Discrete Logarithms and Local Units

Oliver Schirokauer

Philosophical Transactions: Physical Sciences and Engineering, Volume 345, Issue 1676, Theory and Applications of Numbers without Large Prime Factors (Nov. 15, 1993), 409-423.

Stable URL:
http://links.jstor.org/sici?sici=0962-8428%2819931115%29345%3A1676%3C409%3ADLALU%3E2.0.CO%3B2-1

Discrete logarithms and local units

By Oliver Schirokauer

Department of Mathematics, Oberlin College, Oberlin, Ohio 44074, U.S.A.

Let K be a number field and $\mathcal{O}_K$ its ring of integers. Let l be a prime number and e a positive integer. We give a method to construct l^eth powers in $\mathcal{O}_K$ using smooth algebraic integers. This method makes use of approximations of the l-adic logarithm to identify l^eth powers. One version we give is successful if the class number of K is not divisible by l and if the units in $\mathcal{O}_K$ which are congruent to 1 modulo l^{e+1} are l^eth powers. A second version only depends on Leopoldt's conjecture.

We use the technique of constructing l^eth powers to find discrete logarithms in a finite field of prime order. Our method for computing discrete logarithms is closely modelled after Gordon's adaptation of the number field sieve to this problem. We conjecture that the expected running time of our algorithm is

$$L_p[1/3; (64/9)^{1/3} + o(1)] \quad \text{for} \quad p \to \infty,$$

where

$$L_p[s; c] = \exp(c(\log q)^s (\log\log q)^{1-s}).$$

This is the same running time as is conjectured for the number field sieve factoring algorithm.

1. Introduction

Smooth numbers appear in a vast assortment of number theoretic algorithms, and generally in one of two roles. Either the algorithm depends on the smoothness of a particular number for its success or it makes use of many smooth elements to construct a useful multiplicative relation. The elliptic curve factoring method is an example of the former. In this case, a factor p of an integer n is found as soon as an elliptic curve with smooth order over $\mathbb{F}_p$ is found. An example of the latter type is the quadratic sieve. In this method, n is factored by finding a and b in $\mathbb{Z}/n\mathbb{Z}$ such that $a^2 = b^2$ but $a \neq \pm b$. Constructing the relation $a^2 = b^2$ requires finding many smooth elements and using the easily proved fact that any set of B-smooth integers whose cardinality is greater than the number of primes less than or equal to B contains elements whose product is a square.

In this paper, we contribute to the collection of algorithms of the second sort. The problem we address is that of computing logarithms in a finite field of prime order p. Let t and v be two elements in $\mathbb{F}_p^*$, and assume that v is in the subgroup of $\mathbb{F}_p^*$ generated by t. Our goal is to find the smallest integer $x \in \{0, \ldots, p-2\}$ such that $t^x = v$. We write $x = \log_t v$ and call x the discrete logarithm of v.

Let $\prod l^{e_l}$ be the prime factorization of $p-1$. To compute $\log_t v$ it is sufficient to find, for each prime l, a relation in $\mathbb{F}_p$ of the form $a^{t^{e_l}} = t^{x_l} v$, for in this case $\log_t v \equiv -x_l \bmod l^{e_l}$. The technique we give to find such a relation requires solving the following computational problem in the ring of integers $\mathcal{O}_k$ of a number field K.

Phil. Trans. R. Soc. Lond. A (1993) **345**, 409–423
Printed in Great Britain

Given a prime number l, a positive integer e, and a set $M \subset \mathcal{O}_K$ whose elements have smooth norm, find integers $y(m) \in \{0, \ldots, l^e - 1\}$ such that the product

$$\prod_{m \in M} m^{y(m)}$$

is an l^eth power in $\mathcal{O}_K$. We give a solution to this problem which works if $|M|$ is sufficiently large and which relies on l-adic logarithms to detect l^eth powers in $\mathcal{O}_K$. One version produces the desired exponents if the class number of K is prime to l and the set of l^eth powers in $\mathcal{O}_K$ contains all the units in $\mathcal{O}_K$ which are 1 modulo l^{e+1}. A second version succeeds if Leopoldt's conjecture is true.

The method used to find the set M mentioned above and the use of l^eth powers in $\mathcal{O}_K$ to find a relation in $\mathbb{F}_p$ are based on ideas found in the number field sieve factoring algorithm (Buhler *et al.* 1993; Lenstra *et al.* 1993) and in Gordon's adaptation of the number field sieve to the problem of discrete logarithms (1993). The close relationship between the number field sieve and our algorithm is reflected in the fact that they have the same conjectured expected running time of

$$L_p[1/3; (64/9)^{1/3} + o(1)] \quad \text{for} \quad p \to \infty,$$

where

$$L_p[s; c] = \exp(c (\log p)^s (\log \log p)^{1-s}).$$

Moreover, though the algorithm we give has yet to be implemented, it should be as practical as the general number field sieve factoring algorithm.

We begin in the next section with some background information on smooth numbers, sieving and linear algebra. We then describe in §3 the technique for constructing l^eth powers in $\mathcal{O}_K$ using smooth algebraic integers and in §4 present an algorithm for finding discrete logarithms. The running time analysis of the discrete logarithm algorithm follows in §5. The paper then concludes in §6 with a solution to a problem which arises in §3 and §4, namely that of constructing valuations in a number field.

2. Preliminaries

(*a*) *Smooth numbers*

Let B be a positive real number. We call an integer B-smooth if all its prime factors are at most B. We call an algebraic integer B-smooth if its norm to $\mathbb{Q}$ is B-smooth in $\mathbb{Z}$ or equivalently if the principal ideal it generates is divisible only by prime ideals of norm at most B. In the algorithms in this paper, we often search for smooth elements. The following result establishes the likelihood of finding such an element. We denote by $\psi(x, B)$, the number of integers less than x which are B-smooth.

Theorem 2.1. (Canfield *et al.* 1983.) *Let ϵ be a positive constant. Then*

$$x/\psi(x, B) = u^{u(1+o(1))} \tag{2.2}$$

uniformly in the region $x \geqslant 10$ and $B \geqslant (\log x)^{1+\epsilon}$, where $u = (\log x)/\log B$ and the limit implicit in the $o(1)$ is for $u \to \infty$.

The conclusion of Theorem 2.1 looks particularly simple if we express the quantities involved as numbers of the form $L_p[s; c]$. In fact, with $x = L_p[s; c]$ and with $B = L_p[s'; c']$, the right-hand side of (2.2) becomes

$$L_p[s - s'; (s - s') c/c' + o(1)],$$

where the limit is for $p \to \infty$.

(*b*) *Sieving*

A sieve can be used to detect smooth values of a polynomial. We describe briefly how this is done for a polynomial in one variable and for a homogeneous polynomial in two variables. Let $f \in \mathbb{Z}[X]$ be of degree n, and assume that we are interested in finding those integers a in the interval $[-\frac{1}{2}C, \frac{1}{2}C]$ for which $f(a)$ is B-smooth. Let N be the number of primes less than or equal to B and label these primes $q_1, \ldots, q_N$. We initialize the sieve be setting $d_0(a) = f(a)$. Next we compute $d_i(a)$ for $i > 0$ inductively. To do this we find all the values of a between $-\frac{1}{2}(q_i - 1)$ and $\frac{1}{2}(q_i - 1)$ for which $q_i | f(a)$ by solving f modulo q_i. We then note that q_i divides $f(a)$ if and only if q_i divides $f(a + q_i)$ and so find all a in the range $[-\frac{1}{2}C, \frac{1}{2}C]$ for which $f(a)$ is divisible by q_i by adding $\pm q_i$ to the roots of $f \bmod q_i$. We now divide $d_{i-1}(a)$ by the highest power of q_i dividing it and call the quotient $d_i(a)$. Clearly those a for which $f(a)$ is B-smooth are those for which $d_N(a) = \pm 1$.

