20

Finite fields

This chapter develops some of the basic theory of finite fields. As we already know (see Theorem 9.7), every finite field must be of cardinality p^w, for some prime p and positive integer w. The main results of this chapter are:

- for any prime p and positive integer w, there exists a finite field of cardinality p^w, and
- any two finite fields of the same cardinality are isomorphic.

20.1 Preliminaries

In this section, we prove a few simple facts that will be useful in this and later chapters; also, for the reader's convenience, we recall a few basic algebraic concepts that were discussed in previous chapters, but which will play important roles in this chapter.

Theorem 20.1. *Let F be a field, and let k, ℓ be positive integers. Then $\mathtt{X}^k - 1$ divides $\mathtt{X}^\ell - 1$ if and only if k divides ℓ.*

Proof. Let $\ell = kq + r$, with $0 \leq r < k$. We have

$$\mathtt{X}^\ell \equiv \mathtt{X}^{kq}\mathtt{X}^r \equiv \mathtt{X}^r \pmod{\mathtt{X}^k - 1},$$

and $\mathtt{X}^r \equiv 1 \pmod{\mathtt{X}^k - 1}$ if and only if $r = 0$. □

Theorem 20.2. *Let $a \geq 2$ be an integer and let k, ℓ be positive integers. Then $a^k - 1$ divides $a^\ell - 1$ if and only if k divides ℓ.*

Proof. The proof is analogous to that of Theorem 20.1. We leave the details to the reader. □

One may combine these two theorems, obtaining:

Theorem 20.3. *Let $a \geq 2$ be an integer, k, ℓ be positive integers, and F a field. Then $\mathtt{X}^{a^k} - \mathtt{X}$ divides $\mathtt{X}^{a^\ell} - \mathtt{X}$ if and only if k divides ℓ.*

Proof. We have $\mathtt{X}^{a^k} - \mathtt{X}$ divides $\mathtt{X}^{a^\ell} - \mathtt{X}$ iff $\mathtt{X}^{a^k-1} - 1$ divides $\mathtt{X}^{a^\ell-1} - 1$, and by Theorem 20.1, this happens iff $a^k - 1$ divides $a^\ell - 1$, which by Theorem 20.2 happens iff k divides ℓ. □

Let F be a field. A polynomial $f \in F[\mathtt{X}]$ is called **square-free** if it is not divisible by the square of any polynomial of degree greater than zero. Using formal derivatives, we obtain the following useful criterion for establishing that a polynomial is square-free:

Theorem 20.4. *If F is a field, and $f \in F[\mathtt{X}]$ with $\gcd(f, \mathbf{D}(f)) = 1$, then f is square-free.*

Proof. Suppose f is not square-free, and write $f = g^2 h$, for $g, h \in F[\mathtt{X}]$ with $\deg(g) > 0$. Taking formal derivatives, we have

$$\mathbf{D}(f) = 2g\mathbf{D}(g)h + g^2\mathbf{D}(h),$$

and so clearly, g is a common divisor of f and $\mathbf{D}(f)$. □

We end this section by recalling some concepts discussed earlier, mainly in §17.1, §17.5, and §17.6.

Suppose F is a field, and E is an extension field of F; that is, F is a subfield of E, or F is embedded in E via some canonical embedding, and we identify elements of F with their images in E under this embedding. We may naturally view E as an F-vector space. Assume that as an F-vector space, E has finite dimension $\ell > 0$. This dimension ℓ is called the degree of E over F, and is denoted $(E : F)$; moreover, E is called a finite extension of F.

We may also naturally view E as an F-algebra, either via the inclusion map or via some canonical embedding. Let E' be another field extension of F, and let $\rho : E \to E'$ be a ring homomorphism (which in fact, must be injective). Then ρ is an F-algebra homomorphism if and only if $\rho(a) = a$ for all $a \in F$.

For any $\alpha \in E$, the set $F[\alpha] = \{g(\alpha) : g \in F[\mathtt{X}]\}$ is a subfield of E containing F. Moreover, there exists a non-zero polynomial g of degree at most ℓ such that $g(\alpha) = 0$. The monic polynomial ϕ of least degree such that $\phi(\alpha) = 0$ is called the minimal polynomial of α over F, and this polynomial is irreducible over F. The field $F[\mathtt{X}]/(\phi)$ is isomorphic, as an F-algebra, to $F[\alpha]$, via the map that sends $[g]_\phi \in F[\mathtt{X}]/(\phi)$ to $g(\alpha) \in F[\alpha]$. We have $(F[\alpha] : F) = \deg(\phi)$, and this value is called the degree of α over F. If E' is

an extension field of F, and if $\rho : F[\alpha] \to E'$ is an F-algebra homomorphism, then the action of ρ is completely determined by its action on α; indeed, for any $g \in F[\mathtt{X}]$, we have $\rho(g(\alpha)) = g(\rho(\alpha))$.

