SHORT STICKELBERGER CLASS RELATIONS AND APPLICATION TO IDEAL-SVP

RONALD CRAMER[1,2], LÉO DUCAS[1] AND BENJAMIN WESOLOWSKI[3]

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

Abstract. The hardness of finding short vectors in ideals of cyclotomic number fields (Ideal-SVP) serves as the worst-case hypothesis underlying the security of numerous cryptographic schemes, including key-exchange, public-key encryption and fully-homomorphic encryption. A series of recent works has shown that, for large approximation factors, Principal Ideal-SVP is not as hard as finding short vectors in general lattices. Namely, there exists a quantum polynomial time algorithm for an approximation factor of $\exp(\tilde{O}(\sqrt{n}))$, even in the worst-case. Some schemes were broken, but more generally this exposed an unexpected hardness gap between general lattices and some structured ones, and called into question the exact security of various assumption over structured lattices.

In this work, we generalize the previous result to general ideals. We show an efficient way of finding a close enough principal multiple of any ideal by exploiting the classical theorem that, in our setting, the class-group is annihilated by the (Galois-module action of) the so-called Stickelberger ideal. Under some plausible number-theoretical hypothesis, we conclude that worst-case Ideal-SVP in this same set-up — choice of ring, and approximation factor $\exp(\tilde{O}(\sqrt{n}))$ — is also solvable in quantum polynomial time.

Although it is not yet clear whether the security of further cryptosystems is directly affected, we contribute novel ideas to the cryptanalysis of schemes based on structured lattices. Moreover, our result shows a deepening of the gap between general lattices and structured one.

1. Introduction

The problem of finding the shortest vector of a Euclidean lattice (the shortest vector problem, or SVP) is a central hard problem in complexity theory. Approximated versions of this problem (approx-SVP) have become the theoretical foundation for many cryptographic constructions thanks to the average-case to worst-case reductions of Ajtai [Ajt99] — a classical reduction from approx-SVP to the Short Integer Solution (SIS) problem — and Regev [Reg09] — a quantum reduction from approx-SVP to Learning with Errors (LWE).

For efficiency reasons, it is tempting to rely on structured lattices, in particular lattices arising as ideals or modules over certain rings, the earliest example being the NTRUencrypt[1] proposal from Hoffstein et al. [HPS98]. Later on, variations on

Date: September 8, 2016.

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

these foundations were also considered, namely the worst-case hardness of approx-SVP in lattices that are ideals of a certain ring (Ideal-SVP). Under the worst-case Ideal-SVP assumption were proven the average-case hardness of Ring-SIS and Ring-LWE, two very versatile problems to build efficient cryptographic schemes upon [Mic07, SSTX09, LPR13]. The typical choices of rings are the integer rings of the cyclotomic number fields $\mathbb{Q}(\omega_m)$, of degree $n = \varphi(m)$, where ω_m is a primitive m-th root of unity. This choice further ensures the hardness of the decisional version of Ring-LWE under the same worst-case Ideal-SVP hardness assumption.

For some time, it was generally assumed that the structured versions of lattice problems should be just as hard to solve as the unstructured ones: only some (almost) linear-time advantages were known. This assumption was challenged by a claim of Campbell et al. [CGS14]: a quantum polynomial-time attack against a few schemes (Soliloquy, and the fully-homomorphic encryption scheme of [SV10] and cryptographic multilinear maps [GGH13, LSS14], all using *principal ideals*). Parts of this claim were quickly supported by numerical experiments [Sch15]. And indeed, Biasse and Song [BS16] proved that the Principal Ideal Problem could be efficiently solved using a quantum computer. In other words, given a principal ideal one may recover an arbitrary generator in quantum polynomial time. Analyzing the geometry of cyclotomic units, Cramer et al. [CDPR16] proved that from an arbitrary generator, one could asymptotically recover a short one.

Whereas some cryptosystems were broken by this quantum attack, the current limitations this approach to tackle more standard problems as Ring-LWE are three-fold.

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

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

1.1. **Contribution.**

This work provides strong evidence that the general case of Ideal-SVP is not harder than the principal case for large — yet non-trivial — approximation factors. As a consequence, the approximation factors reachable in quantum polynomial time appear to be significantly better in arbitrary ideals of cyclotomic fields of prime-power conductor than known for general lattices.

