

Algorithms for Finite Fields

Notes for September 19th and 21st

Salman Butt

September 25, 2006

Recall our setup from last week: $f(x) \in \mathbb{F}_q[x]$ is a square-free polynomial; $u(x) \in \mathbb{F}_q[x]$ is a splitting polynomial for f , so it satisfies

$$u^q \equiv u \pmod{f(x)} \quad u \not\equiv c \pmod{f(x)} \text{ for any } c \in \mathbb{F}_q;$$

$C = \{c \in \mathbb{F}_q : (u - c, f) \neq 1\}$ is the set of splitting values for f and u ; and $h(y) \in \mathbb{F}_q[y]$ is defined to be

$$h(y) = \prod_{c \in C} (y - c).$$

Let us clarify a misleading point from last week: In general, $h_1(y) = \text{Res}_x(u(x) - y, f(x))$ may not be equal to $h(y)$. We proved last time that the zeros of $h_1(y)$ are indeed the elements of C and $\deg(h_1(y)) \leq \deg(f(x))$ by construction. But generally h_1 factors as

$$h_1(y) = \prod_{c \in C} (y - c)^{\alpha_c}$$

where $\alpha_c \geq 1$. Thus h_1 may not have simple roots, though h does.

Example. Let $f(x) = x^4 + 3x^2 + 2 = (x^2 + 1)(x^2 + 2)$ be a polynomial in $\mathbb{F}_7[x]$. $u(x) = x^2$ is a splitting polynomial for f and $C = \{1, -2\}$. (All this was proven in a previous class.) Thus $h(y) = y^2 + 3y + 2$. Let's compute $h_1(y)$. An elementary property of the resultant is that it satisfies

$$\text{Res}(f, g) = a_n^m \prod_{i=1}^n g(x_i)$$

where x_i are the roots of f , a_n is the leading coefficient of f , and $m = \deg(g)$ [See Dummit and Foote, p. 621]. In our case we have $h_1(y) = \text{Res}_x(x^2 - y, f(x)) = f(\sqrt{y})f(-\sqrt{y}) = h(y)^2$.

Theorem 1. $h(y)$ is the monic polynomial of smallest degree such that $f(x)$ divides $h(u(x))$.

Proof. Observe that $I = \{g(y) \in \mathbb{F}_q[x] : f(x)|g(u(x))\}$ is an ideal in $\mathbb{F}_q[x]$, which is a principal ideal domain and hence I is generated by a single element (of minimal degree). Our claim is equivalent to the claim that $I = (h)$. First observe that $h \in I$: recall that $f = g_1 \cdots g_r$ where each g_i is irreducible and that $u \equiv c_i \pmod{g_i}$ for some $c_i \in \mathbb{F}_q$. (For this last claim, recall $u^q \equiv u \pmod{f}$, so $u^q \equiv u \pmod{g_i}$ for all i . Since the g_i 's are irreducible, $\mathbb{F}_q[x]/(g_i)$ is a finite field containing a copy of \mathbb{F}_q as the solutions of $x^q \equiv x \pmod{g_i}$. Thus $u \equiv c_i \pmod{g_i}$ for some $c_i \in \mathbb{F}_q \subseteq \mathbb{F}_q[x]/(g_i)$.) Thus g_i divides $u - c_i$ and hence also $\prod_{c \in \mathbb{F}_q} (u - c) = h(u)$. Thus every g_i divides $h(u)$, so f does as well.

Now assume $I = (k)$ for some $k \in \mathbb{F}_q[y]$, $k \neq h$. Since $h \in I$, we know k divides h , so

$$k(y) = \prod_{c \in C'} (y - c)$$

for $C' \subsetneq C$. So there exists a $c_i \in C$ such that $c_i \notin C'$. Recall that we know $u \equiv c_i \pmod{g_i}$ for some i . We claim for this fixed i , g_i does not divide $k(u)$. Assume this is not true, so g_i divides $\prod_{c \in C'} (u - c)$. But g_i is irreducible and hence a prime in $\mathbb{F}_q[x]$, so g_i divides $u - c$ for some $c \in C'$, so $u \equiv c \pmod{g_i}$, but $u \equiv c_i \pmod{g_i}$. Thus $c = c_i$, but this contradicts the fact that $c_i \notin C'$. Thus for some i , g_i does not divide $k(u)$, hence neither does f , so $k \notin I$, a contradiction. Hence k must be h , so $I = (h)$. \square

Problem: Given a square-free $f(x) \in \mathbb{F}_q[x]$ that splits completely in \mathbb{F}_q (i.e. $f(x)$ divides $x^q - x$), find the (distinct) roots of $f(x)$.

