



HAL
open science

Mildly Short Vectors in Cyclotomic Ideal Lattices in Quantum Polynomial Time

Ronald Cramer, Léo Ducas, Benjamin Wesolowski

► **To cite this version:**

Ronald Cramer, Léo Ducas, Benjamin Wesolowski. Mildly Short Vectors in Cyclotomic Ideal Lattices in Quantum Polynomial Time. *Journal of the ACM (JACM)*, Association for Computing Machinery, 2021, 68 (2), pp.1-26. 10.1145/3431725 . hal-03102234

HAL Id: hal-03102234

<https://hal.archives-ouvertes.fr/hal-03102234>

Submitted on 7 Jan 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

MILDLY SHORT VECTORS IN CYCLOTOMIC IDEAL LATTICES IN QUANTUM POLYNOMIAL TIME

RONALD CRAMER^{1,2}, LÉO DUCAS¹ AND BENJAMIN WESOLOWSKI^{3,4}

¹Cryptography Group, CWI, Amsterdam, The Netherlands

²Mathematical Institute, Leiden University, The Netherlands

³Univ. Bordeaux, CNRS, Bordeaux INP, IMB, UMR 5251, F-33400, Talence, France

⁴INRIA, IMB, UMR 5251, F-33400, Talence, France

ABSTRACT. In this paper, we study the geometry of units and ideals of cyclotomic rings, and derive an algorithm to find a mildly short vector in any given cyclotomic ideal lattice in quantum polynomial time, under some plausible number-theoretic assumptions. More precisely, given an ideal lattice of the cyclotomic ring of conductor m , the algorithm finds an approximation of the shortest vector by a factor $\exp(\tilde{O}(\sqrt{m}))$. This result exposes an unexpected hardness gap between these structured lattices and general lattices: the best known polynomial time generic lattice algorithms can only reach an approximation factor $\exp(\tilde{O}(m))$. Following a recent series of attacks, these results call into question the hardness of various problems over structured lattices, such as Ideal-SVP and Ring-LWE, upon which relies the security of a number of cryptographic schemes.

NOTE. This article is an extended version of the conference paper [CDW17]. The results are generalised to arbitrary cyclotomic fields. In particular, we also extend some results of [CDPR16] to arbitrary cyclotomic fields. In addition, we prove the numerical stability of the method of [CDPR16]. These extended results appeared in the Ph.D. dissertation of the third author [Wes18a].

1. INTRODUCTION

1.1. Cyclotomic ideal lattices. Fix an integer $m > 2$ and a primitive m -th root of unity $\zeta_m \in \mathbb{C}$. Let $K = \mathbb{Q}(\zeta_m)$ be the cyclotomic field of conductor m . By the cyclotomic ring of conductor m , we shall mean $\mathcal{O}_K = \mathbb{Z}[\zeta_m]$, the ring of integers of K . The trace $\text{Tr} : K \rightarrow \mathbb{Q}$ induces an inner product on K as $\langle a, b \rangle = \text{Tr}(ab^\tau)$, where τ is complex conjugation. The field K is then a Hermitian vector space, and ideals in \mathcal{O}_K are Euclidean lattices, which are referred to as *cyclotomic ideal lattices*. In this article, we consider the problem of finding short vectors in such ideal lattices.

The problem of finding short vectors of a Euclidean lattice (the shortest vector problem, SVP, or its approximated version, approx-SVP) is a central hard problem in complexity theory. It is presumed to be hard even for quantum algorithms, and thanks to the worst-case to average-case reductions of Ajtai [Ajt99] and Regev [Reg09], it has become the theoretical foundation for many ‘post-quantum’ cryptographic constructions — cryptosystems that are meant to resist an adversary equipped with a quantum computer. Instantiations of these problems over

ideal lattices have attracted particular attention, as they allow very efficient implementations. The Ring-SIS [Mic07, LM06, PR06] and Ring-LWE [SSTX09, LPR13, PRSD17] problems were introduced, and shown to reduce to worst-case instances of Ideal-SVP (the specialisation of approx-SVP to ideal lattices). Both problems Ring-SIS and Ring-LWE have shown very versatile problems for building efficient cryptographic schemes. Typically, Ring-SIS, Ring-LWE and Ideal-SVP are instantiated over cyclotomic rings.

For some time, it seemed plausible that the ideal 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 series of works, initiated by Campbell et al. [CGS14], and followed by [BS16] and [CDPR16]. They show that in a cyclotomic ring of prime-power conductor, given a principal ideal, one can retrieve a short generator in quantum polynomial time. As a consequence, some cryptographic schemes were broken [SV10, GGH13, LSS14, CGS14], but it had a limited impact on the more general Ideal-SVP since principal ideals are a very sparse family of ideals for those fields.

1.2. Main result. In this paper, we tackle the general case of Ideal-SVP, for arbitrary ideal lattices in any cyclotomic ring. Studying the geometry of units and ideals of cyclotomic rings, we devise a quantum algorithm that given an ideal lattice of the cyclotomic ring of conductor m , finds an approximation of one of the shortest non-zero vectors (henceforth, *the shortest vector*, abusing language) by a factor $\exp(\tilde{O}(\sqrt{m}))$. Under some plausible (and carefully justified) number-theoretic assumptions, the algorithm runs in polynomial time. This is our main result, formalised as Theorem 5.1. In contrast, the best known polynomial time generic lattice algorithms can only reach an approximation factor $\exp(\tilde{O}(m))$. This unexpected hardness gap between approx-SVP in generic lattices and in cyclotomic ideal lattices is illustrated in Figure 1.

1.3. Overview. An integer $m > 2$ is fixed for the entire paper, as well as a primitive m -th root of unity $\zeta_m \in \mathbb{C}$. The cyclotomic field of conductor m is $K = \mathbb{Q}(\zeta_m)$, and the cyclotomic ring of conductor m is $\mathcal{O}_K = \mathbb{Z}[\zeta_m]$, the ring of integers of K . The degree of K over \mathbb{Q} is $\varphi(m)$ (where φ is Euler's totient function). Let Δ_K be the absolute value of the discriminant of K . The field $K^+ = \mathbb{Q}(\zeta_m + \zeta_m^{-1})$ is the maximal real subfield of K . The algebraic norm of an ideal \mathfrak{h} is denoted $N(\mathfrak{h})$. Let Cl_K be the class group of \mathcal{O}_K . The class of an ideal \mathfrak{h} is denoted $[\mathfrak{h}]$, and if two ideals \mathfrak{h} and \mathfrak{h}' are in the same class, we write $\mathfrak{h} \sim \mathfrak{h}'$. Let G denote the Galois group of the extension K/\mathbb{Q} , and let $\tau \in G$ be the complex conjugation of K .

1.3.1. Short vectors in ideal lattices. The field K is a Hermitian vector space over \mathbb{Q} for the inner product $\langle a, b \rangle = \text{Tr}(ab^\tau)$. The corresponding Euclidean norm is denoted $\|a\|$, and coincides with the ℓ_2 -norm induced by the Minkowski embedding

$$K \longrightarrow \mathbb{C}^{\varphi(m)} : a \longmapsto (a^\sigma)_{\sigma \in G}.$$

We also denote the ℓ_1 -norm and ℓ_∞ -norm by $\|a\|_1$ and $\|a\|_\infty$. The volume of an ideal \mathfrak{h} as a lattice relates to its algebraic norm by $\text{Vol}(\mathfrak{h}) = \sqrt{|\Delta_K|}N(\mathfrak{h})$. The length $\lambda_1(\mathfrak{h})$ of the shortest vector of \mathfrak{h} is determined by its algebraic norm up to a polynomial factor:

$$(1) \quad \frac{1}{\text{poly}(m)}N(\mathfrak{h})^{1/\varphi(m)} \leq \lambda_1(\mathfrak{h}) \leq \text{poly}(m)N(\mathfrak{h})^{1/\varphi(m)}.$$

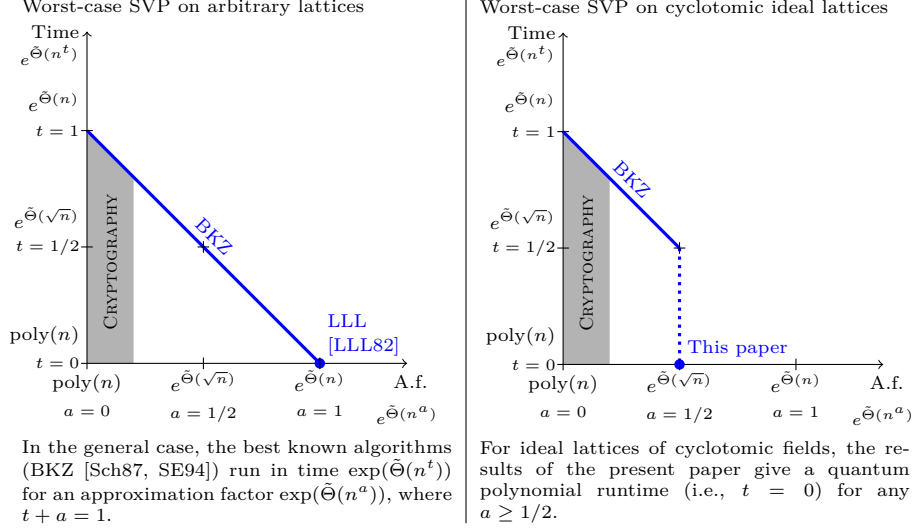


FIGURE 1. Best known (quantum) time–approximation factor tradeoffs to solve approx-SVP in arbitrary lattices (on the left) and in cyclotomic ideal lattices (on the right), in the worst case. The integer n is the dimension of the lattice; for a cyclotomic field of conductor m , this dimension is $n = \varphi(m)$. The approximation factors (A.f.) upon which the security of cryptographic schemes relies are typically between polynomial $\text{poly}(n)$ and quasi-polynomial $\exp(\text{polylog}(n))$ (represented as the grey area).

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 a shortest non-zero vector v of \mathfrak{h} is a multiple (a sub-ideal) of \mathfrak{h} , and that $\text{Vol}(v\mathcal{O}_K) \leq \|v\|^{\varphi(m)}$. The approximated Ideal-SVP in \mathfrak{h} for some approximation factor α consists in finding a vector in \mathfrak{h} of length at most $\alpha\lambda_1(\mathfrak{h})$. Our method is divided into two main steps. First, we show how to find a short vector in the case where the ideal is principal; then, we show how to reduce the general case to the principal case.

1.3.2. *Approx-SVP for principal ideals.* The principal case is dealt with in Section 3, via a study of the geometry of cyclotomic units. Following in the footsteps of [CDPR16], we generalise their method to cyclotomic fields of arbitrary conductor. Let \mathfrak{a} be a principal ideal in \mathcal{O}_K . The idea is to find a short *generator* of \mathfrak{a} (rather than just a short element). First, the algorithms of [BS16] allow to find an arbitrary generator g of \mathfrak{a} in quantum polynomial time. This generator g is typically extremely long, but it provides us with a convenient search space: the set of all generators of \mathfrak{a} is $g\mathcal{O}_K^\times$. We are looking for a short element of $g\mathcal{O}_K^\times$. The logarithmic embedding (see Definition 3.1) allows to transform this into a lattice problem: the image $\text{Log}(\mathcal{O}_K^\times)$ is a lattice of dimension $\varphi(m) - 1$, and the logarithmic embedding of $g\mathcal{O}_K^\times$ is the translation

$$\text{Log}(g) + \text{Log}(\mathcal{O}_K^\times)$$

of this lattice. We exhibit a full-rank set of short elements in $\text{Log}(\mathcal{O}_K^\times)$, which can be used to find a short vector in the translated lattice $\text{Log}(g) + \text{Log}(\mathcal{O}_K^\times)$. This short vector gives rise to a short element of $g\mathcal{O}_K^\times$, i.e., a short generator. We show in Theorem 3.7 that this method allows to find in quantum polynomial time an approximation of the shortest vector of \mathfrak{a} for the subexponential approximation factor $\exp(\tilde{O}(\sqrt{m}))$.

We note that this approach involves manipulating real numbers, leading to delicate numerical stability considerations. While [CDPR16] glosses over this issue, we provide a full, rigorous analysis.

1.3.3. The close principal multiple problem. To reduce the problem from arbitrary ideals to principal ideals, we introduce the *close principal multiple problem* (or CPM): given an arbitrary ideal \mathfrak{a} , find an integral ideal \mathfrak{b} such that $\mathfrak{a}\mathfrak{b}$ is principal, and $N(\mathfrak{b})$ is small. Suppose one can solve CPM with $N(\mathfrak{b}) \leq \exp(\tilde{O}(m^{1+c}))$, for some constant $c > 0$. Then, one can apply the aforementioned results to find a generator g of the principal ideal $\mathfrak{a}\mathfrak{b}$ such that

$$\|g\| \leq N(\mathfrak{a}\mathfrak{b})^{1/\varphi(m)} \exp(\tilde{O}(\sqrt{m})) \leq N(\mathfrak{a})^{1/\varphi(m)} \exp\left(\tilde{O}\left(m^{\max(1/2,c)}\right)\right).$$

Since $g \in \mathfrak{a}\mathfrak{b} \subset \mathfrak{a}$, one has found an approximation of the shortest vector of \mathfrak{a} for an approximation factor $\exp(\tilde{O}(m^{\max(1/2,c)}))$. This is asymptotically as good as the principal case when $c = 1/2$, and better than LLL for any $c < 1$.

