March 2023

Recovering generators of principal ideals using subfield structure and applications to cryptography

William Youmans
University of South Florida

Recovering Generators of Principal Ideals Using Subfield Structure and Applications to Cryptography

by

William Youmans

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Mathematics & Statistics
College of Arts and Sciences
University of South Florida

Major Professor: Jean-François Biasse, Ph.D.
Dmytro Savchuk, Ph.D.
Nataša Jonoska, Ph.D.
Giacomo Micheli, Ph.D.
Attila Yavuz, Ph.D.

Date of Approval:
March 29, 2023

Keywords: Ideal lattice, cyclotomic field, approximate short vector problem, ideal class group, S-unit group

ACKNOWLEDGMENTS

I am grateful to my advisor, Dr. Jean-François Biasse, for his invaluable guidance and unwavering support throughout the entire process of researching and writing this dissertation. His insightful feedback, expertise, and patience have been instrumental in shaping my research and helping me overcome various challenges along the way.

I would also like to thank Dr. Claus Fieker for his invitation to Kaiserslautern, Germany to work on material included in this dissertation, as well as Dr. Tommy Hofmann, who was an incredibly gracious host during my stay and patiently answered my endless questions. I am also indebted to Dr. Aurel Page for helpful discussions and invaluable feedback. All three were involved in the development of this work and their support has been crucial.

I am also indebted to my colleagues and friends who have provided me with encouragement and a sense of community throughout my graduate studies. Finally, I thank my family for their love, support, and understanding during this journey.

TABLE OF CONTENTS

LIST OF TABLES

ABSTRACT

The principal ideal problem (PIP) is the problem of determining if a given ideal of a number field is principal, and if so, of finding a generator. Algorithms for resolving the PIP can be efficiently adapted to solve many hard problems in algebraic number theory, such as the computation of the class group, unit group, or S-unit group of a number field. The PIP is also connected to the search for approximate short vectors, known as the γ-Shortest Vector Problem (γ-SVP), in certain structured lattices called ideal lattices, which are prevalent in cryptography. We present an algorithm for resolving the PIP that leverages the norm relation techniques of Biasse, Fieker, Hofmann, and Page [20] to efficiently reduce the PIP in certain number fields to instances of the PIP in subfields. Our algorithm is focused on practical performance and we demonstrate its viability by resolving instances of the PIP in cyclotomic fields of degree up to 1800. We further adapt this technique to the problem of finding mildly short vectors, solutions to γ-SVP for $\gamma = 2^{\tilde{O}(\sqrt{n})}$, in an ideal lattice of a cyclotomic field. Cramer, Ducas, and Wesolowski [28, 29] show that the search for mildly short vectors in such a lattice reduces efficiently to the PIP on a quantum computer. We describe a classical variant of this reduction that applies to non-cyclic cyclotomic fields, demonstrating that our technique implies a polynomial improvement over the state of the art in almost all cyclotomic fields. We further show that there are infinite families of cyclotomic fields where this approach achieves a superpolynomial improvement over the state of the art.

CHAPTER 1
INTRODUCTION

For hundreds of years mathematicians have been interested in the algebraic properties of numbers for their own sake. That these properties would come to be a cornerstone of modern cryptography is no surprise, as their study has produced a large number of well-known, computationally hard problems. In 1987 Zassenhaus [61] summarized four central tasks of algorithmic algebraic number theory. Tasks three and four on this list are the computation of the unit group and class group of a number field respectively (tasks one and two are the computation of its Galois group and ring of integers). Both tasks three and four are closely connected to the so-called Principal Ideal Problem (PIP), the problem of determining if a given ideal is principal, and if so, recovering a generator.

1.1 The Principal Ideal Problem

Let K be a number field with ring of integers $\mathcal{O}_K$ and discriminant Δ_K. In general, resolution of the PIP for ideals of $\mathcal{O}_K$ relies on the computation of the ideal class group of $\mathcal{O}_K$, denoted $\mathrm{Cl}(\mathcal{O}_K)$. A classical result of algebraic number theory is that $\mathrm{Cl}(\mathcal{O}_K)$ is a finite abelian group, but determining its order and structure (decomposition into a direct product of cyclic subgroups) appears to be a hard problem.

The study of class groups goes back to Gauss [37] who studied the group of equivalence classes of certain binary quadratic forms, which we now know as another interpretation for the class group of a quadratic field. In 1801 Gauss gave an algorithm for computing the class number $h_K = |\mathrm{Cl}(\mathcal{O}_K)|$ in time $O(\Delta_K^{1/2})$ for imaginary quadratic fields [37, Article 305]. Approximately 170 years later, Shanks [70, 71] improved this to $O(\Delta_K^{1/4+\epsilon})$, or $O(\Delta_K^{1/5+\epsilon})$ under the assumption of the Extended Riemann Hypothesis (ERH) (for a complexity analysis see [46, 69] or [40, Theorem 10.6]). However, these results are limited to quadratic fields and take time exponential in the bit-length of the discriminant. Furthermore, it requires more work to then determine the structure of the class group from it's order.

In 1989 Hafner and McCurly [38] described an algorithm for computing the class group of an imaginary quadratic field in time subexponential in the bit-length of the discriminant. This was subsequently generalized by Buchmann [22] to the case of an arbitrary number field with fixed degree with practical improvements by Cohen, Diaz y Diaz, and Olivier [27]. In 2014 Biasse and Fieker [14] gave a heuristic subexponential algo-

rithm for computing the class group in all classes of number fields, with an improved asymptotic complexity when specialized to the case of cyclotomic fields [18]. Despite these improvements class group computation remains a hard problem and advancements show no signs of breaking the subexponential barrier in the general case.

In the particular case of multiquadratic fields, fields of the form $\mathbb{Q}(\sqrt{d_1}, \ldots, \sqrt{d_n})$ for $d_i \in \mathbb{Z}$, it was demonstrated by Bauch, Bernstein, de Valence, Lange, and Van Vredendaal that it was possible to reduce an instance of the PIP to multiple instances of the PIP in subfields [5]. This was generalized to the computation of the class group of multiquadratics (by way of the S-unit group) by Biasse and Van Vredendaal [17], who showed that for certain choices of d_i the class group could be computed in polynomial time. A similar approach was used by Lesavourey, Plantard, and Susilo [47] in the case of multicubic fields, number fields generated by cube roots of integers.

In 2020, Biasse, Fieker, Hofmann, and Page [20] generalized the methods of [5, 17, 47] to arbitrary non-cyclic number fields using an object known as a *norm relation*. They showed how to compute invariants of a number field by exploiting its subfield structure, and using norm relations they computed the class group structure (as an abstract abelian group) of a cyclotomic field of degree 1728 in just 4 hours on a single core machine, a massive improvement over the state-of-the-art. In Chapter 3 we recall important details on norm relations from [20] and provide some examples which hopefully illuminate their utility. In Chapter 4 we show how norm relations can be used to solve the PIP in various ways, and describe a variant that is optimized for practical performance.

1.2 Ideal Lattices and Cryptography

The problem of integer factorization is widely believed to be hard on a classical computer. This presumed difficulty is at the core of the security of the widely used RSA public key cryptosystem [64]. However, in 1994 Shor [73] proposed the first polynomial time algorithm for factoring integers, and more generally, solving the discrete logarithm problem, with a quantum computer. This discovery renders RSA (among other cryptosystems) useless in the post-quantum setting, motivating the research of new problems and security assumptions able to resist an adversary equipped with a quantum computer.

The seminal work of Regev [63] describes a quantum reduction from certain worst-case problems in lattices to a learning problem called Learning With Errors (LWE). In other words, cryptosystems whose security is based on the hardness of LWE are then based on the *worst-case* quantum hardness of lattice problems such as the decisional version of the Short Vector Problem (SVP). Hence LWE provides a strong foundation for the security of lattice-based cryptosystems. In an effort to increase efficiency of LWE-based

cryptosystems Lyubashevsky, Peikert, and Regev defined a variant of LWE called Ring Learning With Errors (Ring-LWE) [49]. In contrast to the Euclidean lattices of LWE, Ring-LWE is instantiated over cyclotomic integer rings and relies on the hardness of the approximate search variant of SVP, γ-SVP, in ideals of a cyclotomic field, for polynomial approximation factor γ.

Furthermore, γ-SVP in an ideal lattice and the PIP are closely related. In 2016 Biasse and Song [16] gave an efficient quantum algorithm for resolving the PIP in a arbitrary number field. Cramer, Ducas, Peikert, and Regev showed that γ-SVP in a principal ideal of a prime-power conductor cyclotomic field with $\gamma \in 2^{\tilde{O}(\sqrt{n})}$ efficiently reduces to the PIP [30] on a quantum computer. Following the convention of [28] we will refer to a solution to γ-SVP with $\gamma = 2^{O(\sqrt{n})}$ as a *mildly short vector*. Combining the above results we see that there is an efficient quantum algorithm for finding mildly short vectors in principal ideal lattices of prime-power conductor cyclotomic fields. Cramer, Ducas, and Wesolowski [29] extended this work by exhibiting a quantum reduction from the search for a mildly short vector in a *general* (i.e. not necessarily principal) ideal of $\mathbb{Q}(\zeta_{2^k})$ to an instance of the PIP, and further generalized this to arbitrary cyclotomic fields in [28]. We refer to this reduction as CDW.

As we alluded to earlier, the security proof of Ring-LWE relies on the hardness of γ-SVP for a polynomial approximation factor γ. As a mildly short vector is only a solution to γ-SVP for subexponential γ these results have no impact on the security of Ring-LWE. However, in the case of an arbitrary Euclidean lattice the BKZ algorithm [65] finds a mildly short vector in time $2^{\tilde{O}(\sqrt{n})}$, and as of right now this is the best one can hope for in an arbitrary lattice. Thus the existence of an efficient quantum algorithm for finding mildly short vectors in cyclotomic ideal lattices demonstrates a gap in the hardness of lattice problems in structured lattices versus general lattice.

1.3 Outline

The goal of this work is to employ the norm relations of [20] in the resolution of the PIP and the search for mildly short vectors in ideal lattices. In Chapter 2 we introduce preliminary results. We recall some technical details regarding the saturation of multiplicative groups in Section 2.3. In Section 2.4 we give a concise description of compact representations of field elements alongside a complexity analysis, and in Section 2.5 we discuss the heuristic subexponential algorithms for resolution of the PIP and related problems. Chapter 3 summarizes the main results on norm relations of [20]. We provide a slightly simplified exposition and examples that demonstrate the utility of norm relations, as well as families of fields where norm relations can yield impressive asymptotic improvements compared to algorithms over general number fields.

In Chapter 4 we investigate the application of norm relations to resolution of the PIP. We begin with a simple method in Section 4.1 where we analyze the asymptotic complexity. We find that in certain families of fields, norm relations allow resolution of the PIP in time $2^{n^{o(1)}}$. In Section 4.2 we further develop saturation techniques to the particular case of finding roots modulo multiplicative groups. We use this in Section 4.3 to reduce the problem of ideal decomposition to subfields with the same asymptotic cost as resolution of the PIP. In Section 4.4 we give an alternative algorithm for resolution of the PIP optimized for practical performance and demonstrate its viability by solving instances of the PIP in cyclotomic fields of degree up to 1800.

We then consider the applications of these results to finding mildly short vectors in ideal lattices of a cyclotomic field in Chapter 5. In Section 5.1 we review the case of principal ideals, where the search for mildly short vectors efficiently reduces to the PIP [30, 81]. In Section 5.2 we recall the core components of the efficient quantum reduction of [28, 29] in the case of an arbitrary ideal. We give an algorithm for computing generators of the minus part of the class group in Section 5.3, used in the description of a classical analogue of the CDW reduction which we provide in Section 5.4. We show that, using the results of Chapter 4, our classical variant can find mildly short vectors in arbitrary ideals in time $2^{n^{o(1)}}$ in certain families of fields. In Section 5.6 we provide numerical evidence in support of heuristic assumptions made in [28, 29] used to justify the asymptotic cost of the CDW reduction. Finally, we implemented our classical variant of the CDW reduction and report our findings in Section 5.6.3.

Much of Chapter 4 is based on unpublished joint work with Jean-François Biasse, Claus Fieker, and Tommy Hofmann. Chapter 5 and Section 4.3 are based on the following joint work with Jean-François Biasse, Muhammed Rashad Erukulangara, Claus Fieker, and Tommy Hofmann [19]:

J.-F. Biasse et al. “Mildly Short Vectors in Ideals of Cyclotomic Fields Without Quantum Computers”. In: *Mathematical Cryptology* 2.1 (Nov. 2022), pp. 84–107. URL: https://journals.flvc.org/mathcryptology/article/view/132573

CHAPTER 2
BACKGROUND

We use Bachmann-Landau notation to compare the asymptotic behavior of functions. Given functions f and g, we say $f(x) \in O(g(x))$ (or by abuse of notation, $f(x) = O(g(x))$) when $|f(x)| \leq Cg(x)$ for some constant $C \in \mathbb{R}_{>0}$ and sufficiently large x. We say $f(x) \in o(g(x))$ when for *every* $c \in \mathbb{R}_{>0}$, $|f(x)| \leq cg(x)$ for sufficiently large x. Furthermore, $f(x) \in \tilde{O}(g(x))$ if $f(x) \in O(g(x)\log^k(g(x)))$ for some $k \in \mathbb{N}$. In other words, $\tilde{O}$ hides logarithmic factors. We say a function is polynomial in x when it is in $O(x^k)$ for some $k \in \mathbb{N}$. We will sometimes denote this $\text{Poly}(x)$.

2.1 Lattices and Hard Lattice Problems

A lattice $\mathcal{L}$ is an additive, discrete subgroup of an m-dimensional $\mathbb{R}$-vector space of the form $\mathcal{L} = \mathbb{Z}b_1 + \cdots + \mathbb{Z}b_n$ where $(b_1, \ldots, b_n)$ are linearly independent vectors. The tuple $(b_1, \ldots, b_n)$ is called a *basis* of $\mathcal{L}$ and is not unique. We say $\mathcal{L}$ has *full rank* if $m = n$. Let B be the $m \times n$ matrix with columns $b_1, \ldots, b_n$. The determinant of $\mathcal{L}$ is defined as $\det(\mathcal{L}) = \sqrt{\det(B^T B)}$ and is sometimes referred to as the volume $\text{Vol}(\mathcal{L})$.

We denote by $\lambda_1(\mathcal{L})$ the l_2-norm of the shortest non-zero vector of $\mathcal{L}$. More generally, the *successive minima* of $\mathcal{L}$ are $\lambda_1, \ldots, \lambda_n \in \mathbb{R}$ such that λ_i is minimal and that there are independent vectors $v_1, \ldots, v_n \in \mathcal{L}$ such that $\|v_j\|_2 \leq \lambda_i$ for $1 \leq j \leq i$. There are a number of hard problems surrounding the search for short vectors in a lattice. We formally state some of these problems that we will be referencing throughout this work.

Definition 2.1 (Shortest Vector Problem (SVP))**.** Given a basis $(b_1, \ldots, b_n)$ for a lattice $\mathcal{L}$, find $v \in \mathcal{L}$ such that $\|v\|_2 = \lambda_1(\mathcal{L})$.

Definition 2.2 (Shortest Independent Vector Problem (SIVP))**.** Given a basis $(b_1, \ldots, b_n)$ for a lattice $\mathcal{L}$, find n linearly independent vectors $v_1, \ldots v_n \in \mathcal{L}$ such that $\|v_i\|_2 \leq \lambda_n(\mathcal{L})$.

Definition 2.3 (γ-Shortest Vector Problem (γ-SVP))**.** Given a basis $(b_1, \ldots, b_n)$ for a lattice $\mathcal{L}$, find $v \in \mathcal{L}$ such that $\|v\|_2 \leq \gamma\lambda_1(\mathcal{L})$.

Definition 2.4 (γ-Shortest Independent Vector Problem (γ-SIVP))**.** Given a basis $(b_1, \ldots, b_n)$ for a lattice $\mathcal{L}$, find n linearly independent vectors $v_1, \ldots v_n \in \mathcal{L}$ such that $\|v_i\|_2 \leq \gamma\lambda_n(\mathcal{L})$.

The LLL algorithm solves γ-SVP for $\gamma \in 2^{O(n)}$ in time polynomial in n. Exponential algorithms such as sieve methods [2] can solve exact SVP in time $2^{O(n)}$, while the BKZ algorithm [65] allows one to solve γ-SVP for $\gamma \in 2^{O(n/k)}$ in time $2^{O(k)}$. In particular, the time to solve γ-SVP for $\gamma \in 2^{\tilde{O}(\sqrt{n})}$ is in $2^{\tilde{O}(\sqrt{n})}$. In [28], solutions of γ-SVP for $\gamma \in 2^{\tilde{O}(\sqrt{n})}$ are referred to as *mildly short vectors* and we adopt this terminology.

2.2 Number Theory

Let K be an algebraic number field of degree n, that is, a finite extension of the rational numbers $\mathbb{Q}$ with $n = [K : \mathbb{Q}]$. Let $\alpha \in K$. Its minimal polynomial is the unique monic irreducible polynomial $f \in \mathbb{Q}[x]$ with $f(\alpha) = 0$, and we call α *integral* if the minimal polynomial is in $\mathbb{Z}[x]$. The set of all integral elements of K is a subring of K, which is called the *ring of integers* of K and which is denoted by $\mathcal{O}_K$. The ring of integers $\mathcal{O}_K$ as well as all non-zero ideals of $\mathcal{O}_K$ are free $\mathbb{Z}$-modules of rank n. For a non-zero ideal $\mathfrak{a}$ of $\mathcal{O}_K$ the quotient $\mathcal{O}_K/\mathfrak{a}$ is a finite abelian group, whose order is called the *norm* of $\mathfrak{a}$ and which we denote by $\mathrm{N}(\mathfrak{a})$. By setting $\mathrm{N}(\{0\}) = 0$ for the zero ideal, the norm becomes a multiplicative map on the set of all ideals of $\mathcal{O}_K$. Any $\mathbb{Z}$-basis of $\mathcal{O}_K$ is called an *integral basis*. Given such an integral basis $\omega_1, \ldots, \omega_n$, we denote by $\Delta_K = \det((\mathrm{Tr}_{K/\mathbb{Q}}(\omega_i\omega_j))_{1\leq i,j\leq n}$ the *discriminant* of K, whose value is independent of the chosen integral basis. The degree of any number field K is polynomial in $\log|\Delta_K|$ and we will use this to simplify complexity statements accordingly [58]. Some of the results we describe will be dependent on the Generalized Riemann Hypothesis (GRH), a statement about the zeros of the Dedekind zeta function ζ_K. We will clearly state when this is the case.

For a non-zero prime ideal $\mathfrak{p}$ of $\mathcal{O}_K$ we will denote by $K_\mathfrak{p}$ the $\mathfrak{p}$-adic completion of K, by $v_\mathfrak{p}$ the $\mathfrak{p}$-adic valuation and by $k_\mathfrak{p} = \mathcal{O}_K/\mathfrak{p} \cong \mathcal{O}_{K_\mathfrak{p}}/\mathfrak{p}\mathcal{O}_{K_\mathfrak{p}}$ the residue field at $\mathfrak{p}$. We use a bar notation to denote cosets of various multiplicative groups, and $\langle X\rangle$ to denote the subgroup generated by X. A *fractional ideal* of K is a non-zero $\mathcal{O}_K$-submodule of K, or equivalently, a set of the form $\frac{\mathfrak{a}}{d}$, where $\mathfrak{a}$ is a non-zero ideal of $\mathcal{O}_K$ and $d \in \mathbb{Z}$, $d \neq 0$. The set I_K of fractional ideals of K is an abelian group with respect to multiplication with neutral element $\mathcal{O}_K$. The inverse of a fractional ideal $\mathfrak{a}$ is given by $\mathfrak{a}^{-1} = \{x \in K \mid x\mathfrak{a} \subseteq \mathcal{O}_K\}$. The group of fractional ideals is free on the set of non-zero prime ideals of $\mathcal{O}_K$ and therefore the norm map extends uniquely to a group homomorphism $I_K \to \mathbb{Q}$, which we also denote by N. Of particular interest are *principal fractional ideals*, which are the fractional ideals of the form $\alpha\mathcal{O}_K$ with $\alpha \in K^\times$. The set P_K of principal fractional ideals is a subgroup of the abelian group I_K. The quotient group I_K/P_K is the *(ideal) class group* of $\mathcal{O}_K$ or K, which we denote by $\mathrm{Cl}(\mathcal{O}_K)$. A classical theorem of algebraic number theory asserts that $\mathrm{Cl}(\mathcal{O}_K)$ is a finite group [54, Theorem 6.3]. We denote its order by h_K and call it the *class number* of $\mathcal{O}_K$ or K. When two fractional ideals $\mathfrak{a}$ and $\mathfrak{b}$ are in the same ideal class, we denote this as $\mathfrak{a} \sim \mathfrak{b}$. We call

$K = \mathbb{Q}(\zeta_m)$ the *cyclotomic field* of conductor m where ζ_m is the m-th primitive root of unity, and denote its totally real subfield by $K^+ := \mathbb{Q}(\zeta_m + \zeta_m^{-1})$. Define h_m^+ as the class number of K^+ and the minus part of the class group $\mathrm{Cl}^-(\mathcal{O}_K)$ as the kernel of the relative norm map $\mathrm{N}_{K/K^+}\colon \mathrm{Cl}(\mathcal{O}_K) \mapsto \mathrm{Cl}(\mathcal{O}_{K^+})$ which maps $[\mathfrak{a}] \mapsto [\mathfrak{a}\overline{\mathfrak{a}}]$, where $\bar{\cdot}$ denotes complex conjugation. The cardinality of $\mathrm{Cl}^-(\mathcal{O}_K)$ is denoted by h_m^- and we have $h_m = h_m^+ h_m^-$.

Since K is of degree n, there are n embeddings $\sigma\colon K \to \mathbb{C}$ which can be classified as follows: If the image of σ is contained in the real numbers $\mathbb{R}$, we call σ a *real embedding* of K. Otherwise σ is called a *complex embedding*. Because of complex conjugation, the complex embeddings of K always come in pairs. If $2s$ denotes the number of complex embeddings and r the number of real embeddings, then (r, s) is the *signature* of K. The *Minkowski embedding* is the map $K \to \mathbb{R}^r \oplus \mathbb{C}^s$, or equivalently $K \to \mathbb{R}^n$ induced by the embeddings of K modulo complex conjugation. Denote by $\sigma_1, \ldots, \sigma_r$ the real embeddings and by $\sigma_{r+1}, \overline{\sigma}_{r+1}, \ldots, \sigma_{r+s}, \overline{\sigma}_{r+s}$ the complex embeddings. We call K *totally real* if all embeddings are real and *totally complex* if all embeddings are complex.

As an ideal $\mathfrak{a} \subset \mathcal{O}_K$ is free $\mathbb{Z}$-module with a $\mathbb{Z}$-basis $(\alpha_1, \ldots, \alpha_n)$, we can view it under the Minkowski embedding as a full rank lattice in $\mathbb{R}^n$ called the *ideal lattice* with volume $\mathrm{Vol}(\mathfrak{a}) = \sqrt{|\Delta_K|}\mathrm{N}(\mathfrak{a})$. By [28] we have the following relationship between the shortest vector of an ideal and its algebraic norm.

$$\frac{1}{\mathrm{Poly}(n)}\mathrm{N}(\mathfrak{a})^{1/n} \leq \lambda_1(\mathfrak{a}) \leq \mathrm{Poly}(n)\mathrm{N}(\mathfrak{a})^{1/n}.$$

Given an embedding σ of K, $\alpha \mapsto \log|\sigma(\alpha)|$ is a group morphism mapping $K^\times \to \mathbb{R}$. We obtain the *logarithmic embedding*

$$\mathrm{Log}\colon K^\times \longrightarrow \mathbb{R}^{r+s},\ \alpha \longmapsto (\log(|\sigma_1(\alpha)|), \ldots, \log(|\sigma_{r+s}(\alpha)|)).$$

The celebrated theorem of Dirichlet asserts that $\mathrm{Log}(\mathcal{O}_K^\times)$ is a lattice of rank $r + s - 1$ and $\ker(\mathrm{Log})$ is equal to the torsion units of K. In particular $\mathcal{O}_K^\times \cong \mathbb{Z}^{r+s-1} \times T$, where T are the torsion units of K. For a set S of prime ideals of $\mathcal{O}_K$ we denote by $\mathcal{O}_{K,S}$ the ring of S-integers, that is, the elements $x \in K$ with $v_\mathfrak{p}(x) \geq 0$ for all $\mathfrak{p} \notin S$, and $\mathcal{O}_{K,S}^\times$ the group of S-units, i.e., the elements $x \in K^\times$ such that $v_\mathfrak{p}(x) = 0$ for all $\mathfrak{p} \notin S$. Equivalently, $\mathcal{O}_{K,S}^\times$ is the multiplicative group of all α that generate a principal ideal of the form $(\alpha)\mathcal{O}_K = \prod_{\mathfrak{p} \in S} \mathfrak{p}^{x_\mathfrak{p}}$.

2.3 Saturation of Multiplicative Groups

In the course of computing the unit group or S-unit group of a number field K we often first compute a finite index subgroup from which we wish to deduce the full group. More generally, assume we are given a subgroup V of W where $W \subseteq K^\times$. If $[W : V] = m$ then for some d dividing m there exists an element $\alpha \in W$ such that $\alpha^d \in V$ but $\alpha \notin V$. By finding such elements, taking their d-th roots, and enlarging V by adding these roots to its set of generators, we can reduce the index $[W : V]$. We refer to this process as *saturation*.

We define the *d-saturation* W of V as the smallest group $W \subseteq K^\times$ with $V \subseteq W$ where $K^\times/W$ is d-torsion free. We say the group V is *d-saturated* if it equals its d-saturation.

Lemma 2.5 ([20, Lemma 4.3])**.** *Let $V \subseteq K^\times$ be finitely generated. Then the following hold.*

1. *The d-saturation of V contains the d-torsion of $K^\times$.*
2. *The group V is d-saturated if and only if V is p-saturated for all primes p dividing d.*
3. *For a prime p the group V is not p-saturated if and only if there exists $\alpha \in K^\times \setminus V$ with $\alpha^p \in V$. In this case p divides the index $[\langle V, \alpha\rangle : V]$.*
4. *Let p be a prime and assume that V contains the p-torsion of $K^\times$. Then V is p-saturated if and only if $V \cap (K^\times)^p = V^p$.*

In practice, saturation techniques employ local computations to detect global powers. This is a well known technique in computational algebraic number theory, used for example in the class and unit group computation of number fields ([60, Section 5.7]) or the number field sieve ([1]). Note that, in contrast to previous applications of this technique, in our case the number d is in general not a prime. However, by Lemma 2.5 we may assume that d is a prime power, since it is enough to p-saturate at each prime p dividing d.

Let S be a finite set of prime ideals of $\mathcal{O}_K$. Recall that $K_\mathfrak{p}$ is the completion of K at the prime $\mathfrak{p}$. Consider the following map

$$K^\times/(K^\times)^d \longrightarrow \prod_{\mathfrak{p}\notin S} K_\mathfrak{p}^\times/(K_\mathfrak{p}^\times)^d.$$

If this map is an injection then it follows that any element that is locally a d-th power for each $\mathfrak{p} \notin S$ is a d-th power globally. This is a variation of the Hasse local-global principle. It turns out that by the following theorem of Grunwald–Wang (see [3, Chapter X] or [55, Chapter IX, §1]) this is not always true, and there is an obstruction to this local-global principle in certain cases.

Theorem 2.6 (Grunwald–Wang)**.** *Consider the canonical map*

$$K^\times/(K^\times)^d \longrightarrow \prod_{\mathfrak{p}\notin S} K_\mathfrak{p}^\times/(K_\mathfrak{p}^\times)^d.$$

Either this map is injective or the kernel is cyclic of order 2.

If the map of Theorem 2.6 is injective we say we are in the "good case", otherwise, using the terminology of [19], we say we are in the "bad case". This is also referred to as the "special case" in the literature. For $k \in \mathbb{Z}_{\geq 1}$ denote by ζ_k a primitive k-th root of unity and set $\eta_k = \zeta_k + \zeta_k^{-1}$. Let $s \geq 2$ be an integer such that $\eta_s \in K$ but $\eta_{s+1} \notin K$. We have a straightforward criterion for determining which case we are in. Recall that d is a prime power. We are in the bad case when the following conditions are simultaneously satisfied [3, Chapter X, Theorem 1]:

1. $d = 2^t$ with $t > s$.
2. The elements $-1, 2+\eta_s$ and $-(2+\eta_s)$ are non-squares in K.
3. We have $\{\mathfrak{p} \mid 2 \in \mathfrak{p}$ and $-1, 2+\eta_s$ and $-(2+\eta_s)$ are non-squares in $K_\mathfrak{p}\} \subseteq S$.

Given d, it is straightforward to test conditions (1) and (2). To test condition (3), it is sufficient to determine all prime ideals $\mathfrak{p}$ lying over 2 such that $-1, 2+\eta_s$ and $-(2+\eta_s)$ are non-squares in $K_\mathfrak{p}$. Being locally a square can be checked using the so-called quadratic defect [56, §63.A], which can be computed using an efficient algorithm due to Kirschmer [41, Algorithm 3.1.3]. Thus given K, d and S, we can always check whether we are in the good case or not. Although the conditions for being in the bad case look rather complicated, this situation is not as rare as it might appear. More precisely, if K is linear disjoint from the cyclotomic field $\mathbb{Q}(\zeta_8)$, then we are always in the bad case for $d = 2^t$, $t \geq 3$. Thus for almost all fields we are in the bad case at the prime 2.

2.3.1 The Good Case

Finding d-th powers in the good case can be done exclusively by detecting local d-th powers modulo a set of prime ideals. Recall that for a set S of prime ideals of $\mathcal{O}_K$ we denote by $\mathcal{O}_{K,S}$ the ring of S-integers, that is, the elements $x \in K$ with $v_\mathfrak{p}(x) \geq 0$ for all $\mathfrak{p} \notin S$, and $\mathcal{O}_{K,S}^\times$ the group of S-units, i.e., the elements $x \in K^\times$ such that $v_\mathfrak{p}(x) = 0$ for all $\mathfrak{p} \notin S$. Recall that $k_\mathfrak{p} = \mathcal{O}_K/\mathfrak{p} \cong \mathcal{O}_{K_\mathfrak{p}}/\mathfrak{p}\mathcal{O}_{K_\mathfrak{p}}$ is the residue field of K at $\mathfrak{p}$. We use the following proposition to detect local powers.

Proposition 2.7 ([20, Proposition 4.5])**.** *Assume that $\mathfrak{p}$ is a non-zero prime ideal with $d \notin \mathfrak{p}$ and let $\varpi \in K$ be a local uniformizer at $\mathfrak{p}$, that is, an element with $v_\mathfrak{p}(\varpi) = 1$. Then the map*

$$K_\mathfrak{p}^\times/(K_\mathfrak{p}^\times)^d \longrightarrow \mathbb{Z}/d\mathbb{Z} \times k_\mathfrak{p}^\times/(k_\mathfrak{p}^\times)^d, \quad \bar{x} \longmapsto (\overline{v}, \overline{x\varpi^{-v}}) \text{ where } v = v_\mathfrak{p}(x),$$

is an isomorphism.

By Proposition 2.7 if we restrict to degree 1 prime ideals then we are working in $k_\mathfrak{p} = \mathbb{F}_p$, which simplifies the computations needed. Denote the map of Proposition 2.7 by $\chi_\mathfrak{p}$. Note that $\chi_\mathfrak{p}$ induces a map on V/V^d for any $V \subseteq K^\times$. The following propositions show that we can compute the d-saturation of $V \subseteq K^\times$ by computing sufficiently many kernels of the $\chi_\mathfrak{p}$.

