g : Y \to Z$ are morphisms of schemes such that g is separated. Show that if $g \circ f$ is proper then f is proper.

Exercise 8.8.16. Let X be a normal projective $\mathbb{K}$-variety and let Y be a projective $\mathbb{K}$-variety. Suppose we have a rational map $\phi : X \dashrightarrow Y$. Prove that ϕ can be extended to each point in X whose closure is a codimension 1 subvariety of X. Deduce that if U is the largest locus where ϕ can be defined then the complement $X \backslash U$ has codimension ≥ 2 in X.

Exercise 8.8.17. Find a (non-regular) curve C defined over a field $\mathbb{K}$, a proper $\mathbb{K}$-scheme X, and a rational map $f : C \dashrightarrow X$ which does not extend to a morphism.

Chapter 9

Sheaves of modules

The motivation for studying sheaves of modules comes from the geometric theory of vector bundles. Recall that a vector bundle $\pi : \mathcal{V} \to X$ is a locally trivial fibration whose fibers can be identified with $\mathbb{R}^r$. Given two vector bundles $\pi : \mathcal{V} \to X$ and $\rho : \mathcal{U} \to X$, a map of vector bundles is a commuting diagram

$$\begin{tikzcd} \mathcal{V} \arrow[rr, "f"] \arrow[dr, "\pi"'] & & \mathcal{U} \arrow[dl, "\rho"] \\ & X & \end{tikzcd}$$

such that f restricts to a linear map f_x over each point $x \in X$. In general the rank of the linear map f_x need not be constant as x varies. Thus, there is no natural notion of a "kernel" or a "image" of a vector bundle map unless the rank happens to be constant. The nicest situation is when we have a short exact sequence of vector bundles, i.e. a sequence of vector bundle maps

$$0 \to \mathcal{V} \xrightarrow{f} \mathcal{U} \xrightarrow{g} \mathcal{W} \to 0$$

such that over every point in x the fibers yield a short exact sequence of vector spaces. It turns out that every SES of vector bundles splits, i.e. there is a vector bundle map $s : \mathcal{W} \to \mathcal{U}$ such that $g \circ s = id_{\mathcal{W}}$.

One can develop a similar theory of vector bundles in the setting of schemes using the Zariski topology in place of the usual topology and using $\mathbb{A}^r_{\mathbb{Z}}$ in place of $\mathbb{R}^r$. Since the Zariski topology is much coarser than the topologies used in geometry, there are a few key differences. For example, it is much harder for a map of vector bundles to split in the Zariski topology: a surjection of vector bundles no longer needs to admit a splitting.

9.0.1 Vector bundles and modules

In this chapter we will take an alternative approach: we will expand the category of vector bundles to the larger category of sheaves of modules on X (known as the category of $\mathcal{O}_X$-

modules). This is an abelian category that combines the algebraic techniques from module theory with the geometric intuition behind vector bundles. It has one essential advantage: one can always take kernels and cokernels of a map of $\mathcal{O}_X$-modules (although the resulting objects need not be vector bundles). Since the connection between vector bundles and sheaves of modules is not obvious, we will spend a little time developing it.

Suppose that X is a scheme. Recall that a vector bundle $\pi : \mathcal{V} \to X$ is "locally isomorphic" to a product in the Zariski topology. In other words, for any sufficiently small open affine $U \subset X$ with ring of functions $R = \mathcal{O}_X(U)$ and preimage $\mathcal{V}_U = \pi^{-1}U$ there is an isomorphism

$$\psi_U : \mathcal{V}_U \cong U \times \mathbb{A}^r_{\mathbb{Z}} = \mathrm{Spec}(R[x_1, \ldots, x_r]).$$

We now will make an essential change in perspective. Instead of working with the entire ring $R[x_1, \ldots, x_r]$, it suffices just to remember the degree 1 part of the ring, i.e. the module

$$M_U := Rx_1 \oplus Rx_2 \oplus \ldots \oplus Rx_r.$$

This has the key advantage of replacing a "large" R-module by a finitely generated R-module. As we vary the open set U, the resulting modules M_U will naturally yield a sheaf $\mathcal{M}$ of abelian groups on X. Thus we have the additional advantage that our construction now yields an object in an abelian category.

Let's analyze the geometric meaning of this transition. Remember, the ring $R[x_1, \ldots, x_r]$ is the ring of functions on the scheme $\mathcal{V}_U$. A section $\sigma : U \to \mathcal{V}_U$ will be defined by an R-algebra homomorphism $R[x_1, \ldots, x_r] \to R$, or equivalently, by an R-module homomorphism $Rx_1 \oplus \ldots \oplus Rx_r \to R$. Note that such a homomorphism is *not* an element of the module M_U, but rather of the dual module $M_U^\vee$. We can then glue the $M_U^\vee$ together to yield a sheaf $\mathcal{M}^\vee$. Thus:

We systematically replace a vector bundle $\pi : \mathcal{V} \to X$ by its sheaf of sections $\mathcal{M}^\vee$.

Although the presence of the dual is a bit surprising, it is there for a good reason: it represents the fact that the space of linear functions on a vector space is naturally identified with the dual space.

Note that we can always recover the vector bundle $\mathcal{V}_U$ from the module $M_U = Rx_1 \oplus \ldots \oplus Rx_r$ by taking the Spec of the symmetric algebra $\mathrm{Sym}(M_U)$. But most of the time we won't need to bother reconstructing the "geometric" vector bundle $\mathcal{V}$.

9.0.2 Algebra preliminaries

Definition 9.0.1. Let R be a ring. An R-module M is called:

(1) Finitely generated, if there is a positive integer b and a surjection $R^{\oplus b} \to M$.

(2) Finitely presented, if there are positive integers a, b and an exact sequence $R^{\oplus a} \to R^{\oplus b} \to M \to 0$.

(3) Coherent, if M is finitely generated and for *any* positive integer c and morphism $R^{\oplus c} \to M$ the kernel is finitely generated.

When R is a Noetherian ring these three definitions are all equivalent, but in general they are successively stronger.

Warning 9.0.2. When we leave the realm of Noetherian rings, these notions can exhibit pathological behavior. For example, R need not be a coherent R-module. Even worse, an arbitrary ring need not have any non-zero coherent modules at all (see [Zho]).

The notion of coherent is mostly useful for rings R which are coherent over themselves; for such rings coherent is the same as finitely presented.

When we leave the setting of Noetherian rings, the notion of finitely generated no longer behaves well. Here is a summary of some of the nice properties of coherent modules.

Suppose R is a ring and M, N are R-modules which satisfy one of our three properties. We first discuss which operations on M, N preserve these properties.

finitely generated	finitely presented	coherent	if M and N have property P then:
✓	✓	✓	does $M \otimes N$?
✗	✗	✓	does $\mathrm{Hom}(M, N)$?
✗	✗	✓	does $\ker(\phi : M \to N)$?
✓	✓	✓	does $\mathrm{cok}(\phi : M \to N)$?
✓	✓	✓	does an extension of M by N?

Table 9.1: Finiteness properties over fixed ring

In particular, the coherent R-modules always form an abelian subcategory of $R-\mathbf{Mod}$. In fact, we can classify when the other types of modules form abelian categories:

Theorem 9.0.3 ([CEH])**.** *Let R be a ring. The following are equivalent:*

(1) The finitely generated R-modules form an abelian category.

(2) An R-module is finitely generated if and only if it is coherent.

(3) R is Noetherian.

Theorem 9.0.4 ([Rob])**.** *Let R be a ring. The following are equivalent:*

(1) The finitely presented R-modules form an abelian category.

(2) An R-module is finitely presented if and only if it is coherent.

(3) R is coherent over itself.

We next discuss how these properties behave when we change the base ring R.

Definition 9.0.5. We say that a property P is:

(1) preserved by localization, if for every R-module M that satisfies P and every multiplicatively closed set S we have that $S^{-1}M$ satisfies P as an $S^{-1}R$-module.

(2) determined locally, if for every set of elements $\{r_i\}$ which generate the unit ideal in R we have that an R-module M satisfies P if and only if for every index i the localization M_{r_i} satisfies P as an R_{r_i}-module.

(3) preserved by base change, if for every R-module M that satisfies P and every ring homomorphism $R \to T$ we have that $M \otimes_R T$ satisfies P as a T-module.

finitely generated	finitely presented	coherent	is the property P:
✓	✓	✓	preserved by localization?
✓	✓	✓	determined locally?
✓	✓	✗	preserved by base change?

Table 9.2: Finiteness properties over varying ring

Note that the only property on our list that fails for coherent modules is preservation under base change. Finally, we mention one important property of finitely presented modules:

Theorem 9.0.6. *Let R be a ring, S a multiplicatively closed subset of R. Let M, N be R-modules such that M is finitely presented. Then*

$$\mathrm{Hom}_R(M, N) \otimes S^{-1}R \cong \mathrm{Hom}_{S^{-1}R}(S^{-1}M, S^{-1}N).$$

9.1 $\mathcal{O}_X$-modules

Definition 9.1.1. Let $(X, \mathcal{O}_X)$ be a ringed space. A presheaf of $\mathcal{O}_X$-modules is a presheaf of abelian groups $\mathcal{F}$ on X such that:

(1) for every open set U, $\mathcal{F}(U)$ carries the structure of an $\mathcal{O}_X(U)$-module, and

(2) for every inclusion of open sets $V \subset U$, the restriction map $\rho_{\mathcal{F},U,V} : \mathcal{F}(U) \to \mathcal{F}(V)$ is compatible with the module structures via the restriction map $\rho_{\mathcal{O}_X,U,V} : \mathcal{O}_X(U) \to \mathcal{O}_X(V)$. In other words, we have a commuting diagram of module actions

$$\begin{tikzcd}
\mathcal{O}_X(U) \times \mathcal{F}(U) \arrow[r] \arrow[d, "\rho"] & \mathcal{F}(U) \arrow[d, "\rho"] \\
\mathcal{O}_X(V) \times \mathcal{F}(V) \arrow[r] & \mathcal{F}(V)
\end{tikzcd}$$

A sheaf of $\mathcal{O}_X$-modules (or more briefly, an $\mathcal{O}_X$-module) is a presheaf of $\mathcal{O}_X$-modules which is also a sheaf.

A morphism of $\mathcal{O}_X$-modules $\phi : \mathcal{F} \to \mathcal{G}$ is a morphism of sheaves such that for every open set U the map $\phi(U)$ is a homomorphism of $\mathcal{O}_X(U)$-modules.

Remark 9.1.2. We will of course mostly be interested in the case when X is a scheme, but it is helpful to consider this more general situation.

Our goal in this section is simply to verify some basic properties of the category $\mathcal{O}_X - \mathbf{Mod}(X)$ of $\mathcal{O}_X$-modules. We will see some systematic ways of constructing $\mathcal{O}_X$-modules in the next sections.

Theorem 9.1.3. *The category of $\mathcal{O}_X$-modules is an abelian category when equipped with the constructions $\oplus, \ker, \operatorname{cok}$ of sheaves of abelian groups.*

Proof. Suppose that $\phi : \mathcal{F} \to \mathcal{G}$ is a morphism of $\mathcal{O}_X$-modules. It is clear that the kernel, presheaf image, and presheaf cokernel of ϕ are each presheaves of $\mathcal{O}_X$-modules. Note that if we have an open set U, an open cover $\{V_i\}$ of U, and sections $s_i \in \mathcal{F}(V_i)$ that agree on overlaps, then we can define the action of any $f \in \mathcal{O}_X(U)$ on the various s_i via the restriction function. Thus the sheafification process preserves the existence of an action by the various section rings $\mathcal{O}_X(U)$ and its clear by construction that the resulting object is an $\mathcal{O}_X$-module. The rest of the argument is straightforward. $\square$

Given two $\mathcal{O}_X$-modules $\mathcal{F}, \mathcal{G}$, we denote by $\mathrm{Hom}_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G})$ the abelian group consisting of $\mathcal{O}_X$-module morphisms from $\mathcal{F}$ to $\mathcal{G}$. In fact, $\mathrm{Hom}_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G})$ carries the structure of an $\mathcal{O}_X(X)$-module: given a morphism ϕ and an open set U, the action of $\mathcal{O}_X(X)$ on $\phi(U)$ is defined via the restriction map. There is a map of $\mathcal{O}_X(X)$-modules $\mathrm{Hom}_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G}) \to \mathrm{Hom}_{\mathcal{O}_X(X)}(\mathcal{F}(X), \mathcal{G}(X))$ but in general this map is not an isomorphism.

9.1.1 Sheaf hom

Definition 9.1.4. Suppose that $\mathcal{F}$ and $\mathcal{G}$ are $\mathcal{O}_X$-modules. For any open set U, the restrictions $\mathcal{F}|_U$ and $\mathcal{G}|_U$ are $\mathcal{O}_X|_U$-modules. As discussed above, $\mathrm{Hom}_{\mathcal{O}_X|_U}(\mathcal{F}|_U, \mathcal{G}|_U)$ will have the structure of a $\mathcal{O}_X(U)$-module.

Given an inclusion of open subsets $V \subset U$, we can restrict a $\mathcal{O}_X|_U$-homomorphism $\phi : \mathcal{F}|_U \to \mathcal{G}|_U$ to V. These maps are clearly compatible with the restriction maps for the sheaf $\mathcal{O}_X$.

This construction defines a presheaf of $\mathcal{O}_X$-modules denoted by $\mathcal{H}om_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G})$.

Lemma 9.1.5. *The presheaf $\mathcal{H}om_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G})$ is a sheaf.*

Proof. This follows from the gluing property for morphisms of sheaves in Corollary 7.6.4. □

Warning 9.1.6. Suppose U is an open subset of X. It is tempting, but incorrect, to identify $\mathcal{H}om_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G})(U)$ with $\mathrm{Hom}_{\mathcal{O}_X(U)}(\mathcal{F}(U), \mathcal{G}(U))$. This latter object is not even a presheaf: if $V \subset U$ how does a map $\mathcal{F}(U) \to \mathcal{G}(U)$ yield a map $\mathcal{F}(V) \to \mathcal{G}(V)$?

Warning 9.1.7. Although $\mathcal{H}om_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G})_x$ admits a morphism to $\mathrm{Hom}_{\mathcal{O}_{X,x}}(\mathcal{F}_x, \mathcal{G}_x)$ these two modules are *not* isomorphic in general. This is a consequence of the failure of module Hom to commute with localizations; see Exercise 9.1.21.

9.1.2 Sheaf tensor product

Definition 9.1.8. Suppose that $\mathcal{F}$ and $\mathcal{G}$ are $\mathcal{O}_X$-modules. The tensor product presheaf is the assignment

$$U \mapsto \mathcal{F}(U) \otimes_{\mathcal{O}_X(U)} \mathcal{G}(U)$$

equipped with the restriction maps induced by the universal property of tensor products.

The tensor product $\mathcal{F} \otimes_{\mathcal{O}_X} \mathcal{G}$ is the sheafification of the tensor product presheaf.

Example 9.1.9. In general the tensor product presheaf is not a sheaf. For example, Exercise 9.3.14 shows that $\mathcal{O}_{\mathbb{P}^1}(1) \otimes \mathcal{O}_{\mathbb{P}^1}(1) \cong \mathcal{O}_{\mathbb{P}^1}(2)$. On the other hand, the tensor product of $\mathcal{O}_{\mathbb{P}^1}(1)(\mathbb{P}^1) \cong \mathbb{K}^2$ with itself is different from $\mathcal{O}_{\mathbb{P}^1}(2)(\mathbb{P}^1) \cong \mathbb{K}^3$.

Exercise 9.1.10. Suppose that $\mathcal{F}, \mathcal{G}$ are $\mathcal{O}_X$-modules. Prove that for any point $x \in X$ we have $(\mathcal{F} \otimes_{\mathcal{O}_X} \mathcal{G})_x \cong \mathcal{F}_x \otimes_{\mathcal{O}_{X,x}} \mathcal{G}_x$.

Remark 9.1.11. In a similar way we can define other tensor product constructions, such as the kth symmetric power or the kth exterior power of an $\mathcal{O}_X$-module. We will not need these constructions but please remember that they are defined in the natural way.

The Hom and $\otimes$ constructions satisfy an adjointness property, just as they do for modules.

Theorem 9.1.12. *Let $(X, \mathcal{O}_X)$ be a ringed space and let $\mathcal{F}$, $\mathcal{G}$, $\mathcal{H}$ be $\mathcal{O}_X$-modules. There is a canonical isomorphism*

$$\mathcal{H}om_{\mathcal{O}_X}(\mathcal{F} \otimes_{\mathcal{O}_X} \mathcal{G}, \mathcal{H}) \to \mathcal{H}om_{\mathcal{O}_X}(\mathcal{F}, \mathcal{H}om_{\mathcal{O}_X}(\mathcal{G}, \mathcal{H}))$$

which is functorial in all entries.

Exercise 9.1.13. Let $(X, \mathcal{O}_X)$ be a ringed space and let $\mathcal{F}$ be an $\mathcal{O}_X$-module. Prove that $\mathcal{H}om_{\mathcal{O}_X}(\mathcal{F}, -)$ is a left exact functor. Prove that $\mathcal{F} \otimes_{\mathcal{O}_X} -$ is a right exact functor.

9.1.3 Functors associated to morphisms

Let $f : X \to Y$ be a morphism of ringed spaces (that is, a continuous map and a morphism of sheaves of rings $f^\sharp : \mathcal{O}_Y \to f_*\mathcal{O}_X$). The most common way to pass sheaves of modules between X and Y is the following.

- If $\mathcal{F}$ is an $\mathcal{O}_X$-module, then $f_*\mathcal{F}$ is a $f_*\mathcal{O}_X$-module. Thus $f_*\mathcal{F}$ obtains the structure of an $\mathcal{O}_Y$-module via $f^\sharp : \mathcal{O}_Y \to f_*\mathcal{O}_X$. We continue to call $f_*\mathcal{F}$ the pushforward of $\mathcal{F}$. Note that f_* defines a functor from the category of $\mathcal{O}_X$-modules to the category of $\mathcal{O}_Y$-modules.
- If $\mathcal{G}$ is an $\mathcal{O}_Y$-module then $f^{-1}\mathcal{G}$ is an $f^{-1}\mathcal{O}_Y$-module. To obtain a $\mathcal{O}_X$-module we need to make a further modification.

Definition 9.1.14. Let $f : X \to Y$ be a morphism of ringed spaces. For any $\mathcal{O}_Y$-module $\mathcal{F}$ we define the pullback $f^*\mathcal{F}$ to be

$$f^*\mathcal{F} = f^{-1}\mathcal{F} \otimes_{f^{-1}\mathcal{O}_Y} \mathcal{O}_X.$$

With this definition f^* is a functor from the category of $\mathcal{O}_Y$-modules to the category of $\mathcal{O}_X$-modules. By combining Lemma 7.5.8 and Exercise 9.1.10 we see that for any point $x \in X$ we have $f^*\mathcal{F}_x = \mathcal{F}_{f(x)} \otimes_{\mathcal{O}_{Y,f(x)}} \mathcal{O}_{X,x}$.

Exercise 9.1.15. Suppose that $f : X \to Y$ is a morphism of ringed spaces. Prove that f_* and f^* are adjoint functors: for any $\mathcal{O}_X$-module $\mathcal{F}$ and $\mathcal{O}_Y$-module $\mathcal{G}$ we have an isomorphism

$$\mathrm{Hom}_{\mathcal{O}_X}(f^*\mathcal{G}, \mathcal{F}) \cong \mathrm{Hom}_{\mathcal{O}_Y}(\mathcal{G}, f_*\mathcal{F})$$

that is natural in both entries.

As a consequence of Exercise 9.1.15, there are canonical maps $\mathcal{G} \to f_*f^*\mathcal{G}$ and $f^*f_*\mathcal{F} \to \mathcal{F}$. This exercise also implies:

Corollary 9.1.16. *The pushforward functor f_* is left-exact and the pullback functor f^* is right-exact.*

Remark 9.1.17. Suppose that $f : X \to Y$ is a morphism of schemes. Then we always have $f^*\mathcal{O}_Y \cong \mathcal{O}_X$. On the other hand, the structural map $f^\sharp : \mathcal{O}_Y \to f_*\mathcal{O}_X$ need not be an isomorphism, and in fact the condition that this map be an isomorphism is a very interesting one to study.

Example 9.1.18. Let $f : X \to Y$ be a morphism of schemes. Suppose that $\mathcal{G}$ is an $\mathcal{O}_Y$-module. As in Exercise 9.1.20 we will identify global sections $s \in \mathcal{G}$ with sheaf morphisms $s : \mathcal{O}_Y \to \mathcal{G}$. Since f^* is a functor, any section $s : \mathcal{O}_Y \to \mathcal{G}$ yields a section $f^*s : \mathcal{O}_X = f^*\mathcal{O}_Y \to f^*\mathcal{G}$. This map $\Gamma(Y, \mathcal{G}) \to \Gamma(X, f^*\mathcal{G})$ is known as "pulling back sections."

9.1.4 Summary

In summary, after including the global sections functor $\Gamma(X, -)$ introduced in Definition 7.4.7 we have identified the following functors:

- Functors from $\mathcal{O}_X-\mathbf{Mod}$ to $\mathcal{O}_X(X)-\mathbf{Mod}$ or $\mathcal{O}_X(X)-\mathbf{Mod}^{op}$:

$$\Gamma(X, -) \qquad \mathrm{Hom}_{\mathcal{O}_X}(\mathcal{F}, -) \qquad \mathrm{Hom}_{\mathcal{O}_X}(-, \mathcal{F})$$

- Functors from $\mathcal{O}_X-\mathbf{Mod}$ to $\mathcal{O}_X-\mathbf{Mod}$ or $\mathcal{O}_X-\mathbf{Mod}^{op}$:

$$\mathcal{H}om_{\mathcal{O}_X}(\mathcal{F}, -) \qquad \mathcal{H}om_{\mathcal{O}_X}(-, \mathcal{F}) \qquad - \otimes_{\mathcal{O}_X} \mathcal{F}$$

- Given a morphism $f : X \to Y$,

$$f_* : \mathcal{O}_X-\mathbf{Mod} \to \mathcal{O}_Y-\mathbf{Mod} \qquad f^* : \mathcal{O}_Y-\mathbf{Mod} \to \mathcal{O}_X-\mathbf{Mod}$$

Much of the rest of the course will be dedicated to understanding these various functors. Since most of these functors are left or right exact, they are suitable candidates for the theory of derived functors.

Remark 9.1.19. In this section we kept subscripts of $\mathcal{O}_X$ everywhere to emphasize our underlying category. In upcoming sections we will bravely drop the $\mathcal{O}_X$-subscripts.

9.1.5 Exercises

Exercise 9.1.20. Let X be a scheme and let $\mathcal{F}$ be a $\mathcal{O}_X$-module. Prove that $\mathcal{H}om_{\mathcal{O}_X}(\mathcal{O}_X, \mathcal{F})$ is isomorphic to $\mathcal{F}$. Deduce that there is a bijection between morphisms $s : \mathcal{O}_X \to \mathcal{F}$ and sections $s \in \mathcal{F}(X)$.

Exercise 9.1.21. Let X be a scheme with a closed point x. Let $i : U \to X$ be the inclusion of the complement of x and let $\mathcal{F} = i_!\mathcal{O}_U$. Show that $\mathcal{F}$ is an $\mathcal{O}_X$-module. Show that $\mathcal{H}om_{\mathcal{O}_X}(\mathcal{F}, \mathcal{F})_x \neq \mathrm{Hom}_{\mathcal{O}_{X,x}}(\mathcal{F}_x, \mathcal{F}_x)$.

Exercise 9.1.22. Let X be a scheme. Let $x \in X$ be a point, let M be a $\kappa(x)$-module, and let $\mathcal{G}$ denote the skyscraper sheaf at x with value M. Show that $\mathcal{G}$ is an $\mathcal{O}_X$-module and that

$$\mathcal{H}om_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G}) \cong \mathrm{Hom}_{\mathcal{O}_{X,x}}(\mathcal{F}_x, M)$$

Exercise 9.1.23. Let $f : X \to Y$ be a morphism of schemes. Prove that f^* commutes with tensor product. (Hint: this relates to the module computation $(M \otimes_S R) \otimes_R (N \otimes_S R) \cong (M \otimes_S N) \otimes_S R$.)

Show that f_* need not commute with tensor product. (One option is to use Exercise 9.3.14.) However, construct for every pair of $\mathcal{O}_X$-modules $\mathcal{F}, \mathcal{G}$ a natural map $f_*\mathcal{F} \otimes f_*\mathcal{G} \to f_*(\mathcal{F} \otimes \mathcal{G})$.

Exercise 9.1.24. Let X be a $\mathbb{K}$-scheme and let $\mathcal{F}$ be an $\mathcal{O}_X$-module. Let $\mathbb{L}/\mathbb{K}$ be an extension and let $f : X_{\mathbb{L}} \to X$ be the morphism induced by base change. Prove that $f^*\mathcal{F}(X_{\mathbb{L}}) \cong \mathcal{F} \otimes_{\mathbb{K}} \mathbb{L}$. (To make this argument work, the key is that X should be quasicompact quasiseparated.)

Exercise 9.1.25. Find an example of a morphism of schemes $f : X \to Y$ such that f^* is not a left exact functor. (Note the contrast with the inverse image functor f^{-1}.)

9.2 $\mathcal{O}_X$-modules on affine schemes

Suppose that $X = \mathrm{Spec}(R)$ is an affine scheme. The goal of this section is to build up a dictionary that relates R-modules with $\mathcal{O}_X$-modules. Just as we constructed the structure sheaf using localizations of R, we will construct $\mathcal{O}_X$-modules using localizations of R-modules.

9.2.1 Modules and localization

When we constructed the structure sheaf of an affine $\mathbb{K}$-scheme, the key step was Proposition 1.11.4 which verified that localizations were compatible with gluing. Our first result will be the analogous statement for modules.

Proposition 9.2.1. *Let R be a ring and let M be an R-module. Fix a finite set of elements $\{g_i\}_{i=1}^r$ which generate R. Then there is an exact sequence*

$$0 \to M \xrightarrow{\phi} \prod_i M_{g_i} \xrightarrow{\psi} \prod_{i,j} M_{g_i g_j}$$

where ϕ is the product of the localization maps $M \to M_{g_i}$ and ψ sends a tuple $(m_i \in M_{g_i})$ to $(\frac{m_i}{1} - \frac{m_j}{1} \in M_{g_i g_j})$.

The proof is essentially the same as the proof of Proposition 1.11.4 (and for good reason – the two statements are saying essentially the same thing).

Proof. We let $\rho_i : M \to M_{g_i}$ and $\rho_{i,j} : M_{g_i} \to M_{g_i g_j}$ be the localization maps. It is clear that the image of the leftmost map $\prod_i \rho_i$ is contained in the set of compatible elements

$$(m_i \in M_{g_i} | \rho_{i,j}(m_i) = \rho_{j,i}(m_j) \forall i, j)$$

and we must show this map is an isomorphism.

First we show injectivity. Suppose that $m \in M$ is mapped to 0. In other words, for every index i there is some positive integer k_i such that $mg^{k_i} = 0$. Set $N = \sup_i k_i$. Since $R = (g_1, \ldots, g_r)$, we also have $R = (g_1^N, \ldots, g_r^N)$. We deduce that $m = 0$.

Next we show surjectivity. We write $m_i = n_i/g_i^{k_i}$. By assumption we have $n_i/g_i^{k_i} = n_j/g_j^{k_j}$ as elements in $M_{g_i g_j}$. Thus for any pair of indices $i \neq j$ there is some non-negative integer t_{ij} such that

$$n_i g_j^{k_j+t_{ij}} g_i^{t_{ij}} = n_j g_i^{k_i+t_{ij}} g_j^{t_{ij}}.$$

We define $N = \sup_i k_i + \sup_{j \neq i} t_{ij}$. Then we can rewrite

$$m_i = \frac{n_i}{g_i^{k_i}} = \frac{n_i g_i^{N-k_i}}{g_i^N}.$$

For notational convenience we will write $p_i = n_i g_i^{N-k_i}$. The advantage of this change is that for $i \neq j$ we have the simpler relation

$$p_i g_j^N = n_i g_j^{k_j}(g_i^{N-k_i} g_j^{N-k_j}) = n_j g_i^{k_i}(g_i^{N-k_i} g_j^{N-k_j}) = p_j g_i^N.$$

Since $R = (g_1, \ldots, g_r)$, we also have $R = (g_1^N, \ldots, g_r^N)$. In particular we have an equality $1 = \sum_{i=1}^r c_i g_i^N$ for appropriate choices of c_i. Define $m = \sum_{i=1}^r c_i p_i$. We claim that $m \in R$ maps to $m_j \in R_{g_j}$ under the localization map. Indeed, we have

$$m g_j^N = \sum_{i=1}^r c_i p_i g_j^N = \sum_{i=1}^r c_i p_j g_i^N = p_j.$$

$\square$

Exercise 9.2.2. Prove the analogue of Proposition 9.2.1 with no assumption on the finiteness of the g_i.

9.2.2 The $\widetilde{\ }$-functor

Construction 9.2.3. Let $X = \mathrm{Spec}(R)$ be an affine scheme. Suppose that M is an R-module. We define a $\mathcal{O}_X$-module $\widetilde{M}$ as follows.

Consider the base $\mathcal{B}$ of the topology on X consisting of distinguished open affines D_f. We assign to D_f the module obtained by localizing M along all the functions g such that $V_g \cap D_f = \emptyset$. The restriction maps are defined using the universal property of localization.

Exercise 9.2.4. Use Proposition 9.2.1 to show that the definition above defines a $\mathcal{B}$-sheaf. (Hint: mimic the proof of Corollary 1.11.6.)

We define $\widetilde{M}$ to be the sheaf on $\mathrm{Spec}(R)$ associated to this $\mathcal{B}$-sheaf. It is clear that this construction defines a functor from $R-\mathbf{Mod}$ to $\mathcal{O}_X-\mathbf{Mod}$.

The characterizing properties of the $\mathcal{O}_X$-module $\widetilde{M}$ are:

- We have $\widetilde{M}(D_f) \cong M_f$ (and in particular $\widetilde{M}(\mathrm{Spec}(R)) = M$).
- For any point $x \in \mathrm{Spec}(R)$ corresponding to the prime ideal $\mathfrak{p}$ we have $\widetilde{M}_x \cong M_{\mathfrak{p}}$.
- For any open set $U \subset X$ and any open cover $\{D_{f_i}\}$ of U by distinguished open affines, we have an exact sequence

$$0 \to \widetilde{M}(U) \to \prod_i M_{f_i} \to \prod_{i,j} M_{f_i f_j}.$$

Let's study the properties of the $\widetilde{\ }$-functor in more depth.

Proposition 9.2.5. *Let $X = \operatorname{Spec}(R)$ be an affine scheme. The* $\widetilde{\ }$*-functor $R-\mathbf{Mod} \to \mathcal{O}_X-\mathbf{Mod}$ is a left adjoint to the global sections functor $\mathcal{O}_X-\mathbf{Mod} \to R-\mathbf{Mod}$. In other words, for any $\mathcal{O}_X$-module $\mathcal{F}$ and any R-module M we have an isomorphism*

$$\operatorname{Hom}_{\mathcal{O}_X}(\widetilde{M}, \mathcal{F}) \cong \operatorname{Hom}_R(M, \mathcal{F}(X))$$

that is natural in both entries.

Proof. Given a morphism $\psi : M \to \mathcal{F}(X)$ and an element $g \in R$, we can compose the localization of ψ with restriction to obtain

$$\psi_g : M_g \to \mathcal{F}(X) \otimes_R R_g \to \mathcal{F}(D_g).$$

If $\mathcal{B}$ denotes the base of the topology of X consisting of distinguished open affines then we obtain a morphism from the $\mathcal{B}$-sheaf associated to $\widetilde{M}$ to the $\mathcal{B}$-sheaf associated to $\mathcal{F}$. By Theorem 7.6.2 we obtain a morphism from $\widetilde{M}$ to $\mathcal{F}$, and it is clear that the induced map on global sections agrees with the map we started with. The rest of the proof is straightforward. □

Proposition 9.2.6. *Let $X = \operatorname{Spec}(R)$ be an affine scheme. The* $\widetilde{\ }$*-functor defines an exact fully faithful functor from $R-\mathbf{Mod}$ to $\mathcal{O}_X-\mathbf{Mod}$.*

Proof. We first show the exactness of the functor. If we have an exact sequence of R-modules, then the sequence obtained by localizing at a prime $\mathfrak{p} \subset R$ is still exact. Since exactness of a sequence of $\mathcal{O}_X$-modules can be checked on the level of stalks, we see that the corresponding sequence of $\mathcal{O}_X$-modules is still exact.

To say that the functor is fully faithful means that the natural map

$$\operatorname{Hom}_R(M, N) \to \operatorname{Hom}_{\mathcal{O}_X}(\widetilde{M}, \widetilde{N})$$

is a bijection. This follows from Proposition 9.2.5. □

It is important to note that the $\widetilde{\ }$-functor is not essentially surjective – there are many $\mathcal{O}_X$-modules that are not isomorphic to $\widetilde{M}$ for any M. (However, Proposition 9.2.6 does imply that the category of R-modules is equivalent to the subcategory of $\mathcal{O}_X-\mathbf{Mod}$ defined by the image of the $\widetilde{\ }$-functor.)

Example 9.2.7. Let R be a ring and suppose that $i : U \to \operatorname{Spec}(R)$ is the inclusion of an open subscheme. Then $i_!(\mathcal{O}_U)$ is a $\mathcal{O}_X$-module that is usually not isomorphic to a sheaf in the image of the $\widetilde{\ }$-functor. For example, if $\operatorname{Spec}(R)$ is irreducible then $i_!(\mathcal{O}_U)(X) = 0$ but it need not be the zero sheaf. Sheaves obtained by $i_!$ (and in particular the exact sequence of Exercise 7.5.15) are an important source of counterexamples.

The following important result gives conditions guaranteeing that the global sections functor is exact. Note that it goes in the "opposite direction" of Proposition 9.2.6.

Proposition 9.2.8. *Let X be an affine scheme. Suppose that we have an exact sequence of $\mathcal{O}_X$-modules*

$$0 \to \mathcal{F}_1 \xrightarrow{\phi} \mathcal{F}_2 \xrightarrow{\psi} \mathcal{F}_3 \to 0.$$

Assume that $\mathcal{F}_1$ is in the image of the $\widetilde{\ }$-functor. Then the sequence of global sections

$$0 \to \mathcal{F}_1(X) \xrightarrow{\phi(X)} \mathcal{F}_2(X) \xrightarrow{\psi(X)} \mathcal{F}_3(X) \to 0$$

is also exact.

The assumption on $\mathcal{F}_1$ guarantees that for any open set U and any distinguished open affine D_f we have $\mathcal{F}_1(U \cap D_f) = \mathcal{F}_1(U)_f$. In other words, for any section $s \in \mathcal{F}_1(U \cap D_f)$ we can find some positive integer n such that $f^n s$ extends to a section defined on all of U.

Here is a sketch of the argument. Fix a section $t \in \mathcal{F}_3(X)$. Since ψ is surjective, we can find an open cover $\{U_i\}$ of X and "local liftings" $s_i \in \mathcal{F}_2(U_i)$ of t. We would like to glue the s_i to get a global section s, but of course there is no reason to expect the s_i to agree on overlaps. However, we know that $s_i|_{U_i \cap U_j} - s_j|_{U_i \cap U_j}$ is in $\mathcal{F}_1(U_i \cap U_j)$. Using the assumption on $\mathcal{F}_1$, we can "extend" these differences to larger open sets (at the cost of multiplying by a power of a function). Using these extensions we can hope to cancel these obstructions (on a global level) to obtain functions which glue.

Proof. Since the global sections functor is left-exact, it suffices to prove exactness of the sequence of global sections on the right. Suppose $t \in \mathcal{F}_3(X)$. Let $f \in \mathcal{O}_X(X)$ be any function such that $t|_{D_f}$ is the image of an element $s \in \mathcal{F}_2(D_f)$ under $\psi(D_f)$. The first step of the proof is to a weaker lifting property: there is some positive integer N such that $f^N t \in \mathcal{F}_3(X)$ is in the image of $\psi(X)$.

Since ψ is surjective, we can find a finite open cover $\{U_i\}_{i=1}^r$ of X by distinguished open affines $U_i = D_{g_i}$ and elements $s_i \in \mathcal{F}_2(U_i)$ such that $\psi(U_i)(s_i) = t|_{U_i}$. Define $V_i = U_i \cap D_f$. Since the two sections $s|_{V_i}$ and $s_i|_{V_i}$ both map to $t|_{V_i}$ under $\psi(V_i)$, their difference lies in $\mathcal{F}_1(V_i)$. By assumption on $\mathcal{F}_1$, there is some positive integer M_i such that $f^{M_i}(s|_{V_i} - s_i|_{V_i})$ is the restriction of an element in $\mathcal{F}_1(U_i)$. As there are only finitely many i, we can set $M = \sup_i M_i$ and use this one constant M for every open set U_i. We let $u_i \in \mathcal{F}_1(U_i)$ denote the section which restricts to $f^M(s|_{V_i} - s_i|_{V_i})$. We then define $s_i' = f^M s_i + u_i$ in $\mathcal{F}_2(U_i)$.

We have now achieved that these sections $s_i' \in \mathcal{F}_2(U_i)$ are local lifts of $f^M t$ with the further property that their restrictions to $U_i \cap D_f$ agree with $f^M s|_{U_i \cap D_f}$. But we do not know that the s_i' agree on the overlaps $U_i \cap U_j$. We next increase the power of f to achieve this overlap condition.

Set $U_{ij} = U_i \cap U_j$. Since $s_i'|_{U_{ij}}$ and $s_j'|_{U_{ij}}$ both map to $f^M t|_{U_{ij}}$, their difference lies in $\mathcal{F}_1(U_{ij})$. Furthermore, by construction the restriction of their difference to $U_{ij} \cap D_f$ is zero. By assumption on $\mathcal{F}_1$, this means that there is some constant L_{ij} such that $f^{L_{ij}}(s_i'|_{U_{ij}} - s_j'|_{U_{ij}})$ is the zero element in $\mathcal{F}_1(U_{ij})$. If we set $L = \sup_{i,j} L_{ij}$ and set $s_i^+ = f^L s_i'$, we see that each $s_i^+ \in \mathcal{F}_2(U_i)$ maps to $f^{M+L}t|_{U_i}$ under $\psi(U_i)$ and also $s_i^+|_{U_{ij}} = s_j^+|_{U_{ij}}$. Gluing, we find a global section of $\mathcal{F}_2(X)$ whose image is $f^{M+L}t \in \mathcal{F}_3(X)$.

Since X is quasicompact, there is a finite set of functions $\{f_k\}$ such that the D_{f_k} form an open cover of X and for every k we have that $t|_{D_{f_k}}$ is in the image of $\psi(D_{f_k})$. Applying the argument above for each f_k, we find that there are positive constants T_k such that $f_k^{T_k}t$ is in the image of $\psi(X)$. However, since the D_{f_k} form an open cover the various $f_k^{T_k}$ generate the unit ideal in $\mathcal{O}_X(X)$. We conclude that t is in the image of $\psi(X)$ as claimed. $\square$

9.2.3 Properties of the $\widetilde{\ }$-functor

When a module operation commutes with localization, we can expect the $\widetilde{\ }$-functor to be compatible with the module operation.

Exercise 9.2.9. Prove that the $\widetilde{\ }$-functor commutes with arbitrary direct sums.

Proposition 9.2.10. *Let $X = \operatorname{Spec}(R)$ be an affine scheme and let M, N be R-modules. Then:*

(1) $\widetilde{M} \otimes_{\mathcal{O}_X} \widetilde{N} \cong \widetilde{M \otimes_R N}$.

(2) If M is a finitely presented R-module, then $\mathcal{H}om(\widetilde{M}, \widetilde{N}) \cong \widetilde{\operatorname{Hom}(M, N)}$.

In part (2) the finite presentation of M is essential – only in this situation are we guaranteed that Hom commutes with localization.

Proof. (1) Recall that $\widetilde{M} \otimes_{\mathcal{O}_X} \widetilde{N}$ is constructed by sheafifying the presheaf assigning to the open set U the abelian group $\widetilde{M}(U) \otimes_{\mathcal{O}_X(U)} \widetilde{N}(U)$. Since the global sections of this presheaf are $M \otimes_R N$, we obtain an R-module homomorphism $M \otimes_R N \to (\widetilde{M} \otimes_{\mathcal{O}_X} \widetilde{N})(X)$. By Proposition 9.2.5 we obtain a morphism $\widetilde{M \otimes_R N} \to \widetilde{M} \otimes_{\mathcal{O}_X} \widetilde{N}$.

To show that this map is an isomorphism, it suffices to check that it induces an isomorphism on the level of stalks. The stalk of the tensor product presheaf at a prime ideal $\mathfrak{p}$ is isomorphic to $M_{\mathfrak{p}} \otimes_{R_{\mathfrak{p}}} N_{\mathfrak{p}}$. By Exercise 7.2.6 we see that ϕ induces the map on stalks

$$(M \otimes N)_{\mathfrak{p}} \to M_{\mathfrak{p}} \otimes_{R_{\mathfrak{p}}} N_{\mathfrak{p}}$$

which is an isomorphism.

(2) Proposition 9.2.5 yields a morphism $\widetilde{\operatorname{Hom}(M, N)} \to \mathcal{H}om(\widetilde{M}, \widetilde{N})$. To show that this map is an isomorphism, it suffices to check that it induces an isomorphism on the level

of stalks. For a point x corresponding to a prime $\mathfrak{p}$ we have

$$\begin{aligned}\mathcal{H}om(\widetilde{M}, \widetilde{N})_x &= \varinjlim_{f \notin \mathfrak{p}} \operatorname{Hom}(M_f, N_f) \\ &= \varinjlim_{f \notin \mathfrak{p}} \operatorname{Hom}(M, N)_f \\ &= \operatorname{Hom}(M, N)_{\mathfrak{p}}\end{aligned}$$

where we used the fact that M is finitely presented to commute the Hom and the localization at the second step. This proves the result. □

9.2.4 Exercises

Exercise 9.2.11. Let $X = \operatorname{Spec}(R)$ be an affine scheme and let $\mathcal{F}$ be an $\mathcal{O}_X$-module. Show that $\mathcal{F}$ is in the image of the $\widetilde{\ }$-functor if and only if for every distinguished open affine $D_g \subset X$ we have $\mathcal{F}(D_g) \cong \mathcal{F}(X)_g$.

Exercise 9.2.12. Find an affine scheme X and an exact sequence of $\mathcal{O}_X$-modules

$$0 \to \mathcal{F}_1 \xrightarrow{\phi} \mathcal{F}_2 \xrightarrow{\psi} \mathcal{F}_3 \to 0$$

such that the map $\psi(X)$ on global sections is not surjective.

9.3 Quasicoherent sheaves

We now "globalize" our discussion of the $\widetilde{\ }$-functor from the previous section.

Definition 9.3.1. Let X be a scheme. We say that a $\mathcal{O}_X$-module $\mathcal{F}$ is a quasicoherent sheaf if there is a cover of X by open affines $\{U_i\}$ and $\mathcal{O}_X(U_i)$-modules M_i such that $\mathcal{F}|_{U_i}$ is isomorphic to $\widetilde{M_i}$ for every i.

We let $\mathbf{Qcoh}(X)$ denote the full subcategory of $\mathcal{O}_X-\mathbf{Mod}$ whose objects are quasicoherent sheaves.

A key property of quasicoherence is that it does not depend upon the choice of affine cover. More precisely:

Proposition 9.3.2. *Let X be a scheme. If $\mathcal{F}$ is a quasicoherent sheaf, then for every open affine U in X there is a $\mathcal{O}_X(U)$-module M such that $\mathcal{F}|_U \cong \widetilde{M}$.*

Proof. Suppose that $\{V_i\}$ is an open cover of X consisting of open affines such that $\mathcal{F}|_{V_i}$ is equal to $\widetilde{M_i}$ for some $\mathcal{O}_X(V_i)$-module M_i. Applying Nike's Lemma (Lemma 8.3.4) and refining our open cover, we may suppose that some subset $\{V_j\}_{j\in J}$ forms an open cover of U and furthermore that each V_j is a distinguished open affine in U corresponding to some function $f_j \in \mathcal{O}_X(U)$.

If we apply the gluing property of sheaves to the open cover $\{V_j\}$ of U we obtain an exact sequence

$$0 \to \mathcal{F}(U) \to \prod_j \mathcal{F}(V_j) \to \prod_{j,k} \mathcal{F}(V_j \cap V_k).$$

In other words, we have an exact sequence of $\mathcal{O}_X(U)$-modules

$$0 \to \mathcal{F}(U) \to \prod_j M_j \to \prod_{j,k} (M_j)_{f_k}.$$

where the first map is restriction and the second map is a product of localization maps.

Suppose that $W \subset U$ is some distinguished open affine defined by a function $g \in \mathcal{O}_X(U)$. Since localization is exact, we obtain a sequence

$$0 \to \mathcal{F}(U)_g \to \prod_j (M_j)_g \to \prod_{j,k} (M_j)_{f_k g}.$$

Using the fact that $\mathcal{F}|_{V_j} \cong \widetilde{M_j}$, we see that $(M_j)_g \cong \mathcal{F}(W \cap V_j)$ and $(M_j)_{f_k g} \cong \mathcal{F}(W \cap V_j \cap V_k)$. Furthermore, under these identifications the second map in the sequence above is still induced by restriction. By comparing this sequence to the exact sequence expressing the gluing property for $\mathcal{F}|_W$ with respect to the open cover $\{W \cap V_j\}$, we conclude that $\mathcal{F}(U)_g \cong \mathcal{F}(W)$. Since this is true for any $g \in \mathcal{O}_X(U)$, Exercise 9.2.11 shows that $\mathcal{F}|_U$ is equal to $\widetilde{\mathcal{F}(U)}$. $\square$

Proposition 9.3.2 implies that a quasicoherent sheaf on an affine scheme $\mathrm{Spec}(R)$ must be isomorphic to a sheaf of the form $\widetilde{M}$ for some R-module M. In other words:

Corollary 9.3.3. *Let $X = \mathrm{Spec}(R)$ be an affine scheme. Then the $\widetilde{\ }$-functor and the global sections functor define an equivalence of categories between R−**Mod** and **Qcoh**(X).*

This also gives us a useful criterion for checking whether an $\mathcal{O}_X$-module is quasicoherent. By combining Proposition 9.3.2 and Exercise 9.2.11 we obtain:

Corollary 9.3.4. *Let X be a scheme and let $\mathcal{F}$ be an $\mathcal{O}_X$-module. Then $\mathcal{F}$ is quasicoherent if and only if for any open affine $U \subset X$ and any $g \in \mathcal{O}_X(U)$ the map*

$$\mathcal{F}(U)_g \to \mathcal{F}(D_g)$$

is an isomorphism. (Here the map is obtained by applying the universal property of localization to the restriction map $\mathcal{F}(U) \to \mathcal{F}(D_g)$.)

9.3.1 The category of quasicoherent sheaves

Our next goal is to show that the category of quasicoherent sheaves is an abelian category. The key is to study how quasicoherence behaves in exact sequences.

Proposition 9.3.5. *Let X be a scheme. Suppose that $0 \to \mathcal{F}_1 \xrightarrow{\phi} \mathcal{F}_2 \xrightarrow{\psi} \mathcal{F}_3 \to 0$ is an exact sequence of $\mathcal{O}_X$-modules. If two of the $\mathcal{F}_i$ are quasicoherent, then the third one is as well.*

In particular, if $\phi : \mathcal{F} \to \mathcal{G}$ is a morphism of quasicoherent sheaves then the kernel and cokernel of ϕ are also quasicoherent.

Proof. It suffices to consider the case when X is affine. We will write $M_i = \mathcal{F}_i(X)$ for $i = 1, 2, 3$.

First suppose that $\mathcal{F}_1, \mathcal{F}_2$ are quasicoherent. Since the $\widetilde{\ }$-functor is fully faithful, there is a module homomorphism $q : M_1 \to M_2$ such that when we take $\widetilde{\ }$-images we get ϕ. Since $\widetilde{\ }$ is exact, the cokernel of ϕ is equal to $\widetilde{K}$ where K is the cokernel of q. This forces $\mathcal{F}_3 \cong \widetilde{K}$ so that $\mathcal{F}_3$ is quasicoherent. The argument when $\mathcal{F}_2, \mathcal{F}_3$ are quasicoherent is exactly the same.

Finally suppose that $\mathcal{F}_1, \mathcal{F}_3$ are quasicoherent. By Proposition 9.2.8 we know that $0 \to M_1 \to M_2 \to M_3 \to 0$ is exact. By Proposition 9.2.6 $\widetilde{\ }$ will preserve the exactness of this sequence. Using the adjointness property of $\widetilde{\ }$ as in Proposition 9.2.5 we obtain a commuting diagram

$$\begin{tikzcd}
0 \arrow[r] & \widetilde{M_1} \arrow[r] \arrow[d] & \widetilde{M_2} \arrow[r] \arrow[d] & \widetilde{M_3} \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \mathcal{F}_1 \arrow[r] & \mathcal{F}_2 \arrow[r] & \mathcal{F}_3 \arrow[r] & 0
\end{tikzcd}$$

By hypothesis the left and right vertical arrows are isomorphisms, so the middle vertical arrow is also an isomorphism by the 5-lemma. ∎

Theorem 9.3.6. *Let X be a scheme. Then $\mathbf{Qcoh}(X)$ is an abelian category.*

Proof. It suffices to check that the structures $\oplus$, ker, cok which give $\mathcal{O}_X-\mathbf{Mod}$ the structure of an abelian category preserve quasicoherence. In particular, it suffices to verify this when X is an affine scheme. The fact that ker and cok preserve quasicoherence follows Proposition 9.3.5, and the compatibility of $\oplus$ with the $\widetilde{\ }$-functor was checked in Exercise 9.2.9. $\square$

9.3.2 Pushforward and pullback

We next study the behavior of quasicoherent sheaves under pushforward and pullback. The first step is to analyze the special case of morphisms between affine schemes.

Lemma 9.3.7. *Let $f : \mathrm{Spec}(R) \to \mathrm{Spec}(S)$ be a morphism of affine schemes.*

(1) For any S-module N we have $f^*\widetilde{N} = \widetilde{N \otimes_S R}$.

(2) For any R-module M we have $f_*\widetilde{M} = \widetilde{{}_SM}$ where ${}_SM$ means we think of M as an S-module via $f^\sharp$.

Proof. (1) By tracing through the inverse image presheaf and tensor product presheaf constructions, we see that there is an R-module homomorphism $N \otimes_R S \to f^*\widetilde{N}(\mathrm{Spec}(R))$. By Proposition 9.2.5 we obtain a $\mathcal{O}_X$-morphism $\widetilde{N \otimes_R S} \to f^*\widetilde{N}$. It suffices to show that this morphism induces isomorphisms of stalks, which follows from Lemma 7.5.8 and Exercise 9.1.10.

(2) Via the identification $f_*\widetilde{M}(\mathrm{Spec}(S)) \cong {}_SM$ we obtain a map $\widetilde{{}_SM} \to f_*\widetilde{M}$. It suffices to show this is an isomorphism on stalks. For any $g \in S$ we have $f^{-1}(D_g) = D_{f^\sharp g}$ and thus $(f_*\widetilde{M})_\mathfrak{p} \cong M \otimes_R S_\mathfrak{p}$. $\square$

Theorem 9.3.8. *Let $f : X \to Y$ be a morphism of schemes.*

(1) For any quasicoherent sheaf $\mathcal{G}$ on Y the pullback $f^*\mathcal{G}$ is quasicoherent.

(2) Suppose f is quasicompact quasiseparated. Then for any quasicoherent sheaf $\mathcal{F}$ on X the pushforward $f_*\mathcal{F}$ is quasicoherent.

Proof. (1) Follows immediately from Lemma 9.3.7.(1).

(2) It suffices to prove the theorem when $Y = \mathrm{Spec}(S)$ is affine. By hypothesis X is covered by a finite union of open affines $\{U_i\}$ and each intersection $U_i \cap U_j$ is also covered by a finite union of open affines $\{U_{ijk}\}$. Note that to check whether local sections on U_i and U_j agree when restricted to $U_i \cap U_j$, we can equally well check whether they agree when restricted to every U_{ijk}. Thus, for any open set $V \subset Y$ we have an exact sequence

$$0 \to \mathcal{F}(f^{-1}V) \to \oplus_i \mathcal{F}(f^{-1}V \cap U_i) \to \oplus_{i,j,k}\mathcal{F}(f^{-1}V \cap U_{ijk})$$

(Note that we can use a direct sum instead of a product in the second and third modules since there are only finitely many terms involved in each product.) Since the maps in this exact sequence commute with restriction, we obtain an exact sequence

$$0 \to f_*\mathcal{F} \to \oplus_i (f|_{U_i})_*\mathcal{F}|_{U_i} \to \oplus_{i,j,k}(f|_{U_{ijk}})_*\mathcal{F}|_{U_{ijk}}.$$

Each of the sheaves $(f|_{U_i})_*\mathcal{F}|_{U_i}$ and $(f|_{U_{ijk}})_*\mathcal{F}|_{U_{ijk}}$ is quasicoherent by Lemma 9.3.7.(2) and their direct sums are quasicoherent by Exercise 9.2.9. Thus the kernel $f_*\mathcal{F}$ is also quasicoherent by Proposition 9.3.5. □

Note that the adjointness of the functors f^* and f_* still holds if we restrict them to categories of quasicoherent sheaves. In particular f^* is left-exact and f_* is right exact.

Exercise 9.3.9. Find a morphism of schemes $f : X \to Y$ and an injective morphism $\phi : \mathcal{F} \to \mathcal{G}$ of quasicoherent sheaves on Y such that $f^*\phi$ is not injective.

9.3.3 Computations

One of the nice features of quasicoherent modules is that our main constructions can be understood on the level of open affines.

(1) A morphism $\phi : \mathcal{F} \to \mathcal{G}$ of quasicoherent sheaves is surjective if and only if for every open affine V the map $\phi(V) : \mathcal{F}(V) \to \mathcal{G}(V)$ is surjective. Similarly exactness can be checked on the level of open affines. (This follows from Proposition 9.2.8.)

(2) Suppose that $\mathcal{F}, \mathcal{G}$ are quasicoherent sheaves. We can compute $\mathcal{F} \otimes \mathcal{G}$ by choosing an open affine cover, applying module $\otimes$ to the module of sections on each open affine, and then gluing the resulting sheaves. (This follows from Proposition 9.2.10.)

   Under suitable finite presentation hypotheses, we can do a similar procedure for $\mathcal{H}om(\mathcal{F}, \mathcal{G})$. (See Exercise 9.4.21.)

(3) Given a morphism $f : X \to Y$, the functor f^* can also be computed on the level of open affines using module operations. That is, if we choose an open cover of Y by open affines V and choose an open cover of X open affines U contained in the preimage of open affines of V, the the restriction of $f^*\mathcal{G}$ to U is defined by $\mathcal{O}_X(U) \otimes_{\mathcal{O}_Y(V)} \mathcal{G}(V)$. (This follows from Lemma 9.3.7.)

   However, f_* cannot be computed locally: given an open affine $V \subset Y$, the value of $f_*\mathcal{F}(V)$ really depends on the entire preimage over V and can't be reduced to a local computation on X. This again highlights why the pushforward is so hard to work with. (Note however that when f is affine then $f_*\mathcal{F}$ can be computed using local algebra.)

These properties make computations with quasicoherent sheaves much easier.

Warning 9.3.10. Although a tensor product of quasicoherent sheaves is quasicoherent, a $\mathcal{H}om$ of quasicoherent sheaves need not be quasicoherent without a finite presentation hypothesis on the first entry; see Exercise 9.4.21.

9.3.4 Quasicoherent ideal sheaves

Finally, we finish the discussion of closed embeddings and quasicoherent ideal sheaves begun in Section 8.3.3. The first step is to finish the proof of Proposition 8.3.8:

Lemma 9.3.11. *Suppose that $f : X \to Y$ is a closed embedding (that is, f is a homeomorphism to a closed subset of Y and $f^\sharp$ is surjective). Then for every open $V \subset Y$ the preimage $f^{-1}V$ is affine and the induced map $f^\sharp(V) : \mathcal{O}_Y(V) \to \mathcal{O}_X(f^{-1}V)$ is surjective.*

Proof. Evidently f is a quasicompact morphism. We also showed in Exercise 8.7.16 that f is separated. Thus f is quasicompact quasiseparated. By Theorem 9.3.8 $f_*\mathcal{O}_X$ is a quasicoherent sheaf on Y. The surjectivity of $f^\sharp$ on open affines follows from Proposition 9.2.8. $\square$

As in Definition 8.3.11 a quasicoherent ideal sheaf on a scheme X is simply a quasicoherent subsheaf of $\mathcal{O}_X$. We can now prove the key theorem:

Theorem 9.3.12. *Let X be a scheme. There is a bijection between closed subschemes of X and quasicoherent ideal sheaves $\mathcal{I}$ on X.*

Proof. Exercise 8.3.12 explains how to construct a closed subscheme Z from a quasicoherent ideal sheaf $\mathcal{I}$. Conversely, suppose we have a closed subscheme $i : Z \subset X$. As shown in Lemma 9.3.11 $i_*\mathcal{O}_Z$ is a quasicoherent sheaf. By Proposition 9.3.5 the kernel of the map $\mathcal{O}_X \to i_*\mathcal{O}_Z$ is a quasicoherent ideal sheaf. These two constructions are inverses: they are connected by the property that on any affine $U \subset X$ the intersection $Z \cap U$ is the vanishing locus of $\mathcal{I}(U)$. $\square$

In summary, if we have a closed subscheme $i : Z \to X$ then we obtain an exact sequence of quasicoherent sheaves

$$0 \to \mathcal{I}_Z \to \mathcal{O}_X \to i_*\mathcal{O}_Z \to 0.$$

On each open affine U, this sequence simply expresses the structure sheaf of Z as the quotient of the ideal defining $U \cap Z$. We will frequently appeal to this exact sequence in the future.

Example 9.3.13. Let $i : Z \to \mathbb{P}^n$ be the inclusion of a degree d hypersuface defined by $f \in \mathcal{O}_{\mathbb{P}^n}(d)$. The ideal sheaf of Z is the same as the image of the map $\phi : \mathcal{O}_{\mathbb{P}^n}(-d) \to \mathcal{O}_{\mathbb{P}^n}$ induced by multiplication by f. Thus $\mathcal{I}_Z \cong \mathcal{O}_{\mathbb{P}^n}(-d)$ and we have an exact sequence

$$0 \to \mathcal{O}_{\mathbb{P}^n}(-d) \xrightarrow{\cdot f} \mathcal{O}_{\mathbb{P}^n} \to i_*\mathcal{O}_Z \to 0.$$

9.3.5 Exercises

Exercise 9.3.14. Let $\mathbb{K}$ be a field and consider $\mathbb{P}^n_{\mathbb{K}}$.

(1) Show that the sheaf $\mathcal{O}_{\mathbb{P}^n_{\mathbb{K}}}(d)$ defined in Example 7.1.8 is a quasicoherent $\mathcal{O}_X$-module.

(2) Show that $\mathcal{H}om_{\mathcal{O}_{\mathbb{P}^n}}(\mathcal{O}_{\mathbb{P}^n_{\mathbb{K}}}(a), \mathcal{O}_{\mathbb{P}^n_{\mathbb{K}}}(b)) \cong \mathcal{O}_{\mathbb{P}^n_{\mathbb{K}}}(b-a)$.

(3) Show that $\mathcal{O}_{\mathbb{P}^n_{\mathbb{K}}}(a) \otimes_{\mathcal{O}_{\mathbb{P}^n_{\mathbb{K}}}} \mathcal{O}_{\mathbb{P}^n_{\mathbb{K}}}(b) \cong \mathcal{O}_{\mathbb{P}^n_{\mathbb{K}}}(a+b)$.

Exercise 9.3.15. Let $f : \mathbb{P}^1 \to \mathbb{P}^d$ be a rational normal curve of degree d. Show that $f^*\mathcal{O}_{\mathbb{P}^d}(m)$ is isomorphic to $\mathcal{O}_{\mathbb{P}^1}(md)$.

Exercise 9.3.16. Let $f : \mathbb{P}^n \times \mathbb{P}^m \to \mathbb{P}^{nm+n+m}$ be the Segre embedding. Show that $f^*\mathcal{O}(1)$ is equal to $\pi_1^*\mathcal{O}(1) \otimes \pi_2^*\mathcal{O}(1)$ where π_1, π_2 are the projection maps.

Exercise 9.3.17. Consider $\mathbb{P}^n \times \mathbb{P}^m$ equipped with the two projection maps π_1, π_2. Compute $\pi_{2*}\pi_1^*\mathcal{O}_{\mathbb{P}^n}(d)$.

Exercise 9.3.18. Consider the squaring map $f : \mathbb{P}^1 \to \mathbb{P}^1$ defined by $[s : t] \mapsto [s^2 : t^2]$. Compute $f_*\mathcal{O}(d)$.

Exercise 9.3.19. Let X be the disjoint union of countably many copies of $\operatorname{Spec}(\mathbb{Z})$. Consider the map $f : X \to \operatorname{Spec}(\mathbb{Z})$ which is the identity on each component. Prove that $f_*\mathcal{O}_X$ is not a quasicoherent sheaf. (Hint: Apply Corollary 9.3.4.)

Exercise 9.3.20. Let X be a scheme and let $f \in \mathcal{O}_X(X)$. Define X_f to be the complement of the vanishing locus $V(f)$ defined in Exercise 8.2.18. Let $\mathcal{F}$ be a quasicoherent sheaf on X. The following statements generalize Exercise 8.5.17.

(1) Suppose that X is quasicompact. Suppose that $a \in \mathcal{F}(X)$ is an element whose restriction to X_f is 0. Prove that for some $n > 0$ we have $f^n a = 0$ in $\mathcal{F}(X)$.

(2) Suppose that X is quasicompact quasiseparated. Let $b \in \mathcal{F}(X_f)$. Show that for some $n > 0$ the element $f^n b$ is the restriction of an element of $\mathcal{F}(X)$.

(3) Suppose that X is quasicompact quasiseparated. Then $\mathcal{F}(X_f) = \mathcal{F}(X)_f$.

9.4 Coherent sheaves

In this section we discuss "finite" quasicoherent sheaves. For a non-Noetherian ring there are three different reasonable definitions for "finite" modules: finitely generated, finitely presented, and coherent (see Definition 9.0.1). For Noetherian rings these three definitions all coincide, but in general they have slightly different properties (see Table 9.0.2 and Table 9.0.2) and the right notion will depend on the context.

Warning 9.4.1. For a non-Noetherian ring R the set of coherent R-modules might be pathologically small; see Warning 9.0.2. Coherent modules are most useful for rings R which are coherent over themselves, in which case coherent is the same as finitely presented.

Definition 9.4.2. Let X be a scheme and let $\mathcal{F}$ be a quasicoherent sheaf on X. We say that $\mathcal{F}$ is coherent (resp. finitely generated, finitely presented) if there is an open cover of X by open affines U_i such that $\mathcal{F}(U_i)$ is a coherent (resp. finitely generated, finitely presented) $\mathcal{O}_X(U_i)$-module for every i.

A key property of coherent sheaves is that the coherent condition can be verified locally:

Lemma 9.4.3. *Let X be a scheme and let $\mathcal{F}$ be a coherent (resp. finitely generated, finitely presented) sheaf on X. Then for every open affine U we have that $\mathcal{F}(U)$ is a coherent (resp. finitely generated, finitely presented) $\mathcal{O}_X(U)$-module.*

Proof. This follows from the fact that all three conditions are determined locally (as in Definition 9.0.5). □

Warning 9.4.4. It is not true that if $\mathcal{F}$ is coherent then $\mathcal{F}(U)$ is a finitely generated $\mathcal{O}_X(U)$-module for every open set U; see Exercise 9.4.22. (This is related to the fact that the global sections of the structure sheaf on a quasiprojective $\mathbb{K}$-variety need not be finitely generated over $\mathbb{K}$.)

9.4.1 Category of coherent sheaves

Let $\mathbf{Coh}(X)$ denote the category of coherent sheaves on X. The following results verify that $\mathbf{Coh}(X)$ is a well-behaved subcategory of $\mathbf{QCoh}(X)$.

Theorem 9.4.5. *Let X be a scheme. Suppose we have an exact sequence $0 \to \mathcal{F}_1 \to \mathcal{F}_2 \to \mathcal{F}_3$ of $\mathcal{O}_X$-modules. If two of the $\mathcal{F}_i$ are coherent, then so is the third.*

Proof. This follows from the corresponding fact for coherent R-modules combined with Proposition 9.3.5. □

Corollary 9.4.6. *Let X be a scheme. Then $\mathbf{Coh}(X)$ is an abelian category.*

Furthermore if we have two coherent modules on a Noetherian scheme then their $\mathcal{H}om$ or tensor product is again a coherent sheaf.

Remark 9.4.7. Note that the analogue of Theorem 9.4.5 fails for finitely generated or finitely presented sheaves since these properties are not preserved by taking kernels. In fact sheaves of these types need not form an abelian category.

9.4.2 Pullbacks and pushforwards

Although coherent modules are not preserved under base change, finitely generated modules are. Thus:

Lemma 9.4.8. *Let $f : X \to Y$ be a morphism of schemes. Suppose $\mathcal{G}$ is a finitely generated quasicoherent sheaf on Y. Then $f^*\mathcal{G}$ is finitely generated. In particular, if X is locally Noetherian then the pullback of a coherent sheaf on Y is coherent on X.*

It is much rarer for the pushforward of a coherent sheaf $\mathcal{F}$ to be coherent. Even for quasiprojective $\mathbb{K}$-schemes coherence is often not preserved by pushforward – for example, consider the pushforward of the structure sheaf under the map $\mathbb{A}^1_{\mathbb{K}} \to \mathrm{Spec}(\mathbb{K})$.

Suppose that $f : X \to Y$ is a morphism of schemes. Since $\mathcal{F}$ is coherent, $f_*\mathcal{F}$ will be a finite $f_*\mathcal{O}_X$-module. The key question is: when can we expect $f_*\mathcal{O}_X$ to be a finite $\mathcal{O}_Y$-module?

Perhaps the best general answer to this question is "when f is proper". We have seen before that the global sections of a proper scheme satisfy various types of finiteness. Thus, it is reasonable to hope that for any proper morphism $f : X \to Y$ and any open affine $V \subset Y$ the ring $f_*\mathcal{O}_X(V)$ is a finite $\mathcal{O}_V(V)$-module.

Theorem 9.4.9. *Let $f : X \to Y$ be a proper morphism to a locally Noetherian scheme Y. For any coherent sheaf $\mathcal{F}$ on X the pushforward $f_*\mathcal{F}$ is a coherent sheaf on Y.*

We will prove a special case – when f is a projective morphism – in Corollary 12.3.4.

9.4.3 Rank of coherent sheaves

Definition 9.4.10. Let X be a scheme and let $x \in X$. For any $\mathcal{O}_X$-module $\mathcal{F}$, we define the $\kappa(x)$-vector space $\mathcal{F}(x)$ by tensoring $\mathcal{F}_x$ by $\mathcal{O}_{X,x}/\mathfrak{m}_x$. $\mathcal{F}(x)$ is known as the fiber of $\mathcal{F}$ at x.

We define the rank of $\mathcal{F}$ at x to be $\mathrm{rk}_x(\mathcal{F}) := \dim_{\kappa(x)} \mathcal{F}(x)$.

Be careful not to confuse the fiber of $\mathcal{F}$ at x with the stalk of $\mathcal{F}$ at x! The stalk captures the behavior of $\mathcal{F}$ on arbitrarily small open neighborhoods of x via localization; the fiber records the behavior of $\mathcal{F}$ at x via quotienting. Alternatively, if $i : x \to X$ is the inclusion then $\mathcal{F}(x)$ is $i^*\mathcal{F}$ while $\mathcal{F}_x$ is $i^{-1}\mathcal{F}$. (In the next section we will see the definition of a "vector bundle" in algebraic geometry; it is this construction which motivates the "fiber" terminology.)

Warning 9.4.11. Suppose we have an exact sequence of coherent sheaves on X. It is not necessarily true that the sum of the ranks of the two outer terms at a point $x \in X$ is the rank of the middle term; see Exercise 9.4.25.

We have already seen that the stalk $\mathcal{F}_x$ captures information about $\mathcal{F}$ on "small" open neighborhoods of x. Our next goal is to show that when $\mathcal{F}$ is finitely generated, the fiber $\mathcal{F}(x)$ also captures certain features of $\mathcal{F}$ on an open neighborhood of x. The following lemma is a first step in this direction.

Lemma 9.4.12. *Let X be a scheme and let $\mathcal{F}$ be a finitely generated quasicoherent sheaf on X. Suppose that for some point $x \in X$ we have $\mathcal{F}_x = 0$. Then there is an open neighborhood U of x such that $\mathcal{F}|_U = 0$.*

Proof. Since $\mathcal{F}$ is quasicoherent, it suffices to consider the case when $X = \operatorname{Spec}(R)$ is an affine scheme so that $\mathcal{F} = \widetilde{M}$ for some coherent module M. Suppose that $\{m_i\}_{i=1}^r$ are the generators of M. By assumption, the prime ideal $\mathfrak{p}$ corresponding to x satisfies $M_\mathfrak{p} = 0$. This means that for each of the generators m_i, there is an element $f_i \in R \backslash \mathfrak{p}$ such that $f_i m_i = 0$. Then $M_{f_1 \cdots f_r} = 0$. □

As you might guess, Nakayama's Lemma is the key for relating the stalk and the fiber of $\mathcal{F}$ at a point. The following statement is the geometric version of Nakayama's Lemma:

Theorem 9.4.13. *Let X be a scheme and let $\mathcal{F}$ be a finitely generated quasicoherent sheaf on X. Let $x \in X$. Suppose that $f_1, \dots, f_r \in \mathcal{F}_x$ have images in $\mathcal{F}(x)$ which span this $\kappa(x)$-vector space. Then there is an open neighborhood U of x such that $f_1, \dots, f_r \in \mathcal{F}(U)$ and the map*

$$\phi : \mathcal{O}_X^{\oplus r}|_U \xrightarrow{(f_1, \dots, f_r)} \mathcal{F}|_U$$

is surjective.

In particular, if the fiber $\mathcal{F}(x)$ is zero then $\mathcal{F}$ is identically zero on an open neighborhood of x.

Proof. We can apply the usual Nakayama's lemma to see that $\mathcal{F}_x$ is generated by $f_1, \dots, f_r$ as an $\mathcal{O}_{X,x}$-module. Since each f_i is defined by a section of $\mathcal{F}$ on a neighborhood of x, by taking intersections we find a neighborhood U of x such that every $f_i \in \mathcal{F}(U)$. For this choice of U, the cokernel of the map ϕ is a finitely generated quasicoherent sheaf whose stalk at x vanishes. By Lemma 9.4.12, we may shrink U further to ensure that the cokernel of ϕ is the 0 sheaf □

Corollary 9.4.14. *Let X be a scheme and let $\mathcal{F}$ be a finitely generated quasicoherent sheaf on X. Then the function $x \mapsto \operatorname{rk}_x(\mathcal{F})$ is upper semicontinuous. In particular, if $\operatorname{rk}_x(\mathcal{F}) = 0$ then $\mathcal{F}$ is identically zero on an open neighborhood of x.*

9.4.4 Support

Recall that the support of a sheaf $\mathcal{F}$ of abelian groups is the set of points where the stalk is not zero. We next discuss the support of finitely generated sheaves.

Exercise 9.4.15. Let X be a scheme and let $\mathcal{F}$ be a finitely generated quasicoherent sheaf on X. Show that for any point $x \in X$ we have $x \in \operatorname{Supp}(\mathcal{F})$ if and only if $\mathcal{F}(x) \neq 0$.

Lemma 9.4.16. *Let X be a scheme and let $\mathcal{F}$ be a finitely generated quasicoherent sheaf. For any open affine U we have*

$$\operatorname{Supp}(\mathcal{F}) \cap U = V(\operatorname{Ann}(\mathcal{F}(U))).$$

In particular, the support of $\mathcal{F}$ is a closed subset of X.

Proof. Set $R = \mathcal{O}_X(U)$ and $M = \mathcal{F}(U)$. For any prime $\mathfrak{p} \subset R$, Lemma 9.4.12 shows that $M_{\mathfrak{p}} = 0$ if and only if there is some $f \in R$ such that $\mathfrak{p} \not\subset V(f)$ and $M_f = 0$. The first condition on f says that $f \notin \mathfrak{p}$ and the second says that $f \in \operatorname{Ann}(M)$. Thus, we see that $M_{\mathfrak{p}} = 0$ if and only if $\mathfrak{p} \not\subset \operatorname{Ann}(M)$.

The last statement follows from the fact that the closedness of a set can be tested locally. $\square$

Corollary 9.4.17. *Let X be a scheme and let $\mathcal{F}$ be a finitely generated quasicoherent sheaf. Then there is a closed embedding $i : Z \to X$ such that the set-theoretic image of i is* $\operatorname{Supp}(\mathcal{F})$ *and there is a finitely generated quasicoherent sheaf $\mathcal{G}$ on Z with $i_*\mathcal{G} \cong \mathcal{F}$.*

Proof. Annihilators of finitely generated modules are compatible with localization, in the sense that for a ring R, a finitely generated R-module M, and a multiplicatively closed subset S we have $S^{-1}\operatorname{Ann}(M) = \operatorname{Ann}(S^{-1}M)$. Thus the locally defined ideals $\operatorname{Ann}(\mathcal{F}(U))$ define a quasicoherent ideal sheaf on X, hence a closed subscheme Z. The last statement follows from the fact that an R-module M can equally well be thought of as an $R/\operatorname{Ann}(M)$-module. $\square$

Definition 9.4.18. Let X be a scheme and let $\mathcal{F}$ be a finitely generated quasicoherent sheaf. The scheme-theoretic support of $\mathcal{F}$ is the subscheme Z defined locally by the annihilator of $\mathcal{F}(U)$.

When X is a Noetherian scheme, then the scheme theoretic support of a finitely generated quasicoherent sheaf $\mathcal{F}$ can be studied using the theory of associated primes.

9.4.5 Torsion sheaves

Definition 9.4.19. Let X be an scheme and let $\mathcal{F}$ be an $\mathcal{O}_X$-module. We say that $\mathcal{F}$ is a torsion sheaf if for every point $x \in X$ the stalk $\mathcal{F}_x$ is a torsion $\mathcal{O}_{X,x}$-module.

When R is an integral domain, we have a nice theory of torsion modules: an R-module M is torsion if and only if it vanishes upon tensoring with $\mathrm{Frac}(R)$.

Lemma 9.4.20. *Let X be an integral scheme and let $\mathcal{F}$ be a finitely generated quasicoherent sheaf on X. The following are equivalent:*

(1) For every $x \in X$ the stalk $\mathcal{F}_x$ is a torsion $\mathcal{O}_{X,x}$-module.

(2) For every open affine $U \subset X$ we have that $\mathcal{F}(U)$ is a torsion $\mathcal{O}(U)$-module.

(3) The stalk of $\mathcal{F}$ at the generic point of X is zero.

(4) $\mathrm{Supp}(\mathcal{F}) \subsetneq X$.

Proof. (1) $\Leftrightarrow$ (2): follows from the fact that for a domain R and an R-module M we have that M is a torsion module if and only if $M_{\mathfrak{p}}$ is torsion for every prime ideal $\mathfrak{p} \subset R$.

(1) $\Rightarrow$ (3): clear

(3) $\Leftrightarrow$ (4): follows from Lemma 9.4.16.

(4) $\Rightarrow$ (2): Lemma 9.4.16 shows that for any open affine U we have $\mathrm{Ann}(\mathcal{F}(U)) \neq 0$. $\square$

Every finitely generated R-module M fits into an exact sequence $0 \to M_{tors} \to M \to M_{tf} \to 0$ where M_{tors} is the torsion submodule. This construction sheafifies to give an exact sequence

$$0 \to \mathcal{F}_{tors} \to \mathcal{F} \to \mathcal{F}_{tf} \to 0$$

where $\mathcal{F}_{tors}$ is the torsion subsheaf. We say that $\mathcal{F}$ is torsion-free if its torsion subsheaf is zero.

9.4.6 Exercises

Exercise 9.4.21. Let X be a scheme. Suppose that $\mathcal{F}$ is a coherent sheaf on X and that $\mathcal{G}$ is a quasicoherent sheaf. Prove that $\mathcal{H}om(\mathcal{F}, \mathcal{G})$ is quasicoherent.

Let R be the DVR $\mathbb{K}[x]_{(x)}$, let $X = \mathrm{Spec}(R)$, and let $\mathcal{F} = \widetilde{\oplus_{i=1}^{\infty} R}$. Show that $\mathcal{H}om_{\mathcal{O}_X}(\mathcal{F}, \mathcal{O}_X)$ is not quasicoherent.

Exercise 9.4.22. Let $X = \mathbb{P}^2$ and consider the coherent sheaf $\mathcal{F} = i_*\mathcal{O}_{\mathbb{P}^1}$ where $i : \mathbb{P}^1 \to \mathbb{P}^2$ is the inclusion of a line ℓ. Let U be the open set in $\mathbb{P}^2$ which is the complement of a point in ℓ. Show that $\mathcal{F}(U)$ is not a finitely generated $\mathcal{O}_X(U)$-module.

Exercise 9.4.23. Let $f : X \to Y$ be a morphism of schemes and let $\mathcal{F}$ be a finitely generated quasicoherent sheaf on Y. Prove that $\mathrm{Supp}(f^*\mathcal{F}) = f^{-1}(\mathrm{Supp}(\mathcal{F}))$.

Exercise 9.4.24. Find an example of a scheme X and a quasicoherent sheaf $\mathcal{F}$ such that $\mathrm{Supp}(\mathcal{F})$ is not closed. (Hint: consider the sheaf $\oplus_{a\in\mathbb{C}}\mathcal{F}_a$ on $\mathbb{A}^1_{\mathbb{C}}$ where $\mathcal{F}_a$ is the skyscraper sheaf with value $\mathbb{C}$ at the point $(x - a)$.)

Exercise 9.4.25. Suppose that we have a scheme X, a point $x \in X$, and an exact sequence of coherent sheaves

$$0 \to \mathcal{F} \to \mathcal{G} \to \mathcal{H} \to 0.$$

(1) Prove that $\mathrm{rk}_x(\mathcal{F}) + \mathrm{rk}_x(\mathcal{H}) \geq \mathrm{rk}_x(\mathcal{G})$.

(2) Find an example where we have a strict inequality.

(3) Suppose that $\mathcal{H}$ is locally free. Prove that $\mathrm{rk}_x(\mathcal{F}) + \mathrm{rk}_x(\mathcal{H}) = \mathrm{rk}_x(\mathcal{G})$.

Exercise 9.4.26. Let X be a regular curve over a field.

(1) Prove that every torsion free coherent sheaf $\mathcal{F}$ on X is locally free, i.e. X admits an open cover by open affines U such that the restriction of $\mathcal{F}$ to U is defined by a free $\mathcal{O}_X(U)$-module.

(2) Prove that for every torsion sheaf $\mathcal{T}$ there is some closed subscheme $i : Z \to X$ such that $\mathcal{T} \cong i_*\mathcal{O}_Z$.

9.5 Locally free sheaves

As described in the introduction to the chapter, in algebraic geometry there are certain $\mathcal{O}_X$-modules which play the role of vector bundles. This section is dedicated to exploring these quasicoherent sheaves.

Definition 9.5.1. Let X be a scheme. A quasicoherent sheaf $\mathcal{F}$ on X is locally free if X admits a covering by open affines U such that $\mathcal{F}|_U$ is defined by a free $\mathcal{O}_X(U)$-module.

We say that a locally free sheaf $\mathcal{F}$ on X has rank r if the rank of $\mathcal{F}$ is r at every point, or equivalently, if X admits an open cover $\{U_i\}$ such that $\mathcal{F}|_{U_i}$ is isomorphic to $\mathcal{O}_{U_i}^{\oplus r}$. A locally free sheaf of rank 1 is called an invertible sheaf.

We will focus on locally free sheaves of finite rank.

Warning 9.5.2. In contrast to some of our other constructions, locally free does *not* mean that the restriction of $\mathcal{F}$ to every open affine is free. Rather, the correct affine analogue of a "finite rank locally free sheaf" is a projective module. More precisely, over any ring R a finitely presented R-module M is projective if and only if the localization of M along every prime ideal is free.

It is interesting to ask which types of rings have the property that every projective module is free. Easier examples include local rings and PIDs. The Quillen-Suslin Theorem (which contributed to Quillen winning a Fields Medal) shows that projective modules are free for polynomial rings over fields – in other words, every locally free coherent sheaf on affine space is defined by a free module.

For finitely presented quasicoherent sheaves local freeness is a condition that can be checked on stalks.

Proposition 9.5.3. *Let X be a scheme and let $\mathcal{F}$ be a finitely presented quasicoherent sheaf. Suppose that $x \in X$ is a point such that $\mathcal{F}_x$ is a locally free $\mathcal{O}_{X,x}$-module. Then there is some open neighborhood U of x such that $\mathcal{F}|_U$ is a locally free quasicoherent sheaf on U.*

Proof. Choose a finite set of elements $m_1, \ldots, m_r$ which generate $\mathcal{F}_x$ as a $\mathcal{O}_{X,x}$-module. Applying Theorem 9.4.13 (Geometric Nakayama's Lemma) we obtain an open neighborhood V of x and a surjection $\phi : \mathcal{O}_X^r|_V \to \mathcal{F}|_V$. Since $\mathcal{F}$ is finitely presented, the kernel of this map is a finitely generated $\mathcal{O}_V$-module $\mathcal{G}_V$. Since the stalk of $\mathcal{G}_V$ at x is trivial, it is trivial on an open neighborhood U of x in V by Lemma 9.4.12. On U the map ϕ is an isomorphism. □

Corollary 9.5.4. *Let X be a scheme and let $\mathcal{F}$ be a finitely presented quasicoherent sheaf on X. If $\mathcal{F}_x$ is locally free for every point $x \in X$ then $\mathcal{F}$ is locally free.*

In fact, if we assume a bit more about X then local freeness can even be detected on fibers.

Theorem 9.5.5. *Let X be a reduced scheme. Suppose that $\mathcal{F}$ is a finitely presented quasicoherent sheaf on X such that the rank of $\mathcal{F}$ is constant at all the points of X. Then $\mathcal{F}$ is locally free.*

We need to assume that X is reduced because the rank does not "see" the non-reduced structure of X along the generic point of a component of X. (Consider for example the skyscraper sheaf $\mathbb{K}$ on $\mathrm{Spec}(\mathbb{K}[x]/(x^2))$.)

Proof. It suffices to check that for an open affine U the restriction $\mathcal{F}|_U$ is locally free. We let r denote the rank of $\mathcal{F}$.

Fix a point $x \in U$ and choose r elements $f_1, \dots, f_r \in \mathcal{F}_x$ which generate the fiber $\mathcal{F}(x)$ as a $\kappa(x)$-vector space. By Theorem 9.4.13 after perhaps shrinking U the f_i extend to U and we have a surjection

$$\phi : \mathcal{O}_X^{\oplus r}|_U \xrightarrow{(f_1,\dots,f_r)} \mathcal{F}|_U$$

Write $U \cong \mathrm{Spec}(R)$ and let $R^{\oplus R} \to M$ denote the map of modules induced by ϕ. Suppose that there is a non-zero element $(r_1, \dots, r_r)$ in the kernel. In particular there is some r_j which is non-zero; since R is reduced, this means that there is some prime $\mathfrak{p}$ such that $r_j \notin \mathfrak{p}$. If $y \in U$ denotes the corresponding point, then $\phi_y : \mathcal{O}_{X,y}^{\oplus r} \to \mathcal{F}_y$ has a kernel. Since it is surjective by construction, we conclude that the rank of $\mathcal{F}$ at y is smaller than r, contradicting our assumption. $\square$

Locally free sheaves interact well with many of the module constructions we have dealt with so far. Just as for vector bundles, the natural functorial operation for locally free sheaves is the pullback. The exercises explain these compatibilities in more detail. We will prove one important statement here that is frequently used to compute pushforwards.

