Algebraic Geometry

(2016/2017 edition)

Eduard Looijenga

Preface

These notes accompany my course Algebraic Geometry I. Every time I taught that course, I revised the text and although I do not expect drastic changes anymore, this is a process that will probably only stop when I cease teaching it. One reason is that these notes are tailored to what I think are the needs of the course and this changes with time. This makes me all the more aware of its deficiencies (by sometimes not giving a topic the treatment it deserves) and omissions (by skipping a nearby point of interest that would have merited discussion). Occasionally I try to make up for this by including some remarks in a smaller font.

As I hope will become clear (and even more so in its sequel, Algebraic Geometry II), much of commutative algebra owes its existence to algebraic geometry and vice versa, and this is why there is no clear border between the two. This also explains why some familiarity with commutative algebra is a prerequisite, but as a service to students lacking such background, I occassionally recall basic facts from that area and from Galois theory (all of it being standard fare in a first course on these subjects), also in a smaller font. Otherwise these notes are essentially self-contained.

On <www.staff.science.uu.nl/∼looij101/> I maintain a web page of this course, where among other things, I briefly explain what this field is about and list some books for further reading. To repeat a recommendation that is made there, I strongly encourage you to buy a (preferably paper!) text book as a companion to use with the course, for such a book generally covers more ground and also tends to do so in a more balanced manner. And it may be consulted, even long after these notes have perished. A good choice is Hartshorne's book (though certainly not the only one), which has the additional benefit that it can also serve you well for a sequel to this course. (That the content of these notes have a substantial overlap with Chapter 1 of that volume is unlikely to be a coincidence.)

You may occasionally find in the text forwarding references to course notes of Algebraic Geometry II. These indeed exist, but as they are in a much more tentative and preliminary form, I have not included them here.

Some conventions. Rings are always supposed to be commutative and to possess a unit and a ring homomorphism is required to take unit to unit. We allow that $1 = 0$, but in that case we get of course the zero ring $\{0\}$ and there cannot be any ring homomorphism going from this ring to a nonzero ring, as it must take unit to unit. Since a prime ideal of a ring is by definition not the whole ring, the zero ring has no prime ideals and hence also no maximal ideals. When R and R' are two rings, then $R \times R'$ is also one for componentwise addition and multiplication, the unit being $(1, 1)$. The projections onto its factors are admitted as ring homomorphisms, but an inclusion obtained by putting one coordinate zero is not, as this is not unital, unless in that coordinate we have the zero ring (in other words, "$\times$" defines a categorical product but not a categorical sum).

We say that a ring is a *domain*([1]) if its zero ideal is a prime ideal, in other words, if the ring is not the zero ring ($1 \neq 0$) and has no zero divisors.

[1] Since we assume all our rings to be commutative and with unit, this is the same notion as *integral domain*.

Given a ring R, then an *R-algebra* is a ring A endowed with a ring homomorphism $\phi : R \to A$. When is ϕ is understood, then for every $r \in R$ and $a \in A$, the product $\phi(r)a$ is often denoted by ra. In case R is a field, ϕ will be injective so that R may be regarded as a subring of A, but this need not be so in general. We say that A is *finitely generated as an R-algebra* if we can find $a_1, \ldots, a_n$ in A such that every element of A can be written as a polynomial in these elements with coefficients in R; in other words, if the R-algebra homomorphism $R[x_1, \ldots, x_n] \to A$ which sends the variable x_i to a_i is onto. This is not to be confused with the notion of finite generation of an R-module M which merely means the existence of a surjective homomorphism of R-modules $R^n \to M$ for some $n \geq 0$.

Similarly, a field L is said to be *finitely generated as a field* over a subfield K if there exist $b_1, \ldots, b_n$ in L such that every element of L can be written as a fraction of two polynomials in these elements (the denominator being nonzero of course) with coefficients in K.

We denote the multiplicative group of the invertible elements (units) of a ring R by $R^\times$.

Contents

Preface 3

Chapter 1. Affine varieties 7
1. The Zariski topology 7
2. Irreducibility and decomposition 10
3. Finiteness properties and the Hilbert theorems 16
4. The affine category 20
5. The sheaf of regular functions 27
6. The product 30
7. Function fields and rational maps 31
8. Finite morphisms 36
9. Normalization 42
10. Dimension 48
11. Smoothness 51
12. Differentials and derivations. 60
13. The notion of a variety 63

Chapter 2. Projective varieties 67
1. Projective spaces 67
2. The Zariski topology on a projective space 69
3. The Segre embeddings 73
4. Blowing up and projections 74
5. Elimination theory and projections 77
6. The Veronese embeddings 80
7. Grassmannians 82
8. Fano varieties and the Gauß map 86
9. Multiplicities of modules 88
10. Hilbert functions and Hilbert polynomials 92
11. Projective curves 97

Bibliography 115

CHAPTER 1

Affine varieties

Throughout these notes k stands for an algebraically closed field, unless we explicitly state otherwise. Recall that this means that every polynomial $f \in k[x]$ of positive degree has a root $x_1 \in k$: $f(x_1) = 0$. This implies that f is divisible by $x - x_1$ with quotient a polynomial of degree one less than f. Continuing in this manner we then find that f decomposes simply as $f(x) = c(x - x_1) \cdots (x - x_d)$ with $c \in k^\times = k \smallsetminus \{0\}$, $d = \deg(f)$ and $x_1, \ldots, x_d \in k$. Since an algebraic extension of k is obtained by the adjunction of certain roots of polynomials in $k[x]$, this also shows that the property in question is equivalent to: every algebraic extension of k is equal to k.

A first example you may think of is the field of complex numbers $\mathbb{C}$, but as we proceed you should become increasingly aware of the fact that there are many others: it is shown in a standard algebra course that for any field F an algebraic closure $\bar{F}$ is obtained by adjoining to F the roots of every polynomial $f \in F[x]$ ([1]). So we could take for k an algebraic closure of the field of rational numbers $\mathbb{Q}$, of the finite field $\mathbb{F}_q$, where q is a prime power or even of the field of fractions of any domain such as $\mathbb{C}[x_1, \ldots, x_r]$.

1. The Zariski topology

Any $f \in k[x_1, \ldots, x_n]$ determines in an evident manner a function $k^n \to k$. In such cases we prefer to think of k^n not as vector space—its origin and vector addition will be irrelevant to us—but as a set with a weaker structure. We shall make this precise later, but it basically amounts to only remembering that elements of $k[x_1, \ldots, x_n]$ can be understood as k-valued functions on it. For that reason it is convenient to denote this set differently, namely as $\mathbb{A}^n$ (or as $\mathbb{A}^n_k$, if we feel that we should not forget about the field k). We refer to $\mathbb{A}^n$ as the *affine n-space over k*. A k-valued function on $\mathbb{A}^n$ is then said to be *regular* if it is defined by some $f \in k[x_1, \ldots, x_n]$. We denote the zero set of such a function by $Z(f)$ and its complement (the nonzero set) by $\mathbb{A}^n_f \subseteq \mathbb{A}^n$.

A *principal* subset of $\mathbb{A}^n$ is any subset of the form $\mathbb{A}^n_f$ and a *hypersurface* of $\mathbb{A}^n$ is any subset of the form $Z(f)$, where in the last case we ask that f be nonconstant (that is, $f \notin k$).

EXERCISE 1. Prove that $f \in k[x_1, \ldots, x_n]$ is completely determined by the regular function it defines. (Hint: do first the case $n = 1$.) So the ring $k[x_1, \ldots, x_n]$ can

[1] This can not be done in one step: it is an infinite process which involves in general many choices. This is reflected by the fact that the final result is not canonical, although it is unique up to a (in general nonunique) isomorphism; whence the use of the indefinite article in 'an algebraic closure'.

be regarded as a ring of functions on $\mathbb{A}^n$ under pointwise addition and multiplication. Show that this fails to be so for the finite field $\mathbb{F}_q$ (which is not algebraically closed).

EXERCISE 2. Prove that a hypersurface is neither empty, nor all of $\mathbb{A}^n$.

It is perhaps somewhat surprising that in this rather algebraic context, the language of topology proves to be quite effective: algebraic subsets of $\mathbb{A}^n$ shall appear as the closed sets of a topology, albeit a rather peculiar one.

Lemma-definition 1.1. The collection of principal subsets of $\mathbb{A}^n$ is a basis of a topology on $\mathbb{A}^n$, called the *Zariski topology*. A subset of $\mathbb{A}^n$ is closed for this topology if and only if it is an intersection of zero sets of regular functions.

PROOF. Recall that a collection $\mathfrak{U}$ of subsets of a set X may serve as a basis for a topology on X (and thus determines this topology) if and only if the intersection of any two its members is a union of members of $\mathfrak{U}$. As the collection of principal subsets is even closed under finite intersection: $\mathbb{A}^n_{f_1} \cap \mathbb{A}^n_{f_2} = \mathbb{A}^n_{f_1 f_2}$, the first assertion follows. Since an open subset of $\mathbb{A}^n$ is by definition a union of subsets of the form $\mathbb{A}^n_f$, a closed subset must be an intersection of subsets of the form $Z(f)$. □

EXAMPLE 1.2. The Zariski topology on $\mathbb{A}^1$ is the cofinite topology: its closed subsets $\neq \mathbb{A}^1$ are the finite subsets.

EXERCISE 3. Show that the diagonal in $\mathbb{A}^2$ is closed for the Zariski topology, but not for the product topology (where each factor $\mathbb{A}^1$ is equipped with the Zariski topology). So $\mathbb{A}^2$ does not have the product topology.

We will explore the mutual relationship between the following two basic maps:

$$\begin{array}{ccc} \{\text{subsets of } \mathbb{A}^n\} & \xrightarrow{\ I\ } & \{\text{ideals of } k[x_1,\ldots,x_n]\} \\ \cup & & \cap \\ \{\text{closed subsets of } \mathbb{A}^n\} & \xleftarrow{\ Z\ } & \{\text{subsets of } k[x_1,\ldots,x_n]\}. \end{array}$$

where for a subset $X \subseteq \mathbb{A}^n$, $I(X)$ is the ideal of $f \in k[x_1,\ldots,x_n]$ with $f|X = 0$ and for a subset $J \subseteq k[x_1,\ldots,x_n]$, $Z(J)$ is the closed subset of $\mathbb{A}^n$ defined by $\cap_{f\in J} Z(f)$. Observe that

$$I(X_1 \cup X_2) = I(X_1) \cap I(X_2) \quad \text{and} \quad Z(J_1 \cup J_2) = Z(J_1) \cap Z(J_2).$$

In particular, both I and Z are inclusion reversing. Furthermore, the restriction of I to closed subsets defines a section of Z: if $Y \subseteq \mathbb{A}^n$ is closed, then $Z(I(Y)) = Y$. We also note that by Exercise 1 $I(\mathbb{A}^n) = (0)$, and that any singleton $\{p\} \subseteq \mathbb{A}^n$ is closed, as it is the common zero set of the degree one polynomials $x_1 - p_1, \ldots, x_n - p_n$.

EXERCISE 4. Prove that $I(\{p\})$ is equal to the ideal generated by these degree one polynomials and that this ideal is maximal.

EXERCISE 5. Prove that the (Zariski) closure of a subset Y of $\mathbb{A}^n$ is equal to $Z(I(Y))$.

Given $Y \subseteq \mathbb{A}^n$, then $f, g \in k[x_1,\ldots,x_n]$ have the same restriction to Y if and only if $f - g \in I(Y)$. So the quotient ring $k[x_1,\ldots,x_n]/I(Y)$ (a k-algebra) can be regarded as a ring of k-valued functions on Y. Notice that this k-algebra does not change if we replace Y by its Zariski closure.

DEFINITION 1.3. Let $Y \subseteq \mathbb{A}^n$ be closed. The k-algebra $k[x_1, \ldots, x_n]/I(Y)$ is called the *coordinate ring* of Y and we denote it by $k[Y]$. A k-valued function on Y is said to be *regular* if it lies in this ring.

So $k[\mathbb{A}^n] = k[x_1, \ldots, x_n]$. Given a closed subset $Y \subseteq \mathbb{A}^n$, then for every subset $X \subseteq \mathbb{A}^n$ we have $X \subseteq Y$ if and only if $I(X) \supseteq I(Y)$, and in that case $I_Y(X) := I(X)/I(Y)$ is an ideal of $k[Y]$: it is the ideal of regular functions on Y that vanish on X. Conversely, an ideal of $k[Y]$ is of the form $J/I(Y)$, with J an ideal of $k[x_1, \ldots, x_n]$ that contains $I(Y)$, and such an ideal defines a closed subset $Z(J)$ contained in Y. So the two basic maps above give rise to such a pair on Y:

$$\begin{array}{ccc} \{\text{subsets of } Y\} & \xrightarrow{I_Y} & \{\text{ideals of } k[Y]\} \\ \cup & & \cap \\ \{\text{closed subsets of } Y\} & \xleftarrow{Z_Y} & \{\text{subsets of } k[Y]\}. \end{array}$$

Exercise 5 tells us what $Z \circ I$ does. We now ask this question for $I \circ Z$. In particular, which ideals of $k[x_1, \ldots, x_n]$ are of the form $I(Y)$ for some Y? Clearly, if $f \in k[x_1, \ldots, x_n]$ is such that some positive power vanishes on Y, then f vanishes on Y. In other words: if $f^m \in I(Y)$ for some $m > 0$, then $f \in I(Y)$. This suggests:

Proposition-definition 1.4. Let R be a ring (as always commutative and with 1) and let $J \subseteq R$ be an ideal. Then the set of $a \in R$ with the property that $a^m \in J$ for some $m > 0$ is an ideal of R, called the *radical* of J and denoted $\sqrt{J}$.

We say that J is a *radical ideal* if $\sqrt{J} = J$.

We say that the ring R is *reduced* if the zero ideal (0) is a radical ideal (in other words, R has no nonzero nilpotents: if $a \in R$ is such that $a^m = 0$, then $a = 0$).

PROOF. We show that $\sqrt{J}$ is an ideal. Let $a, b \in \sqrt{J}$ so that $a^m, b^n \in J$ for certain positive integers m, n. Then for every $r \in R$, $ra \in \sqrt{J}$, since $(ra)^m = r^m a^m \in J$. Similarly $a - b \in \sqrt{J}$, for $(a-b)^{m+n}$ is an R-linear combination of monomials that are multiples of a^m or b^n and hence lie in J. □

EXERCISE 6. Show that a prime ideal is a radical ideal.

Notice that J is a radical ideal if and only if R/J is reduced. The preceding shows that for every $Y \subseteq \mathbb{A}^n$, $I(Y)$ is a radical ideal, so that $k[Y]$ is reduced. The dictionary between algebra and geometry begins in a more substantial manner with

Theorem 1.5 (Hilbert's Nullstellensatz)**.** For every ideal $J \subseteq k[x_1, \ldots, x_n]$ we have $I(Z(J)) = \sqrt{J}$.

We inclusion $\supseteq$ is clear; the hard part is the opposite inclusion (which says that if $f \in k[x_1, \ldots, x_n]$ vanishes on $Z(J)$, then $f^m \in J$ for some positive integer m). We postpone its proof and first discuss some of the consequences.

Corollary 1.6. Let $Y \subseteq \mathbb{A}^n$ be closed. Then the maps I_Y and Z_Y define inclusion reversing bijections that are each others inverse

$$\{\text{closed subsets of } Y\} \leftrightarrow \{\text{radical ideals of } k[Y]\}.$$

Via this bijection $\emptyset \leftrightarrow k[Y]$, $Y \leftrightarrow \sqrt{(0)}$ and

$$\{\text{points of } Y\} \leftrightarrow \{\text{maximal ideals of } k[Y]\},$$

where $I_Y(\{p\}) = \mathfrak{m}_p/I(Y)$, with $\mathfrak{m}_p := (x_1 - p_1, \ldots, x_n - p_n)$.

PROOF. We first prove this for $Y = \mathbb{A}^n$. We already observed that for every closed subset X of $\mathbb{A}^n$ we have $Z(I(X)) = X$. The Nullstellensatz says that for a radical ideal $J \subseteq k[x_1, \ldots, x_n]$, we have $I(Z(J)) = J$. This establishes the bijection. It is clear that via this bijection, $\emptyset \leftrightarrow k[Y]$ and $Y \leftrightarrow \sqrt{(0)}$. It then also follows that the smallest nonempty closed subsets of Y must correspond to the biggest proper radical ideals of $k[x_1, \ldots, x_n]$. Since a singleton is closed and a maximal ideal is a radical ideal (as it is a prime ideal), we thus obtain a bijection between the points of $\mathbb{A}^n$ and the maximal ideals of $k[x_1, \ldots, x_n]$. We already noticed that for a given $p \in \mathbb{A}^n$, $(x_1 - p_1, \ldots, x_n - p_n)$ is a maximal ideal whose zero set is $\{p\}$.

The general case now also follows, because an ideal of $k[Y]$ is of the form $J/I(Y)$ with $J \supseteq (Y)$ and this is a radical ideal if and only if J is one; a maximal ideal of $k[Y]$ corresponds to a maximal ideal of $\mathbb{A}^n$ which contains $I(Y)$. □

Via this (or a very similar) correspondence, algebraic geometry seeks to express geometric properties of Y in terms of algebraic properties of $k[Y]$ and vice versa. In the end we want to forget about the ambient $\mathbb{A}^n$ (for the same reason as one wants to study manifolds in their own right and not as embedded in some $\mathbb{R}^n$).

2. Irreducibility and decomposition

We introduce a property which for most topological spaces is of little interest, but as we will see, is useful and natural for the Zariski topology.

Proposition-definition 2.1. Let Y be a topological space. We say that Y is *irreducible* if (i) it is nonempty and (ii) Y is not the union of two closed proper subsets (or equivalently, any nonempty open subset of Y is dense in Y).

We call a maximal irreducible subset of Y an *irreducible component* of Y.

The proof that the two characterizations of irreducible are indeed equivalent is left as an exercise. It follows from the definition that if $A_1 \subseteq A_2 \subseteq A_3 \subseteq \cdots \subseteq Y$ is a nested sequence of irreducible subsets of the space Y, then $\cup_{i=1}^{\infty} A_i$ is irreducible([2]). It follows that every irreducible subset of Y is contained in an irreducible component of Y and that the irreducible components of Y cover Y. Clearly two distinct irreducible components of Y cannot obey an inclusion relation.

EXERCISE 7. Prove that an irreducible Hausdorff space must consist of a single point. Prove also that an infinite set with the cofinite topology is irreducible.

EXERCISE 8. Let $Y_1, \ldots, Y_s$ be closed subsets of a topological space Y whose union is Y. Prove that every irreducible subset of Y is contained in some Y_i. Deduce that $\{Y_i\}_{i=1}^s$ is the collection of irreducible components of Y if each Y_i is irreducible and $Y_i \subseteq Y_j$ implies $Y_i = Y_j$.

Lemma 2.2. Let Y be a topological space. If Y is irreducible, then every nonempty open subset of Y irreducible. Conversely, if $C \subseteq Y$ is an irreducible subspace, then $\overline{C}$ is also irreducible. In particular, an irreducible component of Y is always closed in Y.

PROOF. Suppose Y is irreducible and let $U \subseteq Y$ be open and nonempty. A nonempty open subset of U is dense in Y and hence also dense in U. So U is irreducible.

[2]This is more generally true for a directed system of irreducible subsets.

Let now $C \subseteq Y$ be irreducible (and hence nonempty). Let $V \subseteq \overline{C}$ be nonempty and open in $\overline{C}$. Then $V \cap C$ is nonempty. It is also open in C and hence dense in C. But then $V \cap C$ is also dense in $\overline{C}$ and so V is dense in $\overline{C}$. So $\overline{C}$ is irreducible. □

Here is what irreducibility means in the Zariski topology.

Proposition 2.3. A closed subset $Y \subseteq \mathbb{A}^n$ is irreducible if and only if $I(Y)$ is a prime ideal (which we recall is equivalent to: $k[Y] = k[x_1, \ldots, x_n]/I(Y)$ is a domain).

PROOF. Suppose Y is irreducible and $f, g \in k[x_1, \ldots, x_n]$ are such that $fg \in I(Y)$. Then $Y \subseteq Z(fg) = Z(f) \cup Z(g)$. Since Y is irreducible, Y is contained in $Z(f)$ or in $Z(g)$. So $f \in I(Y)$ or $g \in I(Y)$, proving that $I(Y)$ is a prime ideal.

Suppose that Y is the union of two closed subsets Y_1 and Y_2 that are both $\neq Y$. Then $I(Y)$ is not a prime ideal: since $Y_i \neq Y$ implies that there exist $f_i \in I(Y_i) \smallsetminus I(Y)$ $(i = 1, 2)$ and then $f_1 f_2$ vanishes on $Y_1 \cup Y_2 = Y$, so that $f_1 f_2 \in I(Y)$. □

One of our first aims is to prove that the irreducible components of any closed subset $Y \subseteq \mathbb{A}^n$ are finite in number and have Y as their union. This may not sound very surprising, but we will see that this reflects some nonobvious algebraic properties. Let us first consider the case of a hypersurface. Since we are going to use the fact that $k[x_1, \ldots, x_n]$ is a unique factorization domain, we begin with recalling that notion.

2.4. UNIQUE FACTORIZATION DOMAINS. Let us first observe that in a ring R without zero divisors two nonzero elements a, b generate the same ideal if and only if b is a unit times a.

DEFINITION 2.5. A ring R is called a *unique factorization domain* (often abbreviated as UFD) if it has no zero divisors and every principal ideal $(a) := Ra$ in R which is neither the zero ideal nor all of R is in unique manner an (unordered) product of principal prime ideals: $(a) = (p_1)(p_2)\cdots(p_s)$ (so the ideals $(p_1), \ldots, (p_s)$ are unique up to order).

Note that last property amounts to the statement that every $a \in R \smallsetminus \{0\}$ is a unit or can be written as a product $a = p_1 p_2 \cdots p_s$ such that each p_i generates a prime ideal and this is unique up to order and multiplication by units: if $a = q_1 q_2 \cdots q_t$ is another such way of writing a, then $t = s$ and $q_i = u_i p_{\sigma(i)}$, where $\sigma \in \mathcal{S}_n$ is a permutation and $u_1 u_2 \cdots u_s = 1$.

For a field (which has no proper principal ideals distinct from (0)) the imposed condition is empty and hence a field is automatically a unique factorization domain. A more substantial example (that motivated this notion in the first place) is $\mathbb{Z}$: a principal prime ideal of $\mathbb{Z}$ is of the form (p), with p a prime number. Every integer $n \geq 2$ has a unique prime decomposition and so $\mathbb{Z}$ is a unique factorization domain.

A basic theorem in the theory of rings asserts that if R is a unique factorization domain, then so is its polynomial ring $R[x]$. This implies (with induction on n) that $R[x_1, \ldots, x_n]$ is one. For a field F, the units of $F[x_1, \ldots, x_n]$ are those of F and a nonzero principal ideal of this ring is prime precisely when it is generated by an irreducible polynomial of positive degree. So then every $f \in F[x_1, \ldots, x_n]$ of positive degree then can be written as a product of irreducible polynomials: $f = f_1 f_2 \cdots f_s$, a factorization that is unique up to order and multiplication of each f_i by a nonzero element of F.

The following proposition connects two notions of irreducibility.

Proposition 2.6. Let $f \in k[x_1, \ldots, x_n]$ have positive degree. Then f is irreducible if and only if $Z(f)$ is. More generally, if $f = f_1 f_2 \cdots f_s$ is a factoring of f into irreducible polynomials, then $Z(f_1), \ldots, Z(f_s)$ are the irreducible components of $Z(f)$ and their union equals $Z(f)$ ([3]). In particular, a hypersurface is the union

[3] But we are not claiming that the $Z(f_i)$'s are pairwise distinct.

of its irreducible components; these irreducible components are hypersurfaces and finite in number.

PROOF. For the first assertion, note $f \in k[x_1, \ldots, x_n]$ is irreducible if and only if f generates a prime ideal. By Proposition 2.3 this is equivalent to $Z(f)$ being an irreducible hypersurface.

It follows that if $f = f_1 f_2 \cdots f_s$ is as in the proposition, then $Z(f) = Z(f_1) \cup \cdots \cup Z(f_s)$ with each $Z(f_i)$ irreducible. To see that $\{Z(f_i)\}_{i=1}^s$ is the collection of irreducible components of $Z(f)$, it suffices, in view of Exercise 8, to prove that any inclusion relation $Z(f_i) \subseteq Z(f_j)$ is necessarily an identity. The inclusion $Z(f_i) \subseteq Z(f_j)$ implies $f_j \in \sqrt{(f_i)}$. Since f_i is irreducible it generates a prime ideal and hence a radical ideal, so that $f_j \in (f_i)$. But f_j is irreducible also and so f_j is a unit times f_i. This proves that $Z(f_j) = Z(f_i)$. □

We continue the discussion of irreducibility with the somewhat formal

Lemma 2.7. For a partially ordered set $(A, \leq)$ the following are equivalent:

(i) $(A, \leq)$ satisfies the *ascending chain condition*: every ascending chain $a_1 \leq a_2 \leq a_3 \leq \cdots$ becomes stationary: $a_n = a_{n+1} = \cdots$ for n sufficiently large.

(ii) Every nonempty subset $B \subseteq A$ has a maximal element, that is, an element $b_0 \in B$ such that there is *no* $b \in B$ with $b > b_0$.

PROOF. (i)⇒(ii). Suppose $(A, \leq)$ satisfies the ascending chain condition and let $B \subseteq A$ be nonempty. Choose $b_1 \in B$. If b_1 is maximal, we are done. If not, then there exists a $b_2 \in B$ with $b_2 > b_1$. We repeat the same argument for b_2. We cannot indefinitely continue in this manner because of the ascending chain condition.

(ii)⇒(i). If $(A, \leq)$ satisfies (ii), then the set of members of any ascending chain has a maximal element, in other words, the chain becomes stationary. □

If we replace $\leq$ by $\geq$, then we obtain the notion of the *descending chain condition* and we find that this property is equivalent to: every nonempty subset $B \subseteq A$ has a minimal element. These properties appear in the following pair of definitions.

DEFINITIONS 2.8. We say that a ring R is *noetherian* if its collection of ideals satisfies the ascending chain condition.

We say that a topological space Y is *noetherian* if its collection of closed subsets satisfies the descending chain condition.

EXERCISE 9. Prove that a subspace of a noetherian space is noetherian. Prove also that a ring quotient of a noetherian ring is noetherian.

EXERCISE 10. Prove that a noetherian space is quasi-compact: every covering of such a space by open subsets contains a finite subcovering.

The interest of the noetherian property is that it is one which is possessed by almost all the rings we encounter and that it implies many finiteness properties without which we are often unable to go very far.

But let us give a nonexample first. The ring $\mathcal{H}(\mathbb{D})$ of holomorphic functions on the unit disk $\mathbb{D} \subseteq \mathbb{C}$ is not noetherian: choose $f \in \mathcal{H}(\mathbb{D})$ such that f has *simple* zeroes in a sequence $(z_i \in \mathbb{D})_{i \geq 1}$ whose terms are pairwise distinct (e.g., $\sin(\pi/(1-z))$). Put $f_n := f(z)(z - z_1)^{-1} \cdots (z - z_n)^{-1}$. Then $f_n = (z - z_{n+1}) f_{n+1}$

and so the ideal in $\mathcal{H}(\mathbb{D})$ generated by f_n is strictly contained in the ideal generated by f_{n+1}. We thus obtain a strictly ascending chain of ideals in $\mathcal{H}(\mathbb{D})$.

On the other hand, the ring of convergent power series $\mathbb{C}\{z\}$ is noetherian (we leave this as a little exercise). Obviously, a field is noetherian. The ring $\mathbb{Z}$ is noetherian: if $I_1 \subseteq I_2 \subseteq \cdots$ is an ascending chain of ideals in $\mathbb{Z}$, then $\cup_{s=1}^{\infty} I_s$ is an ideal of $\mathbb{Z}$, hence of the form (n) for some $n \in \mathbb{Z}$. But if s is such that $n \in I_s$, then clearly the chain is stationary as of index s. (This argument only used the fact that any ideal in $\mathbb{Z}$ is generated by a single element, i.e., that $\mathbb{Z}$ is a principal ideal domain.) That most rings we encounter are noetherian is a consequence of the following theorem.

Theorem 2.9 (Hilbert's basis theorem)**.** If R is a noetherian ring, then so is $R[x]$.

As with the Nullstellensatz, we postpone the proof and discuss some of its consequences first.

Corollary 2.10. If R is a noetherian ring (for example, a field) then so is every finitely generated R-algebra. Also, every closed subset of $\mathbb{A}^n$ is noetherian.

PROOF. The Hilbert basis theorem implies (with induction on n) that the ring $R[x_1, \ldots, x_n]$ is noetherian. By Exercise 9, every quotient ring $R[x_1, \ldots, x_n]/I$ is then also noetherian. But a finitely generated R-algebra is (by definition) isomorphic to some such quotient and so the first statement follows.

Suppose $\mathbb{A}^n \supseteq Y_1 \supseteq Y_2 \supseteq \cdots$ is a descending chain of closed subsets. Then $(0) \subseteq I(Y_1) \subseteq I(Y_2) \subseteq \cdots$ is an ascending chain of ideals. As the latter becomes stationary, so will become the former. □

Proposition 2.11. Let Y be a noetherian space. Then its irreducible components are finite in number and their union equals Y.

PROOF. We first show that every closed subset can be written as a finite union of closed irreducible subsets. First note that the empty set has this property (despite the fact that an irreducible set is nonempty by definition), for a union with empty index set is empty. Let B be the collection of closed subspaces of Y for which this is not possible, i.e., that can *not* be written as a finite union of closed irreducible subsets. Suppose that B is nonempty. According to 2.7 this collection has a minimal element, Z, say. This Z must be nonempty and cannot be irreducible. So Z is the union of two proper closed subsets Z' and Z''. The minimality of Z implies that neither Z' nor Z'' is in B and so both Z' and Z'' can be written as a finite union of closed irreducible subsets. But then so can Z and we get a contradiction.

In particular, there exist closed irreducible subsets $Y_1, \ldots, Y_s$ of Y whose union is Y (if $Y = \emptyset$, take $s = 0$). We may of course assume that no Y_i is contained in some Y_j with $j \neq i$. An application Exercise 8 then shows that the Y_i's are the irreducible components of Y. □

If we apply this to $\mathbb{A}^n$ (endowed as always with its Zariski topology), then we find that every subset $Y \subseteq \mathbb{A}^n$ has a finite number of irreducible components, the union of which is all of Y. If Y is closed in $\mathbb{A}^n$, then so is every irreducible component of Y and according to Proposition 2.3 such an irreducible component is defined by a prime ideal. This allows us to recover the irreducible components of a closed subset $Y \subseteq \mathbb{A}^n$ from its coordinate ring:

Corollary 2.12. Let $Y \subseteq \mathbb{A}^n$ be a closed subset. If C is an irreducible component of Y, then the image $I_Y(C)$ of $I(C)$ in $k[Y]$ is a minimal prime ideal of $k[Y]$ and any minimal prime ideal of $k[Y]$ is so obtained: we thus get a bijective correspondence between the irreducible components of Y and the minimal prime ideals of $k[Y]$.

PROOF. Let C be a closed subset of Y and let $I_Y(C)$ be the corresponding ideal of $k[Y]$. Now C is irreducible if and only if $I(C)$ is a prime ideal of $k[x_1, \ldots, x_n]$, or what amounts to the same, if and only if $I_Y(C)$ is a prime ideal of $k[Y]$. It is an irreducible component if C is maximal for this property, or what amounts to the same, if $I_Y(C)$ is minimal for the property of being a prime ideal of $k[Y]$. □

EXAMPLE 2.13. First consider the set $C := \{(t, t^2, t^3) \in \mathbb{A}^3 \mid t \in k\}$. This is a closed subset of $\mathbb{A}^3$: if we use (x, y, z) instead of (x_1, x_2, x_3), then C is the common zero set of $y - x^2$ and $z - x^3$. Now the inclusion $k[x] \subset k[x, y, z]$ composed with the ring quotient $k[x, y, z] \to k[x, y, z]/(y - x^2, z - x^3)$ is clearly an isomorphism. Since $k[x]$ has no zero divisors, $(y - x^2, z - x^3)$ must be a prime ideal. So C is irreducible and $I(C) = (y - x^2, z - x^3)$.

We now turn to the closed subset $Y \subseteq \mathbb{A}^3$ defined by $xy - z = 0$ and $y^3 - z^2 = 0$. Let $p = (x, y, z) \in Y$. If $y \neq 0$, then we put $t := z/y$; from $y^3 = x^2$, it follows that $y = t^2$ and $z = t^3$ and $xy = z$ implies that $x = t$. In other words, $p \in C$ in that case. If $y = 0$, then $z = 0$, in other words p lies on the x-axis. Conversely, any point on the x-axis lies in Y. So Y is the union of C and the x-axis and these are the irreducible components of Y.

We begin with recalling the notion of localization and we do this in the generality that is needed later.

2.14. LOCALIZATION. Let R be a ring and let S be a *multiplicative subset* of R: $1 \in S$ and S closed under multiplication. Then a ring $S^{-1}R$, together with a ring homomorphism $R \to S^{-1}R$ is defined as follows. An element of $S^{-1}R$ is by definition written as a formal fraction r/s, with $r \in R$ and $s \in S$, with the understanding that $r/s = r'/s'$ if and only if $s''(s'r - sr') = 0$ for some $s'' \in S$. This is a ring indeed: multiplication and subtraction is defined as for ordinary fractions: $r/s.r'/s' = (rr')/(ss')$ and $r/s - r'/s' = (s'r - sr')/(ss')$; it has $0/1$ as zero and $1/1$ as unit element and the ring homomorphism $R \to S^{-1}R$ is simply $r \mapsto r/1$. Observe that the definition shows that $0/1 = 1/1$ if and only if $0 \in S$, in which case $S^{-1}R$ is reduced to the zero ring. We also note that any $s \in S$ maps to an invertible element of $S^{-1}R$, the inverse of $s/1$ being $1/s$ (this is also true when $0 \in S$, for 0 is its own inverse in the zero ring). In a sense (made precise in part (b) of Exercise 11 below) the ring homomorphism $R \to S^{-1}R$ is universal for that property. This construction is called the *localization away from* S.

Of special interest is when $S = \{s^n \mid n \geq 0\}$ for some $s \in R$. We then usually write $R[1/s]$ for $S^{-1}R$. Notice that the image of s in $R[1/s]$ is invertible and that $R[1/s]$ is the zero ring if and only if s is nilpotent.

It is clear that if S does not contain zero divisors, then $r/s = r'/s'$ if and only if $s'r - sr' = 0$; in particular, $r/1 = r'/1$ if and only if $r = r'$, so that $R \to S^{-1}R$ is then injective. If we take S maximal for this property, namely take it to be the set $S(R)$ of nonzero divisors of R (which is indeed multiplicative), then $S(R)^{-1}R$ is called the *total fraction ring* $\mathrm{Frac}(R)$ of R. When R is a domain, $S(R) = R \setminus \{0\}$ and so $\mathrm{Frac}(R)$ is a field, the *fraction field* of R. This gives the following corollary, which hints to the importance of prime ideals in the subject.

Corollary 2.15. An ideal $\mathfrak{p}$ of a ring R is a prime ideal if and only if it is the kernel of a ring homomorphism from R to a field.

PROOF. It is clear that the kernel of a ring homomorphism from R to a field is always a prime ideal. Conversely, if $\mathfrak{p}$ is a prime ideal, then it is the kernel of the composite $R \to R/\mathfrak{p} \hookrightarrow \mathrm{Frac}(R/\mathfrak{p})$. □

EXERCISE 11. Let R be a ring and let S be a multiplicative subset of R.

(a) What is the the kernel of $R \to S^{-1}R$?
(b) Prove that a ring homomorphism $\phi : R \to R'$ with the property that $\phi(s)$ is invertible for every $s \in S$ factors in a unique manner through $S^{-1}R$.
(c) Consider the polynomial ring $R[x_s \,:\, s \in S]$ and the homomorphism of R-algebras $R[x_s : s \in S] \to S^{-1}R$ that sends x_s to $1/s$. Prove that this homomorphism is surjective and that its kernel consists of the $f \in R[x_s \,:\, s \in S]$ which after multiplication by an element of S lie in the ideal generated the degree one polynomials $sx_s - 1$, $s \in S$.

EXERCISE 12. Let R be a ring and let $\mathfrak{p}$ be a prime ideal of R.

(a) Prove that the complement $R \smallsetminus \mathfrak{p}$ is a multiplicative system. The resulting localization $(R \smallsetminus \mathfrak{p})^{-1}R$ is called the *localization at* $\mathfrak{p}$ and is usually denoted $R_{\mathfrak{p}}$.
(b) Prove that $\mathfrak{p}R_{\mathfrak{p}}$ is a maximal ideal of $R_{\mathfrak{p}}$ and that it is the only maximal ideal of $R_{\mathfrak{p}}$. (A ring with a unique maximal ideal is called a *local ring*.)
(c) Prove that the localization map $R \to R_{\mathfrak{p}}$ drops to an isomorphism of fields $\mathrm{Frac}(R/\mathfrak{p}) \to R_{\mathfrak{p}}/\mathfrak{p}R_{\mathfrak{p}}$.
(d) Work this out for $R = \mathbb{Z}$ and $\mathfrak{p} = (p)$, where p is a prime number.
(e) Same for $R = k[x, y]$ and $\mathfrak{p} = (x)$.

Lemma 2.16. Let I be an ideal of a ring R. Then the intersection of all the prime ideals of R containing I equals $\sqrt{I}$. In particular, for every nonnilpotent $a \in R$, there exists a ring homomorphism from R to a field that is nonzero on a.

PROOF. By passing to R/I we are reduced to the case when $I = (0)$: we must then show that the intersection of all the prime ideals of R is the ideal of nilpotent elements. It is easy to see that a nilpotent element lies in every prime ideal. Now for a nonnilpotent $a \in R$ consider the homomorphism $R \to R[1/a]$. The ring $R[1/a]$ is nonzero, hence has a maximal ideal([4]) $\mathfrak{m}$ so that $F := R[1/a]/\mathfrak{m}$ is a field. Then the kernel of the composite $\phi : R \to R[1/a] \to F$ is a prime ideal and a is not in this kernel (for $\phi(a) \in F$ is invertible with inverse the image of $1/a$). □

EXERCISE 13. Let R be a ring. Prove that the intersection of all the maximal ideals of a ring R consists of the $a \in R$ for which $1 + aR \subseteq R^\times$ (i.e., $1 + ax$ is invertible for every $x \in R$). You may use the fact that every proper ideal of R is contained in a maximal ideal.

We can do better if R is noetherian. The following proposition is the algebraic counterpart of Proposition 2.11. Note the similarity between the proofs.

Proposition 2.17. Let R be a noetherian ring. Then for every ideal $I \subset R$, the minimal elements of the collection of prime ideals containing I are finite in number (with every prime ideal containing I containing one of these) and their intersection

[4]Every nonzero ring has a maximal ideal. For noetherian rings, which are our main concern, this is obvious, but in general this follows with transfinite induction, the adoption of which is equivalent to the adoption of the axiom of choice.

is $\sqrt{I}$. In particular, the minimal prime ideals of R are finite in number and their intersection is the ideal of nilpotents $\sqrt{(0)}$.

PROOF. We first make the rather formal observation that R is a radical ideal and indeed appears as a finite (namely empty) intersection of prime ideals. So the collection B of the *radical* ideals $I \subseteq R$ that can *not* be written as an intersection of finitely many prime ideals does not contain R. We prove that B is empty. Suppose otherwise. Since R is noetherian, B will have a maximal member $I_0 \neq R$, say. We then derive a contradiction as follows.

Since I_0 is not a prime ideal, there exist $a_1, a_2 \in R \smallsetminus I_0$ with $a_1 a_2 \in I_0$. Consider the radical ideal $J_i := \sqrt{I_0 + Ra_i}$. The ideal J_i strictly contains I_0 and so cannot belong to B. In other words, J_i is an intersection of finitely many prime ideals. We next show that $J_1 \cap J_2 = I_0$, so that I_0 is an intersection of finitely many prime ideals also, thus arriving contradiction. The inclusion $\supseteq$ is obvious and $\subseteq$ is seen as follows: if $a \in J_1 \cap J_2$, then for $i = 1, 2$, there exists an $n_i > 0$ such that $a^{n_i} \in I_0 + Ra_i$. Hence $a^{n_1+n_2} \in (I_0 + Ra_1)(I_0 + Ra_2) \subseteq I_0$, so that $a \in I_0$.

In particular, for an arbitrary ideal I, $\sqrt{I} = \mathfrak{p}_1 \cap \cdots \cap \mathfrak{p}_s$ for certain prime ideals $\mathfrak{p}_i$. We may of course assume that no $\mathfrak{p}_i$ contains some $\mathfrak{p}_j$ with $j \neq i$ (otherwise, omit $\mathfrak{p}_i$). It now remains to prove that every prime ideal $\mathfrak{p} \supset I$ of R contains some $\mathfrak{p}_i$. If that is not the case, then choose $a_i \in \mathfrak{p}_i \smallsetminus \mathfrak{p}$ $(i = 1, \ldots, s)$. But then $a_1 a_2 \cdots a_s \in \mathfrak{p}_1 \cap \cdots \cap \mathfrak{p}_s = \sqrt{I} \subseteq \mathfrak{p}$ and since $\mathfrak{p}$ is a prime ideal, some factor a_i lies in $\mathfrak{p}$. This is clearly a contradiction. □

EXERCISE 14. Let R be a ring, $S \subseteq R$ be a multiplicative system and denote by $\phi : R \to S^{-1}R$ the natural homomorphism. Prove that the map which assigns to every prime ideal of $S^{-1}R$ its preimage in R under ϕ defines a bijection between the prime ideals of $S^{-1}R$ and the prime ideals of R disjoint with S. Prove also that if S has no zero divisors, then the preimage of the ideal of nilpotents of $S^{-1}R$ is the ideal of nilpotents of R.

3. Finiteness properties and the Hilbert theorems

The noetherian property in commutative algebra is best discussed in the context of modules, even if one's interest is only in rings. We fix a ring R and first recall the notion of an R-module.

The notion of an R-module is the natural generalization of a K-vector space (where K is some field). Let us observe that if M is an (additively written) abelian group, then the set $\mathrm{End}(M)$ of group homomorphisms $M \to M$ is a ring for which subtraction is pointwise defined and multiplication is composition (so if $f, g \in \mathrm{End}(M)$, then $f - g : m \in M \mapsto f(m) - g(m)$ and $fg : m \mapsto f(g(m))$); clearly the zero element is the zero homomorphism and the unit element is the identity. It only fails to obey our convention in the sense that this ring is usually noncommutative. We only introduced it in order to be able state succinctly:

DEFINITION 3.1. An *R-module* is an abelian group M, equipped with a ring homomorphism $R \to \mathrm{End}(M)$.

So any $r \in R$ defines a homomorphism $M \to M$; we usually denote the image of $m \in M$ under this homomorphism simply by rm. If we write out the properties of an R-module structure in these terms, we get: $r(m_1 - m_2) = rm_1 - rm_2$, $(r_1 - r_2)m = r_1 m - r_2 m$, $1.m = m$, $r_1(r_2 m) = (r_1 r_2)m$. If R happens to be field, then we see that an R-module is the same thing as an R-vector space.

The notion of an R-module is quite ubiquitous, once you are aware of it. A simple example is an ideal $I \subseteq R$. Any abelian group M is in a natural manner a $\mathbb{Z}$-module. And a $\mathbb{R}[x]$-module can be understood as an real linear space V (an $\mathbb{R}$-module) endowed with an endomorphism (the image of x in $\mathrm{End}(V)$). A more involved example is the following: if X is a manifold, f is a C^∞-function on X and ω a C^∞-differential p-form on X, then $f\omega$ is also a C^∞ differential p-form on X. Thus the linear space of C^∞-differential forms on X of a fixed degree p is naturally a module over the ring of C^∞-functions on X.

Here are a few companion notions, followed by a brief discussion.

3.2. In what follows M is an R-module. A map $f : M \to N$ from M to an R-module N is called a *R-homomorphism* if it is a group homomorphism with the property that $f(rm) = rf(m)$ for all $r \in R$ and $m \in M$. If f is also bijective, then we call it an *R-isomorphism*; in that case its inverse is also a homomorphism of R-modules.

For instance, given a ring homomorphism $f : R \to R'$, then R' becomes an R-module by $rr' := f(r)r'$ and this makes f a homomorphism of R-modules.

A subset $N \subseteq M$ is called an *R-submodule* of M if it is a subgroup and $rn \in N$ for all $r \in R$ and $n \in N$. Then the group quotient M/N is in a unique manner a R-module in such a way that the quotient map $M \to M/N$ is a R-homomorphism: we let $r(m+N) := rm+N$ for $r \in R$ and $m \in M$. Notice that a R-submodule of R (here we regard R as a R-module) is the same thing as an ideal of R.

Given a subset $S \subseteq M$, then the set of elements $m \in M$ that can be written as $r_1 s_1 + \cdots + r_k s_k$ with $r_i \in R$ and $s_i \in S$ is a R-submodule of M. We call it the *R-submodule of M generated by S* and we shall denote it by RS. If there exists a finite set $S \subseteq M$ such that $M = RS$, then we say that M is *finitely generated* as an R-module.

DEFINITION 3.3. We say that an R-module M is *noetherian* if the collection of R-submodules of M satisfies the ascending chain condition: any ascending chain of R-submodules $N_1 \subseteq N_2 \subseteq \cdots$ becomes stationary.

It is clear that then every quotient module of a noetherian module is also noetherian. The noetherian property of R as a ring (as previously defined) coincides with this property of R as an R-module.

The following two propositions provide the passage from the noetherian property to finite generation:

Proposition 3.4. An R-module M is noetherian if and only if every R-submodule of M is finitely generated as an R-module.

PROOF. Suppose that M is a noetherian R-module and let $N \subseteq M$ be a R-submodule. The collection of finitely generated R-submodules of M contained in N is nonempty. Hence it has a maximal element N_0. If $N_0 = N$, then N is finitely generated. If not, we run into a contradiction: just choose $x \in N \smallsetminus N_0$ and consider $N_0 + Rx$. This is a R-submodule of N. It is finitely generated (for N_0 is), which contradicts the maximal character of N_0.

Suppose now that every R-submodule of M is finitely generated. If $N_1 \subseteq N_2 \subseteq \cdots$ is an ascending chain of R-modules, then the union $N := \cup_{i=1}^{\infty} N_i$ is a R-submodule. Let $\{s_1, \ldots, s_k\}$ be a finite set of generators of N. If $s_\kappa \in N_{i_\kappa}$, and $j := \max\{i_1, \ldots, i_k\}$, then it is clear that $N_j = N$. So the chain becomes stationary as of index j. □

Proposition 3.5. Suppose that R is a noetherian ring. Then every finitely generated R-module M is noetherian.

PROOF. By assumption $M = RS$ for a finite set $S \subseteq M$. We prove the proposition by induction on the number of elements of S. If $S = \emptyset$, then $M = \{0\}$ and

there is nothing to prove. Suppose now $S \neq \emptyset$ and choose $s \in S$, so that our induction hypothesis applies to $M' := RS'$ with $S' = S \smallsetminus \{s\}$: M' is noetherian. But so is M/M', for it is a quotient of the noetherian ring R via the surjective R-module homomorphism $R \to M/M'$, $r \mapsto rs + M'$.

Let now $N_1 \subseteq N_2 \subseteq \cdots$ be an ascending chain of R-submodules of M. Then $N_1 \cap M' \subseteq N_2 \cap M' \subseteq \cdots$ becomes stationary, say as of index j_1. Hence we only need to be concerned for $k \geq j_1$ with the stabilization of $(N_k/(N_{j_1} \cap M'))_{k \geq j_1}$. But for $k \geq j_1$, $N_k/(N_{j_1} \cap M' = N_k/(N_k \cap M') \cong (N_k + M')/M'$, so that this can be regarded as an ascending chain in M/M'. Since M/M' is noetherian, this chain stabilizes (say as of index j_2). So the original chain stabilizes as of index j_2. □

We are now sufficiently prepared for the proofs of the Hilbert theorems. They are gems of elegance and efficiency.

We will use the notion of initial coefficient of a polynomial, which we recall. Given a ring R, then every nonzero $f \in R[x]$ is uniquely written as $r_d x^d + r_{d-1}x^{d-1} + \cdots + r_0$ with $r_d \neq 0$. We call $r_d \in R$ the *initial coefficient* of f and denote it by $\mathrm{in}(f)$. For the zero polynomial, we simply define this to be $0 \in R$. Notice that when $\mathrm{in}(f)\,\mathrm{in}(g)$ nonzero, then it is equal to $\mathrm{in}(fg)$.

PROOF OF THEOREM 2.9. The assumption is here that R is a noetherian ring. In view of Proposition 3.4 we must show that every ideal I of $R[x]$ is finitely generated. Consider the subset $\mathrm{in}(I) := \{\mathrm{in}(f) : f \in I\}$ of R. We first show that this is an ideal of R. If $r \in R$, $f \in I$, then $r\,\mathrm{in}(f)$ equals $\mathrm{in}(rf)$ or is zero and since $rf \in I$, it follows that $r\,\mathrm{in}(f) \in I$. If $f, g \in I$, then $\mathrm{in}(f) - \mathrm{in}(g)$ equals $\mathrm{in}(x^{\deg g} f - x^{\deg f} g)$ or is zero. So $\mathrm{in}(I)$ is an ideal as asserted.

Since R is noetherian, $\mathrm{in}(I)$ is finitely generated: there exist $f_1, \ldots, f_k \in I$ such that $\mathrm{in}(I) = R\,\mathrm{in}(f_1) + \cdots + R\,\mathrm{in}(f_k)$. Let $d_0 := \max\{\deg(f_1), \ldots, \deg(f_k)\}$ and $R[x]_{<d_0}$ the set of polynomials of degree $< d_0$. So $R[x]_{<d_0}$ is the R-submodule of $R[x]$ generated by $1, x, \ldots, x^{d_0-1}$. We claim that

$$I = R[x]f_1 + \cdots + R[x]f_k + (I \cap R[x]_{<d_0}),$$

in other words, that every $f \in I$ is modulo $R[x]f_1 + \cdots + R[x]f_k$ a polynomial of degree $< d_0$. We prove this with induction on the degree d of f. Since for $d < d_0$ there is nothing to prove, assume that $d \geq d_0$. We have $\mathrm{in}(f) = r_1\,\mathrm{in}(f_1) + \cdots + r_k\,\mathrm{in}(f_k)$ for certain $r_1, \ldots r_k \in R$, where we may of course assume that every term $r_i\,\mathrm{in}(f_i)$ is nonzero and hence equal to $\mathrm{in}(r_i f_i)$. Since $\mathrm{in}(f)$ is nonzero, it then equals $\sum_i \mathrm{in}(r_i f_i) = \mathrm{in}(\sum_i r_i f_i x^{d-\deg(f_i)})$. So $f - \sum_i r_i f_i x^{d-\deg(f_i)}$ is an element of I of degree $< d$ and hence lies in $R[x]f_1 + \cdots + R[x]f_k + (I \cap R[x]_{<d_0})$ by our induction hypothesis. Hence so does f.

Our claim implies the theorem: $R[x]_{<d_0}$ is a finitely generated R-module and so a noetherian R-module by Proposition 3.5. Hence the R-submodule $I \cap R[x]_{<d_0}$ is a finitely generated R-module by Proposition 3.4. If $\{f_{k+1}, \ldots, f_{k+l}\}$ is a set of R-generators of $I \cap R[x]_{<d_0}$, then $\{f_1, \ldots, f_{k+l}\}$ is a set of $R[x]$-generators of I. □

For the Nullstellensatz we need another finiteness result.

Proposition 3.6 (Artin-Tate)**.** Let R be a noetherian ring, B an R-algebra and $A \subseteq B$ an R-subalgebra. Assume that B is finitely generated as an A-module. Then A is finitely generated as an R-algebra if and only if B is so.

PROOF. By assumption there exist $b_1, \ldots, b_m \in B$ such that $B = \sum_{i=1}^m Ab_i$.

If there exist $a_1, \dots, a_n \in A$ which generate A as an R-algebra (which means that $A = R[a_1, \dots, a_n]$), then $a_1, \dots, a_n, b_1, \dots, b_m$ generate B as an R-algebra.

Suppose, conversely, that there exists a finite subset of B which generates B as a R-algebra. By adding this subset to $b_1, \dots, b_m$, we may assume that $b_1, \dots, b_m$ also generate B as an R-algebra. Then every product $b_i b_j$ can be written as an A-linear combination of $b_1, \dots, b_m$:

$$b_i b_j = \sum_{k=1}^m a_{ij}^k b_k, \quad a_{ij}^k \in A.$$

Let $A_0 \subseteq A$ be the R-subalgebra of A generated by all the (finitely many) coefficients a_{ij}^k. This is a noetherian ring by Corollary 2.10. It is clear that $b_i b_j \in \sum_k A_0 b_k$ and so $\sum_k A_0 b_k$ is an R-subalgebra of B. Since the $b_1, \dots, b_m$ generate B as an R-algebra, it then follows this is all of B: $B = \sum_k A_0 b_k$. So B is finitely generated as an A_0-module. Since A is an A_0-submodule of B, A is also finitely generated as an A_0-module by Proposition 3.4. It follows that A is a finitely generated R-algebra. □

This has a consequence for field extensions:

Corollary 3.7. A field extension L/K is finite if and only if L is finitely generated as a K-algebra.

PROOF. It is clear that if L is a finite dimensional K-vector space, then L is finitely generated as a K-algebra.

Suppose now $b_1, \dots, b_m \in L$ generate L as a K-algebra. It suffices to show that every b_i is algebraic over K. Suppose that this is not the case. After renumbering we can and will assume that (for some $1 \leq r \leq m$) $b_1, \dots, b_r$ are algebraically independent over K and $b_{r+1}, \dots, b_m$ are algebraic over the quotient field $K(b_1, \dots, b_r)$ of $K[b_1, \dots, b_r]$. So L is a finite extension of $K(b_1, \dots, b_r)$. We apply Proposition 3.6 to $R := K$, $A := K(b_1, \dots, b_r)$ and $B := L$ and find that $K(b_1, \dots, b_r)$ is as a K-algebra generated by a finite subset of $K(b_1, \dots, b_r)$. If $g \in K[b_1, \dots, b_r]$ is a common denominator for the elements of this subset, then clearly $K(b_1, \dots, b_r) = K[b_1, \dots, b_r][1/g]$. Since $K(b_1, \dots, b_r)$ strictly contains $K[b_1, \dots, b_r]$, g must have positive degree. In particular, $g \neq 1$, so that $1/(1-g) \in K(b_1, \dots, b_r)$ can be written as f/g^N, with $f \in K[b_1, \dots, b_r]$. Here we may of course assume that f is not divisible in $K[b_1, \dots, b_r]$ by g. From the identity $f(1-g) = g^N$ we see that $N \geq 1$ (for the left hand side has positive degree). But then $f = g(f + g^{N-1})$ shows that f is divisible by g. We thus get a contradiction. □

Corollary 3.8. Let A be a finitely generated k-algebra. Then for every maximal ideal $\mathfrak{m} \subseteq A$, the natural map $k \to A \to A/\mathfrak{m}$ is an isomorphism of fields.

PROOF. Since $\mathfrak{m}$ is maximal, $A/\mathfrak{m}$ is a field that is also finitely generated as a k-algebra. By corollary 3.7, $k \to A/\mathfrak{m}$ is then a finite extension of k. Since k is algebraically closed, this extension will be the identity. □

EXERCISE 15. Prove that a field which is finite generated as a ring (i.e., is isomorphic to a quotient of $\mathbb{Z}[x_1, \dots, x_n]$ for some n) is finite.

We deduce from the preceding corollary the Nullstellensatz.

PROOF OF THE NULLSTELLENSATZ 1.5. Let $J \subseteq k[x_1, \dots, x_n]$ be an ideal. We must show that $I(Z(J)) \subseteq \sqrt{J}$. This amounts to: for every $f \in k[x_1, \dots, x_n] \smallsetminus \sqrt{J}$

there exists a $p \in Z(J)$ for which $f(p) \neq 0$. Consider $k[x_1, \ldots, x_n]/J$ and denote by $\bar{f} \in k[x_1, \ldots, x_n]/J$ the image of f. Since $\bar{f}$ is not nilpotent,

$$A := (k[x_1, \ldots, x_n]/J)[1/\bar{f}].$$

is not the zero ring and so has a maximal ideal $\mathfrak{m} \subseteq A$. Observe that A is a finitely generated k-algebra (we can take the images of $x_1, \ldots, x_n$ and $1/\bar{f}$ as generators) and so the map $k \to A/\mathfrak{m}$ is by Corollary 3.8 an isomorphism. Denote by $\phi : k[x_1, \ldots, x_n] \to A \to A/\mathfrak{m} \cong k$ the corresponding surjection and put $p_i := \phi(x_i)$ and $p := (p_1, \ldots, p_n) \in \mathbb{A}^n$. So if we view x_i as a function on $\mathbb{A}^n$, then $\phi(x_i)$ is the value of x_i at p. The fact that ϕ is a homomorphism of k-algebras implies that it is then given as 'evaluation in p': for any $g \in k[x_1, \ldots, x_n]$ we have $\phi(g) = g(p)$. Since the kernel of ϕ contains J, every $g \in J$ will be zero in p, in other words, $p \in Z(J)$. On the other hand, $f(p) = \phi(f)$ is invertible, for it has the image of $1/\bar{f}$ in $A/\mathfrak{m} \cong k$ as its inverse. In other words, $f(p) \neq 0$. □

4. The affine category

We begin with specifying the maps between closed subsets of affine spaces that we wish to consider.

DEFINITION 4.1. Let $X \subseteq \mathbb{A}^m$ and $Y \subseteq \mathbb{A}^n$ be closed subsets. We say that a map $f : X \to Y$ is *regular* if the components $f_1, \ldots, f_n$ of f are regular functions on X (i.e., are given by the restrictions of polynomial functions to X).

Composition of a regular function on Y with f yields a regular function on X (for if we substitute in a polynomial of n variables $g(y_1, \ldots, y_n)$ for every variable y_i a polynomial $f_i(x_1, \ldots, x_m)$ of m variables, we get a polynomial of m variables). So f then induces a k-algebra homomorphism $f^* : k[Y] \to k[X]$. This property is clearly equivalent to f being regular. The same argument shows that if $f : X \to Y$ and $g : Y \to Z$ are regular maps, then so is their composite $gf : X \to Z$. So we have a category (with objects the closed subsets of some affine space $\mathbb{A}^n$ and regular maps as defined above). In particular, we have a notion of isomorphism: a regular map $f : X \to Y$ is an *isomorphism* if is has a two-sided inverse $g : Y \to X$ which is also a regular map. This implies that $f^* : k[Y] \to k[X]$ has a two-sided inverse $g^* : k[X] \to k[Y]$ which is also an homomorphism of k-algebras, and hence is an isomorphism of k-algebras.

There is also a converse:

Proposition 4.2. Let be given closed subsets $X \subseteq \mathbb{A}^m$ and $Y \subseteq \mathbb{A}^n$ and a k-algebra homomorphism $\phi : k[Y] \to k[X]$. Then there is a unique regular map $f : X \to Y$ such that $f^* = \phi$.

PROOF. The inclusion $j : Y \subseteq \mathbb{A}^n$ defines a k-algebra homomorphism $j^* : k[y_1, \ldots, y_n] \to k[Y]$ with kernel $I(Y)$. Put $f_i := \phi j^*(y_i) \in k[X]$ $(i = 1, \ldots, n)$ and define $f = (f_1, \ldots, f_n) : X \to \mathbb{A}^n$, so that $f^* y_i = f_i = \phi j^* y_i$. Since the k-algebra homomorphisms $f^*, \phi j^* : k[y_1, \ldots, y_n] \to k[X]$ coincide on the generators y_i, they must be equal: $f^* = \phi j^*$. It follows that f^* is zero on the kernel $I(Y)$ of j^*, which means that f takes its values in $Z(I(Y)) = Y$, and that the resulting map $k[Y] \to k[X]$ equals ϕ. The proof of uniqueness is left to you. □

In particular, an isomorphism of k-algebras $k[Y] \to k[X]$ comes from a unique isomorphism $X \to Y$. In the special case of an inclusion of a closed subset $Z \subseteq Y$,

the induced map $k[Y] \to k[Z]$ is of course the formation of the quotient algebra $k[Z] = k[Y]/I_Y(X)$. So $f : X \to Y$ is an isomorphism of X onto a closed subset of Y (we then say that f is a *closed immersion*) if and only if $f^* : k[Y] \to k[X]$ is a surjection of k-algebras (with $\ker(f^*)$ being the ideal defining the image of f).

