

Introduction to finite fields

Topics in Finite Fields (Fall 2013)
Rutgers University
Swastik Kopparty
Last modified: Monday 16th September, 2013

Welcome to the course on finite fields! This is aimed at graduate students in mathematics and theoretical computer science¹. We will see a number of important classical and modern themes in the study of finite fields.

1 Course Info

- Instructor: Swastik Kopparty (swastik.kopparty@rutgers.edu, swastik.kopparty@gmail.com).
- Meeting times: Monday, Wednesday: 5:00 - 6:20.
- NOTE: Some classes will be cancelled, but we make up for this by extending the usual time of 5:00 - 6:20 to 5:00 - 8:00 on some Mondays. More info on this as the semester proceeds.
- Recommended texts: Finite Fields (Lidl and Niederrieter), Equations over Finite Fields (Schmidt), Additive Combinatorics (Tao and Vu).
- Problem sets: There will be problem sets and problems scattered through the lecture notes. Each problem will be worth some number of points (between 1 (easy) and 10 (open problem)). You should turn in 20 points.

2 Finite field basics

We will start by reviewing some of the basics of field theory, and then get a complete classification of all finite fields.

Recall that a field is a set \mathbb{F} equipped with two operations, addition (+) and multiplication (\cdot), and two special elements 0, 1, satisfying:

- $(\mathbb{F}, +)$ is an abelian group with identity element 0.
- (\mathbb{F}^*, \cdot) is an abelian group with identity element 1 (here \mathbb{F}^* denotes $\mathbb{F} \setminus \{0\}$).
- For all $a \in \mathbb{F}$, $0 \cdot a = a \cdot 0 = 0$.
- Distributivity: for all $a, b, c \in \mathbb{F}$, we have $a \cdot (b + c) = a \cdot b + a \cdot c$.

A finite field is a field which is, well, finite.

Recall the notion of a vector space over a field. Note that an n -dimensional vector space over a finite field of cardinality q has cardinality q^n (and in particular, is finite).

¹Students from the Department of Agriculture should see me after class. There might have been a misunderstanding.

2.1 \mathbb{F}_p

The simplest example of a finite field is as follows. Take a prime $p \in \mathbb{Z}$. Let $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$ (the quotient of the ring \mathbb{Z} mod the ideal $p\mathbb{Z}$). Very explicitly, $\mathbb{F}_p = \{0, 1, \dots, p-1\}$, and the operations are addition and multiplication of integers mod p .

To see that this is a field, the main step is to verify that every $a \in \mathbb{F}_p^*$ has a multiplicative inverse. Since $a \in \mathbb{F}_p^*$ and p is prime, we have that $\text{GCD}(a, p) = 1$, and so by Euclid, we know that there exist integers x, y s.t. $ax + py = 1$. Then $x \pmod p$ is a^{-1} .

2.2 The prime subfield

Let \mathbb{F} be a finite field. For a positive integer r , consider the r -fold sum $s_r = 1 + 1 + \dots + 1$. Since \mathbb{F} is finite, some s_r must equal 0. Let p be the smallest positive p for which s_p equals 0. Observe that if p exists, it must be prime; for if $p = a \cdot b$ with $a, b < p$, then by distributivity we have $0 = s_p = s_a \cdot s_b$, and so one of s_a, s_b must equal 0, contradicting the minimality of p . This p is called the characteristic of the field \mathbb{F} .

Now observe that the subset $\{0, s_1, s_2, \dots, s_{p-1}\} \subseteq \mathbb{F}$ is itself a field, isomorphic to \mathbb{F}_p . This is called the prime subfield of \mathbb{F} .

The key to the full classification of all finite fields is the observation that \mathbb{F} is a finite dimensional vector space over \mathbb{F}_p . Let $n = \dim_{\mathbb{F}_p}(\mathbb{F})$. Then we have $|\mathbb{F}| = p^n$. In particular, the cardinality of a finite field must be a prime power.

2.3 $\mathbb{F}[X]$

Let \mathbb{F} be a field. Consider $\mathbb{F}[X]$, the ring of 1-variable polynomials over \mathbb{F} . Because of the division algorithm $\mathbb{F}[X]$, we have:

- $\mathbb{F}[X]$ is a principal ideal domain: every ideal in $\mathbb{F}[X]$ is generated by a single $g(X) \in \mathbb{F}[X]$.
- $\mathbb{F}[X]$ is a unique factorization domain.
- The remainder theorem: $P(X) \in \mathbb{F}[X]$ vanishes at a point $\alpha \in \mathbb{F}$ if and only if $(X - \alpha)$ divides $P(X)$.