Proposition 9.5.6 (Projection formula)**.** *Let $f : X \to Y$ be a morphism of schemes. Suppose that $\mathcal{F}$ is an $\mathcal{O}_X$-module and that $\mathcal{G}$ is a locally free $\mathcal{O}_Y$-module of finite rank. Then*

$$f_*(\mathcal{F} \otimes f^*\mathcal{G}) \cong f_*\mathcal{F} \otimes \mathcal{G}.$$

Proof. By tensoring the counit map $f^*f_*\mathcal{F} \to \mathcal{F}$ by $f^*\mathcal{G}$ we get an induced morphism $f^*f_*\mathcal{F} \otimes f^*\mathcal{G} \to \mathcal{F} \otimes f^*\mathcal{G}$. Since tensor product commutes with pullback, we can identify the left-hand side as $f^*(f_*\mathcal{F} \otimes \mathcal{G}) \cong f^*f_*\mathcal{F} \otimes f^*\mathcal{G}$. Thus by adjunction we obtain a morphism $\phi : f_*\mathcal{F} \otimes \mathcal{G} \to f_*(\mathcal{F} \otimes f^*\mathcal{G})$.

To show that ϕ is an isomorphism, it suffices to argue locally. Thus we may assume that $\mathcal{E} = \mathcal{O}_Y^{\oplus r}$. Using $f^*\mathcal{O}_Y^{\oplus r} \cong \mathcal{O}_X^{\oplus r}$, we obtain an isomorphism as desired. $\square$

9.5.1 Vector bundles

We next clarify how a locally free sheaf of finite rank is the "same thing" as a vector bundle in algebraic geometry. The following construction will allow us to construct a scheme from a locally free sheaf.

Definition 9.5.7. Let X be a scheme. A quasicoherent sheaf of $\mathcal{O}_X$-algebras is a quasicoherent sheaf of $\mathcal{O}_X$-modules that is simultaneously a sheaf of rings on X.

Construction 9.5.8. Let X be a scheme and let $\mathcal{A}$ be a quasicoherent sheaf of $\mathcal{O}_X$-algebras. For any open affine U in X, the sections $\mathcal{A}(U)$ form a ring which in turn defines an affine scheme $A_U := \operatorname{Spec}(\mathcal{A}(U))$. For any inclusion of open affines $V \subset U$ the restriction map $\rho_{U,V} : \mathcal{A}(U) \to \mathcal{A}(V)$ defines an inclusion $\psi_{V,U} : \operatorname{Spec}(\mathcal{A}(V)) \to \operatorname{Spec}(\mathcal{A}(U))$.

As we vary over all open affines U, the resulting schemes A_U equipped with the maps $\psi_{V,U}$ satisfy the cocycle condition. By gluing over all identifications obtained via the maps $\psi_{V,U}$, we obtain a scheme $\underline{\operatorname{Spec}}(\mathcal{A})$ equipped with a morphism $\pi : \underline{\operatorname{Spec}}(\mathcal{A}) \to X$. By construction π is an affine morphism.

The construction $\underline{\operatorname{Spec}}$ is known as the "relative Spec" construction. Its defining feature is that $\pi_*\mathcal{O}_{\underline{\operatorname{Spec}}(\mathcal{A})} = \mathcal{A}$.

Exercise 9.5.9. Let X be a scheme and let $\mathcal{A}$ be a quasicoherent sheaf of $\mathcal{O}_X$-algebras. Show that the fiber of $\underline{\operatorname{Spec}}(\mathcal{A}) \to X$ over a point x is isomorphic to $\operatorname{Spec}(\mathcal{A}(x))$.

Definition 9.5.10. Let X be a scheme. Given any locally free sheaf $\mathcal{F}$ of constant rank r, the symmetric algebra $\operatorname{Sym}(\mathcal{F}^\vee)$ is a quasicoherent sheaf of $\mathcal{O}_X$-algebras. Set $\mathcal{V} = \underline{\operatorname{Spec}}(\operatorname{Sym}(\mathcal{F}^\vee))$. We call $\pi : \mathcal{V} \to X$ the vector bundle associated to the sheaf $\mathcal{F}$. Note that we can recover $\mathcal{F}$ from $\mathcal{V}$ by defining $\mathcal{F} = (\pi_*\mathcal{O}_{\mathcal{V}})^\vee$.

As discussed in the introduction to the chapter, this construction realizes $\mathcal{F}$ as the sheaf of sections on $\mathcal{V}$. We included the dual $\vee$ in the definition to ensure that we obtained in this relationship. (This argument is made rigorous in Exercise 9.5.25.)

Note that the fiber of $\underline{\operatorname{Spec}}(\mathcal{F}^\vee) \to X$ over a point x is isomorphic to $\mathbb{A}^r_{\kappa(x)}$ by Exercise 9.5.9. In fact, $\underline{\operatorname{Spec}}(\mathcal{F}^\vee) \to X$ is an example of a "geometric vector bundle" as defined in Exercise 9.5.24. The converse is also true; Exercise 9.5.24 shows that every geometric vector bundle over X is defined by a locally free sheaf on X so that there is a bijection between the two objects.

Example 9.5.11. Let's analyze the tautological bundle $\mathcal{T}$ of $\mathbb{P}^n$. By definition $\mathcal{T}$ is the subbundle of the trivial bundle $\mathbb{A}^{n+1} \times \mathbb{P}^n$ that associates to each point in projective space the corresponding line in $\mathbb{A}^{n+1}$. (Note that we can also think of $\mathcal{T}$ as the blow-up of the origin in $\mathbb{A}^{n+1}$.) We will use coordinates x_i on $\mathbb{P}^n$ and y_i on $\mathbb{A}^{n+1}$.

Over the affine chart $U_i = D_{+,x_i}$, the tautological bundle $\mathcal{T}$ is the vanishing locus of the ideal

$$\left(\tfrac{x_0}{x_i}y_i - y_0, \dots, \tfrac{x_n}{x_i}y_i - y_n\right) \subset \mathbb{K}\left[y_0, \dots, y_n, \tfrac{x_0}{x_i}, \dots, \tfrac{x_n}{x_i}\right]$$

Note that the quotient is a free $\mathbb{K}[\frac{x_0}{x_i}, \dots, \frac{x_n}{x_i}]$-module of rank 1 generated by y_i. Since y_i matches up with the degree 1 homogeneous function x_i, by comparing against Example 7.1.8 we see that the pushforward of the structure sheaf of $\mathcal{T}$ is the bundle $\mathcal{O}_{\mathbb{P}^n}(1)$. We can also see this using transition functions; the sequence of maps

$$\mathcal{O}_{U_i}|_{U_i \cap U_j} \to \pi_*\mathcal{O}_{\mathcal{T}}|_{U_i \cap U_j} \to \mathcal{O}_{U_j}|_{U_i \cap U_j}$$

sends $1 \mapsto y_i$ and $y_j \mapsto 1$. Since on $U_i \cap U_j$ we have the equation $x_j y_i = x_i y_j$, the composition sends $1 \mapsto \frac{x_i}{x_j}$ and this defines $\mathcal{O}_{\mathbb{P}^n}(1)$ by Exercise 7.6.6.

Taking a dual as in Definition 9.5.10, we see that the tautological bundle $\mathcal{T}$ on projective space has sheaf of sections $\mathcal{O}_{\mathbb{P}^n}(-1)$. This coheres with the well-known geometric fact that the tautological bundle does not admit any global sections; indeed, there are no non-constant holomorphic maps $\mathbb{CP}^1 \to \mathbb{C}^2$. (In contrast, sections of $\mathcal{O}_{\mathbb{P}^n}(1)$ are linear functions on $\mathbb{A}^{n+1}$ and thus correspond to points of the dual space.)

9.5.2 Transition functions

Suppose that $\mathcal{F}$ is a locally free sheaf of rank r on a scheme X. Then there is an open cover $\{U_i\}$ of X such that we have isomorphisms $\phi_i : \mathcal{F}|_{U_i} \to \mathcal{O}_{U_i}^{\oplus r}$. Furthermore the maps ψ_{ij} defined by the compositions

$$\psi_{ij} : \mathcal{O}_{U_i \cap U_j}^{\oplus r} \xrightarrow{\phi_i^{-1}} \mathcal{F}|_{U_i \cap U_j} \xrightarrow{\phi_j} \mathcal{O}_{U_i \cap U_j}^{\oplus r}$$

satisfy the cocycle conditions. We will call these sheaf morphisms ψ_{ij} the transition functions for $\mathcal{F}$. (Note that "the" transition functions depend upon the choice of isomorphisms ϕ_i. Since an isomorphism of $\mathcal{O}_X(U_i)$ is multiplication by a unit in $\mathcal{O}_X(U_i)^\times$, two possible choices of transition function will also differ locally by multiplication by a unit. More precisely, if we choose units $u_i \in \mathcal{O}_X(U_i)^\times$ then the transition functions $\psi'_{ij} = \psi_{ij} \cdot u_i^{-1} \cdot u_j$ should be considered equivalent to the transition functions ψ_{ij}.)

Conversely, starting from the sheaves $\mathcal{O}_{U_i}^{\oplus r}$ and "local gluing data" $\psi_{ij} : \mathcal{O}_{U_i \cap U_j}^{\oplus r} \to \mathcal{O}_{U_i \cap U_j}^{\oplus r}$ that satisfies the cocycle condition, we can use Corollary 7.6.3 to construct a locally free sheaf on X.

Example 9.5.12. In Example 7.6.5 and Exercise 7.6.6 we showed that the locally free sheaves $\mathcal{O}(d)$ on $\mathbb{P}^n$ can be constructed by taking the structure sheaves on our standard charts and gluing them via the isomorphisms ψ_{ij} defined by multiplying by $(\frac{x_i}{x_j})^d$.

It is important to note that the global sections s of a locally free sheaf $\mathcal{F}$ admit the same transition functions as $\mathcal{F}$ does. Namely, if we take $s \in \mathcal{F}$ and restrict to a trivializing open set U_i we can identify $s|_{U_i}$ with an element $s_i \in \mathcal{O}_X(U_i)^{\oplus r}$. Given two open sets U_i, U_j, the map ψ_{ij} will take $s_i|_{U_i \cap U_j}$ to $s_j|_{U_j \cap U_j}$. In other words, the global sections of $\mathcal{F}$ are the same as systems of local sections whose behavior on the overlaps are described by the transition functions.

9.5.3 Zero loci of sections

In Chapter 2 we associated to any homogeneous polynomial $f \in \mathcal{O}_{\mathbb{P}^n}(d)$ the vanishing locus $V_+(f) \subset \mathbb{P}^n$. As described in Example 9.3.13, the vanishing locus of f is the closed subscheme defined by the image of the map $\mathcal{O}_{\mathbb{P}^n}(-d) \xrightarrow{\cdot f} \mathcal{O}_{\mathbb{P}^n}$.

More generally, given any scheme X equipped with a locally free sheaf $\mathcal{F}$ and any global section $s \in \mathcal{F}(X)$, we can define the vanishing locus of s in an analogous way. This construction is more important than it appears at first: the zero loci of sections of locally free sheaves play a special role in many situations (for example in intersection theory).

Definition 9.5.13. Suppose that X is a scheme and $\mathcal{F}$ is a locally free sheaf of rank r on X. Suppose that $s \in \mathcal{F}(X)$ is a global section, or equivalently, $s : \mathcal{O}_X \to \mathcal{F}$ is a morphism. The zero locus of the section s is the closed subscheme of X defined by the quasicoherent ideal sheaf which is the image of $s^\vee : \mathcal{F}^\vee \to \mathcal{O}_X$.

Example 9.5.14. When $\mathcal{F}$ is the structure sheaf $\mathcal{O}_X$ then the zero locus of a section f is the same as $V(f)$. When $\mathcal{F} \cong \mathcal{O}(d)$ on $\mathrm{Proj}(S)$ then the zero locus of a section f is the same as $V_+(f)$.

The zero locus of a section of a locally free sheaf is a very natural geometric notion. The following exercise gives an alternative definition (see also Exercise 9.5.27).

Exercise 9.5.15. Let X be a scheme and let $\mathcal{F}$ be a locally free sheaf of rank r on X. Suppose $s \in \mathcal{F}(X)$. For any point $x \in X$, there is an open affine neighborhood U such that we have an isomorphism $\psi_U : \mathcal{F} \to \mathcal{O}_U^{\oplus r}$. In particular, $\psi_U(s|_U)$ is an r-tuple of functions $(f_1, \ldots, f_r)$. Show that the intersection of $Z(s)$ with U is defined by the ideal $(f_1, \ldots, f_r) \subset \mathcal{O}_X(U)$. (Why does the choice of ψ not affect the resulting ideal sheaf?)

Exercise 9.5.16. Let X be a scheme, $\mathcal{F}$ a locally free sheaf, $s \in \mathcal{F}(X)$. Prove that a point x is contained in $Z(s)$ if and only if the restriction of s to the fiber $\mathcal{F}(x)$ is zero.

By Krull's PIT we can "expect" the vanishing locus of a section of a locally free sheaf of rank r to have codimension r (although as we have seen it is certainly possible for the number of equations to differ from the codimension). This relationship is particularly important for locally free sheaves of rank 1; we will return to this topic in depth in Chapter 10.

9.5.4 Exercises

Exercise 9.5.17. Let X be a scheme and let $\mathcal{F}$ be a locally free sheaf on X. Show that the functor $\mathcal{F} \otimes -$ is an exact functor.

Exercise 9.5.18. Let X be a scheme. Suppose that $\mathcal{F}$ and $\mathcal{G}$ are locally free sheaves of finite rank on X. Prove that $\mathcal{F} \otimes \mathcal{G}$, $\mathcal{H}om(\mathcal{F}, \mathcal{G})$, and $\mathcal{F} \oplus \mathcal{G}$ are also locally free sheaves of finite rank.

Exercise 9.5.19. Let X be a scheme and let $\mathcal{F}$ be a locally free sheaf of finite rank on X. Prove that $\mathcal{F}^\vee$ is again a locally free sheaf and that $(\mathcal{F}^\vee)^\vee \cong \mathcal{F}$. Show by example that $\mathcal{F}^\vee$ need not be isomorphic to $\mathcal{F}$.

Exercise 9.5.20. (1) Let $f : X \to Y$ be a morphism. Show that the pullback of a locally free sheaf on Y is locally free on X.

(2) Let $f : X \to \mathbb{A}^2$ be the blow-up of the origin and let $\mathcal{F}$ be the ideal sheaf of the exceptional divisor in X. Show that $\mathcal{F}$ is locally free but that $f_*\mathcal{F}$ is not locally free.

Exercise 9.5.21. Let X be a scheme. Suppose that $\mathcal{F}$ is a finite rank locally free sheaf and that $\mathcal{G}, \mathcal{H}$ are $\mathcal{O}_X$-modules. Prove that there is an isomorphism $\mathcal{H}om(\mathcal{G} \otimes \mathcal{F}, \mathcal{H}) \cong \mathcal{H}om(\mathcal{G}, \mathcal{F}^\vee \otimes \mathcal{H})$.

Exercise 9.5.22. Let X be a scheme and let $\mathcal{F}$ be a locally free sheaf of finite rank on X. Fix an open cover $\{U_i\}$ that trivializes $\mathcal{F}$ and suppose that ψ_{ij} are the transition functions for $\mathcal{F}$ with respect to the cover.

(1) Prove that $\{U_i\}$ also trivializes $\mathcal{F}^\vee$ with the transition functions ψ_{ij}^{-1}.

(2) Suppose that $\mathcal{G}$ is an invertible sheaf on X that is also trivialized by $\{U_i\}$ with transition functions μ_{ij}. Show that $\{U_i\}$ trivializes $\mathcal{F} \otimes \mathcal{G}$ with transition functions defined by $\psi_{ij} \cdot \mu_{ij}$.

Exercise 9.5.23. There are some important ways in which vector bundles behaves differently in algebraic geometry and in topology. For example, consider the surjective map of bundles $\phi : \mathcal{O}_{\mathbb{P}^1}(-1) \oplus \mathcal{O}_{\mathbb{P}^1}(-1) \to \mathcal{O}_{\mathbb{P}^1}$ defined by multiplication by x on the first factor and y on the second. Show that ϕ is a surjective map of constant rank but that it does not admit any splitting.

Exercise 9.5.24. Let X be a scheme. A geometric vector bundle on X of rank r is a morphism $\pi : \mathcal{V} \to X$ such that there is an open cover $\{U_i\}$ of X satisfying the following properties:

(1) For every i there is an isomorphism $\phi_i : \mathbb{A}^r_{\mathbb{Z}} \times U_i \to \pi^{-1}U_i$ that commutes with the morphisms to U_i.

(2) Suppose that $V \subset U_i \cap U_j$ is an open affine defined by the ring R so that $\mathbb{A}^r_{\mathbb{Z}} \times V \cong \operatorname{Spec}(R[x_1, \ldots, x_n])$. Then the isomorphism $\phi_j \circ \phi_i^{-1} : \operatorname{Spec}(R[x_1, \ldots, x_n]) \to \operatorname{Spec}(R[x_1, \ldots, x_n])$ is defined by an invertible linear R-algebra homomorphism.

Given a geometric vector bundle π, prove that $\pi_*\mathcal{O}_V$ is a locally free sheaf of rank r. Deduce that there is a bijection between locally free sheaves of rank r (up to isomorphism) and geometric vector bundles of rank r (up to isomorphisms compatible with the structure maps to X).

Exercise 9.5.25. Let X be a scheme and let $\mathcal{A}$ be a quasicoherent sheaf of $\mathcal{O}_X$-algebras. Let Y be an X-scheme, that is, a scheme Y equipped with a structural morphism $Y \to X$. Show that there is a bijection

$$\operatorname{Hom}_{\mathcal{O}_X-alg}(\mathcal{A}, \pi_*\mathcal{O}_Y) \cong \operatorname{Hom}_{\mathbf{Sch}/S}(Y, \underline{\operatorname{Spec}}(\mathcal{A}))$$

Suppose now that $\mathcal{F}$ is a locally free sheaf on X and $\pi : \mathcal{V} \to X$ the corresponding vector bundle. Prove that for any open set U there is a bijection between sections $s \in \mathcal{F}(U)$ and geometric sections $s : U \to \mathcal{V}|_U$ of π over U. (Hint: show that $s \in \mathcal{F}(U)$ defines a map of algebras $\mathrm{Sym}(\mathcal{F}|_U^\vee) \to \mathcal{O}_U$ via contraction.)

Exercise 9.5.26. Show that if $f : X \to Y$ is an affine morphism then there is some quasicoherent sheaf $\mathcal{A}$ of $\mathcal{O}_Y$-algebras such that $X \cong \underline{\mathrm{Spec}}(\mathcal{A})$ and f is identified with the canonical map $\pi : \underline{\mathrm{Spec}}(\mathcal{A}) \to Y$.

Exercise 9.5.27. Let X be a scheme and let $\mathcal{F}$ be a locally free sheaf of rank r on X defining the geometric vector bundle $\mathcal{V} = \underline{\mathrm{Spec}}(\mathcal{F}^\vee)$. Note that there is a closed embedding $z : X \to \mathcal{V}$ defined by taking the zero section of $\mathcal{V}$.

Fix a global section $s \in \mathcal{F}(X)$; by Exercise 9.5.25 we get a morphism $s : X \to \mathcal{V}$. Show that the zero locus $Z(s)$ is the same as the pullback of the closed embedding z under the map $s : X \to \mathcal{V}$.

Exercise 9.5.28. Let X be a scheme, let $\mathcal{L}$ be an invertible sheaf on X, and let $s \in \mathcal{L}(X)$. Define X_s to be the complement of the zero locus $Z(s)$. Let $\mathcal{F}$ be a quasicoherent sheaf on X. The following statements give our final generalization of Exercise 8.5.17.

(1) Suppose that X is quasicompact. Suppose that $a \in \mathcal{F}(X)$ is an element whose restriction to X_s is 0. Prove that for some $n > 0$ we have $s^n a = 0$ as a section of $(\mathcal{F} \otimes \mathcal{L}^n)(X)$.

(2) Suppose that X is quasicompact quasiseparated. Let $b \in \mathcal{F}(X_s)$. Show that for some $n > 0$ the element $s^n b$ is the restriction of an element of $(\mathcal{F} \otimes \mathcal{L}^n)(X)$.

(3) Suppose that X is quasicompact quasiseparated. Then $\mathcal{F}(X_s) \cong \oplus_{n \geq 0}(\mathcal{F} \otimes \mathcal{L}^n)(X)$.

Exercise 9.5.29. Let X be a normal integral Noetherian scheme. Suppose that $U \subset X$ is an open subset such that $X \backslash U$ has codimension ≥ 2. Prove that for any locally free sheaf $\mathcal{F}$ on X the restriction map $\mathcal{F}(X) \to \mathcal{F}(U)$ is an isomorphism. (Hint: use the S2 property locally.)

9.6 Graded modules and the Proj construction

Suppose that S is a $\mathbb{Z}_{\geq 0}$-graded ring. In this section we will study the correspondence between graded S-modules and quasicoherent sheaves on $\mathrm{Proj}(S)$. Recall that for any $\mathbb{Z}$-graded S-module N, we define $N(d)$ to be the same module shifted in degree by d: $N(d)_m := N_{d+m}$.

A $\mathbb{Z}_{\geq 0}$-graded ring S is said to be generated in degree 1 if S is generated by S_1 as an S_0-algebra. This is an important technical assumption that will show up several times in this section. In practice this assumption is not too restrictive: if S is a finitely generated S_0-algebra then Exercise 2.7.10 shows that some Veronese subalgebra $S^{(m)}$ that is generated in degree 1, and Proposition 2.7.7 verifies that $\mathrm{Proj}(S) \cong \mathrm{Proj}(S^{(m)})$.

One special feature of degree 1 elements in S is the following result:

Exercise 9.6.1. Let N be a $\mathbb{Z}$-graded S-module. If $f \in S$ has degree 1, prove that $(N_f)_0 \cong N/(f-1)N$ under the bijection $n/f^{\deg(n)} \leftrightarrow n + (f-1)N$.

As a consequence we have

Corollary 9.6.2. *Let N be a $\mathbb{Z}$-graded S-module. If $f \in S$ has degree 1, prove that $(N_f)_0 \cong N \otimes_S (S_f)_0$.*

9.6.1 Constructing a sheaf

We will define a functor $\widetilde{}^{+}$ which associates to any $\mathbb{Z}$-graded S-module N a quasicoherent sheaf $\widetilde{N}^+$.

Construction 9.6.3. Let S be a $\mathbb{Z}_{\geq 0}$-graded ring. Suppose that N is a $\mathbb{Z}$-graded S-module. For every homogeneous $f \in S_+$ the $(S_f)_0$-module $(N_f)_0$ defines a quasicoherent sheaf $\widetilde{(N_f)_0}$ on $D_{+,f}$. These modules satisfy a natural compatibility under localization maps; applying Corollary 7.6.3 to glue the resulting sheaves we obtain a quasicoherent sheaf $\widetilde{N}^+$ on $\mathrm{Proj}(S)$.

Exercise 9.6.4. Show that for any point $\mathfrak{p} \in \mathrm{Proj}(S)$ the stalk of $\widetilde{N}^+$ at $\mathfrak{p}$ is the degree 0 piece of the homogeneous localization $N_{(\mathfrak{p})}$ (that is, the localization of N along all homogeneous elements not contained in $\mathfrak{p}$).

Exercise 9.6.5. Show that $\widetilde{}^{+}$ defines an exact functor from the category of $\mathbb{Z}$-graded S-modules to the category of quasicoherent sheaves on $\mathrm{Proj}(S)$.

Unlike the affine case, $\widetilde{}^{+}$ is not an equivalence of categories. The issue is one that we have observed before in the context of homogeneous ideals: if two graded S-modules N, N' agree in sufficiently high degrees then they define the same quasicoherent sheaf on $\mathrm{Proj}(S)$.

There is one more important property satisfied by the $\widetilde{\ }^{+}$ functor. Suppose that M and N are $\mathbb{Z}$-graded S-modules. For any homogeneous $f \in S_+$ we have a natural map

$$\begin{aligned}(M_f)_0 \otimes_{(S_f)_0} (N_f)_0 &\to ((M \otimes_S N)_f)_0 \\ \frac{m}{f^{\deg(m)/deg(f)}} \otimes \frac{n}{f^{\deg(n)/deg(f)}} &\mapsto \frac{m \otimes n}{f^{\deg(n+m)/deg(f)}}\end{aligned}$$

These local maps induce a global morphism $\widetilde{M}^+ \otimes_{\mathcal{O}_{\mathrm{Proj}(S)}} \widetilde{N}^+ \to \widetilde{M \otimes_S N}^+$.

Proposition 9.6.6. *Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is generated in degree 1. For any two $\mathbb{Z}$-graded S-modules M, N the map*

$$\widetilde{M}^+ \otimes_{\mathcal{O}_X} \widetilde{N}^+ \to \widetilde{M \otimes_S N}^+$$

is an isomorphism.

Proof. By assumption $\mathrm{Proj}(S)$ is covered by distinguished open affines of the form $D_{+,f}$ where f is homogeneous of degree 1. It suffices to verify that the globally-defined morphism discussed above restricts to an isomorphism on each of these open sets. Using Corollary 9.6.2 we see that

$$\begin{aligned}(M_f)_0 \otimes_{(S_f)_0} (N_f)_0 &\cong (M \otimes_S (S_f)_0) \otimes_{(S_f)_0} (N \otimes_S (S_f)_0) \\ &\cong (M \otimes_S N) \otimes_S (S_f)_0\end{aligned}$$

which verifies the desired isomorphism over $D_{+,f}$. □

9.6.2 Invertible sheaves

There are certain quasicoherent sheaves which play a particularly important role in our graded theory.

Definition 9.6.7. Let S be a $\mathbb{Z}_{\geq 0}$-graded ring. We denote the quasicoherent sheaf $\widetilde{S(d)}^+$ by $\mathcal{O}_{\mathrm{Proj}(S)}(d)$ (or just $\mathcal{O}(d)$ when the context is understood).

Example 9.6.8. For $\mathbb{P}^n_{\mathbb{K}}$ Definition 9.6.7 agrees with the definition of the invertible sheaves $\mathcal{O}_{\mathbb{P}^n}(d)$ given in Example 7.1.8 so there is no conflict in notation.

One might expect that the quasicoherent sheaves $\mathcal{O}(d)$ have a natural interaction with the grading structure of graded S-modules N. The following exercise verifies this relationship under our usual hypothesis:

Exercise 9.6.9. Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is generated in degree 1. Let N be a $\mathbb{Z}$-graded S-module. Show that $\widetilde{N}^+ \otimes \mathcal{O}(d) \cong \widetilde{N(d)}^+$. In particular, show that $\mathcal{O}(n) \otimes \mathcal{O}(m) \cong \mathcal{O}(n+m)$.

In keeping with the previous exercise, we define:

Definition 9.6.10. Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is generated in degree 1. Let $\mathcal{F}$ be a quasicoherent sheaf on $\mathrm{Proj}(S)$. We define $\mathcal{F}(n) := \mathcal{F} \otimes \mathcal{O}(n)$.

We now turn to the study of the sheaves $\mathcal{O}(d)$. Just as for $\mathbb{P}^n_{\mathbb{K}}$, these sheaves are often invertible sehaves.

Lemma 9.6.11. *Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is generated in degree 1. Then $\mathcal{O}(d)$ is an invertible sheaf on $\mathrm{Proj}(S)$.*

The restriction on S is necessary; see Exercise 9.6.23.

Proof. By assumption $\mathrm{Proj}(S)$ is covered by the distinguished open affines $D_{+,f}$ where $f \in S$ has degree 1. For such functions f we have that $(S_f)_d$ is the free $(S_f)_0$-module generated by f^d. In this way we see that $\mathcal{O}(d)$ is obtained by gluing together locally free sheaves of rank 1 on an open cover, hence it is itself a locally free sheaf of rank 1. □

Remark 9.6.12. In particular, when S is generated in degree 1 by elements $\{x_1, \ldots, x_m\}$ the proof above shows that on the open affine cover D_{+,f_i} the transition functions for the invertible sheaf $\mathcal{O}(d)$ are defined by $(\frac{x_i}{x_j})^d$.

By analogy with the situation for $\mathbb{P}^n_{\mathbb{K}}$, one might expect the global sections of $\mathcal{O}(d)$ on $\mathrm{Proj}(S)$ to be isomorphic to the dth graded piece of S. This is not true in general, even if we assume that S is generated in degree 1. The following proposition clarifies the situation:

Proposition 9.6.13. *Let S be a $\mathbb{Z}_{\geq 0}$-graded ring which is an integral domain and is finitely generated in degree 1 by elements $x_0, \ldots, x_n$. Then we have an equality of graded S-modules*

$$\bigoplus_{i \geq 0} \Gamma(\mathrm{Proj}(S), \mathcal{O}(d)) \cong \bigcap_{i=0}^{n} (S_{x_i})$$

inside the field $\mathbb{F}$ obtained by localizing S along all non-zero homogeneous elements.

The right hand side will always contain S, but need not be equal to it in general. When every x_i is prime – for example, if S is a UFD – then the right hand side is equal to S. (The RHS is reminiscent of the intersection over height 1 primes that characterizes integrally closed Noetherian domains, and indeed when S is Noetherian then every element in the RHS will be integral over S.)

Proof. We can compute the global sections of $\mathcal{O}(d)$ by applying the gluing axiom to the open cover of $\mathrm{Proj}(S)$ by D_{+,x_i}. We obtain our usual exact sequence

$$0 \to \Gamma(\mathrm{Proj}(S), \mathcal{O}(d)) \to \oplus_{i=0}^{n} (S_{x_i})_d \to \oplus_{i,j=0}^{n} (S_{x_i x_j})_d$$

Taking sums, we get

$$0 \to \bigoplus_{i \geq 0} \Gamma(\mathrm{Proj}(S), \mathcal{O}(d)) \to \oplus_{i=0}^{n}(S_{x_i}) \to \oplus_{i,j=0}^{n}(S_{x_i x_j})$$

Since all the localizations of S are contained in $\mathbb{F}$, we see that the kernel of the righthand map is the set of all elements $(s_i \in S_{x_i})$ which are identified with the same element in $\mathbb{F}$. This implies the desired statement. □

Exercise 9.6.14. Let S be a $\mathbb{Z}_{\geq 0}$-graded ring which is finitely generated in degree 1 by elements $x_0, \ldots, x_n$. Prove that there is a "canonical" map $S \to \oplus_{d \geq 0} \Gamma(\mathrm{Proj}(S), \mathcal{O}(d))$.

The true importance of the invertible sheaves $\mathcal{O}(d)$ won't become clear until Section 10.6 when we discuss ample line bundles. A first step in this direction is the geometric version of Hilbert's Syzygy Theorem:

Theorem 9.6.15 (Hilbert's Syzygy Theorem)**.** *If $\mathcal{F}$ is a coherent sheaf on $\mathbb{P}^n_{\mathbb{K}}$ then there is an exact sequence*

$$0 \to \mathcal{E}_r \to \mathcal{E}_{r-1} \to \ldots \to \mathcal{E}_0 \to \mathcal{F} \to 0$$

where each $\mathcal{E}_i$ is a direct sum of line bundles $\mathcal{E}_i = \oplus_{j=1}^{k_i} \mathcal{O}(d_{i,j})$ and where $r \leq n+1$.

In particular we can try to understand arbitrary coherent sheaves on $\mathbb{P}^n$ using finite exact sequences of sheaves we understand very well.

9.6.3 Constructing a graded module

Having defined a functor $\widetilde{}^{+} : S{-}\mathbf{GrMod} \to \mathbf{QCoh}(\mathrm{Proj}(S))$, we now define a functor $\Gamma_\bullet$ in the reverse direction. In contrast to the situation for affine schemes, we do not obtain an equivalence between the category of graded S-modules and the category of quasicoherent sheaves on $\mathrm{Proj}(S)$. Nevertheless, the resulting correspondence is quite nicely behaved.

Construction 9.6.16. Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is generated by a finite set of elements in degree 1 as an S_0-algebra. Given any quasicoherent sheaf $\mathcal{F}$ on $\mathrm{Proj}(S)$ we define

$$\Gamma_n(\mathcal{F}) := \Gamma(\mathrm{Proj}(S), \mathcal{F}(n)).$$

We also set

$$\Gamma_\bullet(\mathcal{F}) := \bigoplus_{n=-\infty}^{\infty} \Gamma_n(\mathcal{F})$$

$\Gamma_\bullet(\mathcal{F})$ has the structure of a graded $\Gamma_\bullet(\mathcal{O}_{\mathrm{Proj}(S)})$-module via the multiplication maps

$$\Gamma(\mathrm{Proj}(S), \mathcal{F}(n)) \otimes \Gamma(\mathrm{Proj}(S), \mathcal{O}(d)) \to \Gamma(\mathrm{Proj}(S), \mathcal{F}(n+d))$$

induced by tensor product. Via the map $S \to \Gamma_\bullet(\mathcal{O}_{\mathrm{Proj}(S)})$ from Exercise 9.6.14 we see that $\Gamma_\bullet(\mathcal{F})$ also has the structure of a graded S-module.

Definition 9.6.17. Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is generated by a finite set of elements in degree 1 as an S_0-algebra. We call the functor $\Gamma_\bullet$ from $\mathbf{QCoh}(\mathrm{Proj}(S))$ to the category of graded S-modules the saturation functor. We say that a graded S-module is saturated if it lies in the image of this functor.

Exercise 9.6.18. Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is generated by a finite set of elements in degree 1 as an S_0-algebra. Show that for any graded S-module N there is a canonical morphism of graded S-modules $N \to \Gamma_\bullet(\widetilde{N}^+)$. This function is called the "saturation map". (Hint: send m to the collection of elements $m/1$ in the distinguished open affines D_{+,f_i}.)

When we are working with projective $\mathbb{K}$-schemes, the saturation map can be described more explicitly (see Exercise 9.6.26). We next analyze what happens if we compose the two functors $\Gamma_\bullet$ and $\widetilde{\ }^+$ in the opposite order. We will continue to assume that S is finitely generated in degree 1. We first construct a morphism

$$\phi : \widetilde{\Gamma_\bullet(\mathcal{F})}^+ \to \mathcal{F}.$$

For any $f \in S_1$, a section of $\widetilde{\Gamma_\bullet(\mathcal{F})}^+$ along $D_{+,f}$ is a fraction m/f^d where $m \in \Gamma(X, \mathcal{F}(d))$. We can associate to m/f^d the section

$$m|_{D_{+,f}} \otimes f^{-d} \in (\mathcal{F}(d) \otimes \mathcal{O}(-d))(D_{+,f}) \cong \mathcal{F}(D_{+,f}).$$

This module-theoretic map defines the function $\phi|_{D_{+,f}}$, and by gluing we obtain the map ϕ.

Theorem 9.6.19. *Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is generated in degree 1 and finitely generated as an S_0-algebra. Let $\mathcal{F}$ be a quasicoherent sheaf on $\mathrm{Proj}(S)$. Then the map $\phi : \widetilde{\Gamma_\bullet(\mathcal{F})}^+ \to \mathcal{F}$ is an isomorphism.*

In particular, this shows that the functor $\Gamma_\bullet$ is essentially surjective onto the category of quasicoherent sheaves on $\mathrm{Proj}(S)$.

Proof. It suffices to prove that each $\phi|_{D_{+,f}}$ is an isomorphism. Exercise 9.5.28 shows that every element of $\mathcal{F}(D_{+,f})$ has the form $m \otimes f^{-d}$ for some positive integer d and some $m \in \mathcal{F}(X)$. In particular, ϕ is surjective.

Conversely, suppose that m is a section of $\Gamma(X, \mathcal{F}(d))$ such that $m \otimes f^{-d} \in \mathcal{F}(D_{+,f})$ is zero. Then Exercise 9.5.28 shows that there is some positive integer n such that $m \otimes f^{-d+n} \in \mathcal{F}(n)(D_{+,f})$ is zero. Thus $mf^n/f^{-d+n} = m/f^{-d}$ will be zero in $(\Gamma_\bullet(\mathcal{F})_f)_0$, and we conclude that ϕ is injective. $\square$

9.6.4 Adjoint pairs

In sum, given a $\mathbb{Z}_{\geq 0}$-graded ring S that is finitely generated in degree 1, we have the following functors

$$\begin{tikzcd}
\mathbf{QCoh}(\mathrm{Proj}(S)) \arrow[dr, "\Gamma_\bullet"'] & \text{graded } S\text{-modules} \arrow[l, "{}^{\sim +}"] \arrow[d, bend right, "\Gamma_\bullet({}^{\sim +})"'] \\
& \text{saturated } S\text{-modules} \arrow[u, bend right, "\text{inclusion}"']
\end{tikzcd}$$

The following results describe the relationships between these various functors.

Theorem 9.6.20. *Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is generated in degree 1 and finitely generated as an S_0-algebra. For any graded S-module N and any quasicoherent sheaf $\mathcal{F}$ on* $\mathrm{Proj}(S)$ *we have a natural bijection*

$$\mathrm{Hom}(\widetilde{N}^+, \mathcal{F}) \cong \mathrm{Hom}(N, \Gamma_\bullet(\mathcal{F})).$$

In particular, ${}^{\sim +}$ and $\Gamma_\bullet$ form an adjoint pair between the categories $\mathbf{QCoh}(\mathrm{Proj}(S))$ *and the category of graded S-modules.*

Theorem 9.6.21. *Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is generated in degree 1 and finitely generated as an S_0-algebra. Then $\Gamma_\bullet$ defines an equivalence of categories between* $\mathbf{QCoh}(\mathrm{Proj}(S))$ *and the category of saturated graded S-modules.*

Theorem 9.6.22. *Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is generated in degree 1 and finitely generated as an S_0-algebra. Then the saturation functor $\Gamma_\bullet({}^{\sim +})$ is the left adjoint of the inclusion functor as functors between the categories of graded S-modules and saturated graded S-modules.*

9.6.5 Exercises

Exercise 9.6.23. Let S be the graded ring $\mathbb{K}[x]$ where x has degree 2. Show that $\mathcal{O}_{\mathrm{Proj}(S)}(1)$ is not an invertible sheaf.

Exercise 9.6.24. Let S be the graded subalgebra of $\mathbb{K}[s,t]$ generated by s^4, s^3t, st^3, t^4. Let $X = \mathrm{Proj}(S)$. Prove that $\mathcal{O}_X(1)$ is locally free and that $\dim_{\mathbb{K}} \Gamma(X, \mathcal{O}_X(1)) = 5$. Show that $\Gamma_\bullet(\mathcal{O}_{\mathrm{Proj}(S)})$ is isomorphic to the Veronese subalgebra $\mathbb{K}[s,t]^{(4)}$.

Exercise 9.6.25. Show that if S is a Noetherian $\mathbb{Z}_{\geq 0}$-graded algebra and N is a finitely generated graded S-module then $\widetilde{N}^+$ is coherent.

Exercise 9.6.26. Suppose that S is a finitely generated graded $\mathbb{K}$-algebra with $S_0 \cong \mathbb{K}$ and that S is generated in degree 1. Let N be a finitely generated graded S-module. Prove that the canonical map $N \to N_{sat}$ is an isomorphism in all sufficiently high degrees.

Exercise 9.6.27. Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is generated in degree 1 and finitely generated as an S_0-algebra. Show that $\mathrm{Proj}(S)$ is isomorphic to $\mathrm{Proj}(\Gamma_\bullet(\widetilde{S}^+))$.

9.7 Flat morphisms

A central idea in topology and geometry is that of a fibration – loosely speaking, this is a map whose fibers look "the same". Such a condition is too stringent in algebraic geometry, as it fails even in very simple examples (such as a pencil of cubic curves in the plane). In this chapter, we will introduce a different notion of a morphism whose fibers are "varying in a nice way."

Recall that an R-module M is said to be flat if $M \otimes_R -$ is an exact functor. Flatness is a stalk-local property – that is, it can be checked after localizing at prime ideals – so it generalizes naturally to quasicoherent sheaves. We will need one important result about flat ring extensions:

Theorem 9.7.1 (Going Down)**.** *Let $S \subset R$ be a flat ring extension. Suppose that for positive integers $1 \leq m < n$ we have a chain of prime ideals $\mathfrak{q}_1 \supsetneq \mathfrak{q}_2 \supsetneq \ldots \supsetneq \mathfrak{q}_n$ in S and a chain of prime ideals $\mathfrak{p}_1 \supsetneq \ldots \supsetneq \mathfrak{p}_m$ such that $\mathfrak{p}_i \cap S = \mathfrak{q}_i$. Then we can find a prime $\mathfrak{p}_{m+1} \subsetneq \mathfrak{p}_m$ such that $\mathfrak{p}_{m+1} \cap S = \mathfrak{q}_{m+1}$.*

9.7.1 Flat morphisms

Definition 9.7.2. Let $f : X \to Y$ be a morphism of schemes. We say that a $\mathcal{O}_X$-module $\mathcal{F}$ is flat over Y at a point $x \in X$ if $\mathcal{F}_x$ is a flat $\mathcal{O}_{Y,f(x)}$-module. We say that $\mathcal{F}$ is flat over Y if it is flat at every point.

A morphism $f : X \to Y$ is flat if $\mathcal{O}_X$ is flat over Y.

Example 9.7.3. A finitely generated module over a local ring is flat if and only if it is free. Thus, a coherent sheaf $\mathcal{F}$ on X is flat (with respect to the identity map) if and only if it is locally free.

This definition was introduced by Serre, who noticed its usefulness in algebraic geometry. It is perhaps surprising that this algebraic definition has important geometric ramifications. This section is dedicated to the geometry of flat morphisms; later on we will see how such maps also have nice cohomological properties.

Proposition 9.7.4. *Let $f : X \to Y$ be a morphism, $\mathcal{F}$ a quasicoherent sheaf on X. Then $\mathcal{F}$ is flat over Y if and only if for every open affine $V \subset Y$ and every open affine $U \subset f^{-1}V$ we have that $\mathcal{F}(U)$ is a flat $\mathcal{O}_Y(V)$-module.*

In particular, this means that the map $\mathrm{Spec}(S) \to \mathrm{Spec}(R)$ is flat when S is a flat R-algebra – for example, when S is a localization of R, a free R-module, or (when R is Noetherian) a completion along an ideal.

Proof. This reduces to an algebraic statement: M is a flat S-module if and only if $M_{\mathfrak{p}}$ is a flat $S_{\mathfrak{p}}$-module for every prime ideal $\mathfrak{p} \subset S$. $\square$

Theorem 9.7.5. *(1) Let $f : X \to Y$ be a morphism and suppose that $\mathcal{F}$ is an $\mathcal{O}_X$-module that is flat over Y. For any morphism $g : Z \to Y$, the pullback of $\mathcal{F}$ to $X \times_Y Z$ is flat over Z.*

(2) Suppose that $f : X \to Y$ and $g : Y \to Z$ are morphisms of schemes. If $\mathcal{F}$ is an $\mathcal{O}_X$-module that is flat over Y and g is a flat morphism then $\mathcal{F}$ is flat over Z.

Proof. (1) Since flatness can be verified locally, this follows from the algebraic fact that if M is a flat S-module and $S \to R$ is any homomorphism then $M \otimes_S R$ is a flat R module.

(2) Since flatness can be verified locally, this follows from the algebraic fact that if M is a flat S-module and S is a flat R-algebra then M is a flat R-module. □

As an immediate corollary, we see:

Corollary 9.7.6. *Flat morphisms are preserved under composition, stable under base change, and local on the target.*

A key property of flat morphisms (that follows immediately from Proposition 9.7.4) is:

Corollary 9.7.7. *Let $f : X \to Y$ be a flat morphism. For any exact sequence of quasicoherent sheaves on Y, the pullback under f is an exact sequence of quasicoherent sheaves on X.*

9.7.2 Flatness and open sets

We next show that flatness interacts well with open sets.

Exercise 9.7.8. Prove that open embeddings are flat.

Proposition 9.7.9. *Let $f : X \to Y$ be a flat morphism of locally finite presentation. Then f is topologically open.*

In fact, since both properties of f are preserved by base change, f is even "universally open" in the sense that every base change of f is still topologically open.

Just as the closedness of a finite map follows from the Going Up theorem, the openness of a flat map follows from the Going Down theorem. Note that the statement is *not* true if we remove the finite presentation hypothesis: consider the inclusion of the generic point $\mathrm{Spec}(\mathbb{K}(x)) \to \mathrm{Spec}(\mathbb{K}[x])$ or even the inclusion $\mathrm{Spec}(\widehat{\mathbb{K}[x]_{(x)}}) \to \mathrm{Spec}(\mathbb{K}[x])$. Nevertheless even in these examples flat morphisms retain information about a "dense" subset.

Proof. We will only prove the statement when X, Y are Noetherian schemes. By Exercise 9.7.8, it suffices to prove that $f(X)$ is an open subset of Y. By Chevalley's Theorem $f(X)$ is a constructible subset. Thus it suffices to show that if y_1, y_2 are points of Y such that $y_1 \in \overline{y_2}$ and we have $y_1 \in f(X)$ then $y_2 \in f(X)$ as well.

Note that $y_1 \in \overline{y_2}$ if and only if y_2 is in the image of the natural map $\mathrm{Spec}(\mathcal{O}_{Y,y_1}) \to Y$. Thus it suffices to show that for any point $x \in X$ the induced morphism $\mathrm{Spec}(\mathcal{O}_{X,x}) \to \mathrm{Spec}(\mathcal{O}_{Y,f(x)})$ is surjective.

By assumption the morphism $f_x^\sharp : \mathcal{O}_{Y,f(x)} \to \mathcal{O}_{X,x}$ is a flat local homomorphism. A standard argument in commutative algebra implies that it is faithfully flat, hence injective. Indeed, suppose that K is the kernel of $f_x^\sharp$. Since $\mathcal{O}_{X,x}$ is flat, if we tensor by it we get the exact sequence

$$K \otimes_{\mathcal{O}_{Y,f(x)}} \mathcal{O}_{X,x} \to \mathcal{O}_{X,x} \to \mathcal{O}_{X,x} \otimes_{\mathcal{O}_{Y,f(x)}} \mathcal{O}_{X,x}$$

where the map on the right send $f \mapsto 1 \otimes f$. In particular the map on the right is injective, since if we compose it with the homomorphism $g_1 \otimes g_2 \to g_1 g_2$ we get the identity map. This means that $K \otimes_{\mathcal{O}_{Y,f(x)}} \mathcal{O}_{X,x}$ is the zero module. If K were non-zero, then by Nakayama's Lemma $K/\mathfrak{m}_{f(x)}K$ would also be non-zero. Again appealing to flatness, the quotient map would lead to a surjection

$$\begin{aligned} K \otimes_{\mathcal{O}_{Y,f(x)}} \mathcal{O}_{X,x} &\to K/\mathfrak{m}_{f(x)}K \otimes_{\mathcal{O}_{Y,f(x)}} \mathcal{O}_{X,x} \\ &\cong K/\mathfrak{m}_{f(x)}K \otimes_{\mathcal{O}_{Y,f(x)}/\mathfrak{m}_{f(x)}} \mathcal{O}_{X,x}/\mathfrak{m}_x \end{aligned}$$

However, the RHS is clearly non-zero, since we are tensoring a non-zero vector space by a field extension. This contradicts the fact that the LHS should vanish. We conclude that K is zero and that $f_x^\sharp$ is an inclusion. We can thus apply the Going Down Theorem (Theorem 9.7.1) to see that the induced map $\mathrm{Spec}(\mathcal{O}_{X,x}) \to \mathrm{Spec}(\mathcal{O}_{Y,f(x)})$ is surjective. □

There is another key result relating flatness to open sets:

Theorem 9.7.10. *Let $f : X \to Y$ be a finite type morphism such that Y is integral. Suppose that $\mathcal{F}$ is a finitely generated quasicoherent sheaf on X. Then there is a non-empty open set $V \subset Y$ such that $\mathcal{F}$ is flat over V.*

Note that we are allowing the preimage of V to be empty – for example, if f is a closed embedding then we should take U to be the complement of X. One can also weaken the hypotheses a little bit; see [Sta15, Tag 052B].

Proof. We will just prove the statement under the additional hypothesis that Y (and hence X) is Noetherian. It suffices to prove the statement when $Y = \mathrm{Spec}(S)$ is affine. Note that since Y is integral S is a domain.

Since f is quasicompact, we can choose a finite cover of X by open affines $U_i = \mathrm{Spec}(R_i)$ for $i = 1, \ldots, r$. For each index i, we apply Grothendieck's Generic Freeness (Theorem 9.7.11) to find an element $g_i \in S$. Then $\mathcal{F}$ will be flat over the open set $D_{g_1 g_2 \ldots g_r}$ of S. □

The proof relied on the following algebraic fact:

Theorem 9.7.11 (Grothendieck's Generic Freeness)**.** *Suppose that B is a finitely generated algebra over a Noetherian domain A. For any finitely generated B-module M, there is an element $g \in A$ such that M_g is a free A_g-module.*

There is a companion fact which we will not prove:

Theorem 9.7.12. *Let $f : X \to Y$ be a finite type morphism of Noetherian schemes and let $\mathcal{F}$ be a finitely generated quasicoherent sheaf on X. Then the set of points $x \in X$ such that $\mathcal{F}$ is flat over Y at x is an open subset.*

9.7.3 Flatness criteria

When R is a Dedekind domain, we can classify the flat R-modules in an explicit way: an R-module M is flat if and only if it is torsion free. (This follows from the localization criterion for flat modules and the structure theorem for modules over a PID.) The following result is a scheme-theoretic version.

Theorem 9.7.13. *Let $f : X \to Y$ be a morphism of Noetherian schemes such that Y is integral, regular, and has dimension 1. Then f is flat if and only if every associated point of X maps to the generic point of Y.*

The forward implication is always true (see Exercise 9.7.23). However, the reverse implication relies upon the fact that the base has dimension 1.

Proof. Let $V \subset Y$ be an open affine and let $U \subset f^{-1}V$ be an open affine. Since $\mathcal{O}_Y(V)$ is a Dedekind domain, we see that $\mathcal{O}_X(U)$ will be a flat $\mathcal{O}_Y(V)$-module if and only if $\mathcal{O}_X(U)$ is torsion free as a $\mathcal{O}_Y(V)$-module.

First suppose that every associated point of X maps to the generic point of Y. Note that an element $f \in \mathcal{O}_X(U)$ will be a zero divisor if and only if there is a prime ideal $\mathfrak{p}$ representing an associated point of U such that $V(f) \supset V(\mathfrak{p})$. In particular, for any element $g \in \mathcal{O}_Y(V)$ we have that $V(f^\sharp(V)(g)) = f^{-1}(V(g))$ and thus by assumption the vanishing locus cannot contain any associated point of X. Thus no element of $\mathcal{O}_Y(V)$ maps to a zero divisor in $\mathcal{O}_X(U)$. Varying U and V, we deduce that f is flat.

Conversely, suppose that f is flat. Since $g \in \mathcal{O}_Y(V)$ cannot map to a zero divisor in $\mathcal{O}_X(U)$, it is impossible for $V(g)$ to contain the f-image of an associated point of X. Since this is true for every $g \in \mathcal{O}_Y(V)$, the f-image of an associated point of X must be the generic point. $\square$

Remark 9.7.14. An important consequence is the ability to take "flat limits". Suppose we have an integral curve C parametrizing a flat family of closed subschemes of $\mathbb{P}^n$. In other words, suppose we have a closed subscheme $\mathcal{U} \subset C \times \mathbb{P}^n$ such that the projection map $p_1 : \mathcal{U} \to C$ is flat. If we take a projective closure $\overline{C}$, we can extend the family over $\overline{C}$ simply by taking the closure of $\mathcal{U}$ in $\overline{C} \times \mathbb{P}^n$. Since taking closures does not affect the associated points, the new family $\overline{\mathcal{U}}$ will still be flat over $\overline{C}$.

Another situation in which flatness admits a nice description is when we are working with finite morphisms.

Exercise 9.7.15. Suppose that $f : X \to Y$ is a finite morphism and $\mathcal{F}$ is a coherent sheaf on X. Show that $\mathcal{F}$ is flat over Y iff $f_*\mathcal{F}$ is locally free.

As we saw in Theorem 4.2.10, flatness is the key ingredient which allows us to identify the degree of f with the degree of the fibers of f. Indeed, if f is a finite flat morphism then the proof of Theorem 4.2.10 shows that both of these numbers will coincide with the rank of $f_*\mathcal{O}_X$. The following useful theorem identifies examples of finite morphisms which are guaranteed to be flat.

Theorem 9.7.16 (Miracle Flatness, [Sta15] Tag 00R3)**.** *Let $f : X \to Y$ be a surjective finite morphism. If X is Cohen-Macaulay and Y is regular then f is flat.*

In particular, a finite morphism between regular schemes is always flat. In fact, the statement of Miracle Flatness still holds true if we only know that the fibers of f have constant dimension.

9.7.4 Fibers of flat families

Finally, we turn to the key property of flatness: the fibers of a flat family will "vary nicely". We will only prove one statement in this direction.

Theorem 9.7.17. *Let $f : X \to Y$ be a flat morphism of irreducible $\mathbb{K}$-schemes. Then every non-empty fiber of f has dimension* $\dim(X) - \dim(Y)$.

There is a similar statement for arbitrary Noetherian schemes, except that (as usual) one needs to compare codimensions instead.

Proof. By Proposition 9.7.9 $f(X)$ is an open subset of Y. Since flatness is preserved by base change, we may replace Y by the reduced scheme underlying the open set $f(X)$ and X by the base change over this set. Since the desired statement is local on the target we may also assume that Y is affine. In this case for every open affine $U \subset X$ the map $f^\sharp : \mathcal{O}_Y(Y) \to \mathcal{O}_X(U)$ is injective

Fix a point $x \in X$. Set $y = f(x)$ and let F be an irreducible component of the fiber over y that contains x. Let U be an open affine neighborhood of x. If $\mathfrak{p}_0 \subset \mathcal{O}_Y(Y)$ corresponds to the point y, there is a chain of prime ideals

$$\mathfrak{p}_0 \supsetneq \mathfrak{p}_1 \supsetneq \cdots \supsetneq \mathfrak{p}_{\dim Y}$$

Since $\mathcal{O}_X(U)$ is flat over $\mathcal{O}_Y(Y)$, the Going Down theorem for flat extensions shows that we have a chain

$$\mathfrak{q}_0 \supsetneq \mathfrak{q}_1 \supsetneq \cdots \supsetneq \mathfrak{q}_{\dim Y}$$

where $\mathfrak{q}_0$ is the prime ideal defining $F \cap U$ and where $(f^\sharp)^{-1}(\mathfrak{q}_i) = \mathfrak{p}_i$. We can extend this chain on the left using prime ideals contained in $\mathfrak{q}_0$ so that the total length is (one more than) $\dim(F) + \dim(Y)$. This shows that $\dim(X) \geq \dim(F) + \dim(Y)$. Since the reverse inequality was proved in Theorem 4.4.9 (and Remark 4.4.10), this finishes the proof. $\square$

We finish with one more statement showing how the fibers of a flat morphism "vary nicely." Suppose that Y is an integral scheme and let X be a closed subscheme of $Y \times \mathbb{P}^n$. For any point $y \in Y$, the fiber of f over y is a closed subscheme of $\mathbb{P}^n_{\kappa(y)}$. We can thus associate to any point of y the Hilbert polynomial of the fiber over y.

Theorem 9.7.18. *Let Y be an integral scheme and let X be a closed subscheme of $Y \times \mathbb{P}^n$. Then the projection map $f : X \to Y$ is flat if and only if the Hilbert polynomial of the fibers of f is constant over the points of Y.*

Thus important result shows that the Hilbert polynomial is exactly the right tool for understanding flat families of projective schemes. We will prove a more general statement in Theorem 12.7.7.

9.7.5 Exercises

Exercise 9.7.19. Prove that the following morphisms are not flat using the definition directly:

(1) The map $\operatorname{Spec}(\mathbb{K}[x, y]/(xy)) \to \operatorname{Spec}(\mathbb{K}[x])$.

(2) The map $\operatorname{Spec}(\mathbb{K}[x, y]/(y^2, xy)) \to \operatorname{Spec}(\mathbb{K}[x])$.

(3) The blow-up of a point $\phi : X \to \mathbb{P}^2$.

How do these morphisms violate the properties of flat morphisms described earlier?

Exercise 9.7.20. Prove that the normalization of the cuspidal curve $\operatorname{Spec}(\mathbb{K}[x, y]/(y^2 - x^3))$ is an open morphism that is not flat. (More generally, explain why the normalization of a non-normal $\mathbb{K}$-variety will never be flat.)

Exercise 9.7.21. Let $f : X \to Y$ be a flat morphism. Show that every locally free sheaf on X is flat over Y.

Exercise 9.7.22. Prove that a morphism $f : X \to Y$ is flat if and only if $f^* : \mathbf{QCoh}(X) \to \mathbf{QCoh}(Y)$ is exact.

Exercise 9.7.23. Let $f : X \to Y$ be a flat morphism of Noetherian schemes. Show that the f-image of an associated point of X is an associated point of Y. (Hint: use the fact that the maximal ideal in a Noetherian local ring R is an associated prime of R if and only if every element contained in it is a zero-divisor.)

More generally, if $\mathcal{F}$ is a coherent sheaf on X that is flat over Y then every associated point of $\mathcal{F}$ maps to an associated point of Y.

Exercise 9.7.24. Let $U = \mathbb{A}^1 \backslash \{0\}$. Consider the family of subschemes $W \subset \mathbb{A}^3 \times U$ where the fiber of W over the point $p \in U$ is the union of the lines $y = z = 0$ and $x - p = z = 0$. Note that the map $\pi : W \to U$ is flat.

By Remark 9.7.14 there is a subscheme $\overline{W} \subset \mathbb{A}^3 \times \mathbb{A}^1$ which is flat over $\mathbb{A}^1$ and agrees with W over U. Compute the flat limit of this family: what is the fiber of $\overline{W}$ over 0? (Hint: it is not the reduced union of the x and y axes.)

Exercise 9.7.25. Consider the double plane $X = V((w, x) \cup (y, z))$ in $\mathbb{A}^4$. Consider the map $\mathbb{A}^4 \to \mathbb{A}^2$ given by $s \mapsto w + y, t \mapsto x + z$ and let $f : X \to \mathbb{A}^2$ be the restriction to X. Show that f is finite but not flat. (Hint: compute the degrees of f above various points.) Using Theorem 9.7.16 deduce that X is not Cohen-Macaulay.

Exercise 9.7.26. A morphism $f : X \to Y$ is said to be faithfully flat if it is flat and surjective.

(1) Suppose that X and Y are affine schemes. Prove that f is faithfully flat if and only if $f^\sharp$ realizes $\mathcal{O}_X(X)$ as a faithfully flat $\mathcal{O}_Y(Y)$-algebra.

(2) Prove that faithfully flat morphisms are preserved under composition, stable under base change, and local on the target.

(3) Prove that if $f : X \to Y$ is a faithfully flat morphism then a sequence of quasicoherent sheaves on Y is exact if and only if its pullback to X is exact.

Chapter 10

Line bundles

Line bundles play an important role in every area of geometry. This is particularly true in algebraic geometry. Indeed, the fundamental example of a scheme is projective space and the key geometric feature of projective space is the tautological line bundle. Thus it should come as no surprise that the tautological line bundle and its dual $\mathcal{O}(1)$ feature prominently in the study of projective schemes.

Suppose that X is a scheme. Given a morphism $f : X \to \mathbb{P}^n$, we can obtain an invertible sheaf $\mathcal{L}$ on X by pulling back $\mathcal{O}_{\mathbb{P}^n}(1)$. The invertible sheaf $\mathcal{L}$ has one key property: since the global sections of $\mathcal{O}(1)$ don't simultaneously vanish at any point, the global sections of $\mathcal{L}$ also do not vanish at any point.

In the first section of the chapter, we show that this construction can be reversed: starting from an invertible sheaf $\mathcal{L}$ and a set of global sections, we can construct a (rational) map to projective space. (The construction of a rational map from a graded homomorphism of homogeneous coordinate rings is a special case of this construction.) This result provides a fundamental shift in perspective: we can analyze the category of quasiprojective $\mathbb{K}$-schemes by instead studying invertible sheaves and their sections.

We next discuss how to classify invertible sheaves on a scheme X. It turns out that the study of invertible sheaves is closely tied to the study of codimension 1 subvarieties. This correspondence leads us to the study of Weil divisors (formal sums of codimension 1 integral subschemes) and of Cartier divisors (Weil divisors locally defined by a single equation). After introducing these new concepts, we explain how in some cases they can be used to classify invertible sheaves.

Finally, we turn to the question of which invertible sheaves on X can be the pullback of $\mathcal{O}(1)$ under a morphism $f : X \to \mathbb{P}^1$. In fact, we ask for something even stronger: we are mainly interested in the case when f is a closed embedding. This leads to the notion of a very ample invertible sheaf (which is the pullback of $\mathcal{O}(1)$ under a closed embedding) and its "stable" analogue, an ample invertible sheaf. In the last two sections we discuss these essential notions, first in the "absolute" setting, then in the "relative" setting.

10.0.1 Algebraic preliminaries

Let X be a Noetherian scheme and let Y be an irreducible closed subset. As discussed in Section 8.5.3, we define the codimension of Y to be the maximal integer r such that we have an increasing chain of irreducible closed subsets

$$Y = X_0 \subsetneq X_1 \subsetneq \ldots \subsetneq X_r = X.$$

While we have an inequality $\dim(X) \geq \dim(Y) + \operatorname{codim}_X(Y)$, in general we do not obtain equality.

The following theorems summarize the properties of codimension for Noetherian schemes:

Theorem 10.0.1 (Krull's Prinicipal Ideal Theorem)**.** *Let X be an irreducible Noetherian scheme. Let $U \subset X$ be an open subset and suppose that $f \in \mathcal{O}_X(U)$. Then $V(f)$ is a closed subset of U of codimension ≤ 1.*

Theorem 10.0.2 (Dimension of fibers)**.** *Let $f : X \to Y$ be a morphism of Noetherian schemes. Suppose that p is a point of X and $q = f(p)$. Then we have*

$$\operatorname{codim}_X(\overline{p}) \leq \operatorname{codim}_Y(\overline{q}) + \operatorname{codim}_{\overline{f^{-1}(q)}}(\overline{p}).$$

10.1 Invertible sheaves and maps to projective space

An invertible sheaf (or equivalently, a line bundle) on a scheme X is a locally free sheaf $\mathcal{L}$ of rank 1. Let's recall some of the basic properties of invertible sheaves:

(1) Given any morphism $f : X \to Y$, the pullback of an invertible sheaf on Y is an invertible sheaf on X. (See Exercise 9.5.20.)

(2) For any invertible sheaf $\mathcal{L}$ on X, there is an open cover $\{U_i\}$ of X such that $\mathcal{L}|_{U_i} \cong \mathcal{O}_{U_i}$. Conversely, given an open cover $\{U_i\}$ of X and transition maps $\phi_{ij} : \mathcal{O}_{U_i \cap U_j} \to \mathcal{O}_{U_i \cap U_j}$ that satisfy the cocycle condition we can construct an invertible sheaf $\mathcal{L}$. (See Section 9.5.2.)

(3) Given two invertible sheaves $\mathcal{L}, \mathcal{M}$ their tensor product $\mathcal{L} \otimes \mathcal{M}$ is also an invertible sheaf.

We will need a couple more basic results about invertible sheaves.

Proposition 10.1.1. *Let X be a scheme and let $\mathcal{L}$ be an invertible sheaf on X. Then $\mathcal{L}^\vee \otimes \mathcal{L} \cong \mathcal{O}_X$.*

In fact this result motivates the term "invertible sheaf": it turns out that the invertible sheaves are exactly the same as the $\mathcal{O}_X$-modules which admit an "inverse" under the tensor product.

Proof. Let R be a ring and let M be a rank 1 free R-module. There is a canonical isomorphism $\phi_R : M^\vee \otimes_R M \to R$ given by sending $\psi \otimes m \mapsto \psi(m)$. Note that this map is compatible with localization, in the sense that for any element $f \in R$ the map ϕ_{R_f} is $(\phi_R)_f$.

Choose an open cover of X by open affines U_i such that $\mathcal{L}|_{U_i}$ is isomorphic to $\mathcal{O}_{U_i}$. The map $\phi_{\mathcal{O}_X(U_i)}$ defines an isomorphism $\phi_i : (\mathcal{L}^\vee \otimes \mathcal{L})|_{U_i} \to \mathcal{O}_X|_{U_i}$. By Corollary 7.6.4, these local ϕ_i glue together to give a global morphism $\phi : \mathcal{L}^\vee \otimes \mathcal{L} \to \mathcal{O}_X$. Since ϕ induces isomorphisms of stalks at every point $x \in X$, it is an isomorphism. □

Recall that for any global section s of an invertible sheaf $\mathcal{L}$ we can define a closed subscheme of X known as the vanishing locus of s; its ideal is the image of the map $\mathcal{L}^\vee \to \mathcal{O}_X$ corresponding to s.

Proposition 10.1.2. *Let X be a scheme. Let $\mathcal{L}$ be an invertible sheaf and let $s \in \mathcal{L}(X)$. Denote by U the complement of the zero locus $Z(s)$. Then there is an isomorphism $\psi : \mathcal{L}|_U \to \mathcal{O}_U$ such that $\psi(U)(s) = 1$.*

Proof. Let $V \subset U$ be any open affine such that there is an isomorphism $\phi_V : \mathcal{L}|_V \to \mathcal{O}_V$. Define $t = \phi_V(V)(s)$. By construction t is an element of $\mathcal{O}_X(V)$ that does not vanish anywhere on V; by Exercise 8.2.18 we see that t is a unit in $\mathcal{O}_X(V)$. Thus we can define

$\mu_V : \mathcal{O}_V \to \mathcal{O}_V$ by multiplication by t^{-1}. Finally, we set ψ_V to be the composition $\mu_V \circ \phi_V : \mathcal{L}|_V \to \mathcal{O}_V$.

If V' is an open affine contained in V, it is clear that $\psi_V|_{V'} = \psi_{V'}$. By Corollary 7.6.4 the various ψ_V glue together to give a global map $\psi : \mathcal{L}|_U \to \mathcal{O}_U$. This is an isomorphism since it defines an isomorphism of stalks at every point. Using the condition $\psi_V(V)(s) = 1$ and gluing we see that $\psi(U)(s) = 1$. □

10.1.1 Picard group

Definition 10.1.3. Let X be a scheme. The Picard group $\mathrm{Pic}(X)$ is the set of invertible sheaves on X up to isomorphism. $\mathrm{Pic}(X)$ is a group under $\otimes$: the identity is $\mathcal{O}_X$ and the inverse of $\mathcal{L}$ is $\mathcal{L}^\vee$.

Note that if we have a morphism $f : X \to Y$ then the pullback defines a homomorphism $f^* : \mathrm{Pic}(Y) \to \mathrm{Pic}(X)$.

Example 10.1.4. We claim that every invertible sheaf on $\mathbb{P}^1_{\mathbb{K}}$ is isomorphic to $\mathcal{O}(d)$ for some integer d. Using the classification of finitely generated modules over a PID, we see that the restriction of any invertible sheaf to the affine charts $D_{+,x}, D_{+,y}$ will be isomorphic to the structure sheaf. Thus we only need to specify the gluing data: every invertible sheaf on $\mathbb{P}^1_{\mathbb{K}}$ is determined by an isomorphism $\psi : \mathcal{O}_U \to \mathcal{O}_U$ where $U = \mathrm{Spec}(\mathbb{K}[t, t^{-1}])$. Since the units of $\mathbb{K}[t, t^{-1}]$ all have the form ct^d for some $c \in \mathbb{K}^\times$ and some integer d, every isomorphism of this ring has the form $t \mapsto ct^d$. The resulting invertible sheaf is isomorphic to $\mathcal{O}(d)$.

Using the relation $\mathcal{O}(d) \otimes \mathcal{O}(e) \cong \mathcal{O}(d + e)$ we see that $\mathrm{Pic}(\mathbb{P}^1_{\mathbb{K}}) \cong \mathbb{Z}$.

Exercise 10.1.5. Let K be a number field and let $\mathcal{O}_K$ be the ring of integers in K. Show that the Picard group of $\mathrm{Spec}(\mathcal{O}_K)$ is isomorphic to the class group of $\mathcal{O}_K$.

10.1.2 Globally generated sheaves

We next introduce the notion of a globally generated sheaf.

Definition 10.1.6. Let X be a scheme and let $\mathcal{F}$ be an $\mathcal{O}_X$-module. We say that $\mathcal{F}$ is globally generated if there is a surjection $\mathcal{O}_X^{\oplus I} \to \mathcal{F}$ for some index set I.

As indicated by the name, global generation should be thought of as a "global" condition on the sheaf $\mathcal{F}$ and not a "local" one.

Exercise 10.1.7. Show that every quasicoherent sheaf on an affine scheme is globally generated.

Example 10.1.8. The sheaf $\mathcal{O}(-1)$ on $\mathbb{P}^n$ is not globally generated. Indeed, since $\mathcal{O}(-1)$ has no global sections, it admits no non-zero morphisms from $\mathcal{O}_{\mathbb{P}^n}$.

The following definition identifies

Definition 10.1.9. Given a point $x \in X$, we say that $\mathcal{F}$ is globally generated at x if there is a map $\mathcal{O}_X^{\oplus I} \to \mathcal{F}$ such that the induced map on stalks at x is surjective.

Warning 10.1.10. To say that $\mathcal{F}$ is globally generated at x does not mean that there is a surjection $\mathcal{O}_{X,x}^{\oplus I} \to \mathcal{F}_x$. (Indeed, we basically always have such a surjection for any $\mathcal{F}$.) We need a *global* map that induces the surjection.

Remark 10.1.11. An $\mathcal{O}_X$-module $\mathcal{F}$ is globally generated if and only if it is globally generated at every point $x \in X$. Indeed, since the indexing set for our direct sum is arbitrary we can define a surjection $\mathcal{O}_X^{\oplus I} \to \mathcal{F}$ by taking the union of the indexing sets for all the points in X.

Often Definition 10.1.6 is rephrased using the correspondence between the space of morphisms $\mathcal{O}_X \to \mathcal{F}$ and the space $\mathcal{F}(X)$ of global sections. For example, $\mathcal{F}$ is globally generated at x if and only if $\mathcal{F}_x$ is generated as an $\mathcal{O}_{X,x}$-module by the restrictions of global sections. In particular:

Definition 10.1.12. Let X be a scheme and let $\mathcal{L}$ be an invertible sheaf on X. Fix global sections $s_1, \ldots, s_n \in \mathcal{L}(X)$. We say that these sections generate $\mathcal{L}$ if the induced map $\oplus_{i=1}^n \mathcal{O}_X \to \mathcal{L}$ is surjective.

Exercise 10.1.13. Let X be a quasicompact scheme and let $\mathcal{F}$ be a finitely generated quasicoherent sheaf. Prove that $\mathcal{F}$ is globally generated if and only if there is a finite number of global sections which define a surjection $\mathcal{O}^{\oplus r} \to \mathcal{F}$.

A key feature of global generation is that it is an open property for finitely generated quasicoherent sheaves:

Exercise 10.1.14. Let X be a scheme and let $\mathcal{F}$ be a finitely generated quasicoherent sheaf on X. Fix a point $x \in X$. Suppose that the fiber $\mathcal{F}(x) := \mathcal{F}_x/\mathfrak{m}_x\mathcal{F}_x$ is spanned (as a $\kappa(x)$-vector space) by the restrictions of global sections of $\mathcal{F}$. Prove that there is an open neighborhood U of x such that $\mathcal{F}$ is globally generated at every point of U.

We will exclusively be interested in the globally generated property when we have a finitely generated quasicoherent sheaf $\mathcal{F}$ and we are looking at a finite number of global sections. In fact, for the rest of this chapter we focus on invertible sheaves $\mathcal{L}$. For invertible sheaves we can measure the failure of the global generation condition using a closed subscheme known as the base locus.

Definition 10.1.15. Let X be a scheme and let $\mathcal{L}$ be an invertible sheaf on X. Suppose we fix a finite set $\{s_i\}_{i=1}^r$ of global sections of $\mathcal{L}$. Consider the evaluation map

$$\mathcal{O}_X^{\oplus r} \to \mathcal{L}.$$

After tensoring by $\mathcal{L}^\vee$, we obtain a map $(\mathcal{L}^\vee)^{\oplus r} \to \mathcal{O}_X$. The image of this map is a quasicoherent ideal sheaf. The corresponding closed subscheme of X is known as the base locus of $\mathcal{L}$.

In other words, the base locus is the scheme-theoretic intersection of the zero loci $Z(s_i)$. Note that $\mathcal{L}$ is globally generated by the s_i if and only if the corresponding base locus is the empty set.

10.1.3 Maps to projective space

The following fundamental theorem describes the universal property of projective space.

Theorem 10.1.16. *Fix a ring R and let X be a $\mathrm{Spec}(R)$-scheme. There is a bijection between the set of $\mathrm{Spec}(R)$-morphisms $f : X \to \mathbb{P}^n_R$ and the isomorphism classes of tuples*

$$(\mathcal{L}, s_0, \ldots, s_n)$$

where $\mathcal{L}$ is a invertible sheaf on X and the $\{s_i\}_{i=0}^n$ are a set of global sections which generate $\mathcal{L}$.

Here we say that two tuples $(\mathcal{L}, s_0, \ldots, s_n)$ and $(\mathcal{L}', t_0, \ldots, t_n)$ are isomorphic if there is an isomorphism $\phi : \mathcal{L} \to \mathcal{L}'$ such that $\phi(s_i) = t_i$. In brief, Theorem 10.1.16 is saying that we can define a map $X \to \mathbb{P}^n_R$ via the equations $(s_0 : s_1 : \ldots : s_n)$.

Proof. First suppose given an R-morphism $f : X \to \mathbb{P}^n_R$. Define $\mathcal{L} = f^*\mathcal{O}_{\mathbb{P}^n_R}(1)$. The surjection $\mathcal{O}_{\mathbb{P}^n_R}^{\oplus n+1} \to \mathcal{O}_{\mathbb{P}^n_R}(1)$ defined by the tuple of sections $x_0, x_1, \ldots, x_n$ of $\mathcal{O}_{\mathbb{P}^n_R}(1)$ pulls back to give a surjection $\mathcal{O}_X^{n+1} \to \mathcal{L}$. We let s_i be the section of $\mathcal{L}$ defined by the ith direct summand.

Conversely, suppose given a line bundle $\mathcal{L}$ with sections $\{s_i\}$. Let U_i be the complement of $Z(s_i)$; since the s_i generate $\mathcal{L}$ we see that the U_i give an open cover of X. By Proposition 10.1.2 we have isomorphisms $\psi_i : \mathcal{L}|_{U_i} \to \mathcal{O}_{U_i}$ with $\psi_i(U_i)(s_i) = 1$. We define $f_i : U_i \to D_{+,x_i} \subset \mathbb{P}^n_R$ by the ring map

$$\begin{aligned} R\left[\frac{x_0}{x_i}, \ldots, \frac{x_n}{x_i}\right] &\to \mathcal{O}_X(U_i) \\ \frac{x_j}{x_i} &\mapsto \psi_i(U_i)(s_j|_{U_i}) \end{aligned}$$

It is clear that the f_i agree on overlaps and thus define a global morphism $f : X \to \mathbb{P}^n_R$.

These two constructions are inverse to each other. $\square$

The main geometric content is the choice of line bundle $\mathcal{L}$ and the R-submodule of $\mathcal{L}(X)$ generated by the s_i. Let's briefly discuss how the map f changes as we vary the choice of sections:

- Suppose that $(\mathcal{L}, s_0, \ldots, s_n)$ and $(\mathcal{L}, s_0', \ldots, s_n')$ are two tuples such that the $\mathcal{O}_X(X)$-submodule of $\mathcal{L}(X)$ generated by the s_i' coincides with the submodule generated by the s_i. Then the two maps $f, f' : X \to \mathbb{P}^n_R$ vary by an automorphism of $\mathbb{P}^n_R$ induced by an element of $GL_{n+1}(R)$.

- Suppose that we add in a section s_{n+1} that is contained in the R-submodule generated by the sections $s_0, \ldots, s_n$ so that $s_{n+1} = \sum_{i=0}^n r_i s_i$. Then the image of the new map $X \to \mathbb{P}^{n+1}_R$ is the composition of the original map $X \to \mathbb{P}^n_R$ with the inclusion of $\mathbb{P}^n_R \to \mathbb{P}^{n+1}_R$ as a hyperplane defined by the equation $x_{n+1} = \sum_{i=0}^n r_i x_i$.

  Conversely, if we can remove the section s_n without changing the corresponding R-submodule of $\mathcal{L}(X)$, the resulting map $X \to \mathbb{P}^{n-1}_R$ is the composition of the original map $X \to \mathbb{P}^n_R$ with projection away from the locus $x_0 = \ldots = x_{n-1} = 0$.

As an immediate consequence of Theorem 10.1.16, we obtain:

Corollary 10.1.17. *Let X be a* $\mathrm{Spec}(R)$*-scheme, let $\mathcal{L}$ be an invertible sheaf on X, and let $s_0, \ldots, s_n$ be global sections of $\mathcal{L}$. Let U denote the complement of the base locus of $\{s_i\}_{i=0}^n$. This collection of data induces a morphism $U \to \mathbb{P}^n_R$ over* $\mathrm{Spec}(R)$.

Theorem 10.1.16 is particularly important when we are working in the category of quasiprojective $\mathbb{K}$-schemes. Given a quasiprojective $\mathbb{K}$-scheme X, any morphism $f : X \to Y$ can be composed with an injection $Y \hookrightarrow \mathbb{P}^n_{\mathbb{K}}$ to yield a morphism $\widehat{f} : X \to \mathbb{P}^n_{\mathbb{K}}$. In this way, understanding the set of all morphisms from X is roughly equivalent to understanding the set of all morphisms from X to $\mathbb{P}^n_{\mathbb{K}}$. Via Theorem 10.1.16, the latter set is determined by the set of line bundles on X and their sections.

In particular, computing the Picard group of X can be seen as the first step toward understanding all the morphisms from X.

10.1.4 Exercises

Exercise 10.1.18. Let X be a scheme. Suppose that $\mathcal{F}$ and $\mathcal{G}$ are globally generated at a point $x \in X$. Prove that $\mathcal{F} \otimes \mathcal{G}$ is also globally generated at x.

Exercise 10.1.19. Let X be a scheme and let $\mathcal{L}$ be an invertible sheaf on X. Fix a point $x \in X$. Prove that $\mathcal{L}$ is globally generated at x if and only there exists a section $s \in \mathcal{F}(X)$ whose zero locus does not contain x.

Exercise 10.1.20. Let A be a ring. Prove that the set of A-valued points of $\mathbb{P}^n_A$ is the same as the set of rank 1 projective modules M equipped with a surjection $A^{n+1} \to M$ up to isomorphism. Here, an isomorphism of two sets of data is a commuting diagram

$$\begin{tikzcd} A^{n+1} \arrow[r, "id"] \arrow[d] & A^{n+1} \arrow[d] \\ M \arrow[r, "\cong"] & M' \end{tikzcd}$$

In particular, show that if $a_0, \dots, a_n$ are elements of A which generate the unit ideal then we obtain an A-valued point of $\mathbb{P}^n_A$ and two such $(n+1)$-tuples define the same point if and only if they can be identified by rescaling by a unit in A. Find an example of a ring (e.g. $A = \mathbb{Z}[\sqrt{-5}]$) such that there are A-valued points of $\mathbb{P}^n_A$ which do not have this form.

Exercise 10.1.21. Let X be the affine line with the doubled origin. Prove that $\mathrm{Pic}(X) \cong \mathbb{Z}$. Identify which line bundles on X are globally generated.

Exercise 10.1.22. Let $f : \mathbb{P}^1 \to \mathbb{P}^d$ be a rational normal curve. What is the function $f^* : \mathrm{Pic}(\mathbb{P}^d) \to \mathrm{Pic}(\mathbb{P}^1)$?

10.2 Cartier divisors

Suppose that $\mathcal{L}$ is an invertible sheaf on a scheme X and $s \in \mathcal{L}(X)$. The zero locus $Z(s)$ is locally defined by the vanishing of a single equation. As we have seen when discussing Krull's PIT, subschemes defined by principal ideals have special properties not shared by arbitrary schemes, and thus they are worthy of further attention.

In this section we study Cartier divisors. These objects are locally defined by a single *rational* function (i.e. a function which is allowed to have poles). In the special case when the local data consists of *regular* functions, our construction will yield a subscheme locally defined by principal ideals (see Exercise 10.2.20). The advantage of working with *rational* functions is that we will be able to define a group structure.

10.2.1 Cartier divisors on integral schemes

We start by defining Cartier divisors in a special situation – when X is an integral scheme. In this setting we have a function field $K(X)$ equal to the residue field of the generic point of X (or equivalently, the fraction field for any open affine in X). We let $\mathcal{K}(X)$ denote the locally constant sheaf on X corresponding to the abelian group $K(X)$ and let $\mathcal{K}(X)^\times$ denote the locally constant sheaf on X corresponding to the abelian group $K(X)^\times$. Note that there is an injection of sheaves of abelian groups $\mathcal{O}_X^\times \to \mathcal{K}(X)^\times$.

Definition 10.2.1. Let X be an integral scheme. A Cartier divisor L on X is a global section of the sheaf $\mathcal{K}(X)^\times/\mathcal{O}_X^\times$.

It is hard to see the geometric meaning of a Cartier divisor directly from the definition, so frequently another construction is used. Since $\mathcal{K}(X)^\times \to \mathcal{K}(X)^\times/\mathcal{O}_X^\times$ is surjective, we can lift sections of $\mathcal{K}(X)^\times/\mathcal{O}_X^\times$ locally to sections of $\mathcal{K}(X)^\times$. These local rational functions need not agree on overlaps, but they will agree up to multiplication by an element in $\mathcal{O}_X^\times$.

Construction 10.2.2. Let X be an integral scheme and let $\{U_i\}$ be an open cover of X. Suppose that for each U_i we choose an element $f_i \in \mathcal{K}^\times$ in such a way that $f_i/f_j \in \mathcal{O}_X^\times(U_i \cap U_j)$. Then the set of data $\{(U_i, f_i)\}$ defines a Cartier divisor on X. If we set $t_{ij} = f_i/f_j$ in $\mathcal{O}_X^\times(U_i \cap U_j)$, then the t_{ij} satisfy the cocycle conditions

$$t_{ii} = 1 \qquad t_{ij} = t_{ji}^{-1} \qquad t_{ik} = t_{jk}t_{ij}.$$

Note that the representation of a Cartier divisor is not unique; $\{(U_i, f_i)\}$ and $\{(V_j, f'_j)\}$ define the same Cartier divisor if $f_i f_j'^{-1} \in \mathcal{O}_X(U_i \cap V_j)^\times$ for all i, j.

Remark 10.2.3. The local gluing data for Cartier divisors should remind us of the local gluing data for line bundles. Indeed, our main goal in this section is to investigate the close relationship between these two types of objects. One can conceptualize a Cartier divisor as a "twisted" version of a rational function, just as an invertible sheaf is a "twisted" version of the structure sheaf.

As global sections of a sheaf of abelian groups, Cartier divisors admit the structure of an abelian group $\mathrm{CDiv}(X)$ (whose operation we denote using $+$). Precisely, if refine our open cover $\{U_i\}$ so that L_1 is represented by (U_i, f_i) and L_2 is represented by (U_i, f_i') then $L_1 + L_2$ is represented by $(U_i, f_i f_i')$. The inverse $-L_1$ is represented by $\{(U_i, f_i^{-1})\}$.

The Cartier divisors which admit no "twisting" play an important role in the theory.

Definition 10.2.4. Let X be an integral scheme. A Cartier divisor which can be represented as (X, f) for some $f \in K(X)^\times$ is called a principal Cartier divisor. The principal Cartier divisors form a subgroup of $\mathrm{CDiv}(X)$.