We complete the picture by showing that any finitely generated *reduced* k-algebra A is isomorphic to some $k[Y]$; the preceding then shows that Y is unique up to isomorphism. Since A is finitely generated as a k-algebra, there exists a surjective k-algebra homomorphism $\phi : k[x_1, \ldots, x_n] \to A$. If we put $I := \mathrm{Ker}(\phi)$, then ϕ induces an isomorphism $k[x_1, \ldots, x_n]/I \cong A$. Put $Y := Z(I) \subseteq \mathbb{A}^n$. Since A is reduced, I is a radical ideal and hence equal to $I(Y)$ by the Nullstellensatz. It follows that ϕ factors through a k-algebra isomorphism $k[Y] \cong A$.

We may sum up this discussion in categorical language as follows.

Proposition 4.3. The map which assigns to a closed subset of some $\mathbb{A}^n$ its coordinate ring defines an anti-equivalence between the category of closed subsets of affine spaces (whose morphisms are the regular maps) and the category of reduced finitely generated k-algebras (whose morphisms are k-algebra homomorphisms). It makes closed immersions correspond to epimorphisms of such k-algebras.

EXAMPLE 4.4. Consider the regular map $f : \mathbb{A}^1 \to \mathbb{A}^2$, $f(t) = (t^2, t^3)$. The maps $\mathbb{A}^1$ bijectively onto the hypersurface (curve) C defined by $x^3 - y^2 = 0$: the image is clearly contained in C and the inverse sends $(0,0)$ to 0 and is on $C \smallsetminus \{(0,0)\}$ given by $(x, y) \mapsto y/x$. The Zariski topology on $\mathbb{A}^1$ and C is the cofinite topology and so this is even a homeomorphism. In order to determine whether the inverse is regular, we consider f^*. We have $k[C] = k[x, y]/(x^3 - y^2)$, $k[\mathbb{A}^1] = k[t]$ and $f^* : k[C] \to k[t]$ is given by $x \mapsto t^2, y \mapsto t^3$. This algebra homomorphism is not surjective for its image misses $t \in k[t]$. In fact, f identifies $k[C]$ with the subalgebra $k + t^2k[t]$ of $k[t]$. So f is not an isomorphism.

EXAMPLE 4.5. An *affine-linear transformation* of k^n is of the form $x \in k^n \mapsto g(x) + a$, where $a \in k^n$ and $g \in \mathrm{GL}(n, k)$ is a linear transformation. Its inverse is $y \mapsto g^{-1}(y-a) = g^{-1}(y)-g^{-1}(a)$ and so of the same type. When we regard such an affine linear transformation as a map from $\mathbb{A}^n$ to itself, then it is regular: its coordinates $(g_1, \ldots, g_n)$ are polynomials of degree one. So an affine-linear transformation is also an isomorphism of $\mathbb{A}^n$ onto itself. When $n \geq 2$, there exist automorphisms of $\mathbb{A}^n$ not of this form. For instance $\sigma : (x, y) \mapsto (x, y + x^2)$ defines an automorphism of $\mathbb{A}^2$ with inverse $(x, y) \mapsto (x, y - x^2)$ (see also Exercise 17). This also shows that the group of affine-linear translations in $\mathbb{A}^n$ is not a normal subgroup, for conjugation by σ takes the transformation $(x, y) \mapsto (x + y, y)$ to an automorphism that is not affine-linear (check this). Hence the group of affine-linear transformations of $\mathbb{A}^n$ is not a "natural" subgroup of the automorphism group of $\mathbb{A}^n$ (this makes that the name *affine n-space* for $\mathbb{A}^n$ is a bit unfortunate).

EXERCISE 16. Let $C \subseteq \mathbb{A}^2$ be the 'circle', defined by $x^2 + y^2 = 1$ and let $p_0 := (-1, 0) \in C$. For every $p = (x, y) \in C \smallsetminus \{p_0\}$, the line through p_0 and p has slope $f(p) = y/(x + 1)$. Denote by $\sqrt{-1} \in k$ a root of the equation $t^2 + 1 = 0$.

(a) Prove that when $\mathrm{char}(k) \neq 2$, f defines an isomorphism([5]) onto $\mathbb{A}^1 \smallsetminus \{\pm\sqrt{-1}\}$.
(b) Consider the map $g : C \to \mathbb{A}^1$, $g(x,y) := x + \sqrt{-1}y$. Prove that when $\mathrm{char}(k) \neq 2$, g defines an isomorphism of C onto $\mathbb{A}^1 \smallsetminus \{0\}$.
(c) Prove that when $\mathrm{char}(k) = 2$, the defining polynomial $x^2 + y^2 - 1$ for C is the square of a degree one polynomial so that C is a line.

EXERCISE 17. Let $f_1, \ldots, f_n \in k[x_1, \ldots, x_n]$ be such that $f_1 = x_1$ and $f_i - x_i \in k[x_1, \ldots, x_{i-1}]$ for $i = 2, \ldots, x_n$. Prove that f defines an isomorphism $\mathbb{A}^n \to \mathbb{A}^n$.

EXAMPLE 4.6. QUADRATIC HYPERSURFACES IN CASE $\mathrm{char}(k) \neq 2$. Let $H \subseteq \mathbb{A}^n$ be a hypersurface defined by a polynomial of degree two:

$$f(x_1, \ldots, x_n) = \textstyle\sum_{1 \leq i \leq j \leq n} a_{ij} x_i x_j + \sum_{i=1}^n a_i x_i + a_0.$$

By means of a linear transformation the quadratic form $\sum_{1 \leq i \leq j \leq n} a_{ij} x_i x_j$ can be brought in diagonal form (this involves splitting off squares, hence requires that $2 \in k^\times$). This means that we can make all the coefficients a_{ij} with $i \neq j$ vanish. Another diagonal transformation (which replaces x_i by $\sqrt{a_{ii}} x_i$ when $a_{ii} \neq 0$) takes every nonzero coefficient a_{ii} to 1 and then renumbering the coordinates (which is also a linear transformation) brings f into the form $f(x_1, \ldots, x_n) = \sum_{i=1}^r x_i^2 + \sum_{i=1}^n a_i x_i + a_0$ for some $r \geq 1$. Splitting off squares once more enables us to get rid of $\sum_{i=1}^r a_i x_i$ so that we get

$$f(x_1, \ldots, x_n) = \textstyle\sum_{i=1}^r x_i^2 + \sum_{i=r+1}^n a_i x_i + a_0.$$

We now have the following cases.

If the nonsquare part is identically zero, then we end up with the equation $\sum_{i=1}^r x_i^2 = 0$ for H.

If the linear part $\sum_{i=r+1}^n a_i x_i$ is nonzero (so that we must have $r < n$), then an affine-linear transformation which does not affect $x_1, \ldots, x_r$ and takes $\sum_{i=r+1}^n a_i x_i + a_0$ to $-x_n$ yields the equation $x_n = \sum_{i=1}^r x_i^2$. This is the graph of the function $\sum_{i=1}^r x_i^2$ on $\mathbb{A}^{n-1}$ and so H is then isomorphic to $\mathbb{A}^{n-1}$.

If the linear part $\sum_{i=r+1}^n a_i x_i$ is zero, but the constant term a_0 is nonzero, then we can make another diagonal transformation which replaces x_i by $\sqrt{-a_0} x_i$) and divide f by a_0: then H gets the equation $\sum_{i=1}^r x_i^2 = 1$.

In particular, there are only a finite number of quadratic hypersurfaces up to isomorphism. (This is also true in characteristic two, but the discussion is a bit more delicate.)

The previous discussion (and in particular Proposition 4.3) leads us to associate to any finitely generated k-algebra A in a direct manner a space (which we shall denote by $\mathrm{Spm}(A)$) which for $A = k[X]$ yields a space homeomorphic to X. Since in that case the points of X correspond to maximal ideals of $k[X]$, we simply choose the underlying set of $\mathrm{Spm}(A)$ to be the collection of maximal ideals of A. For $x \in \mathrm{Spm}(A)$, we shall denote the corresponding maximal ideal of A by $\mathfrak{m}_x$. Since A is finitely generated as a k-algebra, $A/\mathfrak{m}_x$ can be identified with k by Corollary 3.8. We denote the resulting k-algebra homomorphism $A \to k$ by ρ_x. It is clear that any

[5] We have not really defined yet what is an isomorphism between two nonclosed subsets of an affine space. Interpret this here as: f^* maps $k[x,y][1/(x+1)]/(x^2+y^2-1)$ (the algebra of regular functions on $C \smallsetminus \{p_0\}$) isomorphically onto $k[t][1/(t^2+1)]$ (the algebra of regular functions on $\mathbb{A}^1 \smallsetminus \{\pm\sqrt{-1}\}$). This will be justified by Proposition 4.8.

k-algebra homomorphism $A \to k$ has a maximal ideal of A as its kernel and so we may also think of $\mathrm{Spm}(A)$ as the set of k-algebra homomorphisms $A \to k$.

Any $f \in A$ defines defines a 'regular function' $\bar{f} : \mathrm{Spm}(A) \to k$ which takes in $x \in \mathrm{Spm}(A)$ the value $\rho_x(f) \in k$. So its zero set $Z(f) \subseteq \mathrm{Spm}(A)$ is the set of $x \in \mathrm{Spm}(A)$ with $f \in \mathfrak{m}_x$. We denote the complement $\mathrm{Spm}(A) \smallsetminus Z(f)$ by $\mathrm{Spm}(A)_f$. We have $Z(ff') = Z(f) \cup Z(f')$ (for $\rho_x(ff') = \rho_x(f)\rho_x(f')$) and hence $\mathrm{Spm}(A)_{ff'} = \mathrm{Spm}(A)_f \cap \mathrm{Spm}(A)_{f'}$. So the collection of $\{\mathrm{Spm}(A)_f\}_{f\in A}$ is the basis of a topology on $\mathrm{Spm}(A)$. Note that a subset $\mathrm{Spm}(A)$ is closed precisely if it is an intersection of subsets of the form $Z(f)$; this is equal to the common zero set of the set of functions defined by an ideal of A. For $A = k[x_1, \ldots, x_n]/I$, the above discussion shows that $\mathrm{Spm}(A)$ can be identified with $Z(I) \subseteq \mathbb{A}^n$ as a topological space and that under this identification, $\overline{A} := A/\sqrt{(0)}$ becomes the ring of regular functions on $Z(I)$. More generally, for any finitely generated k-algebra A, the map $f \mapsto \bar{f}$ maps A onto a subalgebra of the algebra of k-valued functions on $\mathrm{Spm}(A)$ with kernel the ideal of nilpotents (Exercise 19).

The space $\mathrm{Spm}(A)$ is called the *maximal ideal spectrum*([6]) of R (but our notation for it is less standard). In case A is a finitely generated k-algebra, we refer to $\mathrm{Spm}(A) = \mathrm{Spm}(\overline{A})$ as an *affine variety* (we will give a more complete definition later). We then recover A as its algebra of regular functions.

We observe for later reference:

Lemma 4.7. The maximal ideal spectrum $\mathrm{Spm}(A)$ is quasi-compact: every open covering of $\mathrm{Spm}(A)$ admits a finite subcovering.

PROOF. It suffices to verify this for an open covering by principal open subsets. So let $S \subseteq A$ be such that $\mathrm{Spm}(A) = \cup_{g\in S} \mathrm{Spm}(A)_g$. This means that $\cap_{s\in S} Z(g) = \emptyset$. So the ideal generated by S is not contained in any maximal ideal and hence must be all of A. In particular, $1 = \sum_{i=1}^n f_i g_i$ for certain $f_i \in A$ and $g_i \in S$. It follows that $\{g_i\}_{i=1}^n$ generates A so that $\mathrm{Spm}(A) = \cup_{i=1}^n \mathrm{Spm}(A)_{g_i}$. □

A homomorphism $\phi : A \to B$ of finitely generated k-algebras gives rise to a map $\mathrm{Spm}(\phi) : \mathrm{Spm}(B) \to \mathrm{Spm}(A)$: if $y \in \mathrm{Spm}(B)$, then the composite homomorphism $\rho_y\phi : A \to k$ is the identity map when restricted to k so that $\phi^{-1}\mathfrak{m}_y$ is a maximal ideal of A with residue field k. We thus get a map

$$\mathrm{Spm}(\phi) : \mathrm{Spm}(B) \to \mathrm{Spm}(A).$$

characterized by $\rho_{\mathrm{Spm}(\phi)(y)} = \rho_y\phi$. Note that for $f \in A$, $\overline{\phi(f)}$ is the composite

$$\mathrm{Spm}(B) \xrightarrow{\mathrm{Spm}(\phi)} \mathrm{Spm}(A) \xrightarrow{\overline{f}} k.$$

In particular, the preimage of $Z(f)$ under $\mathrm{Spm}(\phi)$ is $Z(\phi(f))$ and hence the preimage of $\mathrm{Spm}(A)_f$ is $\mathrm{Spm}(B)_{\phi(f)}$. This shows that $\mathrm{Spm}(\phi)$ is continuous. We call the resulting pair $(\mathrm{Spm}(\phi), \phi)$ a *morphism*.

Proposition 4.8. Let A be a finitely generated k-algebra. Then for every $g \in A$, $A[1/g]$ is a finitely generated k-algebra (which is reduced when A is) and the natural k-algebra homomorphism $A \to A[1/g]$ induces a homeomorphism of $\mathrm{Spm}(A[1/g])$ onto $\mathrm{Spm}(A)_g = X \smallsetminus Z(g)$. Moreover, for $g, g' \in A$ the following are equivalent:

[6] I. Gelfand was presumably the first to consider this, albeit in the context of functional analysis: he characterized the Banach algebras that appear as the algebras of continuous $\mathbb{C}$-valued functions on compact Hausdorff spaces. So it might be appropriate to call this the Gelfand spectrum.

(i) $\mathrm{Spm}(A)_g \subseteq \mathrm{Spm}(A)_{g'}$,
(ii) g' divides some positive power of g,
(iii) there exists a A-homomorphism $A[1/g'] \to A[1/g]$ (which must then be unique).

PROOF. It is clear that $A[1/g]$ is a k-algebra and is as such finitely generated (just add to a generating set for A the generator $1/g$). We show that if A is reduced, then so is $A[1/g]$. For this we may suppose that g is not nilpotent (otherwise $A[1/g]$ is the zero ring). Suppose that $f/g^r \in A[1/g]$ is nilpotent: $(f/g^r)^m = 0/1$ for some $m \geq 1$. This means that there exists an $n \geq 0$ such that $f^m g^n = 0$. Then $(fg^n)^m = 0$ and since A is reduced it follows that $fg^n = 0$. So $f/g^r = fg^n/g^{r+n} = 0$ in $A[1/g]$.

A point of $A[1/g]$ is given by a k-algebra homomorphism $A[1/g] \to k$. This is the same thing as to give a k-algebra homomorphism $A \to k$ that is nonzero on g, in other words a point of $\mathrm{Spm}(A)_g$. So the map $A \to A[1/g]$ induces an injection of $\mathrm{Spm}(A[1/g])$ in $\mathrm{Spm}(A)$ with image $\mathrm{Spm}(A)_g$. The map $\mathrm{Spm}(A[1/g]) \to \mathrm{Spm}(A)$ is a morphism and hence continuous. To see that it is also open, note that a principal open subset of $\mathrm{Spm}(A[1/g])$ is of the form $\mathrm{Spm}(A[1/g])_{f/g^n}$, with $f \in A$. By the preceding discussion we may identify this with $\mathrm{Spm}(A[1/g][g^n/f]) = \mathrm{Spm}(A[1/(fg)])$ and so its image in $\mathrm{Spm}(A)$ is the open subset $\mathrm{Spm}(A)_{fg}$.

We check the equivalence of the three conditions.

$(i) \Rightarrow (ii)$ If $\mathrm{Spm}(A)_g \subseteq \mathrm{Spm}(A)_{g'}$, then $Z(g) \supseteq Z(g')$ and so by the Nullstellensatz, $g \in \sqrt{(g')}$. This implies that we can write $g^n = fg'$ for some $f \in A$ and some $n \geq 1$ and (ii) follows.

$(ii) \Rightarrow (iii)$ If $g^n = fg'$ for some $f \in A$, then we have an A-homomorphism $A[1/g'] \to A[1/(fg')] = A[1/g^n] = A[1/g]$ that is easily checked to be independent of the choices made for n and f' and so (iii) follows.

$(iii) \Rightarrow (i)$ If we have an A-homomorphism $A[1/g'] \to A[1/g]$, then we get a morphism $\mathrm{Spm}(A[1/g]) \to \mathrm{Spm}(A[1/g'])$ whose composition with the identification of $\mathrm{Spm}(A[1/g'])$ with the open subset $\mathrm{Spm}(A)_{g'} \subseteq \mathrm{Spm}(A)$ yields the identification of $\mathrm{Spm}(A[1/g])$ with the open subset $\mathrm{Spm}(A)_g \subseteq \mathrm{Spm}(A)$. This means that $\mathrm{Spm}(A)_g \subseteq \mathrm{Spm}(A)_{g'}$ and shows at the same time that such an A-homomorphism is unique. □

From now on we identify the principal open subset $\mathrm{Spm}(A)_g$ with the maximal ideal spectrum $\mathrm{Spm}(A[1/g])$.

EXERCISE 18. Give an example of ring homomorphism $\phi : S \to R$ and a maximal ideal $\mathfrak{m} \subseteq R$, such that $\phi^{-1}\mathfrak{m}$ is not a maximal ideal of S. (Hint: take a look at Exercise 12.)

EXERCISE 19. Prove that if A is a finitely generated k-algebra, then the map $f \in A \mapsto \bar{f}$ is a k-algebra homomorphism from A onto the algebra of k-valued regular functions on $\mathrm{Spm}(A)$ with kernel $\sqrt{(0)}$. Show that for every subset $X \subseteq \mathrm{Spm}(A)$, the set $I(X)$ of $f \in A$ with $\bar{f}|X = 0$ is a radical ideal of A.

Let $f : X \to Y$ be a morphism between affine varieties. Since f is continuous, a fiber $f^{-1}(y)$, or more generally, the preimage $f^{-1}Z$ of a closed subset $Z \subseteq Y$, will be closed in X. It is the zero set of the ideal in $k[X]$ generated by $f^*I(Z)$. In fact, any ideal $I \subseteq k[X]$ can arise this way, for if $(f_1, \ldots, f_r) \in k[X]$ generate I, then take $f = (f_1, \ldots, f_r) : X \to \mathbb{A}^r = Y$ and $y = 0$. We will later see that the failure of

$f^*I(Z)$ to be a radical ideal is sometimes a welcome property, as it can be exploited to define multiplicity. Here is a very simple example.

EXAMPLE 4.9. Let $f : X = \mathbb{A}^1 \to \mathbb{A}^1 = Y$ be defined by $f(x) = x^2$. Then $f^* : k[y] \to k[x]$ is given by $f^*y = x^2$. If we assume k not to be of characteristic 2, and we take $a \in Y \smallsetminus \{0\}$, then the fiber $f^{-1}(a)$ is defined by the ideal generated by $f^*(y-a) = x^2 - a$. It consists of two distinct points that are the two roots of $x^2 = a$, denoted $\pm\sqrt{a}$ and the pair of evaluation maps $(\rho_{\sqrt{a}}, \rho_{-\sqrt{a}})$ identifies the coordinate ring $k[x]/(x^2-a)$ with $k \oplus k$. However, the fiber over $0 \in Y = \mathbb{A}^1$ is the singleton $\{0\} \subset X = \mathbb{A}^1$ and the ideal generated by $f^*y = x^2$ is not a radical ideal. This example indicates that there might good reason to accept nilpotent elements in the coordinate ring of $f^{-1}(0)$ by endowing $f^{-1}(0)$ with the ring of functions $k[f^{-1}(0)] := k[x]/(x^2)$. Since this is a k-vector space of dimension 2 (a k-basis is defined by the pair $\{1, x\}$), we thus retain the information that two points have come together and the fiber should be thought of as a point with multiplicity 2.

Example 4.4 shows that a continuous bijection (and even a homeomorphism) of affine varieties need not be an isomorphism. We next discuss a class of examples of an entirely different nature. It involves a notion that plays a central role in algebraic geometry when the base field k has positive characteristic.

EXAMPLE 4.10 (THE FROBENIUS MORPHISM). Assume that k has positive characteristic p and consider the morphism $\Phi_p : \mathbb{A}^1 \to \mathbb{A}^1$, $a \mapsto a^p$. If we remember that $\mathbb{A}^1$ can be identified with k, then we observe that under this identification, Φ_p is a field automorphism: $\Phi_p(a-b) = (a-b)^p = a^p - b^p = \Phi_p(a) - \Phi_p(b)$ (and of course $\Phi_p(ab) = (ab)^p = \Phi_p(a)\Phi_p(b)$). This shows that Φ_p is injective. Since k is algebraically closed, every element of k has a pth root and so Φ_p is also surjective. But the endomorphism Φ_p^* of $k[x]$ induced by Φ_p sends $\sum_{i=0}^n c_i x^i$ to $\sum_{i=0}^n c_i x^{pi}$ and has therefore image $k[x^p]$. Clearly, Φ_p^* is not surjective.

The fixed point set of Φ_p (the set of $a \in \mathbb{A}^1$ with $a^p = a$) is via the identification of $\mathbb{A}^1$ with k just the prime subfield $\mathbb{F}_p \subset k$ and we therefore denote it by $\mathbb{A}^1(\mathbb{F}_p) \subset \mathbb{A}^1$. Likewise, the fixed point set $\mathbb{A}^1(\mathbb{F}_{p^r})$ of Φ_p^r is the subfield of k with p^r elements. Since the algebraic closure $\overline{\mathbb{F}}_p$ of $\mathbb{F}_p$ in k is the union of the finite subfields of k, the affine line over $\overline{\mathbb{F}}_p$ equals $\cup_{r\geq 1}\mathbb{A}^1(\mathbb{F}_{p^r})$. This generalizes in a straightforward manner to higher dimensions: by letting Φ_p act coordinatewise on $\mathbb{A}^n$, we get a morphism $\mathbb{A}^n \to \mathbb{A}^n$ (which we still denote by Φ_p) that is also a bijection. The fixed point of Φ_p^r is $\mathbb{A}^n(\mathbb{F}_{p^r})$ and $\mathbb{A}^n(\overline{\mathbb{F}}_p) = \cup_{r\geq 1}\mathbb{A}^n(\mathbb{F}_{p^r})$.

EXERCISE 20. Assume that k has positive characteristic p. Let $q = p^r$ be a power of p with $r > 0$ and denote by $\mathbb{F}_q \subset k$ the subfield of $a \in k$ satisfying $a^q = a$. We write Φ_q for $\Phi_p^r : a \in \mathbb{A}^n \mapsto a^q \in \mathbb{A}^n$.

(a) Prove that $f \in k[x_1, \ldots, x_n]$ is in $\mathbb{F}_q[x_1, \ldots, x_n]$ if and only if $\Phi_q f = f^q$.
(b) Prove that an affine-linear transformation of $\mathbb{A}^n$ with coefficients in $\mathbb{F}_q$ commutes with Φ_q.
(c) Let $Y \subseteq \mathbb{A}^n$ be the common zero set of a subset of $\mathbb{F}_q[x_1, \ldots, x_n] \subset k[x_1, \ldots, x_n]$. Prove that Φ_q restricts to a bijection $\Phi_{Y,q} : Y \to Y$ and that the fixed point set of $\Phi_{Y,q}^m$ is $Y(\mathbb{F}_{q^m}) := Y \cap \mathbb{A}^n(\mathbb{F}_{q^m})$.
(d) Suppose that k is an algebraic closure of $\mathbb{F}_p$. Prove that every closed subset $Y \subseteq \mathbb{A}^n$ is defined over a finite subfield of k and hence is invariant under some positive power of Φ_p.

REMARK 4.11. After this exercise we cannot resist to mention the Weil zeta function. This function and its relatives—among them the Riemann zeta function—codify arithmetic properties of algebro-geometric objects in a very intricate manner. In the situation of Exercise 20, we can use the numbers $|Y(\mathbb{F}_{q^m})|$ (= the number of fixed points of Φ^m in Y) to define a generating series $\sum_{m\geq 1}|Y(\mathbb{F}_{q^m})|t^m$. It appears to be more convenient to work with the *Weil zeta function*:

$$Z_Y(t) := \exp\Big(\sum_{m=1}^{\infty}|Y(\mathbb{F}_{q^m})|\frac{t^m}{m}\Big),$$

which has the property that $t\frac{d}{dt}\log Z_Y$ yields the generating series above. This series has remarkable properties. For instance, a deep theorem due to Bernard Dwork (1960) asserts that it represents a rational function of t. Another deep theorem, due to Pierre Deligne (1974), states that the roots of the numerator and denominator have for absolute value a nonpositive half-integral power of q and that moreover, these powers have an interpretation in terms of an 'algebraic topology for algebraic geometry', as was predicted by André Weil in 1949. (This can be put in a broader context by making the change of variable $t = q^{-s}$. Indeed, now numerator and denominator have their zeroes when the real part of s is a nonnegative half-integer and this makes Deligne's result reminiscent of the famous conjectured property of the Riemann zeta function.)

EXERCISE 21. Compute the Weil zeta function of affine n-space relative to the field of q elements.

REMARK 4.12. The Frobenius morphism as defined above should not be confused with pth power map $F_A : a \in A \mapsto a^p \in A$ that we have on any commutative $\mathbb{F}_p$-algebra A and that is sometimes referred to as the *absolute Frobenius*. This is an $\mathbb{F}_p$-algebra endomorphism, but in case A is in fact a k-algebra, *not* a k-algebra endomorphism, for it is on k also the pth power map (the usual Frobenius F_k) and so not the identity. We can in a sense remedy this by replacing the ring homomorphism $i : k \hookrightarrow A$ that makes A a k-algebra by its precomposite with F_k, i.e., by replacing $i : k \hookrightarrow A$ by $iF_k : k \hookrightarrow A$ (so that $F_k i(\lambda) = i(\lambda)^p$). Precisely, we define the *Frobenius twist* of A as $A^{(p)} := k\ {}_{F_k}\!\otimes_i A$, where the tensor product is relative to $F_k : k \to k$ on the first factor and the structural map $i : k \hookrightarrow A$ on the second, so that in $A^{(p)}$ we have $\tau \otimes \lambda a = \tau\lambda^p \otimes a$ for all $\lambda, \tau \in k$ and $a \in A$. The map

$$\tilde{F}_A := F_k \otimes \mathbf{1}_A : A \to A^{(p)},\ \lambda \otimes_i a \mapsto \lambda^p\ {}_{F_k}\!\otimes_i a$$

is an isomorphism of $\mathbb{F}_p$-algebras and via this identification, A gets indeed endowed with the structural map $iF_k : k \hookrightarrow A$. Note that $F_{A^{(p)}/k} := F_A\tilde{F}_A^{-1} : A^{(p)} \to F_{A^{(p)}/k}A$ is now a homomorphism of k-algebras which sends $\tau \otimes a$ to τa^p. So the top row of the commutative diagram below is a factorization of F_A:

$$\begin{array}{ccccc} A & \xrightarrow{\tilde{F}_A} & A^{(p)} & \xrightarrow{F_{A^{(p)}/k}} & A \\ {\scriptstyle i}\uparrow & & {\scriptstyle 1\otimes i}\uparrow & & {\scriptstyle i}\uparrow \\ k & \xrightarrow{F_k} & k & \xrightarrow{=} & k \end{array}$$

The induced map $\mathrm{Spm}(A^{(p)}) \to \mathrm{Spm}(A)$ is the identity, but the map $\mathrm{Spm}(A) \to \mathrm{Spm}(A^{(p)})$ induced by $F_{A^{(p)}/k}$ is in general not. Observe that the other composite $(F_k \otimes 1)F_{A^{(p)}/k} : A^{(p)} \to A^{(p)}$ is just the pth power map $F_{A^{(p)}}$ of $A^{(p)}$. So if we put $B := A^{(p)}$ and write $B^{(1/p)}$ for A, then this composite can be written as $F_B = \tilde{F}_B^{-1}F_{B/k} : B \to B^{(1/p)} \to B$.

Iterating this r times yields a k-algebra $A^{(q)}$ with $q = p^r$. If it so happens that A is given to us as obtained from a $\mathbb{F}_q$-algebra A_o by extension of scalars: $A = k \otimes_{\mathbb{F}_q} A_o$ (where we

have identified $\mathbb{F}_q$ with a subfield of k), then $A^{(q)} = k^{(q)} \otimes_{\mathbb{F}_q} A_o$, and since the qth power map $k \to k^{(q)}$ is a field isomorphism, we can use that isomorphism to identify $A^{(q)}$ with A *as a k-algebra*. So the qth power map then *does* determine a k-algebra homomorphism $A \to A$ (given by $c \otimes_{\mathbb{F}_q} a_o \mapsto c \otimes_{\mathbb{F}_q} a_o^q$). It is called the *geometric Frobenius*. For $A_o = \mathbb{F}_q[x_1, \ldots, x_n]$ (and more generally for the coordinate ring of an Y as above), this yields our Φ_q^*.

5. The sheaf of regular functions

In any topology or analysis course you learn that the notion of continuity is *local*: there exists a notion of continuity at a point so that a function is continuous if it is so at every point of its domain. We shall see that in algebraic geometry the property for a function to be regular is also local in nature.

Let X be a topological space, $x \in X$ and R a set. Consider pairs (U, ϕ), where U is a neighborhood of x in X and $\phi : U \to R$ a map. We define a relation on such pairs by: $(U, \phi) \sim (U', \phi')$ if there exists a neighborhood of x in $U \cap U'$ to which ϕ and ϕ' have the same restriction. This is clearly an equivalence relation. An equivalence class is called a *germ of an R-valued map at x*. We may of course restrict ourselves here to U belonging to a given neighborhood basis of x. Such a germ has a well-defined value in x, but not in general in any other point of X. Note that when R is a ring, then the germs of R-valued maps at x make up an R-algebra.

We use this notion in the situation where $X = \mathrm{Spm}(A)$, with A a reduced finitely generated k-algebra, $R = k$, and the maps in question are regular functions on principal neighborhoods of x. So we represent a germ of a regular function at x by a pair (X_g, ϕ), where $g \in A \smallsetminus \mathfrak{m}_x$ and $\phi \in A[1/g]$ and $(X_g, \phi) \sim (X_{g'}, \phi')$ precisely when there exists a neighborhood U of x in $X_g \cap X_{g'}$ such that $\phi|U = \phi'|U$. The germs of regular functions on X at x form a k-algebra, which we shall denote by $\mathcal{O}_{X,x}$. In fact $\mathcal{O}_{X,x}$ is nothing but the localization $A_{\mathfrak{m}_x} = (A \smallsetminus \mathfrak{m}_x)^{-1}A$ (so that $\mathcal{O}_{X,x}$ is a local ring): any $\phi \in A_{\mathfrak{m}_x}$ is represented by a fraction f/g with $f \in A$ and $g \in A \smallsetminus \mathfrak{m}_x$, hence comes from a regular function on the principal neighborhood X_g of x. And if ϕ is also given as f'/g', then we have $(fg' - f'g)g''$ for some $g'' \in A \smallsetminus \mathfrak{m}_x$, which just means that f/g and f'/g' define the same element of $A[1/(gg'g'')]$, where we note that $gg'g'' \notin \mathfrak{m}_x$ so that $\mathrm{Spm}(A[1/(gg'g'')]) = X_{gg'g''}$ is a principal neighborhood of x in X.

Observe that ρ_x defines a surjective 'evaluation homomorphism' $\rho_{X,x} : \mathcal{O}_{X,x} \to k$: it takes any $\phi \in \mathcal{O}_{X,x}$ as above to $\rho_x(f)/\rho_x(g)$. So ϕ is invertible in $\mathcal{O}_{X,x}$ precisely when $\rho_x(f) \neq 0$, or equivalently, when $\rho_{X,x}(\phi) \neq 0$ (with its inverse represented by g/f) and hence the kernel of $\rho_{X,x}$ is the maximal ideal $\mathfrak{m}_{X,x}$ of the local ring $\mathcal{O}_{X,x}$.

DEFINITION 5.1. We say that a k-valued function ϕ defined on an open subset U of X is *regular* at $x \in U$ if its restriction to some principal neighborhood of x in X is so and thus defines an element of $\mathcal{O}_{X,x}$. We denote by $\mathcal{O}(U)$ the set of k-valued functions $U \to k$ that are regular at every point of U .

Note that is $\mathcal{O}(U)$ is in fact a k-algebra. We would like to call an element of $\mathcal{O}(U)$ a *regular function on U*, but we have that notion already defined in case $U = X$, or more generally, when U is a principal open subset. Fortunately, there is no conflict here:

Proposition 5.2. Let X be an affine variety. Then for every principal open subset $X_g \subseteq X$, the natural k-algebra homomorphism $k[X][1/g] \to \mathcal{O}(X_g)$ is an isomorphism. A map $F : X \to Y$ between affine varieties is a morphism if and only if F is continuous and for any $\phi \in \mathcal{O}(V)$ (with $V \subseteq Y$ open) we have $F^*\phi = \phi F \in \mathcal{O}(F^{-1}V)$.

PROOF. Recall that X_g is affine with $k[X_g] = [X][1/g]$. So for the proof of the first statement we may without loss of generality assume that $X_g = X$. The map $k[X] \to \mathcal{O}(X)$ is injective: an element in the kernel is an element is an element $k[X]$ which is zero as a function on X and since X is affine, it is then zero in $k[X]$.

To prove surjectivity, let $\phi \in \mathcal{O}(X)$. We must show that ϕ is representable by some $f \in k[X]$. By assumption there exist for every $x \in X$, a $g_x \in k[X] \smallsetminus \mathfrak{m}_x$ and an $f_x \in k[X]$ such that $\phi|X_{g_x}$ is representable as f_x/g_x.

Since X is quasicompact (Lemma 4.7), the covering $\{X_{g_x}\}_{x\in X}$ of X has a finite subcovering $\{X_{g_{x_i}}\}_{i=1}^N$. Let us write f_i for f_{x_i} and g_i for g_{x_i}. Then f_i/g_i and f_j/g_j define the same regular function on $X_{g_i} \cap X_{g_j} = X_{g_ig_j}$ and so $g_if_j - g_jf_i$ is annihilated by $(g_ig_j)^{m_{ij}}$ for some $m_{ij} \geq 0$. Let m be the maximum of these exponents m_{ij}, so that $g_i^{m+1}g_j^m f_j = g_j^{m+1}g_i^m f_i$ for all i, j. Upon replacing f_i by $f_ig_i^m$ and g_i by g_i^{m+1}, we may then assume that in fact $g_if_j = g_jf_i$ for all i, j.

Since $\cup_i X_{g_i} = X$, we have $\cap_i Z(g_i) = \emptyset$. In other words, the ideal $(g_1, \ldots, g_N) \subset k[X]$ is not contained in a maximal ideal and so must be all of $k[X]$: $1 = \sum_{i=1}^N h_ig_i$ for certain $h_i \in k[X]$. Now consider $f := \sum_{i=1}^N h_if_i \in k[X]$. We have for every j,

$$fg_j = \textstyle\sum_{i=1}^N h_if_ig_j = \sum_{i=1}^N h_ig_if_j = f_j$$

and so the restriction of f to X_{g_j} is equal to f_j/g_j. As this is also the restriction of ϕ to X_{g_j} and $\cup_j X_{g_j} = X$, it follows that ϕ is represented by f.

The last statement is left as an exercise. □

Let us denote by $\mathcal{O}_X$ the collection of the k-algebras $\mathcal{O}(U)$, where U runs over all open subsets of X. The preceding proposition says that $\mathcal{O}_X$ is a *sheaf* of k-valued functions on X, by which we mean the following:

DEFINITION 5.3. Let X be a topological space and R a ring. A *sheaf $\mathcal{O}$ of R-valued functions*([7]) on X assigns to every open subset U of X an R-subalgebra $\mathcal{O}(U)$ of the R-algebra of R-valued functions on U with the property that

(i) for every inclusion $U \subseteq U'$ of open subsets of X, 'restriction to U' maps $\mathcal{O}(U')$ in $\mathcal{O}(U)$ and
(ii) given a collection $\{U_i\}_{i\in I}$ of open subsets of X, then a function $f : \cup_{i\in I}U_i \to R$ lies in $\mathcal{O}(\cup_i U_i)$ if and only if $f|U_i \in \mathcal{O}(U_i)$ for all i.

If $(X, \mathcal{O}_X)$ and $(Y, \mathcal{O}_Y)$ are topological spaces endowed with a sheaf of R-valued functions, then a continuous map $f : X \to Y$ is called a *morphism* if for every open $V \subseteq Y$, composition with f takes $\mathcal{O}_Y(V)$ to $\mathcal{O}_X(f^{-1}V)$.

This definition simply expresses the fact that the functions we are considering are characterized by a local property—just as we have a sheaf of continuous $\mathbb{R}$-valued functions on a topological space, a sheaf of differentiable $\mathbb{R}$-valued functions on a manifold and a sheaf of holomorphic $\mathbb{C}$-valued functions on a complex

[7] We give the general definition of a sheaf later. This will do for now. A defect of this definition is that a sheaf of R-valued functions on a space X need not restrict to one on a subspace of X.

manifold. With the notion of a morphism, we have a category of topological spaces endowed with a sheaf of R-valued functions. In particular, we have the notion of isomorphism: this is a homeomorphism $f : X \to Y$ which for every open $V \subseteq Y$ maps $\mathcal{O}_Y(V)$ onto $\mathcal{O}_X(f^{-1}V)$. Note that a sheaf $\mathcal{O}$ of R-valued functions on X restricts to a sheaf $\mathcal{O}|U$ for every open $U \subseteq X$.

We are now ready to introduce the notion of an affine variety in a more proper fashion. We take our cue from the definition of a manifold.

DEFINITION 5.4. A topological space X endowed with a sheaf $\mathcal{O}_X$ of k-valued functions is called an *affine variety* when it is isomorphic to a pair $(\mathrm{Spm}(A), \mathcal{O}_{\mathrm{Spm}(A)})$ as above. We refer to $\mathcal{O}_X(X)$ as its *coordinate ring* and usually denote it by $k[X]$.

We call $(X, \mathcal{O}_X)$ a *quasi-affine variety* if it is isomorphic to an open subset of some pair $(\mathrm{Spm}(A), \mathcal{O}_{\mathrm{Spm}(A)})$.

Thus a reduced finitely generated k-algebra defines an affine variety and conversely, an affine variety determines a reduced finitely generated k-algebra. These two assignments are inverses of each other. It follows from Proposition 5.2 that a map $F : X \to Y$ between affine varieties is a morphism in the old sense if and only if is one as a morphism between spaces endowed with a sheaf of k-valued functions. This exhibits the category of affine varieties as a full subcategory of the category of spaces endowed with a sheaf of k-valued functions.

The present definition lends itself better than the previous one to immediate generalization (e.g., when we will introduce the notion of a variety) and has other technical advantages as well. Here is an example.

EXAMPLE 5.5. Here is an example an affine open subset of an affine variety that is not principal. Take the cuspidal plane cubic curve $C \subset \mathbb{A}^2$ of Example 4.4 defined by $y^2 = x^3$ and assume that k *is of characteristic zero*. As we have seen, the parametrization $f : t \in \mathbb{A}^1 \mapsto (t^2, t^3) \in C$ identifies $k[C]$ with the subalgebra $k + t^2k[t]$ of $k[t]$. Now let $a \in \mathbb{A}^1 \smallsetminus \{0\}$. So $U := C \smallsetminus \{f(a)\}$ is quasi-affine. But U is not a principal open subset: it is not of the form C_g for some $g \in k[x, y]$. For then $f^*(g)$ would have a as its only zero, so that f^*g is a nonzero constant times $(t-a)^n$. But the coefficient of t in $(t-a)^n$ is $n(-a)^{n-1}$, and hence nonzero. This contradicts the fact that $f^*g \in k + t^2k[t]$.

We claim however that U is affine, with $k[U]$ via f^* identified with $k[t^2, t^3, t^2/(t-a)]$. Clearly, $k[t^2, t^3, t^2/(t-a)]$ is a finitely generated k-algebra. Since it is contained in the reduced k-algebra $k[t][1/(t-a)]$, it is also reduced and so it defines an affine variety $\tilde{U}$. The inclusion $k[t^2, t^3] \subseteq k[t^2, t^3, t^2/(t-a)]$ defines a morphism $j : \tilde{U} \to C$. We prove that j is an isomorphism of $\tilde{U}$ onto U in the sense of Definition 5.4 by checking this over two open subsets U_0 and U_a of C that cover C: it will then follow that U is affine.

We let $U_a := C_{x|C}$. The ideal generated by $f^*(x|C) = t^2$ in $k[\mathbb{A}^1]$ defines $\{0\}$ and so $U_a = C \smallsetminus \{f(0)\}$. We note that $k[U_a] = k[C][1/t^2] = k[t^2, t^3, t^{-2}] = k[t, t^{-1}]$ and hence $k[U_a \smallsetminus \{f(a)\}] = k[U_a][1/(t-a)]$. On the other hand,

$$k[j^{-1}U_a] = k[t^2, t^3, t^2/(t-a)][t^{-2}] = \\ = k[t, t^{-1}, 1/(t-a)] = k[U_a][1/(t-a)] = k[U_a \smallsetminus \{f(a)\}].$$

This proves that $j^{-1}U_a$ maps isomorphically onto $U_a \smallsetminus \{f(0)\}$. It remains to show that there is neighborhood U_0 of $(0,0) \in C \smallsetminus \{f(a)\}$ such that j maps $j^{-1}U_0$ isomorphically onto U_0. Take for U_0 the principal open subset $C_{x|C-a^2}$. Then $k[U_0] = k[t^2, t^3][1/(t^2 - a^2)]$ and $k[j^{-1}U_0] = k[t^2, t^3, t^2/(t-a)][1/(t^2 - a^2)]$. But these k-algebras are the same, because $t^2/(t-a) = (t^3 + at^2)/(t^2 - a^2)] \in k[t^2, t^3][1/(t^2 - a^2)]$, and so $j : j^{-1}U_0 \cong U_0$.

Other such examples (among them smooth plane cubic curves) that are also valid in positive characteristic are best understood after we have discussed the Picard group.

6. The product

Let m and n be nonnegative integers. If $f \in k[\mathbb{A}^m] = k[x_1, \dots, x_m]$ and $g \in k[\mathbb{A}^n] = k[y_1, \dots, y_n]$, then we define $f * g \in k[\mathbb{A}^{m+n}] = k[x_1, \dots, x_m, y_1, \dots, y_n]$ by

$$f * g(x_1, \dots, x_m, y_1, \dots, y_n) := f(x_1, \dots, x_m) g(y_1, \dots, y_n).$$

It is clear that $\mathbb{A}^{m+n}_{f*g} = \mathbb{A}^m_f \times \mathbb{A}^n_g$, which shows that the Zariski topology on $\mathbb{A}^{m+n}$ refines the product topology on $\mathbb{A}^m \times \mathbb{A}^n$. Equivalently, if $X \subseteq \mathbb{A}^m$ and $Y \subseteq \mathbb{A}^n$ are closed, then $X \times Y$ is closed in $\mathbb{A}^{m+n}$. We give $X \times Y$ the topology it inherits from $\mathbb{A}^{m+n}$ (which is finer than the product topology when $m > 0$ and $n > 0$). For the coordinate rings we have defined a map:

$$k[X] \times k[Y] \to k[X \times Y], \quad (f, g) \mapsto f * g$$

which is evidently k-bilinear (i.e., k-linear in either variable). We want to prove that the ideal $I(X \times Y)$ defining $X \times Y$ in $\mathbb{A}^{m+n}$ is generated by $I(X) \cup I(Y)$ (viewed as a subset of $k[x_1, \dots, x_m, y_1, \dots y_n]$) and that $X \times Y$ is irreducible when X and Y are. This requires that we translate the formation of the product into algebra. The translation centers around the notion of the tensor product, the definition of which we recall. (Although we here only need tensor products over k, we shall define this notion for modules over a ring, as this is its natural habitat and is the setting that is needed later anyhow.)

If R is a ring and M and N are R-modules, then we can form their *tensor product over* R, $M \otimes_R N$: as an abelian group $M \otimes_R N$ is generated by the expressions $a \otimes_R b$, $a \in M$, $b \in N$ and subject to the conditions $(ra) \otimes_R b = a \otimes_R (rb)$, $(a + a') \otimes_R b = a \otimes_R b + a' \otimes_R b$ and $a \otimes_R (b + b') = a \otimes_R b + a \otimes_R b'$. So a general element of $M \otimes_R N$ can be written like this: $\sum_{i=1}^N a_i \otimes_R b_i$, with $a_i \in M$ and $b_i \in N$. We make $M \otimes_R N$ an R-module if we stipulate that $r(a \otimes_R b) := (ra) \otimes_R b$ (which is then also equal to $a \otimes_R (rb)$). Notice that the map

$$\otimes_R : M \times N \to M \otimes_R N, \quad (a, b) \mapsto a \otimes_R b,$$

is R-bilinear (if we fix one of the variables, then it becomes an R-linear map in the other variable).

In case $R = k$ we shall often omit the suffix k in $\otimes_k$ and so write $\otimes$ instead.

EXERCISE 22. Prove that $\otimes_R$ is universal for this property in the sense that every R-bilinear map $M \times N \to P$ of R-modules is the composite of $\otimes_R$ and a *unique* R-homomorphism $M \otimes_R N \to P$. In other words, the map

$$\mathrm{Hom}_R(M \otimes_R N, P) \to \mathrm{Bil}_R(M, N; P)), \quad f \mapsto f \circ \otimes_R$$

is an isomorphism of R-modules.

EXERCISE 23. Let m and n be nonnegative integers. Prove that $\mathbb{Z}/(n) \otimes_{\mathbb{Z}} \mathbb{Z}/(m)$ can be identified with $\mathbb{Z}/(m, n)$.

If A is an R-algebra and N is an R-module, then $A \otimes_R N$ acquires the structure of an A-module which is characterized by

$$a.(a' \otimes_R b) := (aa') \otimes_R b.$$

For instance, if N is an $\mathbb{R}$-vector space, then $\mathbb{C} \otimes_{\mathbb{R}} N$ is a complex vector space, the *complexification* of N. If A and B are R-algebras, then $A \otimes_R B$ acquires the structure of an R-algebra characterized by

$$(a \otimes_R b).(a' \otimes_R b') := (aa') \otimes_R (bb').$$

Notice that $A \to A \otimes_R B$, $a \mapsto a \otimes_R 1$ and $B \to A \otimes_R B$, $b \mapsto 1 \otimes_R b$ are R-algebra homomorphisms. For example, $A \otimes_R R[x] = A[x]$ as A-algebras (and hence $A \otimes_R R[x_1, \ldots, x_n] = A[x_1, \ldots, x_n]$ with induction).

EXERCISE 24. Prove that $\mathbb{C} \otimes_{\mathbb{R}} \mathbb{C}$ is as a $\mathbb{C}$-algebra isomorphic to $\mathbb{C} \oplus \mathbb{C}$ with componentwise multiplication.

Proposition 6.1. For closed subsets $X \subseteq \mathbb{A}^m$ and $Y \subseteq \mathbb{A}^n$ the bilinear map $k[X] \times k[Y] \to k[X \times Y]$, $(f, g) \mapsto f * g$ induces an isomorphism $\mu : k[X] \otimes k[Y] \to k[X \times Y]$ of k-algebras (so that in particular $k[X] \otimes k[Y]$ is reduced).

If X and Y are irreducible, then so is $X \times Y$, or equivalently, if $k[X]$ and $k[Y]$ are domains, then so is $k[X] \otimes k[Y]$.

PROOF. Since the obvious map

$$k[x_1, \ldots, x_m] \otimes k[y_1, \ldots, y_n] \to k[x_1, \ldots, x_m, y_1, \ldots, y_n]$$

is an isomorphism, it follows that μ is onto. In order to prove that μ is injective, let us first observe that every $\phi \in k[X] \otimes k[Y]$ can be written $\phi = \sum_{i=1}^N f_i \otimes g_i$ such that $g_1, \ldots, g_N$ are k-linearly independent. Given $p \in X$, then the restriction of $\mu(\phi) = \sum_{i=1}^N f_i * g_i$ to $\{p\} \times Y \cong Y$ is the regular function $\phi_p := \sum_{i=1}^N f_i(p) g_i \in k[Y]$. Since the g_i's are linearly independent, we have $\phi_p = 0$ if and only if $f_i(p) = 0$ for all i. In particular, the subset $X(\phi) \subseteq X$ of $p \in X$ for which $\phi_p = 0$, is equal to $\cap_{i=1}^N Z(f_i)$ and hence closed.

If $\mu(\phi) = 0$, then $\phi_p = 0$ for all $p \in X$ and hence $f_i = 0$ for all i. So $\phi = 0$. This proves that μ is injective.

Suppose now X and Y irreducible. We prove that $k[X] \otimes k[Y]$ is a domain so that $X \times Y$ is irreducible. Let $\phi, \psi \in k[X] \otimes k[Y]$ be such that $\phi\psi = 0$. Since the restriction of $\phi\psi = 0$ to $\{p\} \times Y \cong Y$ is $\phi_p \psi_p$ and $k[Y]$ is a domain, it follows that $\phi_p = 0$ or $\psi_p = 0$. So X is the union of its closed subsets $X(\phi)$ and $X(\psi)$. Since X is irreducible we have $X = X(\phi)$ or $X = X(\psi)$. This means that $\phi = 0$ or $\psi = 0$. □

EXERCISE 25. Let A and B be finitely generated k-algebras. Prove that $A \otimes B$ is a finitely generated k-algebra. Define a natural map $\mathrm{Spm}(A \otimes B) \to \mathrm{Spm}(A) \times \mathrm{Spm}(B)$ and show that this is a bijection (hint: do *not* use Proposition 6.1).

EXERCISE 26. Let X and Y be closed subsets of affine spaces. Prove that each irreducible component of $X \times Y$ is the product of an irreducible component of X and one of Y.

It is clear that the projections $\pi_X : X \times Y \to X$ and $\pi_Y : X \times Y \to Y$ are regular. We have observed that the space underlying $X \times Y$ is usually not the topological product of its factors. Still it is the 'right' product in the sense of category theory: it has the following universal property, which almost seems too obvious to mention: if Z is a closed subset of some affine space, then any pair of regular maps $f : Z \to X$, $g : Z \to Y$ defines a regular map $Z \to X \times Y$ characterized by the property that its composite with π_X resp. π_Y yields f resp. g (this is of course (f, g)).

7. Function fields and rational maps

In this section we interpret the total fraction ring of an algebra of regular functions.

Let X be an affine variety. Recall that an element $\phi \in \mathrm{Frac}(k[X])$ is by definition represented by a fraction f/g with $f, g \in k[X]$ and g not a zero divisor in $k[X]$.

When X is irreducible, $k[X]$ is a domain and hence $\mathrm{Frac}(k[X])$ is a field, called the *function field of* X, and will be denoted $k(X)$. We need the following lemma.

Lemma 7.1. For any $g \in k[X]$ the following are equivalent:

(i) X_g is dense in X.
(ii) $Z(g)$ does not contain an irreducible component of X,
(iii) g is not a zero divisor.

Moreover, every open-dense subset of X contains a principal open-dense subset defined by a nonzero divisor.

If $C_1, \ldots, C_r$ are the distinct irreducible components of X, then for every i we have a well-defined restriction map $R_i : \mathrm{Frac}(k[X]) \to k(C_i)$ of k-algebras and taken together these define a k-algebra isomorphism

$$R = (R_i)_{i=1}^r : \mathrm{Frac}(k[X]) \to \oplus_{i=1}^r k(C_i).$$

PROOF. Let $C_1, \ldots, C_r$ be as in the lemma, i.e., the distinct irreducible components of X. Since for every $i = 1, \ldots, r$, $\cup_{j \neq i} C_j$ is a proper closed subset of X, there exists a nonzero $h_i \in k[X]$ which vanishes on $\cup_{j \neq i} C_j$ (so $h_i | C_i \neq 0$).

$(i) \Rightarrow (ii)$ Assume X_g is dense in X. Then for every i, $X_g \cap C_i$ is a nonempty open subset of C_i and so $Z(g)$ cannot contain C_i.

$(ii) \Rightarrow (iii)$ Assume that $Z(g)$ does not contain an irreducible component of X. Suppose there exists a nonzero $g' \in k[X]$ with $gg' = 0$. So if C_i is such that $g'|C_i$ is nonzero, then $C_i \subseteq Z(g)$ and we get a contradiction.

$(iii) \Rightarrow (i)$ Assume g is not a zero divisor. Then for every i, $gh_i \neq 0$ and so $Z(g)$ cannot contain C_i. Hence $X_g \cap C_i$ is nonempty. Since C_i is irreducible, it follows that the closure of X_g must contain C_i. So X_g is dense in X.

To prove the subsequent assertion, let $U \subseteq X$ be open-dense. Then for every i, $C_i \cap U$ is a nonempty open subset of C_i and so contains a nonempty principal open subset X_{g_i}. Then $g_i | C_j$ is nonzero if and only $j = i$ and hence $g := \sum_{j=1}^r g_j$ has the property that $g|C_i = g_i|C_i$. So $C_i \cap X_g = C_i \cap X_{g_i}$ is nonempty and open (hence dense) in C_i and contained in U $(i = 1, \ldots, r)$. It follows that X_g is open-dense in X and contained in U.

As to the last assertion, first note that R_i is well-defined: if $f/g \in \mathrm{Frac}(k[X])$, then g is a nonzero divisor and is therefore not identically zero on C_i. If $R_i(f/g) = 0$ for all i, then $f|C_i = 0$ for all i and hence $f = 0$. So R is injective. To see that R is onto, let for $i = 1, \ldots, r$, $\phi_i/\psi_i \in k(C_i)$, with $\phi_i, \psi_i \in k[C_i]$ and $\psi_i \neq 0$. Let $f_i, g_i \in k[X]$ map to $\phi_i, \psi_i \in k[C_i]$ and put $f = \sum_i h_i f_i$, $g := \sum_i h_i g_i$. Then $f|C_i = h_i|C_i.\phi_i$ and $g|C_i = h_i|C_i.\psi_i$ and the latter is nonzero. By $(ii) \Leftrightarrow (iii)$, g is then not a zero divisor of $k[X]$. So $f/g \in \mathrm{Frac}(k[X])$ and we have $R_i(f/g) = \phi_i/\psi_i$. □

Corollary 7.2. If ϕ is a regular function defined on an open-dense subset U of X, then ϕ determines an element of $\mathrm{Frac}(k[X])$. All of $\mathrm{Frac}(k[X])$ is thus obtained and two pairs (U, ϕ) and (U', ϕ') determine the same element of $\mathrm{Frac}(k[X])$ if and only if $\phi|U \cap U' = \phi'|U \cap U'$.

PROOF. Suppose ϕ is a regular function on an open-dense subset U of X. According to Lemma 7.1, $U \supseteq X_g$ for some nonzero divisor g and hence $\phi|X_g$ is by Proposition 5.2 given by an element of $k[X][1/g]$ and hence determines an element of $\mathrm{Frac}(k[X])$. We check that this element of $\mathrm{Frac}(k[X])$ is unique and prove at the same time that any element of $\mathrm{Frac}(k[X])$ is thus obtained. If $f'/g' \in \mathrm{Frac}(k[X])$, then $X_{g'}$ is open-dense in X by Lemma 7.1 and so f'/g' defines a regular function

on $X_{g'}$. If this regular function is zero on an open-dense subset of $X_{g'}$, then f' is zero on this subset and so $f' = 0$. But then $f'/g' = 0$. □

There is in general no best way to represent a given element of $\mathrm{Frac}(k[X])$ as a fraction (as there is in a UFD), and so we must be content with the corollary above. Note that it is essentially equivalent to the assertion that

$$\mathrm{Frac}(k[X]) = \varinjlim_{g \text{ nonzero divisor in } k[X]} k[X][1/g] = \varinjlim_{U \text{ open-dense in } X} \mathcal{O}_X(U).$$

We call an element of $\mathrm{Frac}(k[X])$ a *rational function* on X. This is the algebro-geometric analogue of a meromorphic function in complex function theory.

We shall now give a geometric interpretation of finitely generated field extensions of k and the k-linear field homomorphisms between them.

DEFINITION 7.3. Let X and Y be affine varieties. A *rational map* from X to Y is given by a pair (U, F), where U is an open-dense subset of X and $F : U \to Y$ is a morphism, with the understanding that a pair (U', F') defines the same rational map if F and F' coincide on an open-dense subset of $U \cap U'$ (then they coincide on all of $U \cap U'$ by continuity, but by formulating it in this way we easier check that we are dealing with an equivalence relation). We denote such a rational map as $f : X \dashrightarrow Y$. We say that the rational map is *dominant* if for a representative pair (U, F), $F(U)$ is dense in Y. (This is then also so for any other representative pair. Why?)

So a rational map $f : X \dashrightarrow \mathbb{A}^1$ is the same thing as a rational function on X.

Observe that for a morphism of affine varieties $f : X \to Y$, $g \in \ker(f^*)$ is equivalent to $f(X) \subseteq Z_g$ and hence equivalent to the closure $f(X)$ being contained in Z_g. It follows that *f is dominant if and only if f^* is injective*.

Proposition 7.4. Every finitely generated field extension of k is k-isomorphic to the function field of an irreducible affine variety.

Let X and Y be irreducible affine varieties. Then a dominant rational map $f : X \dashrightarrow Y$ determines a k-linear field embedding $f^* : k(Y) \hookrightarrow k(X)$.

Conversely, every k-linear field embedding $k(Y) \to k(X)$ is induced by a unique dominant rational map $X \dashrightarrow Y$.

PROOF. Let K/k be a finitely generated field extension of k. This means that there exist $a_1, \ldots, a_n \in K$ such that every element of K can be written as a fraction of polynomials in $a_1, \ldots, a_n$ with coefficients in k. So the k-subalgebra of K generated by $a_1, \ldots, a_n$ is a domain $A \subseteq K$ (since K is a field) that has K as its field of fractions. Since A is the coordinate ring of a closed irreducible subset $X \subseteq \mathbb{A}^n$ (defined by the kernel of the obvious ring homomorphism $k[x_1, \ldots, x_n] \to A$), it follows that K can be identified with $k(X)$.

Suppose we are given a principal open-dense subset $U \subseteq X$ and a morphism $F : U \to Y$ with $F(U)$ dense in Y. As we observed above, then $F^* : k[Y] \to k[U]$ is injective. Its composite with $k[U] \hookrightarrow k(U) = k(X)$ is then an injective homomorphism from a domain to a field and therefore extends to a field embedding $k(Y) \hookrightarrow k(X)$.

It remains to show that every k-linear field homomorphism $\Phi : k(Y) \to k(X)$ is so obtained. For this, choose generators $b_1, \ldots, b_m$ of $k[Y]$. Then $\Phi(b_1), \ldots, \Phi(b_m)$ are rational functions on X and so are regular on a principal nonempty subset $X_h \subseteq X$. Since $b_1, \ldots, b_m$ generate $k[Y]$ as a k-algebra, it follows that Φ maps $k[Y]$

to $k[X_h] = k[X][1/h] \subseteq k(X)$. This k-algebra homomorphism defines a morphism $F : X_h \to Y$ such that $F^* = \Phi | k[Y]$. Since Φ is injective, so is F^* and hence $F(X_h)$ is dense in Y. It is clear that Φ is the extension of F^* to the function fields.

As to the uniqueness: if (U', F') is another solution, then choose a nonempty principal subset $U'' \subseteq U \cap U'$ such that F and F' both define morphisms $U'' \to Y$. These must be equal since the associated k-algebra homomorphisms $k[Y] \to k[U'']$ are the same (namely the restriction of Φ). □

The following exercise explains the focus on irreducible varieties when considering rational maps.

