GALOIS THEORY: LECTURE 24

LEO GOLDMAKHER

1. Constructing finite fields

Although most of the semester we stated and proved theorems about general field extensions L/K, in practice we've barely touched the theory of finite fields. Our primary goal today is to develop this theory, and to explore the Galois theory of finite fields. (We'll finish by discussing some cool additional topics related to things we've explored during the course.) We begin with a motivating question.

Question 1. *Does there exist a field $\mathbb{F}$ with precisely six elements?*

Will suggested $\mathbb{F} := \{0,1\} \times \{0, i, 2i\}$, under some appropriate operation. After playing around with a few possible operations, however, we were unable to make this into a field. So what would a field with 6 elements look like?

To get a better sense of this, we reviewed some true facts about finite fields. First recall that if $\mathbb{F}$ is a finite field, then it must have characteristic p for some prime p; furthermore, $\mathbb{F}$ must contain $\mathbb{F}_p$ as a subfield (see problems **4.1** and **4.2**). In short, $\mathbb{F}/\mathbb{F}_p$. This fact already imposes a strong restriction on $\mathbb{F}$:

Proposition 1.1. *Given any finite field $\mathbb{F}$ with* $\operatorname{char} \mathbb{F} = p$. *Then $|\mathbb{F}| = p^n$ for some $n \in \mathbb{N}$.*

Proof. Let $\mathcal{B}$ be a basis of $\mathbb{F}$ over $\mathbb{F}_p$. Since every element of $\mathbb{F}$ can be expressed in a unique way as a linear combination of the elements of $\mathcal{B}$, we deduce that $|\mathbb{F}| = |\mathbb{F}_p|^{|\mathcal{B}|} = p^{|\mathcal{B}|}$. Since $\mathbb{F}$ is finite, $\mathcal{B}$ must be finite, and we conclude. □

This explains why we struggled to find a field with six elements – there aren't any! By contrast, there could be a field with four elements. Notice that Proposition 1.1 doesn't guarantee the existence of such a field, however. Can we construct one? If it did exist, Proposition 1.1 implies that it must have characteristic 2. Andrew proposed a natural guess: $\mathbb{Z}_2 \times \mathbb{Z}_2 = \{(0,0), (0,1), (1,0), (1,1)\}$. Under addition this is nicely behaved, but unfortunately isn't nice with respect to multiplication. For example, what's the multiplicative inverse of $(0,1)$?

Instead, we go back to the proof of Proposition 1.1: if a field $\mathbb{F}$ with 4 elements exists, then it must consist of linear combinations of $\mathbb{F}_2$. In other words, we must build $\mathbb{F}$ on the skeleton of $\mathbb{F}_2$. We quickly realized that we have a nice way to do this: we use Kronecker's approach! More precisely, if $g \in \mathbb{F}_2[x]$ is a quadratic irreducible in $\mathbb{F}_2$, then

$$\mathbb{F} := \mathbb{F}_2[x]/(g)$$

is a field of degree 2 over $\mathbb{F}_2$, i.e. with 4 elements in it. Playing around a bit, we see that we can take $g(x) = x^2 + x + 1$. (In fact, this is the unique quadratic irreducible in $\mathbb{F}_2[x]$.) This allows us to label the four elements of $\mathbb{F}$:

$$\mathbb{F} := \mathbb{F}_2[x]/(x^2 + x + 1) = \{0, 1, x, 1 + x\}.$$

(Note that we're omitting brackets around the elements, which really should be there but are annoying to write and read.) To build up some intuition about this field, let's write down addition and multiplication tables.

Date: May 14, 2018.
Based on notes by Ben Mygatt.

$+$	0	1	x	$x+1$
0	0	1	x	$x+1$
1	1	0	$x+1$	x
x	x	$x+1$	0	1
$x+1$	$x+1$	x	1	0

$\times$	1	x	$x+1$
1	1	x	$x+1$
x	x	$x+1$	1
$x+1$	$x+1$	1	x

These tables exhibit something potentially counterintuitive about $\mathbb{F}$: unlike $\mathbb{F}_p$, it is *not* cyclic with respect to addition. It is, however, cyclic with respect to multiplication. For example, x generates the group, since $x^2 = x+1$ and $x^3 = x(x+1) = 1$. In hindsight, this is clear on theoretical grounds, since $\mathbb{F}^\times$ has 3 elements and any group with a prime number of elements is cyclic (and generated by any non-identity element).

It turns out that this property generalizes to all finite fields:

Theorem 1.2. *If $\mathbb{F}$ is a finite field, then $\mathbb{F}^\times$ is cyclic.*

Example 1. Consider $\mathbb{F}_7 = \{0,1,2,3,4,5,6\}$. Then $\mathbb{F}_7^\times = \{1,2,3,4,5,6\}$, which is not obviously cyclic. (Unlike the case we considered above, this group doesn't have prime order!) However, some trial and error shows that

$$\mathbb{F}_7^\times = \langle 3 \rangle.$$

But 3 isn't the only generator; 5 also generates all of $\mathbb{F}_7^\times$. Note that 2, 4, and 6 are not generators (since they have order 3, 3, and 2 respectively).

