Short Stickelberger Class Relations and Application to Ideal-SVP

Ronald Cramer[1,2], Léo Ducas[1] and Benjamin Wesolowski[3]

[1] Cryptology Group, CWI, Amsterdam, The Netherlands
[2] Mathematical Institute, Leiden University, The Netherlands
[3] École Polytechnique Fédérale de Lausanne, EPFL IC LACAL, Switzerland

Abstract. The worst-case hardness of finding short vectors in ideals of cyclotomic number fields (Ideal-SVP) is a central matter in lattice based cryptography. Assuming the worst-case hardness of Ideal-SVP allows to prove the Ring-LWE and Ring-SIS assumptions, and therefore to prove the security of numerous cryptographic schemes and protocols — including key-exchange, digital signatures, public-key encryption and fully-homomorphic encryption.

A series of recent works has shown that *Principal* Ideal-SVP is not always as hard as finding short vectors in general lattices, and some schemes were broken using quantum algorithms — the Soliloquy encryption scheme, Smart-Vercauteren fully homomorphic encryption scheme from PKC 2010, and Gentry-Garg-Halevi cryptographic multilinear-maps from Eurocrypt 2013.

Those broken schemes were using a special class of principal ideals, but these works also showed how to solve SVP for principal ideals in the *worst-case* in quantum polynomial time for an approximation factor of $\exp(\tilde{O}(\sqrt{n}))$. This exposed an unexpected hardness gap between general lattices and some structured ones, and called into question the hardness of various problems over structured lattices, such as Ideal-SVP and Ring-LWE.

In this work, we generalize the previous result to general ideals. Precisely, we show how to solve the close principal multiple problem (CPM) by exploiting the classical theorem that the class-group is annihilated by the (Galois-module action of) the so-called Stickelberger ideal. Under some plausible number-theoretical hypothesis, our approach provides a close principal multiple in quantum polynomial time. Combined with the previous results, this solves Ideal-SVP in the worst case in quantum polynomial time for an approximation factor of $\exp(\tilde{O}(\sqrt{n}))$.

Although it does not seem that the security of Ring-LWE based cryptosystems is directly affected, we contribute novel ideas to the cryptanalysis of schemes based on structured lattices. Moreover, our result shows a deepening of the gap between general lattices and structured ones.

1 Introduction

The problem of finding the shortest vector of a Euclidean lattice (the shortest vector problem, or SVP) is a central hard problem in complexity theory. Approximated versions of this problem (approx-SVP) have become the theoretical

foundation for many cryptographic constructions thanks to the average-case to worst-case reductions of Ajtai [Ajt99] — a classical reduction from approx-SVP to the Short Integer Solution (SIS) problem — and Regev [Reg05] — a quantum reduction from approx-SVP to Learning with Errors (LWE).

For efficiency reasons, it is tempting to rely on structured lattices, in particular lattices arising as ideals or modules over certain rings, the earliest example being the NTRUEncrypt[4] proposal from Hoffstein et al. [HPS98]. Later on, variations on these foundations were also considered.

Precisely, the Ring-SIS [Mic02,LM06,PR06] and Ring-LWE [SSTX09,LPR10] problems were introduced, and shown to reduce to worst-case instances of Ideal-SVP, a specialization of SVP to ideals viewed as lattices. Both problems Ring-SIS and Ring-LWE have shown very versatile problems for building efficient cryptographic schemes upon.

The typical choices of rings for Ring-SIS, Ring-LWE and Ideal-SVP are the ring of integers of a cyclotomic number field of conductor m, that is $K = \mathbb{Q}(\omega_m)$, of degree $n = \varphi(m)$, where ω_m is a complex primitive m-th root of unity. This choice further ensures the hardness of the decisional version of Ring-LWE under the same worst-case Ideal-SVP hardness assumption [LPR10].

Attack on principal ideals. For some time, it seemed plausible that the *structured* versions of lattice problems should be just as hard to solve as the unstructured ones: only some (almost) linear-time advantages were known. This was challenged by a claim of Campbell et al. [CGS14]: a quantum polynomial-time attack against their schemes Soliloquy. The attack also applies to the fully-homomorphic encryption scheme of [SV10] and the cryptographic multilinear maps candidates [GGH13,LSS14], as they all share a common key generation procedure, describe below.

For the secret key, choose an integral element $g \in \mathcal{O}_K$ with small distortion, *i.e.* a $g \in \mathcal{O}_K$ such that

$$\frac{\max_\sigma |\sigma(g)|}{\min_\sigma |\sigma(g)|} \leq \text{poly}(n) \tag{1}$$

where σ ranges over the n complex embeddings $K \mapsto \mathbb{C}$. A corresponding public key consists of the ideal $\mathfrak{I} = (g)$, described by a "bad" $\mathbb{Z}$-basis (e.g. a $\mathbb{Z}$-basis in Hermite normal form).

The attack consists of two steps, sketched in [CGS14]. First, using a quantum computer, it should be possible to solve the Principal Ideal Problem (PIP): given $\mathfrak{I} \subset \mathcal{O}_K$ find $h \in \mathcal{O}_K$ such that $\mathfrak{I} = (h)$. Second, a (classical) close-vector algorithm in the log-unit lattice $\text{Log}\,\mathcal{O}_K^\times$ should allow to recover the secret key[5] g from h. Both steps are claimed to be polynomial time.

While the analysis of the quantum step was unclear[6], such a result seemed plausible considering the recent breakthrough on the Hidden Subgroup Problem

[4] Proposal which is not supported by a worst-case hardness argument, but a variant is [SS11].

[5] up to a root of unity.

[6] and even challenged [BS16, Sec. 6].

over $\mathbb{R}^n$ by Eisentrager et al. [EHKS14] including efficient quantum unit-group computation. And indeed Biasse and Song [BS16] generalized [EHKS14] to S-unit-group computation, allowing in particular to solve PIP [BS16, Thm. 1.3].

The claimed correctness of the short generator recovery step also raised questions: unless a particularly orthogonal basis of the log-unit lattice $\mathrm{Log}\,\mathcal{O}_K^\times$ is known, this step should take exponential time. It was already noticed [GGH13, Full version, pp. 43] that the log-unit lattice could be efficiently decoded up to a radius of $n^{-O(\log\log n)}$ thanks to the Gentry-Szydlo algorithm [GS02], but this is far from sufficient. Yet, the claim that it can be done in polynomial time was quickly supported by convincing numerical experiments [Sch15]. And indeed, by analyzing the geometry of cyclotomic units, Cramer et al. [CDPR16, Thm 4.1] proved that the decoding-radius given by a basis of such units is in fact much better.

A second result of Cramer et al. [CDPR16, Thm 6.3] analyses how good of an approximation of the shortest vector is obtained in the worst-case, i.e. without condition (1). Using a variation on the algorithm of [CGS14], they prove that from any generator h of $\mathfrak{I}$, one can efficiently find a generator g of euclidean length $(N\mathfrak{I})^{1/n} \cdot \exp(\tilde{O}(\sqrt{n}))$. Combined with [BS16], this solves in quantum polynomial time the Short Vector Problem over principal ideals in the worst-case for an approximation factor $\gamma = \exp(\tilde{O}(\sqrt{n}))$.

Claim 1 ([BS16, Thm 1.3] combined with [CDPR16, Thm 6.3]) *There exists a quantum polynomial time algorithm* PrincipalIdealSVP($\mathfrak{a}$)*, that given an ideal of $\mathcal{O}_K$ for K a cyclotomic number field of prime power conductor, returns an generator $v \in \mathfrak{a}$ of Euclidean norm $\|v\| \leq (N\mathfrak{a})^{1/n} \cdot \exp(\tilde{O}(\sqrt{n}))$.*

In particular, v is a solution to Ideal-SVP for an approximation factor $\gamma = \|v\|/\lambda_1(\mathfrak{a}) = \exp(\tilde{O}(\sqrt{n}))$ where $\lambda_1(\mathfrak{a})$ denotes the length of the shortest vector of $\mathfrak{a}$.

It is also shown [CDPR16, Lem. 6.2] that this result is tight up to a polylog(n) factor in the exponent: the shortest generator is typically larger than the shortest element by a factor $\exp(\tilde{O}(\sqrt{n}))$.

Impact and limitatioms of the attack on principal ideals. Whereas some cryptosystems were broken by this quantum attack, the current limitations of this approach to tackle more standard problems as Ring-LWE are three-fold.

(i) First, it is restricted to principal ideals, while Ring-SIS and Ring-LWE rely on worst-case hardness of SVP over general ideals.
(ii) Second, the approximation factor $\gamma = \exp(\tilde{O}(\sqrt{n}))$ in the worst-case is asymptotically too large to affect any actual Ring-LWE based schemes even for advanced cryptosystems such as the state of the art fully homomorphic encryption schemes (see [BV11,DM15]).
(iii) Third, Ring-LWE is known to be at least as hard as Ideal-SVP but not known to be equivalent.

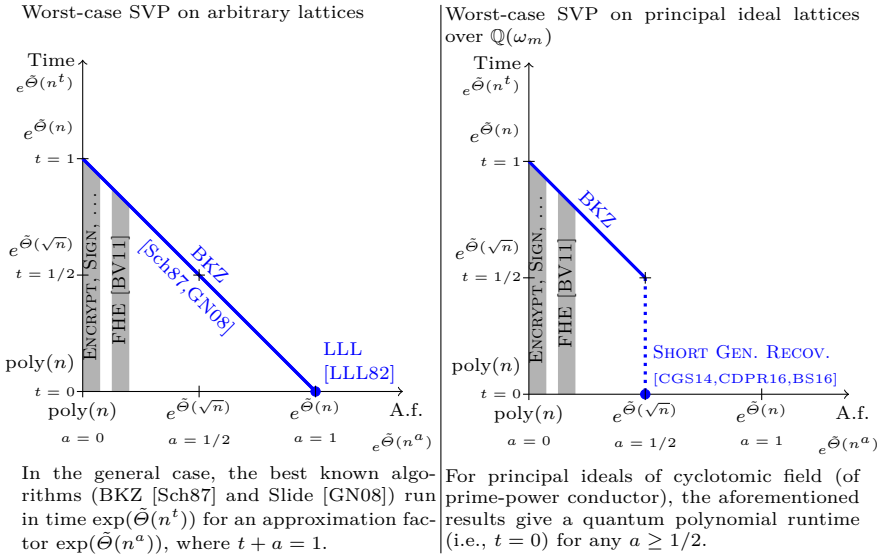

Fig. 1. Best known (quantum) Time–Approximation factor tradeoffs to solve approx-SVP in arbitrary lattices (on the left) and in principal ideal lattices (on the right), in the worst case. The approximation factors of (Ideal)-SVP used to build cryptography upon are typically between polynomial $\text{poly}(n)$ and quasi-polynomial $\exp(\text{polylog}(n))$.