20.2 The existence of finite fields

Let F be a finite field. As we saw in Theorem 9.7, F must have cardinality p^w, where p is prime and w is a positive integer, and p is the characteristic of F. However, we can say a bit more than this. As discussed in Example 9.41, the field $\mathbb{Z}_p$ is embedded in F, and so we may simply view $\mathbb{Z}_p$ as a subfield of F. Moreover, it must be the case that w is equal to $(F : \mathbb{Z}_p)$.

We want to show that there exist finite fields of every prime-power cardinality. Actually, we shall prove a more general result:

> *If F is a finite field, then for every integer $\ell \geq 1$, there exists an extension field E of degree ℓ over F.*

For the remainder of this section, F denotes a finite field of cardinality $q = p^w$, where p is prime and $w \geq 1$.

Suppose for the moment that E is an extension of degree ℓ over F. Let us derive some basic facts about E. First, observe that E has cardinality q^ℓ. By Theorem 9.16, E^* is cyclic, and the order of E^* is $q^\ell - 1$. If $\gamma \in E^*$ is a generator for E^*, then every non-zero element of E can be expressed as a power of γ; in particular, every element of E can be expressed as a polynomial in γ with coefficients in F; that is, $E = F[\gamma]$. Let $\phi \in F[\mathtt{X}]$ be the minimal polynomial of γ over F, which is an irreducible polynomial of degree ℓ. It follows that F is isomorphic (as an F-algebra) to $F[\mathtt{X}]/(\phi)$.

So we have shown that any extension of F of degree ℓ must be isomorphic, as an F-algebra, to $F[\mathtt{X}]/(\phi)$ for some irreducible polynomial $\phi \in F[\mathtt{X}]$ of degree ℓ. Conversely, given any irreducible polynomial ϕ over F of degree ℓ, we can construct the finite field $F[\mathtt{X}]/(\phi)$, which has degree ℓ over F. Thus, the question of the existence of a finite fields of degree ℓ over F reduces to the question of the existence of an irreducible polynomial over F of degree ℓ.

We begin with a simple generalization Fermat's little theorem:

Theorem 20.5. *For any $a \in F^*$, we have $a^{q-1} = 1$, and for any $a \in F$, we have $a^q = a$.*

Proof. The multiplicative group of units F^* of F has order $q-1$, and hence, every $a \in F^*$ satisfies the equation $a^{q-1} = 1$. Multiplying this equation by a yields $a^q = a$ for all $a \in F^*$, and this latter equation obviously holds for $a = 0$ as well. □

Theorem 20.6. *We have*

$$\mathtt{X}^q - \mathtt{X} = \prod_{a \in F} (\mathtt{X} - a).$$

Proof. The polynomial

$$(\mathtt{X}^q - \mathtt{X}) - \prod_{a \in F} (\mathtt{X} - a)$$

has degree less than q, but has q distinct roots (namely, every element of F), and hence must be the zero polynomial. □

The following theorem generalizes Example 17.6:

Theorem 20.7. *Let E be an F-algebra. Then the map $\rho : E \to E$ that sends $\alpha \in E$ to α^q is an F-algebra homomorphism.*

Proof. Recall that E being an F-algebra simply means that E is a ring and that there is a ring homomorphism $\tau : F \to E$, and because F is a field, either τ is injective or E is trivial. Also, recall that ρ being an F-algebra homomorphism simply means that ρ is a ring homomorphism and $\rho(\tau(a)) = \tau(a)$ for all $a \in F$.

Now, if E is trivial, there is nothing to prove. Otherwise, as E contains a copy of F, it must have characteristic p. Since q is a power of the characteristic, the fact that ρ is a ring homomorphism follows from the discussion in Example 9.42. Moreover, by Theorem 20.5, we have $\tau(a)^q = \tau(a^q) = \tau(a)$ for all $a \in F$. □

Theorem 20.8. *Let E be a finite extension of F, and consider the map $\sigma : E \to E$ that sends $\alpha \in E$ to $\alpha^q \in E$. Then σ is an F-algebra automorphism on E. Moreover, if $\alpha \in E$ is such that $\sigma(\alpha) = \alpha$, then $\alpha \in F$.*

Proof. The fact that σ is an F-algebra homomorphism follows from the previous theorem. Any ring homomorphism from a field into a field is injective (see Exercise 9.38). Surjectivity follows from injectivity and finiteness.

For the second statement, observe that $\sigma(\alpha) = \alpha$ if and only if α is a root of the polynomial $\mathtt{X}^q - \mathtt{X}$, and since all q elements of F are already roots of this polynomial, there can be no other roots. □

The map σ defined in Theorem 20.8 is called the **Frobenius map on E over F**. As it plays a fundamental role in the study of finite fields, let us develop a few simple properties right away.