We have a probabilistic algorithm that is quite fast in practice due to Legendre. The idea is to split up \mathbb{F}_q into two disjoint halves and hope that this will induce a splitting of f . After at most $\deg f$ successful splits, we will have found its roots. Here is the procedure:

Algorithm: Suppose first that q is odd. Observe that

$$x^q - x = x(x^{(q-1)/2} - 1)(x^{(q-1)/2} + 1).$$

The last two factors on the right distinguish the squares from the non-squares in \mathbb{F}_q , respectively, and the first factor distinguishes 0. With this in mind, pick a $b \in \mathbb{F}_q$ at random and consider

$$f(x) = (f(x), x - b) \cdot (f(x), (x - b)^{(q-1)/2} - 1) \cdot (f(x), (x - b)^{(q-1)/2} + 1).$$

If this is a non-trivial splitting of f , recurse over these other factors by choosing different b 's. If not, choose a different b and try again.

Now suppose q is even, so say $q = 2^m$. Let $S(x) = x + x^2 + \cdots + x^{2^{m-1}}$. Observe $S(x)(S(x) + 1) = x^q - x$. So for any $a \in \mathbb{F}_q$, $S(a) = \text{Tr}_{\mathbb{F}_q/\mathbb{F}_2}(a) \in \mathbb{F}_2$. Again, select a $b \in \mathbb{F}_q$ at random and consider

$$f(x) = (f(x), S(bx)) \cdot (f(x), S(bx) + 1).$$

If this is a non-trivial splitting of f , again recurse over these factors with new b 's. If not, choose a different b and try again.

Remark 0.1. (q odd.) Computing $(x - b)^{(q-1)/2}$ using modular exponentiation is fast, as is computing the gcd's. Hence the only difficulty is finding a satisfactory b . So how often do we have a “bad” b , i.e. a b for which we get a non-trivial splitting of f ? Suppose $f(x) = \prod_{i=1}^n (x - c_i)$. Then $f(x)$ divides $(x - b)^{(q-1)/2} - 1$ if and only if $x - c_i$ divides $(x - b)^{(q-1)/2} - 1$ for all i if and only if $(c_i - b)^{(q-1)/2} = 1$ for all i . Equivalently, $c_i - b$ is a square for every i . We will prove later that the probability of this event is approximately 2^{-n} where n is the degree of f .

Remark 0.2. (q even.) We will show below that picking a “bad” b in this case also has small probability.

We have the following theorems:

Theorem 2. For $c_1, \dots, c_n \in \mathbb{F}_q$ distinct,

$$\#\{b \in \mathbb{F}_q : b - c_i \text{ is a square for all } i\} = \frac{q}{2^n} + O(n\sqrt{q})$$

Proof. Consider the system of equations

$$x - c_i = y_i^2 \tag{1}$$

for $i = 1, \dots, n$ in the variables x, y_1, \dots, y_n . Any b for which $b - c_i$ is a non-zero square in \mathbb{F}_q for all i gives rise to solutions $x = b, y_i = \pm\sqrt{b - c_i}$. Thus we see that every such $b \neq c_i, \forall i$ yields 2^n solutions to the system (1). Our claim will follow if and only if (1) has $q + O(2^n n\sqrt{q})$ solutions in \mathbb{F}_q . Observe that (1) defines an affine algebraic curve X over \mathbb{F}_q . We have the map $\phi : X \rightarrow \mathbb{A}^1$ given by $(x, y_1, \dots, y_n) \mapsto x$. This extends to a rational map $\bar{X} \rightarrow \mathbb{P}^1$, where \bar{X} is the projective closure of X . (In fact this is a $(2, \dots, 2)$ Galois cover of \mathbb{P}^1 .) The genus of (the normalization of) \bar{X} is $g = 2^{n-2}(n - 3) + 1$, which can be calculated using Hurwitz's formula. Recall the Weil bound:

Lemma 1. If X is a smooth irreducible projective curve over \mathbb{F}_q of genus g , then

$$|\#X(\mathbb{F}_q) - (q + 1)| \leq 2g\sqrt{q}.$$

Thus the estimate $q + O(2^n n\sqrt{q})$ for the number points on X is just the Weil bound. This completes the proof. \square

Theorem 3. For $c_1, \dots, c_n \in \mathbb{F}_q$ distinct,

$$\#\{b \in \mathbb{F}_q : S(bc_i) = 0 \text{ for all } i\} = \frac{q}{2^k} = 2^{m-k}$$

where k is the dimension of the \mathbb{F}_2 -vector space spanned by c_1, \dots, c_n inside \mathbb{F}_q .

Proof. Observe that S is \mathbb{F}_2 -linear: $S(x + y) = S(x) + S(y)$ since we are in characteristic 2. Let $V = \{b \in \mathbb{F}_q : S(bc_i) = 0 \text{ for all } i\}$ be the \mathbb{F}_2 -subspace of \mathbb{F}_q . We want to show that $\dim_{\mathbb{F}_2}(V) = m - k$. We recall the following lemma from field theory:

Lemma 2. *If L/K is a finite separable extension of fields then the map $L \times L \rightarrow K$ given by $(x, y) \mapsto \text{Tr}_{L/K}(xy)$ is a non-degenerate bilinear pairing of K -vector spaces, where $\text{Tr}_{L/K}$ is the trace of L/K .*

