Constructing Nonresidues in Finite Fields and the Extended Riemann Hypothesis

Johannes Buchmann
Universität des Saarlandes; Fb 14 – Informatik; PF 151150
66041 Saarbrücken; Germany
Victor Shoup
Bellcore; 445 South St.; Morristown, NJ 07960; USA

Abstract We present a new deterministic algorithm for the problem of constructing kth power nonresidues in finite fields $\mathbf{F}_{p^n}$, where p is prime and k is a prime divisor of $p^n - 1$. We prove under the assumption of the Extended Riemann Hypothesis (ERH), that for fixed n and $p \to \infty$, our algorithm runs in polynomial time. Unlike other deterministic algorithms for this problem, this polynomial-time bound holds even if k is exponentially large. More generally, assuming the ERH, in time $(n \log p)^{O(n)}$ we can construct a set of elements that generates the multiplicative group $\mathbf{F}_{p^n}^*$.

To appear in *Mathematics of Computation*, 1996. An extended abstract of this paper appeared in *Proc. 23rd Ann. ACM Symp. on Theory of Computing*, 1991.

This research was done while the second author was a postdoctoral fellow at the University of Toronto.

1991 *Mathematics Subject Classification.* Primary 11Y16.

1 Introduction

Consider the following problem: given a finite field $\mathbf{F}_{p^n}$, where p is prime, and a prime divisor k of $p^n - 1$, construct a kth power nonresidue in $\mathbf{F}_{p^n}$, i.e., an element that is not a perfect kth power of any other element in $\mathbf{F}_{p^n}$.

The problem of constructing nonresidues lies at the heart of many deterministic algorithms for fundamental problems in finite fields. For example, the problem of constructing an irreducible polynomial of given degree over a finite field can be reduced in deterministic polynomial time to the problem of constructing nonresidues (see [22]). Furthermore, many deterministic algorithms for various special cases of the problem of factoring polynomials over finite fields can be viewed as deterministic reductions to the problem of constructing nonresidues (see [2, 6, 12, 19, 20, 13]).

We are therefore interested in the deterministic complexity of constructing nonresidues. The problem of *testing* whether a given α in $\mathbf{F}_{p^n}$ is a kth power nonresidue has a trivial solution: just test if $\alpha^{(p^n-1)/k} \neq 1$. If probabilistic algorithms are allowed, then the problem of constructing nonresidues also has a trivial solution: just choose α in $\mathbf{F}_{p^n}$ at random and test whether it is a kth power nonresidue. However, the deterministic complexity of constructing nonresidues is currently unknown, even under the assumption of the Extended Riemann Hypothesis (ERH). We shall show that for any *fixed* value of n, this problem can be solved in deterministic polynomial time assuming the ERH.

Our main result is as follows.

There exists a deterministic algorithm with the following properties. It takes as input a prime p and a positive integer n, and outputs a model for the finite field $\mathbf{F}_{p^n}$ together with a set $S \subset \mathbf{F}_{p^n}^*$. Under the assumption of the ERH, the running time of the algorithm is $(n \log p)^{O(n)}$ and S is a generating set for $\mathbf{F}_{p^n}^*$.

By a *model* for $\mathbf{F}_{p^n}$ we mean an $\mathbf{F}_p$-basis for $\mathbf{F}_{p^n}$, together with information that tells us how to express the product of any two basis elements in this basis. By a *generating set* for $\mathbf{F}_{p^n}^*$ we mean a set of elements with the property that every element in the multiplicative group $\mathbf{F}_{p^n}^*$ can be written as a power product of elements in this set. Notice that both the running time and correctness of this algorithm depend on the ERH.

Since isomorphisms between finite fields can be computed in deterministic polynomial time (as was proved in [15] without any hypothesis), we can construct a generating set for $\mathbf{F}_{p^n}^*$ in any given model of $\mathbf{F}_{p^n}$ in time $(n \log p)^{O(n)}$, assuming the ERH. Since a generating set always contains a kth power nonresidue, we can construct nonresidues within the same time bound.

The following conjecture seems plausible: there exists an absolute constant C such that for any basis $\theta_1, \ldots, \theta_n$ for $\mathbf{F}_{p^n}$ over $\mathbf{F}_p$, the set of all elements $\sum a_i \theta_i \in \mathbf{F}_{p^n}^*$ with $\max |a_i| \leq (n + \log p)^C$ forms a generating set for $\mathbf{F}_{p^n}^*$.

If this conjecture were true, then we could very simply enumerate all of the elements in this generating set in time $(n \log p)^{O(n)}$. However, it is not known how to prove such a conjecture, even assuming the ERH.

We are able to prove the following, somewhat weaker, statement.

Assume the ERH. There exist absolute constants C and D, and a deterministic algorithm with the following properties. The algorithm takes as input a prime p and a positive integer n. It runs in time $(n \log p)^{O(1)}$, and produces as output a model for $\mathbf{F}_{p^n}$ for which the associated basis $\theta_1, \ldots, \theta_n$

has the property that the set of all elements $\sum a_i\theta_i \in \mathbf{F}_{p^n}^*$ *with* $\max|a_i| \leq Cn^{Dn}(\log p)^{\max(n-1,2)}$ *forms a generating set for* $\mathbf{F}_{p^n}^*$.

Using this result, we can very easily enumerate all of the elements in this generating set in time $(n \log p)^{O(n^2)}$. However, to obtain a running time bound of the form $(n \log p)^{O(n)}$, we need to use an algorithm that is a bit more complicated.

Previous Work

There is a deterministic algorithm that will construct a model for $\mathbf{F}_{p^n}$ together with a generating set for $\mathbf{F}_{p^n}^*$ in time $(pn)^{O(1)}$. Indeed, given p and n, we can construct an irreducible polynomial f of degree n over $\mathbf{F}_p$ deterministically in time $(pn)^{O(1)}$, using the algorithm in [23]. This allows us to represent $\mathbf{F}_{p^n}$ as $\mathbf{F}_p[x]/(f)$. Then, with the help of character sum bounds appearing, for example, in [25], it follows by a standard argument that the set of images in $\mathbf{F}_p[x]/(f)$ of all monic polynomials of degree up to $2\log n/\log p + 1$ forms a generating set, and we can clearly enumerate this set in time $(np)^{O(1)}$. Thus, for small p the problem of constructing a generating set can be solved in deterministic polynomial time unconditionally.

Ankeny's Theorem [3] as sharpened in [4] states that, under the assumption of the ERH, the set of positive integers less than $2(\log p)^2$ generates $\mathbf{F}_p^*$. This result was generalized to $n = 2$ in [25], where it is shown that, assuming the ERH, we can construct in deterministic polynomial time a model for $\mathbf{F}_{p^2}$ together with a generating set for $\mathbf{F}_{p^2}^*$. Thus, for $n = 1$ and $n = 2$, the problem of constructing a generating set can be solved in deterministic polynomial time under the ERH.

