FEBRUARY 13 NOTES

1. 13.4: Splitting Fields and Algebraic Closures

Let F be a field and $f(x) \in F[x]$ a polynomial. Then, we know there exists an extension K of F in which $f(x)$ has a root α. (We proved the existence of K for f irreducible, but we can apply that construction to an irreducible factor of any polynomial.)

If $f(x)$ has a root α in a field K, this implies that $f(x)$ factors as $(x-\alpha)g(x)$ in $K[x]$. (Proof: divide $f(x)$ by $x-a$ using polynomial long division and apply the fact that α is a root of $f(x)$ and $x-a$.)

Fields in which polynomials factor completely into linear factors have a special name:

Definition 1.1. Let $f(x) \in F[x]$ be a polynomial. A **splitting field** of $f(x) \in F[x]$ is an extension K/F in which $f(x)$ factors completely into linear factors ('splits completely') in $K[x]$ but does not factor completely over any proper subfield of K containing F.

A preliminary exercise/fact about polynomials:

Exercise 1.2. If $f(x) \in F[x]$ is a polynomial of degree n, then f has at most n roots in F. (Hint: use induction on degree of f.) It has exactly n roots in F (counting multiplicities) if and only if $f(x)$ splits completely over F.

Theorem 1.3. *If F is a field and $f(x) \in F[x]$, a splitting field for $f(x)$ exists.*

Proof. We use induction on the degree of f to show F admits an extension E in which $f(x)$ splits completely. If $f(x)$ has degree 1, then $E = F$ and this is clear. If $n = \deg f > 1$ and all irreducible factors of $f(x)$ have degree 1, then still $E = F$. So, assume $n > 1$ and at least one irreducible factor $p(x)$ of $f(x)$ has degree at least 2. Then, there exists an extension E_1 of F in which $p(x)$ has a root, i.e. $p(x)$ is divisible by $(x-\alpha)$ in $E_1[x]$ for some root α of $p(x)$. Then, $f(x) = (x-\alpha)f_1(x)$ and $\deg f_1 = n-1$ so by induction there is an extension E of E_1 (and hence of F) in which f_1 (and hence f) splits completely. Therefore, an extension E of F containing all roots $\{\alpha_i\}$ of $f(x)$ exists. Let $K = F(\{\alpha_i\}) \subset E$. By definition, K is the smallest subfield that contains F and all roots of $f(x)$ so K is the splitting field of F. $\square$

Definition 1.4. If K is an algebraic extension of F such that K is the splitting field of some collection of polynomials $f(x) \in F[x]$, then K is a **normal extension** of F.

Example 1.5. The splitting field for x^2-2 over $\mathbb{Q}$ is $\mathbb{Q}(\sqrt{2})$, since the roots of x^2-2 are just $\pm\sqrt{2}$, which are both contained in $\mathbb{Q}(\sqrt{2})$.

Example 1.6. The splitting field for $(x^2-2)(x^2-3)$ is $\mathbb{Q}(\sqrt{2},\sqrt{3})$ since it contains all roots $\pm\sqrt{2}, \pm\sqrt{3}$. We have several intermediate subfields between $\mathbb{Q}$ and $\mathbb{Q}(\sqrt{2},\sqrt{3})$, namely $\mathbb{Q}(\sqrt{2}), \mathbb{Q}(\sqrt{3})$, and $\mathbb{Q}(\sqrt{6})$.

Example 1.7. The splitting field of x^3-2 over $\mathbb{Q}$ is not $\mathbb{Q}(\sqrt[3]{2})$: the other two roots of x^3-2 are complex numbers that are not contained in $\mathbb{Q}(\sqrt[3]{2})$ (because it is contained in $\mathbb{R}$). We can compute the other roots, which are $\frac{\sqrt[3]{2}}{2}(-1+\sqrt{-3})$ and $\frac{\sqrt[3]{2}}{2}(-1-\sqrt{-3})$.

To get the splitting field K, we must adjoin all three roots $\alpha_1, \alpha_2, \alpha_3$ to $\mathbb{Q}$. But, note that this field contains $2\alpha_3/\alpha_1 + 1 = \sqrt{-3}$, so K contains $\sqrt{-3}$. Because we can write any of the roots as a combination of $\sqrt[3]{2}$ and $\sqrt{-3}$, this implies that $K = \mathbb{Q}(\sqrt[3]{2}, \sqrt{-3})$ is the splitting field.