Two Cartier divisors L_1, L_2 are said to be linearly equivalent – written $L_1 \sim L_2$ – if $L_1 - L_2$ is a principal divisor. We denote the group of linear equivalence classes of Cartier divisors by $\mathrm{CaCl}(X)$. In other words, $\mathrm{CaCl}(X)$ is the quotient of $\mathrm{CDiv}(X)$ by the subgroup of principal divisors.

Example 10.2.5. Consider the open cover $U_0 = D_{+,x_0}$, $U_1 = D_{+,x_1}$ of $\mathbb{P}^1$. Since the units on the overlap $U_0 \cap U_1$ have the form $c\frac{x_0}{x_1}^d$, one can define a Cartier divisor L on $\mathbb{P}^1$ by choosing rational functions f_0 and f_1 such that $f_1 = c\frac{x_0}{x_1}^d \cdot f_0$ for some integer d. Note that such a divisor is determined by f_0 and by the transition function $c\frac{x_0}{x_1}^d$. (Of course, one could also use a different open cover to define a Cartier divisor.)

If we replace f_0 by some other rational function f_0' and define $f_1' = c'\frac{x_0}{x_1}^d \cdot f_0'$ for some constant c' then the resulting Cartier divisor L' is linearly equivalent to L. In other words, the linear equivalence class of L is determined solely by the exponent d.

There is one more important type of Cartier divisor: those which are locally defined by regular functions (and not just rational functions).

Definition 10.2.6. Let X be an integral scheme. A Cartier divisor L on X is effective if for some (equivalently any) representative consisting of local data $\{(U_i, f_i)\}$ the functions f_i satisfy $f_i \in \mathcal{O}_X(U_i)$ for all i.

We will see that effective Cartier divisors correspond to "locally principal" ideal sheaves. In other words, these are the Cartier divisors with the most geometric significance: they actually come from closed subschemes of X.

10.2.2 Cartier divisors and invertible sheaves

Our first task to associate to every Cartier divisor L an invertible sheaf $\mathcal{O}_X(L)$. One option is to use local gluing data $\{(U_i, f_i)\}$ for L. As in Construction 10.2.2 let $t_{ij} = f_i/f_j$ denote the element of $\mathcal{O}_X(U_i \cap U_j)^\times$ which defines the local gluings. Since the t_{ij} satisfy the cocycle condition, we can glue the sheaves $\mathcal{O}_{U_i}$ along the local isomorphisms ϕ_{ij} defined by multiplication by t_{ij} to obtain an invertible sheaf.

It is more traditional to use the following equivalent formulation. For simplicity, suppose we define L by the gluing data $\{(U_i, f_i)\}$ where the U_i are all open affines. On the set U_i, consider the free rank 1 $\mathcal{O}_{U_i}$-submodule $\mathcal{F}_i$ of $\mathcal{K}(X)|_{U_i}$ generated by f_i^{-1}. Since f_i and f_j only differ by an element of $\mathcal{O}_X(U_i \cap U_j)^\times$ we have an equality

$$\mathcal{F}_i|_{U_i \cap U_j} = \mathcal{F}_j|_{U_i \cap U_j}.$$

Thus the various $\mathcal{F}_i$ glue together (with the identity maps) to give an invertible sheaf $\mathcal{O}_X(L)$ with the transition maps described above. As a consequence, the invertible sheaf $\mathcal{O}_X(L)$ is described explicitly as:

Construction 10.2.7. Let X be an integral scheme and let $L = \{(U_i, f_i)\}$ be a Cartier divisor on X. We define the invertible sheaf $\mathcal{O}_X(L)$ via the prescription

$$\mathcal{O}_X(L)(V) = \{ f \in K(X) \mid f_i f \in \mathcal{O}_X(U_i \cap V)\, \forall i \}$$

where the restriction maps are the inclusions.

Remark 10.2.8. Note that $\mathcal{O}_X(L)$ comes equipped with an inclusion into $\mathcal{K}(X)$.

Remark 10.2.9. At first it is unclear why we would define $\mathcal{O}_X(L)$ using the local functions f_i^{-1} instead of just f_i. Conceptually this choice is similar to the decision about whether a vector bundle should be associated to its sheaf of sections or the dual sheaf. The choice f_i^{-1} guarantees that the sections of $\mathcal{O}_X(L)$ will satisfy the same gluing properties as L (see Proposition 10.2.17).

The following proposition outlines the key properties of this construction.

Proposition 10.2.10. *Let X be an integral scheme and let L_1, L_2 be two Cartier divisors on X.*

(1) We have $\mathcal{O}_X(L_1 + L_2) \cong \mathcal{O}_X(L_1) \otimes \mathcal{O}_X(L_2)$ and $\mathcal{O}_X(-L_1) \cong \mathcal{O}_X(L_1)^\vee$.

(2) We have $\mathcal{O}_X(L_1) \cong \mathcal{O}_X(L_2)$ if and only if L_1 and L_2 are linearly equivalent.

Precisely, if L_1 and L_2 are linearly equivalent then $\mathcal{O}_X(L_1)$ and $\mathcal{O}_X(L_2)$ are isomorphic invertible sheaves equipped with different embeddings into $K(X)$.

Proof. (1) By refining our open cover we may suppose that L_1 corresponds to local data $\{(U_i, f_i)\}$ and L_2 corresponds to local data $\{(U_i, f_i')\}$ where the U_i are affine. Then $L_1 + L_2$ is defined by $\{(U_i, f_i f_i')\}$. Thus for every open affine V we have

$$\mathcal{O}_X(L_1 + L_2)(U_i) = \{ f \in K(X) \mid f_i f_i' f \in \mathcal{O}_X(U_i) \}.$$

On the other hand, for any open affine U_i we have

$$\begin{aligned}(\mathcal{O}_X(L_1) \otimes \mathcal{O}_X(L_2))(U_i) &= \mathcal{O}_X(L_1)(U_i) \otimes_{\mathcal{O}_X(U_i)} \mathcal{O}_X(L_2)(U_i) \\ &= \mathcal{O}_X(U_i) \cdot f_i^{-1} \otimes_{\mathcal{O}_X(U_i)} \mathcal{O}_X(U_i) \cdot f_i'^{-1}\end{aligned}$$

Since quasicoherent sheaves are determined by what happens on an open affine cover, it is clear that these constructions define isomorphic sheaves. The proof of the second statement is similar.

(2) It suffices to show that L is a principal Cartier divisor if and only if $\mathcal{O}_X(L) \cong \mathcal{O}_X$. To see the forward direction, suppose L is defined by the global function $f \in K(X)^\times$. Then the multiplication by f map defines an isomorphism $\mathcal{O}_X(L) \to \mathcal{O}_X$. Conversely, suppose given an isomorphism $\phi : \mathcal{O}_X(L) \to \mathcal{O}_X$. By taking the image of the constant section 1 under ϕ^{-1} we get an element $s \in \mathcal{O}_X(L)(X) \subset K(X)$. If $\{(U_i, f_i)\}$ denotes local data for L, we see that on each U_i we must have that $f_i s \in \mathcal{O}_X(U_i)^\times$. Thus the principal divisor (X, s^{-1}) defines the same Cartier divisor. $\square$

Example 10.2.11. Consider the Cartier divisor L on $\mathbb{P}^1$ defined by f_0, f_1 satisfying $f_1 = c\frac{x_0}{x_1}^d \cdot f_0$ as in Example 10.2.5. Then the corresponding line bundle $\mathcal{O}_X(L)$ will be isomorphic to $\mathcal{O}(d)$. The choice of f_0 (which determines f_1) will only affect how this line bundle is embedded into $\mathcal{K}(X)$.

Proposition 10.2.10 shows that we get an injective group homomorphism $\mathrm{CaCl}(X) \to \mathrm{Pic}(X)$. In our setting of integral schemes, it turns out that this map is an isomorphism.

Theorem 10.2.12. *Let X be an integral scheme. Then the map* $\mathrm{CaCl}(X) \to \mathrm{Pic}(X)$ *is an isomorphism of groups.*

Proof. It suffices to show that for any invertible sheaf $\mathcal{L}$ there is a Cartier divisor L such that $\mathcal{O}_X(L) \cong \mathcal{L}$. Consider the inclusion $\mathcal{O}_X \to \mathcal{K}(X)$ and tensor by $\mathcal{L}$ to get a morphism $\phi : \mathcal{L} \to \mathcal{L} \otimes_{\mathcal{O}_X} \mathcal{K}(X)$. Since $\mathcal{L}$ is locally isomorphic to $\mathcal{O}_X$, the map ϕ is injective.

We claim that $\mathcal{L} \otimes_{\mathcal{O}_X} \mathcal{K}(X) \cong \mathcal{K}(X)$. Indeed, since $\mathcal{L}$ is locally isomorphic to $\mathcal{O}_X$, X admits an open cover by sets U such that $(\mathcal{L} \otimes_{\mathcal{O}_X} \mathcal{K}(X))|_U$ is the locally constant sheaf on U associated to $K(X)$. Since X is irreducible, this implies that the entire sheaf $\mathcal{L} \otimes_{\mathcal{O}_X} \mathcal{K}(X)$ is locally constant, and thus isomorphic to $\mathcal{K}(X)$.

Choose an open cover $\{U_i\}$ which trivializes $\mathcal{L}$. The map ϕ identifies $\mathcal{L}|_{U_i}$ with a free $\mathcal{O}_X$-submodule of $\mathcal{K}(X)|_{U_i}$ generated by some element $g_i \in K(X)^\times$. By the construction we must have that on $U_i \cap U_j$ the quotient g_i/g_j is an element of $\mathcal{O}_X(U_i \cap U_j)^\times$. If we define a Cartier divisor L by the local data $\{(U_i, g_i^{-1})\}$ then $\mathcal{O}_X(L)$ is the same as the ϕ-image of $\mathcal{L}$. $\square$

10.2.3 Cartier divisors and sections

Having associated an invertible sheaf to every Cartier divisor, we now construct a map in the "reverse direction." More precisely, there are many Cartier divisors associated to an invertible sheaf. What extra data on $\mathcal{L}$ do we need to specify to allow us to recover a Cartier divisor uniquely?

Definition 10.2.13. Let X be an integral scheme and let $\mathcal{L}$ be an invertible sheaf on X. Consider a pair (U, s) where U is a dense open subset and $s \in \mathcal{L}(U)$. We say that two pairs (U, s) and (V, t) are equivalent if $s|_{U \cap V} = t|_{U \cap V}$. A rational section of $\mathcal{L}$ is an equivalence class of pairs (U, s). We will often omit the open set U from our notation.

Suppose we are given two invertible sheaves equipped with rational sections $(\mathcal{L}, s)$ and $(\mathcal{L}, s')$. We say that these two pairs are isomorphic if there is an isomorphism $\phi : \mathcal{L} \to \mathcal{L}'$ and an open set U where s, s' are defined such that $\phi(U)(s) = s'$.

In particular, suppose we have two rational sections s, s' of the same line bundle $\mathcal{L}$. By Exercise 9.5.21 we have

$$\mathrm{Hom}(\mathcal{L}, \mathcal{L}) \cong \mathrm{Hom}(\mathcal{O}_X, \mathcal{O}_X) \cong \mathcal{O}_X(X)$$

so every isomorphism $\phi : \mathcal{L} \to \mathcal{L}$ is defined by global multiplication by an element of $\mathcal{O}_X(X)^\times$. Thus $(\mathcal{L}, s)$ and $(\mathcal{L}, s')$ are isomorphic if and only if s and s' are related by an element of $\mathcal{O}_X(X)^\times$.

Construction 10.2.14. Let X be an integral scheme, $\mathcal{L}$ an invertible sheaf on X, s a rational section of $\mathcal{L}$. The divisor of zeros and poles of s is the Cartier divisor defined as follows. First choose an open affine cover $\{U_i\}$ trivializing $\mathcal{L}$ and choose isomorphisms $\phi_i : \mathcal{L}|_{U_i} \to \mathcal{O}_{U_i}$. If s is defined on all of U_i then we take the local datum $(U_i, \phi_i(s))$. If not, then we take the local datum $(U_i, \widehat{\phi}_i(s))$ where $\widehat{\phi}_i$ is the restriction of ϕ_i to a distinguished open affine in U_i where s is defined. (The function $\widehat{\phi}_i$ maps to a localization of $\mathcal{O}_X(U_i)$ and thus to $K(X)$.)

Note that our choices of the open cover U_i and the isomorphisms ϕ_i only affect the local functions $\phi_i(s)$ up to rescaling by a unit in $\mathcal{O}_X(U_i)^\times$. Thus altogether the local data define a Cartier divisor which is independent of all choices.

Remark 10.2.15. The Cartier divisor in Construction 10.2.14 only depends upon the isomorphism type of the pair $(\mathcal{L}, s)$.

Remark 10.2.16. If in Construction 10.2.14 the rational section s is defined on the open set U_i then $\phi_i(s)$ is actually an element of $\mathcal{O}_{U_i}$. In particular, if $s \in \mathcal{L}(X)$ then the resulting Cartier divisor is effective.

Proposition 10.2.17. *Let X be an integral scheme and let $\mathcal{L}$ be an invertible sheaf on X. Then:*

(1) There is a bijection between global sections $s \in \mathcal{L}(X)$ up to rescaling by $\mathcal{O}_X(X)^\times$ and effective Cartier divisors L such that $\mathcal{O}_X(L) \cong \mathcal{L}$.

(2) There is a bijection between rational sections of $\mathcal{L}$ up to rescaling by $\mathcal{O}_X(X)^\times$ and Cartier divisors L such that $\mathcal{O}_X(L) \cong \mathcal{L}$.

Proof. Given a rational section s of $\mathcal{L}$, consider the Cartier divisor L obtained by taking the divisor of zeros and poles. Recall that L is effective whenever s is a global section and that rescaling s by an element of $\mathcal{O}_X(X)^\times$ does not change L. We still need to check that $\mathcal{O}_X(L) \cong \mathcal{L}$. But this follows from the fact that the transition maps for $\mathcal{L}$ are the same as the transition maps for s, which in turn define the transition maps for $\mathcal{O}_X(L)$.

Conversely, suppose given a Cartier divisor L such that $\mathcal{O}_X(L)$ is isomorphic to $\mathcal{L}$. Note that the function $1 \in K(X)$ defines a rational section of $\mathcal{O}_X(L)$; when L is effective, $1 \in K(X)$ is even a global section of $\mathcal{O}_X(L)$. Under our isomorphism $\phi : \mathcal{O}_X(L) \to \mathcal{L}$ the rational section 1 is taken to a rational section s of $\mathcal{L}$. Note that this construction only identifies $(\mathcal{L}, s)$ up to isomorphism.

These two constructions are inverses: the Cartier divisor L corresponds to the rational section 1 of $\mathcal{O}_X(L)$ and this identification is compatible with both constructions. □

10.2.4 Cartier divisors on arbitrary schemes

Our definition of Cartier divisors does not extend to arbitrary schemes since we used the function field $K(X)$ in an essential way. Here is how to extend the notion more generally.

For any scheme X we can define the sheaf of total quotient rings (also denoted $\mathcal{K}(X)^\times$) as follows. For each open affine U let $S(U)$ denote the subset of functions in $\mathcal{O}_X(U)$ whose restriction to the stalk $\mathcal{O}_{X,x}$ is a non-zerodivisor for every $x \in U$. (This condition on $S(U)$ implies, but is stronger than, the condition that each element of $S(U)$ be a non-zerodivisor in $\mathcal{O}_X(U)$.) The assignment $U \mapsto S(U)^{-1}\mathcal{O}_X(U)$ is a presheaf of $\mathcal{O}_X$-algebras. We let $\mathcal{K}(X)^\times$ be the sheafification of this presheaf.

We can define Cartier divisors using $\mathcal{K}(X)^\times$ in place of the locally constant sheaf with value $K(X)^\times$. Unfortunately this construction is somewhat subtle when X is not Noetherian (see [Kle79]). However, in practice this still amounts to defining Cartier divisors via local invertible functions which satisfy a compatibility from $\mathcal{O}_X^\times$ on overlaps. A principal divisor is a Cartier divisor associated to a global section of $\mathcal{K}(X)^\times$, and $\mathrm{CaCl}(X)$ is the group of Cartier divisors modulo the principal ones.

For an arbitrary scheme X we still have an inclusion $\mathrm{CaCl}(X) \hookrightarrow \mathrm{Pic}(X)$. However, this map need not be surjective in general, since there may be invertible sheaves which do not embed in $\mathcal{K}(X)$. It is an isomorphism in most common situations: when X is an integral scheme, or when X is projective over a Noetherian ring, or when X is a Noetherian scheme without embedded points.

For an arbitrary scheme X and an invertible sheaf $\mathcal{L}$, we define a rational section of $\mathcal{L}$ to be (an equivalence class of) a section defined on a scheme-theoretically dense open subset U. (When X is Noetherian, this means that U contains every associated point of X.) Furthermore, we are only interested in those rational sections s which are invertible, in the sense that there is a rational section $s^\vee$ of $\mathcal{L}^\vee$ such that under the pairing $\mathcal{L} \otimes \mathcal{L}^\vee \to \mathcal{O}_X$

we have $s \otimes s^\vee \mapsto 1$. (When X is Noetherian, this amounts to saying that the restriction of s to the closure of any associated point does not vanish.) With these definitions the group of Cartier divisors is isomorphic to the group of pairs $(\mathcal{L}, s)$ of an invertible sheaf $\mathcal{L}$ and an invertible rational section s up to isomorphism; see [Sta].

10.2.5 Exercises

Exercise 10.2.18. Consider the standard open covering $U_0, \ldots, U_n$ of $\mathbb{P}^n$ by affine charts. We will define a Cartier divisor L on $\mathbb{P}^n$ using this open covering. To U_0 we associate the function f_0 and to U_1 we associate the function $f_1 = c\frac{x_0}{x_1}^d \cdot f_0$. Explain why this data uniquely determines the values of L on every other open chart U_i.

(This exercise is an avatar of the fact that a function or invertible sheaf on a normal variety is determined uniquely over the complement of a codimension 2 closed subset.)

Exercise 10.2.19. Suppose that $f : X \to Y$ is a dominant morphism of integral schemes. In particular we have an induced injection $f^\sharp : K(Y) \to K(X)$. Given a Cartier divisor $L = \{(V_i, g_i)\}$ on Y, we define the pullback f^*L by the prescription $f^*L := \{(f^{-1}V_i, f^\sharp(g_i))\}$.

(1) Prove that f^*L is a Cartier divisor.

(2) Prove that f^* preserves linear equivalence.

(3) Prove that $\mathcal{O}_X(f^*L) \cong f^*\mathcal{O}_Y(L)$.

In fact, we can define pullbacks in much more general situations. Given any morphism of integral schemes $f : X \to Y$, we can define the pullback of a Cartier divisor L so long as $f(X)$ is not contained in the closure of the vanishing locus of the denominator of one of the functions g_i defining L (and the results of this exercise will still hold).

Exercise 10.2.20. Let X be an integral scheme and let $L = \{U_i, f_i\}$ be an effective Cartier divisor. For each index i let $\mathcal{I}_{U_i}$ be the quasicoherent ideal sheaf on U_i generated by f_i. (That is, for every open affine $V \subset U_i$ we have $\mathcal{I}_{U_i}(V) = \widetilde{(f_i|_V)}$.)

(1) Explain why the various $\mathcal{I}_{U_i}$ can be glued to yield a quasicoherent ideal sheaf $\mathcal{I}$ on X.

(2) Prove the important relation $\mathcal{I} \cong \mathcal{O}_X(-L)$.

It is quite common to implicitly identify the effective Cartier divisor L with the closed subscheme Z defined by $\mathcal{I}$.

Exercise 10.2.21. Let X be an integral affine scheme and let L be an effective Cartier divisor on X. The support of L is defined to be the closed subset underlying the corresponding subscheme defined in Exercise 10.2.20. Prove that the complement of the support of L is again an affine scheme. (Hint: show that the inclusion $U \to X$ is an affine morphism by arguing locally.)

10.3 Weil divisors

Throughout this section X will denote an integral normal Noetherian scheme. Recall that a scheme is normal if all of the stalks of its structure sheaf are integrally closed; for integral schemes, this is the same as saying that every affine open is defined by an integrally closed ring.

Definition 10.3.1. Let X be an integral normal Noetherian scheme. A prime divisor on X is an integral closed subscheme of codimension 1.

A Weil divisor is a formal sum of prime divisors, i.e. a finite sum $\sum_{i=1}^{n} a_i Y_i$ where $a_i \in \mathbb{Z}$ and Y_i is a prime divisor. The group of Weil divisors is denoted $\mathrm{WDiv}(X)$. The multiplicity of a Weil divisor D along a prime divisor Y, denoted by $\mathrm{mult}_Y(D)$, is the coefficient of Y in D. The support of a Weil divisor D, denoted $\mathrm{Supp}(D)$, is the union of all prime divisors Y such that $\mathrm{mult}_Y(D) \neq 0$.

A Weil divisor D is said to be effective if every coefficient a_i is non-negative; we denote this condition by writing $D \geq 0$.

Exercise 10.3.18 shows how any closed subscheme $Z \subset X$ of pure codimension 1 yields an effective Weil divisor which records only the "codimension 1 information" of Z.

Construction 10.3.2. Given a prime divisor Y, let η_Y denote its generic point. Since X is normal and Y has codimension 1, $\mathcal{O}_{X,\eta_Y}$ will be an integrally closed local Noetherian ring with Krull dimension 1. In other words, $\mathcal{O}_{X,\eta_Y}$ is a discrete valuation ring yielding a discrete valuation val_Y on the function field $K(X)$.

To any $f \in K(X)^\times$ we can associate a Weil divisor via the prescription

$$\mathrm{div}(f) := \sum \mathrm{val}_Y(f) Y.$$

We call $\mathrm{div}(f)$ the divisor of zeros and poles of f – prime divisors with positive valuations are known as "zeros" of f and prime divisors with negative valuations are known as "poles" of f. The fact that $\mathrm{div}(f)$ is a well-defined Weil divisor is a result of the following claim.

Claim 10.3.3. Let X be an integral normal Noetherian scheme with function field $K(X)$. For any $f \in K(X)^\times$ there are only finitely many prime divisors Y such that $\mathrm{val}_Y(f) \neq 0$.

Proof. Let $U \subset X$ be an open affine such that $f \in \mathcal{O}_X(U)$. The complement $X \backslash U$ is a proper closed subscheme. Since any prime divisor contained in $X \backslash U$ must be an irreducible component of $X \backslash U$, this set can only contain finitely many prime divisors. Thus, it suffices to prove that there are only finitely many prime divisors Y which intersect U and have $\mathrm{val}_Y(f) \neq 0$.

Since $f \in \mathcal{O}_X(U)$, any prime divisor Y which intersects U will have $\mathrm{val}_Y(f) \geq 0$. Furthermore, $\mathrm{val}_Y(f) > 0$ if and only if $Y \subset V(f)$. Since $V(f)$ has only finitely many irreducible components, we obtain the desired statement. $\square$

Remark 10.3.4. Lemma 10.3.10 will show that the kernel of the map $\operatorname{div} : K(X)^\times \to \mathrm{WDiv}(X)$ is $\mathcal{O}_X(X)^\times$.

Definition 10.3.5. A divisor of the form $\operatorname{div}(f)$ for $f \in K(X)^\times$ is called a principal divisor. The set of all principal divisors forms a subgroup of $\mathrm{WDiv}(X)$ (since div is a homomorphism).

Two Weil divisors D_1, D_2 are said to be linearly equivalent – written $D_1 \sim D_2$ – if their difference is a principal divisor. The group of linear equivalence classes of Weil divisors is known as the class group and is denoted $\mathrm{Cl}(X)$. In other words, $\mathrm{Cl}(X)$ is the quotient of $\mathrm{WDiv}(X)$ by the subgroup of principal divisors.

Example 10.3.6. Let's analyze the group $\mathrm{Cl}(\mathbb{P}^n)$. We can define a homomorphism $\deg : \mathrm{WDiv}(X) \to \mathbb{Z}$ via the prescription

$$\deg\left(\sum a_i Y_i\right) = \sum a_i \deg(Y_i).$$

Recall that $K(\mathbb{P}^n_{\mathbb{K}})$ is the set of quotients $\frac{f}{g}$ where f, g are homogeneous polynomials in $\mathbb{K}[x_0, \ldots, x_n]$ which have the same degree. Then $\operatorname{div}(\frac{f}{g})$ is the Weil divisor $V_+(f) - V_+(g)$. In particular, every principal Weil divisor on $\mathbb{P}^n$ has degree 0.

Conversely, suppose $D = \sum a_i Y_i$ is a Weil divisor of degree 0. Recall that every prime divisor on $\mathbb{P}^n$ is defined by a single irreducible homogeneous equation (by applying Proposition 4.4.8 on an affine chart). Suppose that $Y_i = V_+(f_i)$. If we set $f = \prod f_i^{a_i}$, then f is an element of $K(\mathbb{P}^n_{\mathbb{K}})$ with $\operatorname{div}(f) = D$. We conclude that the principal divisors are exactly the same as the divisors of degree 0. Thus $\mathrm{Cl}(\mathbb{P}^n) \cong \mathbb{Z}$.

Example 10.3.7. Let R be a Noetherian UFD. Then $\mathrm{Spec}(R)$ is an integral normal Noetherian scheme. We claim that $\mathrm{Cl}(\mathrm{Spec}(R)) = 0$. Indeed, Proposition 4.4.8 shows that every prime divisor Y is the vanishing locus of a principal ideal (f). Then $\operatorname{div}(f) = Y$ and so every prime divisor is principal.

It turns out that the converse is also true: if R is an integrally closed Noetherian domain with $\mathrm{Cl}(\mathrm{Spec}(R)) = 0$ then R is a UFD. The proof is not too hard but we will not need it.

10.3.1 Cartier divisors and Weil divisors

In Construction 10.3.2 we saw how to associate a Weil divisor to any element in $K(X)^\times$. In fact this construction can be extended to arbitrary Cartier divisors as follows.

Construction 10.3.8. Let X be an integral normal Noetherian scheme and let $\{(U_i, f_i)\}$ be local data determining a Cartier divisor L. For any prime divisor Y, we define the multiplicity of L along Y by choosing an open set U_i such that $Y \cap U_i \neq 0$ and setting

$$\mathrm{mult}_Y(L) := \mathrm{val}_Y(f_i).$$

Referring to the discussion of representations of a Cartier divisor in Construction 10.2.2, it is clear that this definition is independent of the choice of open set U_i and the choice of local data $\{(U_i, f_i)\}$ determining L. We then define the Weil divisor associated to L as

$$\mathrm{div}(L) = \sum_Y \mathrm{mult}_Y(L)Y.$$

Note that this sum is finite since we can choose a finite cover of opens and apply Claim 10.3.3 on each.

We thus obtain a map $\mathrm{div} : \mathrm{CDiv}(X) \to \mathrm{WDiv}(X)$. Note that this map sends effective Cartier divisors to effective Weil divisors: if the local data for L is defined by equations in $\mathcal{O}_X(U_i)$, then all the valuations of f_i for divisors intersecting U_i are non-negative and thus $\mathrm{div}(L) \geq 0$.

Example 10.3.9. Suppose that L is the Cartier divisor on $\mathbb{P}^n$ corresponding to a rational section f/g of $\mathcal{O}_{\mathbb{P}^n}(d)$ where $\deg(f) - \deg(g) = d$. Then we have $\mathrm{div}(L) = V_+(f) - V_+(g)$.

Lemma 10.3.10. *Let X be an integral normal Noetherian scheme. The map* $\mathrm{div} : \mathrm{CDiv}(X) \to \mathrm{WDiv}(X)$ *is injective.*

Proof. Suppose $L = \{(U_i, f_i)\}$ is in the kernel of div. After refining the cover we may suppose that each U_i is affine. Then the valuations of f_i along every prime divisor intersecting U_i is zero. In particular, $f_i|_U$ is contained in every localization of $\mathcal{O}_X(U)$ along a height one prime. Since $\mathcal{O}_X(U_i)$ is integrally closed, this implies that $f_i \in \mathcal{O}_X(U_i)$. By the same argument $f_i^{-1} \in \mathcal{O}_X(U_i)$ so that f_i is a unit on U_i. But the only Cartier divisor defined locally by units is the trivial Cartier divisor. □

By quotienting out by the group of prinicipal divisors we obtain an injection $\mathrm{CaCl}(X) \to \mathrm{Cl}(X)$. Note that a Weil divisor D will be linearly equivalent to a Cartier divisor L if and only if D is itself Cartier.

Loosely speaking, the image of the map $\mathrm{div} : \mathrm{CDiv}(X) \to \mathrm{WDiv}(X)$ has image the Weil divisors which can locally be defined by a single equation. The following example illustrates that the map $\mathrm{CDiv}(X) \to \mathrm{WDiv}(X)$ need not be surjective. As we will see in Theorem 10.4.2, this failure of surjectivity is related to the singularities of X.

Example 10.3.11. Consider the quadric cone $X := \mathrm{Spec}(\mathbb{K}[x, y, z]/(xy - z^2))$. The prime ideal (x, z) defines one of the lines through the origin; we call this Weil divisor D. In Exercise 5.1.16 we used the Zariski tangent space to show that D is not locally principal – there is no open neighborhood of the origin along which the ideal of D is defined by a single equation. Thus D is not in the image of the map div.

However $\mathrm{div}(x) = 2D$ is a Cartier divisor (since $2D$ is locally defined by $x = 0$). In fact, it turns out that the cokernel of the map $\mathrm{CDiv}(X) \to \mathrm{WDiv}(X)$ is the abelian group $\mathbb{Z}/2\mathbb{Z}$ and this cokernel is generated by D. (This is related to the fact that X is the quotient of $\mathbb{A}^2$ by the $\mathbb{Z}/2\mathbb{Z}$-action sending $(x, y) \mapsto (-x, -y)$.)

10.3.2 The sheaf associated to a Weil divisor

Suppose that $D = \sum a_i Y_i$ is a Weil divisor. We would like to associate to D a sheaf $\mathcal{O}_X(D)$ consisting of all rational functions whose poles are "at worst $-D$".

Definition 10.3.12. Let X be an integral normal Noetherian scheme and let $D = \sum a_i Y_i$ be a Weil divisor on X. We define the sheaf $\mathcal{O}_X(D)$ via the prescription

$$\begin{aligned}\mathcal{O}_X(D)(U) &= \{f \in K(X) | (\mathrm{div}(f) + D)|_U \geq 0\} \\ &= \{f \in K(X) | \operatorname{val}_Y(f) \geq \operatorname{mult}_Y(D)\ \forall Y \text{ s.t. } Y \cap U \neq \emptyset\}.\end{aligned}$$

and whose restriction maps are the inclusion maps.

Since multiplying a rational function by an element of $\mathcal{O}_X(U)$ can only increase the valuations along divisors which intersect U, we see that $\mathcal{O}_X(D)$ is actually an $\mathcal{O}_X$-module. In fact it turns out that:

Lemma 10.3.13. *Let X be an integral normal Noetherian scheme and let $D = \sum a_i Y_i$ be a Weil divisor on X. Then $\mathcal{O}_X(D)$ is a coherent sheaf on X whose rank at the generic point of X is equal to* 1.

Remark 10.3.14. There is an important subtlety concerning Definition 10.3.12. Suppose that L is a Cartier divisor. We claim that the sheaf $\mathcal{O}_X(\mathrm{div}(L))$ of Definition 10.3.12 is isomorphic to the sheaf $\mathcal{O}_X(L)$ of Construction 10.2.7 (resolving a potential conflict in notation).

The key observation is that since X is normal, if U is an open affine and $f \in K(X)$ is a rational function which is well-defined on all of U then $f \in \mathcal{O}_X(U)$ (as we saw before in the proof of Lemma 10.3.10). Thus, for integral normal Noetherian schemes the two sheaf descriptions in Definition 10.3.12 and Construction 10.2.7 exactly coincide.

However, if X is not normal, we can only define $\mathcal{O}_X(L)$ using Construction 10.2.7 and not Definition 10.3.12; see Exercise 10.3.16.

Theorem 10.3.15. *Let X be an integral normal Noetherian scheme. A Weil divisor D is Cartier if and only if $\mathcal{O}_X(D)$ is an invertible sheaf.*

Proof. Remark 10.3.14 shows that if D is a Cartier divisor then $\mathcal{O}_X(D)$ is an invertible sheaf. Conversely, suppose that D is a Weil divisor such that $\mathcal{O}_X(D)$ is an invertible sheaf. Choose a trivializing cover of open affines $\{U_i\}$ for $\mathcal{O}_X(D)$; for each U_i, there is an element $s_i \in K(X)$ which generates $\mathcal{O}_X(D)(U_i)$ as a $\mathcal{O}_X(U_i)$-submodule of $K(X)$. If we consider the Cartier divisor L defined by $\{(U_i, s_i^{-1})\}$ then it is clear that $\mathrm{div}(L) = D$. $\square$

It turns out that the interaction between Weil divisors and the sheaves $\mathcal{O}_X(D)$ is similar to the interaction between Cartier divisors and the sheaves $\mathcal{O}_X(L)$ (see [Sch]):

(1) The sheaves of the form $\mathcal{O}_X(D)$ are precisely the reflexive sheaves (i.e. the coherent sheaves satisfying $\mathcal{F}^{\vee\vee} \cong \mathcal{F}$) which have rank 1 at the generic point.

(2) Given Weil divisors D_1, D_2 we have $\mathcal{O}_X(D_1 + D_2) \cong (\mathcal{O}_X(D_1) \otimes \mathcal{O}_X(D_2))^{\vee\vee}$ and $\mathcal{O}_X(-D_1) = \mathcal{O}_X(D_1)^{\vee}$.

(3) The set of global sections of a reflexive rank 1 sheaf $\mathcal{D}$ up to rescaling by $\mathcal{O}_X(X)^{\times}$ is in bijection with the set of effective Weil divisors D such that $\mathcal{O}_X(D) \cong \mathcal{D}$.

10.3.3 Exercises

Exercise 10.3.16. Consider the "pinched plane" defined by the subring $R = \mathbb{K}[x^3, x^2, xy, y^2]$ of $\mathbb{K}[x, y]$. Then $\mathrm{Spec}(R)$ is an example of a scheme that is regular in codimension 1 but is not normal.

Since the singularities of $\mathrm{Spec}(R)$ all have codimension 2, we can still define $\mathrm{div} : \mathrm{CDiv}(\mathrm{Spec}(R)) \to \mathrm{WDiv}(\mathrm{Spec}(R))$ in the same way. Show however that for a Cartier divisor L we might have $\mathcal{O}_X(L) \ncong \mathcal{O}_X(\mathrm{div}(L))$.

Exercise 10.3.17. Let X be the cone over a smooth quadric surface, i.e. X is defined by the ring $\mathbb{K}[w, x, y, z]/(wy - xz)$. Consider the prime divisor $D \subset X$ defined by the equation $w = x = 0$. Show that the complement $X \backslash D$ is not an affine variety. (Hint: if it were affine, then its intersection with the plane $y = z = 0$ in X would also be affine. What is this intersection?)

Exercise 10.2.21 shows that the complement of an effective Cartier divisor in an integral affine scheme is always affine. Conclude that no multiple of the prime divisor $D \subset X$ defined by the equation $w = x = 0$ is Cartier.

Exercise 10.3.18. Let X be an integral normal Noetherian scheme and let $Z \subset X$ be a closed subscheme of pure codimension 1. For every irreducible component $Z_i \subset Z$ let $\eta_i \in X$ denote the generic point of Z_i. Let $\mathcal{I}_Z$ denote the ideal sheaf of Z and let $\mathrm{mult}_{\eta_i}(\mathcal{I}_Z)$ denote the minimal power of the maximal ideal $\mathfrak{m}_{\eta_i}$ that is contained in the stalk $\mathcal{I}_{Z,\eta_i}$. We define the Weil divisor underlying Z to be the effective Weil divisor

$$\sum_i \mathrm{mult}_{\eta_i}(\mathcal{I}_Z) Z_i.$$

Show that two closed subschemes Z and Z' define the same Weil divisor if and only if there is a locally closed subscheme $U \subset X$ that is a (set-theoretically) dense open subscheme of both Z and Z'.

10.4 Computing the Picard group

Let X be an integral normal Noetherian scheme. The following diagram summarizes the relationships between the various constructions in the chapter:

$$\begin{tikzcd}
\{(\mathcal{L}, s)\}/\mathrm{iso} \arrow[d] & \mathrm{CDiv}(X) \arrow[l] \arrow[r] \arrow[d] & \mathrm{WDiv}(X) \arrow[d] \\
\mathrm{Pic}(X) & \mathrm{CaCl}(X) \arrow[l] \arrow[r] & \mathrm{Cl}(X)
\end{tikzcd}$$

Here the two left hand arrows are isomorphisms by Theorem 10.2.12 and Proposition 10.2.17 and the two right hand arrows are injections by Lemma 10.3.10. Our next result gives a condition for the right hand arrows to be isomorphisms as well.

Definition 10.4.1. Let X be an integral normal Noetherian scheme. We say that X is locally factorial if every local ring $\mathcal{O}_{X,x}$ is a UFD.

Theorem 10.4.2. *Let X be an integral normal Noetherian scheme that is locally factorial. Then the map* $\mathrm{div} : \mathrm{CDiv}(X) \to \mathrm{WDiv}(X)$ *is an isomorphism.*

In practice the group of Weil divisors is sometimes easier to compute than the group of Cartier divisors, so this result is often applied by computing the group of Weil divisors to deduce facts about Cartier divisors.

Proof. It suffices to show that div is surjective. Let D be a Weil divisor and let $x \in X$ be any point. The intersection of D with $\mathrm{Spec}(\mathcal{O}_{X,x})$ defines a Weil divisor D_x on this scheme. Since $\mathcal{O}_{X,x}$ is a UFD by assumption, this means that there is a Cartier divisor L_x on $\mathrm{Spec}(\mathcal{O}_{X,x})$ such that $\mathrm{div}(L_x) = D_x$. We let $f_x \in K(X)^\times$ be a defining equation for L_x on an open neighborhood of x.

Let $U_x \subset X$ be an open neighborhood of $x \in X$. By shrinking U_x, we may ensure that it only intersects the components of D which contain x. After shrinking further, we may ensure that it only intersects the components of $\mathrm{div}(f_x)$ which contain x. It is then clear that on such an open set U_x we have $\mathrm{div}(f_x) = D \cap U_x$.

As we vary the pairs (U_x, f_x) we obtain a Cartier divisor D on X. Indeed, if $y \in U_x$ then according to Lemma 10.3.10 the functions f_y and f_x on $U_y \cap U_x$ only differ by a unit. Furthermore it is clear that $\mathrm{div}(L) = D$. □

Corollary 10.4.3. *Let X be a regular $\mathbb{K}$-variety. Then* $\mathrm{Pic}(X) \cong \mathrm{Cl}(X)$.

Proof. Since a regular local ring is an integrally closed UFD, Theorem 10.4.2 applies in this situation to show that $\mathrm{CaCl}(X) \cong \mathrm{Cl}(X)$. Theorem 10.2.12 shows that $\mathrm{Pic}(X) \cong \mathrm{CaCl}(X)$. □

10.4.1 Computing the class group

The following theorem facilitates the computation of the class group (and thus, in settings where Theorem 10.4.2 applies, the Picard group).

Theorem 10.4.4. *Let X be an integral normal Noetherian scheme. Let Z be a closed subscheme and let $U = X \backslash Z$. Suppose that $Z_1, \ldots, Z_r$ are the prime divisors contained in Z. Then there is an exact sequence*

$$\bigoplus_{i=1}^{r} \mathbb{Z}Z_i \to \mathrm{Cl}(X) \xrightarrow{\psi} \mathrm{Cl}(U) \to 0$$

where the map ψ sends $D \mapsto D \cap U$.

Proof. If D is a prime divisor on U, then the closure of D in X is a prime divisor whose ψ-image is D. Thus the map on the right is surjective.

To show exactness in the middle, suppose that D is a Weil divisor on X whose image in the class group is in the kernel of ψ. The means that $D \cap U$ is a prinicipal divisor, i.e. $D \cap U = \mathrm{div}(f)$ for some $f \in K(U)^\times = K(X)^\times$. In other words, $D - \mathrm{div}(f)$ is a divisor that is supported on Z, and thus is a linear combination of the Z_i. □

This theorem combines well with Example 10.3.7, which shows that an open affine $U \subset X$ has trivial class group if its ring of functions is a Noetherian UFD.

Example 10.4.5. Let $X = \mathbb{P}^1 \times \mathbb{P}^1$. We will let $\mathcal{O}_X(1,0)$ denote the pullback of $\mathcal{O}_{\mathbb{P}^1}(1)$ under the first projection map, and similarly for $\mathcal{O}_X(0,1)$ and the second projection map. We will denote by $\mathcal{O}_X(a,b)$ the tensor product of a copies of $\mathcal{O}_X(1,0)$ and b copies of $\mathcal{O}_X(0,1)$ (where if a or b is negative we replace these bundles by their duals.)

We claim that $\mathrm{Pic}(X) \cong \mathbb{Z}^2$ is generated by $\mathcal{O}_X(0,1)$ and $\mathcal{O}_X(1,0)$. If we let $U \subset X$ denote the open set $D_{+,s} \times D_{+,u} \cong \mathbb{A}^2$, then the complement $X \backslash U$ is the union $F_1 \cup F_2$ where F_i is a fiber of the ith projection map. Theorem 10.4.4 yields an exact sequence

$$\mathbb{Z}F_1 \oplus \mathbb{Z}F_2 \to \mathrm{Cl}(\mathbb{P}^2) \to \mathrm{Cl}(\mathbb{A}^2) = 0 \to 0.$$

We now translate back to invertible sheaves. Theorem 10.4.2 shows that $\mathrm{Cl}(\mathbb{P}^2) \cong \mathrm{Pic}(\mathbb{P}^2)$ and it is clear that $\mathcal{O}_X(F_1) \cong \mathcal{O}_X(1,0)$ and $\mathcal{O}_X(F_2) \cong \mathcal{O}_X(0,1)$. We then have a surjection

$$\mathbb{Z}\mathcal{O}_X(1,0) \oplus \mathbb{Z}\mathcal{O}_X(0,1) \to \mathrm{Pic}(\mathbb{P}^2)$$

and it suffices to show that this map is injective. Note however that the restriction of $\mathcal{O}_X(a,b)$ to F_1 is isomorphic to $\mathcal{O}_{\mathbb{P}^1}(b)$ and the restriction to F_2 is isomorphic to $\mathcal{O}_{\mathbb{P}^1}(a)$. Thus the sheaves $\mathcal{O}_X(a,b)$ are non-isomorphic when a, b are different.

Exercise 10.4.6. Consider the embedding $\mathbb{P}^1 \times \mathbb{P}^1 \hookrightarrow \mathbb{P}^3$ as a quadric hypersurface. What is the image of the restriction map $\mathrm{Pic}(\mathbb{P}^3) \to \mathrm{Pic}(\mathbb{P}^1 \times \mathbb{P}^1)$?

10.4.2 Lefschetz hyperplane theorem

Suppose that X is a smooth projective $\mathbb{K}$-variety equipped with a closed embedding into $\mathbb{P}^n$. Let H be a general hyperplane on $\mathbb{P}^n$. There is a collection of results known as the Lefschetz Hyperplane Theorems which show that the cohomology of $X \cap H$ is similar to the cohomology of X. This implies a similar comparison result for the Picard group:

Theorem 10.4.7 (Grothendieck-Lefschetz Theorem)**.** *Let $\mathbb{K}$ be an algebraically closed field. Let X be a projective variety equipped with a closed embedding $X \hookrightarrow \mathbb{P}^n$ and let H be a hyperplane in $\mathbb{P}^n$.*

(1) If X and $X \cap H$ are smooth and $\dim(X) \geq 4$ *then the restriction map* $\mathrm{Pic}(X) \to \mathrm{Pic}(X \cap H)$ *is an isomorphism.*

(2) If X and $X \cap H$ are smooth and $\dim(X) \geq 3$ *then the restriction map* $\mathrm{Pic}(X) \to \mathrm{Pic}(X \cap H)$ *is injective.*

(3) If $\mathbb{K}$ has characteristic 0, X is a normal variety of dimension ≥ 4, and H is a general hyperplane on $\mathbb{P}^n$ then there is an isomorphism $\mathrm{Cl}(X) \to \mathrm{Cl}(X \cap H)$*.*

This theorem significantly extends the class of varieties for which we can compute the Picard group; for example, a smooth complete intersection X in $\mathbb{P}^n$ of dimension ≥ 3 will have $\mathrm{Pic}(X) \cong \mathbb{Z}$.

Example 10.4.8. When $\dim(X) = 2$ the restriction map $\mathrm{Pic}(X) \to \mathrm{Pic}(X \cap H)$ will usually not be an isomorphism. For example, a smooth curve in $\mathbb{P}^2$ will only have Picard group isomorphic to $\mathbb{Z}$ when it has degree 2 or 1 (see Example 10.5.11 and Theorem 10.5.16) and the restriction map is only an isomorphism in the latter case.

The situation is a little more delicate when $\dim(X) = 3$.

Example 10.4.9. Consider the restriction map $\mathrm{Pic}(\mathbb{P}^3) \to \mathrm{Pic}(X)$ when X is a smooth degree 4 surface in $\mathbb{P}^3$. It turns out that the restriction map is an isomorphism for very general hypersurfaces of degree 4 where "very general" means "away from a countable union of proper closed subsets". Each hypersurface X is a K3 surface, and the proof involves an analysis of the Picard groups of degree 4 K3 surfaces using Hodge theory.

This example is a special case of the Noether-Lefschetz theorem showing that the restriction map $\mathrm{Pic}(\mathbb{P}^3) \to \mathrm{Pic}(X)$ is an isomorphism when X is a very general hypersurface of degree ≥ 4.

10.4.3 Hodge theory

Our final technique for computing the Picard group comes from Hodge theory. Suppose that X is a smooth projective $\mathbb{C}$-variety. We will implicitly identify X and the holomorphic manifold obtained by taking the closed points of X. Under this identification there is a bijection between isomorphism classes of invertible sheaves on X and holomorphic line bundles on the associated holomorphic manifold.

We can associate to any line bundle $\mathcal{L}$ its first Chern class $c_1(\mathcal{L}) \in H^2(X,\mathbb{Z})$. Recall that line bundles are classified up to *topological* equivalence by their first Chern class; however, the *algebraic* equivalence of line bundles is finer. Thus we get a map

$$c_1 : \mathrm{Pic}(X) \to H^2(X,\mathbb{Z}).$$

The Lefschetz-Hodge theorem (which is a special case of the Hodge Conjecture) describes the image of this map:

Theorem 10.4.10 (Lefschetz-Hodge Theorem)**.** *Let X be a smooth projective $\mathbb{C}$-variety. A class $\alpha \in H^2(X,\mathbb{Z})$ is in the image of $c_1 : \mathrm{Pic}(X) \to H^2(X,\mathbb{Z})$ if and only if its image in $H^2(X,\mathbb{C})$ has type (1,1) under the Hodge Decomposition.*

This suggests that we can take a two-step approach to computing $\mathrm{Pic}(X)$. First, we can compute the $(1,1)$-classes in $H^2(X,\mathbb{C})$ using homological techniques. The second step is to compute the kernel of the map c_1. To execute this second step we will need the following definition:

Definition 10.4.11. An abelian variety over $\mathbb{K}$ is a projective $\mathbb{K}$-variety that carries the structure of a group scheme (as in Exercise 8.4.18).

The definition of an abelian variety is not so enlightening. It turns out that this is a very restrictive condition.

Example 10.4.12. An abelian variety of dimension 1 is the same thing as an elliptic curve – i.e. a curve with trivial tangent bundle equipped with a $\mathbb{K}$-point representing the identity element. The group law on an elliptic curve (in characteristic $\neq 2,3$) can be described by embedding the curve into $\mathbb{P}^2$ as the vanishing locus of an elliptic equation; given three closed points $p,q,r \in C$ we set $p+q+r=0$ if and only if $\mathcal{O}_C(p+q+r) \cong \mathcal{O}_{\mathbb{P}^2}(1)|_C$.

Example 10.4.13. Suppose that X is an abelian variety of dimension n over $\mathbb{C}$. It turns out that the corresponding holomorphic manifold is always isomorphic to a quotient of $\mathbb{C}^n$ by a lattice of rank $2n$. Conversely, it is clear that any quotient of $\mathbb{C}^n$ by a rank $2n$ lattice is a compact holomorphic manifold with a group structure. However, not every such quotient corresponds to a projective variety; projectivity requires an extra property of the lattice known as the “Riemann conditions”.

We are now equipped to describe the kernel of the first Chern class:

Theorem 10.4.14. *Let X be a smooth projective $\mathbb{C}$-variety. The kernel of the first Chern class map $c_1 : \mathrm{Pic}(X) \to H^2(X,\mathbb{Z})$ is an abelian variety (known as the Picard variety) which can be defined as a lattice quotient using sheaf cohomology groups:*

$$\mathrm{Pic}^0(X) = \frac{H^1(X,\mathcal{O}_X)}{H^1(X,\mathbb{Z})}.$$

10.4.4 Exercises

Exercise 10.4.15. Compute the Picard groups of:

(1) Any product of projective spaces.

(2) The Grassmannian $G(k,n)$.

(3) The complement $\mathbb{P}^n \backslash Y$ of a degree d hypersurface $Y \subset \mathbb{P}^n$. (This is related to the fact that $\pi_1(\mathbb{P}^n \backslash Y) = \mathbb{Z}/d\mathbb{Z}$.)

Exercise 10.4.16. Let X be the projective cone over a smooth quadric surface, i.e. $X \subset \mathbb{P}^4_{v,w,x,y,z}$ is defined by the ideal $(wy - xz)$. Suppose that f is a rational function on X. Explain why the intersection of X with the hyperplane $v = 0$ is isomorphic to a quadric surface Q. Explain why the pullback of f to the hyperplane $v = 0$ must define an invertible sheaf of the form $\mathcal{O}(d,d)$ on Q.

Use this to give another proof of the fact that no multiple of the prime divisor $D \subset X$ defined by the equation $w = x = 0$ will be Cartier.

Exercise 10.4.17. Find an example of a threefold X and a closed embedding $X \hookrightarrow \mathbb{P}^n$ such that the restriction map $\mathrm{Pic}(X) \to \mathrm{Pic}(X \cap H)$ is not an isomorphism for any hyperplane H such that $X \cap H$ is smooth.

10.5 Divisors on curves

In this section we will work over a fixed field $\mathbb{K}$. A curve C will mean an integral separated scheme of finite type over a field $\mathbb{K}$ with dimension 1. Thus every point of C beside the generic point will be a closed point. For a closed point p on C the following are equivalent:

(1) C is normal at p.

(2) C is regular at p.

(3) $\mathcal{O}_{C,p}$ is a DVR.

The next statement summarizes some results from Section 9.7.

Proposition 10.5.1. *Let $f : C \to Z$ be a finite morphism of regular $\mathbb{K}$-curves. Then f is flat and $f_*\mathcal{O}_C$ is locally free.*

In this setting the degree of $f : C \to Z$ is equivalently the degree $[K(C) : K(Z)]$, the $\mathbb{K}$-dimension of the Artinian ring defining any fiber of f, and the rank of the locally free sheaf $f_*\mathcal{O}_C$.

10.5.1 Weil divisors on curves

By Theorem 10.4.2 Weil divisors and Cartier divisors are the same things for regular curves. A prime divisor is just a closed point of C. Note however that not all closed points are the same; the basic invariant of a closed point x is $\dim_{\mathbb{K}} \kappa(x)$.

Definition 10.5.2. Let C be a regular proper curve. Let $D = \sum a_i p_i$ be a Weil divisor on X. We define the degree of D to be

$$\deg(D) = \sum a_i \dim_{\mathbb{K}} \kappa(p_i).$$

Note that this definition is compatible with the notion of degree introduced in Section 6.2. In order to study the degree, we will need to study how it behaves with respect to morphisms of curves.

Construction 10.5.3. Suppose that C, Z are regular proper curves and that $f : C \to Z$ is a finite morphism. Given a Weil divisor $D = \sum a_i p_i$ on Z, we define

$$f^*D := \sum a_i f^{-1}(p_i)$$

This definition is compatible with the more general notion of pullback of Cartier divisors discussed in Exercise 10.2.19.

Exercise 10.5.4. Suppose that $f : C \to Z$ is a finite morphism of curves. Prove that for any Cartier divisor L on Z our two notions of pullback coincide, i.e. we have $f^*L = f^*(\mathrm{div}(L))$.

We then have the obvious result:

Proposition 10.5.5. *Let $f : C \to Z$ be a finite morphism of regular proper curves. For any Weil divisor D on Z we have*

$$\deg(f^*D) = \deg(f) \cdot \deg(D).$$

One useful application of this proposition is to study principal divisors on curves.

Exercise 10.5.6. Let C be a regular proper curve and let $g \in K(C)$. Show that there is a finite morphism $f : C \to \mathbb{P}^1$ such that under the induced inclusion $f^\sharp : K(\mathbb{P}^1) \to K(C)$ we have $f^\sharp(\frac{x}{y}) = g$. (Hint: you can construct f explicitly on by identifying two particular open affine subsets of X and mapping them to D_{+,x_0} and D_{+,x_1}. Alternatively, you can appeal to the construction in Theorem 2.6.10.)

Theorem 10.5.7. *Let C be a regular proper curve. Then any principal divisor D on C satisfies* $\deg(D) = 0$.

Proof. Suppose that $D = \mathrm{div}(g)$. By Exercise 10.5.6 there is a finite morphism $f : C \to \mathbb{P}^1$ such that g is the pullback of $\frac{x}{y}$. This means that D is the pullback of the divisor $(0) - (\infty)$ on $\mathbb{P}^1$. By Proposition 10.5.5 we obtain the desired result. $\square$

This implies that the degree function descends to linear equivalence classes of divisors. In particular:

Definition 10.5.8. Let C be a regular proper curve. We define the degree $\deg : \mathrm{Pic}(X) \to \mathbb{Z}$ by descending the degree map on Cartier divisors. More generally, if C is a proper curve, we define $\deg(\mathcal{L})$ to be the degree of the pullback of $\mathcal{L}$ to the normalization of C.

Remark 10.5.9. If C is a Riemann surface then $H^2(C, \mathbb{Z}) \cong \mathbb{Z}$. In this setting we can identify the degree map $\deg : \mathrm{Pic}(X) \to \mathbb{Z}$ with the first Chern class map $c_1 : \mathrm{Pic}(X) \to H^2(C, \mathbb{Z})$.

Since the degree is compatible with the group operation $+$ on $\mathrm{WDiv}(C)$, it is also compatible with the group operation $\otimes$ on $\mathrm{Pic}(C)$ and so deg is a homomorphism. However, in general deg is not surjective.

Example 10.5.10. Consider the conic $x^2 + y^2 + z^2 = 0$ in $\mathbb{P}^2_{\mathbb{Q}}$. We claim that the image of $\deg : \mathrm{Pic}(C) \to \mathbb{Z}$ is $2\mathbb{Z}$.

Note that C has invertible sheaves of even degree. For example, the restriction of $\mathcal{O}_{\mathbb{P}^2}(1)$ to C will have degree 2, since a section will represent the intersection of C with a line. By taking tensor powers of this line bundle and its dual, we obtain invertible sheaves of any even degree.

We need to show that C has no line bundles of odd degree. If it did, then by tensoring by an appropriate line bundle of even degree we obtain a line bundle $\mathcal{L}$ of degree 1. By Exercise 9.1.24 we have

$$\dim_{\mathbb{Q}}(\mathcal{L}(C)) = \dim_{\overline{\mathbb{Q}}}(\mathcal{O}_{\mathbb{P}^1}(1)(\mathbb{P}^1)) \neq 0.$$

But a global section of $\mathcal{L}$ will be an effective Weil divisor of degree 1, or equivalently, a point on X with residue field $\mathbb{Q}$. But C has no $\mathbb{Q}$-points, since on any affine chart it is defined by the equation $x^2 + y^2 + 1 = 0$.

10.5.2 Examples

We next compute the Picard groups of a couple curves.

Example 10.5.11. Let $\mathbb{K}$ be an algebraically closed field of characteristic $\neq 2$. Let X be the regular curve $y^2z = x^3 - xz^2$ in $\mathbb{P}^2$. (This is an example of an elliptic curve.) Since X is regular, we have $\mathrm{Pic}(X) \cong \mathrm{Cl}(X)$ and we will compute the latter group. Note that any Weil divisor on X is a formal sum of closed points and (since $\mathbb{K}$ is algebraically closed) every closed point is isomorphic to $\mathrm{Spec}(\mathbb{K})$.

First, let's observe that our UFD trick no longer works. For example, the intersection of X with the affine chart $D_{+,z}$ is the affine scheme defined by $R = \mathbb{K}[x,y]/(y^2 - x^3 + x)$. Since X is regular, R is an integrally closed domain. We claim that the units in R are just the elements of $\mathbb{K}^\times$. Indeed, let σ denote the involution of R defined by $y \leftrightarrow -y$. For any $r \in R$ we define $N(u) = r \cdot \sigma(r)$. If u is a unit in R, then $N(u)$ is also a unit. However, the subring of R fixed by the involution is simply $\mathbb{K}[x]$. Thus we see that $N(u) \in \mathbb{K}^\times$. It's easy to see directly that this implies that $u \in \mathbb{K}^\times$.

We are now equipped to show that R is not a UFD. It suffices to show that the height 1 prime ideal (x, y) is not principal. Indeed, since the units in R are all contained in $\mathbb{K}^\times$, for degree reasons both x and y are irreducible elements of R and there is no common divisor of both.

Now we return to the computation of $\mathrm{Cl}(X)$. Consider the degree morphism $\deg : \mathrm{Cl}(X) \to \mathbb{Z}$ sending

$$\deg\left(\sum_i a_i p_i\right) = \sum a_i.$$

Let $\mathrm{Cl}^0(X)$ denote the kernel of this map. We will show that there is a bijection *alb* from the closed points of X to $\mathrm{Cl}^0(X)$. (In particular, this will impose a group structure on the points of X, the famous "group law" for an elliptic curve.)

Let o denote the point $(0 : 1 : 0)$ in X. We define the map *alb* by associating to any $p \in X$ the divisor $p - o \in \mathrm{Cl}^0(X)$. We first show that this map is injective. Indeed, suppose that $p, q \in X$ are different points which satisfy $p - o \sim q - o$. This implies that $p \sim q$. In other words, there is a rational function f on X such that $\operatorname{div}(f) = p - q$. Then f defines a morphism $f : X \to \mathbb{P}^1$ such that $f^{-1}(0) = p$ and $f^{-1}(\infty) = q$. This f will be a bijective finite flat morphism of degree 1, hence will be an isomorphism. But X is not isomorphic to $\mathbb{P}^1$. (For example, every open affine subset of $\mathbb{P}^1$ is defined by a UFD.)

Finally, we show that the map $alb : X \to \mathrm{Cl}^0(X)$ is surjective. Note that the line $z = 0$ is triply tangent to X at the point o. Thus, if any line ℓ meets X along the points p, q, r we have $p + q + r \sim 3o$, or equivalently. Similarly, if we take the line connecting r with o and let r' denote the third point of intersection with X, we see that $r + o + r' \sim 3o$. Altogether we have

$$(p - o) + (q - o) \sim -(r - o) \sim (r' - o). \tag{10.5.1}$$

Now, let $D = \sum a_i p_i$ be any divisor in the kernel of deg. Since $\sum a_i = 0$, we have $D \sim \sum a_i(p_i - o)$. Using the two relations in Equation (10.5.1) repeatedly, we can replace D by a linearly equivalent divisor of the form $q - o$. This proves surjectivity of the map $X \to \mathrm{Cl}^0(X)$.

Remark 10.5.12. The map *alb* used in Example 10.5.11 is the Albanese map that is often studied in a course on Riemann surfaces.

Remark 10.5.13. A posteriori we see that there is no open affine U in the elliptic curve X of Example 10.5.11 such that $\mathcal{O}_X(U)$ is a UFD. Indeed, if there were then by Theorem 10.4.4 the class group of X would be a finitely generated abelian group.

Our next example is similar: we look at a cuspidal plane cubic. Although the curve is birational to $\mathbb{P}^1$, the presence of the singularity forces its Picard group to look more similar to the Picard group of an elliptic curve.

Example 10.5.14. Let $\mathbb{K}$ be an algebraically closed field of characteristic $\neq 2$. Let X be the cuspidal curve $y^2 z = x^3$ in $\mathbb{P}^2$. We claim that $\mathrm{Pic}(X)$ fits into an exact sequence

$$0 \to \mathbb{K}^\times \to \mathrm{Pic}(X) \to \mathbb{Z} \to 0.$$

Since X is not normal, we cannot compute the Picard group of X using Weil divisors. However, we know that $\mathrm{CaCl}(X) \cong \mathrm{Pic}(X)$ and thus we can use Cartier divisors instead.

Given any Cartier divisor L on X, its pullback under the normalization map $\nu : \mathbb{P}^1 \to X$ will be a Cartier divisor on $\mathbb{P}^1$. In this way we can define a degree morphism deg : $\mathrm{CaCl}(X) \to \mathbb{Z}$, and we would like to compute the kernel $\mathrm{CaCl}^0(X)$ of this map.

Let s denote the singular point $(0 : 0 : 1)$ of X. Note that $X \backslash \{s\}$ is isomorphic to $\mathbb{P}^1 \backslash \{0\} \cong \mathbb{A}^1$. Let o denote the point $(0 : 1 : 0)$ of X. To each closed point $p \in X^{sm}$ we associate the Cartier divisor which associates to the open set $X^{sm} \cong \mathbb{A}^1$ the ratio of linear functions with a zero at p and a pole at o, and which is the constant function 1 on a small open neighborhood of s. We claim that this map $a : X^{sm} \to \mathrm{CaCl}^0(X)$ is a bijection.

First we show a is injective. Suppose that $p, q \in X^{sm}$ are different points such that $p - o \sim q - o$. Then we see that $p \sim q$, and as in Example 10.5.11 we obtain a birational morphism $f : X \to \mathbb{P}^1$. Although these two curves are birational, there is no birational morphism in this direction, yielding a contradiction.

Next we show that a is surjective. It is clear that any Cartier divisor is linearly equivalent to a divisor whose local function on a neighborhood of s is invertible. Thus we can associate to any class in $\mathrm{CaCl}(X)$ a (not uniquely defined) Weil divisor on X^{sm}. We then mimic the argument of Example 10.5.11. The only additional subtlety is that we must verify that any line through two points of X^{sm} cannot go through s; this is because s is singular so the local intersection number of any line through s with X will be at least 2.

Remark 10.5.15. Example 10.5.14 implicitly gives a group structure on the open set $X \backslash \{s\}$. This corresponds to the standard group structure $\mathbb{G}_a$ on the isomorphic scheme $\mathbb{A}^1$.

10.5.3 Jacobians

In Example 10.5.11 we saw that for a particular elliptic curve X the kernel of the degree map could be associated with the closed points of an algebraic variety (in this case X itself). It turns out that that this statement holds for every "nice" curve.

Theorem 10.5.16. *Let C be a smooth, projective, geometrically integral $\mathbb{K}$-curve that admits a $\mathbb{K}$-point. Then the set of invertible sheaves of degree 0 on X are parametrized by the $\mathbb{K}$-points of an abelian variety known as the Jacobian of C. In fact, for any field extension $\mathbb{L}/\mathbb{K}$ there is a bijection between degree 0 invertible sheaves on $X_{\mathbb{L}}$ and $\mathbb{L}$-points of the Jacobian of C.*

Using the identification of the degree map with the first Chern class, Theorem 10.5.16 is closely related to Theorem 10.4.14. This result gives an essentially complete description of $\mathrm{Pic}(C)$: it is a disjoint union of copies of $\mathrm{Jac}(C)$ indexed by the degree $\mathbb{Z}$.

Example 10.5.17. Suppose that C is an elliptic curve equipped with a $\mathbb{K}$-point. Then $\mathrm{Pic}(\mathbb{Z}) \cong C \times \mathbb{Z}$ and the Jacobian of C is isomorphic to C itself. We checked this by hand for a particular elliptic curve in Example 10.5.11.

10.5.4 Numerical equivalence

Let X be a projective $\mathbb{K}$-variety of arbitrary dimension. Given any curve $C \subset X$ and any invertible sheaf $\mathcal{L}$ on X, we define the intersection product

$$\mathcal{L} \cdot C := \deg(\nu^* \mathcal{L})$$

where $\nu : C^\nu \to X$ is the normalization of C. In this way we obtain for each curve C a homomorphism $\mathrm{Pic}(X) \to \mathbb{Z}$.

Definition 10.5.18. Let X be a projective $\mathbb{K}$-variety. We say that two line bundles $\mathcal{L}_1, \mathcal{L}_2$ are numerically equivalent, and write $\mathcal{L}_1 \equiv \mathcal{L}_2$, if for every curve C we have

$$\mathcal{L}_1 \cdot C = \mathcal{L}_2 \cdot C.$$

When X is a curve, then two line bundles are numerically equivalent if and only if they have the same degree. Thus one can think of numerical equivalence as a "generalization" of the degree map to higher dimension.

Note that the set of line bundles that are numerically equivalent to the trivial line bundle $\mathcal{O}_X$ is a subgroup of $\operatorname{Pic}(X)$. We will denote the quotient of $\operatorname{Pic}(X)$ by this subgroup by $N^1(X)$.

Theorem 10.5.19. *Let X be a projective variety over an algebraically closed field $\mathbb{K}$. Then:*

(1) *The group $N^1(X)$ of invertible sheaves up to numerical equivalence is a free finitely generated abelian group.*

(2) *Suppose that X is regular. Then the kernel of the map $\operatorname{Pic}(X) \to N^1(X)$ is a disjoint union of abelian varieties indexed by a finite group.*

Loosely speaking, the map $\operatorname{Pic}(X) \to N^1(X)$ is similar to taking a line bundle and sending it to its first Chern class. In other words, the space $N^1(X)$ acts like a "homology space" for invertible sheaves. The following example clarifies this loose analogy for $\mathbb{C}$-varieties.

Example 10.5.20. Let X be a smooth projective $\mathbb{C}$-variety. Then the quotient map $\operatorname{Pic}(X) \to N^1(X)$ factors through the first Chern class map $c_1 : \operatorname{Pic}(X) \to H^2(X, \mathbb{Z})$. In fact, it turns out that $N^1(X)$ is obtained by taking the image of c_1 and quotienting out by the torsion subgroup.

10.5.5 Exercises

Exercise 10.5.21. Let C be a projective $\mathbb{K}$-curve equipped with a closed embedding $f : C \to \mathbb{P}^n$. Show that the degree of $f^*\mathcal{O}(1)$ is the same as the degree of C as defined in Section 6.2.

Exercise 10.5.22. Let $X = \mathbb{P}^1 \times \mathbb{P}^1$. Prove that the map $\operatorname{Pic}(X) \to N^1(X)$ is an isomorphism.

Exercise 10.5.23. Let X be a projective $\mathbb{K}$-variety. Prove that $\operatorname{Pic}(X)$ is not a torsion group. (Hint: first use Exercise 10.5.21 to show that it is non-zero.)

Exercise 10.5.24. Let X be a projective $\mathbb{K}$-variety. Suppose that $\mathcal{L}$ is a basepoint free invertible sheaf defining a morphism $f : X \to \mathbb{P}^n$. Show that a closed curve $C \subset X$ is contracted by f if and only if $\deg(\mathcal{L}|_C) = 0$.

Exercise 10.5.25. Suppose that $f : C \to D$ is a finite morphism of regular $\mathbb{K}$-curves. Suppose we fix a point $p \in D$ and let $\{q_i\}_{i=1}^s$ be the points in $f^{-1}p$. We know that $f^{-1}p$ is defined by an Artinian ring over $\kappa(p)$ of length $\deg(f)$. In this exercise we compute the local contributions of the various points q_i to the length.

(1) Suppose we fix a uniformizer t of $\mathcal{O}_{D,p}$ (that is, a rational function on D which generates the maximal ideal $\mathfrak{m}_p \subset \mathcal{O}_{D,p}$). Fix a point q_i in the fiber over p. Show that the valuation $\mathrm{val}_{q_i}(f^*t)$ is the same as the integer e such that $f_p^\sharp(\mathfrak{m}_p) = \mathfrak{m}_q^e$ under the map $f_p^\sharp : \mathcal{O}_{D,p} \to \mathcal{O}_{C,q_i}$.

(2) Show that we have
$$\deg(f) = \sum_{i=1}^{s} \mathrm{val}_{q_i}(f^*t)\deg(\kappa(q_i)/\kappa(p)).$$

Of course this same computation works whenever f is a morphism of schemes which are locally defined by Dedekind domains. In particular, there is a similar formula using inertia degrees that controls the splitting of primes in rings of integers.

Exercise 10.5.26. Let X be the nodal cubic $zy^2 = x^3 - x^2$ in $\mathbb{P}^2_{\mathbb{K}}$ (where $\mathrm{ch}(\mathbb{K}) \neq 2, 3$) and let $f : \mathbb{P}^1_{\mathbb{K}} \to X$ denote the normalization map. Show that the data of an invertible sheaf $\mathcal{L}$ on X is the same as the data of the pullback $f^*\mathcal{L}$ on $\mathbb{P}^1$ and an isomorphism of fibers $\phi : \mathcal{L}((x-1)) \xrightarrow{\cong} \mathcal{L}((x+1))$. Use this identification to prove that $\mathrm{Pic}(X)$ fits into an exact sequence
$$0 \to \mathbb{K}^\times \to \mathrm{Pic}(X) \to \mathbb{Z} \to 0.$$

10.6 Ample line bundles

Let X be a complex manifold and let $\mathcal{L}$ be a holomorphic line bundle on X. The line bundles $\mathcal{L}$ which have positive curvature – that is, whose Chern class is represented by a Kähler metric – have particularly nice properties. In this section we develop the analogous notion of a "positive line bundle" in algebraic geometry – an ample invertible sheaf.

It turns out that the "positive" line bundles are also the ones which have the best behavior with regards to global generation. In general the global generation of a sheaf is a subtle issue, but ample invertible sheaves give us a systematic way of finding many globally generated sheaves. Thus ampleness plays a key role in the theory of schemes as developed by Serre and Grothendieck.

10.6.1 Very ample line bundles

Definition 10.6.1. Let X be a proper scheme over a Noetherian ring A and let $\mathcal{L}$ be an invertible sheaf on X. We say that $\mathcal{L}$ is very ample if there exists a closed embedding $f : X \to \mathbb{P}^n_A$ such that $\mathcal{L} \cong f^*\mathcal{O}(1)$. Equivalently, $\mathcal{L}$ is very ample if there is a finite set of global sections $\{s_0, \dots, s_n\}$ which generate $\mathcal{L}$ such that the induced morphism $X \to \mathbb{P}^n_A$ is a closed embedding.

The main example of very ample invertible sheaves comes from the $\mathcal{O}(d)$ construction.

Proposition 10.6.2. *Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is finitely generated in degree* 1 *over a Noetherian ring S_0. Then $\mathcal{O}_{\mathrm{Proj}(S)}(d)$ is a very ample invertible sheaf on* $\mathrm{Proj}(S)$ *for every $d > 0$.*

Proof. Proposition 2.7.7 shows that $\mathrm{Proj}(S) \cong \mathrm{Proj}(S^{(d)})$. Furthermore $S^{(d)}$ is still finitely generated in degree 1 over S_0 so that there is a surjection $S_0[x_0, \dots, x_n] \to S^{(d)}$. This induces a closed embedding $\mathrm{Proj}(S) \hookrightarrow \mathbb{P}^n_{S_0}$ and the pullback of $\mathcal{O}_{\mathbb{P}^n_{S_0}}(1)$ under this map is $\mathcal{O}_{\mathrm{Proj}(S)}(d)$. □

Using the relative Segre embedding, we see that very ample invertible sheaves are compatible with tensor products:

Lemma 10.6.3. *Let X be a proper scheme over a Noetherian ring A. Suppose that $\mathcal{L}$ is a very ample invertible sheaf on X and that $\mathcal{M}$ is a globally generated invertible sheaf on X. Then $\mathcal{L} \otimes \mathcal{M}$ is again a very ample invertible sheaf.*

Note that if $\mathcal{L}$ is very ample, then in particular it must be globally generated so that this lemma also shows that very ampleness is preserved by tensor products.

Proof. Let $f : X \to \mathbb{P}^n_A$ be the closed embedding defined by $\mathcal{L}$ and $g : X \to \mathbb{P}^m$ the morphism defined by $\mathcal{M}$. Then the induced map $(f, g) : X \to \mathbb{P}^n_A \times_{\mathrm{Spec}(A)} \mathbb{P}^m_A$ is a closed embedding by Proposition 8.6.6. Under the Segre embedding $\mathbb{P}^n_A \times_{\mathrm{Spec}(A)} \mathbb{P}^m_A \to \mathbb{P}^{nm+n+m}_A$

the pullback of $\mathcal{O}_{\mathbb{P}_A^{nm+n+m}}(1)$ is equal to $\pi_1^*\mathcal{O}_{\mathbb{P}_A^n}(1) \otimes \pi_2^*\mathcal{O}_{\mathbb{P}_A^m}(1)$. Thus the pullback of $\mathcal{O}_{\mathbb{P}_A^{nm+n+m}}(1)$ under the closed embedding $X \to \mathbb{P}_A^{nm+n+m}$ is isomorphic to $\mathcal{L} \otimes \mathcal{M}$. $\square$

10.6.2 Ample line bundles

Many theorems involving very ample line bundles require us to first take a high tensor product. (We have seen the first instances of this principle in the fact that if S is a finitely generated S_0-algebra then some Veronese subring will be generated in degree 1.) It is thus very natural to expand our attention to a more general class of invertible sheaves.

Definition 10.6.4. Let X be a proper scheme over a Noetherian ring A. We say that an invertible sheaf $\mathcal{L}$ on X is ample if $\mathcal{L}^{\otimes n}$ is very ample for some positive integer n.

Warning 10.6.5. An ample invertible sheaf need not be globally generated, and in fact, need not have any sections at all; see Example 12.6.8.

Remark 10.6.6. Kodaira's Theorem shows that ample invertible sheaves are the algebraic analogue of "positive" line bundles in complex geometry.

The following fundamental theorem shows that ample divisors can be characterized in several different ways.

Theorem 10.6.7. *Let X be a proper scheme over a Noetherian ring A. Suppose $\mathcal{L}$ is an invertible sheaf on X. The following are equivalent:*

(1) $\mathcal{L}$ is ample.

(2) There is some constant M such that $\mathcal{L}^{\otimes m}$ is very ample for all $m > M$.

(3) For every finitely generated quasicoherent sheaf $\mathcal{F}$ there is a constant M such that $\mathcal{L}^{\otimes m} \otimes \mathcal{F}$ is globally generated for every $m \geq M$.

(4) As we vary over all positive integers n and all global sections $s \in \mathcal{L}^{\otimes n}(X)$ the open sets $X \backslash Z(s)$ form a base for the topology on X.

Note that even if $\mathcal{L}$ is very ample, we cannot avoid using the exponent m in Conditions (3) and (4).

Remark 10.6.8. Let us briefly explain why we might expect Condition (4) to hold. Recall that the complement of every hypersurface $V_+(f)$ in $\mathbb{P}_A^n$ is affine and that as we vary f these distinguished open affines form a base for the topology. If $f : X \to \mathbb{P}_A^n$ is a closed embedding, then the intersections of these open sets with X will again be open affines which form a base for the topology. We can interpret these open sets as complements of certain sections of the line bundles $f^*\mathcal{O}(m)$. Condition (4) says that this is a defining property of an ample line bundle.

Proof. (1) $\Rightarrow$ (3): We first prove this under the additional assumption that $\mathcal{L}$ is very ample. Let $f : X \to \mathbb{P}^n_A$ be a closed embedding defined by $\mathcal{L}$. Since f is a closed embedding, $f_*\mathcal{F}$ is a coherent sheaf on $\mathbb{P}^n_A$. It suffices to show that $(f_*\mathcal{F})(m)$ is globally generated for sufficiently large m; indeed, the global sections of $\mathcal{F} \otimes \mathcal{L}^{\otimes m}$ and of $(f_*\mathcal{F})(m)$ are the same.

Consider the standard open cover $\{D_{+,x_i}\}$ of $\mathbb{P}^n_A$. For every i there is a finitely generated $\mathcal{O}_X(U)$-module M_i such that $\mathcal{F}|_{D_{+,x_i}} \cong \widetilde{M_i}$. Choose a generating set $\{t_{ij}\}$ of M_i. Since X is quasicompact quasiseparated, Exercise 9.5.28 shows that for every pair of indices i, j there some exponent k such that the section $x_i^k t_{ij}$ extends to a global section of $\mathcal{F}$. Letting m_0 denote the supremum over all such powers k, we find a collection of sections $\widetilde{t}_{ij} \in (i_*\mathcal{F})(m_0)(\mathbb{P}^n_A)$ such that the restriction of $\widetilde{t}_{ij}$ to D_{+,x_i} is a power of x_i times t_{ij}. It is clear that for any $m \geq m_0$ the sections $x_i^{m-m_0}\widetilde{t}_{ij}$ generate $(i_*\mathcal{F})(m)$, or equivalently, generate $\mathcal{F} \otimes \mathcal{L}^{\otimes m}$.

To prove the statement when $\mathcal{L}$ is only ample, choose an integer N such that $\mathcal{L}^{\otimes N}$ is very ample. We apply the argument above to the very ample sheaf $\mathcal{L}^{\otimes N}$ and to the coherent sheaves $\mathcal{F} \otimes \mathcal{L}^{\otimes i}$ for $i = 1, 2, \ldots, N-1$. For each i, we find a power m_i of $\mathcal{L}^{\otimes n}$ that guarantees global generation of $\mathcal{F} \otimes \mathcal{L}^{\otimes i+mN}$ for $m \geq m_i$. We then set $M = \sup\{i + m_i N\}$.

(3) $\Rightarrow$ (4): We want to show that for any open set $U \subset X$ and any point $x \in U$ there is an open neighborhood of x contained in U defined by a section of $\mathcal{L}^{\otimes m}$. Let $Z = X \backslash U$ denote the closed subscheme of X with the reduced structure. Since X is Noetherian the ideal sheaf $\mathcal{I}_Z$ is a coherent sheaf on X. Thus there is some positive integer m such that $\mathcal{I}_Z \otimes \mathcal{L}^{\otimes m}$ is globally generated. Choose a section s which does not vanish at p. Then the image of s under the inclusion $(\mathcal{I}_Z \otimes \mathcal{L}^{\otimes m})(X) \to \mathcal{L}^{\otimes m}(X)$ has the desired property.

(4) $\Rightarrow$ (1): We first note that we can find a base of the topology consisting of affine open subsets of the form $X \backslash Z(s)$. Indeed, if U is an open affine subset of X and $X \backslash Z(s) \subset U$ then $X \backslash Z(s) \cong U \backslash Z(s)$ is also affine by Exercise 10.2.21.

Choose a finite open cover of X by open affine sets of the form $U_i := X \backslash Z(s_i)$. By replacing each s_i by a power, we may suppose that all the s_i are sections of the same power $\mathcal{L}^{\otimes N}$. Since U_i is affine and X is finite type, the ring $\mathcal{O}_X(U_i)$ is finitely generated as an A-algebra. Furthermore, each generator has the form $t_{ij}/s_i^{r_{ij}}$ for some positive integer r_{ij} and some $t_{ij} \in \mathcal{L}^{\otimes N r_{ij}}$. Set $R = \sup_{ij} r_{ij}$. Using the relation $t_{ij}/s_i^{r_{ij}} = t_{ij} s_i^{R - r_{ij}}/s_i^R$, we see that $\mathcal{O}_X(U_i)$ is generated by elements of the form $\widetilde{t_{ij}}/s_i^R$ where $\widetilde{t_{ij}} \in \mathcal{L}^{\otimes NR}$.

We claim that $\mathcal{L}^{\otimes NR}$ is very ample. Since s_i^R is a section, we see that $\mathcal{L}^{\otimes NR}$ is globally generated, yielding a morphism $f : X \to \mathbb{P}^n_A$. If x_i is the coordinate on $\mathbb{P}^n_A$ that corresponds to s_i^R, then the map $U_i \to D_{+,x_i}$ is a closed embedding since it corresponds to a surjection of the defining rings. Since X is proper, this implies that f is a closed embedding.

(1) + (3) $\Leftrightarrow$ (2): suppose that $\mathcal{L}^{\otimes m}$ is very ample and that $\mathcal{L}^{\otimes n}$ is globally generated for every $n \geq n_0$. Then by Lemma 10.6.3 $\mathcal{L}^{\otimes q}$ is very ample for every $q \geq m + N$. For the reverse direction, (2) $\implies$ (1) is clear and thus (2) $\implies$ (3) by the argument above. $\square$

Remark 10.6.9. We used the Noetherian assumption in the implication (3) $\implies$ (4) to see that $\mathcal{I}_Z$ is coherent. The statement is still true without the Noetherian assumption but the argument is harder.

Remark 10.6.10. If X is any quasicompact quasiseparated scheme, a similar argument shows that the following conditions on an invertible sheaf $\mathcal{L}$ are equivalent:

(1) Set $S = \oplus_{n\geq 0}\Gamma(X, \mathcal{L}^{\otimes n})$. Then as we vary $s \in \Gamma(X, \mathcal{L}^{\otimes n})$ the corresponding open sets $X\backslash Z(s)$ cover X and the induced map $X \to \mathrm{Proj}(S)$ is an open embedding.

(2) As we vary over all positive integers n and all sections $s \in \mathcal{L}^{\otimes n}(X)$, the open sets $X\backslash Z(s)$ form a base of the topology on X.

(3) For every finitely generated quasicoherent sheaf $\mathcal{F}$ there is a constant N such that $\mathcal{L}^{\otimes n} \otimes \mathcal{F}$ is globally generated for every $n \geq N$.

We use these three conditions to define ampleness in this more general setting. For example, every invertible sheaf on an affine scheme is ample under this definition (since every finitely generated quasicoherent sheaf is globally generated).

Exercise 10.6.11. Let X be a proper scheme over a Noetherian ring A and suppose that $\mathcal{L}$ is an ample invertible sheaf. Prove that for any invertible sheaf $\mathcal{M}$ there is a constant N such that $\mathcal{L}^n \otimes \mathcal{M}$ is very ample for every $n \geq N$.

Use this fact to deduce that if X carries an ample invertible sheaf then every invertible sheaf $\mathcal{M}$ is isomorphic to $\mathcal{A}_1 \otimes \mathcal{A}_2^\vee$ for some very ample invertible sheaves $\mathcal{A}_1, \mathcal{A}_2$.

10.6.3 Serre twisting sheaves

As we saw in Section 9.6, the invertible sheaves $\mathcal{O}_{\mathbb{P}^n}(d)$ play a key role in the theory of coherent sheaves on projective space via Hilbert's Syzygy Theorem. Although we do not obtain similarly strong statements for arbitrary projective schemes, these line bundles still play a crucial role.

Proposition 10.6.12 (Serre's Theorem A)**.** *Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is finitely generated in degree* 1 *over a Noetherian ring S_0. Let $\mathcal{F}$ be any finitely generated quasicoherent sheaf on* Proj(S)*. Then there is some positive integer n such that $\mathcal{F}(n)$ is globally generated by a finite set of global sections.*

The only reason we insist that S_0 be Noetherian is because our definition of ampleness only holds for proper schemes over Noetherian rings. If we instead use Remark 10.6.10, then we can drop this condition (both here and in the following corollary).

Proof. Follows from Proposition 10.6.2 and Theorem 10.6.7. □

We will often leverage this result via the following consequence:

Corollary 10.6.13. *Let S be a $\mathbb{Z}_{\geq 0}$-graded ring that is finitely generated in degree 1 over a Noetherian ring S_0. Let $\mathcal{F}$ be any finitely generated quasicoherent sheaf on* $\mathrm{Proj}(S)$*. Then there is a surjection*

$$\mathcal{O}(d)^{\oplus r} \to \mathcal{F}$$

for some integer d.