If we allow probabilistic algorithms, we can find the roots of f over $\mathbb{Z}/q_i\mathbb{Z}$ in time bounded by $(n + \log q_i)^{O(1)}$ (see Lenstra Jr 1990). Once f is solved modulo q_i, the time required to find the a for which $f(a)$ is divisible by q_i and to compute $d_i(a)$ for these a is $O(C/q_i)$. Summing over the q_i, we get a bound of

$$\pi(B)\,(n + \log B)^{O(1)} + O(C \log\log B)$$

on the time required for the sieving process, where $\pi(B)$ is the number of primes less than or equal to B.

Assume now that $f \in \mathbb{Z}[X_1, X_2]$ is homogeneous of degree n and that we are interested in smooth values of $f(a_1, a_2)$ for a_1 and a_2 both ranging from $-\frac{1}{2}C$ to $\frac{1}{2}C$. We follow the same procedure used in the one variable case. In particular, for each q_i, we find those pairs (a_1, a_2) in the given range for which $q_i | f(a_1, a_2)$. To do this we let $\tilde{f} \in \mathbb{Z}[Y]$ be the polynomial obtained by dividing f by ${X_2}^n$ and letting $Y = X_1/X_2$, and then we find the roots of $\tilde{f}$ modulo q_i. For each b in $[-\frac{1}{2}(q_i - 1), \frac{1}{2}(q_i - 1)]$ such that $q_i | \tilde{f}(b)$ and for each value of $a_2 \in [-\frac{1}{2}C, \frac{1}{2}C]$, we obtain values $f(a_1, a_2)$ divisible by q_i by finding $a_1 \in [-\frac{1}{2}C, \frac{1}{2}C]$ such that $a_1 \equiv b a_2 \bmod q_i$. We see then that for each q_i, the time required to detect all the pairs (a_1, a_2) such that $q_i | f(a_1, a_2)$ is bounded by

$$(n + \log q_i)^{O(1)} + O(C^2/q_i),$$

the first term being the time needed to solve $\tilde{f}$ modulo q_i and the second to mark off all appropriate values of a_1 and a_2. We conclude that the total running time for the sieve is

$$\pi(B)\,(n + \log B)^{O(1)} + C^2 \log\log B.$$

(*c*) *Linear algebra*

Let A be an $n \times n$ matrix and v an n-dimensional vector, both with entries in $\mathbb{Z}$. We consider the problem of solving the equation $Ax = v$ over $\mathbb{Z}/l^c\mathbb{Z}$, where l is prime. For $c = 1$, we can use gaussian elimination, or in the case that A is sparse, Wiedemann's coordinate recurrence method (Wiedemann 1986). The first method runs in time $O(n^3)$, the latter in time $O(n^2)$. For $c > 1$, we let x_1 be a vector such that $Ax_1 \equiv v \bmod l$ and define x_i, for $1 < i \leqslant c$, inductively as follows. Assume $A(x_1 - lx_2 - \ldots - l^{i-2}x_{i-1}) \equiv v \bmod l^{i-1}$ so that $A(x_1 - lx_2 - \ldots - l^{i-2}x_{i-1}) - v = l^{i-1}w_{i-1}$ for some vector w_{i-1}. Now let x_i be a vector such that $Ax_i \equiv w_{i-1} \bmod l$. Then $A(x_1 - lx_2 - \ldots - l^{i-1}x_i) \equiv v \bmod l^i$. When $i = c$ we have constructed a solution to $Ax = v$ over $\mathbb{Z}/l^c\mathbb{Z}$. As is evident, finding this solution requires solving c equations

of the form $Ax = v$ over $\mathbb{Z}/l\mathbb{Z}$. We conclude that the original problem can be solved in time $O(cn^3)$ in general and in time $O(cn^2)$ if A is sparse.

3. Constructing powers in a number ring

Let $K = \mathbb{Q}(\alpha)$ be a number field of degree n over $\mathbb{Q}$ and denote by $\mathcal{O}_K$ the ring of integers of K. Let S be a finite set of prime ideals of $\mathcal{O}_K$. Call an element $\delta \in K$ an S-unit if $\mathrm{ord}_{\mathfrak{p}}(\delta) = 0$ for all prime ideals $\mathfrak{p} \notin S$, and let K_S be the set of all S-units in K. In this section we address the following problem.

Problem 3.1. *Given a positive integer e, a prime number l, a finite set $M \subset (\mathcal{O}_K \cap K_S)$ such that $|M| > |S|$, and an S-unit $m_0 \in \mathcal{O}_K$, determine a map $y: M \to \{0, \ldots, l^e - 1\}$ such that*

$$m_0 \prod_{m \in M} m^{y(m)}$$

is an l^eth power in $\mathcal{O}_K$.

Though neither this problem nor the algorithm we present to solve it refer to smooth numbers, usually S is taken to be the set of prime ideals of norm less than some bound, for in this case a search through the elements of $\mathcal{O}_K$ of norm less than a second bound is most likely to yield S-units. In other words, in our applications of the results given here, M will contain smooth algebraic integers.

Before addressing Problem 3.1, we give a solution to a second problem. We will see later in the section how these two problems are related.

Problem 3.2. *Given positive integers c and σ, a prime number l which does not ramify in K, a set $M \subset (\mathcal{O}_K \cap K_S)$ such that $|M| > |S| + \sigma n$, and an S-unit $m_0 \in \mathcal{O}_K$, determine a map $y: M \to \{0, \ldots, l^e - 1\}$ such that the element*

$$\gamma = m_0 \prod_{m \in M} m^{y(m)}$$

satisfies:

(i) $\mathrm{ord}_{\mathfrak{p}} \gamma \equiv 0 \bmod l^e$ *for all* $\mathfrak{p} \in S$ *and*

(ii) $\lambda_\sigma(\gamma) = 0$.

Our solution to Problem 3.2 depends on a sequence of logarithmic maps λ_i which we define as follows. Assume l does not ramify in K. Let Γ_1 be the multiplicative subset of $\mathcal{O}_K$ consisting of those elements with norm not divisible by l. For each prime ideal ℓ lying above l in $\mathcal{O}_K$, let $\epsilon_\ell = |(\mathcal{O}_K/\ell)^*|$, and denote by ϵ the least common multiple of the ϵ_ℓ. Then for all $\gamma \in \Gamma_1$,

$$\gamma^\epsilon \equiv 1 \bmod l. \tag{3.3}$$

We now define λ_1 to be the map from Γ_1 to $l\mathcal{O}_K/l^2\mathcal{O}_K$ given by the equation

$$\lambda_1(\gamma) = (\gamma^\epsilon - 1) + l^2\mathcal{O}_K.$$

For $i > 1$, we let $\Gamma_i = \{\gamma \in \Gamma_{i-1} \mid \lambda_{i-1}(\gamma) = 0\}$ and define $\lambda_i: \Gamma_i \to l^{2^{i-1}}\mathcal{O}_K/l^{2^i}\mathcal{O}_K$ to be the map which sends $\gamma \in \Gamma_i$ to $(\gamma^\epsilon - 1) + l^{2^i}\mathcal{O}_K$. Notice that $l^{2^{i-1}}\mathcal{O}_K/l^{2^i}\mathcal{O}_K$ is a module of rank n over $\mathbb{Z}/l^{2^{i-1}}\mathbb{Z}$. Therefore, if we fix a basis $\{b_j l^{2^{i-1}} + l^{2^i}\mathcal{O}_K\}$, where $1 \leqslant j \leqslant n$, then λ_i is given by the maps $\lambda_{i,j}: \Gamma_i \to \mathbb{Z}/l^{2^{i-1}}\mathbb{Z}$ determined by the congruence

$$\gamma^\epsilon - 1 \equiv \sum_{j=1}^{n} \lambda_{i,j}(\gamma)\, b_j\, l^{2^{i-1}} \bmod l^{2^i}.$$