1.3.4. Existence of close principal multiples. Before searching for a solution to the CPM problem, let us discuss whether a $\exp(\tilde{O}(m^{1+c}))$ -close principal multiple exists in general. A positive answer follows from the results of [JW15, Corollary 6.5] refining [JMV09, Corollary 1.3] for large degree number fields. setting a factor base of prime ideals $\mathfrak{B} = \{\mathfrak{p} \mid N\mathfrak{p} \leq m^{4+o(1)}\}$, for any class $C \in \text{Cl}_K$, there exists a non-negative small solution $e \in \mathbb{Z}_{\geq 0}^{\mathfrak{B}}$ to the class equation $[\prod \mathfrak{p}^{e_{\mathfrak{p}}}] = C$, of ℓ_1 -norm $\|e\|_1 \leq O(m^{1+o(1)})$. This proves, assuming the generalised Riemann hypothesis (GRH), the existence of a solution $\mathfrak{b} = \prod \mathfrak{p}^{e_{\mathfrak{p}}}$ to the CPM problem as small as $\exp(\tilde{O}(m^{1+c}))$ for $c = o(1)$.

This argument is based on the analysis of the expander properties of certain Cayley graphs on the class group. For our purpose, existence is not enough, as we wish to efficiently find a close principal multiple.

1.3.5. Intermezzo: short discrete logarithms and lattices. Our approach to solving the CPM problem is based, in part, on a judicious application of the following generic paradigm.

Let H be a finite Abelian group (denoted multiplicatively). Let k be a positive integer and let $\mathcal{G} = (g_1, \dots, g_k) \in H^k$ be a vector such that g_1, \dots, g_k constitute a generating set for H as a \mathbb{Z} -module. Consider the surjective \mathbb{Z} -linear map

$$\begin{aligned} \phi_{\mathcal{G}} : \mathbb{Z}^k &\longrightarrow H, \\ \alpha = (\alpha_1, \dots, \alpha_k) &\longmapsto \mathcal{G}^\alpha = \prod_{i=1}^k g_i^{\alpha_i}. \end{aligned}$$

The *representation problem* for H (given \mathcal{G}) is to find, given an arbitrary $h \in H$, some vector $\alpha \in \mathbb{Z}^k$ such that

$$\phi_{\mathcal{G}}(\alpha) = \mathcal{G}^\alpha = h.$$

Note that this is in fact a discrete logarithm problem in H , with respect to the generators g_1, \dots, g_k ; this kind of problem can typically be solved in quantum polynomial time following Shor [Sho97] and subsequent generalizations [EHKS14, BS16].

Now, suppose such solution is also required to be “short”: this is the *short representation problem*. This new constraint translates into a lattice problem as follows. Write $\Lambda_{\mathcal{G}} = \text{Ker } \phi_{\mathcal{G}}$. We have that $H \cong \mathbb{Z}^k / \Lambda_{\mathcal{G}}$ and that $\Lambda_{\mathcal{G}}$ is a full-rank lattice in \mathbb{Z}^k . If one can find $\beta \in \Lambda_{\mathcal{G}}$ such that $(\alpha - \beta) \in \mathbb{Z}^k$ is “small”, where $\alpha \in \mathbb{Z}^k$ is just any solution to the representation problem, then $\alpha - \beta$ is a solution to the short representation problem since

$$\mathcal{G}^{\alpha - \beta} = \mathcal{G}^{\alpha} = h.$$

Finding this β is known as the *close vector problem* in the lattice $\Lambda_{\mathcal{G}}$ with respect to α . In all generality, this problem is not known to be solvable in (quantum) polynomial time. However, it can be solved efficiently if a “short” basis of $\Lambda_{\mathcal{G}}$ (or a full-rank sublattice) is known¹.

The special case we have identified as being particularly relevant to our purposes is the following. Suppose G is a finite group (also written multiplicatively) acting on H , i.e., there is a morphism mapping G to the automorphism group of H . Then H may be viewed naturally as a module over the group ring $\mathbb{Z}[G]$. Note that $\mathbb{Z}[G]$ is a free \mathbb{Z} -module of finite rank $|G|$, the elements of G forming a basis. For simplicity, say G is Abelian (so that $\mathbb{Z}[G]$ is commutative).

For $g \in H$ and $\sigma \in G$, write the group action as $\sigma \cdot g = g^{\sigma}$. Let $\lambda = \sum_{\sigma \in G} \lambda_{\sigma} \sigma \in \mathbb{Z}[G]$, with coefficients $\lambda_{\sigma} \in \mathbb{Z}$, and write

$$g^{\lambda} := \prod_{\sigma \in G} (g^{\lambda_{\sigma}})^{\sigma}.$$

Suppose temporarily that H is cyclic as a $\mathbb{Z}[G]$ -module and suppose $g \in H$ is a generator. Note that the latter is equivalent to H being generated as a \mathbb{Z} -module by $\{g^{\sigma} \mid \sigma \in G\}$. Suppose furthermore that some nontrivial ideal $S \subset \mathbb{Z}[G]$ is given that *annihilates* H , i.e., $h^s = 1_H$ for any $h \in H$ and $s \in S$. Writing $\mathcal{G} = (g^{\sigma})_{\sigma \in G}$ (i.e., a vector defined by the G -orbit of $g \in H$), it follows that S naturally defines a sublattice of $\Lambda_{\mathcal{G}}$, by sending each element of S to its integer coordinate vector (under the stated basis of $\mathbb{Z}[G]$).

If this annihilating sublattice S is full-rank and if a short basis can be found, the strategy described above may be applicable so that, for any given $h \in H$, an $\alpha \in \mathbb{Z}[G]$ with “short” coordinate vector may be found such that $g^{\alpha} = h$, or, equivalently, a “short” $\alpha \in \mathbb{Z}^G$ (identifying the previous with its coordinate-vector) such that $\mathcal{G}^{\alpha} = h$, as desired.

This generalizes to the case where G is generated as a $\mathbb{Z}[G]$ -module by a larger set of elements (so H need not be $\mathbb{Z}[G]$ -cyclic), instead of a single one, essentially by applying the same strategy for each generator.

1.3.6. Application to CPM: class groups and the Stickelberger Theorem. Our idea for solving CPM is based on an application of the above paradigm to the class group Cl_K of the Galois number fields K of our interest. For the discussion, it is

¹Note that if a lattice is not of full rank, no close-vector algorithm can guarantee any distance bound, as any fundamental domain is unbounded.

convenient to recall that the Galois group G of the number field K acts on \mathcal{O}_K as well as on Cl_K in the natural way, so that each may be viewed as a *Galois module*, i.e., as a $\mathbb{Z}[G]$ -module.

In fact, we first solve CPM under the condition that the class of the target \mathcal{O}_K -ideal of CPM is an element of the subgroup Cl_K^- , the *minus-part* of the class group Cl_K . This choice is beneficial for technical reasons that will become clear shortly. Afterwards, we show how this can be extended to solving the general case, using additional ideas. Along the way, we make several assumptions that are only informally stated, so as not to detract from the conceptual aspects of our approach. These assumptions will be discussed in detail, together with the precise bounds achieved, in Section 5.

Let Cl_{K^+} be the class group of the maximal real subfield $K^+ = K \cap \mathbb{R}$. The relative norm map $N_{K/K^+} : \text{Cl}_K \rightarrow \text{Cl}_{K^+}$ on ideal classes (which sends the class of \mathfrak{h} to the class of $\mathfrak{h}^{1+\tau} \cap K^+$) is a surjection, and its kernel is the relative class group Cl_K^- . By definition, if \mathfrak{a} is an \mathcal{O}_K -ideal, it holds that $[\mathfrak{a}] \in \text{Cl}_K^-$ if and only if the \mathcal{O}_{K^+} -ideal $N_{K/K^+}(\mathfrak{a})$ is principal. Furthermore, for each \mathcal{O}_K -ideal \mathfrak{a} , it holds that $\mathfrak{a} \cdot \mathfrak{a}^\tau = N_{K/K^+}(\mathfrak{a}) \cdot \mathcal{O}_K$.² Hence, if $[\mathfrak{a}] \in \text{Cl}_K^-$, then $[\mathfrak{a}]^{-1} = [\mathfrak{a}^\tau]$. This property is very useful when one wants to work only with integral ideals of “small” norm: if \mathfrak{a} is integral, \mathfrak{a}^τ is an integral representative of $[\mathfrak{a}]^{-1}$ of same norm as \mathfrak{a} . Finding integral inverses in Cl_K with a good control on the norm is in general more difficult.

An immediate application to the representation problem in Cl_K^- is the following. Suppose we have an identity

$$\prod_{i=1}^m [\mathfrak{p}_i]^{\alpha_i} = [\mathfrak{a}]^{-1},$$

where all the classes are in Cl_K^- , the \mathfrak{p}_i 's are integral \mathcal{O}_K -ideals, and the α_i 's are in \mathbb{Z} . Then, for each index i such that $\alpha_i < 0$, we may redefine \mathfrak{p}_i as the integral \mathcal{O}_K -ideal \mathfrak{p}_i^τ (having the same norm as \mathfrak{p}_i) and redefine α_i by taking its absolute value. Thus, without loss of generality, this gives us the result that the \mathcal{O}_K -ideal

$$\mathfrak{a} \cdot \prod_{i=1}^m \mathfrak{p}_i^{\alpha_i}$$

is principal, and $\prod_{i=1}^m \mathfrak{p}_i^{\alpha_i}$ is integral ($\alpha_i \geq 0$). Also note that, if the exponent vector is “short” and if the respective algebraic norms of the \mathfrak{p}_i 's are “small,” then by multiplicativity of norm, $\prod_{i=1}^m \mathfrak{p}_i^{\alpha_i}$ is also “small” and would constitute a solution to CPM with target \mathfrak{a} . This is how, in essence, the short representation problem connects to the CPM problem.

Now, Biasse and Song [BS16] have shown that there is a quantum algorithm for computing discrete logarithms in the class group in quantum polynomial time. Therefore it remains to discuss

²For ideals of this form, the claim follows from Dedekind factorization in combination with basic theory of Galois extensions.

- how to get a small number of small-norm generators for Cl_K^- , i.e., their classes are presented by small-norm \mathcal{O}_K -ideals, and
- how to extract a short solution from a solution of the representation problem as produced by the algorithm of [BS16].

We start with the latter. The crucial property of Cl_K^- that furthers our purpose is the following. Consider the group ring $\mathbb{Z}[G]$ and the so-called *Stickelberger ideal* $S \subset \mathbb{Z}[G]$. It has the property that it annihilates Cl_K and, therefore, it annihilates Cl_K^- as well. Moreover, when considered as a lattice (in the way discussed in Section 1.3.5), it can be shown to have a relatively short basis which can be effectively computed. However, this lattice does not satisfy our earlier requirement to be of full-rank $|G|$, which presents an impediment to using close-vector algorithms with respect to S .

Fortunately, this issue can be solved by working with a quotient of $\mathbb{Z}[G]$. By construction, the element $1 + \tau \in \mathbb{Z}[G]$ acts trivially on Cl_K^- . So Cl_K^- is, in fact, a $\mathbb{Z}[G]/(1 + \tau)$ -module. The latter can be shown to be a free \mathbb{Z} -module of rank $|G|/2$. For instance, take a maximal subset of G such that no two elements are conjugate under τ ; then their classes form a \mathbb{Z} -basis. Furthermore, the image S^- of the Stickelberger ideal S under the canonical ring morphism $\mathbb{Z}[G] \rightarrow \mathbb{Z}[G]/(1 + \tau)$ annihilates Cl_K^- . But this time, S^- can be shown to be of full-rank $|G|/2$ when considered as a lattice. Moreover, effective computation of a relatively short basis can be shown to carry over from that of S .

Putting everything together, we get a solution for CPM in the case that the class of the target \mathfrak{a} is in Cl_K^- , assuming that we are given a small generating set of Cl_K^- as a $\mathbb{Z}[G]$ -module consisting of classes represented by small-norm ideals.

Concluding formally on which value of c can be achieved by the CPM algorithm sketched above is not straightforward, as it relies on the structure of Cl_K^- as a $\mathbb{Z}[G]$ -module. The structure of Cl_K^- remains in general quite elusive but it appears that it admits a small minimum number of generators as a $\mathbb{Z}[G]$ -module. For instance, Schoof [Sch98] computed that for all prime conductors $m \leq 509$, it is $\mathbb{Z}[G]$ -cyclic, i.e., it is generated by a single element as a $\mathbb{Z}[G]$ -module. In fact, we introduce a hypothesis stating that, asymptotically, few generators should suffice. Furthermore, we show that if Cl_K^- is generated by r classes, then $r \cdot \text{polylog}(n)$ many uniformly randomly selected classes in Cl_K^- will generate it as a $\mathbb{Z}[G]$ -module with overwhelming probability. As it is expected that classes of small random ideals behave similarly to uniformly random classes, that would settle the issue. Based on these heuristic arguments, we strongly believe that $c = 1/2$ is reachable at least for a dense family of conductors m , if not all. This leads to the main result of this paper: Approx-SVP in arbitrary ideals, Theorem 5.1.

Finally, to remove the condition that the target \mathcal{O}_K -ideal \mathfrak{a} satisfies $[\mathfrak{a}] \in \text{Cl}_K^-$, our goal is to first find a small-norm integral \mathcal{O}_K -ideal \mathfrak{a}' such that $[\mathfrak{a} \cdot \mathfrak{a}'] \in \text{Cl}_K^-$, followed by application of the strategy above. So it remains to discuss this first step. We exploit the fact that random walks in certain Cayley graphs of Cl_K , where the generators are ideals of small norm, will land rather quickly in Cl_K^- with high probability if the index of Cl_K^- in Cl_K is “small.” This index is equal to h_K^+ , the class number of K^+ . Taking the class of \mathfrak{a} as the starting point for the random walk, the desired \mathfrak{a}' is found efficiently (under GRH) as a product of random ideals

of small norm. Note that there is strong theoretical and computational evidence that h_K^+ is indeed small. In fact, there is a well-known conjecture that for certain classes of fields it equals 1, in which case Cl_K^- and Cl_K are identical, and the random walks can be omitted altogether.