The strategy consists in reducing the problem over general ideals to that over principal ideals, for cyclotomic fields of prime-power conductor m. We show that under some heuristic assumptions on the structure of the class group, it is possible to solve the *close principal multiple* (CPM) problem in quantum polynomial time for an interesting approximation factor. More precisely, the CPM problem consists

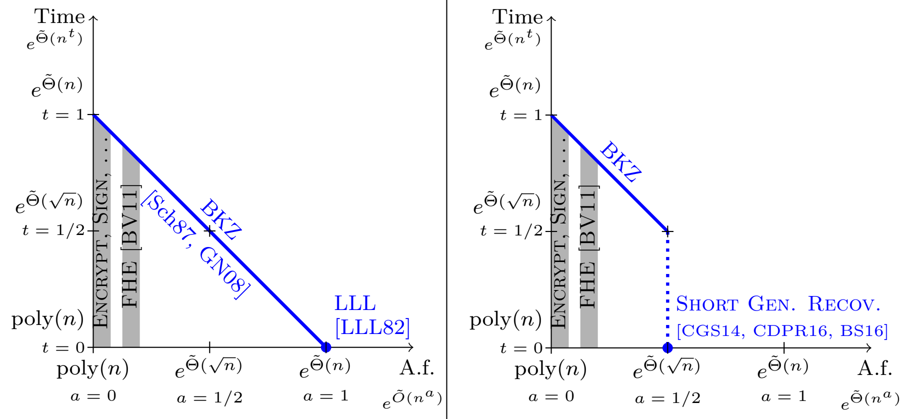

Worst-case SVP on arbitrary lattices

In the general case, the best known algorithms (BKZ [Sch87] and Slide [GN08]) run in time $\exp(\tilde{\Theta}(n^t))$ for an approximation factor $\exp(\tilde{\Theta}(n^a))$, where $t + a = 1$.

Worst-case SVP on principal ideal lattices over $\mathbb{Q}(\omega_m)$

For principal ideals of cyclotomic rings (of prime-power conductor), the aforementioned results give a quantum polynomial runtime (i.e., $t = 0$) for any $a \geq 1/2$.

FIGURE 1. Best known (quantum) Time–Approximation factor tradeoffs to solve approx-SVP in arbitrary lattices (on the left) and in principal ideal lattices (on the right), in the worst case. Approximation factors used in cryptography are typically between polynomial poly(n) and quasi-polynomial exp(polylog(n)).

in finding a principal ideal $\mathfrak{c} \subset \mathfrak{a}$ for an arbitrary ideal $\mathfrak{a}$, such that the algebraic norm of $\mathfrak{c}$ is not much larger than the norm of $\mathfrak{a}$, say up to a factor $\exp(\tilde{O}(n^{1+c}))$.

Our main tool to solve CPM is the Stickelberger ideal, an object providing explicit class relations between an ideal and its Galois conjugates. An important fact is that this ideal has many short elements and that these can be explicitly constructed (see for example [Sch10]). This leads to a quantum polynomial time algorithm to solve CPM for a factor $\exp(\tilde{O}(n^{1+c}))$, where the constant c depends on the structure of the class group. It remains to apply the short generator recovery to $\mathfrak{c}$ to find a short vector of $\mathfrak{a}$, approximating the shortest vector by a factor $\exp(\tilde{O}(n^{\max(1/2,c)}))$.

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

Concluding formally on which value of c can be achieved is not straightforward, as it relies on the structure of the class group Cl_K^- as a $\mathbb{Z}[G]$-module (see Section 2.2). Based on computations of the class group structure of Schoof [Sch98] and a heuristic argument, we strongly believe it is plausible that $c = 1/2$ is reachable at least for a dense family of conductors m, if not all.

1.2. **Impact, open questions and recommendations.** To the best of our knowledge, this new result does not immediately lead to an attack on any proposed scheme, since most of them are based on Ring-LWE and that obstacles (ii) and (iii) remain. This leaves two crucial open cryptanalytic questions.

- The first question is whether the approximation factors in those attacks can be improved, potentially increasing the running time. One could for example consider many CPM solutions rather than just one, and hope that one of them leads to a much shorter vector.
- The second is whether an oracle for Ideal-SVP (an approx-SVP oracle for modules of rank 1) can be helpful to solve Ring-LWE, which can be formulated as an "unusually-Short Vector Problem" over a module of rank 3. The natural approach of generalizing LLL over other rings than $\mathbb{Z}$ [Nap96] fails since only a few cyclotomic rings of small degree are Euclidean [LJ75].

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

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

2. Overview

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

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

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

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

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

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

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

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

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

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

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

Because $g \in \mathfrak{c} \subset \mathfrak{a}$, one has found a short vector of $\mathfrak{a}$, larger than the shortest vector of $\mathfrak{a}$ by a sub-exponential approximation factor $\exp(\tilde{O}(n^{\max(1/2,c)}))$.

2.2.2. *CPM as a close vector problem.* Even before trying to solve this problem, one may wonder if such a close principal multiple should exist in general. A positive answer follows from the results of [JMV09, JW15]: setting a prime factor base $\mathfrak{B} = \{\mathfrak{p}, N\mathfrak{p} \leq n^{4+o(1)}\}$, for any class $C \in \mathrm{Cl}_K$, there exists a non-negative small solution $e \in \mathbb{Z}_{\geq 0}^{\mathfrak{B}}$ to the class equation $[\prod \mathfrak{p}^{e_\mathfrak{p}}] = C$, of ℓ_1-norm $\|e\|_1 \leq O(n^{1+o(1)})$. This proves the existence of a solution $\mathfrak{b} = \prod \mathfrak{p}^{e_\mathfrak{p}}$ to the CPM problem as small as $\exp(\tilde{O}(n^{1+c}))$ for $c = o(1)$.

