



On the Shortness of Vectors to Be Found by the Ideal-SVP Quantum Algorithm

Léo Ducas¹(✉), Maxime Plançon², and Benjamin Wesolowski¹

¹ Cryptology Group, CWI, Amsterdam, The Netherlands
ducas@cwi.nl

² Ecole Normale Supérieure Paris-Saclay, Cachan, France

Abstract. The hardness of finding short vectors in ideals of cyclotomic number fields (hereafter, Ideal-SVP) can serve as a worst-case assumption for numerous efficient cryptosystems, via the average-case problems Ring-SIS and Ring-LWE. For a while, it could be assumed the Ideal-SVP problem was as hard as the analog problem for general lattices (SVP), even when considering quantum algorithms.

But in the last few years, a series of works has led to a quantum algorithm for Ideal-SVP that outperforms what can be done for general SVP in certain regimes. More precisely, it was demonstrated (under certain hypotheses) that one can find in quantum polynomial time a vector longer by a factor at most $\alpha = \exp(\tilde{O}(n^{1/2}))$ than the shortest non-zero vector in a cyclotomic ideal lattice, where n is the dimension.

In this work, we explore the constants hidden behind this asymptotic claim. While these algorithms have quantum steps, the steps that impact the approximation factor α are entirely classical, which allows us to estimate it experimentally using only classical computing. Moreover, we design heuristic improvements for those steps that significantly decrease the hidden factors in practice. Finally, we derive new provable effective lower bounds based on volumetric arguments.

This study allows to predict the crossover point with classical lattice reduction algorithms, and thereby determine the relevance of this quantum algorithm in any cryptanalytic context. For example we predict that this quantum algorithm provides shorter vectors than BKZ-300 (roughly the weakest security level of NIST lattice-based candidates) for cyclotomic rings of rank larger than about 24000.

Keywords: Quantum cryptanalysis · Cyclotomic ideal lattices

L. Ducas—Supported by a Veni Innovational Research Grant from NWO under project number 639.021.645 and by the European Union Horizon 2020 Research and Innovation Program Grant 780701.

M. Plançon—Part of this work was realized during an internship at the Cryptology Group, CWI, Amsterdam.

B. Wesolowski—Supported by the ERC Advanced Investigator Grant 740972 (ALGSTRONGCRYPTO).

© International Association for Cryptologic Research 2019

A. Boldyreva and D. Micciancio (Eds.): CRYPTO 2019, LNCS 11692, pp. 322–351, 2019.

https://doi.org/10.1007/978-3-030-26948-7_12

1 Introduction

The shortest vector problem (hereafter, SVP), that is the problem of finding the shortest vector of a Euclidean lattice, is a central hard problem in complexity theory. An approximated version (approx-SIVP) can serve as a theoretical foundation for many cryptographic constructions thanks to the worst-case to average-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-SIVP to Learning with Errors (LWE).

However, the efficiency of schemes based on plain SIS and LWE remains unsatisfactory, and one may prefer to rely on certain structured lattices, namely lattices that are also modules over certain rings, as done by the NTRU cryptosystem [HPS98], and more recently by many more cryptosystems based on Ring-SIS and Ring-LWE. The Ring-SIS and Ring-LWE problems also enjoy worst-case to average-case reductions from a variant of approx-SIVP¹ for lattices that are ideals in some ring [Mic07, SSTX09, LPR10, SS11, PRSD17]. The typical choice of ring is the integer ring of a cyclotomic number field $\mathbb{Q}(\omega_m)$, of degree $n = \varphi(m)$, where ω_m is a primitive m -th root of unity. One notable exception is the NTRU Prime cryptosystem [BCLvV17], which was designed to mitigate the potential cryptanalytic risk that we are about to discuss.

The assumption that approx-SIVP is as hard in ideal lattices as in general lattices was challenged by Campbell *et al.* [CGS14], who sketched a quantum polynomial-time attack against a few schemes using so-called *principal ideals* (Soliloquy, and the fully-homomorphic encryption scheme of [SV10]). Following the claims of Campbell *et al.*, 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 in the log-unit lattice, Cramer *et al.* [CDPR16] also confirmed that the secret key (a short generator) of the few aforementioned schemes could be recovered exactly, due to their specific distribution.

Furthermore, it is also proven in [CDPR16] that from an arbitrary generator, one could asymptotically recover a short one, with an approximation factor of $\alpha = \exp(\tilde{O}(n^{1/2}))$. A generalization to all ideals (i.e., not necessarily principal) was provided in [CDW17]. They showed that by exploiting the Stickelberger class relations, one could efficiently find a sub-ideal $\mathfrak{b} \subset \mathfrak{a}$ (i.e., an integral multiple) such that \mathfrak{b} is principal, and such that $|\mathfrak{b}/\mathfrak{a}| \leq \exp(\tilde{O}(n^{3/2}))$ (i.e., the sub-ideal is not much sparser than the original lattice). Putting both results together leads to an approximation factor of $\alpha = \exp(\tilde{O}(n^{1/2}))$ also for non-principal ideals.

Nevertheless, the work of [CDW17] still leaves two obstacles for cryptanalytic applications of their algorithm to the widespread hardness assumptions Ring-SIS, Ring-LWE and NTRU:

1. The approximation factor in the worst-case is asymptotically too large to affect any actual Ring-LWE based scheme, which typically rely on polynomial approximation factors $\alpha = \text{poly}(n)$.

¹ For cyclotomic ideal lattices, approx-SVP and approx-SIVP are trivially equivalent.

2. Ring-SIS and Ring-LWE are known to be at least as hard as Ideal-SIVP [Mic07, SS11, LPR13] but not known to be equivalent. In fact, problems like Ring-SIS, Ring-LWE and NTRU, are naturally phrased as short vector problems in *module* lattices of rank $k \geq 2$. An approach for such a converse reduction would be to generalize LLL over other rings than \mathbb{Z} , but this seems to fail since only a few cyclotomic rings of small degree are Euclidean [Len75].

This work. In this work, we are interested in precisely quantifying the obstacle 1 above. It is proven in [CDPR16] that the short generator that can be recovered is asymptotically close to optimal, yet it is unclear how this asymptotic statement translates in practice. One could fear that the hidden factors in $\alpha = \exp(\tilde{O}(n^{1/2}))$ make α small enough in practice to threaten concrete cryptosystems (assuming obstacle 2 can also be tackled). Or, on the contrary, one could doubt that this algorithm is ever to give better results than classical methods for reasonable dimensions, given how small the Hermite factor $\eta = 1.022^n$ of LLL [LLL82] already is in practice [NS06].²

After some preliminaries in Sect. 2, we recall in Sect. 3 the main steps for solving Ideal-SVP [CGS14, EHKS14, CDPR16, BS16, CDW17]. We discuss the slackness of the bounds derived in these works, and we identify the more meaningful quantities that should be studied to predict more precisely the Hermite factor achieved by the algorithm. We note that we do not need to resort to a quantum computer to experiment with those meaningful quantities, at least by making an informed assumption on the input distribution of the relevant steps (see Assumption 8 and the subsequent discussion). All working hypotheses are summarized in Sect. 3.4.

We then propose in Sect. 4 several heuristic techniques, designed to improve in practice the meaningful quantities determined above. First, we propose to exploit the knowledge of $\Theta(d^2)$ many short vectors of both the Stickelberger and log-unit lattices and go beyond what can be done with just a basis (where d is the dimension). To properly exploit a large number of short vectors, we propose to use an approximate Voronoi-cell-based algorithm [MV10, Laa16, DLdW19], adapted to our specific setting, where we wish to minimize some carefully determined meaningful quantities rather than the Euclidean distance.

In Sect. 5, we discuss our implementation, and report on the experimental behavior of both the original algorithm, and our heuristically improved variant. We observe that the experimental behavior asymptotically matches with the upper bound, and we experimentally determine the hidden constants. We also note that our heuristic variant indeed improves these hidden constants, especially for the Approx-CVP step in the log-unit lattice.