1.4. Related works. The idea of exploiting the logarithmic unit-lattice to obtain shorter generators of a principal ideal can be traced back to Rekaya, Belfiore and Viterbo [RBV04]. The first cryptanalytic claim exploiting this idea [CGS14] concerned a specific distribution of principal ideals, for which an exceptionally small generator (the secret key) could be recovered exactly. This was proved for cyclotomic number fields of prime-power conductor in [CDPR16], by showing that the standard bases of logarithmic cyclotomic units is good enough to solve instances of the bounded-distance decoding (BDD) problem. The second result of [CDPR16] treats the case of arbitrary principal ideal (as opposed to principal ideals that admit an exceptionally small generator), by relating it to the close-vector problem (CVP) in the same lattice. Note that in this article, we only generalized this second result of [CDPR16], as the first one does not play a role in our final result. The first result has also been the object of a generalization to more (but not all) cyclotomic fields [HWB17]. The case of multi-quadratic fields was also studied in [BBdV⁺17], for which no quantum algorithms are required.

Since the conference version [CDW17] of the present article, this line of research has developed in several directions. By allowing exponential time pre-computation depending only on the number field, new trade-offs between the efficiency of the algorithm and the shortness of the resulting vector were obtained [PHS19]. This article considers both the classical and quantum settings, and applies to any number field. It was also shown that finding short vectors in module lattices of rank 2 could be reduced to the close vector problem in lattices constructed in a similar fashion but with a much larger dimension [LPSW19].

The results of [CDPR16, CDW17] have recently also been the object of a more concrete study, both using non-asymptotic analysis of the regulators and class numbers for formal lower bounds, and numerical experiments for empirical upper bounds [DPW19].

1.5. Roadmap. The rest of the paper is structured as follows. A few preliminaries on lattice and number theoretic algorithms are recalled in Section 2. In Section 3, we show how to exploit the geometry of cyclotomic units to solve Approx-SVP for principal ideals, generalising the results of [CDPR16]. In Section 4, we study the geometry of the Stickelberger ideal, with applications to the CPM problem. Our main result on Approx-SVP for arbitrary cyclotomic ideals is the object of Section 5. The Galois module structure of Cl_K^- is studied in Section 6.

2. PRELIMINARIES

2.1. Close vector algorithms. If $B = [b_1, \dots, b_k] \in \mathbb{R}^{n \times k}$ is a matrix composed of linearly independent column vectors $b_i \in \mathbb{R}^n$, we denote by $\tilde{B} = [\tilde{b}_1, \dots, \tilde{b}_k] \in \mathbb{R}^{n \times k}$ its Gram-Schmidt orthogonalization. Moreover, we denote by $\mathcal{P}(B)$ the centered parallelepiped spanned by B , defined as

$$\mathcal{P}(B) = B \cdot [-1/2, 1/2]^k = \left\{ \sum x_i b_i \mid x_i \in [-1/2, 1/2] \right\}.$$

We recall that if B is the basis of a full-rank lattice $L \subset \mathbb{R}^n$, both $\mathcal{P}(B)$ and $\mathcal{P}(\tilde{B})$ are fundamental domains for L on \mathbb{R}^n . These fundamental domains admit polynomial-time reduction algorithms. For the latter fundamental domain $\mathcal{P}(\tilde{B})$, this algorithm is referred to as *Size-Reduction* or as the *Nearest-Plane algorithm* [LLL82, Bab86].

Lemma 2.1. *There is a classical deterministic polynomial time algorithm $\text{NP}(B, t)$, that given the basis $B \in \mathbb{Q}^{n \times k}$ of a lattice $L \subset \mathbb{R}^n$, and a target $t \in L \otimes \mathbb{Q}$, outputs a pair (v, d) where $v \in \mathbb{Z}^k$, $d \in \mathcal{P}(\tilde{B})$ and $t = Bv + d$.*

Given a short basis B of a lattice L , the above algorithm can be used to find a close lattice point v to any target t . In fact, it is even sufficient to know a set of short vectors of L that span $L \otimes \mathbb{R}$.

Corollary 2.2. *There is a classical deterministic polynomial time algorithm $\text{CV}(W, t)$, that given a finite set W of k vectors of a lattice $L \subset \mathbb{Q}^n$ that spans $L \otimes \mathbb{Q}$, and a target $t \in L \otimes \mathbb{Q}$, outputs a vector $v \in \mathbb{Z}^k$ such that*

$$(2) \quad \|W \cdot v - t\| \leq 1/2 \cdot \sqrt{n} \cdot \max_{w \in W} \|w\|$$

$$(3) \quad \|W \cdot v - t\|_1 \leq 1/2 \cdot n \cdot \max_{w \in W} \|w\|.$$

Proof. First, we construct a maximal set of linearly independent vectors $C \subset W$, which can be done in deterministic polynomial time in a greedy manner. It holds that C generates a full-rank sub-lattice of L , in particular setting $(v', d) = \text{NP}(C, t)$ it holds that $Cv' - t = d \in \mathcal{P}(\tilde{C})$, and by Euclidean additivity

$$\begin{aligned} \|d\|^2 &\leq 1/4 \cdot \sum_i \|\tilde{c}_i\|^2 \\ &\leq 1/4 \cdot n \cdot \max_i \|\tilde{c}_i\|^2 \\ &\leq 1/4 \cdot n \cdot \max_{w \in W} \|w\|^2. \end{aligned}$$

It remains to pad v' to v with appropriately placed zeros to conclude the proof of the first item. The second item is simply derived from the first by Cauchy-Schwartz inequality. \square

We also require an algorithm to find a close vector with respect to the ℓ_∞ -norm, yet in the worst-case the algorithm may not provide a close enough vector. This can be improved by resorting to a probabilistic approach, thanks to the following proposition.

Proposition 2.3. *If $X \in \mathbb{R}^{n \times k}$ has orthogonal rows, and if x is uniformly distributed over $\mathcal{P}(X)$, then*

$$\|x\|_\infty \leq \tau \cdot \max_i \|x_i\|$$

holds except with probability at most $2n \cdot \exp(-2\tau^2)$.

Proof. First, let us write $(, 0) = QD$ where D is a diagonal matrix with coefficients $(\|x_1\|, \dots, \|x_k\|, 0, \dots, 0)$ and Q is an orthogonal matrix (i.e., $QQ^t = Q^tQ = I$).

We write $x = QDy$ where y is uniform in $[-1/2, 1/2]^k$. In particular, for each j we have $x_j = \sum_i Q_{j,i} D_{i,i} y_i$. Hoeffding's bound states that the probability that $|x_j| \geq s$ is less than $2 \exp(-2s^2 / \sum_i (Q_{j,i} D_{i,i})^2)$. Note that $\sum_i (Q_{j,i} D_{i,i})^2 \leq \max_i \|x_i\|^2$. Taking $s = \tau \cdot \max_i \|x_i\|$, one concludes by the union bound over all j 's. \square

Lemma 2.4. *There is a classical randomized polynomial time algorithm $\text{CV}_\infty(W, t)$ that given a set W of k vectors of a lattice $L \subset \mathbb{R}^n$ that spans $L \otimes \mathbb{Q}$, and a target $t \in L \otimes \mathbb{Q}$, outputs a vector $v \in \mathbb{Z}^k$ such that*

$$(4) \quad \|W \cdot v - t\|_\infty \leq \sqrt{2 \cdot \ln(8n)} \cdot \max_{w \in W} \|w\|$$

with probability at least $1/2$.

Proof. First, we construct a set of linearly independent vectors $C \subset W$, and consider the lattice L' generated by C . Sample a uniform $p \in \mathcal{P}(\tilde{C})$, compute $(v, d) = \text{NP}(C, t + p)$. Note that $\|W \cdot v - t\|_\infty \leq \|W \cdot v - (t + p)\|_\infty + \|p\|_\infty$. Because p is uniform over a fundamental domain of L' , it holds that $t + p \bmod L'$ is uniform, therefore $(W \cdot v - (t + p))$ is uniform over $\mathcal{P}(\tilde{C})$.

We apply Proposition 2.3 to p (resp. $W \cdot v - (t + p)$) with $\tau = \sqrt{1/2 \cdot \ln(8n)}$: $\|p\|_\infty$ (resp. $\|W \cdot v - (t + p)\|_\infty$) is less than $\sqrt{1/2 \cdot \ln(8n)} \cdot \max_{w \in W} \|w\|$ except with probability at most $1/4$. One concludes by a union bound. \square

2.2. Representation of elements of \mathcal{O}_K . The standard representation of an element $\alpha \in \mathcal{O}_K$ is the vector $\vec{\alpha} = (\alpha_0, \dots, \alpha_{\varphi(m)-1})$ in the standard power \mathbb{Z} -basis of \mathcal{O}_K , i.e., the sequence of coefficients of the polynomial $\alpha = \sum \alpha_i X^i \bmod \Phi_m(X)$ where Φ_m denotes the m -th cyclotomic polynomial. A fractional element $\alpha \in K$ is uniquely represented as $\frac{1}{q} \cdot \vec{\alpha}'$ where q is a positive integer coprime to the greatest common divisor of the coefficients of α' .

Often, algorithms for \mathcal{O}_K have to manipulate very large elements, so large that a standard representation would have an exponential length. It is the case for instance for the quantum polynomial time algorithms of [BS16]. This issue is resolved by using a *compact representation*: a compact representation of an element $\alpha \in K$ is a sequence of elements in $\gamma_1, \dots, \gamma_\ell \in \mathcal{O}_K$ in the standard representation and integers k_1, \dots, k_ℓ such that $\alpha = \prod_{i=1}^\ell \gamma_i^{k_i}$.

If it is guaranteed that $\alpha \in K$ is short, one can efficiently recover a standard representation from a compact one. In [CDPR16], this is dealt with by resorting to floating-point approximations, yet Biase [Bia18] suggested to instead perform fast modular exponentiation.

Lemma 2.5 (Formalized from [Bia18]). *Given elements $\gamma_1, \dots, \gamma_\ell \in K$ in standard representation and integers $k_1, \dots, k_\ell, q, B \in \mathbb{Z}$, assuming that $\alpha = \prod_{i=1}^\ell \gamma_i^{k_i}$ satisfies $q\alpha \in \mathcal{O}_K$ and $\|\alpha\| \leq B$, one can compute α in standard representation in polynomial time in the size of the input.*

Proof. Choose $Q \geq 2qB$, and compute $q\alpha = q \prod_{i=1}^\ell \gamma_i^{k_i} \bmod Q$ using fast modular exponentiation. Recover $q\alpha$ as the representative of $q\alpha \bmod Q$ with coefficients in $[-qB, qB]$. \square

2.3. 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, BEF⁺17]). Yet, building on the recent advances on quantum algorithms for the Hidden Subgroup Problem in large dimensions [EHKS14, dBDF20], Biase 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.

Theorem 2.6 ([BS16, Theorem 1.3]). *There is a quantum algorithm for deciding if an ideal $\mathfrak{a} \subseteq \mathcal{O}$ of an order \mathcal{O} in a number field K is principal, and for computing $\alpha \in \mathcal{O}$ in compact representation such that $\mathfrak{a} = (\alpha)$, in polynomial time in the parameters $\log(N(\mathfrak{a}))$ and $\log(|\Delta|)$, where Δ is the discriminant of \mathcal{O} .*

The Biasse-Song [BS16] algorithm for S -unit group computations also allows to solve the class group discrete logarithm problem³: 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} .

Proposition 2.7 ([BS16]). *Let \mathfrak{B} be a set of prime ideals generating a subgroup H of Cl_K . There exists a quantum algorithm $\text{CIDL}_{\mathfrak{B}}$ which, when given as input any ideal \mathfrak{a} in \mathcal{O}_K such that $[\mathfrak{a}] \in H$, outputs a vector $y \in \mathbb{Z}^{\mathfrak{B}}$ such that $\prod \mathfrak{p}^{y_{\mathfrak{p}}} \sim \mathfrak{a}$, and runs in polynomial time in $\max_{\mathfrak{p} \in \mathfrak{B}} \log(N\mathfrak{p})$, $\log(N\mathfrak{a})$, $|\mathfrak{B}|$ and $\log(\Delta_K)$, where Δ_K is the absolute value of the discriminant of K .*

Proof. Given Theorem 1.1 of [BS16] the proof of this corollary is standard, and known as the linear-algebra step of index calculus methods.

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

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

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 $\{\vec{v}_i\}_{i=1}^c$ 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 $\vec{v}_i = ((w_{i,\mathfrak{p}})_{\mathfrak{p} \in \mathfrak{B}}, (v_{i,1}, \dots, v_{i,k}))$. 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 $\vec{x} \in \mathbb{Z}^c$ which can be computed in polynomial time. We obtain

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

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

3. THE GEOMETRY OF CYCLOTOMIC UNITS

In this section, we study the geometry of cyclotomic units, and as an application, we provide a quantum algorithm for approx-SVP in principal ideals, Algorithm 2. Suppose g is a generator of some principal ideal \mathfrak{a} . Then, $g\mathcal{O}_K^\times$ is the set of all generators of \mathfrak{a} . Generators of short Euclidian norm can be studied and found by investigating the geometry of the unit group \mathcal{O}_K^\times , and more specifically of the lattice $\text{Log}(\mathcal{O}_K^\times)$ obtained via the logarithmic embedding. The main results exploit

³Proposition 2.7 is a corollary of [BS16, Theorem 1.1]. We only include technical details for completeness.

the subgroup $C \subset \mathcal{O}_K^\times$ of cyclotomic units, whose corresponding lattice $\text{Log}(C)$ admits an efficiently computable set of short generators.