Notice also that $\lambda_i(\gamma\gamma') = \lambda_i(\gamma)+\lambda_i(\gamma')$, and that $\lambda_{i,j}(\gamma\gamma') = \lambda_{i,j}(\gamma)+\lambda_{i,j}(\gamma')$. We conclude that the maps λ_i and $\lambda_{i,j}$ are homomorphisms on the group of units of $\mathcal{O}_K$.

Though the formulation of λ_i given above is computationally advantageous, it is worth noting that λ_i can alternatively be thought of as an approximation of the l-adic logarithm. To see this, let ℓ again denote a prime ideal of $\mathcal{O}_K$ lying above l, let K_ℓ be the completion of K at the valuation induced by ℓ, let $\mathcal{O}_\ell$ be the ring of integers of K_ℓ, and let

$$\mathcal{O}_l = \prod_{\ell | l} \mathcal{O}_\ell = \mathcal{O}_K \otimes \mathbb{Z}_l.$$

Notice that Γ_1 contains precisely those elements mapped to $\mathcal{O}_l^*$ under the embedding of $\mathcal{O}_K$ into $\mathcal{O}_l$. In this way Γ_1 can be thought of as the subset of $\mathcal{O}_K$ of local units at l. Now define $\log_l$ to be the map from $\mathcal{O}_l^*$ to $l\mathcal{O}_l$ given by the usual ℓ-adic logarithms on each component. It is easily seen from the definitions that $\log_l$ maps Γ_i^ϵ to $l^{2^{i-1}}\mathcal{O}_l$. Since $\log_l(\gamma^\epsilon) = \epsilon \log_l(\gamma)$ and ϵ is not divisible by l, we conclude that $\log_l$ in fact maps all of Γ_i to $l^{2^{i-1}}\mathcal{O}_l$. Let θ_i be the map from Γ_i to $l\mathcal{O}_K/l^{2^i}\mathcal{O}_K$ given by the sequence

$$\Gamma_i \hookrightarrow \mathcal{O}_l^* \xrightarrow{\log_l} l^{2^{i-1}}\mathcal{O}_l \longrightarrow l^{2^{i-1}}\mathcal{O}_l/l^{2^i}\mathcal{O}_l \simeq l^{2^{i-1}}\mathcal{O}_K/l^{2^i}\mathcal{O}_K.$$

Then

$$\lambda_i(\gamma) = \theta_i(\gamma^\epsilon) = \epsilon\theta_i(\gamma).$$

If λ_i were defined with ϵ replaced by a multiple of ϵ which is congruent to 1 modulo $l^{2^{i-1}}$, then λ_i would actually equal θ_i.

If l ramifies, then (3.3) does not hold. To overcome this difficulty, we can make use of the fact that, whether l ramifies or not,

$$\gamma^\epsilon \equiv 1 \bmod \mathfrak{a},$$

where $\mathfrak{a}$ is the product of the prime ideals lying above l in $\mathcal{O}_K$. We can, therefore, replace λ_1 with the map $\mu_1 : \Gamma_1 \to \mathfrak{a}/l\mathfrak{a}$ which sends γ to $\gamma^\epsilon - 1 + l\mathfrak{a}$, and similarly, for $i > 1$, we can define μ_i just as we did λ_i except that $l^{2^{i-1}}\mathcal{O}_K/l^{2^i}\mathcal{O}_K$ is replaced by $l^{2^{i-2}}\mathfrak{a}/l^{2^{i-1}}\mathfrak{a}$. With μ_i playing the role of λ_i, all the results of this paper can be modified to hold in the case that l ramifies. We choose, however, to continue with the maps λ_i and to restrict to the case that l does not ramify for ease of exposition.

The following algorithm solves Problem 3.2. We continue with the notation introduced in that problem.

Algorithm 3.4. *Step 1.* To each element $m \in M$ associate a vector v_m whose entries are the exponents in the ideal factorization

$$(m) = \prod_{\mathfrak{p} \in S} \mathfrak{p}^{e_\mathfrak{p}}$$

together with the values $\lambda_{1,j}(m)$, for $1 \leqslant j \leqslant n$. Define v_{m_0} similarly. These vectors all have length $|S|+n$.

Step 2. Let M' be a subset of M containing $|S|+n$ many elements, and label those elements of M which remain $m_1, m_2, \ldots, m_{c_1}$. Let $c_0 = \sigma n$ and notice that $c_1 = c_0 - n$. Now let A be the square matrix whose columns are the vectors v_m for $m \in M'$, and for $k = 0, \ldots, c_1$, solve the matrix equation

$$Ax = -v_{m_k} \tag{3.5}$$

over the ring $\mathbb{Z}/l^c\mathbb{Z}$.

It may, of course, turn out that (3.5) cannot be solved, for v_{m_k} might not be in the column space of A. If this is the case, A should be replaced by another square matrix composed of column vectors v_m. We do not address here the question of the likelihood that A is suitable, but only point out that in the applications that follow l is large, making it unlikely that l divides $\det(A)$ and so unlikely that A is singular modulo l.

Step 3. For $k = 0, \ldots, c_1$, let w_k be the solution found to (3.5), and denote by $w_k(m)$ the entry in w_k corresponding to the column v_m in A. Let

$$b_k = m_k \prod_{m \in M'} m^{w_k(m)}.$$

Then

(i) $\operatorname{ord}_{\mathfrak{p}}(b_k) \equiv 0 \bmod l^e$ for all $\mathfrak{p} \in S$ and

(ii) $\lambda_1(b_k) = 0$.

If $c_1 = 0$ the algorithm now terminates. In this case we are in the situation that $\sigma = 1$ and $M = M'$. Letting $y(m) = w_0(m)$ accomplishes our goal.

If $c_1 > 0$ we proceed inductively. Let $w_{k,1} = w_k$ and $b_{k,1} = b_k$. For $r > 1$, let $c_r = c_{r-1} - n$. We compute $b_{k,r}$, where $k = 0, \ldots, c_r$, as follows. Let $v_{k,r}$ be the vector of length n consisting of the values $\lambda_{r,j}(b_{k,r-1})$ for $j = 1, \ldots, n$. Let A_r be the $n \times n$ matrix whose ith column is v_{c_r+i}. For $k = 0, \ldots, c_r$, solve

$$A_r x = -v_{k,r} \tag{3.6}$$

modulo $l^{2^{r-1}}$. Denote by $w_{k,r}$ the solution to (3.6) and by $w_{k,r}(i)$ the ith entry in $w_{k,r}$. Now let

$$b_{k,r} = b_k \prod_{i=1}^{n} b_{c_r+i}^{w_{k,r}(i)}.$$

Then for $k = 0, \ldots, c_r$,

(i) $\operatorname{ord}_{\mathfrak{p}}(b_{k,r}) \equiv 0 \bmod l^e$ for all $\mathfrak{p} \in S$ and

(ii) $\lambda_r(b_{k,r}) = 0$.

