MATHEMATICS OF COMPUTATION
VOLUME 56, NUMBER 193
JANUARY 1991, PAGES 329–347

FINDING ISOMORPHISMS BETWEEN FINITE FIELDS

H. W. LENSTRA, JR.

ABSTRACT. We show that an isomorphism between two explicitly given finite fields of the same cardinality can be exhibited in deterministic polynomial time.

1. INTRODUCTION

Every finite field has cardinality p^n for some prime number p and some positive integer n. Conversely, if p is a prime number and n a positive integer, then there exists a field of cardinality p^n, and any two fields of cardinality p^n are isomorphic. These results are due to E. H. Moore (1893) [10]. In the present paper we are interested in an algorithmic version of his theorem, in particular of the uniqueness part.

We say that a finite field is *explicitly given* if, for some basis of the field over its prime field, we know the product of any two basis elements, expressed in the same basis. Let, more precisely, p be a prime number and n a positive integer. Then by *explicit data* for a finite field of cardinality p^n we mean a system of n^3 elements $(a_{ijk})_{i,j,k=1}^{n}$ of the prime field $\mathbf{F}_p = \mathbf{Z}/p\mathbf{Z}$ such that $\mathbf{F}_p^n$ becomes a field with the ordinary addition and multiplication by elements of $\mathbf{F}_p$, and the multiplication determined by

$$e_i e_j = \sum_{k=1}^{n} a_{ijk} e_k ,$$

where $e_1, e_2, \ldots, e_n$ denotes the standard basis of $\mathbf{F}_p^n$ over $\mathbf{F}_p$. For example, if we know an irreducible polynomial $f \in \mathbf{F}_p[X]$ of degree n, then such explicit data are readily calculated, since $\mathbf{F}_p[X]/f\mathbf{F}_p[X]$ is a field of cardinality p^n. Conversely, given explicit data for a field of cardinality p^n, one can find an irreducible polynomial $f \in \mathbf{F}_p[X]$ of degree n by means of a polynomial-time algorithm (see Theorem (1.1) below). By *polynomial-time* we mean that the time used by the algorithm—i.e., the number of bit operations that it performs—is bounded by a polynomial function of $\log p$ and n. It is supposed that

Received October 17, 1989; revised April 6, 1990.
1980 *Mathematics Subject Classification* (1985 *Revision*). Primary 11T30.
Key words and phrases. Finite field, algorithm.
Research supported by NSF contract DMS 87-06176.

0025-5718/91 $1.00 + $.25 per page

the elements of $\mathbf{F}_p$ are represented in the conventional way, so that the field operations in $\mathbf{F}_p$ can be performed in time $(\log p)^{O(1)}$.

It is not known whether there exists a polynomial-time algorithm that, given p and n, constructs explicit data for a finite field of cardinality p^n. If the generalized Riemann hypothesis is valid, then such an algorithm exists [1, 4]. Also, V. Shoup has shown [11] that the problem can be reduced to the problem of factoring polynomials in one variable over finite fields into irreducible factors. For the latter problem, no polynomial-time algorithm is known, even if the generalized Riemann hypothesis is assumed; there does exist an algorithm that runs in time $(pn)^{O(1)}$ (see [5, §4.6.2]), so for small p the problem is solved. If random algorithms are allowed, then both the problem of constructing finite fields and the problem of factoring one-variable polynomials over finite fields have perfectly satisfactory solutions, both from a practical and a theoretical point of view (see [7]).

Theorem (1.1). *There exists a polynomial-time algorithm that, given a prime number p, a positive integer n, and any of* (a), (b), (c), *constructs the two others*:

(a) *explicit data for a field of cardinality p^n;*
(b) *an irreducible polynomial in $\mathbf{F}_p[X]$ of degree n;*
(c) *for each prime number r dividing n, an irreducible polynomial in $\mathbf{F}_p[X]$ of degree r.*

The only nontrivial assertion of this theorem is that (c) suffices to construct (a) and (b). If for each prime number r that is *at most* n, an irreducible polynomial in $\mathbf{F}_p[X]$ of degree r were known, then (a) and (b) could be constructed using auxiliary cyclotomic extensions of $\mathbf{F}_p$. In our proof, which is given in §9, we work with auxiliary cyclotomic *ring* extensions of $\mathbf{F}_p$, which can be constructed without any hypothesis. The other assertions of the theorem are proved in §2.

We now come to the uniqueness part of Moore's theorem. Suppose that two finite fields of the same cardinality are explicitly given, can one find an isomorphism between them in polynomial time? The isomorphism is to be represented by means of its matrix on the given bases of the fields over the prime field.

For this second problem, the same results have been obtained as for the first problem. Thus, a polynomial-time algorithm exists if the generalized Riemann hypothesis is true, as was shown by S. A. Evdokimov [4]. Also, the problem can be reduced to factoring polynomials in one variable over finite fields. To see this, write the first field as $\mathbf{F}_p[X]/f\mathbf{F}_p[X]$; then finding an isomorphism is equivalent to finding a zero of f in the other field. This solves the problem if p is small, and also if random algorithms are allowed, as is the case in practice. In the present paper we prove the same result without any restriction.

Theorem (1.2). *There exists a polynomial-time algorithm that, given explicit data for two finite fields of the same cardinality, finds an isomorphism between them.*

The proof uses the same technique as the proof of Theorem (1.1). The result of Evdokimov that we just mentioned depends on auxiliary cyclotomic extensions of $\mathbf{F}_p$, and it is to construct these that the generalized Riemann hypothesis is needed. In our proof we use *ring* extensions, which can be obtained for free.

The contents of this paper are as follows. In §2 we discuss what can be done if explicit data for a finite field are available, and we define what is meant by explicit data for field extensions and field homomorphisms. In §3 we show how *normal bases* can be found in polynomial time. Normal bases are not absolutely vital for our purposes, but they provide an elegant solution to a technical problem that comes up later (see (5.6)), and the result is of interest in itself as well. In §§4, 5, and 6, we do not deal with algorithms at all. Section 4 is devoted to algebraic properties of certain cyclotomic ring extensions that need not be fields. A special role is played by the *Teichmüller subgroup* of the group of units of such a ring extension. In §5 we show that knowing an extension of given prime degree of a finite field is equivalent to knowing a generator of this Teichmüller subgroup. Conversely, such a generator can be used to make prime power degree extensions, as we show in §6. It is clear that such results can be used to make prime power degree extensions out of prime degree extensions and thus complete the proof of Theorem (1.1). Before we carry this through, we have to deal with certain exceptional cases. The case that the given prime equals the characteristic of the field is dealt with, by well-known techniques, in §7. A second exceptional case is considered in §8. In this section we show that techniques from linear algebra can in certain cases be used to solve problems of a multiplicative nature. As an application we solve, in a theoretical sense, a minor problem that comes up in primality testing. Finally, in §9 we formulate and prove theorems that are slightly more general than Theorems (1.1) and (1.2).

Although the algorithms presented in this paper are not necessarily inefficient, I do not expect that in practice they can compete with the probabilistic algorithms referred to above. Accordingly, I have refrained from estimating the running times of the various algorithms precisely, and from optimizing the algorithms from either a theoretical or a practical point of view.

2. Explicit data

Let p be a prime number, n a positive integer, and $(a_{ijk})_{i,j,k=1}^{n}$ explicit data for a field of cardinality p^n. Denote by E the field with underlying set $\mathbf{F}_p^n$ that is determined by these data, as described in the introduction. We say in this situation also that $(a_{ijk})_{i,j,k=1}^{n}$ are explicit data for the field E. By $e_1, \ldots, e_n$ we denote the standard basis of $\mathbf{F}_p^n$ over $\mathbf{F}_p$.