3.1. The logarithmic embedding and cyclotomic units. Recall that G denotes the Galois group $\text{Gal}(K/\mathbb{Q})$, which is isomorphic to $(\mathbb{Z}/m\mathbb{Z})^\times$, and $\tau \in G$ is complex conjugation.

Definition 3.1 (Logarithmic embedding). *The logarithmic embedding of K is*

$$\begin{aligned} \text{Log} : K^\times &\longrightarrow \mathbb{R}[G]/(1-\tau) \\ a &\longmapsto \sum_{\sigma \in G} \log(|a^\sigma|) \cdot \sigma^{-1}. \end{aligned}$$

It is easy to check that this is a morphism of $\mathbb{Z}[G]$ -modules. The ring $\mathbb{R}[G]/(1-\tau)$ also has a geometric structure: given any set $B \subset G$ of representatives of $G/\langle\tau\rangle$, the projection of B to $\mathbb{R}[G]/(1-\tau)$ forms an \mathbb{R} -basis (which does not actually depend on the choice of B) and we consider the (vector space) norms on $\mathbb{R}[G]/(1-\tau)$ coming from the induced isomorphism with $\mathbb{R}^{\varphi(m)/2}$.

The kernel of the logarithmic embedding restricted to \mathcal{O}_K^\times is the subgroup generated by -1 and ζ_m . Dirichlet's unit theorem implies that $\text{Log}(\mathcal{O}_K^\times)$ is a full-rank lattice in the linear subspace of $\mathbb{R}[G]/(1-\tau)$ orthogonal to $s(G) = \sum_{\sigma \in G} \sigma$.

Definition 3.2 (Cyclotomic units). *Let V be the multiplicative group generated by*

$$\{\pm\zeta_m\} \cup \{1 - \zeta_m^j \mid j = 1, \dots, m-1\}.$$

The group of cyclotomic units of K is the intersection $C = V \cap \mathcal{O}_K^\times$.

Theorem 3.3. *The lattice $\text{Log}(C)$ has full rank in $\text{Log}(\mathcal{O}_K^\times)$.*

Proof. From [Sin78], the group $C^+ = C \cap K^+$ has finite index in the group of real units $E^+ = \mathcal{O}_K^\times \cap K^+$. Let W be the multiplicative group generated by -1 and ζ_m . From [Was12, Th. 4.12], the group WE^+ has index 1 or 2 in \mathcal{O}_K^\times . Since W is the kernel of $\text{Log} : \mathcal{O}_K^\times \rightarrow \mathbb{R}[G]$, we get

$$\begin{aligned} [\text{Log}(\mathcal{O}_K^\times) : \text{Log}(C^+)] &= [\text{Log}(\mathcal{O}_K^\times) : \text{Log}(E^+)] [\text{Log}(E^+) : \text{Log}(C^+)] \\ &= [\mathcal{O}_K^\times : WE^+] [E^+ : C^+], \end{aligned}$$

which is finite. Therefore $[\text{Log}(\mathcal{O}_K^\times) : \text{Log}(C)]$ is also finite. \square

3.2. Short generating vectors of the cyclotomic units. We are interested in finding short generators of the lattice $\text{Log}(C)$. Let $m = p_1^{\alpha_1} \dots p_k^{\alpha_k}$ be the prime factorization of m , and for any index i let $m_i = mp_i^{-\alpha_i}$. For $0 < j < m$, let

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

Theorem 3.4 ([Kuĉ92, Th. 4.2]). *The lattice $\text{Log}(C)$ is generated by a subset of $\{\text{Log}(v_j) \mid 0 < j < m\}$.*

Lemma 3.5. *For any integer j not divisible by m , we have $\|\text{Log}(1 - \zeta_m^j)\| = O(\sqrt{m})$.*

Proof. Write $j = ab$, where a divides m and $(b, m/a) = 1$.

$$\|\mathrm{Log}(1 - \zeta_m^j)\|^2 = \sum_{i \in (\mathbb{Z}/m\mathbb{Z})^\times} (\log|1 - \zeta_m^{ij}|)^2 = \sum_{i \in (\mathbb{Z}/m\mathbb{Z})^\times} (\log|1 - \zeta_{m/a}^{ib}|)^2.$$

The natural group homomorphism $(\mathbb{Z}/m\mathbb{Z})^\times \rightarrow (\mathbb{Z}/(m/a)\mathbb{Z})^\times$ is a surjection, so its kernel has cardinality $\varphi(m)/\varphi(m/a)$, and we obtain

$$\begin{aligned} \|\mathrm{Log}(1 - \zeta_m^j)\|^2 &= \frac{\varphi(m)}{\varphi(m/a)} \sum_{i \in (\mathbb{Z}/(m/a)\mathbb{Z})^\times} (\log|1 - \zeta_{m/a}^i|)^2 \\ &= \frac{\varphi(m)}{\varphi(m/a)} \sum_{i \in (\mathbb{Z}/(m/a)\mathbb{Z})^\times} (\log|2 \sin(\pi ia/m)|)^2 \\ &\leq 2a \sum_{i=1}^{\lfloor m/2a \rfloor} (\log(2 \sin(\pi ia/m)))^2 \\ (5) \qquad &= 2a \sum_{i=1}^{\lfloor m/2a \rfloor} f(ia/m), \end{aligned}$$

where $f : [0, 1/2] \rightarrow \mathbb{R}$ is defined as $f(x) = (\log(2 \sin(\pi x)))^2$. The inequality $\varphi(m)/\varphi(m/a) \leq a$ follows from the observation that for any positive integers x and y , we have $x/\varphi(x) \leq (xy)/\varphi(xy)$ (it follows from the simple case where y is prime). Since $f(x) \leq (\log 2)^2$ for $1/6 \leq x \leq 1/2$, the terms in Equation (5) coming from $i > \lfloor m/6a \rfloor$ sum to at most $O(m)$. It remains to estimate the contribution of the remaining terms. Since $\sin(\pi x) \geq 2x$ for $0 \leq x \leq 1/2$, we have

$$2a \sum_{i=1}^{\lfloor m/6a \rfloor} f(ia/m) \leq 2a \sum_{i=1}^{\lfloor m/6a \rfloor} (\log(4ia/m))^2 \leq 2a \frac{m}{a} \int_0^{1/6} (\log(4x))^2 dx = O(m),$$

where the last equality follows from $\int_0^y (\log x)^2 dx = y((\log y)^2 - 2 \log y + 2)$. \square

3.3. Finding a short generator of a principal ideal.

Theorem 3.6. *There is a randomized algorithm SHORTGENERATOR (Algorithm 1) that for any $g \in \mathcal{O}_K$ (in compact representation), finds an element $h \in \mathcal{O}_K$ (in compact representation) such that $g\mathcal{O}_K = h\mathcal{O}_K$ and*

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

and runs in polynomial time in the size of the input.

Proof. A technical hurdle for this algorithm is the need to resort to approximate computations. We sketch here the proof ignoring this issue, by assuming that all operations on \mathbb{R} can be performed in polynomial time. The full proof accounting for precision issues is deferred to Section 3.4.

Recall that g is given in compact representation $(\gamma_i, k_i)_{i=1}^\ell$, where $g = \prod_{i=1}^\ell \gamma_i^{k_i}$. Using the notation from Algorithm 1, $t = t' - t''$ is the orthogonal projection of $t' = \sum_{i=1}^\ell k_i \mathrm{Log}(\gamma_i)$ on the subspace $\mathrm{Log}(\mathcal{O}_K^\times) \otimes \mathbb{R}$. Let $W = (w_1, \dots, w_{m-1})$ where $w_i = \mathrm{Log}(v_i)$. From Theorem 3.4, W is a set of generators of $\mathrm{Log}(C)$, and each w_i writes either as $\mathrm{Log}(1 - \zeta_m^i)$ or $\mathrm{Log}(1 - \zeta_m^i) - \mathrm{Log}(1 - \zeta_m^{m-i})$; Applying Lemma 3.5

Algorithm 1 SHORTGENERATOR(g): finds a short generator of $g\mathcal{O}_K$.

Require: An element $g \in \mathcal{O}_K$ in compact representation $(\gamma_i, k_i)_{i=1}^\ell$.

Ensure: The compact representation of a short element generating $g\mathcal{O}_K$.

- 1: $W = (w_1, \dots, w_{m-1})$ where $w_i = \text{Log}(v_i)$; $s(G) = \sum_{\sigma \in G} \sigma \in \mathbb{R}[G]/(1 - \tau)$;
 - 2: $t' \leftarrow \sum_{i=1}^\ell k_i \text{Log}(\gamma_i)$;
 - 3: $t'' \leftarrow 1/\varphi(m) \cdot \log(N(g)) \cdot s(G)$;
 - 4: $t \leftarrow t' - t'' \in \text{Log}(\mathcal{O}_K^\times) \otimes \mathbb{R}$;
 - 5: **repeat**
 - 6: $x \leftarrow \text{CV}_\infty(W, t)$; {randomized, see Lemma 2.4}
 - 7: **until** $\|W \cdot x - t\|_\infty \leq \sqrt{2 \cdot \log(4\varphi(m))} \cdot \max_{w \in W} \|w\|$
 - 8: **return** concatenation of $(\gamma_i, k_i)_{i=1}^\ell$ and $(v_i, -x_i)_{i=1}^{m-1}$.
-

once or twice for each w_i , we get $\max_{w \in W} \|w\| = O(\sqrt{m})$. Calls to the randomized algorithm $\text{CV}_\infty(W, t)$ are repeated until the output x satisfies

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

According to Lemma 2.4, this procedure terminates in average polynomial time. Let h be the element with compact representation $(\gamma_i, k_i)_{i=1}^\ell \widehat{\ } (v_i, -x_i)_{i=1}^{m-1}$ (where $\widehat{\ }$ denotes the concatenation of sequences). We have

$$\begin{aligned} \|h\|_\infty &\leq \exp(\|\text{Log}(g) - W \cdot x\|_\infty) \\ &\leq \exp(\|t + t'' - W \cdot x\|_\infty) \\ &\leq \exp(\|t''\|_\infty) \cdot \exp(\|t - W \cdot x\|_\infty) \\ &\leq \exp(\|1/\varphi(m) \cdot \log(N(g)) \cdot s(G)\|_\infty) \cdot \exp\left(\sqrt{2 \cdot \log(4\varphi(m))} \cdot \max_{w \in W} \|w\|\right) \\ &\leq N(g)^{1/\varphi(m)} \cdot \exp\left(O\left(\sqrt{m \log m}\right)\right). \end{aligned}$$

We conclude from the inequality $\|h\| \leq \sqrt{\varphi(m)} \|h\|_\infty$. □

Theorem 3.7 (Approx-SVP for cyclotomic, principal ideals). *There is a quantum algorithm PRINCIPALIDEALSVP (Algorithm 2) that, when given a principal ideal \mathfrak{a} in the cyclotomic ring of conductor m , finds a generator of Euclidean norm $\exp(O(\sqrt{m \log m})) \cdot N(\mathfrak{a})^{1/\varphi(m)}$, in expected polynomial time in m and $\log N(\mathfrak{a})$. This generator approximates SVP in the lattice \mathfrak{a} with an approximation factor $\exp(O(\sqrt{m \log m}))$.*

Algorithm 2 PRINCIPALIDEALSVP(\mathfrak{a}): solves Approx-SVP in a principal ideal \mathfrak{a} .

Require: A principal ideal \mathfrak{a} of \mathcal{O}_K .

Ensure: The compact representation of a short generator of \mathfrak{a} .

- 1: $g \leftarrow \text{PIP}(\mathfrak{a})$; {PIP algorithm [BS16, Theorem 1.3]}
 - 2: $h \leftarrow \text{SHORTGENERATOR}(g)$; {Algorithm 1}
 - 3: **return** h .
-

Proof. First apply the quantum algorithm of [BS16, Theorem 1.3] on \mathfrak{a} . Since \mathfrak{a} is principal, it returns an element $g \in \mathcal{O}_K$ in compact representation such that

$\mathfrak{a} = g\mathcal{O}_K$, in polynomial time in $\log(N(\mathfrak{a}))$ and m . From Theorem 3.6, Algorithm 1 returns another generator h of \mathfrak{a} such that

$$\|h\| = \exp\left(O\left(\sqrt{m \log m}\right)\right) \cdot N(\mathfrak{a})^{1/\varphi(m)},$$

also in polynomial time. It follows from (1) that h approximates SVP in \mathfrak{a} with an approximation factor $\exp(O(\sqrt{m \log m}))$. \square

3.4. Numerical stability. In this section, we prove that we can round all the logarithms $\text{Log}(\gamma_i)$ and $\text{Log}(v_i)$ to \mathbb{Q} with polynomially many bits of precision and still obtain a small generator h . Let $n = \varphi(m)/2$ and suppose without loss of generality that the first $n - 1$ vectors w_1, \dots, w_{n-1} are linearly independent.