With the ERH assumed, the algorithm of Huang [13] as generalized by Evdokimov [11] allows us to deterministically construct a kth power nonresidue in $\mathbf{F}_{p^n}$ in time $k^A \cdot (n \log p)^{O(1)}$ for some positive constant A. The precise value of A is not worked out in [13] or [11], but is certainly at least 1. So for $k = (n \log p)^{O(1)}$ the problem of constructing kth power nonresidues can be solved in deterministic polynomial time under the ERH.

Related to the problem of constructing a generating set is that of searching for a primitive root, i.e., a single element that generates $\mathbf{F}_{p^n}^*$. It is not known how to efficiently test (either deterministically or probabilistically) if a given element in $\mathbf{F}_{p^n}$ is a primitive root (unless the factorization of $p^n - 1$ is known); however, one can still ask the question of how to deterministically enumerate a set that is guaranteed to contain a primitive root.

It is shown in [25] that for any irreducible polynomial $f \in \mathbf{F}_p[X]$ of degree n, there exists a monic polynomial $g \in \mathbf{F}_p[X]$ (itself irreducible) of degree at most $1 + O(\log n/\log p)$ such that $(g \bmod f)$ is a primitive root for $\mathbf{F}_p[X]/(f) \cong \mathbf{F}_{p^n}$. So for small p we can search for a primitive root in polynomial time.

It was shown by Wang [26] that under the ERH, for any p there is a positive integer that is bounded by $(\log p)^{O(1)}$ whose image in $\mathbf{F}_p$ is a primitive root. In [25], it is shown that, assuming the ERH, we can construct in time $(\log p)^{O(1)}$ a model for $\mathbf{F}_{p^2}$ that has an $\mathbf{F}_p$-basis for which there exists a primitive root for $\mathbf{F}_{p^2}$ whose coordinates in this basis are bounded in absolute value by $(\log p)^{O(1)}$.

So for $n = 1$ and $n = 2$, we can search for a primitive root in polynomial time, assuming the ERH. Unfortunately, it does not seem that the techniques of the present paper can extend these results, even to $\mathbf{F}_{p^3}$.

We mention also the recent result of Perel'muter and Shparlinsky [17] which states that for any

$n > 1$ and $\epsilon > 0$, there exists a p_0, depending on n and ϵ, such that for all primes $p > p_0$ and any $\alpha \in \mathbf{F}_{p^n}$ of degree n over $\mathbf{F}_p$, there exists a nonnegative integer $t < p^{1/2+\epsilon}$ with $\alpha + t$ a primitive root for $\mathbf{F}_{p^n}$.

Applications

We mention three applications of our main result. In these applications, n is a *fixed* positive integer, and we assume the ERH.

1. *Taking kth roots in $\mathbf{F}_{p^n}$.* Combining our result with the algorithms in [2], [13] and [18], we can take kth roots in $\mathbf{F}_{p^n}$ in deterministic time $\sqrt{k}$ times a polynomial in the input size.
2. *Factoring polynomials over $\mathbf{F}_p$.* Combining our result with techniques in [24], [6] and [7], we can factor polynomials over $\mathbf{F}_p$ in deterministic time $\sqrt{k}$ times a polynomial in the input size, where k is the largest prime dividing $\Phi_n(p)$, and Φ_n is the nth cyclotomic polynomial.
3. *Constructing primitive roots in $\mathbf{F}_{p^n}$.* Our result implies that, given the prime factorization of $p^n - 1$, we can construct a primitive root for $\mathbf{F}_{p^n}$ in deterministic polynomial time.

Previous to this work, these statements had been proven only for the special cases $n = 1$ and $n = 2$.

Overview

If $n = 1$ or $p \mid n$ (and so in particular $p \leq n$), the problem of constructing a generating set can be solved by results mentioned previously, so we will assume that $n > 1$ and $p \nmid n$.

In §2, we describe our model for $\mathbf{F}_{p^n}$ and how to construct it. We represent $\mathbf{F}_{p^n}$ as $\mathbf{O}/p\mathbf{O}$, where $\mathbf{O}$ is the ring of integers of a certain number field K, which is a Galois extension of $\mathbf{Q}$ of degree n contained in $\mathbf{R}$. The constructions in this section rely on the ERH. Each element of $\mathbf{O}$ is represented as a *coordinate vector*, contained in $\mathbf{Z}^n$, with respect to a certain integral basis.

Any element $\alpha \in K$ corresponds to a *conjugate vector*, contained in $\mathbf{R}^n$, whose components consist of the images of α under each of the n automorphisms on K. In §3, we discuss the relationship between coordinate and conjugate vectors.

In §4, we show that there is a set of elements of $\mathbf{O}$ whose conjugate vectors lie in a certain geometrically defined region of $\mathbf{R}^n$ and whose images in $\mathbf{O}/p\mathbf{O}$ form a generating set. This relies on the ERH.

In §5, we use the results of §3 to derive an algorithm that enumerates all of the coordinate vectors of the elements of $\mathbf{O}$ whose conjugate vectors lie in the region of $\mathbf{R}^n$ given in §4, and thus enumerates a generating set.

In §6, we briefly indicate an alternative method for constructing a generating set, also based on the methods of §§3 and 4, which is faster, but much less elegant and also less space-efficient.

Before continuing, we define some terms.

Let $R \subset S$ be rings, where S is a free R-module with basis $s_1, \ldots, s_m$. Then by a *multiplication table* for this basis we mean a collection $\{a_{ijk} : 1 \leq i, j, k \leq m\}$ of m^3 elements in R such that for $1 \leq i, j \leq m$

$$s_i s_j = \sum_{k=1}^{m} a_{ijk} s_k.$$

For a finite field $\mathbf{F}_{p^n}$, where p is prime, by a *model* for this field we mean a multiplication table for some $\mathbf{F}_p$-basis for $\mathbf{F}_{p^n}$. Moreover, the entries in this table are integers representing residue classes modulo p. This definition of a model for a finite field comes from [15] (in that paper the term "explicit data" is used, rather than "model").

By the ERH we mean the following assertion: the Dedekind zeta-function of any number field has no zeros in the half-plane $\mathrm{Re}(s) > 1/2$. We refer the reader to [4] for more on the ERH.

All statements of running times in this paper are in terms of bit operations.

2 Constructing a model

We now describe the model for $\mathbf{F}_{p^n}$ that we will use in the rest of the paper. If the ERH is true, this model can be quickly constructed, and it will also enjoy certain properties that will be exploited later.