The previous argument is based on the analysis of the expander properties of certain Caley graphs on the class group. To achieve an algorithmic solution rather than an existential one, we instead write the class group using lattices. If the factor base $\mathfrak{B}$ generates the whole class group, then one may rewrite $\mathrm{Cl}_K \simeq \mathbb{Z}^{\mathfrak{B}}/\Lambda$ where Λ is the lattice of class relations: $\Lambda = \{e \in \mathbb{Z}^{\mathfrak{B}}, [\prod \mathfrak{p}^{e_\mathfrak{p}}] = [\mathcal{O}_K]\}$. In fact, it will be enough to consider any full-rank sublattice $\Gamma \subset \Lambda$ of class relations.

The CPM problem can now be rephrased as a close vector problem: given a class $C = [\mathfrak{a}]^{-1} \in \mathrm{Cl}_K$, one first computes a representative of that class[3] $\alpha \in \mathbb{Z}^{\mathfrak{B}}$ in base $\mathfrak{B}$, and then searches for a close lattice vector $\beta \in \Gamma$. This provides a solution[4] $\mathfrak{b} = \prod \mathfrak{p}^{\beta_\mathfrak{p} - \alpha_\mathfrak{p}}$, of norm at most $B^{\|\beta-\alpha\|_1}$, where B is a bound such that $N\mathfrak{p} \leq B$ for every $\mathfrak{p} \in \mathfrak{B}$. It is therefore sufficient to find an appropriate factor base together with a good basis of the lattice of relations Γ to attack this problem.

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

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

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

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

A technical issue is that the Stickelberger ideal is not of full rank in $\mathbb{Z}[G]$ as a $\mathbb{Z}$-module, so needs to be extended in order to serve as the lattice of relations Γ. This can be resolved by working only with the *minus* part Cl_K^- of the class group, i.e., the relative class group of K over the maximal real subfield K^+, which is annihilated by the augmented Stickelberger ideal $S' = S + (1+\tau)\,\mathbb{Z}[G]$.

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

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

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

The second assumption that we know of a factor basis of the form $\mathfrak{B} = \{\mathfrak{p}^\sigma = \sigma(\mathfrak{p}), \sigma \in G\}$ for a single ideal $\mathfrak{p}$ of small norm $N\mathfrak{p} = \mathrm{poly}(n)$ can also be relaxed.

[3] Computing such a representative is a class group discrete logarithm problem, and is not so straightforward. But it was recently proved that such problems can be solved in quantum polynomial time [BS16], see Proposition 3.1.

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

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

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

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

Therefore, the crux of the matter is about how small of a factor basis $\mathfrak{B}$ can be built[5]. The structure of the class group Cl_K^- remains quite elusive, but it appears that it admits a very small minimum number of generators as a $\mathbb{Z}[G]$-module. Schoof [Sch98] computed that for all prime conductors $m \leq 509$, Cl_K^- is $\mathbb{Z}[G]$-cyclic (i.e., it is generated by a single element as a $\mathbb{Z}[G]$-module). This property is more than enough to argue that one can efficiently find a small generating set and reach $c = 1/2$, when heuristically considering that classes of small random ideals behave similarly to uniformly random classes. Even if the minimal number of generators is not always 1 but still small, say $O(n^\epsilon)$ for some $\epsilon > 0$, this argument allows to reach $c = 1/2 + \epsilon$.

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]). Yet, building on the recent breakthrough on quantum algorithms for the Hidden Subgroup Problem in large dimensions [EHKS14], Biasse and Song [BS16] introduced quantum algorithms to perform S-unit group computations, class group computations, and to solve the principal ideal problem (PIP) in quantum polynomial time. In particular, we derive from their work a quantum polynomial time algorithm to solve the following class group discrete logarithm problem: given a basis $\mathfrak{B}$ of ideals generating a subgroup of the class group Cl_K containing the class of $\mathfrak{a}$, express the class of $\mathfrak{a}$ as a product of ideals in $\mathfrak{B}$. More precisely:

Proposition 3.1 (Class group discrete logarithm, Corollary of [BS16, Theorem 1.1])**.** *Let $\mathfrak{B}$ be a set of prime ideals generating a subgroup H of Cl_K. There exists a quantum algorithm $\mathrm{ClDL}_{\mathfrak{B}}$ which, when given as input any ideal $\mathfrak{a}$ in $\mathcal{O}_K$ such that $[\mathfrak{a}] \in H$, outputs a vector $\mathbf{y} = \mathrm{ClDL}_{\mathfrak{B}}(\mathfrak{a}) \in \mathbb{Z}^{\mathfrak{B}}$ such that $\prod \mathfrak{p}^{y_\mathfrak{p}} \sim \mathfrak{a}$, and runs in polynomial time in $n = \deg(K)$, $\max_{\mathfrak{p}\in\mathfrak{B}}\{\log(N\mathfrak{p})\}$, $\log(N\mathfrak{a})$, and $|\mathfrak{B}|$.*

