25 Field Extensions

25.1 Primary Fields

We have the following useful fact about fields:

Fact 25.1
Every field is a (possibly infinite) extension of either $\mathbb{Q}$, or $\mathbb{F}_p$ for a prime p. These are called the **primary fields**.

Proof. In general, for any ring R, there is a unique ring homomorphism $\mathbb{Z} \to R$ — we must have $1 \mapsto 1_R$, so then $n \mapsto \underbrace{1_R + \cdots + 1_R}_{n} = n_R$ for positive integers n, and $-n \mapsto -n_R$.

The image of the homomorphism is a quotient of $\mathbb{Z}$ — it's either $\mathbb{Z}$ or $\mathbb{Z}/n\mathbb{Z}$. Now consider the kernel of this homomorphism. If R is an integral domain (note that all fields are domains), then either the homomorphism is one-to-one, or its kernel is (p) for a prime p — otherwise, the image would be $\mathbb{Z}/n\mathbb{Z}$ for composite n, which is not a domain (as it has zero divisors).

Now taking $R = F$ to be a field, if the kernel is zero, then $\mathbb{Z}$ is a subring of F. But then $\mathbb{Q} = \text{Frac}(\mathbb{Z})$ must be inside F as well (since we can invert elements in a field) — in our original notation, the copy of $\mathbb{Q}$ in F is the fractions of the form n_R/m_R.

On the other hand, if the kernel is (p), then we have a copy of $\mathbb{Z}/p\mathbb{Z}$ in F, and we're done. □

Definition 25.2
The generator of the kernel (as in the above proof) is called the **characteristic** of the field.

So fields of characteristic 0 contain $\mathbb{Q}$, and fields of characteristic p contain $\mathbb{Z}/p\mathbb{Z}$ (and these are the only possible characteristics).

25.2 Algebraic Elements

Last time, we defined algebraic elements in a field extension L/K:

Definition 25.3
An element $\alpha \in L$ is **algebraic** over K if $P(\alpha) = 0$ for some nonzero $P \in K[x]$.

As stated last class, α is algebraic if and only if $K(\alpha)/K$ is finite (since a polynomial in α is the same as a linear relation between powers of α).

We also looked at towers of extensions $E/F/K$ — here E/K is called the **composite** extension, while E/F and F/K are called **intermediate** extensions. In particular, we saw the following theorem:

Theorem 25.4
We have
$$[E : K] = [E : F] \cdot [F : K].$$
In particular, E/K is finite if and only if both E/F and F/K are finite.

This has some useful corollaries regarding algebraic elements.

Corollary 25.5
If $\alpha, \beta \in L$ are algebraic over K, then $\alpha + \beta$, $\alpha\beta$, and $\frac{\alpha}{\beta}$ are also algebraic.

Proof. If α and β are algebraic, then $K(\alpha)/K$ and $K(\alpha, \beta)/K(\alpha)$ are both finite — since β satisfies a polynomial relation with coefficients in K, it satisfies the same polynomial relation with coefficients in $K(\alpha)$. So we can conclude that $K(\alpha, \beta)/K$ is finite, and therefore any element in it is algebraic. □

Corollary 25.6
Given an arbitrary extension, the set of elements in L which are algebraic over K form a subfield of L, called the **algebraic closure** of K in L.

For example, the algebraic closure of $\mathbb{Q}$ in $\mathbb{C}$ is called the **algebraic numbers**.

This is an abstract argument that doesn't exactly tell us how to construct the polynomial; but it's possible to come up with a procedure to write down an equation as well.

Example 25.7
Let $\alpha = \sqrt{2}$ and $\beta = \sqrt{3}$, and $\gamma = \alpha + \beta$. How can we write down a polynomial equation for γ?

One possible method is that by Corollary 25.5, we know that $1, \gamma, \gamma^2, \ldots$ must be linearly dependent. In this case, they are all linear combinations of $1, \sqrt{2}, \sqrt{3}$, and $\sqrt{6}$ with coefficients in $\mathbb{Q}$ — so they lie in a vector space of dimension at most 4. Then $1, \gamma, \ldots, \gamma^4$ are five elements in a four-dimensional vector space, so they must be linearly dependent; and using linear algebra, it's possible to explicitly calculate this linear relation.