EXERCISE 27. Let X and Y be an affine varieties with distinct irreducible components $X_1, \ldots, X_r$ resp. $Y_1, \ldots, Y_s$. Prove that to give a rational map $f : X \dashrightarrow Y$ is equivalent to giving a rational map $f_i : X_i \dashrightarrow Y_{j_i}$ for $i = 1, \ldots, r$. Show that f is dominant if and only if for each $j \in \{1, \ldots, s\}$, there exists an $i_j \in \{1, \ldots, r\}$ such that f maps X_{i_j} to Y_j as a dominant map.

EXERCISE 28. Let $f \in k[x_1, \ldots, x_{n+1}]$ be irreducible of positive degree. Its zero set $X \subseteq \mathbb{A}^{n+1}$ is then closed and irreducible. Assume that the degree d of f in x_{n+1} is positive.

(a) Prove that the projection $\pi : X \to \mathbb{A}^n$ induces an injective k-algebra homomorphism $\pi^* : k[x_1, \ldots, x_n] \to k[X] = k[x_1, \ldots, x_n]/(f)$.
(b) Prove that π is dominant and that the resulting field homomorphism $k(x_1, \ldots, x_n) \to k(X)$ is a finite extension of degree d.

Corollary 7.5. Two dominant maps $f : X \dashrightarrow Y$ and $g : Y \dashrightarrow Z$ between irreducible affine varieties can be composed to yield a dominant map $gf : X \dashrightarrow Z$ so that we have a category with the irreducible affine varieties as objects and the rational dominant maps as morphisms. Assigning to an irreducible affine variety its function field makes this category anti-equivalent to the category of finitely generated field extensions of the base field k.

PROOF. The dominant maps yield k-linear field extensions $f^* : k(Y) \hookrightarrow k(X)$ and $g^* : k(Z) \hookrightarrow k(Y)$ and these can be composed to give a k-linear field extension $f^*g^* : k(Z) \hookrightarrow k(X)$. Proposition 7.4 says that this is induced by a unique rational map $X \dashrightarrow Z$. This we define to be gf. The rest of the corollary now follows. □

Proposition-definition 7.6. A rational map $f : X \dashrightarrow Y$ is an isomorphism in the above category (that is, induces a k-linear isomorphism of function fields) if and only if there exists a representative pair (U, F) of f such that F maps U *isomorphically* onto an open subset of Y. If these two equivalent conditions are satisfied, then f is called a *birational map*. If a birational map $X \dashrightarrow Y$ merely exists (in other words, if there exists a k-linear field isomorphism between $k(X)$ and $k(Y)$), then we say that X and Y are *birationally equivalent*.

PROOF. If f identifies a nonempty open subset of X with one of Y, then $f^* : k(Y) \to k(X)$ is clearly a k-algebra isomorphism.

Suppose now we have a k-linear isomorphism $k(Y) \cong k(X)$. Represent this isomorphism and its inverse by (U, F) and (V, G) respectively. Then $F^{-1}V$ is a nonempty open subset of U, it contains a nonempty principal open subset X_g. Now $X_g \xrightarrow{GF} X$ is morphism between affine varieties which induces the identity on $k(X) = k(X_g)$ and hence induces the inclusion $k[X] \subset k[X_g]$. This implies that

$GF|X_g$ is the inclusion $X_g \subseteq X$. Since X_g is dense in $F^{-1}V$, and GF is continuous, GF is the inclusion $F^{-1}V \subseteq X$. In particular, F maps $F^{-1}V$ injectively to $G^{-1}U$. For the same reason, G maps $G^{-1}U$ injectively to $F^{-1}V$. Both composites are the identity and so F defines an isomorphism $F : F^{-1}V \cong G^{-1}U$. □

EXERCISE 29. Prove that the curve in $\mathbb{A}^2$ defined by $x_1^2 + x_2^2 = 1$ is birationally equivalent to the affine line $\mathbb{A}^1$ when $\mathrm{char}(k) \neq 2$ (hint: take a look at Exercise 16). What happens when $\mathrm{char}(k) = 2$? Same questions for the quadrics in $\mathbb{A}^{n+1}$ defined by $x_1^2 + x_2^2 + \cdots + x_{n+1}^2 = 1$ and $x_1x_2 + x_3^2 + \cdots + x_{n+1}^2 = 1$.

In case $k(X)/k(Y)$ is a finite extension, one may wonder what the geometric meaning is of the degree d of that extension, perhaps hoping that this is just the number of elements of a general fiber of the associated rational map $X \dashrightarrow Y$. We will see that this is often true (namely when the characteristic of k is zero, or more generally, when this characteristic does not divide d), but not always, witness the following example.

EXAMPLE 7.7. Suppose k has characteristic $p > 0$. We take $X = \mathbb{A}^1 = Y$ and take for f the Frobenius morphism: $\Phi_p : a \in \mathbb{A}^1 \mapsto a^p \in \mathbb{A}^1$. We have seen in Example 4.10 that Φ_p is homeomorphism, but that $\Phi_p^* : k[Y] = k[y] \to k[x] = k[X]$ is given by $y \mapsto x^p$ and so induces the field extension $k(y) = k(x^p) \subset k(x)$, which is of degree p. From the perspective of Y, we have enlarged its algebra of regular functions by introducing a formal pth root of its coordinate y (which yields another copy of $\mathbb{A}^1$, namely X). From the perspective of X, $k[Y]$ is just the subalgebra $k[x^p] \subset k[x]$.

This is in fact the basic example of a *purely inseparable* field extension, i.e., an algebraic field extension L/K with the property that every element of L has a minimal polynomial in $K[T]$ that has precisely one root in L. If $L \neq K$, then the characteristic p of K must be positive and such a polynomial of the form $T^{p^r} - c$, with $c \in K$ and $r > 0$. For the polynomial to be minimal, c cannot be a p-th power of an element of K (for if $c = a^p$, then $T^{p^r} - c = (T^{p^{r-1}} - a)^p$); in particular, the Frobenius map $F_p : a \in K \mapsto a^p \in K$ cannot be surjective. Purely inseparable extensions have trivial Galois group as there is only one root to move around and hence are not detected by Galois theory.

EXERCISE 30. Let $f : X \dashrightarrow Y$ be a dominant rational map of irreducible affine varieties which induces a purely inseparable field extension $k(X)/k(Y)$. Prove that there is an open-dense subset $V \subseteq Y$ such that f defines a homeomorphism $f^{-1}V \to V$. (Hint: show first that it suffices to treat the case when $k(X)$ is obtained from $k(Y)$ by adjoining the pth root of an element $f \in k(Y)$. Then observe that if f is regular on the affine open-dense $V \subseteq Y$, then Y contains as an open dense subset an open sense subset of the locus of $(x,t) \in V \times \mathbb{A}^1$ satisfying $t^p = f$.)

Much of the algebraic geometry in the 19th century and early 20th century was of a birational nature: birationally equivalent varieties were regarded as not really different. This sounds rather drastic, but it turns out that many interesting properties of varieties are an invariant of their birational equivalence class.

Here is an observation which not only illustrates how affine varieties over algebraically nonclosed fields can arise when dealing with affine k-varieties, but one that also suggests that we ought to enlarge the maximal ideal spectrum. Let

$f : X \to Y$ be a dominant morphism of irreducible affine varieties. This implies that $f^* : k[Y] \to k[X]$ is injective and that $f(X)$ contains an open-dense subset of Y. Then we may ask whether there exists something like a general fiber: is there an open-dense subset $V \subseteq Y$ such that the fibers $f^{-1}(y)$, $y \in V$ all "look the same"? The question is too vague for a clear answer and for most naive ways of making this precise, the answer will be no. For instance, we could simply refuse to specify one such V by allowing it to be arbitrarily small, but if we then want to implement this idea by taking the (projective) limit $\varprojlim_{V \text{ open-dense in } Y} f^{-1}V$, then we end up with the empty set unless Y is a singleton. However, its algebraic counterpart, which amounts to making all the nonzero elements of $k[Y]$ in $k[X]$ invertible, is nontrivial. To be precise, we have an isomorphism

$$\varinjlim_{V \text{ open-dense in } Y} k[f^{-1}V] = \varinjlim_{g \neq 0} k[X][1/f^*g] \cong k(Y) \otimes_{k[Y]} k[X]$$

The latter is in fact a reduced finitely generated $k(Y)$-algebra and this is a hint that an adequate geometric description requires that we include more points. First of all, we would like to regard the maximal ideal spectrum of $k(Y) \otimes_{k[Y]} k[X]$ as an affine variety over the (algebraically nonclosed) field $k(Y)$ so that every regular function on X which comes from Y is now treated as a scalar (and will be invertible when nonzero) ([8]). And secondly, in order to give this a geometric content, we would like that every irreducible variety Z defines a point η_Z (its *generic point*) with 'residue field' $k(Z)$, which for a singleton must give us back its unique element with the field k. For we then can extend f to the points defined by closed irreducible subsets $Z \subseteq X$ by putting $f(\eta_Z) := \eta_{\overline{f(Z)}}$. Then as a set, the generic fiber of f is the fiber of this extension over η_Y, i.e., the set of η_Z for which $f|Z : Z \to Y$ is dominant. Such considerations directly lead to the notion of a scheme that we shall discuss later.

8. Finite morphisms

In this section A is a ring and B is a A-algebra and let p be a prime number. In other words, we are given a ring homomorphism $A \to B$ (that is sometimes denoted by B/A). We do not assume that $A \to B$ is injective. Nevertheless we often make no notational distinction between an element of A and its image in B. When $A \to B$ is injective (so that we may regard A as subring of B), we say that B is an *extension* of A. We say that B is *finite* resp. *finite prime-to-p* over A if B is a finitely generated A-module (resp. when we can do this with a number of generators that is not divisible p). So when we say that B/A is a finite extension, we mean that $A \to B$ is injective and B is a finitely generated A-module.

Proposition-definition 8.1. We say that $b \in B$ is *integrally dependent* on A if one the following equivalent properties is satisfied.

(i) b is a root of a monic polynomial $x^n + a_1 x^{n-1} + \cdots + a_n \in A[x]$,
(ii) $A[b]$ is finitely generated as an A-module,
(iii) b is contained in a A-subalgebra $C \subseteq B$ which is finitely generated as an A-module.

[8] Our notion of affine variety required that we work over an algebraically closed field. This is of course arranged by choosing an algebraic closure L of $k(Y)$. The maximal ideal spectrum of $L \otimes_{k[Y]} k[X]$ is then an affine L-variety, and yields a notion of a *general fiber* that is even closer to our geometric intuition.

The three properties are still equivalent when 'taken prime to p': in (i) take n prime to p and in (ii) resp. (iii), assume that $A[b]$ resp. C is finite prime-to-p over A. When these equivalent conditions are satisfied, we say that b is *prime-to-p integral over A*.

PROOF. $(i) \Rightarrow (ii)$. If b is a root of $x^n + a_1 x^{n-1} + \cdots + a_n \in A[x]$, then clearly $A[b]$ is generated as a A-module by $1, b, b^2, \ldots, b^{n-1}$.

$(ii) \Rightarrow (iii)$ is obvious.

$(iii) \Rightarrow (i)$. Suppose that C is as in (iii). Choose an epimorphism $\pi : A^n \to C$ of A-modules and denote the standard basis of A^n by $(e_1, \ldots, e_n)$. We may (and will) assume that $\pi(e_1) = 1_B$. By assumption, $b\pi(e_i) = \sum_{j=1}^n a_{ij}\pi(e_j)$ for certain $a_{ij} \in A$. We regard the $n \times n$-matrix $\sigma := (b\delta_{ij} - a_{ij})_{i,j}$ with entries in $A[b]$ as an $A[b]$-endomorphism of $A[b]^n$. Note that $\det(\sigma)$ is a monic polynomial in b of degree n with coefficients in A. So it suffices to show that $\det(\sigma) = 0$.

Since $b \in C$, π extends to an epimorphism $\tilde{\pi} : A[b]^n \to C$ of $A[b]$-modules. Then $\tilde{\pi}\sigma(e_i) = 0$ for all i, in other words, $\tilde{\pi}\sigma = 0$. Now Cramer's rule can be understood as stating that if σ' is the matrix of cofactors of σ, then $\sigma\sigma' = \det(\sigma)1_n$. We thus find that in B,

$$\det(\sigma) = \det(\sigma)\pi(e_1) = \tilde{\pi}(\det(\sigma)e_1) = \tilde{\pi}(\sigma\sigma'(e_1)) = (\tilde{\pi}\sigma)(\sigma'(e_1)) = 0.$$

The proof of the prime-to-p version is the same. □

Corollary-definition 8.2. The elements of B that are integrally dependent on A make up an A-subalgebra of B; we call this subalgebra the *integral closure* of A in B (and denote it by $\overline{A}^B$).

The elements of B that are prime-to-p integral over A also make up an A-subalgebra of B; we call this the *prime-to-p integral closure* of A in B.

PROOF. It is enough to prove that if $b, b' \in B$ are integrally dependent over A, then so is every element of $A[b, b']$. Or what amounts to the same: if $A[b]$ and $A[b']$ are finitely generated A-modules, then so is $A[b, b']$. This is clear: if $\{b^k\}_{k=0}^{n-1}$ generates $A[b]$ and $\{b'^k\}_{k=0}^{n'-1}$ generates $A[b']$, then the nn' elements $\{b^k b'^{k'}\}_{k=0,k'=0}^{n,n'}$ generate $A[b, b']$.

The proof of the prime-to-p version is the same. □

DEFINITION 8.3. We say that B is *integral* (resp. *prime-to-p integral*) over A if every element of B is integral (resp. prime-to-p integral) over A.

So if in addition the given homomorphism $A \to B$ is injective, then we say that B is an *integral extension* (resp. a *prime-to-p integral extension*) of A.

The characterization (iii) of Proposition 8.1 shows that when B is finite over A (resp. a finite extension of A), then B is integral over A (resp. an integral extension of A). An important class of example appears in algebraic number theory: if L is a finite field extension of $\mathbb{Q}$ (also called an *algebraic number field*), then the integral closure of $\mathbb{Z} \subset \mathbb{Q}$ defines a subring of L, called the *ring of integers* of L. This ring is often denoted by $\mathcal{O}_L$.

EXERCISE 31. Prove that 'being integral over' is transitive: if B is an A-algebra integral over A, then any B-algebra that is integral over B is as an A-algebra integral over A.

Proposition 8.4. Let $A \subseteq B$ be an integral extension and suppose B is a domain. Then $\mathrm{Frac}(A)B = \mathrm{Frac}(B)$ and (hence) $\mathrm{Frac}(B)$ is an algebraic field extension of $\mathrm{Frac}(A)$. This extension is finite whenever B is a finite extension of A.

If $\mathrm{Frac}(B)$ is of characteristic p and $A \subseteq B$ is a prime-to-p integral, then $\mathrm{Frac}(B)$ is a separable algebraic extension $\mathrm{Frac}(A)$.

PROOF. To prove that $\mathrm{Frac}(B) = \mathrm{Frac}(A)B$, it is enough to show that $1/b \in \mathrm{Frac}(A)B$ for any $b \in B \smallsetminus \{0\}$. By assumption, such a b satisfies an equation $b^n + a_1 b^{n-1} + \cdots + b a_{n-1} + a_n = 0$ with $a_i \in A$ and $a_n \neq 0$. So $1/b = -1/a_n.(b^{n-1} + a_1 b^{n-2} + \cdots + a_{n-1}) \in \mathrm{Frac}(A)B$.

Let now $c \in \mathrm{Frac}(B)$ and write $c = a/b$ with $a \in A$ and $b \in B \smallsetminus \{0\}$ as above. Then the identity $c^n + (a a_{n-1}/a_n)c^{n-1} + \cdots + (a^{n-1}a_1/a_n)c + a^n/a_n = 0$ shows that c is algebraic over $\mathrm{Frac}(A)$. This proves that $\mathrm{Frac}(B)/\mathrm{Frac}(A)$ is algebraic. As any finite set of A-module generators of B is also a (finite) set of $\mathrm{Frac}(A)$-vector space generators of $\mathrm{Frac}(A)B = \mathrm{Frac}(B)$, the last assertion also follows.

The last statement is clear. □

There are two simple ways of producing new integral extensions out of a given one, namely reduction and localization (with Proposition 8.4 being in fact a special case of the latter). Suppose $A \subseteq B$ is integral. Then for every ideal $J \subseteq B$, $J \cap A$ is (clearly) an ideal of A and $A/J \cap A \subseteq B/J$ is an integral extension. And if $S \subseteq A$ is a multiplicative subset, then the induced ring homomorphism $S^{-1}A \to S^{-1}B$ is integral. Both appear in the proof of the 'Going up theorem' below. For this we will also need:

Lemma 8.5 (Nakayama's Lemma)**.** Let R be a local ring with maximal ideal $\mathfrak{m}$ and M a finitely generated R-module. Then a finite subset $S \subseteq M$ generates M as a R-module if (and only if) the image of S in $M/\mathfrak{m}M$ generates the latter as a $R/\mathfrak{m}$-vector space. In particular (take $S = \emptyset$), $\mathfrak{m}M = M$ implies $M = 0$.

PROOF. The special case $S = \emptyset$ is in fact the general case, for we reduce to it by passing to M/RS, as our assumptions then say that then $M = \mathfrak{m}M$ and we must show that $M = 0$. Let $\pi : R^n \to M$ be an epimorphism of R-modules and denote the standard basis of R^n by $(e_1, \ldots, e_n)$. By assumption there exist $r_{ij} \in \mathfrak{m}$ such that $\pi(e_i) = \sum_{j=1}^{s} r_{ij}\pi(e_j)$. So if $\sigma := (\delta_{ij} - r_{ij})_{i,j} \in \mathrm{End}_R(R^n)$, then $\pi\sigma = 0$. Notice that $\det(\sigma) \in 1 + \mathfrak{m}$. Since $1 + \mathfrak{m}$ consists of invertible elements, Cramer's rule shows that σ is invertible. So $\pi = \pi(\sigma\sigma^{-1}) = (\pi\sigma)\sigma^{-1} = 0$ and hence $M = 0$. □

Proposition 8.6 (Going up)**.** Let $A \subseteq B$ be an integral extension and let $\mathfrak{p} \subseteq A$ be a prime ideal of A. Then the *going up* property holds: $\mathfrak{p}$ is of the form $\mathfrak{q} \cap A$, where $\mathfrak{q}$ is a prime ideal of B. If also is given is a prime ideal $\mathfrak{q}'$ of B with the property that $\mathfrak{p} \supseteq \mathfrak{q}' \cap A$, then we can take $\mathfrak{q} \supseteq \mathfrak{q}'$. Moreover the *incomparability* property holds: two distinct prime ideals of B having the same intersection with A cannot obey an inclusion relation.

PROOF. The localization $A \to A_{\mathfrak{p}}$ yields a local ring with maximal ideal $\mathfrak{p}A_{\mathfrak{p}}$ and the prime ideals of $A_{\mathfrak{p}}$ correspond (by taking the preimage in A) to the prime ideals of A that contain $\mathfrak{p}$. The localization $A_{\mathfrak{p}}B$ (as a A-module) is, by the observation above, an integral extension of $A_{\mathfrak{p}}$. If we find a prime ideal $\tilde{\mathfrak{q}}$ of $A_{\mathfrak{p}}B$ with $\tilde{\mathfrak{q}} \cap A_{\mathfrak{p}} = \mathfrak{p}A_{\mathfrak{p}}$, then the preimage $\mathfrak{q}$ of $\tilde{\mathfrak{q}}$ in B is a prime ideal of B with the property that $\mathfrak{q} \cap A$ is the preimage of $\mathfrak{p}A_{\mathfrak{p}}$ in A and so this is just $\mathfrak{p}$. Hence for the first assertion there is no loss in generality in assuming that A is a local ring and $\mathfrak{p}$ is its unique maximal ideal $\mathfrak{m}_A$.

We claim that $\mathfrak{m}_A B \neq B$. Suppose this is not so: $\mathfrak{m}_A B = B$. Then write $1 \in B$ as an $\mathfrak{m}_A$-linear combination of a finite set elements of B. Denote by B' the A-subalgebra of B generated by this finite set. Since B is an integral extension of A, B' is finite over A. Since $1 \in \mathfrak{m}_A B'$, we have $B' = \mathfrak{m}_A B'$, and it then follows from Nakayama's Lemma 8.5 that $B' = 0$. Hence $1 = 0$ so that A is the zero ring. This contradicts our assumption that A has a maximal ideal.

Since $\mathfrak{m}_A B \neq B$, we can take for $\mathfrak{q}$ any maximal ideal of B which contains the ideal $\mathfrak{m}_A B$: then $\mathfrak{q} B \cap A$ is a maximal ideal of A, hence equals $\mathfrak{m}_A$.

For the refinement we can, simply by passing to $A/(\mathfrak{q}' \cap A) \subseteq B/\mathfrak{q}'$ (which is still an integral extension by the observation above), assume that $\mathfrak{q}' = 0$. This reduces the refinement to the case already treated.

For the incomparability property we must show that if $\mathfrak{q}' \subseteq \mathfrak{q}$ and $\mathfrak{q}' \cap A = \mathfrak{q} \cap A = \mathfrak{p}$, then $\mathfrak{q}' = \mathfrak{q}$. By passing to $A/(\mathfrak{q}' \cap A) \subseteq B/\mathfrak{q}'$ we reduce to the case when B is a domain and $\mathfrak{q}' = 0$ and we must then show that $\mathfrak{q} = 0$. To see this, suppose that $b \in \mathfrak{q} \setminus \{0\}$. Then $b^n + a_1 b^{n-1} + \cdots + a_{n-1} b + a_n = 0$ for certain $a_i \in A$, where we can assume that $a_n \neq 0$ (otherwise divide by b; remember that B is a domain). Then $a_n \in Bb \cap A \subseteq \mathfrak{p}$ and so $a_n = 0$, and we arrive at a contradiction. □

REMARK 8.7. Let us agree to call a *prime chain* (of length n) in a ring R a strictly ascending sequence $\mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \cdots \subsetneq \mathfrak{p}_n$ of prime ideals in R. The going up property may then be restated as saying that for an integral extension $A \subseteq B$, any prime chain in A is the intersection with A of a prime chain in B, where we even may prescribe the first member of the latter in advance([9]). The incomparability property says that the intersection a prime chain in B with A is a prime chain in A.

DEFINITION 8.8. We say that a morphism of affine varieties $f : X \to Y$ is *finite* if the k-algebra homomorphism $f^* : k[Y] \to k[X]$ is finite, i.e., makes $k[X]$ a finitely generated $k[Y]$-module (so $k[X]$ is then integral over $k[Y]$).

So $f^* : k[Y] \to k[X]$ is a finite extension if and only if f is finite and dominant.

EXERCISE 32. Prove that if Y is an affine variety, then the disjoint union of its irreducible components is finite over Y.

Propositions 8.4 and 8.6 give in the algebro-geometric setting:

Corollary 8.9. Let $f : Y \to X$ be a finite, dominant morphism of affine varieties. Then for every closed irreducible subset $P \subseteq X$, the collection of closed irreducible subsets $Q \subseteq Y$ with $f(Q) = P$ is nonempty (in particular, f is onto) with no two members satisfying an inclusion relation. If in addition Y is irreducible, then so is X and $f^* : k(X) \to k(Y)$ is a finite algebraic extension of fields.

PROOF. Since $\mathfrak{p} := I(P)$ is a prime ideal, Proposition 8.6 implies that $\mathfrak{p} = (f^*)^{-1}\mathfrak{q}$ for some prime ideal $\mathfrak{q}$. Then $Q := Z(\mathfrak{q})$ is irreducible and f^* induces an injective morphism $k[P] = k[X]/\mathfrak{p} \to k[Y]/\mathfrak{q} = k[Q]$. So $f(Q) \subseteq P$ and is in fact is dense in P. If we apply this to the morphism $Q \xrightarrow{f} P$ (instead of f) and take for P a singleton $\{p\} \subseteq P$, we find that $f^{-1}(p)$ is nonempty. Hence $f(P) = Q$.

Conversely, if $Q \subseteq Y$ is irreducible with $f(Q) = P$, then the prime ideal $\mathfrak{q}$ defining Q has the property that $(f^*)^{-1}\mathfrak{q}$ defines P. The incomparability property then implies that no two such Q satisfy an inclusion relation.

The last assertion follows from Proposition 8.4. □

[9] This is indeed what 'going up' refers to.

EXERCISE 33. Let $f : Y \to X$ be a finite morphism of affine varieties. Prove that f is closed and that every fiber $f^{-1}(x)$ is finite (possibly empty). Prove that if $W \subseteq X$ is closed or a principal open subset of X, then the restriction $f : f^{-1}W \to W$ is also a finite morphism.

EXAMPLE 8.10. This simple example shows that in the situation of Corollary 8.9, not every irreducible component of $f^{-1}P$ need to dominate P. Consider the morphism $f : \mathbb{A}^2 \to \mathbb{A}^4$ defined by $f(x, y) = (x(x-1), x^2(x-1), xy, y)$. Then f is finite and identifies the two basis vectors $e_1 = (1, 0)$ and $e_2 = (0, 1)$, but nothing else. Since f is a finite morphism, its image $X \subset \mathbb{A}^4$ is closed (by the above exercise). We regard f now as mapping to X so that it is then a dominant morphism between irreducible varieties. Let $P \subset X$ to be the image of the y-axis (this the just the last coordinate axis of $\mathbb{A}^4$). Then $f^{-1}P$ is the union of the y-axis and the singleton $\{e_1\}$.

Elementary dominant morphisms. We describe here some dominant morphisms $f : X \to Y$, where X and Y are irreducible affine varieties that are of an elementary type in the sense that any dominant morphism is a composition of such.

Elementary immersions. By this we mean f maps X isomorphically onto a principal open subset of Y: there exists a nonzero $g \in k[Y]$ such that f maps X isomorphically onto Y_g. Hence f^* factors as $k[Y] \hookrightarrow k[Y][1/g] \cong k[X]$. In particular, f is birational.

Elementary projections. By this we mean that there exists an isomorphism $X \cong Y \times \mathbb{A}^1$ having f as its first component. This identifies $k[X]$ with $k[Y \times \mathbb{A}^1] = k[Y] \otimes k[t] = k[Y][t]$ as a $k[X]$-algebra and hence makes $k(X)/k(Y)$ a simple transcendental extension.

Elementary finite separable morphisms. We here assume that there exists an isomorphism of X onto a hypersurface $H \subset Y \times \mathbb{A}^1$ which is given by monic polynomial $g \in k[X][t]$, $g(x, t) = t^d + a_1(x)t^{d-1} + \cdots a_d(x)$ which is both irreducible and separable and such that the first component is f. Then f is a finite morphism, and $k(X)/k(Y)$ is a separable extension of degree d. If $\alpha_1, \ldots \alpha_d$ are the distinct roots of g in some extension of $k(Y)$, then the discriminant $D(g) := \prod_{i \neq j}(\alpha_i - \alpha_j)$ (there are $d(d-1)$ factors) is not identically zero. It is in fact a polynomial in the coefficients $a_1, \ldots, a_d$, hence lies in $k[Y]$. So $Y_{D(g)}$ is a principal open-dense subset of Y with the property that every fiber over a point of $Y_{D(g)}$ has exactly d points.

Elementary finite inseparable morphisms. We here assume that the characteristic p of k is positive and that there exists an isomorphism of X onto a hypersurface $H \subset Y \times \mathbb{A}^1$ which is given by a polynomial $g \in k[Y][t]$ of the form $t^p - a$ with $a \in k[Y]$ not a pth power (this ensures that g is irreducible) and such that the first component is f. Then the projection $f : X \to Y$ is a homeomorphism, but the associated field extension is purely inseparable of degree p. Note that if we replace p by p^r, $r \geq 1$, then f is obtained as an r-fold iteration of such morphisms.

Theorem 8.11. Let $f : X \to Y$ be a dominant morphism of affine irreducible varieties. Then X contains a nonempty principal open subset U such that $f|U$ is open and can be factored into elementary dominant morphisms.

PROOF. We can always arrange that X is a closed subset of $\mathbb{A}^n \times Y$ and $f = \pi_Y|X$: first identify X with a closed subset of $\mathbb{A}^n$ and then identify X with the graph of f in $\mathbb{A}^n \times X$. We then proceed with induction on n. For $n = 0$, the assertion

is trivial. When $n > 0$, we factor f as $X \to \mathbb{A}^1 \times Y \to Y$ by retaining the last coordinate only. Let $X' \subseteq \mathbb{A}^1 \times Y$ denote the closure of the image of first morphism $X \to \mathbb{A}^1 \times Y$ and factor f as $X \xrightarrow{f'} X' \xrightarrow{f''} Y$. Our induction hypothesis produces a nonempty principal open $U' \subseteq X$ such that $f'|U'$ is open and can be factored into elementary dominant morphisms. If we can also find a principal open $U'' \subseteq X'$ such that the projection $U'' \to Y$ is open and can be factored into elementary dominant morphisms, then $U := U' \cap f'^{-1}U''$ is as desired. In other words, we may assume $n = 1$.

When $X = \mathbb{A}^1 \times Y$, we have an elementary simple projection. Otherwise $I(X) \subseteq k[\mathbb{A}^1] \otimes k[Y] = k[Y][t]$ contains a nonzero $g \in k[Y][t]$. Since X is irreducible, we may (and will) take g irreducible. Write $g = a_0 t^N + a_1 t^{N-1} + \cdots + a_N$ with $a_i \in k[Y]$ and a_0 nonzero. We may replace Y by a nonempty principal open subset V of Y and X by $f^{-1}V$. By passing to $V = Y_{a_0}$ and then dividing g by a_0, we may without loss of generality assume that g is monic. This implies that $Z(g)$ is finite over Y and contains X as a closed subset. If we apply Corollary 8.9 to the projection $Z(g) \to Y$, we find that the inclusion $X \subseteq Z(g)$ must be an equality. If g is a separable polynomial, then $X \to Y$ is an elementary finite separable morphism.

It remains to do the case when g is not separable. This can only happen when the characteristic p of k is positive. Let $K/k(Y)$ be an algebraic extension which contains all the roots of g, when g is regarded as an element of $k(Y)[t]$ so that $g = \prod_{i \in I}(t - \alpha_i)$ with $\alpha_i \in K$. If a root α_i occurs with a multiplicity q, then all roots with that same multiplicity determine a factor of g in $k(Y)$. Since g is irreducible, this means that all roots have multiplicity q. We will have $q = p^r$ for some $r \geq 1$ and $g(t) = h(u^q)$ with $h \in k[Y][u]$ a separable irreducible monic polynomial. This means that we have a factorization $f : Z(g) \to Z(h) \to Y$, where $Z(h) \to Y$ is an elementary finite separable morphism. Write a for the image of u in $k[Z(h)]$ and identify $Z(g)$ with the closed subset of $Z(h) \times \mathbb{A}^1$ defined by $t^q = a$. Then a cannot be a pth power in $k[Z(h)]$ (for $Z(g)$ is irreducible) and so the projection $Z(g) \to Z(h)$ is an r-fold iteration of elementary finite inseparable morphisms. □

An interesting application is given in the following exercise.

EXERCISE 34. Let $f : X \to Y$ be a dominant morphism of irreducible varieties for which $k(X)$ is a finite extension of $k(Y)$. Let d be the degree of the separable closure of $k(Y)$ in $k(X)$. Prove that there exists a principal nonempty open subset $V \subseteq Y$ such that $f^{-1}V \to V$ is a finite morphism, all of whose fibers have d points.

Here is another application. As the following simple example shows, the image of a morphism of affine varieties need not be locally closed.

EXAMPLE 8.12. Consider the morphism $f : \mathbb{A}^2 \to \mathbb{A}^2$, $(x_1, x_2) \mapsto (x_1, x_1 x_2)$. A point $(y_1, y_2) \in \mathbb{A}^2$ is of the form $(x_1, x_1 x_2)$ if and only if y_2 is a multiple of y_1. This is the case precisely when $y_1 \neq 0$ or when $y_1 = y_2 = 0$. So the image of f is the union of the open subset $y_1 \neq 0$ and the singleton $\{(0, 0)\}$. This is not a locally closed subset, but the union of two such. This turns out to represent the general situation and the following definition will help us to express this fact.

DEFINITION 8.13. A subset of a topological space X is called *constructible* if it can be written as the union of finitely many locally closed subsets of X.

An exercise in set theory shows that we then can always take this union to be disjoint. But the resulting decomposition need not be unique.

Proposition 8.14. Let $f : X \to Y$ be a morphism of affine varieties. Then f takes constructible subsets of X to constructible subsets of Y. In particular, $f(X)$ is constructible.

PROOF. A constructible subset is a finite union of subvarieties and so it is clearly enough to prove that the image of each of these is constructible. In other words, we only need to show that $f(X)$ is constructible.

It then suffices to show that X contains a proper closed subset $X_1 \subsetneq X$ such that $f(X \smallsetminus X_1)$ is locally closed in Y, for if $X_1 \neq \emptyset$, then the same argument applied to $f|X_1$ shows that there exists a proper closed subset $X_2 \subsetneq X_1$ such that $f(X_1 \smallsetminus X_2)$ is locally closed in Y and since X is a noetherian space, this process will terminate and so $f(X) = \cup_{i\geq 0} f(X_i \smallsetminus X_{i+1})$ (we put $X_0 := X$) is then written as a finite union of locally closed subsets.

For the existence of X_1, we may of course also assume that X is irreducible and upon replacing Y by the closure of $f(X)$, that Y is irreducible and f is dominant. If we then choose $U \subset X$ as in Theorem 8.11, then $X_1 := X \smallsetminus U$ has the desired property. □

9. Normalization

The following theorem has important geometric consequences. Recall that a field K is called *perfect* if it is either of characteristic zero or of positive characteristic $p > 0$ and every element of K has a pth root. For example, every finite field is perfect, but if k is of chacteristic $p > 0$, then $k(t)$ is not.

Theorem 9.1 (Noether normalization)**.** Given a field K, then every finitely generated K-algebra A is a finite (hence integral) extension of a polynomial K-algebra: there exist an integer $r \geq 0$ and an injection $K[x_1, \ldots, x_r] \hookrightarrow A$ of K-algebras such that A is finite over $K[x_1, \ldots, x_r]$.

In case K is perfect of characteristic $p > 0$ and the pth power map F_A in A is injective, then we can also arrange that A is finite prime-to-p over $K[x_1, \ldots, x_r]$.

REMARK 9.2. If in this situation A is a domain, then according to Proposition 8.4, $\mathrm{Frac}(A)$ will be a finite extension of $K(x_1, \ldots, x_r)$ and so r must be the transcendence degree of $\mathrm{Frac}(A)/K$. In particular, r is an invariant of A as a K-algebra.

The proof will be with induction and the following lemma provides the induction step.

Lemma 9.3. Let $\phi : K[x_1, \ldots, x_m] \to A$ be an epimorphism of K-algebras which is not an isomorphism. Then there exists a K-algebra automorphism σ of $K[x_1, \ldots, x_m]$ which fixes x_m and is such that $K[x_1, \ldots, x_{m-1}] \xrightarrow{\phi\sigma^{-1}} A$ is integral.

In case K is perfect of characteristic $p > 0$ and the pth power map F_A is injective, then we can also arrange that A is finite prime-to-p over A'.

PROOF. We define for any positive integer s a K-algebra automorphism σ_s of $K[x_1, \ldots, x_m]$ by $\sigma_s(x_i) := x_i + x_m^{s^{m-i}}$ when $i \leq m-1$ and $\sigma_s(x_m) = x_m$. This is indeed an automorphism with inverse given by $\sigma_s^{-1}(x_i) = x_i - x_m^{s^{m-i}}$ for $i < m-1$ and $\sigma_s(x_m) = x_m$. Note that if $I = (i_1, \ldots, i_m) \in \mathbb{Z}^m_{\geq 0}$, then

$$\sigma_s(x_1^{i_1} \cdots x_m^{i_m}) = (x_1 + x_m^{s^{m-1}})^{i_1} \cdots (x_{m-1} + x_m^s)^{i_{m-1}} x_m^{i_m}.$$

When viewed as an element of $K[x_1,\dots,x_{m-1}][x_m]$, this is a monic polynomial in x_m of degree $p_I(s) := i_1 s^{m-1} + i_2 s^{m-2} + \cdots + i_m$. Now give $\mathbb{Z}^m_{\geq 0}$ the lexicographic ordering. Then $I > J$ implies $p_I(s) > p_J(s)$ for s large enough. Choose a nonzero $f \in \ker(\phi)$. If $I \in \mathbb{Z}^m_{\geq 0}$ is the largest multi-exponent of a monomial which appears in f with nonzero coefficient, then for s large enough, $\sigma_s(f)$ is a nonzero constant times a monic polynomial in x_m of degree $p_I(s)$ with coefficients in $K[x_1,\dots,x_{m-1}]$. This ensures that the image of x_m in $K[x_1,\dots,x_m]/(\sigma_s(f))$ is integral relative to the homomorphism $\phi\sigma^{-1}|K[x_1,\dots,x_{m-1}]$. Hence $\sigma := \sigma_s$ is as desired.

Assume now that K is perfect of characteristic $p > 0$ and F_A is injective and take in the above argument $f \in \ker(\phi)$ of minimal degree (but nonzero). Since degree of integral dependence for x_m with respect to $\phi\sigma_s^{-1}|K[x_1,\dots,x_{m-1}]$ is equal to $p_I(s)$, it suffices to show that for infinitely many $s > 0$, $p_I(s)$ is prime to p. If this is not the case, i.e., when $p_I(s) \equiv 0 \pmod p$ for s sufficiently large, then a little exercise shows that the image if p_I in $\mathbb{F}_p[s]$ is zero: all the coefficients of the polynomial $p_I(s)$ must be divisible by p. This means that x^I is a pth power. So every monomial appearing in $f \in \ker(\phi)$ is a pth power. Since K is perfect, the coefficient of such a monomial is also a pth power. Hence $f = g^p$ for some $g \in K[x_1,\dots,x_{m-1}]$. So the image $\bar{g}$ in A satisfies $\bar{g}^p = 0$. Since F_A is injective, it follows that $\bar{g} = 0$. But this means that $g \in \ker(\phi)$ and as g has smaller degree than f, this contradicts our assumption. □

PROOF OF NOETHER NORMALIZATION. By assumption there exists an epimorphism $\phi : K[x_1,\dots,x_m] \to A$ of K-algebras. We then proceed with induction on m. When ϕ is an isomorphism, there is nothing to show. Otherwise, Lemma 9.3 tells us that there exists a K-algebra homomorphism $\phi' : K[x_1,\dots,x_m] \to A$ such that $\phi'(x_m)$ is integral over $A' := \phi'(K[x_1,\dots,x_{m-1}])$. Since the K-algebra A' has $\leq m-1$ generators, it is by induction a finite extension of some polynomial algebra $K[x_1,\dots,x_r]$. Hence so is A.

This induction argument also works in the prime-to-p setting. □

Corollary 9.4. For every affine variety X there exists an integer $r \geq 0$ and a finite surjective morphism $f : X \to \mathbb{A}^r$ such that $k(X)$ is separable over $k(\mathbb{A}^r)$. □

This corollary gives us a better grasp on the geometry of X, especially when X is irreducible, for it shows that X can then be 'spread' in a finite-to-one manner over affine r-space such that it is like a covering of degree $[k(X) : k(\mathbb{A}^r)]$ over a nonempty open subset of $\mathbb{A}^r$. Here is another important consequence.

Corollary 9.5. Every affine variety is birationally equivalent to a hypersurface.

PROOF. We let $f : X \to \mathbb{A}^r$ be as in Corollary 9.4 so that $k(X)$ is a finite separable extension of $k(x_1,\dots,k_r)$. By the theorem of the primitive element, $k(X)$ is then generated over $k(x_1,\dots,k_r)$ by a single element g. Let $G \in k(x_1,\dots,k_r)[x_{r+1}]$ be the minimal polynomial of g. By multiplying this by a common denominator of the coefficients, we arrange that the coefficients of G lie in $k[x_1,\dots,k_r]$, but have no common denominator. Then G is irreducible as an element of $k[x_1,\dots,x_{r+1}]$ and hence defines an irreducible hypersurface $Z(G) \subset \mathbb{A}^{r+1}$. It is clear that $k(Z(G)) = k(X)$. □

Proposition 8.4 has a kind of converse, also due to Noether.

Theorem 9.6. Let A be a noetherian domain and let $L/\operatorname{Frac}(A)$ be a separable finite extension of fields. Then the integral closure of A in L is finite over A.

PROOF. Let us write K for $\operatorname{Frac}(A)$, d for the degree of K/L and $B \subset L$ for the integral closure of A in L. By Proposition 8.4 we have $L = KB$ and so there exists a K-basis $(e_1, \dots, e_d)$ of L with each e_i in B. We recall that if $b \in L$ has minimal polynomial $x^n + c_1 x^{n-1} + \cdots + c_n$ in $K[x]$, then its K-trace, $tr_{L/K}(b)$, is defined as $(-1)^n c_n$. Note that when $b \in B$, its trace lies in A. The *trace form*

$$(x, y) \in L \times L \to tr_{L/K}(xy) \in K$$

is symmetric bilinear, when L is considered as a d-dimensional K-vector space. Since L/K is separable, it is nondegenerate and so we have a K-basis $(f_1, \dots, f_d)$ of L dual to $(e_1, \dots, e_d)$: characterized by $tr_{L/K}(e_i f_j) = \delta_{ij}$.

We prove that B is contained in the A-submodule generated by $(f_1, \dots, f_d)$. This will suffice, for since A is noetherian, it then follows that B is also finitely generated. Given $b \in B$, write $b = \sum_{i=1}^{d} a_i f_i$ with $a_i \in K$ so that $a_i = tr_{L/K}(be_i)$. Since $be_i \in B$, it follows that $a_i \in A$. □

Of special interest is the case $L = \operatorname{Frac}(A)$. Then $\widehat{A} := \overline{A}^L$ is called the *normalization* of A. So by the above theorem the normalization is a finite extension of A when the latter is finitely generated over a field. We say that A is *normal* if it is integrally closed in its fraction field, i.e., if $\widehat{A} = A$.

In the case of interest to us we can get rid of the separability assumption and then this theorem has a remarkable geometric interpretation:

Corollary 9.7. Let X be an irreducible affine variety. Then for ever finite field extension $L/k(X)$, the integral closure of $k[X]$ in L is finite over $k[X]$. So this defines a normal irreducible affine variety X_L and a finite surjective morphism $X_L \to X$ which induces the given field extension $k(X) \subseteq L$.

PROOF. When $L/k(X)$ is separable, this is immediate from Theorem 9.6. We reduce to that case as follows.

By Proposition 8.4 there exists a $k(X)$-basis of L (so of size $d = [L : k(X)]$) consisting of elements integral over $k(X)$. Then the k-subalgebra of L generated by this basis is finite over $k[X]$; it is the coordinate ring $k[Y]$ of an affine variety Y which maps in a finite dominant manner to X and for which $k(Y) = L$. The integral closure of $k[X]$ in L is then the integral closure $\widehat{k[Y]}$ of $k[Y]$ in $k(Y)$. Since $k[Y]$ is finite over $k[X]$, it now suffices to prove that $\widehat{k[Y]}$ is finite over $k[Y]$. By Corollary 9.4 there exists a finite surjective morphism $Y \to \mathbb{A}^r$ such that $k(Y)$ is separable over $k(\mathbb{A}^r)$. Since $\widehat{k[Y]}$ is the integral closure of $k[\mathbb{A}^r]$ in $k(Y)$, it is finite over $k[\mathbb{A}^r]$ and hence finite over $k[Y]$. □

So every finite field extension of $k(X)$ is canonically realized by a finite morphism of irreducible affine varieties. We write $\widehat{X} = \operatorname{Spm}(\widehat{k[X]})$ for X_L when $L = k(X)$. The associated morphism $\widehat{X} \to X$ is finite and birational, and is called the *normalization* of X. If $k[X]$ is normal (so that $\widehat{X} = X$), then we say that X is *normal*. This is functorial: if $f : X \to Y$ is a dominant morphism with X irreducible, then we may regard $k(X)$ as an extension of $k(Y)$ and hence $\widehat{k[X]}$ will contain $\widehat{k[Y]}$. This means that f determines a morphism $\widehat{f} : \widehat{X} \to \widehat{Y}$.

The affine space $\mathbb{A}^n$ is normal. More generally:

Lemma 9.8. Any unique factorization domain is normal.

PROOF. Let A be a UFD. Any $b \in \mathrm{Frac}(A)$ integral over A obeys an equation $b^d + a_1 b^d + \cdots + a_d = 0$ with $a_i \in A$. Write $b = r/s$ with $r, s \in A$ such that r and s are relatively prime. The identity $r^d + a_1 r s^{d-1} + \cdots + a_d s^d = 0$ shows that any prime divisor which divides s must divide r^d and hence also r. As there is no such prime, it follows that s is a unit so that $b \in A$. □

The Riemann extension theorem asserts that a meromorphic function on an open subset $U \subset \mathbb{C}$ which is locally bounded on U is in fact holomorphic on U. We may understand the normality of an irreducible affine variety X as an algebraic formulation of this property. For let $f, g \in k[X]$ be such that $g \neq 0$ and $\phi := f/g$ is integral over $k[X]$, i.e., $\phi^n + a_1\phi^{n-1} + \cdots + a_n = 0$ for certain n and $a_i \in k[X]$. Now assume that $k = \mathbb{C}$ and let $U \subset X$ be an open, relatively compact (=bounded) subset for the Hausdorff topology. Since the a_i's are continuous functions on X for the Hausdorff topology, they are bounded on U([10]). It then follows that ϕ is also bounded on U. But ϕ is also univalued on $U \cap X_g$. It can be shown that every Hausdorff open subset of X meets X_g. So although ϕ may not be everywhere defined on X, it is (for the Hausdorff topology) locally bounded on X. One can also verify the converse: if a rational function on X is locally bounded on X for the Hausdorff topology, then it is integral over $k[X]$. So X is normal if and only if the Riemann extension theorem holds on X: every locally bounded rational function ϕ on X is in fact regular on X.

The following example and the subsequent exercise illustrate an interesting property of normalization, which when phrased in terms that have not been defined but still have some intuitive appeal, amounts to the removal of singularities in codimension one and (in particular) the separation of local branches.

EXAMPLE 9.9 (Example 4.4 continued). We claim that the morphism $f : \mathbb{A}^1 \to C$ in 4.4 is a normalization, in particular, that C is not normal. This amounts to asserting that the integral closure $\widehat{k[C]}$ of $k[C] \cong k[t^2, t^3]$ in $k(C) \cong k(t)$ is $k[\mathbb{A}^1] \cong k[t]$. Since $k[t]$ is integrally closed in $k(t)$ (it is a UFD), we must have $\widehat{k[C]} \subseteq k[t]$. But $k[t]$ is as a $k[t^2, t^3]$-module spanned by 1 and t and the latter is integral over $k[t^2, t^3]$ as its square lies in this subring. So we also have $\widehat{k[C]} \supseteq k[t]$.

EXERCISE 35. Consider the curve in $\mathbb{A}^2$ defined by $y^2 = x^3 + x^2$. Prove that this curve is irreducible and determine its normalization. Show in particular that the normalization is not a bijection. Draw the locus in $\mathbb{R}^2$ defined by this equation.

The following two lemmas show that normality is a local in nature:

Lemma 9.10. Let R be a domain and $S \subset R$ a multiplicative subset. Then $S^{-1}\widehat{R}$ is the normalization of $S^{-1}R$.

PROOF. An equation of integral dependence for an element $u \in \mathrm{Frac}(R)$ over $S^{-1}R$ is of the form $u^n + a_1 u^{n-1} + \cdots + a_n = 0$ with $a_i \in S^{-1}R$. If $s \in S$ is a common denominator of $a_1, \ldots, a_n$, then we have $(su)^n + sa_1(su)^{n-1} + \cdots + s^n a_n = 0$. This is an equation of integral dependence for su over R and hence $u \in S^{-1}\widehat{R}$. □

[10] The roots of a polynomial can be bounded in terms of its coefficients. For if $(t - z_1) \cdots (t - z_n) = t^n + a_1(z)t^{n-1} + \cdots + a_n(z)$ with $z_i \in \mathbb{C}$, then $z \in \mathbb{C}^n \mapsto \max_i |a_i(z)|$ is a continuous function which is never zero on the compact set $\max_i |z_i| = 1$ and so has there a minimum $c > 0$. Since a_i is homogeneous of degree i, this implies that $\max_i |z_i| \leq c^{-1} \max_i |a_i|^{1/i}$.

Proposition 9.11. An irreducible affine variety X is normal if and only if each local ring $\mathcal{O}_{X,p}$ ($p \in X$) is normal.

PROOF. Assume X is normal. Then $\mathcal{O}_{X,p}$, being a localization of $k[X]$, is also normal by Lemma 9.10.

Now assume that every $\mathcal{O}_{X,p}$ is normal and let $f \in k(X)$ satisfy an equation of integral dependence over $k[x]$. For every $p \in X$ this gives an equation of integral dependence over $\mathcal{O}_{X,p}$ and so by assumption $f \in \mathcal{O}_{X,p}$. This proves that f is a regular function on X so that $f \in k[X]$. □

The preceding also helps us to gain some geometric understanding of the algebraic closure of a function field. Let be given an irreducible affine variety Y and let $L/k(Y)$ be an algebraic closure of $k(Y)$. Then $L/k(Y)$ will not be a finite extension, unless Y is a singleton, but L can be written as a monotone union of finite field extensions: $L = \cup_{i=1}^{\infty} L_i$ with $L_i \subseteq L_{i+1}$ and L_{i+1}/L_i finite. This yields a sequence of finite surjective morphisms

$$Y \twoheadleftarrow Y_{L_1} \twoheadleftarrow Y_{L_2} \twoheadleftarrow Y_{L_3} \twoheadleftarrow \cdots$$

of which the projective limit can be understood as a "pro-affine variety" (a point of this limit is given by a sequence $(y_i \in Y_{L_i})_{i=1}^{\infty}$ such that y_{i+1} maps to y_i for all i). Its algebra of regular functions is $\cup_{i=1}^{\infty} \overline{k[Y]}^{L_i} = \overline{k[Y]}^{L}$ (which is usually not a finitely generated k-algebra) and its function field is L.

Of special interest is when we have a finite *normal*([11]) field extension of a function field. Let us first recall this notion.

NORMAL EXTENSIONS. We recall that an algebraic field extension L/K is *normal* if the minimal polynomial $f \in K[x]$ of an element of L has *all* its roots in L. A Galois extension L/K is the same thing as a normal separable extension (this property can be used as a definition); this requires in addition that all these roots are distinct. Note however that a purely inseparable extension L/K is also normal, for then such an f of degree > 1 is of the form $x^{p^r} - a$, with p the characteristic p of K (which must be positive), $r \geq 1$ and $a \in K$ not a pth power in K. Clearly an algebraic closure $\overline{K}$ of K is normal.

Part of Galois theory still works for normal extensions. If L/K is normal, then all the K-linear field embeddings $L \hookrightarrow \overline{K}$ have the same image and so this image is invariant under the full Galois group of $\overline{K}/K$. The latter then restricts to the group of K-linear field automorphisms of L (the Galois group of L/K) and this group permutes the roots of a minimal polynomial in $K[x]$ of any element of L transitively.

For an arbitrary algebraic field extension F/K, one defines its *normal closure* in $\overline{K}$ as the smallest normal extension of K in $\overline{K}$ that admits a K-linear embedding of F into it. It is obtained as the subfield of $\overline{K}$ generated by the roots of all the irreducible polynomials of $K[x]$ that have a root in F. When F is finite over K, then so is its normal closure in $\overline{K}$.

We begin with the corresponding result in commutative algebra. This has also important applications in algebraic number theory.

Proposition 9.12. Let A be a normal domain and $L/\operatorname{Frac}(A)$ be a finite normal extension with Galois group G. Then G leaves invariant the integral closure $\overline{A}^L$ of A in L and has the property that for every prime ideal $\mathfrak{p} \subseteq A$ it permutes the set of prime ideals $\mathfrak{q} \subseteq \overline{A}^L$ with $\mathfrak{q} \cap A = \mathfrak{p}$ transitively.

[11] As the statement of Proposition 9.12 illustrates, this adjective is a bit overused in mathematics: a normal field extension should not be confused with the *normality* of a ring.

For the proof we need:

Lemma 9.13 (The prime avoidance lemma)**.** Any ideal of a ring that is contained in a finite union of prime ideals is contained in one of them.

PROOF. Let R be a ring, $\mathfrak{q}_1, \dots, \mathfrak{q}_n$ prime ideals in R and $I \subseteq R$ an ideal contained in $\cup_{i=1}^n \mathfrak{q}_i$. We prove with induction on n that $I \subseteq \mathfrak{q}_i$ for some i. The case $n = 1$ being trivial, we may assume that $n > 1$ and that for every $i = 1, \dots, n$, I is not contained in $\cup_{j \neq i} \mathfrak{q}_j$. Choose $a_i \in I \smallsetminus \cup_{j \neq i} \mathfrak{q}_j$. So then $a_i \in \mathfrak{q}_i$. Consider $a := a_1 a_2 \cdots a_{n-1} + a_n$. Then $a \in I$ and hence $a \in \mathfrak{q}_i$ for some i. If $i < n$, then $a_n = a - a_1 a_2 \cdots a_{n-1} \in \mathfrak{q}_i$ and we get a contradiction. If $i = n$, then $a_1 a_2 \cdots a_{n-1} = a - a_n \in \mathfrak{q}_n$ and hence $a_j \in \mathfrak{q}_n$ for some $j \leq n - 1$. This is also a contradiction. □

PROOF OF PROPOSITION 9.12. That any $g \in G$ leaves $\overline{A}^L$ invariant is clear, for g fixes the coefficients of an equation of integral dependence over A.

By the Going-up theorem 8.6, there exists a prime ideal $\mathfrak{q}$ of $\overline{A}^L$ that lies over $\mathfrak{p}$. Let $\mathfrak{q}'$ be another. We show that $\mathfrak{q}' \subseteq \cup_{\sigma \in G} \sigma(\mathfrak{q})$. This suffices, for then the prime avoidance lemma implies that $\mathfrak{q}' \subseteq \sigma(\mathfrak{q})$ for some $\sigma \in G$. As both $\mathfrak{q}'$ and $\sigma(\mathfrak{q})$ lie over $\mathfrak{p}$, we must have $\mathfrak{q}' = \sigma(\mathfrak{q})$ by incomparability.

Let $b \in \mathfrak{q}'$ be nonzero and let $f(x) = x^n + c_1 x^{n-1} + \cdots + c_n \in \mathrm{Frac}(A)[x]$ be a minimum polynomial for b. Since $L/\mathrm{Frac}(A)$ is a normal extension, f completely factors in L with roots in Gb: $f(x) = (x - \sigma_1(b)) \cdots (x - \sigma_n(b))$ for certain $\sigma_i \in G$. Since $\sigma_i(b) \in \overline{A}^L$ we have $\sigma_1(b) \cdots \sigma_n(b) \in \overline{A}^L$. On the other hand, $\sigma_1(b) \cdots \sigma_n(b) = (-1)^n c_n \in \mathrm{Frac}(A)$ and since A is normal it follows that $\sigma_1(b) \cdots \sigma_n(b) \in A$. One of the factors is b and so $\sigma_1(b) \cdots \sigma_n(b) \in A \cap (b) \subseteq A \cap \mathfrak{q}' = \mathfrak{p} \subseteq \mathfrak{q}$. Since $\mathfrak{q}$ is a prime ideal, some factor $\sigma_i(b)$ lies in $\mathfrak{q}$ and hence $b \in \cup_{\sigma \in G} \sigma(\mathfrak{q})$. □

We translate this into geometry:

Corollary 9.14. Let Y be an irreducible affine variety that is normal. Given a normal finite extension $L/k(Y)$, then the Galois group G of $L/k(Y)$ acts naturally on Y_L. For every closed irreducible subset $Z \subseteq Y$, the preimage of Z in Y_L is nonempty and G acts transitively on its set irreducible components. As a topological space, Y may be identified with the G-orbit space of X.

PROOF. Everything is clear except perhaps the last assertion. The map $\pi : Y_L \to Y$ is continuous and closed (Exercise 33). By Proposition 9.12, every fiber of π is a (nonempty) G-orbit and so π_L is topologically the formation of the G-orbit space. □

Exercise 35 provides an example of an irreducible Y for which $Y_L = \widehat{Y} \to Y$ is not a bijection and so we cannot drop in Proposition 9.12 the assumption that Y be normal (here $L = k(Y)$ so that G is trivial).

Here is a supplement of the going up property for normal domains:

Corollary 9.15 (Going down)**.** Let $A \subseteq B$ be a finite extension with B a domain and A normal. Given prime ideals $\mathfrak{p}$ in A and $\mathfrak{q}'$ in B such that $\mathfrak{p} \subseteq \mathfrak{q}' \cap A$, then there exists a prime ideal $\mathfrak{q}$ in B with $\mathfrak{q} \subseteq \mathfrak{q}'$ and $\mathfrak{q} \cap A = \mathfrak{p}$.

PROOF. Put $K := \mathrm{Frac}(A)$ and let L be the normal closure of $\mathrm{Frac}(B)$ in an algebraic closure of $\mathrm{Frac}(B)$. Since B is integral over A, we have $B \subseteq \overline{A}^L$. Now L is a finite normal extension of $\mathrm{Frac}(B)$ (with Galois group G, say), and this brings

us in the situation of Proposition 9.12 above. Since $\mathrm{Frac}(B)$ is finite over K, L is finite over K and so $\overline{A}^L$ is by 9.7 finite over A (and hence also over B).

Put $\mathfrak{p}' := \mathfrak{q}' \cap A$ so that $\mathfrak{p}' \supseteq \mathfrak{p}$. According to Proposition 8.6 we can find a prime ideal $\mathfrak{r}'$ in $\overline{A}^L$ with $\mathfrak{r}' \cap B = \mathfrak{q}'$. The same proposition tells us that there exist nested prime ideals $\tilde{\mathfrak{r}}' \supseteq \tilde{\mathfrak{r}}$ in $\overline{A}^L$ which meet A in $\mathfrak{p}' \supseteq \mathfrak{p}$. Since $\tilde{\mathfrak{r}}'$ and $\mathfrak{r}'$ both meet A in $\mathfrak{p}'$, there exists according to Proposition 9.12 a $\sigma \in G$ such that $\sigma(\tilde{\mathfrak{r}}') = \mathfrak{r}'$. So $\mathfrak{r} := \sigma(\tilde{\mathfrak{r}})$ has the property that $\mathfrak{r} \subseteq \mathfrak{r}'$ and $\mathfrak{r} \cap A = \mathfrak{p}$. So $\mathfrak{q} := \mathfrak{r} \cap B$ is as desired, for it meets A in $\mathfrak{p}$ and is contained in $\mathfrak{r}' \cap B = \mathfrak{q}'$. □

REMARK 9.16. We can rephrase this in the spirit of Remark 8.7 by saying that any prime chain in A is the intersection of prime chain in B for which the *last* member has been prescribed in advance (whence 'going down').

10. Dimension

One way to define the dimension of a topological space X is with induction: agree that the empty set has dimension -1 and that X has dimension $\leq n$ if it admits a basis of open subsets such that the boundary of every basis element has dimension $\leq n-1$. This is close in spirit to the definition that we shall use here (which is however adapted to the Zariski topology; as you will find in Exercise 36, it is useless for Hausdorff spaces).

DEFINITION 10.1. Let X be a nonempty topological space. We say that the *Krull dimension of* X is at least d if there exists an *irreducible chain of length* d in X, that is, a strictly descending chain of closed irreducible subsets $X^0 \supsetneq X^1 \supsetneq \cdots \supsetneq X^d$ of X. The *Krull dimension* of X is the supremum of the d for which an irreducible chain of length d exists and we then write $\dim X = d$. We stipulate that the Krull dimension of the empty set is -1.

Lemma 10.2. For every subspace Z of a topological space X we have $\dim Z \leq \dim X$.