Fact 1 *Let p be a prime not dividing n. Let q be the least prime satisfying the conditions*

$$q \equiv 1 \pmod{2n}, \tag{1}$$

and

$$(\frac{q-1}{f}, n) = 1, \tag{2}$$

where f is the multiplicative order of p modulo q. Then, such a q exists, and if the ERH is true,

$$q = O(n^4(\log(np))^2). \tag{3}$$

Proof. This is proved in [1]. □

Let q be defined as in Fact 1, and let $L = \mathbf{Q}(\xi)$, where ξ is a complex primitive qth root of unity. Then L is a cyclic extension of $\mathbf{Q}$ of degree $q-1$. By (1), L contains a unique subfield K of degree n over $\mathbf{Q}$. Let $\ell = [L : K] = (q-1)/n$. Moreover, ℓ is even, and so K is a subfield of the real numbers (this will be technically convenient, but is not strictly necessary).

Let Δ denote the absolute value of the discriminant of K. Then it is known that $\Delta = q^{n-1}$.

Let $\mathbf{O}$ be the ring of algebraic integers in K. Condition (2) means that $p\mathbf{O}$ is a prime ideal in $\mathbf{O}$. Thus $\mathbf{O}/p\mathbf{O}$ is a finite field of order p^n.

We denote by $\alpha \mapsto \overline{\alpha}$ the residue class map from $\mathbf{O}$ to $\mathbf{O}/p\mathbf{O} = \overline{\mathbf{O}}$.

Let $U = \mathbf{O}^*$, the group of units of $\mathbf{O}$.

Let $T_{L/K}$ be the trace from L to K, and let $\omega = T_{L/K}(\xi)$. Then $K = \mathbf{Q}(\omega)$. The Galois group $G(K/\mathbf{Q})$ is cyclic of order n, and so is isomorphic to the additive group of $\mathbf{Z}/n\mathbf{Z}$. For a residue class $(i \bmod n)$, we denote the corresponding automorphism by $x \mapsto x^{(i)}$. The set

$$\mathbf{\Omega} = \{\omega^{(0)}, \ldots, \omega^{(n-1)}\}$$

is an integral basis for $\mathbf{O}$. Let $\mathbf{M}$ be the multiplication table for this basis.

We shall take as our model of $\mathbf{F}_{p^n}$ the multiplication table $\overline{\mathbf{M}}$, obtained by reducing the entries of $\mathbf{M}$ modulo p.

Subsequent algorithms will take as input the following data describing the field K:

$$\text{the prime } q\text{, and the multiplication table } \mathbf{M}. \tag{4}$$

Fact 2 *The entries in* $\mathbf{M}$ *are bounded by* $\Delta^{O(1)}$ *in absolute value. Furthermore, there is an algorithm that takes as input* p *and* n *as in Fact 1, and produces as output the data* (4) *in time* $q \cdot (\log \Delta)^{O(1)}$*, which is* $(n \log p)^{O(1)}$ *under the assumption of the ERH.*

Proof. For a proof of this, see [5]. $\square$

3 The dual basis

An element $\alpha \in \mathbf{O}$ is represented by the coordinate vector $(a_0, \ldots, a_{n-1}) \in \mathbf{Z}^n$, where $\alpha = \sum_i a_i \omega^{(i)}$. Corresponding to α is its conjugate vector $(\alpha^{(0)}, \ldots, \alpha^{(n-1)}) \in \mathbf{R}^n$. The purpose of this section is to relate coordinate and conjugate vectors. To this end, we use the notion of the dual basis.

For any $\mathbf{Q}$-basis $\theta_1, \ldots, \theta_n$ for K, its *dual basis* $\theta_1^*, \ldots, \theta_n^* \in K$ is determined by the relations

$$T_{K/\mathbf{Q}}(\theta_i \theta_j^*) = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j, \end{cases}$$

where i and j each run from 1 to n.

The next theorem states some properties of the dual basis of $\mathbf{\Omega}$ that will be needed later: first, it gives an explicit formula for the dual basis; second, it gives an explicit linear transformation on $\mathbf{R}^n$ which sends conjugate vectors to coordinate vectors; and third, it shows that this linear transformation does not increase the max-norm of a vector. Before stating the theorem, we need some notation. For a vector $x \in \mathbf{R}^{n\times 1}$, $x = (x_1, \ldots, x_n)^T$, let

$$\|x\| = \max_i |x_i|$$

be the max-norm of x. For a matrix $M \in \mathbf{R}^{n\times n}$, let

$$\|M\| = \sup_{x \neq 0} \frac{\|Mx\|}{\|x\|}.$$

If $M = (m_{ij})$, then

$$\|M\| = \max_i \sum_j |m_{ij}|.$$

Theorem 3 1. *Let*

$$\lambda = \frac{\omega - \ell}{q}.$$

Then $\lambda^{(0)}, \ldots, \lambda^{(n-1)}$ *is the the dual basis of* $\omega^{(0)}, \ldots, \omega^{(n-1)}$.

2. *Let* $A \in \mathbf{R}^{n\times n}$ *be the matrix*

$$A = \begin{pmatrix} \lambda^{(0)} & \lambda^{(1)} & \ldots & \lambda^{(n-2)} & \lambda^{(n-1)} \\ \lambda^{(1)} & \lambda^{(2)} & \ldots & \lambda^{(n-1)} & \lambda^{(0)} \\ \vdots & \vdots & & \vdots & \vdots \\ \lambda^{(n-1)} & \lambda^{(0)} & \ldots & \lambda^{(n-3)} & \lambda^{(n-2)} \end{pmatrix}.$$

For any $\alpha \in K$ expressed as

$$\alpha = \sum_{i=0}^{n-1} a_i \omega^{(i)} \qquad (a_i \in \mathbf{Q}),$$

we have

$$A \begin{pmatrix} \alpha^{(0)} \\ \alpha^{(1)} \\ \vdots \\ \alpha^{(n-1)} \end{pmatrix} = \begin{pmatrix} a_0 \\ a_1 \\ \vdots \\ a_{n-1} \end{pmatrix}.$$

3. *Let A be the matrix defined above. Then*

$$\|A\| = 1.$$

Proof. Recall that $L = \mathbf{Q}(\xi)$, where ξ is a primitive qth root of unity.

First, we claim that for $0 \leq i \leq n-1$

$$T_{K/\mathbf{Q}}(\omega \cdot \omega^{(i)}) = \begin{cases} q - \ell & \text{if } i = 0, \\ -\ell & \text{if } i \neq 0. \end{cases} \tag{5}$$

To prove this, we use the following two easily derived facts:

$$T_{K/\mathbf{Q}}(\alpha) = \ell^{-1} T_{L/\mathbf{Q}}(\alpha) \qquad \text{for } \alpha \in K, \tag{6}$$

$$T_{L/\mathbf{Q}}(\xi^k) = \begin{cases} q - 1 & \text{if } k \equiv 0 \bmod q, \\ -1 & \text{otherwise.} \end{cases} \tag{7}$$

Now, let H be the subgroup of order ℓ in $\mathbf{Z}_q^*$. Then

$$\omega = \sum_{h \in H} \xi^h,$$

and so

$$\omega \cdot \omega^{(i)} = \sum_{\substack{h \in H \\ h' \in H'}} \xi^{h'-h},$$

where H' is a coset of H in $\mathbf{Z}_q^*$. If $i = 0$, then $H' = H$ and $h' - h \equiv 0 \pmod q$ for exactly ℓ pairs (h, h'); otherwise, $h' - h \not\equiv 0 \pmod q$ for all pairs (h, h'). One then obtains (5) from (6), (7), and a simple calculation.