Since the composition of two F-algebra automorphisms is also an F-algebra automorphism, for any $i \geq 0$, the i-fold composition σ^i that sends $\alpha \in E$ to α^{q^i} is also an F-algebra automorphism.

Since σ is an F-algebra automorphism, the inverse function σ^{-1} is also an F-algebra automorphism. Hence, σ^i is an F-algebra automorphism for all $i \in \mathbb{Z}$. If E has degree ℓ over F, then applying Theorem 20.5 to the field E, we see that σ^ℓ is the identity map, from which it follows that $\sigma^{-1} = \sigma^{\ell-1}$. More generally, we see that for any $i \in \mathbb{Z}$, we have $\sigma^i = \sigma^j$, where $j = i \bmod \ell$.

Thus, in considering integer powers of σ, we need only consider the powers $\sigma^0, \sigma^1, \ldots, \sigma^{\ell-1}$. Furthermore, the powers $\sigma^0, \sigma^1, \ldots, \sigma^{\ell-1}$ are all distinct maps. To see this, assume that $\sigma^i = \sigma^j$ for some i, j with $0 \leq i < j < \ell$. Then σ^{j-i} would be the identity map, which would imply that all of the q^ℓ elements of E were roots of the polynomial $\mathtt{X}^{q^{j-i}} - \mathtt{X}$, which is a non-zero polynomial of degree less that q^ℓ, and this yields a contradiction.

The following theorem generalizes Theorem 20.6:

Theorem 20.9. *For $k \geq 1$, let P_k denote the product of all the monic irreducible polynomials in $F[\mathtt{X}]$ of degree k. For all positive integers ℓ, we have*

$$\mathtt{X}^{q^\ell} - \mathtt{X} = \prod_{k \mid \ell} P_k,$$

where the product is over all positive divisors k of ℓ.

Proof. First, we claim that the polynomial $\mathtt{X}^{q^\ell} - \mathtt{X}$ is square-free. This follows immediately from Theorem 20.4, since $\mathbf{D}(\mathtt{X}^{q^\ell} - \mathtt{X}) = q^\ell \mathtt{X}^{q^\ell - 1} - 1 = -1$.

So we have reduced the proof to showing that if f is a monic irreducible polynomial of degree k, then f divides $\mathtt{X}^{q^\ell} - \mathtt{X}$ if and only if $k \mid \ell$. Let $E := F[\mathtt{X}]/(f)$, and let $\eta := [\mathtt{X}]_f \in E$, which is a root of f.

For the first implication, assume that f divides $\mathtt{X}^{q^\ell} - \mathtt{X}$. We want to show that $k \mid \ell$. Now, if $\mathtt{X}^{q^\ell} - \mathtt{X} = fg$, then $\eta^{q^\ell} - \eta = f(\eta)g(\eta) = 0$, so $\eta^{q^\ell} = \eta$. Therefore, if σ is the Frobenius map on E over F, then we have $\sigma^\ell(\eta) = \eta$. We claim that $\sigma^\ell(\alpha) = \alpha$ for all $\alpha \in E$. To see this, recall from Theorem 17.1 that for all $h \in F[\mathtt{X}]$ and $\beta \in E$, we have $\sigma^\ell(h(\beta)) = h(\sigma^\ell(\beta))$. Moreover, any $\alpha \in E$ can be expressed as $h(\eta)$ for some $h \in F[\mathtt{X}]$, and so

$$\sigma^\ell(\alpha) = \sigma^\ell(h(\eta)) = h(\sigma^\ell(\eta)) = h(\eta) = \alpha.$$

That proves the claim.

From the claim, it follows that every element of E is a root of $\mathtt{X}^{q^\ell} - \mathtt{X}$. That is, $\prod_{\alpha \in E}(\mathtt{X} - \alpha)$ divides $\mathtt{X}^{q^\ell} - \mathtt{X}$. Applying Theorem 20.6 to the field E, we see that $\prod_{\alpha \in E}(\mathtt{X} - \alpha) = \mathtt{X}^{q^k} - \mathtt{X}$, and hence $\mathtt{X}^{q^k} - \mathtt{X}$ divides $\mathtt{X}^{q^\ell} - \mathtt{X}$. By Theorem 20.3, this implies k divides ℓ.

For the second implication, suppose that $k \mid \ell$. We want to show that $f \mid \mathtt{X}^{q^\ell} - \mathtt{X}$. Since f is the minimal polynomial of η, and since η is a root of $\mathtt{X}^{q^k} - \mathtt{X}$, we must have that f divides $\mathtt{X}^{q^k} - \mathtt{X}$. Since $k \mid \ell$, and applying Theorem 20.3 once more, we see that $\mathtt{X}^{q^k} - \mathtt{X}$ divides $\mathtt{X}^{q^\ell} - \mathtt{X}$. That proves the second implication, and hence, the theorem. □

For $\ell \geq 1$, let $\Pi(\ell)$ denote the number of monic irreducible polynomials of degree ℓ in $F[\mathtt{X}]$.