Given such explicit data, the unit element 1 of E is characterized by the property $1 \cdot e_1 = e_1$. If we write $1 = \sum_i z_i e_i$, with $z_i \in \mathbf{F}_p$, then it follows that

$(z_i)_{i=1}^n$ is the unique solution of the system of linear equations

$$\sum_{i=1}^{n} a_{i1k} z_i = \begin{cases} 1 & \text{if } k = 1, \\ 0 & \text{if } k \neq 1 \end{cases}$$

over $\mathbf{F}_p$. This system can be solved in polynomial time by the usual techniques from linear algebra. The divisions in the field $\mathbf{F}_p$ that are needed by these techniques can be performed by means of the extended Euclidean algorithm [5, §4.5.2]. It follows that the unit element of E can be determined in polynomial time.

Once the unit element is determined, we can in a similar way find the inverse of any given nonzero element $\alpha \in E$ as the solution of $x\alpha = 1$, which can again be viewed as a system of n linear equations over $\mathbf{F}_p$. We conclude that the field operations in E can all be performed in polynomial time.

By repeated squarings and multiplications, we can calculate α^k for any $\alpha \in E$ and any positive integer k in time $(n + \log p + \log k)^{O(1)}$. This leads to an alternative method to calculate 1 and α^{-1}, since $1 = e_1^{p^n-1}$ and $\alpha^{-1} = \alpha^{p^n-2}$ for $\alpha \neq 0$.

If m is a positive integer, and $(b_{ijk})_{i,j,k=1}^m$ are explicit data for a field F of cardinality p^m, then by *explicit data* for a field homomorphism from E to F we mean a matrix $(c_{ij})_{1\leq i\leq m,\, 1\leq j\leq n}$ with entries from $\mathbf{F}_p$ such that the map $\mathbf{F}_p^n \to \mathbf{F}_p^m$ sending $(x_j)_{j=1}^n$ to $(\sum_{j=1}^n c_{ij}x_j)_{i=1}^m$ is a field homomorphism $\phi\colon E \to F$. We say in this situation also that $(c_{ij})_{1\leq i\leq m,\, 1\leq j\leq n}$ are explicit data for the field homomorphism ϕ. For example, explicit data for the unique field homomorphism $\mathbf{F}_p \to E$ are readily derived from the coordinates z_i of the unit element of E.

Calculating $e_1^p, \ldots, e_n^p$, we can find in polynomial time explicit data for the Frobenius automorphism $\sigma\colon E \to E$ that sends each $\alpha \in E$ to α^p. Likewise, explicit data can be found for each power of σ.

We next determine the subfields of E. These are in one-to-one correspondence with the divisors d of n. Notice that these divisors can all easily be found in time $n^{O(1)}$. Let d be a divisor of n. Then we can calculate the matrix of the $\mathbf{F}_p$-linear map $E \to E$ that sends each $\alpha \in E$ to $\sigma^d\alpha - \alpha$, and using techniques from linear algebra, we can find a basis for the kernel of this map, which is precisely the unique subfield of E of cardinality p^d. Expressing the product of any two basis elements of this subfield as a linear combination of the same basis, we then obtain explicit data for a field of cardinality p^d, as well as for the inclusion map of this field to E. All this can be done in polynomial time.

Let r be a prime number and t a positive integer such that r^t divides n. Applying the above to the divisors r^t and r^{t-1} of n, we can find bases of the subfields of degree r^t and r^{t-1} over $\mathbf{F}_p$. Checking the basis elements of the former field one by one, we can find an element β of the field of degree r^t that is not in the field of degree r^{t-1}. Then β has degree r^t over $\mathbf{F}_p$, so

$\beta^{r^t} = \sum_{i=0}^{r^t-1} c_i \beta^i$ for certain uniquely determined $c_i \in \mathbf{F}_p$, which can be found by solving a system of linear equations. The polynomial $X^{r^t} - \sum_{i=0}^{r^t-1} c_i X^i$ is the irreducible polynomial of β over $\mathbf{F}_p$. It is irreducible in $\mathbf{F}_p[X]$ and of degree r^t. Taking $t = 1$, we see that, in Theorem (1.1), we can construct (c) from (a) in polynomial time.

Let d be any divisor of n, and write d as a product of prime powers r^t that are pairwise relatively prime. For each r, let $\beta = \beta_r$ be an element of degree r^t, as above. It is well known that the degree of $\gamma = \sum_r \beta_r$ over $\mathbf{F}_p$ is then equal to $\prod_r r^t = d$. (It clearly *divides* d; to show that it actually equals d, it suffices to remark that for each r the degree r^t of $\beta_r = \gamma - \sum_{r' \neq r} \beta_{r'}$ divides the lcm of the degrees of γ and the $\beta_{r'}$.) As above, we can use γ to determine an irreducible polynomial in $\mathbf{F}_p[X]$ of degree d. Applying this to $d = n$, we see that (a) in Theorem (1.1) can be used to construct (b).

We already saw in the introduction how (b) in Theorem (1.1) can be used to construct (a), and once one has (a) one can construct (c) as above. The remaining part of the proof of Theorem (1.1), namely how to construct (a) and hence (b) starting from (c), is given in §9.

In the following section we shall see that explicit data for a finite field can also be used to determine a normal basis for that field over a subfield in polynomial time. This is done by means of an algorithm that, as many algorithms in this paper, depends heavily on techniques from linear algebra. These techniques allow one to deal with problems of an additive nature. Multiplicative problems, such as recognizing or determining primitive roots, and computing discrete logarithms [8, §3], are much harder, and no good way is known to solve them, even if random algorithms are allowed.

There is another, even more fundamental, algorithmic problem concerning explicit data for finite fields for which currently no polynomial-time algorithm is known. This is the problem of deciding, given positive integers p and n with $p \geq 2$ and a system of n^3 elements $(a_{ijk})_{i,j,k=1}^{n}$ of $\mathbf{Z}/p\mathbf{Z}$, whether these form explicit data for a field of cardinality p^n. For $n = 1$ this problem is equivalent to *primality testing*: given an integer $p \geq 2$, decide whether p is prime. For this problem no polynomial-time algorithm is known. There is one if the generalized Riemann hypothesis is assumed, and also if random algorithms are allowed [8, §5]. Using the techniques of this section, one can show that primality testing is the *only* obstacle: there is a polynomial-time algorithm that, given p, n, (a_{ijk}) as above, either proves that they do not form explicit data for a field of cardinality p^n, or proves that if p is prime they do.

It is convenient to have *relative* versions of our definitions, in which the base field is an explicitly given finite field E as above, rather than $\mathbf{F}_p$. Let l be a positive integer. By *explicit data* for an lth degree field extension of E we mean a system of l^3 elements $(c_{ijk})_{i,j,k=1}^{l}$ of $E = \mathbf{F}_p^n$ such that E^l becomes a field with the ordinary E-vector space structure and the multiplication determined

by

$$e'_i e'_j = \sum_{k=1}^{l} c_{ijk} e'_k ,$$

where $e'_1, e'_2, \ldots, e'_l$ denotes the standard basis of E^l over E. Denote this field by F. As above, we can determine the unit element of F, and consequently view E as a subfield of F. We shall refer to the explicit data a_{ijk} for the field E together with the c_{ijk} as explicit data for the field extension $E \subset F$. The notion of explicit data for E-homomorphisms—i.e., field homomorphisms between extensions of E that are the identity on E—is defined in the obvious way.

