Field Theory, Part 2: Splitting Fields; Algebraic Closure

Jay Havaldar

2.1 Splitting Fields

As we saw, if $f(x)$ is any polynomial in $F[x]$, then there exists an extension K of F in which $f(x)$ has a root α. Equivalently, f has a factor $x - \alpha$ in $K[x]$. This motivates the following definition.

Definition: An extension K of F is called a **splitting field** for the polynomial $f(x) \in F[x]$ if f factors completely into linear factors in $K[x]$, but does not factor completely over any proper subfield of K containing F.

Theorem 1

For any field F and $f(x) \in F[x]$, there exists a splitting field for $f(x)$.

We proceed by induction on the degree n of f. If $n = 1$, $E = F$ is a splitting field.

If not, then f either splits completely ($E = F$ again), or else it has a reducible factor $p(x)$ of degree at least 2. In Part 1, we showed there is an extension E_1 of F containing a root α of $p(x)$. Thus, in E_1, $f(x)$ has an irreducible factor of degree at most $n - 1$.

By induction, there exists an extension E of E_1 containing all the roots of $f(x)$ other than α. Since $\alpha \in E_1 \subseteq E$, E is an extension of F in which $f(x)$ splits completely. Now, let K be the intersection of all subfields of E which contain F and also all the roots of $f(x)$; K is the splitting field.

We use the terminology ``the'' splitting field; indeed we shall show it is unique.

Definition: If K is an algebraic extension of F which is a splitting field over F for a collection of polynomials in $F[x]$, then K is called a **normal extension** of F.

``Splitting field'' and ``normal extension'' are used more or less interchangeably.

Proposition 1

A splitting field of a polynomial of degree n over F is of degree at most $n!$ over F

We can adjoin one root of $f(x)$ to generate an extension of degree at most n (equal iff f is irreducible). Over this field, $f(x)$ has at least one linear factor. Thus, adjoining another root yields an extension of degree at most $n - 1$. By the multiplicativity of extension degrees, the result follows.

Example: Cyclotomic Fields An important example that will be studied later is that of a **cyclotomic field**. We consider the splitting field of the polynomial:

$$x^n - 1$$

Over $\mathbb{Q}$. The roots are called the nth roots of unity. With multiplication, they form a cyclic group; indeed this group is precisely $\mathbb{Z}/n\mathbb{Z}$.

Definition: A **primitive** nth root of unity is a generator for the cyclic group of all nth roots of unity

We use ζ_n to denote a primitive nth root of unity. Evidently, ζ_n^a is also a primitive root, if a is relatively prime to n.

Since ζ_n generates this entire group, $x^n - 1$ splits completely over the field $\mathbb{Q}(\zeta_n)$.

Definition: The field $\mathbb{Q}(\zeta_n)$ is called the **cyclotomic field** of nth roots of unity.

A particular case is when $n = p$ is prime. Then the factorizatino is given by:

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

And the latter term (denoted Φ_p) is irreducible, which follows from substituting in $(x+1)$ for x and using Eisenstein's Criterion. Thus, we have that Φ_p is the minimal polynomial of ζ_p over the rationals, so that:

$$[\mathbb{Q}(\zeta_p) : \mathbb{Q}] = p - 1$$

Next, we show that indeed the splitting field is unique.

Theorem 2

> Let $\varphi : F \to F'$ be an isomorphism of fields. Let $f(x) \in F[x]$ and denote its image under φ as $f'(x)$ (obtained by applying φ to the coefficients).
>
> Let E be a splitting field for $f(x)$ over F; let E' be a splitting field for $f'(x)$ over F'. Then we can extend φ to an isomorphism $\sigma : E \to E'$.

We proceed by induction on n, the degree of $f(x)$. If $f(x)$ splits completely in F, then $f'(x)$ splits completely in F', and we are done.

