

7. MORE ABOUT SHEAVES

We present a detailed study of sheaves on a scheme X , in particular sheaves of O_X -modules. For any presheaf \mathcal{F}' on X there is an associated sheaf \mathcal{F} that describes “the same objects as \mathcal{F}' but with the conditions on the sections made local”. This allows us to define sheaves by constructions that would otherwise only yield presheaves. We can thus construct e. g. direct sums of sheaves, tensor products, kernels and cokernels of morphisms of sheaves, as well as push-forwards and pull-backs along morphisms of schemes.

A sheaf of O_X -modules is called quasi-coherent if it is induced by an R -module on every affine open subset $U = \text{Spec } R$ of X . Almost all sheaves that we will consider are of this form. This reduces local computations regarding these sheaves to computations in commutative algebra.

A quasi-coherent sheaf on X is called locally free of rank r if it is locally isomorphic to $O_X^{\oplus r}$. Locally free sheaves are the most well-behaved sheaves; they correspond to vector bundles in topology. Any construction and theorem valid for vector spaces can be carried over to the category of locally free sheaves. Locally free sheaves of rank 1 are called line bundles.

For any morphism $f : X \rightarrow Y$ we define the sheaf of relative differential forms $\Omega_{X/Y}$ on X relative Y . The most important case is when Y is a point, in which case we arrive at the sheaf Ω_X of differential forms on X . It is locally free of rank $\dim X$ if and only if X is smooth. In this case, its top alternating power $\Lambda^{\dim X} \Omega_X$ is a line bundle ω_X called the canonical bundle. On a smooth projective curve it has degree $2g - 2$, where g is the genus of the curve.

On every smooth curve X the line bundles form a group which is isomorphic to the Picard group $\text{Pic } X$ of divisor classes. A line bundle together with a collection of sections that do not vanish simultaneously at any point determines a morphism to projective space.

If $f : X \rightarrow Y$ is a morphism of smooth projective curves, the Riemann-Hurwitz formula states that the canonical bundles of X and Y are related by $\omega_X = f^* \omega_Y \otimes O_X(R)$, where R is the ramification divisor. For any smooth projective curve X of genus g and any divisor D the Riemann-Roch theorem states that $h^0(D) - h^0(K_X - D) = \deg D + 1 - g$, where $h^0(D)$ denotes the dimension of the space of global sections of the line bundle $O(D)$ associated to D .


7.1. Sheaves and sheafification. The first thing we have to do to discuss the more advanced topics mentioned in section 6.6 is to get a more detailed understanding of sheaves. Recall from section 2.2 that we defined a sheaf to be a structure on a topological space X that describes “function-like” objects that can be patched together from local data. Let us first consider an informal example of a sheaf that is not just the sheaf of regular functions on a scheme.

Example 7.1.1. Let X be a smooth complex curve. For any open subset $U \subset X$, we have seen that the ring of regular functions $O_X(U)$ on U can be thought of as the ring of complex-valued functions $\varphi : U \rightarrow \mathbb{C}, P \mapsto \varphi(P)$ “varying nicely” (i. e. as a rational function) with P .

Now consider the “tangent sheaf” T_X , i. e. the sheaf “defined” by

$$T_X(U) = \{\varphi = (\varphi(P))_{P \in U} ; \varphi(P) \in T_{X,P} \text{ “varying nicely with } P \text{”}\}$$

(of course we will have to make precise what “varying nicely” means). In other words, a section $\varphi \in T_X(U)$ is just given by specifying a tangent vector at every point in U . As an example, here is a picture of a section of $T_{\mathbb{P}^1}(\mathbb{P}^1)$:



As the tangent spaces $T_{X,P}$ are all one-dimensional complex vector spaces, $\phi(P)$ can again be thought of as being specified by a single complex number, just as for the structure sheaf \mathcal{O}_X . The important difference (that is already visible from the definition above) is that these one-dimensional vector spaces vary with P and thus *have no canonical identification with the complex numbers*. For example, it does not make sense to talk about “the tangent vector 1” at a point P . Consequently, there is no analogue of “constant functions” for sections of the tangent sheaf. In fact, we will see in lemma 7.4.15 that every global section of $T_{\mathbb{P}^1}$ has two zeros, so there is really no analogue of constant functions. (In the picture above, the north pole of the sphere is a point where the section of $T_{\mathbb{P}^1}$ would be ill-defined if we do not choose a section in which the lengths of the tangent vectors approach zero towards the north pole.) Hence we have seen that the tangent sheaf of \mathbb{P}^1 is a sheaf that is not isomorphic to the structure sheaf $\mathcal{O}_{\mathbb{P}^1}$ although its sections are given locally by “one complex number varying nicely”.

(We should mention that the above property of \mathbb{P}^1 is purely topological: there is not even a continuous nowhere-zero tangent field on the unit sphere in \mathbb{R}^3 . This is usually called the “hairy ball theorem” and stated as saying that “you cannot comb a hedgehog (i. e. a ball) without a bald spot”.)

Let us now get more rigorous. Recall that a presheaf of rings \mathcal{F} on a topological space X was defined to be given by the data:

- for every open set $U \subset X$ a ring $\mathcal{F}(U)$,
- for every inclusion $U \subset V$ of open sets in X a ring homomorphism $\rho_{V,U} : \mathcal{F}(V) \rightarrow \mathcal{F}(U)$ called the restriction map,

such that

- $\mathcal{F}(\emptyset) = 0$,
- $\rho_{U,U}$ is the identity map for all U ,
- for any inclusion $U \subset V \subset W$ of open sets in X we have $\rho_{V,U} \circ \rho_{W,V} = \rho_{W,U}$.

The elements of $\mathcal{F}(U)$ are then called the sections of \mathcal{F} over U , and the restriction maps $\rho_{V,U}$ are written as $f \mapsto f|_U$. The space of global sections $\mathcal{F}(X)$ is often denoted $\Gamma(\mathcal{F})$.

A presheaf \mathcal{F} of rings is called a sheaf of rings if it satisfies the following glueing property: if $U \subset X$ is an open set, $\{U_i\}$ an open cover of U and $f_i \in \mathcal{F}(U_i)$ sections for all i such that $f_i|_{U_i \cap U_j} = f_j|_{U_i \cap U_j}$ for all i, j , then there is a unique $f \in \mathcal{F}(U)$ such that $f|_{U_i} = f_i$ for all i . In other words, sections of a sheaf can be patched from compatible local data.

The same definition applies equally to categories other than rings, e. g. we can define sheaves of Abelian groups, k -algebras, and so on. For a ringed space (X, \mathcal{O}_X) , e. g. a scheme, we can also define **sheaves of \mathcal{O}_X -modules** in the obvious way: every $\mathcal{F}(U)$ is required to be an $\mathcal{O}_X(U)$ -module, and these module structures have to be compatible with

the restriction maps in the obvious sense. For example, the tangent sheaf of example 7.1.1 on a curve X is a sheaf of \mathcal{O}_X -modules: “sections of the tangent sheaf can be multiplied with regular functions”.

Example 7.1.2. Let $X \subset \mathbb{P}^N$ be a projective variety over an algebraically closed field k , and let $S(X) = S = \bigoplus_{d \geq 0} S^{(d)}$ be its homogeneous coordinate ring. For any integer n , let $K(n)$ be the n -th graded piece of the localization of S at the non-zero homogeneous elements, i. e.

$$K(n) = \left\{ \frac{f}{g} ; f \in S^{(d+n)}, g \in S^{(d)} \text{ for some } d \geq 0 \text{ and } g \neq 0 \right\}.$$

Now for any $P \in X$ and open set $U \subset X$ we set

$$\mathcal{O}_X(n)_P = \left\{ \frac{f}{g} \in K(n) ; g(P) \neq 0 \right\} \quad \text{and} \quad \mathcal{O}_X(n)(U) = \bigcap_{P \in U} \mathcal{O}_X(n)_P.$$

For $n = 0$ this is precisely the definition of the structure sheaf, so $\mathcal{O}_X(0) = \mathcal{O}_X$. In general, $\mathcal{O}_X(n)$ is a sheaf of \mathcal{O}_X -modules whose sections can be thought of as “functions” of degree n in the homogeneous coordinates of X . For example:

- (i) Every homogeneous polynomial of degree n defines a global section of $\mathcal{O}_X(n)$.
- (ii) There are no global sections of $\mathcal{O}_X(n)$ for $n < 0$.
- (iii) In \mathbb{P}^1 with homogeneous coordinates x_0, x_1 , we have

$$\frac{1}{x_0} \in \mathcal{O}_{\mathbb{P}^1}(-1)(U)$$

for $U = \{(x_0 : x_1) ; x_0 \neq 0\}$.

Note that on the distinguished open subset X_{x_i} (where x_i are the coordinates of \mathbb{P}^N) the sheaf $\mathcal{O}_X(n)$ is isomorphic to the structure sheaf \mathcal{O}_X : for every open subset $U \subset X_{x_i}$ the maps

$$\mathcal{O}_X(U) \rightarrow \mathcal{O}_X(n)(U), \varphi \mapsto \varphi \cdot x_i^n \quad \text{and} \quad \mathcal{O}_X(n)(U) \rightarrow \mathcal{O}_X(U), \varphi \mapsto \frac{\varphi}{x_i^n}$$

give an isomorphism, hence $\mathcal{O}_X(n)|_{X_{x_i}} \cong \mathcal{O}_X|_{X_{x_i}}$. So $\mathcal{O}_X(n)$ is *locally* isomorphic to the structure sheaf, but not *globally*. (This is the same situation as for the tangent sheaf of a smooth curve in example 7.1.1.)

The sheaves $\mathcal{O}(n)$ on a projective variety (or more generally on a projective scheme) are called the **twisting sheaves**. They are probably the most important sheaves after the structure sheaf.

If we want to deal with more general sheaves, we certainly need to set up a suitable category, i. e. we have to define morphisms of sheaves, kernels, cokernels, and so on. Let us start with some simple definitions.

Definition 7.1.3. Let X be a topological space. A **morphism** $f : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ of presheaves of abelian groups (or rings, sheaves of \mathcal{O}_X -modules etc.) on X is a collection of group homomorphisms (resp. ring homomorphisms, $\mathcal{O}_X(U)$ -module homomorphisms etc.) $f_U : \mathcal{F}_1(U) \rightarrow \mathcal{F}_2(U)$ for every open subset $U \subset X$ that commute with the restriction maps, i. e. the diagram

$$\begin{array}{ccc} \mathcal{F}_1(U) & \xrightarrow{\rho_{U,V}} & \mathcal{F}_1(V) \\ \downarrow f_U & & \downarrow f_V \\ \mathcal{F}_2(U) & \xrightarrow{\rho_{U,V}} & \mathcal{F}_2(V) \end{array}$$

is required to be commutative.

Example 7.1.4. If $X \subset \mathbb{P}^N$ is a projective variety and $f \in k[x_0, \dots, x_N]$ is a homogeneous polynomial of degree d , we get morphisms of sheaves of \mathcal{O}_X -modules

$$\mathcal{O}_X(n) \rightarrow \mathcal{O}_X(n+d), \quad \varphi \mapsto f \cdot \varphi$$

for all n .

Definition 7.1.5. If $f : X \rightarrow Y$ is a morphism of topological spaces and \mathcal{F} is a sheaf on X , then we define the **push-forward** $f_*\mathcal{F}$ of \mathcal{F} to be the sheaf on Y given by $f_*\mathcal{F}(U) = \mathcal{F}(f^{-1}(U))$ for all open subsets $U \subset Y$.

Example 7.1.6. By definition, a morphism $f : X \rightarrow Y$ of ringed spaces comes equipped with a morphism of sheaves $\mathcal{O}_Y \rightarrow f_*\mathcal{O}_X$. This is exactly given by the data of the pull-back morphisms $\mathcal{O}_Y(U) \rightarrow \mathcal{O}_X(f^{-1}(U))$ for all open subsets $U \subset Y$ (see definition 5.2.1).

Definition 7.1.7. Let $f : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ be a morphism of sheaves of e. g. Abelian groups on a topological space X . We define the **kernel** $\ker f$ of f by setting

$$(\ker f)(U) = \ker(f_U : \mathcal{F}_1(U) \rightarrow \mathcal{F}_2(U)).$$

We claim that $\ker f$ is a sheaf on X . In fact, it is easy to see that $\ker f$ with the obvious restriction maps is a presheaf. Now let $\{U_i\}$ be an open cover of an open subset $U \subset X$, and assume we are given $\varphi_i \in \ker(\mathcal{F}_1(U_i) \rightarrow \mathcal{F}_2(U_i))$ that agree on the overlaps $U_i \cap U_j$. In particular, the φ_i are then in $\mathcal{F}_1(U_i)$, so we get a unique $\varphi \in \mathcal{F}_1(U)$ with $\varphi|_{U_i} = \varphi_i$ as \mathcal{F}_1 is a sheaf. Moreover, $f(\varphi_i) = 0$, so $(f(\varphi))|_{U_i} = 0$ by definition 7.1.3. As \mathcal{F}_2 is a sheaf, it follows that $f(\varphi) = 0$, so $\varphi \in \ker f$.

What the above argument boils down to is simply that the property of being in the kernel, i. e. of being mapped to zero under a morphism, is a *local* property — a function is zero if it is zero on every subset of an open cover. So the kernel is again a sheaf.

Remark 7.1.8. Now consider the dual case to definition 7.1.7, namely cokernels. Again let $f : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ be a morphism of sheaves of e. g. Abelian groups on a topological space X . As above we define a presheaf $\text{coker}' f$ by setting

$$(\text{coker}' f)(U) = \text{coker}(f_U : \mathcal{F}_1(U) \rightarrow \mathcal{F}_2(U)) = \mathcal{F}_2(U)/\text{im } f_U.$$

Note however that $\text{coker}' f$ is *not* a sheaf. To see this, consider the following example. Let $X = \mathbb{A}^1 \setminus \{0\}$, $Y = \mathbb{A}^2 \setminus \{0\}$, and let $i : X \rightarrow Y$ be the inclusion morphism $(x_1) \mapsto (x_1, 0)$. Let $i^\# : \mathcal{O}_Y \rightarrow i_*\mathcal{O}_X$ be the induced morphisms of sheaves on Y of example 7.1.6, and consider the presheaf $\text{coker}' i^\#$ on Y . Look at the cover of Y by the affine open subsets $U_1 = \{x_1 \neq 0\} \subset Y$ and $U_2 = \{x_2 \neq 0\} \subset Y$. Then the maps

$$k \left[x_1, \frac{1}{x_1}, x_2 \right] = \mathcal{O}_Y(U_1) \rightarrow \mathcal{O}_X(U_1 \cap X) = k \left[x_1, \frac{1}{x_1} \right]$$

and

$$k \left[x_1, x_2, \frac{1}{x_2} \right] = \mathcal{O}_Y(U_2) \rightarrow \mathcal{O}_X(U_2 \cap X) = 0$$

are surjective, hence $(\text{coker}' i^\#)(U_1) = (\text{coker}' i^\#)(U_2) = 0$. But on global sections the map

$$k[x_1, x_2] = \mathcal{O}_Y(Y) \rightarrow \mathcal{O}_X(X) = k \left[x_1, \frac{1}{x_1} \right]$$

is not surjective, hence $(\text{coker}' i^\#)(Y) \neq 0$. This shows that $\text{coker}' i^\#$ cannot be a sheaf — the zero section on the open cover $\{U_1, U_2\}$ has no *unique* extension to a global section on Y .