Note also that $[\mathbb{Q}(\sqrt[3]{2}) : \mathbb{Q}] = 3$, and $[\mathbb{Q}(\sqrt[3]{2}, \sqrt{-3}) : \mathbb{Q}(\sqrt[3]{2})] > 1$ because $\mathbb{Q}(\sqrt[3]{2})$ was *not* the splitting field. Because $\sqrt{-3}$ satisfies the polynomial $x^2 + 3 = 0$ over $\mathbb{Q}(\sqrt[3]{2})$, we have that $[\mathbb{Q}(\sqrt[3]{2}, \sqrt{-3}) : \mathbb{Q}(\sqrt[3]{2})] = 2$, so together this implies $[\mathbb{Q}(\sqrt[3]{2}, \sqrt{-3}) : \mathbb{Q}] = 6$. In particular, the degree of the splitting field is strictly greater than the degree of the polynomial.

In general, if $f(x) \in F[x]$ has degree n, then adjoining one root of $f(x)$ to F generates an extension of degree at most n, and then $f(x)$ has at most one linear factor, so adjoining the next root generates an extension of degree at most $n-1$, etc, to conclude that:

Proposition 1.8. *A splitting field of a polynomial of degree n over F has degree at most $n!$.*

Splitting fields of degree n polynomials can have degree anywhere between 1 and $n!$.

Example 1.9. One example of splitting field that we understand relatively well is that of *cyclotomic polynomials.* The polynomial $x^n - 1 \in \mathbb{Q}[x]$ is called a **cyclotomic polynomial** and the roots of this polynomial are called the **nth roots of unity**. We can express all of the roots as

$$\zeta_n^k := e^{2\pi k i/n} = \cos(2\pi k/n) + i\sin(2\pi k/n), \quad k = 0, 1, \ldots, n-1.$$

Geometrically, these are n equally spaced points around the unit circle starting with $(1, 0)$.

Because we have found n roots of $x^n - 1$, these must be all of the roots of this polynomial and, because they are all contained in $\mathbb{C}$, there exists a splitting field for $x^n - 1$ over $\mathbb{Q}$ contained in $\mathbb{C}$, namely $K = \mathbb{Q}(\zeta_n, \ldots, \zeta_n^{n-1})$.

From group theory, recall that the nth roots of unity form a **group** under multiplication, isomorphic to the cyclic group $\mathbb{Z}_n$. A generator of this group is called a **primitive** nth root of unity, and by abuse of notation, is denoted by ζ_n. By the isomorphism $\mathbb{Z}_n$ with the roots of unity given by $k \mapsto \zeta_n^k$, we see that the generators for this group are precisely ζ_n^a where $\gcd(a, n) = 1$, i.e. there are $\phi(n)$ different primitive nth roots of unity, where $\phi(n)$ is the Euler ϕ-function.

Because every root of unity is obtained as a power of ζ_n for a primitive nth root, the field $\mathbb{Q}(\zeta_n)$ contains all nth roots of unity and is therefore the splitting field of $x^n - 1$. This field is called the **cyclotomic field**.

The degree of the cyclotomic field over $\mathbb{Q}$ will turn out to be $\phi(n)$, which we will prove later. This can be nontrivial to compute; for example, the polynomial $x^n - 1$ is reducible for all $n > 1$ (because $x - 1$ is a factor). As a preliminary exercise to try: if p is prime, prove that the minimal polynomial of ζ_p is $x^{p-1} + x^{p-2} + \cdots + x + 1$, so $[\mathbb{Q}(\zeta_n) : \mathbb{Q}] = p - 1$.

Example 1.10. Let p be prime and consider the polynomial $x^p - 2$. We find that the roots are precisely $\zeta_p^k \sqrt[p]{2}$, where ζ_p is a primitive pth root of unity. The splitting field therefore contains $\sqrt[p]{2}$ (corresponding to $k = 0$, and ζ_p because it is the ratio of any two roots. Therefore, the splitting field contains $\mathbb{Q}(\sqrt[p]{2}, \zeta_p)$, but every root of $x^p - 2$ lies in this field, so the splitting field must be exactly $\mathbb{Q}(\sqrt[p]{2}, \zeta_p)$. This field contains two subfields, $\mathbb{Q}(\sqrt[p]{2})$ (which has degree p over $\mathbb{Q}$) and $\mathbb{Q}(\zeta_p)$ (which has degree $p-1$ over $\mathbb{Q}$). So, $\mathbb{Q}(\sqrt[p]{2}, \zeta_p)$ has degree at most $p(p-1)$ but the degree is divisible by both p and $p-1$ which are relatively prime, so we have $[\mathbb{Q}(\sqrt[p]{2}, \zeta_p) : \mathbb{Q}] = p(p-1)$. In particular, $x^p - 2$ must be *irreducible* over $\mathbb{Q}(\zeta_p)$ because this implies that the degree $[\mathbb{Q}(\sqrt[p]{2}, \zeta_p) : \mathbb{Q}(\zeta_p)] = p$. This is highly non-obvious!