Now, assume that $p(x)$ is an irreducible factor of $f(x)$ which has degree at least 2. Let $p'(x)$ be the image in $F'(x)$. Then if $\alpha \in E$ is a root of $p(x)$ and $\beta \in E'$ is a root of $p'(x)$, then by Part 1 Theorem 3, we can extend φ to an isomorphism σ':

$$\sigma' : F(\alpha) \to F'(\beta) \sigma' : \alpha \mapsto \beta$$

Denote $F_1 = F(\alpha)$ and $F_1' = F'(\beta)$; we have just constructed an isomorphism of fields $\sigma' : F_1 \to F_1'$. Over F_1, we can write:

$$f(x) = (x - \alpha) f_1(x) f'(x) = (x - \beta) f_1'(x)$$

Where each of f_1, f_1' has degree $n - 1$.

Notice that E is a splitting field for f_1 over F_1; if f_1 splits in any subfield, then we have found a subfield of E in which $f(x)$ splits. Similarly, E' is a splitting for f_1' over F_1'.

Thus, since f_1, f_1' have degree $n - 1$, by induction there is a map $\sigma : E \to E'$ which extends the isomorphism $\sigma' : F_1 \to F_1'$.

Thus, we have shown that σ extends σ' which in turn extends φ; thus we have extended an isomorphism of fields to an isomorphism of splitting fields.

Corollary

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

Take the proof above and let φ be the identity from F to itself.

Now that we have looked at field extensions of F which contains the root of a particular polynomial of degree n over F (which necessarily exist and have degree at most $n!$), we ask the question of whether there is an extension of F which contains the roots of all polynomials over F. As you might expect, there are going to be some Zorn's Lemma shenanigans.

Definition: $\overline{F}$ is called an algebraic closure of F if $\overline{F}$ is algebraic over F and if every polynomial $f(x) \in F[x]$ splits completely over $\overline{F}$.

Thus, in a way $\overline{F}$ contains all the elements which are algebraic over F.

Definition: A field K is said to be algebraically closed if every polynomials with coefficients in K has a root in K.

We'd better hope that algebraically closed fields exist (the complex numbers should be one). We also should hope that algebraic closures exist for arbitrary fields. Finally, one should expect that an algebraic closure is itself algebraically closed.

Proposition 2

Let $\overline{F}$ be an alegbraic closure of F. Then $\overline{F}$ is algebraically closed.

Let $f(x)$ be a polynomial in $\overline{F}[x]$ with a root α. Then α generates an algebraic extension $\overline{F}(\alpha)$ of F. However, since $\overline{F}(\alpha)/\overline{F}$ is algebraic, and $\overline{F}/F$ is algebraic, $\overline{F}(\alpha)/F$ is algebraic. But then $\alpha \in \overline{F}$, since α is algebraic over F, so that $\overline{F}$ is algebraically closed.

Proposition 3

For a field F there exists an algebraically closed field K which contains F.

This proof is not too enlightening; continue at your own risk.

For each non-constant monic polynomial f with coefficients in F, we associate an indeterminate x_f. Let S denote the set of all such indeterminates; it is in bijection with the set of polynomials in $F[x]$ with degree at least 1.

$$F[S] = F[\dots, x_f, \dots]$$

In this polynomial ring, consider the ideal I generated by the polynomials $f(x_f)$. We claim that this is a proper ideal.

If the ideal is not proper, then in particular 1 is an element. So there exists a finite linear combination:

$$g_1 f_1(x_{f_1}) + \cdots + g_n f_n(x_{f_n}) = 1$$

Where each $g_i \in F[S]$. For convenience, we denote x_i instead of x_{f_i}. Finally, let $x_{n+1}, \dots x_m$ denote the remaining letters occuring in the polynomials g_i. We can rewrite the above relation:

$$g_1(x_1, \dots x_m)f_1(x_1) + \cdots + g_n(x_1, \dots, x_m)f_n(x_n) = 1$$

Now, let F' be a (finite) extension of F which contains a root α_i of each $f_i(x)$. Then if we let $x_i = \alpha_i$ and set $x_{n+1} = \cdots = x_m = 0$, then the equation above reads $0 = 1$, which is impossible. Thus, the ideal above must be proper.