Proof. The prime factorization $\mathfrak{a} = \mathfrak{q}_1^{a_1} \dots \mathfrak{q}_k^{a_k}$ can be obtained [EH10] in polynomial time in n, $\log(\Delta_K)$ and $\log(N\mathfrak{a})$, by adapting Shor's quantum factoring algorithm [Sho97]. 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

[5] Note that, as a computational problem, this task is *non-uniform*. That is, it must be ran once for each conductor m of interest, but does not need to be re-ran for each CPM instance in $\mathcal{O}_K$. A proof of existence of such a factor basis would already have a consequence in a complexity theoretic perspective. We however heuristically argue in Section 7 that a good basis can actually be found efficiently.

of real embeddings of K, and n_2 is the number of pairs of complex embeddings. Consider the homomorphism

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

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

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

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

4. Close multiple in the relative class group

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

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

4.1. Random walks to the relative class group.

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

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

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

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

$$N\mathfrak{b} = \prod_{i=1}^{\ell} N\mathfrak{p}_i \leq B^\ell = 2^{\tilde{O}(\log h)} = 2^{\tilde{O}(n)},$$

If $[\mathfrak{a}]$ is the starting point of the random walk in the graph, the endpoint $[\mathfrak{ab}]$ falls in Cl_K^- with probability at least $1/(2h_K^+)$, and therefore an ideal $\mathfrak{b}$ such that $[\mathfrak{ab}] \in \mathrm{Cl}_K^-$ can be found in probabilistic polynomial time in h_K^+.

4.2. **The real class number.**

The running time of this step relies on the real class number h_K^+ being polynomial in n. The literature on h_K^+ provides strong theoretical and computational evidence that it is indeed small enough. For any integer k, let $h^+(k)$ be the class number of the maximal totally real subfield of the cyclotomic field of conductor k. First, the authors of [BPR04] formulate and argue in favor of the following conjecture, based on Cohen-Lenstra heuristics.

Conjecture 4.1 (Buhler, Pomerance, Robertson)**.** *For all but finitely many pairs (ℓ, e), where ℓ is a prime and e is a positive integer, we have $h^+(\ell^{e+1}) = h^+(\ell^e)$.*

Weber conjectured long ago the particular case that $h^+(2^e) = 1$ for all e. A direct consequence is that for fixed ℓ and increasing e, $h^+(\ell^e)$ is $O(1)$. But $h^+(\ell)$ itself is also small: in [Sch03], Schoof computed all the values of $h^+(\ell)$ for $\ell < 10,000$ (correct under heuristics of type Cohen-Lenstra, and Miller proved in [Mil15] that it is correct 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. Short relations in Cl_K^- via the Stickelberger ideal

Consider any ideal $\mathfrak{f}$ of $\mathcal{O}_K$ such that $[\mathfrak{f}] \in \mathrm{Cl}_K^-$, and $\mathfrak{F} = G(\mathfrak{f})$ its orbit for the action of the Galois group. Let R be the group ring $\mathbb{Z}[G]$. It projects to $\mathbb{Z}^{\mathfrak{F}}$, via the map sending σ to $\mathbf{1}_{\mathfrak{f}^\sigma}$. The goal of this section is to a explicit a lattice of class relations in $\mathbb{Z}^{\mathfrak{F}}$ with an explicit set of generators. Our main tool will be the *Stickelberger ideal.* This will allow to reduce the representation of a given class expressed in $\mathfrak{F}$, in Subsection 5.3

Recall that the Galois group G is canonically isomorphic to $(\mathbb{Z}/m\mathbb{Z})^*$ via $a \mapsto \sigma_a = \zeta_m \mapsto \zeta_m^a$. In the following, the elements of $(\mathbb{Z}/m\mathbb{Z})^*$ are canonically identified with the positive integers $0 < a_1 < a_2 < \cdots < a_n < m$ such that each a_i is coprime to m. The elements of G are indexed as $(\sigma_{a_1}, \dots \sigma_{a_n})$. Define the extra element $a_{n+1} = m + a_1$, and note that $a_2 \leq 3$ and that $a_{i+1} - a_i \leq 2$ for any i. The norms $\|\cdot\|$ and $\|\cdot\|_1$ over R are defined via the isomorphism $\mathbb{Z}[G] \cong_{\mathbb{Z}} \mathbb{Z}^n$ sending σ_{a_i} to the i-th canonical vector.