Example 2. Next, let's consider a more complicated finite field, with nine elements. We construct it as before, using Kronecker's approach:

$$\mathbb{F}_9 := \mathbb{F}_3[x]/(x^2+1) = \{0,1,2,x,x+1,x+2,2x,2x+1,2x+2\}.$$

The operations here are addition and multiplication (mod 3), with the additional restriction that $x^2 = 2$. We see that $\mathbb{F}_9$ isn't cyclic group with respect to addition. Is $\mathbb{F}_9^\times$ cyclic? Note that x isn't a generator, since

$$x^2 = 2 \qquad \text{and} \qquad 2^2 = 1,$$

whence x has order 4. A bit more thought shows that $\mathbb{F}_9^\times = \langle x+1 \rangle$. Indeed,

$$(x+1)^2 = 2x \qquad \text{and} \qquad (2x)^2 = 2,$$

so $x+1$ doesn't have order 2 or 4. On the other hand, Lagrange's theorem guarantees that the order of any element in $\mathbb{F}_9^\times$ must have order dividing 8. It follows that $x+1$ must have order 8, and is therefore a generator. More generally, here's a multiplication table of $\mathbb{F}_9^\times$:

$\times$	1	2	x	$x+1$	$x+2$	$2x$	$2x+1$	$2x+2$
1	1	2	x	$x+1$	$x+2$	$2x$	$2x+1$	$2x+2$
2	2	1	$2x$	$2x+2$	$2x+1$	x	$x+2$	$x+1$
x	x	$2x$	2	$x+2$	$2x+2$	1	$x+1$	$2x+1$
$x+1$	$x+1$	$2x+2$	$x+2$	$2x$	1	$2x+1$	2	x
$x+2$	$x+2$	$2x+1$	$2x+2$	1	x	$x+1$	$2x$	2
$2x$	$2x$	x	1	$2x+1$	$x+1$	2	$2x+2$	$x+2$
$2x+1$	$2x+1$	$x+2$	$x+1$	2	$2x$	$2x+2$	x	1
$2x+2$	$2x+2$	$x+1$	$2x+1$	x	2	$x+2$	1	$2x$

2. Proof of Theorem 1.2

Recall that Theorem 1.2 asserts that $\mathbb{F}^\times$ is cyclic for any finite field $\mathbb{F}$. Even in the familiar case of $\mathbb{F}_p$ this is far from obvious (as our example with $\mathbb{F}_7$ shows). Indeed, it remains a fascinating open problem to determine an algorithm for producing a generator of $\mathbb{F}_p$ which is more efficient than trial-and-error; see Appendix A for a brief overview of some relevant results and conjectures.

There are many proofs of Theorem 1.2, none of them easy. The simplest I've seen relies on a sharper version of Lagrange's theorem for abelian groups. Before stating this, recall that the *exponent* of a group G, denoted $\exp(G)$, is the largest order of any element of G.

Lemma 2.1. *If G is abelian, then* $\operatorname{ord}(g) \mid \exp(G)$ *for all* $g \in G$.

A proof is sketched below. Taking the result on faith for now, we can prove Theorem 1.2 fairly easily.

Proof of Theorem 1.2. For brevity, set $\epsilon := \exp(\mathbb{F}^\times)$; we will show that $\epsilon = |\mathbb{F}^\times|$. Lemma 2.1 implies that every element of $\mathbb{F}^\times$ is a root of $x^\epsilon - 1$, and since $\mathbb{F}$ is a field, we know this polynomial has at most ϵ roots (see problem **5.4**). It follows that $|\mathbb{F}^\times| \leq \epsilon$. On the other hand, Lagrange's theorem implies that $\epsilon \mid |\mathbb{F}^\times|$, whence $\epsilon \leq |\mathbb{F}^\times|$. Thus, we've proved that

$$\exp(\mathbb{F}^\times) = \epsilon = |\mathbb{F}^\times|.$$

This implies that the order of some element of $\mathbb{F}^\times$ is the order of $\mathbb{F}^\times$, whence $\mathbb{F}^\times$ must be cyclic. □

Exercise 1. The goal of this exercise is to prove Lemma 2.1. Throughout let G be a finite abelian group.

(a) Prove that if $\operatorname{ord}(g)$ and $\operatorname{ord}(h)$ are coprime, then $\operatorname{ord}(gh) = \operatorname{ord}(g) \cdot \operatorname{ord}(h)$.
(b) Given arbitrary $g, h \in G$, prove that there exist positive integers k and ℓ such that

$$\operatorname{ord}(g^k h^\ell) = [\operatorname{ord}(g), \operatorname{ord}(h)],$$

where $[a, b]$ denotes the least common multiple of a and b. [*Hint: try working this out in the case that* $\operatorname{ord}(g) = 60$ *and* $\operatorname{ord}(h) = 630$.]
(c) Prove Lemma 2.1.
(d) Is the hypothesis that G be abelian necessary?

3. Characterizing finite fields

We've proved that any finite field $\mathbb{F}$ must have p^n elements. Does the converse hold? In other words, given a prime power p^n, must there exist a finite field with p^n elements? Our work with Kronecker's method suggests an affirmative answer. In fact, we will prove more:

Theorem 3.1. *For any prime power p^n, there exists a field $\mathbb{F}$ with precisely p^n elements. Moreover, this field is unique (up to isomorphism).*

This theorem, combined with Proposition 1.1, completely characterizes finite fields. As a first step, we prove:

Proposition 3.2. *If $\mathbb{F}$ is a finite field with p^n elements, then it is a splitting field of $x^{p^n} - x$ over $\mathbb{F}_p$.*

Proof. Given a finite field $\mathbb{F}$ with p^n elements. By Proposition 1.1, $\mathbb{F}$ is a field extension of $\mathbb{F}_p$. Theorem 1.2 implies that every element of $\mathbb{F}^\times$ is a root of $x^{p^n-1} - 1$; it follows that every element of $\mathbb{F}$ is a root of $x^{p^n} - x \in \mathbb{F}_p[x]$. Thus $\mathbb{F}$ is a splitting field of $x^{p^n} - x$ over $\mathbb{F}_p$. □

Remark. Suppose $\mathbb{F}$ and $\mathbb{F}'$ are finite fields with $|\mathbb{F}| = |\mathbb{F}'|$. Then Proposition 1.1 implies they both have cardinality p^n, and Proposition 3.2 implies they're both splitting fields of the same polynomial over $\mathbb{F}_p$. Since splitting fields are unique up to isomorphism by the Isomorphism Lifting Lemma, we deduce that $\mathbb{F} \simeq \mathbb{F}'$. In other words, we've shown that there is at most one finite field of any given cardinality (up to isomorphism).