But it does show an asymptotic gap between the search of mildly short vectors in general lattices and in certain structured lattices (see Figure 1), and calls for a more thorough study of the hardness assumption over structured lattices. This work addresses the first of them.

1.1 Contributions

This work provides strong evidence that the general case of Ideal-SVP is not harder than the principal case for similar approximation factors. As a consequence, the approximation factors reachable in quantum polynomial time appear to be significantly smaller in arbitrary ideals of cyclotomic fields of prime-power conductor than known for general lattices, dropping from $\exp(\tilde{\Theta}(n))$ to $\exp(\tilde{\Theta}(\sqrt{n}))$.

Main Result (Under GRH, Assumptions 1 and 2) *There exists a quantum polynomial time algorithm* IdealSVP$(\mathfrak{a})$*, that given an ideal of $\mathcal{O}_K$ for K a cyclotomic number field of prime power conductor, returns an element $v \in \mathfrak{a}$ of Euclidean norm $\|v\| \leq (N\mathfrak{a})^{1/n} \cdot \exp(\tilde{O}(\sqrt{n}))$.*

In other words, Ideal-SVP is solvable in quantum polynomial time in cyclotomic number fields for an approximation factor $\gamma = \exp(\tilde{O}(\sqrt{n}))$.

The strategy consists in reducing the problem over general ideals to that over principal ideals, for cyclotomic fields of prime-power conductor m. We show that

under some number-theoretic assumptions, it is possible to solve the *close principal multiple* (CPM) problem in quantum polynomial time for an a good enough approximation factor. More precisely, the CPM problem consists in finding a principal ideal $\mathfrak{c} \subset \mathfrak{a}$ for an arbitrary ideal $\mathfrak{a}$, such that the algebraic norm of $\mathfrak{c}$ is not much larger than the norm of $\mathfrak{a}$, say up to a factor $\exp(\tilde{O}(n^{1+c}))$. We will argue that one can reach $c = 1/2$, yet, any $c < 1$ will provide a better time-approximation factor tradeoff than the generic algorithms LLL and BKZ.

Our main tool to solve CPM is the classical theorem that the class-group is annihilated by the Galois-module action of the so-called Stickelberger ideal: it provides explicit class relations between an ideal and its Galois conjugates. An important fact is that this Stickelberger ideal has many short elements and that these can be explicitly constructed (see for example [Sch10]). This leads to a quantum polynomial time algorithm to solve CPM for a factor $\exp(\tilde{O}(n^{1+c}))$, where the constant c depends on how many Galois orbits of prime ideals are used to generate the (minus part of the) class group. It remains to apply the short generator recovery to $\mathfrak{c}$ to find a short vector of $\mathfrak{a}$, approximating the shortest vector by a factor $\exp(\tilde{O}(n^{\max(1/2,c)}))$.

We follow the notations of Figure 1. If the exponent c can be made strictly smaller than 1, this gives a non-trivial result compared to generic lattice algorithms (see [Sch87,GN08]): we get $t = 0$ for any $a \geq \max(1/2, c)$, and in particular $a + t < 1$, against $a + t = 1$ for generic algorithms. If c can be made as small as $1/2$, then the asymptotic tradeoffs for Ideal-SVP are as good as the tradeoffs for Principal-Ideal-SVP.

Concluding formally on which value of c can be achieved is not straightforward, as it relies on the structure of the class group Cl_K as a $\mathbb{Z}[G]$-module (see Section 2.3). Based on computations of the class group structure of Schoof [Sch98] and a heuristic argument, we strongly believe it is plausible that $c = 1/2$ is reachable at least for a dense family of conductors m, if not all. This leads to the main result stated above.

1.2 Impact, open questions and recommendations

To the best of our knowledge, this new result does not immediately lead to an attack on any proposed scheme, since most of them are based on Ring-LWE: obstacles (ii) and (iii) remain. Each of this obstacle leaves a crucial open cryptanalytic questions.

- The first question is whether the $\gamma = \exp(\tilde{O}(\sqrt{n}))$ approximation factors can be improved, potentially increasing the running time. One could for example consider many CPM solutions rather than just one, and hope that one of them leads to a much shorter vector.
- The second is whether an oracle for Ideal-SVP (an approx-SVP oracle for modules of rank 1) can be helpful to solve Ring-LWE, which can be summarized as an "unusually-Short Vector Problem" over a module of rank 3. Note that the natural approach of using LLL generalized to other rings as done by Napias [LLL82,Nap96] fails since only the ring of integers of a few cyclotomic fields of small conductor are Euclidean [Len75].

Despite those two serious obstacles to attack Ring-LWE based schemes by the algebraic approach developed in [CGS14,BS16,CDPR16] and in this paper, it seems a reasonable precaution to start considering weaker structured lattice assumptions, such as Module-LWE [LS15] (i.e., an "unusually-Short Vector Problem" in a module of larger rank over a smaller ring), which provides an intermediate problem between ring-LWE and general LWE.

It is also possible to consider other rings, as done in [BCLvV16]. Yet, the latter proposal surprisingly relies on the seemingly stronger NTRU assumption ("unusually-Short Vector Problem" over modules of rank 2). In the current state of affairs [KF16], there seems to be an asymptotic hardness gap between NTRU and Ring-LWE, whatever the ring[7], and down to quite small polynomial approximation factors. Should the concrete security claims of [BCLvV16] not be directly affected, the same reasonable precaution principle should favor weaker assumptions, involving modules of a larger rank.

2 Overview

2.1 Notations and reminders.

Throughout this paper, let m be a prime power, $\omega_m \in \mathbb{C}$ be a complex primitive m-th root of unity, and $K = \mathbb{Q}(\omega_m)$ be the cyclotomic number field of conductor m. It is a number field of degree $n = \varphi(m) = \Theta(m)$. Let G denote its Galois group over $\mathbb{Q}$ and $\tau \in G$ denotes the complex conjugation. We recall that the discriminant Δ_K of K asymptotically satisfies $\log |\Delta_K| = O(n \log n)$.

Ideals as lattices. The field K is endowed with a canonical Hermitian vector space structure via its Minkowsky embedding. Concretely, its inner product is defined via the trace map $\mathrm{Tr} : K \to \mathbb{Q}$ by $\langle a, b \rangle = \mathrm{Tr}(a\tau(b))$, and the associated Euclidean norm is denoted $\| \cdot \| : a \mapsto \langle a, a \rangle = \mathrm{Tr}(a\tau(a))$.

The ring of integers of K is denoted $\mathcal{O}_K$ and in the cyclotomic case is simply given by $\mathcal{O}_K = \mathbb{Z}[\omega_m]$. Any ideal $\mathfrak{h}$ of $\mathcal{O}_K$ can be viewed as a Euclidean lattice via the above inner-product. The algebraic norm of an ideal $\mathfrak{h}$ is written $N\mathfrak{h}$. The volume of $\mathfrak{h}$ as a lattice relates to its algebraic norm by $\mathrm{Vol}(\mathfrak{h}) = \sqrt{|\Delta_K|} N\mathfrak{h}$. The length $\lambda_1(\mathfrak{h})$ of the shortest vector of $\mathfrak{h}$ is determined by its algebraic norm up to a polynomial factor:

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

The right inequality is an application of Minkowsky's second theorem, whereas the left one follows from the fact that the ideal $v\mathcal{O}_K$ generated by the shortest vector v of $\mathfrak{h}$ is a multiple (a sub-ideal) of $\mathfrak{h}$, and that $\mathrm{Vol}(v\mathcal{O}_K) \leq \|v\|^n$.

[7] This actually seems to hold even without any commutative ring structure, i.e., when comparing "matrix-NTRU" to regular LWE.

Class group. The class group $\mathrm{Cl}_K = \mathscr{I}_K/\mathscr{P}_K$ of K is the quotient of the (abelian) multiplicative group of fractional ideals $\mathscr{I}_K$ by the subgroup of fractional principal ideals. We denote $[\mathfrak{h}] \in \mathrm{Cl}_K$ the class of an ideal $\mathfrak{h}$. The trivial class $[\mathcal{O}_K]$ is the class of principal ideals. Given two ideals $\mathfrak{h}$ and $\mathfrak{f}$, we write $\mathfrak{h} \sim \mathfrak{f}$ if they have the same class. The class group is written multiplicatively.

The class number $h_K = |\mathrm{Cl}_K|$ is the order of the class group. Loosely speaking, the class group measures the lack of principality of the ring $\mathcal{O}_K$. In particular, the class group is trivial ($h_K = 1$) if and only if $\mathcal{O}_K$ is a principal ideal domain. This holds only for finitely many conductors $m \geq 1$ and, more precisely, we know that $\log h_K = \Theta(n \log m)$ [Was12, Thm 4.20].

2.2 Overview

It has been shown [CGS14,BS16,CDPR16] (under reasonable assumptions) that given an arbitrary principal ideal $\mathfrak{a} \subset \mathcal{O}_K$, one can recover in quantum polynomial time an element $g \in \mathfrak{a}$ (in fact a generator of $\mathfrak{a}$, i.e. such that $\mathfrak{a} = g\mathcal{O}_K$) such that $\|g\| \leq (N\mathfrak{a})^{1/n} \cdot \exp(\tilde{O}(n^{1/2}))$. Our goal is to reduce the case of general ideals to the case of principal ideals.

The close principal multiple problem (CPM) To do so, a folklore approach is to search for a reasonably close multiple $\mathfrak{c} = \mathfrak{a}\mathfrak{b}$ of $\mathfrak{a}$ that is *principal*; in other words, one searches for a small integral ideal $\mathfrak{b}$ such that $\mathfrak{b} \sim \mathfrak{a}^{-1}$. If such an ideal $\mathfrak{b}$ with norm less than $\exp(\tilde{O}(n^{1+c}))$ for some constant $c > 0$ is found, this implies, by the aforementioned results, that one can find a generator g of $\mathfrak{c}$ such that

$$\begin{aligned}\|g\| &\leq (N\mathfrak{c})^{1/n} \cdot \exp\left(\widetilde{O}\left(n^{1/2}\right)\right) \\ &\leq (N\mathfrak{a})^{1/n} \cdot (N\mathfrak{b})^{1/n} \cdot \exp\left(\widetilde{O}\left(n^{1/2}\right)\right) \\ &\leq (N\mathfrak{a})^{1/n} \cdot \exp\left(\widetilde{O}\left(n^{\max(1/2,c)}\right)\right).\end{aligned}$$

Because $g \in \mathfrak{c} \subset \mathfrak{a}$, one has found a short vector of $\mathfrak{a}$, larger than the shortest vector of $\mathfrak{a}$ by a sub-exponential approximation factor $\exp(\tilde{O}(n^{\max(1/2,c)}))$. This is asymptotically as good as the principal case when $c = 1/2$, and better than LLL for any $c < 1$.

