Panoramas & Synthèses
36, 2012, p. 1–17

ALGORITHMS FOR FINITE FIELDS

by

Karim Belabas

***Abstract.* —** This series of lectures concentrates on deterministic algorithms for finite fields. The emphasis is not on developing algorithms for practical use, but on viewing the quest for polynomial-time algorithms as a challenge of our structural understanding of finite fields. The topics treated include: representing finite fields, recognizing finite fields, constructing finite fields, constructing maps between finite fields and factoring polynomials.

***Résumé* (Algorithmes pour les corps finis). —** Ce cours traite des algorithmes déterministes relatifs aux corps finis. L'accent n'est pas mis sur des algorithmes efficaces en pratique, mais sur le défi posé par la quête d'algorithmes en temps polynomial à notre compréhension de cette structure élémentaire. Les sujets traités incluent : représenter un corps fini, reconnaître un corps fini, construire un corps fini, construire des applications entre corps finis, et factoriser les polynômes.

1. Construction of finite fields

We start with a theorem originating in Galois (around 1820), but more accurately attributed to E. H. Moore (1893) [**5**].

Theorem 1.1. — *The map*

$$\{\textit{finite fields}\}/\simeq \longrightarrow \{\textit{prime numbers}\} \times \mathbb{Z}_{>0}$$
$$k \longmapsto (\operatorname{char} k, [k : \mathbb{F}_p])$$

is bijective.

2000 *Mathematics Subject Classification.* — 11T5, 11Y16.
***Key words and phrases.* —** Finite fields, Polynomial-time algorithms.

This article is based on lectures given by Hendrik Lenstra at Institut Henri Poincaré (11/2004).

Denote the reverse map by $(p,n) \mapsto \mathbb{F}_{p^n}$. We are concerned with the following *open* problem: is there a polynomial time algorithm that, given a prime number p and a positive integer $n > 0$, constructs an explicit model for $\mathbb{F}_{p^n}$?

First let us make the question more precise. For the concepts borrowed to computer science, we shall content ourselves with the following pseudo-definitions: an *algorithm* is a computer program; a *deterministic* algorithm is forbidden to use a random number generator. Rigorous-minded readers can think of a Turing machine or a Random Access Machine (see [**6**]); both suit our purpose. An algorithm is *polynomial time*, or poly-time for short, if its run-time is polynomially bounded in terms of the combined lengths of its input and output. In the above problem, we ask whether there exists c and C such that:

$$\forall (p,n), \quad \text{run-time}(p,n) < C(n + \log p)^c.$$

Algorithms will be deterministic unless specified otherwise but we later discuss *probabilistic algorithms*, allowed to use a random bit generator. Their *expected run-time* is, for *fixed* input, the average over all possible runs of the program.

Finally, an *explicit model* for $\mathbb{F}_{p^n}$ is a system of n^3 numbers $(a_{i,j,k})_{1 \leqslant i,j,k \leqslant n}$, $a_{i,j,k} \in \mathbb{F}_p$, such that the additive group $\mathbb{F}_p^n$, endowed with multiplication

$$(x_i)_{i \leqslant n} \cdot (y_j)_{j \leqslant n} = \Big(\sum_{i,j} a_{i,j,k} x_i y_j \Big)_{k \leqslant n},$$

becomes a field. Alternatively, it is a system of n numbers $c_0, \dots, c_{n-1} \in \mathbb{F}_p$ such that the polynomial $X^n + c_{n-1}X^{n-1} + \cdots + c_0$ is irreducible in $\mathbb{F}_p[X]$. We shall see in Section §4 that these are equivalent models up to poly-time transformations.

Remark 1.2. — Let $p > 2$, $n = 2$. The problem is equivalent to finding a quadratic non-residue modulo p. Allowing probabilistic algorithms, we solve the problem in an average of two trials, making it totally uninteresting from a mathematical point of view.

The following partial results are known:

Partial #1. — There is a probabilistic algorithm with polynomial expected run-time that, upon input (p,n), constructs $\mathbb{F}_{p^n}$.