Finally, we study in Sect. 6 the volumetric lower-bounds for the CVP instances. We determine the effective asymptotic behaviour of those lower bounds (i.e., without hidden constants). We note that our bound for the log-unit lattice is not only effective, but also asymptotically better than the one of [CDPR16]. We also perform numerical experiments, which show that the convergence to the asymptotic behaviour is sufficiently fast to allow reliable use of the estimates.

² In the rest of this work, we prefer to use the so called Hermite factor η instead of the approximation factor α ; this is justified in Remark 1.

We conclude in Sect. 7 with a summary of our effective asymptotic predictions. Combining these results, we compare the predicted performance with that of the classical lattice reduction algorithms LLL and BKZ, in Fig. 5. For a concrete example, we predict that the crossover point between the original algorithm and LLL happens for cyclotomic ideals of rank around 4000, and our heuristic improvement brings this crossover point down to rank 1000. We conclude our work by summarizing the limits of the conclusions that can be drawn from this work regarding cryptanalytic concerns.

Concurrent work. Recently, Pellet–Mary, Hanrot and Stehlé proposed [PMHS19] a heuristic algorithm that should reach even lower approximation factors than discussed above, but at the cost of a pre-computation using exponential time and memory, and a computation using sub-exponential time and memory. It also makes use of the approx-CVP algorithm of Laarhoven [Laa16, DLdW19, SD19], but in a different regime, and in a lattice with much less structure. It would be interesting to find an efficient simulation of their precomputation phase, so as to be able to run more extensive experiments and estimate the hidden constants, possibly using the heuristic improvements introduced in this paper.

2 Preliminaries

Vectors are to be read as column-vectors. Matrices are denoted by capital letters. We write a matrix B as $B = (b_1, \dots, b_n)$ where b_i is the i -th column vector of B . We denote by $B^* = (b_0^*, \dots, b_{n-1}^*)$ the Gram-Schmidt orthogonalization of the matrix B .

2.1 Geometry

Norms, asymmetric norms, pseudo-norms. We will use the ℓ_1 , ℓ_2 (Euclidean) and ℓ_∞ norms, respectively defined by $\|x\|_1 = \sum |x_i|$, $\|x\|_2 = \sqrt{\sum x_i^2}$ and $\|x\|_\infty = \max |x_i|$. Beware that, contrary to some of the literature, the notation $\|\cdot\|$ *does not* refer by default to the Euclidean norm, but is a place-holder for any norm, asymmetric norm or pseudo-norm (defined below).

We will make use of two weakened notions of norm during this paper. We recall that a norm $\|\cdot\| : V \rightarrow [0, +\infty)$ on a real vector space V is a function satisfying the three following axioms:

1. Sub-additivity: $\|x + y\| \leq \|x\| + \|y\|$ for all $x, y \in V$,
2. Absolute homogeneity: $\|ax\| = |a| \cdot \|x\|$ for all $a \in \mathbb{R}$, $x \in V$,
3. Positive definiteness: $\|x\| = 0 \Rightarrow x = 0$ for all $x \in V$.

An asymmetric norm $\|\cdot\| : V \rightarrow [0, +\infty)$ is a function verifying axioms 1 and 3 and the following positive homogeneity axiom:

4. Positive homogeneity: $\|ax\| = a \cdot \|x\|$ for all $a \geq 0$, $x \in V$.

Finally, in this article we will call a pseudo-norm a function $\|\cdot\| : V \rightarrow [0, +\infty)$ verifying the axiom 1 and the following linear monotonicity axiom:

5. Linear monotonicity: $\|ax\| \geq \|x\|$ for all $a \geq 1$.

Lattices. A lattice Λ is a discrete subgroup of a finite-dimensional Euclidean vector space \mathbb{R}^n (or Hermitian vector space $\mathbb{C}^{n/2} \simeq \mathbb{R}^n$). A lattice admits a basis, that is a matrix $B \in \mathbb{R}^{d \times n}$ such that $\Lambda = B \cdot \mathbb{Z}^n$ for some $n \leq d$; n is called the dimension of the lattice. Its volume is defined by $\text{Vol}(\Lambda) = \sqrt{|\det(B^t B)|}$ for any basis B of Λ (this measure is independent of the choice of the basis).

To quantify the shortness of a vector $v \in \Lambda$, we use the so-called Hermite factor $\eta = \|v\|_2 / \text{Vol}(\Lambda)^{1/n}$ (where $n = \dim(\Lambda)$) instead of the approximation factor $\alpha = \|v\|_2 / \lambda_1(\Lambda)$ (where $\lambda_1(\Lambda) = \min_{w \in \Lambda \setminus \{0\}} \|w\|_2$), as the minimal length of a lattice is typically not known exactly. Note that this choice does not affect the comparison of reduction performances between different algorithms.

Remark 1. While the latter approximation factor α is often preferred in worst-case complexity theory, the former Hermite factor η is typically more relevant and convenient for average-case cryptanalysis. Note that from Minkowski’s theorem, we have $\lambda_1(\Lambda) \leq (1 + O(1/n)) \sqrt{2n/\pi e} \text{Vol}(\Lambda)^{1/n}$; moreover, for cyclotomic ideal lattices we also have $\lambda_1(\Lambda) \geq \text{Vol}(\Lambda)^{1/n}$. Therefore, the ratio between both measure is reasonably well controlled: $\alpha/\eta \in [1, (1 + O(1)) \sqrt{2n/\pi e}]$. The extreme case $\alpha/\eta = 1$ is reached by orthogonal lattices, and for random lattices the Gaussian heuristic predicts $\alpha/\eta \approx \sqrt{n/2\pi e}$.

Close vector algorithm. We recall from [Bab86] two polynomial time algorithms RoundOff and NearestPlane (as Algorithms 1 and 2) for solving the close vector problem given a basis of short vectors. The output v is guaranteed to lie in the parallelepiped $t + \mathcal{P}(B)$ for RoundOff and $t + \mathcal{P}(B^*)$ for NearestPlane, where

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

This allows to bound $\|v - t\|$, depending on the quality of the basis B , and of considered norm $\|\cdot\|$.

Algorithm 1. RoundOff(B, t)

Require: A basis B of a full-rank lattice $L \subset \mathbb{R}^n$, a target point $t \in \mathbb{R}^n$.

Ensure: A lattice vector $v \in L$ close to t : $v - t \in \mathcal{P}(B)$

- 1: $x \leftarrow B^{-1}t$
 - 2: $y \leftarrow (\lfloor x_1 \rfloor, \dots, \lfloor x_n \rfloor)$
 - 3: $v \leftarrow By$
 - 4: **return** v
-

2.2 Number Theory

Cyclotomic number fields. Throughout this paper, m denotes the power of a prime, ω_m is a primitive m -th root of unity, and $K = \mathbb{Q}(\omega_m)$ is the m -th cyclotomic number field. It is a number field of degree $n = \varphi(m) = \Theta(m)$.

Algorithm 2. NearestPlane(B, t)

Require: A basis B of a full-rank lattice $L \subset \mathbb{R}^n$, a target point $t \in \mathbb{R}^n$.

Ensure: A lattice vector $v \in L$ close to t : $v - t \in \mathcal{P}(B^*)$

```

1:  $f \leftarrow t$ 
2:  $v \leftarrow 0$ 
3: for  $i = n$  downto 1 do
4:    $y \leftarrow \langle t, b_i^* \rangle / \|b_i^*\|^2$ 
5:    $z_i = \lfloor y \rfloor$ 
6:    $f \leftarrow f - z_i b_i$ 
7:    $v \leftarrow v + z_i b_i$ 
8: end for
9: return  $v$ 

```

We denote by G its Galois group over \mathbb{Q} , while $\tau \in G$ denotes complex conjugation. We recall that $G \simeq (\mathbb{Z}/m\mathbb{Z})^\times$, by constructing the automorphism $\sigma_i \in G : \omega \mapsto \omega^i$ for any $i \in (\mathbb{Z}/m\mathbb{Z})^\times$. Complex conjugation corresponds to -1 , i.e., $\tau = \sigma_{-1}$. The norm of an element $x \in K$ is given by $Nx = \prod_{\sigma \in G} \sigma(x)$, and it holds that $Nx \in \mathbb{Q}$ for any element $x \in K$.

