Elliptic Curves

Version of May 17, 2005.

1 Introduction

An elliptic curve is an object which lives over a field. We will typically write E for such a curve and k for the field. Which field one uses largely determines the areas of mathematics that will be most relevant for studying our curves.

field k considered	**relevant area of mathematics**
$k = \mathbb{R}, \mathbb{C}$	analysis
$k = \overline{k}$	geometry
any k	algebra
$[k : \mathbb{Q}] < \infty$ $[k : \mathbb{Q}_p] < \infty$	arithmetic
$\#k < \infty$	algorithms cryptography

In recent years the theory of modular forms and modular elliptic curves has produced some great results (such as Wiles' proof of Fermat's Last Theorem), but this is a topic too advanced to deal with in an introductory course.

As an introduction to the theory of elliptic curves we shall consider a problem that was originally posed and solved by Diophantus in the third century AD. We don't have Diophantus' solution, only references to it. The first known solution to this problem is due to Fermat in the seventeenth century.

Theorem 1.1 (Diophantus). *Every* number *that is a difference of two* cubes *is also a sum of two* cubes.

For Diophantus a *number* is what we now call a positive rational number. A *cube* is a *number* that can be written as the third power of another *number*. These days we reserve the term *Diophantine equation* for equations to be solved in positive integers, but in this case Diophantus himself actually used rationals.

To get a feel for the kind of thing we can do to solve this problem, let's consider a simple case. We start with the equation

$$2^3 - 1^3 = 7.$$

With the benefit of negative numbers (which were not known to Diophantus and even in the time of Fermat not really accepted as 'actual' numbers), we can rewrite this as

$$2^3 + (-1)^3 = 7.$$

We see that our objective is to raise the number -1 a little bit, so that it becomes positive. In order to do this, we shall have to lower the other number, 2, by some amount t. The amount that -1 is raised by should be some simple expression in t, say, $5t$. From this we get the equation

$$(2 - t)^3 + (-1 + 5t)^3 = 7,$$

or, if we expand everything

$$8 - 12t + 6t^2 - t^3 - 1 + 15t - 75t^2 + 125t^3 = 7.$$

What we end up with is a cubic equation in t, of which we know one root, namely $t = 0$. What remains is a quadratic equation. In general, the roots of such an equation, will not be rational, so that the new solutions we find are not *numbers* in Diophantus' sense.

To resolve this problem we have to exploit the freedom we have to replace the $5t$ term by something more suitable. For instance, if we replace it by just t, then the equation becomes

$$(2-t)^3 + (-1+t)^3 = 7,$$

which expands to

$$8 - 12t + 6t^2 - t^3 - 1 + 3t - 3t^2 + t^3 = 7.$$

The degree three terms cancel, so we are left with a quadratic equation, of which one of the roots is 0, so the other is rational. In fact, we see that the other root is 3. This is rather unfortunate, because if we plug in $t = 3$ we see that we in fact get the original equation back with the numbers reversed.

Fortunately for us there is another way to solve the problem. Rather than making the degree three terms cancel, we can try to make the linear terms cancel. We will then still have a cubic equation, but 0 will now be a double root. To do this, substitute $4t$ in the place of $5t$ in the original equation. This yields

$$(2-t)^3 + (-1+4t)^3 = 7$$

or

$$8 - 12t + 6t^2 - t^3 - 1 + 12t - 48t^2 + 64t^3 = 7.$$

As promised, the linear terms cancel and we are left with

$$63t^3 = 42t^2.$$

Hence $t = 2/3$ is a solution. This gives us the equality

$$\left(\frac{4}{3}\right)^3 + \left(\frac{5}{3}\right)^3 = 7.$$

It is clear that we can replace the numbers in the example above by variables. Suppose u is a positive rational. We are looking for pairs (x, y) of rational numbers that satisfy

$$x^3 + y^3 = u.$$

Suppose we have one such pair (x_0, y_0). The example gives us a procedure to make a new solution (x_1, y_1). The question is, if we begin with a 'bad' pair, with $y_0 < 0$, will we get a 'good' pair, with $x_1 > 0$ and $y_1 > 0$ out of it? It is an easy exercise to see that y_1 will always be positive, but x_1 can take either sign. Fermat's insight was that the process can be repeated and that it eventually will produce a good solution.

To get a better feel for just what is going on, it is most helpful to consider the geometry of the problem, i.e., to draw a picture.

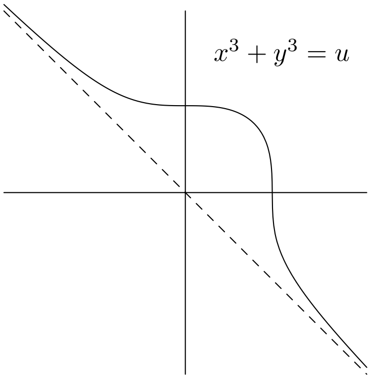

The picture shows the curve in $\mathbb{R}^2$. It is better to consider the curve in the projective plane, $\mathbb{P}^2(\mathbb{R})$. This is the set of all triples $(x : y : z)$ with $x, y, z \in \mathbb{R}$ not all zero, modulo the equivalence

$$(x : y : z) = (u : v : w) \quad \Longleftrightarrow \quad \exists \lambda \in \mathbb{R} : \lambda x = u, \lambda y = v \text{ and } \lambda z = w.$$

We find $\mathbb{R}^2$ embedded in $\mathbb{P}^2(\mathbb{R})$ as the set of all tuples $(x : y : 1)$. The complement of $\mathbb{R}^2$ is a projective line, $\mathbb{P}^1(\mathbb{R})$, at infinity. Each point on this line corresponds to a pair of opposite directions in the plane. For example, the asymptotic direction of our curve is $(-1, 1)$, which corresponds to the point $(-1 : 1 : 0)$ at infinity.

In order to find the equation of our curve in the projective plane we have to make it homogeneous. This means adding factors z to all terms in order to make them degree three. The equation then becomes

$$x^3 + y^3 = uz^3.$$

The points that satisfy this equation are precisely the points of the form $(x : y : 1)$ that were in our plane curve, plus the point $(-1 : 1 : 0)$ at infinity, corresponding to the asymptotic direction of our curve.

There are several reasons for considering the projective curve rather than the plane one. First of all, from a topological point of view the projective curve is simpler. It is in fact simply a circle. From the picture we see that the original curve is just a line and we have to imagine that in the projective space the endpoints are glued together at infinity, making it into a circle.

A very useful theorem concerning curves in the projective plane is Bézout's theorem. It tells us that under certain mild conditions the number of points in which two curves intersect is the product of the degrees of these two curves. Looking at the line l in the following picture, we see that this need not always hold if we just consider the plane curve.

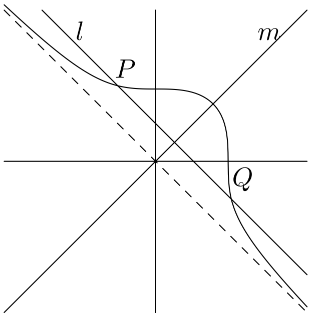

If we look at this line in the projective space, we see that there is a third intersection point lying at infinity. Still we have to be a bit careful. Look at the line m. It only intersects the curve in one point, even if we look at it in projective space. The reason for this is that the other two solutions are in fact imaginary. To make Bézout's theorem work, we really need to be in an algebraically closed field like $\mathbb{C}$. Making pictures in $\mathbb{C}^2$ is not possible, so we stick to $\mathbb{R}^2$ and remember we have to be careful.

The next picture illustrates another subtlety in Bézout's theorem. These apparent counterexamples show us that we need to count intersection points with multiplicities. A tangent point like T has multiplicity two, an inflection point like R has multiplicity three.

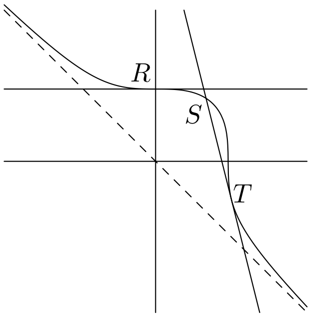

The third reason we like the projective curve better than the plane curve is that the former has more structure than the latter.

Theorem 1.2. *Let $u \in \mathbb{R}_{>0}$. The projective curve $C = \{(x : y : z) \in \mathbb{P}^2(\mathbb{R}) : x^3 + y^3 = ux^3\}$ is an abelian group with zero element $O = (-1 : 1 : 0)$ the point at infinity and such that for all projective lines that intersect C in three point P, Q and R with listed multiplicities, we have $P + Q + R = O$.*

For a few interesting examples we refer back to the points that are marked in the last two pictures. We see that the line l intersects the curve in P, Q and O, which gives the equality $P + Q + O = O$ or $P = -Q$ in the group. The point R is an inflection point, so $3R = O$. Finally we see that $2T + S = O$.

Next we express Diophantus' construction in terms of the group. The process of replacing x by $x - t$ and y by $y + at$ corresponds in the picture to looking at the line through (x, y) with direction $(-1, a)$. Taking the value of a in such a way that the linear term vanishes corresponds to taking the tangent line at the point (x, y). We just saw that for a tangent line at the point T which also intersects the curve at a point S the relation $S = -2T$ holds in the group. So going from the point (x_0, y_0) to (x_1, y_1) corresponds to multiplying by -2 in the group. Since multiplying by -1 simply comes down to swapping x and y we might as well consider the map that multiplies by 2 instead.

Now we need a result from analysis, which tells us that the curve isn't just a circle topologically, but also as a group. Taking $\mathbb{T} = \{\omega \in \mathbb{C}^* : |\omega| = 1\}$ the circle group in the complex numbers we get the following isomorphism

$$\begin{array}{ccc} C & \longrightarrow & \mathbb{T} \\ O & \longmapsto & 1 \\ (x : y : 1) & \longmapsto & \exp\left(2\pi i \int_{-\infty}^{x} \frac{dt}{\sqrt[3]{u - t^3}^2} \Big/ \int_{-\infty}^{\infty} \frac{dt}{\sqrt[3]{u - t^3}^2}\right) \end{array}$$

By looking at the picture of our curve, we see that the intersection points of the curve with the axes are precisely the inflection points. These are points with order 3 in the group and since only two of those exist, we know the image of the part of the curve that lies in the positive quadrant. It is the third of the circle that lies between $\zeta = \frac{-1+i\sqrt{3}}{2}$ and $\overline{\zeta}$.

Notice that if we take a point, like P and start squaring it (the group operation in $\mathbb{T}$ is written multiplicatively) the angle keeps doubling until we end up in the area between ζ and $\overline{\zeta}$. The point can't 'jump' over this area since $\zeta^2 = \overline{\zeta}$. We conclude that the point on the curve will, after a finite number of steps, jump into the positive quadrant. This means that both x and y are positive, so that we have written u as a sum of two positive cubes.

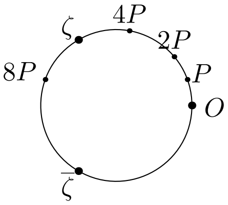

2 Categories and functors

This section is a translation of the section on categories and functors from the Leiden algebra syllabus.

Much of the so called 'conceptual mathematics' can be phrased short and precise in terms of categories and functors. This is more an efficient language than a theory in its own right en the categorial notions are justified largely by the number of concrete examples we find in all areas of math. The mathematical content of this section therefore lies mostly in the numerous examples.

Definition 2.1. A category $\mathcal{C}$ consists of a collection of object and for each pair A, B of objects in $\mathcal{C}$ a set $\mathrm{Hom}_{\mathcal{C}}(A,B)$ of morphisms from A to B. The sets of morphisms are pairwise disjoint an for every three objects A, B and C in $\mathcal{C}$ there is a composition of morphisms

$$\circ : \mathrm{Hom}_{\mathcal{C}}(A,B) \times \mathrm{Hom}_{\mathcal{C}}(B,C) \longrightarrow \mathrm{Hom}_{\mathcal{C}}(A,C).$$

The following criteria are satisfied.

1. For every $A \in \mathcal{C}$ the set $\mathrm{Hom}_{\mathcal{C}}(A,A)$ contains an identity id_A which acts as a neutral element with respect to composition.

2. For morphisms $A \xrightarrow{f} B \xrightarrow{g} C \xrightarrow{h} D$ it holds that $(h \circ g) \circ f = h \circ (g \circ f)$.

The *morphisms* in $\mathcal{C}$ are also refered to as the *arrows* or *maps* in $\mathcal{C}$. One usually writes $\mathrm{Hom}(A,B)$ for $\mathrm{Hom}_{\mathcal{C}}(A,B)$ if it is clear which category is meant.

To avoid certain set theoretic paradoxes like the 'set of all sets' one does not demand that the objects of $\mathcal{C}$ form a set. They are in fact a *class* in the sense of set theory. We shall not deal with such logical finesses, which one usually avoids by working with *small categories* inside a suitable *universe.*

The existence of an identity morphism for every object enables us to talk about inverses of morphisms and with that of *isomorphisms*, morphisms with a two-sided inverse. Morphisms in $\mathrm{Hom}(A,A)$ are called *endomorphisms* of A and isomorphisms in this set are called *automorphisms.* The automorphisms form a group $\mathrm{Aut}(A)$ under composition.

Note that there is no mention of elements anywhere in the definition of a category; we don't even assume the objects consist of elements. Also it is possible for the set $\mathrm{Hom}(A,B)$ to be empty for some A and B.

Categories are often refered to in terms of their objects, for example, the category $\mathfrak{Ab}$ of abelian groups or the category $\mathfrak{Mod}_R$ of R-modules. The reader is expected to understand that the morphisms in the category are the 'obvious' ones. In the case of $\mathfrak{Ab}$ this are the group homomorphisms and in the case of $\mathfrak{Mod}_R$ the R-module homomorphisms.

The category $\mathfrak{Ab}$ is in a natural way a *subcategory* of the category $\mathfrak{Grp}$ of all groups. In general, one calles a category $\mathcal{C}$ a subcategory of $\mathcal{D}$ if the objects in $\mathcal{C}$ are also objects in $\mathcal{D}$ and for every two objects A, B in $\mathcal{C}$ there is an inclusion $\mathrm{Hom}_{\mathcal{C}}(A,B) \subset \mathrm{Hom}_{\mathcal{D}}(A,B)$. If in fact $\mathrm{Hom}_{\mathcal{C}}(A,B) = \mathrm{Hom}_{\mathcal{D}}(A,B)$ for all A, B in $\mathcal{C}$ then $\mathcal{C}$ is called a *full subcategory* of $\mathcal{D}$.

Examples 2.2.

Before we proceed we give some of the numerous examples. Every reader can extend the following list in her favourite direction.