PROOF. Let Y be closed in Z. The closure $\overline{Y}$ of Y in X is the intersection of all the closed subsets of X containing Y, hence meets Z in Y: $\overline{Y} \cap Z = Y$. We also know that if $Y \subset Z$ is irreducible, then so is $\overline{Y}$. So if we have an irreducible chain of length d in Z, then the closures of the members of this chain yield an irreducible chain of length d in X. This proves that $\dim Z \leq \dim X$. □

EXERCISE 36. What is the Krull dimension of a nonempty Hausdorff space?

EXERCISE 37. Let U be an open subset of the space X. Prove that for an irreducible chain $Y^\bullet$ in X of length d with $U \cap Y^d \neq \emptyset$, $U \cap Y^\bullet$ is an irreducible chain of length d in U. Conclude that if $\mathcal{U}$ is an open covering of X, then $\dim X = \sup_{U \in \mathcal{U}} \dim U$.

EXERCISE 38. Suppose that X is a noetherian space. Prove that the dimension of X is the maximum of the dimensions of its irreducible components. Prove also that if all the singletons (= one element subsets) in X are closed, then $\dim(X) = 0$ if and only if X is finite.

It is straightforward to translate this notion into algebra:

DEFINITION 10.3. The *Krull dimension* $\dim R$ of a ring R is the supremum of the integers d for which there exists an *prime chain of length* d in R, where we stipulate that the zero ring (i.e., the ring which has no prime ideals) has Krull dimension -1.

It is clear from Proposition 2.3 that for a closed subset $X \subset \mathbb{A}^n$, $\dim k[X] = \dim X$. Since any prime ideal of a ring R contains the ideal $\sqrt{(0)}$ of nilpotents, R and its 'reduction' $R_{\mathrm{red}} := R/\sqrt{(0)}$ have the same Krull dimension. So the Krull dimension of a finitely generated k-algebra A is that of the affine variety $\mathrm{Spm}(A)$.

Remark 8.7 shows immediately:

Lemma 10.4. The Krull dimension is invariant under integral extension: if B is an integral extension of A, then A and B have the same Krull dimension.

REMARK 10.5. For a domain A the zero ideal (0) is a prime ideal and so $\dim(A) = 0$ if and only if the ideal (0) is maximal, i.e., A is a field. We say that a domain A is a *Dedekind domain* if it is noetherian, normal and of dimension ≤ 1, in other words, if every nonzero prime ideal of A is maximal. For instance, a unique factorization domain (such as $\mathbb{Z}$ and $K[X]$ with K a field) is a Dedekind domain. The importance of this notion comes from the fact that a converse holds on the level of ideals: any ideal of a Dedekind domain is uniquely written a product of prime ideals. Lemma 10.4 shows that the integral closure of a Dedekind domain (such as $\mathbb{Z}$ or $K[X]$) in a finite extension of its field of fractions is a Dedekind domain. Hence the ring of integers of an algebraic number field L (this is by definition the integral closure of $\mathbb{Z}$ in L) is a Dedekind domain.

The Krull dimension was easy to define, but seems difficult to compute in concrete cases. How can we be certain that a given prime chain has maximal possible length? It is not even clear how to tell whether the Krull dimension of a given ring is finite. We will settle this in a satisfactory manner for a domain B containing a field K over which it is a finitely generated: we show that a length of a prime chain in B is bounded by the transcendence degree $\mathrm{Frac}(B)/K$ and that we have equality when the prime chain is maximal (so that the length of *any* maximal prime chain is the Krull dimension).

Theorem 10.6. Let K be a field and B a finitely generated K-algebra without zero divisors. Then the Krull dimension of B equals the transcendence degree of $\mathrm{Frac}(B)/K$ and every maximal prime chain in B (i.e., one that cannot be extended to a longer prime chain) has length $\dim B$.

PROOF. By Noether normalization there exists an integer $r \geq 0$ and a K-algebra monomorphism $K[x_1, \ldots, x_r] \hookrightarrow B$ such that B is finite over $K[x_1, \ldots, x_r]$. We put $A := K[x_1, \ldots, x_r]$. Then $\mathrm{Frac}(B)$ is a finite extension of $\mathrm{Frac}(A) = K(x_1, \ldots, x_r)$ and so the transcendence degree of $\mathrm{Frac}(B)/K$ is r. In A we have the length r prime chain $(0) \subsetneq (x_1) \subsetneq (x_1, x_2) \subsetneq \cdots \subsetneq (x_1, x_2, \ldots, x_r)$. By Remark 8.7 this is the intersection of A with a prime chain in B and so the Krull dimension of B is at least r.

When $r = 0$, B is a domain that is finite over a field, hence itself a field. Then both the Krull dimension of B and its transcendence degree over K are zero. Assume therefore $r > 0$ and the theorem proved for lower values of r. Let $\mathfrak{q}_\bullet := ((0) = \mathfrak{q}_0 \subsetneq \mathfrak{q}_1 \subsetneq \cdots \subsetneq \mathfrak{q}_m)$ be a prime chain in B. We prove that $m \leq r$ with equality when $\mathfrak{q}_\bullet$ is a maximal prime chain.

By the incomparability property of 'going up' (Proposition 8.6), $\mathfrak{p}_\bullet := \mathfrak{q}_\bullet \cap A$ will be a prime chain in A, also of length m. The idea is to show that $\mathfrak{p}_1$ defines a closed subset of $\mathbb{A}^r_K$ of dimension $\leq m-1$. Choose an irreducible $f \in \mathfrak{p}_1$. After renumbering the coordinates, we may assume that f does not lie in $K[x_1, \ldots, x_{r-1}]$. So if we write $f = \sum_{i=0}^N g_i x_r^i$ with $g_i \in K[x_1, \ldots, x_{r-1}]$ and $g_N \neq 0$, then $N \geq 1$. Since f is irreducible, $A/(f)$ is a domain. The image of x_r in the field $\mathrm{Frac}(A/(f))$

is a root of the monic polynomial $t^N + \sum_{i=0}^{N-1}(g_i/g_N)t^i \in K(x_1,\dots,x_{r-1})[t]$, and so $\mathrm{Frac}(A/(f))$ is a finite extension of $K(x_1,\dots,x_{r-1})$. In particular, $\mathrm{Frac}(A/(f))$ has transcendence degree $r-1$ over K. By our induction induction hypothesis, $A/(f)$ has then Krull dimension $r-1$. Since the image of $\mathfrak{p}_1 \subsetneq \cdots \subsetneq \mathfrak{p}_m$ in $A/(f)$ is a prime chain of length $m-1$ (it is strictly ascending, because it is so in $A/\mathfrak{p}_1$), it follows that $m-1 \leq r-1$. Hence $m \leq r$.

Now assume that $\mathfrak{q}_\bullet$ is maximal. Since A is a normal domain, the 'going down' Corollary 9.15 applies and so there exists a prime ideal $\mathfrak{q} \subseteq \mathfrak{q}_1$ such that $\mathfrak{q} \cap A = (f)$. The maximality of $\mathfrak{q}_\bullet$ then implies that $\mathfrak{q} = \mathfrak{q}_1$ and hence that $(f) = \mathfrak{p}_1$. Since $\mathrm{Frac}(B/\mathfrak{q}_1)$ is a finite extension of $\mathrm{Frac}(A/(f))$ (which in turn is a finite extension of $K(x_1,\dots,x_{r-1})$), it has transcendence degree $r-1$ over K. As $\mathfrak{q}_1 \subsetneq \mathfrak{q}_2 \cdots \subsetneq \mathfrak{q}_m$ defines a maximal prime chain in $B/\mathfrak{q}_1$, it follows from our induction hypothesis that $m-1 = r-1$ and so $m = r$. □

REMARK 10.7. At first sight the preceding theorem may look somewhat surprising, as the definition of the Krull dimension of a ring does not mention a possible subfield (it is an 'absolute notion'), whereas the transcendence degree of a field extension clearly does. Corollary 3.7 explains why there is no contradiction here: if K is a field and a finitely generated K-algebra L is in fact a field, then L is finite over K and has therefore the same transcendence degree over a subfield K.

Corollary 10.8. In the situation of Theorem 10.6, let $\mathfrak{m}$ be a maximal ideal of B and $S \subseteq B \setminus \{0\}$ a multiplicative subset generated by a finite subset of B. Then the Krull dimension of the localizations $B_\mathfrak{m} = (B \setminus \mathfrak{m})^{-1}B$ and $S^{-1}B$ are that of B.

PROOF. Any prime chain in $B_\mathfrak{m}$ is a prime chain in B contained in $\mathfrak{m}$. So by Theorem 10.6, the Krull dimension of $B_\mathfrak{m}$ is finite. If such a chain is maximal for this property, then it will end with $\mathfrak{m}$ and will also be maximal in B (for there is no prime ideal strictly containing $\mathfrak{m}$) and again by Theorem 10.6 its length is then the Krull dimension of B.

If S is multiplicatively generated by $s_1,\dots,s_n$, then $S^{-1}B$ is the quotient of the K-algebra $B[x_1,\dots,x_n]$ by the ideal $(x_1s_1-1,\dots,x_ns_n-1)$. In particular, $S^{-1}B$ is a finitely generated K-algebra so that Theorem 10.6 applies to it: $\dim(S^{-1}B)$ is the transcendence degree of $\mathrm{Frac}(S^{-1}B) = \mathrm{Frac}(B)$ over K. This is just $\dim(B)$. □

Corollary 10.9. Let X be an irreducible affine variety and let d be the transcendence degree of $k(X)/k$. Then $\dim X = d$ and every maximal irreducible chain in X has length d. Moreover, d is also the Krull dimension of every nonempty open $U \subseteq X$ and of every local ring $\mathcal{O}_{X,p}$, $p \in X$.

REMARK 10.10. If $S \subseteq R$ is a multiplicative set, then the preimage of a prime ideal $\tilde{\mathfrak{q}}$ of $S^{-1}R$ under the ring homomorphism $R \to S^{-1}R$ is a prime ideal $\mathfrak{q}$ of R which does not meet S and we have $\tilde{\mathfrak{q}} = S^{-1}\mathfrak{q}$. This sets up a bijection between the prime ideals of $S^{-1}R$ and those of R not meeting S. If we take $S = R \setminus \mathfrak{p}$ with $\mathfrak{p}$ a prime ideal, then this implies that $\dim R_\mathfrak{p}$ is the supremum of the prime chains in R which end with $\mathfrak{p}$. Since the prime chains in R which begin with $\mathfrak{p}$ correspond to prime chains in $R/\mathfrak{p}$, it follows that $\dim R_\mathfrak{p} + \dim R/\mathfrak{p}$ is the supremum of the prime chains in R having $\mathfrak{p}$ as a member. So when R is a domain which is finitely generated over a subfield, then this is by Theorem 10.6 equal to $\dim R$.

An affine variety C all of whose irreducible components have dimension 1 is called a *curve*. When C is irreducible this amounts to the property that $k(C)$ is

of transcendence degree one and is equivalent to the property that every nonzero every prime ideal is a maximal ideal. By Remark 10.5, C is an irreducible normal affine curve if and only if $k[C]$ is a Dedekind domain $\neq k$. The fact that rings of integers of number fields and coordinate rings of such curves are both Dedekind domains accounts for their similarity. We will see in the next section that for a curve, normality amounts to the absence of 'singular points'.

EXERCISE 39. Prove that a hypersurface in $\mathbb{A}^n$ has dimension $n-1$.

EXERCISE 40. Let X be an irreducible affine variety and $Y \subseteq X$ a closed irreducible subset. Prove that the codimension of Y in X ($= \dim X - \dim Y$) is equal to the Krull dimension of $k[X]_{I(Y)}$.

Exercise Let X be an irreducible variety of dimension d and let $f_1, \ldots, f_r$ be $r \leq d$ be elements of $k[X]$. Prove that $\operatorname{codim}(Z(f_1, \ldots, f_r)) \leq r$ and show that if we have equality, then for every i the irreducible components of $Z(f_1, \ldots, f_i)$ have codimension i.

EXERCISE 41. Prove that when X and Y are irreducible affine varieties, then $\dim(X \times Y) = \dim X + \dim Y$. (Hint: Embed each factor as a closed subset of some affine space. You may also want to use the fact that the equality to be proven holds in case $X = \mathbb{A}^m$ and $Y = \mathbb{A}^n$.)

11. Smoothness

In this section we focus on the local properties of an affine variety $X = \operatorname{Spm}(A)$ (so $A := k[X]$ is here a reduced finitely generated k-algebra) at a point p. Therefore a central role will be played by the local algebra $\mathcal{O}_{X,p} = A_{\mathfrak{m}_p}$ whose maximal ideal is $\mathfrak{m}_{X,p} = (A \smallsetminus \mathfrak{m}_p)^{-1}\mathfrak{m}_p$.

If $k = \mathbb{C}$ and $X \subseteq \mathbb{C}^n$ is a closed subset of dimension d, then we hope that there is a nonempty open subset of X where X is 'smooth', i.e., where X looks like a complex submanifold of complex dimension d. Our goal is to define smoothness in algebraic terms (so that it make sense for our field k) and then to show that the set of smooth points of a variety is open and dense in that variety.

Our point of departure is the implicit function theorem. One version states that if $U \subseteq \mathbb{R}^n$ is open, $p \in \mathbb{R}^n$ and $f = (f_1, \ldots, f_{n-d}) : U \to \mathbb{R}^{n-d}$ a differentiable map such that the total differentials at p, $df_1(p), \ldots, df_{n-d}(p)$ are linearly independent in p (this is equivalent to: the Jacobian matrix of $(\partial f_j/\partial x_i)(p))_{i,j}$ has maximal rank $n-d$), then $f^{-1}f(p)$ is a submanifold of dimension d at p whose tangent space at p is the common zero set of $df_1(p), \ldots, df_{n-d}(p)$. In fact, one shows that this solution set is near p the graph of a map: we can parametrize this set by means of d coordinates out of $(x_1, \ldots, x_n)$, expressing the $n-d$ remaining ones as differentiable functions in terms of them. Conversely, any submanifold of $\mathbb{R}^n$ at p of dimension d is locally thus obtained.

We begin with the observation that for any ring R, partial differentiation of a polynomial $f \in R[x_1, \ldots, x_n]$ (where the elements of R are treated as constants) is well-defined and produces another polynomial. The same goes for a fraction $\phi = f/g$ in $R[x_1, \ldots, x_n][g^{-1}]$: a partial derivative of ϕ also lies in $R[x_1, \ldots, x_n][g^{-1}]$

(in this case with denominator g^2). We then define the *total differential* of a rational function $\phi \in R(x_1, \ldots, x_n)$ as usual:

$$d\phi := \sum_{i=1}^{n} \frac{\partial \phi}{\partial x_i}(x) dx_i,$$

where for now, we do not worry about how to interpret the symbols dx_i: we think of $d\phi$ simply as a regular map from an open subset of $\mathbb{A}^n$ to a k-vector space of dimension n with basis $dx_1, \ldots, dx_n$, leaving its intrinsic characterization for later. However, caution is called for when R is a field of positive characteristic:

EXERCISE 42. Prove that $f \in k[x]$ has zero derivative if and only if f is constant or (when $\mathrm{char}(k) = p > 0$) a pth power of some $g \in k[x]$.

Generalize this to: given $f \in k[x_1, \ldots, x_n]$, then $df = 0$ if and only if f is constant or (when $\mathrm{char}(k) = p > 0$) a pth power of some $g \in k[x_1, \ldots, x_n]$.

We should also be aware of the failure of the inverse function theorem:

EXAMPLE 11.1. Let $C \subseteq \mathbb{A}^2$ be the curve defined by $y^2 = x^3 + x$. By any reasonable definition of smoothness we should view the origin $o := (0, 0)$ as a smooth point of C. Indeed, when $k = \mathbb{C}$, the projection $f : C \to \mathbb{A}^1$, $(x, y) \mapsto y$, would be a local-analytic isomorphism at o. But the map is not locally invertible within our category: the inverse requires us to find a rational function $x = f(y)/g(y)$ which solves the equation $y^2 = x^3 + x$. This is impossible: we may assume that f and g are relatively prime and from $y^2 g^3 = f^3 + f g^2$, we see that g divides f^3. This can only happen when $g \in k^\times$ and so we may assume that $g = 1$: $y^2 = f^3 + f$. But then $f \in k[y]$ must divide y^2, and hence be equal to a constant times y or y^2 and neither case provides a solution.

We can however solve for x formally: $x = \phi(y) = y^2 + c_3 y^3 + c_4 y^4 + \cdots$, where it is important to note that the coefficients are all integers so that this works for every characteristic. By this we mean that if we put $\phi_n = y^2 + c_3 y^3 + \cdots + c_n y^n$, then $y^2 - \phi_n(y)^3 - \phi_n(y) \in (y)^{n+1}$ for all $n \geq 2$.

Somewhat related to this is an issue illustrated by the following example.

EXAMPLE 11.2. Consider the curve $C \subseteq \mathbb{A}^2$ defined by $xy = x^3 + y^3$. The polynomial $x^3 + y^3 - xy$ is irreducible in $k[x, y]$, so that $k[C]$ is without zero divisors and C' is irreducible. Hence the local ring $\mathcal{O}_{C,o} \subseteq k(C)$ is also without zero divisors. But C seems to have two branches at o which apparently can only be recognized formally: there exists a formal power series $\phi(t) = t^2 + c_3 t^3 + c_4 t^4 + \cdots$ such that one such branch is given by $y = \phi(x)$ and the other by interchanging the roles of x and y: $x = \phi(y)$. To be precise, if for an integer $n \geq 2$, we put $\phi_n := t^2 + c_3 t^3 + c_4 t^4 + \cdots + c_n t^n$, then $(x - \phi_n(y))(y - \phi_n(x)) \equiv xy - x^3 - y^3 \mod (x, y)^{n+1}$, If we use $\xi := x - \phi(y)$ and $\eta := y - \phi(x)$ as new 'formal coordinates', then C is simply given at 0 by the reducible equation $\xi\eta = 0$.

These examples make it clear that for a local understanding of a variety X at o, the local ring $\mathcal{O}_{X,o}$ still carries too much global information. One way to get rid of this overload is by passing formal to power series. This is accomplished by what is known as formal completion([12]).

[12]Another approach would be to allow 'algebraic' functions of the type that we encountered in the two examples above, but then we would have to address the question what the domain of such a

11.3. FORMAL COMPLETION. Let R be a ring and $I \subseteq R$ an ideal. For every R-module M, the descending sequence of submodules $M \supseteq IM \supseteq I^2M \supseteq \cdots \supseteq I^nM \supseteq \cdots$ gives rise to a sequence of surjective R-homomorphisms

$$0 = M/M \twoheadleftarrow M/IM \twoheadleftarrow M/I^2M \twoheadleftarrow M/I^3M \twoheadleftarrow \cdots \twoheadleftarrow M/I^nM \leftarrow \cdots$$

from which we can form the R-module $\hat{M}_I := \varprojlim_n M/I^nM$, called the *$I$-adic completion* of M. So any $\hat{a} \in \hat{M}_I$ is uniquely given by a sequence $(\alpha_n \in M/I^nM)_{n\geq 0}$ whose terms are compatible in the sense that α_n is the reduction of α_{n+1} for all n. In this way $\hat{M}_I$ can be regarded as an R-submodule of $\prod_{n\geq 0}(M/I^nM)$. The natural R-homomorphisms $M \to M/I^nM$ combine to define a R-homomorphism $M \to \hat{M}_I$. Its kernel is $\cap_{n=0}^{\infty} I^nM$ and this turns out to be trivial in many cases of interest (see Lemma 11.4). If we do this for the ring R, we get a ring $\hat{R}_I$ (for each R/I^n is one and the reduction maps are ring homomorphisms) and $R \to \hat{R}_I$ is then a ring homomorphism.

Since M/I^nM is naturally a R/I^n-module, the R-module structure on $\hat{M}_I$ factors through a $\hat{R}_I$-module structure: for any $\hat{r} = (\rho_n \in R/I^n)_{n=0}^{\infty} \in \hat{R}_I$, $\hat{r}\hat{a}$ simply as given by the sequence $(\rho_n\alpha_n)_{n\geq 0}$ (note that $\rho_n\alpha_n$ is indeed the reduction of $\rho_{n+1}\alpha_{n+1}$). Any R-homomorphism $\phi : M \to N$ of R-modules sends M/I^nM to N/I^nN, and the resulting homomorphisms $M/I^nM \to N/I^nN$ are compatible in the sense that they determine a $\hat{R}_I$-homomorphism $\hat{\phi}_I : \hat{M}_I \to \hat{N}_I$. We have thus defined a functor from the category of R-modules to the category of $\hat{R}_I$-modules.

Lemma 11.4. If R is a noetherian local ring and $I \neq R$, then any finitely generated R-module M is Hausdorff for the I-adic topology: $\cap_{n\geq 0} I^nM = 0$.

PROOF. Put $M_\infty := \cap_{n\geq 0} I^nM$. Since M is noetherian, its submodule M_∞ is a finitely generated R-module. It is clear that $M_\infty = IM_\infty$. Since I is contained in the maximal ideal of R, Nakayama's lemma 8.5 then implies that $M_\infty = 0$. □

EXAMPLE 11.5. When A is a ring, $R = A[x_1, \ldots, x_n]$, then the $(x_1, \ldots, x_n)$-adic completion of R is just the ring of formal power series $A[[x_1, \ldots, x_n]]$.

As a variation on this example, first observe that the natural homomorphism $k[x_1, \ldots, x_n] \to \mathcal{O}_{\mathbb{A}^n,p}/\mathfrak{m}_p^r$ is surjective with kernel $(x_1 - p_1, \ldots, x_n - p_n)^r$. This implies that $k[[x_1 - p_1, \ldots, x_n - p_n]]$ can be identified with the $\mathfrak{m}_p$-adic completion of $\mathcal{O}_{\mathbb{A}^n,p}$. In the same spirit, we will find that for (C, o) in Example 11.1 resp. 11.2 the $\mathfrak{m}_{C,o}$-adic completion of $\mathcal{O}_{C,o}$ is isomorphic to $k[[x]]$ resp. $k[[x, y]]/(xy)$.

EXAMPLE 11.6. Take the ring $\mathbb{Z}$. Its completion with respect to the ideal (n), n an integer ≥ 2, yields the ring of n-adic integers $\hat{\mathbb{Z}}_{(n)}$: an element of $\hat{\mathbb{Z}}_{(n)}$ is given by a sequence $(\rho_i \in \mathbb{Z}/(n^i))_{i=1}^{\infty}$ with the property that ρ_i is the image of ρ_{i+1} under the reduction $\mathbb{Z}/(n^{i+1}) \to \mathbb{Z}/(n^i)$.

EXERCISE 43. Prove that if $\phi : M \to N$ is a surjection of R-modules, then $\hat{\phi}_I : \hat{M}_I \to \hat{N}_I$ is surjection (of $\hat{R}_I$-modules). (This need not be true for injections, but we will see that this is so in the noetherian setting.)

11.7. ADIC TOPOLOGIES. We can understand an I-adic completion as completion with regard to a topology. This often helps to clearify its dependence on I

function should be. This can not be achieved by refining the Zariski topology. Rather, this forces us to revisit the very notion of a topology, leading up to what is called the *étale topos*. Despite its somewhat abstract nature this is closer to our geometric intuition than the Zariski topology.

(which is weaker than one might be inclined to think). The topology on an R-module M is the *I-adic topology* on M, of which a basis is the collection of additive translates of the submodules I^nM, i.e., the collection of subsets $a + I^nM$, $a \in M$, $n \geq 0$. This is a topology indeed: given two basic open subsets $a+I^nM$, $a'+I^{n'}M$, then for any element c in their intersection, the basic open subset $c + I^{\max\{n,n'\}}M$ is also in their intersection. So a sequence $(a_n \in M)_{n\geq 1}$ converges to $a \in M$ precisely when for every integer $s \geq 0$, $a_n \in a + I^sM$ for n large enough. The fact that our basis is translation invariant implies that with this topology, M is a topological abelian group ($(a, b) \in M \times M \mapsto a - b \in M$ is continuous). If we endow R also with the I-adic topology and $R \times M$ with the product topology, then the map $(r, a) \in R \times M \to ra \in M$ which gives the module structure is also continuous. It is clear that any R-module homomorphism is continuous for the I-adic topology.

The I-adic topology on $\hat{M}_I$ is Hausdorff: if $0 \neq a \in \hat{M}_I$, then a has a nonzero component in M/I^nM for some n and then $I^n\hat{M}_I$ and $a+I^n\hat{M}_I$ are disjoint neighborhoods of 0 and a. Since the topology on M comes from one on $\hat{M}_I$ in the sense that the open subsets of M are pre-images of open subsets of $\hat{M}_I$, M is Hausdorff precisely when $\ker(M \to \hat{M}_I) = \cap_{n\geq 0}I^nM$ is trivial.

For an ideal $J \subseteq I$, the J-adic topology clearly refines the I-adic topology, but if also $J \supseteq I^r$ for some $r \geq 0$, then the two topologies on an R-module are the same. When I is finitely generated (which is always so when R is noetherian), then $I \subset (\sqrt{I})^r$ for some r, and so the I-adic topology (and hence the associated completion) does not change when passing to $\sqrt{I}$, so that the natural map $\hat{R}_I \to \hat{R}_{\sqrt{I}}$ resp. $\hat{M}_I \to \hat{M}_{\sqrt{I}}$ is an isomorphism of topological rings resp. modules. Also, for any $n_0 \geq 0$, the collection $\{I^{n+n_0}M\}_{n\geq 0}$ is a neighborhood basis of 0 in M and hence still defines the I-adic topology.

For instance, if in Example 11.6 above, the prime decomposition of n is $n = p_1^{k_1} \cdots p_s^{k_s}$, then $\hat{\mathbb{Z}}_{(n)} = \hat{\mathbb{Z}}_{(p_1p_2\cdots p_s)}$. By the Chinese remainder theorem, the natural map $\mathbb{Z}/((p_1p_2 \cdots p_s)^m) \to \prod_i \mathbb{Z}/(p_i^m)$ is an isomorphism and via these isomorphisms $\hat{\mathbb{Z}}_{(n)}$ is identified with $\prod_i \hat{\mathbb{Z}}_{(p_i)}$.

When M is Hausdorff, then its topology is even metrizable: if $\phi : \mathbb{Z}_+ \to (0, \infty)$ is any function with $\phi(n + 1) \leq \phi(n)$ and $\lim_{n\to\infty} \phi(n) = 0$ (one often takes $\phi(n) = u^{-n}$ for some $u > 1$), then a metric δ on M is defined by

$$\delta(a, a') := \inf\{\phi(n) \,:\, a - a' \in I^nM\}.$$

This metric is *nonarchimedean* in the sense that $\delta(a, a'') \leq \max\{\delta(a, a'), \delta(a', a'')\}$. A sequence $(a_n \in M)_{n=0}^{\infty}$ is then a Cauchy sequence if and only if for every integer $k \geq 0$ all but finitely many terms lie in the same coset of I^kM in M; in other words, there exists an index $n_k \geq 0$ such that $a_m - a_n \in I^kM$ for all $m, n \geq n_k$. This makes it clear that the notion of Cauchy sequence is independent of the choice of ϕ. Such a Cauchy sequence defines a compatible sequence of cosets $(\alpha_n \in M/I^nM)_{n\geq 0}$ and hence an element of $\hat{M}_I$. Recall that a metric space is said to be *complete* if every Cauchy sequence in that space converges. A standard construction produces a completion of every metric space M: its points are represented by Cauchy sequences in M, with the understanding that two such sequences represent the same point if the distance between the two nth terms goes to zero as $n \to \infty$. In the present situation we thus recover $\hat{M}_I$. Note that the homomorphism $M \to \hat{M}_I$ is a continuous injection with image dense in $\hat{M}_I$.

If M is an R-module, then the inclusion $M' \subseteq M$ of any submodule is continuous for the I-adic topology. The Artin-Rees lemma says among other things that

in the noetherian setting this is in fact a closed embedding (so that M' has the induced topology). It is based on the following lemma.

Lemma 11.8. Let R be a noetherian ring and $I \subseteq R$ an ideal. Then the subring $R[It] := \sum_{n=0}^{\infty} I^n t^n$ (where $I^0 := R$) of $R[t]$ is noetherian([13]).

If M is a finitely generated R-module, then $M[It] := \sum_{i=0}^{\infty} I^n M t^n$ is a finitely generated $R[It]$-module.

Any $R[It]$-submodule of $M[It]$ has the form $\sum_{j=0}^{j} N_j t^j$ with $\{N_j\}_{j=0}^{\infty}$ a sequence of R-submodules of M such that $IN_j \subseteq N_{j+1}$ for all j and becomes I-stable, i.e., there exists a j_0 such that this inclusion is an equality for $j \geq j_0$.

PROOF. Since R is noetherian, I has a finite set of generators, say $r_1, \ldots, r_n$. Then the ring homomorphism $R[x_1, \ldots, x_n] \to R[It]$, $x_i \mapsto r_i t$ is onto and it then follows that $R[It]$ is noetherian. The proof that $M[It]$ is finitely generated as a $R[It]$-module is similar: if $m_1, \ldots, m_s$ generate M as a R-module, then $m_1 t, \ldots, m_s t$ generate $M[It]$ as a $R[It]$-module.

It is clear that an $R[It]$-submodule of $M[It]$ has the form $\sum_{j=0}^{\infty} N_j t^j$, where N_j is a R-submodule of M such that $IN_j \subseteq N_{j+1}$ for all n. Since $M[It]$ is noetherian as a $R[It]$-module, $\sum_{j=0}^{\infty} N_j t^j$ has a finite set of $R[It]$-generators, say $a_1 t^{j_1}, \ldots, a_l t^{j_l}$. Then $j_0 := \max_i\{j_i\}$ is as desired. □

Corollary 11.9 (Artin-Rees Lemma)**.** Let R be a noetherian ring, $I \subseteq R$ an ideal, M a finitely generated R-module and $M' \subseteq M$ an R-submodule. Then the sequence $\{M' \cap I^j M\}_{j \geq 0}$ of submodules of M' becomes I-stable: there exists an integer $j_0 \geq 0$ such that $M' \cap I^{j+1} M = I(M' \cap I^j M)$ for $j \geq j_0$.

PROOF. We first observe that $\sum_{j=0}^{\infty}(M' \cap I^j M)t^j$ is an $R[It]$-submodule of $M[It]$. It then remains to apply Lemma 11.8. □

Note that this implies that for every $n \geq 0$, $M' \cap I^{j_0+n} M = I^n(M' \cap I^{j_0} M)$. In particular, $M' \cap I^{j_0+n} M \subseteq I^n M'$. So the inclusion $M' \subseteq M$ is not merely continuous, but M' inherits its I-adic topology from that of M and $\hat{M}'_I$ can be identified with the closure of the image of $M' \subseteq M \to \hat{M}_I$. We will use the Artin-Rees Lemma via this property only.

The first part of the corollary below says that when R is noetherian, I-adic completion is an exact functor on the category of finitely generated R-modules.

Corollary 11.10. In the situation of Corollary 11.9, the homomorphism $\hat{M}'_I \to \hat{M}_I$ induced by the inclusion $M' \subseteq M$ is a closed embedding and $\hat{M}_I / \hat{M}'_I$ can be identified with the I-adic completion of M/M'. In case R is also a local ring and $I \neq R$ (so that by Lemma 11.4 we may regard M resp. M' as a submodule of $\hat{M}_I$ resp. $\hat{M}'_I$), then $M \cap \hat{M}'_I = M'$.

PROOF. We already observed that $\hat{M}'_I \to \hat{M}_I$ is a closed embedding. Consider for $n = 0, 1, 2, \ldots$ the exact sequences

$$0 \to M'/M' \cap I^n M \to M/I^n M \to M/(M' + I^n M) \to 0.$$

Taking the projective limits yields the exactness of $0 \to \hat{M}'_I \to \hat{M}_I \to \widehat{M/M'}_I \to 0$.

The last identity follows from the fact that both sides are equal to the kernel of $M \to \hat{M}_I / \hat{M}'_I = \widehat{M/M'}_I$. □

[13] This is sometimes called the *blowup of I in R* for reasons made clear in Exercise 62.

Let R be a noetherian local ring with maximal ideal $\mathfrak{m}$ and residue field κ. The module $\mathfrak{m}$ is finitely generated and since R acts on $\mathfrak{m}/\mathfrak{m}^2$ via $R/\mathfrak{m} = \kappa$, $\mathfrak{m}/\mathfrak{m}^2$ is a finite dimensional vector space over κ.

DEFINITION 11.11. The *Zariski cotangent space* $T^*(R)$ of R is the κ-vector space $\mathfrak{m}/\mathfrak{m}^2$. The *Zariski tangent space* $T(R)$ of R is its κ-dual, $T(R) := \mathrm{Hom}_\kappa(\mathfrak{m}/\mathfrak{m}^2, \kappa)$ (which is also equal to $\mathrm{Hom}_R(\mathfrak{m}, \kappa)$). The *embedding dimension* $\mathrm{embdim}(R)$ is the dimension of $\mathfrak{m}/\mathfrak{m}^2$ as a vector space over κ.

If X is an affine variety and $p \in X$, then we define the Zariski cotangent space T_p^*X, the Zariski tangent space T_pX and the embedding dimension $\mathrm{embdim}_p X$ of X at p to be that of $\mathcal{O}_{X,p}$.

For instance, the embedding dimension of $\mathbb{A}^n$ at any point $p \in \mathbb{A}^n$ is n. This follows from the fact that the map $d_p : f \in \mathfrak{m}_{\mathbb{A}^n,p} \mapsto df(p) \in k^n$ defines an isomorphism of k-vector spaces $\mathfrak{m}_{\mathbb{A}^n,p}/\mathfrak{m}_{\mathbb{A}^n,p}^2 \cong k^n$. We note in passing that we here have a way of understanding the total differential at $p \in \mathbb{A}^n$ in more intrinsic terms as the map $d_p : \mathcal{O}_{\mathbb{A}^n,p} \to \mathfrak{m}_{\mathbb{A}^n,p}/\mathfrak{m}_{\mathbb{A}^n,p}^2$ which assigns to $f \in \mathcal{O}_{\mathbb{A}^n,p}$ the image of $f - f(p) \in \mathfrak{m}_{\mathbb{A}^n,p}$ in $\mathfrak{m}_{\mathbb{A}^n,p}/\mathfrak{m}_{\mathbb{A}^n,p}^2$. Thus, a differential of f at p can be understood as a k-linear function $df(p) : T_p\mathbb{A}^n \to k$ and $(d_p(x_i) = dx_i(p))_{i=1}^n$ is a basis of $\mathfrak{m}_{\mathbb{A}^n,p}/\mathfrak{m}_{\mathbb{A}^n,p}^2 = T_p^*\mathbb{A}^n$ whose dual basis (of $T_p\mathbb{A}^n$) is represented by $(\partial/\partial x_i\big|_p)_{i=1}^n$.

Let us observe that since embedding dimension and Zariski tangent space of a local ring R are defined in terms of $\mathfrak{m}/\mathfrak{m}^2$, these notions only depend on $\hat{R}_\mathfrak{m}$.

EXERCISE 44. Let $(R', \mathfrak{m}')$ and $(R, \mathfrak{m})$ be local rings with residue fields κ resp. κ' and let $\phi : R' \to R$ be a ring homomorphism with the property that $\phi^{-1}\mathfrak{m} = \mathfrak{m}'$ (we then say that ϕ is a *local homomorphism*). Prove that ϕ induces a field embedding $\kappa' \hookrightarrow \kappa$ and a linear map of κ-vector spaces $T(\phi) : T(R) \to \kappa \otimes_{\kappa'} T(R')$.

An application of Nakayama's lemma to the R-module $\mathfrak{m}$ yields:

Corollary 11.12. The embedding dimension of a noetherian local ring R is the smallest number of generators of its maximal ideal. This is zero if and only if R is a field.

DEFINITION 11.13. A noetherian local ring R is said to be *regular* if its Krull dimension equals its embedding dimension. A point p of an affine variety X is called *smooth* if its local ring $\mathcal{O}_{X,p}$ is so; otherwise it is called *singular*. The corresponding subsets of X are called the *smooth locus* resp. *singular locus* of X and will be denoted X_{reg} resp. X_{sing}. An affine variety without singular points is said to be *smooth*.

We shall see that the regularity of a local ring $\mathcal{O}_{X,p}$ indeed amounts to X being 'like a manifold' at p. We begin with a formal version of the implicit function theorem.

Lemma 11.14. Let $p \in \mathbb{A}^n$ and let $f_1, \ldots, f_{n-d} \in \mathfrak{m}_{\mathbb{A}^n,p}$ be such that the differentials $df_1(p), \ldots, df_{n-d}(p)$ are linearly independent.

Then $\mathfrak{p} := (f_1, \ldots, f_{n-d}) \subseteq \mathcal{O}_{\mathbb{A}^n,p}$ is a prime ideal and $\mathcal{O}_{\mathbb{A}^n,p}/\mathfrak{p}$ is a regular local ring of dimension d whose completion with respect to its maximal ideal is isomorphic to $k[[y_1, \ldots, y_d]]$ as a complete local k-algebra and whose Zariski tangent space, regarded as a subspace of the Zariski tangent space $T_p\mathbb{A}^n$ of p in $\mathbb{A}^n$, is the kernel of the linear surjection $(df_1(p), \ldots, df_{n-d}(p)) : T_p\mathbb{A}^n \to k^{n-d}$.

There exists an affine neighborhood U of p in $\mathbb{A}^n$ on which $f_1, \dots, f_{n-d}$ are regular and generate in $k[U]$ a prime ideal whose zero set is smooth of dimension d.

PROOF. Let us abbreviate $\mathcal{O}_{\mathbb{A}^n,p}$ by $\mathcal{O}$ and its maximal ideal $\mathfrak{m}_{\mathbb{A}^n,p}$ by $\mathfrak{m}$. Extend $f_1, \dots, f_{n-d}$ to a system of regular functions $f_1, \dots, f_n \in \mathfrak{m}$ such that the $df_1(p), \dots, df_n(p)$ are linearly independent. This means that their images in $\mathfrak{m}/\mathfrak{m}^2$ are linearly independent over k. In particular, $f_{n-d+1}, \dots, f_n$ map to a k-basis of $\mathfrak{m}/(\mathfrak{m}^2 + (f_1, \dots, f_{n-d}))$ so that d is the embedding dimension of $\mathcal{O}/\mathfrak{p}$. After an affine-linear transformation, we then may (and will) assume that p is the origin o of $\mathbb{A}^n$ and that $f_i \equiv x_i \pmod{\mathfrak{m}^2}$. Hence the monomials of degree r in $f_1, \cdots, f_n$ map to a k-basis of $\mathfrak{m}^r/\mathfrak{m}^{r+1}$. With induction on r it then follows that the monomials of degree $\leq r$ in $f_1, \cdots, f_n$ make up a k-basis of $\mathcal{O}/\mathfrak{m}^{r+1}$. This amounts to the assertion that the map

$$y_i \in k[[y_1, \dots, y_n]] \mapsto f_i \in \hat{\mathcal{O}}$$

defines an isomorphism $k[[y_1, \dots, y_n]] \cong \hat{\mathcal{O}}$ of complete local rings (a ring isomorphism that is also a homeomorphism). The restriction of its inverse to $\mathcal{O}$ is a topological embedding of $\mathcal{O}$ in $k[[y_1, \dots, y_n]]$ (sending f_i to y_i). Denote by $\mathfrak{p}_i \subset \mathcal{O}$ the ideal generated by $f_1, \dots, f_i$ (so that $\mathfrak{p}_{n-d} = \mathfrak{p}$). This inverse sends the closure $\bar{\mathfrak{p}}_i$ of $\mathfrak{p}_i$ in $\hat{\mathcal{O}}$ to the closure in $k[[y_1, \dots, y_n]]$ of the ideal generated by $y_1, \dots, y_i$. The quotient ring of the latter is the domain $k[[y_{i+1}, \dots, y_n]]$. According to Corollary 11.10, the preimage of $\bar{\mathfrak{p}}_i$ in $\mathcal{O}$ is $\mathfrak{p}_i$ (so that $\mathfrak{p}_i$ is also a prime ideal) and the embedding $\mathcal{O}/\mathfrak{p}_i \hookrightarrow k[[y_1, \dots, y_n]]/(y_1, \dots, y_{n-d}) = k[[y_{i+1}, \dots, y_n]]$ realizes the $\mathfrak{m}$-adic completion of $\mathcal{O}/\mathfrak{p}_i$. So $(0) = \mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \cdots \subsetneq \mathfrak{p}_n$ is a prime chain in $\mathcal{O}$ of length n. By Corollary 10.8, $\mathcal{O}$ has Krull dimension n, so that his prime chain is maximal. Theorem 10.6 then implies that $\mathcal{O}/\mathfrak{p} = \mathcal{O}/\mathfrak{p}_{n-d}$ has Krull dimension d. As this is also the embedding dimension of $\mathcal{O}/\mathfrak{p}$, $\mathcal{O}/\mathfrak{p}$ is a regular local ring of dimension d.

Since the $df_1(p), \dots, df_{n-d}(p)$ are linearly independent, there exist $n-d$ indices $1 \leq \nu_1 < \nu_2 < \cdots < \nu_{n-d} \leq n$ such that $\delta := \det((\partial f_i/\partial x_{\nu_j})_{i,j})$ is nonzero in p. Let $g' \in k[x_1, \dots, x_n]$ be a common denominator for $f_1, \dots, f_{n-d}$ and δ with $g(o) \neq 0$. This ensures that $df_1, \dots, df_{n-d}$ are linearly independent everwhere on $\mathbb{A}^n_g$. The preimage $P' \subseteq k[x_1, \dots, x_n][1/g']$ of $\mathfrak{p}$ under the localization map $k[x_1, \dots, x_n][1/g'] \to \mathcal{O}$ is a prime ideal which contains (the images of) $f_1, \dots, f_{n-d}$ and has the property that its localization at $o \in \mathbb{A}^n$ is $\mathfrak{p}$. Choose a finite set of generators $\phi_1, \dots, \phi_r$ of P'. Then in $\mathcal{O}$ we have $\phi_i = \sum_{j=1}^{n-d} u_{ij} f_j$ for certain $u_{ij} \in \mathcal{O}$. Let $g \in (g')$ with $g(o) \neq 0$ be a common denominator for the u_{ij} and put $U = \mathbb{A}^n_g$. Then $P := P'[1/g]$ is a prime ideal in $k[U]$. It is generated by $\phi_1|U, \dots, \phi_r|U$ and hence also by $f_1|U, \dots, f_{n-d}|U$. So $(U; f_1|U, \dots, f_{n-d}|U)$ is as desired. □

Theorem 11.15. Let $X \subseteq \mathbb{A}^n$ be locally closed and let $p \in X$. Then the local ring $\mathcal{O}_{X,p}$ is regular of dimension d if and only if there exist a set of generators $f_1, \dots, f_{n-d}$ of $\mathcal{I}_{X,p}$ such that $df_1(p), \dots, df_{n-d}(p)$ are linearly independent.

The set of $q \in X$ for which these equivalent conditions hold is open in X; in particular, the smooth locus X_{reg} of X is open in X.

PROOF. One direction is immediate from Lemma 11.14, for it tells us that if $f_1, \dots, f_{n-d}$ are elements of $\mathfrak{m}_{\mathbb{A}^n,p}$ such that $df_1(p), \dots, df_{n-d}(p)$ are linearly independent, then the ideal $\mathcal{I}_p \subset \mathcal{O}_{\mathbb{A}^n,p}$ generated by them is a prime ideal and $\mathcal{O}_{\mathbb{A}^n,p}/\mathcal{I}_p$ is regular of dimension d.

For the converse, suppose that $\mathcal{O}_{X,p}$ is regular of dimension d. Let $\mathcal{I}_{X,p} \subseteq \mathcal{O}_{\mathbb{A}^n,p}$ be the ideal of regular functions at p vanishing on a neighborhood of p in X, in other words, the kernel of $\mathcal{O}_{\mathbb{A}^n,p} \to \mathcal{O}_{X,p}$. We have a short exact sequence

$$0 \to (\mathcal{I}_{X,p} + \mathfrak{m}^2_{\mathbb{A}^n,p})/\mathfrak{m}^2_{\mathbb{A}^n,p} \to \mathfrak{m}_{\mathbb{A}^n,p}/\mathfrak{m}^2_{\mathbb{A}^n,p} \to \mathfrak{m}_{X,p}/\mathfrak{m}^2_{X,p} \to 0.$$

Since the middle term has k-dimension n and $\dim_k(\mathfrak{m}_{X,p}/\mathfrak{m}^2_{X,p}) = d$ by assumption, $(\mathcal{I}_{X,p} + \mathfrak{m}^2_{\mathbb{A}^n,p})/\mathfrak{m}^2_{\mathbb{A}^n,p} \cong \mathcal{I}_{X,p}/(\mathcal{I}_{X,p} \cap \mathfrak{m}^2_{\mathbb{A}^n,p})$ must have k-dimension $n-d$. Let $f_1, \dots, f_{n-d} \in \mathcal{I}_{X,p}$ map to k-basis of $\mathcal{I}_{X,p}/(\mathcal{I}_{X,p} \cap \mathfrak{m}^2_{\mathbb{A}^n,p})$. This means that $df_1(p), \dots, df_{n-d}(p)$ are linearly independent. It suffices to show that $f_1, \dots, f_{n-d}$ generate $\mathcal{I}_{X,p}$.

According to Lemma 11.14, the ideal $\mathfrak{p}_i \subseteq \mathcal{O}_{\mathbb{A}^n,p}$ generated by $f_1, \dots, f_i$ is prime and so we have a prime chain

$$(0) = \mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \cdots \subsetneq \mathfrak{p}_{n-d} \subseteq \mathcal{I}_{X,p}.$$

Since $\dim \mathcal{O}_{X,p} = d$, there also exists a prime chain of length d containing $\mathcal{I}_{X,p}$:

$$\mathcal{I}_{X,p} \subseteq \mathfrak{q}_0 \subsetneq \mathfrak{q}_1 \subsetneq \cdots \subsetneq \mathfrak{q}_d \subseteq \mathcal{O}_{\mathbb{A}^n,p}.$$

As $\mathcal{O}_{\mathbb{A}^n,p}$ has dimension n, these two prime chains cannot make up a prime chain of length $n+1$ and so $\mathfrak{p}_{n-d} = \mathcal{I}_{X,p} = \mathfrak{q}_0$. In particular, $f_1, \dots, f_{n-d}$ generate $\mathcal{I}_{X,p}$.

For the last assertion follows from the last clause of Lemma 11.14, since it shows that if p is a regular point of X of dimension p, then p has in X a neigborhood of such points. □

The criterion of Theorem 11.15 is easily amplified to a more abstract setting:

Corollary 11.16. Suppose that in the situation of Theorem 11.15, X is regular at p and that we are given a closed subset $Y \subset X$ with $p \in Y$. Then Y is regular of dimension d' at p if and only if there exist $g_1, \dots, g_{d-d'} \in \mathfrak{m}_{X,p}$ which define Y in X at p and whose differentials define linearly independent forms on T_pX. In that case $T_pY \subseteq T_pX$ is the common kernel of these linear forms.

PROOF. Choose a lift $\tilde{g}_i \in \mathfrak{m}_{\mathbb{A}^n,p}$ of g_i and apply Theorem 11.15 to Y and the ideal generated by $\tilde{g}_1, \dots, \tilde{g}_{d-d'}, f_1, \dots, f_{n-d}$. □

Proposition 11.17. The smooth locus X_{reg} an affine variety X is dense in X.

PROOF. Without loss of generality we may assume that X is irreducible. Since we already know that X_{reg} is open, it remains to see that it is nonempty. It thus becomes an issue which only depends on $k(X)$. In view of Corollary 9.5 it then suffices to treat the case of a hypersurface in $\mathbb{A}^{r+1}$ so that $I(X)$ is generated by an irreducible polynomial $f \in k[x_1, \dots, x_{r+1}]$. In view of Lemma Lemma 11.14 it then suffices to show that df is not identically zero on X. Suppose otherwise, i.e., that each partial derivative $\partial f/\partial x_i$ vanishes on X. Then each $\partial f/\partial x_i$ must be multiple of f and since the degree of $\partial f/\partial x_i$ is less than that of f, this implies that it is identically zero. But then we know from Exercise 42 that the characteristic p of k must then be positive (so ≥ 2) and that f is of the form g^p. This contradicts the fact that f is irreducible. □

EXERCISE 45. Let X be a smooth variety. Prove that X is connected if and only if it is irreducible.

REMARK 11.18. This enables us to find for an affine variety X of dimension d (with downward induction) a descending chain of closed subsets $X = X^d \supseteq X^{d-1} \supseteq \cdots \supseteq X^0$ such that $\dim X^i \leq i$ and all the (finitely many) connected components of $X^i \smallsetminus X^{i-1}$ are smooth subvarieties of dimension i: if X^i has been defined, then take for X^{i-1} the union of the singular locus X^i_{reg} and the irreducible components of dimension $\leq i-1$. Then $\dim X^{i-1} \leq i-1$ and every connected component of $X^i \smallsetminus X^{i-1}$ is an nonempty open subset of some X^i_{reg} and hence a smooth subvariety of dimension i.

EXERCISE 46. Prove that an affine variety X admits a partition $\mathcal{S}$ into locally closed connected smooth subvarieties such that the closure of every $S \in \mathcal{S}$ is a union of members of $\mathcal{S}$ (such a partition is called a *stratification* of X and its members *strata*).

We now show that normalization converts an irreducible curve into a smooth one. This will follow from:

Proposition-definition 11.19. For a local noetherian domain R the following are equivalent:

(i) R is normal and of dimension one,
(ii) R is regular and of dimension one,
(iii) R is a not a field and every proper ideal of R of the form $\mathfrak{m}^n$ for some $n > 0$, where $\mathfrak{m}$ is the maximal ideal.

If these equivalent conditions are fulfilled, we say that R is a *discrete valuation ring (abbreviated as DVR)*. In that case is R is a PID.

PROOF. In what follows we denote by κ the residue field $R/\mathfrak{m}$.

$(i) \Rightarrow (ii)$. Assume R is normal and of dimension 1. We must show that $\dim_\kappa \mathfrak{m}/\mathfrak{m}^2 = 1$. Since $\dim R = 1$, $(0) \subsetneq \mathfrak{m}$ is a maximal prime chain. So R has no other prime ideals than these two. Let $a \in \mathfrak{m} \smallsetminus \{0\}$. Then $\sqrt{Ra}$ is an intersection of prime ideals and hence must be equal to $\mathfrak{m}$. Let $n \geq 1$ be minimal for the property that $\mathfrak{m}^n \subseteq Ra$, so that $\mathfrak{m}^{n-1} \not\subseteq Ra$. Choose $b \in \mathfrak{m}^{n-1} \smallsetminus Ra$ and put $y := b/a \in \mathrm{Frac}(R)$. Then $\mathfrak{m}y \subseteq R$ (for $\mathfrak{m}b \subset Ra$), but $y \notin R$. In particular, y is not integral over R and hence multiplication with y cannot preserve any (finitely generated) R-submodule of R. So $\mathfrak{m}y \not\subseteq \mathfrak{m}$. By the maximality of $\mathfrak{m}$, it then follows that $\mathfrak{m}y = R$. This proves that $\pi := y^{-1}$ is a generator of $\mathfrak{m}$. Multiplication by π defines a surjection $\kappa \twoheadrightarrow R\pi/R\pi^2 \cong \mathfrak{m}/\mathfrak{m}^2$. This is in fact an isomorphism (so that R is indeed regular and of dimension one): otherwise we would have $\mathfrak{m}^2 = \mathfrak{m}$. But this would imply $\mathfrak{m} = 0$ by Nakayama's lemma and so $R = \kappa$, which contradicts the assumption that R has dimension 1.

$(ii) \Rightarrow (iii)$. Assume R is regular of dimension 1 so that $\dim_\kappa \mathfrak{m}/\mathfrak{m}^2 = 1$. Let $\pi \in \mathfrak{m} \smallsetminus \mathfrak{m}^2$. Then π will map to a generator of $\mathfrak{m}/\mathfrak{m}^2$ and will hence generate $\mathfrak{m}$. Then for every $r \geq 1$, π^r generates $\mathfrak{m}^r$ so that $\dim_\kappa \mathfrak{m}^r/\mathfrak{m}^{r+1} \leq 1$.

We next prove that every $a \in \mathfrak{m} \smallsetminus \{0\}$ generates a power of $\mathfrak{m}$. By Lemma 11.4, $\cap_n \mathfrak{m}^n = \{0\}$ and so there exists an $n \geq 0$ such that $a \in \mathfrak{m}^n \smallsetminus \mathfrak{m}^{n+1}$. Since $\dim_\kappa \mathfrak{m}^n/\mathfrak{m}^{n+1} \leq 1$, it must be generated by the image of a and Nakayama's lemma then implies that $Ra = \mathfrak{m}^n$.

Now let $I \subset \mathfrak{m}$ be a proper ideal of R. Each nonzero member of I generates some positive power of $\mathfrak{m}$, so if r is the smallest such power which thus occurs, then $I = \mathfrak{m}^r$.

$(iii) \Rightarrow (i)$. Assume R is as in (iii). Then $(0) \subsetneq \mathfrak{m}$ is a maximal prime chain, so that $\dim R = 1$. Since every proper ideal of R is of the form $\mathfrak{m}^n$ for a unique n, R is also a UFD and hence normal by Lemma 9.8. □

REMARK 11.20. If R is a DVR, then a generator of $\mathfrak{m}$ is called a *uniformizer* of R. By (iii) every proper ideal is then equal to Rt^n for a unique $n \geq 0$. This implies that every nonzero R-submodule of $K := \mathrm{Frac}(R)$ is of the form Rt^n for a unique $n \in \mathbb{Z}$. We denote this submodule also by $\mathfrak{m}^n$. Note that the map $v : K^\times \to \mathbb{Z}$ which takes the value n on $\mathfrak{m}^n \smallsetminus \mathfrak{m}^{n+1}$ is a surjective homomorphism with kernel $R^\times = R \smallsetminus \mathfrak{m}$. The obvious inequality $v(r+r') \geq \min\{v(r), v(r')\}$ makes it a *discrete (nonrarchimedean) valuation* (this explains the name DVR).

Corollary 11.21. Let C be an irreducible curve. Then $p \in C$ is a normal point of C if and only if it is a regular point of C. This is also equivalent to $\mathcal{O}_{C,p}$ being a discrete valuation ring. □

So for an irreducible curve C, the normalization $\pi : \hat{C} \to C$ is a 'desingularization' in the sense that it provides us with a finite birational morphism whose domain $\hat{C}$ is smooth. This also shows that a finitely generated field extension of k of transcendence degree 1 is k-isomorphic to the function field of a smooth irreducible curve.

REMARK 11.22. Recall that a Dedekind domain is a normal noetherian domain of dimension one. The above proposition implies that when R is a Dedekind domain R and $\mathfrak{p} \subset R$ is a nonzero prime ideal, then $\mathfrak{p}$ is maximal and $R_\mathfrak{p}$ is a discrete valuation ring. Using this, one can show that every nonzero ideal $I \subset R$ has a unique 'prime decomposition': there exists a unique map $[\mathfrak{p}] \in \mathrm{Spm}(R) \mapsto n_\mathfrak{p} \in \mathbb{Z}$ with finite support such that $I = \prod_{[\mathfrak{p}] \in \mathrm{Spm}(R)} \mathfrak{p}^{n_\mathfrak{p}}$. But beware that a maximal ideal of a Dedekind domain need not be principal (and hence the domain need not be an UFD). Examples occur in number theory ('most' rings of integers are not UFD's) and in algebraic geometry: for 'most' affine smooth curves C the set of $p \in C$ such that $I_C(\{p\})$ is *not* generated by a single element is dense in C.

12. Differentials and derivations.

The differential that we defined earlier has an intrinsic, coordinate free description that turns out to be quite useful. We review this briefly. Let us begin with the observation that the formation of the total differential of a polynomial, $\phi \in k[x_1, \ldots, x_n] \mapsto d\phi := \sum_{i=1}^n (\partial\phi/\partial x_i)(p) dx_i$ is a k-linear map which satisfies the Leibniz rule: $d(\phi\psi) = \phi d\psi + \psi d\phi$. This property is formalized with the following definition. Fix a ring R (the *base ring*) and an R-algebra A.

DEFINITION 12.1. Let M be a A-module. An *R-derivation of A with values in M* is an R-module homomorphism $D : A \to M$ which satisfies the Leibniz rule: $D(a_1 a_2) = a_1 D(a_2) + a_2 D(a_2)$ for all $a_1, a_2 \in A$.

The last condition usually prevents D from being an A-module homomorphism. Let us observe that (by taking $a_1 = a_2 = 1$) we must have $D(1) = 0$. Since D is R-linear, it then follows that for every $r \in R$, $D(r) = rD(1) = 0$. Note also that if $b \in A$ happens to be invertible in A, then $0 = D(1) = D(b/b) = D(b)/b + bD(1/b)$ so that $D(1/b) = -D(b)/b^2$ and hence $D(a/b) = \big(D(a)b - aD(b)\big)/b^2$ for every $a \in A$.

Given $a_1, \ldots, a_n \in A$, then the values of D on $a_1, \ldots, a_n$ determine its values on the subalgebra A' by the a_i's, for if $\phi : R[x_1, \ldots, x_n] \to A$ denotes the corresponding R-homomorphism and $f \in R[x_1, \ldots, x_n]$, then

$$D\phi(f) = \sum_{i=1}^{n} \phi\Big(\frac{\partial f}{\partial x_i}\Big) Da_i.$$

If we combine this with the formula for $D(1/b)$, we see that this not only determines D on the R-subalgebra A' of A generated by the a_i's, but also on the biggest localization of A' contained in A. In particular, if we are given a field extension L/K, then a K-derivation of L with values in some L-vector space is determined by its values on a set of generators of L as a field extension of K.

The set of R-derivations of A in M form an R-module: if D_1 and D_2 are R-derivations of A with values in M, and $a_1, a_2 \in A$, then $a_1 D_1 + a_2 D_2$ is also one. We denote this module by $\mathrm{Der}_R(A, M)$.

EXERCISE 47. Prove that if $D_1, D_2 \in \mathrm{Der}_R(A, A)$, then $[D_1, D_2] := D_1 D_2 - D_2 D_1 \in \mathrm{Der}_R(A, A)$. What do we get for $R = k$ and $A = k[x_1, \ldots, x_n]$?

It is immediate from the definition that for every A-module homomorphism $\phi : M \to N$ the composition of a D as above with ϕ is an R-derivation of A with values in N. We can now construct a universal R-derivation of A, $d : A \to \Omega_{A/R}$ (where $\Omega_{A/R}$ must of course be an R-module) with the property that every D as above is obtained by composing d with a unique homomorphism of A-modules $\bar{D} : \Omega_{A/R} \to N$. The construction that is forced upon us starts with the free A-module $A^{(A)}$ which has A itself as a generating set—let us denote the generator associated to $a \in A$ by $\tilde{d}(a)$—which we then divide out by the A-submodule of $A^{(A)}$ generated by the expressions $\tilde{d}(ra) - r\tilde{d}(a)$, $\tilde{d}(a_1 + a_2) - \tilde{d}(a_1) - \tilde{d}(a_2)$ and $\tilde{d}(a_1 a_2) - a_1 \tilde{d}(a_2) - a_2 \tilde{d}(a_2)$, with $r \in R$ and $a, a_1, a_2 \in A$. The quotient A-module is denoted $\Omega_{A/R}$ and the composite of $\tilde{d}$ with the quotient map by $d : A \to \Omega_{A/R}$. The latter is an R-derivation of A by construction. Given an R-derivation $D : A \to M$, then the map which assigns to $\tilde{d}(a)$ the value Da extends (obviously) as an A-module homomorphism $A^{(A)} \to M$. It has the above submodule in its kernel and hence determines an A-module homomorphism of $\bar{D} : \Omega_{A/R} \to M$. This has clearly the property that $D = \bar{D}d$. In other words, composition with d defines an isomorphism of A-modules $\mathrm{Hom}_A(\Omega_{A/R}, M) \xrightarrow{\cong} \mathrm{Der}_R(A, M)$. We call $\Omega_{A/R}$ the module of *Kähler differentials*. We shall see that the map $d : A \to \Omega_{A/R}$ can be thought of as an algebraic version of the formation of the (total) differential.