Theorem 20.10. *For all $\ell \geq 1$, we have*

$$q^\ell = \sum_{k \mid \ell} k\Pi(k). \tag{20.1}$$

Proof. Just equate the degrees of both sides of the identity in Theorem 20.9. □

From Theorem 20.10 it is easy to deduce that $\Pi(\ell) > 0$ for all ℓ, and in fact, one can prove a density result—essentially a "prime number theorem" for polynomials over finite fields:

Theorem 20.11. *For all $\ell \geq 1$, we have*

$$\frac{q^\ell}{2\ell} \leq \Pi(\ell) \leq \frac{q^\ell}{\ell}, \tag{20.2}$$

and

$$\Pi(\ell) = \frac{q^\ell}{\ell} + O\left(\frac{q^{\ell/2}}{\ell}\right). \tag{20.3}$$

Proof. First, since all the terms in the sum on the right hand side of (20.1) are non-negative, and $\ell\Pi(\ell)$ is one of these terms, we may deduce that $\ell\Pi(\ell) \leq q^\ell$, which proves the second inequality in (20.2). Since this holds for all ℓ, we have

$$\ell\Pi(\ell) = q^\ell - \sum_{\substack{k \mid \ell \\ k < \ell}} k\Pi(k) \geq q^\ell - \sum_{\substack{k \mid \ell \\ k < \ell}} q^k \geq q^\ell - \sum_{k=1}^{\lfloor \ell/2 \rfloor} q^k.$$

Let us set

$$S(q, \ell) := \sum_{k=1}^{\lfloor \ell/2 \rfloor} q^k = \frac{q}{q-1}(q^{\lfloor \ell/2 \rfloor} - 1),$$

so that $\ell\Pi(\ell) \geq q^\ell - S(q, \ell)$. It is easy to see that $S(q, \ell) = O(q^{\ell/2})$, which proves (20.3). For the first inequality of (20.2), it suffices to show that

$S(q,\ell) \le q^\ell/2$. One can check this directly for $\ell \in \{1,2,3\}$ (verify), and for $\ell \ge 4$, we have

$$S(q,\ell) \le q^{\ell/2+1} \le q^{\ell-1} \le q^\ell/2. \quad \square$$

We note that the inequalities in (20.2) are tight, in the sense that $\Pi(\ell) = q^\ell/2\ell$ when $q = 2$ and $\ell = 2$, and $\Pi(\ell) = q^\ell$ when $\ell = 1$. The first inequality in (20.2) implies not only that $\Pi(\ell) > 0$, but that the fraction of all monic degree ℓ polynomials that are irreducible is at least $1/2\ell$, while (20.3) says that this fraction gets arbitrarily close to $1/\ell$ as either q or ℓ are sufficiently large.

Exercise 20.1. Starting from Theorem 20.10, show that

$$\Pi(\ell) = \ell^{-1} \sum_{k|\ell} \mu(k) q^{\ell/k},$$

where μ is the Möbius function (see §2.6).

Exercise 20.2. How many irreducible polynomials of degree 30 over $\mathbb{Z}_2$ are there?

20.3 The subfield structure and uniqueness of finite fields

We begin with a result that holds for field extensions in general.

Theorem 20.12. *Let E be an extension of a field F, and let σ be an F-algebra automorphism on E. Then the set $E' := \{\alpha \in E : \sigma(\alpha) = \alpha\}$ is a subfield of E containing F.*

Proof. By definition, σ acts as the identity function on F, and so $F \subseteq E'$. To show that E' is a subring of E, it suffices to show that E' is closed under addition and multiplication. To show that E' is closed under addition, let $\alpha, \beta \in E'$. Then $\sigma(\alpha+\beta) = \sigma(\alpha) + \sigma(\beta) = \alpha + \beta$, and hence $\alpha + \beta \in E'$. Replacing "+" by "$\cdot$" in the above argument shows that E' is closed under multiplication. We conclude that E' is a subring of E.

To complete the proof that E' is a subfield of E, we need to show that if $0 \ne \alpha \in E'$ and $\beta \in E$ with $\alpha\beta = 1$, then $\beta \in E'$. We have

$$\alpha\beta = 1 = \sigma(1) = \sigma(\alpha\beta) = \sigma(\alpha)\sigma(\beta) = \alpha\sigma(\beta),$$

and hence $\alpha\beta = \alpha\sigma(\beta)$; canceling α, we obtain $\beta = \sigma(\beta)$, and so $\beta \in E'$. $\square$

The subfield E' in the above theorem is called **the subfield of E fixed**

by σ. Turning our attention again to finite fields, the following theorem completely characterizes the subfield structure of a finite field.