Partial #2. — There is an algorithm that, given (p,n), constructs $\mathbb{F}_{p^n}$, such that $\exists c, C, \forall (p,n), \text{run-time} < C(n+p)^c$.

Partial #3. — The Generalized Riemann Hypothesis (GRH) implies that

$$\exists c, C \ \forall (p,n), \quad \text{run-time} < C(n + \log p)^c.$$

The specific brand of GRH needed in our case says that for each number field K, and each complex number s in the half-plane $\Re(s) > 1/2$, we have $\zeta_K(s) \neq 0$. Here, ζ_K is the Dedekind zeta function,

$$\zeta_K(s) = \sum_{\substack{\mathfrak{a} \subset \mathcal{O}_K \\ \mathfrak{a} \text{ non-zero ideal}}} N\mathfrak{a}^{-s}, \quad \Re s > 1,$$

extended to a meromorphic function on $\mathbb{C}$.

In the course of our investigations, we shall prove an explicit form of Moore's uniqueness theorem for finite fields:

Theorem 1.3. — *There is a deterministic poly-time algorithm that, given (p, n) and two models for $\mathbb{F}_{p^n}$, finds a field isomorphism between them.*

This isomorphism is given by an $n \times n$ matrix over $\mathbb{F}_p$. All algorithms whose existence is asserted in this course will be made perfectly effective and explicit. Although the deterministic algorithms we present are not necessarily inefficient, I do not expect that in practice they can compete with the probabilistic algorithms referred to above. Accordingly, we will not estimate the running times of the various algorithms precisely.

2. Preliminaries on finite rings

Proposition 2.1. — *There are poly-time algorithms that, given $m \in \mathbb{Z}_{\geqslant 1}$ and $a, b \in \mathbb{Z}/m\mathbb{Z}$, compute $a \pm b$, $a \times b$, and either an element $d \in \mathbb{Z}/m\mathbb{Z}$ with $ad = b$ or an element d' with $bd' \neq 0 = ad'$.*

Note that in the latter case, d' provides an explicit proof that $a \nmid b$. If $a \neq 0$, it also yields a non-trivial factor $\gcd(d', m)$ of the integer m.

Proof. — We restrict to the division in the special case $b = 1$, and leave the rest as an exercise. From Euclid's algorithm, let $xa+ym = \gcd(a, m)$, $x, y \in \mathbb{Z}$. If $\gcd(a, m) = 1$, set $d := x$, otherwise set $d' := m/\gcd(a, m)$. □

Similarly, there are poly-time algorithms that solve any linear-algebra problem over $\mathbb{Z}/m\mathbb{Z}$ *or* find a pair of zero-divisors $a, b \in \mathbb{Z}/m\mathbb{Z}$, $a, b \neq 0$, $ab = 0$. We call the latter situation a *side exit* in the sequel: it forces us to stop, but also allows us to factor m, hence start again over smaller finite rings. In particular, we can

- solve or decide unsolvability of systems of linear equations $Ax = b$.
- find $\mathbb{Z}/m\mathbb{Z}$-bases for the kernel, co-kernel and image of any group homomorphism $(\mathbb{Z}/m\mathbb{Z})^s \to (\mathbb{Z}/m\mathbb{Z})^t$ given by a $t \times s$ matrix.

A *finite ring* is a finite unital commutative ring R. We have $\mathbb{Z}/m\mathbb{Z} \subset R$ for $m = \operatorname{char} R$. The most important example for us is $R = \mathbb{F}_q[X]/(f)$, for some polynomial $f \in \mathbb{F}_q[X]$. All the R's we will look at will actually satisfy $(R, +) \simeq (\mathbb{Z}/m\mathbb{Z})^t$ for some t (otherwise, we take a side exit). Such rings are represented by a multiplicative tensor $(a_{i,j,k})_{1 \leqslant i,j,k \leqslant t}$, $a_{i,j,k} \in \mathbb{Z}/m\mathbb{Z}$.