Proof. By Serre's Theorem A we know that $\mathcal{F}(d)$ is globally generated for some n. This yields a surjection $\mathcal{O}_{\mathrm{Proj}(S)}^{\oplus r} \to \mathcal{F}(d)$, hence a surjection $\mathcal{O}(-d)^{\oplus r} \to \mathcal{F}(n)$. □

10.6.4 Exercises

Exercise 10.6.14. Suppose we have a finite morphism $f : X \to Y$ of proper schemes over a Noetherian ring A. Suppose that $\mathcal{L}$ is an ample invertible sheaf on Y. Prove that $f^*\mathcal{L}$ is an ample invertible sheaf on X.

Exercise 10.6.15. Let X be a proper scheme over a Noetherian ring A. Let $\mathcal{L}$ be an ample invertible sheaf and let $\mathcal{M}$ be an invertible sheaf that is either ample or globally generated. Prove that $\mathcal{L} \otimes \mathcal{M}$ is ample.

Exercise 10.6.16. Let X be a proper $\mathbb{K}$-scheme. Suppose that there exists an invertible sheaf $\mathcal{L}$ such that both $\mathcal{L}$ and $\mathcal{L}^{\vee}$ are ample. Prove that X has dimension 0.

10.7 Relative Proj

Suppose that X is a scheme and that $\mathcal{A}$ is a quasicoherent sheaf of $\mathcal{O}_X$-algebras. In Construction 9.5.8 we discussed the relative Spec construction: an affine morphism $\pi : \underline{\mathrm{Spec}}(\mathcal{A}) \to X$ obtained locally by the Spec construction. In particular, when $\mathcal{A} = \mathrm{Sym}(\mathcal{F}^\vee)$ for a locally free sheaf $\mathcal{F}$, the relative Spec of $\mathcal{A}$ yielded a vector bundle over X. In this section we will discuss the analogous projective construction.

Definition 10.7.1. Let X be a scheme. We will say that $\mathcal{A}$ is a graded $\mathcal{O}_X$-algebras to mean more precisely that:

(1) $\mathcal{A} = \oplus_{n\geq 0}\mathcal{A}_n$ is a graded quasicoherent $\mathcal{O}_X$-module,

(2) $\mathcal{A}$ is a sheaf of rings such that the $\mathcal{O}_X$-module structure gives $\mathcal{A}$ the structure of a $\mathbb{Z}_{\geq 0}$-graded $\mathcal{O}_X$-algebra,

(3) $\mathcal{A}_0 \cong \mathcal{O}_X$.

We will define the relative Proj construction by taking the usual Proj construction over open affines in X and gluing. The following exercise will allow us to glue:

Exercise 10.7.2. Let S be a $\mathbb{Z}_{\geq 0}$-graded ring. Suppose we fix an element f in S_0. Then S_f is also a $\mathbb{Z}_{\geq 0}$-graded ring.

(1) Prove that $\mathrm{Proj}(S_f)$ admits an open embedding into $\mathrm{Proj}(S)$ that realizes it as the pullback $D_f \times_{\mathrm{Spec}(S_0)} \mathrm{Proj}(S)$.

(2) Prove that the restriction of the quasicoherent sheaf $\mathcal{O}_{\mathrm{Proj}(S)}(d)$ to $\mathrm{Proj}(S_f)$ is isomorphic to $\mathcal{O}_{\mathrm{Proj}(S_f)}(d)$.

Construction 10.7.3. Let X be a scheme equipped with a quasicoherent sheaf of $\mathbb{Z}_{\geq 0}$-graded $\mathcal{O}_X$-algebras $\mathcal{A}$. For every open affine $U \subset X$, consider the scheme $\mathrm{Proj}(\mathcal{A}(U))$ equipped with the structure map to U. As we vary the open affine U, Exercise 10.7.2 (and Nike's Lemma) show that the resulting schemes can naturally be glued to obtain a scheme $\underline{\mathrm{Proj}}(\mathcal{A})$ equipped with a morphism to X. Furthermore, $\underline{\mathrm{Proj}}(\mathcal{A})$ comes equipped with quasicoherent sheaves obtained by gluing the local $\mathcal{O}(d)$; we denote this sheaf by $\mathcal{O}_{\underline{\mathrm{Proj}}(\mathcal{A})/X}(d)$.

Definition 10.7.4. From now on we will impose the following simplifying conditions on $\mathcal{A}$:

(1) for every open affine U the ring $\mathcal{A}(U)$ is a finitely generated graded $\mathcal{O}_X(U)$-algebra that is generated in degree 1.

We will call this condition "Condition (*)". In this case the sheaf $\mathcal{O}_{\underline{\mathrm{Proj}}(\mathcal{A})/X}(1)$ will be an invertible sheaf.

Proposition 10.7.5. *Let X be a scheme and $\mathcal{A}$ is a graded $\mathcal{O}_X$-algebra that satisfies Condition (*). Then the map $\underline{\mathrm{Proj}}(\mathcal{A}) \to X$ is proper.*

Proof. Since proper is local on the target, this reduces to the fact that for a $\mathbb{Z}_{\geq 0}$-graded ring that is finitely generated in degree 1 the structural map $\mathrm{Proj}(S) \to \mathrm{Spec}(S_0)$ is proper. □

Exercise 10.7.6. Let X be a scheme and let $\mathcal{A}$ be a graded $\mathcal{O}_X$-algebra that satisfies Condition (*). In particular, this means that there is a surjection of graded $\mathcal{O}_X$-algebras $g^\sharp : \mathrm{Sym}(\mathcal{A}_1) \to \mathcal{A}$. Prove that ϕ induces a closed embedding $g : \underline{\mathrm{Proj}}(\mathcal{A}) \to \underline{\mathrm{Proj}}(\mathrm{Sym}(\mathcal{A}_1))$ that commutes with the structure maps and such that $g^*\mathcal{O}(1) \cong \mathcal{O}(1)$.

The following exercises describes the universal property of $\underline{\mathrm{Proj}}$. Note that it is the "relative" version of the universal property of Proj.

Exercise 10.7.7. Let X be a scheme and let $\mathcal{A}$ be a graded $\mathcal{O}_X$-algebra that satisfies Condition (*). Let $g : Y \to X$ be any morphism. Prove that there is a bijection between morphisms $f : Y \to \underline{\mathrm{Proj}}(\mathcal{A})$ such that $g = \pi \circ f$ and equivalence classes of pairs $(\mathcal{L}, \phi)$ where $\mathcal{L}$ is an invertible sheaf on Y and $\phi : g^*\mathcal{A} \to \mathcal{L}$ is a surjection. (Two such pairs $(\mathcal{L}, \phi)$ and $(\mathcal{L}', \phi')$ are said to be equivalent if there is an isomorphism $\psi : \mathcal{L} \to \mathcal{L}'$ making a commutating diagram with the maps ϕ, ϕ'.)

Suppose that $\mathcal{A}$ satisfies Condition (*). Let $\mathcal{L}$ be an invertible sheaf on X. Then we can define a new graded algebra $\mathcal{A}_\mathcal{L}$ via the description

$$\mathcal{A}_{\mathcal{L},n} := \mathcal{A}_n \otimes \mathcal{L}^{\otimes n}$$

with product induced by the multiplications on $\mathcal{A}$ and by the tensor product on the $\mathcal{L}^{\otimes n}$. Note that $\mathcal{A}_\mathcal{L}$ also satisfies Condition (*).

Theorem 10.7.8. *Let X be a scheme, let $\mathcal{A}$ be a sheaf of $\mathcal{O}_X$-algebras satisfying Condition (*), and let $\mathcal{L}$ be an invertible sheaf on X. Then there is an isomorphism*

$$\phi : \underline{\mathrm{Proj}}(\mathcal{A}) \xrightarrow{\cong} \underline{\mathrm{Proj}}(\mathcal{A}_\mathcal{L})$$

such that $\phi^*\mathcal{O}_{\underline{\mathrm{Proj}}(\mathcal{A}_\mathcal{L})/X}(1) \cong \mathcal{O}_{\underline{\mathrm{Proj}}(\mathcal{A})/X}(1) \otimes \pi^*\mathcal{L}$.

In other words, the two schemes are abstractly isomorphic but the relative $\mathcal{O}(1)$ is modified by tensoring by the pullback of $\mathcal{L}$. This may feel counterintuitive – the $\underline{\mathrm{Spec}}$s of the two algebras certainly need not be the same. The point is that the transition functions for the invertible sheaf $\mathcal{L}$ are just rescaling by a local unit and this rescaling "drops out" when computing the projectivization (but is remembered by the homogeneous coordinate ring).

Proof. Let U be any open affine subset in X such that we have an isomorphism $\psi : \mathcal{L}|_U \xrightarrow{\cong} \mathcal{O}_X|_U$. If we choose a different isomorphism ψ', then there is a unit $f \in \mathcal{O}_X(U)$ such that we have a commuting diagram

$$\begin{tikzcd} \mathcal{L}|_U \arrow[r, "\psi"] \arrow[dr, "\psi'"'] & \mathcal{O}_X|_U \arrow[d, "\times f"] \\ & \mathcal{O}_X|_U \end{tikzcd}$$

These two maps ψ, ψ' induce isomorphisms $\psi_\bullet, \psi'_\bullet : \mathcal{A}_\mathcal{L}|_U \to \mathcal{A}|_U$ and we obtain a commuting diagram using the isomorphism $\mathcal{A}|_U \to \mathcal{A}|_U$ which is multiplication by f^n on the nth graded piece. Taking Proj, we get maps

$$\begin{tikzcd} \mathrm{Proj}(\mathcal{A}_\mathcal{L}|_U) & \mathrm{Proj}(\mathcal{A}|_U) \arrow[l, "\psi_\bullet"'] \\ & \mathrm{Proj}(\mathcal{A}|_U) \arrow[u, "id"] \arrow[ul, "\psi'_\bullet"] \end{tikzcd}$$

The key point is that the upwards map is the identity: multiplying by a degree 0 unit does not change the homogeneous primes or the structure sheaves. This shows that the gluing data for $\underline{\mathrm{Proj}}(\mathcal{A}_\mathcal{L})$ can be canonically identified with the gluing data for $\underline{\mathrm{Proj}}(\mathcal{A})$ and thus the two constructions are isomorphic.

However, when we construct the invertible sheaf $\mathcal{O}(1)$ the transition maps come from the degree 1 part of $\psi_\bullet$, and thus must be modified using the transition functions for $\mathcal{L}$. $\square$

10.7.1 Projective bundles

Definition 10.7.9. Suppose that $\mathcal{F}$ is a locally free sheaf of rank r on X for some $r \geq 2$. We define the projective bundle $\mathbb{P}_X(\mathcal{F})$ to be $\underline{\mathrm{Proj}}(\mathrm{Sym}(\mathcal{F}))$. This comes equipped with a morphism $\pi : \mathbb{P}_X(\mathcal{F}) \to X$.

Note that $\mathrm{Sym}(\mathcal{F})$ satisfies Condition (*) so that $\mathcal{O}_{\mathbb{P}_X(\mathcal{F})/X}(1)$ is an invertible sheaf. Furthermore, the fiber of π over a point $x \in X$ is isomorphic to $\mathbb{P}^{r-1}_{\kappa(x)}$. Since sections of $\mathcal{O}_{\mathbb{P}_X(\mathcal{F})/X}(d)$ over sufficiently small open affines $U \subset X$ locally look like the degree d submodule of $\mathcal{O}_X(U)[x_1, \ldots, x_r]$, we see that

$$\pi_* \mathcal{O}_{\mathbb{P}_X(\mathcal{F})/X}(d) \cong \begin{cases} \mathrm{Sym}^d(\mathcal{F}) \text{ if } d > 0 \\ \mathcal{O}_X \text{ if } d = 0 \\ 0 \text{ if } d < 0 \end{cases}$$

Finally, we note that there is a canonical surjective map $\pi^*\mathcal{F} \to \mathcal{O}_{\mathbb{P}_X(\mathcal{F})/X}(1)$. Using the adjunction between π^* and π_*, this map is locally defined over an open affine $U \subset X$ by the identity map

$$\mathcal{F}|_U \to \Gamma(U, \pi_*\mathcal{O}(1)|_U) \cong \mathcal{F}|_U.$$

Remark 10.7.10. In contrast to Definition 9.5.10, we do *not* dualize $\mathcal{F}$ when making this construction. (Some authors do, but our convention is by far the predominant one.) This may seem surprising, so let's briefly analyze the geometry of the situation.

Consider the vector bundle $\rho : \mathcal{V} \to X$ defined by taking $\underline{\mathrm{Spec}}(\mathrm{Sym}(\mathcal{F}^\vee))$. We would like to say that we obtain $\mathbb{P}_X(\mathcal{F})$ by projectivizing the fibers of ρ. However, there are two possible ways to projectivize:

(1) We can replace each fiber V with the parameter space of 1-dimensional subspaces of V.

(2) We can replace each fiber V with the parameter space of 1-dimensional quotients of V.

Although both of these are abstractly isomorphic to $\mathbb{P}^{r-1}$, they are naturally "dual" to each other: the 1-dimensional subspaces of V are the same as 1-dimensional quotients of $V^\vee$.

With our conventions, $\mathbb{P}_X(\mathcal{F})$ is the space of 1-dimensional quotients of the fibers of $\mathcal{V}$. This is the "Grothendieck convention": every projective space and Grassmannian parametrizes quotients, not subspaces. While we have been content to hide this confusing point thus far, we will be more careful about it from here on out.

Example 10.7.11. Consider the rank 2 vector bundle $\mathcal{E}_e := \mathcal{O} \oplus \mathcal{O}(-e)$ on $\mathbb{P}^1_{\mathbb{K}}$. The corresponding projective bundle is known as the eth Hirzebruch surface $\mathbb{F}_e$. Consider the invertible sheaf $\mathcal{O}_{\mathbb{F}_e/\mathbb{P}^1}(1)$. We have

$$\dim \Gamma(\mathbb{F}_e, \mathcal{O}_{\mathbb{F}_e/\mathbb{P}^1}(1)) = \dim(\mathbb{P}^1, \mathcal{E}_e) = 1$$

so that $\mathcal{O}_{\mathbb{F}_e/\mathbb{P}^1}(1)$ has a unique global section up to rescaling. We will call the zero locus of this section the "contractible section" C_0 of $\mathbb{F}_e$. The name arises from the fact that there is a birational map from $\mathbb{F}_e$ to the cone Z over a degree e rational normal curve which is an isomorphism away from C_0 and contracts C_0 to the cone point of Z.

(We will see later that any rank 2 vector bundle on $\mathbb{P}^1_{\mathbb{K}}$ is isomorphic to the twist of some $\mathcal{E}_e$ by an invertible sheaf. Thus by Theorem 10.7.8 every $\mathbb{P}^1$-bundle over $\mathbb{P}^1$ is isomorphic to a Hirzebruch surface.)

10.7.2 Projective morphisms

The notion of a projective morphism is surprisingly subtle. The following definition does not satisfy the three conditions necessary to be a "well-behaved morphism" without some further finiteness constraints.

Definition 10.7.12. We say that a morphism $f : X \to Y$ is projective if there is a graded $\mathcal{O}_Y$-algebra $\mathcal{A}$ satisfying Condition (*) and an isomorphism $\phi : X \to \underline{\mathrm{Proj}}(\mathcal{A})$ such that f is the composition of ϕ with the structure map $\underline{\mathrm{Proj}}(\mathcal{A}) \to Y$.

In particular, every projective morphism is proper. (As we discussed earlier, the converse is false.) The following variant may be closer to what you were expecting.

Proposition 10.7.13. *Suppose that $f : X \to Y$ is a morphism and that Y admits an ample line bundle $\mathcal{T}$. Then f is projective if and only if f factors as a closed embedding $X \hookrightarrow \mathbb{P}^n_Y$ followed by the projection map.*

Here we can define an ample line bundle either as in Definition 10.6.4 (if Y is proper over a Noetherian scheme) or as in Remark 10.6.10 (if Y is quasicompact quasiseparated).

Proof. The interesting direction is the forward implication. Suppose X is isomorphic to $\underline{\mathrm{Proj}}(\mathcal{A})$ where $\mathcal{A}$ satisfies Condition (*). In particular $\mathcal{A}_1$ is a finitely generated quasicoherent sheaf. Thus $\mathcal{A}_1 \otimes \mathcal{T}^{\otimes k}$ is globally generated for k sufficiently large. This yields a surjection $\mathcal{O}_Y^{\oplus N} \to \mathcal{A}_1 \otimes \mathcal{T}^{\otimes k}$ and hence a closed embedding $\underline{\mathrm{Proj}}(\mathcal{A}_{\mathcal{T}}) \hookrightarrow \mathbb{P}^N_Y$ commuting with the maps to Y. By Theorem 10.7.8 we know that the former scheme is isomorphic to X. $\square$

10.7.3 Exercises

Exercise 10.7.14. Consider the Hirzebruch surface $\mathbb{F}_e := \mathbb{P}_{\mathbb{P}^1}(\mathcal{O} \oplus \mathcal{O}(-e))$. As discussed in Example 10.7.11 the sheaf $\mathcal{O}_{\mathbb{F}_e/\mathbb{P}^1}(1)$ has a unique section up to rescaling, defining a curve C_0. We let $\mathcal{L}$ denote the pullback of $\mathcal{O}_{\mathbb{P}^1}(1)$ under the projection map. Any global section of $\mathcal{L}$ will define a fiber F of the projection map.

(1) Prove that $\mathrm{Pic}(\mathbb{F}_e) \cong \mathbb{Z}^2$ is generated by $\mathcal{O}_{\mathbb{F}_e/\mathbb{P}^1}(1)$ and $\mathcal{L}$.

(2) Show that the only line bundles which have global sections will have the form $\mathcal{O}_{\mathbb{F}_e/\mathbb{P}^1}(1)^{\otimes a} \otimes \mathcal{L}^{\otimes b}$ where $a, b \geq 0$.

(3) Show that $\deg(\mathcal{L}|_{C_0}) = 1$ and $\deg(\mathcal{O}_{\mathbb{F}_e/\mathbb{P}^1}(1)|_{C_0}) = -e$. Show that $\deg(\mathcal{L}|_F) = 0$ and $\deg(\mathcal{O}_{\mathbb{F}_e/\mathbb{P}^1}(1)|_F) = 1$. Conclude that $\mathrm{Pic}(X) \to N^1(X)$ is an isomorphism.

(4) Consider the invertible sheaf $\mathcal{T} = \mathcal{O}_{\mathbb{F}_e/\mathbb{P}^1}(1) \otimes \mathcal{L}^{\otimes e}$. Show that $\mathcal{T}$ is basepoint free. Show that $\deg(\mathcal{T}|_{C_0}) = 0$ and that C_0 is the only curve with this property. Conclude that the sections of $\mathcal{T}$ define a morphism to projective space which contracts C_0. (Hint: use Exercise 10.5.24.)

Exercise 10.7.15. Suppose that $f : X \to Y$ is a finite morphism. Define a $\mathcal{O}_Y$-algebra $\mathcal{A}$ where $\mathcal{A}_0 \cong \mathcal{O}_Y$ and $\mathcal{A}_n \cong f_*\mathcal{O}_X$ for any $n \geq 1$ with the natural multiplication structure. Prove that $\underline{\mathrm{Proj}}(\mathcal{A})$ recovers the map f.

Deduce that every finite morphism is projective.

Exercise 10.7.16. Let X be a scheme and let $\mathcal{F}$ be a locally free sheaf on X. Show that $\underline{\mathrm{Spec}}(\mathrm{Sym}(\mathcal{F}))$ admits an open embedding into $\underline{\mathrm{Proj}}(\mathrm{Sym}(\mathcal{O}_X \oplus \mathcal{F}))$ as a "relative distinguished open affine."

Chapter 11

Cotangent sheaves

Suppose that M is a smooth manifold. There are several approaches one can take to defining the tangent space T_xX at a point $x \in M$:

(1) Chart structure: one can first define the tangent space for points on $\mathbb{R}^n$ and then "transform" these spaces to M using the chart structure. One must verify that the definition does not depend on the choice of chart.

(2) Jets of curves: consider the set of curves $\sigma : J \to M$ which are smooth at x. We can define an equivalence relation on such curves by setting $\sigma_1 \sim \sigma_2$ if for every smooth function $f : M \to \mathbb{R}$ defined on a neighborhood of x the derivatives of $f \circ \sigma_1$ and $f \circ \sigma_2$ coincide. We can define the tangent space to be the set of equivalence classes of such σ.

(3) Derivations: let $\mathcal{C}_x^\infty$ denote the set of germs of smooth real-valued functions near x. A derivation is a linear map $T : \mathcal{C}_x^\infty \to \mathbb{R}$ satisfying the product rule $T(fg) = f(x)T(g) + g(x)T(f)$. Then the vector space of derivations is the tangent space at x.

In algebraic geometry we will use the space of derivations as the foundation for building the theory of tangent spaces. This is motivated by the fact that the fundamental object in algebraic geometry is the space of functions instead of points. In particular in our setting the cotangent space is the most natural object: it allows us to work with derivations directly, whereas the tangent space requires taking a dual.

In contrast to the Zariski cotangent space, the theory of cotangent sheaves is best thought of as a "relative" theory: given a finitely presented morphism $f : X \to Y$, we will construct a cotangent sheaf $\Omega_{X/Y}$ that represents the "relative cotangent space" of X over Y. (We can recover the "absolute" cotangent sheaf of a $\mathbb{K}$-scheme by considering the structure map $X \to \mathrm{Spec}(\mathbb{K})$.)

It turns out that all three of the above definitions of the tangent space admit loose analogues in various settings:

(1) Chart structure: suppose we have a finitely presented morphism $f : X \to Y$. Since the construction of the tangent space is "local", we may reduce to the case where X and Y are schemes. While X and Y need not look anything like affine space, we can give a local "algebraic" construction of $\Omega_{X/Y}$ which can then be glued to form the cotangent sheaf.

    Thus suppose that X and Y are affine. Since f is finitely presented we see that X is a closed subscheme of $\mathbb{A}^n_Y$ defined by a finite set of equations $\{f_i\}$. If we were working in the usual geometric setting, we would know how to compute the cotangent space of X: it can be defined as the kernel (or cokernel, depending on your conventions) of the Jacobian matrix defined by the $\{f_i\}$. This computation will work equally well in our setting.

(2) Jets: suppose that X is a $\mathbb{K}$-scheme. As discussed in Definition 5.3.1, given an $\mathbb{L}$-point $x \in X$ we can construct the tangent space of x at X using the space of morphisms from $\mathrm{Spec}(\mathbb{L}[t]/(t^2))$ to X with image x. Conceptually, any such morphism represents taking the "first order jet" of a curve through the point x.

(3) Derivations: as discussed above this will be our foundational definition. We will give a purely algebraic description of the space of derivations associated to a ring map $B \to A$. Since this construction is compatible with localization, it will extend to define the relative cotangent sheaf of a morphism of schemes.

11.1 Modules of differentials

Let $B \to A$ be a ring homomorphism and let M be an A-module. A B-linear derivation from A into M is a B-module homomorphism $d : A \to M$ satisfying the Leibniz rule

$$d(fg) = f \cdot d(g) + g \cdot d(f)$$

It is important to note that we have $d(b) = 0$ for any $b \in B$ (i.e. the ring B acts like the "scalars" for our derivations). This is a formal consequence of the B-module structure and the Leibniz rule via the computation

$$b \cdot d(f) = d(bf) = b \cdot d(f) + f \cdot d(b).$$

As discussed in the introduction to the chapter, our definition of B-linear derivations mimics the analogous notion in a geometric setting.

Warning 11.1.1. Even though M is an A-module, the map d is usually *not* an A-module homomorphism. We only require it to be linear in B. (The argument above shows that d is only an A-module homomorphism when it is the zero map.)

Definition 11.1.2. Let $B \to A$ be a ring homomorphism and let M be an A-module. We let $\mathrm{Der}_B(A, M)$ denote the set of B-linear derivations from A into M. Although a derivation is not an A-module map, the set $\mathrm{Der}_B(A, M)$ carries the structure of an A-module where the action of a sends $d \mapsto a \cdot d$.

11.1.1 Module of differentials

Let $B \to A$ be a ring homomorphism. The module of relative differentials $\Omega_{A/B}$ is a "universal" construction of a B-linear derivation for A.

Construction 11.1.3. Let $B \to A$ be a ring homomorphism. We define the A-module $\Omega_{A/B}$ by taking the free module generated by the symbols da for $a \in A$ and then imposing the relations

(1) $d(a + a') = da + da'$ for every $a, a' \in A$.

(2) $d(aa') = a \cdot da' + a' \cdot da$ for every $a, a' \in A$.

(3) $db = 0$ for every $b \in B$.

The module $\Omega_{A/B}$ comes with a canonical B-linear derivation $d : A \to \Omega_{A/B}$ sending $a \mapsto da$.

Based on the construction, it should be no surprise that $\Omega_{A/B}$ satisfies a universal property.

Proposition 11.1.4. *Let $B \to A$ be a ring homomorphism and let M be an A-module. Every B-linear derivation $\phi : A \to M$ can be written in a unique way as the composition of the universal map $d : A \to \Omega_{A/B}$ with an A-module homomorphism $\psi : \Omega_{A/B} \to M$ and the induced map $\mathrm{Der}_B(A, M) \to \mathrm{Hom}_A(\Omega_{A/B}, M)$ is an isomorphism of A-modules.*

In other words, the functor $M \mapsto \mathrm{Der}_B(A, M)$ is representable by $\Omega_{A/B}$.

Exercise 11.1.5. Prove the previous proposition.

Although the definition of $\Omega_{A/B}$ looks a little intimidating (as it uses a module with infinitely many generators) in practice it behaves quite well. In particular, it behaves compatibly with finite generation. (See also Lemma 11.1.12.)

Lemma 11.1.6. *Let $B \to A$ be a ring homomorphism. Suppose that A is finitely generated over B. Then $\Omega_{A/B}$ is a finitely generated A-module.*

Proof. Suppose that $a_1, \dots, a_r$ are the generators for A over B. Then $\Omega_{A/B}$ is generated by $da_1, \dots, da_r$. Indeed, using the Leibniz axiom repeatedly we see that the d-image of any polynomial in the a_i with coefficients in B will be in the submodule generated by the da_i. □

11.1.2 Compatibilities

In this section we prove two basic compatibilities of the module of relative differentials with algebraic operations.

First, the construction of $\Omega_{A/B}$ is compatible with localization in a strong sense where we allow ourselves to localize A and B along different multiplicatively closed sets. In geometric language, the result shows the compatibility of relative differentials with localization upon passing to open subsets (both in the domain and in the target).

Proposition 11.1.7. *Let $B \to A$ be a ring homomorphism, let $T \subset B$ be a multiplicatively closed subset, and let $S \subset A$ be a multiplicatively closed subset containing the image of T. Then there is an isomorphism of $S^{-1}A$-modules*

$$S^{-1}A \otimes \Omega_{A/B} \cong \Omega_{S^{-1}A/T^{-1}B}.$$

Explicitly the isomorphism is the map $S^{-1}A \otimes \Omega_{A/B} \to \Omega_{S^{-1}A/T^{-1}B}$ given by sending $\frac{1}{s} \otimes da \mapsto d(\frac{a}{s})$. We will describe the inverse map in the proof below, and you should check carefully that these two explicit descriptions are inverse to each other.

Proof. Suppose that M is an $S^{-1}A$-module. Given any derivation $d \in \mathrm{Der}_{T^{-1}B}(S^{-1}A, M)$, by precomposing d with $A \to S^{-1}A$ we obtain an element of $\mathrm{Der}_B(A, M)$. This function is injective: the Leibniz rule shows that $d(\frac{a}{s}) = \frac{1}{s} \cdot da - \frac{a}{s^2} \cdot ds$ and so d is determined by its action on A. This function is surjective: given a B-linear derivation $\widetilde{d} : A \to M$

we extend it to $S^{-1}A$ by sending $d(\frac{a}{s}) = \frac{1}{s} \cdot \widetilde{da} - \frac{a}{s^2} \cdot \widetilde{ds}$. (Note that this extension is automatically a derivation over $T^{-1}B$ and not just B.) We can equip $\mathrm{Der}_B(A, M)$ the "naive" $S^{-1}A$-module structure where the action of $\frac{a}{s}$ sends $d \mapsto \frac{a}{s} \cdot d$. With this definition the map above is an isomorphism of $S^{-1}A$-modules.

We now appeal to the universal property of the module of differentials. For any $S^{-1}A$-module M we can identify $\mathrm{Der}_{T^{-1}B}(S^{-1}A, M)$ as the R-module $\mathrm{Hom}_R(\Omega_{S^{-1}A/T^{-1}B}, M)$. Since as an A-module $\mathrm{Der}_B(A, M)$ is isomorphic to $\mathrm{Hom}_A(\Omega_{A/B}, M)$ we see that as an $S^{-1}A$-module it is isomorphic to

$$\mathrm{Hom}_{S^{-1}A}(S^{-1}A \otimes \Omega_{A/B}, M).$$

Since these two Hom-modules are isomorphic for every choice of M, Yoneda's lemma implies the desired statement. □

Second, the construction of $\Omega_{A/B}$ is compatible with base change.

Proposition 11.1.8. *Let $B \to A$ and $B \to S$ be ring homomorphisms. Set $R = S \otimes_B A$. Then there is an isomorphism of R-modules*

$$R \otimes_A \Omega_{A/B} \cong \Omega_{R/S}.$$

Explicitly the isomorphism is the map $R \otimes_A \Omega_{A/B} \to \Omega_{R/S}$ given by sending $r \otimes da \mapsto r \cdot d(a)$. We will describe the inverse map in the proof below, and you should check carefully that these two explicit descriptions are inverse to each other.

Proof. Fix an R-module M. Given any derivation $d \in \mathrm{Der}_S(R, M)$, its restriction to A is an element of $\mathrm{Der}_B(A, M)$. The restriction map is injective: since d acts trivially on S its action on R is entirely determined by its action on A via the Leibniz rule. The restriction map is surjective: given a B-linear derivation $d' : A \to M$, we extend it to R by defining $d(s \otimes a) = s \cdot d'a$. Note that the restriction map is in fact an isomorphism of R-modules if we equip $\mathrm{Der}_B(A, M)$ with the "naive" R-module structure where the action of r sends $d \mapsto r \cdot d$.

We now appeal to the universal property of the module of differentials. For any R-module M we can identify $\mathrm{Der}_S(R, M)$ as the R-module $\mathrm{Hom}_R(\Omega_{R/S}, M)$. Furthermore, as an A-module $\mathrm{Der}_B(A, M)$ is isomorphic to $\mathrm{Hom}_A(\Omega_{A/B}, M)$, which means that as an R-module it is isomorphic to

$$\mathrm{Hom}_R(R \otimes_A \Omega_{A/B}, M).$$

Since these two Hom-modules are isomorphic for every choice of M, Yoneda's lemma implies the desired statement. □

11.1.3 Exact sequences

Suppose we have a sequence of ring homomorphisms $C \to B \to A$. We might then hope that $\Omega_{A/C}$ is "built up" out of $\Omega_{A/B}$ and $\Omega_{B/C}$ in some way.

Proposition 11.1.9 (Cotangent sequence)**.** *Let $C \to B \to A$ be ring homomorphisms. Then there is an exact sequence of A-modules*

$$A \otimes_B \Omega_{B/C} \xrightarrow{\psi} \Omega_{A/C} \xrightarrow{\phi} \Omega_{A/B} \to 0$$

Here the map on the right sends $da \mapsto da$; the map on the left sends $a \otimes db \mapsto a \cdot db$.

Proof. It is clear that the map $\phi : \Omega_{A/C} \to \Omega_{A/B}$ sending $da \mapsto da$ is a surjective A-module homomorphism. To show that $\ker(\phi) = \operatorname{im}(\psi)$, note that $\Omega_{A/C}$ and $\Omega_{A/B}$ are defined by quotienting the same free module by the same relations except for the additional relations $db = 0$ in $\Omega_{A/B}$. Thus the kernel is generated by elements in the image of $\Omega_{B/C}$. □

While is a bit frustrating that the cotangent sequence is not exact on the left, there is one special situation where we can continue the exact sequence by adding an explicit kernel on the left. Suppose we have ring homomorphisms $C \to B \to A$ and that $B \to A$ is surjective. This easily implies that $\Omega_{A/B} = 0$, so the rightmost entry in Proposition 11.1.9 vanishes. In this case we can extend by one term on the left:

Proposition 11.1.10 (Conormal sequence)**.** *Let $C \to B \to A$ be ring homomorphisms and suppose that $B \to A$ is surjective. Then there is an exact sequence of A-modules*

$$A \otimes_B I \xrightarrow{\eta} A \otimes_B \Omega_{B/C} \xrightarrow{\psi} \Omega_{A/C} \to 0$$

where I is the kernel of the surjection $B \to A$ and the map $\eta = 1 \otimes d$ sends $a \otimes i \mapsto a \otimes di$.

It is more traditional to write I/I^2 for the leftmost term, where we use the identification

$$A \otimes_B I \cong B/I \otimes_B I \cong I/I^2$$

and thus we will use this notation henceforth.

Proof. Recall that $A \otimes_B \Omega_{B/C}$ is the A-module generated by the symbols db subject to the relations described by the additivity and Leibniz rules and by the requirement $dc = 0$. Since A is a quotient of B, $\Omega_{A/C}$ can be defined in the same way but subject to the extra relations $di = 0$ as we vary $i \in I$. But this is exactly the image of the map on the left. □

Remark 11.1.11. In general, it is natural to look for algebraic constructions which further continue the exact sequence of Proposition 11.1.9 on the left. This goal is achieved via the theory of the cotangent complex.

11.1.4 Jacobians

If A is a finitely presented B-module, then we can compute $\Omega_{A/B}$ using the Jacobian matrix. We will usually restrict ourselves to situations where the module of differentials can be computed in this way.

Lemma 11.1.12. *Let $B \to A$ be a ring homomorphism. Suppose that*

$$A = B[x_1, \dots, x_n]/(f_1, \dots, f_r)$$

is a finitely presented B-algebra. Then $\Omega_{A/B}$ is the finitely presented A-module which is the cokernel of the map $A^{\oplus r} \to A^{\oplus n}$ defined by the Jacobian matrix

$$\mathrm{Jac}_{f_1,\dots,f_r}(x) = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_r}{\partial x_1} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_1}{\partial x_n} & \cdots & \frac{\partial f_r}{\partial x_n} \end{bmatrix}.$$

Note that since we are working cotangent spaces, our definition of the Jacobian is the transpose of the more commonly used convention.

Proof. We first consider the case when $A = B[x_1, \dots, x_n]$. Then we claim that $\Omega_{A/B}$ is freely generated by $dx_1, \dots, dx_n$. As in Lemma 11.1.6 we see that these symbols generate and we just need to show that the dx_i are independent. This follows from the fact that the map $\frac{d}{dx_i} : A \to B$ is an example of a B-linear derivation that sends $dx_i \mapsto 1$ and every other $dx_j \mapsto 0$.

In general, consider the sequence of ring homomorphisms $B \to B[x_1, \dots, x_n] \to A$ and let I be the ideal defining A. By Proposition 11.1.10 we obtain an exact sequence

$$I/I^2 \xrightarrow{\phi} A\,dx_1 \oplus \dots \oplus A\,dx_n \to \Omega_{A/B} \to 0.$$

The map ϕ sends a generator f_j to df_j, which by the Leibniz rule is the same as $\sum \frac{\partial f}{\partial x_i} dx_i$. Thus the image of I/I^2 under ϕ is the same as the image of the Jacobian matrix. □

11.1.5 Exercises

Exercise 11.1.13. Let $\mathbb{K}$ be a field and set $A = \mathbb{K}[x, y]/(x^2 + y^2)$. Compute $\Omega_{A/\mathbb{K}}$. (Hint: it will depend on the characteristic of $\mathbb{K}$.)

Exercise 11.1.14. Suppose that $B \to A$ is a ring homomorphism. Let I be the ideal in $A \otimes_B A$ which is the kernel of the ring map $A \otimes_B A \to A$ that sends $a_1 \otimes_B a_2 \mapsto a_1 a_2$.

(1) Verify that I is generated by elements of the form $1 \otimes a - a \otimes 1$.

(2) Show that the map $\widetilde{d} : A \to I/I^2$ sending $a \mapsto 1 \otimes a - a \otimes 1$ is a derivation.

(3) Show that the A-module homomorphism $\phi : \Omega_{A/B} \to I/I^2$ induced by $\widetilde{d}$ is an isomorphism. (Hint: to see that ϕ is injective, show that the map $A \otimes A \to \Omega_{A/B}$ sending $x \otimes y \mapsto x \cdot dy$ yields a one-sided inverse when descended to I/I^2.) In particular, we can use I/I^2 equipped with the derivation $\widetilde{d}$ as an alternative definition of the module of relative differentials.

Exercise 11.1.15. Let $\mathbb{K}$ be a field. Show that $\mathrm{Der}_{\mathbb{K}}(A, A)$ carries the structure of a Lie algebra where $[d_1, d_2] = d_1 \circ d_2 - d_2 \circ d_1$.

11.2 Field extensions

In this section we will systematically analyze the module of differentials for extensions of fields $\mathbb{L}/\mathbb{K}$. These form the foundation for some of our later results.

11.2.1 Basic extensions

Every finitely generated extension of fields can be decomposed into a sequence of extensions of three types: finite separable extensions, finite purely inseparable extensions, and finitely generated purely transcendental extensions. We begin by studying each type separately.

Proposition 11.2.1. *If $\mathbb{K} \to \mathbb{L}$ is a separable algebraic extension then $\Omega_{\mathbb{L}/\mathbb{K}} = 0$.*

Proof. Let $a \in \mathbb{L}$ and let P be its minimal polynomial; since the extension is separable we know that the formal derivative P' of P evaluated at a is non-zero. Taking the d-image of the equation $P(a) = 0$ in $\Omega_{\mathbb{L}/\mathbb{K}}$ we obtain the relation $P'(a) \cdot da = 0$. We conclude that every element da is torsion in the $\mathbb{L}$-module $\Omega_{\mathbb{L}/\mathbb{K}}$. Since $\mathbb{L}$ is a field this implies that $da = 0$ for every a. □

Proposition 11.2.2. *Suppose that $\mathbb{K} \to \mathbb{L}$ is a purely inseparable finite extension. Let r denote the minimal number of elements in $\mathbb{L}$ needed to generate $\mathbb{L}$ over $\mathbb{K}$. Then $\Omega_{\mathbb{L}/\mathbb{K}}$ is an r-dimensional $\mathbb{L}$-vector space.*

Proof. Let p denote the characteristic and define $\mathbb{F} = \mathbb{L}^p\mathbb{K}$. Let s be the minimal number of elements in $\mathbb{L}$ needed to generate $\mathbb{L}$ over $\mathbb{F}$. We claim that $s = r$. It is clear that $s \leq r$. Conversely, suppose that $\mathbb{L} = \mathbb{L}^p\mathbb{K}(\alpha_1, \ldots, \alpha_s)$. Then arguing inductively we see that

$$\mathbb{L} = \mathbb{L}^p\mathbb{K}(\alpha_1, \ldots, \alpha_s) = \mathbb{L}^{p^2}\mathbb{K}(\alpha_1, \ldots, \alpha_s) = \ldots = \mathbb{L}^{p^d}\mathbb{K}(\alpha_1, \ldots, \alpha_s)$$

Since $\mathbb{L}/\mathbb{K}$ is a finite purely inseparable extension, we have $\mathbb{L}^{p^d}\mathbb{K} = \mathbb{K}$ for some sufficiently large d. In this way we see that $s \geq r$.

We next claim that $\Omega_{\mathbb{L}/\mathbb{K}} \cong \Omega_{\mathbb{L}/\mathbb{F}}$. Indeed, by Proposition 11.1.9 we have an exact sequence

$$\mathbb{L} \otimes_{\mathbb{F}} \Omega_{\mathbb{F}/\mathbb{K}} \to \Omega_{\mathbb{L}/\mathbb{K}} \to \Omega_{\mathbb{L}/\mathbb{F}} \to 0.$$

Note that the image of the leftmost module is 0 (since d vanishes on $\mathbb{K}$ and on pth powers in $\mathbb{L}$). This shows the claim.

It only remains to compute $\Omega_{\mathbb{L}/\mathbb{F}}$. As before write $\mathbb{L} = \mathbb{F}(\alpha_1, \ldots, \alpha_r)$. By construction the pth power of every element in $\mathbb{L}$ is contained in $\mathbb{F}$, so

$$\mathbb{L} = \mathbb{F}[x_1, \ldots, x_r]/(x_1^p - \alpha_1^p, \ldots, x_r^p - \alpha_r^p).$$

We conclude the desired statement by Lemma 11.1.12. □

Corollary 11.2.3. *If $\mathbb{K} \to \mathbb{L}$ is a non-separable finite extension then $\Omega_{\mathbb{L}/\mathbb{K}} \neq 0$.*

Exercise 11.2.12 shows that this statement may fail for non-finite extensions.

Proof. In this case the extension $\mathbb{L}/\mathbb{K}$ factors as a separable extension $\mathbb{F}/\mathbb{K}$ followed by a purely inseparable extension $\mathbb{L}/\mathbb{F}$. Proposition 11.1.9 shows that there is a surjection $\mathbb{L}/\mathbb{K} \to \mathbb{L}/\mathbb{F}$, yielding the result. □

Proposition 11.2.4. *If $\mathbb{L} = \mathbb{K}(x_1, \ldots, x_n)$ is a finitely generated purely transcendental extension, then $\Omega_{\mathbb{L}/\mathbb{K}}$ is the free $\mathbb{L}$-module generated by the dx_i.*

Proof. Let $A = \mathbb{K}[x_1, \ldots, x_n]$. We have shown that $\Omega_{A/\mathbb{K}}$ is the free A-module generated by the dx_i. Proposition 11.1.7 implies the desired statement for the localization $\mathbb{L}$ of A. □

11.2.2 Finitely generated field extensions

We next combine these results to address the behavior of $\Omega_{\mathbb{L}/\mathbb{K}}$ for any finitely generated extension $\mathbb{L}/\mathbb{K}$.

Lemma 11.2.5. *Suppose that $\mathbb{K} \to \mathbb{L}$ is a finitely generated extension. Suppose that $\mathbb{F}$ is an intermediate field such that $\mathbb{L}/\mathbb{F}$ is finite separable. Then there is an isomorphism*

$$\mathbb{L} \otimes_{\mathbb{F}} \Omega_{\mathbb{F}/\mathbb{K}} \cong \Omega_{\mathbb{L}/\mathbb{K}}$$

defined by the leftmost map in the cotangent sequence.

Proof. Consider the cotangent sequence

$$\mathbb{L} \otimes_{\mathbb{F}} \Omega_{\mathbb{F}/\mathbb{K}} \to \Omega_{\mathbb{L}/\mathbb{K}} \to \Omega_{\mathbb{L}/\mathbb{F}} \to 0.$$

By Proposition 11.2.1 we know that $\Omega_{\mathbb{L}/\mathbb{F}} = 0$ and so the leftmost map is surjective. To show injectivity, it suffices to check that any $\mathbb{K}$-linear derivation $d : \mathbb{F} \to M$ extends to a derivation $\widetilde{d} : \mathbb{L} \to M \otimes_{\mathbb{F}} \mathbb{L}$. Indeed, using the identifications $\mathrm{Der}_{\mathbb{K}}(\mathbb{F}, M) \cong \mathrm{Hom}_{\mathbb{F}}(\Omega_{\mathbb{F}/\mathbb{K}}, M)$ and similarly for $\mathbb{L}$, a surjection of spaces of derivations leads to an injection on the modules of differentials.

Choose a primitive element α for $\mathbb{L}/\mathbb{F}$ and let $f(x) = \sum c_j x^j$ denote its minimal polynomial in $\mathbb{F}$. We start by defining a derivation $D : \mathbb{F}[x] \to M \otimes_{\mathbb{F}} \mathbb{L}$ by setting $D(c) = d(c)$ for any $c \in \mathbb{F}$, defining

$$D(x) = -\frac{\sum d(c_j)\alpha^j}{f'(\alpha)}$$

which is well-defined by our separability assumption, and then extending D to all of $\mathbb{F}[x]$ via the Leibniz rule. (Here we are realizing $M \otimes_{\mathbb{F}} \mathbb{L}$ as an $\mathbb{F}[x]$-module by letting x act via α.) Since we have the relation

$$D(f(x)) = D\left(\sum c_j x^j\right) = \left(\sum d(c_j)\alpha^j\right) + f'(\alpha)D(x) = 0$$

this derivation D descends to the quotient $\mathbb{F}[x]/f(x) \cong \mathbb{L}$. □

Theorem 11.2.6. *Suppose that $\mathbb{K} \to \mathbb{L}$ is a finitely generated separable extension. Then $\Omega_{\mathbb{L}/\mathbb{K}}$ is a free $\mathbb{L}$-module of rank* $\operatorname{trdeg}(\mathbb{L}/\mathbb{K})$*.*

Proof. Since $\mathbb{L}$ is a separable extension of $\mathbb{K}$, we can find an intermediate field $\mathbb{F}$ such that $\mathbb{L}/\mathbb{F}$ is separable algebraic and $\mathbb{F}/\mathbb{K}$ is purely transcendental. The desired statement follows from Lemma 11.2.5 and Proposition 11.2.4. □

Corollary 11.2.7. *Suppose that $\mathbb{K} \to \mathbb{L}$ is a finitely generated extension. Then $\Omega_{\mathbb{L}/\mathbb{K}} = 0$ if and only if $\mathbb{L}$ is a finite separable extension of $\mathbb{K}$.*

Proof. The reverse implication has been proved already. To see the forward implication, suppose that $\alpha_1, \ldots, \alpha_r$ is a minimal generating set for $\mathbb{L}$ over $\mathbb{K}$. For ease of notation we set $\mathbb{F}_j = \mathbb{K}(\alpha_1, \ldots, \alpha_j)$ and $\mathbb{F}_0 = \mathbb{K}$. Let j be the largest index such that $\mathbb{L}$ is not a finite separable extension of $\mathbb{F}_j$ and suppose for a contradiction that $j \neq 0$. Since $\mathbb{L}$ is a finite separable extension of $\mathbb{F}_{j+1}$, we see that the extension $\mathbb{F}_{j+1}/\mathbb{F}_j$ is either inseparable or a transcendental. In either case $\Omega_{\mathbb{F}_{j+1}/\mathbb{F}_j} \neq 0$, and using the surjection in the cotangent sequence this implies that $\Omega_{\mathbb{F}_{j+1}/\mathbb{K}} \neq 0$. By Lemma 11.2.5 the map $\mathbb{L} \otimes \Omega_{\mathbb{F}_{j+1}/\mathbb{K}} \to \Omega_{\mathbb{L}/\mathbb{K}}$ is an isomorphism, giving a contradiction. □

The situation for finitely generated non-separable extensions is very similar to Proposition 11.2.2.

Theorem 11.2.8. *Suppose that $\mathbb{K} \to \mathbb{L}$ is a finitely generated non-separable extension in characteristic p. Let $\alpha_1, \ldots, \alpha_r$ be a minimal set of generators of $\mathbb{L}$ over $\mathbb{L}^p\mathbb{K}$. Then $\Omega_{\mathbb{L}/\mathbb{K}}$ is a free $\mathbb{L}$-module generated by the $d\alpha_i$.*

Note that the α_i may or may not be transcendental elements over $\mathbb{K}$. The proof is exactly the same as Proposition 11.2.2 but easier since we can omit the first paragraph. In turn this implies:

Corollary 11.2.9. *Suppose that $\mathbb{K} \to \mathbb{L}$ is a finitely generated extension. Then*

$$\dim_{\mathbb{L}}(\Omega_{\mathbb{L}/\mathbb{K}}) \geq \operatorname{trdeg}(\mathbb{L}/\mathbb{K})$$

with equality if and only if $\mathbb{L}/\mathbb{K}$ is a separable extension.

Proof. We have already proved the case when $\mathbb{L}/\mathbb{K}$ is separable, so it suffices to consider the situation when $\mathbb{L}$ is non-separable over $\mathbb{K}$. Let $\alpha_1, \ldots, \alpha_r$ be as in Theorem 11.2.8 and set $\mathbb{F} = \mathbb{K}(\alpha_1, \ldots, \alpha_r)$. We claim that $\mathbb{L}$ is a finite separable extension of $\mathbb{F}$. Indeed, since $d\alpha_1, \ldots, d\alpha_r$ also generate $\Omega_{\mathbb{F}/\mathbb{K}}$ the leftmost map in the cotangent sequence is surjective. We deduce that $\Omega_{\mathbb{L}/\mathbb{F}} = 0$ and so Corollary 11.2.7 shows that $\mathbb{L}$ is a finite separable extension of $\mathbb{F}$. In particular, since the number of generators must be strictly larger than the transcendence degree (as $\mathbb{L}$ is non-separable over $\mathbb{K}$) we see that

$$\operatorname{trdeg}(\mathbb{L}/\mathbb{K}) = \operatorname{trdeg}(\mathbb{F}/\mathbb{K}) < r = \dim_{\mathbb{L}}(\Omega_{\mathbb{L}/\mathbb{K}}).$$

□

Remark 11.2.10. The previous corollary does not hold for non-finitely generated extensions. [ML39, Section 10] gives an example of a field extension $\mathbb{L}/\mathbb{K}$ of transcendence degree 2 such that $\Omega_{\mathbb{L}/\mathbb{K}}$ is 1-dimensional.

11.2.3 Artinian rings

Our work thus far allows us to analyze the module of differentials for any Artinian ring.

Proposition 11.2.11. *Let A be an Artinian ring over a field $\mathbb{K}$. Then $\Omega_{A/\mathbb{K}} = 0$ if and only if $\operatorname{Spec}(A)$ is a finite disjoint union $\sqcup_{i=1}^r \operatorname{Spec}(\mathbb{L}_i)$ such that each extension $\mathbb{L}_i/\mathbb{K}$ is finite separable.*

Proof. The reverse implication follows from Proposition 11.2.1. To prove the forward implication, since $\Omega_{A/\mathbb{K}}$ is compatible with localization it suffices to consider the case when $\operatorname{Spec}(A)$ is irreducible and thus has a unique prime ideal $\mathfrak{m}$. Let $\mathbb{L}$ denote the residue field of the unique point of $\operatorname{Spec}(A)$. Proposition 11.1.10 shows that we have an exact sequence of $\mathbb{L}$-modules

$$\mathfrak{m}/\mathfrak{m}^2 \to \mathbb{L} \otimes_A \Omega_{A/\mathbb{K}} \to \Omega_{\mathbb{L}/\mathbb{K}} \to 0.$$

Since the middle term vanishes, so does the rightmost term. By Corollary 11.2.3 we deduce that the finite extension $\mathbb{L}/\mathbb{K}$ must be separable.

We claim that the sequence above is exact on the left as well. Suppose that $f_1, \ldots, f_r$ form a basis for $\mathfrak{m}/\mathfrak{m}^2$. We need to show that for any f_i the element df_i in $\Omega_{A/\mathbb{K}}$ does not vanish upon tensoring with $\mathbb{L}$. To do this, we will construct an explicit derivation. First consider the $\mathbb{K}$-linear quotient map $g : A \to A/\mathfrak{m}^2$. As a $\mathbb{K}$-vector space $A/\mathfrak{m}^2$ is isomorphic to $\mathbb{L} \oplus \mathbb{L}f_1 \oplus \ldots \oplus \mathbb{L}f_r$. Thus we can define a $\mathbb{K}$-linear function $h : A/\mathfrak{m}^2 \to \mathbb{L}$ by sending the basis vector $f_i \mapsto 1$ and every other basis vector to 0. The composition $d = h \circ g$ is a $\mathbb{K}$-linear derivation, since if we choose $b_i, c_i \in \mathbb{L}$ so that $g(a_1) = b_0 + \sum b_i f_i$ and $g(a_2) = c_0 + \sum c_i f_i$ then

$$\begin{aligned} d(a_1 a_2) = h(g(a_1)g(a_2)) &= h(b_0c_0 + \sum (c_0 b_i + b_0 c_i) f_i) \\ &= c_0 b_i + b_0 c_i = a_1 \cdot d(a_2) + a_2 \cdot d(a_1). \end{aligned}$$

Since our construction of this derivation respected the $\mathbb{L}$-vector space structure, we see that it does not vanish upon tensoring with $\mathbb{L}$, proving injectivity on the left.

Since $\Omega_{A/\mathbb{K}} = 0$ by assumption, we conclude that $\mathfrak{m}/\mathfrak{m}^2 = 0$, or equivalently, that A is a field. The desired statement then follows from Corollary 11.2.7. □

11.2.4 Exercises

Exercise 11.2.12. Set $\mathbb{K} = \mathbb{F}_p(T)$. Consider the chain of finite extensions

$$\mathbb{K} = \mathbb{L}_0 \subset \mathbb{L}_1 \subset \mathbb{L}_2 \subset \ldots$$

where $\mathbb{L}_1$ is generated by a pth root of T and each subsequent $\mathbb{L}_i$ is generated by a pth root of the previous generator. Set $\mathbb{L} = \cup_i \mathbb{L}_i$. Prove that $\Omega_{\mathbb{L}/\mathbb{K}} = 0$.

11.3 Cotangent sheaves

Suppose that $f : X \to Y$ is a smooth morphism of differentiable manifolds. We then obtain a pullback map $df^* : f^*\Omega_Y \to \Omega_X$. The cokernel of this map is known as the relative cotangent sheaf and is denoted by $\Omega_{X/Y}$. In this section we will develop a similar picture in the setting of algebraic geometry.

11.3.1 Relative cotangent sheaf

We first globalize the construction of the module of differentials from the previous section.

Definition 11.3.1. Let $f : X \to Y$ be a morphism of schemes. Suppose that $V \subset Y$ is an open affine and that $U \subset f^{-1}V$ is an open affine. The ring map $f^\sharp : \mathcal{O}_Y(V) \to \mathcal{O}_X(U)$ induces an $\mathcal{O}_X(U)$-module of relative differentials $\Omega_{U/V}$. As demonstrated in Proposition 11.1.7 these modules are compatible with localization. Thus, as we vary U and V the sheaves $\widetilde{\Omega_{U/V}}$ glue together to yield a quasicoherent sheaf $\Omega_{X/Y}$ on X.

$\Omega_{X/Y}$ is called the relative cotangent sheaf for the morphism f.

Remark 11.3.2. Note that (again appealing to Proposition 11.1.7) the stalk of $\Omega_{X/Y}$ at a point x can be identified with $\Omega_{\mathcal{O}_{X,x}/\mathcal{O}_{Y,f(x)}}$.

If we are working in the category of S-schemes (particularly when $S = \mathrm{Spec}(\mathbb{K})$) we will think of $\Omega_{X/S}$ as the "absolute" cotangent sheaf. In this setting we think of the relative cotangent sheaf $\Omega_{X/Y}$ as expressing the "difference" between the two cotangent sheaves $\Omega_{Y/S}$ and $\Omega_{X/S}$. This intuition is captured by the following basic property of $\Omega_{X/Y}$.

Proposition 11.3.3. *Let $f : X \to Y$ and $g : Y \to Z$ be morphisms of schemes. Then we have an exact sequence*

$$f^*\Omega_{Y/Z} \to \Omega_{X/Z} \to \Omega_{X/Y} \to 0.$$

Proof. Follows from Proposition 11.1.9 and compatibility with localization. □

We will usually work with relative cotangent sheaves for finitely presented morphisms. In this case, the relative cotangent sheaf can be computed locally using the Jacobian matrix.

Remark 11.3.4. Suppose that X is a $\mathbb{K}$-scheme and that $x \in X$ is a $\mathbb{K}$-point. Then the Zariski cotangent space $T^\vee_{X,x}$ is the same as the fiber of $\Omega_{X/\mathbb{K}}$ at the point x. Indeed, by Corollary 5.3.5 and Lemma 11.1.12 both spaces can be computed via the Jacobian.

However, Example 5.1.9 shows that for arbitrary closed points x the Zariski cotangent space $T^\vee_{X,x}$ and the fiber $\Omega_{X/\mathbb{K}}(x)$ can be different. The fiber $\Omega_{X/\mathbb{K}}(x)$ is defined using derivations on a neighborhood of x, while $T^\vee_{X,x}$ is defined using the derivations of $\mathfrak{m}_x$. We discuss this comparison more in Section 11.4.1.

Before moving on, we record another basic property of the relative cotangent sheaf.

Proposition 11.3.5. *Let $f : X \to Y$ be a morphism of schemes. Fix a morphism $Z \to Y$ and let $g : X \times_Y Z \to X$ denote the projection map. Then we have $\Omega_{X \times_Y Z/Z} = g^* \Omega_{X/Y}$.*

Proof. Follows from Proposition 11.1.8 and compatibility with localization. □

11.3.2 Relative tangent sheaf

We define the relative tangent sheaf by dualizing.

Definition 11.3.6. Given a morphism $f : X \to Y$, we define the relative tangent sheaf $T_{X/Y}$ to be the dual $\Omega_{X/Y}^\vee$.

If $\Omega_{X/Y}$ is locally free then the relative cotangent sheaf can be recovered from the relative tangent sheaf by dualizing again. However, if $\Omega_{X/Y}$ is not locally free then $T_{X/Y}$ can "carry less information" than $\Omega_{X/Y}$; for example, any information about the torsion submodule of $\Omega_{X/Y}$ is lost upon dualizing. For this reason it is much less common to work with the relative tangent sheaf unless we include a local freeness hypothesis.

11.3.3 Differential forms

Given a morphism $f : X \to Y$, we define the sheaf of relative differential p forms to be

$$\Omega^p_{X/Y} = \bigwedge^p \Omega_{X/Y}.$$

Just as with the tangent sheaf, the sheaves of differential forms are most useful when $\Omega_{X/Y}$ is locally free (in which case the $\Omega^p_{X/Y}$ are also locally free). Using the sheaves of differential forms one can develop an "algebraic de Rham theory" which closely parallels the geometric theory.

The most important example of a sheaf of differential forms is the top exterior power $\Omega^n_{X/\mathbb{K}}$ when X is a $\mathbb{K}$-scheme of dimension n. We will return to this example in Section 11.5.4.

11.3.4 Projective space

Since projective space $\mathbb{P}^n$ is obtained by gluing together copies of $\mathbb{A}^n$, the conormal sheaf of projective space $\mathbb{P}^n$ is a locally free sheaf of rank n. However this locally free sheaf cannot be decomposed into simpler sheaves (e.g. it will not split into a direct sum of invertible sheaves). The best way of working with $\Omega_{\mathbb{P}^n}$ is a fundamental exact sequence known as the Euler sequence.

Proposition 11.3.7. *Let $Y = \operatorname{Spec}(R)$ be an affine scheme. Then the conormal sheaf of $\mathbb{P}^n_R$ over Y fits into an exact sequence*

$$0 \to \Omega_{\mathbb{P}^n_R/Y} \to \mathcal{O}(-1)^{\oplus n+1} \to \mathcal{O} \to 0.$$

The rightmost map is defined by multiplying the ith summand of $\mathcal{O}(-1)^{\oplus n+1}$ *by the variable* x_i.

Our strategy to construct $\Omega_{\mathbb{P}^n_R/Y}$ is to glue together the constant bundles over the various affine charts using the derivatives of the transition maps.

Proof. Set $S = R[x_0, \ldots, x_n]$. We define $E = S(-1)^{\oplus n+1}$ and let e_i denote the ith basis vector of E (so that e_i has degree 1). Consider the graded homomorphism $\phi : E \to S$ sending e_i to x_i and let M denote the kernel of this map. Note that although ϕ is not surjective, it is surjective in degree ≥ 1. Thus we obtain an exact sequence of sheaves

$$0 \to \widetilde{M^+} \to \mathcal{O}(-1)^{\oplus n+1} \to \mathcal{O} \to 0.$$

Suppose we localize at x_j. Using the fact that the kernel of ϕ is generated by the various $x_i e_j - x_j e_i$, we see that the localized kernel M_{x_j} is the free S_{x_j}-module generated by the various $e_i - \frac{x_i}{x_j} e_j$. Taking degree 0 parts, we see that $(M_{x_j})_0$ is a free module generated by the various $\frac{1}{x_j} e_i - \frac{x_i}{x_j^2} e_j$.

We now construct an isomorphism $\psi : \Omega_{\mathbb{P}^n_R/Y} \to \widetilde{M^+}$. On the open affine $U_j := D_{+,x_j}$ note that $\Omega_{\mathbb{P}^n_R/Y}$ is generated by $d(\frac{x_0}{x_j}), \ldots, d(\frac{x_n}{x_j})$. We define

$$\begin{aligned} \psi_j : \Omega_{\mathbb{P}^n_R/Y}|_{U_j} &\to \widetilde{M^+}|_{U_j} \\ d(\tfrac{x_i}{x_j}) &\mapsto \frac{1}{x_j} e_i - \frac{x_i}{x_j^2} e_j \end{aligned}$$

Since both sides are represented by free modules of the same rank, this map is an isomorphism. (Conceptually speaking e_i is representing "dx_i" and the isomorphism is just the quotient rule for $\frac{x_i}{x_j}$.)

We claim that as we vary i the ψ_i glue to give a global isomorphism ψ. Indeed, on $U_i \cap U_j$ we have the identification $\frac{x_k}{x_i} = \frac{x_j}{x_i} \cdot \frac{x_k}{x_j}$. Taking derivatives and rearranging we obtain the coordinate transformation rule for projective space:

$$d(\tfrac{x_k}{x_j}) = \tfrac{x_i}{x_j} \cdot d(\tfrac{x_k}{x_i}) - \tfrac{x_i x_k}{x_j^2} \cdot d(\tfrac{x_j}{x_i}).$$

These transform in the same way as the basis vectors for $\widetilde{M^+}|_{U_j}$:

$$\left(\frac{1}{x_j} e_k - \frac{x_k}{x_j^2} e_j\right) = \frac{x_i}{x_j} \left(\frac{1}{x_i} e_k - \frac{x_k}{x_i^2} e_i\right) - \frac{x_i x_k}{x_j^2} \left(\frac{1}{x_i} e_j - \frac{x_j}{x_i^2} e_i\right)$$

$\square$

The dual sequence is also known as the Euler sequence:

$$0 \to \mathcal{O} \to \mathcal{O}(1)^{\oplus n+1} \to T_{\mathbb{P}^n_R/Y} \to 0$$

Proposition 11.3.7 constructs the cotangent sheaf of projective space using a "gluing" perspective. We can also construct it using a "quotient" perspective:

Example 11.3.8. Consider the quotient map $f : \mathbb{A}^{n+1}\backslash\{0\} \to \mathbb{P}^n$. Note that the vector field $\frac{d}{dx_i}$ does not descend to $\mathbb{P}^n$ since it is not rescaling invariant. However, for a linear function ℓ the vector field $\ell\frac{d}{dx_i}$ will descend to $\mathbb{P}^n$. Since $\mathcal{O}(1)$ on $\mathbb{P}^n$ parametrizes the linear functions on $\mathbb{A}^{n+1}$, the sheaf $\mathcal{O}(1)^{\oplus n+1}$ parametrizes the linear vector fields on $\mathbb{A}^{n+1}\backslash\{0\}$ and we get a map $\phi : \mathcal{O}(1)^{\oplus n+1} \to T_{\mathbb{P}^n}$.

Suppose that $\sum f_i\frac{d}{dx_i}$ is a linear vector field in the kernel of the pushforward map. This means that it should be "radial", i.e. proportional to the vector field $\sum x_i\frac{d}{dx_i}$ obtained by taking the normal vectors to spheres. Thus we see that the kernel of the map ϕ is isomorphic to $\mathcal{O}$ and the map $\mathcal{O} \to \mathcal{O}(1)^{\oplus n+1}$ is exactly multiplication by x_i on the ith coordinate vector.

Let's make this more precise using sheaves.

Construction 11.3.9. The quotient morphism $f : \mathbb{A}^{n+1}\backslash\{0\} \to \mathbb{P}^n$ is locally described by coordinate maps

$$\mathbb{K}[\tfrac{x_0}{x_i}, \dots, \tfrac{x_n}{x_i}] \to \mathbb{K}[x_0, \dots, x_n]_{x_i}.$$

Note that the right hand side is simply the polynomial ring over the left hand side obtained by adjoining the single variable x_i (reflecting the fact that the preimage of D_{+,x_i} is isomorphic to $\mathbb{A}^1 \times D_{+,x_i}$). Consider the cotangent sequence

$$0 \to f^*\Omega_{\mathbb{P}^n} \to \Omega_{\mathbb{A}^{n+1}\backslash\{0\}} \to \Omega_{\mathbb{A}^{n+1}\backslash\{0\}/\mathbb{P}^n} \to 0.$$

Over the chart D_{+,x_i} the leftmost map map sends

$$d\left(\tfrac{x_j}{x_i}\right) \mapsto \frac{1}{x_i}dx_j - \frac{x_j}{x_i^2}dx_i.$$

The cokernel $\Omega_{\mathbb{A}^{n+1}\backslash\{0\}/\mathbb{P}^n}$ is spanned by dx_i; in fact, the "natural" generator is $\frac{dx_i}{x_i}$ since this expression is invariant as we change charts. The rightmost map in the cotangent sequence is the quotient assigning $d(\frac{x_j}{x_i}) = 0$ for $j \neq i$, or in other words, the map sending $dx_j \mapsto x_j\frac{dx_i}{x_i}$.

We now pushforward the cotangent sequence to $\mathbb{P}^n$ to get

$$0 \to \Omega_{\mathbb{P}^n} \otimes f_*\Omega_{\mathbb{A}^{n+1}\backslash\{0\}} \to f_*\Omega^{\oplus n+1}_{\mathbb{A}^{n+1}\backslash\{0\}} \to f_*\Omega_{\mathbb{A}^{n+1}\backslash\{0\}/\mathbb{P}^n}$$

Note that $f_*\mathcal{O}_{\mathbb{A}^{n+1}\backslash\{0\}} \cong \oplus_{d\in\mathbb{Z}}\mathcal{O}(d)$. Furthermore, as explained above the maps in the cotangent sequence are naturally graded: the leftmost map has image in degree -1 and the rightmost map increases the degree by 1. Thus after adjusting the grading we get an exact sequence of graded $\mathcal{O}_{\mathbb{P}^n}$-algebras

$$0 \to \oplus_{d\in\mathbb{Z}}\Omega_{\mathbb{P}^n}(d) \to \oplus_{d\in\mathbb{Z}}\mathcal{O}(d-1)^{\oplus n+1} \to \oplus_{d\in\mathbb{Z}}\mathcal{O}(d).$$

The degree 0 part is the Euler sequence.

11.3.5 Exercises

Exercise 11.3.10. Let X be a $\mathbb{K}$-scheme. Show that for every point $x \in X$ we have that $\operatorname{rk}_x \Omega_{X/\mathbb{K}}$ is at least the local dimension of x in X. (Hint: reduce to the case when X is a variety and consider the generic point of X.)

Exercise 11.3.11. Let X be a scheme and let $\mathcal{E}$ be a locally free sheaf of rank r on X. Let $\mathbb{P}(\mathcal{E})$ denote the projective bundle associated to $\mathcal{E}$ equipped with the structure map $\pi : \mathbb{P}(\mathcal{E}) \to X$. Prove that there is a "relative Euler sequence"

$$0 \to \Omega_{\mathbb{P}(\mathcal{E})/X} \to \mathcal{O}(-1) \otimes \pi^*\mathcal{E} \to \mathcal{O}_{\mathbb{P}(\mathcal{E})} \to 0.$$

Exercise 11.3.12. Suppose we have morphisms $f : X \to Z$ and $g : Y \to Z$. Construct an isomorphism $\Omega_{X\times_Z Y/Z} \cong p_1^*\Omega_{X/Z} \times p_2^*\Omega_{Y/Z}$ where p_1, p_2 are the two projection maps. Show that the cotangent sequence

$$p_1^*\Omega_{X/Z} \to \Omega_{X\times_Z Y/Z} \to \Omega_{X\times_Z Y/Y} \to 0$$

is exact on the left and is split exact (and similarly if we swap the roles of X and Y).

11.4 Conormal sheaves

Suppose that $f : X \to Y$ is a closed embedding of smooth manifolds. The cokernel of the map $T_X \to f^*T_Y$ is called the normal bundle of X in Y and is denoted by $N_{X/Y}$. Dually, the conormal bundle is the kernel of the map $f^*\Omega_Y \to \Omega_X$. Our next goal is to transport this construction to the realm of algebraic geometry.

Definition 11.4.1. Let $f : Z \to X$ be a closed embedding and let $\mathcal{I}$ denote the quasicoherent ideal sheaf defining Z. The quotient $\mathcal{I}/\mathcal{I}^2$ can be thought of as a quasicoherent sheaf on Z via the identification $\mathcal{I}/\mathcal{I}^2 \cong f^*\mathcal{I}$. With this identification $\mathcal{I}/\mathcal{I}^2$ is called the conormal sheaf of Z in X and is frequently denoted by $N^\vee_{Z/X}$.

The dual $\mathcal{H}om(N^\vee_{Z/X}, \mathcal{O}_Z)$ is called the normal sheaf and is denoted by $N_{Z/X}$. Note however that (just as for tangent sheaves) we may be "losing" information when we dualize so that the conormal sheaf is the more fundamental construction. When $N^\vee_{Z/X}$ is locally free we call it the conormal bundle of Z in X, and the dual $N_{Z/X}$ the normal bundle.

Example 11.4.2. Let D be an effective Cartier divisor in X. Then the ideal sheaf for D is the invertible sheaf $\mathcal{O}_X(-D)$. This means that the conormal sheaf for D is the invertible sheaf $\mathcal{O}_X(-D)|_D$ on D. It is important to note that this formula holds no matter what kinds of singularities D has.

Our terminology is motivated by the following result.

Theorem 11.4.3. *Let $f : Z \to X$ be a closed embedding of S-schemes. We have an exact sequence*

$$N^\vee_{Z/X} \to f^*\Omega_{X/S} \to \Omega_{Z/S} \to 0.$$

The sequence of Theorem 11.4.3 is known as the conormal sequence.

Proof. Follows immediately from Proposition 11.1.10 and compatibility with localization. □

Remark 11.4.4. Based on the analogy with the theory of manifolds we might expect the conormal sequence to be exact on the left. Although this can fail to be true (see Exercise 11.4.13), it turns out that it is true if we impose a smoothness hypothesis (see Theorem 11.5.8).

11.4.1 Zariski cotangent space

We next discuss the relationship between the Zariski cotangent space and the cotangent sheaf for a $\mathbb{K}$-scheme X.

Suppose that $x \in X$ is a closed point. Then the conormal sheaf of x at X is the sheaf $\mathfrak{m}/\mathfrak{m}^2$, i.e. the Zariski cotangent space at x. By Theorem 11.4.3 we get an exact sequence

$$\mathfrak{m}/\mathfrak{m}^2 \to \Omega_{X/\mathbb{K}}(x) \to \Omega_{\kappa(x)/\mathbb{K}} \to 0$$

where as usual $\Omega_{X/\mathbb{K}}(x)$ denotes the fiber of $\Omega_{X/\mathbb{K}}$ at x. In other words, we get a morphism from the Zariski cotangent space to the fiber of the cotangent sheaf.

Example 11.4.5. Let $\mathbb{K} = \mathbb{F}_p(u)$ and let $X = \mathrm{Spec}(\mathbb{K}[x,y]/(y^2 - x^p + u))$. We claim that the Zariski cotangent space of the point $\mathfrak{m} = (y, x^p - u)$ is not isomorphic to the fiber of the cotangent sheaf. As an $\mathbb{L}$-vector space the quotient $\mathfrak{m}/\mathfrak{m}^2$ is just $\mathbb{L}y$, hence one-dimensional. However $\Omega_{X/\mathbb{K}}(\mathfrak{m})$ is 2-dimensional, since the Jacobian

$$\mathrm{Jac}_f(x) = \begin{bmatrix} 0 \\ 2y \end{bmatrix}.$$

evaluates to 0 at the point $\mathfrak{m}$. The difference between these two computations is explained by the fact that $\Omega_{\kappa(\mathfrak{m})/\mathbb{K}} \neq 0$.

The previous example shows that the map $T^\vee_{X,x} \to \Omega_{X/\mathbb{K}}(x)$ need not be an isomorphism. However, it will be an isomorphism under certain circumstances:

Proposition 11.4.6. *Let X be a $\mathbb{K}$-scheme. Suppose that $x \in X$ is a regular closed point such that $\kappa(x)$ is separable over $\mathbb{K}$. Then the morphism from the Zariski cotangent space $T^\vee_{X,x}$ to the fiber $\Omega_{X/\mathbb{K}}(x)$ is an isomorphism.*

Proof. Consider the morphism $i : \mathrm{Spec}(\mathcal{O}_{X,x}) \to X$. By Proposition 11.1.7 we have $\Omega_{\mathrm{Spec}(\mathcal{O}_{X,x})/\mathbb{K}} \cong i^*\Omega_{X/\mathbb{K}}$. Consider the conormal sequence for the closed point x in $\mathrm{Spec}(\mathcal{O}_{X,x})$:

$$\mathfrak{m}/\mathfrak{m}^2 \to \Omega_{X/\mathbb{K}}(x) \to \Omega_{\kappa(x)/\mathbb{K}} \to 0$$

Proposition 11.2.1 shows that the term on the right is zero, so we get a surjection $\phi : \mathfrak{m}/\mathfrak{m}^2 \to \Omega_{X/\mathbb{K}}(x)$. However Exercise 11.3.10 shows that

$$\dim_{\kappa(x)} \Omega_{X/\mathbb{K}}(x) \geq \dim(X) = \dim_{\kappa(x)} \mathfrak{m}/\mathfrak{m}^2$$

showing that ϕ is an isomorphism. □

Remark 11.4.7. In general the kernel of $\mathfrak{m}/\mathfrak{m}^2 \to \Omega_{X/\mathbb{K}}(x)$ can be identified by extending the conormal sequence to the left. The next term in the sequence turns out to be $\Gamma = \ker(\Omega_{\mathbb{K}/\mathbb{Z}} \otimes \kappa(x) \to \Omega_{\kappa(x)/\mathbb{Z}})$.

When $x \in X$ is not a closed point, the comparison between the Zariski cotangent space and the fiber of the cotangent sheaf is a little different. Consider the morphism $i : \mathrm{Spec}(\mathcal{O}_{X,x}) \to X$. Just as before, the conormal sheaf of x in $\mathrm{Spec}(\mathcal{O}_{X,x})$ is equal to the

Zariski cotangent space at x. Since the cotangent sheaf $\Omega_{\mathrm{Spec}(\mathcal{O}_{X,x})/\mathbb{K}}$ is the same as $i^*\Omega_{X/\mathbb{K}}$ by Proposition 11.1.7, just as before we get a map $\mathfrak{m}/\mathfrak{m}^2 \to \Omega_{X/\mathbb{K}}(x)$ from the conormal sequence of x in $\mathrm{Spec}(\mathcal{O}_{X,x})$. However, now in the exact sequence

$$\mathfrak{m}/\mathfrak{m}^2 \to \Omega_{X/\mathbb{K}}(x) \to \Omega_{\kappa(x)/\mathbb{K}} \to 0$$

the rightmost term is no longer zero. If $x \in X$ is a regular point and $\kappa(x)$ is separable over X, then the leftmost term has dimension $\mathrm{codim}(\overline{x})$ and the rightmost term has dimension $\dim(\overline{x})$ so this sequence is exact by a similar argument as before.

11.4.2 Regular embeddings

Definition 11.4.8. A closed embedding $f : Z \to X$ is a regular embedding of codimension r at a point $x \in Z$ if the stalk of $\mathcal{I}_Z$ at the point x is defined by a regular sequence of length r in the local ring $\mathcal{O}_{X,x}$.

A regular embedding of codimension r is a closed embedding $f : Z \to X$ which is regular at every point of Z.

Example 11.4.9. If X is a locally Noetherian scheme, then a regular embedding of codimension 1 is the same thing as a Cartier divisor on X. Indeed, in this situation every ideal sheaf $\mathcal{I}$ is locally finitely generated. If the stalk of $\mathcal{I}_x$ is generated by a single element then using Geometric Nakayama's lemma we see that $\mathcal{I}$ is also principal on an open neighborhood of X.

The terminology is motivated by the following important example.

Theorem 11.4.10. *Let X be a regular Noetherian scheme. Suppose that Z is a closed subscheme of X that is also regular. Then the closed embedding $f : Z \to X$ is a regular embedding.*

Proof. Fix a point $x \in Z$. Let $A = \mathcal{O}_{X,x}$ be the local ring at x with maximal ideal $\mathfrak{m}$. Let I be the ideal in $\mathcal{O}_{X,x}$ defining Z, and define $B := A/I$ with the maximal ideal $\mathfrak{n}$. By assumption B is a regular local ring; we must show that I is generated by a regular sequence.

Suppose that $\dim(A) = n$ and that $\dim(B) = d$. We can lift generators of $\mathfrak{n}$ to A to get d linearly indepedent elements in $\mathfrak{m}\backslash\mathfrak{m}^2$. Choose any $n-d$ elements in the kernel of $\mathfrak{m}/\mathfrak{m}^2 \to \mathfrak{n}/\mathfrak{n}^2$ which complete the basis; by lifting these elements to A we obtain elements $f_1, \ldots, f_{n-d}$ in I. Since A is regular, the sequence $f_1, \ldots, f_{n-d}$ is a regular sequence in A. Indeed, for any $i < n-d$ Krull's PIT shows that $\dim(A/(f_1, \ldots, f_i)) \geq n-i$ and Exercise 5.1.6 shows that the Zariski tangent space of this ring has dimension $\leq n-i$. By Theorem 5.2.1 we see that both quantities must be equal to $n-i$. Thus $A/(f_1, \ldots, f_i)$ is regular, hence a domain by Proposition 5.2.5. So f_{i+1} cannot be a zero divisor in this ring.

We claim that $I = (f_1, \ldots, f_{n-d})$. Note that there is a surjection $A/(f_1, \ldots, f_{n-d}) \to B$. Since both are regular local rings, they are both domains. But a surjection of two integral domains which have the same dimension must be an isomorphism. $\square$

We next show that the normal sheaf for a regular embedding has particularly nice properties.

Proposition 11.4.11. *Let A be a Noetherian ring, $I \subset A$ an ideal generated by a regular sequence $(f_1, \ldots, f_r)$. Then the map $\phi : (A/I)^{\oplus r} \to I/I^2$ sending $(a_1, \ldots, a_r) \mapsto \sum a_i f_i$ is an isomorphism of A/I-modules. In particular, I/I^2 is a free A/I-module of rank r.*

Proof. It is clear that ϕ is surjective. To see that ϕ is injective, we consider the analogous map $A^{\oplus r} \to I/I^2$ and analyze the kernel. More precisely, we claim that if $(a_1, \ldots, a_r)$ is in the kernel of ϕ then $a_r \in I$. Indeed, since $\sum a_i f_i = 0$ we see that $a_r f_r = 0$ in $A/(f_1, \ldots, f_{r-1}) + I^2$. In other words, we have an equality $a_r f_r = b f_r^2$ in the ring $A/(f_1, \ldots, f_{r-1})$. Since the f_i form a regular sequence f_r is not a zero divisor in the quotient ring $A/(f_1, \ldots, f_{r-1})$ and we deduce that $a_r - b f_r \in (f_1, \ldots, f_{r-1})$. We conclude that $a_r \in I$.

Since A is Noetherian the order of elements can be permuted in a regular sequence. We conclude that the kernel of $A^{\oplus r} \to I/I^2$ consists exactly of the tuples of elements in I, showing that ϕ is injective. □

Applying the previous proposition locally, we obtain:

Corollary 11.4.12. *Let $f : Z \to X$ be a regular embedding of Noetherian schemes of codimension r. Then the conormal sheaf $N^\vee_{Z/X}$ is a locally free sheaf of rank r.*

11.4.3 Exercises

Exercise 11.4.13. Let $Z \subset \mathbb{A}^2_{\mathbb{K}}$ be the subscheme defined by (x^2, y^2). Show that the map $I/I^2 \to \Omega_{\mathbb{A}^2/\mathbb{K}}|_Z$ is not injective. What is its kernel?

Exercise 11.4.14. Let $f : X \to Y$ be morphism of schemes. Prove that the conormal sheaf $\mathcal{I}/\mathcal{I}^2$ of $\Delta_{X/Y}$ is isomorphic to the relative cotangent sheaf $\Omega_{X/Y}$. (Note that $\Delta_{X/Y}$ is in general only a locally closed subscheme and not necessarily a closed subscheme. To define the "conormal sheaf" of a locally closed subscheme, we first choose an open set U so that our subscheme is closed in U. One can show that the resulting sheaf is independent of the choice of U.)

This construction is often used as the definition of the relative cotangent sheaf. It should be compared to the well-known fact that if X is a smooth manifold then the normal bundle of Δ_X in $X \times X$ is isomorphic to T_X.

Exercise 11.4.15. Suppose that $f : X \to Y$ and $g : Y \to Z$ are both closed immersions. Prove that there is an exact sequence of sheaves on X given by

$$f^* N^\vee_{Y/Z} \to N^\vee_{X/Z} \to N^\vee_{X/Y} \to 0$$

Find an example where the leftmost map is not injective.

11.5 Smooth varieties

If X is an n-dimensional $\mathbb{K}$-variety then Exercise 11.3.10 shows that the rank of $\Omega_{X/\mathbb{K}}$ is at least n at every point. Just as with the Zariski cotangent space, we can expect the cotangent sheaf to have rank $> n$ when X is "singular" in some way. We will use this intuition to define smooth varieties.

Definition 11.5.1. Suppose that X is a finite type scheme over a field $\mathbb{K}$. We will say that X is smooth (over $\mathrm{Spec}(\mathbb{K})$) at a point x if $\Omega_{X/\mathbb{K}}$ is locally free at x with rank equal to the local dimension of x in X.

We say that X is smooth if it is smooth at every point – that is, $\Omega_{X/\mathbb{K}}$ is locally free and its rank at every point is the local dimension.

Note that smoothness does not require any separability hypothesis (although we will often add it in unnecessarily since the focus of this section is $\mathbb{K}$-schemes).

Remark 11.5.2. The Zariski-Lipman conjecture predicts that a complex variety is smooth if and only if its *tangent* sheaf $T_{X/\mathbb{C}}$ is locally free. This has been verified in many cases but is still open in general.

Smoothness behaves well with respect to change of base field. (This is one of the key properties that distinguishes smoothness from regularity.)

Proposition 11.5.3. *Let X be a $\mathbb{K}$-scheme. Suppose that $\mathbb{L}/\mathbb{K}$ is an extension of fields. Then X is smooth if and only if the base change $X_{\mathbb{L}}$ is smooth.*

Proof. Let $\rho : X_{\mathbb{L}} \to X$ be the base change morphism. By arguing as in Exercise 4.3.23 we see that the local dimension is preserved by base change. Proposition 11.3.5 shows that $\Omega_{X_{\mathbb{L}}/\mathbb{L}} \cong \rho^*\Omega_{X/\mathbb{K}}$. We claim that $\Omega_{X/\mathbb{K}}$ is locally free of rank r if and only if its ρ-pullback is locally free of rank r. For the forward implication, we simply note that the pullback of a locally free sheaf is still locally free. Conversely, if $\rho^*\Omega_{X/\mathbb{K}}$ is locally free, then it is flat over X. Thus the original sheaf $\Omega_{X/\mathbb{K}}$ is flat and coherent, hence locally free. It is clear that the rank is preserved by pullback, finishing the proof. $\square$

11.5.1 Smoothness versus regularity

We next compare smoothness and regularity. The next example shows that these two notions are not the same.

Example 11.5.4. Let $\mathbb{K} = \mathbb{F}_p(u)$ and let $X = \mathrm{mSpec}(\mathbb{K}[x,y]/(y^2 - x^p + u))$. As showed in Example 11.4.5, the Zariski cotangent space of the point $\mathfrak{m} = (y, x^p - u)$ is 1-dimensional while the fiber $\Omega_{X/\mathbb{K}}(\mathfrak{m})$ is 2-dimensional. Since X has dimension 1 by Krull's PIT, we see that $\mathfrak{m}$ is a regular point which is not smooth.

The following result summarizes the relationship between smoothness and regularity:

Theorem 11.5.5. *Let X be a $\mathbb{K}$-scheme. Then:*

(1) If x is a smooth point of X, then x is also a regular point.