Theorem 20.13. *Let E be an extension of degree ℓ of a finite field F, and let σ be the Frobenius map on E over F. Then the intermediate fields E', with $F \subseteq E' \subseteq E$, are in one-to-one correspondence with the divisors k of ℓ, where the divisor k corresponds to the subfield of E fixed by σ^k, which has degree k over F.*

Proof. Let q be the cardinality of F. Let k be a divisor of ℓ. Now, by Theorem 20.6, the polynomial $\mathtt{X}^{q^\ell} - \mathtt{X}$ splits into distinct linear factors over E, and by Theorem 20.3, the polynomial $\mathtt{X}^{q^k} - \mathtt{X}$ divides $\mathtt{X}^{q^\ell} - \mathtt{X}$. Hence, $\mathtt{X}^{q^k} - \mathtt{X}$ also splits into distinct linear factors over E. This says that the subfield of E fixed by σ^k, which consists of the roots of $\mathtt{X}^{q^k} - \mathtt{X}$, has precisely q^k elements, and hence is an extension of degree k over F. That proves the existence part of the theorem.

As for uniqueness, we have to show that any intermediate field is of this type. Let E' be an intermediate field of degree k over F. By Theorem 20.6, we have $\mathtt{X}^{q^k} - \mathtt{X} = \prod_{\alpha \in E'}(\mathtt{X} - \alpha)$ and $\mathtt{X}^{q^\ell} - \mathtt{X} = \prod_{\alpha \in E}(\mathtt{X} - \alpha)$, from which it follows that $\mathtt{X}^{q^k} - \mathtt{X}$ divides $\mathtt{X}^{q^\ell} - \mathtt{X}$, and so by Theorem 20.3, we must have $k \mid \ell$. There can be no other intermediate fields of the same degree k over F, since the elements of such a field would also be roots of $\mathtt{X}^{q^k} - \mathtt{X}$. □

The next theorem shows that up to isomorphism, there is only one finite field of a given cardinality.

Theorem 20.14. *Let E, E' be extensions of the same degree over a finite field F. Then E and E' are isomorphic as F-algebras.*

Proof. Let q be of cardinality F, and let ℓ be the degree of the extensions. As we have argued before, we have $E' = F[\alpha']$ for some $\alpha' \in E'$, and so E' is isomorphic as an F-algebra to $F[\mathtt{X}]/(\phi)$, where ϕ is the minimal polynomial of α' over F. As ϕ is an irreducible polynomial of degree ℓ, by Theorem 20.9, ϕ divides $\mathtt{X}^{q^\ell} - \mathtt{X}$, and by Theorem 20.6, $\mathtt{X}^{q^\ell} - \mathtt{X} = \prod_{\alpha \in E}(\mathtt{X} - \alpha)$, from which it follows that ϕ has a root $\alpha \in E$. Since ϕ is irreducible, ϕ is the minimal polynomial of α over F, and hence $F[\alpha]$ is isomorphic as an F-algebra to $F[\mathtt{X}]/(\phi)$. Since α has degree ℓ over F, we must have $E = F[\alpha]$. □

Exercise 20.3. This exercise develops an alternative proof for the existence of finite fields — however, it does not yield a density result for irreducible polynomials. Let F be a finite field of cardinality q, and let $\ell \geq 1$ be an integer. Let E be a splitting field for the polynomial $\mathtt{X}^{q^\ell} - \mathtt{X} \in F[\mathtt{X}]$ (see

Theorem 17.19), and let σ be the Frobenius map on E over F. Let E' be the subfield of E fixed by σ^ℓ. Show that E' is an extension of F of degree ℓ.

Exercise 20.4. Let E be an extension of degree ℓ over a finite field F of cardinality q. Show that at least half the elements of E have degree ℓ over F, and that the total number of elements of degree ℓ over F is $q^\ell + O(q^{\ell/2})$.

20.4 Conjugates, norms and traces

Throughout this section, F denotes a finite field of cardinality q, E denotes an extension over F of degree ℓ, and σ denotes the Frobenius map on E over F.

Consider an element $\alpha \in E$. We say that $\beta \in E$ is **conjugate to α (over F)** if $\beta = \sigma^i(\alpha)$ for some $i \in \mathbb{Z}$. The reader may verify that the "conjugate to" relation is an equivalence relation. We call the equivalence classes of this relation **conjugacy classes**, and we call the elements of the conjugacy class containing α the **conjugates of α**.

Starting with α, we can start listing conjugates:

$$\alpha, \sigma(\alpha), \sigma^2(\alpha), \ldots.$$

As σ^ℓ is the identity map, this list will eventually start repeating. Let k be the smallest positive integer such that $\sigma^k(\alpha) = \sigma^i(\alpha)$ for some $i = 0, \ldots, k-1$. It must be the case that $i = 0$ — otherwise, applying σ^{-1} to the equation $\sigma^k(\alpha) = \sigma^i(\alpha)$ would yield $\sigma^{k-1}(\alpha) = \sigma^{i-1}(\alpha)$, and since $0 \leq i - 1 < k - 1$, this would contradict the minimality of k.