Since I is a proper ideal, it is contained in some maximal ideal $\mathcal{M}$ (Zorn's Lemma appears here). Then the quotient:

$$K_1 = \frac{F[S]}{\mathcal{M}}$$

Is a field which contains F. Furthermore, each of the polynomials f has a root in K_1 by construction, since $f(x_f) \in I$ and therefore the image of x_f is a root.

We repeat this construction with K_1 to obtain a field K_2 in which all polynomials with coefficients in K_1 has a root. In this way, we get a sequence of fields:

$$F = K_0 \subseteq K_1 \subseteq \cdots \subseteq K_j \subseteq K_{j+1} \subseteq \dots$$

And each polynomial with coefficients in K_j has a root in K_{j+1}. Now, denote:

$$K = \bigcup_{j \geq 0} K_j$$

K is a field which contains F and the cofficients of any polynomial in K lie in some field K_N and thus in K; thus the polynomial has a root in $K_{N+1} \subset K$. Thus, K is algebraically closed.

Proposition 4

> Let K be an algebraically closed field and F a subfield of K. Then the collection of elements $\overline{F}$ of K that are algebraic over F is an algebraic closure of F. The algebraic closure is unique up to isomorphism.

By definition, $\overline{F}$ is an algebraic extension of F. Furthermore, K contains all the roots of polynomials with coefficients in F (indeed, with coefficients in K); so, in $\overline{F}[x]$ every polynomial in $F[x]$ splits completely; thus, $\overline{F}$ is an algebraic closure of F.

Thus, if we can locate a field F as a subfield of an algebraically closed field, then we create an algebraic closure $\overline{F}$ by collecting all alements of K which are algebraic over F.

The uniqueness follows from the fact that the splitting field is unique up to isomorphism (and Zorn's Lemma is involved, as you might expect).

Theorem 3 (Fundamental Theorem of Algebra)

> The field $\mathbb{C}$ is algebraically closed.

This theorem will be proven later using Galois theory.

Corollary

> $\mathbb{C}$ contains an algebraic closure for any of its subfields. In particular, $\overline{\mathbb{Q}}$, the collection of complex numbers which are algebraic over $\mathbb{Q}$, is an algebraic closure of $\mathbb{Q}$.

From the above theorem, we can think of any discussion of F as taking place in the context of the (generally larger) field $\overline{F}$. A composite of any collection of algebraic extensions can be viewed as subfields of an algebraic closure. For example, in $\mathbb{Q}$, all of the computation is ``really'' happening in $\mathbb{C}$.

2.2 Separable Extensions

Let F be a field and $f(x)$ a polynomial. Over a splitting field we can write:

$$f(x) = (x - \alpha_1)^{n_1} \dots (x - \alpha_k)^{n_k}$$

With α_i all distinct roots.

Definition: A polynomial over F is called separable if its roots are all distinct; otherwise, a polynomial is inseparable.

Note that since splitting fields are isomorphic, with an isomorphism that is bijective on the roots, separability is in a sense an ``instrinsic'' property of polynomials irrespective of the splitting field.

Definition: The **formal derivative** of the polynomial:

$$f(x) = a_n x^n + \dots + a_1 x + a_0 Df(x) = n a_n x^{n-1} + \dots + a_1$$

We're not actually taking derivatives here; there are no limits.

Proposition 5

> A polynomial $f(x)$ has a multiple root α iff α is also a root of $Df(x)$. In particular, this means that $f(x)$ and $Df(x)$ are both divisible by the minimal polynomial for α.

Suppose that α is a multiple root. Then over a splitting field:

$$f(x) = (x - \alpha)^n g(x)$$

Taking the derivative:

$$Df(x) = n(x - \alpha)^{n-1} g(x) + (x - \alpha)^n Dg(x)$$

And so there is a common factor of $(x - \alpha)$ as desired.

This tells us that f is separable iff $(f, Df) = 1$.

Conversely, suppose that α is a root of both f and Df. Then we can write:

$$f(x) = (x - \alpha)h(x)$$

Taking the derivative yields:

$$Df(x) = h(x) + (x - \alpha)Dh(x)$$

By assumption, $Df(\alpha) = 0$; thus, $h(\alpha) = 0$, and we are done.

Note that the above proof holds over arbitrary characteristic.