Proposition 2.8 ([19, Proposition 2])**.** *Assume that we are in the good case of Grunwald–Wang. For a multiplicative finitely generated subgroup $V \subseteq K^\times$ we have*

$$(V \cap (K^\times)^d)/V^d = \bigcap_{d \notin \mathfrak{p}} \ker(\chi_\mathfrak{p} \colon V/V^d \to \mathbb{Z}/d\mathbb{Z} \times k_\mathfrak{p}^\times/(k_\mathfrak{p}^\times)^d).$$

Furthermore, there exists $c_0 \in \mathbb{R}_{>0}$ (depending on K, V and d) such that

$$(V \cap (K^\times)^d)/V^d = \bigcap_{d \notin \mathfrak{p}, \mathrm{N}(\mathfrak{p}) \leq c_0} \ker(\chi_\mathfrak{p} \colon V/V^d \to \mathbb{Z}/d\mathbb{Z} \times k_\mathfrak{p}^\times/(k_\mathfrak{p}^\times)^d).$$

Proof. The first part is [20, Proposition 4.6]. As V is finitely generated, V/V^d is a finitely generated $(\mathbb{Z}/d\mathbb{Z})$-module. Thus V/V^d is Artinian and the existence of c_0 follows from the first part. □

Proposition 2.9. *Let $c \in \mathbb{R}_{>0}$, $V \subseteq K^\times$ finitely generated, let d be a power of a prime and let m be the dimension of the intersection*

$$\bigcap_{d \notin \mathfrak{p},\, \mathrm{N}(\mathfrak{p}) \leq c} \ker(\chi_\mathfrak{p} \colon V/V^d \to \mathbb{Z}/d\mathbb{Z} \times k_\mathfrak{p}^\times/(k_\mathfrak{p}^\times)^d) \subseteq V/V^d.$$

Assume it is generated by the classes of $\alpha_1, \ldots, \alpha_m \in V$.

1. *If $m = 0$, then V is d-saturated.*
2. *Assume that V is not d-saturated. Then if c is sufficiently large, there exists $1 \leq i \leq m$ such that α_i is a d-th power.*
3. *Assume that V is d-saturated. Then for c sufficiently large we have $m = 0$.*

Proof. Write W/V^d for the intersection. Since $(V \cap (K^\times)^d)/V^d \subseteq W/V^d$, (1) follows from Lemma 2.5. For (2) and (3) we take $c > c_0$ for the c_0 of Proposition 2.8, so $(V \cap (K^\times)^d)/V^d = W/V^d$. If V is not d-saturated then $m \neq 0$ and there is some $\alpha_i \in V \cap (K^\times)^d$ which is a d-th power. For (3), as d-saturation and p-saturation coincide when d is a prime power the result follows from [20, Proposition 4.8]. □

Proposition 2.9 is adapted from the version of [20, Proposition 4.8] which addresses the p-saturation of some $V \subseteq K^\times$. While d-saturation and p-saturation coincide when d is a prime power (by Lemma 2.5) the distinction is that we can find d-powers directly instead of iteratively working with p-powers. Proposition 2.8 shows we can find an upper bound c_0 on the norms of primes required, so we only need to compute finitely many intersections. In [20] a concrete bound is found under GRH. However, in practice we typically want to explicitly compute the d-saturation of V, for example to find generators for the unit group or S-unit group of K. So once we find a d-th power by Proposition 2.9 we will want to take the root anyway. If this is the case, there is no need to use primes up to c_0. If Proposition 2.9 finds the dimension of the intersection of kernels is 0 then we know the input is d-saturated. Otherwise, we have some *candidate* α_i that may be a d-th power, which we can test by explicitly computing the root. If this succeeds we are done, otherwise we need to increase the number of primes and continue. If we don't want to compute the root and just detect powers then we can use the bound c_0, which can be useful for decisional problems.

2.3.2 The General Case

In general, we may have $d = 2^t$ and be in the bad case of Grunwald–Wang. The solution to this is relatively simple. We can iteratively apply Proposition 2.9 with $d = 2$ (or even $d = 4$). If we find that V is not 2-saturated, we find an element which is a power of 2, compute its root, and add this to our generating set for V. Then we repeat the process. In practice this may be more computationally expensive than in the good case as it involves more root computations, but it does not change the overall asymptotic complexity. This is also the general strategy taken in [20], for any input d.

2.4 Compact Representation

It is evident from the previous section that saturation of multiplicative subgroups of $K^\times$ involves the computation of roots in $K^\times$. More precisely, given $\delta \in K$ and $d \in \mathbb{Z}_{>0}$, we need to decide whether there exists an element $\gamma \in K$ such that $\gamma^d = \delta$ (and if so, then compute γ). Since δ will in general be quite large, we first compute a *compact representation with respect to* d, which amounts to finding small elements $\delta_0, \ldots, \delta_k \in K$ such that

$$\delta = \delta_0 \cdot \delta_1^d \cdot \cdots \cdot \delta_k^{d^k}.$$

The advantage, of course, is that computing a d-th root should now be easier as we only need to compute a root of δ_0.

This presentation was introduced by Thiel [76] for units. An algorithm for finding such a presentation for S-units goes back to lecture notes of Fieker and has subsequently been used in [13, 14, 16, 34]. To the best of our knowledge, nobody has carefully analyzed this algorithm. Given such a presentation, it is clear that δ is a d-th power if and only if δ_0 is a d-th power. For the latter task, we use Hensel lifting of the linear factors of $X^d - \delta_0 \in K[X]$ modulo a non-zero prime ideal $\mathfrak{p}$ of $\mathcal{O}_K$, as described in [33]. To evaluate the performance of the compact representation, we need a notion of size of an element in K.

Recall that the number field K has $n = [K : \mathbb{Q}]$ distinct embeddings $\sigma_i \colon K \to \mathbb{C}$. In the following, for an element $\alpha \in K$, we denote by $\alpha^{(i)} = \sigma_i(\alpha)$, the image under such an embedding. The T_2-norm of α is defined to be

$$\|\alpha\| = \Big(\sum_{1 \le i \le r+2s} |\alpha^{(i)}|^2\Big)^{1/2},$$

which is just the Euclidean norm of α under the embedding $K \to \mathbb{C}^d$. For positive real numbers $(w_i)_{1 \le i \le d}$, we denote by

$$T_{2,(w_i)_i}(\alpha) = \sum_{1 \le i \le r+2s} w_i \cdot |\alpha^{(i)}|^2$$

the square of the weigthed T_2-norm. The value $\|\alpha\|$ is a good measure for the size of an element α. Indeed, as recalled in [15, Sec. 3], the maximum absolute value of a coefficient of $\alpha \in \mathcal{O}_K$ when represented on an LLL-reduced integral basis is less than $2^{3n/2}\|\alpha\|$, and, when $\alpha = \alpha_0/d$ for $d \in \mathbb{Z}_{>0}$ and $\alpha_0 \in \mathcal{O}_K$, the bit size $\mathrm{S}(\alpha)$ of α is less than $n\left(\frac{3n}{2}\log\|\alpha\| + \log(d)\right)$.

Algorithm 1: Compact representation algorithm

Input : $\alpha \in \mathcal{O}_K$, such that $\mathrm{N}(\alpha)$ has known factorization, and an integer $d > 1$
Output: $\alpha_0, \ldots, \alpha_k$ such that $\alpha = \prod_i \alpha_i^{d^i}$

$t \leftarrow \max(\mathrm{N}(\alpha), 2)$;
$A_i \leftarrow \prod_{\mathfrak{p}} \mathfrak{p}^{\lfloor (v_{\mathfrak{p}}(\alpha) \bmod d^{i+1})/d^i \rfloor}$ for $i = 0, \ldots, \log_d \log(t)$ and $A_i \leftarrow \mathcal{O}_K$ for $i > \log_d \log(t)$;
Find D such that $\mathrm{N}(A_i) \le D$ for all $i = 0, \ldots, \log_d \log(t)$;
Find k such that $\sqrt[d^k]{|\alpha^{(j)}|} \le D$ for all $1 \le j \le d$;
$I \leftarrow \mathcal{O}_K$;
while $k > 0$ **do**
  $I \leftarrow I A_k$;
  $w_j \leftarrow \sqrt[d^k]{|\alpha^{(j)}|}$ for $j = 1, \ldots, d$, $w \leftarrow \sqrt[d]{\prod w_i}$;
  Find γ_k a $T_{2,(w_j/w)_j}$-LLL-short element in I^{-1};
  $\alpha \leftarrow \alpha \gamma_k^{d^k}$, $I \leftarrow I \cdot (\gamma)$;
  $k \leftarrow k - 1$;
  $I \leftarrow I^d$;
end
return: $(\gamma_i^{-1})_{i=0,\ldots,k}$;

Theorem 2.10. *Algorithm 1 is correct and runs in time*

$$\mathrm{Poly}(\log|\Delta_K|, d, \mathrm{S}(\alpha)).$$

Proof. First of all, notice that the ideal $I \cdot (\gamma)$ is always integral, so I is always integral. By using LLL, in Step 9, $\mathrm{N}(I \cdot (\gamma))$ is bounded by $C \cdot \sqrt{|\Delta_K|}$ where C is the approximation factor in LLL. Therefore, $\mathrm{N}(I)$ has polynomial bit size. In Step 8, we have $w = \sqrt[n]{\prod_j w_j} = \sqrt[n]{\sqrt[d^k]{|\mathrm{N}(\alpha)|}}$. Now at the beginning of the while loop (before executing Step (7)) we always have $\alpha\mathcal{O}_K = I^{d^k} \prod_{i=0}^{k} A_i^{d^i}$. Thus, $\sqrt[d^k]{|\mathrm{N}(\alpha)|}$ stays bounded, and w as well.

We show next that $\|(w_j)_j\|_2$ is bounded during Steps (6)–(11). This is true at the beginning as k is taken large enough for this to be correct. The update at Step (10) implies that $(w_j)_j \leftarrow ((w_j \cdot |\gamma^{(j)}|)^d)_j$, where $\gamma = \gamma_{k+1}$. We then have

$$\|((w_j \cdot |\gamma^{(j)}|)^d)_j\|_2 \leq \|(w_j \cdot |\gamma^{(j)}|)_j\|_2^d.$$

Now

$$\|(w_j \cdot |\gamma^{(j)}|)_j\|_2^2 = T_{2,(w_j)_j}(\gamma) = w^2 \cdot T_{2,(w_j/w)_j}(\gamma) \leq C'$$

since w is bounded, and $T_{2,(w_j/w)_j}(\gamma)$ is bounded as well because it is LLL-reduced.

Finally, $\gamma \in I^{-1}$, but $\mathrm{N}(I)$ is bounded, hence so is the denominator of γ. From the bound on $\|(w_j)_j\|_2$ and the $T_{2,(w_j)_j}$ we get a polynomial bound on the total size of γ, and therefore on its inverse as well.

We now bound the size of D and k. Assume that α has support S, that is, $\alpha \in \mathcal{O}_{K,S}$. Then $\mathrm{N}(A_i) \leq \prod_{\mathfrak{p} \in S} \mathrm{N}(\mathfrak{p})^{d-1}$, which shows that we can choose $\log(D) = d \cdot |S| \cdot \max_{\mathfrak{p} \in S} \log(\mathrm{N}(\mathfrak{p})) \in \mathrm{Poly}(d, \|\alpha\|) = \mathrm{Poly}(d, \mathrm{S}(\alpha))$. As $|\alpha^{(j)}| \leq \|\alpha\|$, we can choose k such that $\sqrt[d^k]{\|\alpha\|} \leq D$, which is true as soon as we have $(\log\log\|\alpha\| - \log\log(D))/\log(d) \leq k$. Hence $k \in \mathrm{Poly}(d, \log\|\alpha\|) = \mathrm{Poly}(d, \mathrm{S}(\alpha))$. □

Corollary 2.11. *Assume that $\alpha \in \mathcal{O}_K$ is given as $\alpha = \prod_{j=1}^{l} \beta_j^{e_j}$ with $\beta_j \in K$ and $e_j \in \mathbb{Z}$ and known factorization of $\mathrm{N}(\beta_j)$ for all $1 \leq j \leq l$. Then Algorithm 1 runs in time*

$$\mathrm{Poly}(\log|\Delta_K|, d, l, \max_j \mathrm{S}(\beta_j), \max_j \log(e_j)).$$

Proof. Let S be the support of α. Then as in the previous proof we can take D with $\log(D) = d \cdot |S| \cdot \max_{\mathfrak{p}} \log(\mathrm{N}(\mathfrak{p})) \in \mathrm{Poly}(d, \max_j \mathrm{S}(\beta_j))$. Similar, it is sufficient to take k with $\log\log\|\alpha\| \leq k$. As

$$\begin{aligned}\log\log\|\alpha\| &= \log\left(\sum_{j=1}^{l} e_j \log\|\beta_j\|\right) \\ &\leq \log(l) \max_j \log(e_j) \max_j \log\|\beta_j\| \in \mathrm{Poly}(\max_j \log(e_j), \max_j \mathrm{S}(\beta_j)),\end{aligned}$$

the claim follows □

Subfield unit calculations, and subfield resolutions of the PIP are assumed to be followed by a compact representation routine. In both cases, the prime factorization of the input is known in advance, therefore $\mathrm{Fact}(\mathrm{N}(\alpha)) = 0$. Moreover, the product of a polynomial number of terms in compact representation can be kept in a compact representation by direct multiplication of the terms. Therefore, the compact representation algorithm is only executed in subfields. Then, operations on compact representations in field extensions have polynomial run time.

2.5 Subexponential Methods For the PIP and Related Problems

In this section, we recall the main results of the subexponential method for computing the class group and solving the PIP in number fields of large degree of [11, 14]. The first ingredient of this method is a reduction algorithm that takes as input an ideal $\mathfrak{a} \subseteq \mathcal{O}_K$ and returns another ideal of norm bounded by the invariant of the fields only in the same ideal class as $\mathfrak{a}$.

Algorithm 2: BKZ-ideal reduction [14, Algorithm 2]

Require: $\mathfrak{a} \subseteq \mathcal{O}_K$ and block size $k > 0$.
Ensure: $\mathfrak{b} \subseteq \mathcal{O}_K$ and $\alpha \in K$ such that $\mathfrak{b} = (\alpha)\mathfrak{a}$ has bounded norm.
1: $\frac{\mathfrak{c}}{d} \leftarrow \mathfrak{a}^{-1}$ where $\mathfrak{c} \subseteq \mathcal{O}_K$, and $d > 0$.
2: $\gamma \leftarrow$ first element of a BKZ-reduced basis of $\mathfrak{c}$ with block size k.
3: $\mathfrak{b} \leftarrow (\gamma/d)\mathfrak{a}$
4: **return** $\mathfrak{b}, (\gamma/d)$

Proposition 2.12. *Algorithm 2 with block size k runs in time*

$$\mathrm{Poly}(\log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a}))) \cdot 2^{O(k)},$$

and returns $\mathfrak{b}$ such that

- $\mathrm{N}(\mathfrak{b}) \in 2^{\tilde{O}(n^2/k)}\sqrt{|\Delta_K|}$.

- $\log(d), \log\|\gamma\| \in \mathrm{Poly}(\log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a})))$.

Proof. BKZ with block size k finds $\gamma \in \mathfrak{c}$ such that

$$\begin{aligned}\|\gamma\| &\leq k^{n/2k}\,\mathrm{Vol}(\mathfrak{c})^{1/n} \\ &= k^{n/2k}|\Delta_K|^{1/2n}\mathrm{N}(\mathfrak{c})^{1/n}\end{aligned}$$

in time $\mathrm{Poly}([K:\mathbb{Q}], \log(\mathrm{N}(\mathfrak{c}))) \cdot 2^{O(k)}$ [66]. Moreover, we have $d \leq \mathrm{N}(\mathfrak{a})$ and

$$\mathrm{N}(\mathfrak{c}) = \mathrm{N}(d\mathfrak{a}^{-1}) \leq d^n/\mathrm{N}(\mathfrak{a}) \leq d^{n-1} \leq \mathrm{N}(\mathfrak{a})^{n-1} \leq \mathrm{N}(\mathfrak{a})^n.$$

This shows the bounds on the bit size of d and $\|\gamma\|$. Additionally, we have

$$\mathrm{N}(\mathfrak{b}) = \frac{\mathrm{N}(\gamma)}{\mathrm{N}(d)}\mathrm{N}(\mathfrak{a}) = \frac{\mathrm{N}(\gamma)}{\mathrm{N}(d)}\frac{\mathrm{N}(d)}{\mathrm{N}(\mathfrak{c})} \leq \frac{\|\gamma\|^n}{\mathrm{N}(\mathfrak{c})} \leq \frac{2^{\tilde{O}(n^2/k)}\sqrt{|\Delta_K|}\mathrm{N}(\mathfrak{c})}{\mathrm{N}(\mathfrak{c})}.$$

□

Given an input ideal $\mathfrak{a}$ and a set of prime ideals S that generate $\mathrm{Cl}(\mathcal{O}_K)$, we want to return a decomposition of the ideal class of $\mathfrak{a}$ over $\langle S\rangle$. This is done by multiplying short products of primes in S, BKZ-reducing the resulting ideal, and hoping that it decomposes as a product of elements in S. To simplify the analysis, we assume that $S = \{\mathfrak{p} \text{ prime ideals with } \mathrm{N}(\mathfrak{p}) \leq 2^{(\log|\Delta_K|)^{2/3}}\}$. We use the fact that under the GRH, primes of norm up to $12(\log|\Delta_K|)^2$ generate $\mathrm{Cl}(\mathcal{O}_K)$ [4], and that the class of an ideal multiplied by a short product of such primes is almost uniformly distributed in $\mathrm{Cl}(\mathcal{O}_K)$. This procedure is described in Algorithm 3. Thus

Algorithm 3: Ideal decomposition [14, Algorithm 3]

Require: $\mathfrak{a} \in \mathcal{O}_K$.
Ensure: $\alpha \in K$ and $(x_i)_{i\leq l} \in \mathbb{Z}^l$ for some l with $(\alpha)\mathfrak{a} = \prod_i^l \mathfrak{p}_i^{x_i}$ and $\mathrm{N}(\mathfrak{p}_i) \leq 2^{\tilde{O}((\log|\Delta_K|)^{2/3})}$.

$S = \{\mathfrak{p} \text{ prime ideals with } \mathrm{N}(\mathfrak{p}) \leq 2^{(\log|\Delta_K|)^{2/3}}\}$, and $l \leftarrow |S|$.
$S_0 = \{\mathfrak{p} \text{ prime ideals with } \mathrm{N}(\mathfrak{p}) \leq 12(\log|\Delta_K|)^2\}$, and $l_0 \leftarrow |S|$
while true **do**
  $(x_i)_{i\leq l} \xleftarrow{\mathcal{R}} [0, \log|\Delta_K|]^{l_0}$. $\mathfrak{a}' \leftarrow \mathfrak{a}\prod_{i\leq l_0}\mathfrak{p}_i^{x_i}$.
  Compute BKZ-reduced $\mathfrak{b}, \alpha$ with Algorithm 2 on input $\mathfrak{a}'$, $k = n^{2/3}$.
  if $\mathfrak{b}$ is S-smooth **then**
    Compute $\vec{y}$ such that $\mathfrak{b} = \prod_{i\leq l}\mathfrak{p}_i^{y_i}$.
    $\vec{x} \leftarrow \vec{y} - \vec{x}||\vec{0}$.
    return $\alpha, \vec{x}$
  end if
end while

far, the run time of Algorithm 3 has only been analysed heuristically. The probability of $\mathfrak{b} = (\alpha)\mathfrak{a}'$ to be

S-smooth is not rigorously understood at this point, but there are rigorous results mentioned in [14, Sec. 3.1] showing that the proportion of ideals of norm less than ι that are a product of prime ideals of norm less than μ is $e^{-u\log(u)(1+o(1))}$ where $u = \log(\iota)/\log(\mu)$. Heuristic 1 of [14, Sec. 3.1] postulates that this is also the smoothness probability of the reduced ideal $\mathfrak{a}$ of Step 6. Due to the high connectivity of the Caley graph of $\mathrm{Cl}(\mathcal{O}_K)$, we can argue that the ideal class of I' is distributed almost uniformly at random, but so far, there is no rigorous proof of how the multiplication by α obtained with Algorithm 2 influences the smoothness probability.

Conjecture 2.13 (Rephrasing of Heuristic 1 of [14])**.** *Let $k > 0$, $\mathfrak{a}$ be an ideal in a class of $Cl(\mathcal{O}_K)$ that is drawn uniformly at random, and $\mathfrak{a}'$ be the output of Algorithm 2 with input $\mathfrak{a}$ and k. Then the probability of $\mathfrak{a}'$ being a product of prime ideals of norm less than μ is $e^{-u\log(u)(1+o(1))}$ where $u = \log(\mathrm{N}(\mathfrak{a}'))/\log(\mu)$.*

Proposition 2.14 (under GRH and Conjecture 2.13)**.** *Algorithm 3 is correct, has asymptotic complexity in $\mathrm{Poly}\left(\log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a}))\right) \cdot 2^{\tilde{O}\left((\log|\Delta_K|)^{2/3}\right)}$, and returns $\alpha = (\gamma/d)$, $\vec{x} \in \mathbb{Z}^l$ where $\gamma \in \mathcal{O}_K$, $d \in \mathbb{Z}_{>0}$, and $(\gamma/d)\mathfrak{a} = \prod_{i\leq l} \mathfrak{p}_i^{x_i}$ with*

- $\log(d), \log\|\gamma\| \in \mathrm{Poly}\left(\log|\Delta_K|\right), \log(\mathrm{N}(\mathfrak{a}))\right)$.
- $\log\|\vec{x}\| \in \mathrm{Poly}\left(\log|\Delta_K|\right)$.

Proof. We apply Proposition 2.12 to the ideal $\mathfrak{a}' = \mathfrak{a} \cdot \prod_{i\leq l_0} \mathfrak{p}_i^{x_i}$. It satisfies

$$\log(\mathrm{N}(\mathfrak{a}')) \in \mathrm{Poly}\left(\log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a}))\right)$$

which proves the bound on $\|\gamma\|$ and d. Moreover, the runtime is $\mathrm{Poly}\left(\log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a}))\right) \cdot 2^{\tilde{O}(k)}$ where $k = n^{2/3}$ is the block size used for the BKZ reduction, hence showing the cost of one reduction. Then, assuming Conjecture 2.13, the probability that the resulting reduced ideals $\mathfrak{a}'$ whose norms satisfy $\log(\mathrm{N}(\mathfrak{a}')) \in \tilde{O}\left((\log|\Delta_K|)^{4/3}\right)$ be S-smooth is in

$$\frac{1}{2^{\tilde{O}\left(\frac{(\log|\Delta_K|)^{4/3}}{(\log|\Delta_K|)^{2/3}}\right)}} = \frac{1}{2^{\tilde{O}\left((\log|\Delta_K|)^{2/3}\right)}}.$$

This shows that the expected cost to find a relation is in $\mathrm{Poly}\left(\log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a}))\right) \cdot 2^{\tilde{O}\left((\log|\Delta_K|)^{2/3}\right)}$. Finally, the size of the output vector derives from the fact that it is of the form $\vec{y} - \vec{x}||\vec{0}$ where $\log\|\vec{x}\| \in \mathrm{Poly}(\log|\Delta_K|)$ by construction, while $\vec{y}$ is the decomposition of the BKZ-reduced ideal $\mathfrak{a}'$ with respect to S. □

The general strategy to solve the PIP is to first apply Algorithm 3 to $\mathfrak{a} = \mathcal{O}_K$ as many times as it takes in order to compute a basis for the lattice $\mathcal{L}$ of vectors $\vec{x} \in \mathbb{Z}^l$ such that $\prod_i \mathfrak{p}_i^{x_i} \sim (1)$, i.e. the so-called *lattice of*

relations between elements of S. Then the input ideal $\mathfrak{a}$ is decomposed with Algorithm 3, and it is principal if and only if its decomposition vector $\vec{x}$ belongs to $\mathcal{L}$. To justify the run time of Algorithm 4 we need to make an additional heuristic, which corresponds to Heuristic 3 of [14]. It argues that the relations drawn during Step 4 of Algorithm 4 are well-enough distributed among the full lattice of relations between classes of primes in S. Even though Algorithm 3 uses randomization, we have no guarantee on the distribution of the relations we create. In [38, Sec. 3.1], Hafner and McCurley show how to estimate this distribution rigorously in the case of quadratic fields, and they show in [38, Sec. 3.2] that once a sublattice of rank $|S|$ is found, only $|S|^{1+o(1)}$ extra relations need to be found randomly to complete the lattice of relations.

Conjecture 2.15 (Rephrasing of Heuristic 3 of [14])**.** *With probability* $1-1/|\Delta_K|$*, the number of iterations of Steps 4 and 5 of Algorithm 4 is bounded by* $|S|^{1+o(1)}$*.*

Algorithm 4: PIP resolution

Require: $\mathfrak{a} \in \mathcal{O}_K$.
Ensure: If $\mathfrak{a}$ is principal, $\beta_0, \ldots, \beta_k \in K$, and $(x_i)_{i\leq k} \in \mathbb{Z}^k$ with $\mathfrak{a} = \left(\beta_0 \cdot \prod_{i>0} \beta_i^{x_i}\right) \mathcal{O}_K$.

1: $S = \{\mathfrak{p}$ prime ideals with $\mathrm{N}(\mathfrak{p}) \leq 2^{(\log|\Delta_K|)^{2/3}}\}$, and $l \leftarrow |S|$.
2: $\mathcal{L} \leftarrow \{\}$. $M \leftarrow () \in \mathbb{Z}^{0\times l}$, $L_\alpha \leftarrow \{\}$.
3: **while** $\mathcal{L}$ is not the full lattice of relations between primes in S **do**
4: $\quad \alpha, \vec{v} \leftarrow$ output of Algorithm 3 on input $\mathcal{O}_K$.
5: $\quad \mathcal{L} \leftarrow \mathcal{L} + \mathbb{Z}\vec{v}$, $M \leftarrow \left(\frac{M}{\vec{v}}\right)$, $L_\alpha \leftarrow L_\alpha || \{\vec{v}\}$.
6: **end while**
7: $k \leftarrow |L_\alpha|$. Find $U \in \mathrm{GL}_k(\mathbb{Z})$ such that UM is the Hermite Normal Form of M.
8: Let $H \in \mathbb{Z}^{l\times l}$ such that $UM = \left(\frac{H}{(0)}\right)$.
9: $\alpha, \vec{y} \leftarrow$ output of Algorithm 3 on input $\mathfrak{a}$.
10: **if** $\vec{x}H = \vec{y}$ has a solution **then**
11: $\quad \vec{x} \leftarrow (\vec{x}||\vec{0})U$, and $\beta_0 \leftarrow 1/\alpha$, $\beta_i \leftarrow \alpha_i$ for $i \geq 1$.
12: $\quad$ **return** $\beta_0, \ldots, \beta_k \in K$, and $(x_i)_{i\leq k} \in \mathbb{Z}^k$.
13: **else**
14: $\quad$ **return** $\mathfrak{a}$ is not principal.
15: **end if**

Theorem 2.16 (Under GRH and Conjectures 2.13-2.15)**.** *Algorithm 4 is correct, runs in time*

$$\mathrm{Poly}\left(\log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a}))\right) \cdot 2^{\tilde{O}\left((\log|\Delta_K|)^{2/3}\right)}$$

with probability $1-1/|\Delta_K|$*, and its output satisfies*

- $\beta_i = \frac{\gamma_i}{d_i}$ *with* $\gamma_i \in \mathcal{O}_K$, $d_i \in \mathbb{Z}_{>0}$.
- $\log\|\gamma_0\|, \log(d_0) \in \mathrm{Poly}(\log|\Delta_K|, \log(\mathrm{N}(I)))$.
- $\log\|\gamma_i\|, \log(d_i) \in \mathrm{Poly}(\log|\Delta_K|)$*, for* $d > 0$.

- $\log\|\vec{x}\| \in 2^{\tilde{O}\left((\log|\Delta_K|)^{2/3}\right)}$.

Proof. From Conjecture 2.15, the number of iterations of Steps 4 and 5 is in $2^{\tilde{O}\left((\log|\Delta_K|)^{2/3}\right)}$, so we have $l, k \in 2^{\tilde{O}\left((\log|\Delta_K|)^{2/3}\right)}$. The cost of the calculation of the Hermite Normal Form of M is polynomial in the size of M, and the resulting U satisfies $\log\left(\|U\|_\infty\right) \in 2^{\tilde{O}\left((\log|\Delta_K|)^{2/3}\right)}$. This shows the bound on $\|x\|$. The bounds on the sizes of the β_i derive directly from Proposition 2.14. □

Corollary 2.17. *The cost of applying the compact representation method of Algorithm 1 to the output of Algorithm 4 is in*

$$\mathrm{Poly}(\log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a}))) \cdot 2^{\tilde{O}((\log|\Delta_K|)^{2/3})},$$

and it yields a compact representation $\alpha = \prod_{i \le k} \alpha_i^{n^i}$ *satisfying*

$$k, \log\|\gamma_i\|, \log(d_i) \in \mathrm{Poly}(\log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a}))),$$

where $\alpha_i = \frac{\gamma_i}{d_i}$ *with* $\gamma_i \in \mathcal{O}_K$, *and* $d_i \in \mathbb{Z}_{>0}$.

Proof. The important point here is that the decomposition of $(1/\beta_0) \cdot \mathfrak{a}$ is known and has polynomial size (it is a BKZ-reduced smooth ideal). So Algorithm 1 should be applied to $\mathfrak{a}' = (1/\beta_0) \cdot \mathfrak{a}$ generated by $\prod_{i>0} \beta_i^{x_i}$. The size of the β_i is polynomial. The bulk of the effort consists in handling the coefficients x_i which have subexponential bit size in $2^{\tilde{O}((\log|\Delta_K|)^{2/3})}$, hence this dependency in the asymptotic cost. In the end, the first term of the compact representation of the generator of $\mathfrak{a}'$ is multiplied by γ_0, thus yielding the compact representation of $\mathfrak{a}$ with the desired properties. □

Finally, we make a note regarding the S-unit group $\mathcal{O}_{K,S}^\times$. The heuristic subexponential methods of [11, 14] allow us to compute S-unit groups for S a generating set of the class group as well as ideal decompositions and solve the PIP, as we have seen. Combined with Simon's work [74] or Cohen [24, Algorithm 7.4.8] this also allows the computation of S-unit groups for arbitrary sets S.

Lemma 2.18. *Let* K *be a number field of degree* n, $\mathfrak{a}$ *be an ideal of* K *and* S *be a set of prime ideals of* K. *Then, under the heuristics of [14], the computation of the* S*-unit group takes time in*

$$\mathrm{Poly}(\log(|\Delta_K|, \max_{\mathfrak{p} \in S} \log(\mathrm{N}(\mathfrak{p})), |S|) \cdot 2^{\tilde{O}((\log|\Delta|)^{2/3})}.$$

CHAPTER 3
NORM RELATIONS

3.1 Background

Exploiting the subfield structure of number fields to reduce difficult problems to smaller instances is not a new concept. As early as 1842 Dirichlet [31] discovered a relationship between the class number of a biquadratic field $\mathbb{Q}(\sqrt{d_1}, \sqrt{d_2})$ and it's quadratic subfields. In 1899 Weber [80] described a formula expressing the relative class number of the conductor 2^k cyclotomic field $\mathbb{Q}(\zeta_{2^k})$ in terms of subfields, exploiting the repeating subfield structure. Brauer [21] and Kuroda [43] placed these ideas on a single foundation by studying similar relationships on group characters, and Kuroda's resulting class number formula for biquadratic fields directly generalized the identity of Dirichlet. This formula was also studied by Walter [78] who exploited a correspondence between a related group-theoretic interpretation and an identity on relative norms of number field elements to give a simple proof of Kuroda's formula.

These techniques however did not see much application to algorithmic number theory: "Until recently, the use of subfields in algorithmic number theory had been restricted to *ad hoc* tricks and heuristic observations" [20]. One of the first algorithmic applications was to multiquadratic fields of the form $\mathbb{Q}(\sqrt{d_1}, \ldots, \sqrt{d_n})$ for $d_i \in \mathbb{Z}$. In [77] Wada used subfield information to describe an exponential time algorithm for computing the unit group of a multiquadratic field. Bauch, Bernstein, de Valence, Lange and Van Vredendaal [5] expanded on this approach and described more efficient techniques for computing the unit group and resolving the PIP in multiquadratic fields. This was generalized to the recursive computation of S-unit groups in multiquadratics by Biasse and Van Vredendaal [17], and Lesavourey, Plantard, and Susilo [47] further extended these techniques to multicubic fields, number fields that are generated by cube roots of integers.