Example 3. From the proof of Proposition 3.2 we immediately deduce that every element of $\mathbb{F}_p$ is a root of $x^p - x$. In particular,

$$x^p - x = \prod_{\alpha \in \mathbb{F}_p} (x - \alpha).$$

Try proving this directly to appreciate how nice the abstract approach is!

To prove Theorem 3.1, all that remains is to show that for any prime power p^n there exists a finite field with cardinality p^n. There are multiple ways to accomplish this. Probably the most popular is to prove the existence of an irreducible polynomial $\pi \in \mathbb{F}_p[x]$ of degree n, since then

$$\mathbb{F} := \mathbb{F}_p[x]/(\pi)$$

would have cardinality p^n. Instead, we take a more direct approach.

Exercise 2. Let $\mathbb{F}$ be a splitting field of $f(x) := x^{p^n} - x$ over $\mathbb{F}_p$. The goal of this exercise is to prove that $|\mathbb{F}| = p^n$. Throughout, let $\mathcal{R}$ denote the set of all roots of f in $\mathbb{F}$.

(a) Prove that if $x, y \in \mathcal{R}$, then both $x + y$ and xy are also in $\mathcal{R}$.

(b) Prove that $\mathcal{R}$ is a field.

(c) Deduce that $|\mathbb{F}| = p^n$.

This concludes our proof of Theorem 3.1: for each prime power p^n there exists precisely one field (up to isomorphism) with p^n elements. Thus the following notation is well-defined:

Definition. Given q a power of a prime, we denote the field of cardinality q by $\mathbb{F}_q$.

Having characterized finite fields, we're ready to study Galois theory in this context.

4. Galois theory for finite fields

Claim. *Given a finite field $\mathbb{F}$ with* $\operatorname{char} \mathbb{F} = p$. *Then $\mathbb{F}/\mathbb{F}_p$ is a Galois extension.*

Proof. Let $f(x) := x^{p^n} - x$. By Proposition 3.2 we know $\mathbb{F}$ is a splitting field of f over $\mathbb{F}_p$, so to prove that $\mathbb{F}/\mathbb{F}_p$ is Galois it suffices to show that f is separable. This immediately follows from Exercise 2. Alternatively, note that $f' = -1$, so it must be relatively prime to f; this implies that f is separable (see Lecture 17). □

Having established that $\mathbb{F}/\mathbb{F}_p$ is Galois, the natural question is: what is $\operatorname{Gal}(\mathbb{F}/\mathbb{F}_p)$? To start investigating this question, we build up intuition by explicitly determining some of the elements of the Galois group. Of course the trivial map is in there. What's an example of a nontrivial element in $\operatorname{Gal}(\mathbb{F}/\mathbb{F}_p)$? In other words, can we identify a nontrivial isomorphism $\mathbb{F} \to \mathbb{F}$ which fixes every element of $\mathbb{F}_p$? It's not so obvious!

Let's start by thinking about what nontrivial maps fix every element of $\mathbb{F}_p$. This might remind you of a famous result from elementary number theory: Fermat's Little Theorem. Recall that this states that $x^p = x$ for all $x \in \mathbb{F}_p$. (This is a special case of Proposition 3.2.) In other words, the map $x \mapsto x^p$ fixes every element of $\mathbb{F}_p$. How does it behave on $\mathbb{F}$? Is it nontrivial? Is it a homomorphism?

Proposition 4.1. *Given a finite field $\mathbb{F}$ of characteristic p, define the map $\phi_p : \mathbb{F} \to \mathbb{F}$ by $\phi_p(t) := t^p$. Then $\phi_p \in \operatorname{Gal}(\mathbb{F}/\mathbb{F}_p)$.*

Proof. We saw above that ϕ_p fixes all of $\mathbb{F}_p$, so it suffices to show that it is an isomorphism. Let's start by verifying that ϕ_p is a homomorphism:

$$\phi_p(xy) = (xy)^p = x^p y^p = \phi_p(x)\phi_p(y) \qquad \text{and} \qquad \phi_p(x+y) = (x+y)^p = x^p + y^p = \phi_p(x) + \phi_p(x).$$

It remains to check that ϕ_p is a bijection. First we check that ϕ_p is injective. Suppose $\phi_p(x) = \phi_p(y)$. Then

$$(x - y)^p = \phi_p(x - y) = \phi_p(x) - \phi_p(y) = 0,$$

whence $x = y$. Since $\mathbb{F}$ is finite, it follows that ϕ_p must be a bijection. We've therefore proved that ϕ_p is an automorphism which fixes all the elements of $\mathbb{F}_p$. □

Example 4. Recall from the beginning of this document the field of four elements $\mathbb{F}_4 = \{0, 1, x, x + 1\}$. As we saw, we have $\phi_2(x) = x^2 = x + 1$ in this field. Thus, ϕ_2 is a nontrivial automorphism of $\mathbb{F}_4$.

Thus we've discovered one nontrivial map in $\mathrm{Gal}(\mathbb{F}/\mathbb{F}_p)$: the map ϕ_p. Are there others? Well, since $\mathrm{Gal}(\mathbb{F}/\mathbb{F}_p)$ is a group, any power of ϕ_p is also an element of the Galois group. What are some other nontrivial elements of the Galois group? It turns out there aren't any!