CPM as a close vector problem. Before searching for a solution to the CPM problem, let us discuss wether a $\exp(\tilde{O}(n^{1+c}))$-close principal multiple exists in general. A positive answer follows from the results of [JW15, Cor. 6.5][8] setting a prime factor basis $\mathfrak{B} = \{\mathfrak{p} \mid N\mathfrak{p} \leq n^{4+o(1)}\}$, for any class $C \in \mathrm{Cl}_K$, there exists a non-negative small solution $e \in \mathbb{Z}^{\mathfrak{B}}_{\geq 0}$ to the class equation $[\prod \mathfrak{p}^{e_\mathfrak{p}}] = C$,

[8] The earlier result of [JMV09, Cor.1.3] is not sufficient as it does not keep track of the dependence on the degree of the number fields, left hidden in the constants.

of ℓ_1-norm $\|e\|_1 \leq O(n^{1+o(1)})$. This proves, assuming GHR, the existence of a solution $\mathfrak{b} = \prod \mathfrak{p}^{e_\mathfrak{p}}$ to the CPM problem as small as $\exp(\tilde{O}(n^{1+c}))$ for $c = o(1)$.

The previous argument is based on the analysis of the expander properties of certain Caley graphs on the class group. For our purpose, existence is not enough, as we wish to efficiently find a close principal multiple. We instead write the class group using lattices. If the factor basis $\mathfrak{B}$ generates the whole class group, then one may rewrite $\mathrm{Cl}_K \simeq \mathbb{Z}^\mathfrak{B}/\Lambda$ where Λ is the lattice of class relations: $\Lambda = \{e \in \mathbb{Z}^\mathfrak{B} | [\prod \mathfrak{p}^{e_\mathfrak{p}}] = [\mathcal{O}_K]\}$. Otherly said, $\Lambda \subset \mathbb{Z}^\mathfrak{B}$ is the kernel of the surjection $\mu : \mathbb{Z}^\mathfrak{B} \twoheadrightarrow \mathrm{Cl}_K$. In fact, it will be enough to consider any full-rank sublattice $\Gamma \subset \Lambda$ of class relations, i.e. any subgroup $\Gamma \subset \Lambda$ of finite index.

The CPM problem can now be rephrased as a *close vector problem*: given a class $C = [\mathfrak{a}]^{-1} \in \mathrm{Cl}_K$, one first use the Biasse-Song quantum algorithm [BS16] to compute a representative of that class $\alpha \in \mathbb{Z}^\mathfrak{B}$ in base $\mathfrak{B}$ (see Proposition 2), that is an α such that $\mu(\alpha) = C$. Then one reduces this representation, by searching for a lattice vector $\beta \in \Gamma$ close to α. Note that $\mu(\alpha - \beta) = \mu(\alpha) = C$. This provides a solution[9] $\mathfrak{b} = \prod \mathfrak{p}^{\alpha_\mathfrak{p} - \beta_\mathfrak{p}}$, of norm at most $B^{\|\alpha-\beta\|_1}$, where B is a bound such that $N\mathfrak{p} \leq B$ for every $\mathfrak{p} \in \mathfrak{B}$. It is therefore sufficient to find an appropriate factor basis together with a good basis of the lattice of relations Γ to attack this problem. The condition over Γ to be of full-rank is necessary to have any guarantee on the length of the reduced representative $\alpha - \beta$.

The Stickelberger ideal: class relations for free. For this discussion, let us assume for now that the class group can be generated by a single ideal of small norm and its conjugates: $\mathfrak{B} = \{\mathfrak{p}^\sigma = \sigma(\mathfrak{p}) | \sigma \in G\}$ and $N\mathfrak{p} = \mathrm{poly}(n)$.

Stickelberger's theorem will provide *explicit class relations* between any ideal $\mathfrak{h}$ and its conjugates. More precisely, consider the group ring $\mathbb{Z}[G]$, which naturally acts on $\mathcal{O}_K$-ideals as follows:

$$\mathfrak{h}^s = \prod_{\sigma \in G} \mathfrak{h}^{s_\sigma \cdot \sigma} = \prod_{\sigma \in G} \sigma(\mathfrak{h})^{s_\sigma} \quad \text{where } s = \sum_{\sigma \in G} s_\sigma \cdot \sigma \in \mathbb{Z}[G].$$

Stickelberger gave an explicit construction of a $\mathbb{Z}[G]$-ideal $S \subset \mathbb{Z}[G]$ that *annihilates the class group*, i.e. $\mathfrak{h}^s \sim \mathcal{O}_K$ (i.e., $\mathfrak{h}^s$ is principal) for any ideal $\mathfrak{h} \subset \mathcal{O}_K$ and any element $s \in S$. Forgetting the multiplicative structure of $\mathbb{Z}[G]$ directly gives a lattice of class relations $\mu(S) \subset \mathbb{Z}^\mathfrak{B}$ by the canonical morphism of $\mathbb{Z}$-modules $\kappa : \mathbb{Z}[G] \to \mathbb{Z}^\mathfrak{B}$, sending σ to the canonical vector $\mathbf{1}_{\mathfrak{p}^\sigma}$.

A technical issue is that the Stickelberger ideal is not of full rank in $\mathbb{Z}[G]$ as a $\mathbb{Z}$-module, so needs to be extended[10] in order to serve as the lattice of relations Γ. This can be resolved by working only with the *minus* part Cl_K^- of the class group, i.e., the relative class group of K over the maximal real subfield K^+. More formally, Cl_K^- is the kernel of the morphism $\mathrm{Cl}_K \to \mathrm{Cl}_{K^+}$ induced

[9] One notes that this solution is not integral as desired, yet getting rid of negative exponents will be easy, at least in the relative class group Cl_K^-.

[10] if a lattice is not of full rank, no close-vector algorithm can guarantee any distance bound, as any fundamental domain is unbounded.

by the relative norm map $N_{K/K^+} : \mathfrak{h} \mapsto \mathfrak{h}\mathfrak{h}^\tau$. This subgroup $\mathrm{Cl}_K^- \subset \mathrm{Cl}_K$ is annihilated by the *augmented* Stickelberger ideal $S' = S + (1+\tau)\,\mathbb{Z}[G]$. For this discussion, let us just assume that Cl_{K^+} is trivial, so that the whole class group $\mathrm{Cl}_K = \mathrm{Cl}_K^-$ is annihilated by the augmented Stickelberger ideal S'.

The geometry of the Stickelberger ideal. An important fact is that this ideal has many short elements and that these can be explicitly constructed — this remark is certainly not new, at least for prime conductors [Sch10]. Under our simplifying assumption that $\mathfrak{B} = \{\mathfrak{p}^\sigma \mid \sigma \in G\}$ generates Cl_K, and the additional assumption that the plus part of the class group Cl_{K^+} is trivial, this approach will allow to solve the close multiple problem within a norm bound

$$\exp\left(\widetilde{O}\left(n^{3/2}\right)\right).$$

Sufficient conditions. In the result sketched above, we made two simplifying assumptions. We now sketch how those assumptions can be relaxed, and provide evidences for the relaxed assumptions. Those assumptions and their supporting evidences will be detailed in Section 2.3.

Triviality of Cl_{K^+}*.* One assumption was that the plus part Cl_{K^+} of the class group is trivial. In fact, we can rather easily handle a non-trivial plus-part as long as $h_K^+ = |\mathrm{Cl}_{K^+}| = \mathrm{poly}(n)$, using rapid-mixing properties of some Cayley graphs on Cl_{K^+}. And since h_K^+ is the class number of a totally real number field, it is actually expected to be small. This assumption is already present in [CGS14,CDPR16], and is supported by numerical evidences ([Was12, p. 420, Table 4], computed by Schoof [Sch89]), and by arguments based on the Cohen-Lenstra heuristic [BPR04].

Knowledge of a $\mathbb{Z}[G]$*-generator of* Cl_K^-*.* The other assumption was that we know of a factor basis of Cl_K^- of the form $\mathfrak{B} = \{\mathfrak{p}^\sigma = \sigma(\mathfrak{p}) \mid \sigma \in G\}$ for a single ideal $\mathfrak{p}$ of small norm $N\mathfrak{p} = \mathrm{poly}(n)$. Otherly said, we know of a small norm ideal $\mathfrak{p} \subseteq \mathcal{O}_K$ such that $[\mathfrak{p}]$ is a $\mathbb{Z}[G]$-generator of Cl_K^-.

This assumption can also be relaxed. We may allow a few primes and their conjugates in the factor basis. Assuming one knows a factor basis $\mathfrak{B} = \{\mathfrak{p}_i^\sigma \mid \sigma \in G, i = 1, \ldots, d\}$ composed of d Galois orbits, (with $N\mathfrak{p}_i \leq \mathrm{poly}(n)$) that generates Cl_K^-, our approach leads to solving the close principal multiple problem within a norm bound

$$\exp\left(\widetilde{O}\left(d \cdot n^{3/2}\right)\right).$$

This leads to solving approximate Ideal-SVP with a better approximation factor than pure lattice reduction for any class of conductors $m \in \mathbb{Z}$ whenever one can build a factor basis of size $d = \tilde{O}(n^a)$ for an $a < 1/2$.

Therefore, the crux of the matter is about how small of a factor basis $\mathfrak{B}$ can be built[11]. The structure of the class group Cl_K^- remains quite elusive, but it

[11] Note that, as a computational problem, this task is *non-uniform*. That is, it must be ran once for each conductor m of interest, but does not need to be re-run for each

appears that it admits a very small minimum number of generators as a $\mathbb{Z}[G]$-module. Schoof [Sch98] computed that for all prime conductors $m \leq 509$, Cl_K^- is $\mathbb{Z}[G]$-cyclic (i.e., it is generated by a single element as a $\mathbb{Z}[G]$-module). This property is sufficient to argue that one can efficiently find a small generating set and reach $c = 1/2$, under the heuristic that classes of small random ideals behave similarly to uniformly random classes. Even if the minimal number of generators is not always 1 but still small, say $O(n^\epsilon)$ for some $\epsilon > 0$, this heuristic allows to reach $c = 1/2 + \epsilon$.

2.3 Assumptions

Our main result is conditionned on two assumptions concerning the asymptotic structure of the class group, sketched above and stated below. Of course, if those statement were to not hold for all prime power conductors m, our result remains meaningful if both assumptions simultaneously hold for a common infinite class of conductors, such as $\mathcal{M}_\ell = \{m = \ell^e \mid e \geq 0\}$ for a fixed prime ℓ. We also note that the second assumption can be weakened from $d = \mathrm{polylog}(n)$ to $d = n^\epsilon$ for any $\epsilon < 1/2$ to reach a non trivial approximation factor $\gamma = \exp(\tilde{O}(n^{1/2+\epsilon}))$.