3.4.1. The algorithm. Fix a set of n representatives of the cosets $G/\langle\tau\rangle$; they form an \mathbb{R} -basis for $\mathbb{R}[G]/(1 - \tau)$. In this basis, consider the matrices

$$\begin{aligned} L &= (\text{Log}(\gamma_i))_{i=1}^\ell, \text{ and} \\ W &= (w_i)_{i=1}^{n-1} = (\text{Log}(v_i))_{i=1}^{n-1}. \end{aligned}$$

Write $k = (k_i)_{i=1}^\ell$. Set

$$p = \lceil \log_2 (\max(\|L \cdot k\|, \|k\|_\infty, 10\sqrt{n}\|W\|^{2n-3})) \rceil$$

and note that p is polynomial in the size of the input (Lemma 3.5, together with the Cauchy-Schwarz inequality, implies $\|W\| = O(m)$). Let $\varepsilon = 2^{-(p+m^2)}$, and compute an approximation \bar{L} with coefficients in $\varepsilon\mathbb{Z}$ such that $\|L - \bar{L}\|_\infty \leq \varepsilon$. Now, we want an approximation \bar{W} of W with coefficients in $\varepsilon\mathbb{Z}$ such that $\|W - \bar{W}\|_\infty \leq \varepsilon$ and each vector \bar{w}_i still lies in $\text{Log}(\mathcal{O}_K^\times) \otimes \mathbb{R}$. To do so, find some approximation \widetilde{W} such that $\|W - \widetilde{W}\|_\infty \leq \varepsilon/2$, and let

$$\bar{w}_i = \widetilde{w}_i - \frac{1}{n} \sum_{j=1}^n \widetilde{w}_{i,j} s(G) \in \text{Log}(\mathcal{O}_K^\times) \otimes \mathbb{R},$$

which satisfies $\|\bar{W} - \widetilde{W}\|_\infty \leq \varepsilon/2$.

We proceed with the same computation as in the proof of Theorem 3.6, using these approximate values. Compute $\bar{t}' = \bar{L} \cdot k$, and project \bar{t}' orthogonally to $s(G)$, that is decompose $\bar{t}' = \bar{t} + \bar{t}''$ such that $\bar{t} \in \text{Log}(\mathcal{O}_K^\times) \otimes \mathbb{R}$ and $\bar{t}'' \in s(G) \cdot \mathbb{R}$. Repeatedly call the randomized algorithm $\bar{x} \leftarrow \text{CV}_\infty(\bar{W}, \bar{t})$ until the output \bar{x} satisfies

$$(6) \quad \|\bar{W} \cdot \bar{x} - \bar{t}\|_\infty \leq \sqrt{2 \cdot \log(8n)} \cdot \max_{w \in W} \|w\|.$$

According to Lemma 2.4, this procedure terminates in average polynomial time. We output the compact representation $(\gamma_i, k_i)_{i=1}^\ell \wedge (v_i, -\bar{x}_i)_{i=1}^{n-1}$ of h .

3.4.2. Analysis. We now prove that the output h is short. We have

$$\begin{aligned} \|h\|_\infty &\leq \exp(\|L \cdot k - W \cdot \bar{x}\|_\infty) \\ &\leq \exp(\|\bar{t}' - \bar{W} \cdot \bar{x}\|_\infty) \cdot \exp(\|(L - \bar{L}) \cdot k\|_\infty + \|(W - \bar{W}) \cdot \bar{x}\|_\infty) \end{aligned}$$

From Lemma 3.5, we have $\max_{w \in W} \|w\| \leq O(\sqrt{m}) + \sqrt{n}\varepsilon \leq O(\sqrt{m})$, and together with (6) we can bound the first factor as

$$\exp(\|\bar{t}' - \bar{W} \cdot \bar{x}\|_\infty) \leq \exp(\|\bar{t}''\|_\infty) \cdot \exp(O(\sqrt{m \log m})).$$

Secondly, since t'', \bar{t}'' are respectively the projections of $t' = Lk$ and $\bar{t}' = \bar{L}k$, it holds that $\|t'' - \bar{t}''\|_\infty \leq n\|t' - \bar{t}'\|_\infty$. So we get that $\|h\|_\infty$ is at most

$$N(g)^{1/\varphi(m)} \cdot \exp(O(\sqrt{m \log m})) \cdot \exp((n+1)\|(L - \bar{L}) \cdot k\|_\infty + \|(W - \bar{W}) \cdot \bar{x}\|_\infty).$$

Next, note that we have $\|(L - \bar{L}) \cdot k\|_\infty \leq \|k\|_\infty \cdot \varepsilon \leq 2^{-m^2}$, so:

$$(7) \quad \|h\|_\infty \leq N(g)^{1/\varphi(m)} \cdot \exp(O(\sqrt{m \log m})) \cdot \exp(2^{-m^2+o(m)} + \|(W - \bar{W}) \cdot \bar{x}\|_\infty).$$

It remains to bound $\|\bar{x}\|$. For any matrix A , write A^+ for its pseudoinverse.

Lemma 3.8. *We have $\|W^+\| \leq 5 \cdot \|W\|^{\varphi(m)-3}$.*

Proof. The elements w_1, \dots, w_{n-1} generate a sublattice of $\text{Log}(\mathcal{O}_K^\times)$, so

$$\det(W^t W) \geq \det(\text{Log}(\mathcal{O}_K^\times)) = R_K \sqrt{n},$$

where R_K denotes the regulator of the field K . Writing $\lambda_1 \leq \dots \leq \lambda_{n-1} = \|W\|^2$ the eigenvalues of $W^t W$, we have

$$\|(W^t W)^{-1}\| = \frac{1}{\lambda_1} = \frac{\prod_{i=2}^{n-1} \lambda_i}{\det(W^t W)} \leq \frac{\|W\|^{2(n-2)}}{R_K \sqrt{n}}.$$

From [Fri89, Theorem B], we have $R_K > \text{lcm}(2, m)/10$ (except for $m = 10$, for which we have $R_K > 0.96$). Since W has full column rank, $W^+ = (W^t W)^{-1} W^t$. We conclude that

$$\|W^+\| \leq \|W\| \|(W^t W)^{-1}\| \leq \frac{5\|W\|^{2n-3}}{\sqrt{n}} \leq 5 \cdot \|W\|^{2n-3}.$$

□

Lemma 3.9. *We have $\|\bar{W}^+\| \leq 2\|W^+\|$.*

Proof. Let $E = \bar{W} - W$. First observe that $\bar{W} = AW$, with $A = I + EW^+$. Now, we have

$$\|I - A\| = \|EW^+\| \leq \|E\| \|W^+\| \leq \sqrt{n} \|E\|_\infty \|W^+\| \leq \varepsilon \cdot \sqrt{n} \cdot 5 \cdot \|W\|^{\varphi(m)-3}.$$

Our choice of ε ensures that $\|I - A\| < 1/2$, so the matrix A is invertible and $\|A^{-1}\| \leq \frac{1}{1 - \|I - A\|} \leq 2$ (an application of the Neumann series). Therefore, $\|\bar{W}^+\| \leq \|A^{-1}\| \|W^+\| \leq 2\|W^+\|$. □

Note that the above proof shows that \bar{W} has full column rank (because W does, $\bar{W} = AW$, and A has full rank). We deduce that $\bar{W}^+ \bar{W}$ is the identity, and

$$(8) \quad \|\bar{x}\| \leq \|\bar{W}^+\| \|\bar{W} \cdot \bar{x}\| \leq \|\bar{W}^+\| (\|\bar{W} \cdot \bar{x} - \bar{t}\| + \|\bar{t}\|).$$

It follows from the two above lemmata that $\|\bar{W}^+\| \leq 2\tilde{O}^{(m)}$ (recall that $\|W\| = O(m)$). Since we have that $\|\bar{W} \cdot \bar{x} - \bar{t}\| \leq 2^{o(m)}$ and

$$\|\bar{t}\| \leq \|\bar{t}'\| = \|\bar{L} \cdot k\| \leq \|(\bar{L} - L) \cdot k\| + \|L \cdot k\| \leq 2^{p+o(m)},$$

we deduce from (8) that $\|\bar{x}\| \leq 2^{p+\bar{O}(m)}$, and therefore $\|(W-\bar{W})\cdot\bar{x}\|_\infty \leq 2^{-m^2+\bar{O}(m)}$. Applying this inequality to (7), we conclude that

$$\begin{aligned} \|h\|_\infty &\leq N(g)^{1/\varphi(m)} \cdot \exp(O(\sqrt{m \log m})) \cdot \exp(2^{-m^2+\bar{O}(m)}) \\ &\leq N(g)^{1/\varphi(m)} \cdot \exp(O(\sqrt{m \log m})). \end{aligned}$$

4. THE GEOMETRY OF THE STICKELBERGER IDEAL

In this section, we study the geometry of the Stickelberger ideal, and as an application, we provide an algorithm for the close principal multiple problem in any $\mathbb{Z}[G]$ -cycle of the relative class group Cl_K^- , Theorem 4.8.

The group ring $\mathbb{Z}[G]$ acts on the ideals of \mathcal{O}_K as follows: if \mathfrak{a} is an ideal of \mathcal{O}_K , and $\alpha = \sum_{\sigma \in G} \alpha_\sigma \sigma \in \mathbb{Z}[G]$, we write

$$\mathfrak{a}^\alpha = \prod_{\sigma \in G} (\mathfrak{a}^\sigma)^{\alpha_\sigma}.$$

If $\mathfrak{a} \sim \mathfrak{b}$, then $\mathfrak{a}^\alpha \sim \mathfrak{b}^\alpha$, thereby inducing an action of $\mathbb{Z}[G]$ on Cl_K . The norms $\|\cdot\|$ and $\|\cdot\|_1$ denote the usual ℓ_2 (Euclidean) and ℓ_1 norms over $\mathbb{R}^{\varphi(m)}$, and are defined over $\mathbb{Z}[G]$ via the natural isomorphism $\mathbb{Z}[G] \cong_{\mathbb{Z}} \mathbb{Z}^{\varphi(m)}$. The ℓ_1 -norm is of particular interest here, since $N(\mathfrak{a}^\alpha) = N(\mathfrak{a})^{\|\alpha\|_1}$ when the coefficients of α are non-negative.

Given an ideal in the form \mathfrak{a}^α , one could build an equivalent ideal of smaller norm by finding an element $\beta \in \mathbb{Z}[G]$ of smaller ℓ_1 -norm such that $\mathfrak{a}^\alpha \sim \mathfrak{a}^\beta$. One can hope to achieve this in the following way. Suppose Λ is a lattice of full rank in $\mathbb{Z}[G]$ such that for any ideal \mathfrak{h} , and any $\lambda \in \Lambda$, the ideal \mathfrak{h}^λ is principal (we say that Λ is a lattice of class relations). Given \mathfrak{a}^α and a good basis for Λ , one could find an element $\gamma \in \Lambda$ close to α . Then, $\alpha - \gamma$ is small, and $\mathfrak{a}^\alpha \sim \mathfrak{a}^{\alpha-\gamma}$. Therefore we are interested in finding a lattice of class relations with a good basis.

4.1. The Stickelberger ideal. The Galois group G is canonically isomorphic to $(\mathbb{Z}/m\mathbb{Z})^\times$ via $a \mapsto \sigma_a$, where σ_a is the automorphism sending ζ_m to ζ_m^a . 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 $\{x\} = x \pmod{\mathbb{Z}}$; equivalently, $\{x\} = x - \lfloor x \rfloor$.

Definition 4.1 (The Stickelberger ideal). *For any integer $a \in \mathbb{Z}$, let*

$$\theta(a) = \sum_{b \in (\mathbb{Z}/m\mathbb{Z})^\times} \left\{ -\frac{ab}{m} \right\} \sigma_b^{-1} \in \mathbb{Q}[G].$$

Let S' be the \mathbb{Z} -module generated by $\{\theta(a) \mid a \in \mathbb{Z}\}$ in $\mathbb{Q}[G]$. The Stickelberger ideal is defined as $S = \mathbb{Z}[G] \cap S'$. It is an ideal in $\mathbb{Z}[G]$, and we will refer to the Stickelberger lattice when S is considered as a \mathbb{Z} -module.

This is the definition from [Sin78], while some references (such as [Was12]) call *Stickelberger ideal* the smaller ideal $\mathbb{Z}[G] \cap \theta(1)\mathbb{Z}[G]$. Note that the definitions coincide when m is a power of a prime number. The Stickelberger ideal provides some class relations, thanks to the following theorem. A proof can be found in [Wei74].

Theorem 4.2 (Stickelberger's theorem). *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.*

4.2. Short generating vectors of the Stickelberger lattice. For any integer $a \in \mathbb{Z}$, let $v_a = a\theta(1) - \theta(a) \in \mathbb{Q}[G]$.

Lemma 4.3. *The set $\{v_a \mid a = 2, \dots, m\}$ generates the Stickelberger lattice.*

Proof. Let L be the lattice generated by $\{v_a \mid a = 2, \dots, m\}$ in $\mathbb{Q}[G]$. A simple calculation shows that $v_a \in \mathbb{Z}[G]$ for any integer a , so $L \subseteq S$. Let γ be an arbitrary element of S . Since $\theta(0) = 0$, and $\theta(a) = \theta(b)$ for any integers a and b such that $a \equiv b \pmod{m}$, we can write $\gamma = \sum_{a=1}^{m-1} x_a \theta(a)$ where the x_a 's are integers. Now,

$$\begin{aligned} \gamma &= \sum_{a=1}^{m-1} x_a \theta(a) = \sum_{a=1}^{m-1} x_a \left(\sum_{b \in (\mathbb{Z}/m\mathbb{Z})^\times} \left\{ -\frac{ab}{m} \right\} \right) \sigma_b^{-1} \\ &= \sum_{b \in (\mathbb{Z}/m\mathbb{Z})^\times} \left(\sum_{a=1}^{m-1} x_a \left\{ -\frac{ab}{m} \right\} \right) \sigma_b^{-1}. \end{aligned}$$

Therefore, for all b , the coefficient of σ_b^{-1} in γ is $\sum_{a=1}^{m-1} x_a \left\{ -\frac{ab}{m} \right\}$, and it is an integer since γ is in the group ring $\mathbb{Z}[G]$, so the sum $\sum_{a=1}^{m-1} x_a a$ is divisible by m . Let q be the integer such that $\sum_{a=1}^{m-1} x_a a = qm$. We obtain

$$\gamma = \sum_{a=1}^{m-1} x_a \theta(a) = \sum_{a=1}^{m-1} x_a a \theta(1) + \sum_{a=1}^{m-1} x_a (\theta(a) - a\theta(1)) = qv_m - \sum_{a=2}^{m-1} x_a v_a \in L,$$

which concludes the proof. \square

We are now ready to construct our set of short generators for S . Let $w_a = v_a - v_{a-1}$ for $a \in \{2, \dots, m\}$, and let

$$W = \{w_2, \dots, w_m\}.$$

Lemma 4.4. *The set W is a set of short generators of S . More precisely,*

- (1) W generates the Stickelberger lattice S ,
- (2) For any $a \in \{2, \dots, m\}$, $w_a = \sum_{b \in (\mathbb{Z}/m\mathbb{Z})^\times} \epsilon_{a,b} \cdot \sigma_b^{-1}$, with $\epsilon_{a,b} \in \{0, 1\}$,
- (3) For any $w \in W$, we have $\|w\| \leq \sqrt{\varphi(m)}$.

