Math 75 notes, Lecture 11

P. Pollack and C. Pomerance

Finite field notation and presentation

We have seen that for each prime or prime power q, there is a unique (up to isomorphism) finite field with q elements. We denote this finite field $\mathbb{F}_q$. (Some books use the notation $GF(q)$, where the letters GF stand for "Galois field".)

From a practical point of view, it is somewhat glib to say there is just one finite field of a given size, even though it is true! For example, if $f(x)$ and $g(x)$ are two irreducible polynomials in $(\mathbb{Z}/(2))[x]$ of degree 100, and $f(x)$ has just 3 nonzero terms, while $g(x)$ has 50 nonzero terms, it is much more pleasant to do arithmetic in $(\mathbb{Z}/(2))[x]/(f(x))$ than in $(\mathbb{Z}/(2))[x]/(g(x))$, even though both of these finite fields are isomorphic (to $\mathbb{F}_{2^{100}}$).

As we shall see just below, the multiplicative group $F^\times$ of a finite field F always has a generator. It is often desirable as well to have a presentation of the field $\mathbb{F}_{p^d}$ as $\mathbb{F}_p[x]/(f(x))$, where f is irreducible of degree d, f has just a few nonzero terms, and $[x]$ is itself the generator of the multiplicative group $\mathbb{F}_{p^d}^\times$. It is an interesting research problem to show that such polynomials (for which $[x]$ is a cyclic generator and the polynomial is sparse) exist and to efficiently find them.

The multiplicative group of a finite field

In this section we study the abelian group $F^\times$ for a finite field F. We have seen that some finite abelian groups have a generator and some do not. For example, the group of units in $\mathbb{Z}/(8)$ does not have a generator, but the additive group $\mathbb{Z}/(8)$ does. In group theory, we say that a group that is generated from a single element is *cyclic*, and perhaps it is better for us to use this standard terminology. We now prove the following result.

Theorem 1. *If F is a finite field, then $F^\times$ is a cyclic group.*

Proof. Suppose F has q elements, so that the group $F^\times$ has $q-1$ elements. We have seen that each element of $F^\times$ has order some divisor of $q-1$. In particular, let $N(d)$ be the number of elements of $F^\times$ of order d, so that $N(d) = 0$ unless $d \mid q-1$. We are to show that $N(q-1) > 0$. We first prove that

$$\text{if } N(d) > 0, \text{ then } N(d) = \varphi(d), \tag{1}$$

where φ is Euler's function. (Note that $\varphi(d)$ is the number of integers j in the interval $[1, d]$ that are relatively prime to d.) Recall the complete list of all powers of an element a of order d is $a, a^2, \dots, a^d$, so a has exactly d powers. To see (1), first note that if a has order d, the order of any a^j is quite predictable: it is $d/(j,d)$. Indeed, this follows since $jd/(j,d)$ is the first multiple of j that is also a multiple of d (it is the least common multiple of j and d), and a necessary and sufficient condition for a^{jm} to be the identiy is that $d \mid jm$. Thus, there are at least $\varphi(d)$ elements of order d, namely the powers a^j for $j \in [1, d]$ and j relatively prime to d. Now say, there is some other element b of order d. Note then that b is not of the form a^j for any

j (do you see why?). Consider the polynomial $x^d - 1$ in $F[x]$. Note that a is a root, and in fact each power of a is a root. So $x^d - 1$ has d roots among the powers of a. But it also has b as a root, and b is not a power of a, so the polynomial has at least $d + 1$ roots. But it is impossible for a polynomial over a field to have more roots than its degree, so this contradiction completes our proof of (1).

To complete the proof of Theorem 1, we first establish the following pleasant identity from elementary number theory:

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

Indeed, consider the n fractions $1/n, 2/n, \ldots, n/n$ and reduce each to lowest terms. The reduced fractions have denominators running over the various divisors d of n, and each reduced fraction j/d with j relatively prime to d and $0 < j/d \leq 1$ will occur. Thus, there are $\varphi(d)$ reduced fractions with denominator d, which proves our identity.

Combining our identity for φ and (1), we have

$$q - 1 = \sum_{d|q-1} N(d) = \sum_{\substack{d|q-1 \\ N(d)>0}} N(d) = \sum_{\substack{d|q-1 \\ N(d)>0}} \varphi(d) \leq \sum_{d|q-1} \varphi(d) = q - 1.$$

Since this inequality chain begins and ends with $q - 1$, it must actually be an equality all the way through. But then look at the point where we drop the condition $N(d) > 0$; it must be that no terms are dropped! That is, $N(d) > 0$ for every $d \mid q - 1$, which shows in particular that $N(q - 1) > 0$, and in fact is equal to $\varphi(q - 1)$. Thus, there is a generator and $F^\times$ is a cyclic group. □

The automorphism group of a finite field

An *automorphism* of a field F is an isomorphism of F to itself. That is, it is a function $\phi : F \longrightarrow F$ such that ϕ is one-to-one and onto (a bijection) and

$$\phi(a + b) = \phi(a) + \phi(b), \quad \phi(ab) = \phi(ab)$$

for all $a, b \in F$. The automorphisms of F form a group under the operation of composition of functions. (It is easy to check that if ϕ_1, ϕ_2 are both automorphisms of F, then $\phi_1 \circ \phi_2$ and ϕ_1^{-1} are as well.) For some fields, the automorphism group is quite boring. For example, it is easy to see that the only automorphism of $\mathbb{Q}$ is the identity, and the same holds for the finite fields $\mathbb{F}_p$ with p prime. It is a not-so-trivial exercise to prove that the only automorphism of $\mathbb{R}$ is the identity map. The field $\mathbb{C}$ has a nontrivial automorphism, namely, complex conjugation. Does it have any others? This question actually has its roots in the set-theoretic foundations of mathematics as a whole.