The real class number. The first assumption concerns the size h_K^+ of the class group of the real subfield K^+, and is already used in [CGS14,CDPR16]. For any integer m, let $h^+(m)$ be the class number of the maximal totally real subfield of the cyclotomic field of conductor m.

Assumption 1 *For prime powers m, it holds that $h^+(m) \leq \mathrm{poly}(n)$.*

The literature on h_K^+ provides strong theoretical and computational evidence that it is indeed small enough. First, the Buhler, Pomerance, Robertson [BPR04] formulate and argue in favor of the following conjecture, based on Cohen-Lenstra heuristics.

Conjecture 1 (Buhler, Pomerance, Robertson [BPR04]) *For all but finitely many pairs (ℓ, e), where ℓ is a prime and e is a positive integer, we have $h^+(\ell^{e+1}) = h^+(\ell^e)$.*

A stronger version for the case $\ell = 2$ was formulated by Weber.

Conjecture 2 (Weber's class number problem) *For any e, $h^+(2^e) = 1$.*

A direct consequence of Conjecture 1 is that for fixed ℓ and increasing e, $h^+(\ell^e)$ is $O(1)$, implying that Assumption 1 holds over the class $\mathcal{M}_\ell$.

But even for increasing primes ℓ, $h^+(\ell)$ itself is also small: Schoof [Sch03] computed all the values of $h^+(\ell)$ for $\ell < 10,000$ (correct under heuristics of type Cohen-Lenstra, and Miller proved in [Mil15] its correctness under GRH at

CPM instance in $\mathcal{O}_K$. A proof of existence of such a factor basis would already have a consequence in a complexity theoretic perspective. We however heuristically argue in Section 2.3 that a good basis can actually be found efficiently.

least for the primes $\ell \leq 241$). According to this table, for 75.3% of the primes $\ell < 10,000$ we have $h^+(\ell) = 1$ (matching Schoof's prediction of 71.3% derived from the Cohen-Lenstra heuristics). All the non-trivial values remain very small, as $h^+(\ell) \leq \ell$ for 99.75% of the primes.

Constructing small factor bases of $\mathbf{Cl}_K^-$. This assumption is arguably new, and can be read as a strengthened version of a Theorem of Bach [Bac90, Theorem 4] and its generalizations from [JMV09] and [JW15, Cor. 6.5].

Assumption 2 *There are integers $d \leq \mathrm{polylog}(n)$ and $B \leq \mathrm{poly}(n)$ such that the following holds. Choose uniformly at random d prime ideals $\mathfrak{p}_1, \ldots, \mathfrak{p}_d$ among the finitely many ideals $\mathfrak{p}$ satisfying $N\mathfrak{p} \leq B$ and $[\mathfrak{p}] \in \mathrm{Cl}_K^-$. Then, the factor basis $\mathfrak{B} = \{\mathfrak{p}_i^\sigma \mid \sigma \in G, i = 1 \ldots d\}$ generates Cl_K^- with probability at least $1/2$.*

To argue for this assumption, we prove (Proposition 1) that if Cl_K^- can be generated by r ideal classes, then $r \cdot \mathrm{polylog}(n)$ many uniformly random classes in Cl_K^- will generate it.

Proposition 1. *Let K be a cyclotomic field of conductor m, with Galois group G and relative class group Cl_K^-. Let r be the minimal number of $\mathbb{Z}[G]$-generators of Cl_K^-. Let $\alpha \geq 0$ be a parameter, and s be any integer such that*

$$s \geq r(\log_2 \log_2(h_K^-) + \alpha)$$

(note that $\log_2 \log_2(h_K^-) \sim \log_2(n)$). Let $g_1, \ldots, g_s$ be s independent uniform elements of Cl_K^-. The probability that $\{g_1, \ldots, g_s\}$ generates Cl_K^- as a $\mathbb{Z}[G]$-module is at least $\exp\left(-\frac{3}{2^\alpha}\right) = 1 - O(2^{-\alpha})$.

The proof is deferred to Appendix A.

To justify Assumption 2, we first argue that r is admittedly as small as $\mathrm{polylog}(n)$. For the case $m = 2^e$, this can be argued by just looking at the value of $h^-(2^e)$ computed up to $e = 9$ in [Was12, Table 3]. These values are square-free, so Cl_K^- is $\mathbb{Z}$-cyclic and therefore $\mathbb{Z}[G]$-cyclic; in other words, $r = 1$. The case of prime conductors was also studied by Schoof [Sch98]: he proved that Cl_K^- is $\mathbb{Z}[G]$-cyclic for every prime conductor $m \leq 509$; again, $r = 1$.

While it is unclear that this cyclicity should be the typical behavior asymptotically, it seems reasonable to assume that r remains as small as $\mathrm{polylog}(n)$, at least for a dense class of prime power conductors.

Once it is admitted that $r \leq \mathrm{polylog}(n)$, Assumption 2 simply assumes that Proposition 1 remains true when imposing that the random classes $g_1 \ldots g_s$ are chosen as the classes of random ideals of small norm, i.e. $g_i = [\mathfrak{p}_i]$ where $N\mathfrak{p}_i \leq \mathrm{poly}(n)$. This restriction on the norms seems reasonable considering that it has been proven that prime ideals of norm $\mathrm{poly}(n)$ are sufficient to generate Cl_K^-, assuming GRH and Assumption 1 (see [JW15, Cor. 6.5]).

3 Quantum algorithms for class groups

Searching for a principal multiple of the ideal $\mathfrak{a}$ in $\mathcal{O}_K$ will require to perform computations in the class group in an efficient way. Classically, problems related to class group computations remain difficult, and the best known classical algorithms run in sub-exponential time (for example, see [BF14,EFGK16]). Yet, building on the recent advances on quantum algorithms for the Hidden Subgroup Problem in large dimensions [EHKS14], Biasse and Song [BS16] introduced a quantum algorithm to perform S-unit group computations. It implies class group computations, and solution to the principal ideal problem (PIP) in quantum polynomial time.

While it is not made explicit in [BS16], the Biasse-Song algorithm for S-unit group computation also allows to solve the class group discrete logarithm problem: given a basis $\mathfrak{B}$ of ideals generating a subgroup of the class group Cl_K containing the class of $\mathfrak{a}$, express the class of $\mathfrak{a}$ as a product of ideals in $\mathfrak{B}$. We provide a formal statement and a proof for completeness. Given the Theorem 1.1 of [BS16] the proof of this corollary is standard, and known as the linear-algebra step of many index calculus methods.

Proposition 2 (Direct corollary of [BS16, Theorem 1.1]). *Let $\mathfrak{B}$ be a set of prime ideals generating a subgroup H of Cl_K. There exists a quantum algorithm $\mathrm{ClDL}_{\mathfrak{B}}$ which, when given as input any ideal $\mathfrak{a}$ in $\mathcal{O}_K$ such that $[\mathfrak{a}] \in H$, outputs a vector $\mathbf{y} \in \mathbb{Z}^{\mathfrak{B}}$ such that $\prod \mathfrak{p}^{y_\mathfrak{p}} \sim \mathfrak{a}$, and runs in polynomial time in $n = \deg(K)$, $\max_{\mathfrak{p}\in\mathfrak{B}} \log(N\mathfrak{p})$, $\log(N\mathfrak{a})$, and $|\mathfrak{B}|$.*

Proof. The prime factorization $\mathfrak{a} = \mathfrak{q}_1^{a_1} \dots \mathfrak{q}_k^{a_k}$ can be obtained in polynomial time in n, $\log(\Delta_K)$ and $\log(N\mathfrak{a})$, by Shor's algorithm [Sho97,EH10]. Let $\mathfrak{C} = \mathfrak{B} \cup \{\mathfrak{q}_1 \dots, \mathfrak{q}_k\}$, and one can assume without loss of generality that this union is disjoint. Let $r = n_1+n_2-1$, where n_1 is the number of real embeddings of K, and n_2 is the number of pairs of complex embeddings. Consider the homomorphism

$$\psi : \mathbb{Z}^{\mathfrak{B}} \times \mathbb{Z}^k \longrightarrow \mathrm{Cl}_K : ((e_\mathfrak{p})_{\mathfrak{p}\in\mathfrak{B}}, (f_1, \dots, f_k)) \longmapsto \left[\prod_{\mathfrak{p}\in\mathfrak{B}} \mathfrak{p}^{e_\mathfrak{p}}\right] \cdot \left[\prod_{i=1}^{d} \mathfrak{q}_i^{f_i}\right].$$

As described in [BS16, Section 4], solving the $\mathfrak{C}$-unit problem provides a generating set of size $c = r + |\mathfrak{B}| + k$ for the kernel L of ψ. From [BS16, Theorem 1.1] such a generating set $\{\mathbf{v}_i\}_{i=1}^c$ can be found by a quantum algorithm in time polynomial in n, $\max_{\mathfrak{p}\in\mathfrak{C}}\{\log(N\mathfrak{p})\}$, $\log(d_K)$ and $|\mathfrak{C}| = O(|\mathfrak{B}| + \log(N\mathfrak{a}))$. For each i, write $\mathbf{v}_i = ((w_{i,\mathfrak{p}})_{\mathfrak{p}\in\mathfrak{B}}, (v_{i,1}, \dots, v_{i,k}))$. Since $[\mathfrak{a}] \in H$ and $\mathfrak{B}$ generates H, the system of equations $\{\sum_{j=1}^c x_j v_{j,i} = a_i\}_{i=1}^k$ has a solution $\mathbf{x} \in \mathbb{Z}^c$ which can be computed in polynomial time. We obtain

$$0 = \psi\left(\sum_{i=1}^{c} x_i \mathbf{v}_i\right) = \left[\prod_{\mathfrak{p}\in\mathfrak{B}} \mathfrak{p}^{\sum_j x_j w_{j,\mathfrak{p}}}\right] \cdot \left[\prod_{i=1}^{d} \mathfrak{q}_i^{\sum_j x_j v_{j,i}}\right] = \left[\prod_{\mathfrak{p}\in\mathfrak{B}} \mathfrak{p}^{\sum_j x_j w_{j,\mathfrak{p}}}\right] \cdot [\mathfrak{a}].$$

Then, the output of $\mathrm{ClDL}_{\mathfrak{B}}$ is $\mathbf{y} = \left(-\sum_j x_j w_{j,\mathfrak{p}}\right)_{\mathfrak{p}\in\mathfrak{B}}$. □

4 Close multiple in the relative class group

Let $K^+ = \mathbb{Q}(\omega_m + \omega_m^{-1})$ denote the maximal real subfield of K, and Cl_{K^+} the class group of K^+. The relative norm map $N_{K/K^+} : \mathrm{Cl}_K \to \mathrm{Cl}_{K^+}$ on ideal classes (which sends the class of $\mathfrak{a}$ to the class of $\mathfrak{a}\mathfrak{a}^\tau$, where τ is the complex conjugation) is a surjection, and its kernel is the relative class group Cl_K^-. In particular, it induces the isomorphism $\mathrm{Cl}_{K^+} \cong \mathrm{Cl}_K/\mathrm{Cl}_K^-$.