In the above situation, one can identify F with $\mathbf{F}_p^{nl}$, using the basis $(e_i e'_j)_{1 \le i \le n,\, 1 \le j \le l}$ of F over $\mathbf{F}_p$, and one can readily calculate explicit data both for F as a field of cardinality p^{nl} and for the inclusion map $E \to F$. Conversely, if explicit data for a field F of cardinality p^m and for a field homomorphism $\phi\colon E \to F$ are given, then F can be viewed as a field extension of E via ϕ, and one can calculate explicit data for this field extension. The precise formulation and proof we leave to the reader.

In the remainder of this paper our language will be less formal, but not less precise. For example, when we speak of constructing a finite field, or an extension, or a homomorphism, then we mean constructing *explicit data* for a finite field, an extension, or a homomorphism. Likewise, if we say "given a finite field", when we speak about an algorithm, we mean that the algorithm is supplied with explicit data for that finite field. Computing an element of a given finite field means calculating the coordinates of that element on the given basis of the field over the prime field.

3. Finding a normal basis

If $E \subset F$ is a finite Galois extension of fields, with Galois group G, then a *normal basis* of F over E is a basis of F as a vector space over E of the form $(\sigma\alpha)_{\sigma \in G}$. A well-known theorem asserts that such a basis exists [12, §67].

Theorem (3.1). *There exists an algorithm that, given an extension $E \subset F$ of finite fields, finds a normal basis of F over E in time $(\log \#F)^{O(1)}$.*

Proof. Let $E \subset F$ be finite fields, and write $q = \#E$ and $l = [F : E]$. Denote by σ the automorphism of F that maps each $\alpha \in F$ to α^q. This is a generator of the Galois group of F over E.

It is convenient to use the following notation and terminology. It is taken from [9, §1], to which we refer for background information. For $f = \sum_i a_i X^i \in E[X]$ and $\alpha \in F$ we define

$$f \circ \alpha = \sum_i a_i \cdot \sigma^i \alpha .$$

This makes the additive group of F into a module over the polynomial ring

$E[X]$. Let $\alpha \in F$. Then the set $\{f \in E[X]: f \circ \alpha = 0\}$ is an ideal of $E[X]$ containing $X^l - 1$, so it is generated by a uniquely determined divisor of $X^l - 1$ with leading coefficient 1. Let this divisor be denoted by $\mathrm{Ord}(\alpha)$, the *Order* of α. From

$$f_1 \circ \alpha = f_2 \circ \alpha \Leftrightarrow f_1 \equiv f_2 \mod \mathrm{Ord}(\alpha)$$

it follows that the set $E[X] \circ \alpha = \{f \circ \alpha: f \in E[X]\}$ is a vector space over E of dimension $\deg \mathrm{Ord}(\alpha)$. Since it is the same as the E-linear span of $\{\sigma^i \alpha: 0 \leq i < l\}$, it follows that α gives rise to a normal basis of F over E if and only if its Order has degree l, which occurs if and only if $\mathrm{Ord}(\alpha) = X^l - 1$.

Suppose now that the extension $E \subset F$ is explicitly given. For any $\alpha \in F$ the degree of $\mathrm{Ord}(\alpha)$ is the least nonnegative integer k for which $\sigma^k \alpha$ belongs to the E-linear span of $\{\sigma^i \alpha: 0 \leq i < k\}$, and if $\sigma^k \alpha = \sum_{i=0}^{k-1} c_i \sigma^i \alpha$ for that value of k, then $\mathrm{Ord}(\alpha) = X^k - \sum_{i=0}^{k-1} c_i X^i$. This description of $\mathrm{Ord}(\alpha)$ makes it clear that there is a polynomial-time algorithm that determines $\mathrm{Ord}(\alpha)$ for any given $\alpha \in F$.

We now describe an algorithm to find a normal basis of F over E. Let α be any element of F (for example, $\alpha = 0$). Determine $\mathrm{Ord}(\alpha)$ by the method indicated above. $(*)$ If $\mathrm{Ord}(\alpha) = X^l - 1$, then α gives rise to a normal basis, and the algorithm stops. Suppose that $\mathrm{Ord}(\alpha) \neq X^l - 1$. Calculate the element $g = (X^l - 1)/\mathrm{Ord}(\alpha)$ of $E[X]$. As we shall prove below, there exists $\beta \in F$ with $g \circ \beta = \alpha$. Determine such an element β; this can be done by means of techniques from linear algebra, since the equation $g \circ \beta = \alpha$ can be formulated as a system of l linear equations over E. Determine $\mathrm{Ord}(\beta)$. If $\deg \mathrm{Ord}(\beta) > \deg \mathrm{Ord}(\alpha)$, then replace α by β and go to $(*)$. Suppose that $\deg \mathrm{Ord}(\beta) \leq \deg \mathrm{Ord}(\alpha)$. As we shall prove below, there exists a nonzero element $\gamma \in F$ with $g \circ \gamma = 0$, and any such γ has the property $\deg \mathrm{Ord}(\alpha + \gamma) > \deg \mathrm{Ord}(\alpha)$. Determine such an element γ by means of linear algebra, replace α by $\alpha + \gamma$, determine the Order of the new α, and go to $(*)$. This completes the description of the algorithm.

We next prove the assertions made in the description of the algorithm. Let α be any element of F, and let δ be an element that gives rise to a normal basis of F over E. Then there exists $f \in E[X]$ with $f \circ \delta = \alpha$. From $\mathrm{Ord}(\alpha) \circ \alpha = 0$ it follows that $(\mathrm{Ord}(\alpha) f) \circ \delta = 0$, so $\mathrm{Ord}(\alpha) f$ is divisible by $X^l - 1$. Therefore f is divisible by the polynomial $g = (X^l - 1)/\mathrm{Ord}(\alpha)$, and with $f = gh$ we now see that $g \circ (h \circ \delta) = \alpha$. This proves the assertion on the existence of β. Suppose now that $\mathrm{Ord}(\alpha) \neq X^l - 1$. Then $\mathrm{Ord}(\alpha) \circ \delta \neq 0$, and $g \circ (\mathrm{Ord}(\alpha) \circ \delta) = (X^l - 1) \circ \delta = 0$. This proves the assertion on the existence of γ. Let next β, γ be such that $g \circ \beta = \alpha$, $\deg \mathrm{Ord}(\beta) \leq \deg \mathrm{Ord}(\alpha)$, $\gamma \neq 0$, $g \circ \gamma = 0$. We prove that $\deg \mathrm{Ord}(\alpha + \gamma) > \deg \mathrm{Ord}(\alpha)$. From $g \circ \beta = \alpha$ it follows that $\mathrm{Ord}(\alpha)$ divides $\mathrm{Ord}(\beta)$, so the hypothesis $\deg \mathrm{Ord}(\beta) \leq \deg \mathrm{Ord}(\alpha)$ implies that $\mathrm{Ord}(\alpha) = \mathrm{Ord}(\beta)$. From $\mathrm{Ord}(g \circ \beta) = \mathrm{Ord}(\beta)$ it follows that g is relatively prime to $\mathrm{Ord}(\alpha)$, and the same is then true for the divisor $\mathrm{Ord}(\gamma)$

of g. This implies that $\mathrm{Ord}(\alpha+\gamma) = \mathrm{Ord}(\alpha)\mathrm{Ord}(\gamma)$, and from $\gamma \neq 0$ it now follows that $\deg \mathrm{Ord}(\alpha+\gamma) > \deg \mathrm{Ord}(\alpha)$. This proves the assertions made in the algorithm.