The key component underpining these results was the following relationship among relative norms of field elements.

Example 3.1. Let $K = \mathbb{Q}(\sqrt{d_1}, \sqrt{d_2})$ with $d_1, d_2 \in \mathbb{Z}$ coprime and squarefree. Then K has Galois group $\mathrm{Gal}(K/\mathbb{Q}) = \langle\sigma\rangle \times \langle\tau\rangle \cong C_2 \times C_2$. Denote by K_σ, K_τ, and $K_{\sigma\tau}$ the quadratic subfields of K fixed by τ, σ,

and $\sigma\tau$ respectively. Then for $\alpha \in K^\times$ we have the following identity:

$$\alpha^2 = \frac{\mathrm{N}_{K/K_\sigma}(\alpha)\mathrm{N}_{K/K_\tau}(\alpha)\mathrm{N}_{K/K_{\sigma\tau}}(\alpha)}{\mathrm{N}_{K/\mathbb{Q}}(\alpha)}. \tag{3.1}$$

Note that this was originally given in [5, 17] as

$$\alpha^2 = \frac{\mathrm{N}_{K/K_\sigma}(\alpha)\mathrm{N}_{K/K_\tau}(\alpha)}{\sigma(\mathrm{N}_{K/K_{\sigma\tau}}(\alpha))}$$

and it is easy to show they are equivalent.

This particular identity is the same as the one in [5, 17, 77]. A similar identity is at the heart of the multicubic variant of [47]. Biasse, Fieker, Hoffman, and Page [20] generalized the technique to arbitrary non-cyclic number fields via the notion of a *norm relation*, closely related to the relations used previously in the study of class number formulae. In what follows we give a summary of norm relations and their applications as developed in [20].

3.2 Definition

Let G be a group. For $H \leq G$ define the norm element $\mathrm{N}_H = \sum_{h\in H} h \in \mathbb{Q}[G]$.

Definition 3.2 ([20, Definition 2.1])**.** Let G be a finite group, $\mathcal{H}$ a set of subgroups of G. $\mathcal{H}$ will be omitted from the terminology whenever it is the set of all subgroups of G.

1. A *norm relation* of G with respect to $\mathcal{H}$ is an equality of the form

$$1 = \sum_{i=1}^{l} a_i N_{H_i} b_i$$

with $a_i, b_i \in \mathbb{Q}[G]$ and $H_i \in \mathcal{H}, H_i \neq 1$, where the equality holds in the group algebra $\mathbb{Q}[G]$.

2. A *scalar norm relation* of G with respect to $\mathcal{H}$ is an equality of the form

$$0 = \sum_{H\in\mathcal{H}} a_H N_H$$

with $a_H \in \mathbb{Q}$ and $a_1 \neq 0$, where the equality holds in the group algebra $\mathbb{Q}[G]$.

We define the *denominator* of a norm relation as the least common denominator of the coefficients of the $a_i, b_i \in \mathbb{Q}[G]$. By clearing the denominators we can rewrite this relation in the form

$$d = \sum_{i=1}^{l} a_i N_{H_i} b_i$$

where now $a_i, b_i \in \mathbb{Z}[G]$, $d \in \mathbb{Z}_{>0}$. By [20, Remark 2.17] if we have a scalar norm relation of the form $0 = \sum_{H \in \mathcal{H}} a_h N_H$ with $a_H \in \mathbb{Z}$ we can view it as a norm relation of the form $1 = \sum_{H \in \mathcal{H}} \frac{-a_H}{a_1} N_H$. Clearing denominators again we find that a scalar norm relation with denominator d has the form

$$d = \sum_{H \in \mathcal{H}} b_H N_h$$

with $d, b_H \in \mathbb{Z}$ coprime.

Definition 3.2 makes some simplifications from [20, Definition 2.1]. Indeed, we ignore for now the closely related Brauer relations and avoid generalizing to arbitrary group rings, as these are not necessary for our purposes. Now we turn our attention to the question of existence.

Theorem 3.3 ([20, Theorem A])**.** *The group G admits a norm relation if and only if G contains a noncyclic subgroup of order pq, where p and q are primes, or a subgroup isomorphic to $SL_2(\mathbb{F}_p)$ where $p = 2^{2^k} + 1$ is a Fermat prime with $k > 1$.*

When a norm relation does exist it is not always obvious how to determine the corresponding subgroups and coefficients. As they are not necessarily unique, we can also ask how to find a norm relation which is optimal with respect to the denominator and index of the involved subgroups (this notion of an "optimal" norm relation and its motivation will become clearer in Section 3.4.) In fact, it turns out that when G is a finite abelian group there is a simple, closed formula for producing just such a norm relation.

3.3 Norm Relations in Finite Abelian Groups

Let G be a finite abelian group. In this case we can be much more precise about the existence and properties of norm relations. First we recall the explicit construction of [20, Proposition 2.26] for a norm relation on G when one exists.

Proposition 3.4 ([20, Proposition 2.26])**.** *Let μ denote the Möbius function. For $n > 1$ an integer, let $rad(n) = \prod_{p|n} p$.*

Let G be a non-cyclic abelian group.

1. *We have the norm relation* $\mathcal{R}_G$*:*

$$1 = \sum_{C=\langle\chi\rangle \leq \widehat{G} \ cyclic} a_{\ker\chi} N_{\ker\chi},$$

where

$$a_{\ker\chi} = \frac{1}{|\ker\chi|} \sum_{C \leq C' \leq \widehat{G} \ cyclic} \mu([C' : C]).$$

2. *We have*

$$a_{\ker\chi} = \frac{c}{|G|} \prod_{p|c} \left(1 - p^{r_p - 1}\delta_{\chi,p}\right) \prod_{p||G|,\, p\nmid c} \left(-p - p^2 - \cdots - p^{r_p-1}\right)$$

where c *denotes the order of* χ*, where* $\delta_{\chi,p} = 1$ *or* 0 *according as whether there exists* $\chi' \in \widehat{G}$ *such that* $(\chi')^p = \chi$*, and where* $r_p = \dim_{\mathbb{F}_p}(G/G^p)$ *denotes the* p*-rank of* G*.*

3. *The denominator of* $\mathcal{R}_G$ *is* $\frac{|G|}{rad(|G|)} \neq 1$*.*

In [20, Theorem 2.28] this formula is applied to produce a *scalar* norm relation on G which is optimal, in the sense that the denominator and the index of the involved subgroups are as small as possible. It can be considered as two cases.

Theorem 3.5 ([20, Theorem 2.28])**.** *Write* $G \cong C \times Q$ *where* C *is the largest cyclic factor of* G*.*

1. *The group* G *admits a norm relation with denominator* 1 *if and only if* $|Q|$ *is divisible by at least two distinct primes. If the condition is satisfied, then* G *admits a scalar norm relation with* $a_i \in \mathbb{Z}$*, denominator* 1*, and where all* H_i *satisfy that* G/H_i *is a* p_i*-group times a cyclic group, for some prime number* p_i*.*

2. *Assume that* Q *is a* p*-group. Then* G *admits a norm relation if and only if* $Q \neq 1$*. If the condition is satisfied, then* G *admits a scalar norm relation with* $a_i \in \mathbb{Z}$*, denominator a power of* p *and where all* H_i *satisfy that* G/H_i *is a cyclic group.*

Note that we omit some details here from the original formulation of [20, Theorem 2.28]. Let's consider an example where both cases of Theorem 3.5 will appear.

Example 3.6. Let $G = C_2^2 \times C_3^2$ of order 36. Then $G \cong C \times Q$ where $C = C_6, Q = C_6$. This is the smallest example of a finite abelian group with a denominator 1 norm relation, and admits a norm relation with respect to $H_i < G$ where G/H_i is isomorphic to one of $C_2 \times C_3^2, C_2^2 \times C_3, C_2 \times C_2$ or $C_3 \times C_3$, and each H_i has index at most 18 in G.

Example 3.7. Let $G = C_2 \times C_3^2$, one of the quotient groups appearing in the previous example. Then $G \cong C \times Q$ where $C = C_6, Q = C_3$. Now Q is a p-group and admits a norm relation with respect to subgroups H_i of G where G/H_i is cyclic, hence $H_i = G$ or the subgroups of C_3^2.

As evidenced by these examples, when G admits a denominator 1 norm relation with respect to H_i the quotient groups G/H_i admit a p-power denominator norm relation as well. This is an important property that we will take advantage of in Section 3.4.1.

3.4 Norm Relations in Number Fields

Let K/F be a normal extension of algebraic number fields with Galois group G. The connection between norm relations and identities of relative norms like those seen in [5, 17, 47, 77] can be explained by the following. By [20, Proposition 3.5] G admits a norm relation

$$d = \sum_{i=1}^{l} a_i N_{H_i} b_i$$

if and only if for all $x \in K^\times$ we have

$$x^d = \prod_{i=1}^{l} \mathrm{N}_{K/K_i}(x^{b_i})^{a_i} \tag{3.2}$$

where K_i is the fixed field K^{H_i} of K with respect to $H_i < G$. We will alternatively use the term norm relation to refer to equations of the form (3.2) when it is clear from the context.

If K is abelian then we can use the results of the previous section to find a *scalar* norm relation of the form $d = \sum_{i=1}^{l} a_i N_{H_i}$ with $a_i, d \in \mathbb{Z}$ coprime, implying that for all $x \in K^\times$ we have

$$x^d = \prod_{i=1}^{l} \mathrm{N}_{K/K_i}(x)^{a_i}. \tag{3.3}$$

Again, when it is clear from the context we will simply refer to identities of the form (3.3) as a norm relation for K.

Lastly, we note that the existence of a norm relation for K implies an equivalent identity on fractional ideals of K.

Proposition 3.8. *Let $\mathfrak{a}$ be a fractional ideal of K. Then*

$$\mathfrak{a}^d = \prod_{i=1}^{l} \mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i})^{a_i} \mathcal{O}_K. \tag{3.4}$$

The proof is analogous to that of [20, Proposition 3.5]. We will make use of this to describe algorithms for resolving the PIP and ideal decomposition in Chapter 4.

3.4.1 Denominator 1

The existence of a norm relation for a number field allows studying their multiplicative properties by working in subfields, as was demonstrated in [5, 17, 47, 77]. As norm relations are multiplicative the norms of involved elements can grow rapidly, and as the relations involve d-th powers we often need to compute many d-th roots. These computational issues can be mitigated by the use of compact representations discussed in Section 2.4, however it is clear that the ideal scenario is the case $d = 1$.

Example 3.9. Let $K = \mathbb{Q}(\zeta_{63})$ with Galois group $G \cong (\mathbb{Z}/63\mathbb{Z})^\times \cong C_2^2 \times C_3^2$. This is the same group appearing in Example 3.6 and is the smallest abelian group with a denominator 1 norm relation. As the subgroups H_i appearing in the norm relation have index $[G : H_i] \leq 18$, by Galois theory we also have $[K_i : \mathbb{Q}] \leq 18$. In this particular case we have 1 subfield of degree 4, 1 of degree 9, 4 of degree 12, and 3 of degree 18, and every element of $K^\times$ can be written directly as a product of elements in these subfields.

By Theorem 3.5 we know that if K admits a denominator 1 norm relation with respect to $H_i \leq G$ then G/H_i is of the form $C \times Q$ for Q a p-group. By Galois theory $\mathrm{Gal}(K_i/\mathbb{Q}) \cong G/H_i$ hence the fixed fields K_i now admit a p-power denominator norm relation. In practice this means we typically treat the denominator 1 case as a two step process: first we can reduce to subfields $K_i \subset K$ occuring in the denominator 1 norm relation "for free" (no root computation is needed to lift information back to K), and these subfield instances can be further reduced to subfields $K_{i,j} \subset K_i$ by the existence of a nontrivial norm relation in K_i. This also implies that even in the denominator 1 case we will still need to consider the cost of root computations, albeit in subfields of smaller degree.

Example 3.10. Take K and G as defined in Example 3.9. Now take for example one of the degree 18 subfields with $\mathrm{Gal}(K_i/\mathbb{Q}) \cong C_2 \times C_3^2$. As noted in Example 3.7 this admits a denominator 3 norm relation with subgroups $H_{i,j}$ of index $[H_i : H_{i,j}] \leq 6$, hence subfields with degree $[K_{i,j} : \mathbb{Q}] \leq 6$. In fact, the norm relation on the subfields K_i have denominator either 2 or 3 and all involve subfields of degree at most 6.

3.4.2 Fields Admitting Good Norm Relations

Cyclotomic fields are a prime target for norm relation applications. They are well-studied and ubiquitous, and their Galois groups are abelian and easy to describe. Let $K = \mathbb{Q}(\zeta_m)$ be the cyclotomic field with conductor m, with Galois group $G \cong (\mathbb{Z}/m\mathbb{Z})^\times$. Write $G \cong C \times Q$ where C is the largest cyclic factor and $|Q|$ is divisible by 2 distinct primes. Then $|G| = \varphi(n)$ and $|C| = \lambda(m)$ where λ is the Carmichael lambda function. By the discussion in Section 3.4.1 we know K admits a denominator 1 norm relation with respect to subfields K_i, and each K_i admits a (nontrivial denominator) norm relation with respect to subfields $K_{i,j}$. Each $K_{i,j}$ has cyclic Galois group by Theorem 3.5 hence does not admit a norm relation, and it follows $[K_{i,j} : \mathbb{Q}] \leq \lambda(m)$ for all i, j. Hence norm relations provide the largest computational advantage in cyclotomic fields of conductor m when $\lambda(m)$ is small.

Unfortunately, the Carmichael function oscillates a lot, so there is no simple function f yielding a useful bound $\lambda(m) \leq f(m)$. We consider the case $\lambda(m) \leq \varphi(m)^a$ for $a < 3/4$. Note that by [32, Theorem 2] such values m must have negligible density. However, we see in Table 1 that in the practical range ($m < 100,000$) a significant fraction of m do satisfy $\lambda(m) < \varphi(m)^a$ for $a < 3/4$. Hence in cyclotomic fields $\mathbb{Q}(\zeta_m)$ for $m < 100,000$ a significant fraction admit a norm relation with the largest subfield having degree at most $\varphi(m)^{3/4}$.

Table 1. Proportion of conductors in the practical range admitting a good norm relation.

N	$\frac{\log(\lambda(m))}{\log(\varphi(m))} < 3/8$	$\frac{\log(\lambda(m))}{\log(\varphi(m))} < 1/2$	$\frac{\log(\lambda(m))}{\log(\varphi(m))} < 5/8$	$\frac{\log(\lambda(m))}{\log(\varphi(m))} < 3/4$
1000	0.300%	2.000%	11.70%	30.20%
10000	0.180%	1.780%	10.42%	29.45%
100000	0.092%	1.580%	8.830%	26.32%

Example 3.11 ([20, Example 5.3])**.** Let $m = 6552$. $K = \mathbb{Q}(\zeta_m)$ has degree $\varphi(m) = 1728$ and Galois group $G \cong C_{12} \times (C_6)^2 \times (C_2)^2$. The largest cyclic factor is C_{12} hence $\lambda(m) = 12$ and m satisfies $\lambda(m) < \varphi(m)^{3/8}$. K admits a denominator 1 norm relation involving 62 subfields of degree at most 192. These subfields admit norm relations with a power of 2 or 3 denominator involving a total of 672 subfields of degree at most 12.

Finally, we can construct an infinite family of conductors with small Carmichael numbers by using the following theorem of Erdös, Pomerance and Schmutz.

Theorem 3.12 (Erdös–Pomerance–Schmutz [32, Theorem 1 part 2])**.** *There exists an infinite sequence $m_1 < m_2 < \ldots$ of positive integers such that*

$$\lambda(m_k) = (\log(m_k))^{O(\log\log\log(m_k))}.$$

Note again that such integers have negligible density. Nevertheless we can easily construct such a sequence in practice, as follows. Let L be a highly divisible number (for instance, take L to be a product of a few small primes). Then let Q be the set of all primes p such that $p-1$ divides L, and let $m = \prod_{p \in Q} p$. This integer satisfies $\lambda(m) \mid L$, and the proof of Theorem 3.12 shows that for suitable choices of L, the integer m is much larger than L.

Example 3.13. We illustrate the construction by taking L to be the product of the first prime numbers.

1. $L = 2 \cdot 3 = 6$, $m = 2 \cdot 3 \cdot 7 = 42$, $\varphi(m) = 12$, $\lambda(m) = 6$.
2. $L = 2 \cdot 3 \cdot 5 = 30$, $m = 2 \cdot 3 \cdot 7 \cdot 11 \cdot 31 = 14322$, $\varphi(m) = 3600$, $\lambda(m) = 30$.
3. $L = 2 \cdot 3 \cdot 5 \cdot 7 = 210$, $m = 2 \cdot 3 \cdot 7 \cdot 11 \cdot 31 \cdot 43 \cdot 71 \cdot 211 = 9225988926$, $\varphi(m) = 2222640000$, $\lambda(m) = 210$.

We will use the existence of such an infinite sequence of conductors to show that norm relations can lead to a significant asymptotic improvement over state of the art techniques in certain families of fields. Finally, note that while conductors with $\lambda(m)$ small are sparse, almost all cyclotomic fields do admit a norm relation. For $m \neq 2, 4, p^k, 2p^k$ for $p \neq 2$ (i.e. a density 1 subset of conductors) the Galois group of $\mathbb{Q}(\zeta_m)$ is non-cyclic hence there exists a norm relation. Thus in almost all cyclotomic fields we can reduce computations to subfields of degree at least half the degree of initial field.

CHAPTER 4
RESOLUTION OF THE PRINCIPAL IDEAL PROBLEM

The known subexponential methods for solving the PIP for ideals of a number field K rely on the computation of the class group of K (see Section 2.5). The subexponential strategy for the computation of the class group of an imaginary quadratic field was described in 1989 by Hafner and McCurley [38]. The expected running time of this method is

$$L_{\Delta}(1/2, 3/\sqrt{8} + o(1)) = e^{(3/\sqrt{8}+o(1))\sqrt{|\Delta_K| \log\log|\Delta_K|}},$$

where Δ_K is the discriminant of the field assuming the Generalized Riemann hypothesis (GRH) [12]. Buchmann [22] generalized this result to the case of infinite classes of number fields with fixed degree. Practical improvements to Buchmann's algorithm were presented in [26] by Cohen, Diaz Y Diaz and Olivier. In [11, 14], Biasse and Fieker showed that there is a heuristic subexponential algorithm for the computation of the ideal class group in all classes of number fields, and that it could be used to solve the PIP. The methods of [14] can be specialized to the case of cyclotomic fields for a better asymptotic complexity [18] (heuristically in $e^{\tilde{O}(\sqrt{\log|\Delta_K|})}$). This complexity can be brought even further down (as low as $e^{\tilde{O}(\sqrt[3]{\log|\Delta_K|})}$) assuming a one-time subexponential precomputation on the field [10]).

A turning point in the development of algorithms for solving the PIP was achieved when Bauch, Bernstein, de Valence, Lange and van Vredendaal [5] showed how to recursively solve the PIP in fields of the form $K = \mathbb{Q}(\sqrt{d_1}, \ldots, \sqrt{d_n})$ (the multiquadratic fields). Depending on the family of d_i chosen, this method can have asymptotic complexity as low as polynomial in the logarithm of the discriminant of the field. This method was successfully adapted to calculation of S-unit groups and ideal class groups by Biasse and van Vredendaal [17] who proved that it had asymptotic run time in $\mathrm{Poly}(\log(|\Delta_K|))e^{\tilde{O}(\sqrt{\log|d|})}$ with $d = d_1 \cdots d_n$, under GRH and an assumption on the distribution of certain families of characters. The main idea allowing a recursive computation in the subfields of relative degree 2 was to find a norm relation implying that the square of any element in K was the product of elements coming from subfields. In another direction, the method of [5] was adjusted by Lesavourey, Plantard and Susilo [47] to the case of multicubic fields. Recent work from Biasse, Fieker, Hofmann and Page [20] generalized this concept using norm relations. See

Chapter 3 for a summary. Among other computational tasks, they showed how to leverage these relations to compute S-unit groups and ideal class groups recursively using subfields.

We use the norm relations construction technique introduced by Biasse, Fieker, Hofmann and Page [20] to solve the PIP recursively in non-cyclic number fields. This framework includes the prior work of [5] on multiquadratics and extends it to a significantly larger variety of fields. The prior work of [20] allows the recursive computation of S-unit groups from subfields, which in turn can be used to solve the PIP [19]. We describe this approach in Section 4.1 and provide an analysis of its complexity in general, as well as in specific families of fields with good norm relations from Section 3.4.2. In these fields we show it is possible to resolve the PIP asymptotically faster than the standard subexponential methods.

This S-unit based method for resolving the PIP is suboptimal in practice, however. The S-unit group computations involve many saturation steps requiring compact representations and root computations, and the primes in S are generally large compared to the input, increasing the impact of this step. In Section 4.2 we develop some technical results that will be necessary to detect and compute roots of number field elements modulo a multiplicative group $U \subseteq K^\times$. In Section 4.3 we apply this to the recursive decomposition of ideals over a set of primes using norm relations, and show this has the same asymptotic performance as the resolution of the PIP using S-units. Ideal decomposition is useful in its own right and is an important subproblem in many number-theoretic algorithms, and is needed in Chapter 5 as well. We show how this can be used for resolving the PIP with the same asymptotic performance as the S-unit approach.

We provide an alternative algorithm for the PIP in Section 4.4 which avoids computing S-unit groups and efficiently reduces the PIP to the resolution of the PIP in subfields. This can be seen as a variant of the recursive ideal decomposition of Section 4.3 optimized for resolving the PIP, focused on practical performance. We show that while has the same asymptotic performance as the S-unit approach, it behave much better in practice. We provide numerical results in Section 5.6, where we resolve the PIP in fields up to degree 1800. Finally, in Section 4.4.3 we compare our methods with the direct use of the S-unit algorithms of [20] to solve the PIP given in Section 4.1 and [19, Sec. 6], further demonstrating the difference between these approaches in practice.

Throughout this chapter let K denote a Galois number field with Galois group G admitting a norm relation with respect to $\mathcal{H} = \{H_1, \dots, H_l\}, 1 \neq H_i \leq G$. Denote by K_i the fixed field of K with respect to H_i. Then for some $a_i, b_i \in \mathbb{Z}[G]$ and $d \in \mathbb{Z}_{>0}$ we have

$$x^d = \prod_{i=1}^{l} \mathrm{N}_{K/K_i}(x^{b_i})^{a_i}. \tag{4.1}$$

for all $x \in K^\times$.

4.1 PIP With S-Units

In the presence of a norm relation, [20, Algorithm 4.16] describes an algorithm for a polynomial-time reduction of the computation of (generators of) an S-unit group of K to the computation of S-unit groups in the subfields $(K_i)_{i \leq l}$. By reframing an instance of the PIP as an S-unit group computation we get a simple, equivalent reduction for the PIP. The PIP variants of Section 4.3 and 4.4 are focused on practical improvements and have no advantage in terms of asymptotic complexity, so we use this simpler PIP variant as a backdrop for a first discussion of the complexity of our results.

The strategy of this algorithm is to enumerate short elements α of an input ideal $\mathfrak{a}$ until $\mathfrak{p} = (\alpha)/\mathfrak{a}$ is prime. As $\mathfrak{p}$ and $\mathfrak{a}$ are equivalent in the class group, $\mathfrak{p}$ is principal if and only if $\mathfrak{a}$ is. We compute the S-unit group $\mathcal{O}_{K,S}^{\times}$ where S contains all Galois conjugates of $\mathfrak{p}$, then $\mathfrak{p}$ can be decomposed over the generators of $\mathcal{O}_{K,S}^{\times}$ and we find a generator of $\mathfrak{p}$, and deduce a generator for $\mathfrak{a}$.

Algorithm 5: Solving the PIP using S-units

Input : A fractional ideal $\mathfrak{a}$ of K.
Output: A generator of $\mathfrak{a}$ if it exists, otherwise false.

$B \leftarrow$ LLL-reduced basis of $\mathfrak{a}$;
$\alpha \xleftarrow{\mathcal{R}} \mathrm{Span}(B)$;
while $\mathfrak{p} = (\alpha)/\mathfrak{a}$ *is not prime* **do**
  $\alpha \xleftarrow{\mathcal{R}} \mathrm{Span}(B)$;
end
$S \leftarrow \{\mathfrak{p}^\sigma$ for $\sigma \in \mathrm{Gal}(K/\mathbb{Q})\}$;
Compute a basis $(\alpha_1, \ldots, \alpha_{r+s})$ of $\mathcal{O}_{K,S}^{\times}$ using [20, Alg. 4.16] where $r = \mathrm{rank}\, \mathcal{O}_K^{\times}$, $s = |S|$;
Set $M \in \mathbb{Z}^{(r+s)\times s}$ such that row i is the valuations of α_i at the primes in S;
Solve $\vec{x} \cdot M = \vec{y}$ where $\vec{y}$ is the zero vector with a one at the index corresponding to $\mathfrak{p}$;
Return: $\alpha \cdot \prod_i \alpha_i^{-x_i}$

Theorem 4.1 (under GRH)**.** *Algorithm 5 is correct and has complexity*

$$\mathrm{Poly}(\log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a})), l, \max_i \log(a_i)) + l \cdot \mathrm{COST}_{S-units}(\mathrm{Subfields}),$$

where $\mathrm{COST}_{S-units}(\mathrm{Subfields})$ *denotes the cost of computing an S-unit group in a subfield.*

Proof. Let $n = [K : \mathbb{Q}]$. The correctness of the algorithm is clear, so we focus on the complexity statement. Assume $\mathfrak{a}$ has basis matrix $M \in \mathbb{Z}^{n\times n}$ with columns $(m_1, \ldots, m_n)$. An LLL-reduced basis of M can be found in time $\mathrm{Poly}(n, \log(B))$ where $\max_i \|m_i\| \leq B$ [45]. If we assume M is given in Hermite Normal Form then its entries are bounded by $|\det(M)| = \mathrm{N}(\mathfrak{a})$ and $\max_i \|m_i\| \leq B$ where $B = n^{1/2}\mathrm{N}(\mathfrak{a})$. Then $\log(B) = (1/2)\log(n) + \log(\mathrm{N}(\mathfrak{a}))$ and an LLL-reduced basis for $\mathfrak{a}$ can be found in time $\mathrm{Poly}(n, \log \mathrm{N}(\mathfrak{a}))$.

An LLL-reduced basis $(b_1, \ldots, b_n)$ is a solution to γ-SIVP with $\gamma = 2^{O(n)}$ so the basis vectors have size $2^{O(n)}\lambda_n(\mathfrak{a})$. We now estimate $\lambda_n(\mathfrak{a})$. By the properties of an LLL-reduced basis $\prod_{i=1}^{n} \lambda_i(\mathfrak{a}) \leq \prod_{i=1}^{n} \|b_i\| \leq 2^{n(n-1)/4}\,\mathrm{Vol}(\mathfrak{a})$. Recall that $\mathrm{Vol}(\mathfrak{a}) = \sqrt{|\Delta_K|}\mathrm{N}(\mathfrak{a})$. Thus we have

$$\lambda_n(\mathfrak{a}) \leq 2^{n(n-1)/4}\sqrt{|\Delta_K|}\mathrm{N}(\mathfrak{a}) \left(\prod_{i=1}^{n-1} \lambda_i(\mathfrak{a})\right)^{-1}.$$

As $\lambda_1(\mathfrak{a}) \leq \lambda_i(\mathfrak{a})$ for $1 \leq i \leq n$ we have

$$\prod_{i=1}^{n-1} \lambda_i(\mathfrak{a}) \geq \lambda_1(\mathfrak{a})^{n-1} \geq \left(\frac{1}{\mathrm{Poly}(n)}\mathrm{N}(\mathfrak{a})^{1/n}\right)^{n-1}.$$

It follows that

$$\begin{aligned}\lambda_n &\leq 2^{n(n-1)/4}\sqrt{|\Delta_K|}\,\mathrm{Poly}(n)^{n-1}\mathrm{N}(\mathfrak{a})^{1-(n-1)/n} \\ &\in 2^{\tilde{O}(n^2)}\sqrt{|\Delta_K|}\mathrm{N}(\mathfrak{a})^{o(1)}.\end{aligned}$$

so the LLL-reduced basis vectors $(b_1, \ldots, b_n)$ satisfy $\|b_i\| \in 2^{\tilde{O}(n^2)}\sqrt{|\Delta_K|}\mathrm{N}(\mathfrak{a})^{o(1)}$.

If we sample $\alpha \xleftarrow{\mathcal{R}} \mathrm{Span}(B)$ with $\alpha = \sum_{i \leq n} a_i b_i$ for $a_i \in 2^{\tilde{O}(n^2)}$ then $\|\alpha\| \in 2^{\tilde{O}(n^2)}\sqrt{|\Delta_K|}\mathrm{N}(\mathfrak{a})^{o(1)}$. As $\mathrm{N}(\alpha) \leq \|\alpha\|^n$ we have

$$\mathrm{N}(\mathfrak{p}) = \frac{\mathrm{N}(\alpha)}{\mathrm{N}(\mathfrak{a})} \leq \frac{\|\alpha\|^n}{\mathrm{N}(\mathfrak{a})} \in 2^{\tilde{O}(n^3)}|\Delta_K|^{n/2}\mathrm{N}(\mathfrak{a})^{o(n)}.$$

Due to the density of prime numbers, the number of times we expect to need to draw an element of norm $2^{\tilde{O}(n^3)}|\Delta_K|^{n/2}\mathrm{N}(\mathfrak{a})^{o(n)}$ before finding one whose norm is prime is about $\tilde{O}(n^5 \cdot \log|\Delta_K| \cdot \log(\mathrm{N}(\mathfrak{a})))$, so the cost of the loop at Step (3) is $\mathrm{Poly}(n, \log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a})))$.

By [20, Theorem 4.18] computing $\mathcal{O}_{K,S}^{\times}$ in Step (7) is reduced to the computation of S-unit groups in the subfields K_i in time $\mathrm{Poly}(n, \log|\Delta_K|, l, \max_i \log(a_i))$. This also accounts for the cost of compact representation from Algorithm 1.

□

4.1.1 Special Families of Fields

We now examine the complexity of Algorithm 5 in the context of certain families of cyclotomic fields. Recall from Section 3.4.2 that in cyclotomic fields $K = \mathbb{Q}(\zeta_m)$ of conductor m, the degree of K is given by $\varphi(m)$ and the largest cyclic factor C of $G = \mathrm{Gal}(K/\mathbb{Q})$ satisfies $|C| = \lambda(m)$ where λ is the Carmichael lambda function. By the discussion of Section 3.4.1 if K admits a denominator $d \neq 1$ norm relation then Algorithm 5 reduces to the computation of S-unit groups in subfields of degree bounded by $|C|$, and if K admits a denominator 1 norm relation then we apply the 2-step process discussed in Section 3.4.1 and once again reduce computations to subfields of degree at most $|C|$. Thus $\lambda(m)$ is an upper bound on the degree of the largest subfields occuring in Step (3) of Algorithm 5 in either case.

While this 2-step process is the most efficient use of norm relations in practice, we describe a simplification in Algorithm 6 that does not impact the asymptotic analysis. This simplification avoids the 2-step procedure, instead finding a norm relation directly to the smallest possible subfields at the cost of a potentially larger denominator.

Algorithm 6: Norm relation with minimal subfields (abelian case) [19, Algorithm 10]

Require: Non-cyclic abelian number field K with Galois group G.
Ensure: Subgroups $(H_i)_{i\leq l}$ of G, integers $(a_i)_{i\leq l}$, $d > 0$, with $d = \sum_i a_i \mathrm{N}_{H_i}$.
1: $\hat{G} \leftarrow$ dual of G, $\mathcal{H} \leftarrow \{H \leq G \text{ with } G/H \text{ cyclic}\}$.
2: **for** $H_i \in \mathcal{H}$ **do**
3: $\quad C \subseteq \hat{G} \leftarrow \langle \chi \rangle$ where $H_i = \ker(\chi)$.
4: $\quad \frac{n_i}{d_i} \leftarrow \frac{1}{|\ker \chi|} \sum_{C \leq C' \leq \hat{G} \text{ cyclic}} \mu([C' : C])$ where μ is the Möbius function.
5: **end for**
6: Find minimal $(a_i), d$ such that $d = \sum_i a_i \mathrm{N}_{H_i} \Leftrightarrow 1 = \sum_i \frac{n_i}{d_i} \mathrm{N}_{H_i}$.
7: **return** $\mathcal{H}, (a_i), d$.