The algorithm terminates when we have computed $b_{0,\sigma}$, in which case we simply set $y(m)$ equal to the least positive residue modulo l^e of the exponent to which m occurs in the factorization of $b_{0,\sigma}$. This concludes our description of Algorithm 3.4.

If Problem 3.1 is modified so that M is required to have cardinality greater than $|S| + \sigma n$ and if l is assumed not to ramify in K, then Algorithm 3.4 provides a solution to Problem 3.1 if the element $m_0 \prod m^{y(m)}$ which is produced is an l^eth power. We consider two choices for the constants c and σ in Problem 3.2 and Algorithm 3.4, and in each case determine conditions under which $m_0 \prod m^{y(m)}$ is an l^eth power.

Version 3.7. *Let $c = e$ and let σ be the least integer such that $2^\sigma > e$.*

Denote by U the group of units of $\mathcal{O}_K$ and by U_i the subgroup of U containing those units which are congruent to 1 modulo l^i.

Proposition 3.8. *Let e be a positive integer and let σ be such that $2^\sigma > e$. Assume that the class number of K is not divisible by l and that*

$$U_{e+1} \subset U^{l^e}.$$

Let $\gamma \in \Gamma_\sigma$ be such that

(i) $\operatorname{ord}_{\mathfrak{p}}(\gamma) \equiv 0 \bmod l^e$ *for all prime ideals $\mathfrak{p}$ in $\mathcal{O}_K$ and*

(ii) $\lambda_\sigma(\gamma) = 0$.

Then γ is an l^eth power in $\mathcal{O}_K$.

Proof. Condition (i) implies that γ generates the l^eth power of an ideal in $\mathcal{O}_K$. Since l does not divide the class number of K, we conclude that γ generates the l^eth power of a principal ideal. Let δ be a generator of this ideal. Then

$$\gamma = \delta^{l^e} u$$

with $u \in U$. Condition (ii) implies that $\gamma^\epsilon \equiv 1 \bmod l^{2^\sigma}$. Since $2^\sigma \geqslant e+1$ and since $(\delta^{l^e})^\epsilon \equiv 1 \bmod l^{e+1}$, we conclude that $u^\epsilon \in U_{e+1}$. By our assumption then, u^ϵ is an l^eth power which in turn implies that u is, since ϵ is prime to l. We conclude that γ is an l^eth power and the proposition is proved.

We claim that the assumptions of Proposition 3.8 are likely to be met in a number field K. We are not interested in this paper in proving theorems about our algorithms and so do not make this assertion rigorous or prove it. Instead we offer heuristic evidence. With regard to the class number, we rely on the analysis of finite R-modules, where R is the ring of integers of a number field, given by Cohen & Lenstra Jr (1983). Their results provide an explanation for certain experimental data about class groups and suggest that for large l, the probability that l divides the class number of K is at most approximately $1/l$ and that in the case that K is not Galois, this probability is approximately $1/l^{r+1}$, where r is the unit rank. To show that it is likely that $U_{e+1} \subset U^{l^e}$, we consider the map $\mu : U_1/U_1^{l^e} \to l\mathcal{O}_K/l^{e+1}\mathcal{O}_K$ induced by sending u to $u-1$ and argue that μ is likely to be injective. Observe that $U_1/U_1^{l^e} \simeq (\mathbb{Z}/l^e\mathbb{Z})^r$, that $l\mathcal{O}_K/l^{e+1}\mathcal{O}_K \simeq (\mathbb{Z}/l^e\mathbb{Z})^n$, and that the image of μ is contained in $(\mathbb{Z}/l^e\mathbb{Z})^{n-1}$ since U lies in the kernel of the norm map. We now compute the probability P that μ is injective under the assumption that it is a random map from $(\mathbb{Z}/l^e\mathbb{Z})^r$ to $(\mathbb{Z}/l^e\mathbb{Z})^{n-1}$, where by random we mean that the image of a set of generators of $(\mathbb{Z}/l^e\mathbb{Z})^r$ is picked at random. It is easily seen that

$$P = \prod_{i=0}^{r-1} \left(1 - \frac{(l^e)^i (l^{e-1})^{n-1-i}}{(l^e)^{n-1}}\right),$$

from which we get the inequality

$$P > 1 - \sum_{i=0}^{r-1} \frac{(l^e)^i (l^{e-1})^{n-1-i}}{(l^e)^{n-1}} = 1 - \frac{l^{(e-1)(n-1)} \sum_{i=0}^{r-1} l^i}{(l^e)^{n-1}} = 1 - \frac{l^r - 1}{l^{n-1}(l-1)}.$$

Even in the worst case that $r = n-1$, we see that P is close to $1 - 1/l$ for large l. This concludes our description of Version 3.7.

We turn to our second choice for c and σ in Problem 3.2 and Algorithm 3.4.

Version 3.9. *Let $c = e + d + c_K$, where e is the integer given in Problem 3.1, d is such that l^d kills the torsion of $E_1/\overline{U_1}$, and c_K is a bound on the power of l dividing the class number of K. Let σ be the least integer such that $2^\sigma > e + d$.*

We show that with c and σ chosen as above, Algorithm 3.4 produces an l^eth power if Leopoldt's conjecture is true. We give one formulation of the conjecture and refer the reader to Lang (1990) and Washington (1982) for a more thorough discussion. Recall that $\mathcal{O}_l = \mathcal{O}_K \otimes \mathbb{Z}_l$ and let

$$E_1 = \{u \in \mathcal{O}_l^* \mid u \equiv 1 \bmod l\}.$$

E_1 is a topological group and a free $\mathbb{Z}_l$-module of rank n under the action of exponentiation. Furthermore U_1 sits inside of E_1. We let $\overline{U_1}$ denote the closure of U_1 in E_1.

Leopoldt's Conjecture. *The $\mathbb{Z}_l$ rank of $\overline{U_1}$ is equal to the $\mathbb{Z}$ rank of U.*

Proposition 3.10. *Let c and σ be given as in the description of Version* 3.9, *and let γ be an element in Γ_σ such that*

(i) $\operatorname{ord}_{\mathfrak{p}} \gamma \equiv 0 \bmod l^c$ *for all prime ideals $\mathfrak{p}$ in $\mathcal{O}_K$ and*

(ii) $\lambda_\sigma(\gamma) = 0$.

Then if Leopoldt's conjecture is true, γ is an l^eth power.

Proof. Condition (i) implies that γ generates the l^cth power of an ideal in $\mathcal{O}_K$. Since l^{c_K} kills the l part of the class group of $\mathcal{O}_K$ and $c = e+d+c_K$, we conclude that γ generates the l^{e+d}th power of some principal ideal. We can therefore write

$$\gamma = \delta^{l^{e+d}} u$$

with $u \in U$. Condition (ii) implies that $\gamma^e \equiv 1 \bmod l^{2^\sigma}$. Since $2^\sigma \geqslant e+d+1$ and since $(\delta^{l^{e+d}})^e \equiv 1 \bmod l^{e+d+1}$, we conclude that $u^e \in U_{e+d+1}$. It remains to show that $U_{e+d+1} \subset U^{l^e}$.