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

5.1. **The (augmented) Stickelberger ideal.**

Definition 5.1 (The Stickelberger ideal)**.** *The* Stickelberger element $\theta \in \mathbb{Q}[G]$ *is defined as*

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

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

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

Theorem 5.2 (Stickelberger's theorem [Was12, Thm. 6.10])**.** *The Stickelberger ideal annihilates the ideal class group of K. In other words, for any ideal $\mathfrak{h}$ of $\mathcal{O}_K$ and any $s \in S$, the ideal $\mathfrak{h}^s$ is principal.*

We cannot directly use $S \subset R$ as our lattice of class relations since it does not have full rank in R as a $\mathbb{Z}$-module (precisely its $\mathbb{Z}$-rank is $n/2+1$). To solve this issue, we will augment the Stickelberger ideal to a full-rank ideal which still annihilates the minus part Cl_K^- of the class group.

Definition 5.3. *The augmented Stickelberger ideal S' is defined as*

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

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

Lemma 5.4. *The augmented Stickelberger ideal S' annihilates Cl_K^-. In other words, for any ideal $\mathfrak{h}$ of $\mathcal{O}_K$ such that $[\mathfrak{h}] \in \mathrm{Cl}_K^-$ and any $s \in S$, the ideal $\mathfrak{h}^s$ is principal. Moreover, S' has full-rank n inside R as a $\mathbb{Z}$-module.*

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

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

5.2. Short generating vectors of the augmented Stickelberger lattice.

Lemma 5.5. *The Stickelberger lattice is generated by the vectors $v_i = (a_i - \sigma_{a_i})\theta$ for $i \in \{2, \ldots, n+1\}$.*

Proof. This is almost Lemma 6.9 of [Was12]. There, S is considered as an ideal in R, but the proof actually shows that these elements generate S as a $\mathbb{Z}$-module. □

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

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

Lemma 5.6. *We have the following properties:*

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

Proof. We prove each item individually.

(1) First note that $\{w_2, \ldots, w_{n+1}\}$ generates S: this is a direct consequence of Lemma 5.5 and the construction of W. By definition of $R = \mathbb{Z}[G]$, the set $\{(1+\tau)\sigma, \sigma \in G\}$ generates $(1+\tau)R$. One can conclude from the definition of $S' = S + (1+\tau)R$.

(2) We follow the computation in the proof of [Was12, Lemma 6.9]:

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

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

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

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

□

5.3. **Reducing a class representative in an R-cycle of Cl_K^-.** Using the generating set of short class relations above, we wish to reduce an arbitrary representative $\alpha \in R$ of a class C' as $[\mathfrak{h}^\alpha]$ to a small representation $C' = [\mathfrak{h}^\beta]$, for some small $\beta \in R$. We shall rely on the following close vector algorithm.

Proposition 5.7 (Close vector algorithm)**.** *Let $\Gamma \subset \mathbb{R}^k$ be a lattice, and let W be a set generating Γ. There exists a (classical) polynomial time algorithm* CV*, that when given any $y \in \Gamma \otimes \mathbb{R}$ as input, outputs a vector $x = \mathrm{CV}(y, W) \in \Gamma$ such that $\|x - y\|_1 \leq \frac{k}{2} \cdot \max\{\|w\| \mid w \in W\}$.*

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

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

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

Theorem 5.8. *Assume $n \geq 3$. There is an algorithm* Reduce*, that given $\alpha \in R$, finds in polynomial time in n and* $\log(||\alpha||)$*, an element $\beta = \mathrm{Reduce}(\alpha) \in R$ such that $||\beta||_1 \leq n^{3/2}$, and $C^\alpha = C^\beta$ for any $C \in \mathrm{Cl}_K^-$.*

Proof. Let W be the basis for the augmented Stickelberger ideal S' as in Lemma 5.6. From Lemma 5.4, it has full rank in R. So the close vector algorithm from Proposition 5.7 can be applied to find an element $\gamma = \mathrm{CV}(\alpha, W) \in S'$ such that $||\alpha - \gamma||_1 \leq \frac{n}{2} \cdot \max\{\|w\| \mid w \in W\} \leq n^{3/2}$. Let $\beta = \alpha - \gamma$. For any $C \in \mathrm{Cl}_K^-$, Lemma 5.4 implies that $C^\gamma = 0$ and therefore $C^\alpha = C^\beta$. □

6. Close principal multiple within the relative class group

For this section, suppose the ideal $\mathfrak{a}$ is in the relative class group Cl_K^-. We are looking for an integral ideal $\mathfrak{b}$ in $\mathcal{O}_K$ of small norm such that $\mathfrak{ab}$ is principal.