These facts give us the following fundamental theorem.

Theorem 1. *If $P(X) \in \mathbb{F}[X]$ is a nonzero polynomial of degree $\leq d$, then $P(X)$ has $\leq d$ roots in \mathbb{F} .*

We will also be using derivatives of polynomials. Define the derivative map $D : \mathbb{F}[X] \rightarrow \mathbb{F}[X]$ by $D(X^i) = i \cdot X^{i-1}$, and extend it by \mathbb{F} -linearity. We then have the product rule:

$$D(f \cdot g) = f \cdot D(g) + D(f) \cdot g.$$

Sometimes we will denote $D(f(X))$ by $f'(X)$.

Some important remarks:

- Even though derivatives do not have a geometric interpretation as in the \mathbb{R} case, they are very useful for us because they can be used to detect double roots of polynomials: $(X - \alpha)^2$ divides $P(X)$ if and only if $P(\alpha) = 0$ and $P'(\alpha) = 0$. This follows from the product rule (prove it).
- For detecting triple and higher roots of polynomials (which we will be interested in later), it turns out that the higher order derivatives D^i do NOT do the job. This is related to the fact that derivatives can end up being zero too easily: in characteristic 2, for example, $D^2(f(X)) = 0$ for all $f(X)$.
- This issue will be solved by the *Hasse* derivatives, coming soon.

2.4 Algebraic extensions

Let \mathbb{F} be a field and let \mathbb{K} be a subfield. Note that \mathbb{F} is a vector space over \mathbb{K} . We denote $\dim_{\mathbb{K}}(\mathbb{F})$ by $[\mathbb{F} : \mathbb{K}]$. Suppose $[\mathbb{F} : \mathbb{K}] = n$. Now pick any $\alpha \in \mathbb{F}$. Consider the elements $1, \alpha, \alpha^2, \dots, \alpha^n$. Since \mathbb{F} is n -dimensional over \mathbb{K} , these $n + 1$ elements must be linearly dependent over \mathbb{K} . Thus α is the root of some nonzero $P(X) \in \mathbb{K}[X]$. Thus every $\alpha \in \mathbb{F}$ is *algebraic* over \mathbb{K} .

Theorem 2. *Every finite field \mathbb{F} of characteristic p is a finite algebraic extension field of \mathbb{F}_p .*

Conversely, every finite algebraic extension field of \mathbb{F}_p is a finite field. This gives us a characterization of all finite fields.

2.5 Constructing algebraic extensions

Let \mathbb{F} be a field and let \mathbb{K} be a subfield with $[\mathbb{F} : \mathbb{K}]$ finite.

Now consider an element $\alpha \in \mathbb{F}$. Let $M(X) \in \mathbb{K}[X]$ be a nonzero monic polynomial of smallest degree for which $M(\alpha) = 0$. Note that $M(X)$ must be irreducible in $\mathbb{K}[X]$. Also note that if $P(X) \in \mathbb{K}[X]$ is such that $P(\alpha) = 0$, then $M(X) \mid P(X)$ (this follows from the division algorithm). In particular, $M(X)$ is the unique nonzero monic polynomial of smallest degree for which $M(\alpha) = 0$. This $M(X)$ is called the minimal polynomial of α over \mathbb{K} .

Define $\mathbb{K}(\alpha)$ to be the smallest subfield of \mathbb{F} containing both \mathbb{K} and α . Observe that if $d = \deg(M(X))$, then

$$\mathbb{K}(\alpha) = \text{span}_{\mathbb{K}}\{1, \alpha, \alpha^2, \dots, \alpha^{d-1}\}.$$

This equality holds between the two subsets of \mathbb{F} ; in particular it assumes that we already have our hands on \mathbb{F} , and describes $\mathbb{K}(\alpha)$ in terms of the operations of \mathbb{F} .

The key to getting a “construction” of \mathbb{F} using only operations from \mathbb{K} is the following field isomorphism:

$$\mathbb{K}(\alpha) \cong \mathbb{K}[X]/\langle M(X) \rangle.$$

(This is the ring $\mathbb{K}[X]$ mod the ideal generated by $M(X)$). The right hand side is defined only in terms of \mathbb{K} , and is called a primitive field extension. If $d = \deg(M(X))$, then it is a degree d extension of \mathbb{K} .