Consider the sequence

$$U_1 \hookrightarrow \overline{U_1} \hookrightarrow E_1 \to E_1/E_1^{l^{e+d}}. \tag{3.11}$$

Because d kills the torsion of $E_1/\overline{U_1}$, the kernel of the combined map from $\overline{U_1}$ to $E_1/E_1^{l^{e+d}}$ is contained in $\overline{U_1}^{l^e}$. It is a consequence of Leopoldt's conjecture that any $u \in U$ which is an l^eth power in $\overline{U_1}$ is an l^eth power in U_1. Given Leopoldt's conjecture, therefore, any unit mapping to 0 under (3.11) is an l^eth power. Since any element in $U^{l^{e+d+1}}$ is an l^{e+d}th power in E_1 and is therefore in the kernel of (3.11), we conclude that $U_{e+d+1} \subset U^{l^e}$, and the proposition is proved.

Though the success of Version 3.9 depends on fewer assumptions than Version 3.7, we believe that Version 3.7 will succeed in practice and that it is best to begin with this version since the choices made for c and σ are smallest possible and thus entail the least work. Unfortunately, if Version 3.7 fails, it is difficult to determine where the problem is, and it is probably best to increase c and σ slowly. In particular, we cannot simply jump to Version 3.9 because, although it is possible to compute a bound for the power of l dividing the class number of K, we know of no way to determine in advance an explicit bound for the torsion of $E_1/\overline{U_1}$.

4. Discrete logarithms

Let p be a prime number and let $\mathbb{F}_p$ denote the field of p elements. In this section we give an algorithm to solve the discrete logarithm problem in $\mathbb{F}_p$. Our method is fashioned closely after Gordon's solution to this problem (1993). We formulate our goal as follows.

Problem 4.1. *Given two elements t and v in $\mathbb{F}_p^*$ such that v is in the subgroup of $\mathbb{F}_p^*$ generated by t, determine the smallest integer $x \in \{0, \ldots, p-2\}$ such that $t^x = v$.*

We begin by reducing Problem 4.1 to an easier problem. We assume that $\mathbb{F}_p$ is given as $\mathbb{Z}/p\mathbb{Z}$, and for $t \in \mathbb{F}_p$ we let $r(t)$ be the smallest non-negative integer such that $t = r(t) + p\mathbb{Z}$. Recall that an element in $\mathbb{F}_p$ is called primitive if it generates $\mathbb{F}_p^*$.

Problem 4.2. *Let B and n be two positive integers. Given a primitive element $t \in \mathbb{F}_p^*$ such that $r(t)$ is B-smooth, an element $v \in \mathbb{F}_p^*$ such that $r(v) \leqslant p^{1/n}$, and a prime divisor l of $p-1$ which divides $p-1$ exactly e times, determine the least positive residue of $\log_t v$ modulo l^e.*

To reduce Problem 4.1 to Problem 4.2 requires three steps. First we factor $p-1$ into a product $\prod l^{e_l}$ of prime powers using the number field sieve factoring method. By the Chinese Remainder Theorem, it is sufficient to be able to compute discrete logarithms modulo l^{e_l} for each prime l. Next we find an integer $s \in \{0, \ldots, p-1\}$ such that $r(t^s v)$ is $p^{1/n}$-smooth. Then we need only compute $\log_t q$, where $r(q)$ is a prime dividing $r(t^s v)$ and is therefore less than $p^{1/n}$. Finally we find a primitive element t' such that $r(t')$ is B-smooth. To compute $\log_t v$, we compute $\log_{t'} v$ and $\log_{t'} t$ and use the identity

$$\log_t v = \log_{t'} v / \log_{t'} t.$$

The time needed to perform these reductions is estimated in §5, where we also analyse the running time of the following solution to Problem 4.2.

Algorithm 4.3 Let B and n be two positive integers and assume we are provided with the information given in Problem 4.2. Let σ be the least integer such that $2^\sigma > e$. Let $m = 2^h r(v)$ where h is chosen so that $p^{1/n} \leqslant m < 2p^{1/n}$. Let b_0 be the least positive residue of $p \bmod m$.

Step 1. Find the least non-negative integer D such that $|b_0 - Dm|$ is B-smooth. As shown in §2, this can be done using a sieve for the polynomial $b_0 - Xm$.

Step 2. Find a polynomial $f \in \mathbb{Z}[X]$ such that

(i) f is irreducible,
(ii) f is monic,
(iii) the degree of f is less than or equal to n,
(iv) the coefficients of f have absolute value less than $(D+1)m$,
(v) $f(m) \equiv 0 \bmod p$,
(vi) the constant term of f is B-smooth,
(vii) l does not divide the discriminant Δ of f.

To construct $f(X)$ let cp be the smallest multiple of p such that $cp \geqslant m^n$ and write cp in base m. In other words find coefficients $b_i < m$ such that

$$\sum_{i=0}^{n} b_i m^i = cp.$$

Then the polynomial

$$f(X) = \sum_{i=0}^{n} b_i X^i + D(X-m)$$

satisfies (ii)–(vi). We assume that $f(X)$ is irreducible and note that if it is not, any irreducible factor of $f(X)$ which has m as a root modulo p can be used in place of $f(X)$ in what follows. We also assume that l does not divide Δ and therefore is unramified in an extension of $\mathbb{Q}$ generated by a root of f. For comments on how to proceed if l does divide Δ, we refer the reader to the discussion preceding Algorithm 3.4 in §3.

Step 3. Let α be a root of f, let $K = \mathbb{Q}(\alpha)$, and let $\mathcal{O}_K$ be the ring of integers of K. Let $S_{\mathbb{Q}}$ be the set of prime numbers less than or equal to B, and let S_K be the set of prime ideals in $\mathcal{O}_K$ with norm at most B. Search for pairs of co-prime integers a, b such that $a+b\alpha$ and $a+bm$ are both B-smooth and such that l does not divide $N(a+b\alpha)$.

Stop when the number of pairs found is equal to $|S_{\mathbb{Q}}|+|S_K|+\sigma n-1$. Notice that $a+bm=g(a,b)$ where $g(X,Y)=X+Ym$ and that $N(a+b\alpha)=\hat{f}(a,b)$ where $\hat{f}(X,Y)=\sum b_i X^i(-Y)^{n-i}$. Thus we are looking for smooth values of two homogeneous polynomials. As seen in §2, these values can be found using a sieve.

Step 4. Let L be the set of pairs (a,b) found in step 3 and let $M=L\cup\{t\}$. We use a modification of Algorithm 3.4 to find a map $y\colon M\to\{1,\ldots,l^e-1\}$ such that $r(t)^{y(t)}r(v)^{-1}m\prod(a+bm)^{y(a,b)}$ and $\alpha\prod(a+b\alpha)^{y(a,b)}$ are l^eth powers.

