INTERNATIONAL COMPUTER SCIENCE INSTITUTE

1947 Center St. • Suite 600 • Berkeley, California 94704-1198 • (510) 643-9153 • FAX (510) 643-7684

Deciding Properties of Polynomials without Factoring

T. Sander and M.A. Shokrollahi

TR-97-027

August 1997

Abstract

The polynomial time algorithm of Lenstra, Lenstra, and Lovász [17] for factoring integer polynomials and variants thereof have been widely used to show that various computational problems in number theory have polynomial time solutions. Among them is the problem of factoring polynomials over algebraic number fields, which is used itself as a major subroutine for several other algorithms. Although a theoretical breakthrough, algorithms based on factorization of polynomials over number fields are notoriously slow and hard to implement, with running times ranging between $O(n^{12+\varepsilon})$ and $O(n^{18+\varepsilon})$ depending on which variant of the lattice basis reduction is used. Here, n is an upper bound for the maximum of the degrees and the bit-lengths of the coefficients of the polynomials involved. On the other hand, in many situations one does not need the full power of factorization, so one may ask whether there exist faster algorithms in these cases.

In this paper we develop more efficient Monte Carlo algorithms to decide certain properties of roots of integer polynomials, without factoring them. Such problems arise, e.g., when solving systems of algebraic equations. Our methods applied to this situation give thus information about the solutions of such systems of equations.

Assuming the validity of the Extended Riemann Hypothesis, our algorithms run in time $O(n^{6+\varepsilon})$ in worst case, though they usually terminate much faster if the input polynomials do not have the properties the algorithm is testing. Besides this substantial improvement in the running time, our algorithms have the advantage of being conceptually easy. Their building blocks are gcd-computations in polynomial rings over finite fields, and primality tests for integers. However, despite the simplicity

of our algorithms, their analysis is involved and uses tools from algebraic and analytic number theory.

Our methods yield polynomial time algorithms even in cases where the factorization method does not. We exhibit such an example by showing that the language consisting of pairs (g, m) where g is a monic irreducible polynomial such that all its roots are integral linear combinations of mth roots of unity, is in co-$\mathcal{RP}$. Currently, we do not know of any deterministic polynomial time algorithm to decide this problem, even if we assume the validity of the Extended Riemann Hypothesis. We will also show that computing the minimal m such that (g, m) belongs to this language is intractable by means of present methods: we prove that this problem is polynomial time equivalent to that of computing the largest square free divisor of an integer.

1 Introduction

Suppose you are given an irreducible integer polynomial g, and you are asked to decide whether g is *normal*, i.e., whether all the roots of g can be written as rational linear combinations of powers of one of the roots. A situation where such a problem might arise is, e.g., when solving a system of algebraic equations which is known to have only finitely many solutions: standard methods based on elimination theory yield a univariate polynomial with the property that all coordinates of all the solutions to that system are zeros of that polynomial. Once one knows that this polynomial is normal, one can try to express all of its roots as rational linear combinations of powers of one of the roots, say α. This might be advantageous, when one wants to have all the roots of g to a high precision: in this case, one only needs to find a good approximation to α, and plug this new value into the formulas for the other roots.

The normality problem can be decided in polynomial time, by using factorization of polynomials over algebraic number fields [8, 15, 16, 13]: one factors $g \in K[T]$, where $K = \mathbb{Q}[x]/(g)$ is the field generated by a root of g over $\mathbb{Q}$. In this case, one also obtains representations for the roots as rational linear combinations of powers of the root T mod $g(x)$ of g in K.

This algorithm is very elegant, but quite slow, and hard to implement: Assuming that the bit-length of the coefficients and the degrees of the polynomials are simultaneously bounded by n, (an assumption we will adapt in the sequel) the factorization algorithms cited above have cost $O(n^{18})$. This makes the problem hard even for small input sizes. (One may be able to reduce this to $O(n^{12})$ using faster versions of the LLL-algorithm due to Schönhage [23]; however, we did not find an explicit reference to such an improvement in the literature.)

Suppose now that we are given a pool of irreducible integer polynomials, from which we want to extract one that is normal, and in that case, to compute the roots of that polynomial as rational linear combinations of powers of one of the roots. Of course, we don't want to spend too much time on polynomials which turn out to be non-normal at the end of the lengthy factorization algorithm. Assuming that the pool contains much more non-normal than normal polynomials, we would thus need a substantially faster algorithm to quickly sort out the non-normal polynomials, and to pass on a normal polynomial to the factorization module as soon as one is found. The algorithm could also have a constant error probability, if only it does not mistakenly identify a non-normal polynomial as normal.

Such a "normality test" is the prototype of decision problems for which we will develop fast Monte Carlo algorithms in this paper. These algorithms are part of a larger body of research motivated by a question posed by B. Sturmfels: given an integer polynomial with Abelian Galois group, design an efficient algorithm that expresses the roots of the polynomial as radical expressions.

In the sequel let g and h be monic irreducible integer polynomials, and let K, and L be fields generated by a root of g and h, respectively. We say that K is normal (or Galois), if g is normal. We will develop Monte Carlo algorithms for the following decision problems:

(1) Is K normal?

(2) If K is normal, is $K \subset L$?

(3) Is K an Abelian extension of $\mathbb{Q}$?

(4) Is K a cyclic extension of $\mathbb{Q}$?

All these problems can be solved in deterministic polynomial time using factorization of polynomials over number fields: for example, to decide Problem (3), one can factor g over K. Each zero of g yields an automorphism of K fixing $\mathbb{Q}$. One can test in polynomial time whether these automorphisms commute.

Factorization of polynomials yields often much more information than is needed in a specific situation. So, one might ask whether there are problems for which the factorization method does not lead to a polynomial time algorithm, while our method does. In fact, the answer to this question is positive: we will design a fast Monte Carlo algorithm for deciding whether the field K lies in the *cyclotomic field* $\mathbb{Q}(\zeta_m)$ generated over $\mathbb{Q}$ by a primitive mth root of unity ζ_m:

(5) Is K a subfield of $\mathbb{Q}(\zeta_m)$?

Note that m contributes only logarithmically to the input size of this problem. The reason why factorization methods do not help here is that the mth cyclotomic polynomial can have a (dense) representation of exponential size in $\log(m)$. In fact, we do not know of *any* deterministic polynomial time algorithm for deciding the above problem.

The deterministic versions to problems (1)–(4) above all run in time $O(n^{18})$ (or $O(n^{12})$), while our Monte Carlo algorithms use only $O(n^{6+\varepsilon})$ time for problems (1), (2), (4), and *even* (5), if we assume the validity of the Extended Riemann Hypothesis (ERH). Our algorithm for Problem (3) uses $O(n^{10+\varepsilon})$ time. However, if we assume that the bit-length of the coefficients of the polynomial defining K is polylogarithmic in the degree n, our algorithm uses $O(n^{6+\varepsilon})$ time, while factorization algorithms are uneffected by this assumption.

Besides this major improvement in running time, our algorithms are conceptually easier than their deterministic counterparts, and enjoy the main advantages of randomized algorithms in general. Their building blocks are gcd computations in polynomial rings over finite fields and primality testing of integers, both theoretically and practically well understood tasks. (For the latter to be doable in polynomial time, we will assume the validity of ERH.)

As in other randomized algorithms, we will also use the concept of a "witness." For example, we will show in Section 5 that any prime p not dividing the discriminant of g such that g mod p has irreducible factors of different degrees is a witness for the non-normality of g. We will proceed by showing that if g is not normal, then at least an $O(1/n)$-fraction of all primes is a witness for the non-normality of g. To obtain an effective algorithm from results of this type, we prove in Section 4 the Witness-Lemma which is our main theoretical tool for analyzing our algorithms. It gives a quantitative version of the above mentioned density result.

Our algorithms for the other four problems are very similar to the normality test, and also use the Witness-Lemma for their analysis.

The solutions for problems (1), (2), (3), (5) are $\mathcal{RP}$-algorithms, while the solution to Problem (4) is a $\mathcal{BPP}$-algorithm. A speed up of all of these algorithms by a factor of n can be obtained by working with numbers whose primality is assured only up to some degree of certainty, on the cost of turning the $\mathcal{RP}$-algorithms into $\mathcal{BPP}$-algorithms. We will show

how to do that in the example of the normality test only, since the other tests can be handled similarly.

The paper is organized as follows: In Section 3 we will introduce some basic facts from algebraic number theory. Section 4 contains the Witness Lemma, the main tool for analyzing the error probability of our algorithms. It estimates from below the number of rational primes among positive integers less than or equal to some number k, which have a prescribed decomposition pattern in a normal field L. The estimate is based on k and the size of a certain subset of the Galois group of L. The proof of this lemma uses tools from analytic number theory, like effective versions of the Chebotarev Theorem [12, 22], and Odlyzko's bound on the discriminant of number fields [21].

In sections 5–9 we proceed with the description and analysis of our algorithms. Section 10 contains a hardness result: we show that computing the smallest m such that $K \subseteq \mathbb{Q}(\zeta_m)$ is as hard as computing the largest square-free divisor of an integer. The only known algorithms for the latter problem are about as costly as factorization of integers. This shows that the problem of computing a minimal m with the above property is, according to current results, computationally hard.

2 Notation and Elementary Facts

The cost measure of our algorithms will be the bit-complexity, and we will use the "soft O" notation to ignore logarithmic factors: $g = \tilde{\mathcal{O}}(n)$ means that $g = \mathcal{O}(n(\log n)^\ell)$ for some constant ℓ. For the analysis of our algorithms we will use some well known results gathered in this section. The reader can consult [5, Chapters 2 and 3] for a discussion of these topics and further references.

Throughout the paper we will use the following notation: log will denote the logarithm to the basis 2, ln is the natural logarithm, $\mathrm{disc}(g)$ denotes the *absolute value* of the discriminant of the polynomial g, $L_\infty(g)$ is the maximum of the absolute values of the coefficients of g, $\lambda(g)$ denotes $\ln(L_\infty(g))$, O_K is the ring of integers of the number field K, and $\mathbb{Q}(\zeta_m)$ denotes the cyclotomic field generated over $\mathbb{Q}$ by a primitive mth root of unity.