1. The category $\mathfrak{Sets}$ of sets with 'ordinary' maps as morphisms is a standard example of a category. The subcategory $\mathfrak{FSets}$ of finite sets forms a full subcategory f $\mathfrak{Sets}$. For every group G there is a category of G-$\mathfrak{sets}$ of sets with a G action. The morphisms are the G-equivariant maps. A map $f : X \to Y$ between two G-sets is called G-equivariant if for all $x \in X$ and $g \in G$ we have $f(gx) = gf(x)$.

2. The category $\mathfrak{Grp}$ of groups with group homomorphisms contains the category $\mathfrak{Ab}$ of abelian groups as a full subcategory. Similarly the category $\mathfrak{Rng}$ of rings with ring homomorphisms contains the full subcategory $\mathfrak{CRng}$ of commutative rings. These are all 'large' categories and one often works with smaller subcategories such as *finite* abelian groups or *Noetherian* rings.

3. The category $\mathfrak{Vec}_K$ of vector spaces over a field K with K-linear maps as morphisms had a full subcategory $\mathfrak{FVec}_K$ of finite dimensional K-vector spaces.

4. The modules over a ring R together with R-homomorphisms form a category $\mathfrak{Mod}_R$. For commutative R once can do just about every universal construction (fibered sums and products, quotients) in $\mathfrak{Mod}_R$. This makes $\mathfrak{Mod}_R$ into the typical example of an *abelian category.*

If R is taken to be the group ring $K[G]$ of a group G over a field K then $\mathfrak{Mod}_R = \mathfrak{Rep}_K(G)$, the category of K-representations of G.

5. The category $\mathfrak{Top}$ of topological spaces has continuous maps as its morphisms. One often works in full categories of topological spaces which have one or more extra properties (connected, Hausdorff, metric, compact, ...). The topology $\mathfrak{T}_X$ of a topological space X is itself a category. The objects are the open subsets of X and the morphisms are inclusions between these open subsets.

6. From every category $\mathcal{C}$ we can construct the *opposite category* $\mathcal{C}^{\mathrm{opp}}$ by 'inverting all the arrows'. More precisely: the objects of $\mathcal{C}^{\mathrm{opp}}$ are the same as those of $\mathcal{C}$ and the sets of morphisms $\mathrm{Hom}_{\mathcal{C}^{\mathrm{opp}}}(A, B)$ are in bijection with $\mathrm{Hom}_{\mathcal{C}}(B, A)$, say, $f^{\mathrm{opp}} \leftrightarrow f$. The composition of morphisms in $\mathcal{C}^{\mathrm{opp}}$ is defined by $f^{\mathrm{opp}} \circ g^{\mathrm{opp}} = (g \circ f)^{\mathrm{opp}}$.

As we can see in some of the examples above the sets $\mathrm{Hom}_{\mathcal{C}}(A, B)$ sometimes inherit some extra structure from $\mathcal{C}$. In $\mathcal{C} = \mathfrak{Ab}$ we get abelian groups and in $\mathfrak{Mod}_R$ for commutative R we get R-modules.

The morphisms in a category $\mathcal{C}$ form another category $\mathfrak{Mor}(\mathcal{C})$. A morphism $\phi : f \to g$ in $\mathfrak{Mor}(\mathcal{C})$ from $f \in \mathrm{Hom}_{\mathcal{C}}(A, B)$ to $g \in \mathrm{Hom}_{\mathcal{C}}(C, D)$ is an ordered pair $\phi = (\phi_1, \phi_2)$ of morphisms in $\mathcal{C}$ that makes the following diagram commute.

$$\begin{tikzcd}
A \arrow[r, "f"] \arrow[d, "\phi_1"'] & B \arrow[d, "\phi_2"] \\
C \arrow[r, "g"'] & D
\end{tikzcd}$$

One often gets interesting subcategories of $\mathfrak{Mor}(\mathcal{C})$ by considering morphisms to or from a fixed object in $\mathcal{C}$. In the first case we take a fixed $A = C$ in the diagram above and only look at the morphisms $\phi = (\phi_1, \phi_2)$ with $\phi_1 = \mathrm{id}_A$. In the second case we fix $B = D$ and look at morphisms with $\phi_2 = \mathrm{id}_B$. We refed to either case as a category of objects over a fixed base object.

Examples 2.3.

For every commutative ring R one can view the category $\mathfrak{CAlg}_R$ of commutative R-algebras as the category of rings over R. After all, a morphisms of R-algebras $A - 1 \to A - 2$ respects the R-algebra structure and with that is a morphism between the structure maps $f_i : R \to A_i$ in $\mathfrak{CRng}$ which is the identity on R.

An interesting example of topological spaces over a fixed base object is given by the category $\mathfrak{Cov}_X$ of *coverings* of a topological space X. A map $f : Y \to X$ of topological spaces is called a *covering* if every point $x \in X$ has a neighbourhood $U_x \subset X$ such that $f^{-1}[U_x] \xrightarrow{f} U - x$ is the *trivial covering*. This means that the *fibre* $f^{-1}(x)$ above x is discrete in Y and that there is a homeomorphism $f^{-1}(x) \times U_x \to f^{-1}[U_x]$ which composed with f gives the projection on the second co-ordinate (make a picture!). A morphism ϕ from a covering $f_1 : Y_1 \to X$ to $f_2 : Y_2 \to X$ (a so-called *decking transform*) is a continuous map $\phi : Y_1 \to Y_2$ with $f_2 \circ \phi = f_1$.

Definition 2.4. A (co-variant) functor $F : \mathcal{C} \to \mathcal{D}$ is a map that assigns to every object A of $\mathcal{C}$ an object $F(A)$ of $\mathcal{D}$ and to every morphism $f \in \mathrm{Hom}_{\mathcal{C}}(A, B)$ a morphism $f_* = F(f) \in \mathrm{Hom}_{\mathcal{D}}(F(A), F(B))$. Moreover, $(\mathrm{id}_A)_* = \mathrm{id}_{F(A)}$ and $(f \circ g)_* = f_* \circ g_*$.

One often simply says that the construction of $F(A) \in \mathcal{D}$ from $A \in \mathcal{C}$ is 'functorial'. Such a construction has all sorts of nice 'stability properties' that make functorial notions much more managable than non-functorial ones.

Examples 2.5.

1. The forming of the commutator subgroup $[G, G]$ from a group G is a functor from $\mathfrak{Grp}$ to itself. The functor $G \mapsto G^{\mathrm{ab}} = G/[G, G]$ that assigns to every group its largest abelian quotient is a functor $\mathfrak{Grp} \to \mathfrak{Ab}$. Forming the center $Z(G)$ of a group G is *not* a functor $\mathfrak{Grp} \to \mathfrak{Ab}$ since in general a group homomorphism $f : G_1 \to G_2$ does not induce a group homomorphism between the centers.

2. Taking the unit group R^* of a ring R is a functor $U : \mathfrak{Rng} \to \mathfrak{Grp}$. For every positive integer n there is a functor $\mathrm{GL}_n : \mathfrak{CRng} \to \mathfrak{Grp}$ that assigns to a commutative ring R the group $\mathrm{GL}_n(R)$ of invertible $n \times n$-matrices with coefficients from R. Note that U and GL_1 are 'the same' functor.

3. The map $\mathfrak{CRng} \to \mathfrak{CRng}$ that assigns to every commutative ring R its *reduced ring* R/N_R, with N_R the nil-radical of R, is a functor. On the subcategory of reduced rings it is the identity.

4. A *forgetful functor* is a functor that forgets part of the structure of an object. There are, for example, forgetful functors from most of the categories metioned in 2.2 to $\mathfrak{Sets}$ that assign to a group (ring, vector space, etc.) the underlying set. Of the same nature are the functors from $\mathfrak{Rng}$ and $\mathfrak{Vec}_K$ to $\mathfrak{Ab}$ that assign to a ring or a vector space the underlying abelian addition group or the functors $\mathfrak{Rep}_K(G) \to \mathfrak{Vec}_K$ and G-$\mathfrak{sets} \to \mathfrak{Sets}$ that forget the G-action.

5. Constructing the fundamental group $\pi(X)$ of a topological space is *not* a functor $\mathfrak{Top} \to \mathfrak{Grp}$, not even when we restrict ourselves to path connected spaces. Instead we need the category $\mathfrak{Top}_*$ of topological spaces X with a base point $x \in X$. A morphism from (X, x) to (Y, y) is a continuous map $f : X \to Y$ such that $f(x) = y$. Observe that this is the category of topological spaces over a one point space. Assigning the fundamental group $\pi(X, x)$ to a space (X, x) is a functor $\mathfrak{Top}_* \to \mathfrak{Grp}$.

6. In every category $\mathcal{C}$ an object $X \in \mathcal{C}$ gives rise to a *representation functor* $\mathrm{Hom}_{\mathcal{C}}(X, -) : \mathcal{C} \to \mathfrak{Sets}$ given by $A \mapsto \mathrm{Hom}_{\mathcal{C}}(X, A)$. There is also a functor $\mathrm{Hom}_{\mathcal{C}}(-, X)$ but it does not satisfy definition 2.4 because it 'reverses the arrows'.

Definition 2.6. A contra-variant functor $F : \mathcal{C} \to \mathcal{D}$ is a map that assigns to every object $A \in \mathcal{C}$ an object $F(A) \in \mathcal{D}$ and to every morphism $f \in \mathrm{Hom}_{\mathcal{C}}(A, B)$ a morphism $f^* = F(f) \in \mathrm{Hom}_{\mathcal{D}}(F(B), F(A))$. Furthermore, $(\mathrm{id}_A)^* = \mathrm{id}_{F(A)}$ and $(f \circ g)^* = g^* \circ f^*$.

Examples 2.7.

We've already seen that the representation functors $\mathrm{Hom}_{\mathcal{C}}(-, X)$ are contra-variant. Special cases of these functors are the various *duality functors* like $M \mapsto M^* = \mathrm{Hom}_R(M, R)$ in the category $\mathfrak{Mod}_R$ of modules over R or the functor $A \mapsto A^\vee = \mathrm{Hom}(A, \mathbb{Q}$ of $\mathbb{Z})$ on $\mathfrak{Ab}$.

In somewhat greater generality there are many categories $\mathcal{C}$ with contra-variant functors that assign to an object $A \in \mathcal{C}$ a set of 'R-valued functions on A'. R is usually a suitable ring which results in the set of R-valued functions inheriting extra structure from R. As an example we can think of the set $C(X)$ of continuous real valued functions on a topological space X. With the usual point-wise operations it inherits a *ring structure* from $\mathbb{R}$.

In category theory the definition of objects is closely connected to the definition of morphisms of such objects. The collection $\mathfrak{Fun}(\mathcal{C}, \mathcal{D})$ of functors $\mathcal{C} \to \mathcal{D}$ becomes a category if we define the morphisms between functors, also called *natural transformations*, as follows.

Definition 2.8. A natural transformation between functors $F, G : \mathcal{C} \to \mathcal{D}$ is a collection of morphisms $\{\tau_C : F(C) \to G(C)\}_{C \in \mathcal{C}}$ in $\mathcal{D}$ such that for every morphism $f : C \to C'$ in $\mathcal{C}$ the following diagram commutes.

$$\begin{tikzcd}
F(C) \arrow[r, "\tau_C"] \arrow[d, "F(f)"'] & G(C) \arrow[d, "G(f)"] \\
F(C') \arrow[r, "\tau_{C'}"'] & G(C')
\end{tikzcd}$$

If all the morphisms τ_C are isomorphisms than the functors F and G are called naturally equivalent or isomorphic.

Examples 2.9.

1. For the functors $\mathrm{GL}_n : \mathfrak{CRng} \to \mathfrak{Grp}$ and $U : \mathfrak{CRng} \to \mathfrak{Grp}$ defined in example 2.5.2 the determinant map $\det : \mathrm{GL}_n \to U$ is a natural transformation. If $n = 1$ it is an isomorphism of functors.

2. In the category of finite abelian groups the mapping $G \mapsto G^\vee$ that sends each group to its dual $G^\vee = \mathrm{Hom}(G, \mathbb{Q}$ or $\mathbb{Z})$ is a contravariant functor $D : \mathfrak{FAb} \to \mathfrak{FAb}$ that sends every group to an isomorphic group. However, there is no 'natural choice' for an isomorphism $G \xrightarrow{\sim} G^\vee$. We can make this more precise by saying that the co-variant functor $D : \mathfrak{FAb} \to \mathfrak{FAb}^{\mathrm{opp}}$ is not naturally equivalent with the 'identity' $\mathfrak{FAb} \to \mathfrak{FAb}^{\text{opp}}$. The construction $G \mapsto G^{\vee\vee}$ of the double dual on the other hand is a covariant functor $\mathfrak{FAb} \to \mathfrak{FAb}$ that *is* naturally isomorphic with the identity. Another way to put this is to say that a finite group G is *canonically isomorphic* with its double dual $G^{\vee\vee}$. Similar things are true for the construction of the dual vector space in $\mathfrak{FVec}_K$. As with finite groups the finiteness condition is essential to get a canonical isomorphism $V \to V^{**}$.

3. The forgetful functor $\mathfrak{Rng} \to \mathfrak{Sets}$ from rings to sets is isomorphic with the representation functor $\text{Hom}_{\mathfrak{Rng}}(\mathbb{Z}[X], -)$. For every ring R a canonical isomorphism $\text{Hom}_{\mathfrak{Rng}}(\mathbb{Z}[X], R) \xrightarrow{\sim} R$ is given by $f \mapsto f(X)$. In general, a functor $F\mathcal{C} \to \mathfrak{Sets}$ is called a *representable functor* if it is isomorphic to a representation functor. After Grothendieck this concept is of fundamental importance in arithmetic algebraic geometry. Wiles' proof of Fermat's Last Theorem for instance consists largely of showing that certain functors in the theory of elliptic curves are representable.

Definition 2.10. The categories $\mathcal{C}$ and $\mathcal{D}$ are called equivalent if there are functors $F : \mathcal{C} \to \mathcal{D}$ and $G : \mathcal{D} \to \mathcal{C}$ such that $G \circ F$ and $F \circ G$ are isomorphic with the identity on $\mathcal{C}$ and $\mathcal{D}$ respectively. If there are contra-variant functors with this property, $\mathcal{C}$ and $\mathcal{D}$ are called anti-equivalent.

Examples 2.11.

1. Example 2.9.2 shows us that the category $\mathfrak{FAb}$ of finite abelian groups is ani-equivalent with itself under the duality functor D.