The universal derivation of a finitely generated R-algebra A can be constructed in a more direct manner as follows. We first do the case when A is a polynomial algebra $P := R[x_1, \ldots, x_n]$. For any R-derivation $D : P \to M$ we have $Df = \sum_{i=1}^{n} (\partial f/\partial x_i) Dx_i$ and this yields $(Dx_1, \ldots, Dx_n) \in M^n$. Conversely, for any n-tuple $(m_1, \ldots, m_n) \in M^n$, we have an R-derivation $D : P \to M$ defined by $Df = \sum_{i=1}^{n} (\partial f/\partial x_i) m_i$. So Dx_i can be prescribed arbitrarily as an element of M. But to give an element of M^n amounts to giving a P-homomorphism $P^n \to M$ and hence $\Omega_{P/R}$ is the free P-module generated by $dx_1, \ldots, dx_n$. Thus the universal R-derivation $d : P \to \Omega_{P/R}$, which is given by $f \mapsto \sum_{i=1}^{n} (\partial f/\partial x_i) dx_i$, may be regarded as the intrinsic way of forming the total differential.

Next consider a quotient $A := P/I$ of P, where $I \subseteq P$ is an ideal. If M is an A-module and $D' : A \to M$ is an R-derivation, then its composite with the projection $\pi : P \to A$, $D = D'\pi : P \to M$, is an R-derivation of P with the property that $Df = 0$ for every $f \in I$. Conversely, every R-derivation $D : P \to M$ in an A-module M which is zero on I factors through an an R-derivation $D' : A \to M$. Note that for *any* R-derivation $D : P \to M$, its restriction to I^2 is zero, for if $f, g \in I$, then $D(fg) = fDg + gDf \in IM = \{0\}$. Now I/I^2 is a module over $P/I = A$ and so we obtain a short exact sequence of A-modules

$$I/I^2 \to \Omega_{P/R}/I\Omega_{P/R} \to \Omega_{A/R} \to 0.$$

It follows from our computation of $\Omega_{P/R}$ that the middle term is the free A-module generated by $dx_1, \ldots, dx_n$. So if I is generated by $f_1, \ldots, f_m$, then $\Omega_{A/R}$ can be identified with the quotient of $\sum_{i=1}^n A dx_i$ by the A-submodule generated by the A-submodule generated by the $df_j = \sum_{i=1}^n (\partial f_j/\partial x_i) dx_i$, $j = 1, \ldots, m$.

Note that if R is a noetherian ring, then so is A (by the Hilbert basis theorem) and since $\Omega_{A/R}$ is a finitely generated A-module, it is noetherian as an A-module. This applies for instance to the case when $R = k$ and $A = k[X]$ for some affine variety X. We then write $\Omega(X)$ for $\Omega_{k[X]/k}$.

EXERCISE 48. Prove that $\Omega_{A/R}$ behaves well under localization: if $S \subseteq A$ is a multiplicative subset, then every R-derivation with values is some A-module M extends naturally to an R-derivation of $S^{-1}A$ with values in $S^{-1}M$. Prove that we have a natural map $S^{-1}\Omega_{A/R} \to \Omega_{S^{-1}A/R}$ and that this map is a A-homomorphism.

For an affine variety X and $p \in X$, we write $\Omega_{X,p}$ for $\Omega_{\mathcal{O}_{X,p}/k}$. The preceding exercise implies that $\Omega_{X,p}$ is the localization of $\Omega(X)$ at p: $\Omega_{X,p} = (k[X] \smallsetminus \mathfrak{p}_p)^{-1}\Omega(X)$.

EXERCISE 49. Let X be an affine variety and let $p \in X$. Show that the Zariski tangent space of X at p can be understood (and indeed, be defined) as the space of k-derivations of $\mathcal{O}_{X,p}$ with values in k, where we regard k as a $\mathcal{O}_{X,p}$-module via $\mathcal{O}_{X,p}/\mathfrak{m}_{X,p} \cong k$. Prove that this identifies its dual, the Zariski cotangent space, with $\Omega_{X,p}/\mathfrak{m}_{X,p}\Omega_{X,p}$.

EXERCISE 50. Show (perhaps with the help of the preceding exercises) that if $p \in \mathbb{A}^n$, then $\Omega_{\mathcal{O}_{\mathbb{A}^n,p}/k}$ is the free $\mathcal{O}_{\mathbb{A}^n,p}$-module generated by $dx_1, \ldots, dx_n$ and that if X is an affine variety in $\mathbb{A}^n$ that has p a regular point, then $\Omega_{X,p}$ is a free $\mathcal{O}_{X,p}$-module of rank $\dim \mathcal{O}_{X,p}$.

We must be careful with this construction when dealing with formal power series rings. For instance, as we have seen, the completion of the local k-algebra $\mathcal{O}_{\mathbb{A}^1,0}$ with respect to its maximal ideal is $k[[x]]$ and the embedding of $\mathcal{O}_{\mathbb{A}^1,0} \hookrightarrow k[[x]]$ is given by Taylor expansion. A k-derivation D of $k[[x]]$ with values in some $k[[x]]$-module is *not* determined by Dx. All it determines are the values on the k-subalgebra $\mathcal{O}_{\mathbb{A}^1,0}$. This issue disappears however if we require that D is continuous for the (x)-adic topology, for this then means that D commutes with (infinite) formal series summation: $D(\sum_{r=0}^\infty c_r x^r) = \sum_{r=0}^\infty c_r r x^{r-1} Dx$.

We obtained $\Omega_{A/R}$ as a quotient of $A^{(A)}$, but the final result made it clear that $\Omega_{A/R}$ is in fact a quotient of $A \otimes_R A$ via $a \otimes_R b \mapsto adb$. The following exercise shows how this leads to an alternative construction of the universal derivation.

EXERCISE 51. Let R be a ring and A an R-algebra. The maps $(a,b) \in A \times A \mapsto ab \in A$, resp. $(a,b) \in A \times A \mapsto adb \in \Omega_{A/R}$ are both R-bilinear and hence factor uniquely through R-linear maps $\mu : A \otimes_R A \to A$ resp. $\delta : A \otimes_R A \to \Omega_{A/R}$. We regard $A \otimes_R A$ as an R-algebra and also as an A-module with A acting by multiplication on the first factor. This makes both μ and δ A-homomorphisms.

(a) Denote by $I \subset A \otimes_R A$ the kernel of μ. Prove that I is as an A-submodule and that it is as such generated by $\{a \otimes_R 1 - 1 \otimes_R a\}_{a \in A}$
(b) Prove that δ maps I onto $\Omega_{A/R}$, but is zero on I^2 so that we have a surjection $\delta' : I/I^2 \to \Omega_{A/R}$ of A-modules.
(c) Prove that $D : a \in A \mapsto a \otimes_R 1 - 1 \otimes_R a + I^2 \in I/I^2$ is an R-derivation. Denote by $\bar{D} : \Omega_{A/R} \to I/I^2$ the associated A-homomorphism that comes from the universal property of $d : A \to \Omega_{A/R}$.
(d) Show that $\bar{D}\delta'$ is the identity of I/I^2. Conclude that D is a universal derivation.
(e) Check that the other A-module structure on $A \otimes_R A$ (obtained by letting A act by multiplication on the second factor) yields the same A-module structure on I/I^2.

Exercise 51 shows that for an affine variety X, $\Omega(X)$ can be identified with the $k[X]$-module $I(\Delta_X)/I(\Delta_X)^2$, where $\Delta_X \subseteq X \times X$ is the diagonal and $I(\Delta_X) \subseteq k[X \times X] = k[X] \otimes k[X]$ the ideal that defines it, and the $k[X]$-module structure comes from (say) the first projection $X \times X \to X$.

13. The notion of a variety

We begin with a 'predefinition'.

DEFINITION 13.1. A *prevariety* is a topological space X endowed with a sheaf $\mathcal{O}_X$ of k-valued functions such that X can be covered by *finitely many* open subsets U such that $(U, \mathcal{O}_X|U)$ is an affine variety. Given prevarieties $(X, \mathcal{O}_X)$ and $(Y, \mathcal{O}_Y)$, then a *morphism of prevarieties* $f : (X, \mathcal{O}_X) \to (Y, \mathcal{O}_Y)$ is simply a morphism in the category of spaces endowed with a sheaf $\mathcal{O}_X$ of k-valued functions: f is continuous and for every open $V \subseteq Y$, composition with f takes $\mathcal{O}_Y(V)$ to $\mathcal{O}_X(f^{-1}V)$.

We often designate a prevariety and its underlying topological space by the same symbol, a habit which rarely leads to confusion. The composite of two morphisms is evidently a morphism so that we have a category. The prefix '*pre*' in prevariety refers to the fact that we have not imposed a separation requirement which takes the place of the Hausdorff property that one normally imposes on a manifold (see Example 13.2 below).

Let X be a prevariety. By assumption X is covered by finitely many affine open subvarieties $\{U_i\}_{i \in I}$ (I is a finite index set). Suppose κ_i is an isomorphism of U_i onto an affine variety X_i which is given as a closed subset in some $\mathbb{A}^{n_i}$. Then $X_{i,j} := \kappa_i(U_i \cap U_j)$ is an open subset of X_i and $\kappa_{i,j} := \kappa_j \kappa_i^{-1}$ is an isomorphism of $X_{i,j}$ onto $X_{j,i} \subseteq X_j$. We can recover X from the disjoint union $\coprod_{i \in I} X_i$ by means of a gluing process, for if we use $\kappa_{i,j}$ to identify $X_{i,j}$ with $X_{j,i}$ for all i, j we get back X. The collection $\{(U_i, \kappa_i)\}_{i \in I}$ is called an *affine atlas* for X and $\kappa_{i,j}$ is called a *transition map*.

EXERCISE 52. Let $(X, \mathcal{O}_X)$ be a prevariety.

(a) Prove that X is a noetherian space.

(b) Prove that X contains an open-dense subset which is affine.
(c) Let $Y \subseteq X$ be *locally closed* (i.e., the intersection of a closed subset with an open subset). Prove that Y is in natural manner a prevariety in such a manner that the inclusion $Y \subseteq X$ is a morphism of prevarieties.

Much of what we did for affine varieties extends in a straightforward manner to this more general context. Here are some examples.

Rational functions. A *rational function* $f : X \dashrightarrow k$ is defined as before: it is represented by a regular function on a subset of X that is open-dense in its set of closed points and two such represent the same rational function if they coincide on a nonempty open-dense subset in their common domain of definition.

Function field and dimension. When X is irreducible, the rational functions on X form a field $k(X)$, the *function field of* X and for an open nonempty affine open subset $U \subseteq X$, we have $k(X) = k(U) = \mathrm{Frac}(\mathcal{O}(U))$ (but we will see that it is not true in general that $k(X) = \mathrm{Frac}(\mathcal{O}(X))$). In particular, any nonempty affine open subset of X has dimension $\mathrm{trdeg}_k\, k(X)$. According to Exercise 37, this is then also the (Krull) dimension of X.

Rational and dominant maps. Similarly, if X and Y are prevarieties, then a *rational map* $f : X \dashrightarrow Y$ is represented by morphism from a nonempty open-dense subset of X to Y with the understanding that two such define the same map if and only if they coincide on a nonempty open-dense subset. If some representative morphism has dense image in Y, then f is said to be *dominant*. If in addition both X and Y are irreducible, then f induces a field extension $f^* : k(Y) \hookrightarrow k(X)$. Conversely, a k-linear field embedding $k(Y) \hookrightarrow k(X)$ determines a dominant rational map $X \dashrightarrow Y$. If $U \subseteq X$ is open and nonempty, then $k(U) = k(X)$ and the inclusion is a birational equivalence.

Finite morphisms. A morphism $f : X \to Y$ between prevarieties is called *finite* if it is locally so over Y, that is, if we can cover Y by open affine subsets V with the property that $f^{-1}V$ is affine and the restriction $f^{-1}V \xrightarrow{f} V$ is finite. According to Exercise 33 a finite morphism between affine varieties is closed. It then follows that a finite morphism between prevarieties is also closed.

Regular and singular points. Since the notion of a regular point is a local one, it automatically carries over to this setting. The smooth locus of X is an open-dense subset X_{reg} of X.

The product of two prevarieties. Our discussion of the product of closed subsets of affine spaces dictates how we should define the product of two prevarieties X and Y: we give the topology on $X \times Y$ by specifying a basis of open subsets, namely the collection $(U \times V)_h$, where $U \subseteq X$ and $V \subseteq Y$ are affine open and $h \in \mathcal{O}(U) \otimes \mathcal{O}(V)$. The sheaf $\mathcal{O}_{X\times Y}$ is then determined by requiring that $\mathcal{O}_{X\times Y}((U \times V)_h) = (\mathcal{O}(U) \otimes \mathcal{O}(V))[1/h]$, so that such a basis element is affine. It is clear that if let U resp. V run over a finite covering of X resp. Y, then the affine open subsets $U \times V$ run over a finite covering of $X \times Y$.

EXERCISE 53. Prove that this product has the usual categorical characterization: the two projections $X \times Y \to X$ and $X \times Y \to Y$ are morphisms and if Z is a prevariety, then a pair of maps $(f : Z \to X, g : Z \to Y)$ defines a morphism $(f, g) : Z \to X \times Y$ if and only both f and g are morphisms. (If we take $X = Y = Z$ and let f and g be the identity map, we obtain the diagonal morphism $\Delta_X : X \to X \times X$.)

The Hausdorff property is not of a local nature: a non-Hausdorff space can very well be locally Hausdorff. The standard example is the space X obtained from two copies of $\mathbb{R}$ by identifying the complement of $\{0\}$ in either copy by means of the identity map. Then X is locally like $\mathbb{R}$, but the images of the two origins cannot be separated. A topological space X is Hausdorff precisely when the diagonal of $X \times X$ is a closed subset relative to the product topology. As we know, the Zariski topology is almost never Hausdorff. But on the other hand, the selfproduct of the underlying space has not the product topology either and so requiring that the diagonal is closed is not totally unreasonable a priori. In fact, imposing this condition turns out to be the appropriate way of avoiding the pathologies that can result from an unfortunate choice of gluing data. This is illustrated by looking at the algebraic version of the standard example:

EXAMPLE 13.2. Let X be obtained from two copies $\mathbb{A}^1_+$ and $\mathbb{A}^1_-$ of $\mathbb{A}^1$ by identifying $\mathbb{A}^1_+ \smallsetminus \{o\}$ with $\mathbb{A}^1_- \smallsetminus \{o\}$ by means of the identity map. Then $\mathbb{A}^1_+ \times \mathbb{A}^1_-$ is an affine open subset of $X \times X$ which can be identified with $\mathbb{A}^2$. Under this identification, $(\mathbb{A}^1_+ \times \mathbb{A}^1_-) \cap X$ becomes the diagonal of $\mathbb{A}^2$ minus the point (o, o) (we are missing $(o_+, o_-) \in \mathbb{A}^1_+ \times \mathbb{A}^1_-$). So $(o_+, o_-) \in X \times X$ lies in the closure of the diagonal, but is not contained in the diagonal.

DEFINITION 13.3. A k-prevariety X is called a *k-variety* if the diagonal is closed in $X \times X$ (where the latter has the Zariski topology as defined above), or equivalently, when the diagonal morphism $\Delta_X : X \to X \times X$ is closed as a map. A subset of a variety X is called a *subvariety* if it is open in a closed subset of X. We say that a morphism of varieties $f : (X, \mathcal{O}_X) \to (Y, \mathcal{O}_Y)$ is an *immersion* if it defines an isomorphism onto a subvariety of Y, that is, is the composite of such an isomorphism and an inclusion.

Strictly speaking, the last part of this definition only makes sense after we have observed that a subvariety is in a natural manner a variety that makes the inclusion a morphism. We leave this as an exercise (see also Exercise 52).

Proposition 13.4. A locally closed subset of a variety is a variety and so is the product of two varieties.

PROOF. Let X be variety. The first assertion follows from the observation that for any $Z \subset X$, the diagonal of Z in $Z \times Z$ is the intersection of $Z \times Z$ with the diagonal of X and hence is closed in $Z \times Z$. For the second, note that under the obvious isomorphism $(X \times Y)^2 \cong X^2 \times Y^2$, the diagonal of $(X \times Y)^2$ becomes the product of the diagonals of X and Y. □

EXAMPLE 13.5. The diagonal in $\mathbb{A}^n \times \mathbb{A}^n$ is closed, so $\mathbb{A}^n$ is a variety. This implies that the same is true for any quasi-affine subset of $\mathbb{A}^n$. Hence a quasi-affine prevariety is in fact a variety.

EXAMPLE 13.6. Let $f : X \to Y$ be a morphism of varieties. Consider the graph of f, $\Gamma_f := \{(x,y) \in X \times Y\} : x \in U, y = f(x)\}$. It is easy to see that Γ_f is a subvariety of $X \times Y$. The map $x \in X \mapsto (x, f(x)) \in \Gamma_f$ and the projection $\Gamma_f \to X$ are regular and each others inverse. So they define an isomorphism $\Gamma_f \to X$. Notice that via this isomorphism f appears as a projection mapping: $(x,y) \in \Gamma_f \mapsto y \in Y$.

EXERCISE 54. The goal of this exercise is to show that $U := \mathbb{A}^2 \smallsetminus \{(o,o)\}$ is not affine. Let $U_x \subseteq U$ resp. $U_y \subseteq U$ be the complement of the x-axis resp. y-axis so that $U = U_x \cup U_y$.

(a) Prove that U_x is affine and that $\mathcal{O}(U_x) = k[x,y][1/x]$.
(b) Prove that every regular function on U extends to $\mathbb{A}^2$ so that $\mathcal{O}(U) = k[x,y]$.
(c) Show that U is not affine.

This exercise also leads to an interesting example. Let X be obtained as the obvious generalization of Example 13.2 where $\mathbb{A}^1$ is replaced by $\mathbb{A}^2$ so that X is covered by two copies of $\mathbb{A}^2$ (appearing as open affine subsets) whose intersection is a copy of $\mathbb{A}^2 \smallsetminus \{(o,o)\}$ (which is not affine). Hence, on a prevariety the intersection of two affine subsets need not be affine. This cannot happen on a variety and that is one of the reasons why we like them:

Proposition 13.7. Let X be a variety. Then for any pair U, U' of affine open subsets of X, $U \cap U'$ is also an affine open subset of X and $\mathcal{O}_X(U \cap U')$ is as a k-algebra generated by $\mathcal{O}_X(U)|U \cap U'$ and $\mathcal{O}_X(U')|U \cap U'$.

PROOF. First note that $U \times U'$ is an affine open subset of $X \times X$. Since the diagonal $\Delta(X)$ of $X \times X$ is closed, its intersection with $U \times U'$ is a closed subset of $U \times U'$ and hence affine. But the diagonal map sends $U \cap U'$ isomorphically onto this intersection (the inverse being given by one of the projections) and so $U \cap U'$ is affine. Since the diagonal defines a closed embedding $U \cap U' \to U \times U'$ of affine varieties, the map $k[U] \otimes k[U'] \cong k[U \times U'] \to k[U \cap U']$ is onto. Since the image of $f \otimes f' \in k[U] \otimes k[U']$ equals $f_{|U \cap U'} \cdot f'_{|U \cap U'}$, the last assertion follows. □

CHAPTER 2

Projective varieties

1. Projective spaces

Two distinct lines in the plane intersect in a single point or are parallel. In the last case one would like to say that the lines intersect at infinity so that the statement becomes simply: two distinct lines in a plane meet in a single point. There are many more examples of geometric configurations for which the special cases disappear by the simple remedy of adding points at infinity. A satisfactory approach to this which makes no a priori distinction between ordinary points and points at infinity involves the notion of a projective space.

Given a finite dimensional k-vector space V, then we denote by $\mathbb{P}(V)$ the collection of its 1-dimensional linear subspaces. Observe that any linear injection $J : V \to V'$ of vector spaces induces an injection $\mathbb{P}(J) : \mathbb{P}(V) \to \mathbb{P}(V')$ (in general $\mathbb{P}(J)$ only makes sense on $\mathbb{P}(V) \smallsetminus \mathbb{P}(\ker(J))$). In particular, when J is an isomorphism, then $\mathbb{P}(J)$ is a bijection. The following definition makes this notion slightly more abstract by suppressing the vector space as part of the data.

DEFINITION 1.1. A *projective space of dimension n over k* is a set P endowed with an extra structure that can be given by a pair (V, ℓ), where V is k-vector space of dimension $n + 1$ and $\ell : P \to \mathbb{P}(V)$ is a bijection, and where we agree that another such pair (V', ℓ') defines the same structure if and only if there exists a k-linear isomorphism $J : V \to V'$ such that $\ell' = \mathbb{P}(J)\ell'$. (So we assert that thus is defined an equivalence relation on the collection of such pairs and that a projective structure is given by an equivalence class.)

For a finite dimensional k-vector space V, the identity map of $\mathbb{P}(V)$ makes $\mathbb{P}(V)$ in a natural manner a projective space, called the *projective space associated to V*. When $V = k^{n+1}$ we write $\mathbb{P}^n$ or $\mathbb{P}^n_k$ and call it simply *projective n-space (over k)*. The difference between a projectivized vector space and an abstract projective space is perhaps elucidated by the following exercise.

EXERCISE 55. Prove that the linear isomorphism J in Definition 1.1 is unique up to scalar multiplication. Conclude that a projective space P determines a vector space up to scalar multiplication. Illustrate this by showing that for a 2-dimensional vector space V we have a canonical isomorphism $\mathbb{P}(V) \cong \mathbb{P}(V^*)$, but that there is no canonical isomorphism between V and V^*.

In Definition 1.1 we could have restricted ourselves to a fixed V, e.g., k^{n+1}. The projective structure is then given by a bijection $\ell : P \cong \mathbb{P}^n$, agreeing that two such give the same structure if and only if the two bijections differ by composition with a linear transformation in $\mathbb{P}^n$ ([1]). Giving this structure on P by means of a

[1]The structure of an m-dimensional k-vector space on a set W can analogously be given by a bijection of W onto k^m given up to an element of $\mathrm{GL}(m, k)$, but the usual definition taught in a linear

pair (k^{n+1}, ℓ) also has the advantage that it gives rise to a *homogenenous coordinate system* on P as follows: if we denote the coordinates of k^{n+1} by $(T_0, \ldots, T_n)$, then every point $p \in P$ is representable as a ratio $[p_0 : \cdots : p_n]$ of $n+1$ elements of k that are not all zero: choose a generator $\tilde{p}$ of the one-dimensional linear subspace of k^{n+1} defined by $\ell(p)$ and let $p_i = T_i(\tilde{p})$. Any other generator is of the form $\lambda\tilde{p}$ with $\lambda \in k \smallsetminus \{0\}$ and indeed, $[\lambda p_0 : \cdots : \lambda p_n] = [p_0 : \cdots : p_n]$. This is why we call (k^{n+1}, ℓ) (or rather the use of $[T_0 : \cdots : T_n]$) a *homogeneous coordinate system on* P. Note that an individual T_i is not a function on P, but that the ratios $t_{i/j} := T_i/T_j$ are, albeit that for $i \neq j$ they are not everywhere defined. It is clear that any other homogenenous coordinate system is of the form $(\sum_j a_{0j}T_0, \ldots, \sum_j a_{nj}T_n)$, where $(a_{ij})_{i,j} \in \mathrm{GL}(n+1, k)$.

DEFINITION 1.2. Given a projective space P of dimension n over k, then a subset Q of P is said to be *linear subspace of dimension* d if, for some (and hence any) pair (V, ℓ) as above, there exists a linear subspace $V_Q \subseteq V$ of dimension $d+1$ such that $\ell(Q)$ is the collection of 1-dimensional linear subspaces of V_Q.

A map $j : P \to P'$ between two projective spaces over k is said to be *linear morphism* if it comes from a linear map of vector spaces, to be precise, if for structural data (V, ℓ) for P and (V', ℓ') for P' there exists a linear *injection* $J : V \to V'$ such that $\ell' = \mathbb{P}(J)\ell$.

So a linear subspace has itself the structure of a projective space and its inclusion in the ambient projective space is a linear morphism. Conversely, the image of a linear morphism is linear subspace.

A linear subspace of dimension one resp. two is often called a *line* resp. a *plane* and a linear subspace of codimension one (= of dimension one less than the ambient projective space) is called a *hyperplane*. It is now clear that two distinct lines in a plane intersect in a single point: this simply translates the fact that the intersection of two distinct linear subspaces of dimension two in a three dimensional vector space is of dimension one.

We put on a projective space P the structure of a k-prevariety as follows. A homogeneous coordinate system $[T_0, \ldots, T_n]$ for P defines a chart for every $i = 0, \ldots, n$: if $P_{T_i} \subseteq P$ is the hyperplane complement defined by $T_i \neq 0$, then

$$\kappa_i : P_{T_i} \xrightarrow{\cong} \mathbb{A}^n, \quad p \mapsto (t_{0/i}(p), \ldots \widehat{t_{i/i}(p)} \ldots, t_{n/i}(p)),$$

is a bijection (chart) with inverse

$$\kappa_i^{-1} : (a_1, \ldots, a_n) \in \mathbb{A}^n \mapsto [a_1 : \cdots : a_i : 1 : a_{i+1} : \cdots : a_n] \in U.$$

Clearly, $\cup_{i=0}^n P_{T_i} = P$. We show that the collection of charts $\{P_{T_i}, \kappa_i\}_{i=0}^n$ can serve as an affine atlas for P. The coordinate change for a pair of charts, say for $\kappa_n\kappa_0^{-1}$, is as follows: the image of $P_{T_0} \cap P_{T_n}$ under κ_0 resp. κ_n is the open subset $\mathbb{A}^n_{x_n}$ resp. $\mathbb{A}^n_{x_1}$ of $\mathbb{A}^n$ and the transition map is

$$\kappa_n\kappa_0^{-1} : \mathbb{A}^n_{x_n} \to \mathbb{A}^n_{x_1}, \quad (a_1, a_2, \ldots, a_n) \mapsto (1/a_n, a_1/a_n, \ldots, a_{n-1}/a_n),$$

---

algebra course is of course much more convenient (apart from the fact that this would only work for finite dimensional vector spaces). Such a more intrinsic (*synthetic* is the word) definition exists also for a projective space, but is not so simple and would for us not have any clear advantages. Yet the definition given here goes one step towards such a formulation.

and hence an isomorphism of affine varieties with inverse $\kappa_0\kappa_n^{-1}$. An atlas thus obtained from a homogeneous coordinate system $(T_0, \ldots, T_n)$ will be called a *standard atlas* for P; it gives P the structure of a prevariety $(P, \mathcal{O}_P)$: $U \subseteq P$ is open if and only if for $i = 0, \ldots, n$, $\kappa_i(U \cap P_{T_i})$ is open in $\mathbb{A}^n$ and $f \in \mathcal{O}_P(U)$ if and only if $f\kappa_i^{-1} \in \mathcal{O}(\kappa_i(U))$. One can easily check that this structure is independent of the coordinate system and that it is in fact that of a k-variety. We will not do this here as we will give in Section 3 a more direct proof of these assertions.

Any hyperplane $H \subseteq P$ can be given as $T_0 = 0$, where $[T_0 : \cdots : T_n]$ is a homogeneous coordinate system on P and so its complement $U = P \smallsetminus H$ is isomorphic to $\mathbb{A}^n$. This can also (and more intrinsically) be seen without the help of such a coordinate system. Let the projective structure on P be given by the pair (V, ℓ). Then the hyperplane H corresponds to a hyperplane $V_H \subseteq V$ and U corresponds to the set of 1-dimensional linear subspaces of V not contained in H. If $e \in V^*$ is a linear form whose zero set is V_H, then $A = e^{-1}(1)$ is an affine space for V_H (it has V_H as its vector space of translations). Assigning to $v \in A$ the 1-dimensional linear subspace spanned by v defines a bijection $A \cong U$ that puts on U a structure of an affine space. This structure is easily checked to be independent of (V, ℓ, ϕ).

We could also proceed in the opposite direction and start with an affine space A and realize it as the hyperplane complement of a projective space. For this consider the vector space $F(A)$ of affine-linear functions on A and denote by $e \in F(A)$ the function on A that is constant equal to 1. Then $e^{-1}(1)$ is an affine hyperplane in $F(A)^*$. Any $a \in A$ defines a linear form on $F(A)$ by evaluation: $f \in F(A) \mapsto f(a) \in k$. Note that this form takes the value 1 on e so that we get in fact a map $A \to e^{-1}(1)$. It is not hard to check that this is an affine-linear isomorphism and so the projective space $\overline{A} := \mathbb{P}(F(A)^*)$ can serve as the projective completion of A. Following the Renaissance painters, we might say that $\overline{A} \smallsetminus A$ consists of "points at infinity" of A; such a point can be given by an affine line in A with the understanding that parallel lines define the same point at infinity.

2. The Zariski topology on a projective space

We begin with giving a simpler and more natural characterization of the Zariski topology on a projective space. Let P be a projective space of dimension n over k and let $[T_0 : \cdots : T_n]$ be a homogeneous coordinate system for P. Suppose $F \in k[X_0, \ldots, X_n]$ is homogeneous of degree d so that $F(tT_0, \ldots, tT_n) = t^d F(T_0, \ldots, T_n)$ for $t \in k$. The property of this being zero only depends on $[T_0 : \cdots : T_n]$ and hence the zero set of F defines a subset of P. We shall denote this subset by $Z[F]$ and its complement $P \smallsetminus Z[F]$ by P_F. We will show in the next section that P_F is in fact affine.

Proposition 2.1. The collection $\{P_F\}_F$, where F runs over the homogeneous polynomials in $k[X_0, \ldots, X_n]$, is a basis for the Zariski topology on P. This topology is independent of the choice of our homogeneous coordinate system $[T_0, \ldots, T_n]$ and (so) every linear chart is a homeomorphism onto $\mathbb{A}^n$ that identifies the sheaf of regular functions on its domain with $\mathcal{O}_{\mathbb{A}^n}$. If $G \in k[X_0, \ldots, X_n]$ is homogeneous of the same degree as F and nonzero, then F/G defines a regular function on P_G.

PROOF. We first observe that the obvious equality $P_F \cap P_{F'} = P_{FF'}$ implies that the collection $\{P_F\}_F$ is a basis of a topology $\mathcal{T}$ on P. This topology is independent

of the coordinate choice, because a linear substitution transforms a homogeneous polynomial into another one.

Let us verify that this is the Zariski topology defined earlier. First note that the domain of each member $\kappa_i : P_{T_i} \cong \mathbb{A}^n$ of the standard atlas is also a basis element (hence open) for $\mathcal{T}$. So we must show that κ_i defines a homeomorphism onto $\mathbb{A}^n$ when its domain is endowed with the topology induced by $\mathcal{T}$. If $F \in k[T_0, \dots, T_n]$ is homogeneous of degree d, then $\kappa_i(P_F \cap P_{T_i}) = \mathbb{A}^n_{f_i}$, where $f_i(y_1, \dots, y_n) := F(y_1, \dots, y_i, 1, y_{i+1}, \dots, y_n)$ and so κ_i is open. Conversely, if $f \in k[y_1, \dots, y_n]$ is nonzero of degree d, then its homogenization of degree d, $F(T_0, \dots, T_n) := T_i^d f(T_0/T_i, \dots \widehat{T_i/T_i} \dots T_n/T_i)$, has the property that $\kappa_i^{-1}(\mathbb{A}^n_f) = P_F \cap P_{T_i}$. So κ_i is also continuous.

For the last statement first observe that F/G indeed defines a function on P_G (think of it as regular function on $\mathbb{A}^{n+1}_G$ that is constant under multiplication with a nonzero scalar). Its pull-back under κ_i is $f_i(y_1, \dots, y_n)/g_i(y_1, \dots y_n)$, where $g_i(y_1, \dots, y_n) := G(y_1, \dots, 1, \dots y_n)$, which is indeed regular on $\mathbb{A}^n_{g_i}$. So F/G is regular on P_G. □

EXERCISE 56. Let $0 \neq G \in k[X_0, \dots, X_n]$ be homogeneous of degree d. Prove that every regular function $P_G \to k$ has the form F/G^r, with $r \geq 0$ and F homogeneous of the same degree as G^r.

In order to discuss the projective analogue of the (affine) $I \leftrightarrow Z$ correspondence, we shall need the following notions from commutative algebra.

DEFINITION 2.2. A *(nonnegative) grading* of a ring A is a direct sum decomposition $A = \oplus_{k=0}^{\infty} A_d$ such that the product maps $A_d \times A_e$ in A_{d+e}, or equivalently, such that $\sum_{d=0}^{\infty} A_d t^d$ is a subalgebra of $A[t]$ (so that this makes A_0 a subring of A). We then call A a *graded ring*. If we are also given a ring homomorpism $R \to A_0$, then we call A a *graded R-algebra*.

An ideal I of a graded ring A is said to be *graded* if it is the direct sum of its homogeneous parts $I_d := I \cap A_d$ (so that $\sum_{d=0}^{\infty} I_d t^d$ is an ideal of $\sum_{d=0}^{\infty} A_d t^d$).

We shall say that a graded ideal I of A is *proper*([2]) if it is graded and contained in the ideal $A_+ := \oplus_{d \geq 1} A_d$.

If we feel that the presence of a graded structure on A must be expressed by our notation, we shall write $A_\bullet$ for A. Note that if I is a graded ideal of the graded ring A, then $A/I = \oplus_{d=0}^{\infty} A_d/I_d$ is again a graded algebra and that if I is also proper, then A/I has A_0 as its degree zero part. We will be mostly concerned with the case when $A_0 = k$; then A_+ is a maximal ideal (and the only one that is graded).

Lemma 2.3. If I, J are (proper) graded ideals of a graded ring A, then so are $I \cap J$, IJ, $I + J$ and $\sqrt{I}$. Moreover, if $\mathfrak{p} \subseteq A$ is a prime ideal, then so is the graded ideal $\oplus_n(\mathfrak{p} \cap A_n)$; in particular, a minimal prime ideal of R is graded.

PROOF. The proofs of the statements in the first sentence are not difficult and so we omit them. As to the last, let us first agree to denote for any nonzero $a \in A$ by $\mathrm{in}(a)$ the top degree part of a and by $\deg(a)$ its degree and stipulate that when $a = 0$, $\mathrm{in}(a) = 0$ and $\deg(a) = 0$.

Now put $\mathfrak{p}_n := \mathfrak{p} \cap A_n$ let $a, b \in A$ be such that $ab \in \oplus_n \mathfrak{p}_n$. We prove with induction on $\deg(a) + \deg(b)$ that a or b is in $\oplus_n \mathfrak{p}_n$. If $\mathrm{in}(a)\,\mathrm{in}(b) \neq 0$, then it

[2]This is not a generally adopted terminology.

equals $\mathrm{in}(ab) \in \mathfrak{p}_{\deg(a)+\deg(b)} \subseteq \mathfrak{p}$ and since $\mathfrak{p}$ is a prime ideal, we therefore have $\mathrm{in}(a) \in \mathfrak{p}_{\deg(a)}$ or $\mathrm{in}(b) \in \mathfrak{p}_{\deg(b)}$. This is of course also true when $\mathrm{in}(a)\,\mathrm{in}(b) = 0$. Say that $\mathrm{in}(a) \in \mathfrak{p}_{\deg(a)}$. Then $(a - \mathrm{in}(a))b = ab - \mathrm{in}(a)b \in \mathfrak{p}$. We have either $\deg(a - \mathrm{in}(a)) < \deg(a)$ or $\deg(a) = 0$ (so that $a - \mathrm{in}(a) = 0$). Hence by our induction assumption (or trivially in the second case), $a - \mathrm{in}(a)$ or b is in $\oplus_n \mathfrak{p}_n$. So a or b is in $\oplus_n \mathfrak{p}_n$. □

The main example of a graded k-algebra is furnished by a vector space V of finite positive dimension ($n+1$, say), which we consider as an affine variety, but (in contrast to an affine space) one of which we remember that it comes with the action of the multiplicative group of k by scalar multiplication. Let us say that $F \in k[V]$ is *homogeneous of degree* d if we have $F(tv) = t^d F(v)$ for all $v \in V$ and $t \in k$. Such F make up a k-linear subspace $k[V]_d$ of finite dimension and the subspaces thus defined turn $k[V]$ into a graded k-algebra $\oplus_{d\geq 0} k[V]_d$. (A choice of basis $(T_0, \ldots, T_n)$ of V^* identifies V with $\mathbb{A}^{n+1}$ and then $k[V]_d$ becomes the space of homogeneous polynomials in $(T_0, \ldots, T_n)$ of degree d.) We can understand this decomposition as the one into eigenspaces with respect to the action of scalar multiplication. Note that for any $F \in k[V]_d$ the zero set $Z(F) \subseteq V$ is invariant under scalar multiplication. This is still true for an intersection of such zero sets, in other words, for a proper graded ideal $I \subseteq k[V]_+$, $Z(I) \subseteq V$ is a closed subset of V that is invariant under scalar multiplication. Such a closed subset is called an *affine cone*. The origin is called the *vertex* of that cone. Since the vertex is defined by the maximal ideal $k[V]_+$, we always have $0 \in Z(I)$. The intersection of the $Z[F]$, with $F \in \cup_{d\geq 1} I_d$ defines a closed subset $Z[I]$ of $\mathbb{P}(V)$ whose points correspond to the one-dimensional subspaces of V that are contained in $Z(I)$.

It is clear that every closed subset of $\mathbb{P}(V)$ is thus obtained. In fact, given a closed subset $X \subseteq \mathbb{P}(V)$, let for $d \geq 1$, $I_{X,d}$ be the set of $F \in k[V]_d$ for which $X \subseteq Z[F]$ and put $I_{X,0} = 0$. Then $I_{X,d}$ is a k-vector space and $I_{X,d} \cdot k[V]_e \subseteq I_{X,d+e}$ so that $I_X := \oplus_{d\geq 1} I_{X,d}$ is a proper graded ideal of $k[V]$. It is also a radical ideal (exercise) and we have $X = Z[I_X]$. So $Z(I_X)$ is the cone in V that as a set is just the union of the 1-dimensional linear subspaces of V parameterized by X; we denote this cone in V by $\mathrm{Cone}(X)$. Note that the degenerate cone $\{0\} \subseteq V$ corresponds to the empty subset of $\mathbb{P}(V)$ and to the homogeneous maximal ideal $k[V]_+$.

Lemma 2.4. For an affine cone $C \subseteq V$, $I(C)$ is a proper graded radical ideal of $k[V]$ and hence defines a closed subset $\mathbb{P}(C) := Z[I(C)]$ of $\mathbb{P}(V)$.

PROOF. Let $F \in I(C)$. Write $F = \sum_{d\geq 0} F_d$. We must show that each homogeneous component of F_d lies in $I(C)$. As C is invariant under scalar multiplication, the polynomial $F(tv) = \sum_{d\geq 1} t^d F_d(v)$ (as an element of $k[\mathbb{A}^1 \times V]$) vanishes on $\mathbb{A}^1 \times C \subseteq \mathbb{A}^1 \times V$, hence lies in $I(\mathbb{A}^1 \times C)$. But under the identification $k[\mathbb{A}^1 \times V] \cong k[V][t]$, $I(\mathbb{A}^1 \times C) = k[t] \otimes I(C)$ corresponds to $I(C)[t]$ and so it follows that $F_d \in I(C)$ for all d. □

We conclude:

Corollary 2.5. The maps $C \mapsto \mathbb{P}(C)$ and $X \mapsto \mathrm{Cone}(X)$; $C \mapsto I(C)$; $X \mapsto I_X$; $I \mapsto Z(I)$ and $I \mapsto Z[I]$ set up bijections between

(i) the collection of affine cones in V,
(ii) the collection of closed subsets of $\mathbb{P}^n$ and
(iii) the collection of proper graded radical ideals contained in $k[V]$.

This restricts to bijections between (i) the collection of irreducible affine cones in V strictly containing $\{0\}$, (ii) the collection of irreducible subsets of $\mathbb{P}(V)$ and (iii) the collection of graded prime ideals of $k[V]$ strictly contained in $k[V]_+$.

DEFINITION 2.6. The *homogeneous coordinate ring* of a closed subset X of $\mathbb{P}(V)$ is the coordinate ring of the affine cone over X, $k[\mathrm{Cone}(X)] = k[V]/I_X$, endowed with the grading inherited by $k[V]$: $k[\mathrm{Cone}(X)]_d = k[T_0, \dots T_n]_d/I_{X,d}$.

More generally, if Y is an affine variety, and X is a closed subset of $\mathbb{P}(V) \times Y$, then the *homogeneous coordinate ring of X relative to Y* of a closed subset is the coordinate ring of the corresponding closed cone in $V \times Y$ over Y, endowed with the grading defined by the coordinates of V.

EXERCISE 57. Let A be a graded ring.

(b) Prove that if I is a prime ideal in the homogeneous sense: if $rs \in I$ for some $r \in A_k, s \in A_l$ implies $r \in I$ or $s \in I$, then I is a prime ideal.
(c) Prove that the intersection of all graded prime ideals of A is its ideal of nilpotents.

EXERCISE 58. Let S be a graded k-algebra that is reduced, generated by S_1 and has $S_0 = k$ and $\dim_k S_1$ finite.

(a) Prove that S is as a graded k-algebra isomorphic to the homogeneous coordinate ring of a closed subset Y.
(b) Prove that under such an isomorphism, the graded radical ideals contained in the maximal ideal $S_+ := \oplus_{d \geq 1} S_d$ correspond to closed subsets of Y under an inclusion reversing bijection: graded ideals strictly contained in S_+ and maximal for that property correspond to points of Y.
(c) Suppose S a domain. Show that a fraction $F/G \in \mathrm{Frac}(S)$ that is homogeneous of degree zero ($F, G \in S_d$ for some d and $G \neq 0$ defines a function on U_g.

EXERCISE 59. Let Y be an affine variety.

(a) Show that a homogeneous element of the graded ring $k[Y][T_0, \dots, T_n]$ defines a closed subset of $Y \times \mathbb{P}^n$ as its zero set.
(b) Prove that every closed subset of $Y \times \mathbb{P}^n$ is an intersection of zero sets of finitely many homogeneous elements of $k[Y][T_0, \dots, T_n]$.
(c) Prove that we have a bijective correspondence between closed subsets of $Y \times \mathbb{P}^n$ and the homogeneous radical ideals in $k[Y][T_0, \dots, T_n]_+$.

The next proposition is a first illustration of the power of projective methods. Its proof makes interesting use of a homogeneous coordinate ring (and in particular, of Exercise 59).

Proposition 2.7. Let X be a variety and let $Z \subseteq \mathbb{A}^m \times X$ be closed in $\mathbb{P}^m \times X$ (we here identify $\mathbb{A}^m$ with the hyperplane complement $\mathbb{P}^m_{T_0}$). Then the projection $\pi_X|Z : Z \to X$ is a finite morphism.

PROOF. Without loss of generality we may assume that X is affine. Let $I \subseteq k[X][T_0, \dots, T_m]$ be the graded ideal defining Z. Since Z does not meet the zero set of T_0, the ideal $I + (T_0) \subseteq k[X][T_0, \dots, T_m]$ generated by I and T_0 defines the empty set in $\mathbb{P}^n \times X$, or equivalently, $\{0\} \times X$ in $\mathbb{A}^{n+1} \times X$. So $\sqrt{I + (T_0)} = (T_0, \dots, T_m)$ by Hilbert's Nullstellensatz. In particular, there exists an integer $r > 0$ such that $T_i^r \in I_r + (T_0)_r$ for $i \in \{1, \dots, m\}$. Write $T_i^r \equiv T_0 G_i \pmod{I_r}$ with

$G_i \in k[X][T_0, \ldots, T_m]_{r-1}$. We pass to the affine coordinates of $\mathbb{A}^m$ by substituting 1 for T_0 and t_i for T_i. Then G_i yields a $g_i \in k[X][t_1, \ldots, t_m] = k[\mathbb{A}^n \times X]$ of degree $\leq r-1$ in the t-variables and we have $t_i^r \equiv g_i \pmod{I_{\mathbb{A}^n \times X}(Z)}$. So if we write $\bar{t}_i$ for the image $\bar{t}_i$ of t_i in $k[Z]$, then for each i, $\bar{t}_i^r$ is a $k[X]$-linear combination of the monomials $\bar{t}_1^{s_1} \cdots \bar{t}_m^{s_m}$ with $s_1 + \cdots + s_m < r$. Hence $k[Z]$ is as a $k[X]$-module generated by these (finitely many) monomials. This proves that $\pi_X | Z : Z \to X$ is a finite morphism. □

Note the special case when X is a singleton: the proposition then tells us that if a closed subset of $\mathbb{A}^m$ is also closed in $\mathbb{P}^m$, then must be finite. In other words, a closed subset of $\mathbb{A}^m$ of positive dimension, must, when considered as a subset $\mathbb{P}^m$, have a point in the "hyperplane at infinity" $T_0 = 0$ in its closure (in classical language: Z must have an asymptote).

3. The Segre embeddings

First we show how a product of projective spaces can be realized as a closed subset of a projective space. This will imply among other things that a projective space is a variety. Consider the projective spaces $\mathbb{P}^m$ and $\mathbb{P}^n$ with their homogeneous coordinate systems $[T_0 : \cdots : T_m]$ and $[W_0 : \cdots : W_n]$. We also consider a projective space whose homogeneous coordinate system is the set of matrix coefficients of an $(m+1) \times (n+1)$-matrix $[Z_{00} : \cdots : Z_{ij} : \cdots : Z_{mn}]$; this is just $\mathbb{P}^{mn+m+n}$ with an unusual indexing of its homogeneous coordinates.

Proposition 3.1 (The Segre embedding)**.** The map $f : \mathbb{P}^m \times \mathbb{P}^n \to \mathbb{P}^{mn+m+n}$ defined by $Z_{ij} = T_i W_j$, $i = 0, \ldots, m; j = 0, \ldots, n$ is an isomorphism onto a closed subset of $\mathbb{P}^{mn+m+n}$. If $m = n$, then the diagonal of $\mathbb{P}^m \times \mathbb{P}^m$ is the preimage of the linear subspace of $\mathbb{P}^{m^2+2m}$ defined by $Z_{ij} = Z_{ji}$ and hence is closed in $\mathbb{P}^m \times \mathbb{P}^m$.

PROOF. For the first part it is enough to show that for every chart domain $\mathbb{P}^{mn+m+n}_{Z_{ij}}$ of the standard atlas of $\mathbb{P}^{mn+m+n}$, $f^{-1}\mathbb{P}^{mn+m+n}_{Z_{ij}}$ is open in $\mathbb{P}^m \times \mathbb{P}^n$ and is mapped by f isomorphically onto a closed subset of $\mathbb{P}^{mn+m+n}_{Z_{ij}}$. For this purpose we may (simply by renumbering) assume that $i = j = 0$. So then $\mathbb{P}^{mn+m+n}_{Z_{00}} \subseteq \mathbb{P}^{mn+m+n}$ is defined by $Z_{00} \neq 0$ and is parametrized by the coordinates $z_{ij} := Z_{ij}/Z_{00}$, $(i,j) \neq (0,0)$. It is clear that $f^{-1}\mathbb{P}^{mn+m+n}_{Z_{00}}$ is defined by $T_0 W_0 \neq 0$. This is just $\mathbb{P}^m_{T_0} \times \mathbb{P}^n_{W_0}$ and hence is parametrized by $(t_1, \ldots, t_m) := (T_1/T_0, \ldots, T_m/T_0)$ and $(w_1, \ldots, w_n) := (W_1/W_0, \ldots, W_n/W_0)$. In terms of these coordinates, $f : f^{-1}\mathbb{P}^{mn+m+n}_{Z_{00}} \to \mathbb{P}^{mn+m+n}_{Z_{00}}$ is given by $z_{ij} = t_i w_j$, where $(i,j) \neq (0,0)$ and where we should read 1 for t_0 and w_0. So among these are $z_{i0} = t_i$ and $z_{0j} = w_j$ and since these generate $k[\mathbb{A}^m \times \mathbb{A}^n] = k[t_1, \ldots, t_m, w_1, \ldots, w_n]$, f indeed restricts (by Proposition 4.3 of Ch. 1) to a closed immersion $f^{-1}\mathbb{P}^{mn+m+n}_{Z_{00}} \to \mathbb{P}^{mn+m+n}_{Z_{00}}$.

In case $m = n$, we must also show that the condition $T_i W_j = T_j W_i$ for $0 \leq i < j \leq m$ implies that $[T_0 : \cdots : T_m] = [W_0 : \cdots : W_m]$, assuming that not all T_i resp. W_j are zero. Suppose $T_i \neq 0$. Since $W_j = (W_i/T_i).T_j$ for all j, it follows that $W_i \neq 0$ and so $[W_0 : \cdots : W_m] = [T_0 : \cdots : T_m]$. □

Corollary 3.2. A projective space over k is a variety. In particular, a locally closed subset of a projective space is a variety.

PROOF. Proposition 3.1 shows that the diagonal of $\mathbb{P}^m \times \mathbb{P}^m$ is closed. □

DEFINITION 3.3. A variety is said to be *projective* if it is isomorphic to a closed *irreducible* subset of some projective space. A variety is called *quasi-projective* if is isomorphic to an open subset of some projective variety.

Note that we have here required that the variety in question is irreducible.

EXERCISE 60. (a) Prove that the image of the Segre embedding is the common zero set of the homogeneous polynomials $Z_{ij}Z_{kl} - Z_{il}Z_{kj}$.

(b) Show that for every $(p, q) \in \mathbb{P}^m \times \mathbb{P}^n$ the image of $\{p\} \times \mathbb{P}^n$ and $\mathbb{P}^m \times \{q\}$ in $\mathbb{P}^{mn+m+n}$ is a linear subspace.

(c) Prove that the map $\mathbb{P}^n \to \mathbb{P}^{(n^2+3n)/2}$ defined by $Z_{ij} = T_iT_j$, $0 \leq i \leq j \leq n$ is an isomorphism on a closed subset defined by quadratic equations. Find these equations for $n = 2$.

(d) As a special case we find that the quadric hypersurface in $\mathbb{P}^3$ defined by $Z_0Z_1 - Z_2Z_3 = 0$ is isomorphic to $\mathbb{P}^1 \times \mathbb{P}^1$. Identify in this case the two systems of lines on this quadric.

EXERCISE 61 (Intrinsic Segre embedding). Let V and W be finite dimensional k-vector spaces. Describe the Segre embedding for $\mathbb{P}(V) \times \mathbb{P}(W)$ intrinsically as a morphism $\mathbb{P}(V) \times \mathbb{P}(W) \to \mathbb{P}(V \otimes W)$.

4. Blowing up and projections

By way of introduction we first explain the blowup of a linear subspace of a vector space. Let V be a finite dimensional k-vector space of dimension n and $W \subset V$ a linear subspace of codimension c. The first thing to note is that a point of $\mathbb{P}(V/W)$ corresponds to a 1-dimensional subspace in V/W and that a 1-dimensional subspace of V/W amounts to giving a linear subspace of V which contains W as a subspace of codimension one (in other words, as a hyperplane). With this interpretation, the composite map $\pi : V \smallsetminus W \to (V/W) \smallsetminus \{0\} \to \mathbb{P}(V/W)$ can then be understood as assigning to $p \in V \smallsetminus W$ the linear subspace $W + kp \subseteq V$ spanned by W and p.

Now define the *blowup of V along W*, denoted $\mathrm{Bl}_W V$, as the set of pairs $(p, [W']) \in V \times \mathbb{P}(V/W)$, such that $p \in W'$. So when $p \in V \smallsetminus W$, there is precisely one choice of W' such that $(p, [W']) \in \mathrm{Bl}_W V$, namely $W' = W + kp$, but when $p \in W$, $[W'] \in \mathbb{P}(V/W)$ can be arbitrary. So as a set, $\mathrm{Bl}_W V$ is the disjoint union of the graph Γ_π of π and the product $W \times \mathbb{P}(V/W)$. The projection $\mathrm{Bl}_W V \to V$ is an isomorphism over $V \smallsetminus W$ with inverse $p \mapsto (p, \pi(p))$, so that the projection $\mathrm{Bl}_W V \to \mathbb{P}(V/W)$, may be regarded as an extension of π; the latter map is often called the *projection away from W*.

Lemma 4.1. The blowup $\mathrm{Bl}_W V$ of V along W is an irreducible smooth closed subset of $V \times \mathbb{P}(V/W)$ which contains $W \times \mathbb{P}(V/W)$ as a smooth hypersurface (called the *exceptional divisor* of the blowup) so that its complement Γ_π is open dense in $\mathrm{Bl}_W V$.

PROOF. Let $(x_1, \ldots, x_n)$ be a coordinate system for V such that W is defined by $x_1 = \cdots = x_c = 0$. We can use $(x_1, \ldots, x_c)$ as coordinates for V/W, but in order to avoid confusion, we denote the associated homogeneous coordinate system for $\mathbb{P}(V/W)$ by $[T_1 : \cdots : T_c]$. So then $\pi(x_1, \ldots, x_n) = [x_1 : \cdots : x_c]$ and $\mathrm{Bl}_W V$ is the set of $((x_1, \ldots, x_n), [T_1 : \cdots : T_c])$ with $[T_1 : \cdots : T_c] = [x_1 : \cdots : x_c]$ whenever $(x_1, \ldots, x_c) \neq (0, \ldots, 0)$. This is equivalent to: $x_iT_j = x_jT_i$, $1 \leq i < j \leq c$ and so

this indeed defines a closed subset of $V \times \mathbb{P}(V/W)$. Note that $W \times \mathbb{P}(V/W)$ resp. Γ_π is the locus where $(x_1, \dots, x_c)$ equals $(0, \dots, 0)$ resp. does not equal $(0, \dots, 0)$.

Next we show that $\mathrm{Bl}_W V$ is smooth. For this we note that $(\mathrm{Bl}_W V)_{T_1}$ is the set of $((x_1, \dots, x_n), [1 : t_2 : \cdots ; t_c])$ satisfying $x_i = t_i x_1$ for $i = 2, \dots, c$ and hence is parametrized by $\mathbb{A}^n$ via

$$(x_1, t_2, \dots, t_c, x_{c+1}, \dots, x_n) \mapsto ((x_1, t_2 x_1, \dots, t_c x_1, x_{c+1}, \dots, x_n), [1 : t_2 : \cdots ; t_c]).$$

So we have here described an isomorphism $(\mathrm{Bl}_W V)_{T_1} \cong \mathbb{A}^n$. Note that this identifies $(\mathrm{Bl}_W V)_{T_1} \cap (W \times \mathbb{P}(V/W))$ with the coordinate hyperplane in $\mathbb{A}^n$ defined by $x_1 = 0$. All the assertions now follow, as this shows that $\mathrm{Bl}_W(V)$ is covered by the c open subsets $(\mathrm{Bl}_W V)_{T_i}$, $i = 1, \dots, c$, each of which is isomorphic to $\mathbb{A}^n$ and meets $W \times \mathbb{P}(V/W)$ in a smooth hypersurface. □

REMARK 4.2. We here outline the algebraic aspects of this construction, although this is not needed for what follows. The homogeneous coordinate ring of $\mathrm{Bl}_W V$ is the graded $k[V]$-algebra $k[V][T_1, \dots, T_c]$ modulo the ideal generated by the $x_i T_j - x_j T_i$, $1 \le i < j \le c$, with each T_i of degree 1. It admits the following elegant description: consider the homomorphism of graded $k[V]$-algebras

$$k[V][T_1, \dots, T_c]/(x_i T_j - x_j T_i, 1 \le i < j \le c) \to k[V][I(W)t] = \sum_{d \ge 0} I(W)^d t^d, \quad T_i \mapsto x_i t,$$

where $I(W)$ is the ideal defining W (so generated by $x_1, \dots, x_c$). The algebra on the right hand side is an instance of the construction appearing in Lemma 11.8 (so as a $k[V]$-algebra generated by $x_1 t, \dots, x_c t$ and can also be written as $\oplus_{d \ge 0} I(W)^d$. This is in fact an isomorphism whose inverse is defined as follows: a k-basis of the right hand side consists of the monomials $x_1^{d_1} \cdots x_n^{d_n} t^d$ with $d_1 + \cdots + d_c \ge d$ and then the inverse assigns to $x_1^{d_1} \cdots x_n^{d_n} t^d$ the image in the left hand side of a monomial $x_1^{d_1} \cdots x_n^{d_n} T_1^{e_1} \cdots T_c^{e_c}$, where $0 \le e_i \le d_i$ are such that $\sum_i e_i = d$ (check that the image is independent of this choice). So here we have an algebraic description of $\mathrm{Bl}_W V$ in completely coordinate-independent terms. The exceptional divisor is defined by the ideal $I(W) = (x_1, \dots, x_c)$. The associated quotient ring is a graded $k[V]/I(W) = k[W]$-algebra, which in the first description yields $k[W][T_1, \dots, T_c]$ (this is indeed the homogeneous coordinate ring of $W \times \mathbb{P}(V/W)$) and in the second yields $\sum_{d \ge 0} (I(W)^d / I(W)^{d+1}) t^d \cong \oplus_{d \ge 0} I(W)^d / I(W)^{d+1}$. These must of course be isomorphic as $k[W]$-algebras, but the second description is more canonical in the sense that it identifies the exceptional divisor with the projectivized normal bundle of W in V.

The algebra homomorphism $k[V] \subseteq k[V][I(W)t]$ defines the projection $p : \mathrm{Bl}_W V \to V$. If $H \subseteq V$ is a hypersurface, and is defined by $f \in k[V]$ say, then the *total transform of* H is simply $p^{-1}H$ (it is defined as the zero set of f, when regarded as an element of $k[V] \to \sum_{d \ge 0} I(W)^d t^d$), whereas the *strict transform of* H under p is by definition the closure $\tilde{H}$ of $H \setminus H \cap W$ in $\mathrm{Bl}_W V$. A defining equation for $\tilde{H}$ is $f t^{\nu_W(f)} \in k[V][I(W)t]$, where $\nu_W(f)$ is the maximal m for which $f \in I(W)^m$ (also called the *W-valuation* of f). So the strict transform $\tilde{H}$ meets the exceptional divisor $W \times \mathbb{P}(V/W)$ in the locus defined $\mathrm{in}_W(f)$, the image of f in $I(W)^{\nu_W(f)} / I(W)^{\nu_W(f)+1}$ (the *W-initial part* of f).

More generally, if $X \subseteq V$ is a closed subset, then the ideal generated by the $p^* f$, $f \in I(X)$, defines its total transform $p^{-1}X$, whereas the ideal generated by the $f t^{\nu_W(f)}$, $f \in I(X)$, defines the strict transform of X under p, that is, the closure of $X \setminus W$ in $\mathrm{Bl}_W(V)$. (If $X \setminus W$ is dense in X, then this strict transform is called the blowup of X along $X \cap W$ and also denoted $\mathrm{Bl}_{W \cap X} X$.)

EXERCISE 62. Recall that for a ring R and an ideal $I \subseteq R$, the 'blowup of I in R' is the graded R-subalgebra $R[It] = \sum_{d \ge 0} I^d t^d$ of $R[t]$ which appeared in the proof of Lemma 11.8. We regard the graded quotient $R[It]/IR[It] = \sum_{d \ge 0} (I^d / I^{d+1}) t^d \cong \oplus_{d=0}^{\infty} I^d / I^{d+1}$ as

defining its 'exceptional divisor'. If W a closed subset of an affine variety X, then the blowup $\mathrm{Bl}_W(X)$ is defined by applying this to $R = k[X]$ and $I = I(W)$.

Prove that if $I(W)$ has $d > 0$ generators, then $\mathrm{Bl}_W(X)$ is a closed subset of $X \times \mathbb{P}^{d-1}$. Show that if W is nowhere dense in X and $f_1, \ldots, f_d$ generate $I(W)$ in $k[X]$, then $\mathrm{Bl}_W(X)$ is the closure in $X \times \mathbb{P}^{d-1}$ of the graph of the morphism $[f_1 : \cdots : f_d] : X \smallsetminus W \to \mathbb{P}^{d-1}$.

The blowing up of W in V models a projective analogue that we discuss next. Let P be a projective space of dimension n and $Q \subseteq P$ a linear subspace of codimension $c(= \dim P - \dim Q)$. Let us denote by $\mathbb{P}(P;Q)$ the collection of linear subspaces Q' of P which contain Q as a hyperplane (and so are of dimension $\dim Q + 1$).