Using the Schönhage-Strassen algorithm, addition, subtraction, multiplication, and the remainder of the division of an ℓ-bit integer and a k-bit integer can be computed with $\tilde{\mathcal{O}}(\mu)$ operations, where μ is the maximum of ℓ and k. The same is true for computing the greatest common divisor (gcd) of an ℓ-bit and a k-bit integer. In particular, each arithmetic operation in the finite field $\mathbb{F}_p$, p a prime, has cost $\tilde{\mathcal{O}}(\log(p))$. Addition, subtraction, multiplication, and the remainder of the division of polynomials of degree at most n over $\mathbb{F}_p$ all have cost $\tilde{\mathcal{O}}(n\log(p))$. Computing the gcd of polynomials of degree at most n over $\mathbb{F}_p$ has cost $\tilde{\mathcal{O}}(n\log(p))$. In particular, one can compute $\gcd(g, x^p - x)$ over $\mathbb{F}_p$ in time $\tilde{\mathcal{O}}(n\log(p)^2)$: using repeated squaring one computes $x^p - x \bmod g$ with $\tilde{\mathcal{O}}(n\log(p))$ operations, and performs a gcd-computation. A polynomial f of degree n over $\mathbb{F}_p$ can be tested on squarefreeness in time $\tilde{\mathcal{O}}(n\log(p))$, as one has to compute the gcd of f and its derivative.

The *distinct degree factorization* (DDF) of a squarefree polynomial f over $\mathbb{F}_p$ is a set of pairs (g, d), where g is the product of all monic irreducible factors of f of degree d. If f has degree at most n, its DDF can be computed with $\tilde{\mathcal{O}}(n^2\log(p) + n\log(p)^2)$ [10, Algorithm 3.1].

For an integer polynomial g we denote by $L_\infty(g)$ the maximum of the absolute values of the coefficients of g, and by $\lambda(g)$ the quantity $\ln(L_\infty(g))$. If g has degree n, then the absolute value of its discriminant is bounded above by $L_\infty(g)^{2n-1}(n+1)^{(n-1)/2}n^{3n/2}$, hence $\ln(\mathrm{disc}(g)) \leq 2n(\ln(n)+\lambda(g))$. (The bound for the discriminant is obtained by applying the Hadamard inequality to the Sylvester matrix.)

Finally, we mention that, assuming the Extended Riemann Hypothesis (ERH), one can deterministically decide primality of an integer m in time $\tilde{\mathcal{O}}(\log(m)^4)$ using, e.g., Miller's primality testing algorithm [19].

3 Decomposition of Primes and the Frobenius Automorphism

Let L be a Galois number field with group G. Let $p \in \mathbb{Z}$ be a prime unramified in L. Denote by $\mathfrak{P}$ a prime divisor of p in L. Then there exists a unique $\sigma \in G$, called the *Frobenius automorphism* of $\mathfrak{P}$ and denoted by $\left(\frac{L/\mathbb{Q}}{\mathfrak{P}}\right)$ such that $\sigma(x) \equiv x^p \bmod \mathfrak{P}$ for all x in the ring O_L of integers of L. (See, e.g., [11, pp. 125–130].) This automorphism is the link to the Effective Chebotarev Density Theorem which will allow us to obtain quantitative estimates about the number of primes in a given interval which have a certain decomposition behavior.

Let

$$\mathrm{Frob}_L(p) := \left\{ \sigma \in G \;\middle|\; \exists \text{ a prime divisor } \mathfrak{P} \text{ of } p \text{ in } L \text{ such that } \sigma = \left(\frac{L/\mathbb{Q}}{\mathfrak{P}}\right) \right\}.$$

For $\tau \in G$ we have $\left(\frac{L/\mathbb{Q}}{\tau(\mathfrak{P})}\right) = \tau\left(\frac{L/\mathbb{Q}}{\mathfrak{P}}\right)\tau^{-1}$, hence $\mathrm{Frob}_L(p) = \mathrm{Cl}(\sigma)$ for any $\sigma \in \mathrm{Frob}_L(p)$, where $\mathrm{Cl}(\sigma)$ denotes the conjugacy class of σ in G.

The Frobenius automorphism gives information about the decomposition of p in the following way: the order of $\left(\frac{L/\mathbb{Q}}{\mathfrak{P}}\right)$ equals the residue class degree of $\mathfrak{P}$. But we can use this automorphism to obtain information on the decomposition behavior of primes even in the case of a non-normal subfield of L. Let $K \subset L$ be an extension of number fields where L is normal with Galois group $G = \mathrm{Gal}(L/\mathbb{Q})$. Let U be the subgroup of G which fixes K elementwise, and let $\sigma \in G$. The subgroup $\langle\sigma\rangle$ acts via right multiplication on the right cosets $U\tau$ of G by U. The orbits of this action are called *cycles*. The *length* of a cycle is the length of the corresponding orbit. Denote by $f_1, \ldots f_s$ the lengths of the different cycles of the action of $\langle\sigma\rangle$ on the cosets of G by U. We will use the following proposition to study the decomposition of a prime p in K with the help of $\mathrm{Frob}_L(p)$.

Proposition 1. *Let $K \subset L$ be an extension of number fields. Let L be normal, and p be a prime unramified in L and let $\sigma \in \mathrm{Frob}_L(p)$. Denote by $f_1, \ldots f_s$ the lengths of the different cycles of the action of $\langle\sigma\rangle$ on the cosets of G by U. Then p is the product of s prime divisors in K of degrees $f_1, \ldots f_s$.*

Proof. See [11, Proposition 2.7]. □

We call the list $[f_1, \ldots, f_s]$ the *decomposition pattern of p in K*. We call such a decomposition pattern *homogeneous*, if $f_i = f_j$ for $1 \leq i, j \leq s$, and call it *inhomogeneous* otherwise. The proposition shows that the decomposition pattern of an unramified prime

depends only on $\mathrm{Frob}_L(p)$. How can we compute decomposition patterns? This is done by means of the following classical result due to Dedekind (see, e.g. [9, Th. 4.8.13]).

Fact 2. *Suppose K is a number field given by g. Let p be a prime not dividing* $\mathrm{disc}(g)$. *Then pO_K has a decomposition $pO_K = \mathfrak{p}_1 \cdots \mathfrak{p}_s$ with prime ideals $\mathfrak{p}_i$ of O_K of degree f_i if and only if there exist irreducible polynomials $g_1, \ldots, g_s \in \mathbb{F}_p[X]$ with $g \bmod p = g_1 \cdots g_s$ of degrees $f_1, \ldots, f_s$, respectively.*

We say that a separable polynomial $h \in \mathbb{F}_p[x]$, has a *homogeneous factorization pattern*, if all its irreducible factors have the same degree; otherwise we say that h has an *inhomogeneous factorization pattern.*

4 The Witness Lemma

This section contains the main lemma used in the analysis of the randomized algorithms in the next sections. Let M be a number field and denote by L its normal closure. Let G denote the Galois group of L over $\mathbb{Q}$. Let W be a subset of G closed under conjugation. In our applications W will be a subset of G such that all p with $\mathrm{Frob}_L(p) \subseteq W$ are witnesses for some property of M, e.g., not being normal. The basic structure of our algorithms is as follows: we randomly select an integer p between 1 and a number m depending on the field M. If p is prime, we test whether $\mathrm{Frob}_L(p) \subseteq W$. If so, we know that M has the desired property, e.g., is not normal. If not, the algorithm tells us that M probably does not have the desired property. To analyze the error probability of such an algorithm, we need to know

$$\pi_W(x) \;:=\; \#\{p \in \mathbb{P} \mid p \le x, p \text{ unramified in } L, \mathrm{Frob}_L(p) \subseteq W\}.$$

A celebrated theorem of Chebotarev [6] states that asymptotically, as x goes to infinity, $\pi_W(x) \sim (\#W/\#G)x/\ln(x)$. In this section we will be interested in quantitative versions of this result for bounded x.

Lemma 3. *Let M be a number field, L be its normal closure, and denote by d the discriminant of L. Assuming ERH, there exist effectively computable absolute constants α_1, α_2, and α_3 such that if $x \ge \alpha_2 \ln(|d|)^{\alpha_1}$, then*

$$\pi_W(x) \ge \alpha_3 \frac{\#W}{\#G} \frac{x}{\ln(x)}.$$

One may take $\alpha_1 = 5/2$, $\alpha_2 = (13000)^{\alpha_1}$, and $\alpha_3 = 7/10$.

Proof. We may suppose that $M \neq \mathbb{Q}$ since the number $\pi(x)$ of primes less than or equal to x is lower bounded by $x/\ln(x)$ for $x \ge 17$, see [1, Th. 8.8.1].