This gives us the following construction of \mathbb{F} from \mathbb{K} . Set $\mathbb{K}_0 = \mathbb{K}$. Pick $\alpha_1 \in \mathbb{F} \setminus \mathbb{K}_0$, and set $\mathbb{K}_1 = \mathbb{K}_0(\alpha_1)$. If $\mathbb{K}_1 = \mathbb{F}$, then we are done. Otherwise, pick $\alpha_2 \in \mathbb{F} \setminus \mathbb{K}_1$, set $\mathbb{K}_2 = \mathbb{K}_1(\alpha_2)$, etc. At each stage, $[\mathbb{F} : \mathbb{K}_i]$ reduces, and thus the process must stop. This gives us a construction of \mathbb{F} as a finite sequence of primitive extensions of \mathbb{K} . It can be completely specified by the sequence of irreducible polynomials $M_i(X) \in \mathbb{K}_i[X]$, where $M_i(X)$ is the minimal polynomial of α_{i+1} over \mathbb{K}_i .

We will eventually make this even more explicit: we will see that every finite field is a primitive extension of \mathbb{F}_p .

2.6 Existence of a finite field of cardinality p^n

We will first find some properties that any finite field of size p^n must have, and then let that guide our search.

Let \mathbb{F} be a finite field of cardinality $q = p^n$.

Since the multiplicative group of \mathbb{F}^* has cardinality $q - 1$, we have that $\alpha^{q-1} = 1$ for every $\alpha \in \mathbb{F}^*$. Thus $\alpha^q = \alpha$ for every $\alpha \in \mathbb{F}$.

Let us rephrase this. Consider the polynomial $P(X) = X^q - X$. All q elements of \mathbb{F} are roots of this degree q polynomial. Thus $P(X) = \prod_{\alpha \in \mathbb{F}} (X - \alpha)$.

This motivates a construction of \mathbb{F} . First, a lemma.

Lemma 3. *Let \mathbb{K} be a field. Let $P(X) \in \mathbb{K}[X]$. Then there exists a field $\mathbb{L} \supseteq \mathbb{K}$, with $[\mathbb{L} : \mathbb{K}]$ finite, such that $P(X)$ factors into linear factors over $\mathbb{L}[X]$.*

This lemma is proved by iterating the following observation. Let $M(X) \in \mathbb{K}[X]$ be an irreducible polynomial. Then if we consider the extension field $\mathbb{K}_1 = \mathbb{K}[T]/\langle M(T) \rangle$ and let α denote the element corresponding to T in \mathbb{K}_1 , we have that $(X - \alpha) \mid M(X)$ in $\mathbb{K}_1[X]$.

Now let $q = p^n$ with p prime. Set $\mathbb{K} = \mathbb{F}_p$ and $P(X) = X^q - X \in \mathbb{K}[X]$. The lemma tells us that there exists a finite field \mathbb{L} in which $X^q - X$ factors as a product of linear factors:

$$\prod_{i=1}^q (X - \alpha_i),$$

where each $\alpha_i \in \mathbb{L}$.

We first observe that all the α_i are distinct. Consider any $i \in [q]$. If we set $Q(X) = \frac{P(X)}{X - \alpha_i}$, then we have the identity (using the product rule for derivatives):

$$Q(\alpha_i) = P'(\alpha_i) = -1.$$

Thus $Q(X)$ is not divisible by $(X - \alpha_i)$, or equivalently, $P(X)$ is not divisible by $(X - \alpha_i)^2$. Thus all the α_i are distinct.

Set $\mathbb{F} = \{\alpha \in \mathbb{L} \mid P(\alpha) = 0\} = \{\alpha_i \mid i \in [q]\}$, and note that $|\mathbb{F}| = q$. We will now show that \mathbb{F} is a field.

1. Note that $0 \in \mathbb{F}$, since $P(0) = 0$.
2. If $a, b \in \mathbb{F}^*$, then $a^q = a$ and $b^q = b$, and so

$$P\left(\frac{a}{b}\right) = \left(\frac{a}{b}\right)^q - \left(\frac{a}{b}\right) = \frac{a^q}{b^q} - \frac{a}{b} = \frac{a}{b} - \frac{a}{b} = 0.$$

Thus $\frac{a}{b} \in \mathbb{F}^*$, and so \mathbb{F}^* is a group under multiplication.