Corollary

> Over a field of characteristic 0, a polynomial is separable iff it is the product of distinct irreducibles. In particular, an irreducible polynomial is separable.

Suppose that $p(x)$ is irreducible with degree n. Then the derivative has degree $n - 1$, and must thus be relatively prime to $p(x)$. Thus, p is separable. Note also that distinct irreducibles do not have any zeroes in common (since if they did, both divide the minimal polynomial and one must be a factor of the other).

However, in characteristic p, the derivative could simply have degree 0. For example:

$$D(x^{pm}) = pmx^{pm-1} = 0$$

So, the above proof only works if we take an irreducible polynomial whose derivative is nonzero.

However, suppose the derivative of $p(x)$ is zero. Then every exponent (from the above discussion) must be a multiple of p, the characteristic of F. So:

$$p(x) = a_m x^{mp} + a_{m-1}x^{(m-1)p} + \cdots + a_1 x^p + a_0$$

So indeed $p(x)$ is a polynomial in x_p.

Proposition 6

> Let F be a field of characteristic p. Then:
>
> $$(a + b)^p = a^p + b^p (ab)^p = a^p b^p$$
>
> This follows from the binomial theorem. In particular, it tells us that $\varphi(a) = a^p$ is an injective field homomorphism from F to F.

Definition: The map $\varphi(a) = a^p$ for a field of characteristic p is called the **Frobenius endomorphism**.

Corollary

If $\mathbb{F}$ is a finite field of characteristic p, then the Frobenius map is an automorphism, i.e. $\mathbb{F} = \mathbb{F}^p$.

Now, we return to the problem of locating irreducible polynomials over fields of characteristic p.

Proposition 7

A polynomial over a finite field $\mathbb{F}$ is separable iff it is the product of distinct irreducibles. An irreducible polynomial is separable.

Let $\mathbb{F}$ be a finite field and $p(x) \in \mathbb{F}[x]$ is irreducible. If $p(x)$ is inseparable, then its derivative is zero, so from above we can write $p(x) = q(x^p)$ for some polynomial $q(x) \in \mathbb{F}[x]$. Then we let:

$$q(x) = a_m x^m + \cdots + a_1 x + a_0$$

By the Frobenius automorphism, we can denote $a_i = b_i^p$. Thus we can write:

$$\begin{aligned} p(x) = q(x^p) &= a_m (x^p)^m + \cdots + a_1 x^p + a_0 \\ &= b_m^p (x^p) m + \cdots + b_1^p x^p + b_0^p \\ &= (b_m x^m)^p + \cdots + (b_1 x)^p + b_0^p \\ &= (b_m x_m + b_{m-1} x^{m-1} + \cdots + b_1 x + b_0)^p \end{aligned}$$

And thus shows that $p(x)$ is the pth power of a polynomial in $\mathbb{F}[x]$ and hence is not irreducible.

We generalize the concept of the Frobenius automorphism:

Definition: A field K of characteristic p is called perfect if $K = K^p$. Any field of characteristic 0 is called perfect.

As we showed, finite fields are perfect. With the above proof, we proved the more general statement that irreducible polynomials over perfect fields are separable. If K is not perfect, there are inseparable irreducible polynomials.

Example (Existence & Uniqueness of Finite Fields) The polynomial $x^{p^n} - x$ over $\mathbb{F}_{\text{l}}$ has derivative:

$$p^n x^{p^n - 1} - 1 = -1$$

Thus, the derivative has no roots at all. Therefore, this polynomial is separable.

Now, let $n > 0$ and consider the splitting field of the above polynomial. Since it is separable, it has p^n roots exactly. Now, let α, β be any two roots. Then:

$$(\alpha\beta)^{p^n} = \alpha^{p^n} \beta^{p^n} = \alpha\beta(\alpha + \beta)^{p^n} = \alpha^{p^n} + \beta^{p^n} = \alpha + \beta(\alpha^{-1})^{p^n} = \alpha^{-1}$$

Thus, the roots of this polynomial form a subfield of the splitting field; hence it must be the splitting field. Finally, since the number of elements is p^n, we must have:

$$[\mathbb{F} : \mathbb{F}_p] = n$$

This shows the existence of finite fields of size p^n. Now, we show uniqueness.