Next, we set $\lambda = (\omega - 1)/q$, and show that $\lambda^{(0)}, \ldots, \lambda^{(n-1)}$ is dual to $\omega^{(0)}, \ldots, \omega^{(n-1)}$. It will suffice to show that for $0 \leq i \leq n-1$

$$T_{K/\mathbf{Q}}(\lambda \cdot \omega^{(i)}) = \begin{cases} 1 & \text{if } i = 0, \\ 0 & \text{if } i \neq 0. \end{cases} \tag{8}$$

To prove (8), rewrite the left-hand side as

$$q^{-1}[T_{K/\mathbf{Q}}(\omega \cdot \omega^{(i)}) - \ell T_{K/\mathbf{Q}}(\omega^{(i)})],$$

and then (8) follows from a simple calculation, making use of (5) and the fact that $T_{K/\mathbf{Q}}(\omega) = -1$.

This proves assertion (1) of the theorem.

Now let $\alpha \in K$ be expressed as

$$\alpha = \sum_{i=0}^{n-1} a_i \omega^{(i)} \qquad (a_i \in \mathbf{Q}).$$

Then for $0 \le i \le n-1$, we have

$$\begin{aligned} \sum_{j=0}^{n-1} \lambda^{(i+j)} \alpha^{(j)} &= T_{K/\mathbf{Q}}(\lambda^{(i)} \alpha) \\ &= \sum_{j=0}^{n-1} a_j T_{K/\mathbf{Q}}(\lambda^{(i)} \omega^{(j)}) \\ &= a_i. \end{aligned}$$

This proves assertion (2) of the theorem.

To prove assertion (3) of the theorem, first note that for $0 \le i < n$, $\omega^{(i)}$ is a real number, and as it is a sum of ℓ distinct roots of unity, $\omega^{(i)} < \ell$. Thus, $\lambda^{(i)} < 0$, and so it follows that

$$\|A\| = \sum_{0 \le i < n} |\lambda^{(i)}| = -\sum_{0 \le i < n} \lambda^{(i)} = -T_{K/\mathbf{Q}}(\lambda).$$

It is easily seen that the trace $T_{K/\mathbf{Q}}(\lambda)$ is -1, and so $\|A\| = 1$. □

4 A geometric theorem

In this section, assuming the ERH, we show that the images of elements in $\mathbf{O}$ whose conjugate vectors lie in a certain geometrically defined region of $\mathbf{R}^n$ generate the group $(\overline{\mathbf{O}})^*$.

Let $\kappa = 12(\log(\Delta^2 p^n))^2$, and define

$$\mathbf{O}^{(p)} = \mathbf{O} \backslash p\mathbf{O}.$$

The main theorem of this section is the following.

Theorem 4 *Assume the ERH.*

1. *Let*

$$T_1 = \{a \in \mathbf{O}^{(p)} : \prod_{i=0}^{n-1} \max(1, |a^{(i)}|) \le \Delta\},$$

and

$$T_2 = \{a \in \mathbf{O}^{(p)} : \max_{0 \le i < n} |a^{(i)}| \le (\kappa\Delta)^{1/n}\}.$$

Then the image of $T_1 \cup T_2$ in $(\overline{\mathbf{O}})^*$ is a generating set.

2. *Let*

$$T_3 = \{\sum_i a_i\omega^{(i)} \in \mathbf{O}^{(p)} : \max_{0\le i<n} |a_i| \le \max(\Delta^{1/2}, (\kappa\Delta)^{1/n})\}.$$

Then the image of T_3 *in* $(\overline{\mathbf{O}})^*$ *is a generating set.*

The second statement of the theorem immediately gives us a trivial algorithm for enumerating a generating set for $(\overline{\mathbf{O}})^*$: just list all elements in $\mathbf{O}$ of the form $\sum_i a_i\omega^{(i)}$, where each a_i is an integer bounded in absolute value by $\max(\Delta^{1/2}, (\kappa\Delta)^{1/n})$. The quantity $\max(\Delta^{1/2}, (\kappa\Delta)^{1/n})$ is bounded by $Cn^{Dn}(\log p)^{\max(n-1,2)}$ for constants C and D.

While this is a very simple algorithm, its running time is

$$(n \log p)^{O(n^2)},$$

which is not of the desired form $(n\log p)^{O(n)}$. In the next section, we shall use the first statement of Theorem 4 to obtain an algorithm that does run in time $(n\log p)^{O(n)}$.

To prove this theorem, we require some facts relating to the ray class group and to the theory of reduced ideals.

The ray class group

We recall some definitions and facts concerning the ray class group mod p; we refer the reader to [14, §4.1] for background. We denote by $\mathbf{I}$ the group of nonzero (fractional) ideals in K, and $i : K^* \to \mathbf{I}$ is the map that sends $\alpha \in K^*$ to the principal ideal $\alpha\mathbf{O}$. We define

$$\begin{aligned} K^{(p)} &= \{a/b : a, b \in \mathbf{O}^{(p)}\}, \\ K^{(p,1)} &= \{a/b : a, b \in \mathbf{O}^{(p)},\ a \equiv b \bmod p\}. \end{aligned}$$

The domain of the residue class map $\mathbf{O} \to \mathbf{O}/p\mathbf{O}$ extends in a canonical way from $\mathbf{O}$ to $K^{(p)}$.

We let $\mathbf{I}^{(p)}$ denote the subgroup of $\mathbf{I}$ consisting of those ideals that are prime to p. The quotient group $\mathbf{I}^{(p)}/i(K^{(p,1)})$ is a finite group called the *ray class group of* K *mod* p.

We are mainly interested in the subgroup $i(K^{(p)})/i(K^{(p,1)})$. The connection between this group and the group $(\overline{\mathbf{O}})^*$ is given by the following easily derived fact. Recall that U denotes the group of units in $\mathbf{O}$.

Fact 5 *The map*

$$\begin{aligned} i(K^{(p)}) &\to (\overline{\mathbf{O}})^*/\overline{U}, \\ \alpha\mathbf{O} &\mapsto \overline{\alpha}\cdot\overline{U} \end{aligned}$$

is a surjective group homomorphism with kernel $i(K^{(p,1)})$.

We will use the following fact, which requires the assumption of the ERH. For an ideal $\mathbf{A}$, $N(\mathbf{A})$ denotes its norm. Let κ be as defined at the beginning of this section.