Effective versions of Chebotarev due to Lagarias and Odlyzko [12], and Oesterlé [22] give

$$\pi_W(x) \ge \frac{\#W}{\#G}\Big(\mathrm{Li}(x) - \sqrt{x}(2\ln(|d|) + m\ln(x))\Big),$$

where $m = [L : \mathbb{Q}]$, d is the discriminant of L, and $\mathrm{Li}(x) = \int_2^x \mathrm{d}\tau/\tau$ is the logarithmic integral. From Odlyzko's bound on discriminants [21, pp. 381, (1.13)] we deduce that $m \le (254 + \ln(|d|))/20 \le 13\ln(|d|)$, since $L \neq \mathbb{Q}$. It follows that

$$\pi_W(x) \ge \frac{\#W}{\#G}\Big(\mathrm{Li}(x) - 13\sqrt{x}\ln(|d|)\ln(x)\Big).$$

Let $\alpha_1 := 5/2$, $\alpha_2 := (13000)^{\alpha_1}$. Then the above inequality shows that for $x \ge \alpha_2 \ln(|d|)^{\alpha_1}$ we have

$$\pi_W(x) \ge \frac{\#W}{\#G}\left(\mathrm{Li}(x) - 13\sqrt{x}\ln(x)(x/\alpha_2)^{1/\alpha_1}\right).$$

Noting that $\sqrt{x}/\ln^2(x) \ge x^{1/\alpha_1}/100$ for $x \ge 2$, we deduce that

$$\pi_W(x) \ge \frac{\#W}{\#G}\left(\mathrm{Li}(x) - \frac{x}{10\ln(x)}\right).$$

It is easily verified that $\mathrm{Li}(x) > 4x/(5\ln(x))$ for $x \ge 5$, hence the assertion follows with $\alpha_3 := 7/10$. □

Later we will need this theorem in the following two forms:

Corollary 4. *Let K be a number field of degree n given by an irreducible integer polynomial g, and L be its normal closure with Galois group G over $\mathbb{Q}$. Assuming ERH, there exist effectively computable absolute constants $\alpha_1, \alpha_2, \alpha_3$ such that for $x \ge \alpha_2(n!\ln(\mathrm{disc}(g)))^{\alpha_1}$ we have*

$$\#\{p \in \mathbb{P} \mid p \le x, p \text{ unramified in } L,\ p \not| \mathrm{disc}(g), \mathrm{Frob}_L(p) \subseteq W\} \ge \alpha_3 \frac{\#W}{\#G}\frac{x}{\ln(x)}.$$

One may take $\alpha_1 := 5/2$, $\alpha_2 := (13000)^{\alpha_1}$, and $\alpha_3 := 3499/5000$.

Proof. Let us call the number on the left-hand side ℓ. First note that $\ell \ge \pi_W(x) - \log(\mathrm{disc}(g))$. Further $\log(\mathrm{disc}(g)) \le 2(x/\alpha_2)^{1/\alpha_1}/n!$ by our assumption on x. Notice that we have for $x \ge 2$ the inequality $(x/\alpha_2)^{1/\alpha_1} \le x/(10000\ln(x))$, and that $\#W/\#G \ge 1/n!$. As a result, $\log(\mathrm{disc}(g)) \le \#W x/(5000\ln(x)\#G)$. Since the absolute value of the discriminant d of L is at most $\mathrm{disc}(g)^{n!}$, we can apply the last lemma to see that

$$\ell \ge \frac{\#W}{\#G}\left(\frac{7}{10} - \frac{1}{5000}\right)\frac{x}{\ln(x)} = \frac{\#W}{\#G}\frac{3499}{5000}\frac{x}{\ln(x)},$$

which proves the assertion. □

Corollary 5. *Suppose that K and M are fields generated by monic irreducible integer polynomials g and h of degrees m and n, respectively. Furthermore, suppose that K is normal, and let L be the normal closure of the compositum KM. Let $\alpha_1, \alpha_2, \alpha_3$ be the constants from the previous corollary. Assuming ERH, we have for $x \ge \alpha_2(mn!\ln(\mathrm{disc}(g)\mathrm{disc}(h)))^{\alpha_1}$*

$$\#\{p \in \mathbb{P} \mid p \le x,\ p \text{ unramified in } L,\ p \not| \mathrm{disc}(g)\mathrm{disc}(h),\ \mathrm{Frob}_L(p) \subseteq W\} \ge \alpha_3 \frac{\#W}{\#G}\frac{x}{\ln(x)}.$$

PROOF. Again, let us call the number on the right-hand side ℓ. We first note that $\ell \geq \pi_W(x) - \log(\mathrm{disc}(g)\mathrm{disc}(h))$. Further, $\log(\mathrm{disc}(g)\mathrm{disc}(h)) \leq 2(x/\alpha_2)^{1/\alpha_1}/(mn!)$, which, as in the proof of the previous corollary, gives us $\log(\mathrm{disc}(g)) \leq \#Wx/(5000\ln(x)\#G)$: note that $L = K\hat{M}$ since K is normal, where $\hat{M}$ is the normal closure of M, which shows that $\#G \leq [K:\mathbb{Q}][\hat{M}:\mathbb{Q}] \leq mn!$. For the discriminant $d(L)$ of L we have the inequality

$$|d(L)| \leq |d(\hat{M})^m d(K)^{[\hat{M}:\mathbb{Q}]}| \leq \mathrm{disc}(h)^{mn!}\mathrm{disc}(g)^{n!}.$$

(We use the well-known fact that $d(N_1N_2)$ divides $d(N_1)^{[N_2:\mathbb{Q}]}d(N_2)^{[N_1:\mathbb{Q}]}$, which follows from a theorem of Tôyama [25].) Hence, $x \geq \alpha_2(\ln(|d(L)|))^{\alpha_1}$ and, as in the proof of the previous corollary, we can apply Lemma 3 to obtain the assertion. □

Remark 6. The constants given in the last theorem are far from optimal. One could use techniques from [2] to optimize them.

5 Normality Test

In this section we will design a Monte Carlo algorithm with one-sided error to decide whether a given separable monic irreducible integer polynomial g is normal, i.e., whether the field K generated by a root of g contains all roots of g. Equivalently, we want to test whether K is a Galois extension of $\mathbb{Q}$. The idea of the algorithm is quite simple: we randomly select primes p from a certain interval and test whether g mod p has a homogeneous factorization pattern. If not, then K is not normal, and if yes, then this provides some evidence for the normality of K.

Let G be the Galois group of the normal closure L of K over $\mathbb{Q}$, and let U be the subgroup of G fixing K elementwise. Let

$$\begin{aligned} W_1(U) &:= \{\sigma \in G \mid \emptyset \neq \mathrm{Cl}(\sigma) \cap U \neq \mathrm{Cl}(\sigma)\} \\ W(U) &:= \sqrt{W_1(U)} := \{\sigma \in G \mid \exists\, m \in \mathbb{N}\colon \sigma^m \in W_1(U)\}. \end{aligned}$$

The following result states that $W(U)$ is the union of $\mathrm{Frob}_L(p)$ for p ranging over all primes with an inhomogeneous decomposition pattern, and gives an estimate on the size of $W(U)$.

Proposition 7. *Let K be a number field and p be a prime unramified in K. Then the following assertions hold:*

(1) *If K is normal then the decomposition pattern of p in K is homogeneous.*

(2) *If K is not normal, then a prime p has an inhomogeneous decomposition pattern, hence is a witness for the non-normality of K, iff* $\mathrm{Frob}_L(p) \subseteq W(U)$.

(3) *If K is not normal and of degree n over $\mathbb{Q}$, then $\#W(U) \geq \#G/n$.*

PROOF. (1) Assume K is normal and let p decompose as $p = \mathfrak{P}_1 \cdots \mathfrak{P}_s$ in K. The Galois group of K acts transitively on the prime divisors of p and thus $O_K/\mathfrak{P}_i \cong O_K/\mathfrak{P}_j$. So these fields have the same degree over $\mathbb{F}_p$.

(2) If K is not normal, then U is not a normal subgroup of G, hence there exists $\sigma \in U$ such that the conjugacy class of σ is not contained in U. Hence, $W_1(U)$ and $W(U)$ are not empty. Now suppose that $\mathrm{Frob}_L(p) \subseteq W(U)$ and let m be the smallest positive integer such that $\sigma^m \in W_1(U)$ for some $\sigma \in \mathrm{Frob}_L(p)$. Then there exists $\tau \in G$ such that $\tau\sigma^m\tau^{-1} \in U$. It is easily seen that $\tau, \tau\sigma, \ldots, \tau\sigma^{m-1}$ are different cosets modulo U, and that the cycle of the action of $\langle\sigma\rangle$ on $U\tau$ has length m. If all cycles had length m, then there would exist $\tau_1, \ldots, \tau_s$ such that $U\tau_j\sigma^i$, $j = 1, \ldots, s$, $i = 0, \ldots, m-1$ were all the cosets modulo U, which would show that $\tau\sigma^m\tau^{-1} \in U$ for all $\tau \in G$, a contradiction.

Conversely, assume that p has an inhomogeneous decomposition pattern, and let f_1 be the smallest residue class degree of a prime divisor of p in K. Let f_2 be another residue class degree, $f_2 > f_1$, and let $U\tau_1$ and $U\tau_2$ induce cycles of lengths f_1 and f_2, respectively. Then $\tau_1\sigma^{f_1}\tau_1^{-1} \in U$, but $\tau_2\sigma^{f_1}\tau_2^{-1} \notin U$, which shows that $\sigma^{f_1} \in W_1(U)$, i.e., $\sigma \in W(U)$.

(3) Since U is not normal, there exists σ such that $\sigma U\sigma^{-1} =: U^\sigma$ and U are different. Let $H := U \cap U^\sigma$. Then $(U \setminus H) \cup (U^\sigma \setminus H) \subseteq W_1(U)$, and we obtain $\#W_1(U) \geq 2(\#U - \#H)$. Since H is a proper subgroup of U, we have $\#W(U) \geq \#W_1(U) \geq \#U = \#G/n$. □

H. W. Lenstra has given a characterization of all groups G for which the bound in Part (3) of the last proposition is sharp.