The core of the method to find a close principal multiple of an ideal $\mathfrak{a}$ works within the relative class group $\mathrm{Cl}_K^- \subset \mathrm{Cl}_K$. Therefore, as a first step, we need to "send" the ideal $\mathfrak{a} \in \mathrm{Cl}_K$ into this subgroup. More precisely, we want an integral ideal $\mathfrak{b}$ of small norm such that $\mathfrak{a}\mathfrak{b} \in \mathrm{Cl}_K^-$; the rest of the method then works with $\mathfrak{a}\mathfrak{b}$. Let $h_K = |\mathrm{Cl}_K|$ be the class number of K, and $h_K^- = |\mathrm{Cl}_K^-|$ its relative class number. The difficulty of this step is directly related to the index of Cl_K^- inside Cl_K, which is the real class number $h_K^+ = |\mathrm{Cl}_{K^+}|$ of K^+, and is expected to be very small.

4.1 Random walks to the relative class group.

For any $x > 0$, consider the set $\mathcal{S}_x$ of ideals in $\mathcal{O}_K$ of prime norm at most x, and let S_x be the multiset of its image in Cl_K. Let $\mathscr{G}_x$ denote the induced Cayley (multi)graph $\mathrm{Cay}(\mathrm{Cl}_K, S_x)$. From [JW15, Cor. 6.5] (under GRH), for any $\varepsilon > 0$ there is a constant C and a bound

$$B = O\left((n \log \Delta_K)^{2+\varepsilon}\right) = O\left((n^2 \log n)^{2+\varepsilon}\right)$$

such that any random walk in $\mathscr{G}_B$ of length at least $C \log(h_K)/\log\log(\Delta_K)$, for any starting point, lands in the subgroup Cl_K^- with probability at least $1/(2h_K^+)$.

A random walk of length $\ell = \lceil C \log(h_K)/\log\log(\Delta_K)\rceil = \tilde{O}(n)$ is a sequence $\mathfrak{p}_1, ..., \mathfrak{p}_\ell$ of ideals chosen independently, uniformly at random in $\mathcal{S}_B$, and their product $\mathfrak{b} = \prod \mathfrak{p}_i$ has a norm bounded by

$$N\mathfrak{b} = \prod_{i=1}^{\ell} N\mathfrak{p}_i \leq B^\ell = \exp(\mathrm{polylog}(n) \cdot \tilde{O}(\log h_K)) = \exp(\tilde{O}(n)),$$

If $[\mathfrak{a}]$ is the starting point of the random walk in the graph, the endpoint $[\mathfrak{a}\mathfrak{b}]$ falls in Cl_K^- with probability at least $1/(2h_K^+)$, and therefore an ideal $\mathfrak{b}$ such that $[\mathfrak{a}\mathfrak{b}] \in \mathrm{Cl}_K^-$ can be found in probabilistic polynomial time in h_K^+. Note that the PIP algorithm of Biasse and Song [BS16] allows to test the membership $[\mathfrak{a}\mathfrak{b}] \in \mathrm{Cl}_K^-$, simply by testing the principality of $N_{K/K^+}(\mathfrak{a}\mathfrak{b})$ as an ideal of $\mathcal{O}_K^+$.

The procedure is summarized as Algorithm 1, and the effiency is stated below. Under GRH and Assumption 1, this procedure runs in polynomial time.

Lemma 1 (Under GRH). *Algorithm 1 (*WalkToCl$^-(\mathfrak{a})$*) runs in expected time* $O(h_K^+) \cdot \mathrm{poly}(n, \log N\mathfrak{a})$ *and is correct.*

Algorithm 1 WalkToCl$^-$($\mathfrak{a}$): random walk to Cl_K^-

Require: An ideal $\mathfrak{a}$ in $\mathcal{O}_K$
Ensure: An integral ideal $\mathfrak{b}$ such that $[\mathfrak{ab}] \in \mathrm{Cl}_K^-$ and $N\mathfrak{b} \leq \exp(\tilde{O}(n))$
1: $\ell = \tilde{O}(n)$; $B = \mathrm{poly}(n)$
2: **repeat**
3: **for all** $i = 1 \dots \ell$ **do**
4: Choose $\mathfrak{p}_i$ uniformly among the prime ideal of norm less than B
5: **end for**
6: Set $\mathfrak{b} = \prod \mathfrak{p}_i$
7: **until** $N_{K/K^+}(\mathfrak{ab})$ is principal (using the PIP algorithm of [BS16])
8: $\mathfrak{b} \leftarrow \prod_{i=1}^{d} \mathfrak{p}_i$
9: **return** $\mathfrak{b}$

5 Short relations in Cl_K^- via the Stickelberger ideal

Consider any ideal $\mathfrak{f}$ of $\mathcal{O}_K$ such that $[\mathfrak{f}] \in \mathrm{Cl}_K^-$, and its orbit under the action of the Galois group G, denoted $\mathfrak{F} = G(\mathfrak{f})$. Let R be the group ring $\mathbb{Z}[G]$. It projects to $\mathbb{Z}^{\mathfrak{F}}$, via the map sending σ to $\mathbf{1}_{\mathfrak{f}^\sigma}$.

We now show the construction of an explicit full-rank lattice of class relations in $\mathbb{Z}^{\mathfrak{F}}$ with an explicit set of *short* generators. We proceed by augmenting the *Stickelberger ideal*. This allows to reduce the representation of a given class expressed in basis $\mathfrak{F}$, as shown in Subsection 5.3.

Recall that the Galois group G is canonically isomorphic to $(\mathbb{Z}/m\mathbb{Z})^*$ via $a \mapsto \sigma_a = \zeta_m \mapsto \zeta_m^a$. The norms $\|\cdot\|$ and $\|\cdot\|_1$ denote the usuals ℓ_2 (Euclidean) and ℓ_1 norms over $\mathbb{R}^n$, and are defined over $\mathbb{Z}[G]$ via the natural isomorphism $\mathbb{Z}[G] \cong_{\mathbb{Z}} \mathbb{Z}^n$.

The fractional part of a rational $x \in \mathbb{Q}$ is denoted $\{x\}$, it is defined as the unique rational in the interval $[0, 1)$ such that $\{x\} = x \mod \mathbb{Z}$; equivalently, $\{x\} = x - \lfloor x \rfloor$.

5.1 The (augmented) Stickelberger ideal

Definition 1 (The Stickelberger ideal). *The* Stickelberger element $\theta \in \mathbb{Q}[G]$ *is defined as*

$$\theta = \sum_{a \in (\mathbb{Z}/m\mathbb{Z})^*} \left\{ \frac{a}{m} \right\} \sigma_a^{-1}.$$

The Stickelberger ideal *is defined as* $S = R \cap \theta R$. *We will refer to the Stickelberger lattice when* S *is considered as a* $\mathbb{Z}$*-module.*

This ideal $S \subset R$ will provide some class relations in $\mathbb{Z}^{\mathfrak{F}}$, thanks to the following theorem.

Theorem 1 (Stickelberger's theorem [Was12, Thm. 6.10]). *The Stickelberger ideal annihilates the ideal class group of* K. *In other words, for any ideal* $\mathfrak{h}$ *of* $\mathcal{O}_K$ *and any* $s \in S$, *the ideal* $\mathfrak{h}^s$ *is principal.*

We cannot directly use $S \subset R$ as our lattice of class relations since it does not have full rank in R as a $\mathbb{Z}$-module (precisely its $\mathbb{Z}$-rank is $n/2+1$ when $m \geq 2$). Indeed, if the lattice is not full rank, there can be no guarentee of how short of a representant will be obtained by reducing modulo the lattice. To solve this issue, we will augment the Stickelberger ideal to a full-rank ideal which still annihilates the minus part Cl_K^- of the class group.

Definition 2. *The augmented Stickelberger ideal S' is defined as*

$$S' = S + (1+\tau)R. \tag{2}$$

We will refer to the augmented Stickelberger lattice when S' is considered as a $\mathbb{Z}$-module.

Lemma 2. *The augmented Stickelberger ideal S' annihilates Cl_K^-. In other words, for any ideal $\mathfrak{h}$ of $\mathcal{O}_K$ such that $[\mathfrak{h}] \in \mathrm{Cl}_K^-$ and any $s \in S$, the ideal $\mathfrak{h}^s$ is principal. Moreover, $S' \subset R$ has full-rank n as a $\mathbb{Z}$-module.*

Proof. For the annihilation property it suffices to show that both S and $(1+\tau)R$ annihilate Cl_K^-. By Stickelberger's theorem S annihilates Cl_K so it in particular annihilates the subgroup $\mathrm{Cl}_K^- \subset \mathrm{Cl}_K$. The ideal $(1+\tau)R$ also annihilates Cl_K^- since $\mathfrak{h}^{1+\tau} = \mathfrak{h}\bar{\mathfrak{h}} = N_{K/K^+}(\mathfrak{h})$. We conclude from the fact that Cl_K^- is exactly the kernel of the norm map $N_{K/K^+} : \mathrm{Cl}_K \to \mathrm{Cl}_K^+$.

For the rank, consider the ideal $S^- = S \cap (1-\tau)R$. A theorem from Iwasawa (originally published in [Sin80] but reformulated more conveniently in [Was12, Thm. 6.19]) states that S^- is full rank in $(1-\tau)R$. Noting that $2R \subset (1-\tau)R + (1+\tau)R$, we conclude that $S^- + (1+\tau)R$ has full rank in $2R$, and so does S'. □

5.2 Short generating vectors of the augmented Stickelberger lattice

In the following, the elements of $(\mathbb{Z}/m\mathbb{Z})^*$ are canonically identified with the positive integers $0 < a_1 < a_2 < \cdots < a_n < m$ such that each a_i is coprime to m. The elements of G are indexed as $(\sigma_{a_1}, \ldots, \sigma_{a_n})$. Define the extra element $a_{n+1} = m + a_1$, and note that $a_2 \leq 3$ and that $a_{i+1} - a_i \leq 2$ for any i.