Fact 6 *Under the ERH, the group* $\mathbf{I}^{(p)}/i(K^{(p,1)})$ *is generated by the images of those prime ideals* $\mathbf{P} \in \mathbf{I}^{(p)}$ *with* $N(\mathbf{P}) \le \kappa$.

Proof. This is Theorem 4 in [4], specialized to our situation. □

Reduced ideals

We need to make use of the theory of reduced ideals, as described in [8]. Let $\mathbf{A}$ be a fractional ideal in K. A nonzero element $\alpha \in \mathbf{A}$ is called a *minimum* of $\mathbf{A}$ if

$$\{\beta \in \mathbf{A} : |\beta^{(i)}| < |\alpha^{(i)}| \text{ for } 0 \leq i < n\} = \{0\}.$$

The ideal $\mathbf{A}$ is called *reduced* if 1 is a minimum of $\mathbf{A}$.

Every ideal $\mathbf{A}$ contains a minimum, since, e.g., any nonzero element in $\mathbf{A}$ of minimal norm is a minimum of $\mathbf{A}$. If α is a minimum of $\mathbf{A}$, then the ideal $\alpha^{-1}\mathbf{A}$ is reduced. Therefore, every ideal class (in the ordinary class group $\mathbf{I}/i(K^*)$) contains a reduced ideal.

Two minima α_1 and α_2 of $\mathbf{A}$ are called *neighbors* if

$$\{\beta \in \mathbf{A} : |\beta^{(i)}| < \max(|\alpha_1^{(i)}|, |\alpha_2^{(i)}|) \text{ for } 0 \leq i < n\} = \{0\}.$$

If $\mathbf{A}$ is a reduced ideal, and α is a neighbor of 1 in $\mathbf{A}$, then the reduced ideal $\alpha^{-1}\mathbf{A}$ is called a *neighbor* of $\mathbf{A}$. This neighbor relation is easily seen to be symmetric, and we write $\mathbf{A}_1\mathbf{N}\mathbf{A}_2$ when $\mathbf{A}_1$ and $\mathbf{A}_2$ are neighbors. The symbol $\mathbf{N}^*$ denotes the transitive closure of the neighbor relation.

Fact 7

1. *For two reduced ideals $\mathbf{A}_1$ and $\mathbf{A}_2$, we have $\mathbf{A}_1\mathbf{N}^*\mathbf{A}_2$ if and only if $\mathbf{A}_1$ and $\mathbf{A}_2$ belong to the same ideal class.*

2. *For any ideal class, consider the set S of elements in K that are neighbors of 1 in some reduced ideal in the class. Then the subgroup of K^* generated by S contains the unit group U.*

3. *The number of reduced ideals in any one ideal class is bounded by $2^{O(n)}R$, where R is the regulator of K.*

4. *If α is a minimum of an ideal $\mathbf{A}$, then the number of $\beta \in \mathbf{A}$ that are neighbors of α is bounded by $(\log \Delta)^{O(n)}$.*

Proof. These assertions are proved in [8] and [9]. □

Lemma 8

1. *If $\mathbf{A}$ is a reduced ideal, then $\mathbf{A}$ is prime to p.*

2. *If $\mathbf{A}$ is an ideal that is prime to p, and α is a minimum of $\mathbf{A}$, then α is prime to p.*

Proof. Recall that p is a rational prime that is inert in K. To prove the first assertion, suppose $\mathbf{A}$ is a reduced ideal such that $\mathbf{A} = p^e\mathbf{B}$, where $e \neq 0$ and $\mathbf{B}$ is prime to p. Then $p^e \in \mathbf{A}$, and since $1 \in \mathbf{A}$, we must have $e < 0$. However, this contradicts the fact that 1 is a minimum of $\mathbf{A}$.

The second assertion follows from the first, since $\alpha^{-1}\mathbf{A}$ is reduced. □

Fact 9 *Let G be a group, P a set of generators for G, and H a subgroup of G. Let F be a subset of G such that $G = FH$. Then H is generated by its intersection with $F^{-1}PF = \{x^{-1}yz : x, z \in F, y \in P\}$.*

Proof. This is Lemma 6.3 in [16]. □

Theorem 10 *Assume the ERH. Let S_1 be the set of all $\alpha \in K$ such that α is a neighbor of* 1 *in some reduced ideal. Let $\mathbf{A}_1, \ldots, \mathbf{A}_h$ be a set of ideals in $\mathbf{I}^{(p)}$ that form a complete system of representatives for the class group, and let T be the set consisting of all prime ideals $\mathbf{P} \in \mathbf{I}^{(p)}$ with $N(\mathbf{P}) \leq \kappa$, together with $\mathbf{O}$. Let S_2 be any set of elements in K obtained by choosing, for each $1 \leq i \leq h$ and for each $\mathbf{P} \in T$, a minimum of the ideal $\mathbf{A}_i \cdot \mathbf{P}$. Then each element of $S_1 \cup S_2$ is prime to p, and its image in $(\overline{\mathbf{O}})^*$ is a generating set.*

Proof. First of all, the assertion that the elements of $S_1 \cup S_2$ are prime to p follows immediately from Lemma 8.

To prove the theorem, we shall apply Fact 9 with $G = \mathbf{I}^{(p)}/i(K^{(p,1)})$ and $H = i(K^{(p)})/i(K^{(p,1)})$. By Fact 6 the set $P \subset G$ consisting of the images in G of the prime ideals $\mathbf{P} \in \mathbf{I}^{(p)}$ with $N(\mathbf{P}) \leq \kappa$ is a set of generators for G. Furthermore, the set $F \subset G$ consisting of the images of the ideals $\mathbf{A}_1, \ldots, \mathbf{A}_h$ satisfies the property $G = FH$. By Fact 9, the group H is generated by the images of those principal ideals $\alpha\mathbf{O}$ that can be written as

$$\alpha\mathbf{O} = \mathbf{A}_i\mathbf{P}\mathbf{A}_j^{-1}, \tag{9}$$

where $1 \leq i, j \leq h$ and $\mathbf{P}$ is a prime ideal in T. Let $\beta \in S_2$ be a minimum of $\mathbf{P}\mathbf{A}_i$ and $\gamma \in S_2$ be a minimum of $\mathbf{A}_j$, so that $\beta^{-1}\mathbf{P}\mathbf{A}_i$ and $\gamma^{-1}\mathbf{A}_j$ are reduced ideals belonging to the same ideal class. By Fact 7(1) there exists an element $\delta \in K$ that can be expressed as a power product of elements in S_1 such that

$$\beta^{-1}\mathbf{P}\mathbf{A}_i = \delta\gamma^{-1}\mathbf{A}_j. \tag{10}$$

Combining (9) and (10), we obtain

$$\alpha\mathbf{O} = \beta\delta\gamma^{-1}\mathbf{O}.$$

Thus, the image of the set $S_1 \cup S_2$ in H is a generating set. From Fact 5, it follows that the image of $S_1 \cup S_2 \cup U$ in $(\overline{\mathbf{O}})^*$ is a generating set. But by Fact 7(2), U is already generated by S_1, and so the image of $S_1 \cup S_2$ in $(\overline{\mathbf{O}})^*$ is already a generating set. □