What the above argument boils down to is simply that being in the cokernel of a morphism, i. e. of being a quotient in $\mathcal{F}_2(U)/\text{im } f_U$, is *not* a local property — it is a question about finding a global section of \mathcal{F}_2 on U that cannot be answered locally.

Example 7.1.9. Here is another example showing that quite natural constructions involving sheaves often lead to only presheaves because the constructions are not local. Let $X \subset \mathbb{P}^N$ be a projective variety. Consider the tensor product presheaf of the sheaves $\mathcal{O}_X(1)$ and $\mathcal{O}_X(-1)$, defined by

$$(\mathcal{O}_X(1) \otimes' \mathcal{O}_X(-1))(U) = \mathcal{O}_X(1)(U) \otimes_{\mathcal{O}_X(U)} \mathcal{O}_X(-1)(U).$$

As $\mathcal{O}_X(1)$ describes “functions” of degree 1 and $\mathcal{O}_X(-1)$ “functions” of degree -1 , we expect products of them to be true functions of pure degree 0 in the homogeneous coordinates of X . In other words, the tensor product of $\mathcal{O}_X(1)$ with $\mathcal{O}_X(-1)$ should just be the structure sheaf \mathcal{O}_X . However, $\mathcal{O}_X(1) \otimes' \mathcal{O}_X(-1)$ is not even a sheaf: consider the case $X = \mathbb{P}^1$ and the open subsets $U_0 = \{x_0 \neq 0\}$ and $U_1 = \{x_1 \neq 0\}$. On these open subsets we have the sections

$$\begin{aligned} x_0 \otimes \frac{1}{x_0} &\in (\mathcal{O}_X(1) \otimes' \mathcal{O}_X(-1))(U_0) \\ \text{and } x_1 \otimes \frac{1}{x_1} &\in (\mathcal{O}_X(1) \otimes' \mathcal{O}_X(-1))(U_1). \end{aligned}$$

Obviously, both these local sections are the constant function 1, so in particular they agree on the overlap $U_0 \cap U_1$. But there is no global section in $\mathcal{O}_X(1)(X) \otimes_{\mathcal{O}_X(X)} \mathcal{O}_X(-1)(X)$ that corresponds to the constant function 1, as $\mathcal{O}_X(-1)$ has no non-zero global sections at all.

The way out of this trouble is called sheafification. This means that for any presheaf \mathcal{F}' there is an associated sheaf \mathcal{F} that is “very close” to \mathcal{F}' and that should usually be the object that one wants. Intuitively speaking, if the sections of a presheaf are thought of as function-like objects satisfying some conditions, then the associated sheaf describes the same objects *with the conditions on them made local*. In particular, if we look at \mathcal{F}' locally, i. e. at the stalks, then we should not change anything; it is just the global structure that changes. We have done this construction quite often already without explicitly saying so, e. g. in the construction of the structure sheaf of schemes in definition 5.1.11. Here is the general construction:

Definition 7.1.10. Let \mathcal{F}' be a presheaf on a topological space X . The **sheafification** of \mathcal{F}' , or the sheaf **associated to** the presheaf \mathcal{F}' , is defined to be the sheaf \mathcal{F} such that

$$\begin{aligned} \mathcal{F}(U) &:= \{ \varphi = (\varphi_P)_{P \in U} \text{ with } \varphi_P \in \mathcal{F}'_P \text{ for all } P \in U \\ &\text{such that for every } P \in U \text{ there is a neighborhood } V \text{ in } U \\ &\text{and a section } \varphi' \in \mathcal{F}'(V) \text{ with } \varphi_Q = \varphi'_Q \in \mathcal{F}'_Q \text{ for all } Q \in V. \} \end{aligned}$$

(For the notion of the stalk \mathcal{F}'_P of a presheaf \mathcal{F}' at a point $P \in X$ see definition 2.2.7.) It is obvious that this defines a sheaf.

Example 7.1.11. Let $X \subset \mathbb{A}^N$ be an affine variety. Let \mathcal{O}'_X be the presheaf given by

$$\begin{aligned} \mathcal{O}'_X(U) &= \left\{ \varphi : U \rightarrow k ; \text{ there are } f, g \in k[x_1, \dots, x_N] \text{ with } g(P) \neq 0 \right. \\ &\left. \text{and } \varphi(P) = \frac{f(P)}{g(P)} \text{ for all } P \in U \right\} \end{aligned}$$

for all open subsets $U \subset X$, i. e. the “presheaf of functions that are (globally) quotients of polynomials”. Then the structure sheaf \mathcal{O}_X is the sheafification of \mathcal{O}'_X , i. e. the sheaf of functions that are locally quotients of polynomials. We have seen in example 2.1.7 that in general \mathcal{O}'_X differs from \mathcal{O}_X , i. e. it is in general not a sheaf.

Example 7.1.12. If X is a topological space and \mathcal{F} the presheaf of constant real-valued functions on X , then the sheafification of \mathcal{F} is the sheaf of *locally* constant functions on X (see also remark 2.2.5).

The sheafification has the following nice and expected properties:

Lemma 7.1.13. *Let \mathcal{F}' be a presheaf on a topological space X , and let \mathcal{F} be its sheafification.*

- (i) *The stalks \mathcal{F}_P and \mathcal{F}'_P agree at every point $P \in X$.*
- (ii) *If \mathcal{F}' is a sheaf, then $\mathcal{F} = \mathcal{F}'$.*

Proof. (i): The isomorphism between the stalks is given by the following construction:

- An element of \mathcal{F}_P is by definition represented by (U, φ) , where U is an open neighborhood of P and $\varphi = (\varphi_Q)_{Q \in U}$ is a section of \mathcal{F} over U . To this we can associate the element $\varphi_P \in \mathcal{F}'_P$.
- An element of \mathcal{F}'_P is by definition represented by (U, φ) , where $\varphi \in \mathcal{F}'(U)$. To this we can associate the element $(\varphi_Q)_{Q \in U}$ in $\mathcal{F}(U)$, which in turn defines an element of \mathcal{F}_P .

(ii): Note that there is always a morphism of presheaves $\mathcal{F}' \rightarrow \mathcal{F}$ given by $\mathcal{F}'(U) \rightarrow \mathcal{F}(U), \varphi \mapsto (\varphi_P)_{P \in U}$.

Now assume that \mathcal{F}' is a sheaf; we will construct an inverse morphism $\mathcal{F} \rightarrow \mathcal{F}'$. Let $U \subset X$ be an open subset and $\varphi = (\varphi_P)_{P \in U} \in \mathcal{F}(U)$ a section of \mathcal{F} . For every $P \in U$ the germ $\varphi_P \in \mathcal{F}'_P$ is represented by some (V, ψ) with $\psi \in \mathcal{F}'(V)$. As P varies over U , we are thus getting sections of \mathcal{F}' on an open cover of U that agree on the overlaps. As \mathcal{F}' is a sheaf, we can glue these sections together to give a global section in $\mathcal{F}'(U)$. \square

Using sheafification, we can now define all the “natural” constructions that we would expect to be possible:

Definition 7.1.14. Let $f : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ be a morphism of sheaves of e. g. Abelian groups on a topological space X .

- (i) The **cokernel** $\text{coker } f$ of f is defined to be the sheaf associated to the presheaf $\text{coker}' f$.
- (ii) The morphism f is called **injective** if $\ker f = 0$. It is called **surjective** if $\text{coker } f = 0$.
- (iii) If the morphism f is injective, its cokernel is also denoted $\mathcal{F}_2/\mathcal{F}_1$ and called the **quotient** of \mathcal{F}_2 by \mathcal{F}_1 .
- (iv) As usual, a sequence of sheaves and morphisms

$$\cdots \rightarrow \mathcal{F}_{i-1} \rightarrow \mathcal{F}_i \rightarrow \mathcal{F}_{i+1} \rightarrow \cdots$$

is called **exact** if $\ker(\mathcal{F}_i \rightarrow \mathcal{F}_{i+1}) = \text{im}(\mathcal{F}_{i-1} \rightarrow \mathcal{F}_i)$ for all i .

Remark 7.1.15. Let us rephrase again the results of definition 7.1.7 and remark 7.1.8 in this new language:

- (i) A morphism $f : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ of sheaves is injective if and only if the maps $f_U : \mathcal{F}_1(U) \rightarrow \mathcal{F}_2(U)$ are injective for all U .
- (ii) If a morphism $f : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ of sheaves is surjective, this does *not* imply that all maps $f_U : \mathcal{F}_1(U) \rightarrow \mathcal{F}_2(U)$ are surjective. (The converse of this is obviously true however: if all maps $f_U : \mathcal{F}_1(U) \rightarrow \mathcal{F}_2(U)$ are surjective, then $\text{coker}' f = 0$, so $\text{coker } f = 0$.)

This very important fact is the basis of the theory of cohomology, see chapter 8.

Example 7.1.16. Let $X = \mathbb{P}_k^1$ with homogeneous coordinates x_0, x_1 . Consider the morphism of sheaves $f : \mathcal{O}_X(-1) \rightarrow \mathcal{O}_X$ given by the linear polynomial x_0 (see example 7.1.4).

We claim that f is injective. In fact, every section of $\mathcal{O}_X(-1)$ over an open subset of X has the form $\frac{g(x_0, x_1)}{h(x_0, x_1)}$ for some homogeneous polynomials g, h with $\deg g - \deg h = -1$. But $f(\frac{g}{h}) = \frac{gx_0}{h}$ is zero on an open subset of X if and only if $g = 0$ (note that we are not asking

for zeros of $\frac{gx_0}{h}$, but we are asking whether this function vanishes on a whole open subset). As this means that $\frac{g}{h}$ itself is zero, we see that the kernel of f is trivial, i. e. f is injective.

We have seen already in example 7.1.2 that f is in fact an isomorphism when restricted to $U = X \setminus \{P\}$ where $P := (0 : 1)$. In particular, f is surjective when restricted to U . However, f is not surjective on X (otherwise it would be an isomorphism, which is not true as we already know). Let us determine its cokernel.

First we have to compute the cokernel presheaf $\text{coker}' f$. Consider an open subset $U \subset X$. By the above argument, $(\text{coker}' f)(U) = 0$ if $P \notin U$. So assume that $P \in U$. Then we have an exact sequence of $\mathcal{O}_X(U)$ -modules

$$\begin{array}{ccccccc} 0 & \rightarrow & \mathcal{O}_X(-1)(U) & \rightarrow & \mathcal{O}_X(U) & \rightarrow & k \rightarrow 0 \\ & & \downarrow \frac{g}{h} & \mapsto & \downarrow \frac{gx_0}{h} & & \\ & & & & \varphi = \frac{g}{h} & \mapsto & \varphi(P) \end{array}$$

as the functions in the image of $\mathcal{O}_X(-1)(U) \rightarrow \mathcal{O}_X$ are precisely those that vanish on P . So we have found that

$$(\text{coker}' f)(U) = \begin{cases} k & \text{if } P \in U, \\ 0 & \text{if } P \notin U. \end{cases}$$

It is easily verified that $\text{coker}' f$ is in fact a sheaf. It can be thought of as the ground field k “concentrated at the point P ”. For this reason it is often called a **skyscraper sheaf** and denoted k_P .

Summarizing, we have found the exact sequence of sheaves of \mathcal{O}_X -modules

$$0 \rightarrow \mathcal{O}_X(-1) \xrightarrow{\cdot x_0} \mathcal{O}_X \rightarrow k_P \rightarrow 0.$$

Example 7.1.17. Let $\mathcal{F}_1, \mathcal{F}_2$ be two sheaves of \mathcal{O}_X -modules on a ringed space X . Then we can define the direct sum, the tensor product, and the dual sheaf in the obvious way:

- (i) The **direct sum** $\mathcal{F}_1 \oplus \mathcal{F}_2$ is the sheaf of \mathcal{O}_X -modules defined by $(\mathcal{F}_1 \oplus \mathcal{F}_2)(U) = \mathcal{F}_1(U) \oplus \mathcal{F}_2(U)$. (It is easy to see that this is a sheaf already, so that we do not need sheafification.)
- (ii) The **tensor product** $\mathcal{F}_1 \otimes \mathcal{F}_2$ is the sheaf of \mathcal{O}_X -modules associated to the presheaf $U \mapsto \mathcal{F}_1(U) \otimes_{\mathcal{O}_X(U)} \mathcal{F}_2(U)$.
- (iii) The **dual** \mathcal{F}_1^\vee of \mathcal{F}_1 is the sheaf of \mathcal{O}_X -modules associated to the presheaf $U \mapsto \mathcal{F}_1(U)^\vee = \text{Hom}_{\mathcal{O}_X(U)}(\mathcal{F}_1(U), \mathcal{O}_X(U))$.

Example 7.1.18. Similarly to example 7.1.16 consider the morphism $f : \mathcal{O}_X(-2) \rightarrow \mathcal{O}_X$ of sheaves on $X = \mathbb{P}_k^1$ given by multiplication with x_0x_1 (instead of with x_0). The only difference to the above example is that the function x_0x_1 vanishes at two points $P_0 = (0 : 1)$, $P_1 = (1 : 0)$. So this time we get an exact sequence of sheaves

$$0 \rightarrow \mathcal{O}_X(-2) \xrightarrow{\cdot x_0x_1} \mathcal{O}_X \rightarrow k_{P_0} \oplus k_{P_1} \rightarrow 0,$$

where the last morphism is evaluation at the points P_0 and P_1 .