Proposition 4.2 ([19, Proposition 5]). *Let K be an abelian number field of degree n and Galois group $G = C \times Q$ where C is the maximal cyclic subgroup of G and Q is non trivial. Then Algorithm 6 is correct, runs in polynomial time, and returns a norm relation $d = \sum_{i\leq l} a_i \mathrm{N}_{H_i}$ with $a_i \in \mathbb{Z}$ and that satisfies $d, l, |a_i| \leq n$, and $\max_i [K^{H_i} : \mathbb{Q}] \leq |C|$.*

Proof. The number of subgroups H_i is less than $|\hat{G}| = n$. The computation of the n_i/d_i requires the factorization of n which is polynomial in $\log|\Delta_K|$. For all subgroups H_i with cyclic quotient we have $|G/H_i| \,|\, |C|$, which proves the bound on the degrees of the K^{H_i}. Finally, the bound on (a_i) and d comes from $\mu(x) \in \{-1, 0, 1\}$ and [20, Prop. 2.26 (3)]. □

Lemma 4.3 ([19, Lemma 4]). *Let $m \geq 2$, $K = \mathbb{Q}(\zeta_m)$ and $L \subset K$ a subfield. Let $n = [K : \mathbb{Q}]$ and $n' = [L : \mathbb{Q}]$. We have*

1. $|\Delta_L| \leq |\Delta_K|^{n'/n}$.

2. $\log(n) = \log(m) + O(\log\log(m))$ *and* $\log(m) = \log(n) + O(\log\log(n))$.

3. $\log|\Delta_K| = n\left(\log(n) + O(\log\log(n))^2\right)$.

Proof. The first inequality derives directly from the fact that the discriminants satisfy $\Delta_K = \mathrm{N}_{L/\mathbb{Q}}(\Delta_{K/L})\Delta_L^{[K:L]}$ where $\Delta_{K/L}$ is the relative discriminant between K and L. Since K is a cyclotomic field, we have

$$\Delta_K = (-1)^{\varphi(m)/2}\frac{m^{\varphi(m)}}{\prod_{p|m} p^{\varphi(m)/(p-1)}},$$

so $\log|\Delta_K| = \varphi(m)\left(\log(m) - \sum_{p|m}\frac{\log(p)}{p-1}\right)$. Let $d \leq \log_2(m)$ be the number of distinct prime divisors of m:

$$\begin{aligned}
&\sum_{p|m}\frac{\log(p)}{p-1} = \sum_{p|m}\frac{\log(p)}{p} + \sum_{p|m}\frac{\log(p)}{p^2-p}\\
&= \sum_{p|m}\frac{\log(p)}{p} + O(1)\\
&\leq \sum_{k=3}^{d+2}\frac{\log(k)}{k} + O(1) \text{ since } t \mapsto \frac{\log(t)}{t} \text{ is decreasing on } [3,\infty)\\
&\leq \int_1^{\log_2(m)}\frac{\log(t)}{t}dt + O(1)\\
&= O(\log\log(m))^2.
\end{aligned}$$

Moreover we have

$$\begin{aligned}
\log(n) &= \log(\varphi(m))\\
&= \log(m) + \sum_{p|m}\log\left(1-\frac{1}{p}\right)\\
&= \log(m) - \sum_{p|m}\frac{1}{p} + O(1)\\
&= \log(m) + O(\log\log(m)) \text{ by the same argument as above,}
\end{aligned}$$

and therefore $\log(m) = \log(n) + O(\log\log(n))$. This gives

$$\log|\Delta_K| = \varphi(m)\left(\log(m) + O(\log\log(m))^2\right) = n\left(\log(n) + O(\log\log(n))^2\right),$$

as claimed. □

Proposition 4.4 (Under GRH). *Let $a > 0$ and $(m_k)_{k\in\mathbb{Z}_{>0}}$ be a sequence of integers satisfying $\lambda(m_k) \leq \varphi(m_k)^a$ for all k. Then Algorithm 5 applied to the infinite family of fields $K_k := \mathbb{Q}(\zeta_{m_k})$ has asymptotic complexity*

$$\mathrm{Poly}([K:\mathbb{Q}], \log(\mathrm{N}(\mathfrak{a}))) \cdot 2^{\tilde{O}\left([K:\mathbb{Q}]^{2a/3}\right)}.$$

Proof. Fix k and denote by n the degree $\varphi(m_k)$ of K_k and by Δ it's discriminant. Index by i the subfields occuring in the norm relation of K_k, with degree n_i and discriminant Δ_i. By the previous discussion the subfields where we compute S-unit groups have degree bounded by $\lambda(m_k) \leq \varphi(m_k)^a$. The number l of such subfields and the size of the coefficients are bounded by n by Proposition 4.2. The main term of the complexity of the subfield operations is $2^{\tilde{O}((\log|\Delta_i|)^{2/3})}$ by Lemma 2.18, and by Lemma 4.3

$$(\log|\Delta_i|)^{2/3} \leq \left(\frac{n_i}{n}\log|\Delta|\right)^{2/3} \in \tilde{O}(n_i^{2/3}) = \tilde{O}(n^{2a/3}).$$

The subfield S-unit groups are computed with respect to sets $S_i = \{\mathfrak{p} \cap \mathcal{O}_{K_i} \mid \mathfrak{p} \in S\}$, and we have $\mathrm{N}_{K_i/\mathbb{Q}}(\mathfrak{p} \cap \mathcal{O}_{K_i}) = \mathrm{N}_{K_i/\mathbb{Q}}(\mathrm{N}_{K_k/K_i}(\mathfrak{p})) = \mathrm{N}_{K_k/\mathbb{Q}}(\mathfrak{p})$. By the proof of Theorem 4.1 and Lemma 4.3

$$\begin{aligned}\mathrm{N}(\mathfrak{p}) &\in 2^{\tilde{O}(n^3)}|\Delta_K|^{n/2}\mathrm{N}(\mathfrak{a})^{o(n)} \\ &\in 2^{\tilde{O}(n^3)}\mathrm{N}(\mathfrak{a})^{o(n)}\end{aligned}$$

so we have $\log(\mathrm{N}(\mathfrak{p})) \in \tilde{O}(n^3 \log(\mathrm{N}(\mathfrak{a})))$. This proves our complexity bound. □

In Section 3.4.2 we noted that almost all cyclotomic fields admit a norm relation. Applying Proposition 4.4 it follows that in almost all cyclotomic fields Algorithm 13 provides at least a quadratic improvement over the state of the art. However, we are much more interested in determining when we have a superpolynomial improvement, in other words, resolving the PIP in time $2^{\tilde{O}(n^\epsilon)}$ for $\epsilon < 1/2$. In Table 1 we show a significant number of conductors m in the practical range satisfy $\lambda(m) < \varphi(m)^{3/4}$, so in the corresponding cyclotomic fields we achieve the desired superpolynomial improvement. Furthermore, by Theorem 3.12 there exists an infinite family of cyclotomic fields $\mathbb{Q}(\zeta_{m_k})$ admitting norm relations whose subfield degrees are bounded by $\lambda(m_k)$ where

$$\lambda(m_k) = (\log(m_k))^{O(\log\log\log(m_k))}.$$

In these fields we have the following complexity result.

Theorem 4.5 (Under GRH)**.** *There exists an infinite sequence of integers $m_1 < m_2 < \cdots$ such that Algorithm 5 has complexity*

$$\mathrm{Poly}([K:\mathbb{Q}], \log(\mathrm{N}(\mathfrak{a}))) \cdot 2^{(\log(m_k))^{O(\log\log\log(m_k))}}.$$

Proof. Take (m_k) to be the sequence from Theorem 3.12. Let $m = m_k$ be a term in this sequence, and $K_k = \mathbb{Q}(\zeta_{m_k})$ be the corresponding field. Let D be the maximum absolute value of the discriminant of a subfield used by Algorithm 17 applied to K_k. Then by Lemma 4.3 and recalling that $|\Delta_K| \leq m^n$ we have $D \leq m^{\lambda(m)}$, so that

$$\log(D) \leq \lambda(m)\log(m) = (\log(m))^{O(\log\log\log(m))}$$

by Theorem 3.12. Proceeding similarly to the proof of Proposition 4.4 we see the cost for the subfields is $2^{(\log(D))^{O(1)}} = 2^{(\log(m))^{O(\log\log\log(m))}}$. □

Let Δ_k be the discriminant of K_k. Then we have $\log(m_k) = O(\log\log|\Delta_k|)$, so that the second term of the complexity is

$$2^{(\log\log|\Delta_k|)^{O(\log\log\log\log|\Delta_k|)}}.$$

This complexity is not quite quasi-polynomial (which would correspond to $O(1)$ instead of $O(\log\log\log\log|\Delta_k|)$ in the second exponent), but it is strongly subexponential, as can be seen by rewriting it as

$$2^{(\log|\Delta_k|)^{O\left(\frac{\log\log\log|\Delta_k|\,\log\log\log|\Delta_k|}{\log\log|\Delta_k|}\right)}} = 2^{(\log|\Delta_k|)^{o(1)}}.$$

This complexity is in time $2^{n^{o(1)}}$, an impressive improvement over the state of the art. Note again, however, that conductors where we have such an improvement have negligible density and are exceedingly uncommon.

In practice, relying on the computation of S-unit groups as in Algorithm 5 is suboptimal. Recovering the full S-unit group $\mathcal{O}_{K,S}^{\times}$ from subfield S-unit groups potentially requires many root computations in K (or intermediate subfields in the case of denominator 1). This is mitigated by the use of compact representations in theory as discussed in Section 2.4 but it is still desirable to avoid root computations if possible: while they only incur extra polynomial factors in the overall asymptotic complexity they can be a major concern in practice, quickly making larger instances intractable. In Section 4.3 we describe how to take advantage of norm relations to perform recursive ideal decomposition, and in Section 4.4 we go to great lengths to specialize this to the PIP, and remove many of the root computations by avoiding the computation of an S-unit group or even the unit group in K altogether.

4.2 Interlude: Powers Modulo Multiplicative Groups

Recall that by Proposition 3.8 the existence of a norm relation of the form (4.1) for K implies that for a fractional ideal $\mathfrak{a}$ of K the following also holds:

$$\mathfrak{a}^d = \prod_{i=1}^{l} \mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i})^{a_i}\mathcal{O}_K.$$

Assume that we have generators α_i of $\mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i})^{a_i}\mathcal{O}_K$. Then we can deduce a generator $\mathfrak{a}^d = \alpha\mathcal{O}_K$ where $\alpha = \prod_{i=1}^{l}\alpha_i$. If $\mathfrak{a}$ is principal and generated by β then β^d is a generator for $\mathfrak{a}^d$, and we must have $u \in \mathcal{O}_K^\times$ such that $u\alpha = \beta^d$. Thus to recover a generator of $\mathfrak{a}$ from a generator of $\mathfrak{a}^d$ we must find a suitable unit $u \in \mathcal{O}_K^\times$ such that $u\alpha$ is a d-th power. In the case of ideal decomposition we run into a similar but more general problem. We formalize this in Definition 4.6.

Definition 4.6. Let $U \subseteq K^\times$ be a multiplicative group and $\beta \in K^\times$. We say that β is a *d-th power modulo U* or that β *has a d-th root modulo U*, if there exists $u \in U$ such that $u\beta$ is a d-th power, i.e., $u\beta \in (K^\times)^d$.

Let $U \subseteq K^\times$ be a finitely generated multiplicative group and $\beta \in K^\times$. Throughout this section we assume that $U \cap \langle\beta\rangle = \{1\}$. We now investigate how to detect if β is a d-th power modulo U, and if so, compute $u \in U$ such that $u\beta$ is a d-th power. To this end we develop a variation on the saturation techniques described in Section 2.3. First, we show that by Lemma 4.7 it is sufficient to work with d a prime power, so we will assume this is the case.

Lemma 4.7. *Assume that $U \subseteq K^\times$ is a multiplicative group and $d = a \cdot b$ with $\gcd(a, b) = 1$. Then β is a d-th power modulo U if and only if β is an a-th and a b-th power modulo U.*

Proof. Since one of the implications is trivial, let us assume that β is an a-th and a b-th power modulo U, say $\beta = u\gamma^a = u_0\gamma_0^b$. Since a and b are coprime there exist $r, s \in \mathbb{Z}$ with $1 = ra + sb$. Thus

$$\beta = \beta^{ra}\beta^{sb} = (u_0\gamma_0^b)^{ra}(u\gamma^a)^{sb} = u_0^{ra}u^{sb}(\gamma_0^r\gamma^s)^d \in U \cdot (K^\times)^d.$$

□

4.2.1 The Good Case

First we assume that we are in the good case of Grunwald–Wang. Recall that for a non-zero prime ideal $\mathfrak{p}$ of $\mathcal{O}_K$ we denote by $K_\mathfrak{p}$ the $\mathfrak{p}$-adic completion of K, by $v_\mathfrak{p}$ the $\mathfrak{p}$-adic valuation and by $k_\mathfrak{p} = \mathcal{O}_K/\mathfrak{p} \cong \mathcal{O}_{K_\mathfrak{p}}/\mathfrak{p}\mathcal{O}_{K_\mathfrak{p}}$ the residue field at $\mathfrak{p}$. We use a bar notation to denote cosets of various multiplicative groups, and $\langle X \rangle$ to denote the subgroup generated by X. We have the following criterion for detecting whether β is a d-th power modulo U, and if so, computing $u \in U$ such that $u\beta \in (K^\times)^d$. This is an adaptation of Proposition 2.9.

Proposition 4.8. *Let $V = \langle U, \beta \rangle$ where $U = \langle \alpha_1, \dots, \alpha_l \rangle$, and assume $U \cap \langle \beta \rangle = \{1\}$, d is a prime power, and that we are in the good case of Grunwald–Wang. Furthermore let $c \in \mathbb{R}_{>0}$ be arbitrary. Assume that the intersection*

$$\bigcap_{d \nmid \mathfrak{p}, \mathrm{N}(\mathfrak{p}) \leq c} \ker(V/V^d \to \mathbb{Z}/d\mathbb{Z} \times k_\mathfrak{p}^\times/(k_\mathfrak{p}^\times)^d) \subseteq V/V^d$$

is generated by the classes of $\alpha_1\beta^{n_1}, \dots, \alpha_l\beta^{n_l} \in V$ with $\alpha_i \in U$, $n_i \in \mathbb{Z}$.

1. *If $\gcd(d, n_1, \dots, n_l) \neq 1$, then β is not a d-th power modulo U.*
2. *Assume β is not a d-th power modulo U. Then for c sufficiently large we have $\gcd(d, n_1, \dots, n_l) \neq 1$.*
3. *Assume β is a d-th power modulo U. Then for c sufficiently large we have $\gcd(d, n_1, \dots, n_l) = 1$ and that the element $\alpha_1^{k_1} \cdots \alpha_l^{k_l}\beta$ is a d-th power, where $k_i \in \mathbb{Z}$ are integers with $1 = k_0 d + \sum_{i=1}^l k_i n_i$.*

Proof. Let us denote by W/V^d the intersection of the kernels.

(1): Assume that β is a d-th power modulo U, that is, $\alpha\beta \in V \cap (K^\times)^d$ for some $\alpha \in U$. As $(V \cap (K^\times)^d)/V^d \subseteq W/V^d$, there exist integers $0 < k_i < d$ such that

$$\overline{\alpha\beta} = \overline{(\alpha_1\beta^{n_1})^{k_1} \cdots (\alpha_l\beta^{n_l})^{k_l}}$$

in $W/V^d \subseteq V/V^d$. As V is generated by U and β, the group V^d is generated by U^d and β^d. Hence there exists $\alpha_0 \in U$ and $k_0 \in \mathbb{Z}$ such that

$$\alpha\beta = (\alpha_1\beta^{n_1})^{k_1} \cdots (\alpha_l\beta^{n_l})^{k_l} \alpha_0^d (\beta^d)^{k_0}.$$

From $U \cap \langle \beta \rangle = \{1\}$ we get $1 = k_0 d + \sum_{i=1}^l k_i n_i$ i.e. $\gcd(d, n_1, \dots, n_l) = 1$.

(2): Let c_0 be the constant from Proposition 2.8 and assume $c \geq c_0$. In particular it holds $(V \cap (K^\times)^d)/V^d = W/V^d$. Assume $\gcd(d, n_1, \dots, n_l) = 1$. Then there exist $k_i \in \mathbb{Z}$, $0 \leq i \leq l$, such that

$1 = k_0 d + \sum_{i=1}^{l} k_i n_i$. Then the element $\alpha = \alpha_1^{k_1} \cdots \alpha_l^{k_l}$ satisfies

$$\alpha\beta = \alpha\beta^{n_1 k_1} \cdots \beta^{n_l k_l} \beta^{d k_0} = (\alpha_1 \beta^{n_1})^{k_1} \cdots (\alpha_l \beta^{n_l})^{k_l} \beta^{d k_0},$$

that is $\overline{\alpha\beta} \in W/V^d = (V \cap (K^\times)^d)/V^d$ and β is a d-th power modulo U.

(3): Let c_0 be as in Proposition 2.8 and assume $c \geq c_0$. Note that as β is a d-th power modulo U, it follows from (1) that $\gcd(d, n_1, \ldots, n_l) = 1$. The result follows, since

$$\alpha^{k_1} \cdots \alpha^{k_l} \beta = (\alpha_1 \beta^{n_1})^{k_1} \cdots (\alpha_l \beta^{n_l})^{k_l} (\beta^{k_0})^d$$

and for all $1 \leq i \leq l$ we have $\alpha_i \beta^{n_i} \in (K^\times)^d$ (as $c \geq c_0$). □

Algorithm 7 decides whether an element β is a a d-th power modulo U, and if so, it finds an element of u such that $u\beta$ is a d-th power. Proposition 4.8 directly shows its correctness.

Algorithm 7: d-th power modulo units in the good case

Input : $U \subseteq K^\times$ finitely generated, $\beta \in K^\times$ such that $U \cap \langle\beta\rangle = \{1\}$, $d = p^r$ a prime power, such that we are in the good case of Grunwald–Wang

Output: Whether β is a d-th power modulo U and an element $\gamma \in K^\times$ with $\beta/\gamma^d \in U$ in case it exists

1 Let $c \in \mathbb{R}_{>0}$ (chosen arbitrarily);

2 Determine a $(\mathbb{Z}/d\mathbb{Z})$-generating set $\overline{\alpha_1\beta^{n_1}}, \ldots, \overline{\alpha_l\beta^{n_l}}$ of

$$\bigcap_{p \notin \mathfrak{p}, \mathrm{N}(\mathfrak{p}) \leq c} \ker(\langle U, \beta\rangle / \langle U, \beta\rangle^d \to \mathbb{Z}/d\mathbb{Z} \times k_\mathfrak{p}^\times / (k_\mathfrak{p}^\times)^d);$$

if $\gcd(d, n_1, \ldots, n_l) \neq 1$ **then**

3 | **return:** β is not a d-th power modulo U;

4 **else if** $\gcd(d, n_1, \ldots, n_l) = 1$ **then**

5 | Determine $k, k_i \in \mathbb{Z}$, $1 \leq i \leq l$, with $1 = kd + \sum_{i=1}^{l} k_i n_i$;

6 | Test whether the element $\delta = \beta \prod_{i=1}^{l} \alpha_i^{k_i}$ is a d-th power;

7 | **if** *there exists* γ *with* $\gamma^d = \delta$ **then**

8 | | **return:** γ;

9 | **end**

10 Replace c by $2c$ and go to step 2;

4.2.2 The General Case

We now assume that we are in the general case and $d = p^r$ is a prime power. Since this includes the bad case of Grunwald–Wang, in general we cannot detect global d-th powers just using local information. The algorithms are therefore more complicated; the reader can skip this section without significantly affecting their understanding. We state results for an arbitrary p, but for us the only relevant case is $p = 2$. In general, we detect d-th powers for $d = 2^r d'$ with $2 \nmid d'$ by detecting d'-th powers using Algorithm 7, detecting 2^r-th powers using results of this section, and recombining the results using Lemma 4.7. We investigate the situation where U is p-saturated.

Under the GRH, when $c > c_0 = 72d^2(\log|\Delta_K| + 3n\log(p))^2$, iterating [20, Algorithm 4.9] correctly returns the p-saturation of $U \subseteq K^\times$ in polynomial time [20, Theorem 4.11]. We now assume that U is p-saturated, and we show that testing whether β is a p^r-th power modulo U can be reduced to r instances of the problem where the exponent is p (instead of p^r), hence to a situation where we are in the good case of Grunwald–Wang.

Proposition 4.9. *Assume that $U \subseteq K^\times$ is a multiplicative group and $\beta \in K^\times$ a p^r-th power modulo U. Then the following hold:*

1. *The element β is a p^i-th power modulo U for all $1 \leq i \leq r$.*
2. *Assume that U is p-saturated and that there exist $u \in U$, $\gamma_i \in K^\times$ with $u\beta = \gamma_i^{p^i}$ for some $1 \leq i \leq r-1$. Then γ_i is a p^{r-i}-th power modulo U.*

Proof. (1): Trivial. For (2), first note that by assumption there exist $\tilde{u} \in U$, $\gamma \in K^\times$ such that $\tilde{u}\beta = \gamma^{p^r}$. Then

$$\gamma_i^{p^i} = u\beta = \frac{u}{\tilde{u}}\tilde{u}\beta = \frac{u}{\tilde{u}}\gamma^{p^r}.$$

Hence

$$\frac{u}{\tilde{u}} = \frac{\gamma_i^{p^i}}{\gamma^{p^r}} = \left(\frac{\gamma_i}{\gamma^{p^{r-i}}}\right)^{p^i} \in U.$$

As U is p-saturated this implies $\gamma_i/\gamma^{p^{r-i}} \in U$. Thus $(\gamma^{p^{r-i}}/\gamma_i)\gamma_i = \gamma^{p^{r-i}}$ shows that γ_i is a p^{r-i}-th power modulo U. □

Corollary 4.10. *Assume that $U \subseteq K^\times$ is a p-saturated multiplicative group. An element $\beta \in K^\times$ is a p^r-th power modulo U if and only if there exist $u_1, \ldots, u_{r-1} \in U$, $\gamma_1, \ldots, \gamma_r \in K^\times$, $\gamma_1 = \beta$ such that $\gamma_{i+1}^p = u_i\gamma_i$ for all $1 \leq i \leq r-1$.*

Therefore, under the assumption that U is p-saturated, we can check whether β is a p^r-th power modulo U by iteratively checking whether certain elements are p-th powers modulo U. As p is a prime, we are

always in the good case of the Grunwald–Wang theorem and we can use Algorithm 7. We summarize this in Algorithm 8, which is correct according to Corollary 4.10. Finally, note that in Step (4) of Algorithm 8 we must compute γ_i as a p-th root of an element of $K^\times$. This is the primary distinction between the good and bad case in practice: in the good case we can determine if $\beta \in K^\times$ is a d-th power modulo U without any root computation if we take $c > c_0$, or with a single root computation otherwise. In the bad case when $d = p^r$ we may be required to do up to r root computations.

Algorithm 8: d-th power in the bad case

Input : A p-saturated multiplicative group $U \subseteq K^\times$, $r \geq 1$, and $\beta \in K^\times$ with $U \cap \langle \beta \rangle = \{1\}$
Output: Whether β is a p^r-th power modulo U and an element $\gamma \in K^\times$ with $\beta/\gamma^{p^r} \in U$ in case it exists

$\gamma_0 \leftarrow \beta$;
for $i \leftarrow 1$ **to** r **do**
  if γ_{i-1} *is a p-th power modulo* U *using Algorithm 7* **then**
    Compute $\gamma_i \in K^\times$ such that $\gamma_{i-1}/\gamma_i^p \in U$;
  else
    return: that β is not a p^r-th power modulo U;
  end
end
return: γ_r;

4.3 Ideal Decomposition

Given an input ideal $\mathfrak{a}$ whose ideal class in $\mathrm{Cl}(\mathcal{O}_K)$ is known to be a product of powers of the classes of $\mathfrak{g}_1, \ldots, \mathfrak{g}_k$, the task of finding exponents such that $\mathfrak{a} \sim \prod_i \mathfrak{g}_i^{x_i}$ is a core subroutine in number theory referred to as ideal decomposition. As discussed in Section 2.5 the best classical algorithms for ideal class decomposition have the same asymptotic cost as the computation of $\mathrm{Cl}(\mathcal{O}_K)$, which is subexponential. In this section, we show how to leverage norm relations to reduce the decomposition of the class of an input ideal $\mathfrak{a} \subseteq K$ to subfield computations. Again, we assume that K admits a norm relation with subfields $(K_i)_{i \leq l}$ so that ideals of K satisfy Equation 3.4, i.e $\mathfrak{a}^d = \prod_{i=1}^{l} \mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i})^{a_i}\mathcal{O}_K$. We first tackle the case of the decomposition of $\mathfrak{a}$ according to a set S of prime ideals invariant under the action of $\mathrm{Gal}(K/\mathbb{Q})$. Then, given a subgroup $H \subseteq \mathrm{Cl}(\mathcal{O}_K)$, and generators $\mathfrak{g}_1, \ldots, \mathfrak{g}_k$ of order $d_1, \ldots, d_k$ such that $H \simeq \langle \mathfrak{g}_1 \rangle \times \ldots \times \langle \mathfrak{g}_k \rangle$, we show how to find the unique $(x_1, \ldots, x_k) \in \mathbb{Z}/d_1\mathbb{Z} \times \ldots \times \mathbb{Z}/d_k\mathbb{Z}$ such that $\mathfrak{a} \sim \prod_i \mathfrak{g}_i^{x_i}$.

4.3.1 Decomposition With Respect to Primes

Given a set $S = \{\mathfrak{p}_i\}_{i \leq k}$ of non-zero prime ideals of $\mathcal{O}_K$ and an ideal $\mathfrak{a} \subseteq \mathcal{O}_K$, our goal is to find $\vec{x}$ such that $\mathfrak{a} \sim \prod_i \mathfrak{p}_i^{x_i}$. We assume that S is stable under the action of $G = \mathrm{Gal}(K/\mathbb{Q})$, so $\sigma(\mathfrak{p}) \in S$ for any $\sigma \in G, \mathfrak{p} \in S$. Let $\langle S \rangle \subseteq \mathrm{Cl}(\mathcal{O}_K)$ be the subgroup of $\mathrm{Cl}(\mathcal{O}_K)$ generated by the classes of the elements of S. Then, one can recursively find the decomposition of $\mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i})$ in $\mathrm{Cl}(\mathcal{O}_{K_i})$ with respect to the $\mathfrak{p} \cap K_i$ for $\mathfrak{p} \in S$, and deduce the decomposition of $\mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i})^{a_i}\mathcal{O}_K$ in $\mathrm{Cl}(\mathcal{O}_K)$ with respect to S (each $(\mathfrak{p} \cap K_i)\,\mathcal{O}_K$ is a product of elements of S since it is assumed to be stable under the action of G).

At this point, we have a decomposition of $\mathfrak{a}^d$ instead of $\mathfrak{a}$ as desired. In other words, we have $\vec{x}$ such that $\mathfrak{a}^d \sim \prod_i \mathfrak{p}_i^{x_i}$. This information alone is not enough to decompose $\mathfrak{a}$ with respect to the $\mathfrak{p}_i$ in $\mathrm{Cl}(\mathcal{O}_K)$. In particular, we need to use a generator of the principal ideal $\mathfrak{a}^d \prod_i \mathfrak{p}_i^{-x_i}$. To get this information, in each subfield, we can make sure that we obtain an identity of the form

$$\mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i})^{a_i}\mathcal{O}_K = (\alpha_i)\prod_j \mathfrak{p}_j^{x_{i,j}}.$$

This can be done by working exclusively in the subfields (an ideal class decomposition in $\mathrm{Cl}(\mathcal{O}_{K_i})$ followed by solving a PIP in K_i). By recombining all subfield information, we obtain an identity of the form

$$\mathfrak{a}^d = (\alpha)\prod_i \mathfrak{p}_i^{y_i}, \tag{4.2}$$

where $\vec{y} \in \mathbb{Z}^k$ and $\alpha \in K$ is given in product form. We summarize this procedure in Algorithm 9.

Algorithm 9: Decomposition of $\mathfrak{a}^d$ with norm relation

Require: Number field K of unit rank r, norm relation $d = \sum_i a_i \mathrm{N}_{H_i} b_i$, ideal $\mathfrak{a}$ and set S of k prime ideals stable under the action of $G = \mathrm{Gal}(K/\mathbb{Q})$, and with $[\mathfrak{a}] \in \langle S \rangle$.
Ensure: $\vec{y}$ and α such that $\mathfrak{a}^d = (\alpha)\prod_i \mathfrak{p}_i^{y_i}$.
1: Compute a basis $(\beta_i)_{i \leq r+k}$ for the S-unit group (using recursive norm relation techniques), and let $M \in \mathbb{Z}^{(r+k)\times k}$ such that $(\beta_i) = \prod_j \mathfrak{p}_j^{M_{i,j}}$.
2: **for** K_i in the norm relation **do**
3: Compute $\alpha_i, \vec{x}_i$ such that $\mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i})^{a_i}\mathcal{O}_K = (\alpha_i)\prod_j (\mathfrak{p}_j \cap K_i)^{x_{i,j}}$.
4: **end for**
5: Deduce α in product form and $\vec{y}$ such that $\mathfrak{a}^d = (\alpha)\prod_i \mathfrak{p}_i^{y_i}$.
6: **return** $\alpha, \vec{y}$.

Theorem 4.11 (under GRH)**.** *Algorithm 9 is correct and has complexity*

$$\mathrm{Poly}(\log|\Delta_K|, \log \mathrm{N}(\mathfrak{a}), l, \max_i \log a_i) + l \cdot \left(\mathrm{Cost}_{S\text{-}unit}(\textit{subfields}) + \mathrm{Cost}_{Ideal\ Dec}(\textit{subfields})\right),$$

where $\mathrm{Cost}_{S\text{-}unit}(subfields)$ *is the cost of computing* S*-units in the subfields involved in the norm relation* (3.4)*, and* $\mathrm{Cost}_{Ideal\ Dec}(subfields)$ *is the cost of ideal decomposition in the subfields involved in the norm relation* (3.4)*.*

Proof. The cost of computing $\mathcal{O}_{K,S}^{\times}$ in Step (1) is reduced to the computation of S-unit groups in the subfields K_i by [20, Theorem 4.18] in time $\mathrm{Poly}(\log|\Delta_K|, l, \max_i \log(a_i))$. The loop at Step (2) involves an ideal decomposition in each subfield as well, giving us the result. □

Now we also know that since the class of $\mathfrak{a}$ is a product of the classes of S, there must exist $\vec{z} \in \mathbb{Z}^k$ and $\beta \in K$ such that $\mathfrak{a} = (\beta)\prod_i \mathfrak{p}_i^{z_i}$, which means that

$$\mathfrak{a}^d = (\beta^d)\prod_i \mathfrak{p}_i^{dz_i}.$$

Now we have the equality of ideals $(\beta^d)\prod_i \mathfrak{p}_i^{dz_i} = (\alpha)\prod_i \mathfrak{p}_i^{y_i}$, but since α is not necessarily β^d, we don't necessarily have $y_i = dz_i$. However, we know that $\prod_i \mathfrak{p}_i^{y_i} \sim \mathfrak{p}_i^{dz_i}$ so we must have $\vec{y} - d\vec{z} \in \mathcal{L}$ where $\mathcal{L} \subseteq \mathbb{Z}^k$ is the lattice of relations between the $\mathfrak{p}_i$, i.e. the lattice of vectors $\vec{u}$ such that $\prod_i \mathfrak{p}_i^{u_i}$ is a principal ideal. We want to re-write α as $\alpha = \beta'^d \cdot \delta$ where δ is an S-unit with $(\delta)\mathcal{O}_K = \prod_i \mathfrak{p}_i^{u_i}$ such that $\vec{u} + \vec{y} \in d\mathbb{Z}^k$. If this is the case, then $\mathfrak{a}^d = (\beta'^d)\prod_i \mathfrak{p}_i^{dz_i'}$ where $\vec{z'} := \vec{u} + \vec{y}$. Once an S-unit δ_0 such that $(\delta_0) = \prod_i \mathfrak{p}_i^{u_i^0}$ with $\vec{u}^{(0)} + \vec{y} \in d\mathbb{Z}^k$ is found, then any other solution δ is of the form $\delta = \delta_0\delta'$ where δ' is an S-unit satisfying $(\delta')\mathcal{O}_K = \prod_i \mathfrak{p}_i^{u_i'}$ with $\vec{u'} \in d\mathbb{Z}^k$. The set of such δ' is a subgroup of the S-unit group

By [20, Algorithm 4.16] we can compute generators $\alpha_1, \ldots, \alpha_{r+k+1}$ of the S-unit group from subfield computations (where r is the rank of the unit group), together with a matrix $M \in \mathbb{Z}^{(r+k+1)\times k}$ whose rows are the valuations of the α_i according to the primes in S. Thus, there is $\vec{x} \in \mathbb{Z}^{r+k+1}$ such that $\vec{y} = \vec{x}M + d\vec{z}$, i.e.

$$\vec{y} = \vec{x}M \bmod d.$$

This system does not have a unique solution. However, we can put M in row reduced echelon form modulo d and find

1. a solution $\vec{x}^{(0)}$ to $\vec{y} = \vec{x}M \bmod d$,
2. a basis $\vec{x}^{(1)}, \ldots, \vec{x}^{(m)}$ of the left kernel of M mod d.