Lemma 3. *The Stickelberger lattice is generated by the vectors $v_i = (a_i - \sigma_{a_i})\theta$ for $i \in \{2, \ldots, n+1\}$.*

Proof. This is almost [Was12, Lem. 6.9]. There, S is considered as an ideal in R, whereas we need these elements to generate S as a $\mathbb{Z}$-module. Let L be the $\mathbb{Z}$-module generated by the v_i's. First, [Was12, Lem. 6.9] immediately implies that $v_i \in S$ and thereby $L \subseteq S$. Now, let $\left(\sum_{i=2}^{n+1} x_i \sigma_{a_i}\right)\theta$ be an arbitrary element of S, with $a_i \in \mathbb{Z}$. One can prove as in [Was12, Lem. 6.9] that m divides $\sum_{i=2}^{n+1} x_i a_i \in \mathbb{Z}$. Since $m = (m+1) - \sigma_{m+1}$, $m\theta$ is in L, and we deduce that $\left(\sum_{i=2}^{n+1} x_i a_i\right)\theta$ is also in L. Therefore,

$$\left(\sum_{i=2}^{n+1} x_i \sigma_{a_i}\right)\theta = \left(\sum_{i=2}^{n+1} x_i (\sigma_{a_i} - a_i)\right)\theta + \left(\sum_{i=2}^{n+1} x_i a_i\right)\theta \in L.$$

This proves that $S \subseteq L$, hence $L = S$. □

We are now ready to construct our set of short generators for S'. Let $w_2 = v_2$ and $w_{i+1} = v_{i+1} - v_i$ for $i \in \{2, \dots, n\}$, and let

$$W = \{w_2, \dots, w_{n+1}\} \cup \{(1+\tau)\sigma, \sigma \in G\}.$$

Lemma 4. *The set S is a set of short generators of S'. More precisely,*

1. *W generates the augmented Stickelberger lattice S',*
2. *For any $i \in \{3 \dots n+1\}$, $w_i = \sum_{b \in (\mathbb{Z}/m\mathbb{Z})^*} \epsilon_{i,b} \cdot \sigma_b^{-1}$, with $\epsilon_{i,j} \in \{0, 1, 2\}$,*
3. *For any $w \in W$, we have $\|w\| \leq \max(2\sqrt{n}, \sqrt{10})$.*

The second item essentially generalizes [Sch10, Proposition 9.4] from prime conductors to prime-power conductors.

Proof. We prove each item individually.

1. First note that $\{w_2, \dots, w_{n+1}\}$ generates S: this is a direct consequence of Lemma 3 and the construction of W. By definition of $R = \mathbb{Z}[G]$, the set $\{(1+\tau)\sigma, \sigma \in G\}$ generates $(1+\tau)R$. One can conclude from the definition of $S' = S + (1+\tau)R$.
2. We follow the computation in the proof of [Was12, Lemma 6.9]:

$$\begin{aligned} v_i = (a_i - \sigma_{a_i})\theta &= \sum_{b \in (\mathbb{Z}/m\mathbb{Z})^*} \left(a_i \left\{ \frac{b}{m} \right\} - \left\{ \frac{a_i b}{m} \right\} \right) \sigma_b^{-1} \\ &= \sum_{b \in (\mathbb{Z}/m\mathbb{Z})^*} \left\lfloor a_i \left\{ \frac{b}{m} \right\} \right\rfloor \sigma_b^{-1} \end{aligned}$$

   using the identity $x\{y\} - \{xy\} = \lfloor x\{y\} \rfloor$ for any integer x and real number y, since this difference is an integer and the term $\{xy\}$ is in the range $[0, 1)$. It remains to rewrite $w_i = \sum_{b \in (\mathbb{Z}/m\mathbb{Z})^*} \epsilon_{i,b} \sigma_b^{-1}$, where

$$\epsilon_{i,b} = \left\lfloor a_{i+1} \left\{ \frac{b}{m} \right\} \right\rfloor - \left\lfloor a_i \left\{ \frac{b}{m} \right\} \right\rfloor \leq a_{i+1} - a_i \leq 2.$$

3. The property follows from the previous item for any $i > 2$. For $i = 2$, we have $w_2 = v_2 = a_2 - \sigma_{a_2}$, and therefore $\|w_2\| = \sqrt{a_2^2 + 1} \leq \sqrt{3^2 + 1} = \sqrt{10}$. Finally, elements $w \in W$ of the form $(1+\tau)\sigma$ have norm $\|w\| = \sqrt{2} \leq \sqrt{10}$. □

5.3 Reducing a class representative in an R-cycle of Cl_K^-

We now show how to exploit the previously constructed set W of short relations to reduce class representations. More precisely, for any large $\alpha \in R$ we will find a short $\beta \in R$ such that $C^\beta = C^\alpha$, for any class $C \in \mathrm{Cl}_K^-$. We shall rely on the following close vector algorithm.

Proposition 3 (Close vector algorithm). *Let $\Gamma \subset \mathbb{R}^k$ be a lattice, and let W be a set generating Γ. There exists a (classical) polynomial time algorithm* CV*, that when given any $y \in \Gamma \otimes \mathbb{R}$ as input, outputs a vector $x = \mathrm{CV}(y, W) \in \Gamma$ such that $\|x - y\|_1 \leq \frac{k}{2} \cdot \max_{w \in W} \|w\|$.*

Proof. Let first $B \subset W$ be a basis of a full-rank sublattice $\Gamma' \subset \Gamma$ (this is easily built in polynomial time). Let $\tilde{B}$ denote the Gram-Schmidt orthogonalization of B. Let $g = \max_{b \in \tilde{B}} \|\tilde{b}\| \leq \max_{b \in B} \|b\| \leq \max_{w \in W} \|w\|$. Applying the Nearest Plane algorithm leads to $x \in \Gamma$ such that $x - y$ belongs to the fundamental parallelepiped $\{\tilde{B}z, z \in [-1/2, 1/2]\}$. We then have

$$\|x - y\|_2^2 \leq \frac{1}{4} \sum \|\tilde{b}_i\|^2.$$

In particular, $\|x - y\|_2 \leq \sqrt{k} \cdot g/2$ and one concludes $\|x - y\|_1 \leq kg/2$. □

Theorem 2. *Assume $n \geq 3$. There is an algorithm* Reduce*, that given $\alpha \in R$, finds in polynomial time in n and $\log(||\alpha||)$, an element $\beta = \mathrm{Reduce}(\alpha) \in R$ such that $||\beta||_1 \leq n^{3/2}$, and $C^\alpha = C^\beta$ for any $C \in \mathrm{Cl}_K^-$.*

Proof. Let W be the basis for the augmented Stickelberger ideal S' as in Lemma 4. From Lemma 2, it has full rank in R. So the close vector algorithm from Proposition 3 can be applied to find an element $\gamma = \mathrm{CV}(\alpha, W) \in S'$ such that $||\alpha - \gamma||_1 \leq \frac{n}{2} \cdot \max_{w \in W} \|w\| \leq n^{3/2}$. Let $\beta = \alpha - \gamma$. For any $C \in \mathrm{Cl}_K^-$, Lemma 2 implies that $C^\gamma = 0$ and therefore $C^\alpha = C^\beta$. □

6 Close principal multiple within the relative class group

We now show how to solve the CPM problem for ideals sitting in Cl_K^-, given a factor basis $\mathfrak{B}$ of Cl_K^-. The CPM approximation factor will depend on the size of the factor basis $\mathfrak{B}$.

Suppose the ideal $\mathfrak{a}$ is in the relative class group Cl_K^-. We are looking for an integral ideal $\mathfrak{b}$ in $\mathcal{O}_K$ of small norm such that $\mathfrak{ab}$ is principal. Let $\mathfrak{B} = \{\mathfrak{p}_i^\sigma \mid \sigma \in G, i = 1, \ldots, d\}$ be a set generating Cl_K^-, composed of d Galois orbits, such that $N\mathfrak{p}_i \leq \mathrm{poly}(n)$ for all i. To state the algorithm and its correctness, no assumption is made on the factor basis $\mathfrak{B}$. In the final section 7, we will employ Assumption 2 to provide a factor basis with $d = \mathrm{polylog}(n)$ to this algorithm.

Algorithm 2 ClosePrincipalMultiple$^-(\mathfrak{a}, \mathfrak{B})$: close principal multiple in the relative class group

Require: An ideal $\mathfrak{a}$ in $\mathcal{O}_K$ such that $[\mathfrak{a}] \in \mathrm{Cl}_K^-$, a factor basis $\mathfrak{B} = \{\mathfrak{p}_i^\sigma | i = 1 \dots d, \sigma \in G\}$ generating Cl_K^-, such that $N\mathfrak{p}_i \leq \mathrm{poly}(n)$ for all i.
Ensure: An (integral) ideal $\mathfrak{b}$ in $\mathcal{O}_K$ such that $\mathfrak{ab} \sim \mathcal{O}_K$ and $N\mathfrak{b} = \exp\left(\tilde{O}\left(dn^{3/2}\right)\right)$

$\mathbf{y} \leftarrow \mathrm{ClDL}_{\mathfrak{B}}(\mathfrak{a})$
for i = 1 **to** d **do**
  $\alpha_i \leftarrow \sum_{\sigma \in G_i} y_{(\mathfrak{p}_i^\sigma)} \sigma \in \mathbb{Z}[G]$
  $\beta_i \leftarrow \mathrm{Reduce}(\alpha_i)$
  $(\gamma_i^+, \gamma_i^-) \leftarrow$ the pair of elements in $\mathbb{Z}[G]$ with only positive coefficients, such that $\gamma_i^+ - \gamma_i^- = -\beta_i$
  $\mathfrak{b}_i \leftarrow \mathfrak{p}_i^{\gamma_i^+ + \tau\gamma_i^-}$
end for
$\mathfrak{b} \leftarrow \prod_{i=1}^d \mathfrak{b}_i$
return $\mathfrak{b}$

Theorem 3. *Algorithm 2,* ClosePrincipalMultiple$^-$*, runs in quantum polynomial time in $n = \deg(K)$, d and $\log(N\mathfrak{a})$, and is correct.*

Proof. Let $\mathfrak{a}, \mathfrak{B}$ be proper inputs, that is, $\mathfrak{a}$ is an ideal of $\mathcal{O}_K$ such that $[\mathfrak{a}] \in \mathrm{Cl}_K^-$, and $\mathfrak{B}$ is a factor basis $\mathfrak{B} = \{\mathfrak{p}_i^\sigma \mid i = 1 \dots d, \sigma \in G\}$ generating Cl_K^-, such that $N\mathfrak{p}_i \leq \mathrm{poly}(n)$ for all i.