The important difference is that this time the cokernel presheaf is *not* equal to the cokernel sheaf: if we consider our exact sequence on global sections, we get

$$0 \rightarrow \Gamma(\mathcal{O}_X(-2)) \rightarrow \Gamma(\mathcal{O}_X) \rightarrow k \oplus k,$$

where $\Gamma(\mathcal{O}_X(-2)) = 0$, and $\Gamma(\mathcal{O}_X)$ are just the constant functions. But the last morphism is evaluation at P and Q , and constant functions must have the same value at P and Q . So the last map $\Gamma(\mathcal{O}_X) \rightarrow k \oplus k$ is not surjective, indicating that some sheafification is going on. (In example 7.1.16 we only had to evaluate at one point, and the corresponding map was surjective.)

Example 7.1.19. On $X = \mathbb{P}^N$, we have $\mathcal{O}_X(n) \otimes \mathcal{O}_X(m) = \mathcal{O}_X(n+m)$, with the isomorphism given on sections by

$$\frac{f_1}{g_1} \otimes \frac{f_2}{g_2} \mapsto \frac{f_1 f_2}{g_1 g_2}.$$

Similarly, we have $\mathcal{O}_X(n)^\vee = \mathcal{O}_X(-n)$, as the $\mathcal{O}_X(U)$ -linear homomorphisms from $\mathcal{O}_X(n)$ to \mathcal{O}_X are precisely given by multiplication with sections of $\mathcal{O}_X(-n)$.

7.2. Quasi-coherent sheaves. It turns out that sheaves of modules are still too general objects for many applications — usually one wants to restrict to a smaller class of sheaves. Recall that any ring R determines an affine scheme $X = \text{Spec } R$ together with its structure sheaf \mathcal{O}_X . Hence one would expect that any R -module M determines a sheaf \tilde{M} of \mathcal{O}_X -modules on X . This is indeed the case, and almost any sheaf of \mathcal{O}_X -modules appearing in practice is of this form. For computations, this means that statements about this sheaf \tilde{M} on X are finally reduced to statements about the R -module M . But it does not follow from the definitions that a sheaf of \mathcal{O}_X -modules has to be induced by some R -module in this way (see example 7.2.3), so we will say that it is quasi-coherent if it does, and in most cases restrict our attention to these quasi-coherent sheaves. If X is a general scheme, we accordingly require that it has an open cover by affine schemes $\text{Spec } R_i$ over which the sheaf is induced by an R_i -module for all i .

Let us start by showing how an R -module M determines a sheaf of modules \tilde{M} on $X = \text{Spec } R$. This is essentially the same construction as for the structure sheaf in definition 5.1.11.

Definition 7.2.1. Let R be a ring, $X = \text{Spec } R$, and let M be an R -module. We define a sheaf of \mathcal{O}_X -modules \tilde{M} on X by setting

$$\begin{aligned} \tilde{M}(U) &:= \{ \varphi = (\varphi_{\mathfrak{p}})_{\mathfrak{p} \in U} \text{ with } \varphi_{\mathfrak{p}} \in M_{\mathfrak{p}} \text{ for all } \mathfrak{p} \in U \\ &\quad \text{such that “}\varphi \text{ is locally of the form } \frac{m}{r} \text{ for } m \in M, r \in R\text{”} \} \\ &= \{ \varphi = (\varphi_{\mathfrak{p}})_{\mathfrak{p} \in U} \text{ with } \varphi_{\mathfrak{p}} \in M_{\mathfrak{p}} \text{ for all } \mathfrak{p} \in U \\ &\quad \text{such that for every } \mathfrak{p} \in U \text{ there is a neighborhood } V \text{ in } U \text{ and } m \in M, r \in R \\ &\quad \text{with } r \notin \mathfrak{q} \text{ and } \varphi_{\mathfrak{q}} = \frac{m}{r} \in M_{\mathfrak{q}} \text{ for all } \mathfrak{q} \in V \}. \end{aligned}$$

It is clear from the local nature of the definition that \tilde{M} is a sheaf.

The following proposition corresponds exactly to the statement of proposition 5.1.12 for structure sheaves. Its proof can be copied literally, replacing R by M at appropriate places.

Proposition 7.2.2. *Let R be a ring, $X = \text{Spec } R$, and let M be an R -module.*

- (i) *For every $\mathfrak{p} \in X$ the stalk of \tilde{M} at \mathfrak{p} is $M_{\mathfrak{p}}$.*
- (ii) *For every $f \in R$ we have $\tilde{M}(X_f) = M_f$. In particular, $\tilde{M}(X) = M$.*

Example 7.2.3. The following example shows that not all sheaves of \mathcal{O}_X -modules on $X = \text{Spec } R$ have to be of the form \tilde{M} for some R -module M .

Let $X = \mathbb{A}_k^1$, and let \mathcal{F} be the sheaf associated to the presheaf

$$U \mapsto \begin{cases} \mathcal{O}_X(U) & \text{if } 0 \notin U, \\ 0 & \text{if } 0 \in U. \end{cases}$$

with the obvious restriction maps. Then \mathcal{F} is a sheaf of \mathcal{O}_X -modules. The stalk \mathcal{F}_0 is zero, whereas $\mathcal{F}_P = \mathcal{O}_{X,P}$ for all other points $P \in X$.

Note that \mathcal{F} has no non-trivial global sections: if $\varphi \in \mathcal{F}(X)$ then we obviously must have $\varphi_0 = 0 \in \mathcal{F}_0$, which by definition of sheafification means that φ is zero in some neighborhood of 0. But as X is irreducible, φ must then be the zero function. Hence it follows

that $\mathcal{F}(X) = 0$. So if \mathcal{F} was of the form \tilde{M} for some R -module M , it would follow from proposition 7.2.2 (ii) that $M = 0$, hence \mathcal{F} would have to be the zero sheaf, which it obviously is not.

Definition 7.2.4. Let X be a scheme, and let \mathcal{F} be a sheaf of \mathcal{O}_X -modules. We say that \mathcal{F} is **quasi-coherent** if for every affine open subset $U = \text{Spec } R \subset X$ the restricted sheaf $\mathcal{F}|_U$ is of the form \tilde{M} for some R -module M .

Remark 7.2.5. It can be shown that it is sufficient to require the condition of the definition only for every open subset in an affine open cover of X (see e. g. [H] proposition II.5.4). In other words, quasi-coherence is a local property.

Example 7.2.6. On any scheme the structure sheaf is quasi-coherent. The sheaves $\mathcal{O}_X(n)$ are quasi-coherent on any projective subscheme of \mathbb{P}^N as they are locally isomorphic to the structure sheaf. In the rest of this section we will show that essentially all operations that you can do with quasi-coherent sheaves yield again quasi-coherent sheaves. Therefore almost all sheaves that occur in practice are in fact quasi-coherent.

Lemma 7.2.7. Let R be a ring and $X = \text{Spec } R$.

- (i) For any R -modules M, N there is a one-to-one correspondence

$$\{\text{morphisms of sheaves } \tilde{M} \rightarrow \tilde{N}\} \leftrightarrow \{R\text{-module homomorphisms } M \rightarrow N\}.$$
- (ii) A sequence of R -modules $0 \rightarrow M_1 \rightarrow M_2 \rightarrow M_3 \rightarrow 0$ is exact if and only if the sequence of sheaves $0 \rightarrow \tilde{M}_1 \rightarrow \tilde{M}_2 \rightarrow \tilde{M}_3 \rightarrow 0$ is exact on X .
- (iii) For any R -modules M, N we have $\tilde{M} \oplus \tilde{N} = (\tilde{M \oplus N})$.
- (iv) For any R -modules M, N we have $\tilde{M} \otimes \tilde{N} = (\tilde{M \otimes N})$.
- (v) For any R -module M we have $(\tilde{M})^\vee = (\tilde{M^\vee})$.

In particular, kernels, cokernels, direct sums, tensor products, and duals of quasi-coherent sheaves are again quasi-coherent on any scheme X .

Proof. (i): Given a morphism $\tilde{M} \rightarrow \tilde{N}$, taking global sections gives an R -module homomorphism $M \rightarrow N$ by proposition 7.2.2 (ii). Conversely, an R -module homomorphism $M \rightarrow N$ gives rise to morphisms between the stalks $M_{\mathfrak{p}} \rightarrow N_{\mathfrak{p}}$ for all \mathfrak{p} , and therefore by definition determines a morphism $\tilde{M} \rightarrow \tilde{N}$ of sheaves. It is obvious that these two operations are inverse to each other.

(ii): By exercise 7.8.2, exactness of a sequence of sheaves can be seen on the stalks. Hence by proposition 7.2.2 (i) the statement follows from the algebraic fact that the sequence $0 \rightarrow M_1 \rightarrow M_2 \rightarrow M_3 \rightarrow 0$ is exact if and only if $0 \rightarrow (M_1)_{\mathfrak{p}} \rightarrow (M_2)_{\mathfrak{p}} \rightarrow (M_3)_{\mathfrak{p}} \rightarrow 0$ is for all prime ideals $\mathfrak{p} \in R$.

(iii), (iv), and (v) follow in the same way as (ii): the statement can be checked on the stalks, hence it follows from the corresponding algebraic fact about localizations of modules. \square

Example 7.2.8. Let $X = \mathbb{P}^1$ and $P = (0 : 1) \in X$. The skyscraper sheaf k_P of example 7.1.16 is quasi-coherent by lemma 7.2.7 as it is the cokernel of a morphism of quasi-coherent sheaves. One can also check explicitly that k_P is quasi-coherent: if $U_0 = \{x_0 \neq 0\} = \mathbb{P}^1 \setminus \{P\}$ and $U_1 = \{x_1 \neq 0\} = \text{Spec } k[x_0] \cong \mathbb{A}^1$ then $k_P|_{U_0} = 0$ (so it is the sheaf associated to the zero module) and $k_P|_{U_1} \cong \tilde{M}$ where $M = k$ is the $k[x_0]$ -module with the module structure

$$\begin{aligned} k[x_0] \times k &\rightarrow k \\ (f, \lambda) &\mapsto f(0) \cdot \lambda. \end{aligned}$$

Proposition 7.2.9. Let $f : X \rightarrow Y$ be a morphism of schemes, and let \mathcal{F} be a quasi-coherent sheaf on X . Assume moreover that every open subset of X can be covered by finitely many

affine open subsets (this should be thought of as a technical condition that is essentially always satisfied — it is e. g. certainly true for all subschemes of projective spaces). Then $f_*\mathcal{F}$ is quasi-coherent on Y .

Proof. Let us first assume that X and Y are affine, so $X = \text{Spec} R$, $Y = \text{Spec} S$, and $\mathcal{F} = \tilde{M}$ for some R -module M . Then it follows immediately from the definitions that $f_*\mathcal{F} = (M \text{ as an } S\text{-module})^\sim$, hence push-forwards of quasi-coherent sheaves are quasi-coherent if X and Y are affine.

In the general case, note that the statement is local on Y , so we can assume that Y is affine. But it is not local on X , so we cannot assume that X is affine. Instead, cover X by affine opens U_i , and cover $U_i \cap U_j$ by affine opens $U_{i,j,k}$. By our assumption, we can take these covers to be finite.

Now the sheaf property for \mathcal{F} says that for every open set $V \subset Y$ the sequence

$$0 \rightarrow \mathcal{F}(f^{-1}(V)) \rightarrow \bigoplus_i \mathcal{F}(f^{-1}(V) \cap U_i) \rightarrow \bigoplus_{i,j,k} \mathcal{F}(f^{-1}(V) \cap U_{i,j,k})$$

is exact, where the last map is given by $(\dots, s_i, \dots) \mapsto (\dots, s_i|_{U_{i,j,k}} - s_j|_{U_{i,j,k}}, \dots)$. This means that the sequence of sheaves on Y

$$0 \rightarrow f_*\mathcal{F} \rightarrow \bigoplus_i f_*(\mathcal{F}|_{U_i}) \rightarrow \bigoplus_{i,j,k} f_*(\mathcal{F}|_{U_{i,j,k}})$$

is exact. But as we have shown the statement already for morphisms between affine schemes and as finite direct sums of quasi-coherent sheaves are quasi-coherent, the last two terms in this sequence are quasi-coherent. Hence the kernel $f_*\mathcal{F}$ is also quasi-coherent by lemma 7.2.7. \square

Example 7.2.10. With this result we can now define (and motivate) what a closed embedding of schemes is. Note that for historical reasons closed embeddings are usually called closed immersions in algebraic geometry (in contrast to differential geometry, where an immersion is defined to be a *local* embedding).

We say that a morphism $f : X \rightarrow Y$ of schemes is a **closed immersion** if

- (i) f is a homeomorphism onto a closed subset of Y , and
- (ii) the induced morphism $\mathcal{O}_Y \rightarrow f_*\mathcal{O}_X$ is surjective.

The kernel of the morphism $\mathcal{O}_Y \rightarrow f_*\mathcal{O}_X$ is then called the **ideal sheaf** $\mathcal{I}_{X/Y}$ of the immersion.

Let us motivate this definition. We certainly want condition (i) to hold on the level of topological spaces. But this is not enough — we have seen that even isomorphisms cannot be detected on the level of topological spaces (see example 2.3.8), so we need some conditions on the structure sheaves as well. We have seen in example 5.2.3 that a closed immersion should be a morphism that is locally of the form $\text{Spec} R/I \rightarrow \text{Spec} R$ for some ideal $I \subset R$. In fact, this is exactly what condition (ii) means: assume that we are in the affine case, i. e. $X = \text{Spec} S$ and $Y = \text{Spec} R$. As \mathcal{O}_Y and $f_*\mathcal{O}_X$ are quasi-coherent (the former by definition and the latter by proposition 7.2.9), so is the kernel of $\mathcal{O}_Y \rightarrow f_*\mathcal{O}_X$ by lemma 7.2.7. So the exact sequence

$$0 \rightarrow \mathcal{I}_{X/Y} \rightarrow \mathcal{O}_Y \rightarrow f_*\mathcal{O}_X \rightarrow 0$$

comes from an exact sequence of R -modules

$$0 \rightarrow I \rightarrow R \rightarrow S \rightarrow 0$$

by lemma 7.2.7 (ii). In other words, $I \subset R$ is an ideal of R , and $S = R/I$. So indeed the morphism f is of the form $\text{Spec} R/I \rightarrow \text{Spec} R$ and therefore corresponds to an inclusion morphism of a closed subset.