If $\mathbb{F}$ is a finite field of characteristic p, then we denote $n = [\mathbb{F} : \mathbb{F}_p]$, i.e. the degree over its prime subfield. Thus, $\mathbb{F}$ has exactly p^n elements. Since the multiplicative group has order $p^n - 1$, we must have:

$$\alpha^{p^n - 1} = 1$$

For any $\alpha \neq 0$ in $\mathbb{F}$. Therefore, $\alpha^{p^n} = \alpha$, and thus $\mathbb{F}$ is contained in the splitting field for $x^{p^n} - x$; by counting considerations, $\mathbb{F}$ is indeed this splitting field, which is unique up to isomorphism.

We saw that if $p(x)$ is irreducible over a field of characteristic p, then if it is not separable its derivative is zero, hence $p(x) = p_1(x^p)$ for some polynomial $p_1(x)$. Continuing this process, there is a unique power p^k so that:

$$p(x) = p_k(x^{p^k})$$

Where p_k has nonzero derivative.

Thus, if p is an irreducible polynomial over F with char. p, then there is a unique integer $k \geq 0$ and a unique irreducible separable polynomial $p_{\text{sep}}(x) \in F[x]$ such that:

$$p(x) = p_{\text{sep}}(x^{p^k})$$

Definition: Suppose p is an irreducible polynomial over F with char. p. Then we call teh degree of $p_{\text{sep}}(x)$ the separable degree of $p(x)$, and the integer p^k is denoted the inseparable degree of $p(x)$.

Definition: A field K is **separable** over F is every element of K is the root of some separable polynomial over F (or equivalently, if the minimal polynomial of every element is separable).

Corollary

> Every finite extension of a perfect field is separable. In particular, every finite extension of $\mathbb{Q}$ or a finite field is separable.

As we saw before, every finite extension is algebraic. Furthermore, every algebraic element can be realized as the root of some unique minimal polynomial which is irreducible. Finally, by the above discussion, irreducible polynomials over perfect fields are separable.

2.3 Cyclotomic Polynomials and Extensions

We now prove that the cyclotomic extension discussed earlier:

$$\mathbb{Q}(\zeta_n)/\mathbb{Q}$$

Has degree exactly $\varphi(n)$, where φ denotes the totient.

Definition: Let μ_n denote the group of nth roots of unity over $\mathbb{Q}$.

As seen before, this is nothing more than the cyclic group of size n.

If d is a divisor of n, then we have:

$$\mu_d \subseteq \mu_n$$

Conversely, every element of μ_n has an order which is a divisor of n, and thus, we can write:

Definition: Define the nth cyclotomic polynomial $\Phi_n(x)$ to be the polynomial whose roots are the primitive nth roots of unity. Then:

$$\Phi_n(x) = \prod_{\zeta \in \mu_n \text{ primitive}} (x - \zeta) = \prod_{(a,n)=1\ ,\ 1 \leq a < n} (x - \zeta_n^a)$$

The roots of the polynomial $x^n - 1$ gives us the factorization:

$$x^n - 1 = \prod_{d|n} \prod_{\zeta \in \mu_d \text{ primitive}} (x - \zeta)$$

But the inner product is exactly $\Phi_d(x)$ so we can write:

$$x^n - 1 = \prod_{d|n} \Phi_d(x)$$

And taking the degrees we get:

$$n = \sum_{d|n} \varphi(d)$$

Now, we can compute $\Phi_n(x)$ recursively by dividing out the prior cyclotomic polynomials.

Lemma

$\Phi_n(x)$ is a monic polynomial in $\mathbb{Z}[x]$ with degree $\varphi(n)$.

Theorem 4

The cyclotomic polynomial $\Phi_n(x)$ is an irreducible monic polynomial in $\mathbb{Z}[x]$ of degree $\varphi(n)$.

Corollary

$$[\mathbb{Q}(\zeta_n) : \mathbb{Q}] = \varphi(n)$$