Proof of Theorem 4. For $\alpha \in K$, let

$$M(\alpha) = \prod_{i=0}^{n-1} \max(1, |\alpha^{(i)}|).$$

This quantity is sometimes called the *measure* of α. Consider any reduced ideal $\mathbf{A}$. Since $1 \in \mathbf{A}$ it follows that $\mathbf{A}$ is of the form $\mathbf{A} = \mathbf{B}^{-1}$, where $\mathbf{B}$ is an integral ideal. Furthermore, since 1 is a minimum of $\mathbf{A}$, by Minkowski's convex body theorem, $N(\mathbf{B}) \leq \Delta^{1/2}$. Now suppose that α is a neighbor of 1 in $\mathbf{A}$. Again, by Minkowski, we must have $M(\alpha) \leq \Delta^{1/2}N(\mathbf{A})$. Also by Minkowski, we can choose a nonzero $b \in \mathbf{B}$ (prime to p) such that $|b^{(i)}| \leq (\Delta^{1/2}N(\mathbf{B}))^{1/n}$ for $0 \leq i < n$. For this b, we have

$$M(b) \leq \Delta^{1/2}N(\mathbf{B}) \leq \Delta,$$

and

$$M(\alpha b) \leq M(\alpha)M(b) \leq \Delta^{1/2}N(\mathbf{A}) \cdot \Delta^{1/2}N(\mathbf{B}) = \Delta.$$

Since $\alpha = (\alpha b)/b$, and αb and b are in $\mathbf{O}^{(p)}$, it follows that the set S_1 in Theorem 10 is contained in the subgroup of K^* generated by the set T_1.

By Minkowski, we can choose integral ideals $\mathbf{A}_1, \ldots, \mathbf{A}_h$ whose norms are at most $\Delta^{1/2}$ forming a complete system of representatives for the class group. Consider the set T of ideals from Theorem 10. By Minkowski, for $1 \leq j \leq h$, and any $\mathbf{P} \in T$, there exists a minimum $a \in \mathbf{A}_j\mathbf{P}$ with

$$|a^{(i)}| \leq (\Delta^{1/2} N(\mathbf{A}_j) N(\mathbf{P}))^{1/n} \leq (\Delta\kappa)^{1/n} \qquad (0 \leq i < n).$$

Thus, we can take for the set S_2 in Theorem 10 a subset of T_2.

The first assertion of the theorem now follows immediately from Theorem 10.

Now consider the second assertion. As above, each element $\alpha \in S_1$ that is a neighbor of 1 in the reduced ideal $\mathbf{A} = \mathbf{B}^{-1}$ satisfies $M(\alpha) \leq \Delta^{1/2} N(\mathbf{A})$. We can express α in the basis $\mathbf{\Omega}$ as

$$\alpha = \sum_i \frac{a_i}{N(\mathbf{B})} \omega^{(i)},$$

where each of the a_i's are integers. It follows from Theorem 3 that $|a_i/N(\mathbf{B})| \leq \Delta^{1/2} N(\mathbf{A})$, i.e., $|a_i| \leq \Delta^{1/2}$. So we see that $\sum_i a_i \omega^{(i)} \in T_3$. Also, it is clear that $N(\mathbf{B}) \in T_3$, since $N(\mathbf{B}) \leq \Delta^{1/2}$. Thus, α can be written as the ratio of two elements of T_3.

It also follows from Theorem 3 that $T_2 \subset T_3$. This proves the second assertion. $\square$

5 Constructing a generating set

In this section, we use the results of §3 to derive an efficient algorithm that enumerates the coordinate vectors of all elements of $\mathbf{O}$ whose conjugate vectors lie in the region of $\mathbf{R}^n$ defined in §4, and thus enumerates a generating set for $(\overline{\mathbf{O}})^*$.

For $\delta > 0$ and $x \in \mathbf{R}^n$, consider the n-dimensional box

$$\mathcal{B}(x, \delta) = \{y \in \mathbf{R}^n : \|x - y\| \leq \delta\}.$$

For $\delta \geq 1$, consider the region

$$\mathcal{R}(\delta) = \left\{(x_0, \ldots, x_{n-1})^T \in \mathbf{R}^n : \prod_{i=0}^{n-1} \max\{1, |x_i|\} \leq \delta\right\}.$$

Let A be the matrix in Theorem 3(2), which sends conjugate vectors to coordinate vectors, and for $S \subset \mathbf{R}^n$ let $A[S] = \{Ax : x \in S\}$.

Under the assumption of the ERH, Theorem 4, together with Theorem 3, implies that the nonzero images in $\overline{\mathbf{O}}$ of the elements in the set

$$\{\sum_{i=0}^{n-1} a_i \omega^{(i)} \in \mathbf{O} : (a_0, \ldots, a_{n-1})^T \in A[\mathcal{R}(\Delta) \cup \mathcal{B}(0, (\kappa\Delta)^{1/n})] \cap \mathbf{Z}^n\}$$

generate the group $(\overline{\mathbf{O}})^*$. Thus, we have reduced our problem of constructing a generating set to that of enumerating all elements in the set

$$A[\mathcal{R}(\Delta) \cup \mathcal{B}(0, (\kappa\Delta)^{1/n})] \cap \mathbf{Z}^n.$$

Consider the cube

$$\mathcal{Y} = \{(x_0, \ldots, x_{n-1})^T \in \mathbf{R}^n : 0 \le x_i \le 1/2\}.$$

Then the set of translates

$$\mathbf{U} = \{\mathcal{Y} + x : x \in \frac{1}{2}\mathbf{Z}^n\}$$

partitions $\mathbf{R}^n$ into cubes whose sides are of length $1/2$.

Algorithm 11 This algorithm takes as input an integer $B \ge 1$ and the data (4) describing $\mathbf{O}$. It produces as output a set of vectors in $\mathbf{Z}^n$ that contains $A[\mathcal{R}(B)] \cap \mathbf{Z}^n$.

1. Compute rational approximations $\hat{\lambda}_0, \ldots, \hat{\lambda}_{n-1}$ such that

   $$\left|\hat{\lambda}_i - \lambda^{(i)}\right| \le \epsilon,$$

   where

   $$\epsilon = \frac{1}{8(B + 1/4)n}.$$

   Let $\hat{A}$ be the approximation to the matrix A in Theorem 3 obtained by replacing each entry $\lambda^{(i)}$ of A by the approximation $\hat{\lambda}_i$.

2. For each cube $\mathcal{Z}$ in the set $\mathbf{U}$ that intersects $\mathcal{R}(B)$, do the following:

   (a) Let $z \in \mathbf{R}^n$ be the center of $\mathcal{Z}$.