Example 7.2.11. Having studied push-forwards of sheaves, we now want to consider pull-backs, i. e. the “dual” situation: given a morphism $f : X \rightarrow Y$ and a sheaf \mathcal{F} on Y , we want to construct a “pull-back” sheaf $f^*\mathcal{F}$ on X . Note that this should be “more natural” than the push-forward, as sheaves describe “function-like” objects, and for functions pull-back is more natural than push-forward: given a “function” $\varphi : Y \rightarrow k$, there is set-theoretically a well-defined pull-back function $\varphi \circ f : X \rightarrow k$. In contrast, a function $\varphi : X \rightarrow k$ does not give rise to a function $Y \rightarrow k$ in a natural way.

Let us first consider the affine case: assume that $X = \text{Spec } R$, and $Y = \text{Spec } S$, so that the morphism f corresponds to a ring homomorphism $S \rightarrow R$. Assume moreover that the sheaf \mathcal{F} on Y is quasi-coherent, so that it corresponds to an S -module M . Then $M \otimes_S R$ is a well-defined R -module, and the corresponding sheaf on X should be the pull-back $f^*\mathcal{F}$. Indeed, if e. g. $M = S$, i. e. $\mathcal{F} = \mathcal{O}_Y$, then $M \otimes_S R = S \otimes_S R = R$, so $f^*\mathcal{F} = \mathcal{O}_X$: pull-backs of regular functions are just regular functions.

This is our “local model” for the pull-back of sheaves. To show that this extends to the global case (and to sheaves that are not necessarily quasi-coherent), we need a different description though. So assume now that X, Y , and \mathcal{F} are arbitrary. The first thing to do is to define a sheaf of abelian groups on X from \mathcal{F} . This is more complicated than for the push-forward constructed in definition 7.1.5, because $f(U)$ need not be open if U is.

We let $f^{-1}\mathcal{F}$ be the sheaf on X associated to the presheaf $U \mapsto \lim_{V \supset f(U)} \mathcal{F}(V)$, where the limit is taken over all open subsets V with $f(U) \subset V \subset Y$. This notion of limit means that an element in $\lim_{V \supset f(U)} \mathcal{F}(V)$ is given by a pair (V, φ) with $V \supset f(U)$ and $\varphi \in \mathcal{F}(V)$, and that two such pairs (V, φ) and (V', φ') define the same element if and only if there is an open subset W with $f(U) \subset W \subset V \cap V'$ such that $\varphi|_W = \varphi'|_W$. This is the best we can do to adapt definition 7.1.5 to the pull-back case. It is easily checked that this construction does what we want on the stalks: we have $(f^{-1}\mathcal{F})_P = \mathcal{F}_{f(P)}$ for all $P \in X$.

Note that $f^{-1}\mathcal{F}$ is obviously a sheaf of $(f^{-1}\mathcal{O}_Y)$ -modules, but not a sheaf of \mathcal{O}_X -modules. (This corresponds to the statement that in the affine case considered above, M is an S -module, but not an R -module.) We have seen in our affine case what we have to do: we have to take the tensor product over $f^{-1}\mathcal{O}_Y$ with \mathcal{O}_X (i. e. over S with R). In other words, we define the **pull-back** $f^*\mathcal{F}$ of \mathcal{F} to be

$$f^*\mathcal{F} = f^{-1}\mathcal{F} \otimes_{f^{-1}\mathcal{O}_Y} \mathcal{O}_X,$$

which is then obviously a sheaf of \mathcal{O}_X -modules. As this construction restricts to the one given above if X and Y are affine and \mathcal{F} quasi-coherent, it also follows that pull-backs of quasi-coherent sheaves are again quasi-coherent.

It should be stressed that this complicated limit construction is only needed to prove the existence of $f^*\mathcal{F}$ in the general case. To compute the pull-back in practice, one will almost always restrict to affine open subsets and then use the tensor product construction given above.

Example 7.2.12. Here is a concrete example in which we can see again why the tensor product construction is necessary in the construction of the pull-back. Consider the morphism $f : X = \mathbb{P}^1 \rightarrow Y = \mathbb{P}^1$ given by $(s : t) \mapsto (x : y) = (s^2 : t^2)$. We want to compute the pull-back sheaf $f^*\mathcal{O}_Y(1)$ on X .

As we already know, local sections of $\mathcal{O}_Y(1)$ are of the form $\frac{g(x,y)}{h(x,y)}$, with g and h homogeneous such that $\deg g - \deg h = 1$. Pulling this back just means inserting the equations $x = s^2$ and $y = t^2$ of f into this expression; so the sheaf $f^{-1}\mathcal{O}_Y(1)$ has local sections $\frac{g(s^2,t^2)}{h(s^2,t^2)}$, where now $\deg(g(s^2,t^2)) - \deg(h(s^2,t^2)) = 2$.

But note that these sections do not even describe a sheaf of \mathcal{O}_X -modules: if we try to multiply the section s^2 with the function $\frac{t}{s}$ (i. e. a section of \mathcal{O}_X) on the open subset where

$s \neq 0$, we get st , which is not of the form $\frac{g(s^2, t^2)}{h(s^2, t^2)}$. We have just seen the solution to this problem: consider the tensor product with \mathcal{O}_X . So sections of $f^* \mathcal{O}_Y(1)$ are of the form

$$\frac{g(s^2, t^2)}{h(s^2, t^2)} \otimes \frac{g'(s, t)}{h'(s, t)}$$

with $\deg(g(s^2, t^2)) - \deg(h(s^2, t^2)) = 2$ and $\deg g' - \deg h' = 0$. It is easy to see that this describes precisely all expressions of the form $\frac{g''(s, t)}{h''(s, t)}$ with $\deg g'' - \deg h'' = 2$, so the result we get is $f^* \mathcal{O}_Y(1) = \mathcal{O}_X(2)$.

In the same way one shows that $f^* \mathcal{O}_Y(n) = \mathcal{O}_X(dn)$ for all $n \in \mathbb{Z}$ and any morphism $f : X \rightarrow Y$ between projective schemes that is given by a collection of homogeneous polynomials of degree d .


We have seen now that most sheaves occurring in practice are in fact quasi-coherent. So when we talk about sheaves from now on, we will usually think of quasi-coherent sheaves. This has the advantage that, on affine open subsets, sheaves (that form a somewhat complicated object) are essentially replaced by modules, which are usually much easier to handle.

7.3. Locally free sheaves. We now come to the discussion of locally free sheaves, i. e. sheaves that are locally just a finite direct sum of copies of the structure sheaf. These are the most important and best-behaved sheaves one can imagine.

Definition 7.3.1. Let X be a scheme. A sheaf of \mathcal{O}_X -modules \mathcal{F} is called **locally free of rank r** if there is an open cover $\{U_i\}$ of X such that $\mathcal{F}|_{U_i} \cong \mathcal{O}_{U_i}^{\oplus r}$ for all i . Obviously, every locally free sheaf is also quasi-coherent.

Remark 7.3.2. The geometric interpretation of locally free sheaves is that they correspond to “vector bundles” as known from topology — objects that associate to every point P of a space X a vector bundle. For example, the “tangent sheaf” of \mathbb{P}^1 in example 7.1.1 is such a vector bundle (of rank 1). Let us make this correspondence precise.

A **vector bundle** of rank r on a scheme X over a field k is a k -scheme F and a k -morphism $\pi : F \rightarrow X$, together with the additional data consisting of an open covering $\{U_i\}$ of X and isomorphisms $\psi_i : \pi^{-1}(U_i) \rightarrow U_i \times \mathbb{A}_k^r$ over U_i , such that the automorphism $\psi_i \circ \psi_j^{-1}$ of $(U_i \cap U_j) \times \mathbb{A}^r$ is linear in the coordinates of \mathbb{A}^r for all i, j . In other words, the morphism $\pi : F \rightarrow X$ looks locally like the projection morphism $U \times \mathbb{A}_k^r \rightarrow U$ for sufficiently small open subsets $U \subset X$.



We claim that there is a one-to-one correspondence

$$\{\text{vector bundles } \pi : F \rightarrow X \text{ of rank } r\} \leftrightarrow \{\text{locally free sheaves } \mathcal{F} \text{ of rank } r \text{ on } X\}$$

given by the following constructions:

- (i) Let $\pi : F \rightarrow X$ be a vector bundle of rank r . Define a sheaf \mathcal{F} on X by

$$\mathcal{F}(U) = \{k\text{-morphisms } s : U \rightarrow F \text{ such that } \pi \circ s = \text{id}_U\}.$$

(This is called the sheaf of sections of F .) Note that this has a natural structure of a sheaf of \mathcal{O}_X -modules (over every point in X we can multiply a vector with a scalar — doing this on an open subset means that we can multiply a section in $\mathcal{F}(U)$ with a regular function in $\mathcal{O}_X(U)$).

Locally, on an open subset U on which π is of the form $U \times \mathbb{A}_k^r \rightarrow U$, we obviously have

$$\mathcal{F}(U) = \{k\text{-morphisms } s : U \rightarrow \mathbb{A}_k^r\},$$

so sections are just given by r independent functions. In other words, $\mathcal{F}|_U$ is isomorphic to $\mathcal{O}_U^{\oplus r}$. So \mathcal{F} is locally free by definition.

- (ii) Conversely, let \mathcal{F} be a locally free sheaf. Take an open cover $\{U_i\}$ of X such that there are isomorphisms $\psi_i : \mathcal{F}|_{U_i} \rightarrow \mathcal{O}_{U_i}^{\oplus r}$. Now consider the schemes $U_i \times \mathbb{A}_k^r$ and glue them together as follows: for all i, j we glue $U_i \times \mathbb{A}_k^r$ and $U_j \times \mathbb{A}_k^r$ on the common open subset $(U_i \cap U_j) \times \mathbb{A}_k^r$ along the isomorphism

$$(U_i \cap U_j) \times \mathbb{A}_k^r \rightarrow (U_i \cap U_j) \times \mathbb{A}_k^r, \quad (P, s) \mapsto (P, \psi_i \circ \psi_j^{-1}).$$

Note that $\psi_i \circ \psi_j^{-1}$ is an isomorphism of sheaves of \mathcal{O}_X -modules and therefore linear in the coordinates of \mathbb{A}_k^r .

It is obvious that this gives exactly the inverse construction to (i).

Remark 7.3.3. Let $\pi : F \rightarrow X$ be a vector bundle of rank r , and let $P \in X$ be a point. We call $\pi^{-1}(P)$ the **fiber** of F over P ; it is an r -dimensional vector space. If \mathcal{F} is the corresponding locally free sheaf, the fiber can be realized as $i^* \mathcal{F}$ where $i : P \rightarrow X$ denotes the inclusion morphism (note that $i^* \mathcal{F}$ is a sheaf on a one-point space, so its data consists only of one k -vector space $(i^* \mathcal{F})(P)$, which is precisely the fiber F_P).

Lemma 7.3.4. *Let X be a scheme. If \mathcal{F} and \mathcal{G} are locally free sheaves of \mathcal{O}_X -modules of rank r and s , respectively, then the following sheaves are also locally free: $\mathcal{F} \oplus \mathcal{G}$ (of rank $r + s$), $\mathcal{F} \otimes \mathcal{G}$ (of rank $r \cdot s$), and \mathcal{F}^\vee (of rank r). If $f : X \rightarrow Y$ is a morphism of schemes and \mathcal{F} is a locally free sheaf on Y , then $f^* \mathcal{F}$ is a locally free sheaf on X of the same rank. (The push-forward of a locally free sheaf is in general not locally free.)*

Proof. The proofs all follow from the corresponding statements about vector spaces (or free modules over a ring): for example, if M and N are free R -modules of dimension r and s respectively, then $M \oplus N$ is a free R -module of dimension $r + s$. Applying this to an open affine subset $U = \text{Spec } R$ in X on which \mathcal{F} and \mathcal{G} are isomorphic to $\mathcal{O}_U^{\oplus r} = \tilde{M}$ and $\mathcal{O}_U^{\oplus s} = \tilde{N}$ gives the desired result. The statement about tensor products and duals follows in the same way. As for pull-backs, we have already seen that $f^* \mathcal{O}_Y = \mathcal{O}_X$, so $f^* \mathcal{F}$ will be of the form $\mathcal{O}_{f^{-1}(U)}^{\oplus r}$ on the inverse image $f^{-1}(U) \subset X$ of an open subset $U \subset Y$ on which \mathcal{F} is of the form $\mathcal{O}_U^{\oplus r}$. \square

Remark 7.3.5. Lemma 7.3.4 is an example of the general principle that any “canonical” construction or statement that works for vector spaces (or free modules) also works for vector bundles. Here is another example: recall that for any vector space V over k (or any free module) one can define the n -th **symmetric product** $S^n V$ and the n -th **alternating product** $\Lambda^n V$ to be the vector space of formal totally symmetric (resp. antisymmetric) products

$$v_1 \cdot \cdots \cdot v_n \in S^n V \quad \text{and} \quad v_1 \wedge \cdots \wedge v_n \in \Lambda^n V.$$

If V has dimension r , then $S^n V$ and $\Lambda^n V$ have dimension $\binom{n+r-1}{n}$ and $\binom{r}{n}$, respectively. More precisely, if $\{v_1, \dots, v_r\}$ is a basis of V , then

$$\begin{aligned} \{v_{i_1} \cdots v_{i_n} ; i_1 \leq \cdots \leq i_n\} & \text{ is a basis of } S^n V, \text{ and} \\ \{v_{i_1} \wedge \cdots \wedge v_{i_n} ; i_1 < \cdots < i_n\} & \text{ is a basis of } \Lambda^n V. \end{aligned}$$

Using the same construction, we can get symmetric and alternating products $S^n \mathcal{F}$ and $\Lambda^n \mathcal{F}$ on X for every locally free sheaf \mathcal{F} on X of rank r . They are locally free sheaves of ranks $\binom{n+r-1}{n}$ and $\binom{r}{n}$, respectively.