2. Let L/K be a finite Galois extension of fields with Galois group G. The *fundamental theorem of Galois theory* says that the category $\mathfrak{Fld}_{L/K}$ of intermediate fields (with the natural inclusions as morphisms) is anti-equivalent with the category $\mathfrak{Sgrp}_G$ of subgroups of G (also with natural inclusions as morphisms). The functors $\mathfrak{Fld}_{L/K} \to \mathfrak{Sgrp}_G$ and $\mathfrak{Sgrp}_G \to \mathfrak{Fld}_{L/K}$ from definition 2.10 are given by $M \mapsto \text{Aut}(L/M)$ and $H \mapsto L^H$.

3^***.** The fundamental theorem of Galois theory for topological spaces says that the category $\mathfrak{Cov}_X$ of coverings of a path connected topological space X is under mild conditions anti-equivalent with the category $\pi(X)$-$\mathfrak{sets}$ of sets with an action of the fundamental group $\pi(X)$. For every point $x \in X$ there is a *fibre functor* $F_x : \mathfrak{Cov}_X \to \mathfrak{Sets}$ that sends a covering $f : Y \to X$ to $f^{-1}(x)$ and the fundamental group $\pi(X, x)$ acts on the set $f^{-1}(x)$ by taking the image of a $y \in f^{-1}(x)$ under the homotopy class of a closed curve $w \subset X$ in x to be the endpoint of the unique path $w^* \subset Y$ with starting point y that projects to w under f.

We've mentioned before that many of the standard constructions for groups, rings and modules are solutions to certain universal problems in the underlying category. One can rephrase many of known defintions in general categorial terms.

Definition 2.12. A product of a family $\{A_i\}_{i \in I}$ of objects in $\mathcal{C}$ is an object $P \in \mathcal{C}$ together with morphisms $p_i : P \to A_i$ with the property that given an object $T \in \mathcal{C}$ and morphisms $f_i : T \to A_i$ there is a unique morhpism $f : T \to P$ such that $p_i \circ f = f_i$.
A co-product or sum of $\{A_i\}_{i \in I}$ in $\mathcal{C}$ is an object $S \in \mathcal{C}$ together with morphisms $\epsilon_i : A_i \to S$ with the property that given an object $T \in \mathcal{C}$ and morphisms $g_i : A_i \to T$ there is a unique morphism $g : S \to T$ such that $g \circ \epsilon_i = g_i$.

Objects with such a universal property are uniquely determined up to isomorphism, *if* they exist.

Examples 2.13.

1. In the category $\mathfrak{Sets}$ a sum of a family of sets is nothing more than the union of these sets. The product of a familty of sets is the *Carthesian product*. If the sets are groups or rings then this product has a natural group or ring structure. We see that the products in $\mathfrak{Grp}$ and $\mathfrak{Rng}$ are constructed in the familiar way.

2. In the category of topological spaces the sum is the same as the disjoint union. For the product one takes the Carthesian product with the well-known product topology.

3. The construction of sems in $\mathfrak{Grp}$ is not so straightforward. In the category $\mathfrak{Ab}$ they can be constructed following the construction for modules, but in the non-abelian case one gets more complicated groups. The sum of two infinite cyclic groups for example is a 'free (non-abelian) group on two generators'. One can show that the group $\mathrm{SL}_2(Z)/\{\pm 1\}$ is the sum of a subgroup of order 2 and one of order 3!

4. For rings one has the well-known product of rings. In $\mathfrak{CRng}$ the *tensor product* of rings as $\mathbb{Z}$-algebras is a sum.

5. In the category of objects over a fixed base object A the product of $X \to A$ and $Y \to A$ is the *fibered product* of X and Y over A. The sum of $A \to X$ and $A \to Y$ is the *fibered sum* of X and Y.

3 Generalising the introduction

Our aim in this section is to generalise the theorem from the first section about the group structure on a cubic curve.

Let k be a field. A *curve* in $\mathbb{P}^2(k)$ is the zero set of a homogeneous polynomial. Its *degree* is the total degree of the polynomial. When we talk about a curve we really need the polynomial and not just the set. This is particularly relevant if k is finite, because then the polynomial really carries more information than the set of points. For example, if k is $\mathbb{F}_2$ then the projective space has seven points and as many lines:

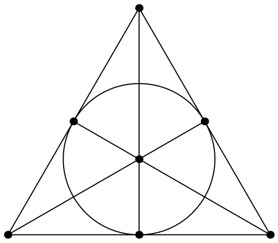

The curves we will consider are cubic curves in $\mathbb{P}^2$. Let $f \in k[X,Y,Z]$ be the homogeneous cubic polynomial

$$f = \sum_{\substack{i+j+k=3 \\ i,j,k \geq 0}} a_{ijk} X^i Y^j Z^k.$$

We define the curve over k to be the set

$$E(k) = \left\{ (x:y:z) \in \mathbb{P}^2(k) \mid f(x,y,z) = 0 \right\}.$$

This is well-defined since f is homogeneous, so $f(x,y,z)$ depends only on the ratios of x, y and z. If $l \supset k$ is an extension field we extend the previous definition by putting

$$E(l) = \left\{ (x:y:z) \in \mathbb{P}^2(l) \mid f(x,y,z) = 0 \right\}.$$

A point $P = (x:y:z) \in E(l)$ is called *singular* if the partial derivatives are all zero at P, i.e., if $f_X(x,y,z) = f_Y(x,y,z) = f_Z(x,y,z) = 0$. Intuitively the curve looks somewhat like this near a singularity:

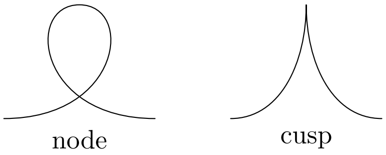

A curve E is called *non-singular* if there are no singular points on $E(l)$ for any algebraic extension $l \supset k$.

Example 3.1. We look at the curve from section 1. Let $u \in k^*$ and put

$$f = X^3 + Y^3 - uZ^3.$$

We see that the partial derivatives are

$$\begin{aligned} f_X &= 3X^2 \\ f_Y &= 3Y^2 \\ f_Z &= 3uZ^2. \end{aligned}$$

If $\mathrm{char}(k) \neq 3$ then there are no singular points in any extension field, since at least one of the coordinates is non-zero for each point in the projective plane. We conclude that in this case the curve is non-singular.
If $\mathrm{char}(k) = 3$ then *every* point is singular. In this case we can factor the polynomial. The cube root of u is in

$$f = (X + Y + \sqrt[3]{u}Z)^3.$$

In general it is true that a *reducible* polynomial f gives rise to a curve with singular points. The factors of f all give rise to curves in $\mathbb{P}^2$. These curves may coindide or one may be contained in another, in which case all the points in the intersection will be singular. Otherwise we are in a situation where we can apply Bézout's theorem to find intersection points in $\mathbb{P}^2(\overline{k})$, which are singular.

The *tangent line* to a curve given by f at a nonsingular point $P = (x : y : z)$ is the line defined by

$$f_X(x, y, z)X + f_Y(x, y, z)Y + f_Z(x, y, z)Z = 0.$$

Note that the point P is on this line, since $f_X + f_Y + f_Z = deg(f)f$.

Now we need to define the *intersection multiplicity* of a line L with the curve E at a point P, $i(L, E; P)$. We would like to have $i(L, E; P) = 0$ if P is not in the intersection and > 0 if it is. Furthermore, we want $i(L, E; P) > 1$ if L is the tangent line to E at P.

We first examine the equivalent notions in k and $\mathbb{P}^1(k)$. Let $g \in k[X] - \{0\}$ and $a \in k$. Recall that there is a map

$$\begin{aligned} \mathrm{ord}_a : k[X] - \{0\} &\longrightarrow \mathbb{Z}_{\geq 0} \\ g &\mapsto \max\left\{i \in \mathbb{Z}_{\geq 0} \mid g \in (X - a)^i\right\}. \end{aligned}$$

We can generalise this to the projective line as follows. For $g \in k[X, Y]$ homogeneous, $g \neq 0$, and $P = (a : b) \in \mathbb{P}^1(k)$ we define

$$\mathrm{ord}_P(g) = \max\left\{i \in \mathbb{Z}_{\geq 0} \mid g \in (bX - aY)^i\right\}.$$

We use this as a motivation for our definition of intersection multiplicity in $\mathbb{P}^2$.

Definition 3.2. Let E be the curve in $\mathbb{P}^2$ given by the polynomial f and let L be the projective line $L = \{(x : y : z) \in \mathbb{P}^2(k) \mid ax + by + cz = 0\}$ with a, b and c in k not all zero. Let $P = (x : y : z)$ be a point in $\mathbb{P}^2(k)$. The intersection multiplicity $i(L, E; P) = 0$ if $P \notin L \cap E(k)$. For points in the intersection it is given by

$$i(L, E; P) = \max\left\{i \in \mathbb{Z}_{\geq 0} \mid f \in (aX + bY + cZ) + (yX - xY, zX - xZ, zY - yZ)^i\right\}.$$

Now we are finally at the point where we can formulate the generalisation of the theorem from section 1.

Theorem 3.3. *Let k be a field, E a non-singular cubic curve over k and $O \in E(k)$. Then there is a unique abelian group structure on $E(k)$ such that O is the zero element and if l and m are any two lines in $\mathbb{P}^2$ that intersect $E(k)$ in three points counting multiplicities then the sum of the points on l is the same as the sum of the points on m.*

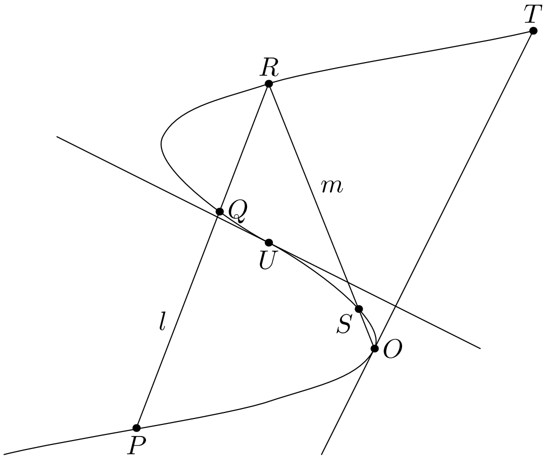

Note that the unit of the group is no longer required to be 'at infinity'. From the theorem we can derive from the picture above how the addition works in this group. Take l the line through P and R and m the line through R and O. The theorem says that $P + Q + R = R + S + O$, or $S = P + Q$. To determine exactly what value the sum of the points on a line is always equal to we look at the tangent at O and conclude that this value is $O + O + T$, that is, T. So if we take an inflection point, like U, as our zero element, the sum of the points on a line will in fact be $U = 0$.

We end this section by stating a famous theorem concerning the structure of this group.

Theorem 3.4 (Mordell-Weil). *Let k be an algebraic number field, E a non-singular cubic curve over k and $O \in E(k)$. With the group law from theorem 3.3, $E(k)$ is finitely generated abelian group.*

Example 3.5. Let E_u the curve defined by the polynomial $X^3 + Y^3 - uZ^3$ that we have seen before. In this case

$$\begin{array}{lcll} E_u(\mathbb{Q}) & \cong & \mathbb{Z}/3\mathbb{Z} & \text{if } u \text{ is a cube} \\ & \cong & \mathbb{Z}/2\mathbb{Z} & \text{if } u \text{ is twice a cube} \\ & \cong & \mathbb{Z}^{r(u)} & \text{otherwise} \end{array}$$

The number $r(u)$ is called the Mordell-Weil rank.

Although there are methods to compute this rank, the known algorithms to do this for an arbitrary elliptic curve are not proven to terminate in finite time. In fact, no one has an algorithm which guarantees to tell you a rational point on a curve for a class of curves including the cubic curves we considered here, if there is such a point, or to prove that there is not. Algorithms for both exist, but it just isn't known if they terminate in finite time.

4 Places and valuations

In the category $\mathfrak{Field}$ of fields, all the morphisms are injective. In a sense, there are not enough of them. *Places* are designed to do something about this. Let k be a field and consider the 'map'

$$\begin{array}{rcl} k(t) & \dashrightarrow & k \\ f & \longmapsto & f(3). \end{array}$$

If $f = g/h$ with $g, h \in k[t]$ coprime, then $f(3) = g(3)/h(3)$, which is only well-defined if h does not have a zero at 3, i.e., $h \notin (t-3)$. Note that if $h \in (t-3)$ then $g \notin (t-3)$, so the image of $f^{-1} = h/g$ is well-defined and is in fact 0 since $h(3) = 0$.

Definition 4.1. Let K, L be fields. A *place from K to L* is a pair (R, f) with $R \subset K$ a subring and $f : R \to L$ a ring homomorphism, such that if $x \in K \setminus R$ then $x^{-1} \in R$ and $f(x^{-1}) = 0$.

In the example one is tempted to say that the value of g/h with $h \in (t-3)$ in 3 is infinite. This motivates the following equivalent definition. Recall that if $P = (a : b) \in \mathbb{P}^1(K)$ then $P^{-1} = (b : a)$.

Definition 4.2. A *place* $K \dashrightarrow L$ is a map $f : \mathbb{P}^1(K) \to \mathbb{P}^1(L)$ such that
(1) $f^{-1}L$ is a subring of K and $f|_{f^{-1}L} : f^{-1}L \to L$ is a ring homomorphism.
(2) $f(P^{-1}) = f(P)^{-1}$ for all $P \in \mathbb{P}^1(K)$.

If K and L are extensions of a field k, then a *place over k* is a place (R, f) with $k \subset R$ and $f|_k = \mathrm{id}_k$. The places over a given field k form a category (see exercise 18).

Definition 4.3. Let K be a field. A *valuation ring of K* is a subring $R \subset K$ with the property that for all $x \in K \setminus R$ one has $x^{-1} \in R$.

Remark 4.4. Any valuation ring of a field K is a local ring, i.e., it has exactly one maximal ideal.

Proof. A ring R is local if and only if the set of non-units is an additive subgroup (exercise). So suppose $x, y \in R$ with $x, y \notin R^*$. We have to prove that $x + y \notin R^*$. This is obvious if $x = 0$ or $y = 0$. Otherwise, either x/y or y/x is in R, say $x/y \in R$. Then $x + y = (x/y + 1)y \in Ry \subsetneq R$. □

Corollary 4.5. If R is a valuation ring of K, then there is a place $K \dashrightarrow R/\mathfrak{m}$ with $\mathfrak{m} \subset R$ the maximal ideal.

Definition 4.6. A *valuation ring of K/k* is a valuation ring R of K with $k \subset R$. A *valuation ring* is a domain that is a valuation ring of its field of fractions.

Suppose A is a unique factorization domain and π is a prime element of A. Then $R = \{\frac{a}{b} \in Q(A) \mid a, b \in A, b \notin \pi A\}$ is a valuation ring. Recall that $Q(A)$ is the field of fractions of A. Actually, it is a *discrete valuation ring*, i.e., a PID with exactly one prime element up to units.