(2) If x is a regular point of X and $\kappa(x)$ is separable over $\mathbb{K}$, then x is a smooth point.

Proof. The proof of the first implication will be deferred to Proposition 11.9.10. To see the second, recall that we have an exact sequence

$$0 \to \mathfrak{m}/\mathfrak{m}^2 \to \Omega_{X/\mathbb{K}}(x) \to \Omega_{\kappa(x)/\mathbb{K}} \to 0$$

The term on the left has rank $\operatorname{codim}(\overline{x})$ and the term on the right has rank $\dim(\overline{x})$, so $Omega_{X/\mathbb{K}}$ has rank $\dim(X)$ at x. By upper semicontinuity the rank of $\Omega_{X/\mathbb{K}}$ is $\leq \dim(X)$ on an open neighborhood of x. Since by Exercise 11.3.10 $\dim(X)$ is a lower bound on the rank, we see that the rank is equal to $\dim(X)$ on an open neighborhood of x.

Since x is regular, it is a reduced point of X. The reduced locus of X is closed, so we conclude that there is an open neighborhood of x which is reduced. Combining with the discussion above, we see that x has a reduced open neighborhood on which $\Omega_{X/\mathbb{K}}$ has constant rank, and thus $\Omega_{X/\mathbb{K}}$ is locally free at x. □

We obtain as an important corollary:

Corollary 11.5.6. *Suppose that X is a smooth $\mathbb{K}$-scheme. Then X is reduced and no two irreducible components of X can intersect.*

The following bullet points summarize the relationship between smoothness and regularity:

- Smoothness is a relative concept – the definition of the relative cotangent sheaf $\Omega_{X/\mathbb{K}}$ implicitly depends on the structure morphism $X \to \operatorname{Spec}(\mathbb{K})$. In contrast, regularity is an absolute concept that does not depend upon a morphism.

- Regularity is more sensitive than smoothness: every smooth variety is regular but a regular variety need not be smooth. Thus theorems with a regularity hypothesis are stronger than theorems with a smoothness hypothesis. On the other hand, smoothness is better behaved than regularity and is easier to work with.

- Smoothness can be thought of as "geometric regularity": Proposition 11.4.6 shows that after a base change to an algebraically closed field the Zariski cotangent space at a closed point can be identified with the fiber of the cotangent sheaf. Thus a $\mathbb{K}$-scheme will be smooth if and only if its base change to an algebraic closure is regular.

11.5.2 Generic smoothness

Recall that a finitely generated field extension is separably generated if it is the composition of a purely transcendental extension followed by a separable algebraic extension.

Proposition 11.5.7 (Generic smoothness for varieties)**.** *Let X be a $\mathbb{K}$-variety. Then there is an open subset $U \subset X$ such that every point in U is smooth if and only if the function field $K(X)$ is separable over $\mathbb{K}$.*

In particular the proposition applies to any variety X defined over a perfect field $\mathbb{K}$.

Proof. Since the rank is an upper semicontinuous function and X is reduced, we see that the rank of $\Omega_{X/\mathbb{K}}$ will be equal to $\dim(X)$ along an open subset of X if and only if the rank of $\Omega_{X/\mathbb{K}}$ at the generic point is equal to $\dim(X)$. We can identify the stalk of $\Omega_{X/\mathbb{K}}$ at the generic point of X with $\Omega_{K(X)/\mathbb{K}}$. Then Corollary 11.2.9 shows that the rank of $\Omega_{X/\mathbb{K}}$ at the generic point is equal to $\dim(X)$ if and only if $K(X)$ is separable over $\mathbb{K}$. $\square$

11.5.3 Conormal bundles of smooth varieties

The conormal sheaf sequence is exact on the left under a smoothness hypothesis.

Theorem 11.5.8. *Let X and Z be smooth $\mathbb{K}$-varieties with a closed embedding $f : Z \to X$. Then the conormal sheaf sequence is exact on the left, that is, the sequence*

$$0 \to N^{\vee}_{Z/X} \to f^*\Omega_{X/\mathbb{K}} \to \Omega_{Z/\mathbb{K}} \to 0$$

is exact.

The assumptions of Theorem 11.5.8 are stronger than necessary in order to simplify the proof. It turns out that it is enough to assume that Z is smooth over $\mathbb{K}$ (see [Sta15, 06CD]).

Proof. By Theorem 11.5.5 a smooth scheme is also regular. Thus Theorem 11.4.10 and Corollary 11.4.12 show that the conormal sheaf of Z in X is locally free of rank $\dim(X) - \dim(Z)$. Note that $f^*\Omega_{X/\mathbb{K}}$ and $\Omega_{Z/\mathbb{K}}$ are also locally free of ranks $\dim(X)$ and $\dim(Z)$ respectively. By comparing ranks, we see that the image of the map $N^{\vee}_{Z/X} \to f^*\Omega_{X/\mathbb{K}}$ must have rank $\dim(X) - \dim(Z)$ at every point. But since $N^{\vee}_{Z/X}$ is locally free this implies that the map is injective. $\square$

11.5.4 Canonical bundle

Instead of working directly with the cotangent sheaf, it is often easier to work with its top exterior power.

Definition 11.5.9. Suppose that X is a smooth $\mathbb{K}$-variety of dimension n. The canonical bundle on X is the top exterior power of the cotangent sheaf:

$$\omega_X = \bigwedge^n \Omega_{X/\mathbb{K}}.$$

In a topological setting, the first chern class of ω_X will coincide with the first chern class of $\Omega_{X/\mathbb{K}}$. Thus one can think of ω_X as a line bundle which captures the "curvature" of X. There are a number of deep conjectures predicting that the positivity (i.e. ampleness) of the canonical bundle controls the geometric features of X, such as the existence of rational points.

Example 11.5.10. The canonical bundle of $\mathbb{P}^n$ is isomorphic to $\mathcal{O}_{\mathbb{P}^n}(-n-1)$. This can be deduced from the Euler sequence

$$0 \to \Omega_{\mathbb{P}^n} \to \mathcal{O}(-1)^{\oplus n+1} \to \mathcal{O} \to 0$$

by taking top exterior powers. More explicitly, $\omega_{\mathbb{P}^n}$ is locally generated on our standard chart D_{+,x_i} by the meromorphic n-form

$$\sigma_i := \frac{x_i}{x_0} d\left(\frac{x_0}{x_i}\right) \wedge \frac{x_i}{x_1} d\left(\frac{x_1}{x_i}\right) \wedge \ldots \wedge \frac{x_i}{x_n} d\left(\frac{x_n}{x_i}\right).$$

The main tool for computing the canonical bundle is called the adjunction formula. The most important case is when we have a smooth Cartier divisor Y in a smooth variety X.

Proposition 11.5.11 (Adjunction)**.** *Let Y be a smooth Cartier divisor in a smooth variety X. Then*

$$\omega_Y \cong (\omega_X \otimes \mathcal{O}_X(Y))|_Y.$$

Proof. The conormal sequence for $Y \subset X$ is left exact by Theorem 11.5.8. Since the leftmost term is simply $\mathcal{O}_X(-Y)|_Y$, the desired formula follows from taking top exterior powers in the conormal sequence. □

11.5.5 Exercises

Exercise 11.5.12 (Jacobian Criterion)**.** Let $X \subset \mathbb{A}^n$ be a closed subscheme defined by a homogeneous ideal $(f_1, \ldots, f_r)$. For any point $x \in X$, show that x is a smooth point of X if and only if the rank of the Jacobian $\mathrm{Jac}_{f_1,\ldots,f_r}(x)$ is equal to $n - \dim_x(X)$. (Hint: for the reverse implication, first show that X is regular, hence reduced.)

This means that the local freeness assumption in our definition of smoothness is redundant and can be left out.

Exercise 11.5.13 (Projective Jacobian Criterion)**.** Let $X \subset \mathbb{P}^n$ be a closed subscheme defined by a homogeneous ideal $(f_1, \ldots, f_r)$. For any point $x \in X$, show that x is a smooth point of X if and only if the rank of the projective Jacobian $(r \times (n+1))$-matrix $\mathrm{Jac}_{f_1,\ldots,f_r}(x)$ is equal to $n - \dim_x(X)$.

Note that the partial derivatives of the projective Jacobian matrix are not well-defined functions on the points of X – however, the rank of the matrix is well-defined.

Exercise 11.5.14. Let $f : X \to Y$ be a closed embedding of irreducible $\mathbb{K}$-schemes. Suppose that X and Y are smooth and have the same dimension. Show that f is an isomorphism from X to Y. (Hint: use the fact that smooth schemes are regular.)

Exercise 11.5.15. Let $\mathbb{K}$ be a perfect field. Let X be a smooth $\mathbb{K}$-variety and let Z be an integral $\mathbb{K}$-variety equipped with a regular embedding $f : Z \to X$. Show that the conormal sheaf sequence

$$0 \to \mathcal{I}/\mathcal{I}^2 \to f^*\Omega_{X/\mathbb{K}} \to \Omega_{Z/\mathbb{K}} \to 0$$

is exact. (Hint: use the fact that Z is generically smooth.)

Exercise 11.5.16. Suppose that X and Y are birationally equivalent smooth projective $\mathbb{K}$-varieties. Show that for every $i \geq 1$ we have an isomorphism

$$\Gamma(X, \Omega^i_{X/\mathbb{K}}) \cong \Gamma(Y, \Omega^i_{Y/\mathbb{K}}).$$

(Hint: fix a birational equivalence $\phi : X \dashrightarrow Y$. By Exercise 8.8.16 there is an open subset U_X whose complement has codimension 2 such that ϕ is a morphism on U. Use this open subset to construct a pullback map $\Gamma(Y, \Omega^i_{Y/\mathbb{K}}) \to \Gamma(X, \Omega^i_{X/\mathbb{K}})$ and apply Exercise 9.5.29.)

11.6 Cotangent bundles of curves

In this section we systematically discuss the cotangent sheaf of projective geometrically integral curves defined over a field. The most important definition in this section is the genus, a fundamental invariant of a curve that we will use many times in the future.

11.6.1 Genus

For a smooth curve C over a field $\mathbb{K}$ the cotangent sheaf $\Omega_{C/\mathbb{K}}$ is an invertible sheaf. The following definition identifies the basic invariant of a smooth projective curve.

Definition 11.6.1. Let C be a smooth projective geometrically integral curve. The genus of C is defined by the formula

$$\deg(\Omega_{C/\mathbb{K}}) = 2g(C) - 2$$

Remark 11.6.2. As always, one should think of the degree as a "topological" invariant. It turns out that complex curves are classified up to topological equivalence (but not up to algebraic equivalence) by their genus.

Exercise 11.6.3. Prove that the genus is preserved by base change: if C is a smooth projective geometrically integral curve over a field $\mathbb{K}$, then for any field extension $\mathbb{L}/\mathbb{K}$ we have $g(C) = g(C_{\mathbb{L}})$.

Later on we will see how to connect this definition with various other notions (including the more traditional definition using the Betti numbers of C). In particular, in Section 12.5.1 we will prove the following fact:

Fact 11.6.4. Let C be a smooth projective geometrically integral curve over a field $\mathbb{K}$. Then the genus of C is a non-negative integer.

This implies that the cotangent sheaf of C always has even degree and that $\deg(\Omega_{C/\mathbb{K}}) \geq -2$. For now we will assume this fact and see how to work with this important definition.

Example 11.6.5. Suppose that C is a smooth geometrically integral plane curve of degree d. We can compute the genus of the curve using the adjunction formula of Proposition 11.5.11:

$$\begin{aligned} \deg(\Omega_{C/\mathbb{K}}) &= \deg(\mathcal{O}(-3+d)|_C) \\ &= 3d - d^2. \end{aligned}$$

In this way we obtain the genus formula for C:

$$g(C) = 1 + \frac{3d - d^2}{2} = \binom{d-1}{2}$$

11.6.2 Hyperelliptic curves

We next give an extended example giving us the opportunity to work with the cotangent sheaf in a hands-on way. In particular, this example illustrates how to construct a curve of arbitrary genus.

Example 11.6.6 (Hyperelliptic curves)**.** Let $\mathbb{K}$ be an algebraically closed field of characteristic $\neq 2$. Let d be a positive even integer and fix a degree d polynomial $f(x)$ with non-zero constant coefficient. Let C be the curve in $\mathbb{A}^2_{\mathbb{K}}$ defined by the equation $y^2 = f(x)$. The sheaf $\Omega_{U/\mathbb{K}}$ will be defined by the quotient

$$A\,dx \oplus A\,dy/(2y\,dy - f'(x)\,dx)$$

where $A = \mathbb{K}[x,y]/(y^2 - f(x))$ is the coordinate ring of U. This sheaf will be locally free so long as y and $f'(x)$ do not simultaneously vanish at any point of U. (Since $y = 0$ implies $f(x) = 0$, this is equivalent to asking that $f(x)$ has no double roots.) With this assumption $\Omega_{U/\mathbb{K}}$ will be generated by $\frac{f'(x)}{2y}\,dx$ when $y \neq 0$ and by $\frac{2y}{f'(x)}\,dy$ when $f'(x) \neq 0$.

Although we could naturally compactify C in $\mathbb{P}^2$, the resulting curve would almost never be smooth. (Check!) Instead, a better way to compactify C is to take the map $f : C \to \mathbb{A}^1$ given by projecting onto the x-coordinate and compactify it to obtain a morphism $f : \overline{C} \to \mathbb{P}^1$. Let's use the coordinates u, x for the base $\mathbb{P}^1$. One way to execute this plan is to think of C as a closed subscheme of $\mathbb{P}^1 \times \mathbb{A}^1_x$ and then take a flat limit in $\mathbb{P}^1 \times \mathbb{P}^1_{u,x}$. It is even better to take the flat limit inside of the Hirzebruch surface $\mathbb{F}_{d/2}$ equipped with the projection map $\mathbb{F}_{d/2} \to \mathbb{P}^1_{u,x}$. The effect of this change is that to work on the other coordinate patch, the appropriate coordinates are

$$u = \frac{1}{x} \qquad\qquad z = \frac{y}{x^{d/2}}.$$

Thus the intersection of $\overline{C}$ with the other coordinate patch $\mathbb{A}^2_{u,z}$ is given by the equation

$$z^2 = u^d f\left(\frac{x}{u}\right)$$

which has the same form as our original equation. In particular $\overline{C}$ is a smooth projective curve.

Let's compute the degree of $\Omega_{\overline{C}/\mathbb{K}}$. To do this, we should compute a rational section and identify its zeros and poles. For example, consider the section dx of $\Omega_{C/\mathbb{K}}$. This will only vanish at points where $\Omega_{C/\mathbb{K}}$ is generated by dy. Since $dx = \frac{2y}{f'(x)}dy$, the section dx will vanish precisely at the points defined by $y = 0$, i.e. at the d distinct roots of the equation $f(x)$. (This should be no surprise; if we think of dx as the pullback of the cotangent bundle on $\mathbb{A}^1_x$ we see that dx will vanish precisely where the curve has vertical tangent lines.) The section dx will not have any poles on the chart C

We next turn to the new chart C'. In the new coordinates we have $dx = -\frac{1}{u^2}du$. We only care about zeros and poles which didn't appear on our earlier chart, or in other words, which appear at the points where $u = 0$. There are two such points (corresponding to two options of the sign of z), and the section $-\frac{1}{u^2}du$ has a double pole at each. Altogether the total contribution of the poles to the degree is -4. Thus the degree of $\Omega_{\overline{C}/\mathbb{K}}$ is $d-4$ and the genus of C is $\frac{d}{2}-1$.

11.6.3 Nodal curves

Suppose that C is a (possibly reducible) reduced curve over an algebraically closed field $\mathbb{K}$. We say that C is nodal at a point p if the completion of the local ring $\mathcal{O}_{C,p}$ along the maximal ideal satisfies

$$\widehat{\mathcal{O}_{C,p}} \cong \mathbb{K}[[x,y]]/(xy).$$

As discussed earlier, this ring represents the "formal analytic neighborhood" of the point p and is roughly analogous to looking at a small open neighborhood in the Euclidean topology. From this formal-local description we can see that C has two local branches through the point p.

The point of this definition is that many of the important local algebraic properties at p can be detected upon passing to the completion. In particular, we can understand how these properties behave for an arbitrary nodal curve by first analyzing a specific example and then using the "completion principle" to generalize to all nodes. Often the most convenient example is the union of the coordinate axes in $\mathbb{A}^2$.

Example 11.6.7. Let C be the union of the two coordinate axes in $\mathbb{A}^2_{\mathbb{K}}$ defined by the ring $R = \mathbb{K}[x,y]/(xy)$. We then know that $\Omega_{C/\mathbb{K}}$ is equal to $R\,dx \oplus R\,dy/(x\,dy + y\,dx)$.

Consider the element $\alpha = x\,dy = -y\,dx$. Then $x\alpha = y\alpha = 0$, so α is a torsion element of $\Omega_{C/\mathbb{K}}$. Taking the quotient $\Omega_{C/\mathbb{K}}/R\alpha$ is the same as adding in the relations $x\,dy = 0$ and $y\,dx = 0$, so that

$$\Omega_{C/\mathbb{K}}/R\alpha \cong R/(y) \cdot dx \oplus R/(x) \cdot dy.$$

Note that the first term in the summand is the same as the cotangent sheaf for the y-axis and the second term is the cotangent sheaf for the x-axis.

This is indicative of the general picture:

Theorem 11.6.8. *Let $\mathbb{K}$ be an algebraically closed field. Let C be a reduced (possibly reducible) nodal curve and let $\nu : C^\nu \to C$ denote the normalization map. Then $\Omega_{C^\nu/C} = 0$. Furthermore, by pushing forward the cotangent sequence under ν the surjection*

$$\Omega_{C/\mathbb{K}} \xrightarrow{\phi} \nu_*\Omega_{C^\nu/\mathbb{K}}$$

has kernel which is the torsion subsheaf of $\Omega_{C/\mathbb{K}}$ consisting of the direct sum of a skyscraper sheaf of dimension 1 at each node.

To prove this, one uses flatness of the completion operation to reduce the computation to the local version done earlier. While it is *not* true that the stalk of $\Omega_{C/\mathbb{K}}$ at a node is always the same up to isomorphism – most curves do not have isomorphic open subsets – the theorem shows that in some sense the behavior of $\Omega_{C/\mathbb{K}}$ at the node is always the same.

11.6.4 Exercises

Exercise 11.6.9. Suppose that C is a smooth complete intersection curve in $\mathbb{P}^3$ – that is, C is scheme-theoretically the intersection of two hypersurfaces H_1, H_2 of degrees d, e respectively. Compute the genus of C in terms of d, e.

Exercise 11.6.10. In this exercise we see another way to construct curves of arbitrary genus.

Consider the surface $\mathbb{P}^1 \times \mathbb{P}^1$ equipped with the two projections maps π_1, π_2.

(1) Show that $\Omega_{\mathbb{P}^1 \times \mathbb{P}^1} \cong \pi_1^* \mathcal{O}(-2) \oplus \pi_2^* \mathcal{O}(-2)$.

(2) Choose positive integers a, b. Show that a general curve in the linear series $|\mathcal{O}(a, b)|$ is smooth.

(3) Using the adjunction formula, show that a general curve in the linear series $|\mathcal{O}(a, b)|$ has genus $(a-1)(b-1)$.

In particular this argument shows that $\mathbb{P}^3$ contains curves of every genus.

Exercise 11.6.11. Let C be the nodal cubic in $\mathbb{A}^2_{\mathbb{C}}$ defined by the equation $y^2 - x^3 - x^2$. Write down the cotangent sequence for the normalization map and explicitly verify Theorem 11.6.8. (How would this problem change in characteristic p?)

Exercise 11.6.12. Let C be the cuspidal cubic in $\mathbb{A}^2_{\mathbb{C}}$ defined by the equation $y^2 - x^3$. Write down the cotangent sequence for the normalization map and interpret each part of the diagram. (How would this problem change in characteristic p?)

Exercise 11.6.13. Let C be the hyperelliptic curve defined by taking the closure of $y^2 = f(x)$ as in Example 11.6.6. Let $g = \frac{\deg(f)}{2} - 1$ denote the genus of C. Let's analyze the space of global sections of $\Omega_{C/\mathbb{K}}$.

(1) Show that dx/y is a well-defined element of $\Omega_{C/\mathbb{K}}(C)$. (That is, show that by expressing the same rational section in a different way we can obtain well-defined expressions along each open set in an open cover of C.)

(2) Show that for every $0 \leq i < g$ the differential $\frac{x^i}{y} dx$ is a well-defined element of $\Omega_{C/\mathbb{K}}(C)$.

(3) Show that the differentials in part b are linearly independent so that the space of global sections satisfies $\dim_{\mathbb{K}} \Omega_{C/\mathbb{K}}(C) \geq g$. (It turns out that it is exactly equal to g.)

Exercise 11.6.14 (Dualizing sheaf)**.** Any nodal curve C carries an invertible sheaf ω_C called the dualizing sheaf which is closely related to the cotangent sheaf. In many situations the dualizing sheaf can be used in place of the cotangent sheaf.

Here is the construction over an algebraically closed field $\mathbb{K}$. We can construct ω_C locally, so we may assume that C has a single node p. Let $\nu : C^\nu \to C$ denote the normalization of C^ν and let $x, y \in C^\nu$ denote the preimages of the nodal point. Consider the exact sequence

$$0 \to \Omega^1_{C^\nu} \to \Omega^1_{C^\nu}(x+y) \to \mathbb{K}(x) \oplus \mathbb{K}(y) \to 0.$$

By pushing forward to C (which is exact since ν is finite), we get an exact sequence

$$0 \to \nu_*\Omega^1_{C^\nu} \to \nu_*\Omega^1_{C^\nu}(x+y) \xrightarrow{\psi} \mathbb{K}(p)^{\oplus 2} \to 0.$$

Then ω_C is the subsheaf of $\nu_*\Omega^1_{C^\nu}(x+y)$ which is the ψ-preimage of the kernel of the addition map $\mathbb{K}(p)^2 \to \mathbb{K}(p)$.

(1) Prove that ω_C is an invertible sheaf fitting into an exact sequence

$$0 \to \nu_*\Omega^1_{C^\nu} \to \omega_C \to \mathbb{K}(p) \to 0.$$

(2) Suppose C is the union of the coordinate axes in $\mathbb{A}^2$. Compute the module defining ω_C and show that the restriction of ω_C to each axis is the ideal sheaf of the origin.

(3) By composing the map in (1) above with the cotangent sequence map we obtain a morphism $\Omega^1_C \to \omega_C$. Identify the kernel and cokernel of this map.

11.7 Smooth morphisms

Suppose $f : X \to Y$ is a differentiable map of smooth manifolds. We say that f is a submersion if the induced maps on tangent bundles $T_X \to f^*T_Y$ is surjective. This class of morphisms has several special properties. For example, the regular value theorem shows that locally a submersion looks like a coordinate projection map. In particular, the fibers of a submersion are submanifolds of dimension $\dim(X) - \dim(Y)$.

11.7.1 Smooth morphisms

The algebro-geometric analogue of a submersion is known as a smooth morphism. The definition is the relative analogue of the "absolute" notion of a smooth $\mathbb{K}$-variety.

Definition 11.7.1. Suppose that $f : X \to Y$ is a morphism of schemes. We say that f is smooth of relative dimension n if:

(1) f is locally of finite presentation,

(2) f is flat of relative dimension n, and

(3) $\Omega_{X/Y}$ is locally free of rank n.

More generally, we say that f is smooth at a point $x \in X$ if it is smooth (of some relative dimension) on some open neighborhood of x, and we say that f is smooth if it is smooth (of some relative dimension) at every point.

Exercise 11.7.2. Prove that this definition of smoothness agrees with Definition 11.5.1 for a $\mathbb{K}$-scheme $X \to \mathrm{Spec}(\mathbb{K})$.

The flatness condition guarantees that the fibers of f "vary nicely". For example, this condition guarantees that X does not consist of a disjoint union of $\mathbb{A}^1$s with one lying over each point of Y. It is not immediately clear how the definition of a smooth morphism relates to the geometric notion of a submersion; it will take us some time to develop the connection.

Remark 11.7.3. Note that smoothness is compatible with open subsets: if $f : X \to Y$ is smooth and U is an open subscheme of X then $f|_U : U \to Y$ is also smooth. On the other hand, closed embeddings are almost never flat, hence almost never smooth.

Remark 11.7.4. Suppose that $f : X \to Y$ is a smooth morphism. Since $\Omega_{X/Y}$ pulls back under base change, we see that for every fiber X_y the cotangent sheaf $\Omega_{X_y/\kappa(y)}$ is locally free of rank $\dim(X_y)$. Thus X_y is a smooth $\kappa(y)$-scheme.

This is a *very* strong condition for a morphism f that will fail to be satisfied even in basic examples (see Example 11.7.13). We will soon show that (under the right hypotheses) the smoothness of fibers of f is equivalent to the smoothness of f.

There is one class of smooth maps that plays a special role in the theory:

Definition 11.7.5. A morphism $f : X \to Y$ is said to be étale if it is smooth of relative dimension 0. In particular we must have $\Omega_{X/Y} = 0$.

We will analyze étale morphisms systematically in a later section.

11.7.2 Alternative definitions

Our first task is to give several alternate characterizations of smoothness. One of these alternate characterizations will rely on the following definition.

Definition 11.7.6. Let $f : \operatorname{Spec}(R) \to \operatorname{Spec}(S)$ be a morphism of affine schemes. We say that f is standard smooth if it is finitely presented and if we use $f^\sharp$ to identify R as $S[x_1, \dots, x_n]/(g_1, \dots, g_r)$ we have that $r \leq n$ and that the Jacobian matrix

$$\operatorname{Jac}_{g_1,\dots,g_r}(x) = \begin{bmatrix} \frac{\partial g_1}{\partial x_1} & \cdots & \frac{\partial g_r}{\partial x_1} \\ \vdots & \ddots & \vdots \\ \frac{\partial g_1}{\partial x_n} & \cdots & \frac{\partial g_r}{\partial x_n} \end{bmatrix}.$$

has an $r \times r$ minor which maps to an invertible element in R.

Recall that the Jacobian criterion for smoothness of an affine $\mathbb{K}$-scheme requires the Jacobian to have full rank. A standard smooth morphism captures the notion of "relative smoothness" by again requiring the Jacobian to have full rank. One can think of this as an "algebraic" definition of smoothness (in contrast to the "geometric" Definition 11.7.1).

Lemma 11.7.7. *Let $R = \mathbb{K}[x_1, \dots, x_n]/(g_1, \dots, g_r)$ be a finitely generated $\mathbb{K}$-algebra and suppose that R is standard smooth over $\mathbb{K}$. Then every closed point of $\operatorname{Spec}(R)$ is regular and every component of $\operatorname{Spec}(R)$ has dimension $n - r$.*

Proof. Fix a closed point $x \in X$ and consider the image in $\mathbb{A}^n$. The local ring $\mathcal{O}_{\operatorname{Spec}(R),x}$ is cut out in the local ring $\mathcal{O}_{\mathbb{A}^n,x}$ by the functions $g_1, \dots, g_r$. Furthermore, the Jacobian condition guarantees that the images of the g_i in Zariski cotangent space of $\mathbb{A}^n$ at x are linearly independent. By Exercise 5.1.6 we see that the Zariski cotangent space of $\mathcal{O}_{\operatorname{Spec}(R),x}$ has dimension $n - r$. On the other hand, Krull's PIT shows that the dimension of $\mathcal{O}_{\operatorname{Spec}(R),x}$ is at least $n - r$. Applying Theorem 5.2.1 we see that $\mathcal{O}_{\operatorname{Spec}(R),x}$ is a regular local ring. This computation also shows that $\dim(\operatorname{Spec}(R)) = n - r$. □

Remark 11.7.8. Since the Auslander-Buchsbaum theorem implies that the localization of a regular ring is still regular, this implies that all the points of a standard smooth $\mathbb{K}$-scheme $\operatorname{Spec}(R)$ are regular. We will give a different proof (not relying on the Auslander-Buchsbaum theorem) in Proposition 11.9.10.

The following theorem is our main result in this section.

Theorem 11.7.9. *Let $f : X \to Y$ be a morphism of schemes. Fix a point $x \in X$ and set $y = f(x)$. The following are equivalent:*

(1) f is smooth at x.

(2) f is flat and locally presented on a neighborhood of x and the fiber X_y is smooth at x.

(3) There is an open affine neighborhood U of x and an open affine neighborhood V of y such that $f(U) \subset V$ and $f|_U : U \to V$ is standard smooth.

Proof. (1) $\implies$ (2): By definition if f is smooth at x then it is flat at x. This implies that $\mathrm{locdim}_x(X_y)$ is the difference in the local codimensions of x in X and y in Y. Furthermore, the restriction of $\Omega_{X/Y}$ to X_y is equal to $\Omega_{X_y/y}$. Since pullbacks preserve local freeness we see that $\Omega_{X_y/y}$ is locally free of rank $\dim(X_y)$ at x.

(2) $\implies$ (3): Since the question is local, we may suppose that $X = \mathrm{Spec}(R)$ and $Y = \mathrm{Spec}(S)$ where R is a finitely presented S-algebra. Thus

$$R \cong S[x_1, \dots, x_n]/(g_1, \dots, g_r)$$

We also let $\mathfrak{q} \subset S$ denote the prime ideal defining y. Note that locally near x the fiber X_y is defined by $S/\mathfrak{q}[x_1, \dots, x_n]/(\overline{g}_1, \dots, \overline{g}_r)$. Let d denote the dimension of X_y. Since X_y is smooth at x, the Jacobian matrix for $g_1, \dots, g_r$ in $S/\mathfrak{q}[x_1, \dots, x_n]$ has rank $n - d$ at x. In particular this implies that $n - d \leq r$. After relabeling we may suppose that the first $(n - d) \times (n - d)$-minor of this Jacobian is non-vanishing.

Consider the morphisms

$$\begin{tikzcd} \mathrm{Spec}(S[x_1, \dots, x_n]/(g_1, \dots, g_r)) \arrow[rr, "h"] \arrow[dr, "f"'] & & \mathrm{Spec}(S[x_1, \dots, x_n]/(g_1, \dots, g_{n-d})) \arrow[dl, "f'"] \\ & \mathrm{Spec}(S) & \end{tikzcd}$$

Note that h is a closed embedding. By construction the local dimension of the fibers of f' and f at x are the same. Furthermore both fibers over y are smooth at x by construction. Thus by Exercise 11.5.14 there is an open neighborhood of x in the fibers on which h induces an isomorphism.

We claim that this isomorphism spreads out to an open neighborhood of x in X. It suffices to show that h induces an isomorphism on stalks of the structure sheaf at x. If we let $\mathfrak{p}$ denote the prime ideal in $S[x_1, \dots, x_n]$ defining the image of x, we must show that h induces an isomorphism

$$S[x_1, \dots, x_n]_{\mathfrak{p}}/(g_1, \dots, g_{n-d}) \to S[x_1, \dots, x_n]_{\mathfrak{p}}/(g_1, \dots, g_r)$$

Considering this as a map of $S_{\mathfrak{q}}$-modules, let K denote the kernel

$$K = \frac{(g_1, \ldots, g_{n-d}) \cdot S[x_1, \ldots, x_n]_{\mathfrak{p}}}{(g_1, \ldots, g_r) \cdot S[x_1, \ldots, x_n]_{\mathfrak{p}}}.$$

In this exact sequence of $S_{\mathfrak{q}}$-modules, the rightmost entry $S[x_1, \ldots, x_n]_{\mathfrak{p}}/(g_1, \ldots, g_r)$ is a flat $S_{\mathfrak{q}}$-module by assumption. Thus the tensor product of this sequence by $S_{\mathfrak{q}}/\mathfrak{q} \cong \kappa(y)$ is still exact. But as observed above we have an isomorphism of fibers over y, showing that $K \otimes \kappa(y) = 0$. By Nakayama's Lemma we see that $K = 0$ as well.

Altogether we see that $r = n - d$. Since by construction the Jacobian matrix for $g_1, \ldots, g_{n-d}$ has rank $n - d$ at x, it will have full rank on an open neighborhood of x and we conclude that X is smooth on a neighborhood of x.

(3) $\implies$ (1): Since all three conditions in Definition 11.7.1 are local on the source and on the target, it suffices to consider the case when $f : \mathrm{Spec}(A) \to \mathrm{Spec}(B)$ is a standard smooth morphism. We will write $A = B[x_1, \ldots, x_n]/(g_1, \ldots, g_r)$. It is clear that f is finitely presented. To prove that f is flat, we will appeal to the following algebraic result:

Theorem 11.7.10 (Slicing criterion for flatness)**.** *Suppose that A is a B-algebra, M is a finitely generated A-module, and $f \in A$ has the property that for every maximal ideal $\mathfrak{n} \subset B$ multiplication by f is injective on $M/\mathfrak{n}M$. If M is B-flat, then M/fM is also B-flat.*

We will prove that by induction on i that the ring $B[x_1, \ldots, x_n]/(g_1, \ldots, g_i)$ is a flat B-algebra. For the base case, it is clear that $B[x_1, \ldots, x_n]$ is flat over B. For the induction step, by appealing to the slicing criterion it suffices to show that for every maximal ideal $\mathfrak{m} \subset B$ the restriction of g_i is not a zero divisor in the quotient $T_{i-1} := B/\mathfrak{m}[x_1, \ldots, x_n]/(g_1, \ldots, g_{i-1})$. Note that T_i defines a standard smooth scheme over the field $B/\mathfrak{m}$, and in particular Lemma 11.7.7 shows that $\mathrm{Spec}(T_{i-1})$ has dimension $n - i - 1$ and every closed point is regular. In particular, the localization of T_i at every maximal ideal is an integral domain; this implies that T_{i-1} must be a product of integral domains. Similarly $\mathrm{Spec}(T_i)$ has dimension $n - i$. Since intersecting with the vanishing locus of g_i must lower the dimension of every component of T_{i-1}, we see g_i cannot vanish identically on any component of $\mathrm{Spec}(T_{i-1})$ and thus cannot be a zero divisor.

Finally, we show that the sheaf of relative differentials of f is locally free. Recall that $\Omega_{A/B}$ can be computed as the cokernel of the Jacobian. By assumption there is some $r \times r$ submatrix of the Jacobian which is invertible over A. This implies that the map $A^{\oplus r} \to A^{\oplus n}$ defined by the Jacobian admits a splitting. Thus the cokernel is a projective A-module, hence locally free. □

Remark 11.7.11. The argument shows that a smooth morphism $f : X \to Y$ of relative dimension d can locally be expressed in the form $S[x_1, \ldots, x_{r+d}]/(g_1, \ldots, g_r)$ for some integer r.

In other words, locally f is the composition of a closed embedding $X \to Y \times \mathbb{A}^{r+d}$ and the projection map $Y \times \mathbb{A}^{r+d} \to Y$ and furthermore the ideal of X in $Y \times \mathbb{A}^{r+d}$ is defined by a regular sequence. One refers to this important property of f by saying that it is a "local complete intersection morphism".

11.7.3 Exercises

Exercise 11.7.12. Let $f : X \to Y$ be a morphism of smooth $\mathbb{K}$-varieties. Suppose that the map $T_{X/\mathbb{K}} \to f^*T_{Y/\mathbb{K}}$ is surjective. Prove that f is a smooth morphism.

(Hint: the most challenging step is to show that f is flat. One option is to show that the fibers of f are equidimensional and appeal to the Miracle Flatness Theorem. Another option is to use the slicing criterion for flatness for the induced maps of regular local rings as in the proof of Theorem 11.7.9.)

Exercise 11.7.13. Let $\mathbb{K}$ be an algebraically closed field. Fix a positive integer $d \geq 2$ and consider the moduli space M_d parametrizing degree d curves in $\mathbb{P}^2$.

(1) Show that the singular curves are parametrized by a hypersurface in M_d. (Hint: use the resultant.)

(2) Suppose that T be a $\mathbb{K}$-variety. Let $\mathcal{C} \subset T \times \mathbb{P}^2$ be a closed subscheme such that the map $f : \mathcal{C} \to T$ is flat over T and has fibers which are plane curves. Show that if f is smooth then $\mathcal{C}$ is the product of a smooth curve with T.

11.8 Properties of smooth morphisms

In this section we continue our study of smooth morphisms. In particular, we are now able to clarify the relationship between smoothness and submersions by identifying a list of analogous properties. We first show that smoothness is a well-behaved property of morphisms.

Proposition 11.8.1. *Smoothness is preserved under composition, stable under base change, and local on the target.*

Proof. Note that both flat and locally finitely presented are well-behaved properties, so it suffices to restrict our attention to the local freeness of the sheaf of differentials. Since $\Omega_{X/Y}$ is pulled back under base change we deduce that smooth morphisms are stable under base change. Similarly, since the local freeness of $\Omega_{X/Y}$ at a point is determined locally we see that smoothness is local on the target (and on the source).

It only remains to show that smoothness is preserved under composition. By Theorem 11.7.9 we may reduce to the situation where $f : \mathrm{Spec}(A) \to \mathrm{Spec}(B)$ and $g : \mathrm{Spec}(B) \to \mathrm{Spec}(C)$ are standard smooth morphisms of affine schemes. We can then write

$$A = B[x_1, \ldots, x_n]/(p_1, \ldots, p_r) = C[y_1, \ldots, y_m, x_1, \ldots, x_n]/(q_1, \ldots, q_s, p_1, \ldots, p_r).$$

Note that the corresponding Jacobian matrix has a block of zeros since each $\partial q_i/\partial x_j = 0$. In particular, we can choose an $(r + s) \times (r + s)$ submatrix whose determinant is the product of the invertible $r \times r$ matrix for A/B and the invertible $s \times s$ matrix for B/C. This concludes the proof. □

11.8.1 Local structure

We next prove a new "local structure" theorem for smooth morphisms. In a geometric setting, a submersion of manifolds has the local structure of a coordinate projection. This is not true for schemes in the Zariski topology, but the following result shows that it is true up to an étale map.

Proposition 11.8.2. *Let $f : X \to Y$ be a morphism. Then f is smooth of relative dimension d at a point $x \in X$ if and only if there is an open neighborhood U of x such that $f|_U$ factors as a composition*

$$U \xrightarrow{g} Y \times \mathbb{A}^d \xrightarrow{h} Y$$

where g is étale and h is the projection map.

Proof. The reverse implication follows from the fact that the composition of two smooth morphisms is smooth. To see the forward implication, we will use the fact that f is standard smooth in an open neighborhood U of x. As in Remark 11.7.11, locally near x the morphism is defined by a ring map $S \to S[x_1, \ldots, x_{r+d}]/(f_1, \ldots, f_r)$ where some $r \times r$

minor of the Jacobian matrix is invertible. After rearranging, we may suppose that it is the last $r \times r$ minor which is invertible. Now consider the factorization

$$S \to S[x_1, \ldots, x_d] \to S[x_1, \ldots, x_{r+d}]/(f_1, \ldots, f_r).$$

The first map defines the projection $\operatorname{Spec}(S) \times \mathbb{A}^d \to \operatorname{Spec}(S)$. The second map is still standard smooth and has relative dimension 0, thus it is étale. Finally, since open immersions are étale we can equally well replace the map $U \to \operatorname{Spec}(S) \times \mathbb{A}^d$ with the map $U \to Y \times \mathbb{A}^d$. □

11.8.2 Left exactness of the cotangent sequence

In a geometric setting, a submersion is defined as a map which induces a surjection of tangent spaces. The following statement proves that a smooth morphism induces an injection of cotangent sheaves.

Proposition 11.8.3. *Suppose that $f : X \to Y$ is a smooth morphism of Z schemes. Then the cotangent sequence is exact on the left:*

$$0 \to f^*\Omega_{Y/Z} \to \Omega_{X/Z} \to \Omega_{X/Y} \to 0.$$

The proof relies on a property of étale morphisms that we will prove in the next section.

Proof. Since $f : X \to Y$ is smooth, Proposition 11.8.2 shows that f locally factors as the composition of an étale morphism $g : X \to Y \times \mathbb{A}^d$ and the projection map. Since the injectivity of a map of sheaves can be checked locally, we may assume we are in this special situation. Consider the commuting diagram with exact rows:

$$\begin{tikzcd}
f^*\Omega_{Y/Z} \arrow[r, "\psi_1"] & \Omega_{X/Z} \arrow[r] & \Omega_{X/Y} \arrow[r] & 0 \\
f^*\Omega_{Y/Z} \arrow[u, "="] \arrow[r, "\psi_2"] & g^*\Omega_{Y\times\mathbb{A}^n/Z} \arrow[u, "\phi"] \arrow[r] & g^*\Omega_{Y\times\mathbb{A}^n/Y} \arrow[u] \arrow[r] & 0
\end{tikzcd}$$

To show that ψ_1 is injective, it suffices to show that ψ_2 and ϕ are injective. Exercise 11.3.12 shows that the bottom row is split exact so that ψ_2 is injective. Theorem 11.9.6 shows that ϕ is an isomorphism (and thus injective). We conclude that ψ_1 is injective. □

Remark 11.8.4. The converse of Proposition 11.8.3 is false: there are many examples of morphisms of smooth $\mathbb{K}$-varieties which are not smooth but do induce an injection of cotangent sheaves (see Proposition 11.10.2). The issue is that when we take the dual of

$$0 \to f^*\Omega_{Y/\mathbb{K}} \to \Omega_{X/\mathbb{K}} \to \Omega_{X/Y} \to 0$$

we only obtain a left-exact sequence

$$0 \to T_{X/Y} \to T_{X/\mathbb{K}} \to f^*T_{Y/\mathbb{K}}.$$

Thus having a surjection of tangent sheaves is a priori stronger than injectivity of the cotangent sequence.

Although having an injection of cotangent sheaves does not imply smoothness, under the right assumptions having a surjection of tangent sheaves does imply smoothness (see Exercise 11.7.12). This highlights again the analogy between submersions and smooth morphisms.

Remark 11.8.5. It turns out that if $f : X \to Y$ is a smooth morphism of Z-schemes then the dual to the cotangent sequence is also exact:

$$0 \to T_{X/Y} \to T_{X/Z} \to f^*T_{Y/Z} \to 0.$$

The easiest way to show this is to appeal to the $\mathcal{E}xt$ sheaves which will be defined in Definition 13.7.1. Since $\Omega_{X/Y}$ is locally free Exercise 13.7.2 shows that $\mathcal{E}xt^1(\mathcal{O}_X, \Omega_{X/Y}) = 0$. Then Proposition 13.7.5 proves the exactness of the sequence above.

11.8.3 Generic smoothness

Just as in the absolute case, we have a generic smoothness for morphisms. There are two versions of the statement, with slightly different hypotheses and conclusions – be sure to keep them straight!

Proposition 11.8.6. *Let $f : X \to Y$ be a dominant finitely presented morphism of integral schemes. Then $K(X)$ is separable over $K(Y)$ if and only if there is a dense open set $U \subset X$ such that $f|_U$ is smooth.*

Proof. Set $r = \dim(X) - \dim(Y)$. We first prove the forward implication. If $K(X)$ is separable over $K(Y)$ then there are elements $t_1, \ldots, t_r$ in $K(X)$ which are algebraically independent over $K(Y)$ such that $K(X)$ is a separable algebraic extension over $K(Y)(t_1, \ldots, t_r)$. Thus $\Omega_{K(X)/K(Y)}$ is a free module of rank r. By upper semicontinuity, $\Omega_{X/Y}$ has rank r over a dense open subset of X. Since X is reduced, this implies that $\Omega_{X/Y}$ is locally free of rank r over a dense open subset. By generic flatness f is flat of relative dimension r over a dense open subset of Y. Taking the intersection of these two open subsets proves the result.

We next prove the reverse implication. Suppose that there is an open set $U \subset X$ such that $f|_U$ is smooth. Since smoothness is preserved by base change we see that the fiber over the generic point $K(Y)$ is smooth. Since smoothness is preserved by localization, we see that $K(X)$ is smooth over $K(Y)$. Since $K(X)$ is finitely generated over $K(Y)$, by combining Corollary 11.2.3 with the cotangent sequence we conclude that $K(X)$ is separable over $K(Y)$. $\square$

To set up our next version of generic smoothness, we will need the following lemma.

Lemma 11.8.7. *Let $\mathbb{K}$ be a field of characteristic 0. Let $f : X \to Y$ be a morphism of $\mathbb{K}$-schemes and define*

$$X_r = \{x \in X \,|\, \mathrm{rk}(T_{X/\mathbb{K}} \to f^*T_{Y/\mathbb{K}})|_x \leq r\}.$$

Then $\dim(\overline{f(X_r)}) \leq r$.

Proof. We replace Y by an irreducible component of the closure of $f(X_r)$ and we replace X by an irreducible component of X_r that dominates Y, giving both of them their reduced structures. Note the rank of the image of the map $T_{X/\mathbb{K}} \to f^*T_{Y/\mathbb{K}}$ can only drop when we replace X by a closed subscheme since a closed embedding induces an injection of tangent space. Thus that after these changes we have $X = X_r$.

Now we have reduced to case where $f : X \to Y$ is a dominant morphism of $\mathbb{K}$-varieties and $X = X_r$. We want to show that $\dim(Y) \leq r$. Since Y is a variety, by generic smoothness we may shrink Y (and replace X by the preimage of this open set) to ensure that Y is smooth. Applying Proposition 11.8.6 we may also replace X by an open subset so that X is smooth and $f|_X$ is smooth (without changing the dimension of the image). By Remark 11.8.5 we have an exact sequence

$$0 \to T_{X/Y} \to T_{X/\mathbb{K}} \to f^*T_{Y/\mathbb{K}} \to 0.$$

Since we are assuming that the rightmost map has rank $\leq r$ at every point of X, we deduce that $\mathrm{rk}(f^*T_{Y/\mathbb{K}}) \leq r$. In particular, the rank of $\Omega_{Y/\mathbb{K}}$ at the generic point of Y is also at most r. In turn this implies that $\dim(Y) \leq r$. □

The previous lemma quickly yields:

Proposition 11.8.8. *Let $\mathbb{K}$ be a field of characteristic 0. Let $f : X \to Y$ be a dominant morphism of $\mathbb{K}$-varieties such that X is smooth. Then there is a dense open subset $V \subset Y$ such that f is smooth over V.*

Proof. Define X_r as in Lemma 11.8.7. By generic smoothness, we can shrink Y to assume that it is smooth. This implies that $T_{Y/\mathbb{K}}$ has rank $\dim(Y)$. Lemma 11.8.7 shows that $f(X_{\dim(Y)-1})$ is contained in a proper closed subset of Y; after removing this closed subset we may ensure that the map

$$(T_{X/\mathbb{K}} \to f^*T_{Y/\mathbb{K}})|_x$$

is surjective for every point $x \in X$. Since Y is reduced we see that $T_{X/\mathbb{K}} \to f^*T_{Y/\mathbb{K}}$ is surjective. But this implies that f is smooth by Exercise 11.7.12 □

Warning 11.8.9. The analogue of Proposition 11.8.8 does not hold in characteristic p even when $K(X)$ is a separable extension of $K(Y)$; see Exercise 11.8.10. The issue is that the map f can have singularities along subvarieties of X whose function fields are not separable over $K(Y)$. We can recover some kinds of generic smoothness if we require that every residue field extension induced by f is separable, but this condition is extremely rare.

11.8.4 Exercises

Exercise 11.8.10. In this exercise we show that the analogue of Proposition 11.8.8 does not hold in characteristic p. Let $\mathbb{K}$ be an algebraically closed field of characteristic $p > 2$, let $Y = \operatorname{Spec}(\mathbb{K}[t])$ and let $X = \operatorname{Spec}(\mathbb{K}[t, x, y]/(y^2 - x^p + t))$.

(1) Show that every fiber of the map $f : X \to Y$ over a point $t \neq 0$ is a non-regular integral curve.

(2) Show that the generic fiber of the map $f : X \to Y$ is an integral curve that is regular but not smooth.

(3) Show that $K(X)$ is separable over $K(Y)$.

In particular, regularity is not an open property for families of $\mathbb{K}$-schemes.

11.9 Étale maps

Recall that an étale morphism is a smooth morphism of relative dimension 0. Étale morphisms are analogous to the notion of a "local isomorphism" in a geometric setting. Thus a finite étale morphism is the algebraic geometer's version of a covering map. For example, just as one the topological fundamental group classifies covering spaces, one can define an "étale fundamental group" which classifies finite étale morphisms.

Étale morphisms play a central role in the theory of étale cohomology. The goal of this theory is to transport results from the setting of singular cohomology into algebraic geometry. Unfortunately the Zariski topology is unsuited to the task – it simply is not sensitive enough to capture the information we want. Grothendieck realized that by systematically using étale maps in place of open subsets, one could construct a "more sensitive" cohomology theory that does match up well with singular cohomology computations.

Proposition 11.9.1. *Let $f : X \to Y$ be a morphism of schemes. Then f is étale if and only if it is flat, locally finitely presented, and for every point $y \in Y$ the fiber over y is a finite disjoint union of $\mathrm{Spec}(\mathbb{L}_i)$ where each $\mathbb{L}_i$ is a finite separable extension of $\kappa(y)$.*

Proof. The forward implication follows from Proposition 11.2.11 and the compatibility of cotangent sheaves with base change. The reverse implication follows from Theorem 11.7.9. □

11.9.1 Local isomorphisms

One of the main themes of this section is that étale morphisms $f : X \to Y$ are analogous to the geometric notion of a local isomorphism. Note that this is not literally true; there need not be any open set in X which is isomorphic to an open set in Y. (However, it is true that an étale morphism of smooth complex varieties will be a local isomorphism in the Euclidean topology.)

Exercise 11.9.2. Find an étale morphism $f : X \to Y$ such that no open set in X is isomorphic to any open set in Y. (Hint: let $\mathbb{K}$ be algebraically closed and choose a smooth plane cubic $E \subset \mathbb{P}^2$. Define a dominant finite morphism $E \to \mathbb{P}^1$ by projecting away from a general point. Show that there is an open subset $Y \subset \mathbb{P}^1$ such that the restriction of f to the preimage of Y is étale. Prove the last claim by appealing to the Picard group.)

As we have seen before, one way to translate topological properties into algebraic geometry is to use the diagonal. This perspective works for étale morphisms as well.

Lemma 11.9.3. *Let $f : X \to Y$ be an étale morphism. Then the diagonal $\Delta_{X/Y} : X \to X \times_Y X$ is an open embedding.*

This diagonal property is shared by local isomorphisms. (However, it does not characterize étale maps; see Exercise 11.9.11.)

Proof. Take open affines $U \subset X$, $V \subset Y$ satisfying $f(U) \subset V$. Since $\Omega_{X/Y} = 0$, Exercise 11.4.14 shows that the ideal $\mathcal{I}$ that defines the diagonal inside of $U \times_V U$ satisfies $\mathcal{I}/\mathcal{I}^2 = 0$. Since $\mathcal{I}$ is finitely generated Nakayama's lemma shows that $\mathcal{I}$ is a principal ideal generated by an idempotent element. Thus on this open subset $\Delta_{X/Y}$ defines the inclusion of some connected components. Since open embeddings are local on the source and the target we conclude that $\Delta_{X/Y}$ is an open embedding everywhere. □

One of the geometric consequences of the fact that a covering map is a local isomorphism is the lifting property: any path on the base admits a unique lift (once we fix the starting point). In algebraic geometry schemes are much more "rigid" and thus finding lifts is more difficult. We should think of the following theorem as a loose analogue of the lifting property.

Theorem 11.9.4. *Let $f : X \to Y$ be a morphism. Suppose that $g : Y \to X$ is a section (that is, $f \circ g$ is the identity map). If f is étale, then g is an open embedding. If f is étale and separated then g is an isomorphism from Y onto a union of some connected components of X.*

Proof. By Lemma 11.9.3 the diagonal $\Delta_{X/Y}$ is an open embedding. Applying cancellation (Proposition 8.6.6) to the diagram

$$\begin{tikzcd} Y \arrow[rr, "g"] \arrow[dr, "id"'] & & X \arrow[dl, "f"] \\ & Y & \end{tikzcd}$$

we see that g is an open embedding as well.

To see the last statement, recall that if f is separated then the diagonal is a closed embedding. Repeating the argument we see that g is an open and closed embedding, thus an isomorphism onto a union of connected components of X. □

Using this result, we can show that a finite étale morphism $f : X \to Y$ is "étale locally" a local isomorphism. (This result is based on the same intuition as Proposition 11.8.2; while the literal analogue of a geometric result need not be true in algebraic geometry, it is often true up to passing to an étale cover.)

Corollary 11.9.5. *Let $f : X \to Y$ be a finite étale morphism of degree d. For any point $y \in Y$, there is an étale map $g : V \to Y$ whose image contains y such that $X \times_Y V$ is the disjoint union of d copies of V.*

Proof. Let $x_1, \ldots, x_d$ be the points in the preimage of y. We construct a sequence of étale base changes

$$V = V_d \to V_{d-1} \to \ldots \to V_1 \to V_0 = Y$$

and points $v_i \in V_i$ such that:

(1) $v_0 = y$ and the image of v_{i+1} in V_i is v_i,

(2) each map $V_{i+1} \to V_i$ is étale, and

(3) for each i and for every $j \leq i$ there is a section $g_j : V_i \to X \times_Y V_i$ whose image contains the preimage of x_j lying above v_i.

We construct the V_i inductively. For the $(i+1)$st step, choose any open neighborhood U_{i+1} of the preimage of x_{i+1} over $v_i \in V_i$. Set $V_{i+1} = U_{i+1}$ and $v_{i+1} = x_{i+1}$. We let $V_{i+1} \to V_i$ be the restriction of the projection map $X \times_Y V_i \to V_i$ to U_{i+1}; in particular, since the projection map is étale (being the base change of an étale map) the map $V_{i+1} \to V_i$ is also étale. By combining the identity map $V_{i+1} \to V_{i+1}$ and the inclusion $V_{i+1} \to X \times_Y V_i$ we obtain a section g of the projection map $X \times_Y V_{i+1} \to V_{i+1}$. Furthermore, the existence of the sections $g_j : V_i \to X \times_Y V_i$ for $j \leq i$ is preserved by base changing to V_{i+1}.

Putting all the steps together, we see that $X \times_Y V_d \to V_d$ admits sections through every point in the fiber above v_d. By Theorem 11.9.4 we obtain the desired statement. □

11.9.2 Cotangent sheaves

Our next property shows another way in which étale morphisms are similar to local isomorphisms.

Theorem 11.9.6. *Suppose that $f : X \to Y$ is an étale morphism of S-schemes. Then $f^*\Omega_{Y/S} \cong \Omega_{X/S}$.*

We will prove this result by interpreting the cotangent sheaf via the diagonal as in Exercise 11.4.14.

Proof. Consider the diagram

$$\begin{tikzcd} X \arrow[r, "\Delta_{X/S}"] \arrow[d, "f"] & X \times_S X \arrow[d, "(f,f)"] \\ Y \arrow[r, "\Delta_{Y/S}"] & Y \times_S Y \end{tikzcd}$$

We let $\mathcal{I}$ denote the ideal sheaf of $\Delta_{Y/S}$ and $\mathcal{J}$ denote the ideal sheaf of $\Delta_{X/S}$. (More accurately, $\mathcal{I}$ and $\mathcal{J}$ are ideal sheaves on suitably chosen open subsets of $Y \times_S Y$ and $X \times_S X$.) Since (f, f) is flat, the pullback $(f, f)^*\mathcal{I}$ is the ideal sheaf of the base change

$$Y \times_{Y \times_S Y} (X \times_S X) \cong X \times_Y X$$

inside of an open subset of $X \times_S X$. Since f is étale, Lemma 11.9.3 shows that $\Delta_{X/Y} : X \to X \times_Y X$ is an open embedding. Furthermore, the composition of $\Delta_{X/Y}$ with $X \times_Y X \to$

$X \times_S X$ is the diagonal $\Delta_{X/S}$. Thus the ideal sheaves $(f,f)^*\mathcal{I}$ and $\mathcal{J}$ agree when pulled back under $\Delta_{X/S}$. We conclude that

$$\begin{aligned} f^*\Omega_{Y/S} = f^*\Delta^*_{Y/S}\mathcal{I} &= \Delta^*_{X/S}(f,f)^*\mathcal{I} \\ &= \Delta^*_{X/S}\mathcal{J} = \Omega_{X/S}. \end{aligned}$$

□

This completes the last step in the proof of Proposition 11.8.3, showing that the cotangent sequence is left exact under a smoothness assumption.

11.9.3 Permanence properties

We have already seen (essentially by definition) that Krull dimension is preserved by étale maps. It turns out that there are many other local properties which are left unchanged by étale morphisms.

Theorem 11.9.7. *Let P be one of the following properties: reduced, regular, normal, Cohen-Macaulay.*

Suppose that $f : X \to Y$ is an étale morphism. Let $x \in X$ and let $y = f(x)$. Then X satisfies property P at x if and only if Y satisfies property P at y.

Note that if we have a finitely presented local homomorphism $\phi : A \to B$ of local rings, then ϕ defines an étale morphism if and only if B is a flat A-algebra, $\phi(\mathfrak{m}_A)B = \mathfrak{m}_B$, and the residue field $B/\mathfrak{m}_B$ is a finite separable extension of the residue field $A/\mathfrak{m}_A$. Using this condition the theorem essentially boils down to algebraic properties of local rings.

Remark 11.9.8. This is our first glimpse of a more general theory of "descent": the study of which properties of a scheme can be passed back and forth over a faithfully flat quasicompact morphism.

We will only prove one special case of Theorem 11.9.7.

Lemma 11.9.9. *Let $f : X \to Y$ be an étale morphism. Let $x \in X$ and let $y = f(x)$. Suppose that y is a regular point of Y such that $\mathcal{O}_{Y,y}$ has dimension n. Then x is also regular and $\mathcal{O}_{X,x}$ has dimension n.*

Proof. We prove the statement by induction on n. The base case is when $n = 0$, in which case the statement follows immediately from Proposition 11.2.11.

For the induction step, we cut down by hyperplanes. Choose any function $f \in \mathfrak{m}_y \backslash \mathfrak{m}_y^2$. By Lemma 5.2.4 the quotient $\mathcal{O}_{Y,y}/(f)$ is a regular local ring of dimension $n-1$. Since étale maps are preserved by base change, $\mathcal{O}_{X,x}/(f)$ is still étale over $\mathcal{O}_{Y,y}/(f)$. The induction assumption implies that $\mathcal{O}_{X,x}/(f)$ is étale of dimension $n-1$. We then apply Lemma 5.2.4 again to see that $\mathcal{O}_{X,x}$ is regular of dimension n. □

By combining the permanence properties of étale morphisms with the local structure theorem for smooth morphisms (Proposition 11.8.2), one obtains permanence properties for smooth morphisms. For example:

Proposition 11.9.10. *Let $f : X \to Y$ be a morphism which is smooth at some point $x \in X$. If $f(y)$ is regular, then x is also regular.*

The converse statement is also true, but we will not prove it.

Proof. By Proposition 11.8.2 there is an open neighborhood U of x and a factoring $U \to Y \times \mathbb{A}^d \to Y$. By Proposition 5.2.8 every point of $Y \times \mathbb{A}^d$ lying over y is regular. By Lemma 11.9.9 we deduce that every point in U lying over y is regular. □

In particular this implies that a smooth point of a $\mathbb{K}$-scheme is regular, finishing the proof of Theorem 11.5.5. A similar argument works for normality, reducedness, etc.

11.9.4 Exercises

Exercise 11.9.11. We say that a morphism $f : X \to Y$ is unramified if it is locally of finite type and $\Omega_{X/Y} = 0$. In other words, we drop the "flatness" condition for étale morphisms and weaken "finitely presented" to "finite type." (The main reason we weaken the definition is to ensure that all closed embeddings are unramified.)

The relationship between unramified morphisms and étale morphisms is similar to the relationship between finite morphisms and finite flat morphisms.

(1) Show that every closed embedding is unramified. Thus unramified morphisms form a larger class than étale morphisms.

(2) Show that f is an unramified map if and only if f is locally of finite type and the diagonal map $\Delta_{X/Y}$ is an open embedding. (Hint: it suffices to argue locally, so we may assume $X = \mathrm{Spec}(R)$. Show that if $I \subset R$ is a finitely generated ideal satisfying $I = I^2$ then $V(I) = D_e$ for some idempotent element e. You will need to use the finite type assumption to show that the ideal sheaf of the diagonal is locally finitely generated.)

(3) Conclude that every monomorphism which is locally of finite type is unramified.

Exercise 11.9.12. Show that the normalization of a nodal curve satisfies $\Omega_{X/Y} = 0$ but is not étale. (It is an example of an unramified morphism.)

Exercise 11.9.13. Suppose that $f : X \to Y$ is a morphism which is smooth at a point $x \in X$. Prove that if $f(x)$ is a normal point of Y then x is also normal.

Exercise 11.9.14. Remark 11.7.11 shows that any étale morphism can locally be described via ring maps $S \to S[x_1, \ldots, x_r]/(g_1, \ldots, g_r)$. It turns out that one can do a little bit better.

Definition 11.9.15. Let S be a ring and suppose that $R \cong S[x]/g$ where $g \in S[x]$ is a monic polynomial. Let $b \in S[x]$ be an element such that image of g' in the localization R_b is a unit. Then the corresponding map $\mathrm{Spec}(R_b) \to \mathrm{Spec}(S)$ is said to be a standard étale morphism.

In other words, a standard étale morphism is defined by a ring extension of the form

$$S[x,t]/(g, tb-1)$$

where g, b satisfy the conditions above. It turns out that every étale morphism locally has the form of a standard étale morphism.

Theorem 11.9.16 ([Sta15] Tag 02GU)**.** *Let $f : X \to Y$ be a morphism of schemes. Then f is étale if and only if there for any point $x \in X$ there is an open affine neighborhood U of X and an open affine neighborhood V of $f(x)$ such that $f(U) \subset V$ and $f|_U : U \to V$ is standard étale.*

(1) Show that every standard étale morphism is étale.

(2) Use Theorem 11.9.16 to deduce the Primitive Element Theorem for separable extensions.

(3) Show that not every étale morphism can be written locally in the form $S[x]/g$ for a single monic polynomial $g \in S[x]$. (Hint: show that this defines a finite morphism. What is an example of an étale morphism that is not "locally finite"?)

11.10 Ramified covers

As discussed earlier, finite étale morphisms are analogous to covering maps and finite flat morphisms are analogous to ramified covers. In this section we will study how cotangent sheaves behave for finite flat morphisms of curves. By "curve" we will mean a (not necessarily geometrically integral) $\mathbb{K}$-variety of dimension 1.

Suppose that $f : C \to Z$ is a finite morphism of smooth projective curves over a field $\mathbb{K}$. The miracle flatness theorem shows that f is automatically flat as well. In this situation we obtain an exact sequence

$$f^*\Omega_{Z/\mathbb{K}} \xrightarrow{\phi} \Omega_{C/\mathbb{K}} \to \Omega_{C/Z} \to 0.$$

In general this sequence need not be exact on the left. For example, if f is the Frobenius morphism then the leftmost map is the zero map and the rightmost map is an isomorphism. We will mainly be interested in the situation where we actually do have an exact sequence.

Definition 11.10.1. Let $f : X \to Y$ be a finite morphism of two $\mathbb{K}$-varieties. We say that f is separable if f is dominant and the field extension $K(X)/K(Y)$ is separable.

Proposition 11.10.2. *Let $f : C \to Z$ be a finite separable morphism of smooth projective curves over a field $\mathbb{K}$. Then $\Omega_{C/Z}$ is a torsion sheaf on C and the cotangent sequence is left exact:*

$$0 \to f^*\Omega_{Z/\mathbb{K}} \xrightarrow{\phi} \Omega_{C/\mathbb{K}} \to \Omega_{C/Z} \to 0.$$

Proof. Note that $f^*\Omega_{Z/\mathbb{K}}$ and $\Omega_{C/\mathbb{K}}$ are both invertible sheaves on C. We first show that ϕ is injective at the generic point. Since the stalk of $\Omega_{C/Z}$ at the generic point is $\Omega_{K(C)/K(Z)}$, the separability assumption implies that the stalk is zero. On the other hand the stalks of $f^*\Omega_{Z/\mathbb{K}}$ and $\Omega_{C/\mathbb{K}}$ are both one-dimensional $K(C)$-vector spaces. Using middle exactness we see that ϕ induces an isomorphism at the generic point, proving the injectivity of ϕ at the generic point.

Spreading out, this means that ϕ is injective on an open subset of C. In particular the kernel of ϕ is a torsion sheaf. But since $f^*\Omega_{Z/\mathbb{K}}$ is torsion-free we see that the kernel of ϕ is zero and thus ϕ is injective.

Since the stalk of $\Omega_{C/Z}$ at the generic point is zero, it must be a torsion sheaf. □

Since $\Omega_{C/Z}$ is supported on a 0-dimensional subscheme, it corresponds to an effective Weil divisor R. The easiest way to construct R explicitly is as follows. After tensoring by $\Omega_{C/\mathbb{K}}^\vee$ we obtain an injection

$$f^*\Omega_{Z/\mathbb{K}} \otimes \Omega_{C/\mathbb{K}}^\vee \to \mathcal{O}_C.$$

We let R denote the closed subscheme of C defined by the image of this injection. Since any invertible sheaf on $\mathrm{Spec}(R)$ is isomorphic to the structure sheaf, we see that

$$\mathcal{O}_R \cong \mathcal{O}_R \otimes \Omega_{C/\mathbb{K}} \cong \Omega_{C/Z}.$$

Definition 11.10.3. Let $f : C \to Z$ be a finite separable morphism of smooth projective curves over a field $\mathbb{K}$. The ramification divisor R is the effective Weil divisor defined by $\Omega_{C/Z}$.

This leads us to our first version of the Riemann-Hurwitz fomula.

Theorem 11.10.4 (Riemann-Hurwitz formula)**.** *Let $f : C \to Z$ be a finite separable morphism of smooth projective geometrically integral curves over a field $\mathbb{K}$. Then we have*

$$\Omega_{C/\mathbb{K}} \cong f^*\Omega_{Z/\mathbb{K}}(R).$$

where R denotes the ramification divisor. Taking degrees, we can relate the genus of C and Z via the formula

$$2g(C) - 2 = \deg(f)(2g(Z) - 2) + \deg(R).$$

Proof. As discussed above, we know that $f^*\Omega_{Z/\mathbb{K}} \otimes \Omega^\vee_{C/\mathbb{K}}$ is isomorphic to the ideal sheaf $\mathcal{O}_X(-R)$ of R. We obtain the first statement by tensoring both sides by $\Omega_{C/\mathbb{K}}(R)$. Recall that $\deg(\mathcal{O}_C(R)) = \deg(R)$. Thus by taking degrees we get

$$\begin{aligned} 2g(C) - 2 &= \deg(f^*\Omega_{Z/\mathbb{K}} \otimes \Omega_C(R)) \\ &= \deg(f^*\Omega_{Z/\mathbb{K}}) + \deg(\Omega_C(R)) \\ &= \deg(f)(2g(Z) - 2) + \deg(R) \end{aligned}$$

where we applied Proposition 10.5.5 to obtain the last line. This rearranges to give the desired formula. □

Here are some important consequences of the Riemann-Hurwitz formula.

Example 11.10.5. Suppose that C and Z are smooth projective geometrically integral curves such that $g(Z) > g(C)$. Then there are no finite separable morphisms $f : C \to Z$. Indeed, in this situation we must have $\deg(f) \geq 2$ and $g(Z) \geq 1$ by Fact 11.6.4. But then the Riemann-Hurwitz formula implies that

$$g(C) - g(Z) = (g(Z) - 1)(\deg(f) - 1) + \frac{\deg(R)}{2} \geq 0.$$

Example 11.10.6. Assume for simplicity that $\mathbb{K}$ is a perfect field. Then $\mathbb{P}^1$ is "algebraically simply connected", in the sense that there are no non-trivial finite étale morphisms $f : C \to \mathbb{P}^1$ from a geometrically integral curve C.

Suppose for a contradiction that there were such a map. Since f is étale C is regular, and thus (due to our assumption on $\mathbb{K}$) a smooth curve. Arguing for each geometrically connected component separately, we may as well suppose that C is integral. Since f is étale it must be a separable morphism, so we may apply the Riemann-Hurwitz formula:

$$2g(C) - 2 = \deg(f) \cdot (-2) + \deg(R).$$

Since f is étale we have $\deg(R) = 0$. But this contradicts Fact 11.6.4 unless $\deg(f) = 1$. Since C is geometrically integral, we see that $C \cong \mathbb{P}^1$ and f is an isomorphism.

11.10.1 Computing the ramification divisor

We next discuss how to compute the ramification divisor explicitly. Our intuition will be guided by the behavior of Riemann surfaces. If $f : C \to Z$ is a degree d morphism of Riemann surfaces, then generically f will look like a d-sheeted covering. Over special points on Z (known as branch points) it is possible for multiple sheets of the covering to "coalesce" along a single point. If k sheets come together at a point $c \in C$, then in local coordinates near c the function f has the form $f(z) = uz^k$ for some local unit u. We can compute the behavior of the derivative using these local coordinates.

For algebraic curves the argument we can make a similar argument using local rings. However our computation will require an additional separability assumption for the point in question.

Construction 11.10.7. Let C be a smooth projective curve and let $c \in C$ be a closed point such that $\kappa(c)$ is separable over $\mathbb{K}$. Consider the discrete valuation ring $R := \mathcal{O}_{C,c}$ and let t be a uniformizer. In this situation we know that $\Omega_{R/\mathbb{K}}$ is a rank 1 free R-module. Since by Proposition 11.4.6 the map $\mathfrak{m}_c/\mathfrak{m}_c^2 \to \Omega_{R/\mathbb{K}}(c)$ sending $t \mapsto dt$ is injective, by applying Nakayama's lemma we see that $\Omega_{R/\mathbb{K}}$ is generated by dt.

Now suppose we have a finite separable morphism $f : C \to Z$ of smooth projective geometrically integral curves. Let $z = f(c)$ and let $S = \mathcal{O}_{Z,z}$ with uniformizer s. Since $\kappa(c)$ is separable over $\mathbb{K}$, so is $\kappa(z)$. Thus both $\Omega_{R/\mathbb{K}}$ and $\Omega_{S/\mathbb{K}}$ are free rank 1 modules generated by dt and ds respectively. We have an exact sequence

$$R \otimes_S \Omega_{S/\mathbb{K}} \xrightarrow{\phi} \Omega_{R/\mathbb{K}} \to (\Omega_{C/Z})_c \to 0$$

and since $\Omega_{C/Z}$ is torsion the stalk at c is the same as the fiber at c. The map $f^\sharp : S \to R$ has the form $s \mapsto ut^e$ for some unit $u \in R$ and some positive integer e. Thus ϕ identifies

$$ds \mapsto t^e\, du + eut^{e-1}\, dt = t^{e-1}(t\, du + eu\, dt).$$

Since dt generates $\Omega_{R/\mathbb{K}}$ we know that $du = w\, dt$ for some function w. Then we can write $ds \mapsto t^{e-1}(tw + eu)\, dt$. Thus the part of $\Omega_{C/Z}$ supported on c depends on the characteristic:

- If $\mathrm{char}(\mathbb{K})$ does not divide e, then $tw + eu$ is a unit. Thus, the part of $\Omega_{C/Z}$ supported at c is isomorphic to R/t^{e-1} and so has dimension $e - 1$ over the residue field $\kappa(c)$. In this situation we say that f is tamely ramified at c.
- If $\mathrm{char}(\mathbb{K})$ divides e, then $tw + eu = tw$ is not a unit. Thus the part of $\Omega_{C/Z}$ supported at c is isomorphic to R/t^r for some $r \geq e$ and so has dimension $\geq e$ over the residue field $\kappa(c)$. In this situation we say that f is wildly ramified at c.

The integer e defined above is called the ramification index of f at the point c. When f is tamely ramified, c will be contained in the support of the ramification locus if and only if its ramification index is ≥ 2.

Putting everything together, we obtain our second version of the Riemann-Hurwitz formula:

Theorem 11.10.8 (Riemann-Hurwitz formula)**.** *Let $f : C \to Z$ be a finite separable morphism of smooth projective geometrically integral curves over a field $\mathbb{K}$. Suppose that f is tamely ramified. Then we have*

$$2g(C) - 2 = \deg(f)(2g(Z) - 2) + \sum_{c \in C} (e_c - 1) \cdot [\kappa(c) : \mathbb{K}]$$

where e_c denotes the ramification index at c.

This theorem pairs naturally with the result of Exercise 10.5.25: if we fix a point $z \in Z$, then we have

$$\deg(f) = \sum_{c \in f^{-1}z} e_c \cdot [\kappa(c) : \kappa(z)].$$

Thus the degree of f imposes a combinatorial bound on the possible contributions of the fiber $f^{-1}z$ to the ramification divisor.

Remark 11.10.9. There is a similar analysis in higher dimensions for canonical bundles that is also known as the Riemann-Hurwitz formula. Suppose that $f : X \to Y$ is a finite separable morphism of smooth projective geometrically integral $\mathbb{K}$-varieties. Then we obtain an equality of canonical bundles

$$\omega_X \cong f^*\omega_Y(R)$$

where R is an effective Weil divisor known as the ramification divisor. Just as for curves, R can be computed explicitly by analyzing the behavior of f along the discrete valuation rings which are the generic points of prime divisors on X.

To prove this, the key is to show that the map $\phi : f^*\Omega_Y \to \Omega_X$ is injective. Since $f^*\Omega_Y$ is locally free the injectivity of this map can be checked on the generic fiber, and we reduce to a local computation as in Proposition 11.10.2. Once we know ϕ is injective, we also obtain injectivity of the map $f^*\omega_Y \to \omega_X$ obtained by taking top exterior products, and the rest of the argument is similar.

11.10.2 Exercises

Exercise 11.10.10. Let $\mathbb{K}$ be an algebraically closed field. Suppose that C is a smooth plane curve of degree d. Note that if we project away from a general point the rational map $\phi : \mathbb{P}^2 \dashrightarrow \mathbb{P}^1$ defines a morphism on C. Use the Riemann-Hurwitz formula to compute the genus of C (and compare against Example 11.6.5).

Exercise 11.10.11. Let $\mathbb{K}$ be a field of characteristic 0. Prove that there are no non-trivial finite étale morphisms $f : C \to \mathbb{A}^1$ where C is a geometrically integral curve.

Prove that if $\mathbb{K}$ has characteristic p then the curve defined in $\mathbb{K}[x, t]$ by the equation $x^p - x - t$ defines an étale cover of $\operatorname{Spec}(\mathbb{K}[t])$. Thus the analogous result does not hold in characteristic p even when $\mathbb{K}$ is perfect. (In fact, a result of Katz shows that in characteristic p every curve admits a morphism to $\mathbb{P}^1$ which is only ramified over a single point, yielding many étale covers of $\mathbb{A}^1$.)

Chapter 12

Čech cohomology

Let X be a topological space. We will associate to any sheaf $\mathcal{F}$ on X the Čech cohomology groups $\check{H}^i(X, \mathcal{F})$ for indices $i \geq 0$. The first couple groups can be understood explicitly: $\check{H}^0(X, \mathcal{F})$ is just the space of global sections $\mathcal{F}(X)$, and $\check{H}^1(X, \mathcal{F})$ detects whether a set of local sections of $\mathcal{F}$ which satisfy the "cocycle condition" can come from a global set of data. However, as i increases it becomes more challenging to identify the geometric content of $\check{H}^i$.

The key property of the Čech cohomology groups is that they control the failure of the global sections functor to be right exact. That is, if the topology on X is "nice enough" (e.g. Hausdorff paracompact) then a SES of sheaves

$$0 \to \mathcal{F} \xrightarrow{\phi} \mathcal{G} \xrightarrow{\psi} \mathcal{H} \to 0$$

leads to a long exact sequence of Čech cohomology groups

$$\begin{aligned} 0 \to \check{H}^0(X, \mathcal{F}) \to \check{H}^0(X, \mathcal{G}) \to \check{H}^0(X, \mathcal{H}) \to \\ \to \check{H}^1(X, \mathcal{F}) \to \check{H}^1(X, \mathcal{G}) \to \check{H}^1(X, \mathcal{H}) \to \\ \to \check{H}^2(X, \mathcal{F}) \to \check{H}^2(X, \mathcal{G}) \to \check{H}^2(X, \mathcal{H}) \to \dots \end{aligned}$$

Let's motivate the construction of $\check{H}^1(X, \mathcal{F})$ from this perspective. Since $\mathcal{G} \to \mathcal{H}$ is a surjective morphism of sheaves, we know that for any $t \in \mathcal{H}(X)$ there is an open cover $\{U_i\}$ such that $t|_{U_i}$ is in the image of $\psi(U_i)$. For each i, choose a lift s_i of $t|_i$. If we knew that the s_i agreed on overlaps, we could glue them to get a global section which maps to t.

Note that we can modify each s_i by an element of $\mathcal{F}(U_i)$ without changing the image in $\mathcal{H}(U_i)$. The key question is: can we choose local elements of $\mathcal{F}(U_i)$ to "fix" the discrepancies $s_i|_{U_i \cap U_j} - s_j|_{U_i \cap U_j}$? Although we have arrived at this question using an exact sequence, in the end we are really interested in a property intrinsic to $\mathcal{F}$: given an open cover $\{U_i\}$ and a collection of sections in the various $\mathcal{F}(U_i \cap U_j)$, is it possible to choose sections of the various $\mathcal{F}(U_i)$ whose differences yield our original set of data?

We can also ask the analogous question where we replace a pair of open sets by a collection of r open sets. Letting r vary, one constructs a complex of abelian groups (the Čech complex) recording the behavior of the restriction maps for the open sets in a fixed open cover. The Čech cohomology of $\mathcal{F}$ is defined to be the limit of the homology groups of this complex as we pass to further and further refinements of our open cover. The first section develops the basic theory of Čech complexes and Čech cohomology groups.

When X is a scheme, then it does not have the necessary nice topological properties to obtain a "clean" theory of Čech complexes. However, for *quasicoherent* sheaves on X, the behavior of restrictions for open affine subsets is defined via localization and is thus sufficiently nice to obtain a well-behaved theory. (Recall that surjectivity of a sheaf morphism can always be detected on the level of open affines for quasicoherent sheaves.) We develop the necessary tools for working with quasicoherent sheaves in the second section.

The remaining sections are dedicated to calculations. The first and most important example is coherent sheaves on $\mathbb{P}^n$. We can calculate the Čech cohomology groups of the sheaves $\mathcal{O}(d)$ "by hand"; using Hilbert's Syzygy Theorem we can derive interesting consequences for arbitrary coherent sheaves on $\mathbb{P}^n$. We obtain the best cohomological behavior when we allow a "twist" by a large invertible sheaf $\mathcal{O}(d)$. For arbitrary projective schemes, the analogous theory is developed by studying ample invertible sheaves and their cohomology. Finally, we focus on the special case of invertible sheaves and curves and prove some classical statements about the behavior of global sections.

12.0.1 Algebraic preliminaries

Let $\mathbf{C}$ be an abelian category. A complex of objects in $\mathbf{C}$ is a sequence

$$\dots \xrightarrow{d^{i-2}} A^{i-1} \xrightarrow{d^{i-1}} A^i \xrightarrow{d^i} A^{i+1} \xrightarrow{d^{i+2}} A^{i+2} \xrightarrow{d^{i+2}} \dots$$

such that the composition of any two consecutive maps $d_i \circ d_{i-1}$ is the zero map. We will denote such a complex using the notation $A^\bullet$; we will always denote the maps inside the complex using d^i. Note that our conventions are compatible with the "cohomological grading".

Definition 12.0.1. Let $A^\bullet$ be a complex of objects in the abelian category $\mathbf{C}$. The ith cohomology of $A^\bullet$ is

$$H^i(A^\bullet) := \frac{\ker(d^i)}{\mathrm{im}(d^{i-1})}.$$

Conceptually, cohomology groups measure the failure of the complex $A^\bullet$ to be an exact sequence.

Definition 12.0.2. A morphism of complexes $\phi^\bullet : A^\bullet \to B^{bullet}$ is a commuting diagram

$$\begin{tikzcd}
\dots \arrow[r, "d^{i-2}"] & A^{i-1} \arrow[r, "d^{i-1}"] \arrow[d, "\phi^{i-1}"] & A^i \arrow[r, "d^i"] \arrow[d, "\phi^i"] & A^{i+1} \arrow[r, "d^{i+1}"] \arrow[d, "\phi^{i+1}"] & A^{i+2} \arrow[r, "d^{i+2}"] \arrow[d, "\phi^{i+2}"] & \dots \\
\dots \arrow[r, "d^{i-2}"] & B^{i-1} \arrow[r, "d^{i-1}"] & B^i \arrow[r, "d^i"] & B^{i+1} \arrow[r, "d^{i+1}"] & B^{i+2} \arrow[r, "d^{i+2}"] & \dots
\end{tikzcd}$$

This yields a category $\mathbf{Com}(\mathbf{C})$ of complexes of objects in $\mathbf{C}$.

Given a morphism of complexes $\phi^\bullet$, we define $\ker(\phi^\bullet)$ by taking the complex whose ith piece is $\ker(\phi^i)$ (equipped with the natural maps between them induced by the d^i for $A^\bullet$). Similarly, we define $\mathrm{cok}(\phi^\bullet)$ to be the complex whose ith piece is $\mathrm{cok}(\phi^i)$. These are part of the data giving $\mathbf{Com}(\mathbf{C})$ the structure of an abelian category.

A morphism of complexes $\phi^\bullet : A^\bullet \to B^\bullet$ induces a collection of morphisms between cohomology groups $\phi_* : H^i(A^\bullet) \to H^i(B^\bullet)$.

Proposition 12.0.3. *Suppose we have an exact sequence of complexes*

$$0 \to A^\bullet \to B^\bullet \to C^\bullet \to 0.$$

Then we obtain a long exact sequence of cohomology groups

$$\begin{aligned}
\dots \to H^i(A^\bullet) \to H^i(B^\bullet) \to H^i(C^\bullet) &\xrightarrow{\delta^i} \\
\to H^{i+1}(A^\bullet) \to H^{i+1}(B^\bullet) \to H^{i+1}(C^\bullet) &\xrightarrow{\delta^{i+1}} \\
\to H^{i+2}(A^\bullet) \to \dots &
\end{aligned}$$

where the connecting morphisms δ^i are defined as in the Snake Lemma.

12.1 Čech cohomology for sheaves

As discussed in the introduction to the chapter, the purpose of Čech cohomology is to measure the failure of the global sections functor to be right exact.

12.1.1 Čech complex for an open cover

Notation 12.1.1. Let X be a topological space and let $\{U_i\}_{i\in I}$ be an open cover of X. In this section we will always assume that the index set I is equipped with a total ordering. We do allow $U_i = U_{i'}$ for different indices i, i'.

Given distinct indices $i_0, \ldots, i_p$ we denote by $U_{i_0\ldots i_p}$ the intersection $U_{i_0} \cap \ldots \cap U_{i_p}$.

Definition 12.1.2. Let X be a topological space and let $\mathcal{F}$ be a presheaf of abelian groups on X. Let $\mathfrak{U} = \{U_i\}_{i\in I}$ be an open cover of X. The Čech complex for $\mathcal{F}$ with respect to $\mathfrak{U}$ is the sequence

$$\check{C}^\bullet(\mathfrak{U}, \mathcal{F}) := \check{C}^0(\mathfrak{U}, \mathcal{F}) \xrightarrow{d^0} \check{C}^1(\mathfrak{U}, \mathcal{F}) \xrightarrow{d^1} \check{C}^2(\mathfrak{U}, \mathcal{F}) \xrightarrow{d^2} \ldots$$

where

$$\check{C}^0(\mathfrak{U}, \mathcal{F}) = \prod_{i_0} \mathcal{F}(U_{i_0})$$
$$\check{C}^1(\mathfrak{U}, \mathcal{F}) = \prod_{i_0<i_1} \mathcal{F}(U_{i_0 i_1})$$

and in general

$$\check{C}^p(\mathfrak{U}, \mathcal{F}) = \prod_{i_0<\ldots<i_p} \mathcal{F}(U_{i_0 i_1 \ldots i_p}).$$

Given an element $\sigma \in \check{C}^p(\mathfrak{U}, \mathcal{F})$ and a multi-index $i_0, \ldots, i_p$, we will denote by $\sigma_{i_0\ldots i_p}$ the image of σ under the $i_0, \ldots, i_p$ projection.