Here is an example of a linear algebra lemma that translates directly into the language of locally free sheaves:

Lemma 7.3.6. *Let $0 \rightarrow U \rightarrow V \rightarrow W \rightarrow 0$ be an exact sequence of vector spaces of dimensions a , $a+b$, and b , respectively. Then $\Lambda^{a+b} V = \Lambda^a U \otimes \Lambda^b W$.*

Proof. Denote the two homomorphisms by $i : U \rightarrow V$ and $p : V \rightarrow W$. Then there is a canonical isomorphism

$$\begin{aligned} \Lambda^a U \otimes \Lambda^b W & \rightarrow \Lambda^{a+b} V \\ (u_1 \wedge \cdots \wedge u_a) \otimes (w_1 \wedge \cdots \wedge w_b) & \mapsto i(u_1) \wedge \cdots \wedge i(u_a) \wedge p^{-1}(w_1) \wedge \cdots \wedge p^{-1}(w_b). \end{aligned}$$

The key remark here is that the $p^{-1}(w_i)$ are well-defined up to an element of U by the exact sequence. But if the above expression is non-zero at all, the u_1, \dots, u_a must form a basis of U , so if we plug in any element of U in the last b entries of the alternating product we will get zero. Therefore the ambiguity in the $p^{-1}(w_i)$ does not matter and the above homomorphism is well-defined. It is obviously not the zero map, and it is then an isomorphism for dimensional reasons (both sides are one-dimensional vector spaces). \square

Corollary 7.3.7. *Let $0 \rightarrow \mathcal{F}_1 \rightarrow \mathcal{F}_2 \rightarrow \mathcal{F}_3 \rightarrow 0$ be an exact sequence of locally free sheaves of ranks a_1, a_2, a_3 on a scheme X . Then $\Lambda^{a_2} \mathcal{F}_2 = \Lambda^{a_1} \mathcal{F}_1 \otimes \Lambda^{a_3} \mathcal{F}_3$.*

Proof. Immediately from lemma 7.3.6 using the above principle. \square

7.4. Differentials. We have seen in proposition 4.4.8 that (formal) differentiation of functions is useful to compute the tangent spaces at the (closed) points of a scheme X . We now want to introduce this language of differentials. The idea is that the various tangent spaces T_P for $P \in X$ should not just be independent vector spaces at every point, but rather come from a global object on X . For example, if X is smooth over \mathbb{C} , so that it is a complex manifold, we know from complex geometry that X has a cotangent bundle whose fiber at a point P is just the cotangent space, i. e. the dual of the tangent space, at P . We want to give an algebro-geometric analogue of this construction. So let us first define the process of formal differentiation. We start with the affine case.

Definition 7.4.1. Let $f : X = \text{Spec } R \rightarrow Y = \text{Spec } S$ be a morphism of affine schemes, corresponding to a ring homomorphism $S \rightarrow R$. We define the R -module $\Omega_{R/S}$, the **module of relative differentials**, to be the free R -module generated by formal symbols $\{dr ; r \in R\}$, modulo the relations:

- $d(r_1 + r_2) = dr_1 + dr_2$ for $r_1, r_2 \in R$,
- $d(r_1 r_2) = r_1 dr_2 + r_2 dr_1$ for $r_1, r_2 \in R$,
- $ds = 0$ for $s \in S$.

Example 7.4.2. Let $S = k$ be a field and $R = k[x_1, \dots, x_n]$, so that we consider the morphism $f : \mathbb{A}_k^n \rightarrow \text{pt}$. Then by the relations in $\Omega_{R/k}$, which are exactly the rules of differentiation with the elements of k being the ‘‘constant’’ functions, it follows that $df = \sum_i \frac{\partial f}{\partial x_i} dx_i$ for all $f \in k[x_1, \dots, x_n]$. So $\Omega_{R/k}$ is just the free R -module generated by the symbols dx_1, \dots, dx_n .

Again let $S = k$, but now let $R = k[x_1, \dots, x_n]/(f_1, \dots, f_m)$ be the coordinate ring of an affine variety. By the same calculation as above, $\Omega_{R/S}$ is still generated as an R -module by dx_1, \dots, dx_n , but the relations f_i give rise to relations $df_i = 0$ in $\Omega_{R/S}$. It is easy to see that these are all relations in $\Omega_{R/S}$, so we have

$$\Omega_{R/S} = (Rdx_1 + \dots + Rdx_n) / \left(\sum_i \frac{\partial f_i}{\partial x_i} dx_i, j = 1, \dots, m \right).$$

In particular, if $X = \text{Spec} R$, k is algebraically closed, and $P \in X$ is a closed point of X corresponding to a morphism $R \rightarrow k$, then by definition 4.4.1 we see that

$$\Omega_{R/S} \otimes_R k = \langle dx_1, \dots, dx_n \rangle / \left(\sum_i \frac{\partial f_i}{\partial x_i}(P) dx_i, j = 1, \dots, m \right)$$

is just the dual $T_{X,P}^\vee$ of the tangent space to X at P .

Example 7.4.3. If Y is not a point, then the difference in the module of differentials is just that all elements of S (i. e. all differentials that come from Y) are treated as “constants”. So then $\Omega_{R/S}$ can be thought of as “the differentials on X modulo pull-backs of differentials on Y ”. We will probably not need this very often.

Of course, if $f : X \rightarrow Y$ is a morphism of general (not necessarily affine) schemes, we want to consider the relative differentials of every restriction of f to affine opens of X and Y , and glue them together to get a quasi-coherent sheaf $\Omega_{X/Y}$. To do this, we have to give a different description of the relative differentials, as the construction given above does not glue very well.

Lemma 7.4.4. *Let $S \rightarrow R$ be a homomorphism of rings. Consider the map $\delta : R \otimes_S R \rightarrow R$ given by $\delta(r_1 \otimes r_2) = r_1 r_2$ and let $I \subset R \otimes_S R$ be its kernel. Then I/I^2 is an R -module that is isomorphic to $\Omega_{R/S}$.*

Proof. The R -module structure of I/I^2 is given by $r \cdot (r_1 \otimes r_2) := rr_1 \otimes r_2 = r_1 \otimes rr_2$, where the second equality follows from

$$rr_1 \otimes r_2 - r_1 \otimes rr_2 = (r_1 \otimes r_2) \cdot (r \otimes 1 - 1 \otimes r) \in I \cdot I$$

if $r_1 \otimes r_2 \in I$. Define a map of R -modules $\Omega_{R/S} \rightarrow I/I^2$ by $dr \mapsto 1 \otimes r - r \otimes 1$. Now we construct its inverse. The R -module $E := R \oplus \Omega_{R/S}$ is a ring by setting $(r_1 \oplus dr'_1) \cdot (r_2 \oplus dr'_2) := r_1 r_2 \oplus (r_1 dr'_2 + r_2 dr'_1)$. It is easy to check that the map $R \times R \rightarrow E$ given by $(r_1, r_2) \mapsto (r_1 r_2, r_1 dr_2)$ is an S -bilinear ring homomorphism, hence gives rise to a map $g : R \otimes_S R \rightarrow E$. As $g(I) \subset \Omega_{R/S}$ by definition and $g(I^2) = 0$, this induces a map $I/I^2 \rightarrow \Omega_{R/S}$. It is easy to see that this is in fact the inverse of the map $\Omega_{R/S} \rightarrow I/I^2$ given above. \square

Remark 7.4.5. It is easy to translate this lemma into the language of schemes: let $X = \text{Spec} R$ and $Y = \text{Spec} S$, so that the ring homomorphism $S \rightarrow R$ corresponds to a map $X \rightarrow Y$. Then $\text{Spec} R \otimes_S R = X \times_Y X$, and $\delta : R \otimes_S R \rightarrow R$ corresponds to the diagonal morphism $X \rightarrow X \times_Y X$. Hence $I \subset R \otimes_S R$ is the ideal of the diagonal $\Delta(X) \subset X \times_Y X$. This motivates the following construction.

Definition 7.4.6. Let $f : X \rightarrow Y$ be a morphism of schemes. Let $\Delta : X \rightarrow X \times_Y X$ be the diagonal morphism, and let $\mathcal{I} = \mathcal{I}_{\Delta(X)/X \times_Y X}$ be its ideal sheaf. Then the **sheaf of relative differentials** $\Omega_{X/Y}$ is defined to be the sheaf $\Delta^*(\mathcal{I}/\mathcal{I}^2)$ on X . If X is a scheme over a field k and $Y = \text{Spec} k$ is a point, then we will usually write $\Omega_{X/Y}$ as Ω_X .

Remark 7.4.7. Here we assume that the diagonal morphism Δ is a closed immersion, which is the case if the schemes in question are separated (this is the analogue of lemma 2.5.3 for schemes). We will always assume this here to avoid further complications.

Remark 7.4.8. It should be stressed that definition 7.4.6 is essentially useless for practical computations. Its only use is to show that a global object $\Omega_{X/Y}$ exists that restricts to the old definition 7.4.1 on affine open subsets. For applications, we will always use definition 7.4.1 and example 7.4.2 on open subsets.

Remark 7.4.9. The sheaf $\Omega_{X/Y}$ is always quasi-coherent: on affine open subsets it restricts to the sheaf associated to the module $\Omega_{R/S}$ constructed above.

Remark 7.4.10. Any morphism $f : X \rightarrow Y$ of schemes over a field induces a morphism of sheaves $f^*\Omega_Y \rightarrow \Omega_X$ on X that is just given by $d\varphi \mapsto d(f^*\varphi) = d(\varphi \circ f)$ for any function φ on Y .

Proposition 7.4.11. *An n -dimensional scheme X (of finite type over an algebraically closed field, e. g. a variety) is smooth if and only if Ω_X is locally free of rank n . (Actually, this is a local statement: $P \in X$ is a smooth point of X if and only if Ω_X is (locally) free in a neighborhood of P .)*

Proof. One direction is obvious: if Ω_X is locally free of rank n then its fibers at any point P , i. e. the cotangent spaces $T_{X,P}^\vee$, have dimension n . By definition this means that P is a smooth point of X .

Now let us assume that X is smooth (at P). As the proposition is of local nature we can assume that $X = \text{Spec } R$ with $R = k[x_1, \dots, x_r]/(f_1, \dots, f_m)$. By example 7.4.2 we then have

$$T_{X,P}^\vee = \langle dx_1, \dots, dx_r \rangle / \left(\sum_i \frac{\partial f_j}{\partial x_i}(P) dx_i, j = 1, \dots, m \right).$$

As this vector space has dimension n , we know that the matrix of differentials $D(P) = \left(\frac{\partial f_j}{\partial x_i}(P) \right)$ at the point P has rank $r - n$. Without loss of generality we can assume that the submatrix of D given by the first $r - n$ columns and rows has non-zero determinant. This means that dx_{r-n+1}, \dots, dx_r form a basis of $T_{X,P}^\vee$.

But the condition for a determinant to be non-zero is an “open condition”, i. e. the set on which it is satisfied is open. In other words, there is a neighborhood U of P in X such that the submatrix of $D(Q)$ given by the first $r - n$ columns and rows has non-zero determinant for all $Q \in U$. Consequently, the differentials dx_{r-n+1}, \dots, dx_r generate $T_{X,Q}^\vee$ for all $Q \in U$. In particular, the dimension of $T_{X,Q}^\vee$ is at most n . But the opposite inequality $\dim T_{X,Q}^\vee \geq n$ is always true; so we conclude that the differentials dx_{r-n+1}, \dots, dx_r actually form a basis of the cotangent space at all points $Q \in U$. So

$$\Omega_X|_U = \mathcal{O}_U dx_{r-n+1} \oplus \cdots \oplus \mathcal{O}_U dx_r,$$

i. e. Ω_X is locally free. □

Remark 7.4.12. There is a similar statement for any quasi-coherent sheaf \mathcal{F} . It says that:

- (i) The dimension of the fibers is an *upper semi-continuous function*. This means that if the dimension of the fiber of \mathcal{F} at a point P is n , then it is at most n in some neighborhood of P .
- (ii) If the dimension of the fibers is constant on some open subset U , then $\mathcal{F}|_U$ is locally free.

The idea of the proof of this statement is very similar to that of proposition 7.4.11.

Definition 7.4.13. Let X be a smooth n -dimensional scheme over an algebraically closed field. The dual bundle Ω_X^\vee of the cotangent bundle is called the **tangent bundle** and is denoted T_X . It is a locally free sheaf of rank n . The top exterior power $\Lambda^n \Omega_X$ of the cotangent bundle is a locally free sheaf of rank 1; it is called the **canonical bundle** ω_X of X .

Remark 7.4.14. The importance of the cotangent / canonical bundles stems from the fact that these bundles are *canonically defined* (hence the name) for any smooth scheme n . This gives e. g. a new method to show that two varieties are not isomorphic: if we have two varieties whose canonical bundles have different properties (say their spaces of global sections have different dimensions), then the varieties cannot be isomorphic.

As an example, let us now compute the cotangent / tangent / canonical bundles of some easy varieties.

Lemma 7.4.15. *The cotangent bundle of \mathbb{P}^n is determined by the exact sequence*

$$0 \rightarrow \Omega_{\mathbb{P}^n} \rightarrow \mathcal{O}(-1)^{\oplus(n+1)} \rightarrow \mathcal{O} \rightarrow 0.$$

(This sequence is usually called the **Euler sequence**.) Consequently, the tangent bundle fits into the dual exact sequence

$$0 \rightarrow \mathcal{O} \rightarrow \mathcal{O}(1)^{\oplus(n+1)} \rightarrow T_{\mathbb{P}^n} \rightarrow 0,$$

and the canonical bundle is $\omega_{\mathbb{P}^n} = \mathcal{O}(-n-1)$.

Proof. We know already from example 7.4.2 that the cotangent bundle $\Omega_{\mathbb{P}^n}$ is generated on the standard open subsets $U_i = \{x_i \neq 0\} \cong \mathbb{A}^n$ by the differentials $d(\frac{x_0}{x_i}), \dots, d(\frac{x_n}{x_i})$ of the affine coordinates. Therefore the differentials $d(\frac{x_i}{x_j})$, where defined, generate all of $\Omega_{\mathbb{P}^n}$. By the rules of differentiation we have to require formally that

$$d\left(\frac{x_i}{x_j}\right) = \frac{x_j dx_i - x_i dx_j}{x_j^2}.$$

Note that the dx_i are not well-defined objects, as the x_i are not functions. But if we formally let the symbols dx_0, \dots, dx_n be the names of the generators of $\mathcal{O}(-1)^{\oplus(n+1)}$, the morphism of sheaves

$$\Omega_{\mathbb{P}^n} \rightarrow \mathcal{O}(-1)^{\oplus(n+1)}, \quad d\left(\frac{x_i}{x_j}\right) \mapsto \frac{1}{x_j} \cdot dx_i - \frac{x_i}{x_j^2} \cdot dx_j$$

is obviously well-defined and injective. It is now easily checked that the sequence of the lemma is exact, with the last morphism given by

$$\mathcal{O}(-1)^{\oplus(n+1)} \mapsto \mathcal{O}, \quad dx_i \mapsto x_i.$$

The sequence for the tangent bundle is obtained by dualizing. The statement about the canonical bundle then follows from corollary 7.3.7. \square

Lemma 7.4.16. *Let $X \subset \mathbb{P}^n$ be a smooth hypersurface of degree d , and let $i : X \rightarrow \mathbb{P}^n$ be the inclusion morphism. Then the cotangent bundle Ω_X is determined by the exact sequence*

$$0 \rightarrow \mathcal{O}_X(-d) \rightarrow i^* \Omega_{\mathbb{P}^n} \rightarrow \Omega_X \rightarrow 0.$$

Consequently, the tangent bundle is determined by the exact sequence

$$0 \rightarrow T_X \rightarrow i^* T_{\mathbb{P}^n} \rightarrow \mathcal{O}_X(d) \rightarrow 0,$$

and the canonical bundle is $\omega_X = \mathcal{O}_X(d-n-1)$.