Let $\mathfrak{B}$ be a set of prime ideals of $\mathcal{O}_K$ generating Cl_K^-. Consider the morphism

$$\phi : \mathbb{Z}^{\mathfrak{B}} \longrightarrow \mathscr{I}_K : (x_\mathfrak{p})_{\mathfrak{p}\in\mathfrak{B}} \longmapsto \prod_{\mathfrak{p}\in\mathfrak{B}} \mathfrak{p}^{x_\mathfrak{p}},$$

where $\mathscr{I}_K$ is the group of fractional ideals in $\mathcal{O}_K$. Let $P = \{p_1, \dots, p_d\}$ be the set of rational primes below the primes of $\mathfrak{B}$. We will show in Subsection 7 that these generators can be chosen such that $N\mathfrak{p} = \mathrm{poly}(n)$. We obviously have $d \leq |\mathfrak{B}| = O(n)$, but much better bounds will be discussed in Subsection 7.

Algorithm 1 Close principal multiple in the relative class group: CPM^-

```
Require: An ideal $\mathfrak{a}$ in $\mathcal{O}_K$ such that $[\mathfrak{a}] \in \mathrm{Cl}_K^-$
Ensure: An (integral) ideal $\mathfrak{b}$ in $\mathcal{O}_K$ such that $\mathfrak{ab} \sim \mathcal{O}_K$ and $N\mathfrak{b} = \exp\left(\tilde{O}\left(dn^{3/2}\right)\right)$
$\mathbf{y} \leftarrow \mathrm{ClDL}_{\mathfrak{B}}(\mathfrak{a})$
for i = 1 to $d$ do
    $\alpha_i \leftarrow \sum_{\sigma\in G_i} y_{(\mathfrak{p}_i^\sigma)}\sigma \in \mathbb{Z}[G]$
    $\beta_i \leftarrow \mathrm{Reduce}(\alpha_i)$
    $(\gamma_i^+, \gamma_i^-) \leftarrow$ the pair of elements in $\mathbb{Z}[G]$ with only positive coefficients, such that $\gamma_i^+ - \gamma_i^- = -\beta_i$
    $\mathfrak{b}_i \leftarrow \mathfrak{p}_i^{\gamma_i^+ + \tau\gamma_i^-}$
end for
$\mathfrak{b} \leftarrow \prod_{i=1}^d \mathfrak{b}_i$
return $\mathfrak{b}$
```

Theorem 6.1. *Algorithm 1,* CPM^-*, runs in quantum polynomial time in* $n = \deg(K)$ *and* $\log(N\mathfrak{a})$*, and is correct.*

Proof. The running time follows immediately from Proposition 3.1, Theorem 5.8, and the fact that d is polynomially bounded in n. Let us now prove the correctness. We have

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

Observe that for each i, $\mathfrak{b}_i \sim \mathfrak{p}_i^{-\beta_i}$, since $\mathfrak{p}_i^{-1} \sim \mathfrak{p}_i^{\tau}$. From Theorem 5.8, we obtain $\mathfrak{p}_i^{\alpha_i}\mathfrak{b}_i \sim \mathcal{O}_K$, which implies that $\phi(\mathbf{y})\mathfrak{b} \sim \prod_{i=1}^d \mathfrak{p}_i^{\alpha_i}\mathfrak{b}_i \sim \mathcal{O}_K$. From Proposition 3.1, we have $\phi(\mathbf{y}) \sim \mathfrak{a}$, and therefore $\mathfrak{ab} \sim \mathcal{O}_K$.

Now, Theorem 5.8 ensures that $||\beta||_1 \leq n^{3/2}$. So $||\gamma_i^+||_1 + ||\gamma_i^-||_1$ is bounded by $n^{3/2}$ and we obtain that $N\mathfrak{b}_i \leq (N\mathfrak{p}_i)^{n^{3/2}}$. Then,

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

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

7. Good generating sets for the relative class group

The norm of the ideal $\mathfrak{b}$ found by Algorithm 1 is bounded by $\exp\left(\tilde{O}\left(dn^{3/2}\right)\right)$, where d is the number of distinct rational primes below the prime ideals in the basis $\mathfrak{B}$ of Cl_K. To optimize the quality of $\mathfrak{b}$, we are thereby interested in finding a basis with a small value of d. We will provide heuristic arguments and computational evidence that $\mathfrak{B}$ can be chosen so that d is polylogarithmic in n.

From [JW15, Corollary 6.5], for any subgroup H of the class group Cl_K, and any constant $\varepsilon > 0$, there is a bound $B_H = O\left([\mathrm{Cl}_K : H]n \log \Delta_K)^{2+\varepsilon}\right)$ such that H is generated by classes of ideals of $\mathcal{O}_K$ of prime norm bounded by B_H. In particular, Cl_K^- is generated by ideals of prime norm bounded by

$$B := B_{\mathrm{Cl}_K^-} = O\left(\left(h_K^+ n \log \Delta_K\right)^{2+\varepsilon}\right) = O\left(\left(h_K^+ n^2 \log n\right)^{2+\varepsilon}\right),$$