3. If $a, b \in \mathbb{F}$, then $a^q = a$ and $b^q = b$, and so

$$P(a - b) = (a - b)^q - (a - b).$$

We now make a very very important observation²:

Observation 4. *If \mathbb{L} is a field of characteristic p , and $a, b \in \mathbb{L}$, then*

$$(a + b)^p = a^p + b^p.$$

Proof. Binomial theorem, along with the fact that $\binom{p}{r} = 0 \pmod{p}$ when $0 < r < p$. □

Iterating this observation, we have that for every m ,

$$(a + b)^{p^m} = a^{p^m} + b^{p^m}.$$

Thus³

$$P(a - b) = a^q + (-b)^q - a + b = a + (-1)^q b - a + b = 0,$$

so $a - b \in \mathbb{F}$. This implies that \mathbb{F} is a group under addition.

Thus \mathbb{F} is a finite field of cardinality q .

²This is sometimes called the freshman's dream, because freshmen prefer working in fields of positive characteristic, especially during Calculus 1.

³Note that if q is even, then \mathbb{L} is of characteristic 2, and so $b + b = (1 + 1) \cdot b = 0$.

2.7 Uniqueness of \mathbb{F}_q

We now show that this field is unique: any two fields of cardinality q are isomorphic.

Lemma 5 (Uniqueness of the splitting field). *Let \mathbb{K} be a field and let $P(X) \in \mathbb{K}[X]$. Let $\mathbb{F}_1, \mathbb{F}_2$ be fields such that for $i \in \{1, 2\}$ (1) $\mathbb{K} \subseteq \mathbb{F}_i$, (2) $P(X)$ factors into linear factors in $\mathbb{F}_i[X]$, and (3) No strict subfield of \mathbb{F}_i satisfies both (1) and (2).*

Then $\mathbb{F}_1 \cong \mathbb{F}_2$.

Proof. This lemma is proved as follows. If $P(X)$ factors into linear factors in \mathbb{K} , then $\mathbb{F}_1 \cong \mathbb{K} \cong \mathbb{F}_2$, and we are done. Otherwise, let $M(X)$ be an irreducible factor of $P(X)$. Since $P(X)$ factors into linear factors in $\mathbb{F}_i[X]$, then so does $M(X)$. Let $\alpha_i \in \mathbb{F}_i$ be a root of $M(X)$. Define $K_i = \mathbb{K}(\alpha_i) \subseteq \mathbb{F}_i$. From the previous section, we have $\mathbb{K}_i \cong \mathbb{K}[X]/\langle M(X) \rangle$. Thus $\mathbb{K}_1 \cong \mathbb{K}_2$, and so we may identify \mathbb{K}_1 with \mathbb{K}_2 , and call this field \mathbb{K}_0 . We thus have that: (1) $\mathbb{K}_0 \subseteq \mathbb{F}_i$, (2) $P(X)$ factors into linear factors in $\mathbb{F}_i[X]$, and (3) No strict subfield satisfies both (1) and (2). We can then proceed by induction. \square

Now suppose $\mathbb{F}_1, \mathbb{F}_2$ are finite fields of cardinality $q = p^n$, where p is prime. Set $\mathbb{K} = \mathbb{F}_p$, and we have that $\mathbb{K} \subseteq \mathbb{F}_i$. Now we know that each \mathbb{F}_i contains all the q roots of $X^q - X$, and no strict subfield of \mathbb{F}_i can (just because it is not big enough to). Thus, by the previous lemma, $\mathbb{F}_1 \cong \mathbb{F}_2$, as desired.

This unique finite field of cardinality q is called THE Galois field $GF(q)$. Once you get to know it better, you may call it simply $GF(q)$, or even \mathbb{F}_q .

The construction of $GF(q)$ as an algebraic extension of a prime field was first done by Galois. The observation that $GF(q)$ s are the only fields was made by E. H. Moore.

Another corollary of the uniqueness, is that if $M(X) \in \mathbb{F}_q[X]$ is an irreducible polynomial of degree d , then the primitive extension $\mathbb{F}_q[X]/\langle M(X) \rangle$ is isomorphic to \mathbb{F}_{q^d} .

Remark If you already know/believe the existence and uniqueness of the algebraic closure of a field, then it simplifies some of what we did above. Take $\overline{\mathbb{F}_p}$, and then $\mathbb{F} = \{\alpha \in \overline{\mathbb{F}_p} \mid \alpha^q = \alpha\}$ is the unique finite field of cardinality q .