Proposition 4.2. *Given finite field $\mathbb{F}$ of characteristic p. Then $\mathrm{Gal}(\mathbb{F}/\mathbb{F}_p)$ is cyclic, and is generated by the map ϕ_p defined in Proposition 4.1.*

Proof. For concreteness, let's say $\mathbb{F} = \mathbb{F}_{p^n}$. It follows that $|\mathrm{Gal}(\mathbb{F}/\mathbb{F}_p)| = [\mathbb{F} : \mathbb{F}_p] = n$, so to prove the claim it suffices to show that

$$\mathrm{ord}(\phi_p) = n.$$

Recall from Theorem 1.2 that $\mathbb{F}^\times$ is cyclic, so $\mathbb{F}^\times = \langle\gamma\rangle$ for some $\gamma \in \mathbb{F}^\times$. Then we have

$$\begin{aligned}
\phi_p(\gamma) &= \gamma^p \\
\phi_p^2(\gamma) &= \phi_p(\phi_p(\gamma)) = \phi_p(\gamma^p) = \gamma^{p^2} \\
&\vdots \\
\phi_p^k(\gamma) &= \gamma^{p^k}.
\end{aligned}$$

Since $\mathrm{ord}(\gamma) = p^n - 1$, we deduce that $\phi_p^k(\gamma) \neq \gamma$ for any $k < n$, whence $\mathrm{ord}(\phi_p) \geq n$. On the other hand, since $\phi_p \in \mathrm{Gal}(\mathbb{F}/\mathbb{F}_p)$ we trivially have $\mathrm{ord}(\phi_p) \leq |\mathrm{Gal}(\mathbb{F}/\mathbb{F}_p)| = n$. This completes the proof. □

Clearly the map ϕ_p plays an important role in the study of finite fields. It therefore deserves a name:

Definition. Given a finite field $\mathbb{F}$ of characteristic p, the *Frobenius map* is the map $\mathbb{F} \to \mathbb{F}$ defined by $t \mapsto t^p$. (We denoted it ϕ_p above.)

Remark. This is just one important object named after Frobenius; we'll see another result of his in the next section. For more on his work, see the wikipedia article *List of things named after Ferdinand Georg Frobenius*.

5. The discriminant and the Galois group

Last lecture, we looked at a couple of tricks that are often used to find the Galois group of a given polynomial. However, in general, it remains an open problem on how exactly to compute the Galois group of *any* polynomial. Here we describe one useful tool. Intuitively, $\mathrm{Gal}(f)$ captures the symmetries among the roots of f. But there's a more basic way to capture these symmetries: the discriminant.

Definition (Discriminant)**.** Given $f \in K[x]$ with roots $r_1, r_2, \ldots, r_n$ (not necessarily distinct). The *discriminant* of f is

$$\mathrm{disc}(f) := \prod_{i<j} (r_i - r_j)^2.$$

Example 5. $\mathrm{disc}(x^2 + bx + c) = b^2 - 4c$. This looks familiar – it's the discriminant in the quadratic formula!

Note that the discriminant is a number, so we can't expect it to be as nuanced as the Galois group. On the other hand, it's straightforward to compute, which offers a major advantage!

Example 6. Here are more examples of discriminants (we include the quadratic example for completeness):

- $\mathrm{disc}(x^2 + ax + b) = a^2 - 4b$.
- $\mathrm{disc}(x^3 + ax + b) = -4a^3 - 27b^2$.
- $\mathrm{disc}(x^4 + ax + b) = -27a^4 + 256b^3$.
- $\mathrm{disc}(x^5 + ax + b) = 256a^5 + 3125b^4$.

It turns out that the pattern continues:

- $\mathrm{disc}(x^n + ax + b) = d_{n-1}a^n + d_n b^{n-1}$, where $d_n := (-1)^{\frac{n(n-1)}{2}} n^n$.

Remark. Recall that given $f \in \mathbb{Q}[x]$ of degree ≤ 4 there exist general formulas for the roots. All of these involve combinations of the coefficients of f, rational numbers, field operations, and nested radicals. Here are the expressions appearing in the innermost nested radicals:

- $x^2 + ax + b \overset{\text{innermost radical}}{\rightsquigarrow} a^2 - 4b$
- $x^3 + ax + b \overset{\text{innermost radical}}{\rightsquigarrow} -\frac{1}{27}(-4a^3 - 27b^2)$
- $x^4 + ax + b \overset{\text{innermost radical}}{\rightsquigarrow} -27(-27a^4 + 256b^3)$

Thus, the discriminant is (up to a rescaling) the quantity appearing under the innermost radical in the formula for the roots! It's unclear how to generalize this to higher degree polynomials, however, since there's no formula for the roots.

Note that for polynomials of the form $x^n + ax + b \in \mathbb{Q}[x]$, the discriminant is always rational. Is there a parallel result for general polynomials $f \in K[x]$?

Proposition 5.1. *Given a separable polynomial $f \in K[x]$. Then* $\text{disc}(f) \in K$.

Proof. Note that any permutation of the roots of f leaves the discriminant fixed; in particular, $\text{disc}(f)$ is fixed by every element of $\text{Gal}(f)$. Galois theory implies that $\text{disc}(f) \in K$. □

It turns out that the discriminant can be used to deduce nontrivial information about the Galois group.

Lemma 5.2. *Given K a field of characteristic $\neq 2$, and suppose $f \in K[x]$ is a separable polynomial of degree n. Then* $\text{Gal}(f) \leq A_n$ *if and only if* $\text{disc}(f)$ *is a perfect square in K.*

Example 7. Consider $f(x) := x^5 + x + 1$. Then $\text{disc}(f) = 3381$, which isn't a perfect square. It follows that $\text{Gal}(f)$ isn't a subgroup of A_5.

WARNING. In this example, we *cannot* conclude that $\text{Gal}(f) \simeq S_5$; all we can say is that $\text{Gal}(f)$ must contain an odd permutation. (Indeed, it turns out that $\text{Gal}(f) \simeq S_3 \times \mathbb{Z}_2$.)

Example 8. A quick computation shows that the polynomial $h(x) := x^4 + 8x + 12$ has discriminant $2^{12} \times 3^4$, hence is a perfect square. It follows that $\text{Gal}(h) \leq A_4$. (In fact, it can be shown that $\text{Gal}(h) \simeq A_4$.)

Since a very small proportion of numbers are perfect squares, the lemma suggests that it's rare for $\text{Gal}(f)$ to be a subgroup of A_n. In fact, this can be quantified: in 1936, van der Waerden proved that 100% of degree n polynomials with rational coefficients have Galois group S_n. Thus, for example, 100% of quintic polynomials in $\mathbb{Q}[x]$ are not solvable in radicals.