We recall that the discriminant Δ_K of cyclotomic number fields K asymptotically satisfies $\log |\Delta_K| = O(n \log n)$ [Was12]. More specifically, for any prime power conductor $m = p^k$, the discriminant of $\mathbb{Q}(\omega_{p^k})$ is $\pm p^{p^{k-1}(pk-k-1)}$.

Ideals of \mathcal{O}_K . The ring of integers of K is denoted $\mathcal{O}_K = \mathbb{Z}[\omega_m]$. An integral ideal $\mathfrak{h} \subset \mathcal{O}_K$ is an additive subgroup closed under multiplication by any element of the ring; more precisely $\forall a \in \mathcal{O}_K, a\mathfrak{h} \subset \mathfrak{h}$. A fractional ideal $f \subset K$ is an ideal of the form $f = \frac{1}{s}\mathfrak{h}$ for some scalar $s \in \mathbb{Z}$. Unless specified to be integral, ideals will be considered to be fractional.

The elements (g_1, \dots, g_r) are generators of the ideal f when $f = \sum_i g_i \mathcal{O}_K$. In particular, when the ideal is generated by a single element, it is called principal. For an integral ideal $\mathfrak{h} \subset \mathcal{O}_K$, the quotient $\mathcal{O}_K/\mathfrak{h}$ is finite and $N\mathfrak{h} = |\mathcal{O}_K/\mathfrak{h}|$ is the norm of the ideal \mathfrak{h} . When \mathfrak{h} is principal, there is an element h such that $\mathfrak{h} = h\mathcal{O}_K$, and the norm of \mathfrak{h} coincides with the algebraic norm of h , i.e., $N\mathfrak{h} = Nh$.

Ideals as lattices. The field K is endowed with a canonical structure of Hermitian vector space via its Minkowski embedding. That is, letting $\zeta_m = \exp(2i\pi/m) \in \mathbb{C}$, and letting $\psi_i : K \rightarrow \mathbb{C}$ be the field morphism sending ω_m to ζ_m^i for each $i \in (\mathbb{Z}/m\mathbb{Z})^\times$ coprime to m , each elements $e \in K$ is identified with the vector $\psi(e) = (\psi_i(e))_{i \in (\mathbb{Z}/m\mathbb{Z})^\times} \in \mathbb{C}^n$. By abuse of notation, we often identify the elements e and $\psi(e)$; in particular, $\|e\|_\alpha$ refers to $\|\psi(e)\|_\alpha$ for $\alpha \in \{1, 2, \infty\}$.

Any ideal \mathfrak{h} of \mathcal{O}_K can be viewed as a Euclidean lattice via the above embedding. The volume of \mathfrak{h} as a lattice relates to its algebraic norm via the equation $\text{Vol}(\mathfrak{h}) = \sqrt{|\Delta_K|} N\mathfrak{h}$.

Class group. The class group $\text{Cl}_K = \mathcal{I}_K/\mathcal{P}_K$ of K is the quotient of the (abelian) multiplicative group of fractional ideals \mathcal{I}_K by the subgroup of principal ideals. We denote by $[\mathfrak{h}]$ the class of the ideal \mathfrak{h} in Cl_K . The trivial class $[\mathcal{O}_K]$ is the class of principal ideals. The class group is written multiplicatively. The minus-part Cl_K^- of the class group is defined as the kernel of the relative norm map $N_{K/K^+} : \text{Cl}_K \rightarrow \text{Cl}_{K^+}$, $[\mathfrak{h}] \mapsto [\mathfrak{h}\mathfrak{h}^\tau]$, where K^+ is the maximal real subfield of K , and \mathfrak{h}^τ denotes the complex conjugation of \mathfrak{h} .

The class number $h_m = |\text{Cl}_K|$ is the order of the class group. Denoting $h_m^+ = |\text{Cl}_{K^+}|$ and $h_m^- = |\text{Cl}_K^-|$ we have $h_m = h_m^+ \cdot h_m^-$.

Galois group ring. The Galois group ring $R = \mathbb{Z}[G]$ is the set of formal linear combinations of elements of G with integral coefficients. The group operation of G is extended to a multiplication law in R , providing R with a ring structure. The ring R acts naturally on the ideals of \mathcal{O}_K as follows: let $s = \sum_{\sigma \in G} s_\sigma \sigma \in R$ and let \mathfrak{h} be an ideal of \mathcal{O}_K , then we define the action of s on \mathfrak{h} as

$$\mathfrak{h}^s = \prod_{\sigma \in G} \sigma(\mathfrak{h})^{s_\sigma}.$$

2.3 Cyclotomic Log-Unit Lattice

We abusively call units of K the elements of the group \mathcal{O}_K^\times . The embeddings of K are all complex, and such that $\overline{\psi_i} = \psi_{-i}$, hence $|\psi_i| = |\psi_{-i}|$, so we define the set of indices $I = (\mathbb{Z}/m\mathbb{Z})^\times / \{\pm 1\}$. The logarithmic embedding

$$\begin{aligned} \text{Log} : K^\times &\rightarrow \mathbb{R}^{n/2} \\ x &\mapsto \mathbf{x} = (\log(|\psi_i(x)|))_{i \in I} \end{aligned}$$

defines a group homomorphism. The Dirichlet Unit Theorem ensures that $\Lambda = \text{Log}(\mathcal{O}_K^\times)$ is a lattice (called the log-unit lattice) of rank $n/2 - 1$. The projection of the log-embedding of an element x on the all-1 vector $\mathbf{1} = (1, \dots, 1)$ directed line is proportional to the logarithm of its algebraic norm $\log(Nx)$. In particular, as the algebraic norm is multiplicative, the algebraic norm of a unit is ± 1 and $\Lambda \perp \text{Span}(\mathbf{1})$. We denote by H the orthogonal complement of $\text{Span}(\mathbf{1})$, the minimal vector space supporting the log-unit lattice Λ . Conversely, we define a reciprocal function to Log , that is

$$\begin{aligned} \text{Exp} : \mathbb{R}^{n/2} &\rightarrow \mathbb{R}^n \\ (x_1, \dots, x_{n/2}) &\mapsto (\exp(x_1), \exp(x_1), \dots, \exp(x_{n/2}), \exp(x_{n/2})) \end{aligned}$$

Up to reordering of coefficients, we have that $(|\psi_i(x)|)_{i \in (\mathbb{Z}/m\mathbb{Z})^\times} = \text{Exp}(\text{Log}(x))$, in particular $\|\text{Exp}(\text{Log}(x))\|_2 = \|x\|_2$.

Not only do we know the structure of the units of \mathcal{O}_K by Dirichlet’s Theorem, but in the case of cyclotomic fields of prime-power conductor we also have an explicit set of relatively short vectors (namely, the $\text{Log } b_{i,j}$ ’s defined below) which generate a finite index sublattice of the log-unit lattice Λ .

More precisely, with $\zeta \in K$ a primitive m -th root of unity, we define the multiplicative group V generated by $\pm\zeta$ and the elements $z_i = \zeta^i - 1$, for $1 \leq i \leq m - 1$. Then, the cyclotomic units are defined as $C = V \cap \mathcal{O}_K^\times$. The elements $b_{ij} = \frac{z_i}{z_j}$ (when only one index is given b_i , we refer to b_{i1}) are units of \mathcal{O}_K . Then, the sublattice $\text{Log } C$ is generated by the vectors $(\text{Log } b_i)_{i \in I \setminus \{1\}}$. The index $[A : \text{Log } C]$ and the length of the vectors $\text{Log } b_{ij}$ are controlled by the following results.