2.8 Subfield structure

Let $q = p^n$, where p is prime.

What are the subfields of \mathbb{F}_q ? Suppose \mathbb{F}_ℓ is a subfield of \mathbb{F}_q . First note that we must have that ℓ is a power of p , say $\ell = p^r$. Also, we must have that \mathbb{F}_ℓ^* is a subgroup of \mathbb{F}_q^* , and so $|\mathbb{F}_\ell^*| = p^r - 1$ should divide $|\mathbb{F}_q^*| = p^n - 1$.

This can only happen if $r \mid n$ (exercise).

Conversely, if $r \mid n$, then $p^r - 1$ divides $p^n - 1$, and so $X^{p^r-1} - 1$ divides $X^{p^n-1} - 1$, and so $X^\ell - X$ divides $X^q - X$.

Thus, since $X^q - X = \prod_{\alpha \in \mathbb{F}_q} (X - \alpha)$, there are ℓ distinct roots of $X^\ell - X$ in \mathbb{F}_q , and by previous discussions, they form the field \mathbb{F}_ℓ . Furthermore, \mathbb{F}_q contains exactly one copy of \mathbb{F}_ℓ .

This completely determines the subfield lattice of a given finite field.

2.9 Conjugates

Let $\mathbb{K} \subseteq \mathbb{F}$ be finite fields, with $|\mathbb{K}| = q$ and $|\mathbb{F}| = q^n$.

Lemma 6. *Let $\alpha \in \mathbb{F}$. Let $P(X) \in \mathbb{K}[X]$. Then $P(\alpha^q) = P(\alpha)^q$.*

Proof. Let $P(X) = \sum_{i=0}^d a_i X^i$. Then

$$P(\alpha^q) = \sum_{i=0}^d a_i \alpha^{qi} = \sum_{i=0}^d a_i^q \alpha^{iq} = \left(\sum_{i=0}^d a_i \alpha^i \right)^q = P(\alpha)^q.$$

(Here we used the fact that $a^q = a$ for each $a \in K$, and $(\alpha + \beta)^q = \alpha^q + \beta^q$ for each $\alpha, \beta \in \mathbb{F}$). \square

Let $\alpha \in \mathbb{F}$, and let $M(X) \in \mathbb{K}[X]$ be the minimal polynomial of α . Let d be the degree of $M(X)$.

By the above lemma, $M(X)$ also has $\alpha^q, \alpha^{q^2}, \dots, \alpha^{q^i}, \dots$ as roots. Let r be the smallest positive integer such that $\alpha^{q^r} = \alpha$. Then $\alpha, \alpha^q, \dots, \alpha^{q^{r-1}}$ are distinct roots of $M(X)$. Thus $r \leq d$.

On the other hand, $\alpha^{q^r} = \alpha$ implies that $\alpha \in \mathbb{F}_{q^r}$. Thus $\mathbb{K}(\alpha) \subseteq \mathbb{F}_{q^r}$. But $\mathbb{K}(\alpha) \cong \mathbb{K}[X]/\langle M(X) \rangle$, and so $\mathbb{K}(\alpha) = \mathbb{F}_{q^d}$. So $d \leq r$.

Thus $d = r$.

To summarize: if $\alpha \in \mathbb{F}_{q^n}$, and $M(X) \in \mathbb{F}_q[X]$ is its minimal polynomial, with $\deg(M(X)) = d$, then:

$$M(X) = (X - \alpha) \cdot (X - \alpha^q) \cdot \dots \cdot (X - \alpha^{q^{d-1}}).$$

Furthermore, $\mathbb{F}_q(\alpha) \cong \mathbb{F}_{q^d}$, and $d \mid n$.

We can also do this in the other direction. Let $M(X) \in \mathbb{F}_q[X]$ be an irreducible polynomial of degree d . Let n be a multiple of d . Then $M(X)$ factors into linear factors in $\mathbb{F}_{q^n}[X]$:

$$M(X) = (X - \alpha) \cdot (X - \alpha^q) \cdot \dots \cdot (X - \alpha^{q^{d-1}}),$$

where $\alpha \in \mathbb{F}_{q^n}$ is any root of $M(X)$. (The only new step to verify here is that $M(X)$ has a root in \mathbb{F}_{q^n} : this follows from the fact that $M(X)$ has a root in $\mathbb{F}_q[X]/\langle M(X) \rangle$, which is isomorphic to \mathbb{F}_{q^d} , which is a subfield of \mathbb{F}_{q^n}).