Proposition 8. *In the situation of the preceding proposition we have: $\#W(U) = \#G/n$ if and only if G is a 2-group, $\#U = 2$, the normalizer of U in G has index 2 in G, and the nontrivial element of U is not the square of an element in G.*

Proof. If the index of the normalizer $N_G(U)$ of U in G is larger than 2, then there exist τ, σ such that U, U^τ, U^σ are all different, and there exists an element $\eta \in U^\tau \setminus (U \cup U^\sigma)$. This element certainly belongs to $W_1(U)$, which shows that in this case $\#W_1(U) \geq \#U + 1$. So, if $\#W_1(U) = \#U$, then we have that $\#H = \#U/2$, and $[G : N_G(U)] = 2$. In this case, H is a normal subgroup of U, which shows that H is the trivial subgroup, as U does not contain any nontrivial normal subgroup of G. (Recall that L is the normal closure of K.) Hence, $\#U = 2$. So, $\#W(U) = \#U$ if and only if $\#U = 2$, $\#W(U) = \#W_1(U)$, and the length of the conjugacy class of the nontrivial element u of U is 2. Suppose that G is not a 2-group. Then there exists a prime $p > 2$ such that p divides the order of $N_G(U)$. Let $x \in N_G(U)$ be an element of order p. Then $xu = ux$ for the nontrivial element u of U and $(xu)^p = u$. Hence, $xu \in W(U) = W_1(U)$, which shows that $xu = \sigma u\sigma^{-1}$ for any σ outside $N_G(U)$. This implies that $(xu)^p = u = \sigma u\sigma^{-1}$, a contradiction. Hence, G is a 2-group.

For proving the sufficiency of the conditions given for $\#W(U) = \#U$, we only need to show that if the nontrivial element u of U is not a square, then $W(U) = W_1(U)$, i.e., u is not a power of an element of $G \setminus U$. Suppose that $\tau \in G \setminus U$ is such that $\tau^m = u$. There exists an odd ℓ such that $m\ell$ modulo $\#G$ divides $\#G$, since G is a 2-group. As a result, $\tau^{ml} = u$, which shows that u is a square, a contradiction. □

Remark 9. The following example, also due to H. W. Lenstra shows that the above lower bound $\#G/n$ for $\#W(U)$ is sharp for infinitely many n. Let $G := D_4 \times C_{2^l}$ where $D_4 = \langle\sigma, \tau \mid \sigma^4 = \tau^2 = 1, \tau\sigma\tau^{-1} = \sigma^{-1}\rangle$ is the dihedral group and C_{2^l} is the cyclic group of order 2^l. Set $U := \langle\tau\rangle \times 1$. A famous theorem of Shafarevich [24] states that every solvable group G is realizable over $\mathbb{Q}$ as the Galois group of a number field L. Hence, we obtain K as the fixed field of U.

HOMOG(g, p)

Input: monic irreducible integer polynomial g, prime p
Output: YES, NO, FAIL
(1) Compute $h := g \bmod p$.
(2) **if** $\gcd(h, h') \neq 1$ **then return** FAIL;
(3) Compute the distinct degree factorization $h = h_1 \cdots h_n$ of h;
(4) **if** there exists $i \neq j$ such that $h_i, h_j \neq 1$ **then return** NO;
(5) **else return** YES

Figure 1: Algorithm HOMOG.

We now proceed with the description of our algorithm which was sketched at the beginning of the section. The homogeneity test (which is deterministic) is given in Figure 1.

Theorem 10. *Algorithm* HOMOG *outputs* NO *if and only if p does not divide* $\mathrm{disc}(g)$ *and the factorization pattern of $h := g \bmod p$ is not homogeneous. The running time of this algorithm is $\tilde{\mathcal{O}}(n\nu + n^2 \log(p) + n \log(p)^2)$, where ν is the maximum of $\lambda(g)$ and $\ln(p)$.*

PROOF. The prime p does not divide $\mathrm{disc}(g)$ if and only if g passes the test in line (2). For this reason we may assume in the following that p is a prime not dividing $\mathrm{disc}(g)$.

Let $h = h_1 \cdots h_n$ be the distinct degree factorization of h. Clearly, h has homogeneous factorization iff only one of the h_i is not equal to 1. Hence, the algorithm returns NO if and only if h does not have a homogeneous factorization pattern.

In line (1) we have to perform n divisions mod p, which takes $\tilde{\mathcal{O}}(n\nu)$ operations. Computing the distinct degree factorization of h takes $\tilde{\mathcal{O}}(n^2 \log(p) + n \log(p)^2)$ operations. This implies our result on the running time of HOMOG. □

Our final normality test is given in Figure 2. The constants α_1, α_2, and α_3 in that test are from Corollary 4.

Theorem 11. *If g is a normal polynomial, then* NORMALITY(g) *returns* NORMAL. *If g is not normal, then, assuming ERH,* NORMALITY(g) *returns* NON-NORMAL *with probability at least $1 - 1/e$ and uses $\tilde{\mathcal{O}}(n\ell^5 + n^2\mu\ell)$ bit operations, where $\ell := \max\{n, \ln(\lambda)\}$, $\mu := \max\{\ell, \lambda\}$, and $\lambda := \lambda(g)$.*

PROOF. If g is normal, then the factorization pattern of g modulo p is homogeneous for all p not dividing the discriminant of g. Hence, HOMOG(g, p) never returns NO, which means that NORMALITY always returns NORMAL.

Suppose that g is not normal. Then the number of "witnesses" p in the interval $[1, k]$, i.e., the number of p unramified in L and not dividing the discriminant of g such that g mod p has an inhomogeneous factorization pattern is at least $\alpha_3(\#W(U)/\#G)(k/\ln(k))$ by Corollary 4. By Proposition 7(3) this number is at least $\alpha_3 k/(n \ln(k))$. Hence, the randomly chosen p is not a witness of normality of g with probability at most $1 - 1/t$, where $t = n \ln(k)/\alpha_3$. Hence, the probability of not choosing a witness in any of the runs of the loop in line (2) is at most $(1 - 1/t)^t < 1/e$.

NORMALITY(g)

Input: monic irreducible integer polynomial g of degree n,
Output: NORMAL, NON-NORMAL
(1) Compute $k := \alpha_2\big(n!(2n)(\ln(n)+\lambda(g))\big)^{\alpha_1}$.
(2) **for** j **from** 1 **to** $n\ln(k)/\alpha_3$ **do**
(3) Choose at random an integer p from $\{1,\ldots,k\}$.
(4) **if** p is prime **then**
(5) **if** HOMOG(g,p) = NO **then return** NON-NORMAL and **stop** .
(6) **return** NORMAL.

Figure 2: Algorithm NORMALITY.

Computing k in line (1) takes $\tilde{\mathcal{O}}(\log(k)) = \tilde{\mathcal{O}}(\ell)$ bit operations. Assuming ERH, testing p for primality in line (4) takes $\tilde{\mathcal{O}}(\ell^4)$ operations. By Theorem 10 HOMOG(g,p) takes $\tilde{\mathcal{O}}(n\mu + n\ell^2)$ operations. As the loop is performed at most $\mathcal{O}(n\ell)$ times, this gives an overall running time of $\tilde{\mathcal{O}}(n\ell^5 + n^2\mu\ell)$, and proves the theorem. □

Remark 12. The major part, $\tilde{\mathcal{O}}(\ell^4)$, of the running time in each run of the loop performed in NORMALITY is contributed by the deterministic primality test in line (4). On the cost of turning our algorithm into a $\mathcal{BPP}$-algorithm, we can use a randomized primality test instead. We then have to analyze the probability of NORMALITY(g) returning NOT-NORMAL if g is normal. This can only happen if at least one of the primality tests in the loop of line (2) does not recognize the composite number p, and if this p passes the routine HOMOG and produces a non-homogeneous factorization pattern. However, if we run a randomized primality test like the Miller-Rabin or the Solovay-Strassen test $\tilde{\mathcal{O}}(\log(t))$ times—where t is the number of times the loop is performed—then with a constant probability all the p passing the primality tests are indeed prime. This turns the algorithm into a $\mathcal{BPP}$-algorithm, with constant two-sided error. The cost of performing the primality tests is $\tilde{\mathcal{O}}(\log(t)\ell^2)$ which clearly equals $\tilde{\mathcal{O}}(\ell^2)$. The whole algorithm would then have a running time of $\tilde{\mathcal{O}}(n\ell^3 + n\mu\ell^2)$.

6 Subfield problem

In this section we present a Monte Carlo algorithm with one-sided error which, given fields K and M with K normal, decides whether K is a subfield of M.

The basis for our algorithm is the following well-known theorem of M. Bauer (see, e.g. [20, Cor. 6.8]): K is a subfield of M if and only if all primes which split a factor of degree one in M do so in K. To apply this theorem, we randomly select primes, and test whether they split a factor of degree one in M. If so, we then test whether that prime splits completely in K. If not, then K is not a subfield of M. If yes, then there is some chance that K is a subfield of M. The algorithm is given in pseudocode in Figure 3. The constants α_1, α_2, α_3 are from Corollary 5.

Let L be the normal closure of KM, and let G be the Galois group of L over $\mathbb{Q}$. Assume that H is the normal subgroup of G corresponding to K and that U is the subgroup of G

corresponding to M. If K is not a subfield of M then any prime p such that $\mathrm{Frob}_L(p)$ is a subset of

$$W(U,H) := \{\sigma \in G \mid \mathrm{Cl}(\sigma) \cap U \neq \emptyset \text{ and } \mathrm{Cl}(\sigma) \cap H = \emptyset\}$$

serves as a witness for this property. The following proposition will be used later to show that our algorithm has small one-sided error, given that K is normal.

Proposition 13. *With the above notation we have:*

(1) *If $K \subset M$ then any prime p having a prime divisor of degree 1 in M also has a prime divisor of degree 1 in K.*

(2) *If $K \not\subseteq M$, then any prime p such that $\mathrm{Frob}_L(p) \subseteq W(U,H)$ splits a factor of degree 1 in L, but is not split in K.*

(3) *If $K \not\subseteq M$, then we have $\#W(U,H) \geq \#G/(2n)$, where $n = [M : \mathbb{Q}]$.*

PROOF. (1) Let $\mathfrak{P}$ be a prime divisor of p of degree 1 in L. Then $\mathfrak{p} := \mathfrak{P} \cap K$ is a prime divisor of p in K and $O_K/\mathfrak{p}$ embeds in $O_L/\mathfrak{P}$.

(2) Since H is a normal subgroup of G, the set $U \setminus (U \cap H)$ is contained in $W(U,H)$.