So all the $\vec{x}$ such that $\vec{y} = \vec{x}M \bmod d$ are of the form $\vec{x} = \vec{x}^{(0)} + \sum_j a_j\vec{x}^{(j)}$, including the one that satisfies $\vec{y} = \vec{x}M + d\vec{z}$ for $\vec{z}$ defined above. We denote by $\vec{x}^{(j)}M_i$ the i-th coefficient of $\vec{x}^{(j)}M$, and by $\alpha_i \in K$ the element that satisfies $\prod_j \mathfrak{p}_j^{M_{i,j}} = (\alpha_i)\mathcal{O}_K$. With the notation previously used, $\delta_0 = \prod_i \alpha_i^{x_i^{(0)}}$, while the subgroup of δ''s is generated by $\delta_i := \prod_i \alpha_i^{x_i^{(j)}}$ for $i = 1, \ldots, m$. Therefore, we have

$$
\begin{aligned}
(\alpha)\prod_i \mathfrak{p}_i^{y_i} &= (\alpha)\prod_i \mathfrak{p}_i^{\vec{x}M_i}\cdot\prod_i \mathfrak{p}_i^{y_i-\vec{x}M_i}\\
&= (\alpha)\prod_i \mathfrak{p}_i^{\vec{x}^{(0)}M_i}\cdot\prod_{j\le m}\left[\prod_i \mathfrak{p}_i^{\vec{x}^{(j)}M_i}\right]^{a_j}\cdot\prod_i \mathfrak{p}_i^{y_i-\vec{x}M_i}\\
&= (\alpha)\left(\prod_i \alpha_i^{x_i^{(0)}}\right)\cdot\left(\prod_{j\le m}\left[\prod_i \alpha_i^{x_i^{(j)}}\right]^{a_j}\right)\cdot\prod_i \mathfrak{p}_i^{y_i-\vec{x}M_i}\\
&= (\alpha')\left(\prod_j \delta_j^{a_j}\right)\cdot\prod_i \mathfrak{p}_i^{dz_i'} \quad \text{for some } z_i'\in\mathbb{Z}
\end{aligned}
$$

where we have a product representation of $\alpha' \in K$ and the $\delta_j \in K$.

So we are looking for $(a_j)_{i\le m}$ such that $\alpha'\cdot\prod_j \delta_j^{a_j} = \beta'^d$ for some $\beta' \in K$. Once we find $(a_j)_{i\le m}$, we derive the corresponding $\vec{x} = \vec{x}^{(0)} + \sum_{j\le m} a_j\vec{x}^{(j)}$ and then $\vec{z'} = \frac{1}{d}(\vec{y} - \vec{x}M)$. This means that we have the identity

$$\mathfrak{a}^d = (\beta'^d)\prod_i \mathfrak{p}_i^{dz_i'}.$$

Such an identity exists at least for $\beta' = \beta$ and $z_i' = z_i$ (with the notation above), but other choices of $(a_i)_{i\le m}$ might lead to other solutions. Once a solution is found, we have $\mathfrak{a} = (\beta')\prod_i \mathfrak{p}_i^{z_i'}$ since an equality of fractional ideals of the form $I^d = J^d$ implies that $I = J$ by uniqueness of prime decomposition. Thus, we are able to conclude that $\mathfrak{a} \sim \prod_i \mathfrak{p}_i^{z_i'}$, which solves the ideal class decomposition problem.

Now the question is how to find the desired $(a_i)_{i\le m}$? Since there is a solution, we know that α' is a d-th power modulo U for $U = \langle\delta_1,\dots,\delta_m\rangle$. By Proposition 4.8 we have $\gcd(d, n_1, \dots, n_m) = 1$, where the n_i are the exponents defined in Proposition 4.8. Let k, and $(a_i)_{i\le m}$ such that $1 = kd + \sum_{i=1}^m a_i n_i$. With this choice of a_i we have that $\alpha'\prod_i \delta_i^{a_i}$ is a d-th power and we can find the decomposition of $\mathfrak{a}$ in $\mathrm{Cl}(\mathcal{O}_K)$. Finally, recall from Section 2.4 that taking the d-th root of $\alpha'\prod_i \delta_i^{a_i}$ can be done efficiently by keeping elements in compact representation. We summarize the procedure in Algorithm 10. While we discussed how to compute roots modulo U in both the good and bad case of Grunwald–Wang, we restrict to the good case here for simplicity. Then Algorithm 10 can be easily extended to the general case.

Theorem 4.12 (under GRH)**.** *Algorithm 10 is correct and has complexity*

$$\mathrm{Poly}(\log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a})), l, \max_i \log(a_i)) + l\cdot \mathrm{COST}_{S-units}(\mathrm{Subfields}),$$

where $\mathrm{COST}_{S-units}(\mathrm{Subfields})$ *denotes the cost of computing an S-unit group in a subfield.*

Algorithm 10: Ideal decomposition from subfields in the good case where $d = p^t$

Require: Number field K of unit rank r, norm relation $d = \sum_i a_i \mathrm{N}_{H_i} b_i$ where d is a prime power in the good case of Grunwald–Wang, ideal $\mathfrak{a}$ and set S of k primes stable under the action of $G = \mathrm{Gal}(K/\mathbb{Q})$, together with $\alpha, \vec{y}$ such that $\mathfrak{a}^d = (\alpha) \prod_i \mathfrak{p}_i^{y_i}$.

Ensure: $\beta', \vec{z'}$ such that $\mathfrak{a} = (\beta') \prod_i \mathfrak{p}_i^{z_i'}$.

1: Compute a basis $(\alpha_i)_{i \le k+r+1}$ for the S-unit group (using recursive norm relation techniques), and let $M \in \mathbb{Z}^{(r+k+1)\times k}$ such that $(\alpha_i) = \prod_j \mathfrak{p}_j^{M_{i,j}}$.

2: Put M in row reduced echelon form mod$d$. Find $\vec{x}^{(0)}$ solution to $\vec{y} = \vec{x}M \bmod d$.

3: Compute $\vec{x}^{(1)}, \ldots, \vec{x}^{(m)}$ basis of the left kernel of M mod d.

4: $\alpha' \leftarrow (\alpha) \left(\prod_i \alpha_i^{x_i^{(0)}} \right)$. For $j \le m$: $\delta_j \leftarrow \prod_i \alpha_i^{x_i^{(j)}}$.

5: $U \leftarrow \langle \delta_1, \ldots, \delta_m \rangle$. Let $c \le c_0$ large enough.

6: Compute a $(\mathbb{Z}/d\mathbb{Z})$-generating set $\overline{\delta_1 \alpha'^{n_1}}, \ldots, \overline{\delta_m \alpha'^{n_m}}$ of

$$\bigcap_{p \notin \mathfrak{p}, \mathrm{N}(\mathfrak{p}) \le c} \ker(\langle U, \alpha' \rangle / \langle U, \alpha' \rangle^d \to \mathbb{Z}/d\mathbb{Z} \times k_{\mathfrak{p}}^{\times} / (k_{\mathfrak{p}}^{\times})^d).$$

7: Compute $k, a_i \in \mathbb{Z}$, $1 \le i \le m$, with $1 = kd + \sum_{i=1}^m a_i n_i$. Let $\vec{x} \leftarrow \vec{x}^{(0)} + \sum_{j \le m} a_j \vec{x}^{(j)}$.

8: **return** $\sqrt[d]{\alpha' \cdot \prod_j \delta_j^{a_j}}, \frac{1}{d}(\vec{y} - \vec{x}M)$.

Proof. Just as in Theorem 4.11 the cost is reduced to the cost of subfield S-unit group computation in time $\mathrm{Poly}(\log|\Delta_K|, l, \max_i \log(a_i))$, and the result follows. □

The following cases need extra care:

1. The case where α is an S-unit (which leads to $U \cap \langle \alpha' \rangle \neq \{1\}$ for $U = \langle \delta_1, \ldots, \delta_m \rangle$).
2. The case of d not a prime power.
3. The bad case of Grunwald–Wang.

Case (1) can be easily avoided by replacing α by $\alpha \cdot x^d$ where x is outside of the S-unit group. The procedure will succeed, and lead to the computation of appropriate exponents $a_1, \ldots, a_m$. For Case (2), assume that d is not a prime power. We rely on the following lemma

Lemma 4.13. *Let a, b be coprime integers such that $d = ab$. Assume that with Algorithm 10 we can find $\beta_x, \vec{x}, \beta_y, \vec{y}$ such that*

$$\mathfrak{a}^d = (\beta_x^a) \prod_i \mathfrak{p}_i^{a x_i} = (\beta_y^b) \prod_i \mathfrak{p}_i^{b y_i}.$$

Then $\mathfrak{a} = (\beta_x^s \beta_y^r) \prod_i \mathfrak{p}_i^{s x_i + r y_i}$.

Proof. Let r, s be such that $ra + sb = 1$. This means that

$$\begin{aligned}\mathfrak{a}^d &= \left(\mathfrak{a}^d\right)^{ra+sb} = \left(\mathfrak{a}^d\right)^{ra}\left(\mathfrak{a}^d\right)^{sb} \\ &= \left((\beta_y)\prod_i \mathfrak{p}_i^{y_i}\right)^{rab}\left((\beta_x)\prod_i \mathfrak{p}_i^{x_i}\right)^{sab} \\ &= \left((\beta_x^s\beta_y^r)\prod_i \mathfrak{p}_i^{sx_i+ry_i}\right)^d.\end{aligned}$$

Therefore, by equality of ideals, we have a $\beta, \vec{z}$ such that $\mathfrak{a} = (\beta)\prod_i \mathfrak{p}_i^{z_i}$. □

This process can be iterated for all prime powers that divide d, thus reducing the case of arbitrary d to that of d being a prime power.

Finally, Case (3) concerns the bad case of Grunwald–Wang. Because of the above consideration, we can assume that d is a prime power, and since the bad case only concerns powers of two, there is t such that $d = 2^t$. Algorithm 10 cannot be applied directly on input d, but we can use it with denominator 2. This leads to the creation of $\beta', \vec{z'}$ such that

$$\mathfrak{a}^{2^{t-1}} = (\beta')\prod_i \mathfrak{p}_i^{z_i'}.$$

This can be iterated t times, yielding the decomposition of $\mathfrak{a}$.

4.3.2 Decomposition With Respect to Elementary Generators

Let H be a subgroup of $\mathrm{Cl}(\mathcal{O}_K)$. In the previous section, we established how to decompose the class of an input ideal $\mathfrak{a}$ with respect to a given set of primes S (if this decomposition exists, which is always the case when we pick S a generating set of $\mathrm{Cl}(\mathcal{O}_K)$). For certain uses this is sufficient. For example, if we know that the primes in S generate the class group then we can use this for solving instances of the PIP, which we detail in 12. However, in Section 5.3 we wish to instead find a decomposition of the class of $\mathfrak{a}$ according to a fixed set of generators $\mathfrak{g}_1, \ldots, \mathfrak{g}_k$ where

$$H \simeq \langle[\mathfrak{g}_1]\rangle \times \cdots \times \langle[\mathfrak{g}_k]\rangle \simeq \mathbb{Z}/d_1\mathbb{Z} \times \cdots \times \mathbb{Z}/d_k\mathbb{Z}.$$

We assume that $S = \{\mathfrak{p}_i\}_{i\leq k}$ is a set of non-zero prime ideals stable under the action of $\mathrm{Gal}(K/\mathbb{Q})$ such that $\langle S\rangle = H$. By using the recursive S-unit group computation, we can find a matrix $M \in \mathbb{Z}^{k\times k}$ such that the rows of M generate the lattice of vectors $\vec{v} \in \mathbb{Z}^k$ such that $\prod_i \mathfrak{p}_i^{v_i} \sim (1)$. This can be done by working in the subfields involved in the norm relation. Then, we compute the Smith Normal Form (SNF) $\mathrm{diag}(d_1, \ldots, d_k)$ of M and unimodular matrices U, V such that $UMV = \mathrm{diag}(d_1, \ldots, d_k)$ (note that some d_i might equal 1).

Given a fractional ideal $\mathfrak{a}$ such that $[\mathfrak{a}] \in H$, we are interested in computing the unique exponents in $\mathbb{Z}/d_1\mathbb{Z} \times \cdots \times \mathbb{Z}/d_k\mathbb{Z}$ of the decomposition of $[\mathfrak{a}]$ according to the generators $(\mathfrak{g}_i)_{i \leq k}$. The previous section shows how to decompose $[\mathfrak{a}]$ according to the primes in S. We can convert this decomposition into one with respect to the $\mathfrak{g}_i$ via linear algebra involving V. Indeed, the conversion back-and-forth between a representation over the $\mathfrak{g}_i$ and one over the $\mathfrak{p}_j$ corresponds to a multiplication of V^{-1} (resp. V) with the vector of exponents:

$$\mathfrak{a} \sim \prod_j \mathfrak{p}_j^{x_j} = \prod_j \left(\prod_i \mathfrak{g}_i^{x_j \cdot V_{i,j}^{-1}} \right) = \prod_i \mathfrak{g}_i^{\sum_j x_j V_{i,j}^{-1}} = \prod_i \mathfrak{g}_i^{\left(V^{-1} \cdot \vec{x}^T\right)}.$$

Therefore $\mathfrak{a} \sim \prod_i \mathfrak{g}_i^{x_i'}$ for $\vec{x'} := V^{-1} \cdot \vec{x}$. By a similar argument, if $\mathfrak{a} \sim \prod_i \mathfrak{g}_i^{y_i'}$, then $\mathfrak{a} \sim \prod_i \mathfrak{p}_i^{y_i}$ where $\vec{y} = V \cdot \vec{y'}$.

Algorithm 11: Conversion of decomposition with respect to primes in S to generators

Require: Number field K, Set S of non-zero primes $(\mathfrak{p}_i)_{i \leq s}$, vector $\vec{x}$ such that $\mathfrak{a} \sim \prod_i \mathfrak{p}_i^{x_i}$, U, V unimodular such that $UMV = \mathrm{diag}(d_1, \ldots, d_k)$ where the rows of M are a basis of the lattice of relations between primes in S.

Ensure: $\vec{x'}$ with $\mathfrak{a} \sim \prod_i \mathfrak{g}_i^{x_i'}$ where $\langle S \rangle = \langle [\mathfrak{g}_1] \times \ldots \times \langle [\mathfrak{g}_k] \rangle$.

1: **return** $\vec{x'} := V^{-1} \cdot \vec{x}$.

4.3.3 Application to the PIP

Ideal decomposition, while useful for many routines in number theory, can also be used for resolution of the PIP. Assume we have computed a basis $(\beta_i)_{i \leq r+k}$ for the S-unit group (using recursive norm relation techniques) where the primes in S generate the class group. Let $M \in \mathbb{Z}^{(r+k) \times k}$ such that $(\beta_i) = \prod_j \mathfrak{p}_j^{M_{i,j}}$. Compute a decomposition of $\mathfrak{a} \subseteq \mathcal{O}_K$ as $\mathfrak{a} = (\beta) \prod_i \mathfrak{p}_i^{y_i}$ for some $\beta, \vec{y}$ by Algorithm 10. Then $\mathfrak{a}$ is principal if and only if $\vec{y}$ belongs to the lattice of relations given by the rows of M. We summarize this in Algorithm 12.

Algorithm 12: Solving the PIP using ideal decomposition

Require: A fractional ideal $\mathfrak{a}$ of K.

Ensure: A generator of $\mathfrak{a}$ if it exists, otherwise false.

1: Compute a set S of non-zero primes $(\mathfrak{p}_i)_{i \leq s}$ that generate $\mathrm{Cl}(\mathcal{O}_K)$.
2: Compute a matrix M whose rows are a basis of the relations between $\mathfrak{p}_i$ in S.
3: Compute $\alpha, \vec{x}$ such that $\mathfrak{a} = (\alpha) \prod_i \mathfrak{p}_i^{x_i}$ with Algorithm 10.
4: Solve $\vec{y}M = \vec{x}$.
5: **return** false if no solution, or $\alpha = \prod_i \beta_i^{y_i}$.

Theorem 4.14 (under GRH)**.** *Algorithm 12 is correct and has complexity*

$$\mathrm{Poly}(\log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a})), l, \max_i \log(a_i)) + l \cdot (\mathrm{COST}_{S-units}(\mathrm{Subfields}) + \mathrm{COST}_{IdealDec}(Subfields)),$$

where $\text{COST}_{S\text{-units}}(\text{Subfields})$ *and* $\text{COST}_{Ideal\ Dec}(\text{Subfields})$ *denote the cost of computing an* S*-unit group in a subfield and the cost of ideal decomposition in a subfield respectively.*

Proof. Steps (1) and (2) can be done recursively of course, reducing to the cost of subfield S-unit group computation in time $\text{Poly}(\log|\Delta_K|, l, \max_i \log(a_i))$. Then the result follows directly from the complexity analysis of Algorithm 9 and Algorithm 10. □

Recall the families of fields with good norm relations of Section 3.4.2. Algorithm 12 enjoys the same asymptotic improvements in these fields as in Algorithm 5.

Proposition 4.15 (Under GRH)**.** *Let* $a > 0$ *and* $(m_k)_{k\in\mathbb{Z}_{>0}}$ *be a sequence of integers satisfying* $\lambda(m_k) \leq \varphi(m_k)^a$ *for all* k*. Then Algorithm 12 applied to the infinite family of fields* $K_k := \mathbb{Q}(\zeta_{m_k})$ *has asymptotic complexity*

$$\text{Poly}([K:\mathbb{Q}], \log(\text{N}(\mathfrak{a}))) \cdot 2^{\tilde{O}([K:\mathbb{Q}]^{2a/3})}.$$

Proof. This follows from the proof of Proposition 4.4, with the only difference being the primes used in the S-unit group computation and the use of ideal decomposition in the subfields.

As the class group is generated by prime ideals with norm bounded by $12\log|\Delta_K|^2 \in O(\log|\Delta_K|^2)$ this is no issue, and by 2.14 the subfield ideal decompositions take time $\text{Poly}(\log(\text{N}(\mathfrak{a}))) \cdot 2^{\tilde{O}(n^{2a/3})}$ using the same argument used for the S-unit group. □

Proposition 4.16 (Under GRH)**.** *There exists an infinite sequence of integers* $m_1 < m_2 < \cdots$ *such that Algorithm 12 has complexity*

$$\text{Poly}([K:\mathbb{Q}], \log(\text{N}(\mathfrak{a}))) \cdot 2^{(\log(m_k))^{O(\log\log\log(m_k))}}$$

which by is in $2^{n^{o(1)}}$.

In practice we can do better than Algorithm 12. Consider what happens if we compute a decomposition of $\mathfrak{a} \subseteq \mathcal{O}_K$ by Algorithm 10 using an empty set S. Then subfield decompositions reduce to the PIP and Algorithm 10 only needs the unit group $\mathcal{O}_K^\times$ for lifting these to a solution $\mathfrak{a} = \alpha\mathcal{O}_K$, resolving the PIP in K. This further avoids computing the S-unit group $\mathcal{O}_{K,S}^\times$ where S generates the class group in Algorithm 12. These changes reduce the amount of saturations, compact representations, and root computations needed, which has a significant impact on practical performance. In Section 4.4 we describe this approach, and show that this can be even further improved by eliminating the need for the full unit group $\mathcal{O}_K^\times$, instead working with a finite index subgroup.

4.4 PIP Without S-Units

We will once again leverage the following identity for ideals $\mathfrak{a} \subseteq \mathcal{O}_K$ in the prescence of a norm relation of the form (4.1).

$$\mathfrak{a}^d = \prod_{i=1}^{l} \mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i})^{a_i} \mathcal{O}_K. \tag{4.3}$$

We will specialize the techniques used to compute ideal decompositions in the prescence of a norm relation developed in Section 4.3 to the task of resolving the PIP. Notably, we avoid not just the computation of S-unit groups in K but generally avoid computing the full unit group $\mathcal{O}_K^\times$ at all, instead requiring only a finite index subgroup. This leads to a significant improvement to the resolution of the PIP in practice. In Section 4.4.2 we give computational evidence supporting this claim and successfully compute the generator of a principal ideal in a cyclotomic field of degree 1800, much larger than previously possible.

Lemma 4.17. *Let $\mathfrak{a}$ be a fractional ideal of K. If $\mathfrak{a}$ is principal, then $\mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i})$ is principal for all $1 \leq i \leq l$. If $d = 1$, then the converse also holds and a generator of $\mathfrak{a}$ is given by $\prod_{i=1}^{l} \alpha_i^{a_i}$, where $\alpha_i \mathcal{O}_K = \mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i})$, $1 \leq i \leq l$.*

The previous lemma shows that if the denominator d is equal to 1, then solving the principal ideal problem in K is equivalent to solving the principal ideal problems in the subfields K_i, $1 \leq i \leq l$. In case the denominator is not equal to 1, the situation is more complicated. Indeed, we can only find α such that $\mathfrak{a}^d = \alpha \mathcal{O}_K$. If α happens to be a d-th power, say $\alpha = \beta^d$, then β generates $\mathfrak{a}$. Otherwise, if α is not a d-th power but $\mathfrak{a}$ is principal, there must be another generator of $\mathfrak{a}^d$ that is a d-th power. Then we need to find a unit $u \in \mathcal{O}_K^\times$ such that $u\alpha = \beta^d$, so α is a d-th power modulo the unit group $\mathcal{O}_K^\times$.

It follows that if $\mathfrak{a}^d = \beta \mathcal{O}_K$, then $\mathfrak{a}$ is principal if and only if β is a root modulo $U = \mathcal{O}_K^\times$. Moreover if $u\beta = \alpha^d$ for some $u \in \mathcal{O}_K^\times$, then $\mathfrak{a} = \alpha \mathcal{O}_K$. Working with the full unit group $U = \mathcal{O}_K^\times$ can be expensive in practice. In the following we improve upon this by showing that in our situation we can often pick a smaller group U generated by subgroups of the unit groups $\mathcal{O}_{K_i}^\times$. We begin by showing that one can restrict to full rank subgroups with index coprime to d.

Lemma 4.18. *Assume that $U \subseteq K^\times$ is a multiplicative group, $\beta \in K^\times$ and $d \in \mathbb{Z}$. Further let $V \subseteq U$ be a subgroup of finite index with $[U : V]$ coprime to d. Then β is a d-th power modulo U if and only if β is a d-th power modulo V.*

Proof. Let $k = [U : V]$ and $a, b \in \mathbb{Z}$ such that $ad + bk = 1$. Assume that there exists $u \in U$ such that $u\beta \in (K^\times)^d$. As $u\beta = (u^a)^d (u^k)^b \beta$ and $v = (u^k)^b \in V$, we have $v\beta = u\beta/(u^a)^d \in (K^\times)^d$, thus showing that β is a d-th power modulo V. The other implication is clear. □

We can now show that in the presence of norm relations, it is sufficient to work with a multiplicative group generated by units from the involved subfields. In fact, not even the full unit groups of the subfields are necessary, but just subgroups with index coprime to d.

Proposition 4.19. *Let $\mathfrak{a}$ be a fractional ideal satisfying (3.4). Assume that the ideal $\mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i}) = \alpha_i \mathcal{O}_{K_i}$ is principal for all $1 \leq i \leq l$ and let $\beta = \prod_{i=1}^{l} \alpha_i^{a_i}$. Consider the multiplicative group $W = (\mathcal{O}_{K_1}^{\times})^{a_1} \cdots (\mathcal{O}_{K_l}^{\times})^{a_l} \subseteq \mathcal{O}_K^{\times}$. Let $V \subseteq W$ be a subgroup of finite index with $[W : V]$ coprime to d and $V_i \subseteq \mathcal{O}_{K_i}^{\times}$ subgroups of finite index with $[\mathcal{O}_{K_i}^{\times} : V_i]$ coprime to d. Then the following are equivalent:*

1. *The ideal $\mathfrak{a}$ is principal.*
2. *The element β is a d-th power modulo $\mathcal{O}_K^{\times}$.*
3. *The element β is a d-th power modulo W.*
4. *The element β is a d-th power modulo V.*
5. *The element β is a d-th power modulo $V_1^{a_1} \cdots V_l^{a_l}$.*

If we have $a_i \in \mathbb{Z}$ for all $1 \leq i \leq l$, then we can use $W = \mathcal{O}_{K_1}^{\times} \cdots \mathcal{O}_{K_l}^{\times}$ in (d) and $V_1 \cdots V_l$ in (e).

Proof. (a) $\Leftrightarrow$ (b): Clear.

(a) $\Leftrightarrow$ (c): Assume $\mathfrak{a} = \alpha \mathcal{O}_K$ is principal. As $\mathrm{N}_{K/K_i}(\alpha^{b_i})\mathcal{O}_{K_i} = \mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i}) = \alpha_i \mathcal{O}_{K_i}$, there exist units $u_i \in \mathcal{O}_{K_i}^{\times}$ such that $\mathrm{N}_{K/K_i}(\alpha^{b_i}) = u_i \alpha_i$. Thus

$$\underbrace{u_1^{a_i} \cdots u_l^{a_l}}_{\in W} \cdot \beta = \prod_{i=1}^{l} (u_i \alpha_i)^{a_i} = \prod_{i=1}^{l} \mathrm{N}_{K/K_i}(\alpha^{b_i})^{a_i} = \alpha^d \in (K^{\times})^d$$

and β is a d-th power modulo W. Conversely if β is a d-th power modulo W, it is also a d-th power modulo $\mathcal{O}_K^{\times}$ and hence $\mathfrak{a}$ is principal.

(c) $\Leftrightarrow$ (d): Lemma 4.18.

(d) $\Leftrightarrow$ (e): Since the V_i have index coprime to d, it follows that $[W : V_1^{a_1} \cdots V_l^{a_l}]$ is coprime to d. Hence the result follows again by Lemma 4.18. □

We now describe an algorithm for resolution of the PIP using norm relations in Algorithm 13. In contrast with Algorithm 5 and Algorithm 12 this approach does not require the computation of S-unit groups.

Remark 4.20. The idea of reducing the principal ideal problems to subfields using relative norms and the existence of d-th powers was already considered in [**BBVLV**] for multiquadratic and in [47] for multicubic fields. While not formulated using the notion of norm relations, in both works criterion 4.19 (b) is used to

Algorithm 13: Solving the PIP without S-units

Input : A fractional ideal $\mathfrak{a}$ of K.
Output: A generator of $\mathfrak{a}$ if it exists, otherwise false.

$y \leftarrow 1$;
for $i \leftarrow 1$ **to** l **do**
  if $\mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i})$ *is principal* **then**
    Find a generator $\alpha_i \in K_i$ of $\mathrm{N}_{K/K_i}(\mathfrak{a}^{b_i})$;
  else
    return: $\mathfrak{a}$ is not principal.
  end
end
$\beta \leftarrow \alpha_1^{a_1} \cdots \alpha_l^{a_l}$. //$\beta$ generates $\mathfrak{a}^d$.;
Compute $U = V_1^{a_1} \cdots V_l^{a_l}$, where the V_i are subgroups $\mathcal{O}_{K_i}^\times$ with index coprime to d. // Details in Section 4.2;
if β *is a d-th power modulo* U *by Algorithm 7 or Algorithm 8* **then**
  return: $\alpha \in K^\times$ such that $\beta/\alpha^d \in U$;
else
  return: $\mathfrak{a}$ is not principal;
end

decide the principal ideal problem in the field K. In particular, the full unit group had to be computed via saturation.

In contrast to the aforementioned papers, the use of Proposition 4.19 (c) allows us to avoid the computation of the full unit group $\mathcal{O}_K^\times$ of K (in most cases, see Section 4.2). Actually Proposition 4.19 (e) allows us to avoid the computation of the full unit groups in the subfields themselves. All that is required are subgroups whose index is finite and coprime to d. This results in a significant practical speed-up.

4.4.1 Asymptotic Analysis

In this section, we show how the cost of our PIP method relates to the cost of the PIP in subfields. We reiterate that the PIP algorithm presented here does not offer any asymptotic improvement over the PIP algorithms described in Section 4.1 and Section 4.3.3.

Theorem 4.21 (under GRH)**.** *Algorithm 13 is correct and has complexity*

$$\mathrm{Poly}(\log|\Delta_K|, \log(\mathrm{N}(\mathfrak{a})), l, \max_i \log(a_i)) + l \cdot \mathrm{PIP}(\mathrm{Subfields}),$$

where PIP(Subfields) *denotes the cost of Step (3) (PIP in a subfield).*

Proof. From Lemma 4.7 it follows that it is sufficient to show that Algorithms 7 and 8 have the claimed complexity. We first consider Algorithm 7. From [20, Theorem 4.11] it follows that the algorithm terminates

as soon as $c > 72d^2(\log|\Delta_K| + 3n\log(d))^2$, hence after a number of steps which is polynomial in the size of the input. As the same holds for the final root computation, this proves the claim for Algorithm 7.

For Algorithm 8, first note that a 2-saturated subgroup $U \subseteq \mathcal{O}_K^\times$ can be computed in polynomial time ([20, Corollary 4.13]). As the successive applications of Algorithm 7 for $p = 2$ have the same complexity, the claim follows. □

For completeness we also state the complexity in the families of fields admitting good norm relations described in Section 3.4.2. The proofs follow from the cost of PIP and compact representation in the subfields given by Corollary 2.17 and the analogous proofs in Sections 4.1.1 and 4.3.3.

Proposition 4.22 (Under GRH). *Let $a > 0$ and $(m_k)_{k\in\mathbb{Z}_{>0}}$ be a sequence of integers satisfying $\lambda(m_k) \leq \varphi(m_k)^a$ for all k. Then Algorithm 13 applied to the infinite family of fields $K_k := \mathbb{Q}(\zeta_{m_k})$ has asymptotic complexity*

$$\mathrm{Poly}([K:\mathbb{Q}], \log(\mathrm{N}(\mathfrak{a}))) \cdot 2^{\tilde{O}\left([K:\mathbb{Q}]^{2a/3}\right)}.$$

Proposition 4.23 (Under GRH). *There exists an infinite sequence of integers $m_1 < m_2 < \cdots$ such that Algorithm 13 has complexity*

$$\mathrm{Poly}([K:\mathbb{Q}], \log(\mathrm{N}(\mathfrak{a}))) \cdot 2^{(\log(m_k))^{O(\log\log\log(m_k))}}$$

which is in $2^{n^{o(1)}}$.