There is another way to find the polynomial equation — right now we'll present it as a guess, but later we'll see that it's part of a more general story.

We'd like to find a polynomial P with γ as a root, so we can try to think about what the other roots of P should be. Suppose P factors as $(x - \gamma_1)(x - \gamma_2)\cdots$, for $\gamma_i \in \mathbb{C}$ — it suffices to choose the γ_i such that P has rational coefficients, and $\gamma_1 = \sqrt{2} + \sqrt{3}$.

We can guess that all of $\pm\sqrt{2} \pm \sqrt{3}$ should be roots — from an algebraic perspective, if $\sqrt{2} + \sqrt{3}$ shows up, we "should" be able to switch the sign of the square root (since there isn't a difference between the two signs). So then we can take

$$\begin{aligned}
\gamma_1 &= \sqrt{2} + \sqrt{3} \\
\gamma_2 &= \sqrt{2} - \sqrt{3} \\
\gamma_3 &= -\sqrt{2} + \sqrt{3} \\
\gamma_4 &= -\sqrt{2} - \sqrt{3}.
\end{aligned}$$

We can expand out the polynomial to see that it does indeed have rational coefficients (essentially, this involves using the equation $a^2 - b^2 = (a - b)(a + b)$ twice).

The main idea we used here is to exploit the symmetry between the roots (there is a group of symmetries acting on the roots, by replacing one of the square roots with its negative); we'll later discuss ways to find these symmetries, using Galois theory.

25.3 Compass and Straightedge Construction

Proposition 25.4 also relates to compass and straightedge constructions. It has the following corollary:

Corollary 25.8
If $E/F/K$ is a tower of finite extensions, then $[F : K] \mid [E : K]$.

The problem of which regular n-gons can be constructed using a compass and straightedge can be rephrased algebraically in the following way (we won't discuss the details here).

Fact 25.9
A regular n-gon is constructible with compass and straightedge if and only if $\zeta_n = e^{2\pi i/n}$ lies in an extension $\mathbb{Q}(\alpha_1, \alpha_n)$ such that $\alpha_i^2 \in \mathbb{Q}(\alpha_1, \ldots, \alpha_{n-1})$ for all i.

This means we have a tower of quadratic extensions, where every step in this tower has degree 2 — more explicitly, we can define $F_i = \mathbb{Q}(\alpha_1, \ldots, \alpha_i)$, with $F_0 = \mathbb{Q}$. Without loss of generality we can assume $\alpha_i \notin F_{i-1}$ (or else adding it to the set of generators would be useless). Then we have the tower of extensions $F_n/F_{n-1}/\cdots/F_1/F_0$ where $[F_i : F_{i-1}] = 2$ for all i.

For convenience, we'll assume n is prime. (The general case involves a few more details, but works very similarly.)

Theorem 25.10

Let $n = p$ be prime. Then a regular p-gon can be constructed if and only if $p = 2^k + 1$.

Primes $p = 2^k + 1$ are called **Fermat primes**. There's only 5 known Fermat primes (3, 17, 257, and 65537); it's conjectured that there are no others, but we don't even know whether there's finitely or infinitely many. (Note that if $2^k + 1$ is prime, then k must be a power of 2 — otherwise, $2^k + 1$ can be factored.)

We'll only show one direction: that if ζ_p is constructible, then p is a Fermat prime. To prove this, the following proposition will be useful:

Proposition 25.11

If p is prime, we have $\deg(\zeta_n) = p - 1$, or equivalently $[\mathbb{Q}(\zeta_p) : \mathbb{Q}] = p - 1$.

The extension $\mathbb{Q}(\zeta_p)$ is called a **cyclotomic extension**.

Proof. We know ζ_p is a root of $x^p - 1$. We can easily factor

$$x^p - 1 = (x - 1)(x^{p-1} + x^{p-2} + \cdots + 1),$$

so it suffices to show that the second factor, which we call $P(x)$, is irreducible. Since the polynomial is primitive, it's enough to show that it's irreducible over $\mathbb{Z}$.

Now we can perform a trick — substitute $t = x - 1$. Then if we write $P(x) = Q(t)$, we have

$$tQ(t) = (t + 1)^p - 1.$$

But by expanding and using the Binomial Theorem, we then have

$$Q(t) = \sum_{i=0}^{p-1} \binom{p}{i+1} t^i.$$

(For example, when $p = 3$, we have $Q(t) = t^2 + 3t + 3$.)