Fact. — There are poly-time algorithms for finding 1, performing $+$, $-$, $\times$ in R, and for finding on input $a \in R$ an element c with $ac = 1$, or $c' \neq 0$ such that $ac' = 0$. Likewise, we can do linear algebra over R in poly-time, or take a side exit and find $a, b \in R - \{0\}$, $ab = 0$.

3. Finite commutative $\mathbb{F}_p$-algebras and factorization

Let R be any finite commutative ring, and consider its radical

$$\sqrt{0_R} = \{x \in R \colon \exists n \in \mathbb{Z}_{>0}, x^n = 0\}$$

(this is an R-ideal). We have

$$\begin{aligned}\sqrt{0_R} &= \bigcap_{\substack{\mathfrak{p}\subset R\\ \mathfrak{p}\text{ prime}}} \mathfrak{p} \quad \text{[also true for infinite rings } R\text{, use the axiom of choice]}\\ &= \bigcap_{\substack{\mathfrak{m}\subset R\\ \mathfrak{m}\text{ maximal}}} \mathfrak{m} \quad \text{[since } R \text{ is finite]}\\ &= \prod_{\substack{\mathfrak{m}\subset R\\ \mathfrak{m}\text{ maximal}}} \mathfrak{m} \;=\; \operatorname{Ker}\Big(R \to \prod_{\mathfrak{m}} R/\mathfrak{m}\Big) \quad \text{[since the product is finite]}.\end{aligned}$$

Fact. — If R is a finite commutative ring then $R/\sqrt{0_R}$ is ring-isomorphic to a finite product of fields. More generally, $R \simeq \prod_{\mathfrak{m}} R_{\mathfrak{m}}$, where the $R_{\mathfrak{m}}$ are finite local rings, such that $\sqrt{0_{R_{\mathfrak{m}}}} = \mathfrak{m}$, and $\mathfrak{m}$ runs through the maximal ideals of R.

Can we find (in poly-time) $\sqrt{0_R}$ and the fields appearing in $R/\sqrt{0_R}$? This seems hopeless in general; for instance, take $R = \mathbb{Z}/m\mathbb{Z}$, then

$$\sqrt{0_R} = \operatorname{rad}(m) \cdot R, \quad \text{where } \operatorname{rad}(m) = \prod_{p\,|\,m \text{ prime}} p.$$

It is widely believed that computing the radical $\operatorname{rad}(m)$ is essentially as difficult as factoring the integer m. In any case, no polynomial time algorithm is known at present, that would solve either problem.

From now on, we assume that R is a finite commutative ring of *prime* characteristic. So $(R,+) \simeq (\mathbb{F}_p)^n$, $n = \dim_{\mathbb{F}_p} R$. Define the Frobenius map

$$\begin{aligned} F : R &\longrightarrow R\\ x &\longmapsto x^p.\end{aligned}$$

The Frobenius is an $\mathbb{F}_p$-algebra endomorphism of R; it is computable in poly-time, as a matrix for an $\mathbb{F}_p$-linear map, using repeated squarings in R, see [**2**, §4.3]. Hence

Theorem 3.1. — *Assume that R is a finite commutative ring of prime characteristic. Then $\sqrt{0_R}$ is computable in poly-time.*

Proof. — Since $\sqrt{0}$ is finitely generated and R is commutative, there exists $N > 0$, such that $\sqrt{0}^N = 0$. So we have a decreasing sequence of $\mathbb{F}_p$-subspaces

$$R \supset \sqrt{0} \supset (\sqrt{0})^2 \supset \cdots \supset (\sqrt{0})^N = 0.$$

If N is chosen minimal, the sequence decreases strictly since $\sqrt{0}^m = \sqrt{0}^{m+1}$ for some $m < N$ implies $\sqrt{0}^m = \sqrt{0}^{m+1} = \cdots = \sqrt{0}^N = 0$. Hence $N \leqslant \dim_{\mathbb{F}_p} R = n$. Pick $t \in \mathbb{Z}$ such that $p^t \geqslant n \geqslant N$; then $F^t(x) = x^{p^t}$, hence $\operatorname{Ker} F^t = \sqrt{0}$. □