Theorem 1 (See [Was12] Thm. 8.2 and Exercise 8.5). *For any prime power $m > 2$, the index of the log-unit lattice A over $\text{Log } C$ is*

$$[A : \text{Log } C] = h_m^+ < \infty.$$

Corollary 2 (Corollary of [CDPR16], Lemma 6.7). *Let $m = p^k$ be a prime power. Then, $\|\text{Log } b_{ij}\| = O(\sqrt{m})$.*

The two above statements allow to establish upper bounds on how well one can solve the close vector problem in this lattice A . Lower bounds can also be established by volumetric arguments, as done in [CDPR16]. In particular, they established that $\text{Vol}(A)^{1/(n/2-1)} \geq \Omega(\sqrt{m}/\log m)$. We provide the following better estimate.

Theorem 3. *For prime powers m , we have $\text{Vol}(A)^{\frac{1}{n/2-1}} \sim \sqrt{m}/2$.*

The proof is deferred to Appendix B.

2.4 Stickelberger Lattice

Let us define the Stickelberger lattice S as the $\mathbb{Z}[G]$ -multiples in $\mathbb{Z}[G]$ of the Stickelberger element

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

where $\{x\}$ denotes the fractional part $x - \lfloor x \rfloor$ of the rational number x . In other words, $S = \mathbb{Z}[G] \cap \theta \mathbb{Z}[G]$.

Theorem 4 ([Was12]). *The Stickelberger ideal S is such that for any fractional ideal \mathfrak{h} of \mathcal{O}_K , and for any $s \in S$, the ideal \mathfrak{h}^s is principal. In other words, S annihilates the ideal class group of K .*

Similarly to the log-unit lattice, we know a generating set of relatively short vectors (namely, the w_i 's defined below) of S . Let us define the vectors v_i , $2 \leq i \leq n + 1$ as $v_i = (a_i - \sigma_{a_i})\theta$. The aforementioned vectors w_i 's, $2 \leq i \leq n + 1$ are defined by $w_{i+1} = v_{i+1} - v_i$, and we have the following inequality on their norms from [CDW17, Wes18].

Fact 5 ([Wes18]). *For any $2 \leq i \leq n + 1$, we have $\|w_i\|_2 \leq 2\sqrt{n}$.*

In the case of prime conductors m , Schoof established in [Sch10] that all the w_i 's have ± 1 coefficients, in particular $\|w_i\|_2 = \sqrt{n}$.

3 Approx-SVP on Cyclotomic Ideals

3.1 Overview

Building upon [CGS14, EHKS14, CDPR16, BS16], the Approx-SVP algorithm for cyclotomic ideals of [CDW17] splits in the following 4 steps given below. A more detailed overview of these recent works is given in [Duc17]. Some details have been simplified by making use of several working hypotheses summarized at the end of this section.

Step 1 (quantum): Class-Group Discrete Logarithm. The first step consists in expressing the class $[\mathfrak{a}]$ of the input ideal \mathfrak{a} in base $\mathfrak{B} = \{\mathfrak{p}^\sigma \mid \sigma \in G\}$ for some prime ideal \mathfrak{p} , using the quantum poly-time algorithm of [BS16]. Under Hypothesis 7, such a decomposition always exists. This algorithm is heavily based on the quantum algorithm for the Hidden Subgroup Problem over \mathbb{R}^n from [EHKS14]. This provides an element $e \in \mathbb{Z}[G]$ such that $[\mathfrak{p}^e] = [\mathfrak{a}]$.

Step 2 (classical): Close Principal Multiple. The second step, introduced in [CDW17] consists in finding a *close principal multiple* of \mathfrak{a} , that is a principal ideal of the form $\mathfrak{b} = \mathfrak{a}\mathfrak{c}$ where $\mathfrak{c} \subset \mathcal{O}_K$ is an integral ideal of reasonably small norm $N\mathfrak{c} \leq F$. This will allow to focus the search of a short vector to the (principal) sublattice $\mathfrak{b} \subset \mathfrak{a}$.

This is done by finding a point $v \in S$ close to e . Setting $w = v - e$ gives a ‘small’ ideal $\mathfrak{c} = \mathfrak{p}^w$ such that $\mathfrak{b} = \mathfrak{a}\mathfrak{c}$ is principal. Indeed, $[\mathfrak{b}] = [\mathfrak{a}][\mathfrak{c}] = [\mathfrak{p}]^e[\mathfrak{p}]^{v-e} = [\mathfrak{p}]^v$, and $[\mathfrak{p}]^v = [\mathcal{O}_K]$ by Stickelberger’s Theorem.

Yet \mathfrak{c} is not necessarily integral as coefficients of $w \in \mathbb{Z}[G]$ can be negative. This is nevertheless easy to solve under Hypothesis 6, as it then holds that $[\mathfrak{p}^{-1}] = [\mathfrak{p}^\tau]$. This gives the desired $\mathfrak{b} \subset \mathcal{O}_K$ of bounded norm $N\mathfrak{b} \leq p^{\|w\|_1}$.

Using the NearestPlane algorithm and an explicit short basis of the augmented Stickelberger lattice $S' := S + (1 + \tau)$, it is shown in [CDW17] that one can find a close vector $v \in S$, at ℓ_1 -distance at most B_2

$$\|w\|_1 = \|v - e\|_1 \leq B_2 = O(n^{3/2}). \tag{1}$$

Assuming $N\mathfrak{p} = \text{poly}(n)$ leads to

$$N\mathfrak{b}/N\mathfrak{a} = (N\mathfrak{p})^{B_2} = \exp(\tilde{O}(n^{3/2})). \tag{2}$$

Step 3 (quantum): Principal Ideal Problem. The next step consists of solving the Principal Ideal Problem (PIP) on the principal ideal \mathfrak{b} , that is, finding a generator h of it: $h\mathcal{O}_K = \mathfrak{b}$. As for the Class-Group Discrete Logarithm Problem, there is a quantum poly-time algorithm [BS16] for this task.

Step 4 (classical): Short Generator Problem. The last step consists in finding a unit $u \in \mathcal{O}_K^\times$ such that $g = uh$ (which also generates \mathfrak{b}) has small norm. As in Step 2, this again can be rephrased as a close-vector problem, this time in the log-unit lattice $\Lambda = \text{Log } \mathcal{O}_K^\times$.

Using a randomized variant of the RoundOff Algorithm with the explicit short basis $\{\text{Log } b_j, i \in I\}$ of the log-unit lattice, it is shown [CDPR16, Theorem 6.3] that for any target $H = \text{Span}(\Lambda)$, one can find a logarithmic unit $l \in \Lambda$ at distance at most $B_4 = O(\sqrt{m \log m})$

$$\|l - t\|_\infty \leq B_4. \tag{3}$$

From any target $t \in H$. Setting t to be the orthogonal projection of $\text{Log } h$ onto H , and u such that $l = \text{Log } u$ leads to a short generator $h = gu$, of norm bounded by

$$\|h\|_\infty \leq (Ng)^{1/n} \cdot \exp(\|l - t\|_\infty) \leq (Ng)^{1/n} \cdot \exp(O(\sqrt{n \log n})). \tag{4}$$

Conclusion. In conclusion, we have found a vector $g \in \mathfrak{b} \subset \mathfrak{a}$ of norm at most:

$$\begin{aligned} \|h\|_2 &\leq \sqrt{n} \|h\|_\infty \leq \sqrt{n} \cdot (Ng)^{1/n} \cdot \exp(B_4) \\ &\leq \sqrt{n} \cdot (N\mathfrak{a})^{1/n} \cdot p^{B_2/n} \cdot \exp(B_4) \\ &\leq \text{Vol}(\mathfrak{a})^{1/n} \cdot \sqrt{n} \cdot \Delta_K^{-1/2n} \cdot p^{B_2/n} \cdot \exp(B_4) \\ &\leq \text{Vol}(\mathfrak{a})^{1/n} \cdot \exp(\tilde{O}(\sqrt{n})), \end{aligned}$$

that is, we have solved approx-SVP on the cyclotomic ideal \mathfrak{a} with an Hermite factor of $\eta = \exp(\tilde{O}(\sqrt{n}))$.