The running time follows immediately from Proposition 2 and Theorem 2. Let us now prove the correctness. We have

$$\phi(\mathbf{y}) = \prod_{\mathfrak{p} \in \mathfrak{B}} \mathfrak{p}^{y_\mathfrak{p}} = \prod_{i=1}^d \prod_{\mathfrak{p} \in \mathfrak{B}_i} \mathfrak{p}^{y_\mathfrak{p}} = \prod_{i=1}^d \prod_{\sigma \in G_i} (\mathfrak{p}_i^\sigma)^{y_{(\mathfrak{p}_i^\sigma)}} = \prod_{i=1}^d \mathfrak{p}_i^{\alpha_i}.$$

Observe that for each i, $\mathfrak{b}_i \sim \mathfrak{p}_i^{-\beta_i}$, since $\mathfrak{p}_i^{-1} \sim \mathfrak{p}_i^\tau$. From Theorem 2, we obtain $\mathfrak{p}_i^{\alpha_i}\mathfrak{b}_i \sim \mathcal{O}_K$, which implies that $\phi(\mathbf{y})\mathfrak{b} \sim \prod_{i=1}^d \mathfrak{p}_i^{\alpha_i}\mathfrak{b}_i \sim \mathcal{O}_K$. From Proposition 2, we have $\phi(\mathbf{y}) \sim \mathfrak{a}$, and therefore $\mathfrak{ab} \sim \mathcal{O}_K$.

Now, Theorem 2 ensures that $||\beta||_1 \leq n^{3/2}$. So $||\gamma_i^+||_1 + ||\gamma_i^-||_1$ is bounded by $n^{3/2}$ and we obtain that $N\mathfrak{b}_i \leq (N\mathfrak{p}_i)^{n^{3/2}}$. Then,

$$N\mathfrak{b} = \prod_{i=1}^d N\mathfrak{b}_i \leq \left(\max_{i=1\dots d} N\mathfrak{p}_i\right)^{dn^{3/2}} = \exp\left(\tilde{O}\left(dn^{3/2}\right)\right),$$

where the last inequality uses the fact that each $N\mathfrak{p}_i$ is polynomially bounded in n. □

7 Main result

We now have all the ingredients to demonstrate our main result:

Main Result (Under GRH, Assumption 1 and 2) *Assuming simultaneously the Generalized Riemann Hypothesis, Assumption 1, and Assumption 2, there exists a quantum polynomial time algorithm* IdealSVP($\mathfrak{a}$)*, that given an ideal of $\mathcal{O}_K$ for K a cyclotomic number field of prime power conductor, returns an element $v \in \mathfrak{a}$ of Euclidean norm $\|v\| \leq (N\mathfrak{a})^{1/n} \cdot \exp(\tilde{O}(\sqrt{n}))$.*

Algorithm 3 IdealSVP($\mathfrak{a}$): finding mildly short vectors in an ideal

Require: An ideal $\mathfrak{a}$ in $\mathcal{O}_K$
Ensure: An element $v \in \mathfrak{a}$ of norm $\|v\| \leq (N\mathfrak{a})^{1/n} \exp(\tilde{O}(\sqrt{n}))$
1: $d = \text{polylog}(n)$; $B = \text{poly}(n)$
2: Set $\mathfrak{M} = \{\mathfrak{p} | N\mathfrak{p} \leq B, [\mathfrak{p}] \in \text{Cl}_K^-\}$
3: Choose $\mathfrak{p}_1, \ldots, \mathfrak{p}_d$ uniformly at random in $\mathfrak{M}$
4: Set $\mathfrak{B} = \{\mathfrak{p}_i^\sigma | i \in \{1 \ldots d\}, \sigma \in G\}$
5: $\mathfrak{b}' =$ WalkToCl$^-(\mathfrak{a})$
6: $\mathfrak{b} =$ ClosePrincipalMultiple$^-(\mathfrak{a}\mathfrak{b}', \mathfrak{B})$
7: $v =$ PrincipalIdealSVP$(\mathfrak{a}\mathfrak{b}\mathfrak{b}')$
8: **return** v

Proof. The algorithm is given as Algorithm 3. Efficiency and correctness follow from the previous statements and assumptions:

- Step 2 is quantum polynomial time since membership in Cl_K^- can be tested by applying the Biasse-Song PIP algorithm [BS16, Thm 1.3] to $N_{K/K^+}(\mathfrak{a}\mathfrak{b})$.
- By Assumption 2, Steps 3 and 4 produce a factor basis $\mathfrak{B}$ generating Cl_K^-. Both steps can trivially be performed in polynomial time.
- By Lemma 1, GRH and Assumption 1, Step 5 is quantum polynomial time, and produces an integral ideal $\mathfrak{b}'$ such that $N\mathfrak{b}' \leq \exp(\tilde{O}(n))$ and $[\mathfrak{a}\mathfrak{b}'] \in \text{Cl}_K^-$.
- By Theorem 3, Step 6 produces (in quantum polynomial time) an integral ideal $\mathfrak{b}$ such that

$$N\mathfrak{b} \leq \exp(\tilde{O}(dn^{3/2})) = \exp(\tilde{O}(n^{3/2}))$$

  and such that $\mathfrak{a}\mathfrak{b}\mathfrak{b}'$ is principal.
- By Claim 1 ([CGS14,BS16,CDPR16]), Step 7 produces in quantum polynomial time a vector $v \in \mathfrak{a}\mathfrak{b}\mathfrak{b}'$ of length $\|v\| \leq (N\mathfrak{a}\mathfrak{b}\mathfrak{b}')^{1/n} \cdot \exp(\tilde{O}(\sqrt{n}))$.

Because $\mathfrak{b}$ and $\mathfrak{b}'$ are integral, $\mathfrak{a}\mathfrak{b}\mathfrak{b}' \subset \mathfrak{a}$, and $v \in \mathfrak{a}$. Finally,

$$\begin{aligned}\|v\| &\leq (N\mathfrak{a})^{1/n}(N\mathfrak{b})^{1/n}(N\mathfrak{b}')^{1/n} \cdot \exp(\tilde{O}(\sqrt{n})) \\ &\leq (N\mathfrak{a})^{1/n} \cdot \exp(\tilde{O}(\sqrt{n})).\end{aligned}$$

□

References

[Ajt99] M. Ajtai. Generating hard instances of the short basis problem. In *ICALP*, pages 1–9. 1999.

[Bac90] E. Bach. Explicit bounds for primality testing and related problems. *Mathematics of Computation*, 55(191):355–380, 1990.

[BCLvV16] D. J. Bernstein, C. Chuengsatiansup, T. Lange, and C. van Vredendaal. NTRU prime. Cryptology ePrint Archive, Report 2016/461, 2016. `http://eprint.iacr.org/2016/461`.

[BF14] J.-F. Biasse and C. Fieker. Subexponential class group and unit group computation in large degree number fields. *LMS J. Comput. Math.*, 17(suppl. A):385–403, 2014.

[BPR04] J. Buhler, C. Pomerance, and L. Robertson. Heuristics for class numbers of prime-power real cyclotomic fields,. In *High primes and misdemeanours: lectures in honour of the 60th birthday of Hugh Cowie Williams*, Fields Inst. Commun., pages 149–157. Amer. Math. Soc., 2004.

[BS16] 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*, pages 893–902. SIAM, 2016.

[BV11] Z. Brakerski and V. Vaikuntanathan. Fully homomorphic encryption from ring-LWE and security for key dependent messages. In *CRYPTO*, pages 505–524. 2011.

[CDPR16] R. Cramer, L. Ducas, C. Peikert, and O. Regev. Recovering short generators of principal ideals in cyclotomic rings. In *EUROCRYPT*, pages 559–585. 2016.

[CGS14] P. Campbell, M. Groves, and D. Shepherd. Soliloquy: A cautionary tale. ETSI 2nd Quantum-Safe Crypto Workshop, 2014. Available at `http://docbox.etsi.org/Workshop/2014/201410_CRYPTO/S07_Systems_and_Attacks/S07_Groves_Annex.pdf`.

[DM15] L. Ducas and D. Micciancio. FHEW: bootstrapping homomorphic encryption in less than a second. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 617–640. Springer, 2015.

[EFGK16] T. Espitau, P.-A. Fouque, A. Gélin, and P. Kirchner. Computing generator in cyclotomic integer rings. Cryptology ePrint Archive, Report 2016/957, 2016. `http://eprint.iacr.org/2016/957`.

[EH10] K. Eisenträger and S. Hallgren. Algorithms for ray class groups and hilbert class fields. In *Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '10, pages 471–483. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2010. ISBN 978-0-898716-98-6.

[EHKS14] K. Eisenträger, S. Hallgren, A. Kitaev, and F. Song. A quantum algorithm for computing the unit group of an arbitrary degree number field. In *STOC*, pages 293–302. ACM, 2014.

[GGH13] S. Garg, C. Gentry, and S. Halevi. Candidate multilinear maps from ideal lattices. In *EUROCRYPT*, pages 1–17. 2013. Full version available at `http://eprint.iacr.org/2012/610`.

[GN08] N. Gama and P. Q. Nguyen. Finding short lattice vectors within mordell's inequality. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 207–216. ACM, 2008.

[GS02] C. Gentry and M. Szydlo. Cryptanalysis of the revised NTRU signature scheme. In *EUROCRYPT*, pages 299–320. 2002.

[HPS98] J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, pages 267–288. 1998.

[JMV09] D. Jao, S. D. Miller, and R. Venkatesan. Expander graphs based on GRH with an application to elliptic curve cryptography. *Journal of Number Theory*, 129(6):1491 – 1504, 2009. ISSN 0022-314X. doi:http://dx.doi.org/10.1016/j.jnt.2008.11.006.

[JW15] D. Jetchev and B. Wesolowski. On graphs of isogenies of principally polarizable abelian surfaces and the discrete logarithm problem. *CoRR*, abs/1506.00522, 2015.

[KF16] P. Kirchner and P.-A. Fouque. Comparison between subfield and straightforward attacks on NTRU. Cryptology ePrint Archive, Report 2016/717, 2016. `http://eprint.iacr.org/2016/717`.

[Len75] H. W. Lenstra Jr. Euclid's algorithm in cyclotomic fields. *J. London Math. Soc*, 10:457–465, 1975.

[LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, December 1982.

[LM06] V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP (2)*, pages 144–155. 2006.

[LPR10] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *Journal of the ACM*, 60(6):43:1–43:35, November 2013. Preliminary version in Eurocrypt 2010.

[LS15] A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography*, 75(3):565–599, 2015.

[LSS14] A. Langlois, D. Stehlé, and R. Steinfeld. GGHLite: More efficient multilinear maps from ideal lattices. In *Advances in Cryptology–EUROCRYPT 2014*, pages 239–256. Springer, 2014.

[Mic02] D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007. Preliminary version in FOCS 2002.

[Mil15] J. C. Miller. Real cyclotomic fields of prime conductor and their class numbers. *Math. Comp.*, 84(295):2459–2469, 2015.

[Nap96] H. Napias. A generalization of the LLL-algorithm over euclidean rings or orders. *Journal de théorie des nombres de Bordeaux*, 8(2):387–396, 1996.

[PR06] C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, pages 145–166. 2006.

[Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009. Preliminary version in STOC 2005.

[Sch87] C.-P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.