The differentials are defined by the following rule: for any $\sigma \in \check{C}^{p-1}(\mathfrak{U}, \mathcal{F})$, the component of $d^{p-1}\sigma$ corresponding to $i_0, \ldots, i_p$ is

$$(d^{p-1}\sigma)_{i_0\ldots i_p} := \sum_{j=0}^{p} (-1)^j (\sigma_{i_0\ldots\widehat{i_j}\ldots i_p})|_{U_{i_0\cdots i_p}}$$

where as usual the notation $i_0 \ldots \widehat{i_j} \ldots i_p$ means we omit the jth element from our indexing set.

We claim that $d^p \circ d^{p-1} = 0$. Indeed, since

$$(d^{p-1}\sigma)_{i_0 \dots i_p} := \sum_{j=0}^{p} (-1)^j (\sigma_{i_0 \dots \widehat{i_j} \dots i_p})|_{U_{i_0 \dots i_p}}$$

we see that that the sums defining $(d^p \circ d^{p-1}(\sigma))_{i_0 \dots i_{p+1}}$ will come in cancelling pairs: if we remove the indices i_j and i_k with $j < k$ then $U_{i_0 \dots \widehat{i_k} \dots i_{p+1}}$ contributes the term $(-1)^{j+k}\sigma_{i_0 \dots \widehat{i_j} \dots \widehat{i_k} \dots i_{p+1}}$ and $U_{i_0 \dots \widehat{i_j} \dots i_{p+1}}$ contributes the term $(-1)^{j+k-1}\sigma_{i_0 \dots \widehat{i_j} \dots \widehat{i_k} \dots i_{p+1}}$. We conclude that $\check{C}^\bullet(\mathfrak{U}, \mathcal{F})$ is a complex of abelian groups. An element $\sigma \in \check{C}^p(\mathfrak{U}, \mathcal{F})$ is said to be a cocycle if it lies in the kernel of d^p and a coboundary if it lies in the image of d^{p-1}. The subgroup of cocycles in denoted by Z^p and the subgroup of coboundaries is denoted by B^p.

Definition 12.1.3. The pth Čech cohomology of $\mathcal{F}$ with respect to $\mathfrak{U}$, denoted by $\check{H}^p(\mathfrak{U}, \mathcal{F})$, is the cohomology group of the sequence $\check{C}^\bullet(\mathfrak{U}, \mathcal{F})$:

$$\check{H}^p(\mathfrak{U}, \mathcal{F}) = Z^p / B^p.$$

Suppose that $\mathcal{F}$ is a sheaf. Given an element $\sigma = (\sigma_i \in \mathcal{F}(U_i))$ in $\check{C}^0(\mathfrak{U}, \mathcal{F})$, its image under d^0 is

$$(d^0\sigma)_{ij} = \sigma_j - \sigma_i.$$

By the gluing axiom for sheaves, the 0th cocycles for $\check{C}^\bullet(\mathfrak{U}, \mathcal{F})$ are exactly the same as the tuples (σ_i) which glue to yield a global section of $\mathcal{F}$. In other words,

$$\check{H}^0(\check{C}^\bullet(\mathfrak{U}, \mathcal{F})) = \mathcal{F}(X).$$

The next map d^1 has the form

$$(d^1\sigma)_{ijk} = \sigma_{jk} - \sigma_{ik} + \sigma_{ij}.$$

Thus the cocycles in $\check{C}^1(\mathfrak{U}, \mathcal{F})$ will be the set of data (σ_{ij}) which "satisfies the cocycle condition". The first homology group $\check{H}^1(\check{C}^\bullet(\mathfrak{U}, \mathcal{F}))$ measures whether such data can be obtained from by taking differences of sections (σ_i). We can continue to give explicit descriptions of the maps for higher values of p, but unfortunately the geometric meaning of these cohomology groups can be a bit obscure.

It is important to note that computing the sequence $\check{C}^\bullet(\mathfrak{U}, \mathcal{F})$ is not much harder than computing the values of $\mathcal{F}$ on open subsets of X. If we are able to accomplish the latter, then we can also accomplish the former using some basic algebra.

12.1.2 Čech cohomology groups

In general the cohomology groups $\check{H}^p(\mathfrak{U}, \mathcal{F})$ will depend heavily on the choice of open cover $\mathfrak{U}$. To obtain an invariant of the sheaf $\mathcal{F}$, we will need to think more carefully about how to choose $\mathfrak{U}$. Suppose we had an exact sequence of sheaves

$$0 \to \mathcal{F} \xrightarrow{\phi} \mathcal{G} \xrightarrow{\psi} \mathcal{H} \to 0.$$

As we have seen before, ψ need not induce a surjection on global sections. However, for any global section s of $\mathcal{H}$ we can choose a sufficiently small open cover $\{U_i\}$ such that $s|_{U_i}$ is in the image of $\psi(U_i)$. This is an indication that the Čech cohomology groups are the most sensitive to geometric information when we use a small open cover.

Definition 12.1.4. Suppose that $\mathfrak{U} = \{U_i\}_{i \in I}$ is an open cover of X (where as always the index set is equipped with an order). A refinement of $\mathfrak{U}$ is an open cover $\mathfrak{V} = \{V_j\}_{j \in J}$ and a monotonic function $\tau : J \to I$ such that $V_j \subset U_{\tau(j)}$ for every $j \in J$. Define a partial ordering on the set of open covers of X by saying that $\mathfrak{V} \leq \mathfrak{U}$ if there is a function τ making $\mathfrak{V}$ a refinement of $\mathfrak{U}$.

Note that if τ makes $\mathfrak{V}$ a refinement of $\mathfrak{U}$ then it induces a map

$$\Phi : \check{C}^\bullet(\mathfrak{U}, \mathcal{F}) \to \check{C}^\bullet(\mathfrak{V}, \mathcal{F})$$

defined by

$$\Phi(\sigma)_{j_0 \ldots j_p} = \sigma_{\tau(j_0) \ldots \tau(j_p)}|_{V_{j_0 \ldots j_p}}.$$

Since the differentials in the Čech complex are also defined by restriction we have $\Phi \circ d^p = d^p \circ \Phi$, showing that Φ induces a homomorphism of cohomology groups $\check{H}^p(\mathfrak{U}, \mathcal{F}) \to \check{H}^p(\mathfrak{V}, \mathcal{F})$.

One can show that this induced map of homology groups does not depend upon the choice of the function τ. In other words, to any refinement $\mathfrak{V} \leq \mathfrak{U}$ we have associated a well-defined map of Čech cohomology groups.

Definition 12.1.5. Let X be a topological space and let $\mathcal{F}$ be a presheaf of abelian groups on X. We define the Čech cohomology groups

$$\check{H}^p(X, \mathcal{F}) = \varinjlim \check{H}^p(\mathfrak{U}, \mathcal{F})$$

where the direct limit is taken over all open covers under the ordering $\leq$. (Technically speaking the collection of all open covers of X need not be indexed by a set; we will ignore this issue.)

Example 12.1.6. When $\mathcal{F}$ is a sheaf our earlier discussion shows that $\check{H}^0(X, \mathcal{F}) = \mathcal{F}(X)$. But when $\mathcal{F}$ is a presheaf then $\check{H}^0(X, \mathcal{F})$ need not be the global sections of $\mathcal{F}$.

12.1.3 Exactness properties

As with other cohomology theories, a key property of Čech cohomology is the existence of a LES in various situations. However, one needs to be a little bit careful: it is not true in full generality that a SES of sheaves induces a LES of cohomology. In fact this defect was one of the motivating factors behind the development of sheaf cohomology, an alternative theory we will develop in Section 13.2.

In this section we discuss a few results that do hold for general topological spaces; we will revisit the existence of a LES in the setting of quasicoherent sheaves in Theorem 12.2.4.

Theorem 12.1.7. *Let X be a topological space. Suppose we have an exact sequence of presheaves*

$$0 \to \mathcal{F} \to \mathcal{G} \to \mathcal{H} \to 0.$$

Then we obtain a long exact sequence of Čech cohomology groups

$$\ldots \to \check{H}^p(X, \mathcal{F}) \to \check{H}^p(X, \mathcal{G}) \to \check{H}^p(X, \mathcal{H}) \to \check{H}^{p+1}(X, \mathcal{F}) \to \ldots$$

When we discuss an "exact sequence of presheaves", we mean a sequence with the presheaf notions of image and cokernel.

Proof. For an exact sequence of presheaves $\mathcal{G}(U) \to \mathcal{H}(U)$ is surjective for every open set U. Thus for any open cover $\mathfrak{U}$ we get an exact sequence of complexes

$$0 \to \check{C}^\bullet(\mathfrak{U}, \mathcal{F}) \to \check{C}^\bullet(\mathfrak{U}, \mathcal{G}) \to \check{C}^\bullet(\mathfrak{U}, \mathcal{H}) \to 0.$$

By taking the usual LES of cohomology coming from an exact sequence of complexes, we see that for any open cover $\mathfrak{U}$ we get a LES of cohomology groups. A direct limit of exact sequences is still exact; taking a limit over all refinements, we obtain the desired statement. □

Corollary 12.1.8. *Let X be a topological space. Suppose we have an exact sequence of sheaves*

$$0 \to \mathcal{F} \to \mathcal{C}^0 \to \mathcal{C}^1 \to \mathcal{C}^2 \to \ldots$$

Suppose furthermore that for every $i > 0$ and for every $j \geq 0$ we have $\check{H}^i(X, \mathcal{C}^j) = 0$. Set $V^k := \check{H}^0(X, \mathcal{C}^k)$. Then $\check{H}^k(X, \mathcal{F})$ is isomorphic to the kth cohomology group of the complex $V^\bullet$.

When we have a complex $\check{C}^\bullet$ whose higher cohomologies vanish as above, it is called an acyclic resolution of $\mathcal{F}$.

Proof. Since the first map is injective, the presheaf image of $\mathcal{F}$ in $\mathcal{C}^0$ is isomorphic to $\mathcal{F}$. Set $\mathcal{Q}^0 := \mathcal{F}$. By taking presheaf images and presheaf cokernels, the LES splits into a sequence of short exact sequences

$$0 \to \mathcal{Q}^i \to \mathcal{C}^i \to \mathcal{Q}^{i+1} \to 0.$$

By Theorem 12.1.7 and the assumption on cohomology vanishing, we obtain

$$\check{H}^k(X,\mathcal{F}) \cong \check{H}^{k-1}(X,\mathcal{Q}^1) \cong \ldots \cong \check{H}^1(X,\mathcal{Q}^{k-1}) \cong \frac{\check{H}^0(X,\mathcal{Q}^k)}{\operatorname{im}\check{H}^0(X,\mathcal{C}^{k-1})}.$$

It only remains to note that the kernel of the map $\check{H}^0(X,\mathcal{C}^k) \to \check{H}^0(X,\mathcal{C}^{k+1})$ is the same as the kernel of the map $\check{H}^0(X,\mathcal{C}^k) \to \check{H}^0(X,\mathcal{Q}^{k+1})$, so that both can be identified with $\check{H}^0(X,\mathcal{Q}^k)$. □

12.1.4 Leray's Theorem

One of the main challenges in computing Čech cohomology is fact that we must take a direct limit over sufficiently refined open covers. The following important theorem allows us to identify a single open cover which computes the Čech cohomology.

Theorem 12.1.9 (Leray's Theorem)**.** *Let X be a topological space and let $\mathcal{F}$ denote a sheaf on X. Let $\mathfrak{U} = \{U_i\}$ be an open cover of X. Suppose that for every $i > 0$ and every non-empty open set V which is a finite intersection of open sets in $\mathfrak{U}$ we have $\check{H}^i(V,\mathcal{F}) = 0$. Then*

$$\check{H}^i(\mathfrak{U},\mathcal{F}) = \check{H}^i(X,\mathcal{F})$$

for all $i \geq 0$.

To apply this theorem, one must first identify criteria which guarantee that the higher cohomology of $\mathcal{F}$ vanishes along open sets of a certain type. The theorem then tells us that using these open sets is "good enough" to compute the Čech cohomology.

To prove Leray's Theorem, we will need a new object called the Čech sheaf complex. This is a "sheafified" version of the Čech complex.

Definition 12.1.10. Let X be a topological space equipped with a sheaf $\mathcal{F}$. Let $\mathfrak{U} = \{U_i\}$ be an open cover of X. We define

$$\mathscr{C}^p(\mathfrak{U},\mathcal{F}) = \prod_{i_0<\ldots<i_p} (j_{i_0\ldots i_p})_*\mathcal{F}|_{U_{i_0\ldots i_p}}$$

where $j_{i_0\ldots i_p}$ denotes the inclusion of $U_{i_0\ldots i_p}$ into X. These groups form a complex of sheaves $\mathscr{C}^\bullet(\mathfrak{U},\mathcal{F})$ under the Čech differential.

The key property of $\mathscr{C}^\bullet$ is that its cohomology groups encode the Čech cohomologies of the intersections of the U_i.

Lemma 12.1.11. *For every p and k, we have $\check{H}^k(X,\mathscr{C}^p(\mathfrak{U},\mathfrak{F})) = \prod_{i_0<\ldots<i_p} \check{H}^k(U_{i_0\ldots i_p},\mathcal{F})$.*

Proof. First suppose we fix an open cover $\mathfrak{V}$ of X. Then

$$\begin{aligned}
\check{C}^k(\mathfrak{V}, \mathscr{C}^p(\mathfrak{U}, \mathcal{F})) &= \prod_{\ell_0 < \ldots < \ell_k} \mathscr{C}^p(\mathcal{U}, \mathcal{F})(V_{\ell_0 \ldots \ell_k}) \\
&= \prod_{\ell_0 < \ldots < \ell_k} \prod_{i_0 < \ldots < i_p} \mathcal{F}(V_{\ell_0 \ldots \ell_k} \cap U_{i_0 \ldots i_p}) \\
&= \prod_{i_0 < \ldots < i_p} \check{C}^k(\mathfrak{V} \cap U_{i_0 \ldots i_p}, \mathcal{F}|_{U_{i_0 \cdots i_p}})
\end{aligned}$$

Since the differentials in the chain complex $\check{C}^\bullet(\mathfrak{V}, \mathscr{C}^p(\mathfrak{U}, \mathcal{F}))$ will respect the product structure in the last statement, we get an equality of cohomology groups

$$\check{H}^k(\mathfrak{V}, \mathscr{C}^p(\mathfrak{U}, \mathcal{F})) = \prod_{i_0 < \ldots < i_p} \check{H}^k(\mathfrak{V} \cap U_{i_0 \ldots i_p}, \mathcal{F}|_{U_{i_0 \cdots i_p}})$$

Taking a direct limit over all refinements of $\mathfrak{V}$ we obtain the desired statement. □

We are now equipped to prove Leray's Theorem.

Proof of Theorem 12.1.9: Consider the complex of sheaves

$$0 \to \mathcal{F} \to \mathscr{C}^0(\mathfrak{U}, \mathcal{F}) \to \mathscr{C}^1(\mathfrak{U}, \mathcal{F}) \to \mathscr{C}^2(\mathfrak{U}, \mathcal{F}) \to \ldots$$

It is a consequence of the sheaf axioms that this sequence is exact. Note also that $\check{H}^0(X, \mathscr{C}^k(\mathfrak{U}, \mathcal{F})) = \check{C}^k(\mathfrak{U}, \mathcal{F})$.

Under the hypotheses of Leray's Theorem, Lemma 12.1.11 shows that this resolution of $\mathcal{F}$ is acyclic. By Corollary 12.1.8 this implies that the kth homology groups of $\mathcal{F}$ agree with the kth homology groups of the global sections functor applied to this sequence:

$$\begin{aligned}
\check{H}^k(X, \mathcal{F}) &= \frac{\ker \check{H}^0(X, \mathscr{C}^k(\mathfrak{U}, \mathcal{F})) \to \check{H}^0(X, \mathscr{C}^{k+1}(\mathfrak{U}, \mathcal{F}))}{\operatorname{im} \check{H}^0(X, \mathscr{C}^{k-1}(\mathfrak{U}, \mathcal{F})) \to \check{H}^0(X, \mathscr{C}^k(\mathfrak{U}, \mathcal{F}))} \\
&= \frac{\ker \check{C}^k(\mathfrak{U}, \mathcal{F}) \to \check{C}^{k+1}(\mathfrak{U}, \mathcal{F})}{\operatorname{im} \check{C}^{k-1}(\mathfrak{U}, \mathcal{F}) \to \check{C}^k(\mathfrak{U}, \mathcal{F})} \\
&= \check{H}^k(\mathfrak{U}, \mathcal{F}).
\end{aligned}$$

□

12.1.5 Geometric applications

Before moving on to the special case of quasicoherent sheaves, we briefly mention a few geometric results. The first result gives conditions on the topological space X which guarantee that a SES of sheaves yields a LES of Čech cohomology groups.

Theorem 12.1.12. *Let X be a Hausdorff paracompact topological space. Suppose we have an exact sequence of sheaves*

$$0 \to \mathcal{F} \to \mathcal{G} \to \mathcal{H} \to 0.$$

Then we obtain a long exact sequence of Čech cohomology groups

$$\ldots \to \check{H}^p(X, \mathcal{F}) \to \check{H}^p(X, \mathcal{G}) \to \check{H}^p(X, \mathcal{H}) \to \check{H}^{p+1}(X, \mathcal{F}) \to \ldots$$

Note that this theorem will almost never apply in the setting of schemes. Our next result shows the equality of Čech cohomology and singular cohomology.

Theorem 12.1.13. *Let X be a topological space that is locally contractible. Fix an abelian group A. Then the Čech cohomology groups $\check{H}^i(X, A_X)$ are isomorphic to the singular cohomology groups $H^i(X, A)$.*

Suppose that X is a complex manifold and let $\mathcal{O}_X$ denote the sheaf of holomorphic functions on X. By composing with the exponential function, we obtain a map $exp : \mathcal{O}_X \to \mathcal{O}_X^\times$. Since the kernel of the exponential map consists of constant multiples of $2\pi i$, we obtain an exact sequence of sheaves of abelian groups

$$0 \to 2\pi i\mathbb{Z}_X \to \mathcal{O}_X \to \mathcal{O}_X^\times \to 0.$$

(Since invertible holomorphic functions need not admit a global logarithm, the map on the right is not surjective on global sections, but it is locally surjective due to the existence of local logarithms.)

Since X is paracompact we obtain a LES of sheaf cohomology. The various pieces of this sequence have interesting geometric significance. For example, we see that $\check{H}^1(X, 2\pi i\mathbb{Z}_X) = H^1(X, \mathbb{Z})$ is the obstruction to the global surjectivity of the exponential map. Next consider the portion of the LES

$$\ldots \to \check{H}^1(X, \mathcal{O}_X) \to \check{H}^1(X, \mathcal{O}_X^\times) \xrightarrow{\rho} \check{H}^2(X, 2\pi i\mathbb{Z}_X) \to \check{H}^2(X, \mathcal{O}_X) \to \check{H}^2(X, \mathcal{O}_X^\times) \to \ldots$$

The term $\check{H}^1(X, \mathcal{O}_X^\times)$ can be identified with the Picard group of holomorphic line bundles up to isomorphism (see Exercise 12.1.18). On the right, the connecting homomorphism ρ is the first Chern class map with values in $H^2(X, \mathbb{Z})$. On the left, we see that we can identify

$$\operatorname{Pic}(X) \cong \frac{\check{H}^1(X, \mathcal{O}_X)}{\check{H}^1(X, 2\pi i\mathbb{Z}_X)} \cong \frac{H^{0,1}(X, \mathbb{C})}{\operatorname{im} H^1(X, \mathbb{Z})}$$

12.1.6 Exercises

Exercise 12.1.14. Suppose that X is an irreducible topological space and that A_X is a locally constant sheaf on X with value A. Prove that $\check{H}^i(X, A_X) = 0$ for every $i > 0$.

Exercise 12.1.15. Prove that the Čech cohomology groups of $\mathcal{F} \oplus \mathcal{G}$ are the direct sums of the cohomology groups of $\mathcal{F}$ and $\mathcal{G}$.

Exercise 12.1.16. Let X be a topological space and suppose that

$$0 \to \mathcal{F} \to \mathcal{G} \to \mathcal{H} \to 0$$

is an exact sequence of sheaves. Prove that we have an exact sequence

$$0 \to \check{H}^0(X, \mathcal{F}) \to \check{H}^0(X, \mathcal{G}) \to \check{H}^0(X, \mathcal{H}) \to \check{H}^1(X, \mathcal{F})$$

without any extra assumptions on X.

Exercise 12.1.17. Compute the cohomology groups of the circle and the sphere with respect to the locally constant sheaf with value $\mathbb{Z}$. (Feel free to use the fact that all the higher cohomology of the locally constant sheaf $\mathbb{Z}_X$ along a contractible topological space will vanish.)

Exercise 12.1.18. Let X be an integral scheme. Prove that the Picard group of X is isomorphic to $\check{H}^1(X, \mathcal{O}_X^\times)$.

(Remark: in fact this theorem holds for any ringed space. While you should prove this theorem "by hand", it may be enlightening to compare this result against the connecting homomorphism in the sequence

$$\begin{tikzcd}
\check{H}^0(X, \mathcal{K}^\times) \arrow[r] \arrow[d, leftrightarrow, "\cong"] & \check{H}^0(X, \mathcal{K}^\times/\mathcal{O}_X^\times) \arrow[r] \arrow[d, leftrightarrow, "\cong"] & \check{H}^1(X, \mathcal{O}_X^\times) \arrow[r] & \check{H}^1(X, \mathcal{K}^\times) \arrow[d, leftrightarrow, "\cong"] \\
K(X)^\times & \mathrm{CDiv}(X) & & 0
\end{tikzcd}$$

Can you prove these isomorphisms?)

12.2 Čech cohomology for quasicoherent sheaves

Since most schemes do not have a well-behaved topology, we can not expect the abstract theory of Čech cohomology to be well-behaved. However, it turns out that Čech cohomology is well-behaved for quasicoherent sheaves. In this section we develop the foundational properties of this theory.

12.2.1 Vanishing for affines

We saw in Proposition 9.2.8 that for quasicoherent sheaves on affine varieties the global sections functor is exact. Thus we might expect that the first cohomology group of a quasicoherent sheaf on an affine scheme is equal to 0. The following result is a step in this direction.

Theorem 12.2.1. *Let $X = \mathrm{Spec}(R)$ be an affine scheme and let $\mathcal{F}$ be a quasicoherent sheaf on X. Then we have $\check{H}^i(X, \mathcal{F}) = 0$ for every $i > 0$.*

Proof. Since $\mathcal{F}$ is quasicoherent, we have $\mathcal{F} = \widetilde{M}$ for some R-module M. Since distinguished open affines form a base for the topology on X, it suffices to show that

$$\check{H}^i(\mathfrak{U}, \mathcal{F}) = 0$$

whenever $\mathfrak{U} = \{U_1, \dots, U_r\}$ is an open cover consisting of distinguished open affines $U_i = D_{f_i}$. Since $\mathcal{F}|_{U_i} = \widetilde{M_{f_i}}$, we must show that the following sequence is exact:

$$0 \to M \xrightarrow{d^0} \prod_i M_{f_i} \xrightarrow{d^1} \prod_{i<j} M_{f_i f_j} \xrightarrow{d^2} \prod_{i<j<k} M_{f_i f_j f_k} \xrightarrow{d^3} \dots$$

Note that we have already proved exactness of the first few terms in the proof of Proposition 9.2.1. The proof of exactness of the other terms is essentially the same.

Let us show exactness at the pth term. Suppose that $\sigma \in \prod_{i_0 < \dots < i_p} M_{f_{i_0} \cdots f_{i_p}}$ is in the kernel of d^p. Since the index set is finite, by passing to common denominators we may write

$$\sigma_{i_0 \dots i_p} = \frac{m_{i_0 \dots i_p}}{(f_{i_0} \dots f_{i_p})^t}$$

for some $m_{i_0 \dots i_p} \in M$ and some non-negative integer t. Since $\sigma \in \ker(d^p)$, for any $p+1$ indices $j_0 < \dots < j_{p+1}$ we have that $d\sigma_{j_0 \dots j_{p+1}} = 0$, or in other words,

$$\sum_{k=0}^{p+1} (-1)^k \frac{f_{j_k}^t m_{j_0 \dots \widehat{j_k} \dots j_{p+1}}}{(f_{j_0} \dots f_{j_{p+1}})^t} = 0$$

as elements of $M_{f_{j_0} \cdots f_{j_{p+1}}}$. We can rewrite this equation inside of $M_{f_{j_0} \cdots f_{j_p}}$ by clearing denominators at the cost of adding in an additional power $f_{j_{p+1}}^l$ for a non-negative integer

l (to put our element in the kernel of the localization map). Since there are only finitely many indices, we may also ensure that l is independent of the choice of indices. In this way we obtain the equation

$$\sum_{k=0}^{p}(-1)^k f_{j_{p+1}}^l \frac{f_{j_k}^t m_{j_0\ldots\widehat{j_k}\ldots j_{p+1}}}{(f_{j_0}\cdots f_{j_p})^t} = (-1)^p f_{j_{p+1}}^l \frac{f_{j_{p+1}}^t m_{j_0}\ldots j_p\widehat{j_{p+1}}}{(f_{j_0}\cdots f_{j_p})^t}. \tag{12.2.1}$$

Since the D_{f_i} form an open cover, we have a relation $1 = \sum(-1)^p g_i f_i^{t+l}$ in R. We define $\tau \in \prod_{i_0<\ldots<i_p} M_{f_{i_0}\cdots f_{i_{p-1}}}$ by setting

$$\tau_{i_0\ldots i_{p-1}} = \sum_i g_i f_i^l \frac{m_{ii_0\ldots i_{p-1}}}{(f_{i_0}\cdots f_{i_{p-1}})^t}$$

where as usual we implicitly allow ourselves to reorder the indices in $m_{ii_0\ldots i_{p-1}}$. We claim that $d\tau = \sigma$. Indeed, we have

$$\begin{aligned}
(d\tau)_{i_0\ldots i_p} &= \sum_{k=0}^{p}(-1)^k \sum_i g_i f_i^l \frac{f_{i_k}^t m_{ii_0\ldots\widehat{i_k}\ldots i_p}}{(f_{i_0}\cdots f_{i_p})^t} \\
&= \sum_i g_i \sum_{k=0}^{p}(-1)^k f_i^l \frac{f_{i_k}^t m_{ii_0\ldots\widehat{i_k}\ldots i_p}}{(f_{i_0}\cdots f_{i_p})^t} \\
&= \sum_i g_i \left((-1)^p f_i^{l+t} \frac{m_{i_0\ldots i_p}}{(f_{i_0}\cdots f_{i_p})^t}\right) \text{ by Equation (12.2.1)} \\
&= \frac{m_{i_0\ldots i_p}}{(f_{i_0}\cdots f_{i_p})^t}
\end{aligned}$$

This concludes the proof of exactness at the pth term. □

12.2.2 Applications

Our first application of Theorem 12.2.1 allows us to find a specific open cover which computes Čech cohomology for quasicoherent sheaves. Leray's Theorem 12.1.9 has the following immediate consequence:

Theorem 12.2.2. *Let X be a quasicompact scheme and let $\mathcal{F}$ denote a quasicoherent sheaf on X. Suppose that $\mathfrak{U} = \{U_i\}$ is an open cover of X by open affines U_i such that all the non-empty intersections $U_{i_0} \cap \ldots \cap U_{i_p}$ are affine. Then $\check{H}^i(\mathfrak{U}, \mathcal{F}) = \check{H}^i(X, \mathcal{F})$.*

Note that the quasicompactness hypothesis is necessary to ensure that we can find a finite open cover of X by open affines. The best case is when X is separated.

Corollary 12.2.3. *Let X be a quasicompact separated scheme and let $\mathcal{F}$ denote a quasicoherent sheaf on X. Then for any open affine cover $\mathfrak{U}$ of X we have $\check{H}^i(\mathfrak{U}, \mathcal{F}) = \check{H}^i(X, \mathcal{F})$.*

Now that we have shown that Čech cohomology can be computed on specific open covers, we can construct a LES of Čech cohomology for quasicoherent sheaves.

Theorem 12.2.4. *Let X be a quasicompact separated scheme. Suppose we have an exact sequence of quasicoherent sheaves*

$$0 \to \mathcal{F} \to \mathcal{G} \to \mathcal{H} \to 0.$$

Then we obtain a long exact sequence of Čech cohomology groups

$$\ldots \to \check{H}^p(X, \mathcal{F}) \to \check{H}^p(X, \mathcal{G}) \to \check{H}^p(X, \mathcal{H}) \to \check{H}^{p+1}(X, \mathcal{F}) \to \ldots$$

Proof. If we let $\mathfrak{U}$ be an open cover by open affines then Corollary 12.2.3 shows that for every quasicoherent sheaf $\mathcal{F}$ we have $\check{H}^i(X, \mathcal{F}) \cong \check{H}^i(\mathfrak{U}, \mathcal{F})$. Since every open set V that is an intersection of elements in our open cover is affine and since all our sheaves are quasicoherent, the map $\mathcal{G}(V) \to \mathcal{H}(V)$ is surjective for every such open set V. Thus for an cover $\mathfrak{U}$ by open affines we have an exact sequence of complexes

$$0 \to \check{C}^\bullet(\mathfrak{U}, \mathcal{F}) \to \check{C}^\bullet(\mathfrak{U}, \mathcal{G}) \to \check{C}^\bullet(\mathfrak{U}, \mathcal{H}) \to 0.$$

The corresponding LES of cohomology groups yields the desired sequence. □

Our next application of Theorem 12.2.1 is an important vanishing result.

Theorem 12.2.5. *Let X be a quasiprojective $\mathbb{K}$-scheme of dimension n and let $\mathcal{F}$ be a quasicoherent sheaf on X. Then $\check{H}^i(X, \mathcal{F}) = 0$ for every $i > n$.*

The idea is to show that X admits an open cover consisting of $(n+1)$ open affines. If X is projective we can find such a cover as follows. First, we take a closed embedding of X into some projective space. Since intersecting a projective scheme by a general hypersurface drops the dimension by 1, we can find $(n+1)$ hypersurfaces H_i such that $X \cap H_1 \cap \ldots \cap H_{n+1}$ is empty. The complements U_i of these hypersurfaces are affine, and the subschemes $X \cap U_i$ give the desired open cover.

Since we are only assuming that X is quasiprojective, we need to be a little more careful: if we repeat the argument above the sets $X \cap U_i$ may not be closed subschemes of U_i and thus are not necessarily affine. Nevertheless a similar argument will work.

Proof. We claim that X admits an open cover $\mathfrak{U}$ consisting of $n+1$ open affine sets. For such a cover it is clear that $\check{H}^i(\mathfrak{U}, \mathcal{F}) = 0$ for $i > n$. Since X is quasicompact separated, this also implies $\check{H}^i(X, \mathcal{F}) = 0$ for $i > n$ by Corollary 12.2.3.

Consider the composition of an open embedding from X into a projective $\mathbb{K}$-scheme $\overline{X}$ followed by a closed embedding $\overline{X} \hookrightarrow \mathbb{P}^n$. We set $Z = \overline{X} \backslash X$. We claim that there is a hypersurface H_1 which contains Z but does not contain any component of $\overline{X}$. Indeed, if I_Z denotes the radical of the homogeneous ideal defining Z and I_j denotes the radical of the homogeneous ideal defining the jth component of $\overline{X}$ then none of the I_j contain I_Z and thus we can let H_1 be the vanishing locus of any homogeneous element of $I_Z \backslash (\cup_j I_j)$. Since H_1 does not contain any component of $\overline{X}$ Exercise 4.4.16 shows that $\dim(H_1 \cap X) = n-1$. Since H_1 does contain Z we see that $X \backslash H_1 \cong \overline{X} \backslash H_1$ is affine (since it is a closed subscheme of the affine scheme $\mathbb{P}^n \backslash H_1$).

We now repeat the argument to find a hypersurface H_2 which contains Z but does not contain any component of $\overline{X} \cap H_1$. Inducting, we find hypersurfaces $H_1, \ldots, H_n$ such that each $X \backslash H_i$ is affine and $\dim(H_1 \cap \ldots \cap H_j \cap X) = \dim(X) - j$. Finally, since $H_1 \cap \ldots \cap H_n \cap X$ is 0-dimensional, we can take one more hypersurface H_{n+1} that contains Z and avoids this 0-dimensional subset. Then the open affines $X \backslash H_i$ form an open cover of X. $\square$

12.2.3 Exercises

Exercise 12.2.6. Let X be the complement of the origin in $\mathbb{A}^2_{\mathbb{K}}$. Using the cover $\mathfrak{U}$ consisting of the complements of the two axes, show that $\check{H}^1(X, \mathcal{O}_X)$ is isomorphic to the $\mathbb{K}$-vector space spanned by $\{x^i y^j | i, j < 0\}$. In particular, this space is infinite dimensional.

Exercise 12.2.7. Let $X \subset \mathbb{P}^2_{\mathbb{K}}$ be a degree d hypersurface. Prove that

$$\dim_{\mathbb{K}} \check{H}^0(X, \mathcal{O}_X) = 1$$
$$\dim_{\mathbb{K}} \check{H}^1(X, \mathcal{O}_X) = (d-1)(d-2)/2$$

Exercise 12.2.8. Let $\mathbb{L}/\mathbb{K}$ be a field extension. Suppose that X is a $\mathbb{K}$-scheme and that $\mathcal{F}$ is a quasicoherent sheaf on X. Let $p : X_{\mathbb{L}} \to X$ be the base change map. Prove that

$$\check{H}^i(X_{\mathbb{L}}, p^*\mathcal{F}) \cong \check{H}^i(X, \mathcal{F}) \otimes_{\mathbb{K}} \mathbb{L}.$$

Exercise 12.2.9. Let $f : X \to Y$ be an affine morphism of Noetherian separated schemes. Prove that for any quasicoherent sheaf $\mathcal{F}$ on X we have isomorphisms

$$\check{H}^i(X, \mathcal{F}) \cong \check{H}^i(Y, f_*\mathcal{F}).$$

In particular, this applies to closed embeddings and finite morphisms.

Exercise 12.2.10. Let X be a quasicompact separated scheme and let $\mathcal{F}$ and $\mathcal{G}$ be quasicoherent sheaves on X. In this exercise we construct a cup product

$$\cup : \check{H}^p(X, \mathcal{F}) \times \check{H}^q(X, \mathcal{G}) \to \check{H}^{p+q}(X, \mathcal{F} \otimes \mathcal{G}).$$

Let $\mathfrak{U}$ denote a finite open affine cover of X. Define the map

$$\cup : C^p(\mathfrak{U}, \mathcal{F}) \times C^q(\mathfrak{U}, \mathcal{G}) \to C^{p+q}(\mathfrak{U}, \mathcal{F} \otimes \mathcal{G})$$

where for any ordered subset $K = \{i_0, \ldots, i_{p+q}\}$ of the index set for $\mathfrak{U}$ we define

$$(\sigma \cup \tau)_K = \sigma_{i_0 \ldots i_p} \otimes \tau_{i_p \ldots i_{p+q}}.$$

Prove the relation $d(\sigma \cup \tau) = (d\sigma \cup \tau) + (-1)^p(\sigma \cup d\tau)$ and use it to show that the function $\cup$ descends to cohomology.

12.3 Cohomology of sheaves on projective space

In this section we prove several foundational results about the cohomology of sheaves on projective space. We will fix a ring A and consider the cohomology groups of sheaves on $\mathbb{P}^n_A := \mathrm{Proj}(A[x_0, \ldots, x_n])$.

12.3.1 Cohomology of invertible sheaves

Our starting point will be to compute the Čech cohomology of the invertible sheaves on projective space.

Theorem 12.3.1. *(1) For $m \geq 0$, we have*

$$\check{H}^i(\mathbb{P}^n_A, \mathcal{O}(m)) \cong \begin{cases} A^{\oplus\binom{n+m}{n}} & \text{if } i = 0 \\ 0 & \text{otherwise} \end{cases}$$

(2) For $-n \leq m \leq -1$, we have $\check{H}^i(\mathbb{P}^n_A, \mathcal{O}(m)) = 0$ for every i.

(3) For $m \leq -n-1$, we have

$$\check{H}^i(\mathbb{P}^n_A, \mathcal{O}(m)) = \begin{cases} A^{\oplus\binom{-m-1}{n}} & \text{if } i = n \\ 0 & \text{otherwise} \end{cases}$$

Note that the cohomology groups for all $i \neq 0, n$ are vanishing. For the remaining degrees $0, n$ the cohomology groups are free and we have a symmetry $\check{H}^0(\mathbb{P}^n_A, \mathcal{O}(m)) \cong \check{H}^n(\mathbb{P}^n_A, \mathcal{O}(-n-1-m))$. For example, on $\mathbb{P}^2_A$ the cohomology groups are free modules of the following ranks

	$\mathcal{O}(-6)$	$\mathcal{O}(-5)$	$\mathcal{O}(-4)$	$\mathcal{O}(-3)$	$\mathcal{O}(-2)$	$\mathcal{O}(-1)$	$\mathcal{O}$	$\mathcal{O}(1)$	$\mathcal{O}(2)$	$\mathcal{O}(3)$
$\check{H}^0$	0	0	0	0	0	0	1	3	6	10
$\check{H}^1$	0	0	0	0	0	0	0	0	0	0
$\check{H}^2$	10	6	3	1	0	0	0	0	0	0

Our strategy is to use the standard covering of $\mathbb{P}^n_A$ by open affine charts

Proof. It is notationally simpler to compute the cohomology of the graded sheaf $\mathcal{F} = \oplus_{m\in\mathbb{Z}}\mathcal{O}(m)$ and use the grading to recover the cohomology of the pieces at the end. Set $U_i = D_{+,x_i}$ and consider the open cover $\mathfrak{U} = \{U_i\}$. By Corollary 12.2.3 we have $\check{H}^k(\mathbb{P}^n_A, \mathcal{F}) = \check{H}^k(\mathfrak{U}, \mathcal{F})$.

The Čech complex in this situation is

$$\prod_{i=0}^{n} R_{x_i} \to \prod_{\substack{I\subset\{0,\ldots,n\}\\|I|=2}} R_{x_I} \to \prod_{\substack{I\subset\{0,\ldots,n\}\\|I|=3}} R_{x_I} \to \ldots \to \prod_{\substack{I\subset\{0,\ldots,n\}\\|I|=n}} R_{x_I} \to R_{x_0x_1\ldots x_n} \to 0$$

where for a subset $I \subset \{0, \ldots, n\}$ we denote by x_I the product of the variables indexed by I. Using our usual localization exact sequence we see that $\check{H}^0(\mathfrak{U}, \mathcal{F}) \cong R$. Since this isomorphism respects the grading, we see that $\check{H}^0(\mathbb{P}^n_A, \mathcal{O}(m)) \cong R_m \cong A^{\binom{n+m}{n}}$.

We next address the top cohomology groups $\check{H}^n$. Consider the last map

$$\prod_{\substack{I \subset \{0,\ldots,n\} \\ |I|=n}} R_{x_I} \to R_{x_0 x_1 \ldots x_n}$$

in the sequence. The module on the right is a free A-module spanned by products of the form $x_0^{k_0} x_1^{k_1} \ldots x_n^{k_n}$ where the exponents $k_i \in \mathbb{Z}$. The image of the module on the left consists of those products which have at least one non-negative exponent. Thus the cokernel of this map is the free A-module with basis

$$\{ x_0^{k_0} \ldots x_n^{k_n} \mid k_i < 0 \ \forall i \}.$$

Note that the degree is the sum of the exponents. Thus there are no basis vectors in degree $m \geq -n$, and in degree $m \leq -n-1$ the number of basis vectors is exactly $\binom{-m-1}{n}$ as desired.

Finally, we address the cohomology groups $\check{H}^i$ for $0 < i < n$. Consider the exact sequence

$$0 \to \mathcal{O}(-1) \xrightarrow{\cdot x_n} \mathcal{O} \to i_* \mathcal{O}_{\mathbb{P}^{n-1}_A} \to 0$$

where $i : \mathbb{P}^{n-1} \to \mathbb{P}^n$ is the inclusion of the hyperplane $V_+(x_n)$. Since $\mathcal{F}$ is free, the exactness is preserved by tensoring by $\mathcal{F}$. By the projection formula we have

$$\mathcal{F} \otimes i_* \mathcal{O}_{\mathbb{P}^{n-1}_A} \cong i_* \left(\mathcal{F}|_{\mathbb{P}^{n-1}_A} \right) \cong i_* \mathcal{F}'$$

where $\mathcal{F}'$ is the sheaf $\oplus_{m \in \mathbb{Z}} \mathcal{O}(m)$ on $\mathbb{P}^{n-1}$. Thus we get a long exact sequence of Čech cohomology groups associated to the SES of sheaves

$$0 \to \mathcal{F}(-1) \to \mathcal{F} \to i_* \mathcal{F}' \to 0.$$

By Exercise 12.2.9 the pushforward i_* under a closed embedding preserves all the cohomology groups. Thus by induction on n we see that $\check{H}^i(\mathbb{P}^n_A, \mathcal{F}(-1)) \to \check{H}^i(\mathbb{P}^n_A, \mathcal{F})$ is an isomorphism for $1 < i < n-1$. We claim that this map is also an isomorphism for $i = 1, n-1$. To show that $\check{H}^1(\mathbb{P}^n_A, \mathcal{F}(-1)) \to \check{H}^1(\mathbb{P}^n_A, \mathcal{F})$ is an isomorphism, it suffices to show that the restriction map $\check{H}^0(\mathbb{P}^n_A, \mathcal{F}) \to \check{H}^0(\mathbb{P}^{n-1}_A, \mathcal{F}')$ is surjective. This is just the map $R \mapsto R/(x_n)$.

To show this map is an isomorphism for $i = n-1$, consider the last piece of the LES

$$\check{H}^{n-1}(\mathbb{P}^{n-1}_A, \mathcal{F}') \xrightarrow{\delta} \check{H}^n(\mathbb{P}^n_A, \mathcal{F}(-1)) \to \check{H}^n(\mathbb{P}^n_A, \mathcal{F}) \to 0.$$

As we described earlier $\check{H}^n(\mathbb{P}^n_A, \mathcal{F})$ is a free A-module over the monomials whose exponents are all negative. The surjection on the right is the map $\cdot x_n$ and its kernel is the set of monomials $x_0^{k_0} \dots x_n^{k_n}$ with $k_n = -1$. The leftmost module is a free A-module over the monomials in $x_0, \dots, x_{n-1}$ with negative exponents. We claim that the leftmost map δ is multiplication by x_n^{-1} and is thus injective. To see this we return to the SES of Čech complexes which define this LES. The relevant portion is

$$\begin{tikzcd}
0 \arrow[r] & \prod_{i=0}^n R(-1)_{x_0 \dots \widehat{x_i} \dots x_n} \arrow[r, "\cdot x_n"] \arrow[d] & \prod_{i=0}^n R_{x_0 \dots \widehat{x_i} \dots x_n} \arrow[r] \arrow[d] & (R/(x_n))_{x_0 \dots x_{n-1}} \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & R(-1)_{x_0 \dots x_n} \arrow[r, "\cdot x_n"] & R_{x_0 \dots x_n} \arrow[r] & 0 &
\end{tikzcd}$$

where the columns are the last two entries in the Čech complex and the rows are induced by the SES of sheaves. Tracing through the construction of the Snake Lemma, suppose we have an element $x_0^{j_0} \dots x_{n-1}^{j_{n-1}}$ in the cohomology at the top right place. One choice of a lift to the top middle place is the $(n+1)$-tuple consisting of this element in the last entry with 0s in the other entries. The image of this choice of lift in the bottom middle place is $x_0^{j_0} \dots x_{n-1}^{j_{n-1}}$. Finally, the preimage in the bottom left is $x_0^{j_0} \dots x_{n-1}^{j_{n-1}} x_n^{-1}$, proving the claim.

Altogether we have shown that multiplication by x_n induces isomorphisms between $\check{H}^i(\mathbb{P}^n_A, \mathcal{F}(-1))$ and $\check{H}^i(\mathbb{P}^n_A, \mathcal{F})$ for $0 < i < n$. It only remains to show that this property forces these cohomology groups to vanish. Suppose we take the Čech complex constructed above and localize it along x_n. The result will be a Čech complex for the sheaf $\oplus_{d \geq 0} \mathcal{O}_{U_i}$ of the affine chart U_n with respect to the open cover given by the various $U_i \cap U_n$. By Corollary 12.2.3 the cohomology of this localized sequence will compute the Čech cohomology groups of the sheaf $\oplus_{d \geq 0} \mathcal{O}_{U_i}$ on U_i and by Theorem 12.2.1 all the higher cohomology groups of this sheaf 0. Since localization is exact, we conclude that $H^i(\mathbb{P}^n_A, \mathcal{F})_{x_n} = 0$. In other words, every element of this group is annihilated by some power of x_n. Since we have also shown that multiplication by x_n is an isomorphism, we conclude that $H^i(\mathbb{P}^n_A, \mathcal{F}) = 0$. $\square$

Remark 12.3.2. In fact the proof shows a little more. Using the functoriality of cohomology, we have a natural pairing

$$\mathcal{H}om(\mathcal{O}(-n-1-d), \mathcal{O}(-n-1)) \times \check{H}^n(\mathbb{P}^n, \mathcal{O}(-n-1-d)) \to \check{H}^n(\mathbb{P}^n, \mathcal{O}(-n-1)) \cong A.$$

The left-most term can be identified with $\check{H}^0(\mathbb{P}^n, \mathcal{O}(d))$. The proof shows that for any $d \geq 0$ this identification is a perfect pairing of free A-modules obtained by multiplication.

12.3.2 Cohomology of other sheaves

Hilbert's Syzygy Theorem guarantees that any sheaf on projective space admits a finite resolution by direct sums of line bundles. We use this structure to prove a key result.

Theorem 12.3.3. *Let A be a Noetherian ring and let X be a projective scheme over* $\mathrm{Spec}(A)$*. Then for any coherent sheaf $\mathcal{F}$ on X, the cohomology group $\check{H}^i(X, \mathcal{F})$ is is a coherent A-module for every i.*

Proof. Since every affine scheme admits an ample line bundle, Proposition 10.7.13 shows that X admits a closed embedding into some projective space $\mathbb{P}^n_A$. Exercise 12.2.9 shows that cohomology groups are unchanged when we push forward by a closed embedding. Thus we may assume that $\mathcal{F}$ is a coherent sheaf on $\mathbb{P}^n_A$. We show that $\check{H}^i(\mathbb{P}^n_A, \mathcal{F})$ is coherent by decreasing induction on the index i. By Corollary 10.6.13 we have an exact sequence

$$0 \to \mathcal{G} \to \mathcal{O}(d)^{\oplus r} \to \mathcal{F} \to 0$$

for some coherent sheaf $\mathcal{G}$. Since $\mathbb{P}^n_A$ is quasicompact and separated, Theorem 12.2.4 shows that we obtain a LES of Čech cohomology. The end of the sequence is

$$\check{H}^n(\mathbb{P}^n_A, \mathcal{O}(d)^{\oplus r}) \to \check{H}^n(\mathbb{P}^n_A, \mathcal{F}) \to 0$$

since by Theorem 12.2.5 all the higher cohomology groups vanish. Theorem 12.3.1 shows that the term on the left is a coherent A-module, thus the term in the middle is also finitely generated and thus (since A is Noetherian) coherent. For $i < n$, we have an exact sequence

$$\check{H}^i(\mathbb{P}^n_A, \mathcal{O}(d)^{\oplus r}) \to \check{H}^i(\mathbb{P}^n_A, \mathcal{F}) \to \check{H}^{i+1}(\mathbb{P}^n_A, \mathcal{G}).$$

The leftmost term is a coherent A-module by Theorem 12.3.1 and thus its image in the middle term is finitely generated. The kernel of the rightmost map is coherent by induction. We conclude that the middle term is a coherent A-module. □

In particular, we can prove a special case of Theorem 9.4.9:

Corollary 12.3.4. *Let $f : X \to Y$ be a projective morphism to a locally Noetherian scheme Y. If $\mathcal{F}$ is a coherent sheaf on X then $f_*\mathcal{F}$ is a coherent sheaf on Y.*

Proof. It suffices to prove this when Y is an affine scheme. In this case the desired statement follows from Theorem 12.3.3. □

Corollary 12.3.5. *Let $f : X \to Y$ be a morphism to a locally Noetherian scheme Y. Then f is projective and affine if and only if f is finite.*

Proof. We have already proved the reverse implication. To prove the forward implication, the projectivity assumption shows that $f_*\mathcal{O}_X$ is a coherent sheaf. Since f is affine, we see that for every open affine $V \subset Y$ the ring $\mathcal{O}_X(f^{-1}V)$ is a finitely generated $\mathcal{O}_Y(V)$-module. □

Corollary 12.3.6. *Let $f : X \to Y$ be a projective morphism with finite fibers. Then f is a finite morphism.*

Proof. We first show that f is an affine morphism. Fix a point $y \in Y$ and let V be an open affine neighborhood of y. Since V admits an ample invertible sheaf, the preimage X_V admits a closed embedding into $\mathbb{P}^n_V$. Since f has finite fibers there is a hyperplane $H \subset \mathbb{P}^n_V$ which does not meet the fiber of f over y. Let $U \subset V$ denote the open neighborhood of V which is the complement of $f(H \cap X)$. Note that U contains y and that $f^{-1}U$ is affine. Since every point y admits an open neighborhood whose preimage is affine, we deduce that f is an affine morphism.

By Corollary 12.3.5 we conclude that f is finite. $\square$

12.3.3 Exercises

Exercise 12.3.7. Recall that the cotangent bundle Ω^1 of projective space $\mathbb{P}^n_A$ fits into an exact sequence

$$0 \to \Omega^1 \to \mathcal{O}(-1)^{\oplus n+1} \to \mathcal{O} \to 0.$$

Use this exact sequence to compute the cohomology groups of Ω^1.

Exercise 12.3.8. Let I be the ideal sheaf of the twisted cubic in $\mathbb{P}^3_{\mathbb{K}}$. Compute the cohomology groups of $I(m)$ for every $m \geq 0$.

(Hint: there are two approaches. One is to use the LES of cohomology associated to the exact sequence

$$0 \to I(m) \to \mathcal{O}_{\mathbb{P}^n}(m) \to i_*\mathcal{O}_C(m) \to 0.$$

In this case one must understand the maps $\check{H}^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(m)) \to \check{H}^0(\mathbb{P}^1, \mathcal{O}_{\mathbb{P}^1}(3m))$ explicitly in some way. Another method is to use the fact that I is generated by three quadrics to construct an exact sequence

$$0 \to \mathcal{O}(-3)^{\oplus 2} \to \mathcal{O}(-2)^{\oplus 3} \to I \to 0$$

where the leftmost map describes the relations between the generators of I.)

Exercise 12.3.9. Recall that every line bundle on $\mathbb{P}^m \times \mathbb{P}^n$ has the form $\mathcal{O}(a, b)$ for some $a, b \in \mathbb{Z}$. Prove a Kunneth formula for the corresponding Čech cohomology groups:

$$\check{H}^k(\mathbb{P}^m \times \mathbb{P}^n, \mathcal{O}(a, b)) \cong \oplus_{i+j=k} \check{H}^i(\mathbb{P}^m, \mathcal{O}(a)) \otimes_{\mathbb{K}} \check{H}^j(\mathbb{P}^n, \mathcal{O}(b)).$$

(Hint: construct the Čech complex on $\mathbb{P}^m \times \mathbb{P}^n$ using the open cover that consists of the products of the standard opens in $\mathbb{P}^m$ and $\mathbb{P}^n$. Show that the complex you get is the tensor product of the complexes from the two factors.)

12.4 Cohomology and ample line bundles

It is easiest to perform geometric computations with sheaves which have vanishing higher cohomology. The most important way to find sheaves with vanishing higher cohomology is to leverage the properties of ample line bundles.

12.4.1 Serre's Criterion

Theorem 12.4.1. *Let A be a Noetherian ring, let X be a proper scheme over* $\mathrm{Spec}(A)$*. Suppose that $\mathcal{L}$ is an invertible sheaf on X. Then the following are equivalent:*

(1) $\mathcal{L}$ is ample.

(2) For every coherent sheaf $\mathcal{F}$ on X there exists a positive integer m_0 such that $\check{H}^i(X, \mathcal{F} \otimes \mathcal{L}^m) = 0$ for every $i > 0$ and $m \geq m_0$.

The forward implication is known as Serre Vanishing; the reverse implication is known as Serre's Criterion.

Proof. ($\implies$): We first prove the forward implication in the case when $\mathcal{L}$ is very ample. Since X is proper and $\mathcal{L}$ is very ample, X must be projective and is thus equipped with a closed embedding $\phi : X \to \mathbb{P}^n_A$ such that $\phi^*\mathcal{O}(1) = \mathcal{L}$. By the projection formula and Exercise 12.2.9 we have

$$\check{H}^i(X, \mathcal{F} \otimes \mathcal{L}^{\otimes m}) = \check{H}^i(\mathbb{P}^n_A, \phi_*\mathcal{F} \otimes \mathcal{O}(m)).$$

Furthermore $\phi_*\mathcal{F}$ is coherent since ϕ is a closed embedding. Thus we may assume henceforth that $X \cong \mathbb{P}^n_A$ and $\mathcal{L} = \mathcal{O}(1)$.

We prove the vanishing of the higher cohomology $\check{H}^i(\mathbb{P}^n_A, \mathcal{F} \otimes \mathcal{O}(m))$ by decreasing induction on i. By Corollary 10.6.13 any coherent sheaf $\mathcal{F}$ on projective space fits into a short exact sequence

$$0 \to \mathcal{K} \to \mathcal{O}_{\mathbb{P}^n_A}(d)^{\oplus r} \to \mathcal{F} \to 0$$

If we choose m_0 large enough, then $\check{H}^i(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n_A}(d+m)) = 0$ for every $m > m_0$. We conclude that $\check{H}^n(\mathbb{P}^n, \mathcal{F}(m)) = 0$, proving the base case of the induction. We also see that for $m > m_0$ we have $\check{H}^i(\mathbb{P}^n, \mathcal{F}(m)) \cong \check{H}^{i+1}(\mathbb{P}^n, \mathcal{G}(m))$ for $0 < i < n$. By the inductive assumption, after possibly increasing m_0 we can ensure that the $(i+1)$ cohomology of $\mathcal{G}(m)$ vanishes. We can then conclude the vanishing of the ith cohomology for $\mathcal{F}(m)$.

Finally suppose that $\mathcal{L}$ is only ample. Choose a positive integer M such that $\mathcal{L}^{\otimes M}$ is very ample. We can then deduce the statement for the original $\mathcal{L}$ by applying the very ample case for $\mathcal{L}^{\otimes M}$ to the sheaves $\mathcal{F} \otimes \mathcal{L}^{\otimes i}$ for $i = 0, 1, 2, \ldots, M-1$.

($\Longleftarrow$): Fix any closed point $x \in X$. Tensor the surjection $\mathcal{O}_X \to i_*\mathcal{O}_x$ by $\mathcal{F}$ and let $\mathcal{K}$ denote the kernel of the resulting map. Since $\mathcal{K}$ is coherent, for sufficiently large exponents m we have $\check{H}^1(X, \mathcal{K} \otimes \mathcal{L}^{\otimes m}) = 0$. We deduce that there is a surjection

$$\check{H}^0(X, \mathcal{F} \otimes \mathcal{L}^{\otimes m}) \to \check{H}^0(X, \mathcal{F} \otimes i^*\mathcal{O}_x \otimes \mathcal{L}^{\otimes m}).$$

Thus for sufficiently large exponents m we have that $\mathcal{F} \otimes \mathcal{L}^{\otimes m}$ is globally generated at x. Applying Exercise 10.1.14 we see that for every point x there is a constant m_x such that for every $m \geq m_x$ there is an open neighborhood $U_{m,x}$ of x such that $\mathcal{F} \otimes \mathcal{L}^{\otimes m}$ is globally generated at every point in $U_{m,x}$. (Note that $U_{m,x}$ depends on the choice of m, so that we do not yet know that we can choose one open neighborhood at x to work for all m. The rest of the argument is devoted to removing this dependence.)

Applying the above argument to $\mathcal{F} = \mathcal{O}_X$, we see there is some power $\mathcal{L}^{\ell_x}$ which is globally generated on a neighborhood V of x. For each $i = 0, 1, \ldots, \ell_x - 1$, let $V_i = U_{m_x+i,x}$ denote an open neighborhood of x along which $\mathcal{F} \otimes \mathcal{L}^{\otimes m_x+i}$ is globally generated. Taking intersections of V and all the V_i, we obtain an open set W_x. Suppose that $m \geq m_x$ and write $m = m_x + i + q\ell_x$ for some integer $q \geq 0$ and some non-negative $i \leq \ell_x - 1$. By Exercise 10.1.18 the sheaf

$$\mathcal{F} \otimes \mathcal{L}^{\otimes m} \cong (\mathcal{F} \otimes \mathcal{L}^{\otimes m_x+i}) \otimes (\mathcal{L}^{\otimes \ell_x})^{\otimes q}$$

is globally generated at every point of W_x.

Since X is quasicompact closed points are dense in X. Again applying quasicompactness, we can choose a finite set of points x_i such that the corresponding sets W_{x_i} cover X. If we set m_0 to be the supremum of the various m_{x_i} we see that $\mathcal{F} \otimes \mathcal{L}^{\otimes m}$ is globally generated for every $m \geq m_0$. This shows that $\mathcal{L}$ is ample. $\square$

12.4.2 Kodaira vanishing

Suppose that X is a smooth projective variety over a Noetherian ring A and that $\mathcal{L}$ is an ample invertible sheaf on X. Theorem 12.4.1 guarantees that the higher cohomology groups $\check{H}^i(X, \mathcal{L}^{\otimes m}) = 0$ vanish when m is sufficiently large. It is natural to ask for an explicit bound on m which guarantees the vanishing of the higher cohomology. As discussed in Remark 12.4.4 such a bound is vital for explicit computations.

Siu's version of Matsusaka's Big Theorem identifies an explicit bound for m in Theorem 12.4.1 using the intersection numbers of $\mathcal{L}$ and of the canonical bundle ω_X. Unfortunately the bound is enormous (for example it is doubly exponential in $\dim(X)$) and thus hard to apply in practice.

Instead of focusing on the constant m, it is usually better to use the following foundational theorem to give explicit bounds on vanishing of cohomology.

Theorem 12.4.2 (Kodaira's Theorem)**.** *Let $\mathbb{K}$ be a field of characteristic 0 and let X be a smooth projective $\mathbb{K}$-scheme. Then for any ample invertible sheaf $\mathcal{L}$ and any $i > 0$ we have*

$$\check{H}^i(X, \omega_X \otimes \mathcal{L}) = 0.$$

There are also vanishing theorems for cohomology which use other sheaves of differential forms Ω^p_X, but Kodaira's Theorem is the most famous and the most frequently used. The assumption on the characteristic is essential – the theorem does not hold in characteristic p – but some of the ideas can be extended to arbitrary characteristic.

Kodaira's Theorem indicates that the theory of ample invertible sheaves works best when we include a "twist" by the canonical line bundle. In fact, it turns out that many features of the geometry of X are controlled by the positivity of the canonical line bundle. For example, if X is a Fano variety (i.e. a variety such that ω_X is antiample) then Kodaira's theorem shows the vanishing of the higher cohomology groups of every ample line bundle on X.

12.4.3 Euler characteristic

In most cohomology theories we obtain an interesting invariant by taking an alternating sum of the dimensions of the cohomology groups. For example, when X is a manifold then the alternating sum of its cohomology groups is known as the Euler characteristic. The Euler characteristic of a manifold satisfies many nice properties that the cohomology groups do not: for example, it is additive for inclusion/exclusion and is multiplicative for fibrations.

The corresponding construction for the Čech cohomology groups is also known as the Euler characteristic.

Definition 12.4.3. Let X be a projective $\mathbb{K}$-scheme and let $\mathcal{F}$ be a coherent sheaf on X. We define the Euler characteristic of $\mathcal{F}$ to be

$$\chi_X(\mathcal{F}) := \sum_{i \geq 0} (-1)^i \dim_{\mathbb{K}} \check{H}^i(X, \mathcal{F}).$$

By dimensional vanishing (Theorem 12.2.5) only finitely many cohomology groups are non-zero and by Theorem 12.3.3 the dimension of each $\check{H}^i$ will be finite. Thus we see that $\chi_X(\mathcal{F})$ will always be an integer.

Remark 12.4.4. Just like its counterpart for manifolds the Euler characteristic of a coherent sheaf is "topological" in nature, depending only on the Chern classes of the sheaf $\mathcal{F}$. (This principle is made precise by the Hirzebruch-Riemann-Roch theorem; we will see an instance of this principle in Theorem 12.5.5.)

Remember, one of our motivating problems is to the computation of the space of global sections of a coherent sheaf. Although this problem is challenging, it is much easier to compute the Euler characteristic. In the best situations the higher cohomology will vanish, in which case the two constructions coincide and we can compute global sections using "topological techniques."

Exercise 12.4.5. Compute the Euler characteristic of $\mathcal{O}(m)$ on $\mathbb{P}^n_{\mathbb{K}}$.

Exercise 12.4.6. Let X be a projective $\mathbb{K}$-scheme and let $0 \to \mathcal{F} \to \mathcal{G} \to \mathcal{H} \to 0$ be an exact sequence of coherent sheaves on X. Prove that

$$\chi_X(\mathcal{G}) = \chi_X(\mathcal{F}) + \chi_X(\mathcal{H}).$$

12.4.4 Exercises

Exercise 12.4.7. Let X be a projective scheme over a Noetherian ring A and let $\mathcal{L}$ be an invertible sheaf on X. Prove that $\mathcal{L}$ is ample if and only if $\mathcal{L}|_{X_{red}}$ is ample. (Hint: for the reverse implication, suppose that X_{red} is defined by the ideal sheaf $\mathcal{I}$. For a coherent sheaf $\mathcal{F}$ on X and consider the filtration

$$0 = \mathcal{I}^r\mathcal{F} \subset \mathcal{I}^{r-1}\mathcal{F} \subset \ldots \subset \mathcal{I}\mathcal{F} \subset \mathcal{F}$$

Show that each quotient $\mathcal{I}^n\mathcal{F}/\mathcal{I}^{n-1}\mathcal{F}$ can be interpreted as a sheaf on X_{red} and thus is a suitable candidate for Serre Vanishing.)

Exercise 12.4.8. Let X be a projective scheme over a Noetherian ring A and let $\mathcal{L}$ be an invertible sheaf on X. Prove that $\mathcal{L}$ is ample if and only if the restriction of $\mathcal{L}$ to each component of X is ample. (Hint: it may be convenient to apply the previous exercise to assume that X is reduced. Then use a similar trick as in the hint to that exercise.)

12.5 Cohomology of line bundles on curves

As discussed in Section 10.1, we can study the morphisms from a projective scheme X by classifying the invertible sheaves on X and then studying their sections. In general the problem of identifying the space of global sections for an invertible sheaf is quite difficult. In this section we discuss this problem for curves.

12.5.1 Riemann-Roch for curves

For the rest of this section our standing assumption is that C is a projective, smooth, geometrically integral curve over a field $\mathbb{K}$. We will condense these assumptions by saying that C is a "good" curve. In Section 10.5 we analyzed the Picard group of a good curve C; the main tool was the degree homomorphism $\deg : \mathrm{Pic}(C) \to \mathbb{Z}$.

In this section we will systematically analyze the Čech homology groups of invertible sheaves on C. We will need one key fact:

Fact 12.5.1 (Serre duality)**.** Let C be a good curve over a field $\mathbb{K}$ and let $\mathcal{L}$ be an invertible sheaf on X. We have a duality

$$\check{H}^1(C, \mathcal{L})^\vee \cong \check{H}^0(C, \Omega_C \otimes \mathcal{L}^\vee).$$

This fact gives a geometric interpretation to the higher cohomology groups of invertible sheaves.

We now can give two new definitions of the genus:

Definition 12.5.2. Let C be a good curve. The geometric genus of C is defined to be $\check{H}^0(C, \Omega_{C/\mathbb{K}})$. The arithmetic genus of C is defined to be $1 - \chi(\mathcal{O}_C)$.

Remark 12.5.3. Both definitions make sense in more general situations (including for higher dimensional varieties). The arithmetic genus is defined for singular curves in exactly the same way. In contrast, the geometric genus is usually defined by first replacing C by its normalization.

Applying Serre duality, we quickly see:

Corollary 12.5.4. *Let C be a good curve over a field $\mathbb{K}$. Then the arithmetic genus and geometric genus coincide and are both equal to* $\dim \check{H}^1(C, \mathcal{O}_C)$.

Proof. Let g denote the geometric genus of C. Then $\dim H^1(C, \mathcal{O}_C) = g$ by Serre duality. Furthermore since C is geometrically integral we see that $\dim H^0(C, \mathcal{O}_C) = 1$. This implies that the arithmetic genus is also g. □

The key theorem in this section is Riemann-Roch which relates the cohomology groups of $\mathcal{L}$ to its degree. As discussed earlier, this is the first instance of the general principle that the Euler characteristic is "topological" in nature.

Theorem 12.5.5 (Riemann-Roch)**.** *Let C be a good curve of genus g over a field $\mathbb{K}$. Then for any invertible sheaf $\mathcal{L}$ on C we have*

$$\begin{aligned}\chi_C(\mathcal{L}) &= \deg(\mathcal{L}) + \chi_C(\mathcal{O}_C)\\ &= \deg(\mathcal{L}) + 1 - g.\end{aligned}$$

This statement is the "easy" version of Riemann-Roch. Sometimes people combine Serre duality with Riemann-Roch to obtain the "hard" version of Riemann-Roch.

Proof. For any closed point p we have an exact sequence

$$0 \to I_p \to \mathcal{O}_C \to i_*\mathcal{O}_p \to 0.$$

Note that the leftmost sheaf is isomorphic to $\mathcal{O}_C(-p)$ since p is a Cartier divisor on C. Tensoring by an invertible sheaf $\mathcal{T}$, we obtain

$$0 \to \mathcal{T}(-p) \to \mathcal{T} \to i_*\mathcal{O}_p \otimes \mathcal{T} \to 0.$$

Since the rightmost term is supported on a single point, we have

$$\begin{aligned}\chi_C(i_*\mathcal{O}_p \otimes \mathcal{T}) &= \dim_{\mathbb{K}} \check{H}^0(C, i_*\mathcal{O}_p \otimes \mathcal{T})\\ &= \dim_{\mathbb{K}} \check{H}^0(p, \mathcal{T}|_p) = \dim_{\mathbb{K}} \check{H}^0(p, \mathcal{O}_p)\end{aligned}$$

Using additivity of the Euler characteristic we conclude that $\chi_C(\mathcal{T}) = \chi_C(\mathcal{T}(-p)) + \deg(p)$.

Now suppose that $\mathcal{L}$ has a rational section $\sum a_i p_i$. By applying the formula above repeatedly while we add on the closed points one-by-one to $\mathcal{O}_C$, we see that $\chi_C(\mathcal{L}) = \chi_C(\mathcal{O}_C) + \sum a_i \deg(p_i)$. But the latter quantity is how we defined $\deg(\mathcal{L})$. □

Finally, we are able to compare the new definitions of genus in Definition 12.5.2 with our original definition in Definition 11.6.1.

Corollary 12.5.6. *Let C be a good curve over a field $\mathbb{K}$ of geometric genus g. Then $\deg(\Omega^1_C) = 2g - 2$.*

Proof. By definition $\check{H}^0(C, \Omega^1_C) = g$ and by Serre duality $\check{H}^1(C, \Omega^1_C) = 1$. We then apply Theorem 12.5.5. □

12.5.2 Cohomology groups of invertible sheaves

Riemann-Roch is most useful when we can understand either $\check{H}^0$ or $\check{H}^1$.

Proposition 12.5.7. *Let C be a good curve over a field $\mathbb{K}$ and let $\mathcal{L}$ be an invertible sheaf on C.*

(1) If $\deg(\mathcal{L}) < 0$ then $\check{H}^0(C, \mathcal{L}) = 0$.

(2) *If* $\deg(\mathcal{L}) > 2g-2$ *then* $\check{H}^1(C,\mathcal{L}) = 0$ *and* $\dim \check{H}^0(C,\mathcal{L}) = \deg(\mathcal{L}) + 1 - g$.

Proof. (1) If $\mathcal{L}$ has a global section, then this section defines an effective Cartier divisor D such that $\mathcal{O}_C(D) \cong \mathcal{L}$. Since there is no effective divisor with negative degree, we deduce that no invertible sheaf of negative degree can have a section. (2) follows from (1) by Serre duality. □

In particular, in these degree ranges the behavior of cohomology is entirely determined by the degree. This is a very useful property!

$\deg(\mathcal{L})$	< 0	0	1	$\dots$	$2g-3$	$2g-2$	$> 2g-2$
$\dim \check{H}^0$	0	?	?	$\dots$	?	?	$\deg(\mathcal{L}) + 1 - g$
$\dim \check{H}^1$	$g - 1 - \deg(\mathcal{L})$	?	?	$\dots$	?	?	0

The cohomology groups in the intermediate range are more subtle. In this range, the cohomology groups will depend upon the isomorphism type of the bundle (and not just the degree). The first instance of this principle can be seen in the "boundary" cases:

Proposition 12.5.8. *Let C be a good curve over a field $\mathbb{K}$.*

(1) *If* $\deg(\mathcal{L}) = 0$ *then* $\check{H}^0(C,\mathcal{L}) = 0$ *unless* $\mathcal{L} \cong \mathcal{O}_C$ *(in which case* $\check{H}^0(C,\mathcal{L}) \cong \mathbb{K}$*).*

(2) *If* $\deg(\mathcal{L}) = 2g-2$ *then* $\dim \check{H}^0(C,\mathcal{L}) = g - 1$ *unless* $\mathcal{L} \cong \Omega^1_C$ *(in which case* $\dim \check{H}^0(C,\mathcal{L}) = g$*).*

Proof. (1) Suppose that $\mathcal{L}$ has a global section. This defines a non-zero morphism $\phi : \mathcal{O}_X \to \mathcal{L}$. This morphism is necessarily injective: since $\mathcal{L}$ is torsion-free the image must be a locally free sheaf of rank 1. This implies that the kernel of ϕ is a torsion sheaf, and since $\mathcal{O}_X$ is torsion-free we conclude that the kernel is 0. Let $\mathcal{K}$ denote the cokernel of $\mathcal{L}$. Then we have

$$\check{H}^0(C,\mathcal{K}) = \chi_C(\mathcal{K}) = \chi_C(\mathcal{L}) - \chi(\mathcal{O}_C) = \deg(\mathcal{L}) - \deg(\mathcal{O}_X) = 0.$$

This means that $\mathcal{K} = 0$ and ϕ must be a bijection. (2) follows from (1) by Serre duality. □

We could continue to discuss the next "boundary" cases explicitly. For example:

Exercise 12.5.9. Let C be a good curve over a field and suppose that $\mathcal{L}$ is an invertible sheaf on C of degree 1. Prove that $\dim \check{H}^0(C,\mathcal{L}) \leq 2$ and that equality is attained if and only if $C \cong \mathbb{P}^1$ and $\mathcal{L} \cong \mathcal{O}(1)$.

12.5.3 Twisting by a fixed point

Suppose that C admits a $\mathbb{K}$-point p. Let $\mathrm{Pic}^d(C)$ denote the set of line bundles of degree d; note that each $\mathrm{Pic}^d(C)$ is a coset of the subgroup $\mathrm{Pic}^0(C)$. Then $\cdot p$ defines a bijection $\mathrm{Pic}^d(C) \to \mathrm{Pic}^{d+1}(C)$.

Here is our plan: suppose we start with an invertible sheaf $\mathcal{L} \in \mathrm{Pic}^0(C)$. We can then study the behavior of the cohomology groups of $\mathcal{L}(kp)$ as we increase the coefficient k. Note that as we vary $\mathcal{L} \in \mathrm{Pic}^0(C)$ the set of twists will exhaust all possible invertible sheaves on C. Thus if we can understand this operation well enough we will in obtain a somewhat complete picture of the cohomology groups of invertible sheaves on C. (Note however that our sequences will depend very heavily on the choice of p – there certainly is not a "canonical" choice!)

Exercise 12.5.10. Suppose that C is a good curve with a $\mathbb{K}$-point p. Let $\mathcal{T}$ be any invertible sheaf on p. Show that either

- $\dim \check{H}^0(C, \mathcal{T}(p)) = \dim \check{H}^0(C, \mathcal{T}) + 1$, or
- $\dim \check{H}^1(C, \mathcal{T}(p)) = \dim \check{H}^1(C, \mathcal{T}) - 1$.

The previous exercise shows that the sequence of cohomology groups of $\mathcal{L}(kp)$ is very simple: at each step either $\dim \check{H}^0$ increases by 1 or $\dim \check{H}^1$ decreases by 1. It turns out that for a general invertible sheaf $\mathcal{L} \in \mathrm{Pic}^0(C)$ we expect to get the simplest sequence possible:

k	0	1	2	$\dots$	$g-2$	$g-1$	g	$\dots$	$2g-3$	$2g-2$
$\dim \check{H}^0(C, \mathcal{L}(kp))$	0	0	0	$\dots$	0	0	1	$\dots$	$g-2$	$g-1$
$\dim \check{H}^1(C, \mathcal{L}(kp))$	$g-1$	$g-2$	$g-3$	$\dots$	1	0	0	$\dots$	0	0

For non-general invertible sheaves $\mathcal{L} \in \mathrm{Pic}^0(C)$ it is possible for both $\check{H}^0$ and $\check{H}^1$ to be larger than this "expected value" by the same amount. The behavior can be a bit subtle, depending on the bundle $\mathcal{L}$ in interesting ways.

Example 12.5.11. Let $\mathbb{K}$ be an algebraically closed field. We say that a good curve C over $\mathbb{K}$ is hyperelliptic if C admits a degree 2 finite dominant morphism $f : C \to \mathbb{P}^1$. (This condition is usually studied only when the genus of C is ≥ 2.)

If C is hyperelliptic, then $\mathcal{L} := f^*\mathcal{O}(1)$ is a degree 2 invertible sheaf on C (sometimes called the "hyperelliptic line bundle"). It is clear that $\mathcal{L}$ is basepoint free and by Exercise 12.5.14 we have $\dim \check{H}^0(C, \mathcal{L}) \geq 2$. In fact, we claim that if the genus is ≥ 1 then we must have $\dim \check{H}^0(C, \mathcal{L}) = 2$. Indeed, since $\mathcal{L}$ is basepoint free, for any point $p \in C$ we have $\dim \check{H}^0(C, \mathcal{L} - p) = \dim \check{H}^0(C, \mathcal{L}) - 1$ and we conclude by Exercise 12.5.9. (Note that the dimension of the space of global sections "differs" from the expected value for a general degree 2 line bundle on C.)

Suppose now that C is hyperelliptic of genus ≥ 2. Since $\mathcal{L}$ is basepoint free, for any positive integer r we have

$$\begin{aligned}\dim \check{H}^0(C, \mathcal{L}^{\otimes r}) &\geq \dim \operatorname{Sym}^r(\check{H}^0(C, \mathcal{L})) \\ &= r + 1\end{aligned}$$

Applying this result with $r = g - 1$, Proposition 12.5.8 shows that $\mathcal{L}^{\otimes g-1} \cong \Omega^1_C$. Furthermore, since in this case we have equality in the equation above we see that $\check{H}^0(C, \Omega^1_C) \cong \operatorname{Sym}^{g-1}(\check{H}^0(C, \mathcal{L}))$ and thus the sections of Ω^1_C define the map which is the composition of $f : C \to \mathbb{P}^1$ with the $(g-1)$-Veronese embedding $\mathbb{P}^1 \to \mathbb{P}^{g-1}$.

Since every section of $\check{H}^0(C, \mathcal{L}^{\otimes g-1})$ comes from a symmetric product of $\check{H}^0(C, \mathcal{L})$, the same must be true for all smaller tensor powers so that we have an equality $\dim \check{H}^0(C, \mathcal{L}^{\otimes r}) = r + 1$ for any $1 \leq r \leq g - 1$.

In fact, Clifford's theorem shows that the "largest possible" space of sections occur for tensor products of the hyperelliptic line bundle as in Example 12.5.11.

Theorem 12.5.12 (Clifford's Theorem)**.** *Let C be a good curve over a field $\mathbb{K}$. If $\mathcal{T}$ is an invertible sheaf on a curve C such that $\check{H}^1(C, \mathcal{T}) > 0$ then*

$$\dim \check{H}^0(C, \mathcal{T}) \leq \frac{\deg(\mathcal{T})}{2} + 1$$

with equality iff $\mathcal{T} = \mathcal{O}_C$, Ω^1_C, or C is hyperelliptic and $\mathcal{T}$ is a positive tensor power of the hyperelliptic line bundle $\mathcal{L}$.

12.5.4 Exercises

Exercise 12.5.13. Let C be a good curve over an algebraically closed field. Prove that if $\mathcal{L}$ is an invertible sheaf on C and $\deg(\mathcal{L}) \geq 2g$ then $\mathcal{L}$ is basepoint free.

Exercise 12.5.14. Let C, Z be good curves over a field $\mathbb{K}$ and suppose that $f : C \to Z$ is a finite morphism. Let $\mathcal{L}$ be an invertible sheaf on Z. Show that we have $\dim \check{H}^0(C, f^*\mathcal{L}) \geq \dim \check{H}^0(Z, \mathcal{L})$. Give an example where a strict inequality is achieved.

Exercise 12.5.15. Let C be a good curve of genus 0 over a field $\mathbb{K}$.

(1) Show that $C \cong \mathbb{P}^1_{\mathbb{K}}$ if and only if C has a $\mathbb{K}$-point.

(2) Show that there is a closed embedding of C into $\mathbb{P}^2$ as a conic curve. In particular show that C has a point of degree ≤ 2.

Exercise 12.5.16. Let C be a good curve of genus 2 over an algebraically closed field.

(1) Show that Ω^1_C is basepoint free but not very ample. Since $H^0(C, \Omega^1_C) = 2$, we obtain a morphism $f : C \to \mathbb{P}^1$.

(2) Show that the morphism f is a finite morphism of degree 2. In particular, every genus 2 curve is hyperelliptic.

Exercise 12.5.17. Let C be a good curve of genus 3 over an algebraically closed field. Suppose that C is not hyperelliptic. Then show that C admits a closed embedding into $\mathbb{P}^2$ as a degree 4 curve.

12.6 Ample line bundles on curves

In this section our goal is to classify the ample line bundles on a smooth projective curve. Along the way we will develop a new criterion for showing that a line bundle is very ample.

12.6.1 Closed embedding criterion

Definition 12.6.1. Let $\mathbb{K}$ be an algebraically closed field. Suppose that $f : X \to Y$ is a morphism of $\mathbb{K}$-schemes. We say that:

(1) f is injective on points, if f is set-theoretically injective on closed points.

(2) f is injective on tangent vectors, if for any closed point $p \in X$ the induced map of Zariski tangent spaces $T_{X,p} \to T_{Y,f(p)}$ is injective. (Note that since $\mathbb{K}$ is algebraically closed we do indeed have a morphism of Zariski tangent spaces.)

The following theorem describes one of the most useful tests for a morphism to be a closed embedding over an algebraically closed field $\mathbb{K}$. It is essentially a repackaging of Theorem 4.2.12 which shows that a finite morphism whose fibers are defined by Artinian rings of dimension 1 will be a closed embedding.

Theorem 12.6.2. *Let $\mathbb{K}$ be an algebraically closed field. Suppose that $f : X \to Y$ is a projective morphism of $\mathbb{K}$-schemes that is injective on points and injective on tangent vectors. Then f is a closed embedding.*

The projective hypothesis is necessary, e.g. to rule out open embeddings. The assumption on the ground field is also necessary: it is only when $\mathbb{K}$ is algebraically closed that the "injective on tangent vectors" condition guarantees that the fibers have degree 1. (If we tried to weaken the hypothesis to allow arbitrary fields, the morphism $\mathrm{Spec}(\mathbb{C}) \to \mathrm{Spec}(\mathbb{R})$ would give a counterexample.)

Proof. Since closed embeddings are local on the target, it suffices to consider the case when Y is affine. We know that every non-empty fiber of f over a closed point of Y has dimension 0. Since fiber dimension is upper semicontinuous, we deduce that every non-empty fiber of f has dimension 0. Then f is a projective morphism with finite fibers, hence a finite morphism by Corollary 12.3.6.

We claim that the the degree of f over every closed point is either 0 or 1. Indeed, suppose that $q \in Y$ is a (reduced) closed point with non-empty fiber. Since f is injective on points, the fiber over q consists of a unique point p. Since f is injective on tangent spaces, the map $\mathfrak{m}_q/\mathfrak{m}_q^2 \to \mathfrak{m}_f/\mathfrak{m}_f^2$ is surjective. By Nakayama's lemma we deduce that the image of the map $\mathcal{O}_{Y,q} \to \mathcal{O}_{X,p}$ contains the maximal ideal. Since $\mathbb{K}$ is algebraically closed the residue field of both p and q is isomorphic to $\mathbb{K}$, and thus we deduce that $\mathcal{O}_{Y,q} \to \mathcal{O}_{X,p}$ is surjective. By tensoring we see that the map $\mathcal{O}_{Y,q}/\mathfrak{m}_q \to \mathcal{O}_{X,p}/\mathfrak{m}_q$ is also surjective, showing that this latter ring has dimension 1 over the residue field.

Since the degree of the fiber is an upper semicontinuous function, we conclude that the degree of f over every point is either 0 or 1. We conclude that f is a closed embedding by Theorem 4.2.12. $\square$

We now translate Theorem 12.6.2 into a criterion for very ampleness.

Definition 12.6.3. Let $\mathbb{K}$ be an algebraically closed field. Let X be a projective $\mathbb{K}$-scheme and let $\mathcal{L}$ be an invertible sheaf on X. Let $V \subset \check{H}^0(X, \mathcal{L})$ be a subspace. We say that:

(1) V separates points, if for any two distinct closed points $p, q \in X$ there is a Cartier divisor D parametrized by V such that $p \in \operatorname{Supp}(D)$ but $q \notin \operatorname{Supp}(D)$.

(2) V separates tangent vectors, if for any closed point $p \in X$ and any vectors t in the Zariski tangent space $T_{X,p}$ there is a Cartier divisor D parametrized by V such that $p \in \operatorname{Supp}(D)$ but t is not in the codimension 1 subspace of $T_{X,p}$ cut out by D.

Corollary 12.6.4. *Let X be an algebraically closed field. Let X be a projective $\mathbb{K}$-scheme and let $\mathcal{L}$ be an invertible sheaf on X. If sections of $\mathcal{L}$ separate points and tangent vectors then $\mathcal{L}$ is very ample.*

Proof. Let $f : X \dashrightarrow \mathbb{P}^N$ be the morphism defined by sections of $\mathcal{L}$. Since the sections of $\mathcal{L}$ separate points, we see that $\mathcal{L}$ is globally generated and the morphism f is injective. We claim that f is also injective on tangent vectors. Indeed, suppose we fix a point $p \in X$ and consider the linear subspace W of $\check{H}^0(X, \mathcal{L})$ parametrizing Cartier divisors containing p. We can then identify $T_{\mathbb{P}^N, f(p)}$ as the set of codimension 1 subspaces of W. Then the map of Zariski cotangent spaces $T_{X,p} \to T_{\mathbb{P}^N, f(p)}$ is defined by sending a tangent vector t to the subspace W_t parametrizing Cartier divisors which contain p and whose tangent planes contain t. Since sections of $\mathcal{L}$ separate tangent vectors we see that the kernel of this map is 0. $\square$

Often it is useful to reformulate Definition 12.6.3 as follows.

Corollary 12.6.5. *Let X be an algebraically closed field. Let X be a projective $\mathbb{K}$ scheme and let $\mathcal{L}$ be an invertible sheaf on X. Suppose that for every closed subscheme $Z \subset X$ which has dimension 0 and length 2 we have that the induced map*

$$\check{H}^0(X, \mathcal{L}) \to \check{H}^0(Z, \mathcal{L}|_Z)$$

is surjective. Then $\mathcal{L}$ is very ample.