   (b) Compute $x = \hat{A}z$.

   (c) Round each coordinate of x to the nearest integer, obtaining the vector $a \in \mathbf{Z}^n$.

   (d) Output a.

□

Theorem 12 *Algorithm 11 works correctly as specified. It can be executed in time*

$$B(n + \log B)^{O(n)} + (\log \Delta + \log B)^{O(1)}.$$

Proof. We first argue that the algorithm is correct.

We can write $\hat{A} = A + D$, where D is a matrix whose entries are bounded by ϵ in absolute value.

Let $u \in A[\mathcal{R}(B)] \cap \mathbf{Z}^n$. We show that u is one of the vectors output by the algorithm.

Let $u = Ay$, with $y \in \mathcal{R}(B)$. Let $\mathcal{Z} \in \mathbf{U}$ be the cube containing y, and let z be its center. We have

$$\begin{aligned} \|\hat{A}z - u\| &= \|(A + D)z - Ay\| \\ &\le \|A(z - y)\| + \|Dz\| \\ &\le 1/4 + \epsilon(B + 1/4)n \\ &= 3/8. \end{aligned}$$

It follows that when the coordinates of the vector $x = \hat{A}z$ are rounded to integers, the resulting vector a must equal u.

We next prove the running time bound.

For step 1, we first obtain the minimal polynomial $g \in \mathbf{Q}[X]$ of λ. It is easy to see from Theorem 3 that the coordinates of λ^j for $1 \le j \le n$ expressed in $\mathbf{\Omega}$ are rational numbers whose numerators and denominators (in reduced terms) are bounded in absolute value by $\Delta^{O(1)}$. Using the multiplication table $\mathbf{M}$, we can compute all of these coordinates in time $(\log \Delta)^{O(1)}$. Next, we can use Gaussian elimination to compute the coefficients of g, also in time $(\log \Delta)^{O(1)}$. We can then obtain approximations to the roots of g in time $(\log \Delta + \log B)^{O(1)}$, using a polynomial-time rootfinding algorithm [21].

Thus, step 1 can be carried out in time $(\log \Delta + \log B)^{O(1)}$.

In what follows, all numerical computations involve "binary" rational numbers of the form $x \cdot 2^{-k}$, where x is an integer with $x = (nB)^{O(1)}$ and $k = O(\log(nB))$.

For step 2, we can enumerate the centers of the cubes $\mathcal{Z} \in \mathbf{U}$ that intersect $\mathcal{R}(B)$ in the following way. To list the cubes that lie in the n-dimensional quadrant $\mathbf{R}^n_{\ge 0}$, we consider the points $z = (z_0, 1/4, \ldots, 1/4)$ as z_0 takes on the values 1/4, 3/4, 5/4, etc., until $\mathcal{B}(z, 1/4)$ falls completely outside $\mathcal{R}(B)$. For each such z, we consider the points $z' = (z_0, z_1, 1/4, \ldots, 1/4)$ as z_1 takes on the values 1/4, 3/4, 5/4, etc., until $\mathcal{B}(z', 1/4)$ falls completely outside $\mathcal{R}(B)$. And so on. We use a similar enumeration scheme for the other n-dimensional quadrants.

With this enumeration scheme, the running time of step 2 is bounded by $N \cdot (\log B + n)^{O(1)}$, where N is the number of cubes $\mathcal{Z}$ in $\mathbf{U}$ such that either $\mathcal{Z}$ itself or a cube adjacent to $\mathcal{Z}$ intersects $\mathcal{R}(B)$. We then have

$$N \le 2^n \mathrm{vol}(\mathcal{R}(B) + \mathcal{B}(0, 1)),$$

where

$$\mathcal{R}(B) + \mathcal{B}(0, 1) = \{u + v : u \in \mathcal{R}(B),\ v \in \mathcal{B}(0, 1)\}.$$

It is easy to see that

$$\mathcal{R}(B) + \mathcal{B}(0, 1) \subset \mathcal{R}(2^n B).$$

In the proof of Theorem 6.5 in [16], it is shown that for any $\delta \ge 1$

$$\mathrm{vol}(\mathcal{R}(\delta)) \le \delta \cdot \frac{2^n (n - 1 + \log \delta)^{n-1}}{(n-1)!}.$$

It follows that

$$N \le B(n + \log B)^{O(n)},$$

and thus step 2 can be carried out in time

$$B(n + \log B)^{O(n)}.$$

□

Theorem 13 *Assume the ERH. There is a deterministic algorithm with the following properties. It takes as input an integer $n > 1$ and a prime $p \nmid n$, along with the data (4) describing $\mathbf{O}$. It produces as output a set of elements in $\mathbf{O}^{(p)}$ whose image in $(\overline{\mathbf{O}})^*$ is a generating set. The algorithm runs in time*

$$\Delta (\log \Delta)^{O(n)} + \Delta \kappa 2^{O(n)} (\log(\Delta \kappa))^{O(1)}.$$

Proof. We use Algorithm 11 to enumerate the set $A[\mathcal{R}(\Delta)] \cap \mathbf{Z}^n$. This gives the first term in the above running time. By Theorem 3, we have

$$A[\mathcal{B}(0, (\kappa\Delta)^{1/n}] \subset \mathcal{B}(0, (\kappa\Delta)^{1/n}).$$

We can list all points in $\mathbf{Z}^n$ that lie in this box in a straightforward fashion. This gives the second term in the above running time. $\square$

We can now prove the following theorem, which is the main result of this paper.

Theorem 14 *Assume the ERH. There is a deterministic algorithm with the following properties. It takes as input an integer $n > 1$ and a prime $p \not| n$. It produces as output a model for $\mathbf{F}_{p^n}$, together with a generating set for $\mathbf{F}_{p^n}^*$. The algorithm runs in time*

$$n^{O(n)}(\log p)^{2n}(\operatorname{loglog} p)^{O(1)}.$$

Proof. We can construct our model for $\mathbf{F}_{p^n}$, using the algorithm of Fact 2, in time $(n \operatorname{loglog} p)^{O(1)}(\log p)^2$. Then by Theorem 13, we can construct a generating set in time

$$(n \operatorname{loglog} p)^{O(n)}(\log p)^{2(n-1)} + n^{O(n)}(\log p)^{2n}(\operatorname{loglog} p)^{O(1)}.$$

In the first term, the factor $(\operatorname{loglog} p)^{O(n)}$ is less than $\log p$ unless $\operatorname{loglog} p = O(n \log n)$. Thus, the total running time is

$$n^{O(n)}(\log p)^{2n}(\operatorname{loglog} p)^{O(1)}.$$

$\square$

6 Another algorithm for constructing a generating set

Theorem 14 is not the best possible. In this section, we very briefly sketch a faster algorithm. This algorithm, however, lacks the elegance and simplicity of the one presented in the previous section.