With every replacement of α, the degree of $\mathrm{Ord}(\alpha)$ increases by at least 1. It follows that the algorithm runs in polynomial time. The correctness of the algorithm is clear. This proves Theorem 3.1. □

If α gives rise to a normal basis of F over E, and σ is as above, then for each divisor d of l the element $\sum_{i=1}^{l/d} \sigma^{id}\alpha$ has degree d over E. This leads to an alternative proof of the part of Theorem (1.1) that was proved in §2.

4. Cyclotomic extensions

Let K denote a field and r a prime number that is different from the characteristic of K. In this section we study an rth cyclotomic ring extension of K. The group of units of a ring R with 1 will be denoted by R^*.

Denote by $K[\zeta]$ the ring

$$K[X]/\left(\sum_{i=0}^{r-1} X^i\right)K[X],$$

and let ζ denote the residue class of X. The dimension of $K[\zeta]$ over K equals $r-1$, a basis being given by $(\zeta^i)_{i=0}^{r-2}$, or, alternatively, by $(\zeta^i)_{i=1}^{r-1}$. Note that ζ has order r in the group $K[\zeta]^*$, and that for each integer a not divisible by r there is a unique ring automorphism ρ_a of $K[\zeta]$ that is the identity on K and for which $\rho_a\zeta = \zeta^a$. The set of all ρ_a's forms a group, which we denote by Δ. Clearly, there is a group isomorphism $\Delta \cong \mathbf{F}_r^*$ that maps ρ_a to $a \bmod r$; so Δ is cyclic of order $r-1$. The group Δ allows us to recover K from $K[\zeta]$, as follows. For a group G acting on a set S, we write $S^G = \{x \in S : \sigma x = x \text{ for all } \sigma \in G\}$.

Proposition (4.1). *We have* $K[\zeta]^\Delta = K$.

Proof. The basis $(\zeta^i)_{i=1}^{r-1}$ of $K[\zeta]$ over K is transitively permuted by Δ. Therefore, an element x of $K[\zeta]$ belongs to $K[\zeta]^\Delta$ if and only if all coefficients of x on that basis are equal. This is the case if and only if x is a K-linear multiple of the element $\sum_{i=1}^{r-1}\zeta^i$, which equals -1. This proves (4.1). □

Let k be a positive integer, and ε an element of a multiplicative group for which $\varepsilon^{r^k} = 1$. If a is an integer, then one easily checks that the element $\varepsilon^{a^{r^k-1}}$ only depends on ε and the residue class of $a \bmod r$; in particular, it does not depend on the choice of k. We write $\varepsilon^{\omega(a)}$ for this element. Note that $\varepsilon^{\omega(a)} = (\varepsilon^{\omega(b)})^{\omega(c)}$ if $a \equiv bc \bmod r$. We define the *Teichmüller subgroup* $T_K \subset K[\zeta]^*$ by

$$T_K = \{\varepsilon \in K[\zeta]^* : \varepsilon \text{ has } r\text{-power order, and } \rho_a\varepsilon = \varepsilon^{\omega(a)} \text{ for all } \rho_a \in \Delta\}.$$

To explain the terminology, we remark that ω is often called the *Teichmüller character*. Notice that $\zeta \in T_K$.

Proposition (4.2). *Every finite subgroup of T_K is cyclic. In particular, if K is finite then T_K is cyclic.*

Proof. Let $\mathfrak{m}$ be any maximal ideal of $K[\zeta]$, and let $L = K[\zeta]/\mathfrak{m}$. This is a field extension of K, so every finite subgroup of L^* is cyclic. Therefore, it suffices to show that the restriction of the natural map $\phi: K[\zeta] \to L$ to T_K is injective. Let $\varepsilon \in T_K$, $\phi(\varepsilon) = 1$. Write $\varepsilon = \sum_i c_i \zeta^i$, with $c_i \in K$, and let $\eta = \phi(\zeta)$; this is a primitive rth root of unity in L. For each $\rho_a \in \Delta$ we have $\sum_i c_i \eta^{ai} = \phi(\rho_a \varepsilon) = \phi(\varepsilon^{\omega(a)}) = \phi(\varepsilon)^{\omega(a)} = 1$. This shows that the polynomial $1 - \sum_i c_i X^i \in L[X]$ vanishes at all primitive rth roots of unity in L, so it is divisible by $\sum_{i=0}^{r-1} X^i$ (in $L[X]$, and hence in $K[X]$). Therefore, $1 - \varepsilon = 0$, so $\varepsilon = 1$, as required. This proves (4.2). □

Let $c \in K[\zeta]$, and let s be a positive integer that is a power of r. We denote by $K[\zeta][c^{1/s}]$ the ring

$$K[\zeta][Y]/(Y^s - c)K[\zeta][Y],$$

and by $c^{1/s}$ the residue class of Y in this ring. It contains $K[\zeta]$, and a basis of $K[\zeta][c^{1/s}]$ as a module over $K[\zeta]$ is given by $((c^{1/s})^i)_{i=0}^{s-1}$. The dimension of $K[\zeta][c^{1/s}]$ over K equals $s(r-1)$.

Assume, moreover, that $c \in T_K$. Then $c^{1/s}$ is an element of r-power order of $K[\zeta][c^{1/s}]^*$, so for each $a \in \mathbf{Z}$ there is a well-defined element $(c^{1/s})^{\omega(a)}$.

Proposition (4.3). *The action of Δ on $K[\zeta]$ can in a unique way be extended to an action of Δ as a group of ring automorphisms of $K[\zeta][c^{1/s}]$ such that each $\rho_a \in \Delta$ maps $c^{1/s}$ to $(c^{1/s})^{\omega(a)}$.*

Proof. Let $a \in \mathbf{Z} - p\mathbf{Z}$. The ring homomorphism $K[\zeta][Y] \to K[\zeta][c^{1/s}]$ that equals ρ_a on $K[\zeta]$ and maps Y to $(c^{1/s})^{\omega(a)}$ has $Y^s - c$ in its kernel, because $c \in T_K$. Therefore, it induces a ring homomorphism from $K[\zeta][c^{1/s}]$ to itself, which we again call ρ_a. This ring homomorphism is clearly uniquely determined by its effect on $K[\zeta]$ and $c^{1/s}$. It follows that ρ_1 is the identity and that $\rho_{a'}\rho_{a''} = \rho_a$ if $a'a'' \equiv a \bmod p$, so that each ρ_a is an automorphism. This proves (4.3). □

Proposition (4.4). *Suppose that c_1, c_2 are elements of T_K of the same order. Then there is a ring isomorphism $K[\zeta][c_1^{1/s}] \to K[\zeta][c_2^{1/s}]$ that is the identity on $K[\zeta]$ and respects the action of Δ.*

Proof. By (4.2), the elements c_1, c_2 generate the same subgroup of T_K. Let $c_1 = c_2^j$, with $\gcd(j, r) = 1$. As in the proof of (4.3), one constructs a ring homomorphism $\phi: K[\zeta][c_1^{1/s}] \to K[\zeta][c_2^{1/s}]$ that is the identity on $K[\zeta]$ and sends $c_1^{1/s}$ to $(c_2^{1/s})^j$. Checking the effect on the basis elements $(c_1^{1/s})^i$ of $K[\zeta][c_1^{1/s}]$ over $K[\zeta]$, one sees that this is an isomorphism. Let $\rho_a \in \Delta$. To prove that $\phi(\rho_a x) = \rho_a \phi(x)$ for all $x \in K[\zeta][c_1^{1/s}]$, one remarks that this is

obvious for $x \in K[\zeta]$ and for $x = c_1^{1/s}$, and that these generate $K[\zeta][c_1^{1/s}]$ as a ring. This proves (4.4). □

The ring $K[\zeta]$ studied in this section need not be a field. It is one if and only if $\sum_{i=0}^{r-1} X^i$ is irreducible in $K[X]$. If K is finite, this is the case if and only if $\#K$ is a primitive root modulo r.