Now $\text{Tr}_{\mathbb{F}_q/\mathbb{F}_2}(x) = S(x)$ for all $x \in \mathbb{F}_q$. Observe that V is the orthogonal complement to c_1, \dots, c_n with respect to the trace pairing since $S(bc_i) = 0$ implies b is orthogonal to c_i with respect to this pairing. Thus $\dim_{\mathbb{F}_2}(V) = \dim_{\mathbb{F}_2}(\mathbb{F}_q) - \dim_{\mathbb{F}_2}(V^\perp) = m - k$ where $k = \dim_{\mathbb{F}_2}(\text{span}(c_1, \dots, c_n))$. This completes the proof. \square

Remark 0.3. (q even.) Recall that we have

$$f = \prod_{i=1}^n (x - c_i)$$

and we also have

$$f(x) = (f(x), S(bx))(f(x), S(bx) + 1).$$

We have already discussed the “bad” b ’s arising from the first factor. For the second factor, consider the set $V_1 : \{b \in \mathbb{F}_q : S(bc_i) = 1 \text{ for all } i\}$. Next observe that if there exists a $b_0 \in \mathbb{F}_q$ such that $S(bc_i) = 1$ for all i , then $V_1 = b_0 + V$ since S is \mathbb{F}_2 -linear, so V_1 is just a translation of V by b_0 . In such a case, we see that $\#V_1 = \#V$. A priori though, there may exist no such b_0 . For observe that $V \cap V_1 = \emptyset$, thus if $V = \mathbb{F}_q$, $V_1 = \emptyset$ and we will have no such b_0 .

Now writing f as

$$\begin{aligned} f(x) &= (f(x), S(bx))(f(x), S(bx) + 1) \\ &= \prod_{c_i \in \mathbb{F}_q, S(bc_i)=0} (x - c_i) \prod_{c_i \in \mathbb{F}_q, S(bc_i)=1} (x - c_i), \end{aligned}$$

we see this is a (non-trivial) splitting of f if and only if $b \notin V \cup V_1$ (otherwise one of the products contains all of the c_i ’s, and hence is f). Now $\#(V \cup V_1) = \#V + \#V_1 \leq 2\#V = q/2^{k-1} \leq q/2$ if $k \geq 2$. So if k is at least 2, then at least half of the b ’s are “good.” For $k = 0$, we see that $\{c_1, \dots, c_n\} = \{0\}$ since they span just 0, which corresponds to the polynomial $f(x) = x$, which is already factored. For $k = 1$, we have that all of the c_i ’s are \mathbb{F}_2 -linear combinations of one $c \in \{c_1, \dots, c_n\}, c \neq 0$. This means f divides $x(x - c)$, hence is a quadratic which we can factor easily. These becomes trivial cases, which we check for, and otherwise we use the above algorithm.

Another significant improvement can be made though. We can make this probabilistic algorithm into a deterministic one via the following proposition:

Proposition 1. *Let b_1, \dots, b_m be a basis for \mathbb{F}_q over \mathbb{F}_2 . Then there exists a j such that $b_j \notin V \cup V_1$ so long as $m \geq 2$.*

In fact, thinking of \mathbb{F}_q as $\mathbb{F}_2/(g(x))$ where $g(x) \in \mathbb{F}_2[x]$ is an irreducible polynomial of degree m , we have the natural basis $\{1, \bar{x}, \bar{x}^2, \dots, \bar{x}^{m-1}\}$ where $\bar{\cdot}$ is reduction modulo g .

Proof of Proposition. We argue by contradiction: suppose every b_j is an element of $V \cup V_1$. Then for all j , $S(b_j c_i)$ are all equal. In particular $S(b_j c_1) = S(b_j c_2)$ for all j . So $S(b_j(c_1 - c_2)) = 0$ for all j by linearity of S . So $c_1 - c_2$ is orthogonal to every basis element. But S is a non-degenerate bilinear pairing, so $c_1 - c_2 = 0$, contradicting the assumption that f has distinct roots. \square

Thus for q even, we have a deterministic polynomial time factoring algorithm by picking a b from a basis of $\mathbb{F}_q/\mathbb{F}_2$ instead of a random b .

Remark 0.4. (q odd.) A similar strategy works for $q = p^m$, p a small odd prime. Define

$$S(x) = x + x^p + \dots + x^{p^{m-1}},$$

which is just $\text{Tr}_{\mathbb{F}_q/\mathbb{F}_p}(x)$. Then

$$f(x) = \prod_{j=0}^{p-1} (f(x), S(bx) - j).$$

We want to find a b such that this is a (non-trivial) splitting. As before, there always exists such a b within a basis of $\mathbb{F}_q/\mathbb{F}_p$. This then gives us a deterministic polynomial time algorithm for factoring f in \mathbb{F}_q where $\text{char}(\mathbb{F}_q)$ is small. This algorithm has running time $O((np \log q)^c)$.

Thus we have deterministic polynomial time algorithms for even q and odd q 's with small characteristic (with respect to $n \log q$). This leaves only \mathbb{F}_q with large characteristic, such as \mathbb{F}_p with p large. We will tackle this next week.