Proof of Lemma. Let L be a splitting field of f over K, and let $r_1, r_2, \ldots, r_n$ denote the roots of f. Set

$$\delta := \prod_{i<j} (r_i - r_j).$$

For any $\sigma \in \text{Gal}(f)$, problem **3.3(f)** implies

$$\sigma(\delta) = \text{sgn}(\sigma)\delta.$$

Thus, $\sigma \in A_n$ if and only if δ is fixed by σ. It follows that

$$\text{Gal}(f) \leq A_n \qquad \Longleftrightarrow \qquad \sigma(\delta) = \delta \ \ \forall \sigma \in \text{Gal}(f) \qquad \Longleftrightarrow \qquad \delta \in K.$$

Since $\delta^2 = \text{disc}(f)$, this concludes the proof. □

Exercise 3. Where in the proof did we use the hypothesis char $K \neq 2$?

6. Cycle types and the Galois group

Recall that one of our most powerful irreducibility tests was to study the local irreducibility. More precisely, given $f \in \mathbb{Z}[x]$ and a prime p not dividing the leading coefficient of f, we showed that if f is irreducible over $\mathbb{F}_p$ then f is irreducible over $\mathbb{Q}$. One advantage of this test is that there are only finitely many irreducible polynomials in $\mathbb{F}_p[x]$ of given degree, so one can determine by brute force whether or not f factors over $\mathbb{F}_p$.

Richard Dedekind, one of the pioneers of algebraic number theory, took this a step further: he discovered a lovely connection between the local factorizations of f and the structure of $\mathrm{Gal}(f)$. To state his result, we shall need the concept of *cycle type*. Recall that any permutation can be written in a unique way (up to order) as the product of disjoint cycles. The cycle type describes the shape of this decomposition:

Definition. We say $\sigma \in S_n$ has *cycle type* $(n_1, n_2, \ldots, n_k) \in \mathbb{N}^k$ iff σ can be written in the form

$$\sigma = \sigma_1 \sigma_2 \cdots \sigma_k$$

where the σ_i are disjoint cycles with $|\sigma_i| = n_i$.

Example 9. The cycle type of the permutation $(1\ 5\ 4)(2\ 6)(7\ 8) \in S_8$ is $(1, 2, 2, 3)$. Note that the cycle type is unique up to order, e.g. it could also be written $(2, 1, 3, 2)$. Furthermore, observe that the sum of the components of the cycle type is 8. This isn't a coincidence: if the cycle type of $\sigma \in S_n$ is $(n_1, n_2, \ldots, n_k)$ then $n_1 + n_2 + \cdots + n_k = n$.

With this concept in hand, we can now state Dedekind's remarkable result. Unfortunately, the proof is too far outside the scope of our course – it requires tools from algebraic number theory – but we'll show some applications to Galois theory following the theorem.

Theorem 6.1 (Dedekind)**.** *Given an irreducible monic polynomial $f \in \mathbb{Z}[x]$. Suppose that over $\mathbb{F}_p$ we have the factorization*

$$f = \pi_1 \pi_2 \cdots \pi_k$$

into irreducible monic polynomials $\pi_i \in \mathbb{F}_p[x]$. Let $d_i := \deg \pi_i$. Then $\mathrm{Gal}(f)$ *contains an element of cycle type $(d_1, d_2, \ldots, d_k)$.*

Example 10. Consider the polynomial $f(x) := x^5 + 4x^4 + 4x^3 - 4x^2 - 2x + 2$. Right away we know $\mathrm{Gal}(f) \leq S_5$. Which subgroup is it? By Eisenstein, f is irreducible over $\mathbb{Q}$. Thus we can apply Dedekind's theorem:

- Over $\mathbb{F}_2$ we have $f(x) = x^5$, whence $\mathrm{Gal}(f)$ contains an element of cycle type (5). In other words, $\mathrm{Gal}(f)$ contains a 5-cycle.
- Over $\mathbb{F}_5$ we have $f(x) = (x^2 - 2)(x - 1)(x - 2)(x - 3)$, whence $\mathrm{Gal}(f)$ contains an element of cycle type $(2, 1, 1, 1)$. In other words, $\mathrm{Gal}(f)$ contains a transposition.

Since S_5 is generated by any 5-cycle and any transposition, we conclude that $\mathrm{Gal}(f) \simeq S_5$. (Thus f isn't solvable in radicals.)

Dedekind's theorem asserts that the degree type of f over $\mathbb{F}_p$ implies the existence of an element in $\mathrm{Gal}(f)$ with corresponding cycle type. One is immediately led to ask about the converse: given that a given cycle type is represented in $\mathrm{Gal}(f)$, does it follow that there exists a prime p such that f has the corresponding degree type over $\mathbb{F}_p$? The following result, discovered by our old friend Frobenius in 1880, gives a very strong affirmative answer to this question.

Theorem 6.2 (Frobenius, 1896)**.** *Given a monic polynomial $f \in \mathbb{Z}[x]$, let $G :=$* $\mathrm{Gal}(f)$*. Then the proportion of primes p for which f has degree type $\vec{d}$ over $\mathbb{F}_p$ is the same as the proportion of elements of* $\mathrm{Gal}(f)$ *with cycle type $\vec{d}$. In other words,*

$$\frac{1}{\#\{p \leq x\}} \#\{p \leq x : f \text{ has degree type } \vec{d} \text{ over } \mathbb{F}_p\} \xrightarrow{x \to \infty} \frac{1}{|G|} \# \left\{\sigma \in G : \sigma \text{ has cycle type } \vec{d}\right\}.$$

Remark. In 1922 Chebotarev proved a beautiful generalization of Frobenius' theorem, called the *Chebotarev Density Theorem*, which plays an important role in modern number theory.

Example 11. Let $f(x) := x^4 - x - 1$. It turns out (see below) that $\mathrm{Gal}(f) \simeq S_4$. Since S_4 has precisely six 4-cycles, Frobenius' theorem predicts that $f(x)$ should be a perfect 4^{th} power (mod p) a quarter of the time. Sure enough, this is what happens! Frobenius' theorem also (correctly) predicts that f factors into a product of four linear factors (mod p) precisely $1/24$ of the time.