4.4.2 Numerical Results

We implemented our algorithm using the algebra package Hecke [35] (written in Julia [9]). We used 55 nodes of 20 cores 2x Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GHz processors with 192GB memory. We focused our attention on examples outside of the reach of the previous techniques by a substantial margin: The field $K^{(1)} = \mathbb{Q}(\zeta_{825})$ of degree 400 and discriminant $\approx 10^{960}$, the field $K^{(2)} = \mathbb{Q}(\zeta_{3276})$ of degree 864 and discriminant $\approx 10^{2369}$, and the field $K^{(3)} = \mathbb{Q}(\zeta_{2387})$ of degree 1800 and discriminant $\approx 10^{5539}$. For each field, we report on the resolution of one instance of the PIP chosen at random. The class groups of the fields $K^{(1)}$ and $K^{(2)}$ have been determined using norm relations in [20]. That computation took less than two hours using a single core, but note that the method employed for the class group computations avoids taking the roots of elements and is therefore computationally much easier than solving the PIP. In particular, it cannot be used to compute S-unit groups or to solve the PIP.

Each of these fields admit a denominator 1 norm relation, so we use the two step approach discussed in Section 3.4.1. First, we find a norm relation of denominator 1 of the form $x = \prod_{i=1}^{l} \mathrm{N}_{K/K_i}(x^{b_i})^{a_i}$ such

that the quantity $\max_{1\leq i\leq l}[K_i : \mathbb{Q}]$ is minimal. In other words, we find a norm relation of denominator 1 where the degrees of the subfields are as small as possible. As a second step, for each subfield K_i, we find a norm relation $x^{d_i} = \prod_{j=1}^{l_i} \mathrm{N}_{K_i/K_{i,j}}(x^{b_{i,j}})^{a_{i,j}}$ with d_i as small as possible, such that $\max_{1\leq j\leq l_i}[K_{i,j} : \mathbb{Q}]$ is bounded by some heuristically chosen constant B. To test whether some fractional ideal $\mathfrak{a}$ of K is principal, Algorithm 13 is now applied first using the norm relation of K of denominator 1 and then again using the norm relations of the K_i when testing whether $\mathrm{N}_{K/K_i}(\mathfrak{a}_i^{b_i})$ is principal. In particular this means that:

1. The largest degree field where we compute roots modulo units is $\max_{1\leq i\leq l}[K_i : \mathbb{Q}]$.

2. The largest degree field where we have to classically solve the principal ideal problem and compute (a saturated subgroup of) the unit group is bounded by B.

Since the norm relations are too large to display, we present several values quantifying the difficulty of solving the PIP in Table 2.

Table 2. Quantification of the hardness of PIP instances in large degree fields.

K	$[K : \mathbb{Q}]$	$l = \#\{K_i\}$	$n = \max_i[K_i : \mathbb{Q}]$	$\#\{K_{i,j}\}$	$m = \max_{i,j}[K_{i,j} : \mathbb{Q}]$
$K^{(1)}$	400	19	100	86	20
$K^{(2)}$	864	38	108	341	12
$K^{(3)}$	1800	131	150	297	30

Recall that n is the maximal degree of the subfields where saturation and root computation needs to take place. Likewise, m is the maximal degree of a subfield where the PIP must be solved with a subexponential method and the column labeled "$\#\{K_{i,j}\}$" denotes the number of these subfields. In particular, we observe that the saturation and root computation to solve the PIP in the field $K^{(3)}$ of degree 1800 only occurs in fields of degree bounded by 150 while the subexponential computations occur in subfields of degree no more than 30.

Results We ran the computation of subexponential PIP instances in the subfields on independent cores. A second layer of parallelization was employed by computing individual roots on independent cores. After picking a principal ideal $\mathfrak{a} = (\alpha)$ of $\mathcal{O}_K$, the main steps of our computations are the following: (1) Finding the initial norm relation to determine the subfields K_i, (2) Finding the norm relation in each of the subfields K_i, (3) computing the subfields $K_{i,j}$, (4) Computing the unit groups of the $K_{i,j}$, (5) Computing the relative norms $\mathrm{N}_{K/K_{i,j}}(\mathfrak{a})$, (6) Computing generators of the ideals $\mathrm{N}_{K/K_{i,j}}(\mathfrak{a})$, (7) Identifying d-powers (without root computation), (8) Compact representation, and (9) Root computation.

Table 3. PIP Runtime in CPU hours in large degree fields.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	Total
$K^{(1)}$	0.46	2.70	0.33	0.25	0.55	0.13	0.05	0.37	1.71	6.55
	7.0%	41.2%	5.0%	3.8%	8.4%	2.0%	0.8%	5.6%	26.1%	
$K^{(2)}$	4.12	15.23	2.90	1.00	11.21	1.54	0.12	2.61	60.19	98.92
	4.2%	15.4%	3.0%	1.0%	11.3%	1.6%	0.1%	2.6%	60.8%	
$K^{(3)}$	66.61	140.00	58.34	102.15	641.92	203.62	3.25	55.19	1634.01	2905.09
	2.3%	4.8%	2.0%	3.5%	22.1%	7.0%	0.1%	1.9%	56.2%	

4.4.3 Comparison With S-Unit Method

Algorithm 13 was designed with the practical performance in mind. This is why we strive to avoid the compact representation and saturation steps as much as possible. In particular, the units used are directly coming from subfields, and seldom need saturation. On the other hand, computing S-units as in Algorithm 5 would require significantly more compact representation and saturation steps. We will give a rough comparison of the expected number of compact representations in either case. The analysis given in the proof of Theorem 4.1 was aimed at estimating the worst-case complexity of the PIP using S-units. To simplify the discussion here we consider what is essentially the best-case scenario for Algorithm 5.

At the beginning, we can assume that we have performed an LLL-reduction on the input and replaced $\mathfrak{a}$ by $(b_1)/\mathfrak{a}$ where b_1 is the first basis vector. This means that we may assume

$$\mathrm{N}(\mathfrak{a}) \leq \lambda^n \sqrt{|\Delta|} \in O\left(2^{O(n^2)} \sqrt{|\Delta|}\right),$$

with $n = \deg(K)$, and $\lambda \sim 2^{O(n)}$ the approximation factor of LLL. Let $b_1, \ldots, b_n$ be an LLL-reduced basis of $\mathfrak{a}$, so that

$$\|b_1\| \leq \lambda |\Delta|^{1/2n} \mathrm{N}(\mathfrak{a})^{1/n} \in 2^{O(n)} |\Delta|^{1/2n} \mathrm{N}(\mathfrak{a})^{1/n}.$$

This means that $\mathrm{N}(b_1) \in 2^{O(n^2)} \sqrt{|\Delta|} \mathrm{N}(\mathfrak{a})$ and

$$\mathrm{N}\left((b_1)/\mathfrak{a}\right) \in 2^{O(n^2)} \sqrt{|\Delta|}$$

Due to the density of prime numbers, the number of times we expect to need to draw an element of norm $2^{O(n^2)} \sqrt{|\Delta|}$ before finding one whose norm is prime is about $O(n^2)$. The chosen strategy for enumeration of short elements in $\mathfrak{a}$ is to draw elements of the form

$$\alpha = b_{i_1} + b_{i_2} + \ldots, + b_{i_c},$$

for a constant c (typically $c = 3$ to ensure that the search space is large enough to find a prime norm) and a choice of c random indices $i_1, \cdots, i_c$. The vectors $b_2, \ldots, b_n$ are at least as long as b_1, but in the best case scenario, all vectors are of the same length. If this were the case, then $\|\alpha\| \sim \sqrt{c}\|b_1\|$. In practice, we expect $\|\alpha\|$ to be in fact larger. In any case, the algebraic norm of α is expected to satisfy

$$\mathrm{N}\left(\frac{(\alpha)}{\mathfrak{a}}\right) \geq c^{n/2}\mathrm{N}\left(\frac{(b_1)}{\mathfrak{a}}\right) \sim c^{n/2}\lambda^n\sqrt{|\Delta|} \sim c^{n/2}\mathrm{N}(\mathfrak{a}).$$

Hence the bit size of the norm of an S-unit, where $S = \{\mathfrak{p}^\sigma \mid \mathfrak{p} = (\alpha)/\mathfrak{a}, \sigma \in \mathrm{Gal}(K/\mathbb{Q})\}$, is expected to be $O(n)$ times larger than the bit size of the norm of the generator of $\mathfrak{a}$. In addition, the cardinality of S is equal $|\mathrm{Gal}(K/\mathbb{Q})| + r \in O(n)$ where r is the rank of the unit group of K. Here we assume that $\mathfrak{p}$ is of degree 1, which happens the majority of the time. In any case, $r \in \Omega(n)$ is a lower bound on $|S|$.

In the compact representation algorithm, the complexity of the calls to LLL is proportional to the $\log(\mathrm{N}(\mathfrak{a}'))$ where $\mathfrak{a}'$ is an ideal whose norm is proportional to $\prod_{\mathfrak{p}} \mathrm{N}(\mathfrak{p})$ where $\mathfrak{p}$ runs over the prime ideals dividing the input element. Hence, when computing compact representations of S-units for S defined above, the execution time of the compact representation algorithm is proportional to $|S|\log(\mathrm{N}(\mathfrak{p}))$ where $\mathfrak{p} = (\alpha)/\mathfrak{a}$. On the other hand, this term becomes $\log(\mathrm{N}(\mathfrak{a}))$ when the compact representation is called on a generator of $\mathfrak{a}$. Therefore, each call to the compact representation on input an S-unit is expected to be $O(n^2)$ times more expensive than the call on input a generator of $\mathfrak{a}$. Moreover, to solve the PIP with Algorithm 5, we need the entire S-unit group, which means that there need to be $|S|$ calls to the compact representation instead of a single one on the generator of $\mathfrak{a}$, hence multiplying the compact representation effort by an $O(n)$ factor. Altogether, we estimate that the S-unit based resolution of the PIP should be $O(n^3)$ times slower than the methods introduced in this paper.

Asymptotically, the slowdown induced by opting for Algorithm 5 does not impact the overall complexity which is strongly subexponential. However, given that this complexity is somewhat close to being polynomial, an n^3 slowdown does impact concrete computations to the point that the large degree calculations presented in this paper are infeasible with the S-unit method. To illustrate the sharp increase of the slowdown, we compared the two methods for small instances of increasing difficulty on a single core Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GHz with 192GB memory. We see in Table 4 that trivial examples are more easily solved by using S-units because in small dimension, the random choice of a small element in $\mathfrak{a}$ almost always directly yields a generator. However, the method described in this paper is already showing its impact on examples taking 30 min (by being twice as fast), and on examples taking several hours on a single core, it is already faster by a significant margin. This sharp increase backs our heuristic estimate of an n^3 speedup.

Table 4. Comparison with the S-unit method.

$[K:\mathbb{Q}]$	This paper	S-unit method
36	5.5sec	0.3 sec
72	2.1 min	1.6 min
144	31.7 min	1.1h
216	3.6h	54h

We further document our estimate by running the following experiment. In the degree 1800 field $K^{(3)}$ of Section 5.6, the chosen PIP challenge has the prime decomposition $\mathfrak{a} = \mathfrak{p}_1\mathfrak{p}_2$ where $\mathrm{N}(\mathfrak{p}_1) = 133673$ and $\mathrm{N}(\mathfrak{p}_2)$ is a 733 bit prime. On the other hand, our experiments showed that random short elements $\alpha \in \mathfrak{a}$ had algebraic norms of about 1,000,000 bits. This is consistent with our estimate that the bit size of the norm of the primes in S is $O(n)$ times larger than the bit size of the norm of $\mathfrak{a}$.

CHAPTER 5

MILDLY SHORT VECTORS IN CYCLOTOMIC IDEAL LATTICES

Given a Euclidean lattice $\mathcal{L}$ and $\gamma \geq 1$, recall that the problem of finding a non-zero $v \in \mathcal{L}$ such that $\|v\| \leq \gamma\lambda_1(\mathcal{L})$ is called the γ-Shortest Vector Problem (γ-SVP). The security of lattice-based cryptosystems such as LWE schemes [62] relies on the hardness of γ-SVP for γ polynomial in the dimension of the lattice. The LLL algorithm [45] solves γ-SVP for $\gamma \in 2^{O(n)}$ in polynomial time in n. Exponential algorithms such as sieve methods [2] can solve exact SVP (i.e. $\gamma = 1$) in time $2^{O(n)}$, while the BKZ algorithm [65] allows one to solve γ-SVP for $\gamma \in 2^{O(n/k)}$ in time $2^{O(k)}$. In particular, the time to solve γ-SVP for $\gamma \in 2^{\tilde{O}(\sqrt{n})}$ is in $2^{\tilde{O}(\sqrt{n})}$. In [28], solutions of γ-SVP for $\gamma \in 2^{\tilde{O}(\sqrt{n})}$ are referred to as *mildly short vectors* and we adopt this terminology. The study of the hardness of γ-SVP is crucial both from a fundamental standpoint and for its applications to cryptology. In particular, there are no efficient algorithms to solve γ-SVP for non-exponential γ. In the subexponential γ regime, any superpolynomial improvement over the state of the art (i.e. the BKZ algorithm) represents a significant step forward.

To gain efficiency, variants of lattice-based cryptosystems using lattices that are ideals in cyclotomic number fields were introduced. This is the case of cryptosystems based on the Ring Learning With Error (RLWE) problem [49]. It can be shown that γ-SVP in the cyclotomic field $\mathbb{Q}(\zeta_m)$ with a polynomial γ reduces to RLWE in this field. The most typical cyclotomic fields used in RLWE cryptosystems are those of the form $\mathbb{Q}(\zeta_{2^l})$ for some l (i.e. the fields with a power-of-two conductor). However, the use of general cyclotomic fields is possible [50]. One of the main security assumptions on which ideal lattice based cryptosystems rely is that γ-SVP in ideals of cyclotomic fields is not significantly easier than in general Euclidean lattices. Because of that, γ-SVP algorithms for ideals of $\mathbb{Q}(\zeta_m)$ that outperform the BKZ reduction method are of particular interest. Indeed, they document the gap between the hardness of this problem in the special case of ideals in cyclotomic fields and in the case of general lattices. However, such improvements do not necessarily imply an attack against RLWE schemes. Indeed, the proof of security of RLWE schemes relies on the hardness of γ-SVP for a polynomial γ. Hence, the hardness of the search for mildly short vectors does not directly impact it. In addition, even an efficient algorithm for the resolution of γ-SVP with a polynomial γ would not necessarily imply the cryptanalysis of RLWE schemes. It would however render the security proof moot.

Because of its close connection to the security proof of RLWE cryptosystems, the investigation of the hardness of γ-SVP in ideal lattices of cyclotomic fields (including the search for mildly short vectors) is a crucial stake in mathematical cryptology. It was heuristically observed by a scientific team from the British Government Communications Headquarter (GCHQ) that the search for short generators of principal ideals of $\mathbb{Q}(\zeta_{2^l})$ should be efficient with a quantum computer [23]. This observation relied on two conjectures: a) Quantum computers allow us to efficiently find generators of principal ideals in number fields, and b) The search for a short generators of a principal ideal in $\mathbb{Q}(\zeta_{2^l})$ efficiently reduces (on a classical computer) to the search for an arbitrary generator. Point a) was proven by Biasse and Song [16], while Point b) was proven by Cramer, Ducas, Peikert and Regev [30]. In particular, it is shown in [30] that a short generator of a principal ideal of $\mathbb{Q}(\zeta_{2^l})$ is a solution to γ-SVP for $\gamma \in 2^{\tilde{O}(\sqrt{n})}$ with $n = [\mathbb{Q}(\zeta_{2^l}) : \mathbb{Q}]$ (i.e. a mildly short vector). This is the first example of a superpolynomial gap between the hardness of γ-SVP in ideal lattices and in general lattices. This line of work was further expanded by Cramer, Ducas and Wesolowski [29] who showed that there was an efficient heuristic quantum reduction from the search for mildly short vectors in general ideal lattices of $\mathbb{Q}(\zeta_{2^l})$ (i.e. not necessarily principal) and the search for generators in principal ideals. This result was later extended to ideal lattices of $\mathbb{Q}(\zeta_m)$ for arbitrary m in [28]. We refer to this heuristic reduction as CDW. To achieve an efficient quantum reduction, the CDW approach relies on two assumptions: a) The class group of the maximal real subfield $\mathbb{Q}(\zeta_m)^+$ of $\mathbb{Q}(\zeta_m)$ is small, and b) The minus part $\mathrm{Cl}^-(\mathcal{O}_K)$ of the ideal class group of $K = \mathbb{Q}(\zeta_m)$ is generated by few prime ideals. Limited numerical data in support of these conjectures is available, and most of it concerns the case $m = p$ a prime.

One of the key aspects of $\mathbb{Q}(\zeta_m)$ that enables the CDW approach is the knowledge of a set of units with good properties: the cyclotomic units. Indeed, units in number fields can be arbitrarily large, but in $\mathbb{Q}(\zeta_m)$, one can efficiently construct a set of small units that generate a subset of finite index of the group of units. This approach has been generalized by Pellet-Mary, Hanrot, and Stehlé who used S-units instead of units for a small enough set of primes S that generates the ideal class group of the field [59]. This method, known as PHS, allows one to solve γ-SVP for γ in $2^{O(n^a)}$ where $a < 1/2$ at the cost of an exponential precomputation on the S-units based on the work of Laarhoven [44]. The PHS approach was further improved [6], but the cost of the precomputation prevents it from solving γ-SVP more efficiently than the benchmark BKZ method. Recent preliminary work from Bernstein and Lange [8] conjectured that S-units of cyclotomic field have properties allowing one to adapt the PHS approach to outperform BKZ in the search for solutions to γ-SVP where $\gamma \in 2^{O(n^a)}$ with $a < 1/2$. To this date, there is no available strong evidence of this conjecture, even if it seems like the lattice of logarithmic embeddings of S-units of cyclotomic fields might not comply with so-called "Gaussian heuristics" which provide estimates for the first minima of random lattices. Independent work of Bernard, Lesavourey, Nguyen and Roux-Langlois [7] aimed

to improve S-unit attacks by investigating sets of small S-units analogously to the case of cyclotomic units. They also attempted to remove the need for quantum computers in the CDW approach, but they were not able to improve the bottleneck of the method which consists in decompositions of ideals in the ideal class group.

We adapt the results of Chapter 4 to the problem of finding a mildly short vector of an ideal lattice in a cyclotomic field. First we review the case of a principal ideal. In particular, we review the efficient reduction of the search for a mildly short vector to the PIP of [30, 81] which, when combined with the results of Chapter 4, result in an asymptotic improvement over BKZ. Afterwards we move to the case of an arbitrary ideal and review the efficient quantum reduction of [28, 29]. We show that the bottleneck of this reduction, the decomposition of ideals in the minus part of the class group $\mathrm{Cl}^-(\mathcal{O}_K)$, can also be improved by the results of Chapter 4. With this we propose a classical variant of the CDW algorithm and analyze the asymptotic cost of this algorithm in Section 5.5. in Section 5.6 we present an implementation of our subfield CDW variant, as well as numerical data in support of the heuristics made to support the runtime of the CDW method.

5.1 Mildly Short Vectors in Principal Ideals

First we recall the main results of [30, 81], the reduction from the search for mildly short vectors in a principal ideal lattice to the PIP. A solution to the mildly short vector problem in a principal ideal $\mathfrak{a}$ is a generator α of $\mathfrak{a}$ such that $\|\sigma(\alpha)\| \leq 2^{\tilde{O}(\sqrt{n})}$ where σ is the Minkowski embedding (note that we often work instead with the logarithmic embedding $\|\mathrm{Log}(\alpha)\|$.) When this happens we simply say α is a short generator of $\mathfrak{a}$. If we are given an arbitrary generator β of $\mathfrak{a}$, it is related to α by $\alpha = u\beta$ where $u \in \mathcal{O}_K^\times$ is a unit. Our task then is: given a generator β of $\mathfrak{a}$, find a $u \in \mathcal{O}_K^\times$ such that $\alpha = u\beta$ is short.

The techniques we are using in this section were originally stated for cyclotomic fields $K = \mathbb{Q}(\zeta_m)$ in the case where $m = 2^k$ [30], but this was recently extended to the case of an arbitrary conductor in [81]. It relies on cyclotomic units, which are the units generated by $\{\pm\zeta_m\} \cup \{1 - \zeta_m^i \mid j = 1, \ldots, m-1\}$. We denote this subgroup of $\mathcal{O}_K^\times$ by C. From [79, Th 4.12], we know that $[\mathrm{Log}(\mathcal{O}_K^\times) : \mathrm{Log}(C)]$ has finite index. Let the p_i be the prime divisors of $m = p_1^{\alpha_1} \ldots p_k^{\alpha_k}$, and let $m_i := m/p^{\alpha_i}$. From [81, Th 6.12], we know that the $\mathrm{Log}(v_j)$ generate $\mathrm{Log}(C)$ and that $\|\mathrm{Log}(v_j)\| \in O(\sqrt{m})$ for

$$v_j = \begin{cases} 1 - \zeta_m^j & \text{if for all indices } i, \text{ we have } m_i \nmid j; \\ \frac{1-\zeta_m^j}{1-\zeta_m^{m_i}} & \text{otherwise, for the unique } i \text{ such that } m_i \mid j. \end{cases}$$

Algorithm 14: Finding a short generator of $g\mathcal{O}_K$

Require: A generator g of $I = g\mathcal{O}_K$.
Ensure: A unit u such that ug is a short generator of I.
1: $\forall i,\ w_i \leftarrow \mathrm{Log}(v_i),\ W \leftarrow (w_1, \ldots, w_{m-1})$.
2: $s(G) \leftarrow \sum_{\sigma \in G} \sigma \in \mathbb{R}[G]/(1-\tau)$.
3: $t' \leftarrow \mathrm{Log}(g),\ t'' \leftarrow \frac{1}{\varphi(m)} \cdot \log(\mathrm{N}(g)) \cdot s(G)$.
4: $t \leftarrow t' - t'' \in \mathrm{Log}(\mathcal{O}_K^\times) \otimes \mathbb{R},\ x \leftarrow (0, \ldots, 0)$.
5: **while** $\|W \cdot x - t\|_\infty > \sqrt{2 \cdot \log(4\varphi(m))} \cdot \max_{w \in W} \|w\|$ **do**
6: $\quad x \leftarrow \mathrm{CV}_\infty(W, t)$.
7: **end while**
8: **return** $u := \prod_i v_i^{-x_i}$.

The procedure $\mathrm{CV}_\infty(W, t)$ in Step 6 is described in [81, Lem 6.6] and finds a vector x in W that is close to t for the infinity norm, given the set of short generators w_i for W that we have as input. One of the technical challenges outlined in [81] is that we need to ensure that we can work with rational approximations of the w_i, and of $\mathrm{Log}(g)$ while ensuring numerical stability. Assume the input g is given as the (non-evaluated) product $g = \prod_{i \le k} \gamma_i^{k_i}$ for $\gamma_i \in \mathcal{O}_K$, $k_i > 0$, and let $p = \max_i \lceil \log(k_i \|\gamma_i\|) \rceil$. In [81, Sec. 6.5.2], fixed point approximations with $p + m^2$ bits of precision were used, that is, the approximation of $x \in \mathbb{R}$ is given by $\overline{x} \in \mathbb{Q}$ of the form $\frac{a_x}{2^{p+m^2}}$ where $a_x \in \mathbb{Z}$. Then, to use $\mathrm{CV}_\infty(W, t)$, we need an approximation $\overline{W}$ of W with $p + m^2$ bits of precision that lies in $\mathrm{Log}(\mathcal{O}_K^\times) \otimes \mathbb{R}$. This is achieved by computing an approximation $\tilde{W}$ of W with $p + m^2 + 1$ bits of precision and setting

$$\overline{w}_i := \tilde{w}_i - \frac{2}{\varphi(m)} \sum_{j=1}^{\varphi(m)/2} \tilde{w_{i,j}} s(G).$$

The matrix $\overline{W}$ satisfies $\|\overline{W} - \tilde{W}\|_\infty \le \frac{1}{2^{p+m^2+1}}$, and $\forall i,\ \overline{w}_i \in (s(G) \cdot \mathbb{R})^\perp = \log(\mathcal{O}_K^\times) \otimes \mathbb{R}$. Then it can be shown that the element $\overline{x}$ such that

$$\|\overline{W} \cdot \overline{x} - \overline{t}\|_\infty \le \sqrt{2 \log(4\varphi(m))} \max_{w \in W} \|\overline{w}\|$$

gives us a short generator $g \cdot \prod_i v_i^{-\overline{x}_i}$ of I.

Theorem 5.1 (Th. 6.15 of [81]). *There is a randomized algorithm that for any $g \in \mathcal{O}_K$ finds an element $h \in \mathcal{O}_K$ such that $g\mathcal{O}_K = h\mathcal{O}_K$ and*

$$\|h\| = e^{O\left(\sqrt{m \log(m)}\right)} \cdot \mathrm{N}(g)^{1/\varphi(m)}.$$

In the quantum setting the PIP can be resolved in polynomial time by [16], so Theorem 5.1 shows that there is a quantum polynomial time algorithm for finding mildly short vectors in principal ideal lattices of a cyclotomic field. In the classical case we can apply the results of Chapter 4 to see that in cyclotomic fields with conductor m satisfying $\lambda(m) \leq \varphi(m)^a$ for $a < 3/4$ we can find mildly short vectors of principal ideal lattices asymptotically faster than BKZ, which is in $2^{\tilde{O}(\sqrt{n})}$. While asymptotically such conductors have negligible density, by Table 1 we see that for conductors in the practical range ($m \leq 100,000$) approximately 25% do satisfy this condition. Furthermore by Theorem 3.12 there exists an infinite sequence of conductors $m_1 < m_2 < \ldots$ satisfying $\lambda(m_k) = (\log(m_k))^{O(\log\log\log(m_k))}$, and in the fields $\mathbb{Q}(\zeta_{m_k})$ we can find mildly short vectors in time $2^{n^{o(1)}}$.

5.2 The CDW Technique

We now consider the search for mildly short vectors where the input ideal is not necessarily principal. To reduce the search for mildly short vectors to the PIP, we first find an ideal $\mathfrak{b} \subseteq \mathcal{O}_K$ such that $\mathfrak{ab}$ is principal, and $\mathrm{N}(\mathfrak{b}) \in 2^{\tilde{O}(n^{3/2})}$ where $n = [K : \mathbb{Q}]$. In [29], this task is referred to as the Close Principal Multiple Problem. Then, the techniques of Section 5.1 yield a short generator of $\mathfrak{ab}$, which is a solution to γ-SVP in $\mathfrak{a}$ for γ in $2^{\tilde{O}(\sqrt{n})}$. This involves three main steps:

1. Multiply $\mathfrak{a}$ by random ideals of small norm $\mathfrak{a}_0$ until the class of $\mathfrak{a}' := \mathfrak{a}_0\mathfrak{a}$ is in $\mathrm{Cl}^-(\mathcal{O}_K) \subseteq \mathrm{Cl}(\mathcal{O}_K)$, the "minus part" of $\mathrm{Cl}(\mathcal{O}_K)$.

2. Decompose the class of $\mathfrak{a}'$ in $\mathrm{Cl}^-(\mathcal{O}_K)$ according to a set $S = \{\mathfrak{p}_1, \ldots, \mathfrak{p}_k\}$ of prime ideals that generate $\mathrm{Cl}^-(\mathcal{O}_K)$.

3. Find a close vector v in a lattice $\mathcal{L}$ (which is known to annihilate $\mathrm{Cl}^-(\mathcal{O}_K)$) to $t \in \mathbb{Z}^k$ such that $\mathfrak{a}'^{-1} \sim \prod_i \mathfrak{p}_i^{t_i}$.

Then the solution to the problem is $\mathfrak{b} := \mathfrak{a}_0 \prod_i \mathfrak{p}_i^{t_i - v_i} \sim \mathfrak{a}^{-1}$ (which has small norm if the vector v found in Step (3) is close enough to t). In [28, 29], Steps (1) and (2) require the quantum polynomial time algorithm of [16]. With classical computers, subexponential algorithms for ideal class group computations and the principal ideal problem (PIP) such as [18] can be used, but they do not provide a better complexity than the BKZ algorithm. Step (3) on the other hand, can be performed efficiently on a classical computer with the methods introduced in [29].

If $\mathfrak{a} \sim \prod_i \mathfrak{p}_i^{x_i}$, then $\mathfrak{b} = \prod_i \mathfrak{p}_i^{x_i'}$ with $x_i' = -x_i \bmod h$ and $h = |\mathrm{Cl}(\mathcal{O}_K)|$ satisfies that $\mathfrak{ab}$ is principal. The issue is that the x_i' can be quite large, thus preventing $\mathfrak{b}$ from satisfying $\mathrm{N}(\mathfrak{b}) \leq 2^{\tilde{O}(n^{3/2})}$. However, the techniques of [28, 29] show how to derive $\mathfrak{b} \sim \prod_i \mathfrak{p}_i^{x_i'}$ with small exponents. We recall the general idea of this

method, and we refer to [28, 29] for the details of the proofs. This task involves the search for close vectors in the so-called Stickelberger lattice, and to bound the runtime, we need to rely on a key conjecture:

Conjecture 5.2 ([28, Assumption 1])**.** *There are integers $l \leq \mathrm{Polylog}(m)$ and $B \leq \mathrm{Poly}(m)$ such that the following holds. Choose uniformly at random l prime ideals $\mathfrak{p}_1, \ldots, \mathfrak{p}_l$ among the primes of norm less than B that lie in $Cl^-(\mathcal{O}_K)$. Then the set $S := \{\mathfrak{p}_i^\sigma \mid \sigma \in G\}$ generates $Cl^-(\mathcal{O}_K)$ with probability at least $1/2$.*

We first compute a short generating set of $\mathrm{Cl}^-(\mathcal{O}_K)$. Then we perform a random walk in the Cayley graph of $\mathrm{Cl}(\mathcal{O}_K)$ whose edges are defined by the primes in S from Conjecture 5.2. In other words, this means that we multiply $\mathfrak{a}$ by random elements of S until we get an ideal $\mathfrak{a}'$ whose class lies in $\mathrm{Cl}^-(\mathcal{O}_K)$. This is described in [28, Alg. 5]. Its cost is in $O\left(h_K^+ \cdot \mathrm{Poly}(m, \log(\mathrm{N}(\mathfrak{a}))) \cdot \mathrm{Cost}(\mathrm{PIP})\right)$ according to [28, Lem. 5.2], where h_K^+ denotes the class number of the totally real subfield $\mathbb{Q}(\zeta_m + \zeta_m^{-1})$. To bound this asymptotic cost, we need to assume that h_K^+ is small enough:

Conjecture 5.3 ([28, Assumption 2])**.** *For any integer m, it holds that $h^+(m) \leq \mathrm{Poly}(m)$.*

So we find small $x_i \leq 0$ such that the class of $\mathfrak{a}' := \mathfrak{a} \cdot \prod_i \mathfrak{p}_i^{x_i}$ is in $\mathrm{Cl}^-(\mathcal{O}_K)$, and then we decompose the ideal class of $\mathfrak{a}'$ according to the set of primes S defined by Conjecture 5.2 to get a vector $\vec{y}$ such that $\mathfrak{a}' \sim \prod_i \mathfrak{p}_i^{y_i}$. Then [28, Sec. 4] constructs a lattice of vectors in $\mathbb{Z}[G]$ that act trivially on $\mathrm{Cl}^-(\mathcal{O}_K)$ from the Stickelberger ideal. The Stickelberger ideal (see [28, Sec. 4.1]) is an ideal of $\mathbb{Z}[G]$ that annihilates $\mathrm{Cl}(\mathcal{O}_K)$ but that does not have full rank as a $\mathbb{Z}[G]$-module. To get a full rank module, we project it to $R = \mathbb{Z}[G]/(1+\tau)$, where τ is the complex conjugation. The action of the resulting lattice L of R-rank $\varphi(m)/2$ annihilates $\mathrm{Cl}^-(\mathcal{O}_K)$ because $\tau + 1$ annihilates $\mathrm{Cl}^-(\mathcal{O}_K)$. The decomposition of $\mathfrak{a}'$ is then split according to each cycle under the action of R:

$$\mathfrak{a}' \sim \left(\prod_{\sigma \in G} (\mathfrak{p}_1^\sigma)^{y_{1,\sigma}}\right)\left(\prod_{\sigma \in G} (\mathfrak{p}_2^\sigma)^{y_{2,\sigma}}\right) \cdots \left(\prod_{\sigma \in G} (\mathfrak{p}_d^\sigma)^{y_{d,\sigma}}\right)$$

Then, we apply [28, Alg. 3] on each cycle $\vec{y_i} := (y_{i,\sigma})_{\sigma \in G}$. According to [28, Th. 4.7], this yields a vector $\vec{y_i'}$ such that $\prod_{\sigma \in G} (\mathfrak{p}_i^\sigma)^{y_{i,\sigma}} \sim \prod_{\sigma \in G} (\mathfrak{p}_i^\sigma)^{y'_{i,\sigma}}$ with $\|\vec{y_i'}\|_1 \leq \frac{1}{4}\varphi(m)^{3/2}$ in polynomial time in $\log\|\vec{y_i}\|$. Then, under Conjecture 5.2, $\mathfrak{b} := \prod_i \prod_\sigma (\mathfrak{p}_i^\sigma)^{y'_{i,\sigma}}$ satisfies $\mathrm{N}(\mathfrak{b}) \in 2^{O(n^{3/2})}$ and $\mathfrak{a}\mathfrak{b}$ is principal, thus solving the Close Principal Multiple Problem.