3.2 Slackness of the Bounds of Step 4

Note that the 4th step from [CDPR16] makes use of a non-tight bound. Indeed the exact length of h can be written as

$$\|h\|_2 = (Ng)^{1/n} \cdot \|\text{Exp}(l - t)\|_2,$$

and [CDPR16] simply considers

$$\|\text{Exp}(l - t)\|_2 \leq \sqrt{n} \cdot \exp(\|l - t\|_\infty).$$

For our concrete analysis, it is therefore more relevant to define the pseudo-norm $\|\cdot\|_l : H \rightarrow [0, +\infty)$

$$\|x\|_l := \ln(\|\text{Exp}(x)\|_2).$$

While it holds that $\|x\|_l \leq \|x\|_\infty + \ln(\sqrt{n})$, the slackness of this inequality is *not* only induced by the typical $\ell_2 - \ell_\infty$ slackness $\|x\|_\infty \leq \|x\|_2 \leq \sqrt{n} \|x\|_\infty$, in the sense that we can have $\|x\|_l \not\asymp \|x\|_\infty$.

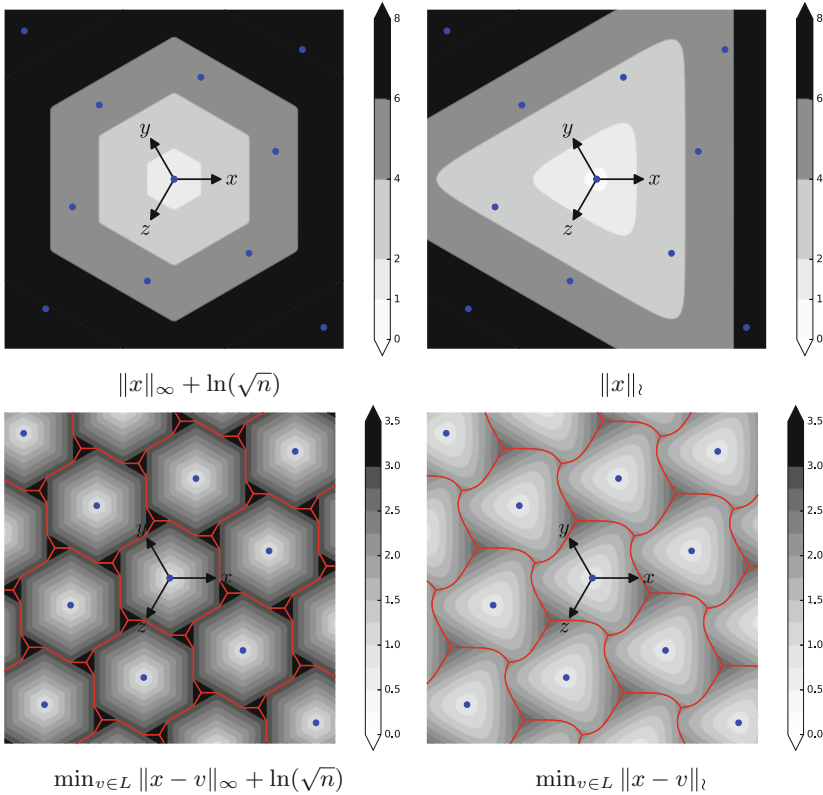


Fig. 1. Greyscale plots comparisons of $\|\cdot\|_\infty$ and $\|\cdot\|_l$ on the space H for $m = 7$ ($n = 6$, $\dim(H) = n/2 - 1 = 2$, $H \subset \mathbb{R}^{n/2}$). The black arrows represent the projection of the canonical axes of \mathbb{R}^3 onto H . The blue dots represent the points of the log-unit lattice $\Lambda = \text{Log } \mathcal{O}_K^\times$. The red cells represent Voronoi partitions. (Color figure online)

Indeed, negative coefficients in x contribute very little to $\|x\|_l$. This is exemplified by having a pathologically negative coefficient: taking $x = \alpha(1 - n/2, 1, \dots, 1) \in H$ where $\alpha > 0$ we have

$$\begin{aligned} \|x\|_\infty &= \alpha(n/2 - 1) \\ \|x\|_l &\leq \alpha + \ln(n). \end{aligned}$$

To represent things more pictorially, let us assume $m = 7$, for which $n = \phi(m) = 6$: the space H has dimension $n/2 - 1 = 2$ and is embedded in \mathbb{R}^3 : $H = \{(x, y, z) | x + y + z = 0\}$. A graphical comparison of $\|\cdot\|_\infty$ and $\|\cdot\|_l$ is provided in Fig. 1. As we can see, not only the $\|\cdot\|_\infty$ is pessimistic, but it can also lead to a wrong choice for the optimal solution: the Voronoi partitioning under $\|\cdot\|_\infty$ and $\|\cdot\|_l$ do differ.

3.3 Concrete Estimation of the Hermite Factor

At the time of writing, the authors do not possess a quantum computer sufficiently powerful to execute the full algorithm. Fortunately, it is nevertheless possible to simulate the behavior of the Hermite factor, since it depends only on the behavior of the classical steps 2 and 4. More precisely, assuming that $e \bmod S'$ and $t \bmod \mathcal{A}$ are uniform and independent (Hypothesis 8), we can study experimentally the average behavior of the whole algorithm.

More precisely, having introduced the appropriate pseudo-norm $\|\cdot\|_\lambda$, we can now write the exact value of the Hermite factor as a function of intermediate values $v - e$ and $l - t$ as follows:

$$\eta = \Delta_K^{-\frac{1}{2n}} \cdot \exp\left(\frac{\ln p}{n} \cdot \|v - e\|_1 + \|l - t\|_\lambda\right). \tag{5}$$

Therefore, we can predict the behavior of η simply by measuring experimentally the distribution of $\|v - e\|_1$ and $\|l - t\|_\lambda$. For comparison with LLL and BKZ, it is more convenient to consider the root Hermite factor $\delta = \eta^{1/n}$. For example, for LLL we have $\delta \approx 1.022$ according to [NS06], and for BKZ with blocksize $\beta \geq 50$, both heuristic arguments and experiments [CN11] give $\delta^{2(\beta-1)} \approx (\beta/(2\pi e))(\beta\pi)^{\frac{1}{\beta}}$.

3.4 Working Hypotheses

Restriction on the conductor. While the algorithm above has been initially studied for all prime-power conductors m in [CDW17, CDPR16], and even generalized to all conductors in [Wes18], the body of this paper will focus only on prime conductors m . This avoids numerous case by case discussions. One may prefer to directly study the case of power of 2 conductors, which is the most common in applications. However, powers of 2 are too sparse to allow for reasonable extrapolation. We therefore defer it to Appendix A, where we will compare it to the prime case.

Number-theoretic hypotheses. Two hypotheses are used in the works of [CDW17, CDPR16] concerning the structure of the class-group. The first is that the size of the *plus-part* of the class group h_m^+ (i.e. the size of the class group of the maximal real subfield) is only polynomial in the conductor m . The second is that one can construct (by random sampling) a small set of small-norm ideals that generate the class group as a $\mathbb{Z}[G]$ -module.

While these assumptions are sufficient for asymptotic results, they are not precise enough for a more effective study such as ours. We will therefore, as a working hypotheses strengthen those assumptions.

Hypothesis 6. *The plus-part of the class group is trivial, i.e. $h_m^+ = 1$.*

Hypothesis 7. *The class group is generated by the ideals above the smallest totally split prime. That is, let $p \in \mathbb{Z}$ be the smallest prime such that $p \equiv 1 \pmod m$, and $\mathfrak{p} \subset \mathcal{O}_K$ such that $N\mathfrak{p} = p$. We assume that $[\mathfrak{p}]$ generates Cl_K as a $\mathbb{Z}[G]$ -module, or equivalently that $\{[\mathfrak{p}^\sigma] \mid \sigma \in G\}$ generates Cl_K as a group.*

We will keep the notation p and \mathfrak{p} as a function of m for the rest of this paper. For our final conclusion, we will need estimates on p . We note that for prime conductors m , we necessarily have $p \geq 2m + 1$. On the other hand, prime density suggest that “on average” over m we have $p \approx m \ln m$.