Example 12. Let $h(x) := x^4+8x+12$, which has Galois group A_4 (see below). Since A_4 contains no 4-cycles, Frobenius' theorem predicts that $h(x)$ is a perfect 4^{th} power (mod p) 0% of the time. Frobenius' theorem also (correctly) predicts that f factors into a product of four linear factors (mod p) precisely $1/12$ of the time.

Here's one immediate consequence of Frobenius' theorem which is super useful for computational Galois theory:

Corollary 6.3. *Given a monic polynomial $f \in \mathbb{Z}[x]$, let*

$$\mathcal{P}_f := \{p \text{ prime} : f \text{ is separable over } \mathbb{F}_p \text{ and splits completely in } \mathbb{F}_p\}.$$

Then $|\mathrm{Gal}(f)| = \lim_{x\to\infty} \dfrac{\#\{p \le x\}}{\#\{p \le x : p \in \mathcal{P}_f\}}$.

Proof. $\mathrm{Gal}(f)$ has precisely one element of cycle type $(1, 1, \ldots, 1)$: the identity. □

In practice this is quite useful! Indeed, it's not too difficult to program a computer to compute the ratio $\dfrac{\#\{p \le x\}}{\#\{p \le x : p \in \mathcal{P}_f\}}$ for any given x. Taking x larger and larger yields the size of Galois group of f!

Combining the discriminant, Dedekind's theorem, and the methods described in Lecture 23, we have a decent set of tools for determining the Galois group of a given polynomial. Although there are plenty of other tricks people have invented, the problem of efficiently determining $\mathrm{Gal}(f)$ remains difficult and largely open.

Exercise 4. Let $f(x) := x^4 - x - 1$. The goal of this exercise is to prove that $\mathrm{Gal}(f) \simeq S_4$.
(a) Prove that $4 \mid \mathrm{Gal}(f)$.
(b) Use the discriminant to prove that $|\mathrm{Gal}(f)| = 4$, 8, or 24.
(c) Use Dedekind's theorem to prove that $3 \mid \mathrm{Gal}(f)$.

Exercise 5. Let $h(x) := x^4 + 8x + 12$.
(a) Prove that h is irreducible over $\mathbb{Q}$.
(b) Prove that $\mathrm{Gal}(h) \simeq A_4$.

7. Cyclotomic extensions

Recall (from a long time ago – Lecture 12) the following results about constructible numbers:

Proposition 7.1. *The set of all constructible numbers forms a field.*

Theorem 7.2. *If $\beta \in \mathbb{C}$ is constructible, then $[\mathbb{Q}(\beta) : \mathbb{Q}]$ is a power of 2.*

As we mentioned in that lecture, the converse to this theorem is false. We've now developed enough theory to see this:

Exercise 6. Let $h(x) := x^4 + 8x + 12$. Prove that all of its roots have degree 4 over $\mathbb{Q}$, but that not all of its roots are constructible. [*Hint: use the fact (proved above) that* $\mathrm{Gal}(h) \simeq A_4$.]

Thus the theorem isn't quite a characterization of constructible numbers, although it's already strong enough to prove the impossibility of constructing certain numbers (see Lecture 12). Employing a bit more care, however, we can obtain an if and only if characterization of constructibility. The first step is the following exercise:

Exercise 7. Suppose $x \in \mathbb{C}$ is constructible. Prove that $\pm\sqrt{x}$ are also constructible. [*Hint: prove this for $x \in \mathbb{R}_{>0}$ first.*]

Lemma 7.3. *Suppose K is a field consisting of constructible numbers, and that L/K is a quadratic extension. Then L consists of constructible numbers.*

Proof. By problem **5.2(c)**, we can write $L = K(\beta)$ for some $\beta \in \mathbb{C}$ satisfying $\beta^2 \in K$. Thus β^2 is constructible, whence (by the exercise above) β must also be constructible. Since constructible numbers form a field, $L = K(\beta)$ must consist of constructible numbers. □

Corollary 7.4. *$\beta \in \mathbb{C}$ is constructible if and only if the extension $\mathbb{Q}(\beta)/\mathbb{Q}$ can be decomposed into a tower of quadratic extensions.*

Proof. The forward direction is a direct consequence of our work in Lecture 12 (at each stage of the construction, we create an extension which is either trivial or quadratic). For the reverse direction, suppose $\mathbb{Q}(\beta)/\mathbb{Q}$ can be decomposed into a tower of quadratic extensions:

$$\mathbb{Q}(\beta) = K_0 \supset K_1 \supset \cdots \supset K_{\ell-1} \supset K_\ell = \mathbb{Q}.$$

Since $\mathbb{Q}$ consists of constructible numbers, Lemma 7.3 implies $K_{\ell-1}$ consists of constructible numbers, which then implies that $K_{\ell-2}$ consists of constructible numbers, etc. Thus, we deduce that $\mathbb{Q}(\beta)$ consists of constructible numbers. In particular, β must be constructible. □

Thus we've come up with a more careful phrasing of Theorem 7.2 which is an equivalence. For certain choices of β, however, Theorem 7.2 itself admits a converse:

Proposition 7.5. *A root of unity ζ_n is constructible if and only if $[\mathbb{Q}(\zeta_n) : \mathbb{Q}]$ is a power of 2.*

Proof. The forward direction follows from the general Theorem 7.2, so it suffices to handle the reverse direction. Suppose $[\mathbb{Q}(\zeta_n) : \mathbb{Q}] = 2^\ell$. Recall (problem **11.3**) that $\mathbb{Q}(\zeta_n)/\mathbb{Q}$ is Galois. The plan is to use the FTGT and Sylow's theorem to produce a chain of quadratic extensions connecting $\mathbb{Q}$ to $\mathbb{Q}(\zeta_n)$.