Thus, $\alpha, \sigma(\alpha), \ldots, \sigma^{k-1}(\alpha)$ are all distinct, and $\sigma^k(\alpha) = \alpha$. Moreover, for any $i \in \mathbb{Z}$, we have $\sigma^i(\alpha) = \sigma^j(\alpha)$, where $j = i \bmod k$, and so $\alpha, \sigma(\alpha), \ldots, \sigma^{k-1}(\alpha)$ are all the conjugates of α. Also, $\sigma^i(\alpha) = \alpha$ if and only if k divides i. Since $\sigma^\ell(\alpha) = \alpha$, it must be the case that k divides ℓ.

With α and k as above, consider the polynomial

$$\phi := \prod_{i=0}^{k-1} (\mathtt{X} - \sigma^i(\alpha)).$$

The coefficients of ϕ obviously lie in E, but we claim that in fact, they lie in F. This is easily seen as follows. Consider the extension of the map σ from E to $E[\mathtt{X}]$ that applies σ coefficient-wise to polynomials. This was discussed in Example 9.48, where we saw that the extended map, which we also denote by σ, is a ring homomorphism from $E[\mathtt{X}]$ into $E[\mathtt{X}]$. Applying σ

to ϕ, we obtain

$$\sigma(\phi) = \prod_{i=0}^{k-1} \sigma(\mathtt{X} - \sigma^i(\alpha)) = \prod_{i=0}^{k-1} (\mathtt{X} - \sigma^{i+1}(\alpha)) = \prod_{i=0}^{k-1} (\mathtt{X} - \sigma^i(\alpha)),$$

since $\sigma^k(\alpha) = \alpha$. Thus we see that $\sigma(\phi) = \phi$. Writing $\phi = \sum_i a_i \mathtt{X}^i$, we see that $\sigma(a_i) = a_i$ for all i, and hence by Theorem 20.8, $a_i \in F$ for all i. Hence $\phi \in F[\mathtt{X}]$. We further claim that ϕ is the minimal polynomial of α. To see this, let $f \in F[\mathtt{X}]$ be any polynomial over F for which α is a root. Then for any integer i, by Theorem 17.1, we have

$$0 = \sigma^i(0) = \sigma^i(f(\alpha)) = f(\sigma^i(\alpha)).$$

Thus, all the conjugates of α are also roots of f, and so ϕ divides f. That proves that ϕ is the minimal polynomial of α. Since ϕ is the minimal polynomial of α and $\deg(\phi) = k$, it follows that the number k is none other than the degree of α over F.

Let us summarize the above discussion as follows:

Theorem 20.15. *Let $\alpha \in E$ be of degree k over F, and let ϕ be the minimal polynomial of α over F. Then k is the smallest positive integer such that $\sigma^k(\alpha) = \alpha$, the distinct conjugates of α are $\alpha, \sigma(\alpha), \ldots, \sigma^{k-1}(\alpha)$, and ϕ factors over E (in fact, over $F[\alpha]$) as*

$$\phi = \prod_{i=0}^{k-1} (\mathtt{X} - \sigma^i(\alpha)).$$

Another useful way of reasoning about conjugates is as follows. First, if $\alpha = 0$, then the degree of α over F is 1, and there is nothing more to say, so let us assume that $\alpha \in E^*$. If r is the multiplicative order of α, then note that any conjugate $\sigma^i(\alpha)$ also has multiplicative order r — this follows from the fact that for any positive integer s, $\alpha^s = 1$ if and only if $(\sigma^i(\alpha))^s = 1$. Also, note that we must have $r \mid |E^*| = q^\ell - 1$, or equivalently, $q^\ell \equiv 1 \pmod{r}$. Focusing now on the fact that σ is the q-power map, we see that the degree k of α is the smallest positive integer such that $\alpha^{q^k} = \alpha$, which holds iff $\alpha^{q^k - 1} = 1$, which holds iff $q^k \equiv 1 \pmod{r}$. Thus, the degree of α over F is simply the multiplicative order of q modulo r. Again, we summarize these observations as a theorem:

Theorem 20.16. *If $\alpha \in E^*$ has multiplicative order r, then the degree of α over F is equal to the multiplicative order of q modulo r.*

Let us define the polynomial

$$\chi := \prod_{i=0}^{\ell-1} (\mathtt{X} - \sigma^i(\alpha)).$$

It is easy to see, using the same type of argument as above, that $\chi \in F[\mathtt{X}]$, and indeed, that

$$\chi = \phi^{\ell/k}.$$

The polynomial χ is called the **characteristic polynomial of α (from E to F)**.