5. Prime-degree extensions

In this section we let E be a finite field, q its cardinality, and r a prime number different from the characteristic of E. By m we denote the order of $(q \bmod r)$ in the group $\mathbf{F}_r^*$, and we let the positive integers t, u be such that $q^m - 1 = ur^t$ and $u \not\equiv 0 \bmod r$. The notation R^*, T_E, $E[\zeta][c^{1/r}]^\Delta$ is explained in the preceding section.

Theorem (5.1). *The group T_E is cyclic of order r^t, and if c generates T_E, then $E[\zeta][c^{1/r}]^\Delta$ is a field extension of E of degree r.*

This theorem is proved at the end of this section. It tells us how to obtain a field extension of degree r from a generator of the Teichmüller group T_E. Our next result tells us, conversely, how to obtain a generator of T_E from a field extension of degree r.

Let F be a field extension of E of degree r, and let α be an element of F that gives rise to a normal basis of F over E (see §3). We define β, $\gamma \in F[\zeta]$ by

$$\beta = \sum_{i=0}^{r-1} \zeta^i \alpha^{q^{im}}, \qquad \gamma = \prod_{a=1}^{r-1} \rho_a^{-1}(\beta^{ua^t}).$$

Below we shall see that $\beta^{ur^{t+1}} = 1$, so the expression a^{r^t} appearing in the definition of γ may be taken modulo r^{t+1}.

Notice that we can view $E[\zeta]$ as a subring of $F[\zeta]$.

Theorem (5.2). *The element $c = \gamma^r$ belongs to $E[\zeta]^*$, and it generates T_E. Moreover, there is a ring isomorphism $E[\zeta][c^{1/r}] \cong F[\zeta]$ that is the identity on $E[\zeta]$, maps $c^{1/r}$ to γ, and respects the action of Δ. It induces a field isomorphism $E[\zeta][c^{1/r}]^\Delta \cong F$.*

Proof. The field F is Galois over E, and its Galois group is generated by the automorphism of F that sends every $x \in F$ to x^q. Denote by τ the mth power of this automorphism. This is still a generator of the Galois group of F over E, because $\gcd(m, r) = 1$. We extend τ to a ring automorphism of $F[\zeta]$ by $\tau\zeta = \zeta$. For $x \in F[\zeta]$ we have

$$\tau x = x \Leftrightarrow x \in E[\zeta]. \tag{5.3}$$

To see this, write $x = \sum_{i=0}^{r-1} c_i \zeta^i$, with $c_i \in F$. Then $\tau x = x$ if and only if $\tau c_i = c_i$ for each i, if and only if $c_i \in E$ for each i, if and only if $x \in E[\zeta]$.

For every $x \in F[\zeta]$ we have

$$\tau x = x^{q^m}. \tag{5.4}$$

For $x \in F$ and for $x = \zeta$ this is clear, and these generate $F[\zeta]$ as a ring.

We can rewrite the definition of β as $\beta = \sum_{t=0}^{r-1} \zeta^t \tau^t \alpha$. From a straightforward computation we find that

$$\tau\beta = \zeta^{-1}\beta. \tag{5.5}$$

We show that

$$\beta \in F[\zeta]^*. \tag{5.6}$$

Since $F[\zeta]$ is finite, it suffices to prove that β is not a zero divisor. Because $(\tau^t\alpha)_{t=0}^{r-1}$ is a basis of F over E, it is also a basis of $F[\zeta]$ over $E[\zeta]$, and therefore $x\beta \neq 0$ for all $x \in E[\zeta]$, $x \neq 0$. To extend this to all $x \in F[\zeta]$, $x \neq 0$, it suffices to prove that every ideal of $F[\zeta]$, in particular the ideal $\{x \in F[\zeta]: x\beta = 0\}$, is generated by an element of $E[\zeta]$; or, equivalently, that every irreducible factor of $\sum_{t=0}^{r-1} X^t$ in $E[X]$ remains irreducible in $F[X]$. This is obvious, because the degree of any such irreducible factor is relatively prime to $[F : E]$. This proves (5.6).

From (5.5), (5.4), and (5.6) it follows that $\beta^{q^m-1} = \zeta^{-1}$, so the element $\delta = \beta^u$ satisfies $\delta^{r^l} = \zeta^{-1}$ and $\delta^{r^{l+1}} = 1$. Using the notation introduced in §4, we can therefore rewrite the definition of γ as

$$\gamma = \prod_{a=1}^{r-1} \rho_a^{-1}(\delta^{\omega(a)}).$$

Using that $\rho_a^{-1}(\zeta^{\omega(a)}) = \zeta$, one finds that

$$\gamma^{r^l} = \zeta. \tag{5.7}$$

From this one sees that γ has order r^{l+1}, and, using (5.4), that

$$\tau\gamma = \zeta^u\gamma. \tag{5.8}$$

An easy computation, which is the multiplicative analogue of the argument that proves (5.5), shows that

$$\rho_b\gamma = \gamma^{\omega(b)} \quad \text{for all } \rho_b \in \Delta,$$

so that $\gamma \in T_F$. Hence, $c = \gamma^r$ also belongs to T_F. It has order r^l. From $(\tau c)/c = c^{q^m-1} = c^{ur^l} = 1$ and (5.3) it follows that $c \in E[\zeta]$, and therefore $c \in T_E$. The order of any element of T_E divides $q^m - 1$, by (5.3), and since it is also a power of r, it actually divides r^l. With (4.2) it follows that c is a generator of T_E. This proves the first two assertions of (5.2).

To prove the remaining assertions, we consider the ring homomorphism $E[\zeta][Y] \to F[\zeta]$ that is the identity on $E[\zeta]$ and sends Y to γ. Clearly,

$Y^r - c$ is in the kernel of this map. We prove that it generates the kernel. For this it suffices to show that $\sum_{i=0}^{r-1} d_i \gamma^i$, with $d_i \in E[\zeta]$, vanishes only if all d_i are zero. Applying all powers of τ to the relation $\sum_{i=0}^{r-1} d_i \gamma^i = 0$, and using (5.8), we find that

$$\sum_{i=0}^{r-1} d_i \zeta^{ij} \gamma^i = 0$$

for all integers $j \pmod r$. Now let $k \in \{0, 1, \ldots, r-1\}$. Multiplying the jth relation by ζ^{-kj} and summing over j, we then see that $r d_k \gamma^k = 0$. Since $r\gamma^k$ is a unit, this implies that $d_k = 0$, as required.

It follows that an injective ring homomorphism $\psi\colon E[\zeta][c^{1/r}] \to F[\zeta]$ is induced. Since both rings are $r(r-1)$-dimensional over E, the map ψ is surjective. This proves the existence of the first ring isomorphism in (5.2).

Let $\rho_a \in \Delta$. For all $x \in E[\zeta]$ one trivially has $\psi(\rho_a x) = \rho_a \psi(x)$, and the same equality holds for $x = c^{1/r}$ because ρ_a raises both $c^{1/r}$ and γ to the power $\omega(a)$.

This proves that ψ respects the action of Δ. Passing to the Δ-invariants and applying (4.1), one concludes that an isomorphism $E[\zeta][c^{1/r}]^\Delta \cong F$ is induced. This proves (5.2). □

The following lemma will be needed in the next section.