where $h_K^+ = |\mathrm{Cl}_{K^+}|$ is the class number of K^+. Let $\mathfrak{G}$ be the set of all prime ideals whose classes are in Cl_K^- and norms are primes bounded by B. It is a generating set for Cl_K^-, but its ideals have too many distinct prime norms to be practical. It is not hard to see that one can find a subset of $\mathfrak{G}$ of size $d_{\max}$ generating Cl_K^- as a group, where $d_{\max}$ is the number of components in the primary decomposition of Cl_K^-. This number $d_{\max}$ is expected to be small, but obtaining provable bounds is a difficult problem. A first approximation comes from the fact that $d_{\max} \leq d'_{\max}$, where $d'_{\max}$ is the number of prime factors of h_K^- counted with multiplicity. The factorizations of the first values of h_K^- can be found in [Was12, Table 3]. To find a much smaller set of generators, we will take a more powerful approach using the $\mathbb{Z}[G]$-module structure of Cl_K^-.

7.1. Generating Cl_K^- with random classes.

Observe that if a set of ideals $\mathfrak{C}$ generates Cl_K^- as a $\mathbb{Z}[G]$-module, then one can choose $\mathfrak{B} = \{\mathfrak{h}^\sigma \mid h \in \mathfrak{C}, \sigma \in G\}$ as a generating set, and obtain $d = |\mathfrak{C}|$. Moreover, the minimal number r of $\mathbb{Z}[G]$-generators of Cl_K^- is expected to be very small, so one should be able to find very small generating sets $\mathfrak{C}$. More precisely, Schoof [Sch98] proved that Cl_K^- is $\mathbb{Z}[G]$-cyclic for every prime conductor $m \leq 509$, i.e., $r = 1$. This cyclicity could be expected to be the typical behavior asymptotically, but more loosely, we make the assumption that r is bounded by $\mathrm{polylog}(n)$. Considering that r is small, we can heuristically find a small generating set $\mathfrak{C}$ by randomly taking ideals of small norm, thanks to the following proposition.

Proposition 7.1. *Let K be a cyclotomic ring of conductor m, with Galois group G and relative class group Cl_K^-. Let r be the minimal number of $\mathbb{Z}[G]$-generators of Cl_K^-. Let $\alpha \geq 0$ be a parameter, and s be any integer at least $r(\log_2 \log_2(h_K^-) + \alpha)$ (note that $\log_2 \log_2(h_K^-) \sim \log_2(n)$). Let $g_1, \ldots, g_s$ be s independent uniform elements of Cl_K^-. The probability that $\{g_1, \ldots, g_s\}$ generates Cl_K^- as a $\mathbb{Z}[G]$-module is at least $\exp\left(-\frac{3}{2^\alpha}\right) = 1 - O(2^{-\alpha})$.*

In other words, a set of $\Theta(r \log(n))$ random ideal classes in Cl_K^- will generate this $\mathbb{Z}[G]$-module with very good probability. Under the heuristic that this behavior remains true when restricting to random ideals of norms bounded by $\mathrm{poly}(n)$, we can conclude that one may efficiently build a basis $\mathfrak{B}$ such that d is bounded by $\mathrm{polylog}(n)$. Let us now prove Proposition 7.1.

Lemma 7.2. *Let $\mathcal{O}$ be a Dedekind domain, and $\mathfrak{h} \subset \mathcal{O}$ be an integral ideal. Let $g_1, \dots, g_s$ be s independent uniform elements from $\mathcal{O}/\mathfrak{h}$. Then, the probability that the set $\{g_1, \dots, g_s\}$ generates $\mathcal{O}/\mathfrak{h}$ as an $\mathcal{O}$-module is*

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

Proof. Let $\mathfrak{h} = \mathfrak{p}_1^{\alpha_1} \cdots \mathfrak{p}_j^{\alpha_j}$ be the prime factorization of $\mathfrak{h}$. The event $E : \mathcal{O}g_1 + \cdots + \mathcal{O}g_s = \mathcal{O}/\mathfrak{h}$ is equivalent to the conjunction $\bigwedge_{i=1}^{j} F_i$, where

$$F_i : \mathcal{O}g_1 + \cdots + \mathcal{O}g_s = \mathcal{O}/\mathfrak{p}_i^{\alpha_i}.$$

Note that F_i holds if and only if there is at least one $k \in \{1, \dots, s\}$ such that g_k is coprime with $\mathfrak{p}_i$. Since the g_k's are chosen independently, we obtain

$$\Pr[F_i] = 1 - (N\mathfrak{p}_i)^{-s} \geq 1 - 2^{-s}.$$

Now, note that the number j of distinct prime factors of $\mathfrak{h}$ is at most $\log_2 N\mathfrak{h}$. We conclude from the fact that the events F_i are independent, by the chinese remainder theorem. □

Lemma 7.3. *Let $\mathcal{O}$ be a Dedekind domain, and M be a finite $\mathcal{O}$-module of cardinality h and let r be the minimal number of $\mathcal{O}$-generators of M. Let $g_1, \dots, g_s$ be s independent uniform elements from M. Then, the probability that the set $\{g_1, \dots, g_s\}$ generates M as an $\mathcal{O}$-module is*