The second item essentially generalizes [Sch10, Prop. 9.4] from prime conductors to arbitrary conductors.

Proof. Point 1 is a direct consequence of Lemma 4.3 and the construction of W . Point 3 follows from Point 2, so we focus on proving Point 2. Similarly to the proof of [Was12, Lem. 6.9], we have

$$\begin{aligned} v_a &= a\theta(1) - \theta(a) = \sum_{b \in (\mathbb{Z}/m\mathbb{Z})^\times} \left(a \left\{ -\frac{b}{m} \right\} - \left\{ -\frac{ab}{m} \right\} \right) \sigma_b^{-1} \\ &= \sum_{b \in (\mathbb{Z}/m\mathbb{Z})^\times} \left[a \left\{ -\frac{b}{m} \right\} \right] \sigma_b^{-1} \end{aligned}$$

using the identity $x\{y\} - \{xy\} = [x\{y\}]$ 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 $w_a = \sum_{b \in (\mathbb{Z}/m\mathbb{Z})^\times} \epsilon_{a,b} \sigma_b^{-1}$, where

$$\epsilon_{a,b} = \left[a \left\{ -\frac{b}{m} \right\} \right] - \left[(a-1) \left\{ -\frac{b}{m} \right\} \right] \leq \left\{ -\frac{b}{m} \right\} + 1 < 2.$$

Therefore $\epsilon_{a,b} \in \{0, 1\}$ for every index a and b . \square

4.3. Class relations for the relative class group. We cannot directly use the Stickelberger ideal $S \subset \mathbb{Z}[G]$ as a lattice of class relations since it does not have full rank in $\mathbb{Z}[G]$ as a \mathbb{Z} -module (precisely, its \mathbb{Z} -rank is $\varphi(m)/2 + 1$ when $m \geq 2$). Indeed, if the lattice is not full rank, a given vector does not necessarily have a short representative modulo the lattice. To resolve this issue, we restrict our attention to the subgroup Cl_K^- .

Recall that $K^+ = \mathbb{Q}(\zeta_m + \zeta_m^{-1})$ is the maximal real subfield of K , with class group Cl_{K^+} , and Cl_K^- is the relative class group (the kernel of the relative norm map $N_{K/K^+} : \text{Cl}_K \rightarrow \text{Cl}_{K^+}$). By construction, the element $1 + \tau \in \mathbb{Z}[G]$ annihilates Cl_K^- , so the action of $\mathbb{Z}[G]$ on Cl_K^- factors through the quotient ring

$$R = \mathbb{Z}[G]/(1 + \tau).$$

The ring R also has a geometric structure. Let $\pi : \mathbb{Z}[G] \rightarrow R$ be the natural projection. Let $B \subset G$ be any set of representatives of $G/\langle \tau \rangle$. Then, the projection $\pi(B)$ forms a \mathbb{Z} -basis of R . The induced isomorphism $R \cong \mathbb{Z}^{\varphi(m)/2}$ naturally induces an ℓ_1 and ℓ_2 -norm on R , and these norms do not actually depend on the choice of B .

Lemma 4.5. *The projected Stickelberger lattice $\pi(S)$ has full rank $\varphi(m)/2$ in R .*

Proof. A generalisation due to Sinnott [Sin78] of a theorem from Iwasawa states that $(1 - \tau)S$ is of full rank in $(1 - \tau)\mathbb{Z}[G]$. We conclude by noting that the projection of $(1 - \tau)\mathbb{Z}[G]$ into R is equal to $2R$, which has full rank. \square

The set of elements $\pi(W)$ has full rank in R . One can easily deduce from Lemma 4.4 that $\|\pi(w)\| \leq 2\sqrt{\varphi(m)}$ for any $w \in W$, but we can show the following slightly stronger bound.

Lemma 4.6. *For any $w \in W$, we have $\|\pi(w)\| \leq \sqrt{\varphi(m)/2}$.*

Proof. Using the notations of Lemma 4.4, we have

$$\pi(w_a) = \sum_{b \in B} (\epsilon_{a,b} \sigma_b^{-1} + \epsilon_{a,-b} \sigma_{-b}^{-1}) = \sum_{b \in B} (\epsilon_{a,b} - \epsilon_{a,-b}) \sigma_b^{-1},$$

hence it is sufficient to show that for any $a \in \{2, \dots, m\}$ and $b \in B$, we have $\epsilon_{a,b} - \epsilon_{a,-b} \in \{-1, 0, 1\}$ since B has cardinality $\varphi(m)/2$. For $a = m$, we have $\epsilon_{a,b} = \epsilon_{a,-b} = 1$, so $\pi(w_m) = 0$. Suppose $a \neq m$. Then, since $ab/m \notin \mathbb{Z}$,

$$\left[a \left\{ -\frac{b}{m} \right\} \right] = \left[a \left(1 - \left\{ \frac{b}{m} \right\} \right) \right] = a + \left[-a \left\{ \frac{b}{m} \right\} \right] = a - \left[a \left\{ \frac{b}{m} \right\} \right],$$

Then, $\epsilon_{a,b} - \epsilon_{a,-b} = 1 - 2\epsilon_{a,-b} \in \{-1, 1\}$. \square

4.4. The close principal multiple problem in a $\mathbb{Z}[G]$ -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 $\alpha \in \mathbb{Z}[G]$ we will find a short $\beta \in \mathbb{Z}[G]$ such that $C^\beta = C^\alpha$, for any class $C \in \text{Cl}_K^-$. We rely on the following close vector algorithm.

Theorem 4.7. *There is an algorithm REDUCE (Algorithm 3), that given $\alpha \in \mathbb{Z}[G]$, finds an element $\beta \in \mathbb{Z}[G]$ such that $\|\beta\|_1 \leq \frac{1}{4} \cdot \varphi(m)^{3/2}$, and $C^\alpha = C^\beta$ for any $C \in \text{Cl}_K^-$, and runs in polynomial time in m and $\log(\|\alpha\|)$.*

Algorithm 3 REDUCE(α): finds a reduction of α .

Require: An element $\alpha \in \mathbb{Z}[G]$.

Ensure: An element $\beta \in \mathbb{Z}[G]$ such that $\|\beta\|_1 \leq \frac{1}{4} \cdot \varphi(m)^{3/2}$, and $C^\alpha = C^\beta$ for any $C \in \text{Cl}_K^-$.

- 1: Let W be the generating set of S as in Lemma 4.4;
 - 2: $v \leftarrow \text{CV}(\pi(W), \pi(\alpha))$; {close vector algorithm of Corollary 2.2}
 - 3: $\gamma \leftarrow \pi(W) \cdot v$;
 - 4: Write $\pi(\alpha) - \gamma = \sum_{\sigma \in B} a_\sigma \pi(\sigma)$ using the basis $\pi(B)$ of R ;
 - 5: $\beta = \sum_{\sigma \in B} a_\sigma \sigma$;
 - 6: **return** β .
-

Proof. Recall that $\pi : \mathbb{Z}[G] \rightarrow R$ is the canonical projection, W is the generating set of S as in Lemma 4.4, and $B \subset G$ is any set of representatives of $G/\langle \tau \rangle$. From Lemma 4.5, $\pi(W)$ has full rank in R . So the close vector algorithm from Corollary 2.2 finds an element v such that $\gamma = \pi(W) \cdot v \in \pi(S)$ satisfies

$$\|\pi(\alpha) - \gamma\|_1 \leq \frac{\varphi(m)}{2} \cdot \max_{w \in W} \|\pi(w)\| \leq \frac{1}{4} \cdot \varphi(m)^{3/2}.$$

The bound on the $\max_{w \in W}$ follows from Lemma 4.6. Then, the element β returned by Algorithm 3 satisfies

$$\|\beta\|_1 = \|\pi(\alpha) - \gamma\|_1 \leq \frac{1}{4} \cdot \varphi(m)^{3/2}.$$

Furthermore, for any $C \in \text{Cl}_K^-$, Stickelberger's theorem implies that $C^\gamma = [\mathcal{O}_K]$, and therefore $C^\alpha = C^\beta$. \square

Theorem 4.8 (Close principal multiple algorithm for $\mathbb{Z}[G]$ -cycles of Cl_K^-). *Let \mathfrak{p} be an ideal such that $[\mathfrak{p}] \in \text{Cl}_K^-$. There is an algorithm CLOSEPRINCIPALMULTIPLE⁻ (Algorithm 4) that given an element $\alpha \in \mathbb{Z}[G]$, finds an integral ideal \mathfrak{b} such that $\mathfrak{p}^\alpha \mathfrak{b}$ is principal and*

$$N(\mathfrak{b}) = N(\mathfrak{p})^{O(\varphi(m)^{3/2})},$$

and runs in polynomial time in m , $\log(N(\mathfrak{p}))$ and $\log(\|\alpha\|)$.

Remark 4.9. If one is given the ideal $\mathfrak{a} = \mathfrak{p}^\alpha \in \mathfrak{p}^{\mathbb{Z}[G]}$ instead of the element α , one could try to recover α by solving a discrete logarithm problem in the relative class group. This is doable in quantum polynomial time (as in Section 5.2), but we choose to have α given in Theorem 4.8 to obtain a classical algorithm.

Proof. Consider β , γ and \mathfrak{b} as in Algorithm 4. From Theorem 4.7, we have $\|\beta\|_1 \leq 0.5 \cdot \varphi(m)^{3/2}$, and $\mathfrak{p}^\alpha \sim \mathfrak{p}^\beta$. Since $[\mathfrak{p}] \in \text{Cl}_K^-$, we have $\mathfrak{p}^{-1} \sim \mathfrak{p}^\tau$, so

$$\mathfrak{p}^\gamma \sim \mathfrak{p}^{\sum_{\sigma \in G} (\tau b_\sigma^+ + b_\sigma^-) \sigma} \sim \mathfrak{p}^{\sum_{\sigma \in G} (-b_\sigma^+ + b_\sigma^-) \sigma} \sim \mathfrak{p}^{-\alpha},$$

hence $\mathfrak{p}^\alpha \mathfrak{b}$ is principal. Since γ has only positive coefficients, the ideal \mathfrak{b} is integral. Finally, $N(\mathfrak{b}) = N(\mathfrak{p})^{\|\gamma\|_1} = N(\mathfrak{p})^{O(\varphi(m)^{3/2})}$. \square

Algorithm 4 CLOSEPRINCIPALMULTIPLE⁻(\mathfrak{p}, α): solves CPM for the ideal \mathfrak{p}^α .

Require: An ideal \mathfrak{p} such that $[\mathfrak{p}] \in \text{Cl}_K^-$, and an element $\alpha \in \mathbb{Z}[G]$.

Ensure: An integral ideal \mathfrak{b} such that $\mathfrak{p}^\alpha \mathfrak{b}$ is principal and $N(\mathfrak{b}) = N(\mathfrak{p})^{O(\varphi(m)^{3/2})}$.

- 1: $\beta \leftarrow \text{REDUCE}(\alpha)$; {Algorithm 3}
 - 2: Write $\beta = \sum_{\sigma \in G} b_\sigma \sigma$;
 - 3: **for all** $\sigma \in G$ **do**
 - 4: $(b_\sigma^+, b_\sigma^-) \leftarrow \begin{cases} (b_\sigma, 0) & \text{if } b_\sigma \geq 0, \\ (0, -b_\sigma) & \text{otherwise;} \end{cases}$
 - 5: **end for**
 - 6: $\gamma \leftarrow \sum_{\sigma \in G} (\tau b_\sigma^+ + b_\sigma^-) \sigma$;
 - 7: **return** $\mathfrak{b} = \mathfrak{p}^\gamma$.
-

5. FINDING SHORT VECTORS IN CYCLOTOMIC IDEALS

Let \mathfrak{a} be an arbitrary ideal in the cyclotomic ring \mathcal{O}_K of conductor m . In this section, we prove the following theorem.

Theorem 5.1 (Approx-SVP for cyclotomic ideals). *Under GRH and Assumption 1, there is a quantum algorithm IDEALSVP (Algorithm 7) that, when given an ideal \mathfrak{a} in the cyclotomic ring of conductor m , finds an element in \mathfrak{a} of Euclidean norm*

$$\exp(\tilde{O}(\sqrt{m})) \cdot N(\mathfrak{a})^{1/\varphi(m)},$$

and runs in polynomial time in m , h_K^+ and $\log N(\mathfrak{a})$, where h_K^+ is the class number of the maximal totally real subfield K^+ . This element approximates SVP in \mathfrak{a} with an approximation factor $\exp(\tilde{O}(\sqrt{m}))$.

The strategy is the following. Suppose that we have a set $\{\mathfrak{p}_1, \dots, \mathfrak{p}_d\}$ of ideals of norm $\text{poly}(m)$ that generate the relative class group Cl_K^- as a $\mathbb{Z}[G]$ -module.