Lemma (5.9). *Let F be a field extension of E of degree r, and let $\varepsilon \in T_F$ be any element satisfying $\varepsilon^{r^t} = \zeta$. Then all conclusions of* (5.2), *with γ replaced by ε, are valid.*

Indeed, all we used about γ was that $\gamma^{r^t} = \zeta$ and $\gamma \in T_F$.

Proof of (5.1). Since E is a finite field, we can choose a field extension F of E of degree r. Applying Theorem (5.2), we find a generator c for T_E, and in the proof we have seen that c has order r^t. Therefore, T_E is cyclic of order r^t. By (4.4), the ring $E[\zeta][c^{1/r}]^\Delta$ does not depend on the choice of the generator c of T_E, up to isomorphism, and by the last assertion of (5.2) it is a field. This proves (5.1). □

6. Prime-power-degree extensions

Let E, q, r, m, t be as in the previous section, let h be a positive integer, and let $s = r^h$. In this section we shall see that the results from the previous section carry over to extensions of degree s, provided that we make the assumption $s = 2$ or $r^t > 2$; thus only the case $r = 2, s \geq 4, q \equiv 3 \bmod 4$ is excluded.

Theorem (6.1). *Suppose that $s = 2$ or $r^t > 2$, and let c be a generator of T_E. Then $E[\zeta][c^{1/s}]^\Delta$ is a field extension of E of degree s.*

The proof is given at the end of this section.

Let F be a field extension of E of degree s, and denote by E' the unique subfield of F with $[F : E'] = r$. Let α be an element of F that gives rise to a normal basis of F over E' (see §3), and let $\beta, \gamma \in F[\zeta]$ be as in the previous section, but with E replaced by E'; so

$$\beta = \sum_{i=0}^{r-1} \zeta^i \alpha^{q^{imst}}, \qquad \gamma = \prod_{a=1}^{r-1} \rho_a^{-1}((\beta^{u'})^{\omega(a)}),$$

where u' is the largest divisor of $\#E'^*$ that is not divisible by r.

Theorem (6.2). *Suppose that $s = 2$ or $r^t > 2$. Then the element $c = \gamma^s$ belongs to $E[\zeta]^*$, and it generates T_E. Moreover, there is a ring isomorphism $E[\zeta][c^{1/s}] \cong F[\zeta]$ that is the identity on $E[\zeta]$, maps $c^{1/s}$ to γ, and respects the action of Δ. It induces a field isomorphism $E[\zeta][c^{1/s}]^\Delta \cong F$.*

Proof. By (5.2) we may assume that s is not prime. Then our hypothesis implies that $r^t > 2$. We consider the chain of fields

$$E = E_0 \subset E_1 \subset \cdots \subset E_{h-1} = E' \subset E_h = F,$$

in which each field has degree r over the preceding one. Let q_i denote the cardinality of E_i. From $q_{i+1} = q_i^r$ it follows that all q_i are congruent modulo r, so they all have the same multiplicative order m modulo r. Also, from $r^t \neq 2$ it follows that the number of factors r in $q_i^m - 1$ equals $t + i$, for $0 \leq i \leq h$. Applying (5.1) to each E_i, we see that the group T_{E_i} is cyclic of order r^{t+i}, so in the sequence of groups

$$T_E = T_{E_0} \subset T_{E_1} \subset \cdots \subset T_{E_{h-1}} = T_{E'} \subset T_{E_h} = T_F$$

each group is of index r in the next one. Applying (5.2) to the extension $E' \subset F$, we find that γ^r is a generator of $T_{E'}$, so for each i the element γ^{r^i} generates $T_{E_{h-i}}$. In particular, the element $c = \gamma^s$ generates T_E.

From (5.9), with $\varepsilon = \gamma^{r^i}$, it now follows that each $E_{h-i}[\zeta]$ is, as a ring, generated by $E_{h-i-1}[\zeta]$ and γ^{r^i}. Combining this for all i, one concludes that $F[\zeta]$ is, as a ring, generated by $E[\zeta]$ and γ. Therefore, the ring homomorphism $E[\zeta][Y] \to F[\zeta]$ that is the identity on $E[\zeta]$ and sends Y to γ is surjective. The element $Y^s - c$ is in the kernel, so a surjective ring homomorphism $E[\zeta][c^{1/s}] \to F[\zeta]$ is induced. Comparing dimensions over E, one concludes that it is an isomorphism. As in the proof of (5.2), one shows that it respects the Δ-action and induces an isomorphism $E[\zeta][c^{1/s}]^\Delta \cong F$. This proves (6.2). □

One derives (6.1) from (6.2) in exactly the same way as (5.1) was derived from (5.2).

7. Artin-Schreier extensions

In this section we deal with extensions of degree equal to the characteristic of the field, using Artin-Schreier theory [6, Chapter VIII, Theorem 6.4]. The following result already appears in [1].

Theorem (7.1). *There is an algorithm that, given a finite field E of characteristic p, constructs a pth-degree field extension F of E in time $(p \log \#E)^{O(1)}$.*

Proof. Let $\wp: E \to E$ be the $\mathbf{F}_p$-linear map sending each $x \in E$ to $x^p - x$. Since $\wp$ maps $\mathbf{F}_p$ to 0, it is not bijective, so there exists $a \in E$ that is not in the image of $\wp$. Also, such an a can be found by applying linear algebra over $\mathbf{F}_p$. Let $f \in E[X]$ be the polynomial $X^p - X - a$. We claim that f is irreducible, so that $F = E[X]/fE[X]$ is an explicitly given extension field of E of degree p.

To prove the claim, let α be a zero of f in an algebraic closure of E. Then all zeros of f are the elements $\alpha + i$, with $i \in \mathbf{F}_p$. Any two zeros of f generate the same field, so they have the same degree over E. Therefore, all irreducible factors of f in $E[X]$ have the same degree. Since f is of prime degree p, this implies that either f is irreducible or splits into p linear factors. The latter possibility is excluded because a was chosen such that f has no zero in E.

This proves Theorem (7.1). □

Theorem (7.2). *There is an algorithm that, given two field extensions F_1, F_2 of degree p of a finite field E of characteristic p, constructs an E-isomorphism $F_1 \to F_2$ in time $(\log \#F_1)^{O(1)}$.*

One way to prove the theorem is to use the reduction to the problem of factoring polynomials in one variable that was mentioned in the introduction. This gives rise to a polynomial-time algorithm because the characteristic is bounded by the degree. I present an alternative solution, which is more in the spirit of the other arguments in this paper.

Proof. Let F_1, F_2 be two explicitly given extensions of E of degree p, and let a, F be as in the proof of (7.1). Since we know that the fields F and F_1 are E-isomorphic, the element a must be in the image of the map $\wp_1: F_1 \to F_1$ sending each x to $x^p - x$. By means of linear algebra over $\mathbf{F}_p$ one can find, in polynomial time, an element $\alpha_1 \in F_1$ with $\alpha_1^p - \alpha_1 = a$. An explicit E-isomorphism $F \to F_1$ is now obtained by sending $X^i \bmod f$ to α_1^i, for $0 \le i < p$. Likewise, one constructs an E-isomorphism $F \to F_2$. Combining these isomorphisms, one obtains the desired E-isomorphism $F_1 \to F_2$. This proves (7.2). □

8. Taking roots

This section is devoted to the case that was excluded in Theorems (6.1) and (6.2). Shoup [11] has a very elegant way to deal with this case. Our approach is less efficient, but it is of interest in itself because it shows that linear algebra

can, in certain situations, be used to take roots in finite fields in polynomial time.