Lemma 4.3. The space $\mathbb{P}(P;Q)$ has in a natural manner the structure of a projective space of dimension $c-1$ (where dimension -1 means empty). Through every $p \in P \smallsetminus Q$ passes exactly one member of $\mathbb{P}(P;Q)$ and this defines a morphism $\pi_Q : P \smallsetminus Q \to \mathbb{P}(P;Q)$. Concretely, if $n := \dim P$ and $[T_0 : \ldots T_n]$ is a system of homogeneous coordinates for P such that Q is given by $T_0 = \cdots = T_{c-1}$, then $[T_0 : \cdots : T_{c-1}]$ defines a system of homogeneous coordinates for $\mathbb{P}(P;Q)$ and π_Q is simply given by $[T_0 : \cdots : T_n] \mapsto [T_0 : \cdots : T_{c-1}]$.

PROOF. Let $\ell : P \cong \mathbb{P}(V)$ be a structural bijection. Then $Q = \ell^{-1}\mathbb{P}(V_Q)$ for some linear subspace $V_Q \subseteq V$ and so the Q' correspond to the linear subspaces $V_{Q'} \subseteq V$ which contain V_Q as a hyperplane. These in turn correspond to the one-dimensional subspaces of V/V_Q and so we get a bijection $\mathbb{P}(P;Q) \cong \mathbb{P}(V/V_Q)$. For another choice of structural bijection (V', ℓ') there must exist a linear isomorphism $V \cong V'$ which then automatically takes V_Q onto V'_Q and so induces a linear isomorphism $V/V_Q \cong V'/V'_Q$. We thus see that the projective space structure on $\mathbb{P}(P;Q)$ is intrinsically defined. The proof of the last assertion is left to you. □

Definition-Lemma 4.4. The *blowup of P along Q*, denoted $\mathrm{Bl}_Q P$, is the set pairs $(p, [Q']) \in P \times \mathbb{P}(P;Q)$ with $p \in Q'$. This is a closed subset of $P \times \mathbb{P}(P;Q)$ (hence is a projective variety) and the projection

$$p_2 : \big(\mathrm{Bl}_Q P, Q \times \mathbb{P}(P;Q)\big) \to \mathbb{P}(P;Q);$$

has the following property: if $U \subseteq \mathbb{P}(P;Q)$ is a hyperplane complement (hence an affine space), then there exist a $[Q'] \in U$ (so Q' contains Q as a hyperplane) and a morphism $r : p_2^{-1}U \to Q'$ which on $Q \times U$ is the projection onto Q and is such that $(r, p_2) : p_2^{-1}U \to Q' \times U$ is an isomorphism[3]. In particular, $\mathrm{Bl}_Q P$ is smooth and irreducible; it contains $Q \times \mathbb{P}(P;Q)$ as a smooth hypersurface (this is called the *exceptional divisor* of the blowup) whose complement is the graph of $\pi_Q : P \smallsetminus Q \to \mathbb{P}(P;Q)$.

PROOF. The set theoretic part is clear: $\mathrm{Bl}_Q P$ contains $Q \times \mathbb{P}(P;Q)$ with complement the graph of $\pi_Q : P \smallsetminus Q \to \mathbb{P}(P;Q)$. To prove the remaining assertions, we use a homogeneous coordinate system $[T_0 : \cdots : T_n]$ for P as above (so that Q is given by $T_0 = \cdots = T_{c-1} = 0$). If we denote the corresponding coordinate system for $\mathbb{P}(P;Q)$ by $[S_0 : \cdots : S_{c-1}]$, then $\mathrm{Bl}_Q P$ is given by the pairs $([T_0 : \cdots : T_n], [S_0 : \cdots : S_{c-1}])$ such that $[T_0 : \cdots : T_{c-1}] = [S_0 : \cdots : S_{c-1}]$ in case $T_0, \ldots, T_c$ are not all zero. In other words, $\mathrm{Bl}_Q P$ is defined by the system of equations $T_i S_j = T_j S_i$ for all $0 \le i < j < c$. So it is a closed subset of $P \times \mathbb{P}(P;Q)$.

[3] We then say that r_2 defines a *trivialization* of the pair $(\mathrm{Bl}_Q P, Q \times \mathbb{P}(P;Q))$ over U; the lemma therefore essentially states that $(\mathrm{Bl}_Q P, Q \times \mathbb{P}(P;Q))$ is *locally trivial* over $\mathbb{P}(P;Q)$.

Now assume our coordinates chosen in such a manner that $U = \mathbb{P}(P,Q)_{S_0}$. We then let $Q' \subseteq P$ be defined by $T_1 = \cdots = T_{c-1} = 0$ so that $[T_0 : T_c : T_{c+1} : \cdots : T_n]$ functions as a homogenenous coordinate system for Q'. It is clear that Q' contains Q as a hyperplane (defined by $T_0 = 0$). Then $p_2^{-1}U = (\mathrm{Bl}_Q P)_{S_0}$ is parametrized by $Q' \times \mathbb{P}(P;Q)_{S_0}$ by means of the morphism

$$([T_0 : T_c : T_{c+1} : \cdots : T_n], [1 : s_1 : \cdots : s_{c-1}]) \in Q' \times U \mapsto$$
$$([T_0 : T_0 s_1 : \cdots : T_0 s_{c-1} : T_c : T_{c+1} : \cdots : T_n], [1 : s_1 : \cdots : s_{c-1}]) \in p_2^{-1}U,$$

This is an isomorphism (the inverse is obvious) which commutes with the projection on U. Since $Q \times U$ is defined by $T_0 = 0$ this isomorphism is the identity on $Q \times U$. All the statements now follow. For example, $Q' \times U$ is smooth and irreducible and hence so is the open subset $p_2^{-1}U$ of $\mathrm{Bl}_Q P$. But then the same is true for $\mathrm{Bl}_Q P$, as it is covered by such open subsets □

Corollary 4.5. Suppose that in the situation of Definition-Lemma 4.4, $Z \subseteq P$ is an irreducible and closed subset such that $Z \cap Q = \emptyset$. Then $\pi_Q|Z : Z \to \mathbb{P}(P;Q)$ is a finite morphism and (so) $\dim Z + \dim Q < \dim P$.

PROOF. We use the notation of Definition-Lemma 4.4. Since p_1 is an isomorphism over $P \smallsetminus Q$, we may identify Z with $p_1^{-1}Z$. Thus Z becomes a closed subset of $\mathrm{Bl}_Q(P)$ which is disjoint with $Q \times \mathbb{P}(P;Q)$ and the projection $\pi_Q|Z$ becomes simply $p_2|Z$. We must show that $p_2|Z$ is finite. This is a local issue on $\mathbb{P}(P;Q)$: we must show that $\mathbb{P}(P;Q)$ can be covered by affine open $U \subset \mathbb{P}(P;Q)$ such that $Z \cap p_2^{-1}U \to U$ is finite. But if we take U as in Lemma 4.4, then this follows from Proposition 2.7 (with Q' identified with $\mathbb{P}^m$ and Q defined by $T_0 = 0$). □

In particular, any linear subspace of P of dimension equal to $\mathrm{codim}(Z)$ must meet Z. On the other hand, for any given Z as in Corollary 4.5, a linear subspace Q of P as in that corollary can always be found:

Proposition 4.6. For every closed subset Z of a projective space P there exists a linear subspace $Q \subseteq P$ of dimension $\mathrm{codim}(Z) - 1$ such that $Q \cap Z = \emptyset$.

PROOF. We may assume that Z is irreducible. We then prove with induction on $i \in \{-1, \ldots, \mathrm{codim}(Z) - 1\}$ that Z misses a linear subspace of dimension i. For $i = -1$, the empty subspace will do. For $i = 0$, we must have $Z \neq P$ and so we can take for our linear subspace any singleton in $P \smallsetminus Z$. When $i > 0$, there exists by induction hypothesis a linear subspace $Q \subseteq P$ of dimension $(i-1)$ which does not meet Z. By Corollary 4.5, $\pi_Q|Z : Z \to \mathbb{P}(P,Q)$ is a finite morphism and so $\dim \pi_Q(Z) = \dim Z < \dim P - i = \dim \mathbb{P}(P,Q)$. Hence there exist a point in $\mathbb{P}(P,Q) \smallsetminus \pi_Q(Z)$. This defines a linear subspace Q' in P of dimension i which passes through Q and misses Z. □

5. Elimination theory and projections

Within a category of reasonable topological spaces (say, the locally compact Hausdorff spaces), the compact ones can be characterized as follows: K is compact if and only if the projection $K \times X \to X$ is closed for every space X in that category. In this sense the following theorem states a kind of compactness property for projective varieties.

Theorem 5.1. Let P be a projective space. Then for any variety X, the projection $\pi_X : P \times X \to X$ is closed.

We derive this theorem from the main theorem of elimination theory, which we state and prove first.

Given an integer $d \geq 0$, let us write V_d for $k[T_0, T_1]_d$, the k-vector space of homogeneous polynomials in $k[T_0, T_1]$ of degree d. The monomials $(T_0^i T_1^{d-i})_{i=0}^d$ form a basis, in particular, $\dim V_d = d + 1$. Given $F \in V_m$ and $G \in V_n$, then

$$u_{F,G} : V_{n-1} \oplus V_{m-1} \to V_{n+m-1}, \quad (A, B) \mapsto AF + BG$$

is a linear map between two k-vector spaces of the same dimension $m + n$. The *resultant* $R(F, G)$ of F and G is defined as the determinant of this linear map with respect to the monomial bases of the summands of $V_{n-1} \oplus V_{m-1}$ and of V_{n+m-1}. So $R(F, G) = 0$ if and only if $u_{F,G}$ fails to be injective. Notice that if $F = \sum_{i=0}^m a_i T_0^i T_1^{m-i}$ and $G = \sum_{j=0}^n b_i T_0^i T_1^{n-i}$, then the matrix of $u_{F,G}$ with respect to the monomial bases is

$$\begin{pmatrix}
a_0 & 0 & 0 & \cdots & 0 & b_0 & 0 & \cdots & \cdots & 0 \\
a_1 & a_0 & 0 & \cdots & 0 & b_1 & b_0 & \cdots & \cdots & 0 \\
a_2 & a_1 & a_0 & \cdots & 0 & b_2 & b_1 & \cdots & \cdots & 0 \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\
a_m & a_{m-1} & * & \cdots & * & * & * & \cdots & \cdots & * \\
0 & a_m & * & \cdots & * & * & * & \cdots & \cdots & * \\
0 & 0 & a_m & \cdots & * & * & * & \cdots & \cdots & * \\
0 & 0 & 0 & \cdots & * & * & * & \cdots & \cdots & * \\
\cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\
0 & 0 & \cdots & \cdots & a_{m-1} & 0 & 0 & \cdots & \cdots & b_{n-1} \\
0 & 0 & 0 & \cdots & a_m & 0 & 0 & 0 & \cdots & b_n
\end{pmatrix}$$

from which we see that its determinant $R(F, G)$ is a polynomial in the coefficients of F and G. So the resultant defines an element of $k[V_m \times V_n] = k[V_m] \otimes k[V_n]$.

Lemma 5.2. $R(F, G) = 0$ if and only if F and G have a common linear factor.

PROOF. If $R(F, G) = 0$, then $u_{F,G}$ is not injective, so that there exist a nonzero $(A, B) \in V_{n-1} \oplus V_{m-1}$ with $AF + BG = 0$. Suppose that $B \neq 0$. It is clear that F divides BG. Since $\deg(B) = m - 1 < m = \deg F$, it follows that F and G must have a common factor.

Conversely, if F and G have a common linear factor L: $F = LF_1$, $G = LG_1$, then $G_1F = F_1G$ and so $(G_1, -F_1) \in V_{n-1} \oplus V_{m-1}$ is a nonzero element of the kernel of $u_{F,G}$. □

PROOF OF THEOREM 5.1. Let $Z \subseteq P \times X$ be closed. It is clear that $\pi_X(Z)$ is closed in X if for every open affine subset $U \subseteq X$, $\pi_X(Z) \cap U$ is closed in U. Since $\pi_X(Z) \cap U = \pi_U(Z \cap (\mathbb{P}^n \times U))$ we may (and will) assume that is X affine. We put $n := \dim P$ and choose a homogeneous coordinate system $[T_0 : \cdots : T_n]$ for P. We proceed with induction on n, starting with the crucial case $n = 1$.

Denote by I_Z the graded ideal in the graded algebra $k[X][T_0, T_1]$ of functions vanishing on Z. Then Z is the common zero set of the members of I_Z (see Exercise 59). For every homogeneous pair $F, G \in \cup_m k[X][T_0, T_1]_m$, we can form their resultant $R(F, G) \in k[X]$ (the coefficients of F and G are regular functions on X and

hence their resultant, which is a polynomial in these coefficients, is also a regular function on X). We claim that $\pi_X(Z)$ is the common zero set $Z(\mathcal{R}) \subseteq X$ of the set of resultants $R(F,G)$ of pairs of homogeneous forms F, G taken in $\cup_m I_{Z,m}$, hence is closed in X.

Suppose that $y \in \pi_X(Z)$. Then $(y,p) \in Z$ for some $p \in \mathbb{P}^1$ and so p is a common zero of each pair F_y, G_y, with $F, G \in \cup_m I_{Z,m}$, where the subscript y refers to substituting y for the first argument. So $R(F,G)(y) = 0$ and hence $y \in Z(\mathcal{R})$.

Next we show that if $y \notin \pi_X(Z)$, then $y \notin Z(\mathcal{R})$. Since $\{y\} \times \mathbb{P}^1$ is not contained in Z, there exists an integer $m > 0$ and a $F \in I_{Z,m}$ with $F_y \neq 0$. Denote by $p_1, \ldots, p_r \in \mathbb{P}^1$ the distinct zeroes of F_y. We show that there exists a $G \in I_{Z,n}$ for some n such that G_y does not vanish in any p_i; this suffices, for this means that $R(F_y, G_y) \neq 0$ and so $y \notin Z(\mathcal{R})$. For any given $1 \leq i \leq r$, $Z \bigcup \cup_{j \neq i} \{(y, p_j)\}$ is closed in $X \times \mathbb{P}^1$, so that there will exist a $G^{(i)} \in \cup_m I_{Z,m}$ with $G^{(i)}_y$ zero in all the p_j with $j \neq i$, but nonzero in p_i. Upon replacing each $G^{(i)}$ by some positive power of it , we may assume that $G^{(1)}, \ldots, G^{(r)}$ all have the same degree n, say. Then $G := G^{(1)} + \cdots + G^{(r)} \in I_{Z,n}$ and $G_y(p_i) = G^{(i)}(p_i) \neq 0$.

Now assume $n \geq 2$. Let $q = [0 : \cdots : 0 : 1]$ and consider the blowup $\tilde{\mathbb{P}}^n := \mathrm{Bl}_{\{q\}} \mathbb{P}^n \to \mathbb{P}^n$. Recall that the projection $\tilde{\mathbb{P}}^n \to \mathbb{P}^{n-1}$ is locally trivial as a $\mathbb{P}^1$-bundle: we can cover $\mathbb{P}^{n-1}$ by affine open subsets U such that over U this is like the projection $\mathbb{P}^1 \times U \to U$. Then the same is true for the projection $\pi_1 : \tilde{\mathbb{P}}^n \times X \to \mathbb{P}^{n-1} \times X$ and so by the case treated above, this projection is closed.

$$\begin{array}{ccccc} \mathbb{P}^n \times X & \xleftarrow{\tilde{\pi}} & \tilde{\mathbb{P}}^n \times X & \hookleftarrow & \mathbb{P}^1 \times U \times X \\ \downarrow{\scriptstyle \pi_X} & & \downarrow{\scriptstyle \pi_1} & & \downarrow \\ X & \xleftarrow{\pi_2} & \mathbb{P}^{n-1} \times X & \hookleftarrow & U \times X \end{array}$$

Denote by $\tilde{\pi} : \tilde{\mathbb{P}}^n \times X \to \mathbb{P}^n \times X$ the projection. Then $\tilde{\pi}^{-1} Z$ is closed and by what we just proved, $\pi_1 \tilde{\pi}^{-1} Z$ is then closed in $\mathbb{P}^{n-1} \times X$. By induction, the image of the latter under the projection $\pi_2 : \mathbb{P}^{n-1} \times X \to X$ is closed. But this is just $\pi_X(Z)$. □

REMARK 5.3. This proof can be adapted to show more, namely that given a closed and irreducible subset $Z \subseteq P \times X$, then for any $x \in \pi_X(Z)$, $Z_x := \{p \in P : (p,x) \in Z\}$ has dimension $\geq \dim Z - \dim \pi_X(Z)$ with equality holding over an open-dense subset of $\pi_X(Z)$.

Here are two corollaries.

Corollary 5.4. Let X be a projective variety. Then for every variety Y, the projection $X \times Y \to Y$ is closed. In particular, every morphism $X \to Y$ is closed (and hence has closed image).

PROOF. Assume that X is closed in $\mathbb{P}^n$. Then $X \times Y$ is closed in $\mathbb{P}^n \times Y$. Since the projection $X \times Y \to Y$ is the composite of the inclusion $X \times Y \subset \mathbb{P}^n \times Y$ and the projection $\mathbb{P}^n \times Y \to Y$ (which is closed by Theorem 5.1), it is closed. Similarly, a morphism $f : X \to Y$ is the composite of the graph morphism $(id_X, f) : X \to X \times Y$ (which is closed) and the closed projection $X \times Y \to Y$ and hence closed. □

It is an elementary result from complex function theory (based on Liouville's theorem) that a holomorphic function on the Riemann sphere is constant. This

implies the corresponding assertion for holomorphic functions on complex projective n-space $\mathbb{P}^n_{\mathbb{C}}$ (to see that a holomorphic function on $\mathbb{P}^n_{\mathbb{C}}$ takes the same value on any two distinct points, simply apply the previous remark to its restriction to the complex projective line passing through them, viewed as a copy of the Riemann sphere). The following corollary is an algebraic version of this fact.

Corollary 5.5. Let X be a projective variety. Then any regular function on X is constant. In particular, any morphism from X to a quasi-affine variety is constant.

PROOF. If $f : X \to Y$ is a morphism to a quasi-affine variety Y, then its composite with an embedding of Y in some affine space $\mathbb{A}^n$ is given by n regular functions on X. So it indeed suffices to prove the special case when $Y = \mathbb{A}^1$. By the previous corollary this image is closed in $\mathbb{A}^1$. But if we think of f as taking its values in $\mathbb{P}^1$ (via the embedding $y \in \mathbb{A}^1 \mapsto [1 : y] \in \mathbb{P}^1$), then we see that $f(X)$ is also closed in $\mathbb{P}^1$. So $f(X)$ cannot be all of $\mathbb{A}^1$ and hence must be finite. Since X is irreducible, it follows that $f(X)$ is a singleton. In other words, f is constant. □

EXERCISE 63. Let P be a projective space of dimension n.

(a) The *dual* $\check{P}$ *of* P is by definition the collection of hyperplanes in P. Prove that $\check{P}$ has a natural structure of a projective space.
(b) Identify the double dual of P with P itself.
(c) The *incidence locus* $I \subseteq P \times \check{P}$ is the set of pairs $(p, q) \in P \times \check{P}$ with the property that p lies in the hyperplane H_q defined by q. Prove that I is a smooth variety of dimension $2n - 1$.
(d) Show that we can find homogeneous coordinates $[Z_0 : \cdots : Z_n]$ for P and $[W_0 : \cdots : W_n]$ for $\check{P}$ such that I is given by $\sum_{i=0}^n Z_i W_i = 0$.

EXERCISE 64. Let $F \in k[X_0, \ldots, X_n]_d$ define a smooth hypersurface H in $\mathbb{P}^n$. Prove that the map $H \to \check{\mathbb{P}}^n$ which assigns to $p \in H$ the projectived tangent space of H at p is given by $[\frac{\partial F}{\partial Z_0} : \cdots : \frac{\partial F}{\partial Z_n}]$. Prove that the image of this map is closed in $\check{\mathbb{P}}^n$ (this image is called *the dual of* H). What can you say in case $d = 2$?

6. The Veronese embeddings

Let be given a positive integer d. We index the monomials in $T_0, \ldots, T_n$ that are homogenous of degree d by their exponents: these are the sequences of non-negative integers $\mathbf{k} = (k_0, \ldots, k_n)$ of length $n + 1$ with sum d. They are $\binom{n+d}{d}$ in number([4]). We use this to label the homogeneous coordinates $Z_{\mathbf{k}}$ of $\mathbb{P}^{\binom{n+d}{d}-1}$.

Proposition 6.1 (The Veronese embedding)**.** The map $f_d : \mathbb{P}^n \to \mathbb{P}^{\binom{n+d}{d}-1}$ defined by $Z_{\mathbf{k}} = T_0^{d_0} \cdots T_n^{d_n}$ is a closed immersion.

PROOF. It is enough to show that for every chart domain $U_{\mathbf{k}} := \mathbb{P}^{\binom{n+d}{d}-1}_{Z_{\mathbf{k}}}$ of the standard atlas of the target space, its preimage $f_d^{-1}U_{\mathbf{k}}$ is open in $\mathbb{P}^n$ and is mapped by f_d isomorphically onto a closed subset of $U_{\mathbf{k}}$. This preimage is defined by $T_0^{k_0} \cdots T_n^{k_n} \neq 0$. Let us renumber the coordinates such that $k_0, \ldots, k_r$ are positive and $k_{r+1} = \cdots = k_n = 0$. Then $f_d^{-1}U_{\mathbf{k}} = \mathbb{P}^n_{T_0 \cdots T_r} \subseteq \mathbb{P}^n_{T_0}$. So if we use the

[4] If we expand $\prod_{i=0}^n (1 - tT_i)^{-1}$ as a power series, we see that the coefficient of t^d is the sum of the monomials in $T_0, \ldots, T_n$ of degree d. So we get the number of such monomials by substituting $T_i = 1$ for all i: it the coefficient of t^d of in $(1 - t)^{-(n+1)}$ and hence the value of $(d/dt)^d(1 - t)^{-(n+1)}/d!$ in $t = 0$, which is $(n + 1)(n + 2) \cdots (n + d)/d! = \binom{n+d}{d}$.

standard coordinates $(t_1, \dots, t_n)$ to identify $\mathbb{P}^n_{T_0}$ with $\mathbb{A}^n$, then $f_d^{-1}U_{\mathbf{k}}$ is identified with $\mathbb{A}^n_{t_1 \cdots t_r}$.

The coordinates on $U_{\mathbf{k}}$ are the functions $Z_{\mathbf{l}}/Z_{\mathbf{k}}$ with $\mathbf{l} \neq \mathbf{k}$. If we write $z_{\mathbf{l}-\mathbf{k}}$ for this function, then f_d is in terms of these coordinates simply:

$$f_d : \mathbb{A}^n_{t_1 \cdots t_r} \cong f_d^{-1}U_{\mathbf{k}} \to U_{\mathbf{k}}, \quad z_{\mathbf{l}-\mathbf{k}} = t_1^{-k_1} \cdots t_r^{-k_r}.t_1^{l_1} \cdots t_n^{l_n},$$

with $(l_1, \dots, l_n)$ running over all the n-tuples of nonnegative integers with sum $\leq d$ and distinct from $(k_1, \dots, k_r, 0, \dots, 0)$. Among the components of this map are $(t_1 \dots t_r)^{-1}$ (take $l_i = k_i - 1$ for $i \leq r$ and $l_i = 0$ for $i > r$) and t_i (take $l_i = k_i + 1$ and $l_j = k_j$ for $j \neq i$; this is allowed because then $l_1 + \cdots + l_n = 1 + k_1 + \cdots + k_n \leq k_0 + k_1 + \cdots + k_n = d$). These generate the coordinate ring $k[t_1, \dots, t_n][1/(t_1 \dots t_r)]$ of $\mathbb{A}^n_{t_1 \cdots t_r}$ and so f_d defines a closed immersion of $\mathbb{A}^n_{t_1 \cdots t_r}$ in $U_{\mathbf{k}}$. □

The following proposition is remarkable for its repercussions in intersection theory.

Proposition 6.2. Let $H \subseteq \mathbb{P}^n$ be a hypersurface. Then $\mathbb{P}^n \smallsetminus H$ is affine and for every closed irreducible subset $Z \subseteq \mathbb{P}^n$ of positive dimension, $Z \cap H$ is nonempty and of codimension ≤ 1 in Z, with equality holding if Z is not contained in H.

PROOF. The hypersurface H is given by a homogeneous polynomial of degree d, say by $\sum_{\mathbf{k}} c_{\mathbf{k}} T_0^{k_0} \cdots T_n^{k_n}$. This determines a hyperplane $\tilde{H} \subseteq \mathbb{P}^{\binom{n+d}{d}-1}$ defined by $\sum_{\mathbf{k}} c_{\mathbf{k}} Z_{\mathbf{k}}$. It is clear that H is the preimage of $\tilde{H}$ under the Veronese morphism and hence the latter identifies $\mathbb{P}^n \smallsetminus H$ with a closed subset of the affine space $\mathbb{P}^{\binom{n+d}{d}-1} \smallsetminus \tilde{H}$. So $\mathbb{P}^n \smallsetminus H$ is affine.

For the rest of the argument we may, by passing to the Veronese embedding, assume that H is a hyperplane. Let c be the codimension of $Z \cap H$ in Z, so that $\dim(H) - \dim(Z \cap H) = (n-1) - (\dim Z - c) = n - \dim Z + c - 1$. By Proposition 4.6 (applied to $Z \cap H \subset H$) there exists then a linear subspace Q of H dimension of $n - \dim Z + c - 2$ which avoids $Z \cap H$. Since is Q also a linear subspace of $\mathbb{P}^n$ which avoids Z, we also have $\dim Q \leq n - \dim(Z) - 1$ by Corollary 4.5. It follows that $c \leq 1$. Clearly, if Z is not contained in H, then $c > 0$. □

REMARK 6.3. A theorem of Lefschetz asserts that if in the situation of Proposition 6.2 above $\dim Z \geq 2$ (so that $\dim(Z \cap H) \geq 1$), then $Z \cap H$ is connected.

EXERCISE 65. Let d be a positive integer. The *universal hypersurface of degree* d is the hypersurface of $\mathbb{P}^n \times \mathbb{P}^{\binom{n+d}{d}-1}$ defined by $F(X, Z) := \sum_{\mathbf{d}} Z_{\mathbf{d}} T_0^{d_0} T_1^{d_1} \cdots T_n^{d_n}$. We denote it by H and let $\pi : H \to \mathbb{P}^{\binom{n+d}{d}-1}$ be the projection. As of item (c) we assume that $d \geq 2$.

(a) Prove that H is smooth.
(b) Prove that projection π is *singular* at (X, Z) (in the sense that the derivative of π at (X, Z) is not a surjection) if and only the partial derivatives of $F_Z \in k[X_0, \dots, X_n]$ have X as a common zero.
(c) Prove that the singular set of π is a smooth subvariety of $\mathbb{P}^n \times \mathbb{P}^{\binom{n+d}{d}-1}$ of codimension $n + 1$.
(d) Prove that the set of $Z \in \mathbb{P}^{\binom{n+d}{d}-1}$ over which π has a singular point is a hypersurface. This hypersurface is called the *discriminant* of π.

(e) For $d = 2$ we denote the coordinates of $\mathbb{P}^{\binom{n+d}{d}-1}$ simply by Z_{ij} (where it is understood that $Z_{ij} = Z_{ji}$). Prove that the discriminant of π is then the zero set of $\det(Z_{ij})$.

7. Grassmannians

Let P be a projective space of dimension n and let $d \in \{0, \dots, n\}$. We want to show that the collection $\mathrm{Gr}_d(P)$ of linear d-dimensional subspaces of P is a smooth projective variety. Let the projective structure on P be defined by the pair (V, ℓ) so that V is a $(n+1)$-dimensional k-vector space and P has been identified with $\mathbb{P}(V)$. This identifies $\mathrm{Gr}_d(P)$ with the collection $\mathrm{Gr}_{d+1}(V)$ of linear $(d+1)$-dimensional subspaces of V.

Lemma 7.1. Let $Q \subseteq P$ be a linear subspace of *codimension* $d+1$. Then the collection $\mathrm{Gr}_d(P)_Q$ of linear d-dimensional subspaces of P contained in $P \smallsetminus Q$ has in a natural manner the structure of an affine space of dimension $(n-d)(d+1)$.

PROOF. The subspace Q determines a linear subspace $V_Q \subseteq V$ of dimension $(n+1) - (d+1) = n - d$ and any $[L] \in \mathrm{Gr}_d(P)_Q$ determines (and is determined by) a linear subspace $V_L \subseteq V$ of dimension $d+1$ with $V_L \cap V_Q = \{0\}$. Since $\dim V_Q + \dim V_L = n+1 = \dim V$, this means that $V_L \oplus V_Q \to V$ is an isomorphism.

We let the vector space $\mathrm{Hom}(V/V_Q, V_Q)$ act on $\mathrm{Gr}_d(P)_Q \cong \mathrm{Gr}_{d+1}(V)$ by stipulating that $\sigma \in \mathrm{Hom}(V/V_Q, V_Q)$ sends V_L to the graph of the map $V_L \subseteq V \to V/V_Q \xrightarrow{\sigma} V_Q$ in $V_L \oplus V_Q \cong V$. This action is simply transitive: given $[L], [L'] \in \mathrm{Gr}_d(P)_Q$, then the projection $V_{L'} \to V/V_Q$ is an isomorphism and so the composite of its inverse $V/V_Q \cong V_{L'}$ with the map $V_{L'} \subset V_V \oplus V_Q \to V_Q$ (the last map is the projection) is the unique element of $\mathrm{Hom}(V/V_Q, V_Q)$ that takes $[L]$ to $[L']$. This makes $\mathrm{Gr}_d(P)_Q$ an affine space over $\mathrm{Hom}(V/V_Q, V_Q)$. It remains to observe that $\dim \mathrm{Hom}(V/V_Q, V_Q) = \dim(V/V_Q) \dim V_Q = (n-d)(d+1)$. □

It can now be shown without much difficulty that $\mathrm{Gr}_d(P)$ admits a unique structure of a variety for which every $\mathrm{Gr}_d(P)_Q$ as in this lemma is affine open and its identification with affine space an isomorphism. We will however proceed in a more direct manner and show in fact that $\mathrm{Gr}_d(P)$ admits the structure of a projective variety.

For this we recall that the exterior algebra $\wedge^\bullet V = \oplus_{p \geq 0} \wedge^p V$ is the quotient of the tensor algebra on V, $\oplus_{p=0}^\infty V^{\otimes p}$ (here $V^{\otimes 0} = k$ by convention), by the two-sided ideal generated by the set of 'squares' $\{v \otimes v\}_{v \in V}$. It is customary to denote the product by the symbol $\wedge$. So we can characterize $\wedge^\bullet V$ as a (noncommutative) associative k-algebra with unit element by saying that is generated by the k-vector space V and is subject to the relations $v \wedge v = 0$ for all $v \in V$. It is a graded algebra ($\wedge^p V$ is the image of $V^{\otimes p}$) and 'graded-commutative' in the sense that if $\alpha \in \wedge^p V$ and $\beta \in \wedge^q V$, then $\beta \wedge \alpha = (-1)^{pq} \alpha \wedge \beta$. If $(\varepsilon_0, \dots, \varepsilon_n)$ is a basis for V, then a basis of $\wedge^p V$ is indexed by the p-element subsets $I \subseteq \{0, \dots, n\}$: if we order the elements of such an I as $0 \leq i_1 < i_2 < \cdots < i_p \leq n$, then to I is associated to the basis element $\varepsilon_I := \varepsilon_{i_1} \wedge \cdots \wedge \varepsilon_{i_p}$ (where the convention is that $\varepsilon_\emptyset := 1 \in k = \wedge^0 V$). So $\dim \wedge^p V = \binom{n+1}{p}$. Notice that $\wedge^{n+1} V$ is one-dimensional and spanned by $\varepsilon_0 \wedge \cdots \wedge \varepsilon_n$, whereas $\wedge^p V = 0$ for $p > n+1$. We also note that if V' and V'' are subspaces of V, then the map

$$\wedge^\bullet V' \otimes \wedge^\bullet V'' \to \wedge^\bullet V, \quad \alpha \otimes \beta \mapsto \alpha \wedge \beta,$$

is a linear map of graded vector spaces which is injective (resp. surjective) when this is so in degree 1, i.e., when $V' \oplus V'' \to V$ is. (The product the left hand side then inherits from the right hand side is $(\alpha_1 \otimes \beta_1).(\alpha_2 \otimes \beta_2) := (-1)^{\deg \beta_1 \deg \alpha_2} \alpha_1 \wedge \alpha_2 \otimes \beta_1 \wedge \beta_2$, where we assume that β_1 and α_2 are homogenenous.)

We say that $\alpha \in \wedge^p V$ is *fully decomposable* if there exist linearly independent $v_1, \ldots, v_p$ in V such that $\alpha = v_1 \wedge \cdots \wedge v_p$. This is equivalent to the existence of a p-dimensional subspace $K \subseteq V$ such that α is a generator of $\wedge^p K$.

Lemma 7.2. For $\alpha \in \wedge^\bullet V$ denote by $K(\alpha)$ the set of $v \in V$ with $v \wedge \alpha = 0$. If $\alpha \in \wedge^p V \smallsetminus \{0\}$, then $\dim K(\alpha) \leq p$ and equality holds if and only if α is fully decomposable and spans $\wedge^p K(\alpha)$.

PROOF. Let $\varepsilon_1, \ldots, \varepsilon_r$ be a basis of $K(\alpha)$ and let $V' \subseteq V$ be a subspace supplementary to $K(\alpha)$ so that $V = K(\alpha) \oplus V'$. Then we have a decomposition

$$\wedge^\bullet V \cong (\wedge^\bullet K(\alpha)) \otimes (\wedge^\bullet V') = \bigoplus_{I \subseteq \{1,\ldots,r\}} \varepsilon_I \otimes (\wedge^\bullet V').$$

The kernel of $\varepsilon_i \wedge : \wedge^\bullet V \to \wedge^\bullet V$ is the subsum of the $\varepsilon_I \otimes (\wedge^\bullet V')$ with $i \in I$. Since $\alpha \in \cap_{i=1}^r \ker(\varepsilon_i \wedge)$, it follows that $\alpha \in \varepsilon_1 \wedge \cdots \wedge \varepsilon_r \otimes (\wedge^{p-r} V')$. Since $\alpha \neq 0$, it follows that $r \leq p$ with equality holding if and only if α is a multiple of $\varepsilon_1 \wedge \cdots \wedge \varepsilon_p$. □

If W is a linear subspace of V of dimension $d+1$, then $\wedge^{d+1} W$ is of dimension 1 and will be thought of as a one dimensional subspace of $\wedge^{d+1} V$. We thus have defined a map $\delta : \mathrm{Gr}_{d+1}(V) \to \mathbb{P}(\wedge^{d+1} V)$, $[W] \mapsto [\wedge^{d+1} W]$. It is called the *Plücker embedding* because of:

Proposition 7.3. Let $0 \leq d \leq \dim V - 1$. Then the map $\delta : \mathrm{Gr}_{d+1}(V) \to \mathbb{P}(\wedge^{d+1} V)$ maps $\mathrm{Gr}_{d+1}(V)$ bijectively onto a closed subset of $\mathbb{P}(\wedge^{d+1} V)$.

PROOF. Let $\alpha \in \wedge^{d+1} V$ be nonzero. According to Lemma 7.2 , $[\alpha]$ is in the image of δ if and only if $K(\alpha)$ is of dimension $d+1$ and if that is the case, then $\delta^{-1}[\alpha]$ has $[K(\alpha)]$ as its unique element. In particular, δ is injective.

The subset $\Sigma_{d+1}(V, \wedge^{d+2} V) \subseteq \mathrm{Hom}(V, \wedge^{d+2} V)$ of linear maps whose kernel is of dimension $\geq d+1$ is (after we have chosen a basis for V) the common zero set of a system of homogeneous equations in $\mathrm{Hom}(V, \wedge^{d+2} V)$, namely the $(n+1-d) \times (n+1-d)$-minors of the corresponding matrices. Consider the linear map

$$\sigma : \wedge^{d+1} V \to \mathrm{Hom}(V, \wedge^{d+2} V), \quad \alpha \mapsto (v \mapsto \alpha \wedge v).$$

Since $\sigma^{-1}\Sigma_{d+1}(V, \wedge^{d+2} V)$ is given by a set of homogeneous equations it defines a closed subset of $\mathbb{P}(\wedge^{d+1} V)$. This is just the image of δ, for by Lemma 7.2, $\sigma^{-1}\Sigma_{d+1}(V, \wedge^{d+2} V) \smallsetminus \{0\}$ is the set of fully decomposable elements of $\wedge^{d+1} V$. □

Proposition 7.3 gives $\mathrm{Gr}_d(P)$ the structure of projective variety. In order to complete the construction, let $Q \subseteq P$ be a linear subspace of codimension d. Let $V_Q \subseteq V$ correspond to Q and choose a generator $\beta \in \wedge^{n-d} V_Q$. Then we have a nonzero linear map to the one-dimensional $\wedge^{n+1} V$:

$$e_\beta : \wedge^{d+1} V \to \wedge^{n+1} V, \quad \alpha \mapsto \alpha \wedge \beta.$$

Its kernel is a hyperplane whose complement defines a principal open subset of $\mathbb{P}(\wedge^{d+1} V)$ that we shall denote by $\mathbb{P}(\wedge^{d+1} V)_Q$. Such principal open subsets cover $\mathbb{P}(\wedge^{d+1} V)$ (to see this, choose a basis $(\varepsilon_0, \ldots, \varepsilon_n)$ of V and observe that if V_Q runs over the codimension d subspaces of V spanned by basis vectors, then $\mathbb{P}(\wedge^{d+1} V)_Q$

runs over a collection of principal open subsets defined by the basis $(\varepsilon_I)_{|I|=d+1}$ of $\wedge^{d+1}V$).

Lemma 7.4. The preimage of $\mathbb{P}(\wedge^{d+1}V)_Q$ under the Plücker embedding δ is the affine space $\mathrm{Gr}_d(P)_Q$ and δ maps this affine space isomorphically onto its image.

PROOF. Let $V_Q \subseteq V$ be the linear subspace defining Q and let β be a generator of $\wedge^{n-d}V_Q$ as above. If $\alpha \in \wedge^{d+1}V$ is fully decomposable and hence generates $\wedge^{d+1}W$ for a unique $(d+1)$-dimensional subspace $W \subseteq V$, then $W \cap V_Q = \{0\}$ if and only if $\alpha \wedge \beta \neq 0$: if $W \cap V_Q$ contains a nonzero vector v then both α and β are divisible by v and so $\alpha \wedge \beta = 0$ and if $W \cap V_Q = \{0\}$, then we have a decomposition $V \cong W \oplus V_Q$ and so $\alpha \wedge \beta \neq 0$. This implies that $\delta^{-1}\mathbb{P}(\wedge^{d+1}V)_\beta = \mathrm{Gr}_d(P)_Q$.

Let us now express the restriction $\delta : \mathrm{Gr}_d(P)_Q \to \mathbb{P}(\wedge^{d+1}V)_Q$ in terms of coordinates. Choose a basis $(\varepsilon_0, \ldots, \varepsilon_n)$ for V such that $(\varepsilon_{d+1}, \ldots, \varepsilon_n)$ is a basis for V_Q and let $\beta := \varepsilon_{d+1} \wedge \cdots \wedge \varepsilon_n$. If $W_0 \subseteq V$ denotes the span of $\varepsilon_0, \ldots, \varepsilon_d$, then $\mathrm{Gr}_d(P)_Q$ is identified with the affine space $\mathrm{Hom}(W_0, V_Q) \cong \mathbb{A}^{(d+1)\times(n-d)}$ of $(d+1) \times (n-d)$-matrices via

$$(a_i^j)_{0\leq i\leq d<j\leq n} \mapsto \text{ } k\text{-span in } V \text{ of the } d+1 \text{ vectors } \{\varepsilon_i + \textstyle\sum_{j=d+1}^n a_i^j \varepsilon_j\}_{i=0}^d\,,$$

so that δ is given by

$$(a_i^j)_{0\leq i\leq d<j\leq n} \mapsto (\varepsilon_0 + \textstyle\sum_{j=d+1}^n a_0^j\varepsilon_j) \wedge \cdots \wedge (e_d + \sum_{j=d+1}^n a_d^j\varepsilon_j).$$

The coefficient of $\varepsilon_{i_0} \wedge \cdots \wedge \varepsilon_{i_d}$ is a determinant of which each entry is 0, 1 or some a_i^j and hence is a polynomial in the matrix coefficients a_i^j. It follows that this restriction of δ is a morphism. Among the components of δ we find the matrix coefficients themselves, for a_i^j appears up to sign as the coefficient of $\varepsilon_0 \wedge \cdots \wedge \widehat{\varepsilon_i} \wedge \cdots \wedge \varepsilon_d \wedge \varepsilon_j$. Since these generate the coordinate ring of $\mathrm{Hom}(W_0, V_Q)$, it follows that δ defines a closed immersion of $\mathrm{Gr}_d(P)_Q$ in $\mathbb{P}(\wedge^{d+1}V)_Q$. □

Corollary 7.5. The Plücker embedding realizes $\mathrm{Gr}_d(P)$ as a smooth irreducible subvariety of $\mathbb{P}(\wedge^{d+1}V)$ of dimension $(n-d)(d+1)$. This structure makes each subset $\mathrm{Gr}_d(P)_Q$ open and isomorphic to affine $(n-d)(d+1)$-space in a way that is compatible with the one obtained in Lemma 7.1.

PROOF. Every two open subsets of the form $\mathrm{Gr}_d(P)_Q$ a have nonempty intersection and so $\mathrm{Gr}_d(P)$ is irreducible. The rest follows from the previous corollary. □

REMARK 7.6. The image of $\mathrm{Gr}_d(P)$ is a closed orbit of the natural $\mathrm{SL}(V)$-action on $\mathbb{P}(\wedge^{d+1}V)$. It lies in the closure of any other $\mathrm{SL}(V)$-orbit([5]).

EXERCISE 66. Let V be a finite dimensional k-vector space. For every linear subspace $W \subseteq V$ we identify $(V/W)^*$ with the subspace of V^* of linear forms on V that are zero on W. Prove that for every $0 \leq r \leq \dim V$ the resulting map $\mathrm{Gr}_r(V) \to \mathrm{Gr}_{\dim V - r}(V^*)$ is an isomorphism of projective varieties.

EXERCISE 67. Let V and W be finite dimensional k-vector spaces and let r be a nonnegative integer $\leq \min\{\dim V, \dim W\}$.

(a) Prove that the subset $\mathrm{Hom}_r(V, W) \subseteq \mathrm{Hom}(V, W)$ of linear maps of rank r is a (locally closed) subvariety of $\mathrm{Hom}(V, W)$.

[5] In representation theory it is shown that $\wedge^{d+1}V$ is an irreducible representation of $\mathrm{GL}(V)$ and that the fully decomposable elements in $\wedge^{d+1}V$ consist of its highest weight vectors.

(b) Prove that the map $\mathrm{Hom}_r(V,W) \to \mathrm{Gr}_{\dim V - r}(V)$ resp. $\mathrm{Hom}_r(V,W) \to \mathrm{Gr}_r(W)$ which assigns to $\phi \in \mathrm{Hom}_r(V,W)$ its kernel resp. image is a morphism.
(c) Prove that the resulting morphism $\mathrm{Hom}_r(V,W) \to \mathrm{Gr}_{\dim V - r}(V) \times \mathrm{Gr}_r(W)$ is trivial over any product of principal open subsets with fiber the general linear group $\mathrm{GL}_r(k)$. Conclude that $\mathrm{Hom}_r(V,W)$ is smooth of codimension $(\dim V - r)(\dim W - r)$.

The Grassmannian of hyperplanes in a projective space is itself a projective space (see Exercise 63). So the simplest example not of this type is the Grassmannian of lines in a 3-dimensional projective space. To see what this is like, let be given a vector space V of dimension 4. On the 6-dimensional space $\wedge^2 V$ we have a homogeneous polynomial $F : \wedge^2 V \to k$ of degree two defined by

$$F(\alpha) := \alpha \wedge \alpha \in \wedge^4 V \cong k$$

(the last identification is only given up to scalar and so the same is true for F). In coordinates F is quite simple: if $e_1, \ldots, e_4$ is a basis for V, then $(e_i \wedge e_j)_{1 \le i < j \le 4}$ is basis for $\wedge^2 V$. So if we label the homogeneous coordinates of $\mathbb{P}(\wedge^2 V)$ accordingly: $[T_{1,2} : \cdots : T_{3,4}]$, then F is given by

$$F(T_{1,2}, \ldots, T_{3,4}) = T_{1,2}T_{3,4} - T_{1,3}T_{2,4} + T_{1,4}T_{2,3}.$$

Notice that F is irreducible. Its partial derivatives are the coordinates themselves (up to sign and order) and so F defines a smooth quadric hypersurface of dimension 4 in a 5-dimensional projective space.

Proposition 7.7. The image of the Plücker embedding of $G_1(\mathbb{P}(V))$ in $\mathbb{P}(\wedge^2 V)$ is the zero set of F.

PROOF. The image of the Plücker embedding is of dimension 4 and so must be a hypersurface. Since the zero set of F is an irreducible hypersurface, it suffices to show that the Plücker embedding maps to the zero set of F. For this, let α be a generator of $\wedge^2 W$ for some linear subspace $W \subseteq V$ of dimension 2. If $e_1, \ldots, e_4$ is a basis of V such that $\alpha = e_1 \wedge e_2$, then it is clear that $\alpha \wedge \alpha = 0$. This proves that the Plücker embedding maps to the zero set of F. □

The smooth quadric hypersurfaces of the same dimension are isomorphic to one another and so this proposition shows that any smooth quadric hypersurface of dimension 4 is isomorphic to the Grassmannian of lines in a three dimensional projective space.

EXERCISE 68. Let P be a projective space dimension 3. Prove for a given $q \in P$, the lines in P through q define a copy of a projective plane $\mathrm{Gr}_1(P;q) \subset \mathrm{Gr}_1(P)$. Prove also that for a given plane $Q \subset P$, the lines in Q define a copy of a projective plane in $\mathrm{Gr}_1(P;Q) \subset \mathrm{Gr}_1(P)$. What is $\mathrm{Gr}_1(P;q) \cap \mathrm{Gr}_1(P;Q)$ like?

REMARK 7.8. The image of the Plücker embedding $\mathrm{Gr}_d(P) \hookrightarrow \mathbb{P}(\wedge^{d+1} V)$ is in fact always the common zero set of a collection of quadratic equations, called the *Plücker relations*. To exhibit these, we first recall that every $\phi \in V^*$ defines a linear 'inner contraction' map $\iota_\phi : \wedge^\bullet V \to \wedge^\bullet V$ of degree -1 characterized by the fact that for $v \in V$, $\iota_\phi(v) = \phi(v) \in k = \wedge^0 V$ and for $\alpha \in \wedge^p V, \beta \in \wedge^\bullet V$, $\iota_\phi(\alpha \wedge \beta) = \iota_\phi(\alpha) \wedge \beta + (-1)^p \alpha \wedge \iota_\phi(\beta)$. Under the natural isomorphism $\mathrm{End}(V,V) \cong V \otimes V^*$, the identity of V defines a tensor in $V \otimes V^*$. The wedge-contraction with this tensor defines a linear map $B_V : \wedge^\bullet V \otimes \wedge^\bullet V \to \wedge^\bullet V \otimes \wedge^\bullet V$ of

bidegree $(1,-1)$. Concretely, if $(e_0,\dots,e_n)$ is a basis of V and $(e_0^*,\dots,e_n^*)$ is the basis of V^* dual to $(e_0,\dots,e_n)$, then

$$B_V(\alpha\otimes\beta) := \textstyle\sum_{r=0}^n (\alpha\wedge e_r)\otimes(\iota_{e_r^*}\beta).$$

Notice that if $W\subseteq V$ is a subspace, then B_W is just the restriction of B_V to $\wedge^\bullet W\otimes\wedge^\bullet W$. So if $\alpha\in\wedge^{d+1}W$ is fully decomposable so that $\alpha\in\wedge^{d+1}W$ for some $(d+1)$-dimensional subspace $W\subseteq V$, then $B_V(\alpha\otimes\alpha)=B_W(\alpha\otimes\alpha)=0$. This is the *universal Plücker relation*.

Conversely, any nonzero $\alpha\in\wedge^{d+1}V$ for which $B_V(\alpha\otimes\alpha)=0$ is fully decomposable. The proof proceeds with induction on n. For $n=0$ there is nothing to show. Assume $n\geq 1$, let $e\in V$ be nonzero and let $V'\subseteq V$ be a hyperplane not containing e. If we write $\alpha=\alpha'+e\wedge\alpha''$ with $\alpha',\alpha''\in\wedge^\bullet V'$, then the component of $B(\alpha\otimes\alpha)$ in $\wedge^\bullet V'\otimes\wedge^\bullet V'$ is $B_{V'}(\alpha'\otimes\alpha')$ and so α' is zero or fully decomposable by our induction hypothesis: there exists a subspace $W'\subseteq V'$ of dimension $d+1$ such that $\alpha'\in\wedge^{d+1}W'$. Then the vanishing of the component of $B(\alpha\otimes\alpha)$ in $\wedge^\bullet V'\otimes e\wedge(\wedge^\bullet V')$ is seen to imply that $\iota_\phi\alpha''=0$ for all $\phi\in(V'/W')^*\subseteq V'^*$. This means that $\alpha''\in\wedge^d W'$. So if we put $M:=ke+W'$, then $\dim M=d+2$ and $\alpha\in\wedge^{d+1}M$. But then $\alpha\in\iota_\phi\wedge^{d+2}M$ for some nonzero $\phi\in M^*$. Then α is a generator of $\wedge^{d+1}\operatorname{Ker}(\phi)$ and hence fully decomposable.

Let us rephrase this in terms of algebraic geometry: every nonzero linear form ℓ on $\wedge^{d+2}V\otimes\wedge^d V$, determines a quadratic form Q_ℓ on $\wedge^{d+1}V$ defined by $\alpha\mapsto\ell(B(\alpha,\alpha))$ whose zero set is a quadratic hypersurface in $\mathbb{P}(\wedge^{d+1}V)$. This hypersurface contains the Plücker locus and the latter is in fact the common zero set of the Q_ℓ, with ℓ running over the linear forms on $\wedge^{d+2}V\otimes\wedge^d V$. It can be shown that the Q_ℓ generate the full graded ideal defined by the Plücker locus. The quadratic forms Q_ℓ are called the Plücker relations([6]).

8. Fano varieties and the Gauß map

The Fano variety of a projective variety is defined in the following proposition.

Proposition-definition 8.1. Let X be a closed subvariety of the projective space P. If d is an integer between 0 and $\dim P$, then the set of d-linear subspaces of P which are contained in X defines a closed subvariety $F_d(X)$ of $\operatorname{Gr}_d(P)$, called the *Fano variety* (of d-planes) of X.

PROOF. An open affine chart of $\operatorname{Gr}_d(P)$ is given by a decomposition $V=W\oplus W'$ with $\dim W=d+1$ and $\dim W'=n-d$ and is then parametrized by $\operatorname{Hom}(W,W')$ by assigning to $A\in\operatorname{Hom}(W,W')$ the graph of A. It suffices to prove that via this identification $F_d(X)$ defines a closed subset of $\operatorname{Hom}(W,W')$.

Choose homogeneous coordinates $[T_0:\cdots:T_n]$ such that W resp. W' is given by $T_{d+1}=\cdots=T_n=0$ resp. $T_0=\cdots T_d=0$. A linear map $\alpha\in\operatorname{Hom}(W,W')$ is then given by $\alpha^*T_{d+i}=\sum_{j=0}^d a_i^jT_j$, $i=1,\dots,n-d$. It defines an element of $F_d(X)$ if and only for all $G\in\cup_{m\geq 0}I_m(X)$, $G(T_0,\dots,T_d,\alpha^*T_{d+1},\dots\alpha^*T_n)$ is identically zero as an element of $k[T_0,\dots,T_d]$. This means that the coefficient of every monomial $T_0^{m_0}\cdots T_d^{m_d}$ in such an expression much vanish. Since this coefficient is a polynomial in the matrix coefficients a_i^j of α, we find that the preimage of $F_d(X)$ in $\operatorname{Hom}(W,V'')$ is the common zero set of a set of polynomials and hence is closed therein. □

EXAMPLE 8.2. Consider the case of a quadratic hypersurface $X\subseteq\mathbb{P}(V)$ and assume for simplicity that $\operatorname{char}(k)\neq 2$. So X can be given by a nonzero quadratic form $F\in k[V]_2$. With F is associated a symmetric bilinear form $B:V\times V\to k$

[6] These show up in the algebro-analytic setting of the Sato Grassmannian (for which both d and $n-d$ are infinity) and are then known as the *Hirota bilinear relations*.

defined by $B(v,v') = F(v+v') - F(v) - F(v')$ so that $B(v,v) = 2F(v)$ (so nonzero, because $\mathrm{char}(k) \neq 2$). Since we have $F(p+tv) - F(v) = tB(p,v) + t^2F(v)$, the partial derivative of F in the v direction is the linear form $B_p \in V^* = k[V]_1$ defined by $v \in V \mapsto B(p,v)$. Let us assume that X is smooth. This means that the partial derivatives of F have no common zero in $\mathbb{P}(V)$. This amounts to $B : V \times V \to k$ being nonsingular in the sense that B_p is zero only when $p = 0$. In other words, $b : p \in V \mapsto B_p \in V^*$ is an isomorphism. A subspace $W \subseteq V$ determines an element of the Fano variety of X precisely when F is zero on W. This implies that B is identically zero on $W \times W$. So b maps W to $(V/W)^* \subseteq V^*$. Since b is injective, this implies that $\dim W \leq \dim(V/W)$, in other words that $\dim W \leq \frac{1}{2}\dim V$.

This condition is optimal. It not difficult to show that we can find coordinates $(T_0, \cdots, T_n)$ such that $F = \frac{1}{2}\sum_{i=0}^n T_iT_{n-i}$ so that $B(v,v') = \sum_{i=0}^n v_iv'_{n-i}$ (the matrix of B is the unit antidiagonal). If for instance $\dim X$ is even, say $2m$ (so that $n = 2m+2$), then let W resp. W' be the linear subspace defined by $T_{m+1} = \cdots = T_{2m+1} = 0$ resp. $T_0 = \cdots = T_m = 0$. Note that $V = W \oplus W'$ and that both $[W]$ and $[W']$ are in $F_m(X)$. The vector space $\mathrm{Hom}(W, W')$ describes an affine open subset of the Grassmannian of m-planes in $\mathbb{P}(V)$. An element $\alpha \in \mathrm{Hom}(W, W')$ is given by $\alpha^*T_{n-i} = \sum_{j=0}^m a_{ij}T_j$, $i = 0, \ldots, m$. The corresponding m-plane is contained in X precisely when $F(T_0, \ldots, T_m, \alpha^*T_{m+1}, \ldots, \alpha_n^*T_n) = \sum_{i,j=0}^m a_{ij}T_iT_j$ is identically zero, i.e., if (a_{ij}) is antisymmetric. It follows that $[W] \in F_m(X)$ has a neighborhood isomorphic to an affine space of dimension $\binom{m+1}{2} = \frac{1}{2}m(m+1)$. In particular, $F_m(X)$ is smooth.

EXERCISE 69. Let X be a quadratic hypersurface $\mathbb{P}(V)$ of odd dimension $2m+1$ and assume that $\mathrm{char}(k) \neq 2$. Prove that $F_{m+1}(X) = \emptyset$ and that $F_m(X) \neq \emptyset$. Prove that $F_m(X)$ is a smooth variety and determine its dimension.

EXERCISE 70. Let $X \subseteq \mathbb{P}^n$ be a hypersurface of degree d and let $0 \leq m \leq n$. Prove that the intersection of $F_m(X)$ with a standard affine subset of $\mathrm{Gr}_m(\mathbb{P}^n)$ is given by $\binom{m+d}{d}$ equations.

EXERCISE 71. Consider the universal hypersurface of degree d in $\mathbb{P}^n$, $H \subseteq \mathbb{P}^n \times \mathbb{P}^{\binom{n+d}{d}-1}$.

(a) For every m-plane $Q \subseteq \mathbb{P}^n$, let Y_z denote the set of $z \in \mathbb{P}^{\binom{n+d}{d}-1}$ for which the corresponding hypersurface H_z contains Q. Prove that Y_z is a linear subspace of $\mathbb{P}^{\binom{n+d}{d}-1}$ of codimension $\binom{m+d}{d}$.

(b) Let $Y \subseteq \mathbb{P}^{\binom{n+d}{d}-1}$ be the set of $z \in \mathbb{P}^{\binom{n+d}{d}-1}$ for which H_z contains an m-plane. Prove that Y is a closed subset of $\mathbb{P}^{\binom{n+d}{d}-1}$ of codimension at most $\binom{m+d}{d} - (m+1)(n-m)$.

(c) Prove that the family of m-planes contained in a generic hypersurface of degree d in $\mathbb{P}^n$ is of dimension $(m+1)(n-m) - \binom{m+d}{d}$ or empty. In particular, this is a finite set when $(m+1)(n-m) = \binom{m+d}{d}$ ([7]).

Let P be a projective space and let X be an irreducible closed subset of P of dimension d. For every smooth point $p \in X$, there is precisely one d-dimensional linear subspace $\hat{T}(X,p)$ of P which contains p and has the same tangent space at

[7] For instance, every cubic surface in $\mathbb{P}^3$ (so here $n = 3$, $d = 3$ and $m = 1$) contains a line. If it is smooth, then it contains in fact exactly 27 lines. This famous result due to Cayley and Salmon published in 1849 is still subject of research.

p as X. In other words, it is characterized by the property that the ideals in $\mathcal{O}_{P,p}$ defining X resp. $\hat{T}(X,p)$ have the same image in $\mathcal{O}_{P,p}/\mathfrak{m}^2_{P,p}$.

Proposition-definition 8.3. The map $G : p \in X_{\mathrm{reg}} \mapsto [\hat{T}(X,p)] \in \mathrm{Gr}_d(P)$ is a morphism, called the *Gauß map*([8]).

PROOF. We assume P identified with $\mathbb{P}(V)$ for some vector space V of dimension $n+1$ so that the Gauß map takes its values in $\mathrm{Gr}_{d+1}(V)$. Let $p_o \in X_{\mathrm{reg}}$ and choose a choose a basis $(T_0,\dots,T_n)$ for V^* such that $p_o \in \mathbb{P}^n_{T_0} \cong \mathbb{A}^n$. We regard X_{T_0} as a closed subset of $\mathbb{A}^n$. Theorem 11.15 of Ch. 1 tells us that there exists a principal neighborhood U of p_o in $\mathbb{P}^n_{T_0}$ and $f_1,\dots,f_{n-d} \in k[U]$ which define $X \cap U$ in U and are such that for all $p \in X \cap U$, T_pX is defined by $df_j(p) = 0$, $j = 1,\dots,n-d$. Then for $p \in X \cap U$, $G(p)$ is the affine subspace of $\mathbb{A}^n$ defined as the common zero set of the $n-d$ affine-linear equations $\sum_{i=1}^n \frac{\partial f_j}{\partial x_i}(p)(t_i - p_i) = 0$, or in homogeneous coordinates, $\sum_{i=1}^n \frac{\partial f_j}{\partial x_i}(p)(T_i - p_iT_0) = 0$ (so this is the intersection of $n-d$ hyperplanes of $\mathbb{P}(V)$).