5.3 Computing the Minus Part of the Class Group

The computation of the minus part of the class group is an essential building block of Step (2) of Section 5.2. Recall that $\mathrm{Cl}^-(\mathcal{O}_K)$ is the kernel of the relative norm $\mathrm{N}_{K/K^+}\colon \mathrm{Cl}(\mathcal{O}_K) \to \mathrm{Cl}(\mathcal{O}_{K^+})$. Let $\mathfrak{g}_1, \ldots, \mathfrak{g}_k$ be such that $\mathrm{Cl}(\mathcal{O}_K) = \langle [\mathfrak{g}_1] \rangle \times \cdots \times \langle [\mathfrak{g}_k] \rangle$, and $\mathfrak{g}'_1, \ldots, \mathfrak{g}'_l$ be such that $\mathrm{Cl}(\mathcal{O}_{K^+}) = \langle [\mathfrak{g}'_1] \rangle \times \cdots \times \langle [\mathfrak{g}'_l] \rangle$. We could compute our norm map by decomposing each $\mathrm{N}_{K/K^+}(\mathfrak{g}_i)$ with respect to the $\mathfrak{g}'_j$ in $\mathrm{Cl}(\mathcal{O}_{K^+})$, however, we only know the $\mathfrak{g}_i$ in a product representation from the primes in $S = \{\mathfrak{p}_1, \ldots, \mathfrak{p}_s\}$ that generate $\mathrm{Cl}(\mathcal{O}_K)$. Evaluating these products would be costly. Instead, it is easier to decompose each $\mathrm{N}_{K/K^+}(\mathfrak{p}_i)$ with respect to the $\mathfrak{g}'_j$ in $\mathrm{Cl}(\mathcal{O}_{K^+})$. Then, since we know how to express each $\mathfrak{g}_i$ with respect to the primes in S, this allows us to associate with each $\mathrm{N}_{K/K^+}(\mathfrak{g}_i)$ a vector $\vec{x} \in \mathbb{Z}/d'_1\mathbb{Z} \times \cdots \times \mathbb{Z}/d'_l\mathbb{Z}$ that corresponds to the exponents of the decomposition of $\mathrm{N}_{K/K^+}(\mathfrak{g}_i)$. Therefore, we get a map

$$\varphi : \mathbb{Z}/d_1\mathbb{Z} \times \cdots \times \mathbb{Z}/d_k\mathbb{Z} \to \mathbb{Z}/d'_1\mathbb{Z} \times \cdots \times \mathbb{Z}/d'_l\mathbb{Z}$$

whose kernel is isomorphic to $\mathrm{Cl}^-(\mathcal{O}_K)$. We summarize this procedure in Algorithm 15.

Algorithm 15: Minus part of the ideal class group

Require: Number field K that admits a norm relation of the form (3.4).
Ensure: $\mathrm{Cl}^-(\mathcal{O}_K)$.
1: Compute a set of non-zero primes S that generate $\mathrm{Cl}(\mathcal{O}_K)$.
2: Compute a set of non-zero primes S_+ that generate $\mathrm{Cl}(\mathcal{O}_{K^+})$.
3: Compute a matrix M whose rows are a basis of the relations between $\mathfrak{p}_i$ in S.
4: Compute a matrix M_+ whose rows are a basis of the relations between $\mathfrak{q}_j$ in S_+.
5: Compute unimodular matrices U, V such that $UMV = \mathrm{diag}(d_1, \ldots, d_k)$.
6: Compute unimodular matrices U', V' such that $U'M_+V' = \mathrm{diag}(d'_1, \ldots, d'_l)$.
7: **for all** $\mathfrak{p}_i$ **do**
8: Find $\vec{x_i}$ such that $\mathrm{N}_{K/K^+}(\mathfrak{p}_i) \sim \prod_j \mathfrak{q}_j^{x_{i,j}}$ with Algorithm 10.
9: Find $\vec{x'_i}$ such that $\mathrm{N}_{K/K^+}(\mathfrak{p}_i) \sim \prod_j \mathfrak{g}'_j{}^{x'_{i,j}}$ with Algorithm 11.
10: **end for**
11: **for all** $\mathfrak{g}_i$ **do**
12: Find $\vec{y_i}$ such that $\mathfrak{g}_i \sim \prod_j \mathfrak{p}_j^{y_{i,j}}$ with the inverse of Algorithm 11.
13: $\vec{y'_i} \leftarrow \sum_j y_{i,j}\vec{x'_j}$ (hence $\mathrm{N}_{K/K^+}(\mathfrak{g}_i) \sim \prod_j \mathfrak{g}'_j{}^{y'_{i,j}}$)
14: **end for**
15: Let φ defined by $(0, \ldots, 0, \underbrace{1}_{i}, 0, \ldots, 0) \in \prod_j \mathbb{Z}/d_j\mathbb{Z} \mapsto \vec{y'_i} \in \prod_j \mathbb{Z}/d'_j\mathbb{Z}$.
16: **return** $\ker(\varphi)$.

Under the Generalized Riemann Hypothesis, there is a polynomial size set S of prime ideals that generate the ideal class group $\mathrm{Cl}(\mathcal{O}_K)$, namely $S := \{\mathfrak{p} \mid \mathrm{N}(\mathfrak{p}) \leq 12\log^2|\Delta_K|\}$, where Δ_K is the discriminant of K (see [4]). We refer to this bound on the norm of the prime ideals as Bach's bound. While this means that Steps (1) and (2) of Algorithm 15 are asymptotically efficient, one can hope to find generating sets of size

at most $O(\log|\Delta_K|)$. However, the effort required might be commensurate with that of computing the ideal class group. A method for class group computations using norm relations is described in [20, Alg. 4.23]. A byproduct of this algorithm is a set of primes S that generate the ideal class group. In a nutshell, it uses the fact that when K admits a norm relation of the form (3.4), the group $\mathrm{Cl}(\mathcal{O}_K)\otimes\mathbb{Z}[1/d]$ is isomorphic to a direct summand of $\bigoplus_{i=1}^{\ell}\mathrm{Cl}(\mathcal{O}_{K_i})\otimes\mathbb{Z}[1/d]$, and the group $\mathrm{Cl}(\mathcal{O}_K)/\mathrm{Cl}(\mathcal{O}_K)[d]$ is isomorphic to a subgroup of $\bigoplus_{i=1}^{\ell}\mathrm{Cl}(\mathcal{O}_{K_i})$. This means that a subset of $\mathrm{Cl}(\mathcal{O}_K)$ (namely $\bigoplus_{i=1}^{\ell}\mathrm{Cl}(\mathcal{O}_{K_i})\otimes\mathbb{Z}[1/d]$) is generated by the prime ideals above the primes that generate the $\mathrm{Cl}(\mathcal{O}_{K_i})$. The rest of the generators are chosen at random (Step (9) of [20, Alg. 4.23]). This probabilistic method relies on subfield computations, and is likely to return a generating set significantly smaller than that obtained from Bach's bound (which is quadratic in $\log(|\Delta_K|)$. Therefore we recommend the use of [20, Alg. 4.23] to perform Steps (1) and (2) of Algorithm 15. Note that there is no direct analogue of [20, Alg. 4.23] to compute the minus part of the class group. Indeed, no formula linking $\mathrm{Cl}^-(\mathcal{O}_K)$ to the class groups of the subfields involved in (3.4) exists, to the best of our knowledge. Therefore, we rely on the recursive ideal decomposition method introduced in Section 4.3 to perform this task.

The computation of the minus part of the class group enables Step (2) of Section 5.2 which consists in calculating a generating set of primes for $\mathrm{Cl}^-(\mathcal{O}_K)$ under Conjecture 5.2. Given the parameters l, B, we construct the set of prime ideals of K whose classes are in the minus part with norm bounded by B, and we repeatedly draw d sets of conjugates until one such subset generates the minus part of the class group. This procedure is summarized in Algorithm 16.

Algorithm 16: Creation of a generating set for $\mathrm{Cl}^-(\mathcal{O}_K)$

Require: Integers $l, B > 0$, number field K, and a norm relation $d = \sum_i a_i \mathrm{N}_{H_i} b_i$.
Ensure: A set $S = \{\mathfrak{p}_i\}_{i\leq k}$ of prime ideals such that $\forall i$, $[\mathfrak{p}_i] \in \mathrm{Cl}^-(\mathcal{O}_K)$, and the classes of $\mathfrak{p}^\sigma$ for $\sigma \in G$ generate $\mathrm{Cl}^-(\mathcal{O}_K)$.

$S_0 \leftarrow \{\}$.
for primes ideal $\mathfrak{p}$ with $\mathrm{N}(\mathfrak{p}) \leq B$ **do**
  if $\mathrm{N}_{K/K^+}(\mathfrak{p})$ is principal (using Algorithm 12) **then**
    $S_0 \leftarrow S_0 \cup \{\mathfrak{p}\}$.
  end if
end for
Compute $\mathrm{Cl}^-(\mathcal{O}_K)$ with Algorithm 15.
while true **do**
  $S \leftarrow d$ elements of S_0 chosen uniformly at random. $S' \leftarrow \{\mathfrak{p}^\sigma \mid \sigma \in G, \mathfrak{p} \in S\}$.
  Compute the S'-unit group and from the finite valuations of a generating set, deduce $\langle S'\rangle \subseteq \mathrm{Cl}(\mathcal{O}_K)$.
  if $\langle S'\rangle = \mathrm{Cl}^-(\mathcal{O}_K)$, **then return** S.
end while

Note that Step (3) only needs the knowledge that an ideal is principal and not a generator. We refer to this as the *decisional PIP*, in contrast to the search PIP. We can optimize Algorithm 12 or Algorithm 13

for this purpose. Assume we have an input ideal $\mathfrak{a}$ and find a generator β for $\mathfrak{a}^d$. Algorithm 7 then uses saturation techniques to find u such that $\alpha = u\beta$ is a d-th power if it exists. By choosing c in Algorithm 7 large enough, by Proposition 4.8 we can gaurantee that if saturation techniques find such a u then β must be a d-th power, and we don't need to actually compute the d-th root.

5.4 Subfield Variant of the CDW Algorithm

The building blocks presented in the previous sections are all that is needed to classically implement the CDW search for mildly short vectors in ideals of $K = \mathbb{Q}(\zeta_m)$ when K admits a norm relation of the form (3.4). For the sake of clarity, we recall the entire procedure in Algorithm 17.

Algorithm 17: Classical CDW search for mildly short vectors from norm relations

Require: Number field $K = \mathbb{Q}(\zeta_m)$ that admits a norm relation of the form (3.4). Ideal $\mathfrak{a} \subseteq \mathcal{O}_K$.
Ensure: A mildly short vector of $\mathfrak{a}$.

Let $S = \{\mathfrak{p} \mid \mathrm{N}(\mathfrak{p}) \leq 12\log^2|\Delta_K|\}$ (i.e. S generates $\mathrm{Cl}(\mathcal{O}_K)$ under GRH).
while true **do**
  Draw a random short product $\prod_i \mathfrak{p}_i^{x_i}$ of elements in S.
  if $\mathrm{N}_{K/K^+}(\mathfrak{a}\prod_i \mathfrak{p}_i^{x_i})$ is principal **then** $\mathfrak{a}_0 \leftarrow \mathfrak{a}\prod_i \mathfrak{p}_i^{x_i}$. **break**
end while
Compute a generating set of primes $S' = \{\mathfrak{p}_i\}_{i\leq s}$ of $\mathrm{Cl}^-(\mathcal{O}_K)$ with Algorithm 16.
Find $\vec{x}$ such that $\mathfrak{a}_0 \sim \prod_i \mathfrak{p}_i^{x_i}$ with Algorithm 10 (on input S').
Use $\vec{x}$ to derive $\mathfrak{b}$ with $\mathrm{N}(\mathfrak{b}) \in 2^{\tilde{O}(n^{3/2})}$ with [28, Alg. 4].
Find a generator α of $\mathfrak{a}_0\mathfrak{b}$ with Algorithm 5, Algorithm 12, or Algorithm 13.
Use Algorithm 14 to derive a short generator α' of $\mathfrak{a}_0\mathfrak{b}$.
return α'.

5.5 Asymptotic Analysis

Now we analyze the asymptotic complexity of Algorithm 17. We show that the cost is dominated (up to polynomial factors) by the cost of S-unit group computation and ideal class decomposition in the subfields involved in (3.4). Note that unless K is cyclic, which happens e.g. when $m = p^l$ is an odd prime power, there is always a norm relation that we can exploit to lower down the cost of ideal decompositions and computation of $\mathrm{Cl}(\mathcal{O}_K)$ and $\mathrm{Cl}^-(\mathcal{O}_K)$. This results in a practical gain for these tasks in almost all cyclotomic fields. Additionally, we observe asymptotic gains over the BKZ algorithm in the cyclotomic fields admitting good norm relations discussed in Section 3.4.2.

Proposition 5.4 (Under GRH). *Under Conjecture 5.2 and Conjecture 5.3 the cost of Algorithm 17 is in*

$$\mathrm{Poly}([K:\mathbb{Q}], \log \mathrm{N}(\mathfrak{a}))\left(\mathrm{Cost}_{S\text{-}unit}(\textit{subfields}) + \mathrm{Cost}_{Ideal\ Dec}(\textit{subfields})\right),$$

where $\text{Cost}_{S\text{-}unit}(subfields)$ *is the cost of computing* S*-units in the subfields involved in the norm relation* (3.4)*, and* $\text{Cost}_{Ideal\ Dec}(subfields)$ *is the cost of ideal decomposition in the subfields involved in the norm relation* (3.4)*.*

Proof. The random walk in Steps (2) to (5) of Algorithm 17 takes time in

$$O\left(h_K^+ \cdot \text{Poly}(m, \log(\text{N}(\mathfrak{a}))) \cdot \text{Cost}(\text{PIP})\right).$$

Under Conjecture 5.3, h_K^+ is polynomial. For the decisional PIP of Step (4) we can simplify the analysis by using a decisional variant of Algorithm 12. The cost of Algorithm 16 requires the computation of S-units (through Steps (3)-(4) of Algorithm 15, and Step (11) of Algorithm 16). It also requires ideal decompositions from Algorithm 9. Then Step (7) is another decomposition with Algorithm 10. Step (8) has an efficient solution by [28, Algorithm 3], as well as Step (10) (with Algorithm 14). For simplicity we will assume Algorithm 5 is used for the PIP in Step (9). Hence, up to polynomial factors, the cost of Algorithm 17 is that of

- the computation of S-units and,
- the decomposition of the ideal class of an ideal according to a set of primes.

We know that computing S-units in K efficiently reduces to computing subfield S-units by [20, Theorem 4.8]. Since we are working in a cyclotomic field, by Proposition 4.2 and Lemma 4.3 $\log|\Delta_K|, l$, and $\max_i \log a_i$ are all in $\text{Poly}([K:\mathbb{Q}])$. Combining this with the cost of recursive ideal decomposition given in Theorems 4.11 and 4.12 we obtain the stated complexity. □

Just as we have done many times already we now analyze the complexity of 17 in the families of fields given in Section 3.4.2. In cyclotomic fields with conductor m satisfying $\lambda(m) \leq \varphi(m)^a$ for $a < 3/4$ we can find mildly short vectors of arbitrary ideal lattices asymptotically faster than BKZ, which is in $2^{\tilde{O}(\sqrt{n})}$. We reiterate that asymptotically such conductors have negligible density, but Table 1 shows that for conductors in the practical range ($m \leq 100,000$) approximately 25% do satisfy this condition.

Proposition 5.5 (Under GRH)**.** *Assume the heuristics of [14], as well as Conjecture5.2, and Conjecture 5.3. Let* $a > 0$ *and* $(m_k)_{k\in\mathbb{Z}_{>0}}$ *be a sequence of integers satisfying* $\lambda(m_k) \leq \varphi(m_k)^a$ *for all* k*. Then Algorithm 17 applied to the infinite family of fields* $K_k := \mathbb{Q}(\zeta_{m_k})$ *has asymptotic complexity*

$$\text{Poly}([K_k:\mathbb{Q}], \log(\text{N}(\mathfrak{a}))) \cdot 2^{\tilde{O}\left([K_k:\mathbb{Q}]^{2a/3}\right)}.$$

Finally we consider the infinite sequence of conductors $m_1 < m_2 < \dots$ given by Theorem 3.12 satisfying $\lambda(m_k) = (\log(m_k))^{O(\log\log\log(m_k))}$. In the fields $\mathbb{Q}(\zeta_{m_k})$ we can find mildly short vectors of an arbitrary ideal lattice in time $\mathrm{Poly}([K_k : \mathbb{Q}], \log(\mathrm{N}(\mathfrak{a}))) \cdot 2^{(\log(m_k))^{O(\log\log\log(m_k))}}$. This complexity is in $2^{n^{o(1)}}$.

Theorem 5.6 (under GRH, Conjecture 5.2 and 5.3)**.** *There exists an infinite sequence of integers $m_1 < m_2 < \cdots$ such that Algorithm 17 has complexity*

$$\mathrm{Poly}([K_k : \mathbb{Q}], \log(\mathrm{N}(\mathfrak{a}))) \cdot 2^{(\log(m_k))^{O(\log\log\log(m_k))}}.$$

5.6 Numerical Results

We implemented the algorithms of this chapter, and use them to compute the structure of minus parts of class groups, provide support for Conjecture 5.2 on the generators of the minus part of the class group, and demonstrate the practicality of Algorithm 17. In [29], some justification in support of Conjecture 5.2 and Conjecture 5.3 is given. Below, we review existing data in the literature, and we discuss the novelty of the data provided via our techniques based on norm relations. We first consider Conjecture 5.3, which asserts that $h^+(m) \leq \mathrm{Poly}(m)$.

5.6.1 Numerical Data on h^+ (Conjecture 2)

Previous efforts The computation of the "plus part" of the class number of a cyclotomic field has been described as "notoriously hard" [67]. Therefore, little data is available in the literature to support Conjecture 5.3. Masley [51] used lower discriminant bounds proved by Odlyzko [57] to compute real class numbers. These results, later extended by Van der Linden [48], yielded the unconditional computation of the class numbers of all real cyclotomic fields of composite conductor $m \leq 200, \phi(m) \leq 72$ and $m \neq 148, 152$.

However, for fields of larger degree, the root discriminant becomes too large for Masley's method to handle. To overcome the problem of large root discriminant, Miller [52] established a lower bound on sums over prime ideals of Hilbert class field, which in turn establishes an upper bound on the class number. According to [52, Th. 1.1], for a composite integer $m \not\equiv 2 \bmod(4)$, the class number of the maximal real subfield of the m-th cyclotomic field $\mathbb{Q}(\zeta_m)$ is $h_m^+ = 1$ if $\phi(m) \leq 116$ and $m \neq 136, 145, 212$. Also, $h_m^+ = 2$ for $m = 136, 145$ and $h_{256}^+ = 1$. Under the GRH, Miller [52] was able to compute $h_{212}^+ = 5$ and $h_{512}^+ = 1$. The method was later extended to prime conductors in [53]. According to [53, Th. 3.1.1], for a prime number p one has $h_p^+ = 1$ if $p \leq 151$. Under GRH we have $h_p^+ = 1$ for $p \leq 263$ and $p \neq 163, 191, 229, 257$. Also $h_{163}^+ = 4$, $h_{191}^+ = 11$, $h_{229}^+ = 3$, and $h_{257}^+ = 3$.

Tables 4.1 and 4.2 of [53] provide the class numbers of the n-th layers of cyclotomic $\mathbb{Z}_p$-extensions over the rationals implying that $h^+_{169} = h^+_{289} = h^+_{361} = 1$. Also, under GRH $h^+_{243} = h^+_{529} = h^+_{841} = h^+_{961} = 1$

Great advances in the direction of computing class numbers of real cyclotomic fields were made by Schoof [67] who presented a table of the orders of certain subgroups of the class groups of the real cyclotomic fields for prime conductors less than 10000. Based on the Cohen–Lenstra heuristics, the probability that the main table presented in [67] is actually a table of class numbers is at least 98%. The largest order in this table is 130473 for the prime conductor 8017. So, according to the Schoof's table, with high probability for prime conductor m the class number of the real cyclotomic field is less than $17m$.

Our results Concrete results on h^+_m (even conditional to GRH) only exist for relatively small degrees, a few sporadic reasonable size degree ($m = 512, 529, 841, 961$), or probabilistically for certain large prime conductors. A Pari/GP [25] implementation of class group computation in abelian fields using norm relations was provided by [20]. Using this, we computed h^+_m (under the GRH) for many conductors for which this invariant was not known before. What is even more interesting about the numerical data we provide is that this method perform better for highly composite conductors, in contrast with the prime conductors for which some probabilistic data is already available. All in all, we were able to compute 149 values of h^+_m that do not appear to be previously known in the literature. We reached a maximum conductor of 2730. Our data supports Conjecture 5.3 which stipulates that h^+_m has moderate size. Besides the support of the CDW heuristics, this data is interesting in its own rights. Given the large number of values of h^+_m we calculated, we chose to disseminate the data in a new online database for invariants of cyclotomic fields, CycloDB [82]. To this day, CycloDB contains 362 values of $h^+(m)$, including the 149 that were not previously known. Note that each entry of the database contains $\mathrm{Cl}(\mathcal{O}_K)$, h_m, the factorization of h_m, h^-_m, h^+_m, and the regulator of the field. We will continue populating it in the future as this data is of general interest.

5.6.2 Numerical Data on the Minus Part (Conjecture 1)

Previous efforts Conjecture 5.2 is an ad-hoc assumption made for the first time in [29] that was not previously studied in the literature. In some sense, the numerical data we provide in this section is the first to ever put Conjecture 5.2 to the test strictly speaking. However, the authors of [28, 29] presented a rationale to justify Conjecture 5.2 based on existing numerical data. In [28, Prop. 6.1], it is proven that if a number s satisfies $s \geq r\left(\log\log_2(h^-) + \alpha\right)$ for a parameter $\alpha \geq 1$ and r the the number of $\mathbb{Z}[G]$-generators of $\mathrm{Cl}^-(\mathcal{O}_K)$, then the probability that s elements of $\mathrm{Cl}^-(\mathcal{O}_K)$ drawn uniformly at random generates $\mathrm{Cl}^-(\mathcal{O}_K)$ is at least $1 - O(2^{-\alpha})$. This means that if we know that the number of (not necessarily prime) generators of $\mathrm{Cl}^-(\mathcal{O}_K)$ is small, then on average few random elements are required to generate $\mathrm{Cl}^-(\mathcal{O}_K)$. The purpose

of [28, Prop. 6.1] is to relate Conjecture 5.2 with existing numerical data from the literature which concerns the number of generators of $\mathrm{Cl}^-(\mathcal{O}_K)$ rather than the number of prime generators of $\mathrm{Cl}^-(\mathcal{O}_K)$ (which is what is needed in Conjecture 5.2). However, to justify Conjecture 5.2 from [28, Prop. 6.1], one needs to make the extra unproven assumption that [28, Prop. 6.1] is still true even if we draw s short prime elements (as opposed to elements chosen uniformly at random). This extra heuristic seems reasonable, but it means that numerical results on the number of generators of $\mathrm{Cl}^-(\mathcal{O}_K)$ do not, on its own, directly support Conjecture 5.2.

Below, we recall known results on $\mathrm{Cl}^-(\mathcal{O}_K)$. Most of the existing literature concerns its cardinality h^-, but not the structure itself. Motivated by the results on divisibility properties of class numbers of cyclotomic fields, Kummer [42] was the first to carry out computations of relative class numbers of cyclotomic fields of prime conductor, for primes below 163. These calculations were extended by Lehmer and Masley [51] in 1978 to the primes $p \leq 509$. According to these results, h_p^- grows rapidly with p. For instance, h_{491}^- already has 138 decimal digits. Later, Fung, Granville and Williams [36] computed all h_p^- for $p \leq 3000$. Then, Shokrollahi [72] extended this result to all $p \leq 10000$.

Regarding the structure of the minus part, in [42], Kummer proved that $\mathrm{Cl}^-(\mathcal{O}_{\mathbb{Q}(\zeta_p)})$ is cyclic for every prime $p \leq 100$ and $p \neq 29, 41$. Furthermore, $\mathrm{Cl}^-(\mathcal{O}_{\mathbb{Q}(\zeta_{29})})$ and $\mathrm{Cl}^-(\mathcal{O}_{\mathbb{Q}(\zeta_{41})})$ are abelian groups of type $(2, 2, 2)$ and $(11, 11)$ respectively. Subsequently, Kummer's methods were refined by Tateyama [75], Horie and Ogura [39] and many other authors. Tateyama was able to compute the structure of $\mathrm{Cl}^-(\mathcal{O}_{\mathbb{Q}(\zeta_p)})$ for prime numbers p smaller than 227 except for seven cases. Horie and Ogura determined the structure of the minus part of any cyclotomic field with conductor less than 100. Later, Schoof [68] determined the structure of $\mathrm{Cl}^-(\mathcal{O}_{\mathbb{Q}(\zeta_p)})$ for $l \leq 509$. As an example, Schoof showed that $\mathrm{Cl}^-(\mathcal{O}_{\mathbb{Q}(\zeta_{491})})$ is isomorphic to a product of 6 cyclic groups. Also, Theorem 3 of [68] roughly states that for prime divisors p of $\ell - 1$, the p-part of $\mathrm{Cl}^-(\mathcal{O}_{\mathbb{Q}(\zeta_\ell)})$ is cyclic whenever it is small.

Our results We present the first experiments that directly test the validity of Conjecture 5.2 without relying on extra assumptions. Additionally, similar to the case of the provision of numerical data on h^+, our methods work for non-cyclic cyclotomic fields, which makes them valuable since all previous data used to justify Conjecture 5.2 was restricted to prime conductors. The results of our experiments are presented in Table 5. For each conductor m for which we tested Conjecture 5.2, we found the minimum B and d for which we could generate $\mathrm{Cl}^-(\mathcal{O}_K)$. Then we repeated 100 time the following experiment: draw d prime ideals of norm less than B uniformly at random, and check whether their conjugates generate $\mathrm{Cl}^-(\mathcal{O}_K)$. We report the corresponding probability. We also report the runtime of the computation of $\mathrm{Cl}^-(\mathcal{O}_K)$ in CPU hours, which is of independent interest. Conjecture 5.2 is of asymptotic nature, and hence difficult to justify

with a finite number of experiments, but the results of Table 5 are clearly consistent with the prediction of a moderate B and d with a high probability of generating $\mathrm{Cl}^-(\mathcal{O}_K)$.

5.6.3 Timings of the Subfield Variant of CDW

In Table 6 we report timings of our implementation of Algorithm 17, i.e. our subfield variant of the CDW method for the computation of mildly short vectors. We assume the generators for the minus part of the class group required in Step (6) of Algorithm 17 have been precomputed. Some of these timings can be seen in Table 5. We selected fields with conductor m ranging between $m = 23$ and $m = 198$. For each field, we report "$\mathrm{lb}\,\mathrm{N}$", the bit size of the algebraic norm of the input ideal, "$\mathrm{lb}\,\mathrm{N}_{\mathrm{svp}}$", the bit size of the algebraic norm of the short generator of the principal ideal found in Step (10) of Algorithm 17, "t_{cpm}", the time to solve the Close Principal Multiple problem, "t_{pip}", the time to solve the Principal Ideal Problem, and "t_{svp}", the time to find the short generator of Step (10). Timings are reported in CPU seconds unless otherwise stated.