The above discussion gives us the following important lemma.

Lemma 7. *Let $I_d \subseteq \mathbb{F}_q[X]$ be the set of monic irreducible polynomials in $\mathbb{F}_q[X]$ of degree exactly d . We have the following polynomial equality:*

$$X^{q^n} - X = \prod_{d \mid n} \prod_{M(X) \in I_d} M(X).$$

2.10 The additive group \mathbb{F}

Let \mathbb{F} be a finite field of cardinality $q = p^n$ (where p is a prime).

By what we have already discussed, the additive group \mathbb{F} is isomorphic to \mathbb{Z}_p^n .

2.11 The multiplicative group \mathbb{F}^*

Let \mathbb{F} be a finite field of cardinality $q = p^n$ (where p is a prime).

We will now show that \mathbb{F}^* is isomorphic to the cyclic group \mathbb{Z}_{q-1} .

The key ingredient is the classification of finite abelian groups:

Theorem 8. *For every finite abelian group G , there is a unique tuple (d_1, \dots, d_k) of positive integers such that $1 < d_1 \mid d_2 \mid d_3 \mid \dots \mid d_k$ and:*

$$G \cong \bigoplus_{i=1}^k \mathbb{Z}_{d_i}.$$

This theorem implies a simple criterion for cyclicity:

Corollary 9. *A finite abelian group G is cyclic if and only if for every $d \geq 1$, the number of $\alpha \in G$ such that $d\alpha = 0$ is at most d .*

Proof. The only if direction is trivial.

For the if direction, let G be a finite abelian group such that for every d , the number of $\alpha \in G$ such that $d\alpha = 0$ is at most d . Let (d_1, \dots, d_k) be the tuple such that $1 < d_1 \mid d_2 \mid d_3 \mid \dots \mid d_k$ and

$$G \cong \bigoplus_{i=1}^k \mathbb{Z}_{d_i}.$$

Then $|G| = \prod_{i=1}^k d_i$, and every element α of G satisfies $d_k \cdot \alpha = 0$. Thus k must equal 1 and so G is cyclic. \square

We now show that this criterion is applicable to \mathbb{F}^* .

Lemma 10. *For every $d \geq 1$, the number of $\alpha \in \mathbb{F}^*$ such that $\alpha^d = 1$ is at most d .*

Proof. Every such α is a root of the nonzero degree d polynomial $P(X) = X^d - 1$. □

We thus get:

Theorem 11. $\mathbb{F}^* \cong \mathbb{Z}_{q-1}$.

In particular, there are elements α of \mathbb{F}^* for which $1, \alpha, \alpha^2, \dots, \alpha^{q-2}$ are all distinct, and $\mathbb{F}^* = \{1, \alpha, \dots, \alpha^{q-2}\}$. Such an α is called a generator for \mathbb{F}_q^* .

2.12 The primitive element theorem

We can now use this to prove the primitive element theorem for finite extensions of \mathbb{F}_q .

Theorem 12. *For every prime power q , for every $n > 1$, there exists $\alpha \in \mathbb{F}_{q^n}$ such that $\mathbb{F}_q(\alpha) = \mathbb{F}_{q^n}$.*

Proof. We want $\alpha \in \mathbb{F}_{q^n}^*$ such that $\alpha^{q^r} \neq \alpha$ for all $r < n$. Equivalently, we want α such that $\alpha^{q^r - 1} \neq 1$ for all $r < n$. Take α to be a generator of $\mathbb{F}_{q^n}^*$.

Another proof follows by noting that all subfields of \mathbb{F}_{q^n} (which contain \mathbb{F}_q) are of size $\leq q^{n/2}$, and there is at most one subfield of size q^r for each $r \leq n/2$. Any α not in the union of these subfields does the job for us. Such an α (in fact many α) exist since $q^n - \sum_{j=1}^{n/2} q^j > 0$. □

Lemma 13. *For every prime power q and every $n > 0$, there exists an irreducible polynomial $P(X) \in \mathbb{F}_q[X]$ of degree exactly n .*

Proof. Let $\alpha \in \mathbb{F}_{q^n}$ such that $\mathbb{F}_q(\alpha) = \mathbb{F}_{q^n}$. Let $P(X)$ be the minimal polynomial of α . □

Later we will exactly count the number of irreducible polynomials of degree n .