[Sch89] R. Schoof. *The structure of the minus class groups of abelian number fields*. Rijksuniversiteit Utrecht. Mathematisch Instituut, 1989.

[Sch98] R. Schoof. Minus class groups of the fields of the ℓ-th roots of unity. *Mathematics of Computation of the American Mathematical Society*, 67(223):1225–1245, 1998.

[Sch03] R. Schoof. Class numbers of real cyclotomic fields of prime conductor. *Mathematics of computation*, 72(242):913–937, 2003.

[Sch10] R. Schoof. *Catalan's conjecture*. Springer Science & Business Media, 2010.
[Sch15] J. Schank. LogCvp, Pari implementation of CVP in $\mathrm{Log}\mathbb{Z}[\zeta_{2^n}]^*$. `https://github.com/jschanck-si/logcvp`, March 2015.
[Sho97] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, October 1997. ISSN 0097-5397. doi:10.1137/S0097539795293172.
[Sin80] W. Sinnott. On the Stickelberger ideal and the circular units of an abelian field. *Inventiones math.*, 62:181–234, 1980.
[SS11] D. Stehlé and R. Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *EUROCRYPT*, pages 27–47. 2011.
[SSTX09] D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, pages 617–635. 2009.
[SV10] N. P. Smart and F. Vercauteren. Fully homomorphic encryption with relatively small key and ciphertext sizes. In *Public Key Cryptography*, pages 420–443. 2010.
[Was12] L. C. Washington. *Introduction to cyclotomic fields*, volume 83. Springer Science & Business Media, 2012.

A Proof of Proposition 1

In this appendix, we provide the proof of Proposition 1 (restated below, used to support Assumption 2).

Proposition 1 *Let K be a cyclotomic field of conductor m, with Galois group G and relative class group Cl_K^-. Let r be the minimal number of $\mathbb{Z}[G]$-generators of Cl_K^-. Let $\alpha \geq 0$ be a parameter, and s be any integer such that*

$$s \geq r(\log_2 \log_2(h_K^-) + \alpha)$$

(note that $\log_2 \log_2(h_K^-) \sim \log_2(n)$). Let $g_1, \ldots, g_s$ be s independent uniform elements of Cl_K^-. The probability that $\{g_1, \ldots, g_s\}$ generates Cl_K^- as a $\mathbb{Z}[G]$-module is at least $\exp\left(-\frac{3}{2^\alpha}\right) = 1 - O(2^{-\alpha})$.

In other words, a set of $\Theta(r \log(n))$ random ideal classes in Cl_K^- will generate this $\mathbb{Z}[G]$-module with very good probability. Let us first establish a few lemmas.

Lemma 5. *Let $\mathcal{O}$ be a Dedekind domain, and $\mathfrak{h} \subset \mathcal{O}$ be an integral ideal. Let $g_1, \ldots, g_s$ be s independent uniform elements from $\mathcal{O}/\mathfrak{h}$. Then, the probability that the set $\{g_1, \ldots, g_s\}$ generates $\mathcal{O}/\mathfrak{h}$ as an $\mathcal{O}$-module is*

$$\Pr\left[\mathcal{O}g_1 + \cdots + \mathcal{O}g_s = \mathcal{O}/\mathfrak{h}\right] \geq \left(1 - 2^{-s}\right)^{\log_2 N\mathfrak{h}}. \tag{3}$$

Proof. Let $\mathfrak{h} = \mathfrak{p}_1^{\alpha_1} \cdots \mathfrak{p}_j^{\alpha_j}$ be the prime factorization of $\mathfrak{h}$. By application of the Chinese Remainder Theorem, (3) is equivalent to $\{g_1 \bmod \mathfrak{p}_i^{\alpha_i}, \ldots, g_s \bmod \mathfrak{p}_i^{\alpha_i}\}$ generating $\mathcal{O}/\mathfrak{p}_i^{\alpha_i}$ for all i. In particular, it is enough that for all $i \in \{0 \ldots j\}$ there exists a k such that the ideal $g_k\mathcal{O}$ is coprime with $\mathfrak{p}_i^{\alpha_i}$, or equivalently coprime with $\mathfrak{p}_i$. For a fixed i, this occurs with probability $1 - (N\mathfrak{p}_i)^{-s} \geq 1 - 2^{-s}$. Also note that those events are indepedents, so the probability of (3) is larger than $(1 - 2^{-s})^j$. We conclude noting that j, the number of distinct prime factors of $\mathfrak{h}$, is at most $\log_2 N\mathfrak{h}$. □

Lemma 6. *Let $\mathcal{O}$ be a Dedekind domain, and M be a finite $\mathcal{O}$-module of cardinality h and let r be the minimal number of $\mathcal{O}$-generators of M. Let $g_1, \ldots, g_s$ be s independent uniform elements from M. Then, the probability that the set $\{g_1, \ldots, g_s\}$ generates M as an $\mathcal{O}$-module is*

$$\Pr\left[\mathcal{O}g_1 + \cdots + \mathcal{O}g_s = M\right] \geq \left(1 - 2^{-\lfloor s/r \rfloor}\right)^{\log_2 h}.$$

Proof. Since M is a torsion module over a Dedekind domain, there exist r ideals $\mathfrak{h}_1, \ldots, \mathfrak{h}_r$ such that $M = \bigoplus_{i=1}^{r} \mathcal{O}/\mathfrak{h}_i$; in particular, $\log_2 h = \sum_{i=1}^{r} \log_2 N\mathfrak{h}_i$. Consider the s' firsts random elements $g_1 \ldots g_{s'}$ where $s' = \lfloor s/r \rfloor$, and their projections $g'_1 \ldots g'_{s'}$ on the first component $\mathcal{O}/\mathfrak{h}_1$. By Lemma 5, $\{g'_1 \ldots g'_{s'}\}$ generates $\mathcal{O}/\mathfrak{h}_1$ with probability at least $(1 - 2^{-s'})^{\log_2 N\mathfrak{h}_1}$.

Suppose that $\{g'_1 \ldots g'_{s'}\}$ indeed generates $\mathcal{O}/\mathfrak{h}_1$. Let $M_1 = \mathcal{O}g_1 + \cdots + \mathcal{O}g_{s'}$. To conclude by induction, it suffices to note that M/M_1 is generated by (at most) $r - 1$ elements. □

Theorem 4. *Let H be a cyclic group, and M a finite, $\mathbb{Z}[H]$-module of cardinality h, and r be the minimal number of $\mathbb{Z}[H]$-generators of M. Let $g_1, \ldots, g_s$ be s independent uniform elements of M. The probability that the set $\{g_1, \ldots, g_s\}$ generates M as a $\mathbb{Z}[H]$-module is*

$$\Pr\left[\mathbb{Z}[H] \cdot g_1 + \cdots + \mathbb{Z}[H] \cdot g_s = M\right] \geq \left(1 - 2^{-\lfloor s/r \rfloor}\right)^{\log_2 h}.$$

Proof. Let t be the order of H. Observe that we have the decomposition

$$\mathbb{Z}[H] \cong \mathbb{Z}[X]/(X^t - 1) \cong \bigoplus_{d|t} \mathbb{Z}[X]/(\Phi_d(X)) \cong \bigoplus_{d|t} \mathbb{Z}[\omega_d].$$

For each $d \mid t$, define $e_d \in \mathbb{Z}[H]$ the idempotent which projects to the unit of $\mathbb{Z}[\omega_d]$ and to zero in all other components of the above direct sum. This is a system of fundamental idempotents of $\mathbb{Z}[H]$, and it induces a decomposition $\mathbb{Z}[H] = \bigoplus_{d|t} e_d\mathbb{Z}[H]$, where $e_d\mathbb{Z}[H] \cong \mathbb{Z}[\omega_d]$. In the following $e_d\mathbb{Z}[H]$ and $\mathbb{Z}[\omega_d]$ are canonically identified. In particular, we have that $M = \bigoplus_{d|t} e_d M$, and $e_d M$ may be viewed as $\mathbb{Z}[\omega_d]$-module.

First, note that $\log_2 h = \sum_{d|t} \log_2 h_d$, where h_d is the cardinality of $e_d M$. Second, each $e_d M$ is generated over $\mathbb{Z}[\omega_d]$ by at most r elements. Noting $\mathbb{Z}[\omega_d]$ is a Dedekind domain, we apply Lemma 6, over each component and conclude

$$\begin{aligned}
\Pr\left[\mathbb{Z}[H] \cdot g_1 + \cdots + \mathbb{Z}[H] \cdot g_s = M\right] &= \prod_{d|t} \Pr\left[\mathbb{Z}[\omega_d] \cdot g_1 + \cdots + \mathbb{Z}[\omega_d] \cdot g_s = e_d M\right] \\
&\geq \prod_{d|t} \left(1 - 2^{-\lfloor s/r \rfloor}\right)^{\log_2 h_d} \\
&= \left(1 - 2^{-\lfloor s/r \rfloor}\right)^{\log_2 h}.
\end{aligned}$$

□

Proof of Proposition 1. Note that G is trivial if and only if $K = \mathbb{Q}$, in which case Cl_K is trivial, and so is the proposition. Otherwise, observe that G splits as $\mathbb{Z}/2\mathbb{Z} \times H$ where H is a cyclic group, and the component $\mathbb{Z}/2\mathbb{Z}$ corresponds to the complex conjugation τ. Note that for any $x \in \mathrm{Cl}_K^-$, the orbits $\mathbb{Z}[G]x$ and $\mathbb{Z}[H]x$ coincide since $\tau \in G$ acts like $-1 \in \mathbb{Z}[H]$ on Cl_K^-. Therefore r is the minimal number of $\mathbb{Z}[H]$-generators of Cl_K^-. We obtain from Theorem 4 that the probability that $\{g_1, \ldots, g_s\}$ generates Cl_K^- as a $\mathbb{Z}[H]$-module is at least $(1 - 2^{-\lfloor s/r \rfloor})^{\log_2 h_K^-}$. For any $0 < x \leq 1/2$, we have $\ln(1-x) > -(3/2)x$. We get

$$\begin{aligned}\left(1 - 2^{-\lfloor s/r \rfloor}\right)^{\log_2 h_K^-} &= \exp\left(\log_2 h_K^- \ln\left(1 - 2^{-\lfloor s/r \rfloor}\right)\right) \\ &\geq \exp\left(-\frac{3}{2}\log_2(h_K^-) 2^{-\lfloor s/r \rfloor}\right).\end{aligned}$$

With $s \geq r(\log_2 \log_2(h_K^-) + \alpha)$, we get $\lfloor s/r \rfloor \geq \log_2 \log_2(h_K^-) + \alpha - 1$ and

$$\left(1 - 2^{-\lfloor s/r \rfloor}\right)^{\log_2 h_K^-} \geq \exp\left(-\frac{3}{2^\alpha}\right).$$

□