But the leading term is 1, and all other terms are divisible by p; and the free term is not divisible by p^2 (in fact, *none* of the terms are divisible by p^2, but we only need to use the free term here).

Now assume for contradiction that Q is reducible, so $Q = Q_1Q_2$ for polynomials Q_1 and Q_2 of degree at least 1. Now consider the reduction mod p, where

$$\overline{Q} = \overline{Q_1Q_2}.$$

But $\overline{Q}$ is now t^{p-1}, and the only way to factor t^{p-1} in $\mathbb{F}_p[x]$ is as $t^i t^{p-1-i}$. But we have $\deg(\overline{Q_1}) = \deg(Q_1) > 1$ (and the same is true for Q_2), since the leading coefficients of Q_1 and Q_2 cannot be divisible by p (their product is the leading coefficient of Q, which is 1). So then we must have $i \neq 0, p - 1$.

But then since Q_1 and Q_2 are t^i and t^{p-1-i} for $0 < i < p - 1$, their free terms must both be divisible by p. So the product of their free terms is divisible by p^2; but this product is the free term of Q, which is *not* divisible by p^2. So this is a contradiction, and Q is irreducible. □

Proof of Necessity in Theorem 25.10. We've seen that $\deg(\zeta_p) = p - 1$. So we have $\deg(\zeta_p) = p - 1$. On the other hand, if $\zeta_p \in F_n$ for a field extension of the form described, then $\deg(\zeta_p)$ must divide $[F_n : \mathbb{Q}]$, which is a power of 2. So $p - 1$ must be a power of 2 as well. □

With our current tools, we can only show one direction — to show the other direction, we need a better extension of which fields can be obtained as the top floor of a tower of quadratic extensions. It's necessary that the degree is a power of 2, but this may not be sufficient. In the case of $\mathbb{Q}(\zeta_p)$, the condition turns out to be sufficient as well (as we'll see later).

25.4 Splitting Fields

We've seen the construction where we start with an irreducible polynomial $P \in F[x]$, and construct the field extension $E = F[x]/(P)$. This is an extension of F of degree $n = \deg(P)$, and we can think of it as adjoining a root of the polynomial.

But there's another construction which also produces a finite extension from a polynomial, which is in some sense harder to control. Here, we do not require the polynomial to be irreducible.

Definition 25.12
For a polynomial $P \in F[x]$, a **splitting field** of P is an extension E/F such that:

1. P splits as a product of linear factors in $E[x]$;
2. $E = F(\alpha_1, \ldots, \alpha_n)$, where the α_i are the roots of P.

The first condition guarantees that P splits completely (so we can find all its roots) in E; the second prevents E from being too large (it only contains the elements which are necessary for P to split).

Proposition 25.13
Given any polynomial P, its splitting field exists, and any two splitting fields of P are isomorphic.

We'll discuss the proof in more detail next time — the main idea is to add one root of P so that it splits partially, then add another root of any remaining irreducible factor, and so on.

Example 25.14
The splitting field of $P(x) = x^3 - 2$ over $F = \mathbb{Q}$ is $E = \mathbb{Q}(\sqrt[3]{2}, \omega\sqrt[3]{2})$, where ω is a primitive 3rd root of unity. We have $[E : \mathbb{Q}] = 6$.

On the other hand, we could start by adjoining ω:

Example 25.15
The splitting field of $P(x) = x^3 - 2$ over $F = \mathbb{Q}(\omega)$ is $E = F(\sqrt[3]{2})$ — the polynomial $x^3 - 2$ remains irreducible, but after adjoining one root, we already have all the others. Here $[E : F] = 3$.

Note that E is the same in both examples (even though F is not).

Example 25.16
The splitting field of $P(x) = x^{p-1} + \cdots + 1$ over $F = \mathbb{Q}$ is $E = \mathbb{Q}(\zeta_p)$ (since all roots are powers of ζ_p), where $[E : F] = p - 1$.

Resource: Algebra II Student Notes
Spring 2022
Instructor: Roman Bezrukavnikov
Notes taken by Sanjana Das and Jakin Ng