We must adjust Algorithm 3.4 to account for the fact that we are interested in finding one set of exponents that creates two l^eth powers simultaneously. The exponent vectors of step 1 must therefore contain information for both constructions. For each $(a,b)\in L$, we have

$$|a+bm|=\prod_{q\in S_{\mathbb{Q}}} q^{e_q} \quad\text{and}\quad (a+b\alpha)=\prod_{\mathfrak{q}\in S_K}\mathfrak{q}^{e_{\mathfrak{q}}}.$$

We associate to each (a,b) the vector $v_{(a,b)}$ of length $|S_{\mathbb{Q}}|+|S_K|+n$ consisting of the exponents e_q, the exponents $e_{\mathfrak{q}}$, and the values $\lambda_{1,j}(a+b\alpha)$, for $1\leqslant j\leqslant n$. For the element t we let v_t be the vector with $\operatorname{ord}_q r(t)$ at all coordinates corresponding to the rational primes and zeros elsewhere. Finally, we define v_0 to be the vector with the value h at the coordinate corresponding to 2, with 0 at the coordinates corresponding to the other primes in $S_{\mathbb{Q}}$, with $\operatorname{ord}_{\mathfrak{q}}\alpha$ at the coordinates corresponding to the prime ideals $\mathfrak{q}\in S_K$, and with the values $\lambda_{1,j}(\alpha)$ in the last n coordinates. We now proceed with Algorithm 3.4, using v_0 in place of v_{m_0}, using $\{v_{(a,b)}\mid(a,b)\in L\}\cup\{v_t\}$ as the set of vectors v_m in step 1, and using the products

$$(a_k+b_k\alpha)\prod_{(a,b)\in M'}(a+b\alpha)^{w_k(a,b)},$$

where (a_k,b_k) is the kth element in $M-M'$, instead of the products $m_k\pi_m{}^{w_k(m)}$ in step 3. Notice that step 3 of Algorithm 3.4 yields the values $y(a,b)$. To find $y(t)$ requires finding the linear combination of the exponents $w_k(t)$ produced in step 2 of Algorithm 3.4 so that $r(t)^{y(t)}\prod(a+bm)^{y(a,b)}$ is an l^eth power in $\mathbb{Z}$. This concludes the description of Algorithm 4.3.

We claim that the integer $y(t)$ found in step 4 above is congruent to $\log_t v$ modulo l^e. To see this, let ϕ be the ring homomorphism from $\mathbb{Z}[\alpha]\to\mathbb{Z}/p\mathbb{Z}$ which sends α to $m+p\mathbb{Z}$. Notice that

$$\phi(m\prod(a+bm)^{y(a,b)})=\phi(\alpha\prod(a+b\alpha)^{y(a,b)})$$

and that therefore

$$\phi(r(t)^{y(t)}r(v)^{-1}m\prod(a+bm)^{y(a,b)})=t^{y(t)}v^{-1}\phi(\alpha\prod(a+b\alpha)^{y(a,b)}).$$

Since both $r(t)^{y(t)}r(v)^{-1}m\prod(a+bm)^{y(a,b)}$ and $\alpha\prod(a+b\alpha)^{y(a,b)}$ are l^eth powers, we conclude that $t^{y(t)}v^{-1}$ is an l^eth power in $\mathbb{F}_p$. It follows that $y(t)\equiv\log_t v \bmod l^e$.

5. Running time analysis

In this section we give evidence in support of the following conjecture.

Conjecture 5.1. *The discrete logarithm problem in $\mathbb{F}_p$ can be solved in time*

$$L_p[1/3;(64/9)^{1/3}+o(1)] \quad\text{for}\quad p\to\infty.$$

We first compute the running time for steps 3 and 4 of Algorithm 4.3 and then show that steps 1 and 2 of this algorithm as well as the reduction of Problem 4.1 to Problem 4.2 can be completed in the same running time. We continue with the notation of §4. Algorithm 4.3 depends on two parameters B and n. We consider these parameters now as functions of p and compute optimal asymptotic values for $p \to \infty$. More specifically, we write

$$B = L_p[s_0; c_0 + o(1)] \quad \text{and} \quad p^{1/n} = L_p[s_1; c_1 + o(1)],$$

in which case

$$n = \left(\frac{1}{c_1 + o(1)} \frac{\log p}{\log\log p}\right)^{1-s_1}. \tag{5.2}$$

All the $o(1)$s in these expressions are for $p \to \infty$. Our reason for considering B and $p^{1/n}$ as functions of the form $L_p[s; c]$ is to simplify our use of Theorem 2.1 in what follows.

In step 3 of Algorithm 4.3 we use a sieve to find pairs (a, b) so that $\hat{f}(a, b)$ and $g(a, b)$ are both B-smooth. Let $\frac{1}{2}C$ be the bound on the absolute value of the a and b that we consider and recall that $\hat{f}$ and g are homogeneous. We conclude from the calculation done in §2 that the running time for step 3 is bounded by

$$\pi(B)\,(n + \log B)^{O(1)} + C^2 \log\log B.$$

Substituting for B and letting $C = L_p[s_2; c_2 + o(1)]$, we obtain the bound

$$n \cdot L_p[s_0; c_0 + o(1)] + L_p[s_2; 2c_2 + o(1)].$$

In step 4 of Algorithm 4.3 we use a variation of Algorithm 3.4 to find the map y. This step takes time

$$B \cdot (n \log p)^{O(1)} + O(nB^2) + O(n^3 \log\log p) \leqslant (n \log p)^{O(1)} \cdot L_p[s_0; 2c_0 + o(1)].$$

The first term on the left represents the time required to find the values of the exponent vectors (see §6), the second term represents the time needed in step 2 of Algorithm 3.4 to do the linear algebra using Wiedemann's algorithm, and the third term represents the time needed for the linear algebra in step 3 of Algorithm 3.4.

Let T be the combined running time of steps 3 and 4 and write $T = L_p[s_T; c_T + o(1)]$. We minimize T subject to the condition that the number of a and b that we test in step 3 must be sufficiently large so that the number of pairs for which $a + b\alpha$ and $a + bm$ are both B-smooth is greater than $|S_{\mathbb{Q}}| + |S_K| + \sigma n - 1$. In other words, the inequality

$$L_p[s_2; 2c_2 + o(1)] \cdot P_{\mathbb{Q}} \cdot P_K > n \cdot L_p[s_0; c_0 + o(1)] \tag{5.3}$$

must hold, where $P_{\mathbb{Q}}$ is the probability that one of the $a + bm$ we test is smooth and P_K is the probability that one of the $a + b\alpha$ is smooth. We rely on the following conjecture to find expressions for these probabilities.

Conjecture 5.4. *Denote by $\psi(x, B)$ the number of positive integers less than x which are B-smooth. Let f be a polynomial in k variables over $\mathbb{Z}$ and assume that $|f(x_1, \ldots, x_k)| < A$ whenever x_i is in the interval $[-\frac{1}{2}C, \frac{1}{2}C]$ for $i \in \{1, \ldots, k\}$. Then the probability that $f(a_1, \ldots, a_k)$ is B-smooth for a_i chosen randomly from $[-\frac{1}{2}C, \frac{1}{2}C]$ is*

$$\psi(A, B)/A.$$

Under the assumption that Conjecture 5.4 is valid, we can use Theorem 2.1 to compute $P_{\mathbb{Q}}$ and P_K. Let $P_{\mathbb{Q}} = L_p[s_{\mathbb{Q}}; c_{\mathbb{Q}} + o(1)]$ and $P_K = L_p[s_K; c_K + o(1)]$. Since

$$a + bm < L_p[s_1; c_1 + o(1)] L_p[s_2; c_2 + o(1)]$$

and

$$\begin{aligned} N_{\mathcal{Q}}^{K}(a+b\alpha) &= (-b)^n f(-a/b) \\ &< (n+1)\cdot D\cdot L_p[s_1; c_1+o(1)]L_p[s_2; c_2+o(1)]^n \\ &= (n+1)\cdot D\cdot L_p[s_1; c_1+o(1)]L_p[s_2+1-s_1; c_2/c_1+o(1)], \end{aligned}$$