Proof. We claim that the exact sequence is given by

$$\begin{array}{ccccccc} 0 & \rightarrow & \mathcal{O}_X(-d) & \rightarrow & i^* \Omega_{\mathbb{P}^n} & \rightarrow & \Omega_X & \rightarrow & 0 \\ & & \varphi & \mapsto & d(f \cdot \varphi), & & & & \\ & & & & d\varphi & \mapsto & d(\varphi|_X), & & \end{array}$$

where f is the equation defining X . In fact, the second map is just the usual pull-back of differential forms as in remark 7.4.10 (which is just a restriction in this case). It is surjective because functions on X are *locally* of the form $\frac{g}{h}$ for some homogeneous polynomials g and h of the same degree, so they are locally obtained by restricting a function on \mathbb{P}^n to X . It

is not an isomorphism though, because we have the identity $f = 0$ on X . Consequently, differentials $d\varphi$ are zero when restricted to X if and only if φ contains f as a factor. This explains the first map of the above sequence.

As in the previous lemma, the statements about the tangent and canonical bundles are obtained by dualizing and applying corollary 7.3.7, respectively. \square

Remark 7.4.17. In general, if $i : X \rightarrow Y$ is a closed immersion of smooth schemes over a field, there is an injective morphism $T_X \rightarrow i^*T_Y$ of sheaves on X . In other words, at points in X the tangent spaces of X are just subspaces of the tangent spaces of Y . The quotient $T_{Y,P}/T_{X,P}$ is called the normal space, and consequently the quotient bundle $N_{X/Y} = i^*T_Y/T_X$ is called the **normal bundle**. This is the same construction as in differential geometry. Thus lemma 7.4.16 just tells us that the normal bundle of a degree- d hypersurface in \mathbb{P}^n is $N_{X/\mathbb{P}^n} = \mathcal{O}_X(d)$.

Example 7.4.18. Let us evaluate lemma 7.4.16 in the simplest cases, namely for curves $X \subset \mathbb{P}^2$ of low degrees d .

- (i) $d = 1$: A linear curve in \mathbb{P}^2 is just isomorphic to \mathbb{P}^1 . We get $\Omega_X = \omega_X = \mathcal{O}(1 - 2 - 1) = \mathcal{O}(-2)$ by lemma 7.4.16. This is consistent with lemma 7.4.15 for $n = 1$.
- (ii) $d = 2$: We know from example 3.3.11 that a smooth plane conic is again just isomorphic to \mathbb{P}^1 by means of a quadratic map $f : \mathbb{P}^1 \rightarrow X \subset \mathbb{P}^2$. Our formula of lemma 7.4.16 gives $\omega_X = \mathcal{O}_X(2 - 2 - 1) = \mathcal{O}_X(-1)$. By pulling this back via f we obtain $\omega_X = \mathcal{O}_{\mathbb{P}^1}(-2)$ by example 7.2.12. So by applying the isomorphism to case (i) we get the same canonical bundle back — which has to be the case, as the cotangent bundle is canonically defined and cannot change with the embedding in projective space.
- (iii) $d = 3$: Here we get $\omega_X = \mathcal{O}(3 - 2 - 1) = \mathcal{O}$, i. e. the canonical bundle is simply isomorphic to the sheaf of regular functions. We can understand this from our representation in proposition 6.5.7 of cubic curves as complex tori of the form \mathbb{C}/Λ for some lattice $\Lambda \subset \mathbb{C}$. If z is the complex coordinate on \mathbb{C} , note that the differential form dz is invariant under shifts in Λ , as $d(z + a) = dz$ for all $a \in \mathbb{C}$. Therefore dz descends to a global differential form on $X = \mathbb{C}/\Lambda$ without zeros or poles. It follows that we have an isomorphism $\mathcal{O}_X \rightarrow \omega_X$ given by $\varphi \mapsto \varphi \cdot dz$.

7.5. Line bundles on curves. We now want to specialize even further and consider vector bundles of rank 1 (also called “line bundles”, because their fibers are just lines) on smooth curves. This section should be compared to section 6.3 where we considered divisors on such curves. We will show that divisor classes and line bundles are essentially the same thing.

Recall that the group $\text{Pic } X$ of divisor classes on a smooth curve X has a group structure in a natural way. So let us first make the set of all line bundles on X into a group as well. In fact, this can be done for any scheme:

Definition 7.5.1. Let X be a scheme. A **line bundle** on X is a vector bundle (i. e. a locally free sheaf) of rank 1. We denote the set of all line bundles on X by $\text{Pic}^l X$. This set has a natural structure of Abelian group, with multiplication given by tensor products, inverses by taking duals, and the neutral element by the structure sheaf.

We will now restrict our attention to smooth curves. To set up a correspondence between line bundles and divisors, we will have to define the divisor of a (rational) section of a line bundle. This is totally analogous to the divisor of a rational function in definition 6.3.4.

Definition 7.5.2. Let \mathcal{L} be a line bundle on a smooth curve X , and let $P \in X$ be a point. Assume that we are given a section $s \in \mathcal{L}(U)$ of \mathcal{L} on some neighborhood U of P . As \mathcal{L} is a line bundle, there is an isomorphism $\psi : \mathcal{L}|_U \rightarrow \mathcal{O}_U$ (possibly after shrinking U). The

order of vanishing $\text{ord}_P s$ of the section s at P is defined to be the order of vanishing of the regular function $\psi(s)$ at P .

Remark 7.5.3. Note that this definition does not depend on the choice of ψ : if $\psi' : \mathcal{L}|_U \rightarrow \mathcal{O}_U$ is another isomorphism, then the composition $\psi' \circ \psi^{-1} : \mathcal{O}_U \rightarrow \mathcal{O}_U$ is an isomorphism of the structure sheaf, which must be given by multiplication with a function φ that is nowhere zero (in particular not at P). So we have an equation of divisors

$$(\psi'(s)) = (\psi' \psi^{-1} \psi(s)) = (\varphi \cdot \psi(s)) = (\varphi) + (\psi(s)) = (\psi(s)),$$

which shows that $\text{ord}_P s$ is well-defined.

Definition 7.5.4. Let \mathcal{L} be a line bundle on a smooth curve X . A **rational section** of \mathcal{L} over U is a section of the sheaf $\mathcal{L} \otimes_{\mathcal{O}_X} \mathcal{K}_X$, where \mathcal{K}_X denotes the “sheaf of rational functions” whose value at every open subset $U \subset X$ is just $K(X)$. In other words, a rational section of a line bundle is given by an ordinary section of the line bundle, possibly multiplied with a rational function.

Now let $P \in X$ be a point, and let s be a rational section of \mathcal{L} in a neighborhood of P . With the same isomorphism ψ as in definition 7.5.2, the order $\text{ord}_P s$ of s at P is defined to be the order of the rational function $\psi(s)$ at P . (This is well-defined for the reason stated in remark 7.5.3.)

If s is a global rational section of \mathcal{L} , we define the divisor (s) of s to be

$$(s) = \sum_{P \in X} \text{ord}_P s \cdot P \in \text{Div } X.$$

Example 7.5.5. Let $X = \mathbb{P}^1$ with homogeneous coordinates x_0, x_1 .

- (i) Consider the global section $s = x_0 x_1$ of $\mathcal{O}_X(2)$. It vanishes at the points $P = (0 : 1)$ and $Q = (1 : 0)$ with multiplicity 1 each, so $(s) = P + Q$.
- (ii) The divisor of the global *rational* section $s = \frac{1}{x_0}$ of $\mathcal{O}_X(-1)$ is $(s) = -P$.

To show that $\text{Pic}' X \cong \text{Pic } X$ for smooth curves we need the following key lemma (which is the only point at which smoothness is needed).

Lemma 7.5.6. *Let X be a curve (over some algebraically closed field), and let $P \in X$ be a smooth point. Then there is a function φ_P in a neighborhood of P such that*

- (i) φ_P vanishes at P with multiplicity 1, i. e. its divisor contains the point P with multiplicity 1.
- (ii) φ_P is non-zero at all points distinct from P .

Proof. We can assume that $X = \text{Spec } R$ is affine, with $R = k[x_1, \dots, x_r]/(f_1, \dots, f_m)$ being the coordinate ring of X . As P is a smooth point of X , its cotangent space

$$T_{X,P}^\vee = \langle dx_1, \dots, dx_r \rangle / \left(\sum_i \frac{\partial f_j}{\partial x_i}(P) dx_i \text{ for all } j \right)$$

is one-dimensional. Let φ_P be any linear function such that $d\varphi_P$ generates this vector space. Then φ_P vanishes at P with multiplicity 1 by construction. We can now pick a neighborhood of P such that φ_P does not vanish at any other point. \square

Remark 7.5.7. If the ground field is \mathbb{C} and one thinks of X as a complex one-dimensional manifold, one can think of the function φ_P of lemma 7.5.6 as a “local coordinate” of X around P , i. e. a function that gives a local isomorphism of X with \mathbb{C} , with P mapping to $0 \in \mathbb{C}$. Note however that this is *not* true in the algebraic category, as the Zariski open subsets are too big.

We are now ready to prove the main proposition of this section.

Definition 7.5.8. A divisor $D = \sum_P a_P P$ on a smooth curve X is called **effective** (written $D \geq 0$) if $a_P \geq 0$ for all P .

Proposition 7.5.9. Let X be a smooth curve. Then there is an isomorphism of Abelian groups

$$\begin{aligned} \text{Pic}' X &\rightarrow \text{Pic} X \\ \mathcal{L} &\mapsto (s) \text{ for any rational section } s \text{ of } \mathcal{L}. \end{aligned}$$

Its inverse is given by

$$\begin{aligned} \text{Pic} X &\rightarrow \text{Pic}' X \\ D &\mapsto O(D), \end{aligned}$$

where $O(D)$ is the line bundle defined by

$$O(D)(U) = \{\varphi \in K(X) ; (\varphi) + D \geq 0 \text{ on } U\}.$$

Proof. We have to check a couple of things:

- (i) If \mathcal{L} is a line bundle, then there is a rational section s of \mathcal{L} : This is obvious, as \mathcal{L} is isomorphic to O on an open subset of X . So we can find a section of \mathcal{L} on this open subset (corresponding to the constant function 1). This will be a rational section of \mathcal{L} on all of X .
- (ii) The divisor class (s) of a rational section s of \mathcal{L} does not depend on the choice of s : If we have another section s' , then the quotient $\frac{s}{s'}$ will be a rational function, which has divisor class zero by definition of $\text{Pic} X$. So $(s) = (\frac{s}{s'} \cdot s') = (\frac{s}{s'}) + (s') = (s')$ in $\text{Pic} X$.
- (iii) If D is a divisor then $O(D)$ is actually a line bundle: let $P \in X$ be a point and choose a neighborhood U of P such that no point of $U \setminus P$ is contained in D . Let n be the coefficient of P in D . Then an isomorphism $\psi : O(D) \rightarrow O$ on U is given by multiplication with φ_P^n , where φ_P is the function of lemma 7.5.6. In fact, a rational function φ in $K(X)$ is by definition a section of $O(D)$ if and only if $\text{ord}_P \varphi + n \geq 0$, which is the case if and only if $\varphi \cdot \varphi_P^n$ is regular at P .
- (iv) If the divisors D and D' define the same element in $\text{Pic} X$ then $O(D) = O(D')$: By assumption we have $D - D' = (\varphi)$ in $\text{Pic} X$ for some rational function φ . Obviously, this induces an isomorphism $O(D) \rightarrow O(D')$ through multiplication with φ .

We have now shown that the maps stated in the proposition are well-defined. Let us now check that the two maps are inverse to each other.

- (v) $\text{Pic}' X \rightarrow \text{Pic} X \rightarrow \text{Pic}' X$: Let s_0 be a rational section of a line bundle \mathcal{L} , and consider $O((s_0)) = \{\varphi \in K(X) ; (\varphi) + (s_0) \geq 0\}$. We have an isomorphism

$$\mathcal{L} \rightarrow O((s_0)), \quad s \mapsto \frac{s}{s_0}.$$

- (vi) $\text{Pic} X \rightarrow \text{Pic}' X \rightarrow \text{Pic} X$: The (constant) rational function 1 defines a rational section of $O(D)$. To determine its order at a point P we have to apply the local isomorphism with O constructed in (iii): the order of this rational section at P is just the order of $1 \cdot \varphi_P^n$, which is n . This is exactly the multiplicity of P in D , so the divisor of our section is precisely D .

Finally, we have to check that the map is a homomorphism of groups. But this is clear: if s and s' are rational sections of \mathcal{L} and \mathcal{L}' , respectively, then ss' is a rational section of $\mathcal{L} \otimes \mathcal{L}'$, and $(ss') = (s) + (s')$. Hence tensor products of line bundles correspond to addition of divisors under our correspondence. \square

Definition 7.5.10. Let X be a smooth curve. From now on we will identify line bundles with divisor classes and call both groups $\text{Pic} X$. In particular, this defines the degree of a


line bundle (to be the degree of the associated divisor class). The divisor class associated to the canonical bundle ω_X is denoted K_X ; it is called the **canonical divisor (class)**.