Because of these strengthened hypotheses, our final claims should be interpreted as a best-case scenario for the efficiency of those algorithms. We remind that various computational results suggest that those assumptions are plausible for a substantial fraction of conductors m [Was12, Sch98, Sch03]. In any case, the failure of those two hypotheses would not invalidate our lower-bound.

Input distribution. In the light of the worst-case to average-case results of [Mic07, SSTX09, LPR10, SS11, PRSD17], it would be interesting to study the worst-case behavior of those algorithms. Alas, finding which input leads to the worst-case is most likely an intractable problem. We therefore instead assume that the inputs will be uniform modulo the respective lattices.

Hypothesis 8. *The input $e \in \mathbb{Z}[G]$ of step 2 is uniform in $\mathbb{Z}[G]/(S + (1 + \tau)\mathbb{Z}[G])$, and the target t in step 4 is uniform in H/Λ .*

Remark 2. The first part of the hypothesis essentially states that the class $[\mathfrak{a}]$ of the input ideal \mathfrak{a} is uniform over $\text{Cl} = \text{Cl}^-$. Interestingly, the main theorem of [JW18] allows to randomize the input so as to ensure its uniformity in the class group, by randomly multiplying it by a few small prime ideals. This only affects its norm by a factor $\exp(\tilde{O}(n))$, which asymptotically has a negligible impact on the final approximation factor. This implies that we can re-randomize any instance (even a worst-case one) to an average case one at a small cost.

The second part of the hypothesis can also be enforced by some randomization of t . A straightforward approach would be to simply add to t a (short) random vector r of H uniformly distributed in H/Λ . Reducing r with the good basis of Λ , this randomization has a limited impact on the final approximation factor. More precisely, we end up with $\|t - l\|_\lambda \leq \|t + r - l\|_\lambda + \|r\|_\lambda$ where both $\|r\|_\lambda$ and $\|t + r - l\|_\lambda$ follow the average case distribution studied in this paper (yet are not independent). In particular, if the average case gives a solution of length less than B with probability greater than $2/3$, we can find solutions of length at most $2B$ in the worst-case, by randomizing on average 3 times.

This loss of a factor 2 should only be read as a preliminary conclusion concerning the worst case. Indeed, heuristically, randomizing the input ideal for the first step will also rerandomize the target of the second step. Making such a statement formal requires generalizing [JW18] to the Arakelov class group; this is beyond the scope of the present article and left as future work.

4 Heuristic Improvement for the Close Vector Steps

In this section, we consider potential heuristic improvements for solving the close vector problems relatively to the log-unit lattice and to the Stickelberger lattice. Indeed, we note that [CDPR16, CDW17] focus on proving worst-case bounds,

and therefore apply simple and easy to analyse close-vector algorithms, namely NearestPlane and RoundOff. There are several reasons to think that this can be improved in practice, as discussed below.

4.1 More Short Vectors to Be Exploited

We note that the NearestPlane and RoundOff algorithms are restricted to use exactly d short vectors for a d -dimensional lattice, while in both cases, we actually know $\Theta(d^2)$ short vectors in these lattices. Indeed, for the log-unit lattice we know the following $n/2(n/2 - 1)$ short units:

$$\text{Log } b_{ij} = \text{Log} \left(\frac{1 - \zeta^i}{1 - \zeta^j} \right), i, j \in I, i \neq j.$$

Similarly, in the Stickelberger lattice, we know the following n^2 short class relations:

$$w_i \sigma, 2 \leq i \leq n + 1, \sigma \in G.$$

This extra knowledge can be exploited by using algorithms that can take advantage of many short vectors to solve CVP. In fact, if one knows the set V of all the *Voronoi relevant vector* of a lattice of dimension d , one can solve exact-CVP in $O(|V| \cdot \text{poly}(d))$ arithmetic operations [MV10,DB15]. This is described as Algorithm 3 (VoronoiCVP). Unfortunately the size of V can be as large as $2^d - 2$, and the best known algorithm [MV10] to determine it takes time $O(4^d)$. Yet it remains possible to run this algorithm with an approximation of the set V' ; this has been proposed and analyzed in [Laa16,DLdW19,SD19]. We cannot apply this analysis in our case because it uses heuristic arguments that are valid for random lattices, and those heuristics are most likely invalid for the lattices at hand which are somewhat close to orthogonal. Furthermore, this analysis is strongly restricted to the Euclidean norm, while we are here interested in other norms, or even pseudo-norms. But we can nevertheless apply a similar strategy and see how it behaves experimentally.

Algorithm 3. VoronoiCVP(V, t) ([MV10, Laa16, DLdW19])

```

c ← 0
while ∃v ∈ ±V such that ||t - c - v||2 < ||t - c||2 do
    c ← c + v
end while
return c

```

This algorithm can be viewed as a discrete gradient descent, and if V is indeed the set of Voronoi relevant vectors this descent will stop at an exact closest vector. Otherwise, the descent can get stuck in a discrete local minima, and therefore it was also proposed in [Laa16,DLdW19] to randomize the starting point $c = 0$ and to take the best results over several descents.

Rather than re-starting from scratch, in practice it seems preferable to continue the search nearby the current local minima: indeed the descent is done on a convex function, it is only because it is discretized that it can get stuck, and we expect the closest point to be not that far from the current point. Proceeding with such a strategy requires care to avoid looping over a cycle; this is easily prevented by keeping track of the points visited so far. At last, we also accelerate the descent by starting from either the `NearestPlane` or `RoundOff` approximation; this also ensures that its result will be at least as good as that of the original algorithm. The resulting algorithm is detailed in Algorithm 4 (`HeuristicCVP`), after the following final tweak.

4.2 Norm Inadequacy

Another source of inefficiency comes from the fact that `NearestPlane`, `RoundOff` and even the above `VoronoiCVP` are attempting to optimize the Euclidean distance, while for our application what we really want to optimize are the ℓ_1 -distance $\|\cdot\|_1$ for Stickelberger lattice, and the pseudo-norm $\|\cdot\|_\lambda$ for the log-unit lattice.

This inadequacy is easily addressed in practice simply by replacing the Euclidean norm $\|\cdot\|_2$ used in our `HeuristicCVP` algorithm by the desired (pseudo)-norm $\|\cdot\|_1$ or $\|\cdot\|_\lambda$.

Algorithm 4. `HeuristicCVP`($B, V, t, S, \|\cdot\|$)

```

 $c \leftarrow \text{NearestPlane}(B, t)$  or  $c \leftarrow \text{RoundOff}(B, t)$ 
 $C \leftarrow \{c\}$ 
for  $i \in \{1, \dots, S\}$  do
     $c \leftarrow \text{argmin}_{c'} \|c' - t\|$  where  $c'$  ranges over  $(c + V) \setminus C$ 
     $C \leftarrow C \cup \{c\}$ 
end for
return  $\text{argmin}_{c'} \|c' - t\|$  where  $c'$  ranges over  $C$ 

```

4.3 Dimension-Halving for Step 2

Because the Stickelberger ideal S is not full rank as a \mathbb{Z} -module in $\mathbb{Z}[G]$, the augmented ideal $S' = S + (1 + \tau)\mathbb{Z}[G]$ was introduced in [CDW17], which also annihilates the class group under the assumption that $h_m^+ = 1$. Alternatively, it is proposed in [Wes18] to instead project the lattice and the target down to the quotient ring $\mathbb{Z}[G]/(1 + \tau)$. More specifically, let $F \subset G$ be such that F and τF form a partition of G . We define a projection morphism

$$\begin{aligned}
 \pi : \mathbb{Z}[G] &\rightarrow \mathbb{Z}^F \\
 f \in F &\mapsto f \\
 f \in G \setminus F &\mapsto -\tau f
 \end{aligned}$$

where $\mathbb{Z}^F \simeq \mathbb{Z}^{n/2}$ is the \mathbb{Z} -module of formal integral sums of elements of F . We note that there is a reciprocal function $\hat{\pi}$ such that for all $x \in \mathbb{Z}^F$ it holds that $\pi(\hat{\pi}(x)) = x$, $\hat{\pi}(x)$ has positive coordinates, and $\|\hat{\pi}(x)\|_1 = \|x\|_1$: any $x \in \mathbb{Z}^F$ can be lifted back to a positive exponent in $\mathbb{Z}[G]$ with the same ℓ_1 norm, as needed to solve the Close Principal Multiple problem.