The algorithm is based directly on Theorem 10, instead of Theorem 4. Consider the sets S_1 and S_2 in Theorem 10. By Fact 7, parts (3) and (4), we can bound the cardinality of these sets as follows:

$$\begin{aligned} \#S_1 &\leq hR(\log\Delta)^{O(n)} \leq \Delta^{1/2}(\log\Delta)^{O(n)}, \\ \#S_2 &\leq h\kappa \leq \kappa\Delta^{1/2}(\log\Delta)^{O(n)}. \end{aligned}$$

Thus,

$$\#(S_1 \cup S_2) \leq (n \operatorname{loglog} p)^{O(n)} \cdot (\log p)^{n+1}.$$

We might hope that we can enumerate the elements of $S_1 \cup S_2$ in time roughly proportional to the cardinality of this set. This is indeed the case.

Using the algorithmic techniques of [8] and [10], we can construct the set of all reduced ideals, along with the set S_1 of all neighbors of 1 in these ideals, in time

$$\Delta^{1/2}(\log\Delta)^{O(n)},$$

assuming the ERH.

From among these reduced ideals, we select a complete system $\mathbf{A}_1, \ldots, \mathbf{A}_h$ of representatives for the class group. Again, using the algorithmic techniques of [8] and [10], we can construct the minima required for the set S_2 in time

$$\Delta^{1/2}(\log \Delta)^{O(n)}\kappa(\log \kappa + \log \Delta)^{O(1)},$$

assuming the ERH.

In the above running time estimates, the ERH is used to allow fast deterministic factorization of polynomials modulo primes (see [19]). These factorizations are needed to factor rational primes in $\mathbf{O}$.

It then follows that, assuming the ERH, we can construct a model for $\mathbf{F}_{p^n}$, together with a generating set for $\mathbf{F}^*_{p^n}$, in time

$$(n \operatorname{loglog} p)^{O(n)} \cdot (\log p)^{n+1}.$$

We note that this faster algorithm requires space for $\Delta^{1/2}(\log \Delta)^{O(n)} + (\log p)^{O(1)}$ bits of working storage (not including the space for the output). The algorithm presented in the previous section requires space for only $(\log \Delta + \log p)^{O(1)}$ bits of working storage.

7 Conclusions and open problems

We have presented a deterministic algorithm that, under the assumption of the ERH, constructs a generating set for $\mathbf{F}^*_{p^n}$ in time $(n \log p)^{O(n)}$. One potential area for future work is improving the cost of finding a guaranteed generating set (assuming the ERH, possibly using a probabilistic algorithm). A second area is removing unproven assumptions (ERH), either in the correctness or complexity analysis. A third area is developing an efficient deterministic search procedure for finding primitive roots in $\mathbf{F}_{p^n}$, for large p and $n > 2$ (assuming the ERH).

Acknowledgement

The authors would like to thank Hendrik Lenstra for generously sharing with us many of his ideas on this topic.

References

[1] L. M. Adleman and H. W. Lenstra Jr. Finding irreducible polynomials over finite fields. In *18th Annual ACM Symposium on Theory of Computing*, pages 350–355, 1986.

[2] L. M. Adleman, K. Manders, and G. L. Miller. On taking roots in finite fields. In *18th Annual Symposium on Foundations of Computer Science*, pages 175–178, 1977.

[3] N. C. Ankeny. The least quadratic nonresidue. *Ann. of Math.*, 55:65–72, 1952.

[4] E. Bach. Explicit bounds for primality testing and related problems. *Math. Comp.*, 55:355–380, 1990.

[5] E. Bach and J. Shallit. Factoring with cyclotomic polynomials. *Math. Comp.*, 52(185):201–219, 1989.

[6] E. Bach and J. von zur Gathen. Deterministic factorization of polynomials over special finite fields. Technical Report 799, Computer Sciences Department, University of Wisconsin–Madison, 1988.

[7] E. Bach, J. von zur Gathen, and H. W. Lenstra. Preprint, 1989.

[8] J. Buchmann. On the computation of units and class numbers by a generalization of Lagrange's algorithm. *J. Number Theory*, 26:8–30, 1987.

[9] J. Buchmann. On the period length of the generalized Lagrange algorithm. *J. Number Theory*, 26:31–37, 1987.

[10] J. Buchmann and H. C. Williams. On principal ideal testing in algebraic number fields. *J. Symbolic Computation*, 4:11–19, 1987.

[11] S. A. Evdokimov. Factoring a solvable polynomial over a finite field and Generalized Riemann Hypothesis. *Zapiski Nauchn. Semin. Leningr. Otdel. Matem. Inst. Acad. Sci. USSR*, 176:104–117, 1989. In Russian.

[12] J. von zur Gathen. Factoring polynomials and primitive elements for special primes. *Theoret. Comput. Sci.*, 52:77–89, 1987.

[13] M. A. Huang. Riemann hypothesis and finding roots over finite fields. In *17th Annual ACM Symposium on Theory of Computing*, pages 121–130, 1985.

[14] G. J. Janusz. *Algebraic Number Fields.* Academic Press, 1973.

[15] H. W. Lenstra. Finding isomorphisms between finite fields. *Math. Comp.*, 56:329–347, 1991.

[16] H. W. Lenstra. Algorithms in algebraic number theory. *Bull. Amer. Math. Soc.*, 26:211–244, 1992.

[17] G. I. Perel'muter and I. E. Shparlinsky. On the distribution of primitive roots in finite fields. *Uspekhi Mat. Nauk*, 45:185–186, 1990. In Russian.

[18] S. Pohlig and M. Hellman. An improved algorithm for computing logarithms over GF(p) and its cryptographic significance. *IEEE Trans. Inf. Theory*, 24:106–110, 1978.

[19] L. Rónyai. Factoring polynomials over finite fields. *J. Algorithms*, 9:391–400, 1988.

[20] L. Rónyai. Galois groups and factoring polynomials over finite fields. In *30th Annual Symposium on Foundations of Computer Science*, pages 99–104, 1989.

[21] A. Schönhage. The fundamantal theorem of algebra in terms of computational complexity. Unpublished manuscript, 1982.

[22] V. Shoup. Removing randomness from computational number theory (Ph. D. thesis). Technical Report 865, Computer Sciences Department, University of Wisconsin–Madison, 1989.

[23] V. Shoup. New algorithms for finding irreducible polynomials over finite fields. *Math. Comp.*, 54(189):435–447, 1990.

[24] V. Shoup. Smoothness and factoring polynomials over finite fields. *Inform. Process. Lett.*, 38:39–42, 1991.

[25] V. Shoup. Searching for primitive roots in finite fields. *Math. Comp.*, 58:369–380, 1992.

[26] Y. Wang. On the least primitive root of a prime. *Scientia Sinica*, 10(1):1–14, 1961.