28 Geometry of Function Fields

Last time, we began discussing finite extensions of the function field $F = \mathbb{C}(t)$; we can write such an extension as $E = F[x]/(P)$ for an irreducible $P \in F[x]$. Since the coefficients of P are rational coefficients, we can clear denominators and then think of P as a polynomial in two variables — and by factoring out the gcd of the coefficients, we can assume that $P \in \mathbb{C}[t, x]$ is primitive, and therefore irreducible in $\mathbb{C}[t, x]$ as well. Then we can consider the ring $R = \mathbb{C}[t, x]/(P)$, so then $E = \text{Frac}(R)$.

A geometric way to think about this situation is that we have

$$X = \text{MSpec}(R) = \{(a, b) \in \mathbb{C}^2 \mid P(a, b) = 0\}$$

(so X consists of the set of zeros of the polynomial), and this maps to

$$Y = \text{MSpec}(\mathbb{C}[t]) = \mathbb{C}.$$

Here X is also called a **Riemann surface**, and we can think of E as a field of rational functions on that Riemann surface.

28.1 Ramified Covers

Recall the simpler example discussed last class:

Example 28.1
Consider $P(x, t) = x^n - t$.

We saw earlier that this corresponds to the following picture. Here the t-plane represents Y, and the x-plane represents X — by definition X corresponds to $(t, x) \in \mathbb{C}^2$ such that $t = x^n$, but such points can just be described by x.

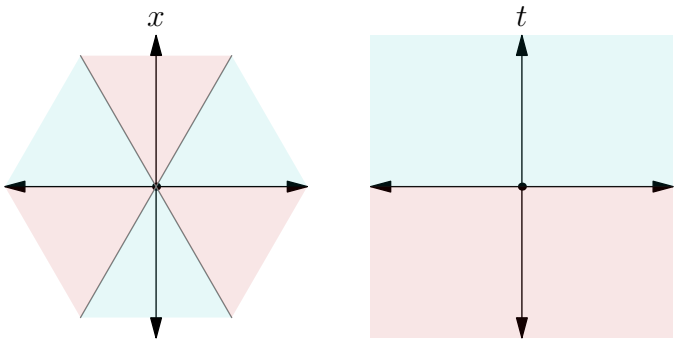

These two maps are related by the map $x \mapsto x^n$ (since when we go from X to Y, we're mapping each point to its corresponding value of t, which here is x^n), which unwraps each smaller angle to a 180° angle.

As mentioned earlier, this is a ramified cover — all but finitely many values of t have the same number of points in their pre-image (and in general, this number of points is equal to the degree of P, as a polynomial in x). But some points give smaller fibers — for the map $x \mapsto x^n$, the point 0 only has one element in its pre-image, and is therefore a ramification point.

Incidentally, the terminology we saw for the behavior of primes, regarding how they split in quadratic extensions, also included *ramified prime*. This isn't a coincidence — they describe the same phenomenon.

Returning to the general situation, let $a_1, \ldots, a_k$ be the ramification points, meaning that $|f^{-1}(a_i)| < n$ for all i, and $|f^{-1}(a)| = n$ for all a not equal to any of the a_i. Now take a generic point a (not equal to any a_i). Then there are n points lying over a (meaning points whose image is a):

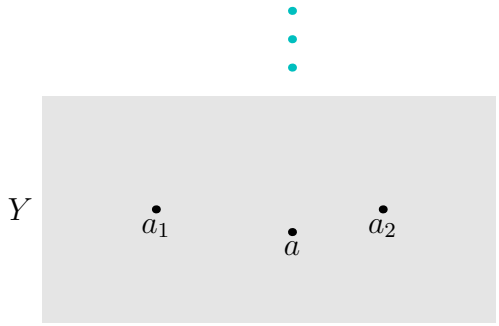

Now consider a closed loop around a, which doesn't pass through any of the ramification points.