Table 5. Experiments on $\mathrm{Cl}^-(\mathcal{O}_K)$

m	n	$\mathrm{Cl}^-(\mathcal{O}_K)$	B	d	prob.	time	m	n	$\mathrm{Cl}^-(\mathcal{O}_K)$	B	d	prob.	time
23	22	[3]	47	2	100.0	0.01	105	48	[13]	211	1	100.0	0.02
46	22	[3]	47	2	100.0	0.01	112	48	[3, 156]	113	3	100.0	0.22
39	24	[2]	13	1	100.0	0.01	130	48	[2, 2, 4, 4]	131	1	100.0	0.18
52	24	[3]	13	1	100.0	0.01	144	48	[13, 39]	433	4	100.0	0.05
56	24	[2]	8	1	100.0	0.01	53	52	[4889]	107	2	100.0	0.24
72	24	[3]	9	1	100.0	0.01	106	52	[4889]	107	2	100.0	0.24
78	24	[2]	13	1	100.0	0.01	81	54	[2593]	163	2	100.0	0.46
29	28	[2, 2, 2]	59	2	100.0	0.01	162	54	[2593]	163	2	100.0	0.34
58	28	[2, 2, 2]	59	2	100.0	0.01	87	56	[8, 8, 24]	523	2	100.0	0.4
31	30	[9]	32	2	100.0	0.01	116	56	[8, 8, 168]	233	2	100.0	0.33
62	30	[9]	32	2	100.0	0.01	174	56	[8, 8, 24]	523	2	100.0	0.38
51	32	[5]	103	1	100.0	0.01	59	58	[41241]	709	2	100.0	3.39
64	32	[17]	193	2	100.0	0.01	118	58	[41241]	709	2	100.0	2.65
68	32	[8]	137	2	100.0	0.02	61	60	[76301]	367	2	100.0	4.9
96	32	[3, 3]	97	2	100.0	0.01	77	60	[4, 4, 4, 20]	463	1	100.0	0.29
102	32	[5]	103	1	100.0	0.01	93	60	[6795]	373	2	100.0	0.09
37	36	[37]	149	2	100.0	0.01	99	60	[31, 93]	199	1	100.0	0.12
57	36	[9]	229	2	100.0	0.01	122	60	[76301]	367	2	100.0	6.34
63	36	[7]	64	1	100.0	0.01	124	60	[2, 22878]	373	3	100.0	0.33
74	36	[37]	149	2	100.0	0.01	154	60	[4, 4, 4, 20]	463	1	100.0	0.31
76	36	[19]	229	1	100.0	0.01	186	60	[6795]	373	2	100.0	0.08
108	36	[19]	109	2	100.0	0.01	198	60	[31, 93]	199	1	100.0	0.09
114	36	[9]	229	2	100.0	0.01	85	64	[6205]	1021	2	100.0	0.14
126	36	[7]	64	1	100.0	0.01	128	64	[359057]	257	2	100.0	45.3
41	40	[11, 11]	83	2	100.0	0.01	170	64	[6205]	1021	2	100.0	0.12
55	40	[10]	11	1	100.0	0.03	192	64	[3, 20451]	193	2	100.0	1.78
75	40	[11]	151	2	100.0	0.01	91	72	[4, 13468]	547	2	100.0	0.07
82	40	[11, 11]	83	2	100.0	0.01	95	72	[107692]	571	2	98.0	2.74
88	40	[55]	89	1	100.0	0.02	135	72	[75961]	271	2	100.0	0.94
100	40	[55]	101	2	100.0	0.02	148	72	[4827501]	593	3	100.0	0.51
110	40	[10]	11	1	100.0	0.03	152	72	[19, 171, 513]	457	2	100.0	0.51
132	40	[11]	397	2	100.0	0.01	190	72	[107692]	571	2	96.0	2.54
150	40	[11]	151	2	100.0	0.01	123	80	[8, 8, 88, 1496]	739	1	100.0	5.16
43	42	[211]	173	2	100.0	0.04	164	80	[11, 7528840]	821	2	100.0	14.43
49	42	[43]	197	2	100.0	0.02	165	80	[92620]	331	2	100.0	7.16
86	42	[211]	173	2	100.0	0.04	176	80	[5, 5874275]	353	1	100.0	5.27
98	42	[43]	197	2	100.0	0.03	129	84	[37821539]	1033	2	100.0	64.45
69	44	[69]	139	1	100.0	0.02	147	84	[5874617]	883	2	100.0	2.23
92	44	[201]	277	1	100.0	0.02	172	84	[2, 396326786]	173	1	100.0	54.66
138	44	[69]	139	1	100.0	0.02	196	84	[82708823]	197	2	100.0	14.79
47	46	[695]	283	2	100.0	0.13	184	88	[67, 22181154]	1289	2	100.0	16.14
94	46	[695]	283	2	100.0	0.17	189	108	[105778197511]	379	1	100.0	2.6
65	48	[2, 2, 4, 4]	131	1	100.0	0.18							

Table 6. Computation of mildly short vectors with our subfield CDW variant.

m	n	lb N	lb N_{svp}	t_{cpm}	t_{pip}	t_{svp}	m	n	lb N	lb N_{svp}	t_{cpm}	t_{pip}	t_{svp}
23	22	51	68	3.05	0.01	0.98	98	42	55	88	3.22	0.1	0.99
46	22	55	72	3.06	0.01	1.04	69	44	54	78	6.08	0.1	0.77
39	24	46	49	1.44	0.01	0.89	92	44	52	94	479.72	0.12	0.71
52	24	52	56	1.17	0.01	0.96	138	44	54	78	5.62	15.1	1.28
56	24	53	56	1.14	0.01	0.92	47	46	55	100	3.26	0.29	1.14
72	24	58	68	3.23	0.01	0.68	94	46	50	104	6.21	0.19	1.12
78	24	52	56	1.22	0.01	0.72	65	48	53	152	1342.92	57.64	1.37
29	28	54	77	2.96	0.03	1.22	112	48	51	264	1894.87	45.02	1.14
58	28	53	75	2.81	0.02	1.01	144	48	53	175	1988.98	71.09	1.58
31	30	56	66	1.04	0.03	1.0	87	56	52	139	5.79	73.53	1.49
62	30	58	79	2.57	0.03	0.84	116	56	56	294	1644.03	81.35	2.07
51	32	55	62	1.48	0.03	0.92	174	56	59	153	7.13	66.76	1.81
64	32	56	81	3.51	0.05	1.27	77	60	56	243	1496.5	69.37	1.3
68	32	56	172	289.47	0.08	0.7	93	60	50	133	7.68	101.88	2.13
96	32	53	79	4.14	0.03	0.73	99	60	55	124	6.76	78.57	1.29
102	32	60	73	3.82	0.03	0.69	124	60	48	147	5.7	124.02	2.25
37	36	52	78	3.15	0.05	1.01	186	60	57	140	7.33	0.47	1.57
57	36	46	70	3.69	0.05	1.12	198	60	58	132	7.62	0.27	0.75
63	36	51	63	4.06	0.04	0.73	85	64	58	155	8.17	1023.02	9.26
74	36	55	81	3.29	0.05	0.98	128	64	50	132	7.92	518.03	8.12
76	36	47	80	202.89	0.05	1.1	192	64	51	131	6.75	1117.83	16.92
108	36	59	83	3.9	0.07	0.76	91	72	51	137	10.63	73.88	2.6
114	36	55	83	4.41	0.06	0.9	95	72	53	139	8.14	528.04	3.21
41	40	57	92	2.78	0.08	1.26	135	72	52	136	8.5	503.78	3.62
55	40	52	69	4.01	0.07	0.81	148	72	54	160	7.24	368.24	3.45
75	40	57	71	1.94	0.07	0.83	152	72	57	148	6.18	950.51	9.39
82	40	55	86	3.39	0.09	1.45	123	80	54	127	12.38	1834.57	11.32
88	40	55	75	3.96	0.06	0.73	164	80	58	153	11.35	1897.85	9.99
100	40	58	97	416.62	0.11	0.69	176	80	55	126	14.97	5500.97	39.56
150	40	55	79	4.59	0.07	0.9	172	84	52	118	13.61	1713.64	7.73
43	42	52	82	4.18	0.15	1.09	196	84	57	144	12.7	18808.78	97.63
49	42	58	73	1.01	0.14	1.35	184	88	58	157	17.82	6611.68	25.21
86	42	57	105	3.74	0.17	1.03							

CHAPTER 6
CONCLUSION

In Chapter 4 we described how to use the norm relations of [20] to reduce instances of the PIP and ideal decomposition in a Galois number field to subfield problems. We first illustrated this in Algorithm 5 where we show how to resolve the PIP using recursive computation of S-unit groups. We analyzed the complexity of Algorithm 5 and showed in Section 4.1.1 that this method achieves a polynomial improvement over the state of the art in almost all cyclotomic fields, and in certain families of cyclotomic fields the improvement is in fact superpolynomial. In Section 4.3 we show how norm relations can be used to perform ideal decomposition by recursively decomposing ideals in subfields with the same asymptotic improvements, and show how this can also be used for resolution of PIP instances in Section 4.3.3. Finally, in Section 4.4 we give a third algorithm for resolving the PIP using norm relations in Algorithm 13. This algorithm avoids the S-unit group computations of the previous PIP algorithms, resulting in significant improvements in practical performance. In Section 4.4.2 we describe our implementation, capable of resolving instances of the PIP in fields of degree up to 1800, breaking all previous records by a significant margin.

In Chapter 5 we moved to cryptographic applications. We recalled the results of [28, 29] which allow for an efficient quantum reduction of the search for a mildly short vector in an arbitrary cyclotomic ideal lattice to the PIP, and applied our variants of PIP and ideal decomposition using norm relations to improve the classical asymptotic cost of this reduction, which is subexponential. In particular, Algorithm 15 describes how to compute the structure of the minus part of the class group of a cyclotomic field using norm relations, and Algorithm 16 allows for the computation of a set of generators for the minus part. In Section 5.4 we describe how to use these results to resolve the Close Principal Multiple problem and provide a classical variant of the CDW reduction whose asymptotic cost is reduced to the cost of S-unit group computation and ideal decomposition in subfields. We implemented the algorithms we described provided numerical data in Section 5.6. We provide evidence for conjectures regarding the class group of cyclotomic fields that are crucial for the efficiency of the CDW reduction, compute minus parts of class groups for cyclotomic fields of degree up to 108, and find mildly short vectors in non-principal cyclotomic ideal lattices in degrees up to 88.

6.1 Future Work

The existence of norm relations allows reducing hard problems such as the PIP to subfield computations. However, when we apply these results to large degree fields admitting norm relations with very small subfields, the large gap in degree means the subfield elements, which are given by relative norms, are often very large. We generally do not evaluate the products of these relative norms, and use compact representations as described in Section 2.4 to reduce the size of elements which we may need to compute a d-th root of, but the sheer size nevertheless takes a toll on the practical performance. Indeed, in Table 3 we give the time for each step in the recursive PIP resolution using norm relations, and in degree 1800 computing subfield unit groups and resolving subfield PIP instances, the most difficult step asymptotically, took only approximately 10% of the total time. However, computing relative norms took 22% of the time, and the compact representations and root computations 56% of the time. Improving these routines, either by optimizing the implementation or theoretical improvements, would have a significant impact on runtime and allow computations in even larger degree fields. This would apply equally as well to ideal decomposition and the search for mildly short vectors in cyclotomic ideal lattices.

REFERENCES

[1] L. M. Adleman. "Factoring numbers using singular integers". In: *Symposium on the Theory of Computing*. 1991.

[2] M. Ajtai, R. Kumar, and D. Sivakumar. "A sieve algorithm for the shortest lattice vector problem". In: *Proceedings on 33rd Annual ACM Symposium on Theory of Computing, July 6-8, 2001, Heraklion, Crete, Greece*. Ed. by Jeffrey Scott Vitter, Paul G. Spirakis, and Mihalis Yannakakis. ACM, 2001, pp. 601–610. DOI: 10.1145/380752.380857. URL: https://doi.org/10.1145/380752.380857.

[3] E. Artin and J. Tate. *Class field theory*. Reprinted with corrections from the 1967 original. AMS Chelsea Publishing, Providence, RI, 2009, pp. viii+194. ISBN: 978-0-8218-4426-7.

[4] E. Bach. "Explicit bounds for primality testing and related problems". In: *Math. Comp.* 55.191 (1990), pp. 355–380.

[5] J. Bauch et al. "Short generators without quantum computers: the case of multiquadratics". In: *Advances in cryptology—EUROCRYPT 2017. Part I*. Vol. 10210. Lecture Notes in Comput. Sci. Springer, Cham, 2017, pp. 27–59. DOI: 10.1007/978-3-319-56620-7_2. URL: https://doi.org/10.1007/978-3-319-56620-7_2.

[6] O. Bernard and A. Roux-Langlois. "Twisted-PHS: using the product formula to solve approx-SVP in ideal lattices". In: *Advances in cryptology—ASIACRYPT 2020. Part II*. Vol. 12492. Lecture Notes in Comput. Sci. Springer, Cham, [2020] ©2020, pp. 349–380. DOI: 10.1007/978-3-030-64834-3_12. URL: https://doi.org/10.1007/978-3-030-64834-3_12.

[7] O. Bernard et al. "Log-S-unit lattices using Explicit Stickelberger Generators to solve Approx Ideal-SVP". In: *IACR Cryptol. ePrint Arch.* (2021), p. 1384. URL: https://eprint.iacr.org/2021/1384.

[8] D. Bernstein and T. Lange. "Non-randomness of S-unit lattices". In: *IACR Cryptol. ePrint Arch.* (2021), p. 1428. URL: https://eprint.iacr.org/2021/1428.

[9] J. Bezanson et al. "Julia: A fresh approach to numerical computing". In: *SIAM review* 59.1 (2017), pp. 65–98. URL: https://doi.org/10.1137/141000671.

[10] J.-F. Biasse. "Approximate short vectors in ideal lattices of $\mathbb{Q}(\zeta_{p^e})$ with precomputation of $\mathrm{Cl}(O_K)$". In: *Selected areas in cryptography—SAC 2017*. Vol. 10719. Lecture Notes in Comput. Sci. Springer, Cham, 2018, pp. 374–393. DOI: 10.1007/978-3-319-72565-9_19. URL: https://doi.org/10.1007/978-3-319-72565-9_19.

[11] J.-F. Biasse. "Subexponential time relations in the class group of large degree number fields". In: *Adv. Math. Commun.* 8.4 (2014), pp. 407–425. ISSN: 1930-5346. DOI: 10.3934/amc.2014.8.407. URL: https://doi.org/10.3934/amc.2014.8.407.

[12] J.-F. Biasse and M. Erukulangara. "A proof of the conjectured run time of the Hafner-McCurley class group algorithm". In: *Advances in Mathematics of Communications* 0.0 (2021), pp. 0-0. ISSN: 1930-5346. DOI: 10.3934/amc.2021055. URL: /article/id/40b2f18b-c250-4134-921a-f130d3278494.

[13] J.-F. Biasse and C. Fieker. "Improved techniques for computing the ideal class group and a system of fundamental units in number fields". In: *ANTS X—Proceedings of the Tenth Algorithmic Number Theory Symposium*. Vol. 1. Open Book Ser. Math. Sci. Publ., Berkeley, CA, 2013, pp. 113–133. DOI: 10.2140/obs.2013.1.113. URL: https://doi.org/10.2140/obs.2013.1.113.

[14] J.-F. Biasse and C. Fieker. "Subexponential class group and unit group computation in large degree number fields". In: *LMS J. Comput. Math.* 17.suppl. A (2014), pp. 385–403. DOI: 10.1112/S1461157014000345. URL: https://doi.org/10.1112/S1461157014000345.

[15] J.-F. Biasse, C. Fieker, and T. Hofmann. "On the computation of the HNF of a module over the ring of integers of a number field". In: *J. Symb. Comput.* 80 (2017), pp. 581–615. DOI: 10.1016/j.jsc.2016.07.027. URL: https://doi.org/10.1016/j.jsc.2016.07.027.

[16] J.-F. Biasse and F. Song. "Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields". In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*. ACM, New York, 2016, pp. 893–902. DOI: 10.1137/1.9781611974331.ch64. URL: https://doi.org/10.1137/1.9781611974331.ch64.

[17] J.-F. Biasse and C. van Vredendaal. "Fast multiquadratic S-unit computation and application to the calculation of class groups". In: *Proceedings of the Thirteenth Algorithmic Number Theory Symposium*. Vol. 2. Open Book Ser. Math. Sci. Publ., Berkeley, CA, 2019, pp. 103–118.

[18] J.-F. Biasse et al. "Computing generator in cyclotomic integer rings". In: *Advances in cryptology—EUROCRYPT 2017. Part I*. Vol. 10210. Lecture Notes in Comput. Sci. Springer, Cham, 2017, pp. 60–88. DOI: 10.1007/978-3-319-56620-7_3. URL: https://doi.org/10.1007/978-3-319-56620-7_3.

[19] J.-F. Biasse et al. “Mildly Short Vectors in Ideals of Cyclotomic Fields Without Quantum Computers”. In: *Mathematical Cryptology* 2.1 (Nov. 2022), pp. 84–107. URL: https://journals.flvc.org/mathcryptology/article/view/132573.

[20] J.-F. Biasse et al. “Norm relations and computational problems in number fields”. In: *J. Lond. Math. Soc. (2)* 105.4 (2022), pp. 2373–2414. ISSN: 0024-6107. DOI: 10.1112/jlms.12563. URL: https://doi.org/10.1112/jlms.12563.

[21] R. Brauer. “Beziehungen zwischen Klassenzahlen von Teilkörpern eines galoisschen Körpers”. In: *Math. Nachr.* 4 (1951), pp. 158–174. ISSN: 0025-584X. DOI: 10.1002/mana.3210040116. URL: https://doi.org/10.1002/mana.3210040116.

[22] J. Buchmann. “A subexponential algorithm for the determination of class groups and regulators of algebraic number fields”. In: *Séminaire de Théorie des Nombres, Paris 1988–1989*. Vol. 91. Progr. Math. Birkhäuser Boston, Boston, MA, 1990, pp. 27–41.

[23] P. Campbell, M. Groves, and D. Shepherd. *SOLILOQUY: a cautionary tale*. online draft available at http://docbox.etsi.org/Workshop/2014/201410_CRYPTO/S07_Systems_and_Attacks/S07_Groves_Annex.pdf.

[24] H. Cohen. *Advanced topics in computational number theory*. Vol. 193. Graduate Texts in Mathematics. Springer-Verlag, New York, 2000, pp. xvi+578. ISBN: 0-387-98727-4. DOI: 10.1007/978-1-4419-8489-0. URL: https://doi.org/10.1007/978-1-4419-8489-0.

[25] H. Cohen. *Pari*. Available at http://pari.math.u-bordeaux.fr/.

[26] H. Cohen, F. Diaz Y Diaz, and M. Olivier. “Subexponential Algorithms for Class Group and Unit Computations”. In: *Journal of Symbolic Computation* 24.3-4 (1997), pp. 433–441. ISSN: 0747-7171. DOI: http://dx.doi.org/10.1006/jsco.1996.0143. URL: http://www.sciencedirect.com/science/article/pii/S0747717196901431.

[27] H. Cohen, F. Diaz y Diaz, and M. Olivier. “Subexponential algorithms for class group and unit computations”. In: vol. 24. 3-4. Computational algebra and number theory (London, 1993). 1997, pp. 433–441. DOI: 10.1006/jsco.1996.0143. URL: https://doi.org/10.1006/jsco.1996.0143.

[28] R. Cramer, L. Ducas, and B. Wesolowski. “Mildly Short Vectors in Cyclotomic Ideal Lattices in Quantum Polynomial Time”. In: *J. ACM* 68.2 (2021), 8:1–8:26. DOI: 10.1145/3431725. URL: https://doi.org/10.1145/3431725.

[29] R. Cramer, L. Ducas, and B. Wesolowski. "Short Stickelberger class relations and application to ideal-SVP". In: *Advances in cryptology—EUROCRYPT 2017. Part I*. Vol. 10210. Lecture Notes in Comput. Sci. Springer, Cham, 2017, pp. 324–348. DOI: 10.1007/978-3-319-56620-7_12. URL: https://doi.org/10.1007/978-3-319-56620-7_12.

[30] R. Cramer et al. "Recovering short generators of principal ideals in cyclotomic rings". In: *Advances in cryptology—EUROCRYPT 2016. Part II*. Vol. 9666. Lecture Notes in Comput. Sci. Springer, Berlin, 2016, pp. 559–585. DOI: 10.1007/978-3-662-49896-5_20. URL: https://doi.org/10.1007/978-3-662-49896-5_20.

[31] G. Lejeune Dirichlet. "Recherches sur les formes quadratiques à coëfficients et à indéterminées complexes. Première partie". In: *J. Reine Angew. Math.* 24 (1842), pp. 291–371. ISSN: 0075-4102. DOI: 10.1515/crll.1842.24.291. URL: https://doi.org/10.1515/crll.1842.24.291.

[32] P. Erdős, C. Pomerance, and E. Schmutz. "Carmichael's lambda function". In: *Acta Arith.* 58.4 (1991), pp. 363–385. ISSN: 0065-1036. DOI: 10.4064/aa-58-4-363-385.

[33] C. Fieker and C. Friedrichs. "On reconstruction of algebraic numbers". In: *Algorithmic number theory (Leiden, 2000)*. Vol. 1838. Lecture Notes in Comput. Sci. Springer, Berlin, 2000, pp. 285–296. DOI: 10.1007/10722028_16. URL: https://doi.org/10.1007/10722028_16.

[34] C. Fieker, T. Hofmann, and C. Sircana. "On the construction of class fields". In: *Proceedings of ANTS XIII*. 2019, pp. 239–255.

[35] C. Fieker et al. "Nemo/Hecke: Computer Algebra and Number Theory Packages for the Julia Programming Language". In: *Proceedings of ISSAC 2017*. 2017, pp. 157–164.

[36] G. Fung, A. Granville, and H. Williams. "Computation of the first factor of the class number of cyclotomic fields". In: *Journal of Number Theory* 42.3 (1992), pp. 297–312. ISSN: 0022-314X. DOI: https://doi.org/10.1016/0022-314X(92)90095-7. URL: https://www.sciencedirect.com/science/article/pii/0022314X92900957.

[37] Carl Friedrich Gauss. *Disquisitiones arithmeticae*. Translated and with a preface by Arthur A. Clarke, Revised by William C. Waterhouse, Cornelius Greither and A. W. Grootendorst and with a preface by Waterhouse. Springer-Verlag, New York, 1986, pp. xx+472. ISBN: 0-387-96254-9.

[38] James L. Hafner and Kevin S. McCurley. "A rigorous subexponential algorithm for computation of class groups". In: *J. Amer. Math. Soc.* 2.4 (1989), pp. 837–850. ISSN: 0894-0347. DOI: 10.2307/1990896. URL: https://doi.org/10.2307/1990896.

[39] K. Horie and H. Ogura. “On the Ideal Class Groups of Imaginary Abelian Fields with Small Conductor”. In: *Transactions of the American Mathematical Society* 347.7 (1995), pp. 2517–2532. ISSN: 00029947. URL: http://www.jstor.org/stable/2154835 (visited on 06/08/2022).

[40] Michael J. Jacobson Jr. and Hugh C. Williams. *Solving the Pell equation*. CMS Books in Mathematics/Ouvrages de Mathématiques de la SMC. Springer, New York, 2009, pp. xx+495. ISBN: 978-0-387-84922-5.

[41] M. Kirschmer. “Definite quadratic and hermitian forms with small class number”. Habilitationsschrift. RWTH Aachen University, 2016.

[42] E. Kummer. “Memoir on the theory of complex numbers composed of roots of unity and integers.” fre. In: *Journal of Pure and Applied Mathematics* (1851), pp. 377–498. URL: http://eudml.org/doc/235621.

[43] S. Kuroda. “Über die Klassenzahlen algebraischer Zahlkörper”. In: *Nagoya Math. J.* 1 (1950), pp. 1–10. ISSN: 0027-7630. URL: http://projecteuclid.org/euclid.nmj/1118764698.

[44] T. Laarhoven. “Sieving for Closest Lattice Vectors (with Preprocessing)”. In: *Selected Areas in Cryptography - SAC 2016 - 23rd International Conference, St. John’s, NL, Canada, August 10-12, 2016, Revised Selected Papers*. Ed. by R. Avanzi and H. Heys. Vol. 10532. Lecture Notes in Computer Science. Springer, 2016, pp. 523–542. DOI: 10.1007/978-3-319-69453-5_28. URL: https://doi.org/10.1007/978-3-319-69453-5%5C_28.

[45] A.K. Lenstra, H.W. Lenstra, and L. Lovász. “Factoring polynomials with rational coefficients”. In: *Mathematische Annalen* 261 (1982), pp. 515–534.

[46] H. W. Lenstra Jr. “On the calculation of regulators and class numbers of quadratic fields”. In: *Number theory days, 1980 (Exeter, 1980)*. Vol. 56. London Math. Soc. Lecture Note Ser. Cambridge Univ. Press, Cambridge, 1982, pp. 123–150.

[47] Andrea Lesavourey, Thomas Plantard, and Willy Susilo. “Short principal ideal problem in multicubic fields”. In: *J. Math. Cryptol.* 14.1 (2020), pp. 359–392. ISSN: 1862-2976. DOI: 10.1515/jmc-2019-0028. URL: https://doi.org/10.1515/jmc-2019-0028.

[48] F. J. van der Linden. “Class Number Computations of Real Abelian Number Fields”. In: *Mathematics of Computation* 39.160 (1982), pp. 693–707. ISSN: 00255718, 10886842. URL: http://www.jstor.org/stable/2007347 (visited on 06/01/2022).

[49] V. Lyubashevsky, C. Peikert, and O. Regev. "On Ideal Lattices and Learning with Errors over Rings". In: *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Monaco / French Riviera, May 30 - June 3, 2010. Proceedings*. Ed. by H. Gilbert. Vol. 6110. Lecture Notes in Computer Science. Springer, 2010, pp. 1–23. DOI: 10.1007/978-3-642-13190-5_1. URL: https://doi.org/10.1007/978-3-642-13190-5%5C_1.

[50] V. Lyubashevsky, C.s Peikert, and O. Regev. "A Toolkit for Ring-LWE Cryptography". In: *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*. Ed. by T.s Johansson and P. Nguyen. Vol. 7881. Lecture Notes in Computer Science. Springer, 2013, pp. 35–54. DOI: 10.1007/978-3-642-38348-9_3. URL: https://doi.org/10.1007/978-3-642-38348-9%5C_3.

[51] J. Masley. "Class numbers of real cyclic number fields with small conductor". In: *Compositio Mathematica* 37 (1978), pp. 297–319.

[52] J. Miller. "Class numbers of real cyclotomic fields of composite conductor". In: *LMS Journal of Computation and Mathematics* 17 (2014), pp. 404–417.

[53] J. Miller. "Class numbers of totally real number fields". PhD thesis. Rutgers University, 2015.

[54] J. Neukirch. *Algebraic number theory*. Comprehensive Studies in Mathematics. ISBN 3-540-65399-6. Springer-Verlag, 1999.

[55] J. Neukirch, A. Schmidt, and K. Wingberg. *Cohomology of number fields*. Second. Vol. 323. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer-Verlag, Berlin, 2008, pp. xvi+825. ISBN: 978-3-540-37888-4. DOI: 10.1007/978-3-540-37889-1. URL: https://doi.org/10.1007/978-3-540-37889-1.

[56] O. T. O'Meara. *Introduction to quadratic forms*. Die Grundlehren der mathematischen Wissenschaften, Bd. 117. Academic Press Inc. Publishers, New York; Springer-Verlag, Berlin-Göttingen-Heidelberg, 1963, pp. xi+342.

[57] A. Odlyzko. "Bounds for discriminants and related estimates for class numbers, regulators and zeros of zeta functions : a survey of recent results". In: *Journal de Theorie des Nombres de Bordeaux* 2 (1990), pp. 119–141.

[58] A. M. Odlyzko. "Lower bounds for discriminants of number fields". In: *Acta Arith.* 29.3 (1976), pp. 275–297. ISSN: 0065-1036. DOI: 10.4064/aa-29-3-275-297. URL: https://doi.org/10.4064/aa-29-3-275-297.

[59] A. Pellet-Mary, G. Hanrot, and D. Stehlé. "Approx-SVP in Ideal Lattices with Pre-processing". In: *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part II*. Ed. by Y. Ishai and V. Rijmen. Vol. 11477. Lecture Notes in Computer Science. Springer, 2019, pp. 685–716. DOI: 10.1007/978-3-030-17656-3_24. URL: https://doi.org/10.1007/978-3-030-17656-3%5C_24.

[60] M. Pohst and H. Zassenhaus. *Algorithmic algebraic number theory*. Vol. 30. Encyclopedia of Mathematics and its Applications. Cambridge University Press, Cambridge, 1989, pp. xiv+465. ISBN: 0-521-33060-2. DOI: 10.1017/CBO9780511661952. URL: https://doi.org/10.1017/CBO9780511661952.

[61] Michael Pohst, ed. *Algorithmic methods in algebra and number theory*. Reprint of J. Symbolic Comput. **4** (1987), no. 1. Academic Press Inc. [Harcourt Brace Jovanovich, Publishers], London, 1987, pp. iv+135. ISBN: 0-12-559190-X.

[62] O. Regev. "On lattices, learning with errors, random linear codes, and cryptography". In: *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*. Ed. by H. Gabow and R. Fagin. ACM, 2005, pp. 84–93. DOI: 10.1145/1060590.1060603. URL: https://doi.org/10.1145/1060590.1060603.

[63] Oded Regev. "On Lattices, Learning with Errors, Random Linear Codes, and Cryptography". In: *J. ACM* 56.6 (Sept. 2009). ISSN: 0004-5411. DOI: 10.1145/1568318.1568324. URL: https://doi.org/10.1145/1568318.1568324.

[64] R. L. Rivest, A. Shamir, and L. Adleman. "A method for obtaining digital signatures and public-key cryptosystems". In: *Comm. ACM* 21.2 (1978), pp. 120–126. ISSN: 0001-0782. DOI: 10.1145/359340.359342. URL: https://doi.org/10.1145/359340.359342.

[65] C. P. Schnorr and M. Euchner. "Lattice Basis Reduction: Improved Practical Algorithms and Solving Subset Sum Problems". In: *Math. Program.* 66.2 (Sept. 1994), pp. 181–199. ISSN: 0025-5610. DOI: 10.1007/BF01581144. URL: http://dx.doi.org/10.1007/BF01581144.

[66] C.P. Schnorr. "A hierarchy of polynomial time lattice basis reduction algorithms". In: *Theoretical Computer Science* 53.2-3 (1987), pp. 201–224. ISSN: 0304-3975. DOI: http://dx.doi.org/10.1016/0304-3975(87)90064-8. URL: http://www.sciencedirect.com/science/article/pii/0304397587900648.

[67] R. Schoof. "Class numbers of real cyclotomic fields of prime conductor". In: *Math. Comput.* 72 (2003), pp. 913–937.

[68] R. Schoof. "Minus class groups of the fields of the l-th roots of unity". In: *Math. Comput.* 67 (1998), pp. 1225–1245.

[69] R. J. Schoof. "Quadratic fields and factorization". In: *Computational methods in number theory, Part II*. Vol. 155. Math. Centre Tracts. Math. Centrum, Amsterdam, 1982, pp. 235–286.

[70] Daniel Shanks. "Class number, a theory of factorization, and genera". In: *1969 Number Theory Institute (Proc. Sympos. Pure Math., Vol. XX, State Univ. New York, Stony Brook, N.Y., 1969)*. Amer. Math. Soc., Providence, R.I., 1971, pp. 415–440.

[71] Daniel Shanks. "Five number-theoretic algorithms". In: *Proceedings of the Second Manitoba Conference on Numerical Mathematics (Univ. Manitoba, Winnipeg, Man., 1972)*. Congressus Numerantium, No. VII. Utilitas Math., Winnipeg, Man., 1973, pp. 51–70.

[72] A. Shokrollahi. "Relative class number of imaginary Abelian fields of prime conductor below 10000". In: *Math. Comput.* 68 (1999), pp. 1717–1728.

[73] Peter W. Shor. "Algorithms for quantum computation: discrete logarithms and factoring". In: *35th Annual Symposium on Foundations of Computer Science (Santa Fe, NM, 1994)*. IEEE Comput. Soc. Press, Los Alamitos, CA, 1994, pp. 124–134. DOI: 10.1109/SFCS.1994.365700. URL: https://doi.org/10.1109/SFCS.1994.365700.

[74] D. Simon. "Équations dans les corps de nombres et discriminants minimaux". PhD thesis. Université Bordeaux I, 1998.

[75] K. Tateyama. "On the ideal class groups of some cyclotomic fields". In: *Proceedings of the Japan Academy, Series A, Mathematical Sciences* 58.7 (1982), pp. 333–335. DOI: 10.3792/pjaa.58.333. URL: https://doi.org/10.3792/pjaa.58.333.

[76] C. Thiel. "On the complexity of some problems in algorithmic algebraic number theory". PhD thesis. Universität des Saarlandes, 1995.

[77] Hideo Wada. "On the class number and the unit group of certain algebraic number fields". In: *J. Fac. Sci. Univ. Tokyo Sect. I* 13 (1966), 201–209 (1966). ISSN: 0368-2269.

[78] C. D. Walter. "Kuroda's class number relation". In: *Acta Arith.* 35.1 (1979), pp. 41–51. ISSN: 0065-1036. DOI: 10.4064/aa-35-1-41-51. URL: https://doi.org/10.4064/aa-35-1-41-51.

[79] Lawrence C. Washington. *Introduction to cyclotomic fields*. Second. Vol. 83. Graduate Texts in Mathematics. Springer-Verlag, New York, 1997, pp. xiv+487. ISBN: 0-387-94762-0. DOI: 10.1007/978-1-4612-1934-7. URL: https://doi.org/10.1007/978-1-4612-1934-7.

[80] Heinrich Weber. "Literatur-Berichte: Lehrbuch der Algebra". In: *Monatsh. Math. Phys.* 9.1 (1898). von Heinrich Weber. II. Band. Braunschweig. Friedrich Vieweg und Sohn 1897, A23–A26. ISSN: 1812-8076. DOI: 10.1007/BF01707905. URL: https://doi.org/10.1007/BF01707905.

[81] B. Wesolowski. "Arithmetic and geometric structures in cryptography". PhD thesis. Lausanne, Switzerland: Ecole Polytechnique Federale de Lausanne, 2018.

[82] W. Youmans. *CycloDB*. https://www.cyclodb.org. [Online; accessed 29 March 2023]. 2023.

APPENDIX A
JMC LICENSE

The material in Chapter 5 was adapted from joint work with Jean-François Biasse, Muhammed Rashad Erukulangara, Claus Fieker, Tommy Hofmann, and myself, published in the Journal of Mathematical Cryptology (JMC) [19]:

J.-F. Biasse et al. "Mildly Short Vectors in Ideals of Cyclotomic Fields Without Quantum Computers". In: *Mathematical Cryptology* 2.1 (Nov. 2022), pp. 84–107. URL: https://journals.flvc.org/mathcryptology/article/view/132573

JMC allows all authors to retain copyright, governed by the Creative-Commons Attribution Only License. The following is taken from the Instruction for Authors file available on the journal website.

DE GRUYTER

Copyright

All authors retain copyright, unless – due to their local circumstances – their work is not copyrighted. The copyrights are governed by the Creative-Commons Attribution Only license (CC-BY) which is compliant with Plan-S. Scanned copy of License to Publish should be sent to the journal, as soon as possible.