Remark 4.7. A Noetherian valuation ring is either a field or a discrete valuation ring. A regular local ring of dimension 0 is a field, one of dimension 1 is a DVR.

An important question in the context of valuation rings is that of extensions. If R is a valuation ring of a field K and L is a finite extension of K, what are the valuation rings of L extending R? Or, phrased in the setting of curves $D \to C$, what are the points of D mapping to a given point of C?

5 Affine varieties over algebraically closed fields

In this section we take k to be an algebraically closed field. Our aim is to show that there is an anti-equivalence of categories between the affine varieties over k and k-algebras of finite type that are domains. Let's begin by specifying what all those things mean.

Definition 5.1.
An *affine variety* is an irreducible algebraic set.
An *algebraic set* is a set of the form $V = \{x = (x_1, \dots, x_n) \in k^n \mid f_1(x) = 0, \dots, f_m(x) = 0\}$, where n and m are non-negative integers and $f_1, \dots, f_m \in k[X_1, \dots, X_n]$.
A non-empty algebraic set V is called *irreducible* if and only if it is not the union of two proper subsets that are also algebraic. Equivalently, V is irreducible if the $f_1, \dots, f_m$ defining V have the property that $(f_1, \dots, f_m) \subset k[X_1, \dots, X_n]$ is a prime ideal.

A morphism between varieties $V \subset k^{n_1}$ and $W \subset k^{n_2}$ is a map $V \to W$ such that there are $g_1, \dots, g_{n_2} \in k[X_1, \dots, X_{n_1}]$ such that $x \in V$ is mapped to $(g_1(x), \dots, g_{n_2}(x)) \in W$. Varieties and their morphisms form a category.

Definition 5.2. Let R be a commutative ring. An *R-algebra* is a commmutative ring A together with a ring homomorphism $R \to A$. A morphism of R-algebras from $R \to A$ to $R \to B$ is a ring homomorphism $A \to B$ such that the following diagram commutes.

$$\begin{tikzcd}
A \arrow[rr] & & B \\
& R \arrow[ul] \arrow[ur] &
\end{tikzcd}$$

An R-algebra A is of *finite type* (f.t.) if there is a finite subset $S \subset A$ such that A is the smallest subring of A containing both S and the image of R. Equivalently: there exists $n \in \mathbb{Z}_{\geq 0}$ and an ideal $I \subset R[X_1, \dots, X_n]$ such that $A \cong R[X_1, \dots, X_n]/I$ as R-algebras.

The anti-equivalence of categories mentioned above comes from the functor F that sends a variety V to $k[X_1, \dots X_n]/I(V)$, where $I(V)$ is the ideal of polynomials that vanish on V:

$$I(V) = \{f \in k[X_1, \dots, X_n] \mid \text{for all } x \in V : f(x) = 0\}.$$

We can also view F as the functor $V \mapsto \mathrm{Hom}_{\mathfrak{Var}}(V, k)$ the set of morphisms from V to k as affine algebraic varieties over k. This enables us to see what F does with morphisms:

$$\begin{tikzcd}
V \arrow[r, "F"] \arrow[d, "g"'] & \mathrm{Hom}(V, k) \\
W \arrow[r, "F"'] & \mathrm{Hom}(W, k) \arrow[u, "F(g)"']
\end{tikzcd}$$

If $\mathcal{C}$ is the category of affine algebraic varieties over k and $\mathcal{D}$ is the category of k-algebras of finite type that are domains, then the anti-equivalence means that there is a contravariant functor $G : \mathcal{D} \to \mathcal{C}$ such that $F \circ G \cong \mathrm{id}_{\mathcal{D}}$ and $G \circ F \cong \mathrm{id}_{\mathcal{C}}$. Explicitly constructing this G is a bit of a hassle, so instead we employ the following theorem from category theory.

Theorem 5.3. *The following are equivalent:*
(1) F is an anti-equivalence $\mathcal{C} \to \mathcal{D}$
(2) There is a functor $G : \mathcal{D} \to \mathcal{C}$ such that $F \circ G \cong \mathrm{id}_{\mathcal{D}}$ and $G \circ F \cong \mathrm{id}_{\mathcal{C}}$.
(3) Every object of $\mathcal{D}$ is $\mathcal{D}$-isomorphic to one of the form $F(V)$ with V an object of $\mathcal{C}$, and for all $V, W \in \mathrm{Ob}(\mathcal{C})$, F gives a bijection $\mathrm{Hom}_{\mathcal{C}}(V, W) \to \mathrm{Hom}_{\mathcal{D}}(F(W), F(V))$.

The first statement of the last property is clear, the second is a bit of a notational nightmare to check, but nothing difficult happens.

6 Algebraic foundation of geometric notions

In the previous section we saw that for an algebraically closed field k there is an anti-equivalence of categories

$$\{\text{affine varieties over } k\} \quad \longleftrightarrow \quad \{k\text{-algebras of finite type that are domains}\}\,.$$

There are two problems with this. Firstly, the fields we want to consider are generally not algebraically closed and secondly, elliptic curves are projective varieties, not affine ones. In this section we shall discuss our strategy to extend our work to fields that need not be algebraically closed. If one relaxes this restriction on k, one runs into some problems in the category on the left hand side. On the right hand side, dropping the restriction does not give the same problems. With this in mind we make the following defintion.

Definition 6.1. For any field k we define an affine variety V over k to be a k-algebra of finite type that is a domain. A morphism of varieties $V \to W$ is a homomorphism of k-algebras going in the opposite direction.

The problem with this definition is that it takes away a lot of the geometric intuition we have. Given a field k and a k-algebra A of finite type that is a domain, how, for instance, do we think about the points of the corresponding variety V? We give two answers to this question:

First answer. Write A as $k[X_1,\dots,X_n]/(f_1,\dots,f_m)$. If L is any field extension of k we define the set of *L-rational points* of V as

$$L(V) = \{x \in L^n \mid f_1(x) = \ldots = f_m(x) = 0\}\,.$$

Remark 6.2. The set $L(V)$ is independent of the chosen representation of A, for we have

$$L(V) \cong \mathrm{Hom}_{k-\mathfrak{Alg}}(A, L).$$

Proof. A k-algebra morphism from $k[X_1,\dots,X_n]$ is uniquely determined by the images of the X_i, so $\mathrm{Hom}(k[X_1,\dots,X_n], L)$ corresponds bijectively with L^n. A morphism $\phi : k[X_1,\dots,X_n] \to L$ factors through A if and only if $(f_1,\dots,f_m) \subset \ker(\phi)$. The action of ϕ on the f_i is simply evaluating them in the point corresponding to ϕ, so we see that ϕ factors through A if and only if the point corresponding to ϕ is mapped to zero by all the f_i. □

Corollary 6.3. From the remark it follows that V is in fact a functor from the category of field extensions of k to the category of sets.

Second answer. We can also view the prime ideals of A as the points of V. We will now show this to be equivalent to the first answer:
Let L be a field extension of k and $g : A \to L$ a k-algebra homomorphism. Then $\ker(g)$ is a prime ideal since the image $g(A)$ is a subring of the field L and therefore is a domain. Conversely, if $\mathfrak{p} \subset A$ is a prime ideal then $A/\mathfrak{p}$ is a domain and putting $L = Q(A/\mathfrak{p})$ equal to the field of fractions we see that there is a natural map of k-algebras $A \to A/\mathfrak{p} \subset L$ of which the kernel is $\mathfrak{p}$.

With this basic question answered, we move on to translating more complicated geometric notions into the algebraic setting. Supposing again that k is any field, A is a k-algebra of finite type that is a domain and V is the corresponding variety, how do we define the dimension $\dim(V)$ of V?

First possibility. We take the dimension to be the transcendence degree of the field of fractions $Q(A)$ of A. That is,

$$\dim(V) = \max\{m \in \mathbb{Z}_{\geq 0} \mid \exists f : k[X_1,\dots,X_m] \to Q(A) \text{ an injective } k\text{-algebra morphism}\}$$

This definition somewhat reflects our intuitive concept of dimension as the number of independent directions in which the object extends. However, it isn't a very good definition to work with in practice.

Second possibility. There is a general concept of dimension for a commutative ring, the *Krull dimension* of the ring. It is the length of the longest chain of prime ideals one can find inside A:

$$\mathrm{Krulldim}(A) = \max \{m \in \mathbb{Z}_{\geq 0} \mid \exists \mathfrak{p}_0, \ldots, \mathfrak{p}_m \text{ prime such that } \mathfrak{p}_0 \subsetneq \mathfrak{p}_1 \subsetneq \ldots \subsetneq \mathfrak{p}_m \subsetneq A\} .$$

There is a theorem that states that for our A (k-algebras of finite type that are domains) these two possibilities will produce the same number.

Examples 6.4. Since we are mainly interested in curves and the points on them, it is convenient to look at the A for which the dimension is 0 or 1. Suppose that A is a domain, then $(0) \subset A$ is a prime ideal. From this we see that

$$\begin{array}{lcl} \mathrm{Krulldim}(A) = 0 & \iff & A \text{ is a field.} \\ \mathrm{Krulldim}(A) = 1 & \iff & A \text{ is not a field and every nonzero prime ideal is maximal.} \end{array}$$

Cases in which this second one is true include $A = \mathbb{Z}$ or $k[X]$ or more generally A a PID, DVR or Dedekind domain.

7 Kähler differentials

In this section we will discuss the tangent space of an affine variety at a point. As usual, we suppose k is a field, A a k-algebra of finite type that is a domain and V the corresponding affine variety.

Given a presentation $A = k[X_1, \ldots, X_n]/(f_1, \ldots, f_m)$ and a k-rational point $P = (x_1, \ldots, x_n)$ of V we define the tangent space $T_P(V)$ of V at P to be the zero set in k^n of

$$\left\{ \sum_{i=1}^{n} \left(\frac{\partial f_j}{\partial X_i}(x) \right) (X_i - x_i) \right\}_{j=1}^{m} .$$

This is a finite dimensional k vector space. Its dimension is in fact

$$\dim_k(T_p(V)) = n - \mathrm{rk} \left(\frac{\partial f_j}{\partial X_i}(x) \right)_{\substack{1 \leq i \leq n \\ 1 \leq j \leq m}} .$$

The point P is called *non-singular* or *smooth* if $\dim_k(T_P(V)) = \dim(V)$.

For a field extension L of k we define the tangent space of an L-rational point in the same way as above, except that everything lives over L now. Note that we now get an L vector space. An L-rational point P is called non-singular or smooth if $\dim_L(T_P(V)) = \dim(V)$. A variety V is called *non-singular* or *smooth* if all of its points over all extensions are non-singular.

The above definition depends rather heavily on a number of choices. A much nicer way to look at it is using *Kähler differentials*.

Definition 7.1. Let $g : R \to A$ be a morphism of commutative rings and let M be an A-module. A map $d : A \to M$ is called an *R-derivation* if

1. d is an additive group homomorphism.
2. For all $a, b \in A$ we have $d(ab) = ad(b) + bd(a)$. (Leibniz' rule)
3. For all $r \in R$ we have $d(g(r)) = 0$.

An R-derivation $d : A \to M$ is called *universal* if for all R-derivations $\delta : A \to N$ there is a unique A-linear map $h : M \to N$ such that $\delta = h \circ d$. In other words, $d : A \to M$ is universal precisely if for all A-modules N we have a bijection between $\mathrm{Hom}_A(M, N)$ and $\mathrm{Der}_R(A, N)$, the set of R-derivations from A to N.

Example 7.2. If $A = R[X]$ is a polynomial ring then the map $A \to A$ that sends a polynomial to its derivative is a universal R-derivation. More generally, if $A = R[X_i \mid i \in I]$ then the map

$$\begin{array}{ccc} A & \longrightarrow & \bigoplus_{i \in I} A \\ f & \longmapsto & \left(\dfrac{\partial f}{\partial X_i}\right)_{i \in I} \end{array}$$

is a universal R-derivation.

Being a universal object, a universal R-derivation, if it exists, is unique up to unique isomorphism. Suppose that $d : A \to M$ and $d' : A \to M'$ are two universal R-derivations. Then by the universality of d there is a unique A-linear map $h : M \to M'$ such that $d' = h \circ d$. Similarly, by the universality of d' there is a unique A-linear map $h' : M' \to M$ such that $d = h' \circ d'$. Now we see that $d = (h' \circ h) \circ d$. We also know that $d = \mathrm{id}_M \circ d$ and by the universality of d these maps must therefore be the same, so $h' \circ h = \mathrm{id}_M$. By symmetry we also have $h \circ h' = \mathrm{id}_{M'}$ so that we see that h is an isomorphism between M and M' and by the universality it is unique.

Theorem 7.3. *For every morphism $R \to A$ of commutative rings there is a universal R derivation $d : A \to \Omega_{A/R}$. The module $\Omega_{A/R}$ is called the module of* Kähler differentials.

Proof. We will not give a full proof here, we shall merely indicate what $\Omega_{A/R}$ and d are.
Fix a representation $A = R[X_i : i \in I]/(f_j : j \in J)$ of A. For an element $f \in R[X_i : i \in I]$ denote by $\overline{f}$ its class in A. We define

$$\Omega_{A/R} = \left(\bigoplus_{i \in I} A\right) \Bigg/ \left(\sum_{j \in J} A \left(\overline{\frac{\partial f_j}{\partial X_i}}\right)_{i \in I}\right).$$

For d we take the map

$$\begin{array}{ccc} A & \longrightarrow & \Omega_{A/R} \\ \overline{f} & \longmapsto & \left(\dfrac{\partial f}{\partial X_i}\right)_{i \in I} \end{array}.$$

□

Example 7.4. Let $R = k$ be a field and $A = k[X, Y]/(X^2 + Y^2 - 1)$. We see that we can take I to be a two element set and J a one element set. The module of Kähler differentials is

$$\Omega_{A/R} = (A \oplus A)/A(2X, 2Y)$$

according to our theorem.
If k does not have characteristic 2 we observe that in A

$$\det \begin{pmatrix} 2X & 2Y \\ \frac{-1}{2}Y & \frac{1}{2}X \end{pmatrix} = 1.$$

We conclude that $(2X, 2Y)$ and $(\frac{-1}{2}Y, \frac{1}{2}X)$ form a basis of $A \oplus A$ and therefore $\Omega_{A/R} \cong A$.

If l/k is a field extension we can wonder what $\Omega l/k$ is. We know that it is an l-module, i.e. a vector space over l, so the thing that is of most interest is the dimension of this space. We won't deal with the general case here, but restrict ourselves to the case where l is *finitely generated* as a field extension of k, that is

$$\exists S \subset l : \#S < \infty \text{ and } l = k(S).$$

Let n be the transcendence degree of l over k. If $\text{char}(k) = 0$ we have $\dim_l(\Omega_{l/k}) = n$ and if $\text{char}(k) = p > 0$ then $\dim_l(\Omega_{l/k}) \geq n$ with equality if and only if l is *separably generated*, i.e., if and only if there are $t_1, \ldots, t_n \in l$ such that the extension $k(t_1, \ldots, t_n) \subset l$ is finite and separable.