(3) Since $U \cap H$ is a proper subgroup of U, its size is at least $\#U/2 = \#G/(2n)$. □

Theorem 14. *Let g and h be monic irreducible integer polynomials of degrees m and n, $m \leq n$, and assume that g is normal. Let K be generated by a root of g. If $K \subseteq M$ for a field M generated by any of the roots of h, then* NORMAL-SUBFIELD(g,h) *returns* SUBFIELD. *Otherwise, assuming ERH,* NORMAL-SUBFIELD(g,h) *returns* NOT-SUBFIELD *with probability at least $1 - 1/e$ and uses $\tilde{\mathcal{O}}(n\ell^5 + n^2\ell\mu)$ bit operations, where $\ell := \max\{n, \ln(\lambda)\}$, $\mu := \max\{\ell, \lambda\}$, and $\lambda := \max\{\lambda(g), \lambda(h)\}$.*

PROOF. Let $K \subseteq M$ and assume that p does not divide $\mathrm{disc}(g)\mathrm{disc}(h)$. Then Fact 2 and Proposition 13(1) imply that if h mod p splits a linear factor over $\mathbb{F}_p$ then so does g mod p. Hence the condition in line (8) is never fulfilled, which means that NORMAL-SUBFIELD always returns SUBFIELD.

Suppose that $K \not\subseteq M$. Then using Corollary 5 and Part (3) of the above proposition, we see that the number of primes p in the interval $[1,k]$ which fail the test in line (8) is at least $\alpha_3 k/(2n\ln(k))$. (Note that $k \geq \alpha_2(n!m\ln(\mathrm{disc}(g)\mathrm{disc}(h))^{\alpha_1}$.) We obtain the assertion on the error probability of our algorithm as in the proof of Theorem 11.

Since $\log(k) = \tilde{\mathcal{O}}(\ell)$, computing k in line (1) takes $\tilde{\mathcal{O}}(\ell)$ operations. The primality test in line (4) takes $\tilde{\mathcal{O}}(\ell^4)$ operations. Computation of $\overline{g}$ and $\overline{h}$ in line (5) is accomplished with $\tilde{\mathcal{O}}(n\mu)$ operations, and the gcd's in line (6) use $\tilde{\mathcal{O}}(n\ell)$ operations. The loop is performed $\tilde{\mathcal{O}}(\ell n)$ times, which gives the desired estimate on the running time. □

Remark 15. Again, if we use a randomized primality test instead of the deterministic test in Line (4), we obtain a $\mathcal{BPP}$-algorithm with running time $\tilde{\mathcal{O}}(n\ell^4 + n^2\mu\ell)$.

NORMAL-SUBFIELD(g, h)

Input: monic irreducible integer polynomials g and h of degrees m and n, g normal, $m \leq n$
Output: SUBFIELD, NOT-SUBFIELD
(1) Compute $\lambda := \max\{\lambda(g), \lambda(h)\}$ and $k := \alpha_2\big(2n!n^2(\ln(n) + \lambda)\big)^{\alpha_1}$.
(2) **for** j **from** 1 **to** $2n\ln(k)/\alpha_3$ **do**
(3) Choose at random an integer p from $\{1, \ldots, k\}$.
(4) **if** p is prime **then**
(5) Compute $\overline{g} = g \bmod p$ and $\overline{h} = h \bmod p$.
(6) **if** $\gcd(\overline{g}, \overline{g}') \neq 1$ **and** $\gcd(\overline{h}, \overline{h}') \neq 1$ **then**
(8) **if** $\gcd(\overline{h}, X^p - X) \neq 1$ **and** $\gcd(\overline{g}, X^p - X) = 1$ **then**
(9) **return** NOT-SUBFIELD and **stop** .
(10) **return** SUBFIELD.

Figure 3: Algorithm NORMAL-SUBFIELD

SUBFIELD-OF-CYCLOTOMIC(g, n)

Input: a monic irreducible integer polynomial g of degree m and a non-negative integer n, $m \leq n$
Output: SUBFIELD, NOT-SUBFIELD
(1) **if** 2 runs of NORMALITY(g) return NORMAL **then**
(2) Compute $k := \alpha_2\big(2mn(\ln(m) + \lambda) + mn\ln(n)\big)^{\alpha_1}$, where $\lambda = \lambda(g)$.
(3) **for** j **from** 1 **to** $2(2\ln(k)/\alpha_3)$ **do**
(4) Choose at random an integer l from $\{1, \ldots, k/n\}$.
(5) Compute $p = ln + 1$
(6) **if** p is prime **then**
(7) Compute $\overline{g} = g \bmod p$.
(8) **if** $\gcd(\overline{g}, \overline{g}') \neq 1$ **then**
(9) **if** $\gcd(\overline{g}, X^p - X) = 1$ **then return** NOT-SUBFIELD and **stop** .
(10) **return** SUBFIELD.

Figure 4: Algorithm SUBFIELD-OF-CYCLOTOMIC.

7 Subfield of a Cyclotomic Field Test

Let us now turn to a problem for which to the best of our knowledge no deterministic polynomial time algorithm is known. We describe a Monte Carlo algorithm showing that the language consisting of the pairs (g, n), g an irreducible monic integer polynomial such that the splitting field K of g is contained in the cyclotomic field $\mathbb{Q}(\zeta_n)$, is in co-$\mathcal{RP}$. The reason that the currently available factorization approach does not apply here is that the nth cyclotomic polynomial can have a (dense) representation of exponential size in $\log(n)$.

The idea of our algorithm is as follows. Assume for the moment that K is normal. Then as in the algorithm NORMAL-SUBFIELD we are going to test if a prime p that splits completely over $\mathbb{Q}(\zeta_n)$ does so over K as well. Fortunately we do know the set of primes which split completely in $\mathbb{Q}(\zeta_n)$: these are exactly the primes $p \equiv 1 \bmod m$ (see, e.g., [26, Chapter 2]). This additional information allows us now to obtain a randomized efficient solution to the problem "$K \subset \mathbb{Q}(\zeta_n)$?" (see Figure 4).

Theorem 16. *Let K be generated by a root of a polynomial g of degree m and let n*

be a non-negative integer. If $K \subset \mathbb{Q}(\zeta_n)$ *then* SUBFIELD-OF-CYCLOTOMIC(g,n) *returns* SUBFIELD. *Assuming ERH, if* $K \not\subset \mathbb{Q}(\zeta_n)$, *then* SUBFIELD-OF-CYCLOTOMIC(g,n) *returns* NOT-SUBFIELD *with probability at least* $1/2$, *and its running time is* $\tilde{\mathcal{O}}(m\ell_m^5 + m^2\mu_m\ell_m + \nu^5 + \mu\nu^2)$, *where* $\ell_m := \max\{m, \ln(\lambda)\}$, $\mu_m := \max\{\ell_m, \lambda\}$, $\mu := \max\{m, \ln(\lambda), \lambda\}$, $\nu := \max\{\ln(\lambda), \ln(n)\}$, *and* $\lambda = \lambda(g)$.

PROOF. Assume $K \subset \mathbb{Q}(\zeta_n)$. Then K is normal as a subfield of an Abelian field. Thus NORMALITY(g) returns NORMAL by Theorem 11. Furthermore if $p \equiv 1 \bmod n$, then p has a prime divisor of degree 1 over K and so SUBFIELD-OF-CYCLOTOMIC(g,n) always returns SUBFIELD.

Let $K \not\subset \mathbb{Q}(\zeta_n)$. Then

$$\begin{aligned}
&\mathbf{Pr}[\text{SUBFIELD-OF-CYCLOTOMIC}(g,n) \text{ returns SUBFIELD} \mid K \not\subset \mathbb{Q}(\zeta_n)] \\
&\quad \leq \mathbf{Pr}[\text{SUBFIELD-OF-CYCLOTOMIC}(g,n) \text{ returns SUBFIELD} \mid K \text{ not normal}] + \\
&\qquad \mathbf{Pr}[\text{SUBFIELD-OF-CYCLOTOMIC}(g,n) \text{ returns SUBFIELD} \mid K \text{ normal and } K \not\subset \mathbb{Q}(\zeta_n)] \\
&\quad \leq 1/e^2 + 1/e^2 \leq 1/4,
\end{aligned}$$

where the first estimate follows from Theorem 11. To see the validity of the second estimate it remains to verify that

$$\begin{aligned}
&\mathbf{Pr}[\text{SUBFIELD-OF-CYCLOTOMIC}(g,n) \text{ returns SUBFIELD} \mid K \text{ normal and } K \not\subset \mathbb{Q}(\zeta_n)] \\
&\qquad\qquad \leq 1/e^2.
\end{aligned}$$

So assume for the rest of the proof that K is normal and $K \not\subset \mathbb{Q}(\zeta_n)$. We start by bounding the discriminant d of the (normal) field $K\mathbb{Q}(\zeta_n)$:

$$|d| \leq \mathrm{disc}(g)^n (n^n)^m.$$

Hence, for k computed in line (2) of the algorithm we have that $k \geq \alpha_2 \ln(|d|)^{\alpha_1}$. Let G be the Galois group of $L := K\mathbb{Q}(\zeta_n)$, U be the normal subgroup of G fixing K, and H be the normal subgroup of G fixing $\mathbb{Q}(\zeta_n)$. Let $W := H \setminus U$. Any prime p not dividing $\mathrm{disc}(g)$ such that $\mathrm{Frob}_L(p) \subseteq W$ is a witness for the fact that K does not lie in $\mathbb{Q}(\zeta_n)$. Further, $\#W \geq \#H/2$, and $\#H = \#G/\varphi(n)$. As in the proof of Corollary 5, the Witness Lemma 3 implies that for k computed in the line (1) of SUBFIELD-OF-CYCLOTOMIC we have

$$\tilde{\pi}_W(k) \geq \frac{\alpha_3}{2} \frac{1}{\varphi(n)} \frac{k}{\ln(k)},$$

where $\tilde{\pi}_W(k)$ is the number of primes p contributing to $\pi_W(k)$ which do not divide $n\mathrm{disc}(g)$. Define $\Omega := \{p \leq k \mid p \equiv 1 \bmod n\}$. Note that any p with $\mathrm{Frob}_L(p) \subseteq H$ belongs to Ω. To find witnesses we can thus sample directly from Ω. We have

$$\frac{\tilde{\pi}_W(k)}{\#\Omega} \geq \frac{\alpha_3}{2} \frac{n}{\varphi(n)} \frac{1}{\ln(k)} \geq \frac{\alpha_3}{2\ln(k)}.$$

The additional factor of 2 for the number of rounds in line (3) decreases the bound on the error probability from $1/e$ to $1/e^2$, which was to be proved.