If E is a finite field of odd cardinality q, then an element $a \in E$ has a square root in E if and only if $a^{(q+1)/2} = a$. It follows that in the case $q \equiv 3 \mod 4$ every square $a \in E$ has $a^{(q+1)/4}$ as one of its square roots. Hence there is a polynomial-time algorithm to take square roots in finite fields of which the cardinality is $3 \mod 4$. The following theorem implies, more generally, that there is a polynomial-time algorithm to take square roots in finite fields whose *characteristic* is $3 \mod 4$.

Theorem (8.1). *There is an algorithm that, given a finite field E of characteristic p, an element $a \in E$ and a positive integer e satisfying*

$$p^h \equiv 1 \mod e, \quad \gcd(e, (p^h - 1)/e) = 1 \quad \textit{for some positive integer } h,$$

decides whether there exists $b \in E$ with $b^e = a$, and constructs such an element b if it exists, in time $(\log(e\#E))^{O(1)}$.

Proof. Let $q = \#E$. We may clearly restrict ourselves to the case that $a \neq 0$. Let it first be assumed that an integer h as in the statement of the theorem is known, with $p^h \leq q$. Let $c = a^{(p^h-1)/e}$. If a is an eth power, then c is a $(p^h - 1)$th power, so there exists a nonzero element x such that $x^{p^h} = cx$. This equation is $\mathbf{F}_p$-linear in x, so by means of linear algebra we can decide whether it has a nonzero solution, and find one if it exists.

If there is no such x, then a is not an eth power. Next suppose that x is nonzero and satisfies the equation. Then

$$x^{p^h-1} = a^{(p^h-1)/e}.$$

Using the extended Euclidean algorithm, one can find integers u, v with $ue + v(p^h - 1)/e = 1$. The element $b = a^u x^{v(p^h-1)/e}$ then satisfies

$$b^e = a^{ue} x^{v(p^h-1)} = a^{ue} a^{v(p^h-1)/e} = a,$$

as required.

To remove the assumption about h, one replaces e by $e' = \gcd(e, q - 1)$ and h by the multiplicative order h' of p modulo e'. From $q = p^n \equiv 1 \mod e'$ it follows that h' divides n, so indeed $p^{h'} \leq q$. We claim that $\gcd(e', (p^{h'} - 1)/e') = 1$. To prove this, note that h' divides h, so $(p^{h'} - 1)/e'$ divides both $(e/e') \cdot (p^h - 1)/e$ and $(q - 1)/e'$. From $\gcd(e/e', (q - 1)/e') = 1$ it follows that $(p^{h'} - 1)/e'$ divides $(p^h - 1)/e$, which is coprime to e and hence to e'. This establishes the claim. If a is an eth power, then it is clearly an e'th power. Conversely, if $a = b^{e'}$, then with $e' = u'e + v'(q - 1)$ we obtain $a = (b^{u'})^e$.

This proves (8.1). □

Corollary (8.2). *There is an algorithm that, given a finite field E of characteristic $p \equiv 3 \bmod 4$ and an element $a \in E$, decides whether there exists $b \in E$ with $b^2 = a$, and constructs such an element b if it exists, in time $(\log \#E)^{O(1)}$.*

Proof. Take $e = 2$, $h = 1$ in (8.1). This proves (8.2). □

Corollary (8.3). *There is an algorithm that, given a finite field E of characteristic $p \equiv 3 \bmod 4$, finds an element of the multiplicative group E^* of E of which the order is the largest power of 2 that divides $\#E^*$, in time $(\log \#E)^{O(1)}$.*

Proof. Starting from $a = -1$, repeat taking square roots until this is no longer possible. This clearly yields an element as desired. The number of iterations equals the number of factors 2 in $\#E^*$, which is less than $(\log \#E)/\log 2$. This proves (8.3). □

Corollary (8.4). *There is an algorithm that, given a finite field E of characteristic $p \equiv 3 \bmod 4$, constructs an extension field of E of degree 2 in time $(\log \#E)^{O(1)}$.*

Proof. If z is the element constructed by the algorithm of Corollary (8.3), then $E[X]/(X^2 - z)E[X]$ is a field extension of E of degree 2. This proves (8.4). □

The following explicit formula is of interest. Let E be a finite field of cardinality q, where $q \equiv 3 \bmod 4$. Then $E(\mathrm{i})$, with $\mathrm{i}^2 = -1$, is a quadratic extension of E. Let the map $f: E(\mathrm{i}) \to E(\mathrm{i})$ be defined by $f(x) = (1 + x)^{(q-1)/2}$. Then for every integer $m \geq 2$ for which 2^m divides $\#E(\mathrm{i})^*$, the element $f^{m-2}(\mathrm{i})$ has multiplicative order 2^m. This follows by induction on m from the fact that $f(x)^2 = x^{-1}$ for all x with $x^{q+1} = 1$.

The final result of this section solves, in a theoretical sense, a problem that comes up in primality testing [3, (11.6)(a); 2, §5].

Corollary (8.5). *There is an algorithm that, given a positive integer p that is $3 \bmod 4$, finds an element $u \in \mathbf{Z}/p\mathbf{Z}$ with the property that, if p is prime, the Legendre symbol $((u^2 + 4)/p)$ equals -1, in time $(\log p)^{O(1)}$.*

Proof. Assume first that p is prime. Using the above formula, one can find an element z of $\mathbf{F}_p(\mathrm{i})^*$ of order equal to the largest power of two dividing $p^2 - 1$. We claim that $u = z - z^{-1}$ has the required property. To see this, notice that z^{p+1} has order 2, so is equal to -1. Hence the irreducible polynomial $(X - z)(X - z^p)$ of z over $\mathbf{F}_p$ equals $X^2 - uX - 1$. Since the polynomial is irreducible, its discriminant $u^2 + 4$ is not a square in $\mathbf{F}_p$.

For general p, the computations leading to the element u can be carried out in $(\mathbf{Z}/p\mathbf{Z})[Y]/(Y^2 + 1)$ instead of $\mathbf{F}_p(\mathrm{i})$. This proves (8.5). □

9. Proofs of the theorems

The following theorem clearly implies Theorem (1.1).

Theorem (9.1). *There exists an algorithm that, given a finite field E of characteristic p, a positive integer n, and any of* (a), (b), (c), *constructs the two others*

in time $(n \log \#E)^{O(1)}$:

(a) *explicit data for a field extension of* E *of degree* n;
(b) *an irreducible polynomial in* $E[X]$ *of degree* n;
(c) *for each prime number* r *that divides* n *but that does not divide the degree* $[E : \mathbf{F}_p]$, *an irreducible polynomial in* $E[X]$ *of degree* r.

The proof that each of (a) and (b) suffices to construct the two others is the same as the proof for the case that the base field is $\mathbf{F}_p$ (see §§1 and 2). In this section we prove that (c) can be used to construct (a) and hence (b). We need the following lemma.

Lemma (9.2). *Given a finite field* E, *a prime number* r, *and a field extension* F *of* E *of degree* r, *one can construct a field extension of* F *of degree* r *in time* $(\log \#F)^{O(1)}$.

Proof. Let p, q denote the characteristic and the cardinality of E, respectively. First suppose that $r \neq p$, and let the case $r = 2$, $q \equiv 3 \bmod 4$ be excluded. Using (3.1), we can construct an element $\alpha \in F$ that gives rise to a normal basis of F over E. Given α, we can calculate the elements β, γ of $F[\zeta]$ that are defined in §5. By (5.2), the element $c = \gamma^r$ is a generator of T_E, and there is a ring isomorphism $E[\zeta][c^{1/r}] \cong F[\zeta]$ that induces an isomorphism $E[\zeta][c^{1/r}]^{\Delta} \cong F$. Also, the ring $F' = E[\zeta][c^{1/r^2}]^{\Delta}$ is a field extension of E of degree r^2, by (6.1). It is clear that explicit data for the field extension $E \subset F'$ are readily calculated from the definition of F'. Since we can view $E[\zeta][c^{1/r}]$ as a subring of $E[\zeta][c^{1/r^2}]$, by identifying $c^{1/r}$ with $(c^{1/r^2})^r$, we can identify F with a subfield of F'. The degree of F over F' equals r, as required.