Example 7.5. For the purpose of this course we are only interested in the case where n is 0 or 1. In these cases the previous condition is not so very difficult.
For $n = 0$ an extension l/k is separably generated is and only if l/k is finite and separable.
For $n = 1$ it is separably generated if and only if there is a $t \in l$ that is transcendental over k such that $k(t) \subset l$ is a finite separable extension.

8 Projective curves

In the previous sections we have seen how to generalise from affine varieties over an algebraically closed field to affine varieties over arbitrary fields. This is mainly a process of translating the geometric notions into algebra, where the generalisation is straightforward. In this section we will look at projective varieties. We do this over arbitrary fields straightaway, working just with the algebraic terminology. We restrict ourselves to the only cases of interest to us in this course, curves and points.

Definition 8.1. Let k be a field. The category of *projective regular varieties of dimension at most* 1 over k is the opposite of the category whose objects are the finitely generated field extensions of k of transcendence degree at most 1 and whose morphisms are the *places* over k. Objects of transcendence degree (tr.deg) 0 are refered to as points and those of tr.deg 1 as curves.

Although this is a perfectly good definition, it appears somewhat unsatisfying. One feels that adjectives like *projective* and *regular* and "*of dimension at most* 1" should have a meaning of their own, making the statement above into a theorem rather than a definition. In algebraic geometry (where the terminology comes from) this is indeed the way one goes about this. However, treating varieties in such generality is not the objective of this course. This being said, we can still provide at least an intuitive feel for the meaning of these notions.
Projective one should think of as being the zero set in $\mathbb{P}^n$ of some set of homogeneous polynomials.
Regular is a technical condition that corresponds roughly to *smooth* or *non-singular*. Indeed smooth implies regular and the converse is also true if k is perfect. In this context the condition for smoothness is relatively easy: $\dim(\Omega_{l/k}) = \text{tr.deg}(l/k)$.
Dimension ≤ 1 corresponds to the $\text{tr.deg}(l/k) \leq 1$ in our definition. It is a condition we impose since this is all we need and it makes life easier for us along the way.

Example 8.2. To show that it is possible to translate a homogeneous equation in three variables into a projective curve in the sense just defined, we give an example. The generalisation to other equations is straightforward.

Suppose we want to look at the variety defined over k by $X^3 + Y^3 - uZ^3 = 0$, where $u \in k^*$ and $\text{char}(k) \neq 3$. This last condition ensures that the curve is non-singular. The corresponding field is then constructed in the following manner. Set one of the unknowns equal to 1, say Z. This gives us a polynomial in two variables. If we divide out the polynomial ring in these variables by that polynomial, we will get a domain. In general this is a consequence of starting with an non-singular equation, in this particular case it is easily verified. The field of fractions of this domain is the field that defines our variety:

$$l = Q\left(k[X, Y]/(X^3 - Y^3 - u)\right).$$

To see that this is a field of the required type we observe that we can also view the field above as

$$l = k(X)[Y]/(Y^3 - (u - X^3)),$$

which is a finite extension of the field $k(X)$ of tr.deg 1. The construction above appears to depend upon a choice of which variable to set to 1 (instead of a variable, it is also possible to fix the value of one arbitrary linear relation in X, Y and Z). A way to define l that is independent of such a choice is given by

$$l = \left\{ \frac{a}{b} \;\middle|\; a, b \in k[X,Y,Z]/(X^3 - Y^3 - uZ^3),\ a, b \text{ homogeneous of the same degree, } b \neq 0 \right\}.$$

Recall that for an affine variety V we can view the domain $A = k[V]$ as an algebra of functions on V. An $f \in A$ defines a function from $V(k)$ to k. Something along the same lines is true for projective varieties. For an $f \in l = k(V)$ we don't get a map from $V(k)$ to k, but one to $\mathbb{P}^1(k)$, or equivalently, a *place* $V(k) \dashrightarrow k$.

Definition 8.3. Let l/k be a finitely generated field of tr.deg $n \in \{0, 1\}$. Write C_l for the corresponding variety. We now give definitions of some important concepts related to this variety.
Rational points. For L/k any field extension we define the set of L-rational points of C_l as

$$C_l(L) = \{\text{places } l \dashrightarrow L \text{ over } k\}.$$

Points of C_l. So far we haven't really assigned any meaning to C_l other than 'the variety corresponding to l'. Yet we want to think of C_l as a set of points. The way to do this is to set

$$C_l = \{\text{valuation rings of } l \text{ over } k\}.$$

Note that if $n = 0$ then there is only one point, we have $C_l = \{l\}$. For $n = 1$ l is still a point, but there are others as well. We call l the *generic point* and the others the *closed points*.

Topology on C_l. We can equip the set C_l with a topology by setting

$$U \subset C_l \text{ is open} \quad \Longleftrightarrow \quad \text{either } C_l \setminus U \text{ is finite and } l \in U, \text{ or } U = \emptyset.$$

For $U \subset C_l$ open we define

$$\mathcal{O}(U) = \begin{cases} 0 & \text{if } U \text{ is empty} \\ \bigcap\limits_{R \in U} R & \text{otherwise.} \end{cases}$$

Example 8.4. Look at $V = \mathbb{P}^1(k) = \{(x : y) \mid x, y \in k\}/k^*$. The function field here is $l = k(t)$ where t is a transcendental that we can think of as $t = x/y$. The following theorem tells us what the points of C_l are.

Theorem 8.5. *Let k be a field and $l = k(t)$ with t transcendental over k. Then*
(a) *l is a valuation ring of l/k.*
(b) *if $f \in k[t]$ is a monic irreducible polynomial then*
$R_f = \left\{\frac{a}{b} \mid a, b \in k[t], b \notin (f)\right\}$ is a valuation ring of l/k.
(c) *$R_\infty = \left\{\frac{a}{b} \mid a, b \in k[t], b \neq 0, \deg(a) \leq \deg(b)\right\}$ is a valuation ring of l/k.*
(d) *each valuation ring l/k is of one of the above types.*

Proof. The first three claims are easily verified.
For (d) we take R a valuation ring of l/k with maximal ideal $\mathfrak{m}$. We know that $k \subset R$. Suppose that we have $t \in R$, then in fact $k[t] \subset R$. It follows that $\mathfrak{p} = \mathfrak{m} \cap k[t]$ is a prime ideal of $k[t]$ and $k[t] \setminus \mathfrak{p} \subset R^*$. We have either $\mathfrak{p} = 0$ or $\mathfrak{p} = (f)$ for some monic irreducible $f \in k[t]$.
If $\mathfrak{p} = 0$ then every $b \in k[t]$ unequal to 0 is in R^*, so $b^{-1} \in R$. We see that $R = k(t) = l$ in this case.
If $\mathfrak{p} = (f)$ then $R_f \subset R \subset l$ and since the R_f are maximal proper subrings of l we see that $R_f = R$.
Now we are left with the case $t \notin R$. We change variables to $u = t^{-1}$. Then $u \in R$ so by the above case there is an $f \in k[u]$ monic irreducible with $R = \left\{\frac{a}{b} \mid a, b \in k[u], b \notin f \cdot k[u]\right\}$. But we know more than this, since $t = u^{-1}$ is not in R, so $u \in f \cdot k[u]$, that is, $f = u$ and $R = R_\infty$. □

Corollary 8.6. If $k = \overline{k}$ then all the irreducible polynomials in $k[t]$ are linear and the map

$$\begin{array}{rcl} \mathbb{P}^1(k) & \longrightarrow & C_{k(t)} \setminus \{\text{generic point}\} \\ \alpha \in k & \longmapsto & R_{t-\alpha} \\ \infty & \longmapsto & R_\infty \end{array}$$

is a bijection.

Before we move on we fix a little bit of notation. For psychological reasons we like to call the points of a variety P, yet this is not a name we commonly use to denote a ring. So if we want to think of the point P in C_l as a ring, we write $\mathcal{O}_P$ for it. Its maximal ideal we denote by $\mathfrak{m}_P$.

Examples 8.7. The first thing to think about is what happens in dimension 0. We have l/k a finite algebraic extension. As we've seen before $C_l = \{l\}$, since every valutation ring is integrally closed. What about the L-rational points for any field extension L of k? We compare two examples:

1. If $l = k$ then the set $C_k(L)$ will always have one element. By definition it consists of the places $k \to L$ over k, but that leaves us with no choice whatsoever, all of k must be sent into L and the morphism should in fact be the identity on k.

2. Suppose we have $i \in \overline{k} \setminus k$ with $i^2 = -1$ and let $l = k(i)$. Again $C_l(L)$ is the set of places $k(i) \to L$ over k, which we can identify with the set of field embeddings $k(i) \to L$ over k. Such an embedding is completely determined by where we send i. It must go to a root of $X^2 + 1$ in L, for which there are either zero or two choices.

From the above example we see that C_k really behaves like a point, but for extension fields the behaviour is a little more complicated. For extensions we have that l/k is smooth if and only if l/k is separable.

Finally we look at the morphisms from any variety V to C_k. If our intuition of C_k being a single point is correct, then there should always be precisely one such morphism. The definition tells us that the morphisms are places $k \to k(V)$ over k, that is, field embeddings $k \to k(V)$ over k, of which there is exacly one.

9 Closed points

Our aim in this section is to understand the closed points of the curves we have just defined. We've already seen what these points look like for the projective space.

In the case that k is algebraically closed, we saw that there is a bijection between the set of closed points and the set $\mathbb{P}^1(k)$. For arbitrary fields something similar holds, except that we can't really 'see' all the points of $\mathbb{P}^1(k)$ anymore. Instead we have

$$\{\text{closed points of } k(t)/k\} \overset{\sim}{\longleftrightarrow} \mathbb{P}^1(k)/\sim_k,$$

where the equivalence relation is given by

$$(x : y) \sim_k (z : w) \Longleftrightarrow \exists \sigma \in \mathrm{Aut}_k(\overline{k}) : (\sigma x, \sigma y) = (z : w).$$

However, the points of our varieties carry additional algebraic structure than we can see from this set-theoretic interpretation. We first state some of the properties that the closed points of the projective line have and then we will show they extend to arbitrary curves.

Theorem 9.1. *Let k be a field. The closed points of $l = k(t)$ (i.e., of $\mathbb{P}^1_k$) have the following properties:*

1. *they are* discrete valuation rings *(DVRs), that is, principle ideal domains with exactly one nonzero prime ideal.*

2. *the* residue class fields $R/\mathfrak{m}_R = \mathcal{O}_P/\mathfrak{m}_P = k(P)$ *are finite extensions of k. For a point P we call* $\deg P = [k(P) : k]$ *the (residue class) degree of P. For $R = R_f$ we have* $\deg R = \deg f$ *and* $\deg R_\infty = 1$.

3. *the following* product formula *is satisfied*

$$\forall x \in l^* : \sum_{\text{Pclosed}} \deg_P \cdot \mathrm{ord}_P x = 0.$$

Proof. 1. This part is immediately verified by inspection.

2. Let R be a valuation ring of $k(t)$. Again the proof goes by inspection. We know that $R = R_f$ from some $f \in k[t]$ monic irreducible or $R = R_\infty$. In the first can ewe can look at the ring homomorphism $R_f \to k[t]/(f)$ sending $\frac{a}{b}$ to $\frac{a(\mathrm{mod} f)}{b(\mathrm{mod} f)} \in k[t]/(f)$. Note that $b(\mathrm{mod} f)$ is invertible by construction of R_f. The kernel of this map is a maximal ideal of R, since the map goes surjectively to a field, and then it must be $\mathfrak{m}_R$ since this is the only maximal ideal of R. We conclude that $R/\mathfrak{m}_R$ is a finite extension of k of degree $\deg(f)$. For R_∞ we make the substitution $u = t^{-1}$, after which we are back at the first case with $f = u$. We conclude that $R/\mathfrak{m}_R$ is in fact equal to k in this case.

3. A proof of the product formula will follow later.

□

Remark 9.2. For $l = k(t)$ we have the following exact sequence of abelian groups:

$$\begin{tikzcd}
1 \arrow[r] & k^* \arrow[r] & l^* \arrow[r] & \bigoplus_P \mathbb{Z} \arrow[r] & \mathbb{Z} \arrow[r] & 0. \\
 & & x \arrow[r, mapsto] & (\mathrm{ord}_P x)_P & & \\
 & & & (n_P)_P \arrow[r, mapsto] & \sum_P (\deg P) n_P &
\end{tikzcd}$$

Our aim in this section is to extend the properties from the first theorem to arbitrary curves. For l is a finite extension of $k(t)$ we have the inclusion map from $k(t)$ to l. This gives us the following morphism of varieties from C_l to $C_{k(t)} = \mathbb{P}^1_k$:

$$\begin{tikzcd}
C_l \arrow[r] & C_{k(t)} \\
T \arrow[u, phantom, sloped, "\in"] \arrow[r, mapsto] & R = T \cap k(t). \arrow[u, phantom, sloped, "\in"]
\end{tikzcd}$$

Our strategy for understanding the curve C_l is to look at the fibers of this map. Their important properties are summarised in the following theorem.

Theorem 9.3. *Let R be a discrete valuation ring and let K be its field of fractions. Let L/K be a field extension of finite degree n and let $\mathcal{T}$ be the set of all valuation rings of L that* lie above *R, that is,*

$$\mathcal{T} = \{T : T \textit{ is a valuation ring of } L \textit{ with } T \cap K = R\}.$$

Then

1. *$\mathcal{T}$ is finite and non-empty. Every $T \in \mathcal{T}$ is a DVR.*

2. *for each $T \in \mathcal{T}$ the* ramification index

$$e(T/R) = [L^* : T^* K^*]$$

is finite.

3. *for each $T \in \mathcal{T}$ the* residue class degree

$$f(T/R) = [T/\mathfrak{m}_T : R/\mathfrak{m}_r]$$

is also finite.

4. *the ramification indices and residue class degrees of all $T \in \mathcal{T}$ satisfy*

$$\sum_{T \in \mathcal{T}} e(T/R) f(T/R) \leq n.$$

Equality holds if and only if the integral closure of R in L is finitely generated as an R-module. This is true if L/K is separable or if K has a subfield k and element t transcendental over k such that $k \subset R$ and $[K : k(t)] < \infty$.