ABELIAN(g)

Input: a monic irreducible integer polynomial g of degree m
Output: ABELIAN, NOT-ABELIAN
(1) Compute $d = \mathrm{disc}(g)$.
(2) **if** SUBFIELD-OF-CYCLOTOMIC($g, |d|$)=SUBFIELD **then return** ABELIAN and **stop** .
(3) **return** NOT-ABELIAN.

Figure 5: Algorithm ABELIAN.

Let us now analyze the running time of the algorithm. Two runs of NORMALITY take $\tilde{\mathcal{O}}(m\ell_m^5 + m^2\mu_m\ell_m)$ operations by Theorem 11. Since $\ln(k) = \tilde{\mathcal{O}}(\nu)$, computation of k in line (2) uses $\tilde{\mathcal{O}}(\nu)$ operations and the primality test in line (6) uses $\tilde{\mathcal{O}}(\nu^4)$ operations. Computing $\overline{g}$ takes $\tilde{\mathcal{O}}(\mu\nu)$ operations. The gcd-computations in lines (8) and (9) use $\tilde{\mathcal{O}}(m\nu)$ operations. The loop in line (3) is performed $\tilde{\mathcal{O}}(\nu)$ times, which gives the running time of $\tilde{\mathcal{O}}(\nu^5 + \mu\nu^2)$ for lines (3) to (13). Putting the pieces together proves the assertion on the running time. □

Remark 17. If g is known to be normal, then the running time of our algorithm SUBFIELD-OF-CYCLOTOMIC can be reduced to $\tilde{\mathcal{O}}(\nu^5 + \mu\nu^2)$. Using a randomized primality test in Line (6) reduces the running time to $\tilde{\mathcal{O}}(\nu^4 + \mu\nu^2)$, on the cost of obtaining a $\mathcal{BPP}$-algorithm.

8 Is K Abelian?

We call a polynomial $g \in \mathbb{Z}[X]$ *Abelian* if its splitting field is an Abelian extension of $\mathbb{Q}$. In this section we will describe a Monte Carlo algorithm with one-sided error to decide whether a monic irreducible polynomial $g \in \mathbb{Z}[X]$ is Abelian. How can we solve this problem with our methods? Let K be an Abelian extension of $\mathbb{Q}$. By the celebrated Theorem of Kronecker and Weber [26, Chapter 14] there exists an integer f such that K is contained in a cyclotomic field $\mathbb{Q}(\zeta_f)$. The smallest such f is called the *conductor* of K. Because every subfield of an Abelian field is Abelian, we see that a field K is Abelian if and only if K is contained in some cyclotomic field $\mathbb{Q}(\zeta_f)$. For a given f we can test this relation with the algorithm SUBFIELD-OF-CYCLOTOMIC(g, f) described in Section 7. Computing f seems the be a hard problem: we will show in Section 10 that the computation of the conductor of an Abelian field is a currently intractable problem. We will prove this by relating this problem to other seemingly hard number-theoretic problems.

How do we get around the conductor? Assume that K is Abelian. By the conductor-discriminant formula quoted before Lemma 27 we know that the conductor of K divides the discriminant of K. So we have proved:

Lemma 18. *A number field K is Abelian if and only if $K \subset \mathbb{Q}(\zeta_{|d|})$, where d is the discriminant of K.*

The desired algorithm now follows, see Figure 5.

Theorem 19. *Let K be generated by a root of a monic irreducible integer polynomial g of degree m. If K is Abelian* ABELIAN(g) *returns* ABELIAN. *Assuming ERH, if K is not Abelian, then* ABELIAN(g) *returns* NOT-ABELIAN *with probability at least $1/2$, and the running time of this algorithm is $\tilde{\mathcal{O}}(m\ell^5 + m^2\mu_1\ell + m^5\lambda^5 + \mu_2 m^4\lambda^4)$, where $\ell := \max\{m, \ln(\lambda)\}$, $\mu_1 := \max\{\lambda, \ell\}$, $\mu_2 := \max\{m, \ln(\lambda), \lambda\}$, and $\lambda = \lambda(g)$.*

PROOF. It only remains to prove that d can be computed in time $\tilde{\mathcal{O}}(m^4\lambda)$. But this follows from the trivial algorithm for this task which computes the determinant of the Sylvester matrix. □

Remark 20. (1) Assuming that $\mathrm{disc}(g) \leq c^{\tilde{\mathcal{O}}(m)}$ for some constant c, we obtain a running time of $\tilde{\mathcal{O}}(m^6)$, which is substantially faster than $\tilde{\mathcal{O}}(m^{18})$ given by factorization routines. Such an inequality for the discriminant is valid if, e.g., λ is polylogarithmic in m.

(2) The running time of the algorithm can be reduced to $\tilde{\mathcal{O}}(m\ell^4 + m^2\mu_1\ell + m^4\lambda^4 + \mu_2 m^4\lambda^4)$ by using randomized primality tests.

(3) If we know that g is normal, we can drop the assumption that it is irreducible! Namely, in that case the splitting field of g is contained in $\mathbb{Q}(\zeta_{|d|})$, d being the discriminant of g, if and only if this is true for all irreducible factors of g.

9 Cyclicity Test

We call a polynomial $g \in \mathbb{Z}[X]$ *cyclic* if its splitting field is a cyclic extension of $\mathbb{Q}$. In this section we describe a Monte Carlo algorithm with *two-sided error* to decide whether a monic irreducible polynomial $g \in \mathbb{Z}[X]$ is cyclic. The idea of the algorithm uses the following: if g is normal, then the Galois group of G is cyclic if and only if it contains an element of order $\#G$. Hence, if we know that g is normal, any p such that $h = g \bmod p$ is irreducible is a witness for g being cyclic by Fact 2 Proposition 1. (Notice that if h is irreducible then in particular $\gcd(h, h') = 1$ and thus $p \not| \mathrm{disc}(g)$ so that we can indeed apply Fact 2.)

Proposition 21. *Let K be an extension of $\mathbb{Q}$ given by a monic irreducible integer polynomial g of degree n with Galois group G.*

(1) *If K is cyclic and $W := \{\sigma \in G \mid \sigma \text{ has order } n\}$, then*

$$\frac{\#W}{\#G} = \frac{\varphi(n)}{n} \geq \frac{1}{5\ln(\ln(n))}.$$

(2) *Assume K is normal. Then K is cyclic if and only $W \neq \emptyset$. Hence, every prime p unramified in K such that $\mathrm{Frob}_K(p) \subseteq W$ is a witness for K being cyclic.*

PROOF. Easy (for a lower bound on $\varphi(n)$ see [1, Th. 8.8.7]). □

Algorithm CYCLIC(g) is given in Figure 6. We first test whether K is normal. If we can assume that K is normal with a sufficiently high probability, then we try to sample a prime

CYCLIC(g)

Input: a monic irreducible integer polynomial g of degree n
Output: CYCLIC, NOT-CYCLIC
(1) **if** 2 runs of NORMALITY(g) return NORMAL **then**
(2) Compute $k := \alpha_2\big(n!(2n)(\ln(n)+\lambda(g))\big)^{\alpha_1}$.
(3) **for** j **from** 1 **to** $2\cdot 5\ln(\ln(n))\ln(k)/\alpha_3$ **do**
(4) Choose at random an integer p from $\{1,\ldots,k\}$.
(5) **if** p is prime **then**
(6) Compute $h := g \bmod p$.
(7) **if** h is irreducible **then**
(8) **return** CYCLIC and **stop** .
(9) **return** NOT-CYCLIC.

Figure 6: Algorithm CYCLIC.

p such that g mod p is irreducible. This implies that any $\sigma \in \mathrm{Frob}_K(p)$ has maximal order n witnessing that G is indeed cyclic. The probability for finding such a prime can be bounded below by the preceding proposition. If we can not find such a prime the algorithm outputs NOT-CYCLIC.

Theorem 22. *Let K be generated by a root of the polynomial g of degree n. Assuming ERH, the probability that* CYCLIC(g) *returns the correct solution to the problem "Is K cyclic?" is at least $3/4$ and the running time of this algorithm is $\tilde{\mathcal{O}}(n\ell^5 + n^2\mu\ell)$ where $\ell := \max\{n, \ln(\lambda)\}$, $\mu := \max\{\ell, \lambda\}$, and $\lambda := \lambda(g)$.*

PROOF. We have

$$\begin{aligned}
&\mathbf{Pr}[\text{CYCLIC}(g) \text{ returns CYCLIC} \mid K \text{ not cyclic}] \\
&\quad \le \mathbf{Pr}[\text{CYCLIC}(g) \text{ returns CYCLIC} \mid K \text{ not normal}] + \\
&\qquad \mathbf{Pr}[\text{CYCLIC}(g) \text{ returns CYCLIC} \mid K \text{ normal and } K \text{ not cyclic }] \\
&\quad \le 1/e^2 + 0 \le 1/4,
\end{aligned}$$

by Theorem 11 and Proposition 21.