An immediate property of an automorphism ϕ of a field F is that

$$\phi(1) = 1.$$

Indeed, if $\phi(1) = 0$, then ϕ must be the 0-function, and so is not an automorphism. But $\phi(1) = \phi(1 \cdot 1) = \phi(1)^2$, so that $\phi(1)$ is a nonzero root of $x^2 - x$, which means it is 1. Using this, one finds that if n is any integer, then

$$\phi(n \cdot 1) = n \cdot 1.$$

Using this, one sees that if F has characteristic some prime p, then ϕ acts like the identity map on the subfield $\mathbb{F}_p$. (If F does not have prime characteristic, which means that 1 does not have an additive order in F, then $\mathbb{Q}$ is seen to be a subfield of F, and ϕ acts like the identity on this subfield.)

Suppose now that p is prime, and consider the finite field $\mathbb{F}_{p^d}$. Note that if $f(x) \in \mathbb{F}_p[x]$ and $f(a) = 0$, where $a \in \mathbb{F}_{p^d}$, then for any automorphism ϕ of $\mathbb{F}_{p^d}$, we have $f(\phi(a)) = 0$. That is, for any root a of f, we have that $\phi(a)$ is another root. To see this, note that if $f(x) = c_n x^n + c_{n-1}x^{n-1} + \cdots + c_0$, then

$$0 = \phi(0) = \phi(f(a)) = \phi(c_n a^n + \cdots + c_0) = \phi(c_n a^n) + \cdots + \phi(c_0) = c_n \phi(a)^n + \cdots + c_0 = f(\phi(a)).$$

Other than the identity, can we spot any automorphisms of $\mathbb{F}_{p^d}$? Let's try the function which takes an element to its pth power, say $\tau(a) = a^p$ for each $a \in \mathbb{F}_{p^d}$. It is quite easy to show that τ is an automorphism, the only possibly tricky thing being that $\tau(a+b) = \tau(a)+\tau(b)$; and this is the bad student's binomial theorem.

If $d > 1$, then τ, known as the *Frobenius automorphism* of $\mathbb{F}_{p^d}$, is not the identity. Indeed, $\tau(a) = a$ if and only if a is a root of $x^p - x$, and this polynomial has exactly p roots in $\mathbb{F}_{p^d}$ (which comprise the subfield $\mathbb{F}_p$).

Let's try composing τ with itself to get more automorphisms. Write τ^j for $\tau \circ \tau \circ \cdots \circ \tau$, with j copies of τ. If you repeatedly raise something to the pth power, and do so j times, then this is raising to the power p^j. That is, $\tau^j(a) = a^{p^j}$ for each $a \in \mathbb{F}_{p^d}$. Note that this cannot be the identity map if $1 \leq j \leq d-1$, since the polynomial $x^{p^j} - x$ cannot have p^d roots. However, τ^d is indeed the identity map, since we've seen that every element of $\mathbb{F}_{p^d}$ is a root of $x^{p^d} - x$. Thus, τ has order d in the automorphism group.

We have just learned that the field $\mathbb{F}_{p^d}$ has d distinct automorphism, namely τ^j for $j = 1, \ldots, d$, where τ is the Frobenius automorphism. Is this the complete story? Yes, let's see why.

Let $f(x) \in \mathbb{F}_p[x]$ be an irreducible polynomial of degree d. (We have seen they must exist.) We know that f splits into linear factors in $\mathbb{F}_{p^d}[x]$, so let a be a root. We claim that $\tau^j(a)$ for $j = 1, \ldots, d$ describe d different elements of $\mathbb{F}_{p^d}$. Indeed, suppose $\tau^i(a) = \tau^j(a)$ where $1 \leq i < j \leq d$. Let $k = d - j$ and take τ^k of both sides of the equation, getting $\tau^{i+k}(a) = \tau^d(a) = a$. This implies that the polynomial $x^{p^{i+k}} - x$ has a as a root, so must be divisible by $f(x)$, which is the minimal polynomial of a. But then $d = \deg(f) \mid i + k$, which cannot occur since $i + k < d$.

So, quite remarkably, we can now write out the complete factorization of $f(x)$ over $\mathbb{F}_{p^d}$. For f in $\mathbb{F}_p[x]$ a monic irreducible of degree d, with $a \in \mathbb{F}_{p^d}$ as a root, we have

$$f(x) = (x - a)(x - a^p) \ldots (x - a^{p^{d-1}}).$$

(Note that $\tau^d(a) = a^{p^d}$ is the same as a itself.)

Now suppose there is some other automorphism ϕ of $\mathbb{F}_{p^d}$ that is not a power of τ. Since $f(\phi(a)) = 0$, we must have that $\phi(a) = \tau^j(a)$ for some j. But, we have seen that $\mathbb{F}_{p^d}$ is isomorphic to $\mathbb{F}_p[x]/(f(x))$, where we let $[x]$ correspond to $a \in \mathbb{F}_{p^d}$. Thus, viewing ϕ as an automorphism of $\mathbb{F}_p[x]/(f(x))$, we have that $\phi([x])$ is the same as $\tau^j([x])$. But every member of $\mathbb{F}_p[x]/(f(x))$ is of the form $g([x])$ where $g \in \mathbb{F}_p[x]$, so τ^j and ϕ are equal on every member of the field. We have just proved that every automorphism of $\mathbb{F}_{p^d}$ is a power of the Frobenius automorphism τ.

We record what we have learned:

Theorem 2. *If τ is the pth power map, then τ is an automorphism of $\mathbb{F}_{p^d}$. There are exactly d automorphisms of this field, namely τ^j for $j = 1, \ldots, d$, where τ^d is the identity automorphism.*