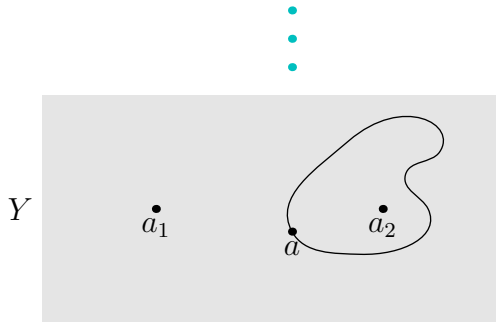

This loop determines a permutation of the fibers (meaning the n points above a): suppose we start at one of these n points. Then we can lift the loop locally in a unique way (since the loop avoids the ramification points). But when we return, we may return to a different one of these points:

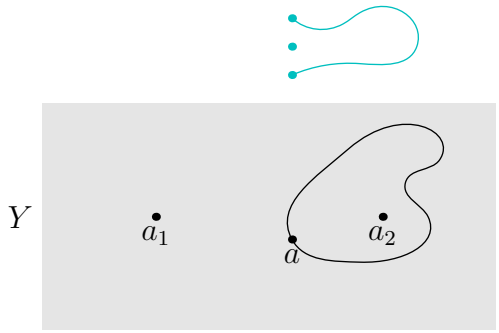

More precisely, if we have a loop $\gamma : [0,1] \to Y$ avoiding the ramification points, with $\gamma(0) = \gamma(1) = a$, then γ defines a permutation σ_γ of the set $f^{-1}(a)$ — to compute $\sigma_\gamma(x)$, we lift γ to a map $\widetilde{\gamma} : [0,1] \to X$ such that $\gamma(0) = x$; then $\gamma(1)$ is another point $y \in f^{-1}(a)$, and we set $\sigma_\gamma(x) = y$.

Example 28.2
Consider the example $P(x,t) = x^n - t$, and let γ be the unit circle (the standard loop).

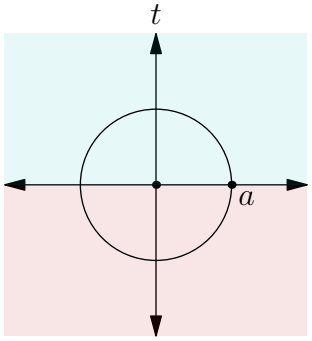

When we walk on the x-plane, we're still walking in a circle, but we walk n times slower (since x is raised to the nth power):

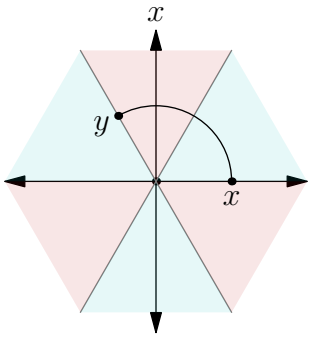

Explicitly, we have $X = \mathbb{C}$ and $Y = \mathbb{C}$, with $f(x) = x^n$. We have that $f^{-1}(1) = \{\exp(2\pi ik/n)\} = \{\zeta_n^k\}$ is the set of nth roots of unity (where $\zeta_n = \exp(2\pi i/n)$). Our loop is defined by $\gamma(t) = \exp(2\pi it)$, and if we start at ζ_n^k, then $\widetilde{\gamma}(t) = \zeta_n^k \exp(2\pi it/n)$. So the permutation is $\sigma_\gamma(\zeta_n^k) = \zeta_n^{k+1}$.

A term for this permutation is the **monodromy**.

A useful fact, which we will not prove, is the following:

Theorem 28.3
If E/F is a splitting field of some polynomial, then σ_γ extends to an automorphism of X which is the identity on Y, coming from an automorphism of E which is the identity on F.

The point is that σ_γ always gives us a permutation of the points in the pre-image of a fixed point a; but in the case of a splitting field, this can be *extended* to an automorphism of all of X.

Example 28.4
In our example situation, where $F = \mathbb{C}(t)$ and $E = \mathbb{C}(t)[x]/(x^n - t)$ (which is a splitting field), then the automorphism corresponding to the unit circle is $x \mapsto \zeta_n x$, which sends $t \mapsto t$.