- (1) First, we find an (integral) ideal \mathfrak{b} of small norm such that the class of $\mathfrak{a}\mathfrak{b}$ is in the relative class group Cl_K^- . This is done via a random walk in the class group in Section 5.1.
- (2) Second, we find $\alpha_1, \dots, \alpha_d \in \mathbb{Z}[G]$ such that $\mathfrak{a}\mathfrak{b} \sim \mathfrak{p}_1^{\alpha_1} \dots \mathfrak{p}_d^{\alpha_d}$, and apply the results of Section 4 to find ideals $\mathfrak{b}_i \sim \mathfrak{p}_i^{-\alpha_i}$ such that $N(\mathfrak{b}_i) = \exp(\tilde{O}(m^{3/2}))$. This is done in Section 5.2.
- (3) Finally, the ideal $\mathfrak{c} = \mathfrak{a}\mathfrak{b}\mathfrak{b}_1 \dots \mathfrak{b}_d$ is principal. Applying the results of Section 3 allows to find an element $g \in \mathfrak{c} \subset \mathfrak{a}$ of norm $\exp(\tilde{O}(d\sqrt{m}))$. This is done in Section 5.3.

The above steps work under GRH, given a generating set $\{\mathfrak{p}_1, \dots, \mathfrak{p}_d\}$ as above. The construction of the generating set is where further number theoretic heuristics are required. Assumption 1 is a statement on the Galois-module structure of Cl_K^- which allows to take $d = \text{polylog}(m)$, and obtain the targeted approximation factor $\exp(\tilde{O}(\sqrt{m}))$. The above three steps together with Assumption 1 lead to the claimed quantum polynomial time algorithm.

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

This assumption is arguably new, and can be read as a strengthened version of a theorem of Bach [Bac90, Theorem 4] and its generalizations to subgroups of the class group [JW15, Corollary 6.5]. This assumption, and its justification, is the object of Section 6.

Note that for the algorithm of Theorem 5.1 to really be efficient, one would also require h_K^+ to be polynomially bounded in m . This Assumption 2 is discussed in Section 5.1.2. Unlike the previous one, this assumption is a well-known question in algebraic number theory, and is related to important conjectures.

5.1. Random walk to the relative class group. As previously, let K^+ denote the maximal real subfield of K , and Cl_{K^+} the class group of K^+ .

The core of the method to find a close principal multiple of an ideal \mathfrak{a} works within the relative class group $\text{Cl}_{\bar{K}} \subset \text{Cl}_K$. Therefore, as a first step, we need to “send” the ideal $\mathfrak{a} \in \text{Cl}_K$ into this subgroup. More precisely, we want an integral ideal \mathfrak{b} of small norm such that $\mathfrak{a}\mathfrak{b} \in \text{Cl}_{\bar{K}}$; the rest of the method then works with $\mathfrak{a}\mathfrak{b}$. Let $h_K = |\text{Cl}_K|$ be the class number of K , and $h_{\bar{K}} = |\text{Cl}_{\bar{K}}|$ its relative class number. The difficulty of this step is directly related to the index of $\text{Cl}_{\bar{K}}$ inside Cl_K , which is the real class number $h_K^+ = |\text{Cl}_{K^+}|$ of K^+ (indeed, the relative norm map $N_{K/K^+} : \text{Cl}_K \rightarrow \text{Cl}_{K^+}$ induces the isomorphism $\text{Cl}_{K^+} \cong \text{Cl}_K/\text{Cl}_{\bar{K}}$), and is expected to be very small.

5.1.1. The random walk algorithm. 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 \mathcal{G}_x denote the induced Cayley (multi)graph $\text{Cay}(\text{Cl}_K, S_x)$. From [JW15, Corollary 6.5] (under GRH), for any $\varepsilon > 0$ there is a constant C and a bound

$$B = O((\varphi(m) \log \Delta_K)^{2+\varepsilon}) = O((\varphi(m)^2 \log \varphi(m))^{2+\varepsilon})$$

such that any random walk in \mathcal{G}_B of length at least $C \log(h_K)/\log \log(\Delta_K)$, for any starting point, lands in the subgroup $\text{Cl}_{\bar{K}}$ 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 $\mathfrak{p}_1, \dots, \mathfrak{p}_\ell$ 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

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

If $[\mathfrak{a}]$ is the starting point of the random walk in the graph, the endpoint $[\mathfrak{a}\mathfrak{b}]$ falls in $\text{Cl}_{\bar{K}}$ with probability at least $1/(2h_K^+)$, and therefore an ideal \mathfrak{b} such that $[\mathfrak{a}\mathfrak{b}] \in \text{Cl}_{\bar{K}}$ can be found in probabilistic polynomial time in h_K^+ . Note that the quantum algorithm of Biasse and Song [BS16] for PIP allows to test the membership $[\mathfrak{a}\mathfrak{b}] \in \text{Cl}_{\bar{K}}$, 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 5, and the efficiency is stated below. Under GRH and Assumption 2, this procedure runs in polynomial time.

Lemma 5.2. *Assuming the GRH, the quantum algorithm WALKTOCl^- (Algorithm 5) runs in expected time $O(h_K^+) \cdot \text{poly}(m, \log N(\mathfrak{a}))$ and is correct.*

Algorithm 5 WALKTOCI⁻(**a**): random walk to CI_K⁻.

Require: An ideal **a** in \mathcal{O}_K .

Ensure: An integral ideal **b** such that $[\mathbf{ab}] \in \text{CI}_K^-$ and $N(\mathbf{b}) \leq \exp(\tilde{O}(m))$.

1: $\ell = \tilde{O}(m)$; $B = \text{poly}(m)$;

2: **repeat**

3: **for** $i = 1$ **to** ℓ **do**

4: Choose \mathfrak{p}_i uniformly among the prime ideal of norm less than B ;

5: **end for**

6: $\mathfrak{b} \leftarrow \prod_{i=1}^{\ell} \mathfrak{p}_i$;

7: **until** $N_{K/K^+}(\mathbf{ab})$ is principal; {using the PIP algorithm of [BS16]}

8: **return** **b**.

5.1.2. *The real class number.* The time complexity of Algorithm 5 has a linear factor h_K^+ the class number of the real subfield K^+ . Assumption 2 below claims that this factor is not a problem. 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 2. For any integer m , it holds that $h^+(m) \leq \text{poly}(m)$.

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 5.3 (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 5.4 (Weber's class number problem). *For any e , $h^+(2^e) = 1$.*

A direct consequence of Conjecture 5.3 is that for fixed ℓ and increasing e , the quantity $h^+(\ell^e)$ is $O(1)$, implying that Assumption 2 holds over the class of cyclotomic fields of conductor a power of ℓ .

But even for increasing primes ℓ , the quantity $h^+(\ell)$ 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 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.

5.2. Close principal multiple algorithm. Combining the random walk from the previous section, the close principal multiple algorithm in CI_K^- from Section 4.4, and the quantum algorithms for class group computations discussed in Section 2.3, one can construct an algorithm for the general close principal multiple problem in \mathcal{O}_K .

Theorem 5.5 (Close principal multiple algorithm). *Under GRH and Assumption 1, there is a quantum algorithm CLOSEPRINCIPALMULTIPLE (Algorithm 6) that given an ideal **a** in the cyclotomic ring of conductor m , finds an integral ideal*

\mathfrak{b} such that $\mathfrak{a}\mathfrak{b}$ is principal and

$$N(\mathfrak{b}) = \exp\left(\tilde{O}\left(m^{3/2}\right)\right),$$

and runs in polynomial time in m , h_K^+ and $\log(N(\mathfrak{a}))$.

Algorithm 6 CLOSEPRINCIPALMULTIPLE(\mathfrak{a}): solves CPM for the ideal \mathfrak{a} .

Require: An ideal \mathfrak{a} in \mathcal{O}_K .

Ensure: An integral ideal \mathfrak{b} such that $\mathfrak{a}\mathfrak{b}$ is principal and $N(\mathfrak{b}) = \exp(\tilde{O}(m^{3/2}))$.

```

1:  $d = \text{polylog}(m)$ ;  $B = \text{poly}(m)$ ;
2:  $\mathfrak{b}' \leftarrow \text{WALKTOCL}^-(\mathfrak{a})$ ; {Algorithm 5}
3:  $\mathfrak{M} \leftarrow \{\mathfrak{p} \mid N(\mathfrak{p}) \leq B, [\mathfrak{p}] \in \text{Cl}_K^-\}$ ;
4: repeat
5:   Choose  $\mathfrak{p}_1, \dots, \mathfrak{p}_d$  uniformly at random in  $\mathfrak{M}$ ;
6:    $\mathfrak{B} \leftarrow \{\mathfrak{p}_i^\sigma \mid \sigma \in G, i = 1 \dots d\}$ ;
7:    $(y_q)_{q \in \mathfrak{B}} \leftarrow \text{CIDL}_{\mathfrak{B}}(\mathfrak{a}\mathfrak{b}')$ , if it exists; {Proposition 2.7}
8: until the discrete logarithm  $(y_q)_{q \in \mathfrak{B}}$  has been found;
9: for  $i = 1$  to  $d$  do
10:   $\alpha_i \leftarrow \sum_{\sigma \in G} y_{\mathfrak{p}_i^\sigma} \sigma \in \mathbb{Z}[G]$ ;
11:   $\mathfrak{b}_i \leftarrow \text{CLOSEPRINCIPALMULTIPLE}^-(\mathfrak{p}_i, \alpha_i)$ ; {Algorithm 4}
12: end for
13:  $\mathfrak{b} \leftarrow \mathfrak{b}' \prod_{i=1}^d \mathfrak{b}_i$ ;
14: return  $\mathfrak{b}$ .
```

Proof. The running time of Algorithm 6 follows from Lemma 5.2, Proposition 2.7 and Theorem 4.8. Note that this algorithm might fail when the chosen \mathfrak{B} does not generate Cl_K^- , but following Assumption 1, it will succeed after an constant expected number of trials. Let us prove that it is correct. The algorithm WALKTOCL^- outputs an integral ideal \mathfrak{b}' such that $[\mathfrak{a}\mathfrak{b}'] \in \text{Cl}_K^-$ and $N(\mathfrak{b}) \leq \exp(\tilde{O}(m))$. When \mathfrak{B} generates Cl_K^- , the algorithm $\text{CIDL}_{\mathfrak{B}}$ finds a sequence of elements $\alpha_1, \dots, \alpha_d \in \mathbb{Z}[G]$ such that $\mathfrak{a}\mathfrak{b}' \sim \prod_{i=1}^d \mathfrak{p}_i^{\alpha_i}$. Now, applying the algorithm from Theorem 4.8 to each $\mathfrak{p}_i^{\alpha_i}$, we obtain ideals $\mathfrak{b}_1, \dots, \mathfrak{b}_d$ such that $\mathfrak{p}_i^{\alpha_i} \mathfrak{b}_i$ is principal and $N(\mathfrak{b}_i) = \exp(\tilde{O}(m^{3/2}))$ for any $i = 1, \dots, d$. It follows that the output $\mathfrak{b} = \mathfrak{b}' \prod_{i=1}^d \mathfrak{b}_i$ has the desired properties. \square

5.3. Proof of Theorem 5.1. The algorithm is summarized in Algorithm 7. The running time and correctness follow from Theorem 3.7 and Theorem 5.5. \square

Algorithm 7 IDEALSVP(\mathfrak{a}): finding mildly short vectors in the ideal \mathfrak{a} .

Require: An ideal \mathfrak{a} in \mathcal{O}_K .

Ensure: An element $v \in \mathfrak{a}$ of norm $\|v\| \leq \exp(\tilde{O}(\sqrt{m})) \cdot N(\mathfrak{a})^{1/\varphi(m)}$.

```

1:  $\mathfrak{b} \leftarrow \text{CLOSEPRINCIPALMULTIPLE}(\mathfrak{a})$ ; {Algorithm 6}
2:  $v \leftarrow \text{PRINCIPALIDEALSVP}(\mathfrak{b})$ ; {Algorithm 2}
3: return  $v$ .
```

6. CONSTRUCTING SMALL FACTOR BASES FOR THE RELATIVE CLASS GROUP

To argue for Assumption 1, we prove (in Proposition 6.1) that if Cl_K^- can be generated by r ideal classes, then $r \cdot \text{polylog}(m)$ uniformly random classes in Cl_K^- will generate it.

Proposition 6.1. *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 $\alpha \geq 1$ be a parameter, and s be any integer such that*

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

(note that $\log_2 \log_2(h_K^-) \sim \log_2(\varphi(m))$)⁴. Let x_1, \dots, x_s be s independent uniform elements of Cl_K^- . The probability that $\{x_1, \dots, x_s\}$ generates Cl_K^- as a $\mathbb{Z}[G]$ -module is at least $\exp(-\frac{3}{2^\alpha}) = 1 - O(2^{-\alpha})$.

In other words, a set of $\Theta(r \log(\varphi(m)))$ random ideal classes in Cl_K^- will generate this $\mathbb{Z}[G]$ -module with very good probability. This proposition is proven at the end of this section.

To justify Assumption 1, we first argue that r is admittedly as small as $\text{polylog}(m)$. 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 $\text{polylog}(m)$, at least for a dense class of prime power conductors.

Once it is admitted that $r \leq \text{polylog}(m)$, Assumption 1 simply assumes that Proposition 6.1 remains true when imposing that the random classes $g_1 \dots g_s$ are chosen as the classes of random ideals of small norm, i.e. $g_i = [\mathfrak{p}_i]$ where $N(\mathfrak{p}_i) \leq \text{poly}(m)$. This restriction on the norms seems reasonable considering that it has been proven that prime ideals of norm $\text{poly}(m)$ that fall in Cl_K^- are sufficient to generate Cl_K^- , assuming GRH and Assumption 2 (see [JW15, Corollary 6.5]). Explicitly, it has been proved in [Wes18b] that prime ideals of norm at most $(2.71h_K^+ \log \Delta_K + 4.13)^2$ are sufficient to generate Cl_K^- .