Assume K given by g is cyclic. Then NORMALITY(g) always returns NORMAL by Theorem 11. The probability that we do not pick a witness p for K being cyclic in one run of the loop is at most $1 - 1/t$, where $t = 5\ln(\ln(n))\ln(k)/\alpha_3$ by Lemma 21 and Corollary 4. The additional factor of 2 in the loop of line (3) implies that

$$\mathbf{Pr}[\text{CYCLIC}(g) \text{ returns NOT-CYCLIC} \mid K \text{ cyclic}] \le 1/e^2 \le 1/4.$$

The assertion on the running time of this algorithm is proved as in Theorem 11. □

Remark 23. (1) A more practical cyclicity test for the polynomial g would be to run the normality test and set a flag as soon as a prime p is found modulo which g is irreducible. After the normality test is finished, the algorithm returns CYCLIC if the normality test returns NORMAL and the flag is set. In all other cases the algorithm returns NON-CYCLIC.

(2) Using randomized primality tests the running time can be reduced to $\tilde{\mathcal{O}}(n\ell^4 + n^2\mu\ell)$.

10 Computing the Conductor of an Abelian Field

Let K be an Abelian field given by a monic irreducible polynomial $g \in \mathbb{Z}[x]$ of degree $n > 1$. How difficult is it to compute the conductor f of K given the polynomial g? In this section we will prove that this problem is polynomial time equivalent to that of computing the largest squarefree divisor of an integer, for which there is no efficient algorithm known at the present time. (For some practical advice on how to compute such a divisor, the reader is referred to Section 7 of [4].)

the model of computation that we will be using is that of a Turing machine. We say that a language $\mathcal{L}$ is polynomial time reducible to a language $\mathcal{L}'$, and denote it by $\mathcal{L} \leq \mathcal{L}'$, if there is a polynomial time reduction from $\mathcal{L}$ to $\mathcal{L}'$ in the usual sense.

10.1 Known Results

In this subsection we will gather some well-known results about the complexity of various computational problems in algebraic number theory.

Fact 24. *The following problems are equivalent under polynomial time reduction:*

(ROI) *Computing the ring of integers of a number field;*

(DISC) *Computing the discriminant of a number field;*

(OPF) *Computing the product of primes dividing an integer to an odd power.*

The equivalence of ROI and OPF is due to Chistov [7]. (The reduction used in this equivalence has been improved by Buchmann and Lenstra [4].) Considering quadratic fields we see that OPF $\leq$ DISC. Furthermore, the definition of the discriminant shows that given the ring of integers, we can easily compute the discriminant in polynomial time which shows that DISC $\leq$ ROI.

Another result we will need later for our investigations is the following one, due to Buchmann and Lenstra [4]. (See also [3, pp. 47–48].)

Fact 25. *Given a number field K and a prime p, one can compute in polynomial time the ramification index of p in K.*

10.2 Largest squarefree divisor and similar problems

For an integer m we denote by lsqf(m) the largest square factor of m, i.e., the integer $\prod_{p,\mathrm{ord}_p(m)>1} p^{\mathrm{ord}_p(m)}$. The largest square-free divisor lsfd(m) is defined as $\prod_{p,\mathrm{ord}_p(m)>0} p$. The *equal order factorization* eof(m) of m is defined as a vector $(s_1, \ldots, s_t)$, where $s_i = \prod_{p,\mathrm{ord}_p(m)=i} p$. Hence, the s_i are squarefree and pairwise co-prime integers and the prime divisors of s_i are exactly those prime divisors of m appearing in m to power i. How hard is it to compute any of these data from m? We will prove in this section that these problems are equivalent under polynomial time reductions. We remark that Landau [14] has also considered the problem EOF below. Her interest has been in showing that this problem can be reduced to that of computing the Euler φ-function of an integer. Though it is likely that the reductions we present here have been already obtained by other people, we choose to give them since they have not appeared in print to the best of our knowledge.

Input: $m \in \mathbb{N}$.
Output: $\mathrm{lsfd}(m)$.
(1) $Q := 1$, $S := m$.
(2) **while** $S \neq 1$ **do**
(3) $t := \mathrm{lsqf}(S)$;
(4) $u := S/t$;
(5) $Q := \mathrm{lcm}(Q, m)$;
(6) $S := \sqrt{t}$;

Figure 7: Reduction of LSFD to LSQF.

Theorem 26. *The following problems are equivalent under polynomial time reductions:*

(LSQF) *Computing the largest square factor of an integer;*

(LSFD) *Computing the largest square-free divisor of an integer;*

(EOF) *Computing the equal order factorization of an integer.*

Proof. Note that $\mathrm{lsfd}(m) = s_1 \cdots s_t$, and that $\mathrm{lsqf}(m) = \prod_{i \text{ even}} s_i^i \cdot \prod_{i \text{ odd}} s_i^{i-1} s_i$. Hence, LSQF $\leq$ EOF and LSFD $\leq$ EOF under polynomial time reduction. Algorithm 7 shows that LSFD $\leq$ LSQF under polynomial time reduction: the number of times the **while**-loop is performed is $\leq \log(n)$, and all the steps inside the loop (except possibly for computing the largest square divisor) can be performed in time polynomial in $\log(n)$. (Note that since t is a square in line (6), its square root can be efficiently computed by any variant of Newton's method.) Algorithm 8 shows that EOF $\leq$ LSFD under polynomial time reduction: once we have found $(\tau_1, \ldots, \tau_r)$ we can compute $\mathrm{eof}(m) = (s_1, \ldots, s_r)$ in polynomial time, since $s_i = \tau_i/\tau_{i+1}$ for $i = 1, \ldots, r-1$, and $s_r = \tau_r$. □

Noting that the problem OPF from the last subsection and LSQF are polynomial time equivalent, the last theorem shows that all of the problems stated there are equivalent to those in the last subsection.

10.3 Computing the conductor

Let X be the character group of K. Each $\chi \in \mathrm{X}$ is a Dirichlet character, and has a period f_χ which we call the conductor of χ. Each $\chi \in \mathrm{X}$ has a unique decomposition $\chi = \prod_p \chi_p$, where the conductor of χ_p is a power of the prime p, and the product formally extends over all prime numbers. (Of course, only finitely many χ_p are nontrivial.) In this case the conductor of χ is just the product of the conductors of the χ_p. Furthermore, the conductor $f(K)$ of K equals the least common multiple of the conductors f_χ, $\chi \in \mathrm{X}$. The conductor-discriminant formula (see [26, Th. 3.11]) relates the discriminant $d(K)$ of K and the conductors of characters in X in the following way:

$$|d(K)| = \prod_{\chi \in \mathrm{X}} f_\chi.$$

In the sequel we will denote for a nonzero integer m and a prime p by $\mathrm{ord}_p(m)$ the largest exponent to which p divides m.

Input: $m \in \mathbb{N}$.
Output: vector $(\tau_1, \ldots, \tau_{\log(m)})$ such that $\tau_i = \prod_{p, \mathrm{ord}_p(m) \geq i} p^{\mathrm{ord}_p(m)}$.
(1) $S := m$;
(2) **for** j **from** 1 **to** $\log(m)$ **do**
(3) $\quad \tau_j := \mathrm{lsfd}(S)$;
(4) $\quad S := S/\tau_j$;

Figure 8: Reduction of EOF to LSFD.

Lemma 27. (1) *If $p \neq 2$ is ramified in K, then the p-part of the conductor $f(K)$ of K equals pd_p, where d_p is the p-part of the ramification index of p in K.*

(2) *If 2 is ramified in K, then the 2-part of $f(K)$ equals 2^{e_2+1} or $2^{\ell+2}$, where $e_2 := 2^\ell$ is the ramification index of 2 in K. The first case occurs if and only if $e_2 \mathrm{ord}_2(d(K))/n = 2^\ell \ell$.*

Proof. Let X_p denote the group $\{\chi_p \mid \chi \in \mathrm{X}\}$. By [26, Th. 3.5] we know that $|\mathrm{X}_p|$ equals the ramification index e_p of p in K. Hence X_p is the character group of a subfield of $\mathbb{Q}(\zeta_{p^\infty})$ of order e_p.

(1) If $p \neq 2$, this subfield is uniquely determined, since the group $\mathbb{Z}_{p^m}^\times$ is cyclic for any m. This subfield has conductor $p^{\mathrm{ord}_p(e_p)+1}$. Hence, $\mathrm{lcm}(f_\chi \mid \chi \in \mathrm{X}_p) = p^{e_p+1}$, and the left hand side of this equation is the p-part of the conductor of K.

(2) X is the direct product of the X_p, hence

$$|\mathrm{X}/\mathrm{X}_p| = \frac{n}{e_p}. \tag{1}$$

Let $\mathrm{X}' := \prod_{p \neq 2} \mathrm{X}_p$. Then $\mathrm{X} = \mathrm{X}_2\mathrm{X}'$, and by the conductor-discriminant formula we have

$$\begin{aligned} \mathrm{ord}_2(d(K)) &= \sum_{\chi \in \mathrm{X}} \mathrm{ord}_2(f_\chi) = \sum_{\chi \in \mathrm{X}_2} \sum_{\psi \in \mathrm{X}'} \mathrm{ord}_2(f_{\chi\psi}) \\ &= \sum_{\chi \in \mathrm{X}_2} \sum_{\psi \in \mathrm{X}'} \mathrm{ord}_2(f_\chi) = \frac{n}{e_2} \sum_{\chi \in \mathrm{X}_2} \mathrm{ord}_2(f_\chi) \\ &= \frac{n}{e_2} \mathrm{ord}_2(d(M)), \end{aligned}$$

where M is a subfield of order $e_2 = 2^\ell$ of $\mathbb{Q}(\zeta_{2^\infty})$ having X_2 as character group. There are three possibilities for M: either $M = \mathbb{Q}(\zeta_{2^{\ell+1}})$, or $M = \mathbb{Q}(\zeta + \zeta^{-1})$, or $M = \mathbb{Q}(\zeta - \zeta^{-1})$, where ζ is a primitive $2^{\ell+2}$th root of unity. In any case, the conductor of M equals the least common multiple of f_χ, $\chi \in \mathrm{X}_2$, which is equal to the 2-part of the conductor of K. In the first case M has conductor $2^{\ell+1}$ and $\prod_2(d(M)) = \ell 2^\ell$, see [26, Prop. 2.1]. In the second and third case M has conductor $2^{\ell+2}$ and the following argument shows that $\mathrm{ord}_2(d(M)) = \sum_{j=0}^{\ell-1} 2^j(j+3)$: in this case X_2 is cyclic, and generated by a character χ, say, of conductor $2^{\ell+2}$. Since $f_{\chi^s} = f_\chi / \gcd(s, 2^{\ell+2})$, we see that X_2 contains $2^{\ell-i}$ characters of conductor $2^{\ell-i+3}$, $i = 1, \ldots, \ell$, and one character of conductor 1. The conductor-discriminant formula shows now our claim for $\mathrm{ord}_2(d(M))$.