In the cases that we excluded, the subfield E of F is not even needed. If $r = p$, then it suffices to apply (7.1) to F instead of E. If $r = 2$ and $q \equiv 3 \bmod 4$, then $p \equiv 3 \bmod 4$, so we may apply (8.4). This proves (9.2). □

Proof of (9.1). Let E and n be given, as well as an irreducible polynomial of degree r in $E[X]$, for every prime number r that divides n but that does not divide $[E : \mathbf{F}_p]$. We construct an nth degree extension of E by induction on the number of primes dividing n, counting multiplicities. We may clearly assume that $n > 1$. Let r be a prime number dividing n, and suppose that a field extension F' of E of degree n/r has been constructed. It will suffice to construct an rth-degree field extension of F'. We distinguish two cases.

In the first case, r divides the degree $[F' : \mathbf{F}_p]$. Then F' has a subfield E' with $[F' : E'] = r$, and E' can be determined by the methods of §2. Applying (9.2) to the extension $E' \subset F'$, we see that we can construct a field extension of F' of degree r, as required.

In the second case, r does not divide $[F' : \mathbf{F}_p]$. Then in particular, r does not divide $[E : \mathbf{F}_p]$, so by hypothesis an irreducible polynomial $f \in E[X]$ is given. Because $[F' : E]$ is not divisible by r either, f is still irreducible in

$F'[X]$. Therefore, $F = F'[X]/fF'[X]$ is the required field extension of F' of degree r.

This proves Theorem (9.1). □

The following theorem clearly implies (1.2).

Theorem (9.3). *There is an algorithm that, given a finite field E, a positive integer n, and two field extensions F_1, F_2 of E of degree n, constructs an E-isomorphism $F_1 \to F_2$ in time $(\log \# F_1)^{O(1)}$.*

We first deal with the case that n is a prime number.

Lemma (9.4). *Given a finite field E, a prime number r, and two field extensions F_1, F_2 of E of degree r, one can construct an E-isomorphism $F_1 \to F_2$ in time $(\log \# F_1)^{O(1)}$.*

Proof. By Theorem (7.2) we may assume that r is different from the characteristic of E. Applying Theorem (5.2), we can, as in the proof of (9.2), construct generators c_1, c_2 of T_E and E-isomorphisms $E[\zeta][c_i^{1/r}]^{\Delta} \cong F_i$, for $i = 1, 2$. Thus, it suffices to construct a ring isomorphism $E[\zeta][c_1^{1/r}] \cong E[\zeta][c_2^{1/r}]$ that is the identity on E and respects the action of Δ. Inspecting the proof of Proposition (4.4), one sees that this can be done if an integer j is known with $c_1 = c_2^j$.

Finding j is done by the following well-known iterative procedure. Let t be such that $\# T_E = r^t$. First put $j = 1$. $(*)$ Determine the smallest nonnegative integer k for which $(c_1/c_2^j)^{r^k} = 1$. If $k = 0$, then one has $c_1 = c_2^j$, and we are done. If $k > 0$, then $(c_1/c_2^j)^{r^{k-1}}$ is an element of order r of T_E, so there is a unique integer $l \in \{1, 2, \ldots, r-1\}$ such that

$$(c_1/c_2^j)^{r^{k-1}} = c_2^{lr^{t-1}}.$$

This integer l can be found by a direct search. Now replace j by $j + lr^{t-k}$, and start again at $(*)$. To justify this algorithm, one remarks that the value of k is initially at most t, and that it decreases by at least 1 in every iteration step. The search among the powers of $c_2^{r^{t-1}}$ is simplified by the fact that they coincide with the powers of ζ, because $c_2^{r^{t-1}} = \zeta$ (see (5.7)). Since also $c_1^{r^{t-1}} = \zeta$, the initial value of k is actually at most $t-1$.

This proves (9.4). □

Proof of (9.3). Let E be a finite field, n a positive integer, and F_1, F_2 two explicitly given field extensions of E of degree n. To find an E-isomorphism $F_1 \to F_2$, one first finds prime numbers r_i such that $n = r_1 r_2 \cdots r_m$, which can easily be done in time $n^{O(1)}$. Next, one determines, by the methods of §2, chains of fields

$$E = E_0 \subset E_1 \subset \cdots \subset E_{m-1} \subset E_m = F_1,$$
$$E = E_0' \subset E_1' \subset \cdots \subset E_{m-1}' \subset E_m' = F_2,$$

such that $[E_i : E_{i-1}] = [E'_i : E'_{i-1}] = r_i$ for $0 < i \leq m$. Using (9.4), one constructs successively E-isomorphisms $E_1 \to E'_1, E_2 \to E'_2, \ldots, E_m \to E'_m$. This proves Theorem (9.3). □

The algorithms given in the proofs of (9.1) and (9.3) can in many cases be made more efficient by working with field extensions of which the degree is a prime power rather than a prime number.

BIBLIOGRAPHY

1. L. M. Adleman and H. W. Lenstra, Jr., *Finding irreducible polynomials over finite fields*, Proc. 18th Annual ACM Sympos. on Theory of Computing (STOC), Berkeley, 1986, pp. 350–355.
2. W. Borho, *Große Primzahlen und befreundete Zahlen: Über den Lucas-Test und Thabit-Regeln*, Mitt. Math. Ges. Hamburg **11** (1983), 232–256.
3. H. Cohen and H. W. Lenstra, Jr., *Primality testing and Jacobi sums*, Math. Comp. **42** (1984), 297–330.
4. S. A. Evdokimov, *Efficient factorization of polynomials over finite fields and generalized Riemann hypothesis*, prepublication, 1986.
5. D. E. Knuth, *The art of computer programming*, vol. 2, second ed., Addison-Wesley, Reading, Mass., 1981.
6. S. Lang, *Algebra*, second ed., Addison-Wesley, Reading, Mass., 1984.
7. A. K. Lenstra, *Factorization of polynomials*, Computational Methods in Number Theory (H. W. Lenstra, Jr. and R. Tijdeman, eds.), Mathematical Centre Tracts 154/155, Mathematisch Centrum, Amsterdam, 1982.
8. A. K. Lenstra and H. W. Lenstra, Jr., *Algorithms in number theory*, Handbook of Theoretical Computer Science (J. van Leeuwen, ed.), North-Holland (to appear).
9. H. W. Lenstra, Jr. and R. J. Schoof, *Primitive normal bases for finite fields*, Math. Comp. **48** (1987), 217–231.
10. E. H. Moore, *A doubly-infinite system of simple groups*, Bull. New York Math. Soc. **3** (1893), 73–78; Math. Papers read at the Congress of Mathematics (Chicago, 1893), Chicago, 1896, pp. 208–242.
11. V. Shoup, *New algorithms for finding irreducible polynomials over finite fields*, Math. Comp. **54** (1990), 435–447.
12. B. L. van der Waerden, *Algebra*, vol. I, seventh ed., Springer-Verlag, Berlin, 1966.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, BERKELEY, CALIFORNIA 94720
E-mail address: hwl@cartan.berkeley.edu