Lemma 3.2. — *The exact sequence of ring homomorphisms*

$$0 \longrightarrow \sqrt{0} \longrightarrow R \longrightarrow R/\sqrt{0} \longrightarrow 0$$

is split. The splitting map $R/\sqrt{0} \to R$ *associates* $(r+\sqrt{0}) \mapsto F^t s$*, for any* s *satisfying* $F^t r = F^{2t} s$*.*

Proof. — If t is as above ($p^t \geqslant \dim_{\mathbb{F}_p} R$), we must prove that $F^t R \oplus \operatorname{Ker} F^t \simeq R$ as $\mathbb{F}_p$-vector spaces. Indeed, $\sqrt{0} = \operatorname{Ker} F^t = \operatorname{Ker} F^{2t}$, hence $F^t R = F^{2t} R$. So for all $r \in R$, there exists $s \in R$ such that $F^t r = F^{2t} s$, hence $r = F^t s +$an element of $\operatorname{Ker} F^t$. □

The lemma identifies $R/\sqrt{0} \simeq F^t R$ with a subring of R, isomorphic to a product of finite fields:

$$R/\sqrt{0} \simeq \prod_{i=1}^{s} \mathbb{F}_{p^{n_i}}.$$

Notice that $F-1$ acts on both sides, with kernel $(\mathbb{F}_p)^s$ on the right-hand side; hence its kernel is isomorphic to $(\mathbb{F}_p)^s$ on $R/\sqrt{0}$. Of course, $F-1$ restricted to $\sqrt{0} = \operatorname{Ker} F^t$ has trivial kernel.

Corollary 3.3. — $s = \#\{$*maximal ideals in* $R\} = \dim_{\mathbb{F}_p} \operatorname{Ker}(F-1)$ *is computable in poly-time.*

Corollary 3.4. — $R \simeq_{\mathbb{F}_p\text{-vector space}} (\mathbb{F}_p)^n$ *is a field if and only if* $\sqrt{0} = 0$ *and* $\operatorname{rank}_{\mathbb{F}_p}(F-1) = n-1$*. Furthermore,* $\sqrt{0} = 0$ *if and only if* $\operatorname{rank}_{\mathbb{F}_p} F = n$*.*

In particular, if R is a finite ring of prime characteristic p, given by a multiplicative tensor $(a_{i,j,k})$, we can test whether it is a field in poly-time (in $\log p$ and $\dim_{\mathbb{F}_p} R$).

Corollary 3.5. — *There is a poly-time algorithm that, given a finite field* k *and an element* $f \in k[X]$*,* $f \notin k$*, tests whether* f *is irreducible in* $k[X]$*.*

Proof. — Test whether $k[X]/(f)$ is a field. □

Let R be a finite commutative $\mathbb{F}_p$-algebra, and let $\dim_{\mathbb{F}_p} R = n < +\infty$ ($n \neq 0$).

Proposition 3.6. — *There is a poly-time algorithm that, given* R *and* $\alpha \in R$*, determines the minimal polynomial of* α *over* $\mathbb{F}_p$*,* i.e.*, the unique monic polynomial in* $\mathbb{F}_p[X]$ *that generates the kernel of the evaluation morphism* $\mathbb{F}_p[X] \to R$*, that sends* X *to* α*.*

Proof. — Use $\mathbb{F}_p$-linear algebra to determine the least $0 < d \leqslant n$ such that $\alpha^d \in \mathbb{F}_p \cdot 1 + \cdots + \mathbb{F}_p \cdot \alpha^{d-1}$. □

The statement of the proposition holds if we replace $\mathbb{F}_p$ by any subring of R.

Remark 3.7. — For instance, let k be a finite field, and $R = k[X]/(f)$, for some polynomial $f \in k[X] \setminus k$. The minimal polynomial of $\bar{X}$ in $R/\sqrt{0_R}$ is $\prod_{g|f} g$, where g is monic and irreducible in $k[X]$. Hence, given k and f, one determines in poly-time the largest (monic) squarefree divisor of f in $k[X]$.