Two functions that are often useful are the "norm" and "trace." The **norm of α (from E to F)** is defined as

$$\mathbf{N}_{E/F}(\alpha) := \prod_{i=0}^{\ell-1} \sigma^i(\alpha),$$

while the **trace of α (from E to F)** is defined as

$$\mathbf{Tr}_{E/F}(\alpha) := \sum_{i=0}^{\ell-1} \sigma^i(\alpha).$$

It is easy to see that both the norm and trace of α are elements of F, as they are fixed by σ; alternatively, one can see this by observing that they appear, possibly with a minus sign, as coefficients of the characteristic polynomial χ—indeed, the constant term of χ is equal to $(-1)^\ell \mathbf{N}_{E/F}(\alpha)$, and the coefficient of $\mathtt{X}^{\ell-1}$ in χ is $-\mathbf{Tr}_{E/F}(\alpha)$.

The following two theorems summarize the most important facts about the norm and trace functions.

Theorem 20.17. *The function $\mathbf{N}_{E/F}$, restricted to E^*, is a group homomorphism from E^* onto F^*.*

Proof. We have

$$\mathbf{N}_{E/F}(\alpha) = \prod_{i=0}^{\ell-1} \alpha^{q^i} = \alpha^{\sum_{i=0}^{\ell-1} q^i} = \alpha^{(q^\ell-1)/(q-1)}.$$

Since E^* is a cyclic group of order $q^\ell - 1$, the image of the $(q^\ell - 1)/(q-1)$-power map on E^* is the unique subgroup of E^* of order $q-1$ (see Theorem 8.31). Since F^* is a subgroup of E^* of order $q-1$, it follows that the image of this power map is F^*. □

Theorem 20.18. *The function $\mathbf{Tr}_{E/F}$ is an F-linear map from E onto F.*

Proof. The fact that $\mathbf{Tr}_{E/F}$ is an F-linear map is a simple consequence of the fact that σ is an F-algebra automorphism (verify). As discussed above, $\mathbf{Tr}_{E/F}$ maps into F. Since the image of $\mathbf{Tr}_{E/F}$ is a subspace of F, the image is either $\{0\}$ or F, and so it suffices to show that $\mathbf{Tr}_{E/F}$ does not map all of E to zero. But an element $\alpha \in E$ is in the kernel of $\mathbf{Tr}_{E/F}$ if and only if α is a root of the polynomial

$$\mathtt{X} + \mathtt{X}^q + \cdots + \mathtt{X}^{q^{\ell-1}},$$

which has degree $q^{\ell-1}$. Since E contains q^ℓ elements, not all elements of E can lie in the kernel of $\mathbf{Tr}_{E/F}$. □

***Example* 20.1.** As an application of some of the above theory, let us investigate the factorization of the polynomial $\mathtt{X}^r - 1$ over F, a finite field of cardinality q. Let us assume that $r > 0$ and is relatively prime to q. Let E be a splitting field of $\mathtt{X}^r - 1$ (see Theorem 17.19), so that E is a finite extension of F in which $\mathtt{X}^r - 1$ splits into linear factors:

$$\mathtt{X}^r - 1 = \prod_{i=1}^{r} (\mathtt{X} - \alpha_i).$$

We claim that the roots α_i of $\mathtt{X}^r - 1$ are distinct—this follows from the Theorem 20.4 and the fact that $\gcd(\mathtt{X}^r - 1, r\mathtt{X}^{r-1}) = 1$.

Next, observe that the r roots of $\mathtt{X}^r - 1$ in E actually form a subgroup of E^*, and since E^* is cyclic, this subgroup must be cyclic as well. So the roots of $\mathtt{X}^r - 1$ form a cyclic subgroup of E^* of order r. Let ζ be a generator for this group. Then all the roots of $\mathtt{X}^r - 1$ are contained in $F[\zeta]$, and so we may as well assume that $E = F[\zeta]$.

Let us compute the degree of ζ over F. By Theorem 20.16, the degree ℓ of ζ over F is the multiplicative order of q modulo r. Moreover, the $\phi(r)$ roots of $\mathtt{X}^r - 1$ of multiplicative order r are partitioned into $\phi(r)/\ell$ conjugacy classes, each of size ℓ; indeed, as the reader is urged to verify, these conjugacy classes are in one-to-one correspondence with the cosets of the subgroup of $\mathbb{Z}_r^*$ generated by $[q]_r$, where each such coset $C \subseteq \mathbb{Z}_r^*$ corresponds to the conjugacy class $\{\zeta^a : [a]_r \in C\}$.

More generally, for any $s \mid r$, any root of $\mathtt{X}^r - 1$ whose multiplicative order is s has degree k over F, where k is the multiplicative order of q modulo s. As above, the $\phi(s)$ roots of multiplicative order s are partitioned into $\phi(s)/k$ conjugacy classes, which are in one-to-one correspondence with the cosets of the subgroup of $\mathbb{Z}_s^*$ generated by $[q]_s$.