we see that $s_{\mathcal{Q}} \geqslant -(\max\{s_1, s_2\}-s_0)$ and $s_K \geqslant -(\max\{s_1, s_2+1-s_1\}-s_0)$. Thus condition (5.3) implies at the very least that

$$s_2 \geqslant \max\{s_1, s_2+1-s_1, s_2\}-s_0.$$

Subject to this inequality, $\max\{s_0, s_2\}$ is minimized when $s_0 = s_2 = \frac{1}{3}$ and $s_1 = \frac{2}{3}$, and we see that s_T is at least $\frac{1}{3}$. We claim that these choices for s_0, s_1 and s_2 actually yield $s_T = \frac{1}{3}$. Formula (5.2) shows that $n = O(\log p)^{1/3}$ when $s_1 = \frac{2}{3}$. The definition of D, Conjecture 5.4, and Theorem 2.1 imply that if $s_1 = \frac{2}{3}$, then $D = L_p[\frac{1}{3}; c_D+o(1)]$ for some constant c_D. It is now an easy matter to check that when $s_0 = s_2 = \frac{1}{3}$ and $s_1 = \frac{2}{3}$ condition (5.3) is met and that $s_T = \frac{1}{3}$. Moreover, in this case we find that

$$\begin{aligned} a+bm &< L_p[\tfrac{2}{3}; c_1+o(1)], \\ N_{\mathcal{Q}}^{K}(a+b\alpha) &< L_p[\tfrac{2}{3}; c_1+c_2/c_1+o(1)], \\ P_{\mathcal{Q}} &= L_p[\tfrac{1}{3}; -(c_1/3c_0)(1+o(1))], \\ P_K &= [\tfrac{1}{3}; -(1/3c_0)(c_1+c_2/c_1)(1+o(1))], \end{aligned}$$

and inequality (5.3) becomes

$$2c_2-(1/3c_0)(c_1+c_2/c_1+c_1) > c_0. \tag{5.5}$$

Minimizing $c_T = \max\{2c_0, 2c_2\}$ subject to (5.5) leads to the conclusion that the optimal values for B and n are

$$L_p[\tfrac{1}{3}; (\tfrac{8}{9})^{1/3}+o(1)] \quad \text{and} \quad ((3+o(1))\log p/\log\log p)^{1/3}.$$

In this case, $p^{1/n} = L_p[\frac{2}{3}; (\frac{1}{3})^{1/3}+o(1)]$, the bound on the absolute value of a and b is $L_p[\frac{1}{3}; (\frac{8}{9})^{1/3}+o(1)]$ and the running time T equals

$$L_p[\tfrac{1}{3}; (64/9)^{1/3}+o(1)].$$

We show now that with B and n equal to the above optimal values, steps 1 and 2 of Algorithm 4.3 as well as the reduction from Problem 4.1 to Problem 4.2 can each be accomplished in time T. The sieve in step 1 requires time

$$\pi(B)(1+\log B)^{O(1)}+D\log\log B,$$

which equals $L_p[\frac{1}{3}; c+o(1)]$ where $c = \max\{(\frac{8}{9})^{1/3}, c_D\}$. Using Conjecture 5.4, we find that $c_D = (1/72)^{1/3}$ and so step 1 runs in time less than T. Step 2 is easily dismissed since expanding cp in base m clearly takes very little time. The reduction to Problem 4.2 is broken into three steps. First the number $p-1$ is factored using the number field sieve. This algorithm has a conjectured expected running time of $L_N[\frac{1}{3}; (64/9)^{1/3}+o(1)]$, where N is the integer being factored. Thus the expected time needed to factor $p-1$ is T. Next s is found such that $r(t^s v)$ is $p^{1/n}$-smooth. We look for s by picking integers at random from $\{0, \ldots, p-1\}$ and expect by Theorem 2.1 to find s within $L_p[\frac{1}{3}; (\frac{1}{3})^{4/3}+o(1)]$ trials. We test for smoothness by factoring candidates with the elliptic curve factoring method. Each trial requires time $L_p[\frac{1}{3}; 2(\frac{1}{3})^{2/3}+o(1)]$ (Lenstra Jr 1987). Thus the expected total time needed to find s is $L_p[\frac{1}{3}; (\frac{1}{3})^{4/3}+2(\frac{1}{3})^{2/3}+o(1)]$ and is therefore less than T. Finally, we find a smooth

primitive element in $\mathbb{F}_p$. The following result asserts that this too can be done in time T.

Proposition 5.6. *Assume the extended Riemann hypothesis. Then there exists a constant c such that, for all primes p, there exists a primitive element $t \in \mathbb{F}_p$ such that $r(t) \leqslant c(\log p)^6$.*

Proof. See Shoup (1990).

6. Constructing valuations

In Algorithm 3.4, and hence in Algorithm 4.3, it is necessary to be able to compute the valuation of an element in a number field at a given prime. In this section, we give a solution to this problem due to H. W. Lenstra Jr. Let $K = \mathbb{Q}(\alpha)$, where α is some algebraic number of degree n over $\mathbb{Q}$. The algorithm we present takes as input an order R in the ring of integers $\mathcal{O}_K$ of K and a prime ideal $\mathfrak{q}$ in R lying over some rational prime q. In particular, we assume a basis for R over $\mathbb{Z}$ and ideal generators for $\mathfrak{q}$ are given, all as polynomials in α. For instance, if α were integral, our input might be the order $\mathbb{Z}[\alpha]$ and the prime $(q, g(\alpha))$ where $g \in \mathbb{Z}[X]$ is any polynomial whose image in $\mathbb{F}_q[X]$ is irreducible and divides the image of the minimum polynomial of α in $\mathbb{F}_q[X]$. The algorithm then outputs the valuation of a given γ at all primes lying above $\mathfrak{q}$.

It is useful to distinguish between two sorts of prime ideals of an order R of K. We call a prime ideal $\mathfrak{q}$ non-singular if the localization of R at $\mathfrak{q}$ is a discrete valuation ring and singular otherwise.

Proposition 6.1. *Let R be an order of a number field K, and let $\mathfrak{q}$ be a prime of R. Then there exists an element γ in $K-R$ such that $\gamma\mathfrak{q} \subset R$. Furthermore, $\mathfrak{q}$ is singular if and only if $\gamma\mathfrak{q} \subset \mathfrak{q}$.*

Proof. The existence of γ is a consequence of Lemma 4.4.3 in Weiss (1969). Assume now that $\gamma\mathfrak{q} = R$. Then there exists $\pi \in \mathfrak{q}$ such that $\gamma\pi = 1$. Let r be an element in $R_\mathfrak{q}$. We can write $r = r_1/r_2$ with $r_1 \in R$ and $r_2 \in R - \mathfrak{q}$. If r is not a unit then $r_1 \in \mathfrak{q}$ and therefore $\gamma r_1 \in R$. Furthermore $r_1 = \pi\gamma r_1$. We conclude that every element of $R_\mathfrak{q}$ which is not a unit is in the ideal generated by π and therefore that $R_\mathfrak{q}$ is a discrete valuation ring. Conversely, assume $R_\mathfrak{q}$ is a discrete valuation ring. Note that for any prime ideal $\mathfrak{p}$ not equal to $\mathfrak{q}$, we can find an element π in $\mathfrak{q}$ but not in $\mathfrak{p}$. Since $\gamma\pi \in R$, we see that γ is integral at all primes other than $\mathfrak{q}$. Therefore, not only is γ not in R, it is not in $R_\mathfrak{q}$. We conclude that $1/\gamma \in \mathfrak{q}R_\mathfrak{q}$ and write

$$1/\gamma = r_1/r_2$$

with $r_1 \in \mathfrak{q}$ and $r_2 \in R - \mathfrak{q}$. But then $\gamma r_1 \notin \mathfrak{q}$, and Proposition 6.1 is proved.