Proof. First suppose that Z is the union of two distinct closed points. Applying the hypothesis to Z, we see that sections of $\mathcal{L}$ separate points. Next fix a tangent vector t at a closed point p and let Z denote the image of the morphism $\operatorname{Spec}(\mathbb{K}[t]/t^2) \to X$ which defines this tangent vector. Applying the hypothesis to Z, we see that sections of $\mathcal{L}$ separate tangent vectors. $\square$

12.6.2 Ample invertible sheaves on curves

We can now classify ample invertible sheaf on curves. We start with a special case:

Lemma 12.6.6. *Let C be a smooth projective integral curve over an algebraically closed field $\mathbb{K}$ and let $\mathcal{L}$ be an invertible sheaf on C. Then $\mathcal{L}$ is ample if and only if* $\deg(\mathcal{L}) > 0$.

Proof. If $\mathcal{L}$ is ample, then for some positive integer r the sheaf $\mathcal{L}^{\otimes r}$ must admit global sections. By Proposition 12.5.7 we see that $0 < \deg(\mathcal{L}^{\otimes r}) < r \deg(\mathcal{L})$.

Conversely, it suffices to show that any invertible sheaf $\mathcal{T}$ of degree $\geq 2g+1$ is very ample. We will apply Corollary 12.6.5. First choose two different points $p, q \in C$. Then we have an exact sequence

$$0 \to \mathcal{T}(-p-q) \to \mathcal{T} \to (i_*\mathcal{O}_p \oplus i_*\mathcal{O}_q) \otimes \mathcal{T} \to 0.$$

Since both p and q must be $\mathbb{K}$-points (as $\mathbb{K}$ is algebraically closed), the degree of the leftmost term is at least $2g-1$. By Proposition 12.5.7 we have $\check{H}^1(C, \mathcal{T}(-p-q)) = 0$. This means that the map on the right induces a surjective map of global sections so that sections of $\mathcal{T}$ separate points. Consider now the sequence

$$0 \to \mathcal{T}(-2p) \to \mathcal{T} \to (i_*\mathcal{O}_{2p}) \otimes \mathcal{T} \to 0.$$

Note that $i_*\mathcal{O}_{2p}$ represents the structure sheaf of a non-reduced scheme at p which records the data of a tangent vector at p. Furthermore since $\mathcal{T}$ is invertible tensoring by $\mathcal{T}$ does not change the isomorphism type of the underlying module. Repeating the argument from above, the rightmost map induces a surjection on global sections. Thus sections of $\mathcal{T}$ separate tangent vectors. □

We deduce the statement over an arbitrary field using a base change argument.

Theorem 12.6.7. *Let C be a smooth projective geometrically integral curve over a field $\mathbb{K}$ and let $\mathcal{L}$ be an invertible sheaf on C. Then $\mathcal{L}$ is ample if and only if* $\deg(\mathcal{L}) > 0$.

Proof. If $\mathcal{L}$ is ample, then for some positive integer r the sheaf $\mathcal{L}^{\otimes r}$ must admit global sections and as before Proposition 12.5.7 implies that $\deg(\mathcal{L}) > 0$.

Conversely, it suffices to show that any invertible sheaf $\mathcal{T}$ of degree $\geq 2g+1$ is very ample. By Lemma 12.6.6 we know that the base change $\mathcal{T}_{\overline{\mathbb{K}}}$ is very ample. Exercise 9.1.24 shows that $H^0(C, \mathcal{T}) \otimes \overline{\mathbb{K}} \cong H^0(C_{\overline{\mathbb{K}}}, \mathcal{T}_{\overline{\mathbb{K}}})$. In particular, this implies that $\mathcal{T}$ is basepoint free if and only if $\mathcal{T}_{\overline{\mathbb{K}}}$ is basepoint free. Furthermore, since $\overline{f} : C_{\overline{\mathbb{K}}} \to \mathbb{P}^n_{\overline{\mathbb{K}}}$ is injective we see that $f : C \to \mathbb{P}^n_{\mathbb{K}}$ is also injective. Finally, since $\overline{\mathbb{K}}$ is a flat $\mathbb{K}$-module we also see that the surjectivity of the sheaf map $\overline{f}^\sharp$ implies the surjectivity of $f^\sharp$. Altogether we see that since $\overline{f}$ is a closed embedding, the map f is as well. □

Example 12.6.8. Let C be a curve of genus 2 over an algebraically closed field and let p, q, r be general points of C. Then $\mathcal{O}_C(p+q-r)$ is an ample invertible sheaf with no sections. Indeed, it is ample since it has positive degree. To see that it has no sections, we need to use the Jacobian. The degree 1 line bundles which admit sections will lie in the image of the map $C \to \mathrm{Jac}(C)$ sending $p \mapsto \mathcal{O}_C(p)$. In particular, since $\dim \mathrm{Jac}(C) = 2$, a general invertible sheaf of degree 1 has no sections. On the other hand, every degree 2 line bundle can be written as $\mathcal{O}(p+q)$ for some choices $p, q \in C$. By generality of p, q, r we obtain the desired statement.

12.6.3 Exercises

Exercise 12.6.9. Let X be an smooth $\mathbb{K}$-variety of dimension 1. Suppose that X is not proper over $\mathbb{K}$. Prove that X is affine. (Hint: embed X in a regular projective curve $\overline{X}$. Identify a very ample divisor on $\overline{X}$ whose support is exactly equal to the complement $\overline{X} \backslash X$.

12.7 Hilbert polynomials

Using the Euler characteristic, we can obtain a general definition of the Hilbert polynomial (extending the definition of Section 6.1).

Definition 12.7.1. Let X be projective $\mathbb{K}$-scheme equipped with a very ample divisor $\mathcal{L}$. Let $\mathcal{F}$ be a coherent sheaf on X. We define the Hilbert polynomial of $\mathcal{F}$ with respect to $\mathcal{L}$ to be the function $P_{\mathcal{F}} : \mathbb{Z} \to \mathbb{Z}$ which sends $d \mapsto \chi_X(\mathcal{F} \otimes \mathcal{L}^{\otimes d})$.

The Hilbert polynomial for X is defined to be the Hilbert polynomial of $\mathcal{O}_X$. We will denote it by P_X.

Since the choice of a very ample divisor $\mathcal{L}$ is equivalent to the choice of a closed embedding $X \hookrightarrow \mathbb{P}^n_{\mathbb{K}}$ (up to change of coordinates), it is also quite common to discuss the Hilbert polynomials of coherent sheaves on closed subschemes of $\mathbb{P}^n_{\mathbb{K}}$.

The following result justifies our nomenclature for $P_{\mathcal{F}}$.

Proposition 12.7.2. *Let X be a projective $\mathbb{K}$ scheme equipped with a very ample divisor $\mathcal{L}$ and let $\mathcal{F}$ be a coherent sheaf on X. Then the function $P_{\mathcal{F}}$ in Definition 12.7.1 is a polynomial function.*

Proof. By pushing forward $\mathcal{F}$ under a closed immersion $X \hookrightarrow \mathbb{P}^n_{\mathbb{K}}$ we reduce to the case where $\mathcal{F}$ is a coherent sheaf on $\mathbb{P}^n_{\mathbb{K}}$. By performing a base change to an algebraically closed field, we may also assume that $\mathbb{K}$ is infinite. In particular, this implies that there is a hyperplane H that doesn't contain any of the (finitely many) associated points of $\mathcal{F}$. Let f be the function defining H.

Tensor the exact sequence

$$0 \to \mathcal{O}(-1) \to \mathcal{O}_{\mathbb{P}^n} \to i_*\mathcal{O}_H \to 0$$

by $\mathcal{F}$ to obtain the exact sequence

$$\mathcal{F}(-1) \to \mathcal{F} \to i_*\mathcal{F}|_H \to 0.$$

We claim that this sequence is exact on the left as well. Indeed, suppose we fix an open affine U and let M denote the module defining $\mathcal{F}$ on U. By assumption the vanishing locus of $f|_U$ does not contain any associated point of M. Equivalently, this means that f is not a zero divisor on M so that $M \xrightarrow{\cdot f} M$ is injective.

Note that $\mathcal{O}(m)$ is flat and that $\mathcal{O}_m|_H$ is identified with $\mathcal{O}(m)$ on the smaller projective space. Thus for every m we have an exact sequence

$$0 \to \mathcal{F}(m-1) \to \mathcal{F}(m) \to i_*(\mathcal{F}|_H(m)) \to 0.$$

Thus we have $P_{\mathcal{F}}(m) - P_{\mathcal{F}}(m-1) = P_{\mathcal{F}|_H}(m)$. By induction on dimension the rightmost term is a polynomial. Since a function whose difference equation is polynomial will itself be polynomial, we conclude the desired statement. $\square$

Our main theorem in this section gives a direct relationship between flatness and the Hilbert polynomial.

Theorem 12.7.3. *Let $f : X \to Y$ be a projective morphism with Y locally Noetherian. Suppose that $\mathcal{F}$ is a coherent sheaf on X. If $\mathcal{F}$ is flat over Y, then the Hilbert polynomial of the fibers $P_{\mathcal{F}|_{X_y}}(d)$ is locally constant. If Y is reduced, then the converse is also true.*

As usual, we need a reducedness assumption when we hope to deduce something about a sheaf from a property of its fibers.

Proof. Since the statement is local on Y, we may assume $Y = \operatorname{Spec}(S)$ is affine and S is a Noetherian ring. In this case f factors through a closed embedding $g : X \to \mathbb{P}^n_S$. Since $\mathcal{F}$ is flat over Y if and only if $g_*\mathcal{F}$ is flat over Y and since Exercise 12.2.9 shows that the Hilbert polynomial can be computed after taking a closed embedding, we reduce to the case where $X = \mathbb{P}^n_S$.

Claim 12.7.4. $\mathcal{F}$ is flat over Y if and only if for m sufficiently large the S-module $\check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m))$ is locally free.

Proof of claim: To prove the forward implication, choose m sufficiently large so that the higher cohomology of $\mathcal{F}(m)$ on $\mathbb{P}^n_S$ vanishes. Consider the Čech complex $\check{C}^\bullet(\mathfrak{U}, \mathcal{F}(m))$ for $\mathcal{F}(m)$ with respect to the standard affine cover of $\mathbb{P}^n_S$. Since $\mathcal{F}$ is flat over $\operatorname{Spec}(S)$, each term in the sequence will be a flat S-module. Furthermore, since the cohomology vanishes this sequence is exact except at the 0th term where the kernel is $\check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m))$. This implies that $\check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m))$ is also flat. In particular, since it is a flat coherent sheaf on $\operatorname{Spec}(S)$ it is locally free.

To prove the reverse implication, choose m_0 sufficiently large so that $\check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m))$ is locally free. Let M be the graded $S[x_0, \ldots, x_n]$-module

$$\bigoplus_{m \geq m_0} \check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m))$$

As a direct sum of flat S-modules, M is flat over A. Furthermore by Theorem 9.6.19 we have that $\widetilde{M^+} \cong \mathcal{F}$. In particular, since flatness is preserved by localization, we see that all the stalks of $\mathcal{F}$ are flat over the corresponding localizations of S. We conclude that $\mathcal{F}$ is flat over Y. $\square$

Since Y is affine, the S-module $\check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m))$ defines the sheaf $f_*\mathcal{F}(m)$. The next step is to compare the stalks $f_*\mathcal{F}(m)_y$ with the fiberwise sections $f_{\mathbb{P}^n_y}|_*(\mathcal{F}|_{\mathbb{P}^n_y})$. (Recall that in general there is no easy comparison between these two objects; this is a very special feature of our specific situation.)

Claim 12.7.5. Fix a point $y \in Y$. Then for sufficiently large m the rank of $\check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m))$ at y is the same as $P_{\mathcal{F}|_{\mathbb{P}^n_y}}(m)$.

Proof. We first prove this when y is a closed point. Since S is Noetherian we have an exact sequence

$$\mathcal{O}_Y^{\oplus r} \to \mathcal{O}_Y \to \kappa(y) \to 0. \tag{12.7.1}$$

This sequence remains exact upon tensoring by $f_*\mathcal{F}$ (since $\otimes$ is right exact). We now apply Serre vanishing to the kernel of $\mathcal{O}_Y \to \kappa(y)$ to see that we get an exact sequence of global sections

$$\check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m)^{\oplus r}) \to \check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m)) \to \check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m)|_{\mathbb{P}^n_y}) \to 0$$

for m sufficiently large. On the other hand, if we tensor Equation 12.7.1 by $\check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m))$ we get

$$\check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m))^{\oplus r} \to \check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m)) \to \check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m)) \otimes \kappa(y) \to 0. \tag{12.7.2}$$

Comparing, we see that for m sufficiently large the fiber of $\check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m))$ is isomorphic to $\check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m)|_{\mathbb{P}^n_y})$. In turn, this latter group has dimension $P_{\mathcal{F}|_{\mathbb{P}^n_y}}(m)$ for m sufficiently large by Serre vanishing.

For arbitrary points y, we first base change over $\mathcal{O}_{Y,y}$. Since this is a flat $\mathcal{O}_Y$-module, the base change preserves exactness of the Čech complex and thus we have the identification

$$\check{H}^0(\mathbb{P}^n_S, \mathcal{F}(m)) \otimes \mathcal{O}_{Y,y} \cong \check{H}^0(\mathbb{P}^n_{\mathcal{O}_{Y,y}}, \mathcal{F}(m)|_{\mathrm{Spec}(\mathcal{O}_{Y,y})})$$

(see Proposition 13.4.4 for more details). We then repeat the argument above. □

By combining the two claims, we immediately see that if $\mathcal{F}$ is flat then the Hilbert polynomial is constant. Conversely, if the Hilbert polynomial is constant then the second claim shows that $f_*\mathcal{F}$ has constant rank. Since Y is reduced, Theorem 9.5.5 shows that $f_*\mathcal{F}$ is locally free and then the first claim shows that $\mathcal{F}$ is flat. □

In particular, this shows:

Corollary 12.7.6. *Let Y be a locally Noetherian scheme. Suppose that X is a closed subscheme of $\mathbb{P}^n_Y$ and consider the projection map $f : X \to Y$. If f is flat, then the dimension and the degree of the fibers X_y is locally constant over Y.*

Proof. As discussed in Chapter 6 both properties of X_y can be detected from the Hilbert polynomial. □

We also point out the following important special case:

Theorem 12.7.7. *Let $f : X \to Y$ be a projective morphism with Y locally Noetherian. Suppose that $\mathcal{F}$ is a coherent sheaf on X that is flat over Y. Then as we vary $y \in Y$ the Euler characteristic of the fibers*

$$\chi_{X_y}(\mathcal{F}|_{X_y}) = \sum_{i \geq 0} (-1)^i \dim \check{H}^i(X_y, \mathcal{F}|_{X_y})$$

is a locally constant function.

12.7.1 Exercises

Exercise 12.7.8. Let $\mathbb{K}$ be an algebraically closed field. Compute the Hilbert polynomials of:

(1) A smooth degree d hypersurface in $\mathbb{P}^n$.

(2) The dth Veronese embedding of $\mathbb{P}^m$.

Exercise 12.7.9. Show that if $X \subset Y$ are two closed subschemes of $\mathbb{P}^n$ then we have $P_X(d) \leq P_Y(d)$ for every sufficiently large integer d.

Exercise 12.7.10. Let $R = \mathbb{K}[t, t^{-1}]$ and consider the closed subscheme $X \subset \mathbb{P}^3_R$ defined by the ideal
$$I = (t^2(xw + w^2) - z^2, tx(x + w) - yz, xz - twy)$$
This is a flat family whose fibers are twisted cubics in $\mathbb{P}^3$. We are interested in taking a flat limit as $t \to 0$.

(1) Show that the limit is set-theoretically contained in the plane $z = 0$.

(2) Compute the Hilbert polynomial of a twisted cubic and the Hilbert polynomial for a nodal plane curve. Show that their Hilbert polynomials are not the same.

(3) Show that the flat limit is defined by the ideal
$$(z^2, yz, xz, y^2w - x^2(x + w)).$$
In other words, the flat limit is a non-reduced planar curve with an associated point at the node whose non-reduced structure points "out of the plane." (By the previous step, we could have predicted the existence of the non-reduced structure from a Hilbert polynomial calculation.)

Chapter 13

Derived functors

Up to this point we have defined many left-exact functors F on the category of $\mathcal{O}_X$-modules (such as the global sections functor, the pushforward, etc.). In this chapter we will study the associated right derived functors: a sequence of functors R^iF which measure the "failure" of F to be exact.

Since the construction of right derived functors is purely formal, our first task is usually to recast these functors in more concrete language. We will ask similar questions in each case: do these right derived functors preserve quasicoherence? What happens if we apply them to an affine scheme?

We will also give two types of applications. First, our results will clarify the geometric behavior of the pushforward functor f_*. More precisely, we will discuss two important theorems – the Cohomology and Base Change theorem and the Theorem on Formal Functions – which sometimes allow us to understand the stalks and fibers of $f_*\mathcal{F}$ more explicitly. When these theorems apply, we can get a better handle on the (somewhat mysterious) sheaf $f_*\mathcal{F}$.

The second application is Serre Duality. One of the fundamental properties of Hodge groups is the duality $H^{p,q}(X) \cong H^{n-p,n-q}(X)^\vee$ obtained via integration of differential forms. Serre Duality translates this duality into the setting of algebraic geometry. Just as in the complex setting, the key player is the canonical bundle $\omega_X = \bigwedge^{\dim X} \Omega^1_X$ and its homology groups.

13.0.1 Algebraic preliminaries

We will briefly review the foundations of derived functors.

Injective objects

Definition 13.0.1. Let $\mathbf{C}$ be an abelian category. An object $I \in \mathbf{C}$ is said to be injective if for every monomorphism $f : X \hookrightarrow Y$ and for every map $g : X \to I$ there admits a "lift" $\widetilde{g} : Y \to I$ fitting into a commutative diagram

$$\begin{tikzcd} X \arrow[r, hook, "f"] \arrow[d, "g"] & Y \arrow[dl, dashed, "\widetilde{g}"] \\ I & \end{tikzcd}$$

Equivalently, I is injective if the functor $\mathrm{Hom}(-, I)$ is exact.

Example 13.0.2. The injective objects in $\mathbf{Ab}$ are the divisible abelian groups.

Injective resolutions

Let $\mathbf{C}$ be an abelian category. An injective resolution of an object $A \in \mathbf{C}$ is an exact sequence

$$0 \to A \to I^0 \xrightarrow{d^0} I^1 \xrightarrow{d^1} I^2 \xrightarrow{d^2} \dots$$

where each I^j is an injective object of $\mathbf{C}$.

Remark 13.0.3. There is a dual theory of projective resolutions of an object A (which extend to the left from A). We won't discuss it here.

Definition 13.0.4. Let $\mathbf{C}$ be an abelian category. We say that $\mathbf{C}$ has enough injectives if every object $A \in \mathbf{C}$ admits a monomorphism into an injective object I.

It is clear that if $\mathbf{C}$ has enough injectives, then every object admits an injective resolution. In fact, even more is true:

Lemma 13.0.5. *Let $\mathbf{C}$ be an abelian category.*

(1) Suppose we have objects $A, B \in \mathbf{C}$ with an arbitrary resolution $0 \to A \to Q^\bullet$ and an injective resolution $0 \to B \to I^\bullet$. Given a morphism $f : A \to B$, there is a lift to a map of complexes $f^\bullet : Q^\bullet \to I^\bullet$ making a commutative diagram

$$\begin{tikzcd} 0 \arrow[r] & A \arrow[r] \arrow[d, "f"] & Q^0 \arrow[r, "d^0"] \arrow[d, "f^0"] & Q^1 \arrow[r, "d^1"] \arrow[d, "f^1"] & Q^2 \arrow[r, "d^2"] \arrow[d, "f^2"] & \dots \\ 0 \arrow[r] & B \arrow[r] & I^0 \arrow[r, "d^0"] & I^1 \arrow[r, "d^1"] & I^2 \arrow[r, "d^2"] & \dots \end{tikzcd}$$

Furthermore, this map of complexes $f^\bullet$ is unique up to chain homotopy.

(2) *Suppose we have an exact sequence* $0 \to A_1 \to A_2 \to A_3 \to 0$ *in* $\mathbf{C}$ *and injective resolutions* $g_1 : A_1 \to I_1^\bullet$ *and* $g_3 : A_3 \to I_3^\bullet$. *Then there is an injective resolution* $g_2 : A_2 \to I_2^\bullet$ *and a commutative diagram with exact rows*

$$\begin{tikzcd} 0 \arrow[r] & A_1 \arrow[r] \arrow[d, "g_1"] & A_2 \arrow[r] \arrow[d, "g_2"] & A_3 \arrow[r] \arrow[d, "g_3"] & 0 \\ 0 \arrow[r] & I_1^\bullet \arrow[r] & I_2^\bullet \arrow[r] & I_3^\bullet \arrow[r] & 0 \end{tikzcd}$$

Right derived functors

Let $\mathbf{C}, \mathbf{D}$ be abelian categories and suppose that $F : \mathbf{C} \to \mathbf{D}$ is a left-exact functor, i.e. for any exact sequence $0 \to A_1 \to A_2 \to A_3 \to 0$ in $\mathbf{C}$ the sequence

$$0 \to F(A_1) \to F(A_2) \to F(A_3)$$

is exact. We can extend this sequence to the right using the right derived functors.

Definition 13.0.6. Let $F : \mathbf{C} \to \mathbf{D}$ be a left-exact functor. Suppose that $\mathbf{C}$ has enough injectives. For each $i \geq 0$ we define the right-derived functor R^iF as follows.

Given an object $A \in \mathbf{C}$, let

$$0 \to A \to I^0 \xrightarrow{d^0} I^1 \xrightarrow{d^1} I^2 \xrightarrow{d^2} \dots$$

be an injective resolution. Then we define $R^iF(A)$ to be the ith cohomology group of the sequence

$$0 \to F(A) \to F(I^0) \xrightarrow{F(d^0)} F(I^1) \xrightarrow{F(d^1)} F(I^2) \xrightarrow{F(d^2)} \dots$$

One can show that the cohomology groups of this complex do not depend on the choice of injective resolution by applying Lemma 13.0.5.(1) several times to $A \xrightarrow{id} A$.

Given a morphism $f : A \to B$ in $\mathbf{C}$, we get an induced morphism $R^iF(A) \to R^iF(B)$ from Lemma 13.0.5.(1).

It is clear that $R^0F = F$. The key property of the right derived functors is that a short exact sequence $0 \to A_1 \to A_2 \to A_3 \to 0$ in A yields a long exact sequence

$$\begin{aligned} 0 &\to F(A_1) \to F(A_2) \to F(A_3) \xrightarrow{\delta_0} \\ &\to R^1F(A_1) \to R^1F(A_2) \to R^1F(A_3) \xrightarrow{\delta_1} \\ &\to R^2F(A_1) \to \dots \end{aligned}$$

In fact, the following result shows that this property can be used to characterize derived functors.

Definition 13.0.7. Let $F : \mathbf{C} \to \mathbf{D}$ be a left exact functor between abelian categories. Suppose we have a sequence of functors $G^i : \mathbf{C} \to \mathbf{D}$ for $i \geq 1$ which satisfy the following properties:

(1) The G^i form a δ-functor: for any short exact sequence $0 \to A_1 \to A_2 \to A_3 \to 0$ in $\mathbf{C}$, there are maps δ_j for $j \geq 0$ which form a long exact sequence

$$\begin{aligned} 0 \to F(A_1) \to F(A_2) \to F(A_3) &\xrightarrow{\delta_0} \\ \to G^1(A_1) \to G^1(A_2) \to G^1(A_3) &\xrightarrow{\delta_1} \\ \to G^2(A_1) \to \dots & \end{aligned}$$

Furthermore, any morphism of SES in $\mathbf{C}$ (that is, morphisms between corresponding entries which form a commuting diagram) induces a morphism of the above LES in $\mathbf{D}$.

(2) The functor G^i is effaceable for every $i > 0$: or every object M of $\mathbf{C}$, there is a monomorphism $\phi : M \to N$ to an object N of $\mathbf{C}$ that satisfies $G^i(N) = 0$.

We say that the G^i form an effaceable δ-functor for F.

Theorem 13.0.8. *Let $F : \mathbf{C} \to \mathbf{D}$ be a left exact functor between abelian categories. Suppose that G^i, H^i are effaceable δ-functors for F. Then G^i and H^i are naturally isomorphic.*

In particular, suppose that $\mathbf{C}$ has enough injectives. Then any effaceable δ-functor is isomorphic to the right derived functors for F.

Acyclic objects

Definition 13.0.9. Let $\mathbf{C}, \mathcal{D}$ be abelian categories and let $F : \mathbf{C} \to \mathbf{D}$ be a left exact functor. Suppose that $\mathcal{C}$ has enough injectives. We say that an object $A \in \mathbf{C}$ is acyclic for F if we have $R^iF(A) = 0$ for all $i > 0$.

One can use acyclic objects in the place of injective objects when computing right derived functors.

Proposition 13.0.10. *Let $F : \mathbf{C} \to \mathbf{D}$ be a left exact functor. Suppose that $\mathbf{C}$ has enough injectives. Given an object $A \in \mathbf{C}$, suppose we have an exact sequence*

$$0 \to A \to J^0 \to J^1 \to \dots$$

where J^i is F-acyclic for every $i \geq 0$. Then there is a natural isomorphism $R^iF(A) \cong H^i(F(J^\bullet))$.

Proof. The argument is the same as the proof of Corollary 12.1.8. $\square$

13.0.2 Spectral sequences

We will avoid spectral sequences for the most part. However, there is one key result: spectral sequences naturally arise from composition of left-exact functors.

Theorem 13.0.11 (Grothendieck spectral sequence)**.** *Let $F : \mathbf{C} \to \mathbf{D}$ and $G : \mathbf{D} \to \mathbf{E}$ be left-exact functors between abelian categories. Suppose that $\mathbf{C}$ and $\mathbf{D}$ have enough injectives. Suppose further that for any injective object I of $\mathbf{C}$ the image $F(I)$ in $\mathbf{D}$ is acyclic for G. Then for every object $A \in \mathbf{C}$ there is a spectral sequence with E_2-page*

$$E_2^{p,q} = R^q G(R^p F(A))$$

that converges to $R^{p+q}(G \circ F)(A)$.

13.1 Injective and projective sheaves

Suppose that $\mathbf{C}, \mathbf{D}$ are abelian categories. When constructing the derived functors of a left exact (resp. right exact) functor $F : \mathbf{C} \to \mathbf{D}$, the key issue is whether the category $\mathbf{A}$ has enough injectives (resp. projectives).

This issue is a little delicate when $\mathbf{C}$ is a category of sheaves on a scheme X. There are many possible categories we could work in – abelian sheaves, $\mathcal{O}_X$-mod, quasicoherent mod, coherent mod – and the injective/projective objects will behave a little differently in each category. Generally speaking, the "looser" the category the easier it should be to construct enough injectives/projectives, but the further away we are from the objects we really care about. In particular, it is a priori possible that the same functor (e.g. the global sections functor) could yield different derived functors as we vary the underlying category.

13.1.1 Enough injectives

Categories of sheaves often have enough injectives.

Theorem 13.1.1. *Let $(X, \mathcal{O}_X)$ be a ringed topological space. Then the category of $\mathcal{O}_X$-modules has enough injectives.*

Applying the result with $\mathcal{O}_X = \mathbb{Z}_X$ we see that in particular the category of sheaves of abelian groups on X has enough injectives. The proof is very similar to the construction given in Exercise 7.4.10.

Proof. Let $\mathcal{F}$ be an $\mathcal{O}_X$-module. For every point x the category of $\mathcal{O}_{X,x}$-modules has enough injectives, so there is an monomorphism $\mathcal{F}_x \to I_x$ to an injective $\mathcal{O}_{X,x}$-module. We let $\mathcal{I}_x$ denote the skyscraper sheaf with value I_x concentrated at the point x.

We claim that $\mathcal{I}_x$ is an injective $\mathcal{O}_X$-module. Indeed, for any $\mathcal{O}_X$-module $\mathcal{G}$ there is a bijection

$$\mathrm{Hom}_{\mathcal{O}_X}(\mathcal{G}, \mathcal{I}_x) \leftrightarrow \mathrm{Hom}_{\mathcal{O}_{X,x}}(\mathcal{G}_x, I_x).$$

Using the lifting property for I_x we obtain the lifting property for $\mathcal{I}_x$.

Note that $\mathcal{F}$ admits a monomorphism into the sheaf defined by $U \mapsto \prod_{x \in U} \mathcal{F}_x$. The sheaf in turn injects into the product sheaf $\prod_{x \in X} \mathcal{I}_x$. Since a product of injective objects is injective, this gives us the desired injection. □

It is also true that the category of quasicoherent sheaves has enough injectives.

Theorem 13.1.2. *Let X be a scheme. The category of quasicoherent sheaves on X has enough injectives.*

When X is a Noetherian scheme, one can use basically the same argument as in Theorem 13.1.1. Namely, given a quasicoherent sheaf $\mathcal{F}$ we fix a cover of open affines $\{U_i\}$, choose a monomorphism from $\mathcal{F}(U_i)$ into an injective $\mathcal{O}_X(U_i)$-module I_i, and then use the sheaf

$\oplus_i f_{i*}(\widetilde{I_i})$. However in general this approach does not work and the theorem is more difficult to prove.

Remark 13.1.3. One possible approach for constructing injective objects in $\mathbf{QCoh}(X)$ is to show that for any injective $\mathcal{O}_X$-module I the corresponding sheaf $\widetilde{I}$ is an injective object in the category of $\mathcal{O}_X$-modules. This approach works well when R is a Noetherian ring (see [Har66, Chapter II Corollary 7.14]). However, for an arbitrary ring R and injective R-module I the sheaf $\widetilde{I}$ may fail to be either injective or flasque. (This is closely related to the fact that the injectiveness of a module need not be preserved under localization for non-Noetherian rings.)

It is important to note that the category of coherent modules almost never has enough injectives. This is true even for affine schemes: if R is an integral domain that is not a field then the only finitely generated injective R-module is the 0 module. By similar logic, injective objects in the categories $\mathbf{QCoh}(X)$ and $\mathcal{O}_X-\mathbf{Mod}$ are almost never finitely generated.

13.1.2 Enough projectives

It is quite rare for a category of sheaves to have enough projectives. (This is unfortunate, since projective objects in the category of R-modules are much easier to understand than injective objects.)

Example 13.1.4. Let us show that the category of coherent sheaves on $\mathbb{P}^1_{\mathbb{K}}$ does not have enough projectives.

Suppose we had a non-zero projective object $\mathcal{P}$ in the category of quasicoherent sheaves. By Corollary 10.6.13 there is a surjection from a direct sum of line bundles onto $\mathcal{F}$:

$$\mathcal{O}(d)^{\oplus r} \to \mathcal{F}.$$

Note that the map $\mathcal{O}(d-1) \oplus \mathcal{O}(d-1) \to \mathcal{O}(d)$ which is multiplication by x_0 on the first factor and multiplication by x_1 on the second is a surjection. Replacing each instance of $\mathcal{O}(d)$ by $\mathcal{O}(d-1)^{\oplus 2}$, we find a surjection

$$\mathcal{O}(d-1)^{\oplus 2r} \xrightarrow{\phi} \mathcal{F}$$

We claim that ϕ does not admit a splitting. Indeed, there is no non-zero morphism $\mathcal{O}(d) \to \mathcal{O}(d-1)$ and since $\mathcal{O}(d)^{\oplus r}$ surjects onto $\mathcal{F}$ we see that $\mathcal{F}$ also does not admit a non-zero morphism to $\mathcal{O}(d-1)$.

In fact, the category of sheaves of abelian groups on a locally connected topological space has enough projectives if and only if every intersection of open sets in X is open (see [Cla]). This condition will almost never hold for schemes.

13.1.3 Comparing functors

As discussed in the introduction to the chapter, one can also compute derived functors using acyclic objects. There is one type of acyclic object (for various functors) which is used quite frequently.

Definition 13.1.5. Let X be a topological space and let $\mathcal{F}$ be a sheaf of abelian groups on X. We say that $\mathcal{F}$ is flasque (or flabby) if for every inclusion of open sets $V \subset U$ the restriction map $\mathcal{F}(U) \to \mathcal{F}(V)$ is surjective.

The key advantage of this definition is that the notion of "flasque" is independent of the underlying category. In particular, flasque sheaves can be useful for comparing the different versions of a right derived functor arising from changing the underlying category of sheaves. On the other hand, it is rare for flasque sheaves to show up in geometric situations – for example, coherent sheaves are rarely flasque – so this notion is mainly useful as a tool for proving theoretical statements.

In view of later applications, we prove a few basic properties of flasque sheaves.

Lemma 13.1.6. *Let X be a topological space. Suppose that*

$$0 \to \mathcal{F} \xrightarrow{\phi} \mathcal{G} \xrightarrow{\psi} \mathcal{H} \to 0$$

is an exact sequence of sheaves and that $\mathcal{F}$ is flasque. Then the sequence of global sections

$$0 \to \mathcal{F}(X) \to \mathcal{G}(X) \to \mathcal{H}(X) \to 0$$

is exact.

Proof. It suffices to show that $\psi(X) : \mathcal{G}(X) \to \mathcal{H}(X)$ is surjective. Fix a section $t \in \mathcal{H}(X)$. Consider the set of all pairs (U, s) such that U is an open subset of X and $\psi(U)(s) = t|_U$. Since t admits local lifts, this is a non-empty set.

We define a partial ordering on these pairs by setting $(U, s) \leq (U', s')$ if $U \subset U'$ and $s'|_U = s$. We claim that any chain of pairs

$$(U_1, s_1) \leq (U_2, s_2) \leq (U_3, s_3) \leq \dots$$

has an upper bound. Indeed, since the sections s_i agree on overlaps, we can glue them to define a section s_V on the open set $V = \cup_i U_i$, and it is clear that $\psi(V)(s_V) = t|_V$. By Zorn's lemma, we conclude that there is a maximal element $(\widetilde{U}, \widetilde{s})$ in the set of pairs.

We claim that $\widetilde{U} = X$. Suppose for a contradiction that $\widetilde{U} \subsetneq X$. Since ψ is surjective, we can find a pair (V, r) such that $V \not\subset \widetilde{U}$. Since $\widetilde{s}|_{V\cap\widetilde{U}} - r|_{V\cap\widetilde{U}}$ is in the kernel of $\psi(V \cap \widetilde{U})$, it is an element of $\mathcal{F}(V \cap \widetilde{U})$. Using the fact that $\mathcal{F}$ is flasque, we can find an element $u \in \mathcal{F}(V)$ whose restriction to $V \cap \widetilde{U}$ is this difference. Then the section $r + u$ on V can be glued to the section $\widetilde{s}$ on $\widetilde{U}$. Since the ψ-image of the resulting section on $V \cup \widetilde{U}$ is the restriction of t, this contradicts the maximality of the pair $(\widetilde{U}, \widetilde{s})$. $\square$

The following proposition shows that there are many flasque sheaves on a topological space.

Proposition 13.1.7. *Let X be a scheme. Suppose that $\mathcal{I}$ is an injective object in the category $\mathcal{O}_X$*$-$**Mod**. *Then $\mathcal{I}$ is flasque.*

The same argument shows that injective sheaves on any topological space are flasque. However, the statement does not hold for quasicoherent sheaves: an injective object in $\mathbf{QCoh}(X)$ need not be flasque in general.

Proof. Suppose we have inclusion of open sets $V \subset U$. Let i_V (respectively i_U) denote the inclusion from V (respectively U) into X. We then have an injection

$$i_{V!}\mathcal{O}_V \to i_{U!}\mathcal{O}_U$$

of $\mathcal{O}_X$-modules (where ! denotes the extension-by-zero functor). Since $\mathcal{I}$ is injective, we get a surjection

$$\mathrm{Hom}(i_{U!}\mathcal{O}_U, \mathcal{I}) \to \mathrm{Hom}(i_{V!}\mathcal{O}_V, \mathcal{I}).$$

However, we can identify

$$\mathrm{Hom}(i_{U!}\mathcal{O}_U, \mathcal{I}) \cong \mathrm{Hom}(\mathcal{O}_U, f^{-1}\mathcal{I}) \cong \mathcal{I}(U)$$

and similarly for V. We conclude there is a surjection $\mathcal{I}(U) \to \mathcal{I}(V)$ as desired. $\square$

As a demonstration of the utility of flasque sheaves, we will analyze the behavior of the right derived functors of the global section functor in several different categories.

Definition 13.1.8. Let X be a topological space. We define $H^i(X, -)$ to be the ith right derived functor of the global sections functor $\Gamma(X, -) : \mathbf{Sh}(X) \to \mathbf{Ab}$.

Note that we define sheaf cohomology in this way even when $\mathcal{F}$ carries additional structure, e.g. when $\mathcal{F}$ is a quasicoherent sheaf. We next show that flasque sheaves of abelian groups are acyclic for the global sections functor.

Lemma 13.1.9. *Let X be a topological space and suppose that $\mathcal{F}$ is a flasque sheaf on X. Then $H^i(X, \mathcal{F}) = 0$ for every $i > 0$.*

Proof. Since the category $\mathbf{Sh}(X)$ has enough injectives, we can find an injection $\phi : \mathcal{F} \to \mathcal{I}$ into an injective sheaf in $\mathbf{Sh}(X)$. Let $\mathcal{G}$ denote the cokernel of ϕ so that we have an exact sequence

$$0 \to \mathcal{F} \to \mathcal{I} \to \mathcal{G} \to 0.$$

Since $\mathcal{F}$ is flasque by assumption and $\mathcal{I}$ is flasque by Proposition 13.1.7, Exercise 13.1.12 shows that $\mathcal{G}$ is flasque as well.

We now prove the statement by induction on i. To prove the base case $i = 1$, we note that $H^1(X, \mathcal{I}) = 0$ (since $\mathcal{I}$ is injective) and the map $H^0(X, \mathcal{I}) \to H^0(X, \mathcal{G})$ is surjective by Lemma 13.1.6. For the induction step, consider the LES of cohomology corresponding to the exact sequence above. Since $\mathcal{I}$ is acyclic, we have $H^i(X, \mathcal{F}) \cong H^{i-1}(X, \mathcal{G})$ for every $i > 1$. Since $\mathcal{G}$ is also flasque, we conclude by induction. $\square$

Altogether, we obtain a fundamental comparison between right-derived functors in different categories:

Theorem 13.1.10. *Let X be a scheme. Then the right derived functors of $\Gamma(X, -)$: $\mathcal{O}_X$−**Mod** $\to$ **Ab** coincide with the functors $H^i(X, -)$.*

Proof. The action of the right derived functors for $\Gamma(X, -) : \mathcal{O}_X$−**Mod** $\to$ **Ab** on an $\mathcal{O}_X$-module $\mathcal{F}$ can be computed by taking an injective resolution $0 \to \mathcal{F} \to \mathcal{I}^\bullet$ in the category of $\mathcal{O}_X$-modules and taking the cohomology of the complex $\Gamma(X, \mathcal{I}^\bullet)$. Since every term of $\mathcal{I}^\bullet$ is flasque, by Lemma 13.1.9 we can compute $H^i(X, \mathcal{F})$ in the same way. $\square$

Warning 13.1.11. It is not in general true that the right derived functors of $\Gamma(X, -)$: **QCoh**$(X) \to$ **Ab** agree with the functors $H^i(X, -)$.

However, when X is Noetherian any injective object in **QCoh**(X) is flasque (see [Har77, Corollary III.3.6] and [Har77, Exercise III.3.6]). Since flasque sheaves are acyclic for every variant of the global sections functor, by repeating the argument of Theorem 13.1.10 we see that when X is a Noetherian scheme the sheaf cohomology of a quasicoherent sheaf agrees with the right derived functors of $\Gamma(X, -)$: **QCoh**$(X) \to$ **Ab**.

13.1.4 Exercises

Exercise 13.1.12. Let X be a topological space. Suppose that $0 \to \mathcal{F} \to \mathcal{G} \to \mathcal{H} \to 0$ is an exact sequence of sheaves. Prove that if $\mathcal{F}, \mathcal{G}$ are flasque then $\mathcal{H}$ is flasque as well.

Exercise 13.1.13. Show that if X is a topological space and $\mathcal{F}$ is a flasque sheaf on X then $\check{H}^i(X, \mathcal{F}) = 0$ for every $i > 0$.

Exercise 13.1.14. Let X be a scheme. Suppose that $\mathcal{I}$ is an injective object in $\mathcal{O}_X$−**Mod** (resp. **Sh**(X)). Show that for any open subset U the restriction $\mathcal{I}|_U$ is an injective object in $\mathcal{O}_U$−**Mod** (resp. **Sh**(U)). (Hint: given an inclusion of sheaves $\mathcal{F} \to \mathcal{G}$ in $\mathcal{O}_U$−**Mod**, consider their image in $\mathcal{O}_X$−**Mod** under the extension-by-zero functor.)

13.2 Global sections functor

Let X be a scheme. Recall that $H^i(X, -)$ denote the right derived functors of the global sections functor $\Gamma(X, -) : \mathbf{Sh}(X) \to \mathbf{Ab}$. While these functors satisfy nice formal properties, unfortunately they are difficult to compute.

The goal of this section is to compare the sheaf cohomology groups $H^i(X, -)$ with the Čech cohomology groups $\check{H}^i(X, -)$. While we cannot expect these two functors to agree in full generality – for example, we know that we don't always get a LES of Čech cohomology – we will show that these two functors agree for most "nice" schemes X. In particular, in "nice" situations we can use Čech cohomology techniques to compute the H^i.

13.2.1 Vanishing for affines

The key in the comparison is to prove the vanishing of the groups $H^i(X, -)$ when X is an affine scheme.

Theorem 13.2.1. *Let $X = \operatorname{Spec}(R)$ be an affine scheme and let $\mathcal{F}$ be a quasicoherent sheaf on X. Then $H^i(X, \mathcal{F}) = 0$ for all $i > 0$.*

Remark 13.2.2. There is a minor subtlety in the statement. We proved in Proposition 9.2.8 that when $X = \operatorname{Spec}(R)$ the functor $\Gamma(X, -) : \mathbf{QCoh}(X) \to R\text{−}\mathbf{Mod}$ is exact. However, since we cannot compute the functors H^i directly in the category $\mathbf{QCoh}(X)$ this fact does not suffice to prove the vanishing of the H^i. (When X is Noetherian, then one can appeal to this approach as discussed in Warning 13.1.11.)

Even though this exactness does not immediately prove Theorem 13.2.1, as we will see it lies at the heart of the proof. In particular, Lemma 13.2.3 can be seen as an analogue for the higher cohomology groups of the fact that we have "local lifts" of global sections for the rightmost term in a SES of quasicoherent sheaves.

Before proving Theorem 13.2.1, we will need the following lemma. Given a topological space X, an open set $h : V \hookrightarrow X$, and a sheaf of abelian groups $\mathcal{F}$ on X, we will denote by $\mathcal{F}_V$ the sheaf $h_* h^{-1}\mathcal{F}$.

Lemma 13.2.3. *Let X be a topological space and let $\mathcal{F}$ be a sheaf on X. Let $\mathcal{U}$ be a base for the topology on X. Suppose that there is a positive integer n such that $H^i(U, \mathcal{F}|_U) = 0$ for every open set U in $\mathcal{U}$ and every $0 < i < n$. Then given any element $\alpha \in H^n(X, \mathcal{F})$ there is an open cover $\mathfrak{U} = \{U_j\}_{j \in J}$ consisting of open sets in $\mathcal{U}$ such that the image of α in $H^n(X, \mathcal{F}_{U_j})$ is zero for every $j \in J$.*

Proof. By Proposition 13.1.7 there is an injection from $\mathcal{F}$ into a flasque sheaf $\mathcal{G}$. Let $\mathcal{H}$ be the cokernel. Note that the restriction of $\mathcal{G}$ to any open subset of X is still flasque. Since flasque sheaves are acyclic for the global sections functor, for any open set $U \subset X$ we have exact sequences

$$0 \to H^0(U, \mathcal{F}|_U) \to H^0(U, \mathcal{G}|_U) \to H^0(U, \mathcal{H}|_U) \to H^1(U, \mathcal{F}|_U) \to 0$$

and isomorphisms $H^i(U, \mathcal{H}|_U) \cong H^{i+1}(U, \mathcal{F}|_U)$.

We prove the statement by induction on n. The base case is when $n = 1$. Since inverse image is exact and pushforward is left-exact, for any open set V we have an inclusion $\mathcal{F}_V \to \mathcal{G}_V$. Denoting the cokernel by $\mathcal{K}^V$, we have a commuting diagram with exact rows:

$$\begin{tikzcd}
0 \arrow[r] & \mathcal{F} \arrow[r] \arrow[d] & \mathcal{G} \arrow[r] \arrow[d] & \mathcal{H} \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \mathcal{F}_V \arrow[r] & \mathcal{G}_V \arrow[r] & \mathcal{K}^V \arrow[r] & 0
\end{tikzcd}$$

Note that any $\alpha \in H^1(X, \mathcal{F})$ is the image of some $\beta \in H^0(X, \mathcal{H})$. Choose the open cover $\mathfrak{U} = \{U_j\}$ of open sets coming from $\mathcal{U}$ so that there are local sections $\gamma_j \in \mathcal{G}(U_j)$ which map to $\beta|_{U_j}$. Using the LES for the bottom row of the diagram above with $V = U_j$ we see that the image of α in $H^1(X, \mathcal{F}_{U_j})$ vanishes.

We next prove the induction step for $n > 1$. For any open subset V of X we have a commuting diagram with exact rows:

$$\begin{tikzcd}
0 \arrow[r] & \mathcal{F} \arrow[r] \arrow[d] & \mathcal{G} \arrow[r] \arrow[d] & \mathcal{H} \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \mathcal{F}_V \arrow[r] & \mathcal{G}_V \arrow[r] & \mathcal{H}_V &
\end{tikzcd}$$

We will need one basic observation:

> Let V be an element of $\mathcal{U}$. Then the morphism $\mathcal{G}_V \to \mathcal{H}_V$ is surjective.

Indeed, suppose that U is an open subset in $\mathcal{U}$ that is contained in V. By assumption $H^1(U, \mathcal{F}|_U) = 0$. Thus we have a surjection

$$H^0(U, \mathcal{G}_V|_U) \cong H^0(U, \mathcal{G}|_U) \to H^0(U, \mathcal{H}|_U) = H^0(U, \mathcal{H}_V|_U).$$

Since such U form a base for the topology on V, we obtain the desired surjection $\mathcal{G}_V \to \mathcal{H}_V$. In particular, since $\mathcal{G}_V$ is still flasque, we have exact sequences relating the sheaf cohomologies of $\mathcal{F}_V$ and $\mathcal{H}_V$.

Using these isomorphisms, we see that $\mathcal{H}$ satisfies the hypotheses for the integer $n-1$. Using the diagram

$$\begin{tikzcd}
H^{n-1}(X, \mathcal{H}) \arrow[r, "\cong"] \arrow[d] & H^n(X, \mathcal{F}) \arrow[d] \\
H^{n-1}(X, \mathcal{H}_V) \arrow[r, "\cong"] & H^n(X, \mathcal{F}_V)
\end{tikzcd}$$

for open sets V contained in $\mathcal{U}$ we see that the induction assumption applied to $\mathcal{H}$ yields the desired statement for $\mathcal{F}$. $\square$

Proof of Theorem 13.2.1: We prove the statement by induction on i. For the base case $i = 1$, since the category of $\mathcal{O}_X$-modules has enough injectives we can find an injection from $\mathcal{F}$ into a flasque $\mathcal{O}_X$-module $\mathcal{G}$. The cokernel $\mathcal{H}$ is also an $\mathcal{O}_X$-module. By Lemma 13.1.6 the map $H^0(X, \mathcal{G}) \to H^0(X, \mathcal{H})$ is surjective. Thus we have an injection $H^1(X, \mathcal{F}) \to H^1(X, \mathcal{G})$ and since the latter group vanishes by Lemma 13.1.9 the former group does as well.

For the induction step, let $\mathcal{U}$ be the basis of X consisting of distinguished open affines. By applying the induction hypothesis to each D_f, we see that the hypotheses of Lemma 13.2.3 are satisfied. Thus given any element $\alpha \in H^i(X, \mathcal{F})$ we can find a finite set of elements $\{f_j\}_{j=1}^m$ in R which generate the unit ideal and satisfy the following property:

> For every j define $\mathcal{F}_j$ to be the pushforward under the inclusion map $D_{f_j} \hookrightarrow X$ of $\mathcal{F}|_{U_i}$. Then the image of α in $H^i(X, \oplus_j \mathcal{F}_j) = 0$.

For a given $\alpha \in H^i(X, \mathcal{F})$, consider the corresponding SES of quasicoherent sheaves

$$0 \to \mathcal{F} \to \oplus_j \mathcal{F}_j \to \mathcal{K} \to 0.$$

By construction α must be in the image of the map $H^{i-1}(X, \mathcal{K}) \to H^i(X, \mathcal{F})$. But by induction the former group vanishes. We conclude that $H^i(X, \mathcal{F}) = 0$. $\square$

In fact, a famous result of Serre shows that the vanishing of higher cohomology is a distinguishing property of affine schemes:

Theorem 13.2.4 ([Sta15] Tag 01XF)**.** *Let X be a quasicompact scheme. Then the following conditions are equivalent:*

(1) X is affine.

(2) $H^i(X, \mathcal{F}) = 0$ for every quasicoherent sheaf $\mathcal{F}$ and every $i > 0$.

13.2.2 Comparing Čech cohomology and sheaf cohomology

We are now prepared to compare Čech cohomology and sheaf cohomology. First note that for any topological space X and any sheaf $\mathcal{F}$ we obtain a morphism $\check{H}^i(X, \mathcal{F}) \to H^i(X, \mathcal{F})$. Indeed, suppose we fix an open cover $\check{U}$ of X. As explained in the proof of Theorem 12.1.9, the Čech cohomology group $\check{H}^i(\mathfrak{U}, \mathcal{F})$ is computed by applying the global sections functor to the complex

$$0 \to \mathcal{F} \to \mathscr{C}^0(\mathfrak{U}, \mathcal{F}) \to \mathscr{C}^1(\mathfrak{U}, \mathcal{F}) \to \dots$$

and taking homology groups of the resulting complex of $\mathcal{O}_X(X)$-modules. If we fix any injective resolution $0 \to \mathcal{F} \to \mathcal{I}^\bullet$ of $\mathcal{F}$, then by Lemma 13.0.5.(1) there is a map of complexes $\mathscr{C}^\bullet(\mathfrak{U}, \mathcal{F}) \to \mathcal{I}^\bullet$ that induces an isomorphism on the 0th homology $\mathcal{F}$. Applying the global sections functor and taking homology, we get a map $\check{H}^i(\mathfrak{U}, \mathcal{F}) \to H^i(X, \mathcal{F})$.

Furthermore, given a refinement $\mathfrak{V}$ of $\mathfrak{U}$, if we choose a suitable injective resolution of $\mathcal{F}$ then the map of complexes assigned to $\mathfrak{U}$ factors through the map of complexes assigned to $\mathfrak{V}$. Thus we have a commuting diagram

$$\begin{tikzcd}
\check{H}^i(\mathfrak{U},\mathcal{F}) \arrow[rr] \arrow[dr] & & \check{H}^i(\mathfrak{V},\mathcal{F}) \arrow[dl] \\
& H^i(X,\mathcal{F}) &
\end{tikzcd}$$

Thus these morphisms descend to the direct limit to yield a morphism $\check{H}^i(X,\mathcal{F}) \to H^i(X,\mathcal{F})$.

Theorem 13.2.5. *Let X be a quasicompact separated scheme and let $\mathcal{F}$ be a quasicoherent sheaf on X. Then for every $i \geq 0$ we have*

$$\check{H}^i(X,\mathcal{F}) \cong H^i(X,\mathcal{F}).$$

In fact more is true – these identifications are functorial in $\mathcal{F}$ and compatible with LES – but we won't worry about these more precise statements.

Proof. Let $\mathfrak{U}$ be an open cover of X by open affines. By Corollary 12.2.3 we have $\check{H}^i(X,\mathcal{F}) = \check{H}^i(\mathfrak{U},\mathcal{F})$. Thus it suffices to show that $\check{H}^i(\mathfrak{U},\mathcal{F}) \cong H^i(X,\mathcal{F})$. More precisely, we will show:

Claim 13.2.6. Let $\mathcal{F}$ be any sheaf such that $H^i(V,\mathcal{F}|_V) = 0$ for every open set V obtained as a (non-empty) intersection of open sets in our cover. Then $\check{H}^i(\mathfrak{U},\mathcal{F}) \cong H^i(X,\mathcal{F})$

We show this claim by induction on i. The base case $i = 0$ has been discussed earlier. To prove the statement when $i = 1$, embed $\mathcal{F}$ into a flasque sheaf $\mathcal{G}$ and let $\mathcal{H}$ denote the cokernel. We have a morphism of exact sequences

$$\begin{tikzcd}
0 \arrow[r] & \check{H}^0(\mathfrak{U},\mathcal{F}) \arrow[r] \arrow[d, "\cong"] & \check{H}^0(\mathfrak{U},\mathcal{G}) \arrow[r] \arrow[d, "\cong"] & \check{H}^0(\mathfrak{U},\mathcal{H}) \arrow[r] \arrow[d, "\cong"] & \check{H}^1(\mathfrak{U},\mathcal{F}) \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & H^0(X,\mathcal{F}) \arrow[r] & H^0(X,\mathcal{G}) \arrow[r] & H^0(X,\mathcal{H}) \arrow[r] & H^1(X,\mathcal{F}) \arrow[r] & 0
\end{tikzcd}$$

which immediately yields the desired isomorphism.

Now we prove the induction step. Suppose that we know the claim in all degrees up to i for some $i \geq 1$. Embed $\mathcal{F}$ into a flasque sheaf $\mathcal{G}$ and let $\mathcal{H}$ denote the cokernel. For every open set V obtained as an intersection of the U_i, we have a SES

$$0 \to \mathcal{F}(V) \to \mathcal{G}(V) \to \mathcal{H}(V) \to 0$$

since by hypothesis $H^1(V,\mathcal{F}|_V) = 0$. This implies that the sequence of Čech complexes

$$0 \to \check{C}^\bullet(\mathfrak{U},\mathcal{F}) \to \check{C}^\bullet(\mathfrak{U},\mathcal{G}) \to \check{C}^\bullet(\mathfrak{U},\mathcal{H}) \to 0$$

is also exact. In particular, we get a long exact sequence of Čech cohomology associated to the open cover $\mathfrak{U}$ and the sheaves $\mathcal{F}$, $\mathcal{G}$, and $\mathcal{H}$.

Since $\mathcal{G}$ is flasque, by Exercise 13.1.13 the higher Čech cohomology of $\mathcal{G}$ vanishes. Thus we get for $i > 0$

$$\check{H}^i(\mathfrak{U}, \mathcal{H}) \cong \check{H}^{i+1}(\mathfrak{U}, \mathcal{F}).$$

Since the higher sheaf cohomology of a flasque sheaf also vanishes by Lemma 13.1.9, we also know that for $i > 0$

$$H^i(X, \mathcal{H}) \cong H^{i+1}(X, \mathcal{F}).$$

Note that since the restriction of $\mathcal{G}$ to any open set in X is still flasque, hence acyclic for the global sections functor by Lemma 13.1.9, it is clear that $\mathcal{H}$ also satisfies the property that $H^i(V, \mathcal{H}|_V) = 0$ for every open set V obtained as an intersection of open sets in our cover. Thus we can deduce the desired isomorphism $\check{H}^{i+1}(\mathfrak{U}, \mathcal{F}) \cong H^{i+1}(X, \mathcal{F})$ by applying our induction hypothesis to $\mathcal{H}$. □

Remark 13.2.7. More generally, there is a spectral sequence relating Čech cohomology and sheaf cohomology; see [Sta15, Tag 01ES].

We end this chapter by stating without proof a general theorem about sheaf cohomology. This result should be compared against the dimension-vanishing result for Čech cohomology in Theorem 12.2.5.

Theorem 13.2.8 ([Har77] Theorem III.2.7)**.** *Let X be a Noetherian topological space of dimension n. Then for all sheaves of abelian groups $\mathcal{F}$ and all $i > n$ we have $H^i(X, \mathcal{F}) = 0$.*

13.3 Higher direct images

Suppose that $f : X \to Y$ is a continuous map of topological spaces. Since the pushforward $f_* : \mathbf{Sh}(X) \to \mathbf{Sh}(X)$ is left exact, we can define its derived functors.

Definition 13.3.1. Let $f : X \to Y$ be a continuous map of topological spaces. The higher direct image functors $R^i f : \mathbf{Sh}(X) \to \mathbf{Sh}(Y)$ are the right derived functors of f_*.

Since right derived functors are somewhat abstract, it would be nice to have a more explicit description of these functors. Thus our first result shows that the higher direct image functors really are a "relative" version of sheaf cohomology.

Proposition 13.3.2. *Let $f : X \to Y$ be a continuous map of topological spaces and let $\mathcal{F}$ be a sheaf on X. For every $i \geq 0$, $R^i f_*(\mathcal{F})$ is isomorphic to the sheafification of the presheaf*

$$V \mapsto H^i(f^{-1}V, \mathcal{F}|_{f^{-1}V}).$$

Proof. We will denote by $S^i(\mathcal{F})$ the sheaf on Y constructed by the description above, i.e. the sheafification of the presheaf constructed from the ith sheaf cohomology groups.

Note that S^0 is the sheafification of the assignment $V \mapsto \mathcal{F}(f^{-1}V)$. Since this assignment already defines a sheaf, we have $S^0(\mathcal{F}) \cong f_*\mathcal{F}$. To show that the other functors agree with the right derived functors of f_*, we will use Theorem 13.0.8.

First, we show that the S^i define a δ-functor. Suppose we have an SES of sheaves

$$0 \to \mathcal{F}_1 \to \mathcal{F}_2 \to \mathcal{F}_3 \to 0.$$

Using the LES of cohomology, it is clear that the assignment $V \mapsto H^i(f^{-1}V, \mathcal{F}|_{f^{-1}V})$ yields a LES of presheaves. Since the sheafification functor is exact, we obtain a LES of shaves

$$\ldots \to S^i(\mathcal{F}_1) \to S^i(\mathcal{F}_2) \to S^i(\mathcal{F}_3) \to S^{i+1}(\mathcal{F}_1) \to \ldots$$

Furthermore, the compatibility with this LES with a morphism of SES of sheaves is a consequence of the functoriality of the LES for sheaf cohomology and the universal property of sheafification.

Second, we need to show that S^i are effaceable. Suppose that $\mathcal{I}$ is an injective object of $\mathbf{Sh}(X)$. By Exercise 13.1.14 the restriction $\mathcal{I}|_{f^{-1}V}$ is also injective for every open set $V \subset Y$. This means that $H^i(f^{-1}V, \mathcal{I}|_{f^{-1}V}) = 0$ and consequently $S^i(\mathcal{I})$ is the 0 sheaf. Since every object of $\mathbf{Sh}(X)$ admits an injection into an injective object, this proves effaceability. $\square$

Since the functors H^i for an $\mathcal{O}_X$-module can be computed in either $\mathbf{Sh}(X)$ or in $\mathcal{O}_X{-}\mathbf{Mod}$, we might expect that the functors $R^i f_*$ have the same property. As before, the key is to show that flasque sheaves are acyclic.

Corollary 13.3.3. *Let $f : X \to Y$ be a continuous map of topological spaces and let $\mathcal{F}$ be a flasque sheaf on X. Then $R^i f_*(\mathcal{F}) = 0$ for all $i > 0$.*

Proof. Follows from Proposition 13.3.2 and the fact that flasque sheaves are acyclic for sheaf cohomology. □

Corollary 13.3.4. *Let $f : X \to Y$ be a morphism of schemes. Then the restriction of the functors $R^i f_*$ to the category of $\mathcal{O}_X$-modules are the same as the right derived functors of $f_* : \mathcal{O}_X-\mathbf{Mod} \to \mathcal{O}_Y-\mathbf{Mod}$.*

Proof. Just as in the proof of Theorem 13.1.10, this follows from the fact that injective objects of $\mathcal{O}_X-\mathbf{Mod}$ are flasque, and thus acyclic for f_* by Corollary 13.3.3. □

Note that $R^i f_*$ need not be the right derived functor of f_* for the category of quasicoherent sheaves (since the H^i are also not right derived functors for the category of quasicoherent sheaves).

13.3.1 Quasicoherent sheaves

Suppose now that $f : X \to Y$ is a morphism of schemes and that $\mathcal{F}$ is a quasicoherent sheaf on X. It is natural to ask: is $R^i f_*\mathcal{F}$ a quasicoherent sheaf on Y? If so, is $R^i f_*\mathcal{F}(V)$ for open affines $V \subset Y$ defined by the cohomology group $H^i(f^{-1}V, \mathcal{F}|_{f^{-1}V})$? (Note that this is not guaranteed by Proposition 13.3.2 since the statement includes a sheafification.)

It turns out that if we impose some finiteness hypotheses on f then we can hope for these statements to be true. We will start with a lemma that proves the analogue of Proposition 13.3.2 for Čech cohomology.

Lemma 13.3.5. *Let $f : X \to Y$ be a quasicompact and separated morphism of schemes. For every $i \geq 0$, $R^i f_*(\mathcal{F})$ is isomorphic to the sheafification of the presheaf*

$$V \mapsto \check{H}^i(f^{-1}V, \mathcal{F}|_{f^{-1}V}).$$

Proof. Let $S^i_{pre}(\mathcal{F})$ denote the presheaf $V \mapsto H^i(f^{-1}V, \mathcal{F}|_{f^{-1}V})$ and let $S^i(\mathcal{F})$ denote its sheafification. We also denote by T^i_{pre} and T^i the corresponding constructions using Čech cohomology. We will consider these constructions as functors from $\mathbf{QCoh}(X) \to \mathcal{O}_Y-\mathbf{Mod}$. Using the Čech-to-sheaf cohomology map, we obtain a natural transformation $\Phi_{pre} : T^i_{pre} \to S^i_{pre}$. This also induces a natural transformation $\Phi : T^i \to S^i$.

For every open affine $V \subset Y$ the preimage $f^{-1}V$ is quasicompact and separated, thus for every quasicoherent sheaf $\mathcal{F}$ we have $\Phi_{pre}(\mathcal{F})(V) : T^i_{pre}(\mathcal{F})(V) \to S^i_{pre}(\mathcal{F})(V)$ is an isomorphism by Theorem 13.2.5. Since such open sets V form a base for the topology and since the sheafification functor is determined by its values on a base, we conclude that $\Phi(\mathcal{F}) : T^i(\mathcal{F}) \to S^i(\mathcal{F})$ is also an isomorphism. We conclude the desired statement by Proposition 13.3.2. □

The following result is the key result of the chapter.

Proposition 13.3.6. *Suppose that $f : X \to Y$ is a quasicompact and separated morphism and that $\mathcal{F}$ is a quasicoherent sheaf on X. Then for any open affine $V \subset Y$ we have*

$$R^i f_* \mathcal{F}(V) \cong H^i(f^{-1}\widetilde{V, \mathcal{F}}|_{f^{-1}V}) \cong \check{H}^i(f^{-1}\widetilde{V, \mathcal{F}}|_{f^{-1}V}).$$

In particular every higher direct image $R^i f_* \mathcal{F}$ is a quasicoherent sheaf on Y.

It turns out that higher direct images are quasicoherent even when f is only quasicompact quasiseparated.

Proof. It suffices to prove the statement when Y is affine. We claim that for any distinguished open affine D_g in Y the natural map

$$\check{H}^i(X, \mathcal{F}) \to \check{H}^i(f^{-1}D_g, \mathcal{F}|_{f^{-1}D_g})$$

is the same as localization by g. To check this, fix a finite open affine cover $\mathfrak{U}$ of X. Then the Čech cohomology of $\mathcal{F}$ can be computed as the cohomology of the Čech complex $\mathcal{C}^\bullet(\mathfrak{U}, \mathcal{F})$. For each open affine U_i in $\mathfrak{U}$, consider the distinguished open affine $V_i \subset U_i$ defined by g. Then the V_i form a finite open cover $\mathfrak{V}$ of $f^{-1}D_g$. Furthermore, the Čech complex of $\mathcal{F}|_{f^{-1}D_g}$ with respect to $\mathfrak{V}$ is simply the localization of the Čech complex for $\mathcal{F}$ on X. This proves the claim.

By Lemma 13.3.5, $R^i f_* \mathcal{F}$ is the sheafification of the presheaf assigning to any open V the ith Čech cohomology group of $\mathcal{F}|_{f^{-1}V}$. But the argument above combined with Exercise 9.2.11 shows that this is just the sheaf $\widetilde{\check{H}^i(X, \mathcal{F})}$. Due to Theorem 13.2.5 we can equally well sheaf cohomology in place of Čech cohomology when computing the value of $R^i f_* \mathcal{F}$ along an open affine. □

13.3.2 Coherent sheaves

Suppose that $f : X \to Y$ is a morphism of sheaves and that $\mathcal{F}$ is a coherent sheaf on X. In general we cannot expect $f_* \mathcal{F}$ to be a coherent sheaf. However, when f is proper and Y is locally Noetherian we asserted (without proof) that $f_* \mathcal{F}$ is coherent, and we proved this statement when f is projective. In this section, we extend these results to the higher derived pushforwards $R^i f_* \mathcal{F}$.

Theorem 13.3.7. *Let $f : X \to Y$ be a proper morphism to a locally Noetherian scheme Y. Suppose that $\mathcal{F}$ is a coherent sheaf on X. Then for every $i > 0$ the sheaf $R^i f_* \mathcal{F}$ is coherent on Y.*

As before, we will only prove this statement when f is projective.

Proof. Assume f is projective. It suffices to consider the case when $Y = \mathrm{Spec}(S)$ is affine. Then Theorem 12.3.3 shows that the Čech cohomology groups $\check{H}^i(X, \mathcal{F})$ are coherent S-modules. We conclude the desired statement from Proposition 13.3.6. □

13.3.3 Leray spectral sequence

One of the fundamental tools for computing higher direct images is the Leray spectral sequence.

Theorem 13.3.8 (Leray spectral sequence)**.** *Let $f : X \to Y$ be a morphism of schemes. For any $\mathcal{O}_X$-module $\mathcal{F}$ there is a spectral sequence with E_2 page given by $H^q(Y, R^p f_* \mathcal{F})$ which abuts to $H^{p+q}(X, \mathcal{F})$.*

This is a special case of the Grothendieck spectral sequence from the introduction to the chapter.

Proof. Our plan is to apply Theorem 13.0.11 to the composition of the left exact functors f_* and $\Gamma(Y, -)$. It only remains to note that if $\mathcal{I}$ is an injective $\mathcal{O}_X$-module then $\mathcal{I}$ is flasque, thus $f_*\mathcal{I}$ is flasque and hence acyclic for $\Gamma(Y, -)$. □

13.3.4 Dimension vanishing

In Theorem 12.2.5 we showed that when X is a quasiprojective $\mathbb{K}$-scheme the Čech cohomology groups $\check{H}^i$ vanish for $i > \dim(X)$. Our next statement is a relative version of dimension vanishing.

Theorem 13.3.9. *Suppose that $f : X \to Y$ is a projective morphism and that Y is locally Noetherian. Let d be the maximum dimension of any fiber of f. Then for any quasicoherent sheaf $\mathcal{F}$ on X we have that $R^i f_* \mathcal{F} = 0$ for every $i > d$.*

Proof. We can check whether a sheaf is 0 locally, so it suffices to consider the case when Y is affine. Since affine varieties always carry ample invertible sheaves, by Proposition 10.7.13 we have a closed embedding of X into $\mathbb{P}^n_Y$.

Fix any point $y \in Y$. Then the fiber $f^{-1}y$ admits a closed embedding into $\mathbb{P}^n_{\kappa(y)}$. Copying the argument of Theorem 12.2.5, we can find distinguished open affines $D_{+,f_1}, \ldots, D_{+,f_{d+1}}$ in $\mathbb{P}^n_{\kappa(y)}$ whose intersections with $f^{-1}y$ form an open cover.

We next "spread out" these open subsets to a neighborhood of y. There is an open affine neigborhood V of y such that each homogeneous function f_i extends to a homogeneous function over all of V. Then the intersection of X with $V_+(f_1, \ldots, f_{d+1})$ defines a closed subset of $\mathbb{P}^n_V$. By properness, the image of this closed subset is a closed subset of V, and by construction this closed subset does not contain y. After shrinking V, we may ensure that $V_+(f_1, \ldots, f_{d+1})$ is disjoint from $f^{-1}V \subset X$. By construction $f^{-1}V$ admits an open cover by $d + 1$ open affine subsets, and thus (again mimicking the argument in Theorem 12.2.5) the cohomology $\check{H}^i$ vanishes for $i > d$. Proposition 13.3.6 allows us to conclude that $R^i f_* \mathcal{F}|_V = 0$ whenever $i > d$. Since this is true for every point y, we see that $R^i f_* \mathcal{F} = 0$ for every $i > d$. □

Warning 13.3.10. In contrast to the absolute case in Theorem 12.2.5, the projectivity assumption for f in Theorem 13.3.9 is crucial. (That is, we used the fact that f is a closed morphism in a crucial way, so the theorem need not hold when f is a "quasiprojective morphism".) For example, Exercise 13.3.12 shows that when f is an open embedding there can be (many) non-vanishing higher direct images.

13.3.5 Exercises

Exercise 13.3.11. Show that if $f : X \to Y$ is an affine morphism and $\mathcal{F}$ is a quasicoherent sheaf on X then $R^i f_* \mathcal{F} = 0$ for every $i > 0$. In particular this is true if f is a finite morphism.

Exercise 13.3.12. Let $X = \mathbb{A}^n - \{0\}$ and consider the open embedding $f : X \to \mathbb{A}^n$. Show that $R^{n-1} f_* \mathcal{O}_X \neq 0$.

Exercise 13.3.13. Let X be a scheme and let $\mathcal{E}$ be a locally free sheaf of rank $r+1$ on X. Consider the projection map $\pi : \underline{\mathrm{Proj}}(\mathcal{E}) \to X$. Compute $R^i \pi_* \mathcal{O}(d)$ for all integers d and for all $i \geq 0$.

Exercise 13.3.14. Let $f : X \to Y$ be a continuous map of topological spaces. Suppose that $\mathcal{F}$ is a sheaf of abelian groups on X such that $R^i f_*(\mathcal{F}) = 0$ for $i > 0$. Prove that $H^i(Y, f_*\mathcal{F}) \cong H^i(X, \mathcal{F})$. (This is a degenerate case of the Leray spectral sequence which shows that $H^p(Y, R^q f_* \mathcal{F}) \implies H^{p+q}(X, \mathcal{F})$.)

Exercise 13.3.15. Let $f : X \to Y$ be a morphism of schemes, let $\mathcal{F}$ be an $\mathcal{O}_X$-module and let $\mathcal{E}$ be a locally free $\mathcal{O}_Y$-module of finite rank. Prove the projection formula

$$R^i f_*(\mathcal{F} \otimes f^* \mathcal{E}) \cong R^i f_* \mathcal{F} \otimes \mathcal{E}.$$

13.4 Cohomology and base change I

Let $f : X \to Y$ be a quasicompact and separated map of schemes and suppose $\mathcal{F}$ is a quasicoherent sheaf on X. By Proposition 13.3.2 $R^i f_*(\mathcal{F})$ assigns to an open set $V \subset Y$ the cohomology of $\mathcal{F}$ along $f^{-1}V$. While this is more explicit than the description as a derived functor, unfortunately it is still a bit hard to understand geometrically.

By taking limits we might hope that the fiber of $R^i f_*(\mathcal{F})$ at a point y is somehow related to the cohomology of the fiber $H^i(f^{-1}y, \mathcal{F}|_{f^{-1}y})$. Of course this can't be true in general – we don't have such an easy description even for the pushforward $f_*\mathcal{F}$. Nevertheless, in situations when we can find an advantageous comparison between these two objects, we will be much better equipped to understand the higher pushforwards explicitly.

In fact, we might hope for even more. Suppose we have a pullback diagram

$$\begin{tikzcd} X \times_Y Z \arrow[r, "\phi"] \arrow[d, "g"] & X \arrow[d, "f"] \\ Z \arrow[r, "\psi"'] & Y \end{tikzcd}$$

Recall that the fibers of $g : X \times_Y Z \to Z$ are "pullbacks" of the fibers of $f : X \to Y$. When the geometry of the higher pushforwards of f has a close relationship with the cohomology of fibers, then we can hope to apply similar logic to g as well.

13.4.1 Comparing cohomology groups

Suppose we have a pullback diagram as before

$$\begin{tikzcd} X \times_Y Z \arrow[r, "\phi"] \arrow[d, "g"] & X \arrow[d, "f"] \\ Z \arrow[r, "\psi"'] & Y \end{tikzcd}$$

where f (and hence g) is quasicompact quasiseparated. Given a quasicoherent sheaf $\mathcal{F}$ on X, there are two ways to obtain a quasicoherent sheaf on Z: we can consider $\psi^* f_*\mathcal{F}$ or $g_*\phi^*\mathcal{F}$. We would like to compare these two sheaves (and also the corresponding higher direct images $\psi^*(R^i f_*\mathcal{F})$ and $R^i g_*(\phi^*\mathcal{F})$). The starting point is to show that there is a natural morphism between them.

Lemma 13.4.1. *Suppose given a diagram as above with f quasicompact and separated where $\mathcal{F}$ is a quasicoherent sheaf on X. Then for every $i \geq 0$ there is a natural morphism*

$$\Theta^i : \psi^*(R^i f_*\mathcal{F}) \to R^i g_*(\phi^*\mathcal{F}).$$

Here as usual "naturality" is meant in the functorial sense – given a morphism of sheaves $\mathcal{F} \to \mathcal{G}$ we should get a commuting square from the various Θ – and as usual we will not check this property carefully. The separated condition is stronger than necessary – quasicompact quasiseparated is enough. We include the separated assumption so that we can understand sheaf cohomology via Theorem 13.2.5.

Proof. We first construct the morphism locally. Assume that $Y = \mathrm{Spec}(S)$ and $Z = \mathrm{Spec}(R)$ are affine so that $X \times_Y Z$ is the base change X_R of X by R. In this case Proposition 13.3.6 shows that we can identify

$$\begin{aligned}\psi^*(R^i f_* \mathcal{F}) &= \psi^*(\widetilde{H^i(X, \mathcal{F})}) = \widetilde{H^i(X, \mathcal{F}) \otimes_S R} \\ R^i g_*(\phi^* \mathcal{F}) &= \widetilde{H^i(X_R, \phi^* \mathcal{F})}\end{aligned}$$

Thus it suffices to construct a homomorphism of R-modules

$$H^i(X, \mathcal{F}) \otimes_S R \to H^i(X_R, \phi^* \mathcal{F})$$

We will compare these by translating into Čech cohomology. Since X is separated over $\mathrm{Spec}(S)$, hence separated, we can compute the cohomology of X by choosing a finite cover by open affines $\mathfrak{U} = \{U_i\}$ and taking the cohomology of the Čech complex $\check{C}^\bullet(\mathfrak{U}, \mathcal{F})$. By taking products against $\mathrm{Spec}(R)$, we find a finite cover of X_R by open affines $\mathfrak{U}_R = \{U_i \times_{\mathrm{Spec}(S)} \mathrm{Spec}(R)\}$. Since X_R is also separated, we can again compute sheaf cohomology using the Čech complex $\check{C}^\bullet(\mathfrak{U}_R, \phi^* \mathcal{F})$.

Since $\mathcal{F}$ is quasicoherent and X is separated, for any finite intersection V of elements of the cover $\mathfrak{U}$ we have

$$H^0(V \times_{\mathrm{Spec}(S)} \mathrm{Spec}(R), \phi^* \mathcal{F}) \cong H^0(V, \mathcal{F}) \otimes_S R.$$

and thus an identification

$$\check{C}^\bullet(\mathfrak{U}_R, \phi^* \mathcal{F}) = \check{C}^\bullet(\mathfrak{U}, \mathcal{F}) \otimes_S R.$$

Then the desired map on cohomology follows from a more general fact: given a complex of S-modules $C^\bullet$ and an S-algebra R there is a natural morphism $H^i(C^\bullet) \otimes_S R \to H^i(C^\bullet \otimes_S R)$.

Finally, we claim that the local construction glues to give a global construction. Recall that to construct a morphism ρ between two sheaves $\mathcal{G}, \mathcal{H}$ on a scheme Z, we can choose an open cover $\{U_i\}$ of Z and construct morphisms $\rho_i : \mathcal{G}|_{U_i} \to \mathcal{H}|_{U_i}$ which agree on the overlaps $U_i \cap U_j$. Choose an open cover of Z consisting of open affines that are contained in the preimage of open affines in Y. The local construction above gives a morphism of sheaves on each open set in this open cover. The compatibility on overlaps boils down to the compatibility of the Čech complex with localization as described in the proof of Proposition 13.3.6. □

Remark 13.4.2. There is an alternative "general nonsense" approach to constructing the morphism $\psi^*(R^i f_* \mathcal{F}) \to R^i g_*(\phi^* \mathcal{F})$.

Let's first construct the map $\psi^* f_* \mathcal{F} \to g_* \phi^* \mathcal{F}$. Using the pullback-pushforward adjunction, we obtain a unit map $\mathcal{F} \to \phi_* \phi^* \mathcal{F}$. We then apply f_* to obtain a morphism

$$f_* \mathcal{F} \to f_* \phi_* \phi^* \mathcal{F} = \psi_* g_* \phi^* \mathcal{F}$$

Finally, using the pullback-pushforward adjunction again, we obtain a morphism $\psi^* f_* \mathcal{F} \to g_* \phi^* \mathcal{F}$. The extension to the higher direct images can be accomplished using a δ-functor argument.

Exercise 13.4.3. Show that the construction of Θ^0 in Lemma 13.4.1 using the Čech complex agrees with the construction of Θ^0 in Remark 13.4.2 using adjunction.

13.4.2 Cohomology and flat base change

Lemma 13.4.1 show that for a pullback diagram the higher direct images are locally related by a tensor product. Thus when we have a flat base change diagram we can expect the higher direct images to be compatible.

Proposition 13.4.4 (Cohomology and Flat Base Change)**.** *Suppose given a pullback diagram*

$$\begin{tikzcd} X \times_Y Z \arrow[r, "\phi"] \arrow[d, "g"] & X \arrow[d, "f"] \\ Z \arrow[r, "\psi"'] & Y \end{tikzcd}$$

where f is quasicompact and separated and ψ is flat. Then for every i the map

$$\Theta^i : \psi^*(R^i f_* \mathcal{F}) \to R^i g_*(\phi^* \mathcal{F}).$$

is an isomorphism.

As usual, it is enough for f to be only quasicompact quasiseparated.

Proof. Returning to the proof of Lemma 13.4.1, it suffices to note that when ψ is flat then in the local situation $Z = \operatorname{Spec}(R)$ and $Y = \operatorname{Spec}(S)$ we have that R is a flat S-module. Thus for any complex $C^\bullet$ of S-modules the tensor products $H^i(C^\bullet) \otimes_S R$ are isomorphic to the cohomology groups $H^i(C^\bullet \otimes R)$. $\square$

Unfortunately Proposition 13.4.4 rarely applies to our motivating example when $Z \to Y$ is the inclusion of a point. In fact, the map Θ^i will frequently fail to be an isomorphism in this situation; we will see some examples in the next section.

Here are a couple easy examples of Cohomology and Flat Base Change in action.

Example 13.4.5. For any inclusion of fields $\mathbb{K} \to \mathbb{L}$ the map $\mathrm{Spec}(\mathbb{L}) \to \mathrm{Spec}(\mathbb{K})$ is flat. Thus if X is a $\mathbb{K}$-scheme and $\mathcal{F}$ is a quasicoherent sheaf on X we obtain an isomorphism

$$H^i(X, \mathcal{F}) \otimes \mathbb{L} \cong H^i(X_{\mathbb{L}}, \mathcal{F}_{\mathbb{L}}).$$

Example 13.4.6. Suppose that $f : X \to Y$ is a finite-type separated morphism to an integral scheme Y. Let $\eta = \mathrm{Spec}(K(Y))$ denote the generic point of Y. Note that the inclusion of the generic point $\eta \to Y$ is a flat morphism. For any quasicoherent sheaf $\mathcal{F}$ on X we obtain an isomorphism

$$R^i f_* \mathcal{F}(\eta) \cong H^i(X_\eta, \mathcal{F}|_{X_\eta})$$

The leftmost term records the rank of $R^i f_* \mathcal{F}$ at the generic point of Y, or equivalently, the rank of $R^i f_* \mathcal{F}$ along a general point of Y. The equation shows that we can determine this rank by computing the cohomology group of the restriction of $\mathcal{F}$ to the generic fiber X_η (which is a scheme over a field).

One can do a similar construction with the flat maps $\mathrm{Spec}(\mathcal{O}_{Y,y}) \to Y$, but it is a little less useful since the fiber over $\mathrm{Spec}(\mathcal{O}_{Y,y})$ is no longer simply a scheme over a field.

13.4.3 Exercises

Exercise 13.4.7. Suppose we are given a commutative diagram of schemes

$$\begin{tikzcd} W \arrow[r, "\phi"] \arrow[d, "g"] & X \arrow[d, "f"] \\ Z \arrow[r, "\psi"'] & Y \end{tikzcd}$$

which is not necessarily a pullback diagram. Show that for any quasicoherent sheaf $\mathcal{F}$ on X we have a natural morphism

$$\psi^*(R^i f_* \mathcal{F}) \to R^i g_*(\phi^* \mathcal{F}).$$

(One option is to first apply Lemma 13.4.1 to reduce to the easier case when $Y = Z$. Another option is to apply a general nonsense approach as in Remark 13.4.2.)

13.5 Cohomology and base change II

We continue the discussion of cohomology and base change from the previous section. In this chapter we will focus exclusively on the following situation. Suppose we have a diagram

$$\begin{tikzcd} X_y \arrow[r, "\phi"] \arrow[d, "g"] & X \arrow[d, "f"] \\ y \arrow[r, "\psi"'] & Y \end{tikzcd}$$

where f is proper and ψ is the inclusion of a point $y \in Y$. Let $\mathcal{F}$ be a coherent sheaf on X that is flat over Y. Then Lemma 13.4.1 gives us a morphism

$$\Theta^i_y : R^i f_* \mathcal{F}(y) \to H^i(X_y, \mathcal{F}|_{X_y}).$$

We would like to understand the properties of Θ^i_y. We are particularly interested in understanding when Θ^i_y will be an isomorphism.

13.5.1 Complexes of modules for Noetherian rings

Recall that the maps Θ^i come from a local construction of the form

$$H^i(C^\bullet) \otimes_S M \to H^i(C^\bullet \otimes_S M)$$

where $C^\bullet$ denotes a suitably chosen Čech complex. In this section we discuss a few features of these maps when S is a Noetherian ring. The key result is:

Lemma 13.5.1 ([Har77] Lemma III.12.3). *Let S be a Noetherian ring. Suppose that $C^\bullet$ is a complex of S-modules such that every cohomology group $H^i(C)$ is a finitely generated S-module and $C^d = 0$ for sufficiently large d. Then there is a complex $K^\bullet$ of finite rank free S-modules such that $K^d = 0$ for sufficiently large d and a map of complexes $\phi : K^\bullet \to C^\bullet$ that induces isomorphisms of all cohomology groups.*

In other words, for a Noetherian ring S a complex with finitely generated homology can be replaced by a "simpler" complex whose entries are finite rank free S-modules. Such complexes are much easier to work with (and in particular, can be understood by applying linear algebra techniques to the differentials). Unfortunately the proof is somewhat technical so we will not give it here. It pairs naturally with the following result:

Lemma 13.5.2 ([Har77] Proposition III.12.4, Proposition III.12.5). *Let S be a Noetherian ring. Fix a complex $K^\bullet$ of finitely generated free S-modules. Fix an integer i and define the functor $T^i : M \mapsto H^i(K^\bullet \otimes M)$. Then:*

(1) T^i is left exact if and only if the cokernel of the map $K^{i-1} \to K^i$ is a projective S-module.