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

Proof. Since M is a torsion module over a Dedekind domain, there exist r ideals $\mathfrak{h}_1, \dots, \mathfrak{h}_r$ such that $M = \bigoplus_{i=1}^{r} \mathcal{O}/\mathfrak{h}_i$; in particular, $\log_2 h = \sum_{i=1}^{r} \log_2 N\mathfrak{h}_i$. Consider the first random elements $g_1 \dots g_{s'}$ where $s' = \lfloor s/r \rfloor$, and their projections $g'_1 \dots g'_{s'}$ on the first component $\mathcal{O}/\mathfrak{h}_1$. By the above Lemma 7.2, $\{g'_1 \dots g'_{s'}\}$ generates $\mathcal{O}/\mathfrak{h}_1$ with probability at least $(1 - 2^{-s'})^{\log_2 N\mathfrak{h}_1}$.

Let $M_1 = \mathcal{O}g_1 + \cdots + \mathcal{O}g_{s'}$. To conclude by induction, it suffices to note that M/M_1 is generated by $r - 1$ (or less) elements. □

Theorem 7.4. *Let H be a cyclic group, and M a finite, $\mathbb{Z}[H]$-module of cardinality h, and r be the minimal number of $\mathbb{Z}[H]$-generators of M. Let $g_1, \dots, g_s$ be s independent uniform elements of M. The probability that the set $\{g_1, \dots, g_s\}$ generates M as a $\mathbb{Z}[H]$-module is*

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

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

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

For each $d \mid t$, identify $\mathbb{Z}[\omega_d]$ with its natural embedding in $\mathbb{Z}[H]$, and let the $e_d \in \mathbb{Z}[H]$ be the idempotent identified with the unit of $\mathbb{Z}[\omega_d]$. It induces the decomposition $M = \bigoplus_{d|t} e_d M$. First, note that $\log_2 h = \sum_{d|t} \log_2 h_d$, where h_d is the cardinality of $e_d M$. Second, each $e_d M$ is generated over $\mathbb{Z}[\omega_d]$ by at most r

elements. Applying Lemma 7.3, we obtain

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

□

Proof of Proposition 7.1. Note that G is trivial if and only if $K=\mathbb{Q}$, in which case Cl_K is trivial, and so is the proposition. Otherwise, observe that G splits as $\mathbb{Z}/2\mathbb{Z}\times H$ where H is a cyclic group, and the component $\mathbb{Z}/2\mathbb{Z}$ corresponds to the complex conjugation τ. Note that for any $x\in\mathrm{Cl}_K^-$, the orbits $\mathbb{Z}[G]x$ and $\mathbb{Z}[H]x$ coincide since $\tau\in G$ acts like $-1\in\mathbb{Z}[H]$ on Cl_K^-. Therefore r is the minimal number of $\mathbb{Z}[H]$-generators of Cl_K^-. We obtain from Theorem 7.4 that the probability that $\{g_1,\ldots,g_s\}$ generates Cl_K^- as a $\mathbb{Z}[H]$-module is at least $(1-2^{-\lfloor s/r\rfloor})^{\log_2 h_K^-}$. For any $0<x\leq 1/2$, we have $\ln(1-x)>-(3/2)x$. We get

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

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

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

□

Acknowledgements

The authors would like to thank René Schoof for helpful and interesting discussions. The second author was supported by a grant through a public-private partnership with NXP Semiconductors. The third author was supported by the Swiss National Science Foundation under grant number 200021-156420.

References

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

[BCLvV16] Daniel J. Bernstein, Chitchanok Chuengsatiansup, Tanja Lange, and Christine van Vredendaal. Ntru prime. Cryptology ePrint Archive, Report 2016/461, 2016. http://eprint.iacr.org/2016/461.

[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.

[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.

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

[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.

[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_CRYPTO/S07_Systems_and_Attacks/S07_Groves_Annex.pdf`.

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

[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 *STOC*, pages 293–302. ACM, 2014.

[GGH13] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In *EUROCRYPT*, pages 1–17, 2013.

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

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

[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.

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

[LJ75] Hendrik W. Lenstra Jr. Euclid's algorithm in cyclotomic fields. *J. London Math. Soc*, 10:457–465, 1975.

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

[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.

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

[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.

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

[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.

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

[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.

[Sch15] John Schank. LogCvp, Pari implementation of CVP in $\text{Log}\mathbb{Z}[\zeta_{2^n}]^*$. `https://github.com/jschanck-si/logcvp`, March 2015.

[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.

[Sin80] Warren Sinnott. On the stickelberger ideal and the circular units of an abelian field. *Inventiones math.*, 62:181–234, 1980.

[SS11] Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *EUROCRYPT*, pages 27–47, 2011.

[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.