While this tweak from [Wes18] was originally mostly aesthetic as it didn't improve the asymptotic analysis, it effectively decreases the dimension of the problem from n to $n/2$; we note experimentally that this trick noticeably improved the average length of $\|v - e\|_1$.

Remark 3. We note during those experiments that the index $\text{Vol}(S') = |\mathbb{Z}[G]/S'|$ (or equivalently the index $\text{Vol}(\pi(S)) = |\mathbb{Z}^F/\pi(S)|$) is not equal to h_m^- , but rather to $2^{n/2-1} \cdot h_m^-$ (at least for all primes $m \leq 1000$): the representation of a class of Cl^- as an element of $\mathbb{Z}[G]/S'$ is not unique. And indeed, only a weaker statement is known, namely the theorem of Iwasawa [Sin80, Was12] stating that $|((1 - \tau)\mathbb{Z}[G])/((1 - \tau) \cap S)| = h_m^-$.

5 Implementation and Experiments

5.1 Implementation Details

Sources. Our implementation was realized in `python3`, and exploits the library `numpy`. It is provided in open-source for repeatability and review of our experiments at <https://github.com/lducas/Cyclotomic-QISVP-Effective>. The algorithms discussed above are implemented in `stickelberger.py` and `logunit.py`. The script `experiments.py` provides a convenient command line interface for running experiments. The script `verifications.py` provides sanity-checks, in particular with respect to the construction of the Stickelberger and log-unit lattices.

Optimizations. The critical computation regarding the performance of Algorithm 4 is the evaluation of the pseudo-norm $\|\cdot\|_l$ of $x + v$. In this loop, $x = c - t$ is fixed, while v varies over the set V , of size $\Theta(n^2)$.

Naively, the efficiency of evaluating the pseudo-norm is pretty terrible: not only does it requires $\Theta(n)$ calls to transcendental functions (`log`, `exp`), but it also requires to run the `for v in V` loop at the `python` level, inducing interpretation overheads. We note the following identity:

$$\exp(\|x + v\|_l^2) = \|\text{Exp}(x + v)\|^2 = \sum e^{2x_i} \cdot e^{2v_i} = \langle \text{Exp}(2x), \text{Exp}(2v) \rangle.$$

Since $y \mapsto \exp(y^2)$ is monotonic over $[0, +\infty)$, this means that we can actually determine the minimizing v using a matrix-vector product $M \cdot \text{Exp}(2x)$, where the rows of M are the row vectors $\{\text{Exp}(-2v)^T | v \in V\}$. Having precomputed M , this step becomes very fast thanks to the optimized linear algebra library included in `numpy`.

Another optimization consists in using a custom hash function \mathcal{H} for testing $c' \notin C$ in Algorithm 4. By making this function linear, we can accelerate the computation of $\mathcal{H}(c') = \mathcal{H}(c) + \mathcal{H}(v)$.

Numerical stability issues. In the experiment reported below, Fig. 2a has been truncated at dimension 800: after this point the behavior started being erratic. We strongly suspect that this is due to numerical stability issues during the Gram-Schmidt Orthogonalization algorithm. Unfortunately, increasing floating point precision seems difficult within our programming set-up, as `python/numpy` does not support more than double precision floats. Perhaps surprisingly, step 4 showed no such issue, at least up to dimension 3000. It may be that matrix inversion is more numerically stable than Gram-Schmidt, but also that the log-unit basis is better conditioned than the Stickelberger basis.

5.2 Experimental Results

We now report on the behavior of the original algorithms of [CDW17, CDPR16] and our heuristic improvements. Our experiments are depicted in Fig. 2. The data points are averaged over 100 samples per prime conductor m . For certain batches of experiments, we may have skipped some conductors so as to obtain data points for larger conductors in reasonable time. The computation ran for about a week, using 8 cores (Intel Xeon E5-2650v3 @2.3 GHz).

Deviation from average. Before commenting on the average behavior, we first note that, apart from the naive algorithms, the deviation from average was extremely small: the standard deviation is smaller than the average by a factor at least 20 for conductors $m \geq 200$, and the gap seems to grow further with the dimension. This may not entirely dismiss the possibility of rare outliers, but the bounds from Sect. 6 will control the probability of outliers.

Experimental effective asymptotics. Our first remark is that the upper bounds from [CDW17, CDPR16] $\|v - e\|_1 = O(n^{3/2})$ and $\|l - t\|_\lambda = O(\sqrt{n \log n})$ seem to be reached in practice, i.e., it seems very plausible that $\|v - e\|_1 = \Theta(n^{3/2})$ and quite plausible $\|l - t\|_\lambda = \Theta(\sqrt{n \log n})$. More precisely, for the original algorithms, for large ranks $n = \varphi(m)$ it seems to hold that:

$$\|v - e\|_1 \approx 0.039 \cdot n^{3/2}, \quad \text{and} \quad \|l - t\|_\lambda \approx 0.32 \cdot \sqrt{n \ln n}. \quad (6)$$

Our heuristically improved variant using `HeuristicCVP` with $n^{3/2}$ iterations yields

$$\|v - e\|_1 \approx 0.032 \cdot n^{3/2}, \quad \text{and} \quad \|l - t\|_\lambda \approx 0.117 \cdot \sqrt{n \ln n}. \quad (7)$$

Increased number of iterations for `HeuristicCVP`. Of course, one would ideally want to estimate how those constants evolve as the number of iterations for `HeuristicCVP` increases. However, such experiments become impractical as this number grows further than $n^{3/2}$.

From Fig. 2c, we note that increasing the number of iterations beyond n does not seem to provide significantly better solutions in the log-unit lattice. No such conclusion can be drawn for the Stickelberger lattice (Fig. 2a). Fortunately, the lower bound studied in Sect. 6.1, Fig. 3, will show that the solution found with $n^{3/2}$ iterations is already quite close to optimal.




Fig. 2. Average distance given by various CVP algorithms for steps 2 & 4, for prime conductors m .

6 Volumetric Lower Bounds

In this section, we provide probabilistic lower bounds using volumetric arguments. More specifically, we compute a lower bound $r := r(L, \|\cdot\|)$ for the covering radius of a given lattice L under a given (asymmetric) norm $\|\cdot\|$. The following proposition states that most points are at a distance almost r from L .

Proposition 1. *Let L be a full-rank lattice in a euclidean vector space V of dimension d , and let $\mathcal{B} = \{x \in V \mid \|x\| < 1\}$ be the open unit ball associated to an (asymmetric) norm $\|\cdot\|$. Let $r = (\text{Vol}(L)/\text{Vol}(\mathcal{B}))^{1/d}$.*

Then, for any $\alpha \in [0, 1]$, and for a random vector x such that $x \bmod L$ is uniformly distributed, the probability that

$$\|x - L\| := \min_{v \in L} \|x - v\| \leq \alpha r$$

is less than α^d . In particular, there exists a vector $x \in V$ such that $\min_{v \in L} \|x - v\| \geq r$.

Proof. We work over the torus V/L , whose total measure is $\text{Vol}(L)$. The probability that $\|x - L\| \leq \alpha r$ is given by

$$P = \frac{\text{Vol}(\alpha r \mathcal{B} \bmod L)}{\text{Vol}(L)}.$$

Note that $\text{Vol}(\alpha r \mathcal{B} \bmod L) \leq \text{Vol}(\alpha r \mathcal{B})$, with equality if and only if the union $\bigcup_{v \in L} v + \alpha r \mathcal{B}$ is disjoint. In particular

$$P \leq \alpha^d r^d \text{Vol}(\mathcal{B}) / \text{Vol}(L) = \alpha^d.$$

Remark 4. When comparing experimental results to those lower bounds, one should keep in mind that a gap does not necessarily imply that the algorithm fails to find the exact closest vector. Indeed, the above bound is tight only for lattices that are a perfect packing with respect to the considered balls.