Example 28.5
In the other example from last class, where $E = \mathbb{C}(t)[x]/(x^2 - t(t+1)(t-\lambda))$, there are three ramification points (0, -1, and λ), and we have a double cover (there's two points above all points except the ramification points). The automorphism is $x \mapsto -x$, which swaps the two points (and again fixes t).

Later, we'll discuss more algebraic ways to construct automorphisms of fields.

Student Question. *Are there always ramification points?*

Answer. *If $Y = \mathbb{C}$, then the answer is yes — this follows from topological reasons, as $\mathbb{C}$ is simply connected. But if Y is some other Riemann surface, there may be no ramification points.*

Note that if γ is deformed continuously, while still avoiding the ramification points and fixing the beginning and end, then the permutation described doesn't change. In our example $P(x,t) = x^n - t$, any closed loop which goes around 0 once will give the same permutation:

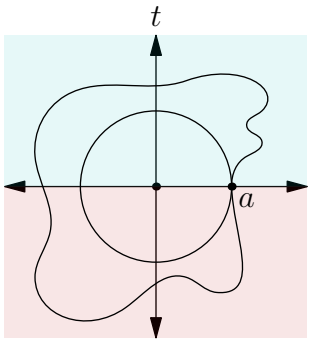

28.2 The Main Theorem of Algebra

There is a proof of the main theorem of algebra using ideas similar to the ones we've seen here.

Theorem 28.6
The field $\mathbb{C}$ is algebraically closed — in other words, every nonconstant polynomial $P \in \mathbb{C}[x]$ has a root.

The proof uses the concept of a **winding number**: suppose we have a continuous map $\gamma : [0,1] \to \mathbb{C} \setminus \{0\}$. Then its winding number, informally, is the number of times γ goes around 0 (counted with sign — going around 0 counterclockwise is counted with positive sign, and clockwise with negative sign). It can actually be formally defined using similar ideas to the ones we've seen here — consider σ_γ for the exponential map where $X = \mathbb{C}$ and $Y = \mathbb{C} \setminus \{0\}$, and we send $x \mapsto \exp(x)$. Then the pre-image of $z \in Y$ is $\exp^{-1}(z) = \{\log z + 2\pi i n\}$ for integers n. So for a loop γ, if we construct the loop $\widetilde{\gamma} : [0,1] \to \mathbb{C}$ as before, so that $\exp(\widetilde{\gamma}(t)) = \gamma(t)$, then we have

$$\gamma(1) = \gamma(0) + 2\pi i n$$

for some integer n. The winding number is defined to be the integer n in this equation.

We use $w(\gamma)$ to denote the winding number of γ.

Lemma 28.7
If $\gamma(t) = \gamma_1(t)\gamma_2(t)$, then $w(\gamma) = w(\gamma_1) + w(\gamma_2)$.

Proof. We have $\widetilde{\gamma}(t) = \widetilde{\gamma_1}(t) + \widetilde{\gamma_2}(t)$, so the discrepancies $2\pi i n$ are added together as well. □

Now we are ready to prove the theorem.

Proof of Theorem 28.6. Consider a polynomial $P(z) \in \mathbb{C}[z]$, and assume for contradiction that $P(z) \neq 0$ for all $z \in \mathbb{C}$. Let

$$P(z) = z^n + a_{n-1}z^{n-1} + \cdots + a_0.$$

Now consider loops

$$\gamma_r(t) = P(re^{2\pi i t})$$

for each $r \geq 0$ (where we take the circle of radius r around the origin, and see what P does to it). Since P has no zeros, this is a loop in $\mathbb{C} \setminus \{0\}$.

Now consider the winding number of each loop γ_r. We can observe three properties, which together lead to a contradiction: first, $w(\gamma_r)$ is independent of r — this is clear because it depends continuously on r (since none of our loops pass through 0).

Second, when $r = 0$, $w(\gamma_0) = 0$ — this is because γ_0 is the constant loop.