Proof. 1. Consider the following diagram, where A is the integral closure of R in L:

$$\begin{array}{ccc} L & \supset & A \\ \cup & & \cup \\ K & \supset & R. \end{array}$$

Since R is a DVR it is also a Dedekind domain and by a well-known theorem on Dedekind domains, it follows that A is also Dedekind. For a past exercise we know that

$$\bigcap_{T \in \mathcal{T}} T = A.$$

Let $T \in \mathcal{T}$ and let $\mathfrak{m}_T$ be its maximal ideal. The prime ideal $\mathfrak{p}_T = \mathfrak{m}_T \cap A$ of A lies above the prime ideal $\mathfrak{p}_T \cap R$ of R and since this is not zero and R is a DVR we see that in fact $\mathfrak{p}_T \cap R = \mathfrak{m}_R$. Since $\mathfrak{p}_T$ is in $\mathfrak{m}_T$ and $A \subset T$ we have $A_{\mathfrak{p}_T} \subset T$. Now $A_{\mathfrak{p}_T}$ is a DVR since A is Dedekind. Since it is a valuation ring, there can be no rings strictly between $A_{\mathfrak{p}_T}$ and its field of fractions L. We know that $T \neq L$, hence $T = A_{\mathfrak{p}_T}$.

It follows that every T is of the from $A_{\mathfrak{p}_T}$ with $\mathfrak{p}_T$ lying above $\mathfrak{m}_R$. So every T is a DVR and there are finitely many. We can find them using the following procedure:

- Determine A.
- Write $\mathfrak{m}_R = \pi R$.
- Decompose πA as a product of prime ideals in A:

  $$\pi A = \prod_{\mathfrak{p}} \mathfrak{p}^{n(\mathfrak{p})}.$$

  This product is finite, i.e., almost all $n(\mathfrak{p})$ are 0.
- Now

  $$\mathcal{T} = \{A_{\mathfrak{p}} : \mathfrak{p} \text{ a non-zero prime ideal of } A \text{ (containing } \pi) \}.$$

2. Since R is a DVR we know that the unique non-zero prime ideal is principal, say πR. Every element $a \in R \setminus \{0\}$ can then be written uniquely as $a = u\pi^n$ where u is a unit in R and n is a non-negative integer. We write $\mathrm{ord}_R(a) = n$. We see that

$$\mathrm{ord}_R : K^*/R^* \longrightarrow \mathbb{Z}$$

is an isomorphism. The converse also holds: if R is a valuation ring inside its field of fractions K and K^*/R^* is infinite cyclic, then R is a DVR.

Let $T \in \mathcal{T}$. The inclusions $K \subset L$ and $R \subset T$ give us a map $K^*/R^* \longrightarrow L^*/T^*$. Suppose that $k, k' \in K^*$ with $kT^* = k'T^*$, then $k^{-1}k' \in T^*$ and since $k^{-1}k' \in K^*$ it follows that $k^{-1}k' \in R^*$, so the map is injective. This allows us to view K^*/R^* as a subgroup of L^*/T^*. As they are both isomorphic to $\mathbb{Z}$ it follows that the index is finite. By similar argument $[T^*K^* : L^*] = [K^*/R^* : L^*/T^*]$ and we conclude that $e(T/R)$ is finite. The group K^*/R^* is generated by π (a generator of the $\mathfrak{m}_R$). Its image in L^*/T^* is $\mathrm{ord}_T \pi = \mathrm{ord}_{\mathfrak{p}_T} \pi = n(\mathfrak{p}_T)$ in the notation of the previous part.

3. This part follows from the next one.

4. Note that there is a bijection between the prime ideals of A containing π and the prime ideals of $A/\pi A$. This ring is in fact an algebra over $R/\pi R$ (which is a field since $\pi R = \mathfrak{m}_R$). We claim that the vector space dimension of $A/\pi A$ over $R/\pi R$ is finite and $\leq [L : K]$.

   For the remaining part of the proof we make one assumption that is not true in general, but does always hold for the DVRs and fields that we are working with, namely that A is a finitely generated R module. From this it follows that

   $$A \underset{R-\text{mod}}{\cong} R/(a_1) \oplus \ldots \oplus R/(a_t),$$

   since R is a PID. Moreover, since A is a domain, it cannot have any zero divisors, so in fact all the a_i must be 0 and

   $$A \underset{R-\text{mod}}{\cong} R^t.$$

   Tensoring with K we see that $L = K \otimes_R A \cong K^t$ as K vector spaces, so $t = [L : K]$. Tensoring with $R/\pi R$ we get $A/\pi A \cong (R/\pi R)^t$. Using the decomposition of πR into primes of R and the Chinese Remainder Theorem we get

   $$A/\pi A = A/\prod_{\mathfrak{p}} \mathfrak{p}^{n(\mathfrak{p})} = \prod_{\mathfrak{p}} A/\mathfrak{p}^{n(\mathfrak{p})}.$$

   Comparing dimensions over $R/\pi R$ we obtain

   $$[L : K] = \sum_{\mathfrak{p}} \dim_{R/\pi R}(A/\mathfrak{p}^{n(\mathfrak{p})}).$$

   Using the filtration $A \supset \mathfrak{p} \supset \mathfrak{p}^2 \supset \ldots \supset \mathfrak{p}^{n(\mathfrak{p})}$ and the fact that for Dedekind domains $\mathfrak{p}^i/\mathfrak{p}^{i+1} \cong A/\mathfrak{p}$ we conclude

   $$[L : K] = \sum_{\mathfrak{p}} n(\mathfrak{p})[A/\mathfrak{p} : R/\pi R] = \sum_{T} e(T/R)f(T/R).$$

□

Remark 9.4. To understand all the fibres of our map, we also have to look at what happens above the generic point of $C_{k(t)}$. The only valuation ring that lies above the generic point of $C_{k(t)}$ (i.e. the valuation ring $k(t)$ of $k(t)$ over k) is the generic point of C_l (i.e. l as a valuation ring of l over k).

Proof. Suppose we have a valuation ring T such that $T \cap k(t) = k(t)$. We know that T is integrally closed and that it contains $k(t)$. So it must contain all elements of l that are integral over $k(t)$. Since this is a field, integral is the same as algebraic and since the extension $l/k(t)$ is finite, all elements of l are algebraic over $k(t)$, so $T = l$ as required. □

Corollary 9.5. The properties 1 and 2 from theorem 9.1 hold for all curves C_l.

Example 9.6. Let $l = k(s)$ with s transcendental over k. Consider the subfield $k(s) \supset k(t)$ with $t = s^2$. Now we know both $C_{k(s)}$ and $C_{k(t)}$ already, so at is interesting to look at which map we get from $C_{k(s)}$ to C_{k_t}. We look at the fibres, i.e. we look at all valuation rings T of $k(s)$ that lie above a given valuation ring R of $k(t)$.

- $R = R_t$. In this case we see that $e(T/R) \cdot \text{ord}_R(t) = \text{ord}_T(t) = 2\text{ord}_T(s)$, that is, $e(T/R) = 2$. It follows that T is unique and that $f(T/R) = 1$.

- $R = R_\infty$. Same as the previous case with $u = t^{-1}$.

- $R = R_f, f \neq t$. Write $k(\alpha) = k[t]/(f)$. If α is not a square in $k(\alpha)$ then there is a unique T that lies above R and we have $f(T/R) = 2$, $e(T/R) = 1$. If α is a square in $k(\alpha)$ then there are two T's, both with $e = f = 1$.

Next we turn our attention to the product formula that was announced for the projective line in theorem 9.1 but not proved there. We will prove it here for any projective curve.

Theorem 9.7. *Let k be a field and $l/k(t)$ a finite extension with $k(t)/k$ transcendental. For every $x \in L^*$ there are only finitely many closed points $P \in C_l$ such that $\mathrm{ord}_P x \neq 0$. The following formula holds:*

$$\sum_{P \text{ closed}} (\deg P)(\mathrm{ord}_P x) = 0.$$

Proof. Suppose that x is algebraic over k. Recall that

$$\begin{aligned} \bigcap_{\substack{R \text{ closed pt} \\ \text{of } l \text{ over } k}} R &= \{\text{integral closure of } k \text{ in } l\} \\ &= \{\text{algebraic closure of } k \text{ in } l\}, \end{aligned}$$

so $x \in R$ for all closed points R, i.e. $\mathrm{ord}_R x \geq 0$ for all R. Since x^{-1} is also algebraic over k it follows by the same argument that $\mathrm{ord}_R x^{-1} \geq 0$ for all R, that is $\mathrm{ord}_R x \leq 0$ for all R. We conclude that $\mathrm{ord}_R x = 0$ for all R and that the theorem holds in this case.

We are left with the case that x is transcendental over k. In this case we can view l as a finite extension of $k(x)$. We first prove that the set of closed points where $\mathrm{ord}_P x < 0$ is finite and that

$$\sum_{\substack{P \text{ closed} \\ \mathrm{ord}_P x < 0}} (\deg P)(\mathrm{ord}_P x) = -[l : k(x)].$$

Suppose T is a valuation ring of l. Then $x \notin T$ if and only if $R = T \cap k(x)$ is R_∞. So the number of such T is indeed finite and we have the formula

$$\sum_{T | R_\infty} e(T/R_\infty) f(T/R_\infty) = [l : k(x)].$$

Note that

$$\mathrm{ord}_T x = e(T/R_\infty) \mathrm{ord}_\infty x = -e(T/R_\infty)$$

and

$$f(T/R_\infty) = [T/\mathfrak{m}_T : R_\infty/\mathfrak{m}_\infty] = [T/\mathfrak{m}_T : k] = \deg T,$$

so the required formula follows.

For all $\alpha \in k$ we get from the above by replacing x with $\frac{1}{x-\alpha}$ that

$$\sum_{\substack{P \text{ closed} \\ \mathrm{ord}_P (x-\alpha) > 0}} (\deg P)(\mathrm{ord}_P x) = [l : k(x)].$$

Combining the original formula with the one for $\alpha = 0$ we conclude that the product formula also holds in this case. □

We conclude this section with a large example of how to apply the theory about closed points when we are given an explicit curve.

Example 9.8. Let k be a field of characteristic not equal to 2 and let $d \in k[t]$ a polynomial that is not divisible by the square of any irreducible polynomial in $k[t]$ and also not a square in k. Let $l = k(t)(\sqrt{d})$. We can think of the variety C_l as defined by the equation $u^2 = d(t)$ in the unknowns u and t.

We are going to look at the closed points of C_l. We will not consider the points that lie above R_∞ in $k(t)$, although our method can easily be extended to deal with this case as well. How do we determine the closed points that lie above $R = R_f$ with $f \in k[t]$ irreducible? Well, first we have to determine the integral closure of R in l. This integral closure is $A = R[\sqrt{d}]$. Since the degree of l over $k(t)$ is 2, we have

$$\sum_{T|R} e(T/R)f(T/R) = 2,$$

so either there are 2 T's each with $e = f = 1$, or there is one T with $e = 2$, or there is one T with $f = 2$. To determine which of these is actually the case for a given R, we have to look at $A/\pi A$ where π is a prime element of R, which we can take to be f. Then we have

$$A/\pi A \cong R[\sqrt{d}]/fR[\sqrt{d}] \cong (R/fR)[X]/(X^2 - d \mathrm{mod} f) \cong k(\alpha)[X]/(X^2 - d(\alpha)),$$

where α is a (formal) root of f. Now it is easy to see what happens:

- $d(\alpha) = 0$ if and only if f divides d and this happens if and only if $e = 2$.
- $d(\alpha)$ is not a square in $k(\alpha)$. This happens if and only if $f = 2$.
- $d(\alpha)$ is a square in $k(\alpha)$. In this case we have two T's both with $e = f = 1$.

If the field k is algebraically closed we see that $k(\alpha) = k$, since k doesn't have any algebraic extensions. So $f = 2$ does not happen. Out of the other two cases the first happens if and only if $f|d$, which is only finitely often. Anyway we see that picking a T over $R = R_{t-\alpha}$ amounts to choosing a $\beta \in k$ such that $\beta^2 = d(\alpha)$. So these closed points are in bijection with the solutions of $u^2 = d(t)$.

Finally we state what happens for $R = R_\infty$. Proving these statements is a nice exercise. We have $e = 2$ if and only if $\deg d$ is odd. For d's with even degree the leading coefficient determines what happens. If it is not a square we have $f = 2$, if it is a square we have two T's, each with $e = f = 1$.

10 Line bundles and divisors

In this section we introduce the concepts of divisors and line bundles on a curve. We shall see that they are essentially two ways to look at the same thing.

Definition 10.1. A *line bundle* $\mathcal{L}$ on $C = C_l$ is a vector $(L_P)_{P \in C}$, where

- $L = L_\xi$ (ξ is the standard notation for the generic point) is a one-dimensional vector space over l,
- $L_P \subset L$ is a free $\mathcal{O}_P$-module of rank 1 for all closed points P,

such that for all $x \in L$ we have that x generates almost all L_P. Equivalently we want that for all $x \in L$ the set $\{P : L_P = \mathcal{O}_P x\}$ is open in C. Note that we can replace the 'for all'-s by 'there exists'-s in the preceding conditions.

One can define morphisms between line bundles, making them into a category, but for our purposes it suffices to say when two line bundles are isomorphic.

Definition 10.2. Two line bundles $\mathcal{L} = (L_P)_P$ and $\mathcal{M} = (M_P)_P$ are isomorphic if and only if there is an l-linear map $\phi : L \to M$ such that $\phi(L_P) = M_P$ for all P.

Definition 10.3. The space of *global sections* of a line bundle $\mathcal{L}$ is defined as

$$\mathcal{L}(C) = \bigcap_P L_P.$$

It is a sub-k-vector space of L. We write $h^0(\mathcal{L})$ for its k-dimension.

We shall see shortly that h^0 of a vector bundle is always finite. The space of global sections is also called $H^0(X, \mathcal{L})$, which explains the somewhat odd notation h^0 for the dimension.

Examples 10.4.

1. Take $L_P = \mathcal{O}_P$ for all P (so $L = l$). Note that $1 \in L$ is a generator for all the L_P's, so they form a line bundle. We call this line bundle $\mathcal{O}$ of $\mathcal{O}_C$. We have

$$\mathcal{O}_C(C) = \bigcap_P \mathcal{O}_P = k',$$

where k' is the algebraic closure of k in l that we have encountered before. We see that $h^0(\mathcal{O}) = [k' : k]$, which will often be 1.

2. Our second example gives us a whole family of line bundles on the projective line. The motivation comes from viewing the function field $l = k(t) = k(\frac{x}{y})$ as a subfield of $k(x, y)$, the field of fractions of $k[x, y]$. This ring splits up into subspaces containing the homogeneous polynomials of every degree:

$$k[x, y] = \bigoplus_{m \geq 0} k[x, y]_m.$$

We can extend this definition by putting $k[x, y]_m = 0$ for $m < 0$.