Let $G := \mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$. By Sylow's theorem (see Lecture 17) there exists an index 2 subgroup $G_1 \leq G$. Applying the FTGT produces a subfield $K_1 \subset \mathbb{Q}(\zeta_n)$. Here's a picture:

$$\begin{array}{ccc} \mathbb{Q}(\zeta_n) & & \{e\} \\ \Big| 2 & & \Big| 2^{\ell-1} \\ K_1 & \xleftrightarrow{\text{Galois correspondence}} & G_1 \\ \Big| 2^{\ell-1} & & \Big| 2 \\ \mathbb{Q} & & G \end{array}$$

Next we apply Sylow's theorem to produce a subgroup $G_2 \leq G_1$ of index 2, and then the FTGT to give a corresponding subfield K_2:

$$\begin{array}{ccc} \mathbb{Q}(\zeta_n) & & \{e\} \\ \Big| 2 & & \Big| 2^{\ell-2} \\ K_1 & & G_2 \\ \Big| 2 & \xleftrightarrow{\text{Galois correspondence}} & \Big| 2 \\ K_2 & & G_1 \\ \Big| 2^{\ell-2} & & \Big| 2 \\ \mathbb{Q} & & G \end{array}$$

Iterating this we obtain the following correspondence:

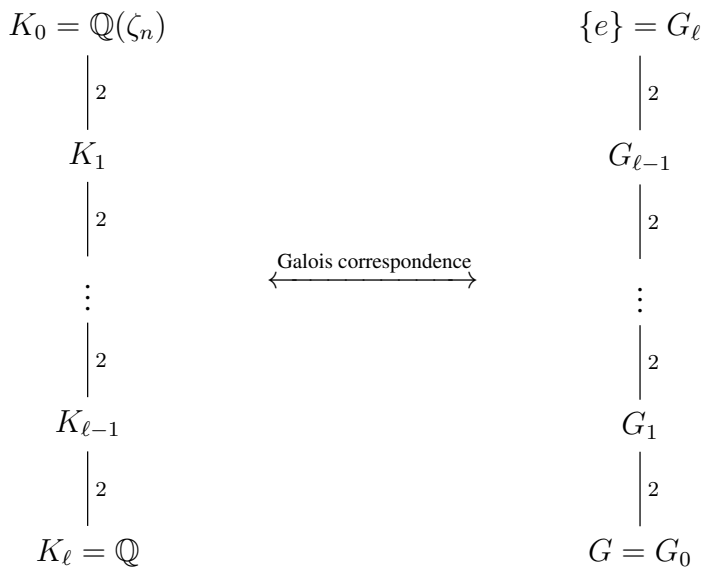

We've thus decomposed $\mathbb{Q}(\zeta_n)/\mathbb{Q}$ into a tower of quadratic extensions, whence Corollary 7.4 implies ζ_n is constructible. □

In view of their importance both in constructibility and in our proof of Galois' solvability criterion, extensions of the form $\mathbb{Q}(\zeta_n)/\mathbb{Q}$ deserve a name:

Definition. A *cyclotomic extension* is any field extension of the form $K(\zeta)/K$, where ζ is a root of unity.

Recall (problem **11.3**) that $\mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q})$ embeds in $(\mathbb{Z}/n\mathbb{Z})^\times$; we used this to prove that $\mathbb{Q}(\zeta_n)/\mathbb{Q}$ is abelian. The same proof works for any cyclotomic extension, with $\mathbb{Q}$ replaced by an arbitrary field K. However, in the special case of $\mathbb{Q}(\zeta_n)/\mathbb{Q}$, more is true:

Proposition 7.6. $\mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) \simeq (\mathbb{Z}/n\mathbb{Z})^\times$.

Remark. This doesn't hold for general cyclotomic extensions, e.g. $\mathrm{Gal}(\mathbb{R}(\zeta_7)/\mathbb{R}) \not\simeq (\mathbb{Z}/7\mathbb{Z})^\times$.

Right away we deduce that

$$[\mathbb{Q}(\zeta_n) : \mathbb{Q}] = |\mathrm{Gal}(\mathbb{Q}(\zeta_n) : \mathbb{Q})| = \varphi(n),$$

where $\varphi(n)$ is Euler's totient function. We're now ready to prove an assertion we made in Lecture 12, originally a combination of work of Gauss and Wantzel:

Theorem 7.7. *The regular n-gon is constructible if and only if n is a product of some power of 2 with any number of distinct Fermat primes.*

Remark. Recall that a *Fermat prime* is any prime of the form $2^k + 1$; as you proved in problem **6.5**, Fermat primes must be of the form $2^{2^k} + 1$.

Proof. First recall that the regular n-gon is constructible if and only if ζ_n is constructible. Proposition 7.5 asserts that this happens if and only if $[\mathbb{Q}(\zeta_n) : \mathbb{Q}]$ is a power of 2. Since $[\mathbb{Q}(\zeta_n) : \mathbb{Q}] = \varphi(n)$, we've proved that the regular n-gon is constructible if and only if $\varphi(n)$ is a power of 2.

Factor n as a product of primes, say $n = 2^e p_1^{e_1} p_2^{e_2} \cdots p_\ell^{e_\ell}$ where $e \geq 0$, $e_i > 0$ for all i, and the p_i are all distinct odd primes. It's well-known that φ is multiplicative, whence

$$\varphi(n) = \varphi(2^e p_1^{e_1} p_2^{e_2} \cdots p_\ell^{e_\ell}) = \varphi(2^e)\varphi(p_1^{e_1})\varphi(p_2^{e_2}) \cdots \varphi(p_\ell^{e_\ell}).$$

Furthermore, it's well-known that $\varphi(p^r) = (p-1)p^{r-1}$ whenever $r \geq 1$, whence

$$\varphi(n) = 2^f (p_1 - 1)(p_2 - 1) \cdots (p_\ell - 1) p_1^{e_1 - 1} p_2^{e_2 - 1} \cdots p_\ell^{e_\ell - 1}$$

for some $f \geq 0$. It's clear that this is a power of 2 if and only if $e_1 = e_2 = \cdots = e_\ell = 1$ and all the p_i are Fermat primes. This concludes the proof. □