2.13 Trace and Norm

Let q be a prime power, $n > 0$, and consider the fields \mathbb{F}_q and \mathbb{F}_{q^n} .

The **trace** from \mathbb{F}_{q^n} to \mathbb{F}_q is the following map:

$$\mathrm{Tr}(\alpha) = \sum_{i=0}^{n-1} \alpha^{q^i}.$$

Tr has the following properties:

1. One verifies that $\mathrm{Tr}(\alpha)^q = \mathrm{Tr}(\alpha)$, and thus the image of Tr is indeed contained in \mathbb{F}_q . This is a very illuminating calculation if you are new to finite fields. Try it!
2. If the minimal polynomial of α is

$$M(X) = X^d + c_1 X^{d-1} + \dots + c_d,$$

then $\mathrm{Tr}(\alpha)$ equals $\frac{n}{d} \cdot c_1$.

3. Observe that Tr is an \mathbb{F}_q -linear map of vector spaces.
4. What is the dimension of the image and of the kernel of Tr ? The image is at most 1 dimensional, and hence the kernel must have dimension either n or $n - 1$.

If the dimension of the kernel of Tr is n , then this means that Tr is 0 on all of \mathbb{F}_{q^n} . But Tr is also a nonzero polynomial of degree at most q^{n-1} , and so it cannot be 0 on more than q^{n-1} points. Thus the dimension of the kernel of Tr equals $n - 1$, and thus the image of Tr equals all of \mathbb{F}_q .

As a consequence, for every $a \in \mathbb{F}_q$, $|\mathrm{Tr}^{-1}(a)| = q^{n-1}$.

5. Out of the q^n \mathbb{F}_q -linear maps from \mathbb{F}_{q^n} to \mathbb{F}_q , Tr is one of them. In fact, for every $\beta \in \mathbb{F}_q$, we have the linear map $T_\beta : \mathbb{F}_{q^n} \rightarrow \mathbb{F}_q$, with $T_\beta(\alpha) = \text{Tr}(\beta \cdot \alpha)$. By the previous discussion, have that T_β is not the identically 0 map for every nonzero β . Since $T_\beta - T_\gamma = T_{\beta-\gamma}$, we have that all the T_β are distinct.

Thus every \mathbb{F}_q -linear map from \mathbb{F}_{q^n} to \mathbb{F}_q is of the form T_β for some $\beta \in \mathbb{F}_{q^n}$. (This condition is equivalent to separability of the extension $\mathbb{F}_{q^n}/\mathbb{F}_q$).

The **norm** from \mathbb{F}_{q^n} to \mathbb{F}_q is the following map:

$$\text{Norm}(\alpha) = \prod_{i=0}^{n-1} \alpha^{q^i} = \alpha^{(q^n-1)/(q-1)}.$$

Norm has the following properties:

1. One verifies that $\text{Norm}(\alpha)^q = \text{Norm}(\alpha)$, and thus the image of Norm is indeed contained in \mathbb{F}_q .
2. If the minimal polynomial of α is

$$M(X) = X^d + c_1X^{d-1} + \dots + c_d,$$

then $\text{Norm}(\alpha)$ equals $c_d^{n/d}$.

3. The only element with Norm equal to 0 is 0.
4. Thus Norm maps $\mathbb{F}_{q^n}^*$ to \mathbb{F}_q^* , and is in fact a multiplicative group homomorphism. Since both groups here are cyclic, we have the following situation: we have the group \mathbb{Z}_{q^n-1} , and the map $\alpha \mapsto \frac{q^n-1}{q-1} \cdot \alpha$, whose image equals the unique subgroup of \mathbb{Z}_{q^n-1} of size $q-1$.

Thus Norm is onto \mathbb{F}_q^* , and every element $a \in \mathbb{F}_q^*$ is such that $|\text{Norm}^{-1}(a)| = (q^n - 1)/(q - 1)$.

The kernels of norm and trace have the following convenient characterizations.