Algorithm 6.2. Let $\{\omega_i\}$ be a $\mathbb{Z}$-basis of R and $\{\delta_j\}$ a set of generators of $\mathfrak{q}$.

Step 1. Find an element $\gamma \notin R$ such that $\gamma\mathfrak{q} \subset R$ as follows. Let $(a_{i,j,k})$ be the three-dimensional matrix determined by the equations

$$\delta_j \omega_i = \sum_k a_{i,j,k}\omega_k,$$

and let $(\bar{a}_{i,j,k})$ be the corresponding matrix obtained by reducing the entries modulo q. For each j the determinant of $(a_{i,j,k})$ is divisible by q, and so the linear

transformation determined by $(\bar{a}_{i,j,,k})$ has a non-trivial kernel which we call B_j. Now choose $\bar{b} \in \bigcap B_j - \{0\}$, and let $b = (b_i)$ be a vector in $\mathbb{Z}^n$ which maps to $\bar{b}$. Then the element

$$\sum b_i w_i / q$$

has the desired properties.

Step 2. Determine whether $\mathfrak{q}$ is non-singular or singular by making use of the fact that $\mathfrak{q}$ is singular if and only if $\gamma\mathfrak{q} \subset \mathfrak{q}$. If $\mathfrak{q}$ is singular, then continue with step 3. If $\mathfrak{q}$ is non-singular, then for all $z \in R$, the valuation corresponding to $\mathfrak{q}$ is calculated by means of the equation

$$\operatorname{ord}_{\mathfrak{q}} z = \max\{m \mid \gamma^m z \in R\},$$

and the algorithm terminates.

Step 3. Let $R' = R[\gamma]$. Note that $\gamma \in \mathcal{O}_K$ since it multiplies the finitely generated group $\mathfrak{q}$ into itself and therefore that $R' = R[\gamma]$ is an order in $\mathcal{O}_K$ strictly containing R. Note also that $\mathfrak{q}$ is an ideal in R' since $R' \cdot \mathfrak{q} \subset \mathfrak{q}$. Let F denote the residue field of R at $\mathfrak{q}$ and $\bar{\gamma}$ the image of γ in $R'/\mathfrak{q}$. Proceed by finding the minimum polynomial $\overline{f'}$ of $\bar{\gamma}$ over F and factoring it into irreducibles over F. To each irreducible factor $\overline{g'}$ of $\overline{f'}$ there corresponds a prime ideal $\mathfrak{q}'$ in R' of the form $\mathfrak{q} + R' \cdot g'(\gamma)$ where $g' \in R[X]$ maps to $\overline{g'}$ in $F[X]$. Furthermore all prime ideals of R' lying over $\mathfrak{q}$ arise in this way.

Step 4. For each prime $\mathfrak{q}' \mid \mathfrak{q}$, apply steps 1, 2, and 3 to the pair R', $\mathfrak{q}'$. This concludes the description of Algorithm 6.2.

Since $|\mathcal{O}_K/R|$ is finite and the number of primes of $\mathcal{O}_K$ lying above $\mathfrak{q}$ is bounded by n, the algorithm terminates. Moreover, if $R = \mathbb{Z}[\alpha]$, then $|\mathcal{O}_K/R| < \sqrt{|\Delta|}$, where Δ is the discriminant of the minimum polynomial of α. In this case the algorithm must be repeated at most $O(n \log |\Delta|)$ times before all the valuations above $\mathfrak{q}$ are constructed. It is easily checked that each implementation of steps 1–3 requires no more than time $(n \log q)^{O(1)}$. We conclude that when $R = \mathbb{Z}[\alpha]$, the total running time for Algorithm 6.2 is less than $(n \log q \log |\Delta|)^{O(1)}$.

I thank Hendrik Lenstra for the many ideas he shared concerning this work.

References

Buhler, J. P., Lenstra, H. W. & Pomerance, C. 1993 Factoring integers with the number field sieve. In *The development of the number field sieve* (ed. A. K. Lenstra & H. W. Lenstra Jr), *Lecture Notes in Math.* Heidelberg: Springer-Verlag.

Canfield, E. R., Erdös, P. & Pomerance, C. 1983 On a problem of Oppenheim concerning 'factorisatio numerorum'. *J. Number Theory* **17**, 1–28.

Cohen, H. & Lenstra, H. W. 1983 Heuristics on class groups of number fields. In *Number theory, Noordwijkerhout 1983* (ed. H. Jager), *Lecture Notes in Math.*, vol. 1068, pp. 33–62. Heidelberg: Springer–Verlag.

Coppersmith, D., Odlyzko, A. M. & Schroeppel, R. 1986 Discrete logarithms in GF(p). *Algorithmica* **1**, 1–15.

Gordon, D. 1993 Discrete logarithms using the number field sieve. *Siam J. Discrete Math.* **6**, 124–138.

Lang, S. 1970 *Algebraic number theory*. Reading, Massachusetts: Addison-Wesley.

Lang, S. 1990 *Cyclotomic fields I and II*. New York: Springer-Verlag.

Lenstra, A. K. & Lenstra, H. W. 1990 Algorithms in number theory. In *Handbook of theoretical computer science* (ed. J. van Leeuwen), vol. A (*Algorithms and complexity*), pp. 673–715. Amsterdam: Elsevier.

Lenstra, A. K., Lenstra, H. W. & Lovász, L. 1983 Factoring polynomials with rational coefficients. *Math. Ann.* **261**, 515–534.

Lenstra, A. K., Lenstra, H. W., Manasse, M. S. & Pollard, J. M. 1993 The number field sieve. In *The development of the number field sieve* (ed. A. K. Lenstra & H. W. Lenstra Jr), *Lecture Notes in Math.* Heidelberg: Springer-Verlag.

Lenstra, H. W. 1987 Factoring integers with elliptic curves. *Ann. Math.* **126**, 649–673.

Lenstra, H. W. 1990 Algorithms for finite fields. In *Number theory and cryptography* (ed. J. H. Loxton), *London Math. Soc. Lecture Note Ser.* **154**, 76–85. Cambridge University Press.

Lenstra, H. W. 1992 Algorithms in algebraic number theory. *Bull. Am. math. Soc.* **26**, 211–244.

McCurley, K. 1990 The discrete logarithm problem. In *Cryptology and computational number theory* (ed. C. Pomerance), *Proc. of Symposia in Applied Mathematics* **42**, 49–74. Providence, Rhode Island: AMS.

Shoup, V. 1990 Searching for primitive roots in finite fields. In *Proc. 22nd Annual ACM Symp. on Theory of Computing* (*STOC*), pp. 546–554. New York: ACM.

Washington, L. C. 1982 *Introduction to cyclotomic fields.* New York: Springer-Verlag.

Weiss, E. 1969 *Cohomology of groups.* New York: Academic Press.

Wiedemann, D. H. 1986 Solving sparse linear equations over finite fields. *IEEE Trans. Info. Theory* **32**, 54–62.