Example 7.5.11. We have seen in lemma 6.3.11 that $\text{Pic } \mathbb{P}^1 = \mathbb{Z}$, i. e. there is exactly one divisor class in every degree. Consequently, there is exactly one line bundle for every degree n , which is of course just $\mathcal{O}(n)$. On the other hand, if $X \subset \mathbb{P}^2$ is a smooth cubic curve we know from corollary 6.3.15 that $\text{Pic } X$ consists of a copy of X in every degree. So on a cubic curve there are (many) more line bundles than just the bundles of the form $\mathcal{O}(n)$.

Remark 7.5.12. The correspondence of proposition 7.5.9 allows us to define the pull-back f^*D of a divisor class D on Y for any (surjective) morphism of smooth curves $f : X \rightarrow Y$: it is just given by pulling back the corresponding line bundle.

In fact, we can even define a pull-back f^*D for any *divisor* $D \in \text{Div } Y$ that induces this construction on the corresponding divisor class: let $P \in X$ be any point, and let $Q = f(P)$ be its image, considered as an element of $\text{Div } Y$. Then the subscheme $f^{-1}(Q)$ of X has a component whose underlying point is P . We define the **ramification index** e_P of f at P to be the length of this component subscheme. In more down to earth terms, this means that we take a function φ_Q as in lemma 7.5.6 that vanishes at Q with multiplicity 1, and define e_P to be the order of vanishing of the pull-back function $f^*\varphi_Q = \varphi_Q \circ f$ at P .

The ramification index has a simple interpretation in complex analysis: in the ordinary topology the curves X and Y are locally isomorphic to the complex plane, so we can pick local coordinates z on X around P and w on Y around Q . Every holomorphic map is now locally of the form $z \mapsto w = uz^n$ for some $n \geq 1$ and an invertible function u (i. e. a function that is non-zero at P). The number n is just the ramification index defined above. It is 1 if and only if f is a local isomorphism at P in complex analysis. We say that f is ramified at P if $n = e_P > 1$, and unramified at P otherwise.



If we now consider a point Q as an element of $\text{Div } Y$, we simply define

$$f^*Q = \sum_{P:f(P)=Q} e_P \cdot P$$

and extend this by linearity to obtain a homomorphism $f^* : \text{Div } Y \rightarrow \text{Div } X$. In other words, f^*D is just obtained by taking the inverse image points of the points in D with the appropriate multiplicities.

Using the correspondence of proposition 7.5.9 it is now easily checked that the induced map $f^* : \text{Pic } Y \rightarrow \text{Pic } X$ on the Picard groups agrees with the pull-back of line bundles.

Example 7.5.13. Let $f : X = \mathbb{P}^1 \rightarrow Y = \mathbb{P}^1$ be the morphism given by $(x_0 : x_1) \mapsto (x_0^2 : x_1^2)$. Then $f^*(1 : 0) = 2 \cdot (1 : 0)$ and $f^*(1 : 1) = (1 : 1) + (1 : -1)$ as divisors in X .

As an application of line bundles, we will now see how they can be used to describe morphisms to projective spaces. This works for all schemes (not just curves).

Lemma 7.5.14. *Let X be a scheme over an algebraically closed field. There is a one-to-one correspondence*

$$\{\text{morphisms } f : X \rightarrow \mathbb{P}^r\} \longleftrightarrow \left\{ \begin{array}{l} \text{line bundles } \mathcal{L} \text{ on } X \text{ together with global} \\ \text{sections } s_0, \dots, s_r \in \Gamma(X, \mathcal{L}) \text{ such that:} \\ \text{for all } P \in X \text{ there is some } s_i \text{ with } s_i(P) \neq 0 \end{array} \right\}$$

Proof. “ \longleftarrow ”: Given $r + 1$ sections of a line bundle \mathcal{L} on X that do not vanish simultaneously, we can define a morphism $f : X \rightarrow \mathbb{P}^r$ by setting $f(P) = (s_0(P) : \dots : s_r(P))$. Note that the values $s_i(P)$ are not well-defined numbers, but their quotients $\frac{s_i}{s_j}(P)$ are (as they are sections of $\mathcal{L} \otimes \mathcal{L}^\vee = \mathcal{O}$, i. e. ordinary functions). Therefore $f(P)$ is a well-defined point in projective space.

“ \longrightarrow ”: Given a morphism $f : X \rightarrow \mathbb{P}^r$, we set $\mathcal{L} = f^* \mathcal{O}_{\mathbb{P}^r}(1)$ and $s_i = f^* x_i$, where we consider the x_i as sections of $\mathcal{O}(1)$ (and thus the s_i as sections of $f^* \mathcal{O}(1)$). \square

Remark 7.5.15. One should regard this lemma as a generalization of lemma 3.3.9 where we have seen that a morphism to \mathbb{P}^r can be given by specifying $r + 1$ homogeneous polynomials of the same degree. Of course, this was just the special case in which the line bundle of lemma 7.5.14 is $\mathcal{O}(d)$. We had mentioned already in remark 3.3.10 that not all morphisms are of this form; this translates now into the statement that not all line bundles are of the form $\mathcal{O}(n)$.

7.6. The Riemann-Hurwitz formula. Let X and Y be smooth projective curves, and let $f : X \rightarrow Y$ be a surjective morphism. We want to compare the sheaves of differentials on X and Y . Note that every projective curve admits a surjective morphism to \mathbb{P}^1 : by definition it sits in some \mathbb{P}^n to start with, so we can find a morphism to \mathbb{P}^1 by repeated projections from points not in X . So if we know the canonical bundle of \mathbb{P}^1 (which we do by lemma 7.4.15: it is just $\mathcal{O}_{\mathbb{P}^1}(-2)$) and how canonical bundles transform under morphisms, we can at least in theory compute the canonical bundles of every curve.

Definition 7.6.1. Let $f : X \rightarrow Y$ be a surjective morphism of smooth projective curves. We define the **ramification divisor** of f to be $R = \sum_{P \in X} (e_P - 1) \cdot P \in \text{Div } X$, where e_P is the ramification index of f at P defined in remark 7.5.12. So the divisor R contains all points at which f is ramified, with appropriate multiplicities.

Proposition 7.6.2. (Riemann-Hurwitz formula) *Let $f : X \rightarrow Y$ be a surjective morphism of smooth projective curves, and let R be the ramification divisor of f . Then $K_X = f^* K_Y + R$ (or equivalently $\omega_X = f^* \omega_Y \otimes \mathcal{O}_X(R)$) in $\text{Pic } X$.*

Proof. Let $P \in X$ be any point, and let $Q = f(P)$ be its image point. Choose local functions φ_P and φ_Q around P (resp. Q) that vanish at P (resp. Q) with multiplicity 1 as in lemma 7.5.6. Then by the definition of the ramification index we have

$$f^* \varphi_Q = u \cdot \varphi_P^{e_P}$$

for some local function u on X with no zero or pole at P . Now pick a global rational section α of ω_Y . If its divisor (α) contains the point Q with multiplicity n , we can write locally

$$\alpha = v \cdot \varphi_Q^n d\varphi_Q,$$


where v is a local function on Y with no zero or pole at Q . Inserting these equations into each other, we see that

$$f^* \alpha = f^* v \cdot (f^* \varphi_Q^n) d(f^* \varphi_Q) = u^n f^* v \cdot \varphi_P^{ne_P} \cdot (\varphi_P^{e_P} du + u e_P \varphi_P^{e_P-1} d\varphi_P).$$

This vanishes at P to order $ne_P + e_P - 1$. Summing this over all points $P \in X$ we see that the divisor of $f^* \alpha$ is $f^*(\alpha) + R$. As $K_X = (f^* \alpha)$ and $f^* K_Y = f^*(\alpha)$, the proposition follows. \square

We will now study the same situation from a topological point of view (if the ground field is \mathbb{C}). Then X and Y are two-dimensional compact manifolds.

For such a space X , we say that a **cell decomposition** of X is given by writing X as a finite disjoint union of points, (open) lines, and discs. This decomposition should be “nice” in a certain topological sense, e. g. the boundary points of every line in the decomposition must be points of the decomposition. It takes some work to make this definition (and the following propositions) bullet-proof. We do not want to elaborate on this, but only remark that every “reasonable” decomposition that one could think of will be allowed. For example, here are three valid decompositions of the Riemann sphere $\mathbb{P}_{\mathbb{C}}^1$:




(In (i), we have only one point (the north pole), no line, and one “disc”, namely \mathbb{P}^1 minus the north pole). We denote by $\sigma_0, \sigma_1, \sigma_2$ the number of points, lines and discs in the decomposition, respectively. So in the above examples we have $(\sigma_0, \sigma_1, \sigma_2) = (1, 0, 1)$, $(2, 2, 2)$, and $(6, 8, 4)$, respectively.

Of course there are many possible decompositions for a given curve X . But there is an important number that is invariant:

Lemma 7.6.3. *The number $\sigma_0 - \sigma_1 + \sigma_2$ depends only on X and not on the chosen decomposition. It is called the (topological) Euler characteristic $\chi(X)$ of X .*

Proof. Let us first consider the case when we move from one decomposition to a “finer” one, i. e. if we add points or lines to the decomposition. For example, in the above pictures (iii) is a refinement of (ii), which is itself a refinement of (i). Note that every refinement is obtained by applying the following steps a finite number of times:

- (i) Adding another point on a line: In this case we raise σ_0 and σ_1 by 1, so the alternating sum $\sigma_0 - \sigma_1 + \sigma_2$ does not change (see the picture below).



- (ii) Adding another line in a disc: In this case we raise σ_1 and σ_2 by 1, so the alternating sum $\sigma_0 - \sigma_1 + \sigma_2$ again does not change (see the picture above).


So we conclude that the alternating sum $\sigma_0 - \sigma_1 + \sigma_2$ does not change under refinements. But it is easily seen that any two decompositions have a common refinement (which is essentially given by taking all the points and lines in both decompositions, and maybe add more points where two such lines intersect. For example, the common refinement of decomposition (ii) above and the same decomposition rotated clockwise by 90 degrees would be (iii)). It follows that the alternating sum is independent of the decomposition. \square

We have already noted in example 0.1.1 that a smooth complex curve is topologically a (real) closed surface with a certain number g of “holes”. The number g is called the genus

of the curve. Let us compute the topological Euler characteristic of such a curve of genus g :

Lemma 7.6.4. *The Euler characteristic of a curve of genus g is equal to $2 - 2g$.*

Proof. Take e. g. the decomposition illustrated in the following picture:



It has $2g + 2$ points, $4g + 4$ lines, and 4 discs, so the result follows. \square

Let us now compare the Euler characteristics of two curves X and Y if we have a morphism $f : X \rightarrow Y$:

Lemma 7.6.5. *Let $f : X \rightarrow Y$ be a morphism of complex smooth projective curves. Let n be the number of inverse image points of any point of Y under f . As in proposition 7.6.2 let R be the ramification divisor of f . Then $-\chi(X) = -n \cdot \chi(Y) + \deg R$.*

Proof. Choose “compatible” decompositions of X and Y , i. e. loosely speaking decompositions such that the inverse images of the points / lines / discs of the decomposition of Y are (finite) unions of points / lines / discs of the decomposition of X , and such that all points / lines / discs of the decomposition of X arise in this way. Moreover, we require that all ramification points of f are points of the decomposition of X . (It is easily seen that this can always be achieved.) Denote by $\sigma_0^X, \sigma_1^X, \sigma_2^X$ the number of points / lines / discs of the decomposition of X , and similarly for Y .

As every point of Y that is not the image of a ramification point has n inverse images under f , it follows that $\sigma_1^X = n\sigma_1^Y$ and $\sigma_2^X = n\sigma_2^Y$. We do not have $\sigma_0^X = n\sigma_0^Y$ however: if P is a ramification point, i. e. $e_P > 1$, then f is locally e_P -to-one around P , i. e. P counts for e_P in $n\sigma_0^Y$, whereas it is actually only one point in the decomposition of X . Hence we have to subtract $e_P - 1$ for any ramification point P from $n\sigma_0^Y$ to get the correct value of σ_0^X . This means that $\sigma_0^X = n\sigma_0^Y - \deg R$ and hence $-\chi(X) = -n\chi(Y) + \deg R$. \square

Corollary 7.6.6. *Let X be a (complex) smooth projective curve. Then $\deg K_X = 2g - 2$.*

Proof. As we have already remarked, any such curve X admits a surjective morphism f to \mathbb{P}^1 by projection. Using that $\deg K_{\mathbb{P}^1} = -\chi(\mathbb{P}^1) = -2$ (by lemma 7.4.15 and lemma 7.6.4) and applying lemma 7.6.5 together with the Riemann-Hurwitz formula 7.6.2, we see that $\deg K_X = -\chi(X)$. The result therefore follows from lemma 7.6.4. \square

7.7. The Riemann-Roch theorem. As in the last section let X be a smooth projective curve of genus g over an algebraically closed field. For any line bundle \mathcal{L} we want to compute the dimensions of the vector spaces $\Gamma(\mathcal{L})$ of global sections of \mathcal{L} . We will denote this dimension by $h^0(\mathcal{L})$ (the reason for this notation will become obvious when we discuss cohomology in chapter 8). By abuse of notation we will also write $h^0(D)$ instead of $h^0(\mathcal{O}(D))$ for any divisor D .

We should remark that this is a classical question that was one of the first problems studied in algebraic geometry: given a smooth projective curve X (resp. a compact one-dimensional complex manifold), points $P_1, \dots, P_r \in X$, and numbers $a_1, \dots, a_r \geq 0$, what is the dimension of the space of rational (resp. meromorphic) functions on X that have poles of order at most a_i at the points P_i and are regular (resp. holomorphic) everywhere else? In our language, this just means that we are looking for the number $h^0(a_1P_1 + \dots + a_rP_r)$.

Example 7.7.1. Let D be a divisor on X with negative degree. Recall that sections of $O(D)$ are just rational functions φ on X such that $(\varphi) + D$ is effective. Taking degrees, this certainly implies that $\deg(\varphi) + \deg D \geq 0$. But $\deg(\varphi) = 0$ by remark 6.3.5 and $\deg D < 0$ by assumption, which is a contradiction. Hence we conclude that $h^0(D) = 0$ if $\deg D < 0$: there are no global sections of $O(D)$ in this case.