We now show a series of results leading to the proof of Proposition 6.1.

Lemma 6.2. *Let R be a finite commutative local ring of cardinality ℓ^n , for some prime number ℓ . A set of s independent uniformly random elements in R generates R as an R -module with probability at least $1 - \ell^{-s}$.*

Proof. An element generates R if and only if it is invertible, meaning that it is not in the maximal ideal of R . This ideal is a fraction at most ℓ^{-1} of R , so an element does not generate R with probability at most ℓ^{-1} . Among s independent elements, the probability that none of them is a generator is at most ℓ^{-s} . \square

Lemma 6.3. *Let R be a finite commutative local ring of cardinality ℓ^n , for some prime number ℓ . Let M be a cyclic R -module. A set of s independent uniformly random elements in M generates M with probability at least $1 - \ell^{-s}$.*

⁴Here, the notation $f \sim g$ denotes the standard asymptotic equivalence $\lim_{x \rightarrow \infty} f(x)/g(x) = 1$.

Proof. Let g be a generator of M , and consider the homomorphism $\varphi : R \rightarrow M : \alpha \mapsto \alpha g$. Let x_1, \dots, x_s be s independent uniformly random element in M . For each i , let α_i be a uniformly random element of the coset $\varphi^{-1}(x_i)$. The elements α_i are independent and uniformly distributed in R , so from Lemma 6.2, they generate R with probability at least $1 - \ell^{-s}$. If the α_i 's generate R , then the x_i 's generate M , and we conclude. \square

Lemma 6.4. *Let R be a finite commutative local ring of cardinality ℓ^n , for some prime number ℓ . Let M be an R -module, and let r be the smallest number of R -generators of M . A set of s independent uniformly random elements in M generates M with probability at least $(1 - \ell^{-\lfloor s/r \rfloor})^r$.*

Proof. Proceed by induction on r . The case $r = 1$ is Lemma 6.3. Suppose that for any R -module M' generated by $r - 1$ elements, and any positive s' , a set of s' random elements in M' generates M' with probability at least

$$\left(1 - \ell^{-\lfloor s'/(r-1) \rfloor}\right)^{r-1}.$$

Choose s independent uniformly random elements x_1, \dots, x_s in M , and let $t = \lfloor s/r \rfloor$. Let g_1, \dots, g_r be a generating set for M . The quotient $M/(Rg_r)$ is generated by $r - 1$ elements, so the first $s - t$ random elements generate it with probability at least

$$\left(1 - \ell^{-\lfloor (s-t)/(r-1) \rfloor}\right)^{r-1} \geq \left(1 - \ell^{-\lfloor s/r \rfloor}\right)^{r-1}.$$

Now assume that these $s - t$ elements indeed generate $M/(Rg_r)$. It remains to show that adding the remaining t random elements allows to generate the full module M with probability at least $1 - \ell^{-\lfloor s/r \rfloor}$. Let $N \subset M$ be the submodule of M generated by the first $s - t$ random elements. Observe that the module M/N is generated by g_r . Indeed, let m be an arbitrary element of M . Since $M/(Rg_r)$ is generated by N , there is an $n \in N$ such that $m + Rg_r = n + Rg_r$. This implies that there is an element $\alpha g_r \in Rg_r$ such that $m + N = \alpha g_r + N$, proving that M/N is generated by g_r . From Lemma 6.3, M/N is generated by the last t random elements with probability at least $1 - \ell^{-\lfloor s/r \rfloor}$. So M is generated by x_1, \dots, x_s with probability at least $(1 - \ell^{-\lfloor s/r \rfloor})^r$. \square

Theorem 6.5. *Let R be a finite commutative ring, M be a finite R -module of cardinality m , and r be the minimal number of R -generators of M . A set of s independent uniformly random elements in M generates M with probability at least $(1 - 2^{-\lfloor s/r \rfloor})^{\log_2 m}$.*

Proof. The ring R decomposes as an internal direct sum $\bigoplus_{i=1}^k R_i$ of finite local subrings R_i . For each i , define $e_i \in R$ the idempotent which projects to the unity of R_i and to zero in all other components of the decomposition (then, $R_i = e_i R$). In particular, we have that $M = \bigoplus_i e_i M$, and $e_i M$ may be viewed as an R_i -module.

Let x_1, \dots, x_s be s independent uniformly random elements in M . They generate M as an R -module if and only if for any i , the projections $e_i x_1, \dots, e_i x_s$ generate M_i as an R_i -module. Let p_i be the probability that $e_i x_1, \dots, e_i x_s$ generate M_i , and let r_i be the minimal number of generators of R_i . From Lemma 6.4, p_i is at least $(1 - 2^{-\lfloor s/r_i \rfloor})^{r_i}$. We have the two bounds $r_i \leq r$ and $r_i \leq \log_2 |M_i|$, and we deduce

$$p_i \geq \left(1 - 2^{-\lfloor s/r \rfloor}\right)^{\log_2 |M_i|}.$$

Therefore x_1, \dots, x_s generate M with probability at least

$$\prod_{i=1}^k p_i = \left(1 - 2^{-\lfloor s/r \rfloor}\right)^{\sum_i \log_2 |M_i|} = \left(1 - 2^{-\lfloor s/r \rfloor}\right)^{\log_2 m},$$

concluding the proof. \square

Proof of Proposition 6.1. Note that a set of elements in Cl_K^- generate it as a $\mathbb{Z}[G]$ -module if and only if they generate it as a $(\mathbb{Z}/h_K^- \mathbb{Z})[G]$ -module. We deduce from Theorem 6.5 that x_1, \dots, x_s generate Cl_K^- with probability 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 have $2^{-\lfloor s/r \rfloor} \leq 2^{-\lfloor \alpha \rfloor} \leq 1/2$, so

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

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

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

proving the proposition. \square

ACKNOWLEDGEMENTS

The authors would like to thank René Schoof for helpful and interesting discussions. We are grateful to Paul Kirchner for pointing out a mistake in an earlier version of this paper. The second author was supported by a Veni Innovative Research Grant from NWO under project number 639.021.645, and by the European Union Horizon 2020 Research and Innovation Program Grant 780701. The third author was partly supported by the Swiss National Science Foundation under grant number 200021-156420. The first and last authors were supported by the ERC Advanced Investigator Grant 740972 (ALGSTRONGCRYPTO).

REFERENCES

- [Ajt99] Miklós Ajtai. Generating hard instances of the short basis problem. In *ICALP*, pages 1–9, 1999.
- [Bab86] László Babai. On Lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986. Preliminary version in STACS 1985.
- [Bac90] E. Bach. Explicit bounds for primality testing and related problems. *Mathematics of Computation*, 55(191):355–380, 1990.
- [BBdV⁺17] Jens Bauch, Daniel J Bernstein, Henry de Valence, Tanja Lange, and Christine Van Vredendaal. Short generators without quantum computers: the case of multi-quadratics. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 27–59. Springer, 2017.
- [BEF⁺17] Jean-François Biasse, Thomas Espitau, Pierre-Alain Fouque, Alexandre Gélín, and Paul Kirchner. Computing generator in cyclotomic integer rings, a subfield algorithm for the principal ideal problem in $L(1/2)$ and application to cryptanalysis of a FHE scheme. To appear at Eurocrypt 2017, 2017.
- [BF14] Jean-François Biasse and Claus Fieker. Subexponential class group and unit group computation in large degree number fields. *LMS J. Comput. Math.*, 17(suppl. A):385–403, 2014.

- [Bia18] Jean-François Biasse. Approximate short vectors in ideal lattices of $\mathbf{Q}(\zeta_{p^e})$ with pre-computation of the class group. In Carlisle Adams and Jan Camenisch, editors, *Selected Areas in Cryptography – SAC 2017*, volume 10719 of *Lecture Notes in Computer Science*, pages 374–393, 2018.
- [BPR04] Joe Buhler, Carl Pomerance, and Leanne 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] Jean-François Biasse and Fang 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.
- [CDPR16] Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. *Recovering Short Generators of Principal Ideals in Cyclotomic Rings*, pages 559–585. Springer Berlin Heidelberg, Berlin, Heidelberg, 2016.
- [CDW17] Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Short stickelberger class relations and application to ideal-svp. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017*, pages 324–348, Cham, 2017. Springer International Publishing.
- [CGS14] Peter Campbell, Michael Groves, and Dan Shepherd. Soliloquy: A cautionary tale. ETSI 2nd Quantum-Safe Crypto Workshop, 2014. Available at http://docbox.etsi.org/Workshop/2014/201410_CRYPT0/S07_Systems_and_Attacks/S07_Groves_Annex.pdf.
- [dBDF20] Koen de Boer, Léo Ducas, and Serge Fehr. On the quantum complexity of the continuous hidden subgroup problem. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 341–370. Springer, 2020.
- [DPW19] Léo Ducas, Maxime Plançon, and Benjamin Wesolowski. On the shortness of vectors to be found by the ideal-svp quantum algorithm. 11692:322–351, 2019.
- [EH10] Kirsten Eisenträger and Sean 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, Philadelphia, PA, USA, 2010. Society for Industrial and Applied Mathematics.
- [EHKS14] Kirsten Eisenträger, Sean Hallgren, Alexei Kitaev, and Fang Song. A quantum algorithm for computing the unit group of an arbitrary degree number field. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 293–302. ACM, 2014.
- [Fri89] Eduardo Friedman. Analytic formulas for the regulator of a number field. *Inventiones mathematicae*, 98(3):599–622, Oct 1989.
- [GGH13] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In *EUROCRYPT*, pages 1–17, 2013.
- [HWB17] Patrick Holzer, Thomas Wunderer, and Johannes A Buchmann. Recovering short generators of principal fractional ideals in cyclotomic fields of conductor $p^\alpha q\beta$. In *International Conference on Cryptology in India*, pages 346–368. Springer, 2017.
- [JMV09] David Jao, Stephen D. Miller, and Ramarathnam Venkatesan. Expander graphs based on GRH with an application to elliptic curve cryptography. *Journal of Number Theory*, 129(6):1491 – 1504, 2009.
- [JW15] Dimitar Jetchev and Benjamin Wesolowski. On graphs of isogenies of principally polarizable abelian surfaces and the discrete logarithm problem. *CoRR*, abs/1506.00522, 2015.
- [Kuč92] Radan Kučera. On bases of the Stickelberger ideal and of the group of circular units of a cyclotomic field. *Journal of Number Theory*, 40(3):284–316, 1992.
- [LLL82] Arjen K. Lenstra, Hendrik W. Lenstra, Jr., and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, December 1982.
- [LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP (2)*, pages 144–155, 2006.
- [LPR13] Vadim Lyubashevsky, Chris Peikert, and Oded 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.

- [LPSW19] Changmin Lee, Alice Pellet-Mary, Damien Stehlé, and Alexandre Wallet. An LLL algorithm for module lattices. Cryptology ePrint Archive, Report 2019/1035, 2019. <https://eprint.iacr.org/2019/1035>.
- [LSS14] Adeline Langlois, Damien Stehlé, and Ron Steinfeld. GGHLite: More efficient multilinear maps from ideal lattices. In *Advances in Cryptology–EUROCRYPT 2014*, pages 239–256. Springer, 2014.
- [Mic07] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007. Preliminary version in FOCS 2002.
- [Mil15] John C. Miller. Real cyclotomic fields of prime conductor and their class numbers. *Math. Comp.*, 84(295):2459–2469, 2015.
- [PHS19] Alice Pellet-Mary, Guillaume Hanrot, and Damien Stehlé. Approx-svp in ideal lattices with pre-processing. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 685–716. Springer, 2019.
- [PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, pages 145–166, 2006.
- [PRSD17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-lwe for any ring and modulus. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 461–473, 2017.
- [RBV04] Ghaya Rekaya, Jean-Claude Belfiore, and Emanuele Viterbo. A very efficient lattice reduction tool on fast fading channels. In *Proceedings of ISITA*, volume 2004, 2004.
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009. Preliminary version in STOC 2005.
- [Sch87] Claus-Peter Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.
- [Sch98] René 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] René Schoof. Class numbers of real cyclotomic fields of prime conductor. *Mathematics of computation*, 72(242):913–937, 2003.
- [Sch10] René Schoof. *Catalan’s conjecture*. Springer Science & Business Media, 2010.
- [SE94] Claus-Peter Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathematical Programming*, 66:181–199, 1994.
- [Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, October 1997.
- [Sin78] Warren Sinnott. On the Stickelberger ideal and the circular units of a cyclotomic field. *Annals of Mathematics*, 108(1):107–134, 1978.
- [SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, pages 617–635, 2009.
- [SV10] Nigel P. Smart and Frederik Vercauteren. Fully homomorphic encryption with relatively small key and ciphertext sizes. In *Public Key Cryptography*, pages 420–443, 2010.
- [Was12] Lawrence C. Washington. *Introduction to cyclotomic fields*, volume 83. Springer Science & Business Media, 2012.
- [Wei74] André Weil. *Sommes de Jacobi et caractères de Hecke*. Nachrichten der Akademie der Wissenschaften in Göttingen, 2, Mathematisch-Physikalische Klasse. Vandenhoeck & Ruprecht, 1974.
- [Wes18a] Benjamin Wesolowski. *Arithmetic and geometric structures in cryptography*. PhD thesis, EPFL, 2018.
- [Wes18b] Benjamin Wesolowski. Generating subgroups of ray class groups with small prime ideals. In *Thirteenth Algorithmic Number Theory Symposium – ANTS-XIII*, 2018. proceedings to appear in the Open Book Series, Mathematical Sciences Publishers.