Let $n \in \mathbb{Z}$ and take

$$L = \left\{ \frac{g}{h} \middle| \exists m \in \mathbb{Z} : h \in k[x, y]_m, h \neq 0, g \in k[x, y]_{m+n} \right\}.$$

For every closed point P we define a polynomial $F_P \in k[x, y]$ by

$$F_P = \begin{cases} f(\frac{x}{y}) y^{\deg(f)} & \text{if } \mathcal{O}_P = R_f \\ y & \text{if } \mathcal{O}_P = R_\infty. \end{cases}$$

Now we take $L_P = \left\{ \frac{g}{h} \middle| \text{ as before and also satisfying } h \notin (F_P) \right\}$. With these choices we get a line bundle $\mathcal{O}(n) = (L_P)_P$ that satisfies $\mathcal{O}(n)(C) = k[x, y]_n$. We see that $h^0(\mathcal{O}(n)) = \max\{0, n + 1\}$.

3. The third and final example concerns the *canonical line bundle* (or *canonical sheaf*) on a non-singular curve. It is denoted ω or $\omega_{C/k}$ and is defined by $L_P = \Omega_{\mathcal{O}_P/k}$ for all P.

For example, in the case our curve is the projective line we have $\Omega_{l/k} = l\ dt$ and if $P \neq \infty$ we have $\Omega_{\mathcal{O}_P/k} = \Omega_{k[t]_{(f)}/k} = \Omega_{k[t]/k} \otimes_{k[t]} k[t]_{(f)} = k[t]\ dt \otimes_{k[t]} k[t]_{(f)} = k[t]_{(f)}\ dt = \mathcal{O}_P\ dt$, by various generalities about tensor products and Kähler differentials. Finally, for $P = \infty$ we change variables to t^{-1}. Now by the previous case we can conclude that $\Omega_{\mathcal{O}_\infty/k} = \mathcal{O}_\infty\ d(t^{-1}) = t^{-2}\mathcal{O}_\infty\ dt$. We say that dt has a pole of order 2 at ∞.

Having identified these spaces we can now compute the global sections of $\omega(\mathbb{P}^1_k)$. It is the intersection of all $\Omega_{\mathcal{O}_P/k}$'s inside $\Omega_{l/k}$, that is

$$\left\{ f\ dt \;\middle|\; f \in l, f \in \left(\bigcap_{P \neq \infty} \mathcal{O}_P \right) = k[t], t^2 f \in \mathcal{O}_\infty \right\}.$$

In other words, any such f is a polynomial in t, while at the same time $t^2 f$ is a polynomial in t^{-1}. The only f for which this holds is 0.

Definition 10.5. The *genus* of a curve C/k is the k-dimension of the global sections of the canonical line bundle, that is,

$$g(C/k) = h^0(\omega_{C/k}(C)) \in \mathbb{Z}_{\geq 0}.$$

For example, we have just computed that the genus of the projective line is 0.

The set of all isomorphism classes of line bundles over a fixed curve C/k is denoted $\mathrm{Pic}(C/k)$ and called the Picard set. Next we focus on describing this set using *divisors*. We shall give $\mathrm{Pic}(C/k)$ a group structure by exhibiting an bijection to a set that already has a group structure. The map then becomes an isomorphism by construction. One can also define the group structure on $\mathrm{Pic}(C/k)$ intrinsically and later show that the bijection with the other group is in fact an isomorphism.

Definition 10.6. A *divisor* on a curve C over k is a finite formal sum of closed points of C, which is to say, an element of

$$\mathrm{Div}(C/k) = \bigoplus_{\substack{P \in C \\ \text{closed}}} \mathbb{Z}.$$

To each nonzero element f of the function field l we can assign a divisor

$$(f) = \sum_P \mathrm{ord}_P(f).$$

These are called the *principal divisors.*

To establish the connection with line bundles we begin by bringing to your attention a fact about discrete valuation rings: let $\mathcal{O}_P$ be a valuation ring and l its field of fractions, then all the sub-$\mathcal{O}_P$-modules of l are of the form $\mathfrak{m}_P^n$ with $n \in \mathbb{Z}$ (where $\mathfrak{m}_P^0 = \mathcal{O}_P$).

Given a divisor $D = (n(P))_{P \text{ closed}}$ we define a line bundle $\mathcal{L}(D)$ by putting $L = l$ and $L_P = \mathfrak{m}_P^{-n(P)}$ for all closed P. We write $L(D) = \mathcal{L}(D)(C)$ for the global sections and $l(D) = h^0(\mathcal{L}(D))$ for their dimension. Conversely, if we have a line bundle $\mathcal{L}$ with $L = l$ then all the L_P are of the form $\mathfrak{m}_P^{-n(P)}$ and since 1 is a generator of almost all the L_P we have $n(P) = 0$ almost everywhere. Hence $D = \sum_{P \text{ closed}} n(P)P$ is a divisor and we have $\mathcal{L}(D) = \mathcal{L}$. For an arbitrary line bundle $\mathcal{L}$ we can find a divisor D such that $\mathcal{L}(D)$ is isomorphic to $\mathcal{L}$ by simply chosing any isomorphism from L to l.

For the above we see that $D \mapsto \mathcal{L}(D)$ is a surjective map from $\mathrm{Div}(C/k)$ to $\mathrm{Pic}(C/k)$. The next step is to identify its kernel:

$$\begin{aligned} \Big(\mathcal{L}(D) = \sum n_P P\Big) \cong \Big(\mathcal{L}(E) = \sum m_P P\Big) &\iff \exists x \in l^* : \forall P \text{ closed} : x\mathfrak{m}_P^{-n_P} = \mathfrak{m}_P^{-m_P} \\ &\iff \exists x \in l^* : \forall P \text{ closed} : \mathrm{ord}_P(x) = n_P - m_P \\ &\iff D - E \text{ is a principal divisor.} \end{aligned}$$

We summarize the previous paragraphs in the following exact sequence of abelian groups:

$$0 \longrightarrow k'^* \longrightarrow l^* \longrightarrow \mathrm{Div}(C/k) \longrightarrow \mathrm{Pic}(C/K) \longrightarrow 0.$$

Note that on $\mathrm{Div}(C/k)$ we have the degree map

$$\begin{array}{cccc} \deg : & \mathrm{Div}(C/k) & \longrightarrow & \mathbb{Z} \\ & \cup & & \cup \\ & \sum n(P)P & \longmapsto & \sum n(P)\deg P. \end{array}$$

The product formula tells us that the image of l^* in $\mathrm{Div}(C/k)$ sits in the kernel of the degree map. As a result, this map factors via the quotient. So we get a degree map from $\mathrm{Pic}(C/k)$ to $\mathbb{Z}$

definied by $\deg(\mathcal{L}(D)) = \deg(D)$. To compute the degree of a line bundle $\mathcal{L}$, take any non-zero $x \in L$ and determine for each closed point P the integer $n(P)$ such that $\mathcal{O}_P x = \mathfrak{m}_P^{n(P)} L_P$; then $\deg \mathcal{L} = \sum n(P) \deg P$.

The exact sequence we have derived can be seen as a generalisation of remark 9.2 for the projective line. Indeed, looking carefully at the maps there we conclude that the degree map on $\mathrm{Pic}(\mathbb{P}^1_k/k)$ is an isomorphism. The line bundles $\mathcal{O}(n)$ we constructed earlier in this section have degree n, so up to isomorphism, these are all the line bundles on $\mathbb{P}^1_k$.

We conclude this section by showing that the space of global sections of a line bundle is always finite-dimensional.

Theorem 10.7. *Let C be a regular projective curve over k and let $\mathcal{L}$ be a line bundle on C. Then* $h^0(\mathcal{L}) \leq \max\{0, \deg \mathcal{L} + [k' : k]\}$.

Proof. If $\mathcal{L}(C) = \{0\}$ then the theorem certainly holds. Otherwise, let $x \in \mathcal{L}(C)$, $x \neq 0$. Now $L = lx$ and the map $l \to L$ sending f to fx is an isomorphism. Furthermore we have $x\mathcal{O}_P \subset L_P$, and $x\mathcal{O}_P = \mathfrak{m}_P^{n(P)} L_P$, hence $n(P) \geq 0$ for all points P. We conclude that we can write

$$\begin{array}{rcl} \mathcal{L}(C) & \cong & \left\{ f \in l \middle| fx \in L_P \text{ for all } P \right\} \\ & = & \left\{ f \in l \middle| \forall P : f \in \mathfrak{m}_P^{-n(P)} \right\} \\ & & \qquad \Big\downarrow \psi \\ & & \displaystyle\bigoplus_{P \text{ closed}} \mathfrak{m}_P^{-n(p)} / \mathcal{O}_P, \end{array}$$

where ψ sends f to $(f + \mathcal{O}_P)_P$. Since ψ is a linear map between k-vector spaces we know that $h^0(\mathcal{L}) = \dim_k \mathcal{L}(C) \leq \dim_k \ker \psi + \dim_k \mathrm{im}\ \psi$. We see that $\ker \psi = \bigcap\limits_P \mathcal{O}_P = k'$, so $\dim_k \ker \psi = [k' : k]$. For the image we use that is can be no larger than the space we are mapping to, so

$$\dim_k \mathrm{im}\ \psi \leq \dim_k \left(\bigoplus_P \mathfrak{m}_P^{-n(P)} / \mathcal{O}_P \right) = \sum_P \dim_k \left(\mathfrak{m}_P^{-n(P)} / \mathcal{O}_P \right) = \sum_P n(P) \deg\ P = \deg\ \mathcal{L}.$$

Putting all this together we get the desired inequality. □

The following picture shows which combinations $(h^0\ \mathcal{L}, \deg\ \mathcal{L})$ fall between the upper bound indicated in the previous theorem and the trivial lower bound $h^0(\mathcal{L}) \geq 0$.

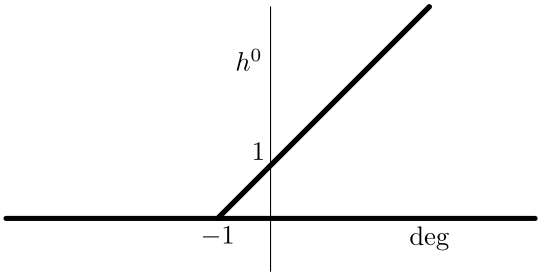

Example 10.8. In this example we shall consider a family of line bundles on certain curves and investigate where each of these line bundles fits in the previous picture. It provides some motivation for the Riemann-Roch theorem that we shall discuss next.

Let k be a field of characteristic different from 2 and consider the curve given by $u^2 = d(t)$ where u and t are unknowns and $d(t) \in k[t]$ a squarefree polynomial of degree $2g+1$ with $g \in \mathbb{Z}_{\geq 0}$. This curve therefor has function field $k(t)(\sqrt{d(t)})$, the field of fractions of $k[t,u]/(u^2 - d(t))$. This is a degree 2 extension of $k(t)$.

First we show that there is exactly one point lying over ∞. Suppose we have such a point Q. We have that $\mathrm{ord}_Q(d(t)) = \mathrm{ord}_\infty(d(t)) e(Q/\infty) = -(2g+1)e(Q/\infty)$ and also $\mathrm{ord}_Q(d(t)) =$

$\mathrm{ord}_Q(u^2) = 2\mathrm{ord}_Q(u)$. The only way this can happen is if $e(Q/\infty)$ is even. Since it is also either 1 or 2, we conclude that $e(Q/\infty) = 2$ and so we must have $f(Q/\infty) = 1$ and there can only be one such Q. The degree of Q is $f(Q/\infty)\deg(\infty) = 1$.

We now look at the divisors nQ with $n \in \mathbb{Z}$. Note that $\deg(nQ) = n$, so that for negative n we have $L(nQ) = \{0\}$ from the previous theorem. For $n \geq 0$ we get from the definitions that

$$L(nQ) = \{f \in l : \mathrm{ord}_Q(f) \geq -n \text{ and } \mathrm{ord}_P(f) \geq 0 \text{ for all closed } P \neq Q\}.$$

This second condition is equivalent to $f \in \bigcap_{P \neq Q} \mathcal{O}_P$ and, since $P \neq Q \iff t \in \mathcal{O}_P$, this is just the integral closure of $k[t]$ in l. This can be determined using methods similar to those used for number fields and because of our assumptions on $d(t)$ one finds that it is equal to $k[t, \sqrt{d(t)} = u]$.

We observe that $L(nQ) \subset L(mQ)$ if $m \leq n$ and that

$$k[t,u] = \bigcup_{n \geq 0} L(nQ).$$

A basis of $k[t,u]$ is given by $\{t^i : i \in \mathbb{Z}_{\geq 0}\} \cup \{ut^i : i \in \mathbb{Z}_{\geq 0}\}$. We compute the order at Q for each of these basis vectors:

$$\begin{array}{rcl} \mathrm{ord}_Q\ t^i & = & -2i \\ \mathrm{ord}_Q\ ut^i & = & -(2g+1) - 2i \end{array}$$

We see that all these numbers are distinct and that we get almost all negative integers this way, except for the first g odd ones. Let $f \in k[t,u]$ and write it as $f = \sum_{i \in I} a_i t^i + \sum_{j \in J} b_j u t^j$. Now we have $f \in L(nQ)$ if and only if all the t^i's and ut^j's are in $L(nQ)$ and using the orders we've computed we see that this happens precisely if $n \geq \max_{i \in I, j \in J}\{2i, 2j+2g+2\}$. We concude that a k-basis for $L(nQ)$ is given by $\{t^i : 0 \leq 2i \leq n\} \cup \{ut^j : j \geq 0,\ 2j + 2g + 1 \leq n\}$.

The following picture shows the $(h^0, \deg)$ pairs that we find for $g = 4$.

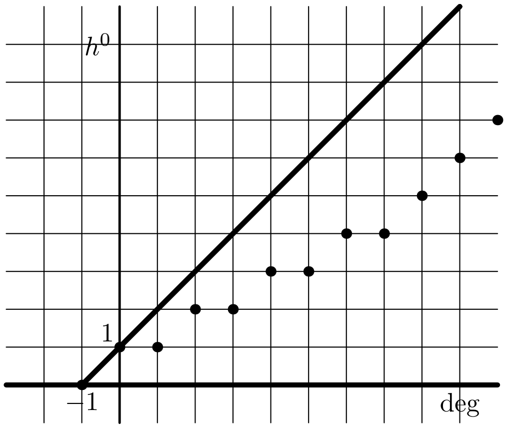

Theorem 10.9 (Riemann-Roch). *Let k be a field and C a regular projective curve with $k' = k$. Then there is $a \geq 0$ and there is a line bundle ω on C such that for all line bundles $\mathcal{L}$ we have*

$$h^0(\mathcal{L}) - h^0(\omega\mathcal{L}^{-1}) = \deg\ \mathcal{L} + 1 - g,$$

or equivalently, such that for all divisors D we have

$$l(D) - l(K - D) = \deg\ D + 1 - g,$$

where K is a divisor such that $\mathcal{L}(K) = \omega$.