Using induction on ℓ we see that $\sum_{j=0}^{\ell-1} 2^j(j+3) > \ell 2^\ell$ for $\ell \geq 1$. Hence, $M = \mathbb{Q}(\zeta_{2^{\ell+1}})$ if and only if $\prod_2(d(K)) = n\ell 2^\ell/e_2$. □

Theorem 28. *The following problems are polynomial time equivalent:*

(COND) *Computing the conductor of an Abelian field;*

(ROI) *Computing the ring of integers of a number field;*

(DISC) *Computing the discriminant of a number field;*

(LSQF) *Computing the largest square factor of an integer;*

(LSFD) *Computing the largest square-free divisor of an integer;*

(EOF) *Computing the equal order factorization of an integer.*

PROOF. We first show that LSQF $\leq$ COND: Given m, we compute the conductor $f(K)$ of $K = \mathbb{Q}(\sqrt{m})$, which, in this case, equals the discriminant of K. If $m' := m/\mathrm{lsqf}(m)$ is congruent to one modulo 4, then $d(K) = m'$. Otherwise, $d(K) = 4m'$. So, we can compute m', and hence $\mathrm{lsqf}(m)$, from $f(K)$ in polynomial time.

Now we show that COND $\leq$ LSFD. By virtue of Fact 24 and Theorem 26 the result follows.

By Facts 24 and 25 we can compute in polynomial time the following invariants of K: the discriminant $d(K)$, for each prime p less than n the ramification index e_p of p in K, and $d'(K)$ which is the product of all prime divisors p of $d(K)$ which are larger than n. Note that for any p the ramification index e_p of p divides n. As a result, all primes larger than n are at most tamely ramified in K, and Lemma 27(1) implies that the p-part of $f(K)$ equals p for any such ramified p. Hence, $f(K) = d'(K)d''(K)$, where $d''(K)$ has only prime divisors $\leq n$ ramified in K. We may thus concentrate on computing $d''(K)$.

Suppose that $p \neq 2$, $p \leq n$, and $e_p \geq 1$. By Lemma 27 we know that $\mathrm{ord}_p(f(K)) = \prod_p(d''(K)) = p^{\mathrm{ord}_p(e_p)+1}$, and this number can be computed in polynomial time. If $p = 2$ and $e_2 := 2^\ell > 1$, we compute $e_2 \mathrm{ord}_2(d(K))/n$ and test whether it equals $\ell 2^\ell$. In this case we know that $\prod_2(d''(K)) = \ell + 1$. Otherwise $\mathrm{ord}_2(d''(K)) = \ell + 2$. □

11 Conclusions and Open Problems

We have developed randomized algorithms for deciding several properties of number fields given by monic irreducible integer polynomials. Our algorithms are orders of magnitude faster than polynomial time deterministic solutions based on factorization of polynomials. We have also found a problem for which the factorization method does not yield a polynomial time algorithm, while our approach gives a fast Monte Carlo algorithm with one-sided error. The main idea of our algorithms is the study of the decomposition behavior of primes in the corresponding number fields. Their basic building blocks are thus primality testing and gcd-computations in polynomial rings over finite fields.

The following questions remain open:

(1) *Does there exist a deterministic polynomial time algorithm for deciding the problem* $K \subseteq \mathbb{Q}(\zeta_m)$*?* Specializing K to a quadratic field given by the polynomial $x^2 - n$, such an algorithm would also decide whether $n/\mathrm{lsqf}(n)$ divides m.

(2) *Is there a fast Monte Carlo algorithm for deciding* $K \subset L$ *for* general K *and* L*?* Our methods only work if K is normal, since for non-normal K Bauer's theorem does not hold.

(3) Algorithm ABELIAN introduced in Section 8 decides whether a normal polynomial g is Abelian. *Can we detect other types of Galois extensions?* For instance, given a normal polynomial g of degree $n!$, can we decide in random polynomial time whether the Galois group of G is the symmetric group? Note that the question whether a degree n polynomial has Galois group S_n or A_n can be decided in polynomial time, see [18, Th. 3.6].

(4) In the algorithm ABELIAN we can drop the assumption that g is irreducible if we know that it is normal. *Is there a way to drop the irreducibility assumption in the other algorithms, such as the normality test, as well?*

(5) Our methods do indeed yield statistical results about orders of different elements in the group. In some cases (as in the cyclicity test) this might be enough. *Is there a way to quickly find a "good guess" for the structure of the Galois group of a normal irreducible polynomial* g*, if we know that the group is Abelian?*

(6) A well-known theorem (which is a consequence of the Chebotarev Density Theorem) states that Galois number fields are uniquely determined by the set of split rational primes (see [20, Cor. 6.9]). *What kind of results can be obtained from this set of primes in random polynomial time?*

12 Acknowledgments

The research of the second author was supported by a Habilitationsstipendium of the Deutsche Forschungsgemeinschaft, Grant Sh–57/1. We would like to thank H. W. Lenstra for stimulating and clarifying discussions as well as for communicating to us Proposition 8 and the example given in Remark 9. Many thanks go to J. Buhler for several useful hints, and to B. Sturmfels for asking the question which got the paper started.

References

[1] E. Bach and J. Shallit. *Algorithmic Number Theory.* MIT Press, Cambridge MA, 1996.

[2] E. Bach and J. Shallit. Explicit bounds for primes in residue classes. *Math. Comp.*, 65:1717–1735, 1996.

[3] J. Buchmann. Complexity of algorithms in algebraic number theory. In R. A. Mollin, editor, *Proceedings of the First Conference of the Canadian Number Theory Association held in Banff, Alberta*, pages 37–53. de Gruyter, Berlin, 1990.

[4] J. Buchmann and H. W. Lenstra. Approximating rings of integers in number fields. *J. de Théorie des Nombres de Bordeaux*, 6:221–260, 1994.

[5] P. Bürgisser, M. Clausen, and M. A. Shokrollahi. *Algebraic Complexity Theory.* Number 315 in Grundlehren der math. Wissenschaften. Springer Verlag, Heidelberg, 1996.

[6] N. G. Chebotarev. Die Bestimmung der Dichtigkeit einer Menge von Primzahlen, welche zu einer gegebenen Substitutionsklasse gehören. *Math. Ann.*, 95:191–228, 1926.

[7] A. L. Chistov. The complexity of constructing the ring of integers of a global field. *Dokl. Akad. Nauk. SSSR*, 306:1063–1067, 1989. English Translation: Soviet Math. Dokl., 39:597–600, 1989.

[8] A.L. Chistov and D. Yu. Grigoreiv. Polynomial-time factoring of multivariable polynomials over a general field. Technical report, USSR Academy of Sciences, Steklov Mathematical Institute Leningrad, 1982.

[9] H. Cohen. *A Course in Computational Algebraic Number Theory*. Springer Verlag, Berlin, 1993.

[10] J. von zur Gathen and V. Shoup. Computing Frobenius maps and factoring polynomials. *Comput. Complexity*, 2:187–224, 1992.

[11] G. J. Janusz. *Algebraic Number Fields*. American Mathematical Society, second edition, 1996.

[12] J. C. Lagarias and A. M. Odlyzko. Effective versions of the Chebotarev density theorem. In A. Fröhlich, editor, *Algebraic Number Fields*, pages 409–464. Academic Press, London, 1977.

[13] S. Landau. Factoring polynomials over algebraic number fields. *SIAM J. Comput.*, 14:184–195, 1985.

[14] S. Landau. Some remarks on computing the square parts of integers. *Inform. and Comput.*, 78:246–253, 1988.

[15] A. K. Lenstra. Factoring polynomials over algebraic number fields. Technical Report IW213/82, Dept. Comp. Science, Stichting Mathematisch Centrum, Amsterdam, 1982.

[16] A. K. Lenstra. Factoring multivariate polynomials over algebraic number fields. *SIAM J. Comput.*, 16(3):591–598, 1987.

[17] A. K. Lenstra, H. W. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261:515–534, 1982.

[18] H. W. Lenstra. Algorithms in algebraic number theory. *Bulletin of the AMS*, 26:211–244, 1992.

[19] G. Miller. Riemann's hypothesis and tests for primality. *J. Comput. System Sci.*, 13:300–317, 1976.

[20] J. Neukirch. *Class Field Theory*. Springer Verlag, Heidelberg, Berlin, 1986.

[21] A. M. Odlyzko. On conductors and discriminants. In A. Fröhlich, editor, *Algebraic Number Fields*, pages 377–407. Academic Press, London, 1977.

[22] J. Oesterlé. Versions effectives du théorème de Chebotarev sous l'hypothèse de Riemann généralisée. *Astérisque*, 61:165–167, 1979.

[23] A. Schönhage. Factorization of univariate integer polynomials by diophantine approximation and an improved basis reduction algorithm. In *Proc. 11th ICALP*, volume LNCS 172, pages 437–447, 1984.

[24] I. R. Shafarevich. Construction of fields of algebraic numbers with given solvable galois group. *Izv. Akad. Nauk SSSR. Ser. Mat.*, 18:525–578, 1954.

[25] H. Tôyama. A note on the different of the composed field. *Kodai Math. Sem. Report*, 7:43–44, 1955.

[26] L. C. Washington. *Introduction to Cyclotomic Fields*. Springer Verlag, New York, second edition, 1997.