Before leaving the subject of cyclotomic extensions, we'd be remiss not to mention a beautiful theorem (which came out of work of Kronecker in 1853, Weber in 1886, and finally Hilbert in 1896):

Theorem 7.8 (Kronecker-Weber)**.** *Every finite abelian extension of $\mathbb{Q}$ is a subfield of $\mathbb{Q}(\zeta_n)$ for some $n \in \mathbb{Z}$. In words: every finite abelian extension is contained in a cyclotomic extension.*

Remark. This is reminiscent of Cayley's theorem (problem **3.6**): every finite group is contained in a symmetric group. Thus the language of symmetric groups is the universal language of finite groups: every finite group can be described in terms of permutations, since every finite group embeds in some S_n. Analogously, the language of $\mathbb{Q}(\zeta_n)$ is the universal language of finite abelian extensions of $\mathbb{Q}$: every finite abelian extension of $\mathbb{Q}$ can be described in terms of rational linear combinations of roots of unity.

In particular, any α which is algebraic over $\mathbb{Q}$ can be expressed as a rational linear combination of roots of unity. For example, $\sqrt{5} = \zeta_5 - \zeta_5^2 - \zeta_5^3 + \zeta_5^4$.

Appendix A. Generating $\mathbb{F}^\times$

As mentioned in Section 2, even though we know that $\mathbb{F}^\times$ is cyclic by Theorem 1.2, it is a notoriously difficult problem to determine a generator in any way which is more efficient than trial-and-error. Of course, we can eliminate some elements from consideration: clearly 1, 0, and -1 will never generate $\mathbb{F}_p^\times$ when $p \geq 5$. A bit more thought shows that a perfect square can never generate $\mathbb{F}_p^\times$. Can you see why not?

Despite phrasing it dismissively, trial-and-error isn't necessarily a bad way to search for a generator; it depends on how small the smallest generator is. In other words, it's possible that starting at 2 and going up, one doesn't have to try too many elements before finding a generator. Is this true? To phrase this more precisely:

Question 2. *Let g_p denote the least positive integer which generates $\mathbb{F}_p^\times$. Can we prove that g_p isn't too large?*

It is conjectured that $g_p \ll \log^2 p$. (Here the notation $f(x) \ll g(x)$ means that there exists a constant $C > 0$ such that $|f(x)| \leq Cg(x)$ for all x.) If this conjecture is true, then trial and error is pretty good: $\log p$ grows very slowly with p (it's proportional to the number of digits of p). Unfortunately, we are very far from being able to prove this conjecture. The best result to date is

Theorem A.1 (Burgess, 1962)**.** *For any fixed $\epsilon > 0$, we have $g_p \ll p^{1/4+\epsilon}$.*

To appreciate just how far this is from the conjecture, note that $\log p \ll p^\epsilon$ for any fixed $\epsilon > 0$. It's worth pointing out that the conjecture is known to hold on average: Burgess and Elliott proved in 1968 that

$$\frac{1}{\pi(x)} \sum_{p \leq x} g_p \ll (\log x)^{2+\epsilon}.$$

(Here $\pi(x)$ denotes the number of primes $p \leq x$.) Of course, this type of average result tells us very little about the size of g_p for any particular prime p.

So much for upper bounds on g_p. Lower bounds are also mysterious. For example, there are a bunch of p for which $g_p = 2$, so the strongest *universal* lower bound one could hope for is $g_p \geq 2$. A more interesting question is:

Question 3. *Is there a nontrivial lower bound on g_p for all sufficiently large p?*

In other words, maybe $g_p = 2$ only for some finite set of primes p, and then eventually starts to grow? No one knows the answer to this question. However, a notorious conjecture asserts that this isn't the case:

Conjecture A.2 (Artin's Conjecture)**.** *Given any integer a which is not -1 or a perfect square, we have*

$$\mathbb{F}_p^\times = \langle a \rangle$$

for infinitely many primes p.

In fact, Artin conjectured a much stronger assertion: whenever $a \not\equiv 1 \pmod 4$ is a prime, it seems to be the case that a generates $\mathbb{F}_p^\times$ for $\approx 37\%$ of primes p. Results of M. Goldfeld, P. J. Stephens, and myself with G. Martin prove that this more precise conjecture is true 'on average' (precisely: we can prove that $\mathbb{F}_p^\times = \langle a \rangle$ for 37% of *pairs* (a, p) with $|a| < p$).

Although there is still no single value of a for which Artin's conjecture is known to hold, we do know that it holds for almost all choices of a. A crazy awesome example of this is the following:

Theorem A.3 (Heath-Brown, 1986)**.** *Let $\mathcal{P}$ denote the set of all primes. There exist $p, q \in \mathcal{P}$ such that Artin's conjecture holds for all $a \in \mathcal{P} \setminus \{p, q\}$.*

Thus, for example, we know that Artin's conjecture must hold for at least one of $3, 5, 7$, but we have no idea which one! (In fact, presumably Artin's conjecture holds for *all* $a \in \mathcal{P}$; Heath-Brown's result simply says that this is true with at most two exceptions.)

Even though it's difficult to produce a generator of $\mathbb{F}_p$, Theorem 1.2 guarantees the existence of one. In fact, we know a bit more:

Theorem A.4 (Gauss)**.** *For any integer $n \geq 2$, $\mathbb{Z}_n^\times$ has either 0 or $\varphi\big(\varphi(n)\big)$ generators. (Here $\varphi(n)$ is Euler's totient function, the number of positive integers less than and relatively prime to n.)*

In other words, if $\mathbb{Z}_n^\times$ is cyclic, then we know precisely how many generators it has. Combining this with Theorem 1.2, we deduce that $\mathbb{F}_p^\times$ has precisely $\varphi(p-1)$ generators. Although a bit more is known (Gauss discovered some identities involving the sum and product of all the generators, for example), the structure of the generators remains largely a mystery.