If C is non-singular then ω is the canonical line bundle and g is the genus of C.

Corollary 10.10.

- If $D = 0 \in \text{Div } C$ then $\mathcal{L}(D) = \mathcal{O}_C$, so $l(D) = \dim_k \mathcal{O}_C(C) = [k' : k] = 1$. From the theorem we get $1 - l(K) = 0 + 1 - g$, or $g = l(K) = h^0(\omega)$.
- If $D = K$ then the formula gives $g - 1 = \deg\ K + 1 - g$, so $\deg\ K = 2g - 2$.
- In general, if we have $D, D' \in \text{Div } C$ that add up to K then the points $p_D = (\deg\ D, l(D))$ and $p_{D'} = (\deg\ D', l(D'))$ have the same distance to the line $\deg = g - 1$ and they are on the same line with slope $\frac{1}{2}$.

Using this last symmetry property, we can make the bounds we computed before a lot stronger, in fact for sufficiently large degree, h^0 is again a function of the degree. In a picture we get

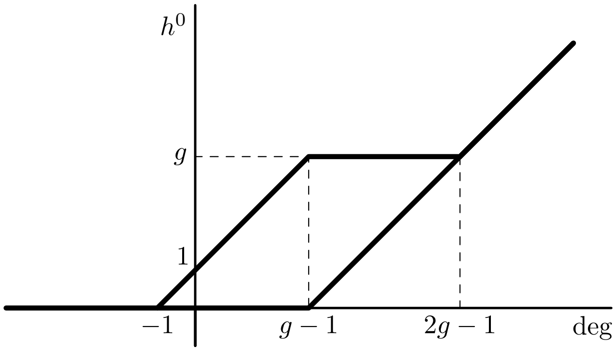

Corollary 10.11. Suppose we have a curve $C = C_l$ of genus 0. The for all D we have $l(D) = \max\{0, 1 + \deg\ D\}$. The degree of the divisor K is $2g - 2 = -2$, so we have $2\mathbb{Z}$ in the image of deg. There are now two possibilities. Either this is the whole image, or deg is surjective. The latter will happen as soon as Div C contains any element of odd degree, for example when $k = \overline{k}$.

Suppose deg is surjective and let D be a divisor of degree 1. We claim that $C = \mathbb{P}^1_k$. From the Riemann-Roch formula we have that $l(D) = 2$, so we can pick and $f \in L(D)$ that is non-zero. Now we have have $(f) + D \geq 0$ (by definition of $L(D)$) and $\deg((f) + D) = 1$ (since the prinicipal divisors are in the kernel of deg). We conclude that $(f) + D = \sum n(P)P$ with $n(P) \geq 0$ and $\sum n(P) \deg(P) = \deg((f) + D) = 1$, so this can only happen if there is a point P of degree 1 such that $(f) + D = P$.

Now we again have $l(P) = 2$ from the Riemann-Roch theorem, so that $L(P)$ properly contains k. This means we can pick a $t \in L(P), t \notin k$. We have that $\text{ord}_P(t) = -1$ and for all $Q \neq P$ we have $\text{ord}_Q(t) \geq 0$. Now we recall that

$$[l : k(t)] = \sum_{\substack{Q \text{ closed} \\ \text{ord}_Q(t) < 0}} -\text{ord}_Q(t) \deg(Q) = -(-1 \cdot 1) = 1,$$

that is, $l = k(t)$ as required.

11 Elliptic Curves

Definition 11.1. An *elliptic curve* over k is a pair $E = (C, O)$ where C is a non-singular projective curve over k of genus 1 and O is a rational point of C, i.e. a closed point P of degree 1.

Suppose we have an elliptic curve E. The Riemann-Roch theorem tells us that for any divisor D on E (by abuse of notation we also say E when we just mean the curve, not the curve-and-point pair) with positive degree we have $l(D) = \deg(D)$. We see that $l(1 \cdot O) = 1$ and $k = L(0 \cdot O) \subset L(1 \cdot O)$, so $L(1 \cdot O) = k$. For every $n > 0$ we have $l((n+1)O) = n+1 = l(nO)+1$ and $L(nO) \subset L((n+1)O)$, so that the quotient is generated as a k-vector space by a single element (any element that is in $L((n+1)O)$ but not in $L(nO)$ will do the trick).

Pick $x \in L(2O), x \notin L(1O)$ and $y \in L(3O), y \notin L(2O)$. Now we have $x^2 \in L(4O), x^2 \notin L(3O)$ and $xy \in L(5O), xy \notin L(4O)$, so $L(5O)$ has a k-basis $\{1, x, y, x^2, xy\}$. Note that both x^3 and y^2 are in $L(6O)$ but not in $L(5O)$, so there must be a unique $\lambda \in k^*$ such that $\lambda x^3 - y^2 \in L(5O)$. Multiplying x and y by λ we see that we can in fact assume $\lambda = 1$, so that $x^3 - y^2 \in L(5O)$, which means we can write it on the basis we have found. This gives an equation in Weierstrass form for our elliptic curve:

$$y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6.$$

Next we will endow elliptic curves with a group structure. Let $\mathrm{Pic}^0(C)$ the kernel of $\deg : \mathrm{Pic}(C) \to \mathbb{Z}$. We will show that the map

$$\begin{array}{ccc} E(k) & \longrightarrow & \mathrm{Pic}^0(C) \\ \cup & & \cup \\ P & \longmapsto & [P - O] \end{array}$$

is a bijection. For $D \in \mathrm{Pic}^0$ we have to show that the set $\{P \in E(k) : [P] = [D + O]\}$ has one element. The condition $[P] = [D + O]$ means that $P = (f) + (D + O)$ for some $f \in l^*$. Since $(f) + D + O$ has degree 1 we see that this divisor is equal to some point P if and only if it is effective. Two f's give rise to the same P if and only if they differ by an element of k^*, so we get

$$\{P \in E(k) : [P] = [D+O]\} \;\cong\; \{f \in l^* : (f) \geq -(D+O)\}/k^* \;=\; (L(D+O) - \{0\})/k^*.$$

From the Riemann-Roch theorem we get that $\dim_k(L(D + O)) = 1$, so we see that this is indeed a set with only one element. Now we get an abelian group structure on $E(k)$ by transport of structure. One tends to write the group operation on an elliptic curve additively. We note that O becomes the zero element of this group. In a similar way one can define $E(L)$ for any field extension L of k, which is also an abelian group with O as its zero element. (One way to go about this is by defining the curve over L using the same equation as one had over k.)

Definition 11.2. A *morphism* between two elliptic curved (E, O) and (E', O') is a morphism of curves from E to E' that sends O to O'.

Suppose $\phi : E \to E'$ is such a morphism. Recall that this gives us a place from l' to l, where l and l' are the function fields of the curves E and E'. So we have a valutation ring $\mathcal{O}$ of l' and a ring homomorphism from $\mathcal{O}$ to l. Since O' is in the image of ϕ, we must have $\mathcal{O}_{O'} \subset \mathcal{O}$, hence either $\mathcal{O} = \mathcal{O}_{O'}$ or $\mathcal{O} = l'$. In case $\mathcal{O} = \mathcal{O}_{O'}$ the only point in the image of ϕ is O' and we get the *zero morphism* sending all of E to $O' \in E'$. If we have $\mathcal{O} = l'$ then the place is in fact a field embedding of l' in l and the condition on the zero elements translates to $\mathcal{O}_O \cap l' = \mathcal{O}_{O'}$. Such maps are called *isogenies*. The curves E and E' are called *isogenous* if there exists an isogeny between them.

Definition 11.3. The *degree* of a morphism $\phi \in \mathrm{Hom}(E, E')$ is defined for an isogeny as $[l : \phi^*l']$, where $\phi^* : l' \to l$ is the inclusion. The degree of the zero morphism is 0.

The degree of an isogeny can be decomposed into a separable and an inseparable part: let l_s the largest separable extension of ϕ^*l' inside l (this means that $l_s = \{\alpha \in l : \alpha \text{ separable over } \phi^*l'\}$) then $\deg_{\mathrm{sep}}(\phi) = [l_s : \phi^*l']$ and $\deg_{\mathrm{ins}}(\phi) = [l : l_s]$. Since the field degree is multiplicative in towers, we have $\deg(\phi) = \deg_{\mathrm{sep}}(\phi)\deg_{\mathrm{ins}}(\phi)$.

Theorem 11.4. *Let $\phi : E \to E'$ be a morphism of elliptic curves and let L be any extension of k. Then*

- *$\phi(L) : E(L) \to E'(L)$ is a group homomorphism.*
- *It has a finite kernel and $\#(\ker \phi(L)) \mid \deg_{\text{sep}} \phi$.*
- *If L is algebraically closed then $\phi(L)$ is surjective and we have equality in the previous statement.*

The set of morphisms $\mathrm{Hom}(E, E')$ *has an additive group structure such that for all field extensions L of k we have*

$$(\phi + \psi)(L)(P) = \phi(L)(P) + \psi(L)(P).$$

The canonical line bundle ω_E is a k vector space of dimension $g_E = 1$. We write $T_O(E)$ for its dual, i.e., $T_O(E) = \mathrm{Hom}_k(\omega_E(E), k)$. This is a one dimensional k vector space. A morphism $\phi : E \to E'$ induces a morphism $T_O(\phi) : T_O(E) \to T_{O'}(E')$. We get a map

$$\begin{tikzcd}
\mathrm{Hom}(E, E') \arrow[r] & \mathrm{Hom}_k(T_O(E), T_{O'}(E)) & \cong k \text{ (not canonical)} \\
\phi \arrow[u, phantom, sloped, "\in"] \arrow[r, mapsto] & T_0(\phi). \arrow[u, phantom, sloped, "\in"] &
\end{tikzcd}$$

This map is an additive group homomorphism. Its kernel consists of the *inseparable* morphisms. A morphism is called inseparable if it is either the zero morphism or has $\deg_{\text{ins}} > 1$. In particular, if char $k = 0$ the map from $\mathrm{Hom}(E, E')$ to k is injective, so this group is torsion-free.

The composition map on $\mathrm{End}(E) = \mathrm{Hom}(E, E)$ is bilinear, so that $\mathrm{End}(E)$ is in fact a ring. Since the composition of two isogenies is again an isogeny, this ring does not have zero divisors. In the case that char $k = 0$ we get from the inclusion $\mathrm{End}(E) \to k$ we just constructed that it is a commutative ring.

We conclude this section with two general theorems on the structure of elliptic curves, first for $k = \mathbb{C}$ and then for general k.

Theorem 11.5. *The category of elliptic curves over $\mathbb{C}$ is equivalent to the category of* lattices *in $\mathbb{C}$. The objects of this category are the lattices in $\mathbb{C}$, i.e., discrete co-compact subgroups and for L, M two lattices the set of morphisms $\mathbb{C}(L, M)$ is the set $\{z \in \mathbb{C} : zL \subset M\}$.*

If E corresponds to L under this equivalence, then there are isomorphisms $E(\mathbb{C}) \cong \mathbb{C}/L$ and $T_O(E) \cong \mathbb{C}$. These isomorphisms are such that if E' is another elliptic curve, corresponding to L', then the composition of the isomorphisms $C(L, L') \to \mathrm{Hom}(T_O(E), T_{O'}(E')) \to \mathbb{C}$ is the natural inclusion map and such that if $\phi : E \to E'$ is the morphism of elliptic curves corresponding to $\alpha \in \mathbb{C}(L, L')$ then the diagram

$$\begin{tikzcd}
E(\mathbb{C}) \arrow[r, "\phi(\mathbb{C})"] \arrow[d] & E'(\mathbb{C}) \arrow[d] \\
\mathbb{C}/L \arrow[r, "\alpha"] & \mathbb{C}/L'
\end{tikzcd}$$

commutes.

Corollary 11.6. For all positive integers n we have $E(\mathbb{C})[n]$, the n-torsion of $E(\mathbb{C})$ is isomorphic to $\mathbb{Z}/n\mathbb{Z} \oplus \mathbb{Z}/n\mathbb{Z}$. So the map $n : E \to E$ (multiplication by n) has degree n^2.

Proof. For a suitable choice of bases the composition of the maps $\mathbb{R} \oplus \mathbb{R} \cong \mathbb{C} \supset L \cong \mathbb{Z} \oplus \mathbb{Z}$ is the standard inclusion of $\mathbb{Z} \oplus \mathbb{Z}$ in $\mathbb{R} \oplus \mathbb{R}$. The means that $E(\mathbb{C}) \cong \mathbb{C}/L \cong (\mathbb{R}/\mathbb{Z}) \oplus (\mathbb{R}/\mathbb{Z})$ and so $E(\mathbb{C})[n] \cong \mathbb{Z}/n\mathbb{Z} \oplus \mathbb{Z}/n\mathbb{Z}$. □

Theorem 11.7. *Let k be a field and E, E' elliptic curves over k. Then the group* $\mathrm{Hom}(E, E')$ *is torsion free and there is a contra-variant function* $\overline{\cdot}$ *from the category of elliptic curves over k to itself such that*

- $\overline{\overline{\cdot}}$ *is the identity functor.*
- $\overline{E} = E$ *for every E.*
- $\overline{\cdot} : \mathrm{Hom}(E, E') \to \mathrm{Hom}(E', E)$ *is a group isomorphism*
- *for all morphisms $\phi : E \to E'$ we have* $\overline{\phi}\phi = \deg(\phi) \in \mathbb{Z} \subset \mathrm{End}(E)$.

Corollary 11.8. The map $\overline{\cdot} : \mathrm{End}(E) \to \mathrm{End}(E)$ is an involution, i.e., an anti ring homomorphism whose square is the identity. We have $\overline{1} = 1$ and so $\deg(n) = \overline{n}n = n^2$. If $k = \overline{k}$ and n is not divisible by $\mathrm{char}\, k$ then $E(k)[n]$ is a group of order n^2: $\deg_{\mathrm{ins}} n$ divides $\mathrm{char}\, k$ since the degree of an inseparable extension always divides the characteristic and $\deg_{\mathrm{ins}} n$ divides $\deg n = n^2$, so it must be 1 since these two numbers are coprime, this means that $\#\ker n = \deg_{\mathrm{sep}} n = \deg n = n^2$. It even holds that $E(k)[n] \cong \mathbb{Z}/n\mathbb{Z} \oplus \mathbb{Z}/n\mathbb{Z}$.