(2) T^i is right exact if and only if the map $H^i(K^\bullet) \otimes M \to H^i(K^\bullet \otimes M)$ is an isomorphism for every S-module M.

13.5.2 Semicontinuity

The semicontinuity theorem analyzes the behavior of the cohomology groups of fibers as we vary the point y. The theorem relies on a reduction step that we will use several times:

Lemma 13.5.3. *Let $f : X \to Y$ be a proper morphism of Noetherian schemes such that $Y = \operatorname{Spec}(S)$ is affine and let $\mathcal{F}$ be a coherent sheaf on X that is flat over Y. Choose a finite open affine cover $\mathfrak{U}$ of X. There is a complex $K^\bullet$ of finite rank free S-modules and a map of complexes $\phi : K^\bullet \to C^\bullet(\mathfrak{U}, \mathcal{F})$ such that for every S-module M the map ϕ induces an isomorphism*

$$H^i(K^\bullet \otimes M) \to H^i(C^\bullet(\mathfrak{U}, \mathcal{F}) \otimes M).$$

Proof. By Theorem 13.3.7 the cohomology groups of $C^\bullet$ are finitely generated S-modules. We can thus apply Lemma 13.5.1 to get a complex $K^\bullet$ of finite rank free S-modules with a map $K^\bullet \to C^\bullet$.

Since $\mathcal{F}$ is flat over Y, every module $C^d(\mathfrak{U}, \mathcal{F})$ is a flat S-module. (Argue this carefully!) Also, by Lemma 13.5.1 the K^d are flat S-modules. This guarantees that for any S-module M the induced map

$$H^i(K^\bullet \otimes M) \to H^i(C^\bullet \otimes M)$$

is an isomorphism. Since both groups are 0 for i sufficiently large, we can argue by decreasing induction on i. It suffices to prove that this map is an isomorphism when M is finitely generated, in which case we have an exact sequence $0 \to R \to A^{\oplus d} \to M \to 0$ for some positive integer d. Since $K^\bullet$ is flat, we get a LES of cohomology of the form $H^i(L^\bullet \otimes R) \to H^i(L^\bullet \otimes A^{\oplus d}) \to H^i(L^\bullet \otimes M)$. We get a similar sequence for $C^\bullet$, and also a morphism of LES from the sequence for $K^\bullet$ to the sequence for $C^\bullet$. Since the cohomology groups of $A^{\oplus d}$ vanish, we can conclude the desired isomorphism of H^i for M by induction from the corresponding isomorphism of H^{i+1} for R. $\square$

We are now prepared to prove the Semicontinuity Theorem.

Theorem 13.5.4 (Semicontinuity)**.** *Let $f : X \to Y$ be a proper morphism of Noetherian schemes and let $\mathcal{F}$ be a coherent sheaf on X that is flat over Y. Fix an index $i \geq 0$ and consider the function $y \mapsto \dim_{\kappa(y)} H^i(f^{-1}y, \mathcal{F}|_{f^{-1}y})$. This function is upper semicontinuous on Y.*

This result pairs naturally with the constancy of the Euler characteristic $\sum(-1)^i \dim H^i$ in flat families; if some H^i "jumps up" on a closed subset, some other cohomology group must also "jump up" to ensure that the Euler characteristic does not change.

Proof. Since the statement is local on the base Y, we may assume that $Y = \operatorname{Spec}(S)$ for a Noetherian ring S. Apply Lemma 13.5.3 to obtain a complex $K^\bullet$. For any point $y \in Y$, let R denote the residue field of y considered as an S-algebra. Then we see that

$$\begin{aligned}\dim_{\kappa(y)} H^i(f^{-1}y, \mathcal{F}|_{f^{-1}y}) &= \dim_R(H^i(K^i \otimes_S R)) \\ &= \dim_R(K^i \otimes_S R) - \dim_R(\operatorname{im} d_R^i) - \dim_R(\operatorname{im} d_R^{i-1})\end{aligned}$$

The dimension of the image of d_R^i is lower semicontinuous in y since there is a closed sublocus of Y parametrizing points where the rank of the matrix d^i drops. The analogous statement for d^{i-1} is also true. This proves the desired statement. □

Remark 13.5.5. Suppose that $f : X \to Y$ is a proper morphism of Noetherian schemes and that $\mathcal{F}$ is a coherent sheaf on X (that is not necessarily flat over Y). Using generic flatness we can stratify Y into locally closed subsets W_i such that the restriction of $\mathcal{F}$ is flat over each W_i. In this way one can show that the map $y \mapsto \dim H^i(X_y, \mathcal{F}|_{X_y})$ is constructible, i.e. for any integer r the set of points with value r is a constructible subset of Y.

Example 13.5.6. Example 13.5.11 gives an example of a flat morphism $p : E \times E \to E$ and an invertible sheaf $\mathcal{L}$ on $E \times E$ such that the cohomology of $\mathcal{L}$ on the general fiber is zero, but both H^0 and H^1 jump up to 1 on a special fiber. Note that both groups must "jump up" simultaneously so that the Euler characteristic is preserved.

Example 13.5.7. Suppose that $f : X \to Y$ is a smooth morphism of projective varieties over $\mathbb{C}$. Since f is flat, Theorem 13.5.4 guarantees that the dimensions of the cohomology groups of the structure sheaf along the fibers are semicontinuous functions. Using Hodge theory one can show a stronger result: the cohomology groups $H^i(f^{-1}y, \mathcal{O}_{f^{-1}y})$ have constant dimension.

13.5.3 Cohomology of fibers

We are now return our attention to the morphisms $\Theta_y^i : (R^i f_*\mathcal{F})_y \to H^i(f^{-1}y, \mathcal{F}|_{f^{-1}y})$. We are interested in finding conditions guaranteeing that this map is a bijection, or conversely, when properties of Θ_y^i tell us something about $R^i f_*(\mathcal{F})$.

Theorem 13.5.8 (Grauert's Theorem)**.** *Let $f : X \to Y$ be a proper morphism to a reduced locally Noetherian scheme and let $\mathcal{F}$ be a coherent sheaf on X that is flat over Y. Suppose that $\dim_{\kappa(y)} H^i(f^{-1}y, \mathcal{F}|_{f^{-1}y})$ is a constant function as we vary $y \in Y$. Then $R^i f_*(\mathcal{F})$ is locally free and Θ_y^i is an isomorphism for every $y \in Y$.*

Proof. We can reduce to the case when $Y = \operatorname{Spec}(S)$ is affine. Let $K^\bullet$ be the complex of finite rank free S-modules constructed by Lemma 13.5.3. We define W^p to be the cokernel of the $(p-1)$-differential. Since tensoring by R is right exact, we see that for any R-module M the cokernel of $K^{p-1} \otimes M \to K^p \otimes M$ is isomorphic to $W^p \otimes M$. Thus we have an exact sequence

$$0 \to H^p(K^\bullet \otimes M) \to W^p \otimes M \to K^{p+1} \otimes M.$$

Arguing as in the proof of Theorem 13.5.4, we see that

$$\dim_{\kappa(y)} H^i(f^{-1}y, \mathcal{F}|_{f^{-1}y}) = \dim_R(K^i \otimes_S R) - \dim_R(\operatorname{im} d_R^i) - \dim_R(\operatorname{im} d_R^{i-1})$$

where R is the residue field at y considered as an S-algebra. By assumption this number is constant as we vary y. Since the image of d_R^p is the same thing as the kernel of the map $W^i \otimes R \to K^{i+1} \otimes R$, we conclude that the sheaves $\widetilde{W^i}$ and $\widetilde{W^{i+1}}$ on Y have constant rank. Since by assumption Y is reduced these sheaves must be locally free.

Let T^i be the functor $M \mapsto H^i(K^\bullet \otimes M)$. Applying Lemma 13.5.2 to W^i, W^{i+1} we see that T^i, T^{i+1} are left exact. Since the K^i are flat S-modules, a SES of modules leads to a LES involving the T^i. We conclude that T^i is also right exact, hence exact. Applying Lemma 13.5.2 again, we see that the maps $H^i(K^\bullet \otimes M) \to H^i(K^\bullet \otimes M)$ are isomorphisms for every module M. Applying this statement when M is the residue field at y we deduce that Θ_y^i is an isomorphism. We also see that $H^i(K^\bullet)$ is a flat S-module. Since it is also finitely generated, it is a projective S-module, showing that $R^i f_* \mathcal{F}$ is locally free. □

Our final result goes in the "opposite direction" of Grauert's Theorem.

Theorem 13.5.9 (Cohomology and Base Change Theorem, [Har77] Theorem III.12.11)**.** *Let $f : X \to Y$ be a proper morphism to a locally Noetherian scheme and let $\mathcal{F}$ be a coherent sheaf on X that is flat over Y. Suppose that $y \in Y$ is a point such that Θ_y^i is a surjection. Then*

(1) There is an open neighborhood U of y such that $\Theta_{y'}^i$ is an isomorphism for every $y' \in U$.

(2) $R^i f_*(\mathcal{F})$ is locally free at y if and only if Θ_y^{i-1} is a surjection. When these equivalent conditions hold, $\dim H^i(f^{-1}z, \mathcal{F}|_{f^{-1}z})$ is constant for z in an open neighborhood of y.

Just as before, the proof relies on Lemma 13.5.1 to reduce the statements to facts about complexes of finitely generated free modules. However the underlying algebraic statement is a bit more complicated and we will not discuss it here.

Let's close with a couple examples.

Example 13.5.10. Let Y be a scheme and let $X = \mathbb{P}(\mathcal{E})$ be the projective bundle $f : X \to Y$ defined by a locally free sheaf $\mathcal{F}$ of finite rank. Consider the sheaf $\mathcal{O}_{X/Y}(1)$ on X. We can then compute $R^i f_* \mathcal{O}_{X/Y}(1)$ directly: over each open affine $V \subset Y$ Theorem 12.3.1 shows that

$$\check{H}^i(f^{-1}V, \mathcal{O}_{X/Y}(1)|_{f^{-1}V}) = \check{H}^i(f^{-1}V, \mathcal{O}_{f^{-1}V/V}(1)) = 0$$

for every $i > 0$. We conclude that the higher direct images are all 0. In this case for every point $y \in Y$ the map Θ_y^i is the isomorphism $0 \to 0$.

Example 13.5.11. Let E be an elliptic curve and consider the product $E \times E$. We will be interested in two Cartier divisors on $E \times E$. First, let F_2 denote the fiber of the second projection map over a point $x \in E$. Second, let Δ denote the diagonal (considered as a closed subscheme of $E \times E$). Define the line bundle $\mathcal{L} = \mathcal{O}_{E\times E}(\Delta - F_2)$. We want to study the higher direct images of $\mathcal{L}$ under the first projection map $p : E \times E \to E$.

We first compute the cohomology groups of the restriction of $\mathcal{L}$ to the fibers of the map. For any closed point $y \in E$, the restriction of $\mathcal{L}$ to the fiber $f^{-1}y \cong E$ is the line bundle $\mathcal{O}_E(y - x)$. This is a degree 0 line bundle on E. As we saw in Proposition 12.5.8, for $i = 0, 1$ we have

$$\dim H^i(E, \mathcal{O}_E(y - x)) = \begin{cases} 1 \text{ if } y = x \\ 0 \text{ if } y \neq x \end{cases}$$

Next we show that $p_*\mathcal{L} = 0$. It suffices to check that $H^0(f^{-1}V, \mathcal{L}|_{p^{-1}V}) = 0$ for every open subset $V \subset E$. If to the contrary there was an effective Cartier divisor D representing this invertible sheaf on $f^{-1}V$ then the restriction of D to a general fiber of p would still be effective. However, the computation above shows that this restriction represents the line bundle $\mathcal{O}_E(y - x)$ which has no global sections at all.

Next we show that $R^1p_*\mathcal{L} = \mathbb{K}(x)$ is the skyscraper sheaf at the point x. Associated to the SES of sheaves

$$0 \to \mathcal{O}_{E\times E}(\Delta - F_2) \to \mathcal{O}_{E\times E}(\Delta) \to \mathcal{O}_{F_2}(x) \to 0.$$

we have the LES of higher direct images

$$\begin{gathered} 0 \to p_*\mathcal{O}_{E\times E}(\Delta - F_2) \to p_*\mathcal{O}_{E\times E}(\Delta) \to p_*\mathcal{O}_{F_2}(x) \to \\ \to R^1p_*\mathcal{O}_{E\times E}(\Delta - F_2) \to R^1p_*\mathcal{O}_{E\times E}(\Delta) \to \dots \end{gathered}$$

We have already shown that the leftmost term is zero. Grauert's Theorem shows that $p_*\mathcal{O}_{E\times E}(\Delta)$ is an invertible sheaf which we call $\mathcal{T}$ and that $R^1p_*\mathcal{O}_{E\times E}(\Delta) = 0$. Furthermore p is an isomorphism between F_2 and E. Thus the exact sequence above simplifies to

$$0 \to \mathcal{T} \to \mathcal{O}_E(x) \to R^1p_*\mathcal{O}_{E\times E}(\Delta - F_2) \to 0.$$

Our desired statement will be proved if we can show that $\mathcal{T} \cong \mathcal{O}_E$. Using the fact that the conormal bundle of Δ is the cotangent bundle of E, we see that $\mathcal{O}_{E\times E}(\Delta)|_\Delta = \mathcal{O}_E$. Thus we have a SES

$$0 \to \mathcal{O}_{E\times E} \to \mathcal{O}_{E\times E}(\Delta) \to \mathcal{O}_\Delta \to 0$$

yielding a LES of higher direct images

$$\begin{gathered} 0 \to \mathcal{O}_E \to \mathcal{T} \to \mathcal{O}_E \to \\ \to \mathcal{O}_E \to 0 \to \dots \end{gathered}$$

Since the connecting homomorphism is surjective it must be an isomorphism, and this shows the claim.

Altogether we see that

- Θ^0_y is the map $0 \to \mathbb{K}$ when $y = x$ and $0 \to 0$ when $y \neq x$.
- Θ^1_y is the map $\mathbb{K} \to \mathbb{K}$ when $y = x$ and $0 \to 0$ when $y \neq x$.

Thus:

- We can apply Grauert's Theorem to the open set $E\backslash\{x\}$ to compute that over this set $p_*\mathcal{L} = 0$ and $R^1p_*\mathcal{L} = 0$. However Grauert's Theorem does not apply at x. Indeed, even though $p_*\mathcal{L} = 0$ is locally free the map Θ^0_x is not an isomorphism.
- We can apply Cohomology and Base Change to Θ^1_y at every point in E. It shows that Θ^1_y is always an isomorphism. However, Θ^0_x is not an isomorphism and correspondingly $R^1p_*\mathcal{L}$ is not locally free at x.

We can find many more examples where the cohomology of fibers "jumps up" in closed subsets by considering a family of line bundles of constant degree on a curve C just as in Example 13.5.11.

13.5.4 Exercises

Exercise 13.5.12. Let $\mathbb{K}$ be an algebraically closed field. Suppose that $f : X \to Y$ is a flat projective morphism of integral $\mathbb{K}$-schemes whose fibers are all integral. Suppose that $\mathcal{L}, \mathcal{M}$ are invertible sheaves such that for every fiber $f^{-1}y$ we have $\mathcal{L}|_{f^{-1}y} \cong \mathcal{M}|_{f^{-1}y}$. Prove that there is an invertible sheaf $\mathcal{N}$ on Y such that $\mathcal{L} \cong \mathcal{M} \otimes f^*\mathcal{N}$.

Exercise 13.5.13. Let X be an integral Noetherian scheme and suppose that P is a projective bundle (of finite rank) over X. Prove that $\mathrm{Pic}(P) \cong \mathrm{Pic}(X) \times \mathbb{Z}$.

Exercise 13.5.14. The following special case of Grauert's Theorem and Cohomology and Base Change is often used.

Let $\pi : X \to Y$ be a proper morphism to a locally Noetherian scheme and let $\mathcal{F}$ be a coherent sheaf on X that is flat over Y. Suppose that $H^i(f^{-1}y, \mathcal{F}|_{f^{-1}y}) = 0$ for every $y \in Y$. Show that $R^if_*(\mathcal{F}) = 0$ and that Θ^{i-1}_y is an isomorphism for every $y \in Y$.

Exercise 13.5.15. For each $a \in \mathbb{C}$, let C_a denote the curve in $\mathbb{P}^4$ which is the image of the map $\mathbb{P}^1 \to \mathbb{P}^4$ defined by $(s^4 : s^3t : as^2t^2 : st^3 : t^4)$. These curves together yield a closed subscheme $\mathcal{C} \subset \mathbb{P}^4 \times \mathbb{A}^1$ such that the fiber over $a \in \mathbb{A}^1$ is C_a.

Let $\mathcal{I}$ denote the ideal sheaf of $\mathcal{C}$. Show that $\mathcal{I}$ is flat over $\mathbb{A}^1$. Prove that if we consider the restriction of $\mathcal{I}$ to the fibers of the projection map $\mathbb{P}^4 \times \mathbb{A}^1 \to \mathbb{A}^1$, the cohomology groups H^0, H^1 are 0 on the fibers over $a \neq 0$ but jump up to 1 over the fiber $a = 0$.

13.6 Theorem on formal functions

In the past couple sections we have been studying higher direct images and how they interact with the cohomology of fibers. In this section, we begin by discussing a new addition to this correspondence. We will then apply our results to an important case: the pushforward of the structure sheaf $\mathcal{O}_X$.

13.6.1 Theorem on formal functions

Suppose that $f : X \to Y$ is a quasicompact and separated morphism of schemes and $\mathcal{F}$ is a quasicoherent sheaf on X. As we have seen, in general there is no relationship between the fibers of $f_*\mathcal{F}$ and the cohomology of $\mathcal{F}$ along the fibers. The problem is that by passing to the fiber we are losing too much information about the local behavior of $\mathcal{F}$ near y.

Suppose we take the "nth infinitesimal neighborhood" of a closed point y: if y is defined by the ideal sheaf $\mathcal{I}$, then consider the closed subscheme defined by $\mathcal{I}^n$. While the underlying set is still just the closed point y, the nilpotent structure now records the "nth order jets" of functions near y. We might hope that this enrichment improves the relationship between the fibers of the cohomology and the cohomology of the fibers.

The following statement shows that if we pass to completions then we get a precise comparison.

Theorem 13.6.1 (Theorem on Formal Functions)**.** *Let $f : X \to Y$ be a proper morphism of Noetherian schemes. Suppose that $\mathcal{F}$ is a coherent sheaf on X. Fix a quasicoherent ideal sheaf $\mathcal{I}$ on Y defining a closed subscheme Z and let Z_n denote the closed subscheme defined by $\mathcal{I}^n$. Then the natural map*

$$\varprojlim \left(R^i f_*\mathcal{F} \otimes \mathcal{O}_Y/\mathcal{I}^{n+1}\right) \to \varprojlim R^i f_* \left(\mathcal{F}|_{X_{Z_n}}\right)$$

is an isomorphism.

Here the "natural map" is induced by applying Lemma 13.4.1 to each nth infinitesimal neighborhood and then appealing to the universal property of the inverse limit.

Suppose for simplicity that $Z = y$ is a closed point. At its heart the Theorem on Formal Functions contrasts two ways of computing the fiber of $R^i f_*\mathcal{F}$ on an infinitesimal neighborhood of y. On the right, the higher direct images coincide with the cohomology groups $H^i(X_{y_n}, \mathcal{F}|_{X_{y_n}})$. On the left, suppose for simplicity that $Y = \mathrm{Spec}(S)$ is affine and that y is defined by the maximal ideal $\mathfrak{m}$. If M is the S-module defining the quasicoherent sheaf $R^i f_*\mathcal{F}$ then the left-hand side of the equation is $M \otimes_S \widehat{S}$ where the completion is taken with respect to $\mathfrak{m}$.

Remark 13.6.2. As mentioned earlier, tensoring by $\lim_{\leftarrow} \mathcal{O}_Y/\mathcal{I}^{n+1}$ is known as "passing to a formal local neighborhood of Z". Even though formal local neighborhoods are not schemes, the Theorem on Formal Functions can be interpreted as an equality of sheaves on the formal local neighborhood of Z.

One often sees the following variant which replaces the closed set Z with a (not necessarily closed) point y.

Theorem 13.6.3. *Let $f : X \to Y$ be a proper morphism of Noetherian schemes. Suppose that $\mathcal{F}$ is a coherent sheaf on X. Fix a point $y \in Y$. Let X_n denote the base-change of f over the natural inclusion $\mathrm{Spec}(\mathcal{O}_{Y,y}/\mathfrak{m}_y^n) \to Y$ and let $\mathcal{F}_n$ denote the pullback of $\mathcal{F}$ to X_n. Let $\widehat{(R^i f_* \mathcal{F})}_s$ denote the completion of the $\mathcal{O}_{Y,y}$-module $(R^i f_* \mathcal{F})_s$ with respect to the maximal ideal $\mathfrak{m}_y$. Then the map*

$$\widehat{(R^i f_* \mathcal{F})}_s \to \varprojlim H^i(X_n, \mathcal{F}_n)$$

is an isomorphism.

Proof. Since the map $\mathrm{Spec}(\mathcal{O}_{Y,y}) \to Y$ is flat, the theorem on Cohomology and Flat Base Change shows that if we form the product diagram the higher direct images commute with pullback. We then apply Theorem 13.6.1 to the map $X \times_Y \mathrm{Spec}(\mathcal{O}_{Y,y}) \to \mathrm{Spec}(\mathcal{O}_{Y,y})$ and the quasicoherent ideal sheaf $\mathfrak{m}_y$. $\square$

As a consequence, we obtain:

Corollary 13.6.4. *Let $f : X \to Y$ be a proper morphism of Noetherian schemes and let $r = \sup_{y \in Y} \dim(f^{-1}y)$. Suppose that $\mathcal{F}$ is a coherent sheaf on X. Then $R^i f_* \mathcal{F} = 0$ for all $i > r$.*

Proof. As in Theorem 13.6.3 let X_n denote the base change of f to $\mathrm{Spec}(\mathcal{O}_{Y,y}/\mathfrak{m}_y^n)$. For any $y \in Y$, the topological space underlying X_n is homeomorphic to the topological space underlying X_y. Thus $H^i(X_n, \mathcal{F}_n) = 0$ for $i > r$ by Theorem 13.2.8. By Theorem 13.6.3 we see that

$$\widehat{(R^i f_* \mathcal{F})}_s = \varprojlim \left(R^i f_* \mathcal{F} \otimes \mathcal{O}_{Y,y}/\mathfrak{m}_y^n\right)$$

is also zero. Recall that for a Noetherian ring the completion operation is exact on the category of finitely generated modules. Since f is proper, $R^i f_* \mathcal{F}$ is a coherent sheaf. Since for $i > r$ every stalk of $R^i f_* \mathcal{F}$ vanishes when taking completions, we conclude that every stalk is zero and hence $R^i f_* \mathcal{F} = 0$. $\square$

13.6.2 Connected fibers

Suppose that $f : X \to Y$ is a proper morphism of Noetherian schemes so that $f_* \mathcal{O}_X$ is a coherent $\mathcal{O}_Y$-module. It is then natural to wonder whether the properties of this module – e.g. being locally free – have geometric consequences. The main question of interest for us is:

Question 13.6.5. When is the map $f^\sharp : \mathcal{O}_Y \to f_* \mathcal{O}_X$ an isomorphism?

Let's start by considering the "absolute situation" where X is a projective $\mathbb{K}$-scheme and we take the structure map $f : X \to \operatorname{Spec}(\mathbb{K})$. In this case we are asking: when is $H^0(X, \mathcal{O}_X) \cong \mathbb{K}$? We have:

- Necessary condition: X is connected. Indeed, if X is disconnected then the ring $H^0(X, \mathcal{O}_X)$ splits up as a product of the sections on each connected component of X.
- Sufficient condition: X is geometrically irreducible and geometrically reduced. In this case we have $H^0(X, \mathcal{O}_X) \otimes \overline{\mathbb{K}} \cong H^0(X_{\overline{\mathbb{K}}}, \mathcal{O}_{X_{\overline{\mathbb{K}}}})$ and Exercise 2.11.12 shows that a variety over an algebraically closed field has global sections equal to $\overline{\mathbb{K}}$.

We now transition to considering the relative situation.

Proposition 13.6.6. *Let $f : X \to Y$ be a proper flat morphism of Noetherian schemes such that* $\dim H^0(X_y, \mathcal{O}_{X_y}) = 1$ *for all $y \in Y$. Then $f^\sharp : \mathcal{O}_Y \to f_*\mathcal{O}_X$ is an isomorphism.*

In particular, the hypotheses hold whenever the fibers of f are geometrically reduced and geometrically connected.

Proof. Consider the map $f^\sharp : \mathcal{O}_Y \to f_*\mathcal{O}_X$. Since f is proper and flat it is surjective, and thus the induced map of fibers $\mathcal{O}_Y(y) \to f_*\mathcal{O}_X(y)$ does not vanish at any point $y \in Y$. We can then compose this map with the map $f_*\mathcal{O}_X(y) \to H^0(X_y, \mathcal{O}_{X_y}) \cong \kappa(y)$ coming from Lemma 13.4.1. Since this composition is surjective, we can apply the Cohomology and Base Change Theorem to $f_*\mathcal{O}_X$. We deduce that $f_*\mathcal{O}_X$ is locally free of rank 1. Since $f^\sharp$ induces an isomorphism of fibers over every point $y \in Y$, Nakayama's Lemma implies that $f^\sharp$ induces isomorphisms of stalks at every point, and thus is an isomorphism. □

The following key result reverses the implication of Proposition 13.6.6.

Theorem 13.6.7. *Suppose that $f : X \to Y$ is a proper morphism of Noetherian schemes such that $f^\sharp : \mathcal{O}_Y \to f_*\mathcal{O}_X$ is an isomorphism. Then for every $y \in Y$ the fiber $f^{-1}y$ is connected.*

Proof. Suppose for a contradiction that $f^{-1}y$ is not connected, say $f^{-1}y = Z' \sqcup Z''$. Then for every positive integer n we have

$$H^0(X_n, \mathcal{O}_{X_n}) \cong H^0(Z'_n, \mathcal{O}_{Z'_n}) \oplus H^0(Z''_n, \mathcal{O}_{Z''_n})$$

where Z'_n and Z''_n denote the nth infinitesimal neighborhoods of Z' and Z'' respectively. By taking a limit and applying Theorem 13.6.1, we see that

$$(f_*\mathcal{O}_X)\widehat{\ }_s = \varprojlim \left(f_*\mathcal{O}_X \otimes \mathcal{O}_{Y,y}/\mathfrak{m}_y^n\right)$$

is also a direct sum of two non-trivial rings. But since by assumption $f_*\mathcal{O}_X = \mathcal{O}_Y$ the right hand side is the same as the completion of $\mathcal{O}_{Y,y}$ along its maximal ideal. Note that the completion of a local ring along its maximal ideal is a local ring. Furthermore, a local ring cannot be a direct sum of two non-trivial rings: if e_1, e_2 are idempotents in a local ring $(A, \mathfrak{m})$ satisfying $e_1e_2 = 0$ then we must have $e_1, e_2 \in \mathfrak{m}$ so that it is impossible for $e_1 + e_2 = 1$. This gives us the desired contradiction. □

Normal varieties interact particularly well with finite maps. The following statement is one of many different versions of Zariski's Main Theorem.

Corollary 13.6.8. *Suppose that $f : X \to Y$ is a birational proper morphism of integral Noetherian schemes such that Y is normal. Then f has connected fibers.*

When Y is not normal, the normalization map frequently provides a counterexample to this statement.

Proof. By Theorem 13.6.7 it suffices to prove that $f_*\mathcal{O}_X$ is isomorphic to $\mathcal{O}_Y$. It suffices to prove this when $Y = \mathrm{Spec}(S)$ is affine. Since f is proper $f_*\mathcal{O}_X$ is coherent, and so $H^0(Y, f_*\mathcal{O}_X)$ is a finitely generated S-module. Furthermore, since f is birational □

13.6.3 Stein factorizations

We next discuss a key application of Theorem 13.6.7. Suppose that $f : X \to Y$ is a proper morphism of Noetherian schemes. Then $f_*\mathcal{O}_X$ is a finite $\mathcal{O}_Y$-algebra. In particular, we can define the scheme $Z = \underline{\mathrm{Spec}}_Y(f_*\mathcal{O}_X)$. Note that Z comes equipped with a finite morphism $g : Z \to Y$.

Furthermore, we claim that the map $f : X \to Y$ factors through $g : Z \to Y$. Indeed, for every open affine $V \subset Y$ and every open affine $U \subset f^{-1}V$ we have a factoring

$$\begin{tikzcd} U \arrow[rr, "f"] \arrow[dr, "h"'] & & V \\ & \mathrm{Spec}(f_*\mathcal{O}_X(V)) \arrow[ur, "g"'] & \end{tikzcd}$$

induced by the ring maps

$$\begin{tikzcd} \mathcal{O}_X(U) & & \mathcal{O}_Y(V) \arrow[ll] \arrow[dl, "f^\sharp(V)"'] \\ & \mathcal{O}_X(f^{-1}V) \arrow[ul, "\rho_{f^{-1}V,U}"] & \end{tikzcd}$$

It is clear that as we vary U, V the resulting maps glue to give a diagram

$$\begin{tikzcd} X \arrow[rr, "f"] \arrow[dr, "h"'] & & Y \\ & Z \arrow[ur, "g"'] & \end{tikzcd}$$

Definition 13.6.9. The diagram above is known as the Stein factorization for the proper morphism $f : X \to Y$ of Noetherian schemes.

The key property of the Stein factorization is that the morphism $g : Z \to Y$ is finite and the morphism $h : X \to Z$ has connected fibers since by construction $\mathcal{O}_Z \cong h_*\mathcal{O}_X$. (In fact, we only used Theorem 13.6.7 to guarantee that the fibers of h are connected.) In some sense this shows that finite morphisms are "orthogonal" to morphisms with connected fibers.

Exercise 13.6.10. Let $f : X \to Y$ be a proper morphism of Noetherian schemes. Suppose that f factors through a finite morphism $g' : Z' \to Y$. Prove that the Stein factorization $g : Z \to Y$ factors through g' as well.

13.7 Ext functors

In this section we discuss two derived functors – the global Ext (derived from the global Hom functor) and the sheafy Ext (derived from the sheafy Hom functor).

13.7.1 Properties of Ext

Definition 13.7.1. Let X be a scheme and let $\mathcal{F}$ be an $\mathcal{O}_X$-module. We define $\operatorname{Ext}^i(\mathcal{F}, -)$ as the right derived functors of $\operatorname{Hom}(\mathcal{F}, -) : \mathcal{O}_X-\mathbf{Mod} \to \mathcal{O}_X(X)-\mathbf{Mod}$.

We define $\mathcal{E}xt^i(\mathcal{F}, -)$ as the right derived functors of $\mathcal{H}om(\mathcal{F}, -) : \mathcal{O}_X-\mathbf{Mod} \to \mathcal{O}_X-\mathbf{Mod}$.

Exercise 13.7.2. Let X be a scheme and let $\mathcal{G}$ be an $\mathcal{O}_X$-module. Show that $\operatorname{Ext}^i(\mathcal{O}_X, \mathcal{G}) \cong H^i(X, \mathcal{G})$.

Show that $\mathcal{E}xt^i(\mathcal{O}_X, \mathcal{G})$ is $\mathcal{G}$ when $i = 0$ and is 0 when $i > 0$.

As in any abelian category the group $\operatorname{Ext}^i(\mathcal{F}, \mathcal{G})$ classifies "ith Yoneda extensions" of $\mathcal{F}$ by $\mathcal{G}$. By far the most commonly used case is Ext^1:

Theorem 13.7.3. *Let X be a scheme and let $\mathcal{F}, \mathcal{G}$ be $\mathcal{O}_X$-modules. Then there is a bijection between elements of $\operatorname{Ext}^1(\mathcal{G}, \mathcal{F})$ and exact sequences*

$$0 \to \mathcal{F} \to \mathcal{K} \to \mathcal{G} \to 0$$

where we consider two sequences with central terms $\mathcal{K}$ and $\mathcal{K}'$ equivalent if there is a isomorphism $\phi : \mathcal{K} \to \mathcal{K}'$ which forms a commutative diagram of SES when combined with the identity maps for $\mathcal{F}, \mathcal{G}$.

This bijection assigns to any extension the image of the identity map under the connecting morphism $\mathcal{H}om(\mathcal{G}, \mathcal{G}) \to \mathcal{E}xt^1(\mathcal{G}, \mathcal{F})$ coming from the LES for $\mathcal{H}om(\mathcal{G}, -)$. In particular, the 0 element in $\operatorname{Ext}^1(\mathcal{G}, \mathcal{F})$ corresponds to the non-split extension $\mathcal{F} \oplus \mathcal{G}$.

It is not true that $\operatorname{Ext}^i(\mathcal{F}, \mathcal{G}) = \mathcal{E}xt^i(\mathcal{F}, \mathcal{G})(X)$ (see Exercise 13.7.2). Instead, the two constructions are related by local-to-global spectral sequence for Ext:

Theorem 13.7.4. *Let X be a scheme and let $\mathcal{F}, \mathcal{G}$ be $\mathcal{O}_X$-modules. Then there is a spectral sequence with E_2-page $H^j(X, \mathcal{E}xt^i(\mathcal{F}, \mathcal{G}))$ which converges to $\operatorname{Ext}^{i+j}(\mathcal{F}, \mathcal{G})$.*

Proof. We want to apply Theorem 13.0.11 to compose the two functors $\mathcal{H}om(\mathcal{F}, -)$ and $\Gamma(X, -)$. We only need to verify that when $\mathcal{I}$ is an injective $\mathcal{O}_X$-module the sheaf $\mathcal{H}om(\mathcal{F}, \mathcal{I})$ is acyclic for $\Gamma(X, -)$.

In fact, we claim that $\mathcal{H}om(\mathcal{F}, \mathcal{I})$ is flasque. Indeed, let U be an open subset of X and let $\phi \in \mathcal{H}om(\mathcal{F}, \mathcal{G})(U)$, that is, $\phi : \mathcal{F}|_U \to \mathcal{I}|_U$. If $j : U \to X$ is the inclusion, we obtain $j_!\phi : j_!\mathcal{F}|_U \to j_!\mathcal{I}|_U$. Furthermore, we have injections $g_1 : j_!\mathcal{F}|_U \to \mathcal{F}$ and $g_2 : j_!\mathcal{I}|_U \to \mathcal{I}$. By composing $j_!\phi$ with g_2 and applying the universal property of injective sheaves to the injection g_1, we obtain a map $\mathcal{F} \to \mathcal{I}$ whose restriction to U is ϕ. Thus $\mathcal{H}om(\mathcal{F}, \mathcal{I})$ is flasque, and hence acyclic for $\Gamma(X, -)$ by Lemma 13.1.9. $\square$

13.7.2 Left-derived-type properties

Note that in contrast to the Hom functor for modules, we can not give the Ext functor the structure of a derived functor in the leftmost entry (since we do not have enough projectives in our category). However, we are able to recover some of this structure using different arguments.

Proposition 13.7.5. *Let X be a scheme and let $0 \to \mathcal{F} \to \mathcal{G} \to \mathcal{H} \to 0$ be an exact sequence of $\mathcal{O}_X$-modules. Then for any $\mathcal{O}_X$-module $\mathcal{L}$ we have a long exact sequence*

$$\begin{aligned} 0 &\to \mathrm{Hom}(\mathcal{H}, \mathcal{L}) \to \mathrm{Hom}(\mathcal{G}, \mathcal{L}) \to \mathrm{Hom}(\mathcal{F}, \mathcal{L}) \to \\ &\to \mathrm{Ext}^1(\mathcal{H}, \mathcal{L}) \to \mathrm{Ext}^1(\mathcal{G}, \mathcal{L}) \to \mathrm{Ext}^1(\mathcal{F}, \mathcal{L}) \to \\ &\to \mathrm{Ext}^2(\mathcal{H}, \mathcal{L}) \to \dots \end{aligned}$$

We have a similar result for $\mathcal{E}xt$ sheaves.

Proof. Choose an injective resolution $0 \to \mathcal{L} \to \mathcal{I}^\bullet$. Since $\mathrm{Hom}(-, \mathcal{I})$ is an exact contravariant functor for any injective sheaf $\mathcal{I}$, we have an exact sequence of complexes

$$0 \to \mathrm{Hom}(\mathcal{H}, \mathcal{I}^\bullet) \to \mathrm{Hom}(\mathcal{G}, \mathcal{I}^\bullet) \to \mathrm{Hom}(\mathcal{F}, \mathcal{I}^\bullet) \to 0.$$

The associated long exact sequence of homology groups gives the long exact sequence of Ext sheaves. The proof for $\mathcal{E}xt$ sheaves is exactly the same. □

The following result is very important for computing Ext sheaves.

Proposition 13.7.6. *Let X be a scheme. Suppose we have an exact sequence of $\mathcal{O}_X$-modules*

$$\dots \to \mathcal{L}_1 \to \mathcal{L}_0 \to \mathcal{G} \to 0$$

where the $\mathcal{L}_i$ are locally free coherent sheaves of finite rank. Then for any $\mathcal{O}_X$-module $\mathcal{F}$ we have

$$\mathcal{E}xt^i(\mathcal{G}, \mathcal{F}) \cong H^i(\mathcal{H}om(\mathcal{L}_\bullet, \mathcal{F})).$$

Proof. Setting $\mathcal{K}_{-1} = \mathcal{G}$, the complex $\mathcal{L}_\bullet \to \mathcal{G} \to 0$ yields a collection of short exact sequences

$$0 \to \mathcal{K}_j \to \mathcal{L}_j \to \mathcal{K}_{j-1} \to 0.$$

For each such short exact sequence Proposition 13.7.5 gives a LES of Ext sheaves, and in particular by Exercise 13.7.12 we have $\mathcal{E}xt^i(\mathcal{K}_{j+1}, \mathcal{F}) = \mathcal{E}xt^{i+1}(\mathcal{K}_j, \mathcal{F})$ for every $i > 0$ and every $j \geq -1$. By truncating the complex $L_\bullet$ at an earlier place and using induction, we see that it suffices to prove the desired statement for $i = 0, 1$. The $i = 0$ case follows easily from the left-exactness of $\mathcal{H}om$.

Finally, we need to prove an isomorphism for $i = 1$. The various LES of cohomology constructed above yield an exact sequence

$$\ldots \to \mathcal{H}om(\mathcal{L}_0, \mathcal{F}) \to \mathcal{H}om(\mathcal{K}_0, \mathcal{F}) \to \mathcal{E}xt^1(\mathcal{G}, \mathcal{F}) \to 0.$$

and an exact sequence

$$0 \to \mathcal{H}om(\mathcal{K}_0, \mathcal{F}) \to \mathcal{H}om(\mathcal{L}_1, \mathcal{F}) \to \mathcal{H}om(\mathcal{K}_1, \mathcal{F}) \to \ldots$$

The second exact sequence shows that $\mathcal{H}om(\mathcal{K}_0, \mathcal{F})$ is the kernel of the first chain map in the complex $\mathcal{H}om(\mathcal{L}_\bullet, \mathcal{F})$ and the first shows that $\mathcal{E}xt^1(\mathcal{G}, \mathcal{F})$ is the quotient of this group by the image of the previous chain map. □

In order to apply Proposition 13.7.6, we need to be able to find a suitable sequence of locally free sheaves of finite rank. (Such a sequence is known as a locally free resolution.) This result is most useful when the scheme X is Noetherian:

- Proposition 13.7.8 shows that $\mathcal{E}xt^i$ can be computed locally, so we may shrink X so that it is a Noetherian affine scheme. For a Noetherian ring R, any finitely generated R-module admits a resolution by finite rank free modules. Constructing such a resolution for $\mathcal{G}(X)$ and then taking $\tilde{\ }$-images, we find a locally free resolution suitable for Proposition 13.7.6.

- When X is a projective scheme over a Noetherian ring and $\mathcal{G}$ is a coherent sheaf on X, Corollary 10.6.13 shows that there is a surjection from a locally free sheaf of finite rank to $\mathcal{G}$. In this case we have a "global" locally free resolution and can use this to compute without passing to an open cover.

Remark 13.7.7. The question of whether every coherent sheaf on a scheme X admits a surjection from a locally free sheaf of finite rank is a delicate one. This question is still open even for normal proper varieties over an algebraically closed field.

It is known that this property holds for quasiprojective varieties over a noetherian ring, or for regular Noetherian integral separated schemes. In these cases one can again use "global" locally free resolutions to compute Ext.

13.7.3 Quasicoherence

Even if $\mathcal{F}$ and $\mathcal{G}$ are quasicoherent, the sheaves $\mathcal{E}xt^i(\mathcal{F}, \mathcal{G})$ need not be quasicoherent. Indeed, this fails even when $i = 0$ unless we impose some finiteness hypothesis on $\mathcal{F}$. In this section we discuss the quasicoherence of these sheaves.

The first step is to show that the $\mathcal{E}xt$ sheaves interact well with passing to open sets. (This should not be surprising since the $\mathcal{H}om$ sheaves are defined using a local construction.)

Proposition 13.7.8. *Let X be a scheme and let U be an open subset. For any $\mathcal{O}_X$-modules $\mathcal{F}, \mathcal{G}$ we have*

$$\mathcal{E}xt^i(\mathcal{F}, \mathcal{G})|_U \cong \mathcal{E}xt^i(\mathcal{F}|_U, \mathcal{G}|_U).$$

Proof. Choose an injective resolution $0 \to \mathcal{G} \to \mathcal{I}^\bullet$. Then $\mathcal{E}xt^i(\mathcal{F}, \mathcal{G})$ is the cohomology of the sequence $\mathcal{H}om(\mathcal{F}, \mathcal{I}^\bullet)$. Since restriction to U is an exact functor and since the restriction of an injective sheaf on X is injective on U by Exercise 13.1.14, we see that both sides are computed by the cohomology of the sequence

$$\mathcal{H}om(\mathcal{F}, \mathcal{I}^\bullet)|_U = \mathcal{H}om(\mathcal{F}|_U, \mathcal{I}|_U^\bullet).$$

□

We can now specialize to the quasicoherent situation.

Proposition 13.7.9. *Let X be a Noetherian scheme, $\mathcal{F}$ a coherent sheaf on X, and $\mathcal{G}$ a quasicoherent sheaf on X. Then $\mathcal{E}xt^i(\mathcal{F}, \mathcal{G})$ is a quasicoherent sheaf on X and for any open affine U we have*

$$\mathcal{E}xt^i(\mathcal{F}, \mathcal{G})|_U \cong \mathrm{Ext}^i(\widetilde{\mathcal{F}(U), \mathcal{G}(U)}).$$

Proof. By Proposition 13.7.8 it suffices to prove the desired statement when $X = \mathrm{Spec}(R)$ is affine. Suppose that $\mathcal{F} \cong \widetilde{M}$ and $\mathcal{G} \cong \widetilde{N}$. Since M is finitely generated, it admits a resolution by finitely generated free R-modules F_i. Set $\mathcal{L}_i = \widetilde{F_i}$. Since each $\mathcal{L}_i$ is coherent we have $\mathcal{H}om(\mathcal{L}_i, \mathcal{G}) \cong \widetilde{\mathrm{Hom}(F_i, N)}$ and we conclude the desired statement by Proposition 13.7.6. □

Corollary 13.7.10. *Let X be a Noetherian scheme, $\mathcal{F}$ and $\mathcal{G}$ coherent sheaves on X. Then $\mathcal{E}xt^i(\mathcal{F}, \mathcal{G})$ is a coherent sheaf on X.*

Corollary 13.7.11. *Let X be a Noetherian scheme, $\mathcal{F}$ a coherent sheaf on X, and $\mathcal{G}$ a quasicoherent sheaf on X. Then for any point $x \in X$ we have*

$$\mathcal{E}xt^i(\mathcal{F}, \mathcal{G})_x \cong \mathrm{Ext}^i_{\mathcal{O}_{X,x}}(\mathcal{F}_x, \mathcal{G}_x).$$

Proof. Since $\mathcal{F}(U)$ is finitely presented for any open affine U, on each open affine localization commutes with taking Homs and the desired statement follows. □

13.7.4 Exercises

Exercise 13.7.12. Let X be a scheme and let $\mathcal{E}$ be a locally free sheaf on X of finite rank. Prove that $\mathcal{E}xt^i(\mathcal{E}, \mathcal{G}) = 0$ for any $\mathcal{O}_X$-module $\mathcal{G}$ for $i > 0$.

Exercise 13.7.13. Let X be a scheme, let $\mathcal{F}$ and $\mathcal{G}$ be $\mathcal{O}_X$-modules and let $\mathcal{E}$ be a locally free sheaf of finite rank.

(1) Show that if $\mathcal{I}$ is an injective $\mathcal{O}_X$-module then $\mathcal{I} \otimes \mathcal{E}$ is also injective.

(2) Show that we have isomorphisms

$$\mathrm{Ext}^i(\mathcal{F} \otimes \mathcal{E}^\vee, \mathcal{G}) \cong \mathrm{Ext}^i(\mathcal{F}, \mathcal{E} \otimes \mathcal{G})$$

and

$$\mathcal{E}xt^i(\mathcal{F} \otimes \mathcal{E}^\vee, \mathcal{G}) \cong \mathcal{E}xt^i(\mathcal{F}, \mathcal{E} \otimes \mathcal{G}) \cong \mathcal{E}xt^i(\mathcal{F}, \mathcal{G}) \otimes \mathcal{E}.$$

13.8 Vector bundles on curves

In this section we analyze locally free sheaves on smooth projective curves over an algebraically closed field $\mathbb{K}$. There are two basic invariants of such a locally free sheaf $\mathcal{E}$: the rank and the degree. We have seen the rank before, but the degree is a new notion for us.

Definition 13.8.1. Let C be a smooth projective geometrically integral curve and let $\mathcal{E}$ be a locally free sheaf on C. We define the degree of $\mathcal{E}$ to be the degree of the invertible sheaf obtained by taking the top exterior power of $\mathcal{E}$:

$$\deg(\mathcal{E}) := \deg\left(\bigwedge^{\mathrm{rk}(\mathcal{E})} \mathcal{E}\right)$$

If we are working with curves defined over $\mathbb{C}$ this definition coincides with the first chern class of $\mathcal{E}$.

Exercise 13.8.2. Suppose that $\mathcal{E} = \mathcal{L}_1 \oplus \ldots \oplus \mathcal{L}_r$ is a direct sum of line bundles. Show that $\bigwedge^r \mathcal{E} = \mathcal{L}_1 \otimes \ldots \otimes \mathcal{L}_r$. In particular the degree of $\mathcal{E}$ is $\deg(\mathcal{L}_1) + \ldots + \deg(\mathcal{L}_r)$.

Exercise 13.8.3. Let C be a smooth projective curve over a field. Suppose that $0 \to \mathcal{F} \to \mathcal{G} \to \mathcal{H} \to 0$ is a SES of locally free sheaves on C. Show that

$$\deg(\mathcal{F}) + \deg(\mathcal{H}) = \deg(\mathcal{G}).$$

We will focus on the following feature of locally free sheaves:

Definition 13.8.4. Let C be a smooth projective geometrically integral curve. A locally free sheaf $\mathcal{E}$ on C is said to be decomposable if there are non-trivial subsheaves $\mathcal{F}, \mathcal{G}$ such that $\mathcal{E} \cong \mathcal{F} \oplus \mathcal{G}$. If no such decomposition exists, we say that $\mathcal{E}$ is indecomposable.

13.8.1 Vector bundles on $\mathbb{P}^1$

We start by proving a fundamental result about $\mathbb{P}^1$ known as the Grothendieck-Birkhoff Theorem.

Theorem 13.8.5. *A locally free sheaf $\mathcal{E}$ on $\mathbb{P}^1$ of rank r is isomorphic to a direct sum of line bundles:*

$$\mathcal{E} \cong \oplus_{i=1}^r \mathcal{O}(a_i).$$

If we insist that $a_i \geq a_{i+1}$ then the a_i are uniquely determined.

Proof. The proof is by induction on the rank r. The base case $r = 1$ was settled in Example 10.1.4.

First, we claim that there is some integer k such that $H^0(\mathbb{P}^1, \mathcal{E}(k)) \neq 0$ but $H^0(\mathbb{P}^1, \mathcal{E}(j)) = 0$ for every $j < k$. Since $O(1)$ is ample, we know that some sufficiently positive twist of

$\mathcal{E}$ is globally generated and in particular has sections. On the other hand, a section of $\mathcal{E}(m)$ yields a non-zero map $\phi : \mathcal{O} \to \mathcal{E}(m)$. Such a map is necessarily injective: since $\mathcal{E}(m)$ is torsion-free ϕ is injective at the generic point, showing that the kernel is a torsion sheaf. But then $\ker(\phi) = 0$ since $\mathcal{O}$ is torsion-free. After twisting, ϕ leads to an injection $\mathcal{O}(-m) \to \mathcal{E}$. Since $H^0(\mathbb{P}^1, \mathcal{E})$ is fixed and $H^0(\mathbb{P}^1, \mathcal{O}(-m))$ goes to ∞ as m becomes more and more negative, the set of m for which we have an inclusion of this type is bounded below. In conclusion, $H^0(\mathbb{P}^1, \mathcal{E}(m))$ is non-zero for sufficiently positive m and is zero for sufficiently negative m, showing the claim.

Let $\mathcal{K}$ denote the cokernel of $\mathcal{O} \to \mathcal{E}(k)$. We claim that $\mathcal{K}$ is locally free. By Exercise 9.4.26, it suffices to show that $\mathcal{K}$ is torsion free. Consider the exact sequence

$$0 \to \mathcal{O}(-1) \to \mathcal{E}(k-1) \to \mathcal{K}(-1) \to 0.$$

If $\mathcal{K}$ were not torsion free, then the torsion part would make a non-zero contribution to $H^0(\mathbb{P}^1, \mathcal{K}(-1))$. Since $H^1(\mathbb{P}^1, \mathcal{O}(-1)) = 0$, we would conclude that $\mathcal{E}(k-1)$ has a section, a contradiction.

By induction we can write $\mathcal{K} \cong \oplus_{i=1}^{r-1} \mathcal{O}(a_i)$. Repeating the exact sequence argument from the previous paragraph, we see that $a_i \leq 0$ for every i. We conclude that $\mathrm{Ext}^1(\mathcal{K}, \mathcal{O}) = 0$, or in other words, $\mathcal{E}(k) \cong \mathcal{O} \oplus \mathcal{K}$. Thus $\mathcal{E}(k)$, and hence also $\mathcal{E}$, is isomorphic to a direct sum of line bundles. $\square$

Exercise 13.8.6. Prove the second claim of Theorem 13.8.5: if $\mathcal{E}_1$ and $\mathcal{E}_2$ are direct sums of line bundles on $\mathbb{P}^1$ then they are isomorphic if and only if the corresponding a_i are the same (up to reordering).

Warning 13.8.7. Theorem 13.8.5 does *not* mean that every short exact sequence of locally free sheaves on $\mathbb{P}^1$ splits. For example, since $H^1(\mathbb{P}^1, \mathcal{O}(-2)) \neq 0$ there is a non-split extension of $\mathcal{O}$ by $\mathcal{O}(-2)$:

$$0 \to \mathcal{O}(-2) \to \mathcal{E} \to \mathcal{O} \to 0.$$

Since this is a nonsplit extension we do not have $\mathcal{E} \cong \mathcal{O}(-2) \oplus \mathcal{O}$; for example, the LES of cohomology shows that $H^0(\mathbb{P}^1, \mathcal{E}) = 0$. In fact it is not hard to see that the only option is $\mathcal{E} \cong \mathcal{O}(-1) \oplus \mathcal{O}(-1)$.

13.8.2 Vector bundles on elliptic curves

The analogue of Theorem 13.8.5 fails for elliptic curves: not every locally free sheaf is a direct sum of line bundles. To show this we will need a basic observation:

Exercise 13.8.8. Let C be a smooth projective geometrically integral curve and let $\mathcal{L}, \mathcal{K}$ be invertible sheaves on C. Suppose that $\deg(\mathcal{L}) > \deg(\mathcal{K})$. Then there is no non-zero morphism $\phi : \mathcal{L} \to \mathcal{K}$.

We will say that a locally free sheaf is indecomposable if it cannot be written as a direct sum of locally free sheaves of smaller rank. In the following example we construct an indecomposable rank 2 locally free sheaf on an elliptic curve.

Example 13.8.9. Let C be an elliptic curve. Since $\dim H^1(C, \mathcal{O}_C) = 1$, there is a coherent sheaf $\mathcal{E}$ that fits into a non-split exact sequence

$$0 \to \mathcal{O}_C \xrightarrow{\phi} \mathcal{E} \xrightarrow{\psi} \mathcal{O}_C \to 0. \tag{13.8.1}$$

(In fact there is a unique such $\mathcal{E}$ up to isomorphism since any two non-zero elements of $H^1(C, \mathcal{O}_C)$ only differ by rescaling.) Since $\mathcal{E}$ has constant rank 2 and C is reduced we see that $\mathcal{E}$ is locally free. Furthermore, using the LES of sheaf cohomology associated to Equation 13.8.1 we see that $H^0(C, \mathcal{E})$ is 1-dimensional.

Suppose that $\mathcal{L}$ is a line bundle on C which admits an injection $\mathcal{L} \hookrightarrow \mathcal{E}$. By composing with ψ we get a morphism $\mathcal{L} \to \mathcal{O}_C$. If this morphism is zero, then $\mathcal{L}$ is the kernel of ψ and hence isomorphic to $\mathcal{O}_C$. If not, then we must have $\deg(\mathcal{L}) \leq \deg(\mathcal{O}) = 0$ by Exercise 13.8.8.

We claim that $\mathcal{E}$ cannot be written as a direct sum of line bundles. Suppose for a contradiction that $\mathcal{E} \cong \mathcal{L} \oplus \mathcal{K}$ for some invertible sheaves $\mathcal{L}, \mathcal{K}$. The argument above shows that both $\mathcal{L}$ and $\mathcal{K}$ have non-positive degree. Note that

$$\mathcal{L} \otimes \mathcal{K} \cong \bigwedge^2 \mathcal{E} = \mathcal{O}_C$$

where the last identification comes from Equation (13.8.1). Thus $\mathcal{L} \cong \mathcal{K}^\vee$ and so we see that both $\mathcal{L}$ and $\mathcal{K}$ must have degree 0.

First suppose that $\mathcal{L} \neq \mathcal{O}_C$. Then $H^0(C, \mathcal{L}) = H^0(C, \mathcal{K}) = 0$ by Proposition 12.5.8. But this contradicts the fact that $\dim H^0(X, \mathcal{E}) = 1$.

Next suppose that $\mathcal{L} \cong \mathcal{O}_C$. Then $\dim H^0(C, \mathcal{L}) = \dim H^0(C, \mathcal{K}) = 1$. But this again contradicts the fact that $\dim H^0(X, \mathcal{E}) = 1$.

A similar argument shows that if $p \in C$ is a closed point then there is a unique non-split extension

$$0 \to \mathcal{O}_C \to \mathcal{G} \to \mathcal{O}_C(p) \to 0$$

and this rank two vector bundle is not isomorphic to a direct sum of two line bundles.

It turns out that every indecomposable rank 2 vector bundle on an elliptic curve is isomorphic to one of these two types described above after twisting by a suitably chosen invertible sheaf (see [Har77, Theorem V.2.15]). In fact, the set of indecomposable locally free sheaves on an elliptic curve was classified by [Ati57, Theorem 7].

Theorem 13.8.10 ([Ati57] Theorem 7)**.** *Let C be an elliptic curve over an algebraically closed field. Let $I(r, d)$ denote the set of isomorphism classes of indecomposable locally free sheaves of rank r and degree d. Then there is a bijection between $I(r, d)$ and the closed points of C.*

We have seen this result in several contexts already. If $r = 1$, we know that the degree d line bundles can be identified with the closed points on the Jacobian $\mathrm{Jac}(C) \cong C$. If $r = 2$ and $d = 0$, then $I(2, 0)$ is the set of twists of the indecomposable bundle in Example 13.8.9 by degree 0 line bundles. If $r = 2$ and $d = 1$, then each element parametrized by $I(2, 1)$ is an extension of a degree 1 line bundle $\mathcal{O}_C(p)$ by $\mathcal{O}_C$.

13.8.3 Rank 2 vector bundles

We now turn our attention to rank 2 locally free sheaves on a curve C of arbitrary genus. The following result shows that the indecomposable rank 2 locally free sheaves on C are "bounded" in some sense.

Theorem 13.8.11. *Let C be a smooth projective curve over an algebraically closed field. Suppose that $\mathcal{E}$ is an indecomposable locally free sheaf on C of rank 2. Then there is an exact sequence*

$$0 \to \mathcal{L}_1 \to \mathcal{E} \to \mathcal{L}_2 \to 0$$

where $\mathcal{L}_1$ and $\mathcal{L}_2$ are invertible sheaves satisfying $|\deg(\mathcal{L}_2) - \deg(\mathcal{L}_1)| \leq 2g(C) - 2$.

There are stronger bounds one can put on the degrees of $\mathcal{L}_1$ and $\mathcal{L}_2$; see for example [Har77, Theorem V.2.12] and [Har77, Exercise V.2.5].

Proof. Exercise 13.8.15 guarantees that there is an invertible sheaf $\mathcal{L}_1 \subset \mathcal{E}$ which has the maximal degree amongst all invertible subsheaves of $\mathcal{E}$. The exercise also shows that the cokernel $\mathcal{L}_2$ is locally free. We verify that this choice of exact sequence has the desired property. (In fact, the argument shows that any exact sequence whose kernel and cokernel are invertible sheaves will satisfy the desired degree bound.)

Since $\mathcal{E}$ is indecomposable, we know that $\mathrm{Ext}^1(\mathcal{L}_2, \mathcal{L}_1) \neq 0$. By Exercise 13.7.13 this is the same as saying that $H^1(C, \mathcal{L}_1 \otimes \mathcal{L}_2^\vee) \neq 0$. By Proposition 12.5.7 we deduce that

$$\deg(\mathcal{L}_1) - \deg(\mathcal{L}_2) = \deg(\mathcal{L}_1 \otimes \mathcal{L}_2^\vee) \leq 2g - 2.$$

The reverse inequality follows from applying the same argument to the dual sequence. □

When we study higher rank vector bundles on curves, usually the focus changes from indecomposable bundles to "stable" bundles (see Exercise 13.8.16). The key advantage of this new perspective is that stable bundles admit a well-behaved moduli space. The study of stable sheaves and their moduli spaces is still an active topic of research today.

Remark 13.8.12. For higher dimensional varieties the set of coherent sheaves can be incredibly complicated. For example, consider the set of coherent sheaves on $\mathbb{P}^2$. Any coherent sheaf on $\mathbb{P}^2$ will have three invariants: the rank, the first chern class, and the second chern class. It is then natural to ask: what are the possible invariants of a stable sheaf on $\mathbb{P}^2$? [DLP85] shows that the invariants of stable sheaves are controlled by a "fractal curve" in the space of invariants.

13.8.4 Exercises

Exercise 13.8.13. Fix an integer n. Classify all the isomorphism types of possible extensions of $\mathcal{O}(n)$ by $\mathcal{O}$ over $\mathbb{P}^1$, or equivalently, all possible middle terms that fit in an exact sequence

$$0 \to \mathcal{O} \to \mathcal{E} \to \mathcal{O}(n) \to 0.$$

Explain how Theorem 13.7.3 relates these extensions to elements of $\mathrm{Ext}^1(\mathcal{O}, \mathcal{O}(n)) \cong H^1(\mathbb{P}^1, \mathcal{O}(n))$.

Exercise 13.8.14. Show that $H^1(\mathbb{P}^2, \Omega_{\mathbb{P}^2/\mathbb{K}}) \neq 0$. Deduce that $\Omega_{\mathbb{P}^2/\mathbb{K}}$ is an indecomposable rank 2 vector bundle on $\mathbb{P}^2$.

Exercise 13.8.15. Let C be smooth projective curve over an algebraically closed field. Let $\mathcal{E}$ be a locally free sheaf on C. Prove that there is an upper bound on the degree of all line bundles $\mathcal{L}$ which admit an injection into $\mathcal{E}$. Show that if $\mathcal{L}$ achieves this maximum value then the cokernel of $\mathcal{L} \to \mathcal{E}$ is torsion free and hence locally free. In particular we get a SES of locally free sheaves

$$0 \to \mathcal{L} \to \mathcal{E} \to \mathcal{K} \to 0.$$

Exercise 13.8.16. For higher genus curves C, one usually shifts attention from the notion of indecomposable vector bundles to the better-behaved notion of stable vector bundles. In this extended exercise we practice working with this notion.

For any locally free sheaf $\mathcal{E}$ on C, we define the slope of $\mathcal{E}$ to be

$$\mu(\mathcal{E}) = \frac{\deg(\mathcal{E})}{\mathrm{rk}(\mathcal{E})}.$$

A locally free sheaf is said to be stable if for every subbundle $\mathcal{F} \subset \mathcal{E}$ we have $\mu(\mathcal{F}) < \mu(\mathcal{E})$.

(1) Suppose we have a SES of locally free sheaves

$$0 \to \mathcal{F} \to \mathcal{G} \to \mathcal{H} \to 0.$$

Prove that $\mu(\mathcal{F}) < \mu(\mathcal{G})$ if and only if $\mu(\mathcal{G}) < \mu(\mathcal{H})$ and the same holds true if $<$ is replaced by $=$ or $>$.

(2) Show that a stable locally free sheaf is indecomposable. Give an example of an indecomposable locally free sheaf that is not stable.

(3) Let $\phi : \mathcal{E} \to \mathcal{F}$ be a morphism of stable locally free sheaves. Prove that if $\mu(\mathcal{E}) > \mu(\mathcal{F})$ then $\phi = 0$. Prove that if $\mu(\mathcal{E}) = \mu(\mathcal{F})$ and ϕ is non-zero then ϕ is an isomorphism.

13.9 Serre duality

One of the most important operations in differential geometry is integration. Recall that for a compact complex manifold X the space $H^{p,q}(X)$ (defined using Dolbeault cohomology) can be identified with the space of harmonic differential forms of type (p,q). Under this identification we have a perfect pairing

$$H^{p,q}(X) \times H^{n-p,n-q}(X) \to H^{n,n}(X) \xrightarrow{\cong} \mathbb{C}$$

where the map on the right is obtained by integration. (Recall that a perfect pairing of two vector spaces V, W is a bilinear map $V \times W \to \mathbb{K}$ which is non-degenerate. In other words, the pairing realizes $V \cong W^\vee$.)

Using the identification $H^{p,q}(X) \cong H^q(X, \Omega_X^p)$, we can rewrite the pairing above as

$$H^q(X, \Omega_X^p) \times H^{n-q}(X, \Omega_X^{n-p}) \to H^n(X, \omega_X) \xrightarrow{\cong} \mathbb{C}$$

where as usual ω_X denotes the canonical bundle $\wedge^{\dim X}\Omega_X$. In fact, we can even enrich this pairing to include the data of a vector bundle E by allowing our differential forms to take values in E.

Serre duality translates this fundamental structure into the setting of algebraic geometry. There are two versions: a general version which holds for arbitrary proper schemes over a field, and a stronger version which holds with some restrictions on the singularities. Throughout we will work over a fixed field $\mathbb{K}$.

13.9.1 Serre duality for projective space

Recall that the canonical bundle on $\mathbb{P}^n$ is isomorphic to $\mathcal{O}(-n-1)$. We computed $H^n(\mathbb{P}^n, \mathcal{O}(-n-1)) \cong \mathbb{K}$ explicitly in Theorem 12.3.1. Returning to the argument there, we see that the elements σ_i define a Čech cocycle in $C^n(\mathfrak{U}, \omega_{\mathbb{P}^n})$ which we associate with $1 \in \mathbb{K}$ under this isomorphism.

Theorem 13.9.1 (Serre duality for $\mathbb{P}^n$)**.** *Let $\mathcal{F}$ be a coherent sheaf on $\mathbb{P}^n$. For every $i \geq 0$ we have an isomorphism*

$$\operatorname{Ext}^i(\mathcal{F}, \omega_{\mathbb{P}^n}) \cong H^{n-i}(\mathbb{P}^n, \mathcal{F})^\vee.$$

Furthermore, when $i = 0$ this isomorphism is induced by the perfect pairing

$$\operatorname{Hom}(\mathcal{F}, \omega_{\mathbb{P}^n}) \times H^n(\mathbb{P}^n, \mathcal{F}) \to H^n(\mathbb{P}^n, \omega_{\mathbb{P}^n}) \cong \mathbb{K}.$$

of Remark 12.3.2.

Due to its prominent role in this theorem, $\omega_{\mathbb{P}^n}$ is sometimes called the "dualizing sheaf" of $\mathbb{P}^n$. When $\mathcal{F}$ is locally free, by Exercise 13.7.13 this reduces to the more commonly used form $H^i(\mathbb{P}^n, \mathcal{F}^\vee \otimes \omega_{\mathbb{P}^n}) \cong H^{n-i}(\mathbb{P}^n, \mathcal{F})$.

Proof. Our plan is to use Theorem 13.0.8 to show that the functor $H^{n-i}(\mathbb{P}^n, -)^\vee$ coincides with the functor $\mathrm{Ext}^i(-, \omega_{\mathbb{P}^n})$ as functors from $\mathbf{Coh}(\mathbb{P}^n)^{op} \to \mathbb{K}{-}\mathbf{Mod}$. In particular, we must show that both form effaceable δ-functors and that their 0th degree functors coincide.

We first show that these two functors coincide when $i = 0$, i.e. that $\mathrm{Hom}(-, \omega_{\mathbb{P}^n})$ is isomorphic to $H^n(\mathbb{P}^n, -)^\vee$. This is done by a direct calculation. First consider the invertible sheaves $\mathcal{O}(d)$. By the functoriality of cohomology we obtain a pairing

$$\mathrm{Hom}(\mathcal{O}(d), \omega_{\mathbb{P}^n}) \times H^n(\mathbb{P}^n, \mathcal{O}(d)) \to H^n(\mathbb{P}^n, \omega_{\mathbb{P}^n})$$

and the proof of Theorem 12.3.1 shows that this pairing is perfect (see Remark 12.3.2). We immediately deduce the desired statement when $\mathcal{F}$ is a direct sum of line bundles. For a general coherent sheaf $\mathcal{F}$, we have an exact sequence

$$\mathcal{E}_1 \to \mathcal{E}_0 \to \mathcal{F} \to 0$$

where both $\mathcal{E}_1$ and $\mathcal{E}_0$ are direct sums of line bundles. Since both $\mathrm{Hom}(-, \omega_{\mathbb{P}^n})$ and $H^n(\mathbb{P}^n, -)^\vee$ are contravariant left exact, the desired statement for $\mathcal{F}$ follows from the established cases $\mathcal{E}_1, \mathcal{E}_0$ and the 5-lemma.

The $H^{n-i}(\mathbb{P}^n, -)^\vee$ have the structure of a δ-functor, since the dual "reverses" the LES for sheaf cohomology. Similarly, the $\mathrm{Ext}^i(-, \omega_{\mathbb{P}^n})$ form a δ-functor by Proposition 13.7.5.

Finally, we must show that these functors are effaceable. Due to the contravariance of the construction, we must show that any coherent sheaf $\mathcal{F}$ admits a surjection from a coherent sheaf $\mathcal{E}$ for which the higher functors vanish. For any coherent sheaf $\mathcal{F}$ there is a direct sum of line bundles $\mathcal{O}(d)^{\oplus r}$ which surjects onto $\mathcal{F}$ by Corollary 10.6.13. Furthermore, for any m there is a surjection $\mathcal{O}(m-1)^{\oplus 2} \to \mathcal{O}(m)$ given by multiplying by x on the first factor and y on the second. Using this modification repeatedly, we obtain a surjection $\mathcal{O}(e)^{\oplus q} \to \mathcal{F}$ where $e < -n-1$. Then all the higher functors vanish on $\mathcal{O}(e)^{\oplus q}$. □

13.9.2 Serre duality for regular varieties

When X is a regular projective $\mathbb{K}$-variety, Serre duality has essentially the same form:

Theorem 13.9.2 (Serre duality for regular varieties)**.** *Let X be a regular projective $\mathbb{K}$-variety of dimension n and let $\mathcal{F}$ be a coherent sheaf on X. For every $i \geq 0$ we have an isomorphism*

$$\mathrm{Ext}^i(\mathcal{F}, \omega_X) \cong H^{n-i}(X, \mathcal{F})^\vee.$$

Furthermore, when $i = 0$ this isomorphism is induced by a perfect pairing

$$\mathrm{Hom}(\mathcal{F}, \omega_X) \times H^n(X, \mathcal{F}) \to H^n(X, \omega_X) \cong \mathbb{K}.$$

Remark 13.9.3. Note that $\Omega_X^p \cong (\Omega_X^{n-p})^\vee \otimes \omega_X$ so that our current formulation really does extend the pairing of Dolbeault cohomology groups described earlier in the section.

Remark 13.9.4. When $\mathcal{F}$ is locally free a stronger statement is true. In this case Exercise 13.7.13 identifies $\mathrm{Ext}^i(\mathcal{F}, \omega_X) \cong H^i(X, \mathcal{F}^\vee \otimes \omega_X)$. Then the isomorphism of Theorem 13.9.2 can be rephrased as the existence of a perfect pairing

$$H^i(X, \mathcal{F}^\vee \otimes \omega_X) \times H^{n-i}(X, \mathcal{F}) \to H^n(X, \omega_X) \cong \mathbb{K}$$

where the pairing is the composition of the cup product in Exercise 12.2.10 and the map $H^n(X, (\mathcal{F} \otimes \mathcal{F}^\vee) \otimes \omega_X) \to H^n(X, \omega_X)$ induced by the canonical map $\mathcal{F} \otimes \mathcal{F}^\vee \to \mathcal{O}_X$.

The isomorphism $H^n(X, \omega_X) \cong \mathbb{K}$ in Theorem 13.9.2 called the trace map. Although we have not defined it explicitly, one should view it as an analogue of map obtained by combining the Dolbeault isomorphism with the integration of a volume form. In some cases we can be more explicit about this map:

Example 13.9.5. Let C be a regular projective curve and let Ω_C denote its cotangent bundle. We will explain how the trace map $H^1(C, \Omega_C) \to \mathbb{K}$ is related to the classical theory of residues of differential forms.

Let $\eta \in C$ be the generic point of C. We will identify the stalk $\Omega_{C,\eta}$ as the set of meromorphic differential forms on C. For a closed point p, the stalk $\Omega_{C,p} \subset \Omega_{C,\eta}$ is the set of meromorphic forms which are regular at p. One can show that for every closed point $p \in C$ there is a $\mathbb{K}$-linear map $\mathrm{res}_p : \Omega_{C,\eta} \to \mathbb{K}$ satisfying the following properties:

(1) If $\tau \in \Omega_{C,p}$ then $\mathrm{res}_p(\tau) = 0$.

(2) If we choose a uniformizer t of the local ring $\mathcal{O}_{C,p}$, then any element $\tau \in \Omega_{C,\eta}$ can be written as $g\,dt$ for some $g \in K(C)$. Then $\mathrm{res}_p(g)$ is the coefficient of t^{-1} in the expansion of g as a rational function of t.

The Residue Theorem states that for any $\tau \in \Omega_{C,\eta}$ we have $\sum_{p \in C} \mathrm{res}_p(\tau) = 0$.

Consider the exact sequence

$$0 \to \mathcal{O}_C \to \mathcal{K}_C \to \mathcal{K}_C/\mathcal{O}_C \to 0$$

where $\mathcal{K}_C$ is the locally constant sheaf with value $K(C)$. The two sheaves on the right are flasque, hence acyclic for the global sections functor. Furthermore, the rightmost term can be identified more explicitly as a direct sum of skyscraper sheaves

$$\mathcal{K}_C/\mathcal{O}_C \cong \oplus_{p \in C}(K(C)/\mathcal{O}_{C,p})(p).$$

If we tensor our exact sequence by Ω_C, the LES of cohomology yields

$$\begin{tikzcd}
H^0(C, \Omega_C \otimes \mathcal{K}_C) \arrow[r] \arrow[d, leftrightarrow, "\cong"] & H^0(C, \Omega_C \otimes \mathcal{K}_C/\mathcal{O}_C) \arrow[r] \arrow[d, leftrightarrow, "\cong"] & H^1(C, \Omega_C) \arrow[r] & 0 \\
\Omega_{C,\eta} & \oplus_{p \in C} \Omega_{C,\eta}/\Omega_p & &
\end{tikzcd}$$

We can define a map $\oplus_{p\in C}\Omega_{C,\eta}/\Omega_{C,p} \to \mathbb{K}$ by taking the sums of the residue maps $\mathrm{res}_p : \Omega_{C,\eta}/\Omega_{C,p} \to \mathbb{K}$. By the Residue Theorem this map vanishes on the image of $\Omega_{C,\eta}$, and thus the residue map descends to define $H^1(X, \Omega_C) \to \mathbb{K}$. This turns out to be the isomorphism we are looking for.

13.9.3 Dualizing sheaves

Definition 13.9.6. Let X be a proper $\mathbb{K}$-scheme of pure dimension n. A coherent sheaf ζ_X is said to be a dualizing sheaf if there is a trace morphism $t : H^n(X, \zeta_X) \to \mathbb{K}$ such that for any coherent sheaf $\mathcal{F}$ on X the composition of the pairing with the trace map

$$\mathrm{Hom}(\mathcal{F}, \zeta_X) \times H^n(X, \mathcal{F}) \to H^n(X, \zeta_X) \xrightarrow{t} \mathbb{K}$$

is a perfect pairing of $\mathbb{K}$-vector spaces.

Exercise 13.9.7. Suppose that ζ_X and ζ'_X are dualizing sheaves for X. Prove that there is a unique isomorphism $\phi : \zeta_X \to \zeta'_X$ such that $t' = t \circ H^n(\phi)$.

Theorem 13.9.8. *Any proper $\mathbb{K}$-scheme of pure dimension n admits a dualizing sheaf.*

When X is a projective $\mathbb{K}$-scheme, one can derive this statement in a very formal way from Serre duality for $\mathbb{P}^n$. First, by embedding X into some projective space and taking generic projections, one obtains a finite morphism $f : X \to \mathbb{P}^n$ for some n. Second, given an affine morphism $f : X \to Y$ one constructs a functor $f^!$ which is right adjoint to the pushforward f_*. (This is surprising: usually we think of f_* as a right adjoint, and in general there is no reason to expect it to be a left adjoint.) Finally, using the formal properties of adjoint functors one shows that the dualizing sheaf ζ_X can be obtained by applying $f^!$ to the dualizing sheaf $\omega_{\mathbb{P}^n}$.

Note that in general we only obtain a pairing of Hom and H^n, not of the higher Ext functors and the H^{n-i}. It turns out that this stronger version holds when X is Cohen-Macaulay. The key is the Miracle Flatness theorem, showing that when X is Cohen-Macaulay the finite morphism $f : X \to \mathbb{P}^n$ constructed above is flat.

13.9.4 Exercises

Exercise 13.9.9. Let $X = \mathbb{P}^n$. Show that for $0 \le p, q \le n$ we have

$$H^q(X, \Omega^p_X) = \begin{cases} 0 \text{ when } p \neq q \\ \mathbb{K} \text{ when } p = q \end{cases}$$

(Compare this result against the singular cohomology groups or Hodge diamond of $\mathbb{P}^n$.)

What's next

The material we have covered is roughly equivalent to the other "standard" introductory texts in algebraic geometry, such as Hartshorne's book or Ravi Vakil's notes. However, there is no standard next step in learning algebraic geometry. I wanted to briefly discuss some good options of a subject and text for a second course. (Of course my recommendations will be based on my own knowledge and thus will leave out some excellent books I am not familiar with.)

First, over the summer you should learn:

- Basic theory of curves and surfaces

  For a quick introduction, I don't know of a better resource than chapters IV and V of Hartshorne's textbook. These are not as comprehensive as other books devoted to the subject but will give you a basic set of examples you will need to know for future work.

Next, if you are planning to work with one of the "classical" algebraic geometers in our department, then there are a couple topics you should learn no matter who you decide to work with. These include:

- Positivity in algebraic geometry

  Lazarsfeld's book "Positivity in Algebraic Geometry I & II" is widely viewed as one of the best "second" texts in algebraic geometry (and is my favorite choice). We have seen that ample invertible sheaves satisfy a number of special properties. Lazarsfeld's book discusses how these ideas can be generalized and extended to other types of divisors and locally free sheaves.

- Intersection theory

  There are two standard texts: Fulton's "Intersection Theory" (which is excellent for understanding the theory but doesn't focus on examples) and Harris and Eisenbud's "3264 and all that" (which is great for examples but doesn't explain the theory). I think it would be best to work through both books together; if you just use one you miss out on a valuable part of this vast subject.

There are a number of additional topics worth learning depending on your eventual interests:

- Riemann surfaces

  Since Riemann surfaces are foundational in most areas of theoretical mathematics, it is worth seeing them from a non-algebraic perspective. Many algebraic geometers prefer Miranda's "Algebraic Curves and Riemann Surfaces" since it retains much of the flavor of the algebraic theory. A more advanced book on algebraic curves is "Geometry of Algebraic Curves" by Arabello, Cornalba, Griffiths, and Harris.

- Algebraic surfaces

  There are two main goals of the theory of algebraic surfaces. First, one would like to study the classification of surfaces – can we describe all surfaces of a certain type? What are the topological constraints on surfaces? Second, one would like to get a deeper understanding of special types of surfaces (K3 surfaces, general type surfaces with $g = q = 0$, etc.). Beauville's book "Complex Algebraic Surfaces" gives a concise introduction to the subject. Other options include "Algebraic Surfaces" by Badescu or "Compact Complex Surfaces" by Barth, Hulek, Peters, van de Ven.

- Topology of algebraic varieties

  Here the goal would be to learn some of the foundations of the topology of complex algebraic varieties – Hodge theory, the Lefschetz theorems, etc. In my view the best option is Voisin's books "Hodge Theory and Complex Algebraic Geometry I & II". Some more introductory options include "Complex Geometry" by Huybrechts and "Principles of Algebraic Geometry" by Griffiths and Harris.

- Toric geometry

  Toric varieties are built up out of simple affine pieces that are glued together in a combinatorial way. (The prototype is projective space where we encode the combinatorics of the gluing data using graded algebra.) Thus toric varieties are one of the few classes of algebraic varieties for which we can easily compute things. The traditional references are "Introduction to Toric Varieties" by Fulton and "Toric Varieties" by Cox, Little, Schenck.

- Minimal model program

  The Minimal Model Program studies how the geometry of algebraic varieties is controlled by the curvature of their canonical bundle. The classical reference is "Birational Geometry of Algebraic Varieties" by Kollár and Mori. The book "Introduction to the Mori Program" by Matsuki is another option which keeps the discussion more elementary. Finally, the book "Higher-Dimensional Algebraic Geometry" by Debarre is a very clear introduction to issues related to the MMP (and also discusses the MMP briefly). See [Alv] for more resources.

- Moduli spaces of curves

  The moduli space of curves is a basic object in algebraic geometry and complex geometry. There is a standard set of tools used to construct and study its properties. One good introduction is "Moduli of curves" by Harris and Morrison, but this book has the unfortunate disadvantage of not explaining the underlying theory in a rigorous way. The book "Rational curves on algebraic varieties" by Kollár carefully discusses some foundational issues about various moduli spaces of curves (including the Chow and Hilbert schemes). The book "Geometric Invariant Theory" by Mumford, Fogarty, Kirwan gives a construction of the moduli space.

- Abelian varieties

  Abelian varieties will be particularly important for the number theorists in our department, but every algebraic geometer would benefit from an understanding of their basic properties. [Moo] is an excellent set of online notes, but doesn't appear to be completed. When I was a graduate student the book "Abelian Varieties" by Mumford was the most popular choice.

- Derived categories of sheaves

  Recently there has been a lot of interest in understanding how derived categories of sheaves reflect the geometry of the ambient variety. I think the best introduction is "Fourier-Mukai Transforms in Algebraic Geometry" by Huybrechts. You can find some more resources at [Ver].

Bibliography

[Alv] Javier Alvarez. Training towards research on birational geometry/minimal model program. MathOverflow. `https://mathoverflow.net/q/114050`.

[Ati57] M. F. Atiyah. Vector bundles over an elliptic curve. *Proc. London Math. Soc. (3)*, 7:414–452, 1957.

[Bra] Martin Brandenburg. Colimits of schemes. MathOverflow. `https://mathoverflow.net/q/9961`.

[CEH] CEH. Why is the category of finitely generated modules over a non-noetherian ring not abelian? Mathematics Stack Exchange. `https://math.stackexchange.com/q/1857330`.

[Cla] Pete Clark. When are there enough projective sheaves on a space x? MathOverflow. `https://mathoverflow.net/q/5378`.

[DLP85] J.-M. Drezet and J. Le Potier. Fibrés stables et fibrés exceptionnels sur $\mathbf{P}_2$. *Ann. Sci. École Norm. Sup. (4)*, 18(2):193–243, 1985.

[EO20] Geir Ellingsrud and John Christian Ottem. Introduction to Schemes. `https://www.uio.no/studier/emner/matnat/math/MAT4215/data/masteragbook.pdf`, 2020.

[Ful89] William Fulton. *Algebraic curves*. Advanced Book Classics. Addison-Wesley Publishing Company, Advanced Book Program, Redwood City, CA, 1989.

[Gat20] Andreas Gathmann. Algebraic Geometry Class Notes. `https://www.mathematik.uni-kl.de/~gathmann/de/alggeom.php`, 2020.

[GW10] Ulrich Görtz and Torsten Wedhorn. *Algebraic geometry I*. Advanced Lectures in Mathematics. Vieweg + Teubner, Wiesbaden, 2010.

[Har66] Robin Hartshorne. *Residues and duality*. Lecture notes of a seminar on the work of A. Grothendieck, given at Harvard 1963/64. With an appendix by P. Deligne. Lecture Notes in Mathematics, No. 20. Springer-Verlag, Berlin-New York, 1966.

[Har77] Robin Hartshorne. *Algebraic geometry*. Springer-Verlag, New York-Heidelberg, 1977. Graduate Texts in Mathematics, No. 52.

[Har95] Joe Harris. *Algebraic geometry*, volume 133 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995. A first course, Corrected reprint of the 1992 original.

[Kat99] Nicholas M. Katz. Space filling curves over finite fields. *Math. Res. Lett.*, 6(5-6):613–624, 1999.

[Kle79] Steven L. Kleiman. Misconceptions about K_x. *Enseign. Math. (2)*, 25(3-4):203–206 (1980), 1979.

[ML39] Saunders Mac Lane. Modular fields. I. Separating transcendence bases. *Duke Math. J.*, 5(2):372–393, 1939.

[Mon17] Pinaki Mondal. When is the intersection of two finitely generated subalgebras of a polynomial ring also finitely generated? *Arnold Math. J.*, 3(3):333–350, 2017.

[Moo] Ben Moonen. Abelian varieties. `http://www.math.ru.nl/~bmoonen/research.html#bookabvar`.

[Rob] Robert. When is the category of finitely presented modules abelian? MathOverflow. `https://mathoverflow.net/q/381984`.

[Sch] Karl Schwede. Generalized divisors and reflexive sheaves. `https://www.math.utah.edu/~schwede/Notes/GeneralizedDivisors.pdf`.

[Sta] Charles Staats. Cartier divisors. `https://math.uchicago.edu/~cstaats/Charles_Staats_III/Notes_and_papers_files/CartierDivisors.pdf`.

[Sta15] The Stacks Project Authors. *Stacks Project*. `http://stacks.math.columbia.edu`, 2015.

[Vak17] Ravi Vakil. The Rising Sea. `http://math.stanford.edu/~vakil/216blog/`, 2017.

[Ver] J Verma. Derived categories of coherent sheaves: suggested references? MathOverflow. `https://mathoverflow.net/q/42463`.

[Zho] Jason Zhou. coherent modules. MathOverflow. `https://mathoverflow.net/q/159836`.