This tells us exactly how $\mathtt{X}^r - 1$ splits into irreducible factors over F. Things are a bit simpler when r is prime, in which case, from the above

discussion, we see that

$$\mathtt{X}^r - 1 = (\mathtt{X} - 1) \prod_{i=1}^{(r-1)/\ell} f_i,$$

where each f_i is an irreducible polynomial of degree ℓ, and ℓ is the multiplicative order of q modulo r.

In the above analysis, instead of constructing the field E using Theorem 17.19, one could instead simply construct E as $F[\mathtt{X}]/(\phi)$, where ϕ is any irreducible polynomial of degree ℓ, and where ℓ is the multiplicative order of q modulo r. We know that such a polynomial ϕ exists by Theorem 20.11, and since E has cardinality q^ℓ, and $r \mid (q^\ell - 1) = |E^*|$, and E^* is cyclic, we know that E^* contains an element ζ of multiplicative order r, and each of the r distinct powers of ζ are roots of $\mathtt{X}^r - 1$, and so this E is a splitting field $\mathtt{X}^r - 1$ over F. □

Exercise 20.5. Let E be a finite extension of a finite field F. Show that for $a \in F$, we have $\mathbf{N}_{E/F}(a) = a^\ell$ and $\mathbf{Tr}_{E/F}(a) = \ell a$.

Exercise 20.6. Let E be a finite extension of a finite field F. Let E' be an intermediate field, $F \subseteq E' \subseteq E$. Show that

(a) $\mathbf{N}_{E/F}(\alpha) = \mathbf{N}_{E'/F}(\mathbf{N}_{E/E'}(\alpha))$, and

(b) $\mathbf{Tr}_{E/F}(\alpha) = \mathbf{Tr}_{E'/F}(\mathbf{Tr}_{E/E'}(\alpha))$.

Exercise 20.7. Let F be a finite field, and let $f \in F[\mathtt{X}]$ be a monic irreducible polynomial of degree ℓ. Let $E = F[\mathtt{X}]/(f) = F[\eta]$, where $\eta := [\mathtt{X}]_f$.

(a) Show that

$$\frac{\mathbf{D}(f)}{f} = \sum_{j=1}^{\infty} \mathbf{Tr}_{E/F}(\eta^{j-1})\mathtt{X}^{-j}.$$

(b) From part (a), deduce that the sequence

$$\mathbf{Tr}_{E/F}(\eta^{j-1}) \quad (j = 1, 2, \ldots)$$

is linearly generated over F with minimal polynomial f.

(c) Show that one can always choose a polynomial f so that sequence in part (b) is purely periodic with period $q^\ell - 1$.

Exercise 20.8. Let F be a finite field, and $f \in F[\mathtt{X}]$ an irreducible polynomial of degree k over F. Let E be an extension of degree ℓ over F. Show that over E, f factors as the product of d distinct irreducible polynomials, each of degree k/d, where $d = \gcd(k, \ell)$.

EXERCISE 20.9. Let E be a finite extension of a finite field F of characteristic p. Show that if $\alpha \in E$ and $0 \neq a \in F$, and if α and $\alpha + a$ are conjugate over F, then p divides the degree of α over F.

EXERCISE 20.10. Let F be a finite field of characteristic p. For $a \in F$, consider the polynomial $f := \mathtt{X}^q - \mathtt{X} - a \in F[\mathtt{X}]$.

(a) Show that if $F = \mathbb{Z}_p$ and $a \neq 0$, then f is irreducible.

(b) More generally, show that if $\mathbf{Tr}_{F/\mathbb{Z}_p}(a) \neq 0$, then f is irreducible, and otherwise, f splits into distinct linear factors over F.

EXERCISE 20.11. Let E be a finite extension of a finite field F. Let $\alpha, \beta \in E$, where α has degree a over F, β has degree b over F, and $\gcd(a, b) = 1$. Show that $\alpha + \beta$ has degree ab over F.

EXERCISE 20.12. Let E be a finite extension of a finite field F. Show that any F-algebra automorphism on E must be a power of a the Frobenius map on E over F.

EXERCISE 20.13. Show that for all primes p, the polynomial $\mathtt{X}^4 + 1$ is reducible in $\mathbb{Z}_p[\mathtt{X}]$. (Contrast this to the fact that this polynomial is irreducible in $\mathbb{Q}[\mathtt{X}]$, as discussed in Exercise 17.39.)

EXERCISE 20.14. This exercise depends on the concepts and results in §19.6. Let F be a finite field and let E be an extension of degree ℓ. Let σ be the Frobenius map on E over F.

(a) Show that the minimal polynomial of σ over F is $\mathtt{X}^\ell - 1$.

(b) Show that there exists $\beta \in E$ such that the minimal polynomial of β under σ is $\mathtt{X}^\ell - 1$.

(c) Conclude that $\beta, \sigma(\beta), \ldots, \sigma^{\ell-1}(\beta)$ is a basis for E over F. This type of basis is called a **normal basis**.