Theorem 14 (Hilbert Theorems 90). *1. For $\alpha \in \mathbb{F}_{q^n}$, $\text{Tr}(\alpha) = 0$ iff there exists $\beta \in \mathbb{F}_{q^n}$ with $\alpha = \beta^q - \beta$.
2. For $\alpha \in \mathbb{F}_{q^n}^*$, $\text{Norm}(\alpha) = 1$ iff there exists $\beta \in \mathbb{F}_{q^n}^*$ with $\alpha = \beta^q / \beta$.*

Proof. Let $S : \mathbb{F}_{q^n} \rightarrow \mathbb{F}_{q^n}$ be the map $S(\beta) = \beta^q - \beta$. The first part can be proved by noting that S is an \mathbb{F}_q -linear map, the image of S is contained in the kernel of Tr (this is a simple calculation: $\text{Tr}(S(\beta)) = \beta^{q^n} - \beta = 0$), the kernel of S has dimension at most 1 (since S is a nonzero degree q polynomial, and hence has at most q zeroes), and thus the image of S has dimension $n - 1$.

The second part is a simple consequence of the cyclicity of $\mathbb{F}_{q^n}^*$. For $d \mid q^n - 1$, an element α is a d th root of 1 if and only if it can be expressed as $\beta^{(q^n-1)/d}$ for some $\beta \in \mathbb{F}_{q^n}^*$. We apply this fact with $d = (q^n - 1)/(q - 1)$. \square

2.14 Computational aspects

Many of the basic theorems about finite fields that we just saw have pretty tricky and indirect proofs. Their trickiness can be measured by the difficulty of making them constructive. Here are some open problems.

Problem 15 (Open problem). *Find a deterministic algorithm, which when given a prime p and an integer n , constructs the finite field \mathbb{F}_{p^n} in time $\text{poly}(\log p, n)$.*

The problem is even open for $n = 2$!

Here construct could be taken to mean “find an irreducible polynomial in $\mathbb{F}_p[X]$ of degree n ”. It is known how to do this with a randomized algorithm, and also how to do this deterministically in time $\text{poly}(p, n)$.

Problem 16 (Open problem). *Find a (possibly randomized) algorithm, which when given a prime p , computes a generator of \mathbb{F}_p^* in time $\text{poly}(\log p)$.*

It is known that under the Extended Riemann Hypothesis, some integer at most $\text{poly}(\log p)$ is a generator of \mathbb{F}_p^* . But unfortunately, even the following is an open problem.

Problem 17 (Open problem). *Find a (possibly randomized) algorithm, which when given a prime p and an integer $a < p$, decides whether a is a generator of \mathbb{F}_p^* in time $\text{poly}(\log p)$.*

These problems are also open for other finite fields too.

2.15 Other general comments about finite fields

One may wonder why the prime fields \mathbb{F}_p got singled out as the easily constructed fields, and we had to struggle so much to construct the other finite fields. Maybe only \mathbb{F}_p s are “natural” and worth studying, and the other fields only exist just because of some annoying accident. But in fact, the study of all finite fields is one unified subject.

Here are some related remarks.

- General finite fields arise as quotients of general number rings. For example, in the ring $\mathbb{Z}[i]$ (where i is a square root of -1), the ideal $\langle 7 \rangle$ has the property that

$$\mathbb{Z}[i]/\langle 7 \rangle \cong \mathbb{F}_{49}.$$

In general, every \mathbb{F}_q arises as the quotient of a number ring by an ideal.

- Some theorems about prime fields \mathbb{F}_p do not hold for general fields \mathbb{F}_q . This often helps us appreciate the subtleties in the proofs of the theorems. We will see a number of theorems, especially when we talk about the sum-product phenomenon, where our proofs will have to differentiate between the different q 's.
- Even if you care only about \mathbb{F}_p , there are some very basic and fundamental theorems about \mathbb{F}_p which we can only prove by considering extensions \mathbb{F}_{p^n} . For example, the Weil theorems on counting \mathbb{F}_p solutions of $Q(X, Y) = 0$, take advantage of information from ALL extensions \mathbb{F}_{p^n} simultaneously.
- There are many analogies that help transfer proof methods and intuition between \mathbb{F}_q 's for different q 's.

A Other characterizations of finite fields

It turns out that finite fields (as opposed to general fields) can be characterized by even simpler conditions. An integral domain is a commutative ring with no nontrivial zero divisors: if $a \cdot b = 0$, then either $a = 0$ or $b = 0$. We have the following simple fact.

Lemma 18. *Every finite integral domain is a field.*

A skew field is almost a field, except that we drop the requirement that \cdot be commutative. Wedderburn's little theorem deals with finite skew fields.

Theorem 19. *Every finite skew field is a field.*