We can now show that $G|X \cap U$ is a morphism. With the help of Exercise 66 we see that the map which assigns to a $(d+1)$-dimensional subspace of V its annihilator in V^* (which is an $(n-d)$-dimensional subspace of V^*) defines an isomorphism of $\mathrm{Gr}_{d+1}(V)$ onto $\mathrm{Gr}_{n-d}(V^*)$. Via this isomorphism, the Gauß map takes its values in $\mathrm{Gr}_{n-d}(V^*)$ and assigns to $p \in X \cap U$ the span of the $n-d$ linearly independent covectors $\sum_{i=1}^n \frac{\partial f_j}{\partial x_i}(p)(T_i - p_iT_0)$, $j = 1,\dots,n-d$. Composed with the Plücker embedding $\mathrm{Gr}_{n-d}(V^*) \to \mathbb{P}(\wedge^{n-d}V^*)$ this gives the map

$$p \in X \cap U \mapsto \Big[\sum_{i=1}^n \frac{\partial f_1}{\partial x_i}(p)(T_i - p_iT_0) \wedge \cdots \wedge \sum_{i=1}^n \frac{\partial f_{n-d}}{\partial x_i}(p)(T_i - p_iT_0)\Big] \in \mathbb{P}(\wedge^{n-d}V^*).$$

Its coordinates are clearly regular functions on $X \cap U$ and so the associated map $X \cap U \to \mathrm{Gr}_{n-d}(V^*)$ is a morphism. □

REMARK 8.4. The closure of the graph of the Gauss map in $X \times \mathrm{Gr}_d(P)$ is called the *Nash blowup* of X. Its projection to X is clearly an isomorphism over the open-dense subset X_{reg} and hence birational. A remarkable property of the Nash blowup is that the Zariski tangent space of each of its points contains a distinguished d-dimensional subspace (prescribed by the second projection to $\mathrm{Gr}_d(P)$) in such a manner that these subspaces extend the tangent bundle of X_{reg} in a regular manner.

9. Multiplicities of modules

Bézout's theorem asserts that two distinct irreducible curves C, C' in $\mathbb{P}^2$ of degrees d and d' intersect in dd' points. Strictly speaking this is only true if C and C' intersect as nicely as possible, but the theorem is true as stated if we count each point of intersection with an appropriate multiplicity. There is in fact a generalization: the common intersection of n hypersurfaces in $\mathbb{P}^n$ has cardinality the product of the degrees of these hypersurfaces, provided that this intersection is finite and each point of intersection is counted with an appropriate multiplicity. One of our

[8] Thus named because it is related to the map that Gauß studied for an oriented surface Σ in Euclidian 3-space $\mathbb{E}^3$: it is then the map $\Sigma \to \mathbb{S}^2$ which assigns to $p \in \Sigma$ the unit outward normal vector of Σ at p.

aims is to define these multiplicities. The tools from commutative algebra that we use for this have an interest in their own right.

DEFINITION 9.1. We say that an R-module *has length* $\geq d$ if there exist a d-step filtration by submodules $M = M^0 \supsetneq M^1 \supsetneq \cdots \supsetneq M^d = \{0\}$. The *length* of M is the supremum of such d (and so may be ∞).

EXERCISE 72. Suppose R is a noetherian local ring with maximal ideal $\mathfrak{m}$ and residue field K. Prove that the length of a finitely generated R-module M is finite precisely when $\mathfrak{m}^d M = 0$ for some d and is then equal to $\sum_{i=0}^{d-1} \dim_K(\mathfrak{m}^i M/\mathfrak{m}^{i+1}M)$.

Prove that if R is a K-algebra, then this is also equal to $\dim_K(M)$.

In the remainder of this section R is a noetherian ring and M a finitely generated (and hence noetherian) R-module.

Recall that if $\mathfrak{p}$ is a prime ideal of R, then $R_\mathfrak{p}$ is a local ring with maximal ideal $\mathfrak{p}R_\mathfrak{p}$ whose residue field can be identified with the field of fractions of $R/\mathfrak{p}$. We define $M_\mathfrak{p} := R_\mathfrak{p} \otimes_R M$. So this is a $R_\mathfrak{p}$-module.

REMARK 9.2. We can describe $M_\mathfrak{p}$ and more generally, any localization $S^{-1}R \otimes_R M$, as follows. Consider the set $S^{-1}M$ of expressions m/s with $m \in M$ and $s \in S$ with the understanding that $m/s = m'/s'$ if the identity $s''s'm = s''sm$ holds in M for some $s'' \in S$ (so we are considering the quotient of $S \times M$ by an equivalence relation). Then the following rules put on $S^{-1}M$ the structure of a R-module:

$$m/s - m'/s' := (s'm - sm')/(ss'), \quad r \cdot m/s := rm/s.$$

The map $S^{-1}R \times M \to S^{-1}M$, $(r/s, m) \to (rm)/s$ is R-bilinear and hence factors through an R-homomorphism $S^{-1}R \otimes_R M \to S^{-1}M$. On the other hand, the map $S^{-1}M \to S^{-1}R \otimes_R M$, $m/s \mapsto 1/s \otimes_R m$ is also defined: if $m/s = m'/s'$, then $s''(s'm = sm)$ for some $s'' \in S$ and so

$$1/s \otimes_R m = 1/(ss's'') \otimes_R s''s'm = 1/(ss's'') \otimes_R s''sm = 1/s' \otimes_R m.$$

It is an R-homomorphism and it immediately verified that it is a two-sided inverse of the map above. So $S^{-1}R \otimes_R M \to S^{-1}M$ is an isomorphism.

This description shows in particular that if $N \subseteq M$ is a submodule, then $S^{-1}N$ may be regarded as submodule of $S^{-1}M$ (this amounts to: S-localization is an exact functor on the category of R-modules).

DEFINITION 9.3. The *multiplicity* of M at a prime ideal $\mathfrak{p}$ of R, denoted $\mu_\mathfrak{p}(M)$, is the length of $M_\mathfrak{p}$ as an $R_\mathfrak{p}$-module.

In an algebro-geometric context we may modify this notation accordingly. For instance, if we are given an affine variety X and an irreducible subvariety Y, then we may write $\mu_Y(\text{-})$ for $\mu_\mathfrak{p}(\text{-})$, where it is understood that $R = k[X]$ and $\mathfrak{p} = I(Y)$.

REMARK 9.4. Let X be a variety, $x \in X$ and $\mathcal{I} \subseteq \mathcal{O}_{X,x}$ an ideal with $\sqrt{\mathcal{I}} = \mathfrak{m}_{X,x}$. So $\mathfrak{m}_{X,x}^r \subseteq \mathcal{I} \subseteq \mathfrak{m}_{X,x}$ for some positive integer r. Then $\dim_k(\mathcal{O}_{X,x}/\mathcal{I})$ is finite (since $\dim_k(\mathcal{O}_{X,x}/\mathfrak{m}_{X,x}^r)$ is) and according to Exercise 72 equal to the length of $\mathcal{O}_{X,x}/\mathcal{I}$ as an $\mathcal{O}_{X,x}$-module and hence to the multiplicity of $\mathcal{O}_{X,x}/\mathcal{I}$ at the maximal ideal $\mathfrak{m}_{X,x}$. In agreement with our convention, we will denote this multiplicity by $\mu_x(\mathcal{O}_{X,x}/\mathcal{I})$. If X is affine and we are given an ideal $I \subseteq k[X]$ whose image in $\mathcal{O}_{X,x}$ is $\mathcal{I}$, then $\mathcal{O}_{X,x}/\mathcal{I}$ is the localization of $k[X]/I$ at x and so $\mu_x(k[X]/I) = \mu_x(\mathcal{O}_{X,x}/\mathcal{I}) = \dim_k(\mathcal{O}_{X,x}/\mathcal{I})$. Note that x is then an isolated point of $Z(I)$.

If X is smooth at x of dimension n and I has exactly n generators $f_1, \dots, f_n$, then we will see that $\mu_p(\mathcal{O}_{X,x}/(f_1, \dots, f_n)) = \dim_k(\mathcal{O}_{\mathbb{A}^n,p}/(f_1, \dots, f_n))$ can be interpreted as the multiplicity of p as a common zero of $f_1, \dots, f_n$.

We wish to discuss the graded case parallel to the ungraded case. This means that when R is graded, $R = \oplus_{i=0}^{\infty} R_i$, then we assume M to be graded as well, that is, M is endowed with a decomposition as an abelian group $M = \oplus_{i\in\mathbb{Z}} M_i$ such that R_j sends M_i to M_{i+j} (we here do *not* assume that $M_i = 0$ for $i < 0$). For example, a graded ideal in R is a graded R-module. In that case we have the notion of *graded length* of M, which is the same as the definition above, except that we only allow chains of *graded* submodules.

CONVENTION 9.5. Given an integer l and a graded module M over a graded ring, then $M[l]$ denotes the same module M, but with its grading shifted over l, meaning that $M[l]_i := M_{l+i}$.

Note that with convention, if M is homogeneous of degree 0, then $M[l]$ is homogeneous of degree $-l$.

Let us call a (graded) R-module *elementary* if it is isomorphic to $R/\mathfrak{p}$ ($(R/\mathfrak{p})[l]$), for some (homogeneous) prime ideal $\mathfrak{p}$ (and some $l \in \mathbb{Z}$).

Given a (graded) R-module M, then every $m \in M$ ($m \in M_l$) defines a homomorphism or R-modules $r \in R \mapsto rm \in M$. Its kernel is a (graded) ideal of R, the annihilator $\mathrm{Ann}(m)$ of m, so that M contains a copy $R/\mathrm{Ann}(m)$ ($R/\mathrm{Ann}(m)[l]$) as a (graded) submodule.

Lemma 9.6. Let M be a finitely generated nonzero (graded) R-module. Then the collection of annihilators of nonzero (homogeneous) elements of M contains a maximal element and any such maximal element is a (homogeneous) prime ideal of R. In particular, M contains an elementary (graded) submodule.

PROOF. We only do the graded case. The first assertion follows from the noetherian property of R. Let now $\mathrm{Ann}(m)$ be a maximal element of the collection (so with $m \in M$ homogeneous and nonzero). It suffices to show that this is a prime ideal in the graded sense (see Exercise 57), i.e., to show that if $a, b \in R$ are homogeneous and $ab \in \mathrm{Ann}(m)$, but $b \notin \mathrm{Ann}(m)$, then $a \in \mathrm{Ann}(m)$. So $bm \neq 0$ and $a \in \mathrm{Ann}(bm)$. Since $\mathrm{Ann}(bm) \supseteq \mathrm{Ann}(m)$, the maximality property of the latter implies that this must be an equality: $\mathrm{Ann}(bm) = \mathrm{Ann}(m)$, and so $a \in \mathrm{Ann}(m)$. □

Corollary 9.7. Every finitely generated (graded) R-module M can be obtained as a successive extension of elementary modules in the sense that there exists a finite filtration by (graded) R-submodules $M = M^0 \supsetneq M^1 \supsetneq \cdots \supsetneq M^d = \{0\}$ such that each quotient M^j/M^{j+1}, $j = 0, \dots, d-1$, is elementary.

PROOF. We do the graded case only. Since M is noetherian, the collection of graded submodules of M which can be written as a successive extension of elementary modules has a maximal member, M', say. We claim that $M' = M$. If M/M' were nonzero, then it contains an elementary submodule by Lemma 9.6. But then the preimage N of this submodule in M is a successive extension of elementary modules which strictly contains M'. This contradicts the maximality of M. □

The *annihilator* of M, $\mathrm{Ann}(M)$, is the set of $r \in R$ with $rM = 0$. It is clearly an ideal of R. We denote by $\mathcal{P}(M)$ the set of prime ideals of R which contain $\mathrm{Ann}(M)$ and are minimal for that property. According to Proposition 2.17 these are

finite in number and their common intersection equals $\sqrt{\mathrm{Ann}(M)}$ (recall that R is noetherian). In the graded setting, $\mathrm{Ann}(M)$ is a graded ideal and then according to Lemma 2.3 the members of $\mathcal{P}(M)$ are all graded.

Proposition 9.8. In the situation of the preceding proposition, let $\mathfrak{p}^{(j)}$ be the prime ideal of R such that $M^j/M^{j+1} \cong R/\mathfrak{p}^{(j)}$. Then $\mathcal{P}(M)$ is the set of minimal members of the collection $\{\mathfrak{p}^{(j)}\}_{j=0}^{d-1}$ and for every $\mathfrak{p} \in \mathcal{P}(M)$, $\mu_{\mathfrak{p}}(M)$ is finite and $\mathfrak{p}$ occurs precisely $\mu_{\mathfrak{p}}(M)$ times in the sequence $(\mathfrak{p}^{(0)}, \dots, \mathfrak{p}^{(d-1)})$.

PROOF. We first show that $\sqrt{\mathrm{Ann}(M)} = \mathfrak{p}^{(0)} \cap \cdots \cap \mathfrak{p}^{(d-1)}$. If $r \in \mathfrak{p}^{(0)} \cap \cdots \cap \mathfrak{p}^{(d-1)}$, then r maps M^{j-1} to M^j and so $r^d \in \mathrm{Ann}(M)$ and hence $r \in \sqrt{\mathrm{Ann}(M)}$. Conversely, if $r \in R$ and $l \geq 1$ are such that $r^l \in \mathrm{Ann}(M)$, then for all j, $r^l \in \mathfrak{p}^{(j)}$ and hence $r \in \mathfrak{p}^{(j)}$. This proves that $\sqrt{\mathrm{Ann}(M)} = \mathfrak{p}^{(0)} \cap \cdots \cap \mathfrak{p}^{(d-1)}$. Since every prime ideal containing $\mathfrak{p}^{(0)} \cap \cdots \cap \mathfrak{p}^{(d-1)}$ contains some $\mathfrak{p}^{(j)}$ it also follows that $\mathcal{P}(M)$ is the collection of minimal members of $\{\mathfrak{p}^{(j)}\}_{i=0}^{d-1}$.

Fix $\mathfrak{p} \in \mathcal{P}(M)$. An inclusion of R-modules induces an inclusion of $R_{\mathfrak{p}}$-modules (but as we will see, a strict inclusion may become an equality). So we have a filtration $M_{\mathfrak{p}} = M_{\mathfrak{p}}^0 \supseteq \cdots \supseteq M_{\mathfrak{p}}^d = \{0\}$ and $M_{\mathfrak{p}}^j/M_{\mathfrak{p}}^{j+1} \cong R_{\mathfrak{p}}/\mathfrak{p}^{(j)}R_{\mathfrak{p}}$. Either $\mathfrak{p}^{(j)} = \mathfrak{p}$, and then the latter is equal to the residue field $R_{\mathfrak{p}}/\mathfrak{p}R_{\mathfrak{p}}$ and hence of length 1. Or $\mathfrak{p}^{(j)} \neq \mathfrak{p}$, and then we cannot have $\mathfrak{p}^{(j)} \subseteq \mathfrak{p}$ by the minimality of $\mathfrak{p}$. So there exists an $r \in \mathfrak{p}^{(j)} \smallsetminus \mathfrak{p}$. This means that $r/1 \in \mathfrak{p}^{(j)}R_{\mathfrak{p}}$ is invertible so that $\mathfrak{p}^{(j)}R_{\mathfrak{p}} = R_{\mathfrak{p}}$, or equivalently $M_{\mathfrak{p}}^j/M_{\mathfrak{p}}^{j+1} = 0$. Following our definition the first case occurs precisely $\mu_{\mathfrak{p}}(M)$ times. □

We can of course pass from the graded case to the nongraded case by just forgetting the grading. But more interesting is the following construction, which we shall use to pass from a projective setting to an affine one and vice versa. The example to keep in mind is when our graded ring is $k[\mathrm{Cone}(X)]$, with X a closed subset $X \subset \mathbb{P}(V)$. If $\mathfrak{p} \subset k[\mathrm{Cone}(X)]$ is the graded prime ideal which defines a point $p \in X$, we then want to express the local ring $\mathcal{O}_{X,p}$ in terms of $k[\mathrm{Cone}(X)]$ and $\mathfrak{p}$ and show that the multiplicity of a graded $k[\mathrm{Cone}(X)]$-module M at this line is the same as that of an associated $\mathcal{O}_{X,x}$-module $\mathcal{M}_p$ at p. See Corollary 9.9 and Example 9.10 below.

Let $\mathfrak{p} \subseteq R$ be a graded prime ideal and let us write $\mathfrak{m}_{\mathfrak{p}}$ for the maximal ideal $\mathfrak{p}R_{\mathfrak{p}}$ of the ring $R_{\mathfrak{p}}$. Given $l \in \mathbb{Z}$, let $R_{\mathfrak{p},l}$ denote the set of homogeneous fractions of degree l in $R_{\mathfrak{p}}$, i.e., that are representable as r/s with $r \in R_{i+l}$ and $s \in R_i \smallsetminus \mathfrak{p}_i$ for some i and put $R_{\mathfrak{p},\bullet} := \oplus_{l \in \mathbb{Z}} R_{\mathfrak{p},l}$ and $\mathfrak{m}_{\mathfrak{p},\bullet} := \mathfrak{m}_{\mathfrak{p}} \cap R_{\mathfrak{p},\bullet}$. Note that $R_{\mathfrak{p},0} \subseteq R_{\mathfrak{p},\bullet} \subseteq R_{\mathfrak{p}}$ are ring inclusions of which $R_{\mathfrak{p},0}$ and $R_{\mathfrak{p}}$ are local rings (the maximal ideal $\mathfrak{m}_{\mathfrak{p},0}$ of $R_{\mathfrak{p},0}$ is obtained by taking in the previous sentence $r \in \mathfrak{p}_i$), but $R_{\mathfrak{p},\bullet}$ has maximal ideals other than $\mathfrak{m}_{\mathfrak{p},\bullet}$ (see below).

Suppose now that $\mathfrak{p}_1 \neq R_1$ and choose $s \in R_1 \smallsetminus \mathfrak{p}_1$ so that $1/s \in R_{\mathfrak{p},-1}$. Then multiplication with s^l defines an $R_{\mathfrak{p},0}$-module isomorphism of $R_{\mathfrak{p},0} \cong R_{\mathfrak{p},l}$ (the inverse is given by multiplication with s^{-l}). So $R_{\mathfrak{p},\bullet}$ is the ring of Laurent polynomials $R_{\mathfrak{p},0}[s, s^{-1}]$ with $\mathfrak{m}_{\mathfrak{p},\bullet}$ corresponding to $\mathfrak{m}_{\mathfrak{p},0}[s, s^{-1}]$. It follows that $R_{\mathfrak{p},\bullet}/\mathfrak{m}_{\mathfrak{p},\bullet} \cong (R_{\mathfrak{p},0}/\mathfrak{m}_{\mathfrak{p},0})[s, s^{-1}]$. But $R_{\mathfrak{p},\bullet}/\mathfrak{m}_{\mathfrak{p},\bullet}$ is a subring of the residue field $R_{\mathfrak{p}}/\mathfrak{m}_{\mathfrak{p}}$ which clearly has the latter as its field of fractions. So we also have a purely transcendental field extension of residue fields:

$$R_{\mathfrak{p}}/\mathfrak{m}_{\mathfrak{p}} \cong (R_{\mathfrak{p},0}/\mathfrak{m}_{\mathfrak{p},0})(s) \supset R_{\mathfrak{p},0}/\mathfrak{m}_{\mathfrak{p},0}.$$

This also makes sense for any graded R-module M by letting $M_{\mathfrak{p},l}$ be the set of fractions m/s with $m \in M_{i+l}$ and $s \in R_i \smallsetminus \mathfrak{p}_i$ for some i. Note that this is a $R_{\mathfrak{p},0}$-module and that the direct sum $M_{\mathfrak{p},\bullet} := \oplus_l M_{\mathfrak{p},l}$ is equal to $R_{\mathfrak{p},\bullet} \otimes_R M$.

An $R_{\mathfrak{p}}$-module $M_{\mathfrak{p}}/\mathfrak{p}M_{\mathfrak{p}}$ is elementary if and only if it is isomorphic to a shift of the big residue field $R_{\mathfrak{p}}/\mathfrak{m}_{\mathfrak{p}}$. And this is equivalent to $M_{\mathfrak{p},0}/\mathfrak{m}_{\mathfrak{p},0}M_{\mathfrak{p},0}$ being isomorphic to a shift of the small residue field $R_{\mathfrak{p},0}/\mathfrak{m}_{\mathfrak{p},0}$, which simply means that the $R_{\mathfrak{p},0}$-module $M_{\mathfrak{p},0}/\mathfrak{m}_{\mathfrak{p},0}M_{\mathfrak{p},0}$ is elementary.

Corollary 9.9. In this situation (so with the setting noetherian and $\mathfrak{p}_1 \neq R_1$) we have $\mu_{\mathfrak{p}}(M) = \mu_{\mathfrak{m}_{\mathfrak{p},0}}(M_{\mathfrak{p},0})$.

PROOF. An iterated extension $M = M^0 \supsetneq M^1 \supsetneq \cdots \supsetneq M^d = \{0\}$ of M by elementary graded R-modules yields an iterated extension of $M_{\mathfrak{p}}$ resp. $M_{\mathfrak{p},0}$ by trivial or by elementary $R_{\mathfrak{p}}$ resp. $R_{\mathfrak{p},0}$-modules. The corollary then follows from the observation that a successive quotient $M^j_{\mathfrak{p}}/M^{j+1}_{\mathfrak{p}}$ is obtained from $M^j_{\mathfrak{p},0}/M^{j+1}_{\mathfrak{p},0}$ by extension of scalars (from the small residue field to the big one). In particular, $M^j_{\mathfrak{p}}/M^{j+1}_{\mathfrak{p}}$ is nonzero if and only $M^j_{\mathfrak{p},0}/M^{j+1}_{\mathfrak{p},0}$ is. □

We use this observation mainly via the following example.

EXAMPLE 9.10. Let V be a vector space of dimension $n+1$, $J \subseteq k[V]$ a homogeneous ideal and $p \in \mathbb{P}(V)$ an isolated point of the closed subset $Z[J] \subseteq \mathbb{P}(V)$ defined by J. We take here $M := k[V]/J$ and take for $\mathfrak{p}$ the graded ideal $I_p \subseteq k[V]$ defining p. Then $k[V]_{I_p,0}$ can be identified with the local k-algebra $\mathcal{O}_{\mathbb{P}(V),p}$. If $\mathcal{J}_p \subseteq \mathcal{O}_{\mathbb{P}(V),p}$ denotes the ideal corresponding to $J_{I_p,0} \subseteq k[V]_{I_p,0}$, then $\sqrt{\mathcal{J}_p} = \mathfrak{m}_{\mathbb{P}(V),p}$ and we can identify $M_{I_p,0}$ with $\mathcal{O}_{\mathbb{P}(V),p}/\mathcal{J}_p$. According to the above discussion $\mu_{I_p}(k[V]/J) = \mu_p(\mathcal{O}_{\mathbb{P}(V),p}/\mathcal{J}_p)$ and by Exercise 72 this is just $\dim_k(\mathcal{O}_{\mathbb{P}(V),p}/\mathcal{J}_p)$.

10. Hilbert functions and Hilbert polynomials

We shall be dealing with polynomials in $\mathbb{Q}[z]$ which take integral values on integers. Such polynomials are called *numerical*. An example is the *binomial function* of degree $n \geq 0$:

$$\binom{z}{n} := \frac{z(z-1)(z-2)\cdots(z-n+1)}{n!}.$$

It has the property that its value in *any* integer i is an integer, for $i \geq n$ this is an ordinary binomial coefficient and hence an integer, for $i \leq -1$ this is so up to sign, for then we get $(-1)^n\binom{n-1-i}{n}$ and for $0 \leq i \leq n-1$ it is 0.

Let $\Delta : \mathbb{Q}[z] \to \mathbb{Q}[z]$ denote the difference operator: $\Delta f(z) := f(z+1) - f(z)$. This is a $\mathbb{Q}$-linear map with kernel $\mathbb{Q}$ and has the property that it decreases the degree of nonconstant polynomials. It clearly sends numerical polynomials to numerical polynomials and a simple verification shows that it maps $\binom{z}{n+1}$ to $\binom{z}{n}$.

Let us say that a function $f : \mathbb{Z} \to \mathbb{Z}$ is *eventually numerical of degree* d if there exists a $P \in \mathbb{Q}[z]$ of degree d such that $f(n) = P(n)$ for n large enough. It is clear that this P is then unique; we call it the *numerical polynomial* associated to f.

Lemma 10.1. Every $P \in \mathbb{Q}[z]$ which is eventually numerical is in fact numerical and a $\mathbb{Z}$-basis of the abelian group of numerical polynomials is provided by the binomial functions.

If $f : \mathbb{Z} \to \mathbb{Z}$ is a function such that Δf is eventually numerical of degree d, then f is eventually numerical of degree $d+1$, unless f is eventually zero.

PROOF. The first assertion is proved with induction on the degree d of P. If $d = 0$, then P is constant and the assertion is obvious. Suppose $d > 0$ and the assertion known for lower values of d. So $\Delta P(z) = \sum_{i=0}^{d-1} c_i \binom{z}{i}$ for certain $c_i \in \mathbb{Z}$. Then $P(z) - \sum_{i=0}^{d-1} c_i \binom{z}{i+1}$ is in the kernel of Δ and hence is constant. As this expression takes integral values on large integers, this constant is an integer. This proves that P is an integral linear combination of binomial functions.

The proof of the second assertion is similar: let $Q \in \mathbb{Q}[z]$ be such that $Q(i) = \Delta f(i) \in \mathbb{Z}$ for large i. By the preceding, $Q(z) = \sum_i c_i \binom{z}{i}$ for certain $c_i \in \mathbb{Z}$. So if we put $P(z) := \sum_i c_i \binom{z}{i+1}$, then P is a numerical polynomial with $\Delta(f - P)(i) = 0$ for large i. This implies that $f - P$ is constant for large i, say equal to $c \in \mathbb{Z}$. So $f(i) = P(i) + c$ for large i and hence $P + c$ is as required. □

We shall see that examples of such functions are furnished by the Hilbert functions of graded noetherian modules.

REMARK 10.2. A function $f : \mathbb{Z} \to \mathbb{Z}$ which is zero for sufficiently negative integers determines a Laurent series $L_f := \sum_{k\in\mathbb{Z}} f(k)u^k \in \mathbb{Z}((u))$. For the function $k \mapsto \max\{0, \binom{k}{n}\}$ this gives

$$\sum_{k\geq n} \frac{k(k-1)\cdots(k-n+1)}{n!} u^k = \frac{u^n}{n!}\frac{d^n}{du^n}\sum_{k\geq 0} u^k = \frac{u^n}{n!}\frac{d^n}{du^n}\frac{1}{1-u} = \frac{u^n}{(1-u)^n} = \Big(\frac{u}{1-u}\Big)^n.$$

So if we also know that for sufficiently large integers f is the restriction of a polynomial function, then Lemma 10.1 implies that $L_f \in \mathbb{Z}[u][\frac{1}{1-u}]$.

In the remainder of this section V is a k-vector space of dimension $n+1$ (we allow $n = -1$). We equip $k[V]$ with the usual grading (for which each linear form on V has degree one) and view it as the homogeneneous coordinate ring of $\mathbb{P}(V)$. A $k[V]$-module is always assumed to be graded and finitely generated.

Let M be a finitely generated graded $k[V]$-module. Then for every $i \in \mathbb{Z}$, M_i is a finite dimensional k-vector space and so we may define the *Hilbert function* of M, $\phi_M : \mathbb{Z} \to \mathbb{Z}$, by $\phi_M(i) := \dim_k M_i$. For example, the Hilbert function of $k[V]$ itself is $i \mapsto \binom{i+n}{n}$ and so is given by a numerical polynomial of degree n.

The graded ideal $\mathrm{Ann}(M)$ defines a closed subset of $\mathbb{P}(V)$ that is called the *(projective) support* of M and denoted $\mathrm{supp}(M)$. It is clear that if N is a graded submodule of M, then $\dim_k M = \dim_k N + \dim_k(M/N)$ and so we have $\phi_M = \phi_N + \phi_{M/N}$. We also observe $\mathrm{Ann}(N) \cap \mathrm{Ann}(M/N)$ has the same radical as $\mathrm{Ann}(M)$ (in fact, $\mathrm{Ann}(M) \subseteq \mathrm{Ann}(N) \cap \mathrm{Ann}(M/N)$ and the square of $\mathrm{Ann}(N) \cap \mathrm{Ann}(M/N)$ is contained in $\mathrm{Ann}(M)$). It follows that $\mathrm{supp}(M) = \mathrm{supp}(N) \cup \mathrm{supp}(M/N)$.

Theorem-definition 10.3 (Hilbert-Serre)**.** Let M be a graded finitely generated $k[V]$-module. Then ϕ_M is eventually numerical of degree $\dim \mathrm{supp}(M)$ (where we agree that the zero polynomial has the same degree as the dimension of the empty set, namely -1). Its associated numerical polynomial is called the *Hilbert polynomial of* M and denoted $P_M \in \mathbb{Q}[z]$.

PROOF. If N is a graded submodule of M such the theorem holds for N and M/N, then by the observations above, it will hold for M. As M is a successive extension of elementary modules, it therefore suffices to do the case $M = A[l]$, where $A = k[V]/\mathfrak{p}$ with $\mathfrak{p}$ a graded prime ideal. But $\phi_{A[l]}(i) = \phi_A(i + l)$ and since the degree of a polynomial does not change after the substitution $z \mapsto z + l$, we only need to do the case $M = A$.

So now $\operatorname{supp}(A) = Z[\mathfrak{p}]$ is the closed irreducible subset of $\mathbb{P}(V)$ defined by the graded ideal $\mathfrak{p}$. We proceed with induction on $\dim Z[\mathfrak{p}]$. When $\dim Z[\mathfrak{p}] = -1$ (or equivalently, $Z[\mathfrak{p}] = \emptyset$), then $\mathfrak{p} = k[V]_+$ and $A = A_0 = k$, and so $A_i = 0$ for $i > 0$. Hence P_A is identically zero, so of degree -1 by convention.

Suppose therefore $\dim Z[\mathfrak{p}] \geq 0$. Then $\mathfrak{p} \neq k[V]_+$, so that there exists a $T \in k[V]_1 = V^*$ that is not in $\mathfrak{p}_1$. Denote by $V' \subseteq V$ its zero hyperplane. Since $k[V]/\mathfrak{p}$ is a domain, multiplication by T induces an injection $A \to A$ (increasing the degree by one) with cokernel $A' := A/TA$ and so

$$\phi_{A'}(i) = \phi_A(i) - \phi_A(i-1) = \Delta\phi_A(i-1).$$

Since $\operatorname{Ann}(A') = \mathfrak{p} + (T)$, we have $\operatorname{supp}(A') = \operatorname{supp}(A) \cap \mathbb{P}(V')$. According to Proposition 6.2 we then have $\dim \operatorname{supp}(A') = \dim \operatorname{supp}(A) - 1$. Our induction hypothesis tells us that $\phi_{A'}$ is eventually numerical of degree $\dim \operatorname{supp}(A')$. Lemma 10.1 then implies that ϕ_A is eventually numerical of degree $\dim \operatorname{supp}(A') + 1 = \dim \operatorname{supp}(A)$. □

REMARK 10.4. For M as in this theorem we may also form the Laurent series $L_M(u) := \sum_i \dim(M_i)u^i$ (this is usually called the *Poincaré series* of M). It follows from Remark 10.2 and Theorem 10.3 that if $P_M(z) = \sum_{i=0}^d c_i \binom{z+i}{i}$, then $L_M(u) - \sum_{i=0}^d c_i (\frac{u}{1-u})^i \in \mathbb{Z}[u]$.

Lemma 10.1 shows that when P_M is nonzero, then its leading term has the form $c_d z^d/d!$, where d is the dimension of $\operatorname{supp}(M)$ and c_d is a positive integer. This observation leads to a notion of degree (which should not be confused with the degree of P_M):

DEFINITION 10.5. If $d = \dim \operatorname{supp}(M)$, then the *(projective) degree* $\deg(M)$ is $d!$ times the leading coefficient of its Hilbert polynomial (an integer, which we stipulate to be zero in case $\operatorname{supp}(M) = \emptyset$). For a closed subset $Y \subseteq \mathbb{P}(V)$, the Hilbert polynomial P_Y resp. the *degree* $\deg(Y)$ of Y is that of $k[V]/I(Y)$ as a $k[V]$-module.

REMARK 10.6. The homogeneous coordinate ring of a singleton $\{y\} \subset \mathbb{P}(V)$ is isomorphic to $k[T]$ and so we we have $\phi_{\{y\}}(i) = 1$ for all $i \geq 0$. Hence $P_{\{y\}}$ is constant equal to 1. In particular, $\{y\}$ has degree 1. So if $Y \subseteq \mathbb{P}(V)$ is nonempty, and $y \in Y$, then $\phi_Y(i) \geq \phi_{\{y\}}(i) = 1$ for all $i \geq 0$. Hence P_Y is nonzero with positive leading coefficient, and so Y has degree ≥ 1.

EXERCISE 73. Let M be a finitely generated $k[V]$-module and $Y \subset \mathbb{P}(V)$ is projective variety.

(a) Suppose M not of finite length. Prove that there is a unique integer $d \geq 0$ such that $i \mapsto \Delta^d \phi_M(i)$ is a nonzero constant for i sufficiently large. Show that d is the dimension of the support of M and that the constant is its degree.

(b) Prove that if $Y \subset \mathbb{P}(V)$ is projective variety, then there exists a nonempty open subset of linear subspaces $Q \subseteq \mathbb{P}(V)$ of dimension equal to the codimension of Y in $\mathbb{P}(V)$ which meet Y in exactly $\deg(Y)$ points. (This characterization is in fact the classical way of defining the degree of Y.)

EXERCISE 74. Compute the Hilbert polynomial and the degree of

(a) the image of the d-fold Veronese embedding of $\mathbb{P}^n$ in $\mathbb{P}^{\binom{n+d}{n}-1}$,

(b) the image of the Segre embedding of $\mathbb{P}^m \times \mathbb{P}^n$ in $\mathbb{P}^{mn+m+n}$.

EXERCISE 75. Let $Y \subseteq \mathbb{P}^m$ and $Z \subseteq \mathbb{P}^n$ be closed and consider $Y \times Z$ as a closed subset of $\mathbb{P}^{mn+m+n}$ via the Segre embedding. Prove that the Hilbert function resp.

polynomial of $Y \times Z$ is the product of the Hilbert functions resp. polynomials of the factors.

We may now supplement Theorem 10.3 as follows. Let M be as in that theorem: a finitely generated graded $k[V]$-module. Recall that $\mathcal{P}(M)$ denotes the set of minimal prime ideals containing $\mathrm{Ann}(M)$. For every $\mathfrak{p} \in \mathcal{P}(M)$ not equal to $k[V]_+$, the associated closed subset $Z[\mathfrak{p}] \subseteq \mathbb{P}(V)$ is an irreducible component of $\mathrm{supp}(M)$ and all irreducible components of $\mathrm{supp}(M)$ are so obtained. Denote by $\mathcal{P}_o(M)$ the set of $\mathfrak{p} \in \mathcal{P}(M)$ that define an irreducible component of $\mathrm{supp}(M)$ of the same dimension as $\mathrm{supp}(M)$.

Proposition 10.7. Let M be a finitely generated graded $k[V]$-module. Then

$$\deg(M) = \sum_{\mathfrak{p} \in \mathcal{P}_o(M)} \mu_{\mathfrak{p}}(M) \deg(Z[\mathfrak{p}]).$$

PROOF. We write M as an iterated extension by elementary modules: $M = M^0 \supsetneq M^1 \supsetneq \cdots \supsetneq M^d = \{0\}$ with $M^j/M^{j+1} \cong k[V]/\mathfrak{p}^{(j)}[l_j]$. Then $P_M(z) = \sum_{j=0}^{d-1} P_{k[V]/\mathfrak{p}^{(j)}}(z + l_j)$. Now $P_{k[V]/\mathfrak{p}^{(j)}}$ is a polynomial of degree equal to the dimension of $\mathrm{supp}(k[V]/\mathfrak{p}^{(j)}) = Z[\mathfrak{p}^{(j)}] \subseteq \mathbb{P}(V)$. This degree does not change if we replace the variable z by $z + l_j$. So we only get a contribution to the leading coefficient of P_M when $\mathfrak{p}^{(j)} \in \mathcal{P}_o(M)$. For any given $\mathfrak{p} \in \mathcal{P}_o(M)$ this happens by Proposition 9.8 exactly $\mu_{\mathfrak{p}}(M_{\mathfrak{p}})$ times. The proposition follows. □

REMARK 10.8. Note the special case when M has finite support: then $\mathcal{P}_o(M) = \mathcal{P}(M) \smallsetminus \{k[V]_+\}$ and this set is in bijective correspondence with the points of $\mathrm{supp}(M)$. For $\mathfrak{p} \in \mathcal{P}_o(M)$, $Z[\mathfrak{p}] \subseteq \mathbb{P}(V)$ is just a singleton $\{p\}$ and so by Remark 10.6, $\deg(Z[\mathfrak{p}]) = 1$. Furthermore, $\mu_{\mathfrak{p}}(M)$ is the length of $M_{\mathfrak{p}}$ as a $k[V]$-module and this is by Example 9.10 equal to $\dim_k \mathcal{M}_p$, where $\mathcal{M}_p := M_{\mathfrak{p},o}$ is a $\mathcal{O}_{\mathbb{P}(V),p} = k[V]_{\mathfrak{p},o}$-module of finite length. So the above formula then says that $\deg(M) = \sum_{p \in \mathrm{supp}(M)} \dim_k(\mathcal{M}_p)$.

EXERCISE 76. Let $Y \subseteq \mathbb{P}(V)$ be closed. Prove that if $Y_1, \ldots, Y_r$ are the distinct irreducible components of Y of maximal dimension ($= \dim Y$), then $\deg(Y) = \sum_{i=1}^{r} \deg(Y_i)$.

We can now state and prove a result of Bézout type.

Proposition 10.9. Let M be a graded $k[V]$-module and $F \in k[V]_d$ such that F is not a zero divisor in M. Then $\deg(M/FM) = d \deg(M)$.

PROOF. Our assumption implies that the sequence

$$0 \to M(-d) \xrightarrow{\cdot F} M \to M/FM \to 0$$

is exact. This shows that $P_{M/FM}(z) = P_M(z) - P_M(z - d)$. Put $e := \dim \mathrm{supp}(M)$ so that $P_M(z) = \sum_{i=0}^{e} a_i z^i/i!$ with $a_e = \deg(M)$. Since we have $z^i/i! - (z-d)^i/i! = dz^{i-1}/(i-1)!+$ lower order terms, we find that $P_{M/FM}(z) = da_e z^{e-1}/(e-1)!+$ lower order terms. So $\deg(M/FM) = da_e = d \deg(M)$. □

Note the special case for which $M = k[V]$ and F is a generator of the ideal defining a hypersurface $H \subseteq \mathbb{P}(V)$. Then $P_M(z) = \binom{n+z}{n}$ and so the degree of M (which is also the degree of $\mathbb{P}(V)$) is 1 and hence the degree of H is d, just as we would expect. Here is a first corollary.

Corollary 10.10. Let P be a projective plane and $C \subset P$ be an irreducible (closed) curve. Then the degree of C is equal to its number of intersection points with a line $L \subset P$ not contained in C when multiplicities are taken into account: if for every $p \in C \cap L$, $f_p \in \mathfrak{m}_{P,p}$ is a local equation for L at p, then $\deg(C) = \sum_{p \in C} \dim_k(\mathcal{O}_{C,p}/(f_p))$.

PROOF. Choose a homogeneous coordinate system on P. We apply Proposition 10.9: we take for M the homogeneous coordinate ring of C and for F a linear form which defines L. Then the proposition tells us that $\deg(C)$ is the degree of M/FM. The latter is a module with support the finite set $C \cap H$ whose degree is computed with the receipe of Remark 10.8. □

We can now also state:

Theorem 10.11 (Theorem of Bézout)**.** Let $H_i \subseteq \mathbb{P}(V)$ be a hypersurface of degree $d_i > 0$ $(i = 1, \dots n)$, and assume that $Z := H_1 \cap \cdots \cap H_n$ is finite. Each H_i determines at $p \in Z$ a principal ideal in $\mathcal{O}_{\mathbb{P}(V),p}$; denote by $\mathcal{I}_p \subseteq \mathcal{O}_{\mathbb{P}(V),p}$ the sum of these ideals. Then

$$d_1 d_2 \cdots d_n = \sum_{p \in Z} \dim_k(\mathcal{O}_{\mathbb{P}(V),p}/\mathcal{I}_p).$$

Here $\dim_k(\mathcal{O}_{\mathbb{P}(V),p}/\mathcal{I}_p)$ should be interpreted as the intersection multiplicity the hypersurfaces $H_1, \dots, H_n$ at p. So the theorem can be paraphrased as saying that $H_1, \dots, H_n$ meet in $d_1 d_2 \cdots d_n$ points, provided we count each such point with its intersection multiplicity([9]).

We shall need the following result which we state without proof.

***Proposition 10.12.** Let R be a *regular* local ring and let $f_1, \dots, f_r \in \mathfrak{m}_R$ with $r \leq \dim R$ be such that $\dim(R/(f_1, \dots, f_r)) = \dim R - r$. Then the image of f_r in $R/(f_1, \dots, f_{r-1})$ is not a zero divisor.

This implies the homogeneous version we shall need:

Corollary 10.13. Let $r \leq m + 1$ and let $F_1, \dots, F_r$ be homogeneous elements of $k[V]$ of positive degree such that $\dim(k[V]/(F_1, \dots, F_r)) = n + 1 - r$. Then the image of F_r in $k[V]/(F_1, \dots, F_{r-1})$ is not a zero divisor.

PROOF. Let $G \in k[V]$ be such that $F_r G \in (F_1, \dots, F_{r-1})$. The homogeneous components of G have then also this property and so it suffices to see that when G is homogeneous, this implies that $G \in (F_1, \dots, F_{r-1})$. The hypotheses of the above proposition are fulfilled by the local ring $k[V]_{\mathfrak{m}_o}$ and the images of $F_1, \dots, F_r$ therein. So the image of F_r in the quotient $k[V]/(F_1, \dots, F_{r-1})$ is not a zero divisor after localization at o. This means that there exists a $H \in k[V]$ with nonzero constant term such that $GH \in (F_1, \dots, F_{r-1})$. By taking at both sides the homogeneous part of degree equal to the degree of G, we then see that $G \in (F_1, \dots, F_r)$. □

PROOF OF THEOREM 10.11. Choose a definining equation $F_i \in k[V]_{d_i}$ for H_i and put $A^i := k[V]/(F_1, \dots, F_i)$ (so that $A^0 = k[V]$). Then Propositions 10.9 and 10.12 imply that $\deg(A^i) = d_i \deg(A^{i-1})$. Since $\deg A^0 = 1$, it follows that $\deg(A^n) = d_1 d_2 \cdots d_n$. The support of A^n is $H_1 \cap \cdots \cap H_n$ and hence finite. Its degree is then also computed as $\sum_{\mathfrak{p} \in \mathcal{P}_o(A^n)} \mu_{\mathfrak{p}}(A^n)$. But according to Remark 10.8 this is just $\sum_{p \in Z} \dim_k(\mathcal{O}_{\mathbb{P}(V),p}/\mathcal{I}_p)$. □

[9] In the language of schemes, Z is a subscheme of $\mathbb{P}(V)$ whose local ring at $p \in Z$ is $\mathcal{O}_{\mathbb{P}(V),p}/\mathcal{I}_p$.

EXAMPLE 10.14. Assume $\mathrm{char}(k) \neq 2$. We compute the intersection multiplicities of the conics C and C' in $\mathbb{P}^2$ whose affine equations are $x^2 + y^2 - 2y = 0$ and $x^2 - y = 0$. There are three points of intersection: $(0,0)$, $(-1,1)$ and $(1,1)$ (so none at infinity). The intersection multiplicity at $(0,0)$ is the dimension of $\mathcal{O}_{\mathbb{A}^2,(0,0)}/(x^2+y^2-2y, x^2-y)$ as a k-vector space. But $\mathcal{O}_{\mathbb{A}^2,(0,0)}/(x^2+y^2-2y, x^2-y) = \mathcal{O}_{\mathbb{A}^1,0}/(x^4 - x^2) = k[x]/(x^2)$ (for (x^2-1) is invertible in $\mathcal{O}_{\mathbb{A}^1,0}$). Clearly $\dim_k(k[x]/(x^2)) = 2$ and so this is also the intersection multiplicity at $(0,0)$. The intersection multiplicities at $(-1,1)$ and $(1,1)$ are easily calculated to be 1 and thus the identity $2 + 1 + 1 = 2 \cdot 2$ illustrates the Bézout theorem.

REMARK 10.15. If $Y \subseteq \mathbb{P}^n$ is closed, then $P_Y(0)$ can be shown to be an invariant of Y in the sense that it is independent of the projective embedding. In many ways, it behaves like an Euler characteristic. (It is in fact the Euler characteristic of $\mathcal{O}_Y$ in a sense that will become clear once we know about sheaf cohomology.) For example, $P_{Y\times Z}(0) = P_Y(0)P_Z(0)$.

We have seen that for a hypersurface $Y \subseteq \mathbb{P}^n$ of degree $d > 0$, $P_Y(z) = \binom{z+n}{n} - \binom{z-d+n}{n}$ and so $P_Y(0) = 1 - \binom{-d+n}{n} = 1 - (-1)^n\binom{d-1}{n}$. For $n = 2$ (so that Y is a curve), we get $P_Y(0) = 1 - \frac{1}{2}(d-1)(d-2)$. The number $1 - P_Y(0) = \frac{1}{2}(d-1)(d-2)$ is then called the *arithmetic genus* of the curve. If the curve is smooth and $k = \mathbb{C}$, then we may regard it as a topological surface (a Riemann surface) and g is then just the genus of this surface (and so $P_Y(0)$ is half its Euler characteristic). We will encounter this in the next section.

11. Projective curves

In this section we will see that a finitely generated field extension of k of transcendence degree 1 is the function field of a projective curve and that this curve is unique up to unique isomorphism. This explains why of properties of and notions associated with such field extensions admit a complete translation into a geometry. We subsequently discuss notions like divisor, Riemann-Roch theorem and Serre duality in terms that are particular to curves. (These can be given a meaning for all projective varieties.)

Function fields of curves. Let C be a smooth curve. We recall from Corollary Ch. 1, 11.21 that for every $x \in C$, $\mathcal{O}_{C,x}$ is a discrete valuation ring. Its fraction field is of course $k(C)$. We recall from Remark 11.20 of Ch. 1, that thenassociated with $\mathcal{O}_{C,x}$ is a valuation $v_x : k(C)^\times \to \mathbb{Z}$: it assigns to any $f \in k(C)^\times$ its 'order of vanishing' $v_x(f)$ at $x \in C$: this is characterized by the property that f is contained in $\mathfrak{m}_{C,x}^{v_x(f)} \setminus \mathfrak{m}_{C,x}^{v_x(f)+1}$. The function $v_x : k(C)\setminus\{0\} \to \mathbb{Z}$ is a *surjective* homomorphism satifying

$$v_x(f+g) \geq \min\{v_x(f), v_x(g)\}.$$

This property continues to hold if we agree that $v_x(0) = +\infty$.

We first use this notion to prove:

Proposition 11.1. Let C be a nonsingular curve. Then every rational map from C to a projective space is in fact a morphism.

PROOF. Let $f : C \dashrightarrow \mathbb{P}^n$ be a rational map. Let $x \in C$. We prove that f is regular at $x \in C$. Observe that x has an affine open neighborhood in C on which there exist rational functions $f_0, \ldots, f_n$ such that f is there of the form $[f_0 : \cdots : f_n]$. Let $r := \min_i v_x(f_i)$. Choose $t \in \mathfrak{m}_{C,x} \setminus \mathfrak{m}_{C,x}^2$ (a uniformizer:

$v_x(t) = 1$). Upon replacing each f_i with $t^{-r}f_i$ we still represent f near x, but we have now arranged that $\min_i v_x(f_i) = 0$. In other words, each f_i is a regular function on a neighborhood of x and at least one of them takes a nonzero value in x. This just means that f is defined on a neighborhood on x. □

Corollary 11.2. Let K/k be a finitely generated field extension of transcendence degree 1. Then there exists a nonsingular projective curve whose function field is k-isomorphic to K. This curve is unique: if C and C' are nonsingular projective curves, then a k-isomorphism $k(C) \cong k(C')$ is induced by a unique isomorphism $C \cong C'$. In particular (take $C' = C$), the Galois group of K/k can be identified with the automorphism group of C.

We need:

Lemma 11.3. Let $\pi : \tilde{C} \to C$ be a finite morphism of irreducible curves with C projective. Then $\tilde{C}$ is projective.

PROOF. By assumption C admits a finite covering $\{U_i\}_{i=1}^m$ by nonempty affine open subsets such that for $i = 1, \dots, m$, $\pi^{-1}U_i$ is affine and finite over U_i. Let $\pi^{-1}U_i \hookrightarrow \mathbb{A}^{n_i}$ be a closed immersion. By Proposition 11.1 this immersion extends to a morphism $f_i : \tilde{C} \to \mathbb{P}^{n_i}$. It has the property that the preimage of $\mathbb{A}^{n_i} = \mathbb{P}^{n_i}_{T_0}$ is $\pi^{-1}U_i$ so that f_i is a closed immersion over $\mathbb{P}^{n_i}_{T_0}$. Consider the morphism

$$f := (f_1, \dots, f_m) : \tilde{C} \to \textstyle\prod_{i=1}^m \mathbb{P}^{n_i}$$

Its restriction to $\pi^{-1}U_i$ is a closed embedding and lands in the open subset $W_i \subset \prod_i \mathbb{P}^{n_i}$ defined by having the ith component lie in $\mathbb{A}^{n_i} \subset \mathbb{P}^{n_i}_{T_0}$. Then $f^{-1}W_i = \pi^{-1}U_i$ and since these open subsets cover $\tilde{C}$, it follows that f is a closed immersion. The Segre embedding shows that $\prod_i \mathbb{P}^{n_i}$ is projective. Hence $\tilde{C}$ is projective. □

PROOF OF COROLLARY 11.2. The uniqueness assertion is immediate from Proposition 11.1. By Proposition 7.4 of Ch. 1, $K \cong k(C^\circ)$ for some irreducible affine curve. Suppose C° closed in $\mathbb{A}^n$. Let C be its closure in $\mathbb{P}^n$ and denote by $\nu : \hat{C} \to C$ its normalization. This means that $\hat{C}$ is smooth. Since ν is finite, Lemma 11.3 implies that $\hat{C}$ is projective. It remains to observe that both $\hat{C} \to C$ and $C^\circ \subseteq C$ induce k-isomorphisms of function fields. □

REMARK 11.4. If C is as in the preceding corollary, then every $x \in C$ defines a nonarchimedian discrete valuation in $K(C)$ and hence on K. One can prove that the converse is also true: every nonarchimedian discrete valuation on $K(C)$ mapping onto $\mathbb{Z} \cup \{\infty\}$ and constant 1 on $k \smallsetminus \{0\}$ is defined by a point of C. So we could define the underlying point set of a projective curve having K as its function field as the set of such valuations and approach the theory of curves form this angle (as some authors do).

We also note that when an element $f \in K$ is regarded as a rational map $C \dashrightarrow \mathbb{P}^1$, then it is by Proposition 11.1 a morphism $\hat{f} : C \to \mathbb{P}^1$. If f is nonconstant, then this morphism is dominant and since K has transcendence degree one over k, the extension of $K/k(\mathbb{P}^1) = K/k(t)$ must be finite. We denote by $\deg(\hat{f})$ its degree.

In the remainder of this section C and C' denote an irreducible smooth projective curves whose function fields we abbreviate by K its function field.

Divisors on a curve. A *divisor* D on C is a $\mathbb{Z}$-valued function on C whose support $\mathrm{supp}(D)$ (the set of $x \in C$ for which $d_x \neq 0$) is finite: D assigns to every $x \in D$ an integer $d_x \in \mathbb{Z}$ such that $d_x = 0$ for all but a finite number of x. In other words, it is a formal integral linear combination of points of C. Note that the divisors on C form an abelian group under pointwise addition: it is the free abelian group $\mathbb{Z}^{(C)}$ generated by the set of points of C.

We denote the generator defined by $x \in C$ (its characteristic function) by (x), so that we may write the divisor D as $\sum_{x\in C} d_x(x)$, the sum being finite. The integer $\sum_{x\in C} d_x$ is called the *degree* of D, denoted $\deg(D)$ and defines a surjective homomorphism $\deg : \mathbb{Z}^{(C)} \to \mathbb{Z}$. The relation $D \geq D'$ (taken pointwise, so that $d_x \geq d'_x$ for all $x \in C$) defines a partial order on the group of divisors. We say that D is *effective* if $D \geq 0$ and we say that D is *positive* (and write $D > 0$) if $D \geq 0$ and is nonzero. Obviously any divisor D can be written as $D_+ - D_-$ with D_+ and D_- effective. When all the nonzero coefficients of D are 1, we say that D is *reduced*.

Every nonzero $f \in K$ has only finitely many zeroes and poles and hence $x \mapsto v_x(f)$ defines a divisor; we denote it by $\mathrm{div}(f)$. Any divisor thus obtained is called a *principal divisor*. If f is a nonzero constant, then clearly $\mathrm{div}(f) = 0$. The converse also holds: if $\mathrm{div}(f) = 0$, then $\hat{f} : C \to \mathbb{P}^1$ is a morphism which takes its values in $\mathbb{P}^1 \smallsetminus \{0, \infty\}$ and since $\hat{f}$ is closed, $\hat{f}$ must be constant.

Notice that $\mathrm{div}(f/g) = \mathrm{div}(f) - \mathrm{div}(g)$ and so div defines a homomorphism $K^\times \to \mathbb{Z}^{(C)}$ from the (multiplicative) group of units of K to the (additive) group of divisors. The cokernel of this homomorphism is called the *Picard group* of C and denoted $\mathrm{Pic}(C)$. So we have an exact sequence

$$1 \to k^\times \to K^\times \to \mathbb{Z}^{(C)} \to \mathrm{Pic}(C) \to 1.$$

We say that two divisors D and D' on C are *linearly equivalent* (and write $D \equiv D'$) if they have the same image in $\mathrm{Pic}(C)$, in other words, if their difference $D - D'$ is a principal divisor.

This construction is functorial in the following sense. Let be given a morphism $\pi : C \to C'$ of smooth projective curves.

Assume first that π is nonconstant. Since π is closed, this implies that $\pi(C) = C'$, in particular, π is dominant. For $x \in C$, we let

$$e_x(\pi) := \dim_k(\mathcal{O}_{C,x}/\pi^*(\mathfrak{m}_{C',\pi(x)})\mathcal{O}_{C,x})$$

(if $t \in \mathcal{O}_{C',\pi(x)}$ is a uniformizer, then this is also equal to $v_x(\pi^* t)$). We call this the *ramification index* of π at $x \in C$. It is clear that $e_x(\pi) \geq 1$ with equality for all but finitely many $x \in X$. For every $y \in C'$, we denote by $\pi^*(y)$ the divisor on C which assigns to $x \in C$ the value zero unless $x \in \pi^{-1}(y)$, in which case it is equal to $e_x(\pi)$. So $\pi^*(y)$ is a positive divisor with support $\pi^{-1}(y)$. We extend this to a homomorphism $\pi^* : \mathbb{Z}^{(C')} \to \mathbb{Z}^{(C)}$. Notice that then for every $k(C')$, we have $\pi^* \mathrm{div}(f') = \mathrm{div}(\pi^* f')$. It follows that we have an induced homomorphism $\mathrm{Pic}(\pi) : \mathrm{Pic}(C') \to \mathrm{Pic}(C)$.

When π is constant, we let $\mathbb{Z}^{(C')} \to \mathbb{Z}^{(C)}$ and $\mathrm{Pic}(\pi)$ be the zero map.

It is clear from the definitions that when $f \in k(C)$ is nonzero, then $\mathrm{div}(f) = \hat{f}^*(0) - \hat{f}^*(\infty) = \hat{f}^*((0) - (\infty))$.

EXERCISE 77. Prove that any two divisors on $\mathbb{P}^1$ of the same degree are linearly equivalent (so that $\deg$ identifies $\mathrm{Pic}(\mathbb{P}^1)$ with $\mathbb{Z}$).

Proposition 11.5. Let $\pi : C \to C'$ be a finite morphism of smooth irreducible curves so that π^* defines a finite field extension K/K'. Then π^* multiplies the degree of a divisor by the degree of this extension: for every divisor D' on C' we have $\deg(\pi^* D') = [K : K'] \deg(D')$.

For the proof we need the following lemmas.

Lemma 11.6. Let Y be an irreducible smooth affine curve and let M be a finitely generated $k[Y]$-module without torsion: $\mathrm{Ann}(m) = 0$ for all nonzero $m \in M$. Then M is a locally free over Y in the sense that we can cover Y by an affine open subsets V such that $k[V] \otimes_{k[Y]} M$ is a free $k[V]$-module.

PROOF. We find such a neighborhood V of any $y \in Y$ as follows. Let $e_1, \ldots, e_r \in M$ be such that $(e_1, \ldots, e_r)$ projects onto a k-basis of $(\mathcal{O}_{Y,y}/\mathfrak{m}_{Y,y}) \otimes_{k[Y]} M$. By Nakayama's lemma, $e_1, \ldots, e_r$ then generate $\mathcal{O}_{Y,y} \otimes_{k[Y]} M$ as a $\mathcal{O}_{Y,y}$-module.

We claim that these elements are in fact a $\mathcal{O}_{Y,y}$-basis of M. Suppose there eyists a nontrivial relation $\sum_{i=1}^r u_i e_i = 0$ with $u_i \in \mathcal{O}_{Y,y}$ We then choose one with $\min_i v(u_i)$ minimal. Since this relation becomes trivial in $(\mathcal{O}_{Y,y}/\mathfrak{m}_{Y,y}) \otimes_{k[Y]} M$, we must have $v(u_i) \geq 1$ for all i. Choose a uniformizer $t \in \mathfrak{m}$. Then $u_i' := u_i/t \in \mathcal{O}_{Y,y}$. Now $\sum_{i=1}^r u_i' e_i$ is annihilated by t and so it must be zero (for M is torsion free and hence so is its localization $\mathcal{O}_{Y,y} \otimes_{k[Y]} M$). As $\min_i v(u_i') = \min_i v(u_i) - 1$, we get a contradiction.

Now let $m_1, \ldots, m_s \in M$ be $k[Y]$-generators. Then $m_i = \sum_{j=1}^r f_i{}^j e_j$ for certain $f_i{}^j \in \mathcal{O}_{Y,y}$. So if we take for V a neighborhood of y on which each $f_i{}^j$ is regular, then V is as desired. □

The following is a special case of what is known as the *approximation lemma*.

Lemma 11.7. Let U be an irreducible smooth affine curve and $S \subset U$ a finite subset. Then the natural k-algebra homomorphism $k[U] \to \oplus_{s \in S} \mathcal{O}_{U,s}/\mathfrak{m}_{U,s}^n$ is onto for all $n \geq 0$.

PROOF. Given $n \geq 0$ and an element of $\oplus_{x \in S} \mathcal{O}_{U,s}/\mathfrak{m}_{U,s}^n$, then represent the latter by $(f_s \in \mathcal{O}_{U,s})_{s \in S}$. Let S' be the union of S and the set of $x \in U$ for which $v_x(f_s) < 0$ for some $s \in S$. This is a finite set. We then observe that here exists for every $s \in S$, a $\phi_s \in k[U]$ which is zero on $S' \smallsetminus \{s\}$ and 1 in s. We can of course replace ϕ_s by $\phi_s' := 1 - (\phi_s - 1)^2 = 2\phi_s - \phi_s^2$. Since $v_s(\phi_s' - 1) = v_s((\phi_s - 1)^2) > v_s(\phi_s - 1)$ we can, by iterating this, also arrange that in addition that $v_s(\phi_s - 1) \geq n$. For every $m \geq 1$, ϕ_s^m still has this property, but will then also vanish of order $\geq m$ at every point of $S' \smallsetminus \{s\}$.

It follows that for $m \gg n$, $f := \sum_{s \in S} \phi_s^m f_s$ is a regular function on U with the property that it has the same image in $\mathcal{O}_{C,x}/\mathfrak{m}_{C,x}^n$ as f_s for all $s \in S$. □

EXERCISE 78. Show that in the situation of Lemma 11.7, the natural k-algebra homomorphism $k[U \smallsetminus S] \to \oplus_{s \in S} k(U)/\mathfrak{m}_{U,s}^n$ is also onto for all $n \geq 0$.

PROOF OF PROPOSITION 11.5. It is enough to show that for every $y \in C'$, $\pi^*(y)$ has degree $[K : K']$. Choose an open affine neighborhood U' of y in C'. Then $U := \pi^{-1}U'$ is also affine and $k[U]$ is finite over $k[U']$. Lemma 11.6 applied to $Y = U'$ and $M = k[U]$ shows that upon replacing U' be a possibly smaller neighborhood of y we can arrange that $k[U]$ is a free $k[U']$-module of finite rank r, say. Proposition 8.4 of Ch. 1 then implies that $r = [K : K']$.