Example 7.7.2. Let \mathcal{L} be the line bundle $O_X(n)$ for some $n \in \mathbb{Z}$. Recall that sections of \mathcal{L} are of the form $\frac{f}{g}$ with f and g homogeneous such that $\deg f - \deg g = n$. Now for global sections g must be a constant function (otherwise we would have a pole somewhere), so we conclude that $\Gamma(\mathcal{L})$ is simply the n -th graded piece of the homogeneous coordinate ring $S(X)$. In other words, $h^0(\mathcal{L})$ is by definition equal to the value $h_X(n)$ of the Hilbert function introduced in section 6.1. We have seen in proposition 6.1.5 that $h_X(n)$ is equal to a linear polynomial $\chi_X(n)$ in n for $n \gg 0$. Moreover, the linear coefficient of $\chi_X(n)$ is the degree of $O_X(n)$, and the constant coefficient is $1 - g$ by definition of g (see example 6.1.10). So we conclude that

$$h^0(D) = \deg D + 1 - g$$

if D is the divisor class associated to a line bundle $O_X(n)$ for $n \gg 0$.

Theorem 7.7.3. (Riemann-Roch theorem for line bundles on curves) Let X be a complex smooth projective curve of genus g . Then for any divisor D on X we have

$$h^0(D) - h^0(K_X - D) = \deg D + 1 - g.$$

Proof. Step 1. Recall that for any point $P \in X$ and any divisor D we have the exact “skyscraper sequence” by exercise 7.8.4

$$0 \rightarrow O(D) \rightarrow O(D+P) \rightarrow k_P \rightarrow 0$$

where the last morphism is given by evaluation at the point P . From this we get an exact sequence of global sections

$$0 \rightarrow \Gamma(O(D)) \rightarrow \Gamma(O(D+P)) \rightarrow \mathbb{C}$$

(where the last map is in general not surjective, see example 7.1.18). Therefore $h^0(D+P) - h^0(D)$ is either 0 or 1. If we denote the left hand side of the Riemann-Roch theorem by $\chi(D) = h^0(D) - h^0(K_X - D)$, we conclude that

$$\chi(D+P) - \chi(D) = (h^0(D+P) - h^0(D)) + (h^0(K_X - D) - h^0(K_X - D - P))$$

is either 0, 1, or 2. (Of course, what we want to prove is that $\chi(D+P) - \chi(D)$ is always equal to 1.)

Step 2. We want to rule out the case that $\chi(D+P) - \chi(D) = 2$. For this we actually have to borrow a theorem from complex analysis.

So assume that $h^0(D+P) - h^0(D) = 1$ and $h^0(K_X - D) - h^0(K_X - D - P) = 1$. The fact that $h^0(D+P) - h^0(D) = 1$ means precisely that there is a global section φ of $O_X(D+P)$ that is not a global section of $O_X(D)$, i. e. that φ is a rational section of $O_X(D)$ that has a simple pole at P and is regular at all other points. Similarly, there is a global section α of $O_X(K_X - D)$ that is not a global section of $O_X(K_X - D - P)$. In other words, α is a global section of $\omega_X \otimes \mathcal{L}^\vee$ that does not vanish at P . By multiplication we see that $\varphi \cdot \alpha$ is a rational section of $\mathcal{L} \otimes (\omega_X \otimes \mathcal{L}^\vee) = \omega_X$ that has a simple pole at P and is regular at all other points. In other words, $\varphi \cdot \alpha$ is a global rational differential form with just a single pole which is

of order 1. But this is a contradiction to the residue theorem of complex analysis: the sum of the residues of any rational (or meromorphic) differential form on a compact Riemann surface is zero, but in our case we have $\sum_{Q \in X} \text{res}_Q(\varphi \cdot \alpha) = \text{res}_P(\varphi \cdot \alpha) \neq 0$.

Step 3. We claim that

$$\chi(D) \geq \deg D + 1 - g$$

for all divisors D . Note that we can choose points P_1, \dots, P_r such that $D + P_1 + \dots + P_r$ is precisely the intersection divisor of X with a certain number n of hyperplanes: for every point in D we just choose a hyperplane through that point and add all other intersection points with X to the P_i . This then means that $O(D + P_1 + \dots + P_r) = O(n)$. By possibly adding more intersection points of X with hyperplanes we can make n arbitrarily large. So by example 7.7.2 we find that

$$h^0(D + P_1 + \dots + P_r) = \deg D + r + 1 - g.$$

Moreover, if n (and thus r) is large enough we see by example 7.7.1 that $h^0(K_X - D - P_1 - \dots - P_r) = 0$ and therefore

$$\chi(D + P_1 + \dots + P_r) = \deg D + r + 1 - g.$$

But by step 2 we know that subtracting a point from the divisor will decrease $\chi(\cdot)$ by 0 or 1. If we apply this r times to the points P_1, \dots, P_r we conclude that $\chi(D) \geq (\deg D + r + 1 - g) - r$, as we have claimed.

Step 4. Replacing D by $K_X - D$ in the inequality of step 3 yields

$$\begin{aligned} -\chi(D) &= h^0(K_X - D) - h^0(D) \geq \deg K_X - \deg D + 1 - g \\ &= -\deg D - 1 + g \end{aligned}$$

as $\deg K_X = 2g - 2$ by corollary 7.6.6. Combining the two inequalities of steps 3 and 4 proves the theorem. \square

Remark 7.7.4. If D is the divisor associated to the line bundle $O(n)$ (for any n), note that $\chi(D)$ is just the value $\chi_X(n)$ of the Hilbert polynomial. So for these line bundles we can reinterpret our main proposition 6.1.5 about Hilbert polynomials as follows: the difference between $h_X(n)$ and $\chi_X(n)$ is simply $h^0(\omega_X \otimes O_X(-n))$. As this vanishes for large n by degree reasons, it follows that $h_X(n) = \chi_X(n)$ for large n .

Example 7.7.5. Setting $D = 0$ in the Riemann-Roch theorem yields $h^0(K_X) = g$. This gives an alternate definition of the genus of a smooth projective curve: one could define the genus of such a curve as the dimension of the space of global differential forms. This definition has the advantage that it is immediately clear that it is well-defined and independent of the projective embedding (compare this to example 6.1.10).

Remark 7.7.6. In general one should think of the Riemann-Roch theorem as a formula to compute $h^0(D)$ for any D , modulo an “unwanted” correction term $h^0(K_X - D)$. In many applications one can make this correction term vanish, e. g. by making the degree of D large enough so that $\deg(K_X - D)$ becomes negative.

Remark 7.7.7. There are numerous generalizations of the Riemann-Roch theorem. In fact, there are whole books on Riemann-Roch type theorems. Let us mention some of the generalizations without proof:

- (i) The requirement that the ground field be \mathbb{C} is not essential. The very same statement holds over any algebraically closed ground field (the proof has to be changed though at step 2 where we invoked complex analysis).
- (ii) The requirement that the curve be projective is not essential either, it only needs to be complete (i. e. “compact”).

- (iii) Instead of a line bundle one can take a vector bundle: if \mathcal{F} is any vector bundle on X of rank r then

$$h^0(\mathcal{F}) - h^0(\omega_X \otimes \mathcal{F}^\vee) = \deg \Lambda^r \mathcal{F} + r(1 - g)$$

(see example 10.4.7).

- (iv) There are versions of the Riemann-Roch theorem for *singular* curves as well. (Note that in the singular case we do not have a canonical bundle, so one needs a new idea here.)
- (v) There are also versions of the Riemann-Roch theorem for varieties of dimension bigger than 1 (see theorem 10.4.5).
- (vi) Finally, the same theorem can be proven (with the same proof actually) in complex analysis, where $h^0(D)$ then denotes the dimension of the space of *meromorphic* functions with the specified zeros and poles. As the resulting dimension does change we conclude that on a projective smooth complex curve *every meromorphic function is in fact rational*. This is an example of a very general result that says that complex analysis essentially reduces to algebraic geometry in the projective case (in other words, we “do not gain much” by allowing holomorphic functions instead of rational ones in the first place).

As an application of the Riemann-Roch theorem let us consider again morphisms to projective spaces. Let X be a smooth projective curve, and let D be a divisor on X . Let s_0, \dots, s_r be a basis of the space $\Gamma(O(D))$ of global sections of $O(D)$. Then we have seen in lemma 7.5.14 that we get a morphism

$$X \rightarrow \mathbb{P}^r, \quad P \mapsto (s_0(P) : \dots : s_r(P))$$

provided that the sections s_i do not vanish simultaneously at any point. Using the Riemann-Roch theorem we can now give an easy criterion when this is the case. Note first however that picking a different basis of section would result in a morphism that differs from the old one simply by a linear automorphism of \mathbb{P}^r . Thus we usually say that the divisor D (or its associated line bundle) determines a morphism to \mathbb{P}^r up to automorphisms of \mathbb{P}^r .

Proposition 7.7.8. *Let X be a smooth projective curve of genus g , and let D be a divisor on X .*

- (i) *If $\deg D \geq 2g$ then the divisor D determines a morphism $X \rightarrow \mathbb{P}^r$ as above.*
- (ii) *If $\deg D \geq 2g + 1$ then moreover this morphism is an embedding (i. e. an isomorphism onto its image).*

Proof. (i): By what we have said above we simply have to show that for every point $P \in X$ there is a global section $s \in \Gamma(O(D))$ that does not vanish at P .

By the degree condition we know that $\deg(K_X - D) \leq 2g - 2 - 2g < 0$ and $\deg(K_X - D + P) \leq 2g - 2 - 2g + 1 < 0$. So by example 7.7.1 we get from the Riemann-Roch theorem that

$$h^0(D) = \deg D + 1 - g \quad \text{and} \quad h^0(D - P) = (\deg D - 1) + 1 - g.$$

In particular we have $h^0(D) - h^0(D - P) = 1$, i. e. there is a section $s \in \Gamma(O(D))$ that is not a section of $O(D - P)$, i. e. that does not vanish at P .

(ii): The idea of the proof is the same as in (i). However, as we have not developed enough powerful techniques yet to prove that a morphism has an inverse, we will restrict ourselves to proving that the morphism determined by D is bijective. So let P and Q be distinct points of X . To prove that they are mapped to different points it suffices to show that there is a section $s \in \Gamma(O(D))$ with $s(P) = 0, s(Q) \neq 0$: the morphism $R \mapsto (s(R) : s'(R) : \dots)$ then maps P to a point with the first coordinate 0, while the first coordinate is non-zero for the image point of Q .

To find this section s , simply apply the argument of (i) to $D - P$ and the point Q : we get $h^0(D - P) - h^0(D - P - Q) = 1$, i. e. there is a section $s \in \Gamma(O(D - P))$ that is not a section of $O(D - P - Q)$, i. e. it is a section of $O(D)$ that vanishes at P but not at Q . \square

Example 7.7.9. If X is a smooth projective curve of genus $g \geq 2$ we get a *canonical* embedding $X \rightarrow \mathbb{P}^r$ into a projective space (up to automorphisms by \mathbb{P}^r) by taking the morphism associated to the divisor $3K_X$. This follows by part (ii) of proposition 7.7.8 as $3(2g - 2) \geq 2g + 1$ if $g \geq 2$. By remark 7.7.7 (ii) the same is true for any complete (i. e. “compact”) curve that is not necessarily given initially as a subvariety of projective space.

7.8. Exercises.

Exercise 7.8.1. Let \mathcal{F}' be a presheaf on a topological space X , and let \mathcal{F} be its sheafification as in definition 7.1.10. Show that

- (i) There is a natural morphism $\theta : \mathcal{F}' \rightarrow \mathcal{F}$.
- (ii) Any morphism from \mathcal{F}' to a sheaf factors uniquely through θ .

Exercise 7.8.2. Let $f : \mathcal{F} \rightarrow \mathcal{G}$ be a morphism of sheaves of abelian groups on a topological space X . Show that f is injective / surjective / an isomorphism if and only if all induced maps $f_P : \mathcal{F}_P \rightarrow \mathcal{G}_P$ on the stalks are injective / surjective / isomorphisms.

Exercise 7.8.3. Let $f : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ be a morphism of locally free sheaves on a scheme X over a field k . Let $P \in X$ be a point, and denote by $(\mathcal{F}_i)_P$ the fiber of the vector bundle \mathcal{F}_i over P , which is a k -vector space. Are the following statements true or false:

- (i) If $\mathcal{F}_1 \rightarrow \mathcal{F}_2$ is injective then the induced map $(\mathcal{F}_1)_P \rightarrow (\mathcal{F}_2)_P$ is injective for all $P \in X$.
- (ii) If $\mathcal{F}_1 \rightarrow \mathcal{F}_2$ is surjective then the induced map $(\mathcal{F}_1)_P \rightarrow (\mathcal{F}_2)_P$ is surjective for all $P \in X$.

Exercise 7.8.4. Prove the following generalization of example 7.1.16: If X is a smooth curve over some field k , \mathcal{L} a line bundle on X , and $P \in X$ a point, then there is an exact sequence

$$0 \rightarrow \mathcal{L}(-P) \rightarrow \mathcal{L} \rightarrow k_P \rightarrow 0,$$

where k_P denotes the “skyscraper sheaf”

$$k_P(U) = \begin{cases} k & \text{if } P \in U, \\ 0 & \text{if } P \notin U. \end{cases}$$

Exercise 7.8.5. If X is an *affine* variety over a field k and \mathcal{F} a locally free sheaf of rank r on X , is then necessarily $\mathcal{F} \cong \mathcal{O}_X^{\oplus r}$?

Exercise 7.8.6. Let X be a scheme, and let \mathcal{F} be a locally free sheaf on X . Show that $(\mathcal{F}^\vee)^\vee \cong \mathcal{F}$. Show by example that this statement is in general false if \mathcal{F} is only quasi-coherent but not locally free.

Exercise 7.8.7. Figure out what exactly goes wrong with the correspondence between line bundles and divisor classes on a curve X if X is singular. Can we still associate a divisor to any section of a line bundle? Can we still construct a line bundle from any divisor?

Exercise 7.8.8. What is the line bundle on $\mathbb{P}^n \times \mathbb{P}^m$ leading to the Segre embedding $\mathbb{P}^n \times \mathbb{P}^m \rightarrow \mathbb{P}^N$ by the correspondence of lemma 7.5.14? What is the line bundle leading to the degree- d Veronese embedding $\mathbb{P}^n \rightarrow \mathbb{P}^N$?

Exercise 7.8.9. Show that any smooth projective curve of genus 2 . . .

- (i) can be realized as a curve of degree 5 in \mathbb{P}^3 ,
- (ii) admits a two-to-one morphism to \mathbb{P}^1 . How many ramification points does such a morphism have?

Exercise 7.8.10. Let X be a smooth projective curve, and let $P \in X$ be a point. Show that there is a rational function on X that is regular everywhere except at P .