Back to some theorems:

Theorem 1.11. *Let $\phi : F \to F'$ be an isomorphism and let $f(x) \in F[x]$. Define $f'(x)$ to be the polynomial obtained by applying ϕ to the coefficients of f. If E is a splitting field for $f(x)$ and E' a splitting field for f', then $\phi : F \to F'$ extends to an isomorphism $\Phi : E \to E'$.*

Proof. (Sketch) Use induction on the degree, noting that we already prove this for adding one root of a polynomial to F and F'. We omit the full details. □

By applying the theorem to the identity map $F \to F$, we see that:

Corollary 1.12. Any two splitting fields for $f(x) \in F[x]$ are isomorphic.

We therefore usually refer to a splitting field as *the* splitting field.

To end this section, we will construct the *algebraic closure* of a field F, a field in which *every polynomial* in $F[x]$ factors completely.

Definition 1.13. A field $\overline{F}$ is called an **algebraic closure** of F if $\overline{F}$ is algebraic over F and every polynomial $f(x) \in F[x]$ splits completely over $\overline{F}$. In other words, $\overline{F}$ contains all of the elements algebraic over F.

A field K is **algebraically closed** if every polynomial with coefficients in K has a root in K. Note that this implies that every polynomial splits completely in K.

Proposition 1.14. *Let $\overline{F}$ be an algebraic closure of F. Then, $\overline{F}$ is algebraically closed.*

Proof. If $f(x) \in \overline{F}[x]$ is a polynomial and α a root of $f(x)$, then $\overline{F}(\alpha)$ is an algebraic extension of $\overline{F}$, but $\overline{F}$ is algebraic over F, so $\overline{F}(\alpha)$ is algebraic over F and hence α is algebraic over F. This implies that $\alpha \in \overline{F}$ by definition of $\overline{F}$, so $\overline{F}$ is algebraically closed. □

We need to show that algebraic closures *exist*. Idea: keep adding roots of polynomials $f(x)$ and then $\overline{F}$ should be the field 'generated' over F by all of the roots. In order for this to make sense, we need there to be *some* field where all of the roots lie (this is needed to say 'generated' over F). We will therefore first construct some huge algebraically closed field containing F, and then look at the appropriate subfield to get the algebraic closure.

Proposition 1.15. *For any field F, there exists an algebraically closed field K containing F.*

Proof. For every nonconstant monic polynomial $f(x) \in F[x]$, let x_f represent a variable. Consider the polynomial ring $R = F[\dots x_f \dots]$ generated over F by all of these variables. Because f is a polynomial, we can plug in the variable x_f. Let I be the ideal generated by all of the polynomials of the form $f(x_f)$.

Claim: I is proper (to be verified next time.

Because I is proper, it is contained in some maximal ideal M. Therefore, the quotient $K_1 = R/M$ is a field containing F and each polynomial f in $F[x]$ has a root in K_1 by construction (we quotiented by an ideal containing $f(x_f)$, so the root is the image of x_f). Now, we want to repeat this procedure to produce a field K_2 in which every polynomial with coefficients in K_1 has a root, and so on, to obtain a sequence of fields

$$F = K_0 \subset K_1 \subset K_2 \subset \dots$$

where each polynomial in K_j has a root in K_{j+1}. We will finish this proof next time! □

For example, something we will prove in the future is:

Theorem 1.16. *The field $\mathbb{C}$ is algebraically closed. The field $\mathbb{C}$ is the algebraic closure of $\mathbb{R}$.*

By the proof of the previous proposition, because $\mathbb{Q} \subset \mathbb{C}$, the algebraic closure of $\mathbb{Q}$ therefore exists and is contained in $\mathbb{C}$. So, whenever we do computations with algebraic elements over $\mathbb{Q}$, we may assume that everything is happening in $\mathbb{C}$.