Bu Lemma 11.7, the natural homomorphism of k-algebras

$$\phi : k[U] \to \oplus_{x\in\pi^{-1}(y)}\mathcal{O}_{C,x}/\mathfrak{m}_y\mathcal{O}_{C,x}.$$

is onto by Lemma 11.7. We claim that $\ker(\phi) = \mathfrak{m}_y k[U]$. The inclusion $\supseteq$ is clear. For the opposite inclusion, choose a uniformizer $t \in \mathfrak{m}_{C',y}$ and let U'_y be an affine neighborhood of y on which t is regular and has no zeroes apart from y. If $f \in \ker(\phi)$, then $v_x(f) \geq v_x(t)$ for all $x \in f^{-1}y$ and hence ft^{-1} is regular on $\pi^{-1}U'_y$. This proves that $\ker(\phi) \subseteq \mathfrak{m}_{C',y} \otimes_{k[U']} k[U]$.

It follows that ϕ induces a k-linear isomorphism

$$\bar{\phi} : (\mathcal{O}_{C,x}/\mathfrak{m}_y) \otimes_{k[U']} k[U] \to \oplus_{x\in\pi^{-1}(y)}\mathcal{O}_{C,x}/\mathfrak{m}_{C',y}\mathcal{O}_{C,x}$$

So r, being the dimension of the right hand, equals the dimension of the left hand side, $\deg \pi^*(y)$. □

Corollary 11.8. Two divisors which are linearly equivalent have the same degree and hence the degree map factors through a homomorphism $\deg : \mathrm{Pic}(C) \to \mathbb{Z}$. Moreover, for any nonconstant $f \in K$, the divisors $\{\hat{f}^*(a)\}_{a\in\mathbb{P}^1}$ have the same degree as $\hat{f}$ and are mutually linearly equivalent.

PROOF. It is enough to show that a principal divisor $\mathrm{div}(f)$ has degree zero. By Proposition 11.5, $\hat{f}^*(0)$ and $\hat{f}^*(\infty)$ have the same degree and so their difference (which is just $\mathrm{div}(f)$) has degree zero. □

The Riemann-Roch theorem. Effective divisors on C arise when we have a morphism $\phi : C \to P$ to a projective space of positive dimension such that the image of ϕ is not contained in a hyperplane of P. Then $\phi(C)$ is of dimension one and any hyperplane $H \subset P$ determines a divisor on C as follows. First note that $\phi^{-1}H$ is a finite set. With every $x \in C$ is associated the intersection number of C and H at x: a local equation f for H at x in $\mathcal{O}_{P,x}$ yields $\phi^*(f) \in \mathcal{O}_{C,x}$ and the intersection number in question is $v_x(\phi^* f)$. This intersection number only depends on C and H and is zero unless $x \in \phi^{-1}H$ (when $\phi(x) \notin H$, then we can take $f = 1$). Denoting this intersection number $(\phi^*H)_x$, then we define the intersection divisor by

$$\phi^*H : x \in C \mapsto (\phi^*H)_x.$$

Lemma 11.9. The map $|\phi|$ which assigns to H the divisor ϕ^*H embeds the projective space $\check{P}$ of hyperplanes in P in the set of positive divisors on C and the image of $|\phi|$ lies in a single linear equivalence class.

In fact, if $H_o \subset P$ is one such hyperplane and L is the vector space of affine-linear functions on the affine space $P \smallsetminus H_o$ (so that $\check{P}$ may be identified with $\mathbb{P}(L)$), then the natural k-linear map $L \to K$ is injective, its image consists of a linear subspace of the k-vector space of rational functions on C with divisor $\geq -\phi^*H_o$, and $|\phi| : \check{P} \cong \mathbb{P}(L) \hookrightarrow {\mathbb{Z}_{\geq 0}}^{(C)}$ is given by the map

$$[f] \in \mathbb{P}(L) \mapsto \phi^*H_o + \mathrm{div}(f).$$

PROOF. $H' \subset P$ is another hyperplane distinct from H, then we can choose a homogeneous coordinate system $[T_0 : \cdots : T_n]$ for P such that $H = Z[T_0]$ and $H' = Z[T_1]$ and $\phi^*(T_1/T_0)$ will be a rational function on C with divisor $\phi^*H' - \phi^*H$ (so that $\phi^*H' \equiv \phi^*H$). To see that $|\phi|$ is injective, suppose $H' \neq H$, but $\phi^*H = \phi^*H'$. Then the divisor of $\phi^*(T_1/T_0)$ is the zero divisor and hence $\phi^*(T_1/T_0)$ is constant, say a on C. This means that $\phi(C)$ would lie in the hyperplane defined by

$T_1 - aT_0 = 0$, thus contradicting our assumption. So $|\phi|$ is injective. This proves the first assertion. The proof of the second assertion is straightforward. □

Note that in the situation of this lemma, the degree of $\phi^* H$ is independent of H (for all such divisors belong to a single linear equivalence class). We refer to this as the *degree* of ϕ.

EXERCISE 79. Prove that L determines ϕ up to a projective-linear transformation: if $(f_0, \dots, f_n)$ is a basis of L, then prove that $[f_0 : \cdots : f_n] : C \to \mathbb{P}^n$ is a morphism which can be identified with ϕ via an isomorphism $\mathbb{P}^n \cong P$.

We now discuss a construction going in the opposite direction. We associate to any divisor D on C two k-vector spaces. The first is

$$L(D) := \{f \in K \,:\, \mathrm{div}(f) \geq -D\}$$

This is indeed a k-vector space. Let us first note that $L(D) = 0$ when $D < 0$: then any element of $L(D)$ is a regular function on C having a zero and hence must be identically zero. We further note that $D \leq D'$ implies $L(D) \subseteq L(D')$. We next prove that $L(D)$ is finite dimensional and give an upper bound for its dimension. Let $S \subset C$ be the finite set of $x \in C$ with $d_x > 0$. Then for any $x \in S$, $f \in L(D)$ has a pole of order at most d_x at x and f has no poles at any point of $C \smallsetminus S$. So the obvious k-linear map

$$L(D) \to \oplus_{x \in S} \mathfrak{m}_{C,x}^{-d_x} / \mathcal{O}_{C,x}$$

has a kernel consisting of regular functions on C. By Corollary 5.5 such functions must be constant and so this kernel has dimension 1. Hence

$$\dim_k L(D) \leq 1 + \textstyle\sum_{x \in C} \dim_k(\mathfrak{m}_{C,x}^{-d_x} / \mathcal{O}_{C,x}) = 1 + \sum_{x \in C} d_x$$

is finite as asserted. The Riemann-Roch theorem to which we are heading provides among other things a lower bound for $\dim_k L(D)$.

Proposition-definition 11.10. A *complete linear system on* C is the set of positive divisors in a linear equivalence class of divisors on C. It has naturally the structure of a projective space: if D is a (not necessarily positive) divisor on C, then the complete linear system it defines (which we shall denote by $|D|$) can be identified as a projective space with $\mathbb{P}(L(D))$. A *linear system on* C is a subset defined by linear subspace of a complete linear system.

PROOF. If $\dim L(D) > 0$ and $(f_0, \dots, f_n)$ is a basis of $L(D)$, then we can use its elements to define a rational map $\phi := [f_0 : \cdots : f_n] : C \dashrightarrow \mathbb{P}^n$. By Proposition 11.1 this is in fact a morphism. Note that $\phi(C)$ is not contained in a hyperplane.

The image of $|\phi|$ consists of all the positive divisors in the linear equivalence class of D: if $f \in L(D) \smallsetminus \{0\}$, then clearly $D + \mathrm{div}(f)$ is effective and every positive member D' of $|D|$ is so obtained. Clearly, $D + \mathrm{div}(f)$ does not change if we replace f by af for some $a \in k^\times$. This is in fact the only ambiguity, for if $D + \mathrm{div}(f) = D + \mathrm{div}(f')$, then $0 = \mathrm{div}(f') - \mathrm{div}(f) = \mathrm{div}(f'/f)$, meaning that f'/f has neither poles nor zeros, so that $f'/f \in k^\times$. This proves that we have a bijection $|D| \cong \mathbb{P}(L(D))$. This structure of a projective space is indeed independent of D: if $D' \equiv D$, so that $D' = D + \mathrm{div}(g)$ for some $g \in K^\times$, then the composite map $\mathbb{P}(L(D')) \cong |D| \cong \mathbb{P}(L(D))$ is induced by the k-linear isomorphism $L(D') \xrightarrow{\cdot g} L(D)$ given by multiplication with g and hence is projective linear isomorphism. □

REMARK 11.11. So a complete linear system on C is naturally a projective space, but as the proof of the above proposition already suggests, there is no canonically defined vector space of which it is the projectivization: a divisor D in the associated linear equivalence class must be chosen to obtain such a vector space. This is why we defined the notion of a projective space as in 1.1.

The definition of the other vector space is somewhat more involved. Central here is the following notion.

DEFINITION 11.12. The ring of *repartitions of* C, which we denote by $\mathcal{K}_C$, is the subring of K^C consisting of the $\mathbf{g} = (g_x \in K)_{x\in C} \in K(C)$ with $g_x \in \mathcal{O}_{C,x}$ for all but a finite number of $x \in C$. We write $\mathcal{O}_C \subset \mathcal{K}_C$ for the subring $\prod_{x\in C} \mathcal{O}_{C,x}$.

The map which assigns to $f \in K$ the element of $\mathcal{K}_C$ that has f in each factor is well-defined (for f has only finitely many poles) and is a ring homomorphism; it is clearly injective and so we regard K as a subring of $\mathcal{K}_C$. This makes $\mathcal{K}_C$ a K-algebra and hence a K-vector space (not of finite dimension).

For every divisor D, we define $\mathcal{O}_C(D) \subset \mathcal{K}_C$ as the set of $\mathbf{g}$ for which $v_x(g_x) \geq -d_x$ for all x (so $\mathcal{O}_C(0) = \mathcal{O}_C$). This is an $\mathcal{O}_C$-submodule of $\mathcal{K}_C$. When $D' \leq D$, then $\mathcal{O}_C(D') \subseteq \mathcal{O}_C(D)$. As k-linear subspaces of $\mathcal{K}_C$, the $\mathcal{O}_C(D)$ are not finite dimensional, but they are still commensurable in size. For example if $D \geq 0$, then

$$\mathcal{O}_C(D)/\mathcal{O}_C = \oplus_{x\in \mathrm{supp}(D)} \mathfrak{m}_{X,x}^{-d_x}/\mathcal{O}_{C,x},$$

which is of (finite) dimension $\deg(D)$. Since every element of $\mathcal{K}_C$ is in some $\mathcal{O}_C(D)$, we have $\mathcal{K}_C = \cup_D \mathcal{O}_C(D)$. In particular, $\mathcal{K}_C/\mathcal{O}_C$ is the *direct sum* $\oplus_{x\in C} K/\mathcal{O}_{C,x}$ (the "space of polar parts on C") and likewise $\mathcal{K}_C/\mathcal{O}_C(D) = \oplus_{x\in C} K/\mathfrak{m}_{C,x}^{-d_x}$.

Observe that $L(D) = K \cap \mathcal{O}(D)$. We put

$$I(D) := \mathcal{K}_C/\big(K + \mathcal{O}_C(D)\big).$$

Note that this is also the cokernel of the diagonal map $\mathrm{Coker}\left(K \to \oplus_{x\in C} K/\mathfrak{m}_{C,x}^{-d_x}\right)$. In particular, $I(0)$ equals $\mathrm{Coker}\left(K \to \oplus_{x\in C} K/\mathcal{O}_{C,x}\right)$, that is, the space of polar parts on C modulo those that are realized by a rational function.

EXERCISE 80. Let D and D' be divisors on C such that $D - D' = \mathrm{div}(f)$ for some nonzero $f \in K$. Prove that multiplication by f defines k-linear isomorphisms $L(D) \cong L(D')$, $\mathcal{O}_C(D) \cong \mathcal{O}_C(D')$ and $I(D) \cong I(D')$.

Our first goal is to show that $I(0)$ is finite dimensional. To this end we consider for a finite nonempty subset $S \subset C$ the natural map

$$\mathcal{O}(C \smallsetminus S) \to \oplus_{x\in S} K/\mathcal{O}_{C,x}.$$

and denote by $\mathcal{P}(S)$ its cokernel (we shall later see that $C \smallsetminus S$ is affine, so that we can also write $k[C \smallsetminus S]$ for $\mathcal{O}(U \smallsetminus S)$). If $S' \subseteq S$, then the preimage of the subsum $\oplus_{x\in S'} K/\mathcal{O}_{C,x} \subseteq \oplus_{x\in S} K/\mathcal{O}_{C,x}$ in $\mathcal{O}(C \smallsetminus S)$ is just $\mathcal{O}(C \smallsetminus S')$. This implies that $\mathcal{P}(S')$ naturally embeds in $\mathcal{P}(S)$. It is also clear that $\cup_S \mathcal{P}(S)$ (or rather $\varinjlim_S \mathcal{P}(S)$) equals $I(0)$, where the union (limit) is taken over all finite subsets of C.

Lemma 11.13. Let $\phi : C \hookrightarrow P$ be a closed immersion in a projective space and $H \subset P$ a hyperplane which does not contain the image of ϕ so that $S = \phi^{-1}H$ is finite. If $P_\phi(z)$ the Hilbert-Serre polynomial of $\phi(C)$, then $\mathcal{P}(S)$ has the (finite) dimension $1 - P_\phi(0)$.

PROOF. The map $\mathcal{O}(C \smallsetminus S) \to \oplus_{x \in S} K/\mathcal{O}_{C,p}$ has as its kernel the functions regular on C and hence consists of the constants. We regard $P \smallsetminus H$ as an affine space which contains $C \smallsetminus S$ as a closed subset (so that $C \smallsetminus S$ is affine). We observe that we may identify the degree r-part of the homogeneous coordinate ring of C with the k-vector space $k[C \smallsetminus S]_r$ of regular functions on $C \smallsetminus S$ that are the restriction of a degree r polynomial on $P \smallsetminus H$. Then for every $r \geq 0$, the above map restricts to an injection

$$k[C \smallsetminus S]_r/k \to \mathcal{O}_C(r\phi^* H)/\mathcal{O}_C.$$

The dimension of the target space is $\deg(r\phi^* H) = r \deg(\phi(C))$. We know that $P_\phi(z)$ is of degree $\dim \phi(C) = 1$ and has leading coefficient $\deg(\phi(C))$ so that $P_\phi(z) = \deg(\phi(C))z + P_\phi(0)$. For r large enough, $\dim_k k[C \smallsetminus S]_r = P_\phi(r)$ and so the cokernel of this map has dimension $1 - P_\phi(0)$. This proves that these cokernels stabilize to $\mathcal{P}(S)$ and that $\dim \mathcal{P}(S) = 1 - P_\phi(0)$. □

Corollary-definition 11.14 (Genus of a curve)**.** For S as in the previous lemma, $\mathcal{P}(S) \to I(0)$ is an isomorphism. In particular, $I(0)$ is of finite dimension $1 - P_\phi(0)$. We call this integer the *genus* of C and denote it by $g(C)$.

PROOF. Since $I(0) = \varinjlim_S \mathcal{P}(S)$, it suffices to show that for every finite $S' \supseteq S$, the natural map $\mathcal{P}(S) \hookrightarrow \mathcal{P}(S')$ is an isomorphism. Given such an S', choose a hypersurface in P which contains $\phi(S')$, but does not contain $\phi(C)$. If r is the degree of this hypersurface, then let $\phi' : C \hookrightarrow P'$ be the r-fold Veronese embedding so that our hypersurface determines a hyperplane H' in P' which does not contain $\phi'(S)$ and for which $S' \subseteq \phi'^{-1} H'$. Then we have injections $\mathcal{P}(S) \hookrightarrow \mathcal{P}(S') \hookrightarrow \mathcal{P}(\phi'^{-1} H')$. Now the Hilbert-Serre polynomial for ϕ' is given by $P_{\phi'}(z) = P_\phi(rz)$ and so $\dim \mathcal{P}(\phi'^{-1} H') = 1 - P_{\phi'}(0) = 1 - P_\phi(0) = \dim \mathcal{P}(S)$. It follows that $\mathcal{P}(S) \hookrightarrow \mathcal{P}(S')$ is an isomorphism. □

EXAMPLE 11.15. We observed in Remark 10.15 that for a smooth plane curve C of degree d, $P_C(0) = 1 - (d-1)(d-2)/2$ and so $g(C) = (d-1)(d-2)/2$.

EXERCISE 81. Prove that $g(\mathbb{P}^1) = 0$.

Corollary 11.16 (Riemann-Roch)**.** For every divisor D on C, $I(D)$ is finite dimensional and we have

$$\dim L(D) - \dim I(D) = 1 - g(C) + \deg(D).$$

In other words, the dimension of the complete linear system defined by a divisor D is equal to $\dim I(D) - g(C) + \deg(D)$.

PROOF OF COROLLARY 11.16. We compare the situation for two divisors D and D' which differ in the simplest possible way: $D' = D + (x)$ for some $x \in C$. We claim that we have an exact sequence of k-vector spaces

$$0 \to L(D) \to L(D') \to \mathcal{O}_C(D')/\mathcal{O}_C(D) \to I(D) \to I(D') \to 0.$$

First note that the term in the middle is a k-vector space of dimension one: if d_x is the coefficient of x in D, then it is just $\mathfrak{m}_{C,x}^{-d_x - 1}/\mathfrak{m}_{C,x}^{-d_x}$. Let us now indicate what the maps are and establish at the same time exactness. The obvious map $L(D) \to L(D')$ is fact an inclusion. It is an equality unless there exists an $f' \in L(D')$ with $v_x(f') = -d_x - 1$ (and then the obvious map $L(D') \to \mathcal{O}_C(D')/\mathcal{O}_C(D)$ is onto).

On the other hand, $\mathcal{O}_C(D')/\mathcal{O}_C(D)$ is also the kernel of $\mathcal{K}_C/\mathcal{O}_C(D) \to \mathcal{K}_C/\mathcal{O}_C(D')$. So we have an evident surjection

$$I(D) = \mathrm{Coker}\left(K \to \mathcal{K}_C/\mathcal{O}_C(D)\right) \longrightarrow \mathrm{Coker}\left(K \to \mathcal{K}_C/\mathcal{O}_C(D')\right) = I(D').$$

whose kernel can be identified with the image of $\mathcal{O}_C(D')/\mathcal{O}_C(D)$ in the left hand side. This image is zero precisely when $\mathcal{O}_C(D') \subset \mathcal{O}_C(D) + K$, i.e., when there exists a $f' \in K \cap \mathcal{O}_C(D') = L(D')$ such that $v_x(f') = -d_x - 1$. Thus our claim follows.

Returning to the proof of the corollary, we note that it is trivially true for $D = 0$. The alternating sum of the dimensions of the terms of a finite exact sequence of finite dimensional vector spaces is zero and so $\dim L(D) - \dim I(D) = \dim L(D') - \dim I(D') - 1$. Since we can restate this as $\dim L(D) - \dim I(D) - \deg(D) = \dim L(D) - \dim I(D) - \deg(D)$, this assertion proves that the Riemann-Roch formula holds for D if and only if it holds for D'. This reduces the assertion to the trivial case $D = 0$. □

REMARK 11.17. We may also regard $\mathcal{O}(D)$ as an abelian sheaf of rational (k-valued) functions on C. For such a sheaf there are defined (sheaf) cohomology groups and $L(D)$ resp. $I(D)$ has then the interpretation of $\mathrm{H}^0(C, \mathcal{O}_C(D))$ resp. $\mathrm{H}^1(C, \mathcal{O}_C(D))$. A very general theorem asserts that the cohomology groups of such sheaves on a projective variety X are finite dimensional k-vector spaces and vanish in degree $> \dim X$. So the Riemann-Roch theorem can be understood as a formula for the alternating sum of the Betti numbers (the Euler characteristic) of the sheaf $\mathcal{O}_C(D)$ and this is the way it has been generalized to an arbitrary projective variety X with (what is called) a *coherent* $\mathcal{O}_X$-module taking the place of $\mathcal{O}_C(D)$.

Residues. The Riemann-Roch theorem becomes much more effective when we use an interpretation of the k-dual of $I(D)$ in terms of the differentials on C. Let us first recall that we have universal k-derivation $d : K \to \Omega_{K/k}$. The target $\Omega_{K/k}$ is a K-vector space of dimension one: if $t \in K$ is such that K is a finite separable extension of $k(t)/k$, then dt is nonzero as an element of $\Omega_{K/k}$ and generates it as K-vector space: indeed, if $\phi \in K$ has minimal polynomial $F = x^n + a_1 x^{n-1} + \cdots + a_n \in k(t)[x]$ (so with $a_i \in k(t)$), then F is separable so that $F'(\phi) \neq 0$ and from

$$0 = d(F(\phi)) = F'(\phi)d\phi + \sum_{i=1}^{n} a_i'(t)x^{n-i}dt.$$

it then follows that $d\phi \in K dt$.

For every $x \in C$ we also have a universal k-derivation $d : \mathcal{O}_{C,x} \to \Omega_{\mathcal{O}_{C,x}/k} =: \Omega_{C,x}$. The universal property of the latter makes that we have a natural homomorphism of $\mathcal{O}_{C,x}$-modules $\Omega_{C,x} \to \Omega_{K/k}$ which extends $d : K \to \Omega_{K/k}$. This homomorphism is nonzero and since $\Omega_{C,x}$ is free $\mathcal{O}_{C,x}$-module of rank one, the resulting K-linear map $K \otimes_{\mathcal{O}_{C,x}} \Omega_{C,x} \to \Omega_{K/k}$ is an isomorphism of K-vector spaces of dimension one. So every $\alpha \in \Omega_{K/k}$ defines a differential on an open-dense subset of C. If $\alpha \neq 0$, we define the *order* $v_x(\alpha)$ of α at x by the property that $\alpha \in \mathfrak{m}_{C,x}^{v_x(\alpha)}\Omega_{C,x} \smallsetminus \mathfrak{m}_{C,x}^{v_x(\alpha)+1}\Omega_{C,x}$. We have a divisor $\mathrm{div}(\alpha)$ on C defined by

$$\mathrm{div}(\alpha) : x \in C \mapsto v_x(\alpha).$$

Any divisor thus obtained is called a *canonical divisor* for C.

Given a divisor D, we define the k-vector space

$$\Lambda(D) := \{\alpha \in \Omega_{K/k} \mid \mathrm{div}(\alpha) \geq -D\}.$$

This is also a finite dimensional vector space, for if $\alpha_o \in \Omega_{K/k}$ is nonzero, then $\Lambda(D)$ is the space of $f\alpha_o$ with $\operatorname{div}(f) \geq -\operatorname{div}(\alpha_o) - D$ and hence we can identify $\Lambda(D)$ with $L(\operatorname{div}(\alpha_o) + D)$. Note that $\Lambda := \Lambda(0)$ is the space of regular differentials on C. It is clear that when $D \geq D'$, then $\Lambda(D) \supseteq \Lambda(D')$ and that $\Omega_{K/k} = \cup_D \Lambda(D)$.

Let $\pi : C \to C'$ be a finite morphism between smooth projective curves. This gives rise to the finite field extension K/K', for which we have defined the trace $\operatorname{Tr}_{K/K'} : K \to K'$. This is K'-linear map which assigns to $f \in K$ the trace of the endomorphism of the (finite dimensional) K'-vector space K defined by multiplication with f. This trace has a counter part for differentials: we have a K'-linear map $\operatorname{Tr}_{K/K'} : \Omega_{K/k} \to \Omega_{K'/k}$ characterized by the property that if $\alpha' \in \Omega_{K'/k}$ and $f \in K$, then $\operatorname{Tr}_{K/K'}(f\pi^*\alpha') = \operatorname{Tr}_{K/K'}(f)\alpha'$. It is easy to check that this is unique: if α' is nonzero, then so is $\pi^*\alpha'$ and hence generates $\Omega_{K/k}$ as a K-vector space. So every $\alpha \in \Omega_{K/k}$ can then be written as $f\pi^*\alpha'$ with $f \in K$. To see this is well-defined, suppose $\beta' \in \Omega_{K'/k}$ is another nonzero element. Then $\beta' = u\beta$ for a unique $u \in K'$ and so if $\alpha = g\pi^*\beta'$, then $f = gu$ and hence

$$\operatorname{Tr}_{K/K'}(f)\alpha' = \operatorname{Tr}_{K/K'}(gu)\alpha' = u \operatorname{Tr}_{K/K'}(g)\alpha' = \operatorname{Tr}(g)\beta'.$$

In what follows we need the residue map, that is, a k-linear map $\operatorname{Res}_x : \Omega_{K/k} \to k$, which has all the properties familiar in the complex case in terms of a uniformizer $t \in \mathfrak{m}_{C,x} \smallsetminus \mathfrak{m}_{C,x}^2$: if we write $\alpha \equiv (\sum_{i=0}^{N} a_{-i}t^{-i})dt/t \pmod{\Omega_{C,x}}$, then $\operatorname{Res}_x(\alpha) = a_0$. That this is well-defined when $k = \mathbb{C}$ is a consequence of the Cauchy residue formula, but such an argument is not immediately available in positive characteristic. The most elegant approach, due to Tate, is explained below. We here show (following Serre [2]) how we can reduce the general case to the complex case. It is based on the fact that an element of $\mathbb{Z}[x_1, \ldots, x_n]$ is zero if and only if the associated function $\mathbb{C}^n \to \mathbb{C}$ has that property.

Proposition 11.18. There is a unique k-linear map $\operatorname{Res}_x : \Omega_{K/k} \to k$ such that

(i) $\operatorname{Res}_x \alpha = 0$ when $\alpha \in \Omega_{C,x}$,
(ii) for any $f \in K^\times$ and $n \in \mathbb{Z}$, $\operatorname{Res}_x f^{n-1}df$ is zero unless $n = 0$, in which case we get the image of $v_x(f) \in \mathbb{Z}$ in k (so zero when $v_x(f)$ is a multiple of the characteristic of k).

PROOF. We fix a uniformizer t of $\mathcal{O}_{C,x}$ as above. If $\alpha \equiv (\sum_{i=0}^{N} a_{-i}t^{-i})dt/t \pmod{\Omega_{C,x}}$, then (i) and (ii) imply that $\operatorname{Res}_x(\alpha) = a_0$. So the uniqueness of the residue is clear. We must show that if Res_x is defined by this formula, then it satisfies (ii), so that this definition is independent of the choice of t.

Let $f \in K^\times$ and let us we write $f = ut^r$ with $r = v_x(f)$ so that $v_x(u) = 0$. After multiplying f with a scalar, we may assume that $u \equiv 1 \pmod{\mathfrak{m}_{C,x}}$. This ensures that if u_l denotes the lth Taylor coefficient of u in its formal expansion: $u \equiv 1 + u_1 t + \cdots + u_l t^l \pmod{\mathfrak{m}_{C,x}^{l+1}}$, then the lth Taylor coefficient of u^{-1} is a universal polynomial with *integral* coefficients in $u_1, \ldots, u_l$. Now

$$f^{n-1}df = (r + u'/u)u^n t^{rn}\frac{dt}{t}$$

and so we need to verify that the coefficient of t^{-rn} in $(r + u'/u)u^n$ is zero for $n \neq 0$ and r for $n = 0$. Now note that we obtain this coefficient as some polynomial

$A \in \mathbb{Z}[U_1, \dots, U_{|rn|}]$ evaluated in $(u_1, \dots, u_{|rn|})$ and then mapped to k. So it suffices to show that A is constant r or 0 according to whether or not $n = 0$. But by the Cauchy residue formula, this is the case when we regard A as a function $\mathbb{C}^{|rn|} \to \mathbb{C}$ (these $|rn|$ coefficients of u can be arbitrary) and so the same is true for A as an element of $k[U_1, \dots, U_{|rn|}]$. □

Corollary 11.19. Let $\pi : C \to C'$ be a finite separable morphism between smooth projective curves. Then for any $y \in C'$ we have

$$\sum_{x \in \pi^{-1}(y)} \operatorname{Res}_x = [K : K'] \operatorname{Res}_y \operatorname{Tr}_{K/K'}$$

(as an identity of maps $\Omega_{K/k} \to k$).

PROOF. Let t be a uniformizer of $\mathcal{O}_{C,y}$. Then any element of $\Omega_{K'/k}$ is as a k-vector space generated by $\Omega_{C',y}$ and the differentials $t^{-n-1}dt$, $n \geq 0$. So we only need to verify this assertion for an element of $\Omega_{C',y}$ and for $t^{-n-1}dt$. The above proposition shows that in each of these cases all terms in our equality vanish, except for $t^{-1}dt$. In that case, the residue of $t^{-1}dt$ in y equals 1, whereas (by 11.18) its residue in $x \in \pi^{-1}(y)$ equals by $v_x(\pi^* t) = e_x(\pi)$. By Proposition 11.5, the sum $\sum_{x \in \pi^{-1}(y)} e_x(\pi)$ equals $[K : K']$ and so the corollary follows. □

For $\alpha \in \Omega_{K/k}$, $\operatorname{Res}_x \alpha$ can only be nonzero if α has a pole at x. So we can form the (finite) sum $\sum_{x \in C} \operatorname{Res}_x \alpha$. We have:

Theorem 11.20 (Residue theorem)**.** For every $\alpha \in \Omega_{K/k}$, $\sum_{x \in C} \operatorname{Res}_x \alpha = 0$.

PROOF. Choose $t \in K$ such that K is finite an separable over $k(t)$. This corresponds to a finite separable morphism $\pi : C \to \mathbb{P}^1$. Corollary 11.19 shows that if the theorem holds for $\mathbb{P}^1$, then it holds for C. But for $\mathbb{P}^1$ this is easy: the forms $(t - t_0)^{n-1}dt$ ($t_0 \in \mathbb{A}^1$ and $n \in \mathbb{Z}$) span $\Omega_{k(t)/k}$ as a k-vector space (we leave this as an exercise) and it is clearly true for them: for $n \neq 0$ all residues are zero and for $n = 0$ we get only two residues: a residue 1 in t_0 and a residue -1 at ∞, which indeed add up to zero. □

A more satisfactory proof, due to Tate [3], is given below.

The residue according to Tate. In this discussion k can actually be any field, it is only when we apply this to curves (which are assumed to be defined over k) that we assume it to be algebraically closed.

Let V be a k-vector space. We say that a k-linear map $f : V \to V$ is *finipotent* if for some $n \geq 0$, $f^n V$ is finite dimensional. Such a map has a trace: since the sequence $\{f^n V\}_{n \geq 0}$ is nonincreasing, it becomes stationary, and we let $\operatorname{Tr}_V(f)$ be the trace of the map f has on the finite dimensional $V_f := \cap_{n \geq 0} f^n V$ (or on any other other f-invariant finite dimensional subspace $V' \subset V$ for which f is nilpotent on V/V'). Some of the usual properties of the trace continue to hold for finipotent endomorphisms. For example, if $f : V \to V$ is finipotent and $W \subset V$ is a f-invariant k-linear subspace, then f induces finipotent maps in W and V/W and we have $\operatorname{Tr}_V(f) = \operatorname{Tr}_W(f) + \operatorname{Tr}_{V/W}(f)$. Also, if $\phi : V \to V'$ and $\phi' : V \to V'$ are k-linear maps of k-vector spaces, and $\phi'\phi : V \to V$ is finipotent, then so is $\phi\phi' : V' \to V'$ and $\phi\phi'$ has the same trace as $\phi'\phi$. This is because ϕ induces an isomorphism $V_{\phi'\phi} \to V'_{\phi\phi'}$ with inverse induced by ϕ'.

More generally, we say that a k-linear subspace $E \subset \operatorname{End}_k(V)$ is finipotent if for some $n \geq 0$, $f_n \cdots f_1 V$ is finite dimensional for all n-tuples $(f_1, \cdots, f_n) \in E^n$. The trace then

defines a k-linear map $E \to k$ and, as with the usual trace, we have $\mathrm{Tr}([f,g]) = 0$ when $f, g \in E$. Tate's approach to the residue exploits the fact that in certain situations $[f,g]$ is finipotent, but may have nonzero trace (so that neither fg nor gf is finipotent).

We wish to consider k-linear subspaces of V up to finite dimensional k-linear subspaces of V: given two such subspaces A, A', we say that A is *not much bigger than* A' (and we write $[A] \leq [A']$) if $A/(A \cap A')$ is finite dimensional. If both $[A] \leq [A']$ and $[A'] \leq [A]$, we say that A and A' are *about the same*. This is clearly an equivalence relation and so if $[A]$ is understood to mean the equivalence class of A, then $\leq$ is indeed a partial order on the set of equivalence classes.

Suppose given a k-linear subspace $A \subset V$. Let E^A resp. E_A denote the space of k-linear maps $f : V \to V$ with $[fV] \leq [A]$ resp. $[fA] \leq [0]$. Then $E^A_A := E_A \cap E^A$ is finipotent and hence the linear form $\mathrm{Tr} : E^A_A \to k$ is defined. We claim that $E^A + E_A$ is the space $E(A)$ of k-linear maps $V \to V$ with $[fA] \leq [A]$. This is easily seen if we choose a supplement A' for A in E: $V = A \oplus A'$, and write any $f \in \mathrm{End}_k(V)$ as a matrix with respect to this decomposition: $f = \left(\begin{smallmatrix} f_{00} & f_{01} \\ f_{10} & f_{11} \end{smallmatrix}\right)$. Then the condition $f \in E^A$ resp. $f \in E_A$ means that the bottom row $\left(\begin{smallmatrix} f_{10} & f_{11} \end{smallmatrix}\right)$ resp. the left column $\left(\begin{smallmatrix} f_{00} \\ f_{10} \end{smallmatrix}\right)$ has finite rank, whereas $f \in E(A)$ means that the bottom left entry f_{10} is of finite rank. It is clear that $E(A)$ is a (possibly noncommutative) k-subalgebra of $\mathrm{End}_k(V)$ which contains E_A and E^A as two-sided ideals. This implies that when $(f,g) \in E^A \times E_A$ or $(f,g) \in E(A) \times E^A_A$, fg and gf lie in E^A_A and will have the same trace.

If $f, g \in E(A)$ happen to commute, then we can take this one step further: Since $E(A) = E_A + E^A$, $E^A \cap (f + E_A) \neq \emptyset$. If we choose $f^A \in E^A \cap (f + E_A)$, then it is clear that $[f^A, g] \in E^A$. But we also have $[f^A, g] \in [f + E_A, g] \subseteq [E_A, g] \subseteq E_A$ and so $[f^A, g] \in E^A_A$. This means that $\mathrm{Tr}_V([f^A, g]) \in k$ is defined. In case $f \in E_A$, we have $f^A \in E^A_A$ and hence $\mathrm{Tr}_V([f^A, g]) = 0$. So $\mathrm{Tr}_V([f^A, g])$ is well-defined, with the dependence on f via its coset $f + E_A$. We will write $\mathrm{Tr}^V_A([f,g])$ for this trace. We may interchange the roles of f and g and conclude that $\mathrm{Tr}^V_A([f,g])$ only depends on the pair of cosets $(f + E_A, g + E_A)$ and is antisymmetric.

It is easy to check that the dependence on A is through $[A]$: changing A by a finite dimensional subspace of V does not change $\mathrm{Tr}^V_A([f,g])$. In particular, $\mathrm{Tr}^V_A([f,g]) = 0$ when A has finite dimension or finite codimension. In the same vein, V hardly matters in the sense that if have a k-linear subspace $V' \subseteq V$ which contains A, and V' is preserved by f and g, then $\mathrm{Tr}^{V'}_A([f,g]) = \mathrm{Tr}^V_A([f,g])$. We may therefore write $\mathrm{Tr}_{[A]}([f,g])$ instead.

Now let R be a commutative k-subalgebra of $E(A)$ (so that $A \subset V$ is an *almost R-submodule*, in the sense that $[fA] \leq [A]$ for all $f \in R$).

Proposition 11.21. For $f, g \in R$, $\mathrm{Tr}_{[A]}([f,g])$ only depends on $fdg \in \Omega_{R/k}$ and hence defines a linear map $\mathrm{Res}_{[A]} : \Omega_{R/k} \to k$. We have $\mathrm{Res}_{[A]}(fdg) = 0$ when f and g preserve A.

PROOF. Recall that $\Omega_{R/k}$ can be obtained as the quotient of $R \otimes_k R$ by the k-linear subspace spanned by the tensors $f \otimes gh - fg \otimes h - fh \otimes g$. So all we need to do is to check that the corresponding identity holds for $\mathrm{Tr}_{[A]}$. Choose f^A, g^A, h^A as above. Then $f^A g^A \in E^A \cap (fg + E_A)$ (and similar for the other products). The required property then follows from the identity $[f^A, g^A h^A] = [f^A g^A, h^A] + [h^A f^A, g^A]$.

When $fA \subseteq A$ and $gA \subseteq A$, then if $\pi \in \mathrm{End}_k(V)$ is a projection onto A, we may take $f^A = f\pi$ and $g^A = g\pi$, so that $[f^A, g^A] = 0$ and hence $\mathrm{Tr}^V_A([f,g]) = 0$. □

It is clear from the definition that if $n \geq 0$, then for any $f \in R$, $\mathrm{Res}_{[A]}(f^n df) = 0$. So if f is invertible in R, then also $\mathrm{Res}_{[A]}(f^{-n-2} df) = \mathrm{Res}_{[A]}(-f^{-n} d(f^{-1})) = 0$. It is an amusing exercise to show that $\mathrm{Res}_{[A]}\, f^{-1} df = \mathrm{Tr}_{[A]}(f, f^{-1}) = \dim_k(A/A \cap fA) - \dim_k fA/(A \cap fA)$.

EXAMPLE 11.22. Let V be a field containing k and let $v : V^\times \to \mathbb{Z}$ be a (surjective) valuation such that the associated DVR $A \subset L$ defined by $v \geq 0$ has k as its residue field. Then A is an almost V-module and we have therefore defined a k-linear map $\mathrm{Res}_{[A]} : \Omega_{V/k} \to k$. We verify the two basic properties that characterize it as a residue:

(i) $\mathrm{Res}_{[A]}$ is zero on the image of $\Omega_{A/k} \hookrightarrow \Omega_{V/k}$,
(ii) for any $f \in V^\times$ and $n \in \mathbb{Z}$, $\mathrm{Res}_{[A]} f^{n-1}df$ is zero unless $n = 0$, in which case we get $v(f)$.

All of this has already been established, except for the assertion that $\mathrm{Res}_{[A]}(df/f) = v(f)$. But since $A \cap fA = \mathfrak{m}^{\max\{0,v(f)\}}$, we indeed get

$$\mathrm{Res}_{[A]}(df/f) = \dim_k \left(A/\mathfrak{m}^{\max\{0,v(f)\}}\right) - \dim_k \left(\mathfrak{m}^{v(f)}/\mathfrak{m}^{\max\{0,v(f)\}}\right) = v(f).$$

Of special interest here is the case when $A = \mathcal{O}_{C,x}$ and (hence) $V = K$. In view of Proposition 11.18 this returns our old residue map: we have $\mathrm{Res}_x = \mathrm{Res}_{[\mathcal{O}_{C,x}]} : \Omega_{K/k} \to k$.

With this definition in hand, it is not hard to prove the residue theorem.

PROOF OF THE RESIDUE THEOREM. Our R is here K and our ambient K-module V will be $\mathcal{K}_C$. First we show that $\sum_{p\in C} \mathrm{Res}_p = \mathrm{Res}_{[\mathcal{O}_C]}$. Let $f, g \in K$ and consider $\alpha = fdg \in \Omega_{K/k}$. We assume $\alpha \neq 0$ and let $S \subset C$ be a finite set which contains the poles of f and g. Put $U := C \smallsetminus S$ and write $\mathcal{K}_C$ as a finite direct sum of K-modules $\mathcal{K}_C = K^S \oplus \mathcal{K}_U$ so that $\mathcal{O}_C = \oplus_{x\in S}\mathcal{O}_{C,x} \oplus \mathcal{O}_U$. It follows that $\mathrm{Res}_{[\mathcal{O}_C]} = \mathrm{Res}_{[\mathcal{O}_U]} + \sum_{p\in C} \mathrm{Res}_p$. Since f and g preserve the summand $\mathcal{O}_U \subset \mathcal{K}_U$, we find that $\mathrm{Res}_{[\mathcal{O}_C]}\, \alpha = \sum_{p\in C} \mathrm{Res}_p\, \alpha$.

So it now remains to show that $\mathrm{Res}_{[\mathcal{O}_C]} = 0$. Since $K \cap \mathcal{O}_C$, being the space of regular functions of C, is equal to k, we have an exact sequence

$$0 \to k \to K \oplus \mathcal{O}_C \to K + \mathcal{O}_C \to 0.$$

This shows that $\mathrm{Res}_{[K+\mathcal{O}_C]} + \mathrm{Res}_{[k]} = \mathrm{Res}_{[K\oplus\mathcal{O}_C]} = \mathrm{Res}_{[K]} + \mathrm{Res}_{[\mathcal{O}_C]}$. But the terms distinct from $\mathrm{Res}_{[\mathcal{O}_C]}$ all vanish: $\mathrm{Res}_{[k]} = 0$ because k is finite dimensional, $\mathrm{Res}_{[K+\mathcal{O}_C]} = 0$ because $K + \mathcal{O}_C \subset \mathcal{K}_C$ is of finite codimension, and $\mathrm{Res}_{[K]} = 0$ because K is a K-submodule of $\mathcal{K}_C$. It follows that $\mathrm{Res}_{[\mathcal{O}_C]} = 0$, also. □

Duality. Given $\mathbf{g} = (g_x \in K)_{x\in C} \in \mathcal{K}_C$ and an $\alpha \in \Omega_{K/k}$, then for only finitely many $x \in C$, $g_x\alpha$ has a pole at $x \in C$ and hence the sum $\sum_{x\in C} \mathrm{Res}_x\, g_x\alpha$ is also finite. The residue theorem tells us that this sum is zero on the main diagonal $K \subset \mathcal{K}_C$ and so we have defined a pairing

$$(\alpha, \mathbf{g} + K) \in \Omega_{K/k} \times (\mathcal{K}_C/K) \mapsto \textstyle\sum_{x\in C} \mathrm{Res}_x\, g_x\alpha$$

If $f \in K$, then $(f\alpha, \mathbf{g} + K)$ and $(\alpha, f\mathbf{g} + K)$ have clearly the same image, and so this pairing factors through $\Omega_{K/k} \otimes_K (\mathcal{K}_C/K)$. When viewed as a pairing between (infinite dimensional) k-vector spaces, it gives rise to a k-linear map

$$R : \Omega_{K/k} \to \mathrm{Hom}_k(\mathcal{K}_C/K, k),$$

which is even K-linear when we let $\mathrm{Hom}_k(\mathcal{K}_C/K, k)$ inherit the structure of a K-vector space via $\mathcal{K}_C/K$ (so if $f \in K$ and $a \in \mathrm{Hom}_k(\mathcal{K}_C/K, k)$, then $fa : \mathcal{K}_C/K \to k$ is obtained by first multiplying in $\mathcal{K}_C/K$ with f and then applying a). So $R(\Omega_{K/k})$ is a K-linear subspace of $\mathrm{Hom}_k(\mathcal{K}_C/K, k)$.

Note that if $\alpha \in \Lambda(D)$, then $R(\alpha)$ vanishes on $\mathcal{O}_C(-D)$. In other words, it induces a k-linear map of *finite dimensional* k-vector spaces

$$R_D : \Lambda(D) \to \mathrm{Hom}_k(I(-D), k) = I(-D)^\vee.$$

Since $\Omega_{K/k}$ is the union of the $\Lambda(D)$, the image of R lies in the union of the finite dimensional subspaces $I(-D)^\vee \subset \mathrm{Hom}_k(\mathcal{K}_C/K, k)$. We call the latter union the *topological dual* of $\mathcal{K}_C/K$ and denote it by $\mathrm{Hom}_k^c(\mathcal{K}_C/K, k)$. Concretely, a k-linear

form $\alpha : \mathcal{K}_C/K \to k$ lies in $\mathrm{Hom}_k^c(\mathcal{K}_C/K, k)$ if and only if it vanishes on the image of $\mathcal{O}_C(-D) \to \mathcal{K}_C/K$ for some divisor D ([10]). This is a K-linear subspace, for if $f \in K$, and α is as above, then $f\alpha : \mathcal{K}_C/K \to k$ vanishes on the image of $\mathcal{O}_C(-D + \mathrm{div}(f)) \to \mathcal{K}_C/K$.

Theorem 11.23 (Duality theorem)**.** The map $R : \Omega_{K/k} \to \mathrm{Hom}_k^c(\mathcal{K}_C/K, k)$ is a K-linear isomorphism and for every divisor D on C, $R_D : \Lambda(D)) \to I(-D)^\vee$ is a k-linear isomorphism.

So the Riemann-Roch theorem can now be stated as:

$$\dim L(D) - \dim \Lambda(-D)) = 1 - g(C) + \deg D.$$

The proof of Theorem 11.23 relies on:

Lemma 11.24. The K-vector space $\mathrm{Hom}_k^c(\mathcal{K}_C/K, k)$ is of dimension ≤ 1.

PROOF. We follow the proof given in Serre [2]. Suppose $a, b \in \mathrm{Hom}_k^c(\mathcal{K}_C/K, k)$ are K-independent. Choose a divisor $D > 0$ such that both a and b lie in $I(-D)^\vee$, in other words, vanish on the image of $\mathcal{O}(-D) \to \mathcal{K}_C/K$ and let E be a positive divisor of degree $n > 3g(C) - 3 + \deg(D)$. Since a and b are K-independent, the k-linear map

$$(f, g) \in L(E) \oplus L(E) \to fa + gb \in \mathrm{Hom}_k^c(\mathcal{K}_C/K, k)$$

is injective. The image lies in $I(-D-E)^\vee$. Since $-D-E < 0$, we have $L(-D-E) = 0$ and hence by the Riemann-Roch theorem,

$$\dim(I(-D-E)^\vee) = \dim I(-D-E) = g(C) - 1 + \deg(D) + n.$$

The Riemann-Roch theorem also tells us that $\dim L(E) \geq 1 - g(C) + n$. It follows that $2(1 - g(C) + n) \leq g(C) - 1 + \deg(D) + n$. But this contradicts our assumption that $n > 3(g(C) - 1) + \deg(D)$. □

PROOF OF THE DUALITY THEOREM 11.23. The K-linear map R is injective, has 1-dimensional source and a target of dimension ≤ 1. So if we prove it to be nonzero, then it will be an isomorphism. For this we fix some $p \in C$ and a nonzero $\alpha \in \Omega_{K/k}$. Then choose for every $x \in C$ a $g_x \in K$ with $v_x(g_x) = -v_x(\alpha)$ except when $x = p$, where we require that $v_p(g_p) = -v_p(\alpha) - 1$. Then $\mathrm{Res}_x(g_x\alpha) = 0$ unless $x = p$ and hence $R(\alpha)(\mathbf{g}) = \mathrm{Res}_p(g_p\alpha) \neq 0$. So R is an isomorphism.

To prove that R_D is an isomorphism, it suffices to show that $R^{-1}I(-D)^\vee = \Lambda(D)$. But this amounts to: if $\alpha \in \Omega_{K/k}$ and for all $x \in C$, $\mathrm{Res}_x\, g_x\alpha = 0$ for all g_x with $v(g_x) \geq d_x$, then $v_x(\alpha) \geq -d_x$ for all $x \in C$ and this is obvious. □

Some consequences of the Riemann-Roch theorem. There are quite a few.

Corollary 11.25. The k-vector space of regular differentials on C has dimension $g(C)$ and the degree of a canonical divisor on C is $2g(C) - 2$.

PROOF. The duality theorem asserts that $\Lambda(0)$ can be identified with the k-dual of $I(0)$ and so its dimension is $g(C)$.

[10] We can make $\mathcal{K}_C/K$ a topological k-vector space by stipulating that the images of the maps $\mathcal{O}(-D) \to \mathcal{K}_C/K$ make up a neighborhood basis of the origin (these are subspaces of finite codimension by Corollary 11.16). If we give k the discrete topology, then $\mathrm{Hom}_k^c(\mathcal{K}_C/K, k)$ is the space of continuous linear forms on $\mathcal{K}_C/K$.

Let D be a canonical divisor. We have $L(D) \cong \Lambda(0)$, while $\Lambda(-D)$ is spanned by α (and is hence of dimension 1). The Riemann-Roch theorem then says that $\deg(D) = \dim L(D) - \dim \Lambda(-D) - 1 + g(C) = 2g(C) - 2$. □

Corollary 11.26. A complete linear system on C of degree $d \geq 2g(C) - 1$ has dimension $d - g(C)$.

PROOF. Let D be a divisor of degree d. In that case a divisor for $\Lambda(-D)$ has degree $\leq 2g(C) - 2 - d \leq -1$, and so $\Lambda(-D) = \{0\}$. The assertion then follows from the Riemann-Roch theorem: $\dim L(D) = 1 - g(C) + d$. □

We say that $p \in C$ is a *fixed point* of a linear system if p is in the support of each of its members. So p is a fixed point of $|D|$ if $D' := D - (p)$ is such that $|D| = (p) + |D'|$. Or equivalently, the inclusion $L(D') \subseteq L(D)$ is an equality.

Corollary 11.27. A complete linear system on C of degree $d \geq 2g(C)$ has no fixed point. If $d > 2g(C)$, then this linear system defines an embedding of C in a projective space.

PROOF. Let D be a divisor of degree $d \geq 2g(C)$ and let $p \in C$. We then have $\dim |D - (p)| = d - 1 - g(C) < d - g(C) = \dim |D|$ and so there exists a member of $|D|$ such that p is not in its support.

Now assume $d \geq 2g(C) + 1$. When $p, q \in C$, then Corollary 11.26 implies that $\dim |D - (p)| > \dim |D - (p) - (q)|$. So if $p \neq q$, then there exists a member D' of $|D|$ with $p \in \operatorname{supp}(D')$ and $q \notin \operatorname{supp}(D')$. This means that for the associated map $\phi : C \to P$, there exists a hyperplane $H \subset P$ such that $\phi(p) \in H$ and $\phi(q) \notin H$. In other words, ϕ is injective. In case $p = q$, this means that there exists a hyperplane $H \subset P$ such that $\phi^* H$ has multiplicity 1 at p, which means that $\phi : \mathfrak{m}_{P,\phi(p)}/\mathfrak{m}^2_{P,\phi(p)} \to \mathfrak{m}_{C,p}/\mathfrak{m}^2_{C,p}$ is onto. Nakayama's lemma then implies that $\phi^* : \mathcal{O}_{P,\phi(p)} \to \mathcal{O}_{C,p}$ is also onto.

In order to show that ϕ is a closed immersion, it now suffices to show that for every $p \in C$, there exists an affine neighborhood $V_p \subset P$ of $\phi(p)$ in P such that ϕ defines a closed immersion of $\phi^{-1}V_p$ into V_p. To prove this, let U_p be an open affine neighborhood of p in C and let $f_1, \ldots, f_r \in k[U_p]$ generate $k[U_p]$ as a k-algebra. Then by the above there exists an affine neighborhood $V_p \subset P$ of $\phi(p)$ in P and $g_i \in k[V_p]$ such that $U_p \supseteq \phi^{-1}V_p$ and $\phi^* g_i = f_i|\phi^{-1}V_p$. Hence ϕ defines a closed immersion of $\phi^{-1}V$ into V. □

REMARK 11.28. The original Riemann-Roch theorem was stated in a complex-analytic setting, where C is a compact connected Riemann surface. The notion of a divisor is then defined as before and the Riemann-Roch theorem and the duality theorem are for spaces of meromorphic functions resp. differentials (rather than for their rational counterparts). It is then proved that the genus is in fact the genus of the underlying topological surface (by showing that every element of $H^1(C; \mathbb{C})$ is uniquely represented as the sum of a holomorphic differential and the complex conjugate of one, so that the first Betti number of C is $2g(C)$). The holomorphic version of Corollary 11.27 tells us that C can be holomorphically embedded in a complex projective space and with a bit more work one may show that the image is in fact Zariski closed. Consequently, C can be endowed with the structure of a smooth complex projective curve for which every meromorphic function is a rational function. It is not hard to show that this structure must be unique. So a compact connected Riemann surface is essentially the same thing as an irreducible smooth complex projective curve and the topological genus of the surface is equal to the genus as defined here.

Corollary 11.29. If $S \subset C$ is finite and nonempty, then $C \smallsetminus S$ is affine.

PROOF. Let $D = \sum_{x \in S}(p)$ and choose an integer $n > 0$ such that $n \deg(D) \geq 2g(C) + 1$. By Corollary 11.27 this defines a closed immersion $\phi : C \hookrightarrow P$ in a projective space and a hyperplane $H \subset P$ such that $nD = \phi^* H$. So ϕ maps $C \smallsetminus S$ isomorphically onto a closed subset of the affine space $P \smallsetminus H$. Hence $C \smallsetminus S$ is affine. □

EXERCISE 82. Prove that if D is a divisor of degree $2g(C) - 2$, then $\dim L(D) = g(C) - 1$, unless D is canonical.

EXERCISE 83. Let C be a smooth irreducible projective curve of genus 1 and let $o \in C$.

(a) Prove that C has a differential with neither poles nor zeroes.
(b) Let $p, q \in C$. Prove that there is a unique $r \in C$ such that $(p) + (q) \equiv (o) + (r)$. Prove that if we define $r := p * q$, then this defines on C the structure of an abelian group having o as its unit element.
(c) Prove that the linear system $|3(o)|$ maps C isomorphically onto a cubic curve in a projective plane with o mapping to a flex point and that $p*q*r = o$ means that $(p) + (q) + (r)$ is the preimage of a line.

The (defining) identity $\dim I(0) = g(C)$ tells us that in order that a given polar part $(f_x + \mathcal{O}_{C,x})_{x \in C} \in \oplus_{x \in C}(K/\mathcal{O}_{C,x}) = \mathcal{K}_C/\mathcal{O}_C$ be the polar part of a rational function, the Laurent coefficients of the $(f_x)_{x \in C}$ must obey $g(C)$ linearly independent linear equations. These linear equations are defined by residue identities. To be precise, we have an exact sequence

$$0 \longrightarrow K \longrightarrow \oplus_{x \in C} K/\mathcal{O}_{C,x} \longrightarrow \Lambda(0)^\vee \longrightarrow 0,$$

where the last map assigns to $(g_x + \mathcal{O}_{C,x})_{x \in C}$ the linear function $\alpha \in \Lambda(0) \mapsto \sum_{x \in C} \operatorname{Res}_x g_x \alpha$. The next corollary gives a similar characterization for the polar parts of differentials. It is simpler, as it involves just one equation:

Corollary 11.30. The sequence

$$0 \longrightarrow \Omega_{K/k} \longrightarrow \oplus_{x \in C} \Omega_{K/k}/\Omega_{C,x} \xrightarrow{\sum_x \operatorname{Res}_x} k \longrightarrow 0$$

is exact.

PROOF. It is clear that the composite map is zero, that the first map is injective and that the last map is surjective. In order to check exactness in the middle, we choose a nonzero $\omega \in \Omega_{K/k}$ and denote by $D = \sum_{x \in C} d_x(x)$ its divisor. It is then clear that $\Lambda(-D)$ is of dimension one and spanned by ω. The duality theorem identifies this with $I(D)^\vee$. This means that for $\mathbf{f} = (f_x)_{x \in C} \in \mathcal{K}_C$ the property $\sum_{x \in C} \operatorname{Res}_x f_x \omega = 0$ is equivalent to $\mathbf{f}$ mapping to zero in $I(D) = \mathcal{K}_C/(K + \mathcal{O}_C(D)$. This, in turn, is equivalent to the existence of an $f \in K$ such that $v_x(f - f_x) \geq -d_x$ for all $x \in C$.

Now let $\boldsymbol{\alpha} \in \oplus_{x \in C} \Omega_{K/k}/\Omega_{C,x}$ be in the kernel of $\sum_x \operatorname{Res}_x$ and represent it by $(\alpha_x)_{x \in C} \in \oplus_{x \in C} \Omega_{K/k}$, so that $\sum_{x \in C} \operatorname{Res}_x \alpha_x = 0$. We can write $\alpha_x = f_x \omega$, with $f_x \in K$. Since for all but finitely many $x \in C$, $v_x(\omega) = v_x(\alpha_x) = 0$, the same is true for $(f_x)_{x \in C}$: we have $v_x(f_x) = 0$ for all but finitely many $x \in C$ and so $(f_x)_{x \in K} \in \mathcal{K}_C$. Since $\sum_{x \in C} \operatorname{Res}_x f_x \omega = 0$, it follows by the above argument that there exists an $f \in K$ such that $v_x(f - f_x) \geq -d_x$ for all $x \in C$. This means

that $v_x(f\omega - \alpha_x) = v_x((f - f_x)\omega) \geq 0$ for all $x \in C$ and so $\boldsymbol{\alpha}$ is the image of $f\omega \in \Omega_{K/k}$. □

Let us see what happens with canonical divisors with respect to a separable morphism. Let $\pi : C \to C'$ be a finite morphism of smooth projective curves. Recall that the ramification index $e_x(\pi)$ of π at $x \in C$ is equal to $v_x(\pi^* t)$, when $t \in \mathcal{O}_{C',\pi(x)}$ is a uniformizer. It is clear that $e_x(\pi) \geq 1$ with equality for all but finitely many $x \in X$ and so we can define the *ramification divisor* R_π of π by $x \in C \mapsto e_x(\pi) - 1$. Note that this divisor is effective.

Lemma 11.31. Let $\pi : C \to C'$ be a separable morphism of projective curves and assume that no ramification index is divisible by the characteristic of k. If D' is a canonical divisor for C', then $\pi^* D' + R_\pi$ is a canonical divisor for C.

PROOF. Let D' be the divisor of a differential α' on C'. It suffices to show that $\pi^* D' + R_\pi$ is the divisor of $\pi^* \alpha'$. Choose $x \in C$ and write e for $e_x(\pi)$. Let $t \in \mathcal{O}_{C',f(x)}$ and $s \in \mathcal{O}_{C,x}$ be uniformizers so that $\pi^* t = us^e$ for some unit $u \in \mathcal{O}_{C,x}$. Then

$$d(\pi^* t) = s^e du + eus^{e-1} ds \equiv eus^{e-1} ds \pmod{\mathfrak{m}_{C,x}^e \Omega_{C,x}}.$$

and so $v_x(d(\pi^* t)) = e - 1$. The corollary follows easily from this. □

So if in the situation of Lemma 11.31 we compare the degrees of the canonical divisors we obtain:

Corollary 11.32 (Riemann-Hurwitz)**.** Let $\pi : C \to C'$ be a separable morphism of projective curves and assume that no ramification index is divisible by the characteristic of k. Then $2g(C) - 2 = \deg(\pi)(2g(C') - 2) + \deg(R_\pi)$.

PROOF. By Corollary 11.25 the degree of D' is $2g(C') - 2$. So by Proposition 11.5, the degree of $\pi^* D' + R_\pi$ is $\deg(\pi)(2g(C') - 2) + \deg(R_\pi)$. This is by Lemma 11.31 the degree of a canonical divisor of C and hence equal to $2g(C) - 2$. □

EXAMPLE 11.33. Assume k not of characteristic 2 and let $\pi : C \to \mathbb{P}^1$ be of degree 2. Then the hypotheses of the above Corollary are fulfilled and and R_π will be a reduced divisor. We find that its degree must be $2g(C) + 2$. Such a curve C is said to be *hyperelliptic*.

Bibliography

[1] M.F. Atiyah, I.G. Macdonald: *Introduction to Commutative Algebra,* Addison-Wesley (1969).

[2] J.-P. Serre: *Algebraic Groups and Class Fields,* Graduate Texts in Mathematics **117**, Springer Verlag (1988) (translated from the French *Groupes algébriques et corps de classes*). 106, 110

[3] J. Tate: *Residues of differentials on curves,* Annales scientifiques de l'É.N.S., 4e série, **1** (1968), 149–159. 107

[4] O. Zariski, P. Samuel: *Commutative algebra. Vol. 1,* Graduate Texts in Mathematics **28**, Springer-Verlag (1975).