Finally, when r is large, we claim $w(\gamma_r) = n$. To prove this, we can write

$$P(z) = z^n \left(1 + \frac{a_{n-1}}{z} + \cdots + \frac{a_0}{z^n}\right).$$

Let z^n correspond to γ_1, and the remaining factor to γ_2. Then z^n runs around a circle n times, so $w(\gamma_1) = n$. Meanwhile, if r is very large, then

$$\left|\frac{a_{n-1}}{z} + \cdots + \frac{a_0}{z^n}\right| < 1,$$

which means γ_2 is trapped inside the circle of radius 1 centered at 1:

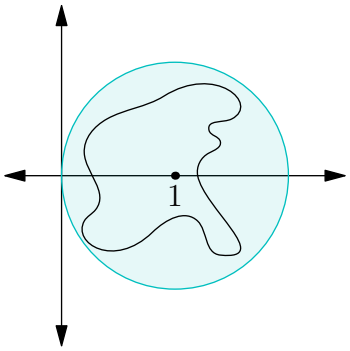

This means it's trapped to the right of the y-axis, so it can't wind at all; so $w(\gamma) = w(\gamma_1) + 0 = n$. □

Note 28.8
The textbook doesn't split $\gamma = \gamma_1\gamma_2$, but it has a nice intuitive explanation — essentially, the loop γ_r is fairly close to the loop which goes around the circle n times. Imagine having a dog on a leash, and walking around a circle n times. As long as the leash is short enough, the dog may run around you in any way it wants, but it will still go around the center of the circle the same number of times that you do. (The terms after 1 in the second factor correspond to the additional movement of the dog.)

28.3 The Primitive Element Theorem

Next class, we will begin discussing Galois theory. The following theorem will be useful:

Theorem 28.9
If E/F is a finite separable extension, then the extension is generated by one element — meaning $E = F(\alpha)$ for some α.

So even if E was defined by adjoining multiple elements (for example, the splitting field construction), it's possible to obtain it just by adjoining one element.

Recall that all finite fields and fields of characteristic 0 are separable; these are the only cases we will work with.

Proof. If F is finite, then E is also finite; so $E = F(\alpha)$ where α is a multiplicative generator of E.

Now assume F is infinite. It's enough to prove this in the case where $E = F(\alpha, \beta)$ is generated by two elements (since we must have $E = F(\alpha_1, \dots, \alpha_n)$ for some finite n, and we can use induction on n).

Define $\gamma_t = \alpha t + \beta$. Then we'll show that for all but finitely many t, γ_t is a generator of E.

Let P be the minimal polynomial of α and Q the minimal polynomial of β, and let $K \supset E$ be a field where both P and Q split completely. Then we can write

$$P(x) = (x - \alpha_1)(x - \alpha_2)\cdots(x - \alpha_m)$$

and

$$Q(x) = (x - \beta_1)(x - \beta_2)\cdots(x - \beta_n),$$

where the α_i are all distinct, and the β_i are all distinct (by separability). Assume $\alpha = \alpha_1$ and $\beta = \beta_1$.

The condition on t we will specify (for $\alpha t + \beta$ to be a generator) is that the nm elements $\alpha_i t + \beta_j$ are all distinct. There are clearly finitely many t for which this isn't true.

It suffices to check that α and β are in $E' = F(\gamma)$ (where $\gamma = \gamma_t$ for some such t). We'll look at what polynomial equations we can write for α over E'. One is obvious — α is a root of $P(x)$. But α is also a root of the polynomial $Q(\gamma - tx)$, which we'll denote by $Q_1(x)$ — this is a polynomial with coefficients in E', and when we plug in α we get $Q(\beta) = 0$.

So then α is a root of the polynomial $S(x)$ which generates the ideal (P, Q) (also known as $\gcd(P, Q)$), working in $E'[x]$. But this gcd doesn't depend on the field in which it's computed — so $S(x)$ is also a generator of (P, Q)

in $K[x]$. And in $K[x]$, both polynomials split completely; and they have a unique common linear factor, namely $x - \alpha$ (by the condition on t).

Then $S(x)$ is a constant times $x - \alpha$; this means $\alpha \in E'$, and then $\beta \in E'$ as well. This means $E' = E$. $\square$

Resource: Algebra II Student Notes

Spring 2022

Instructor: Roman Bezrukavnikov
Notes taken by Sanjana Das and Jakin Ng