For example consider \mathbb{Z}^n , for which CVP is easy to solve in any ℓ_p norm. It is a perfect packing for the ℓ_∞ distance, and we have $r(\mathbb{Z}^n, \|\cdot\|_\infty) = 1/2$, while the average ℓ_∞ distance of a point to \mathbb{Z}^n is $1/2 - o(1)$. Now, consider \mathbb{Z}^n for the ℓ_1 distance, which is far from a perfect packing. We have $r(\mathbb{Z}^n, \|\cdot\|_1) = (n!/2^n)^{1/n} \approx n/2e \approx 0.184 \cdot n$, yet the average ℓ_1 distance is $n/4 = 0.25 \cdot n$.

6.1 Volumetric Bound for Step 2

Before we proceed, we must first discuss whether we should apply the lower bound with or without the dimension halving trick, i.e., whether we should apply it to the augmented Stickelberger lattice $S' = S + (1 + \tau)\mathbb{Z}[G]$, or to the projected one $\pi(S)$. While both have the same volume, the dimension of S' is twice the dimension of $\pi(S)$, which would give a smaller lower bound.

Yet, we note that π can only decrease ℓ_1 distances, so a lower bound for $\pi(S)$ will also apply to S' .

We have that $\dim(\pi(S)) = n/2 =: d$ and $\text{Vol}(\pi(S)) = 2^{d-1}h_m^-$. The volume of the ℓ_1 unit ball in dimension d is given by $\text{Vol}(\mathcal{B}_1) = 2^d/d!$. We need an estimation of h_m^- . Let

$$G(m) = 2m(m/4\pi^2)^{\varphi(m)/4}.$$

Kummer claimed, without publishing a proof, that for any prime m we have $h_m^- \sim G(m)$. Although this is now believed to be unlikely, Lepistö [Lep74] proved a weaker (but sufficient here) explicit bound of the form

$$\left| \log \left(\frac{h_m^-}{G(m)} \right) \right| = O(\log(m)).$$

We deduce that $h_m^- = G(m)e^{O(\log(m))}$, and therefore $h_m^{-1/d} \sim G(m)^{1/d}$. Such an approximation is numerically satisfied up to a 1% error for primes $m \in [100, 2000]$ by our script `verification.py`. Using Stirling's formula, and the facts that $d = n/2 \sim m/2$ (since m is prime) and $(2m)^{1/d} \sim 1$, we conclude that

$$\begin{aligned} r(\pi(S), \|\cdot\|_1) &\sim \left(2^d \cdot m(m/4\pi^2)^{n/4} / (2e/d)^d \right)^{1/d} \\ &\sim (m/4\pi^2)^{1/2} \cdot (d/e) \\ &\sim \frac{1}{4e\pi} \cdot n^{3/2} \approx 0.02927 \cdot n^{3/2} \end{aligned}$$

Adjusting to the integral input setting. While these bounds hold asymptotically, we note that our experiments violate them for dimensions below 200. The reason is that in Step 2, the input is an integral vector, uniform in $\mathbb{Z}^F/\pi(S)$, and not uniform in $\mathbb{R}^F/\pi(S)$ as required by Proposition 1. However, we can rather easily adjust to this setting, by counting integral points $N_{d,b} = |b\mathcal{B}_1 \cap \mathbb{Z}^d|$ in the ball of radius b . Using dynamic programming, $N_{d,b}$ is easily computed in polynomial time thanks to the following recursion:

$$N_{d,0} = 1, \quad N_{1,b} = 2b + 1, \quad N_{d,b} = N_{d-1,b} + 2 \sum_{k=1}^b N_{d-1,b-k}.$$

For our concrete lower bound, we can therefore take r to be the largest integer such that $N_{d,r} \leq |\mathbb{Z}^F/\pi(S)|$. This is depicted in Fig. 3, and compared to the performance of our algorithm `HeuristicCVP`. We note (as expected) that the asymptotic behavior is similar to our continuous volumetric analysis above.




Fig. 3. Numerically computed volumetric lower bounds: maximal r such that $N_{d,r} \leq |\mathbb{Z}^F/\pi(S)|$, compared to the experimental behavior of HeuristicCVP.

Remark 5. In the above analysis, we have accounted for the factor $2^{n/2-1}$ that separates the lattice of (augmented) Stickelberger class relations from the full lattice of class relations (see Remark 3). While we are currently uncertain whether or not this factor is unavoidable, we note that its impact is asymptotically very simple: it contributes a factor 2 to our lower bound. Therefore, one may prefer the rely on a halved lower bound.

6.2 Volumetric Bound for Step 4

We start by noting that we cannot apply Proposition 1 directly to our pseudo-norm, the issue being the lack of homogeneity: $\{x \in H \mid \|x\|_l \leq r\} \neq r \cdot \{x \in H \mid \|x\|_l \leq 1\}$. Fortunately, there seems to be a reasonably close asymmetric norm $\|x\|_{+\infty} = \max_i x_i$ that can be used to bound the pseudo-norm $\|x\|_l$.³ Note that, on the space H , it differs from the usual ℓ_∞ norm by ignoring negative coefficients. For any $x \in H$, we have the inequalities

$$\|x\|_{+\infty} + \ln(\sqrt{n}) \geq \|x\|_l \geq \|x\|_{+\infty} + \ln(\sqrt{2}) \tag{8}$$

The asymmetric unit ball $\mathcal{B}_{+\infty}$ for the $\|\cdot\|_{+\infty}$ asymmetric norm is the $(d-1)$ -simplex whose d vertices that are a permutation of $(1, \dots, 1, 1-d)$. Its volume is given by $\text{Vol}(\mathcal{B}_{+\infty}) = d^{d-1/2}/(d-1)!$, and we have $\text{Vol}(\mathcal{B}_{+\infty})^{1/(d-1)} \rightarrow e$.

On the other hand, According to Theorem 3, the root volume of the log-unit lattice satisfies $\text{Vol}(\Lambda)^{1/(d-1)} \sim \sqrt{n}/2$ for any prime conductors m . Such an approximation is numerically satisfied up to a 1% error for primes

³ To verify that $\|\cdot\|_{+\infty}$ is indeed an asymmetric norm over H , we recall that vector space H is $\{x \in \mathbb{R}^d \mid \sum x_i = 0\}$: there is always one coordinate that is positive.

$m \in [100, 2000]$, as can be verified with the script `verification.py`. We therefore conclude that:

$$r(\Lambda, \|\cdot\|_{+\infty}) \sim \frac{\sqrt{n}}{2e} \approx 0.1839 \cdot \sqrt{n} \tag{9}$$

Remark 6. We note that our concrete lower bound is also asymptotically better than the one given in [CDPR16]. The reason is that it is based on Theorem 3 stating that $\text{Vol}(\Lambda)^{1/(d-1)} \sim \sqrt{n}/2$, while [CDPR16] relied on the inequality $\text{Vol}(\Lambda)^{1/(d-1)} \geq \Omega(\sqrt{n}/\log n)$. This $1/\log(n)$ factor comes from cumulating the approximation factors from Landau’s estimate for L -functions at 1 [Lan27] over all non-trivial character. Our Theorem 3 shows that Landau’s approximations essentially cancel out under geometric average over all characters.




Fig. 4. Numerically computed lower bound $r := r(\Lambda, \|\cdot\|_{+\infty}) + \ln(\sqrt{2})$, compared to the experimental behavior of HeuristicCVP.

7 Conclusion

7.1 Summary

In Table 1 we summarize the asymptotic behavior of the algorithms and lower bounds studied in the previous sections.

Recall from formula (5) that the Hermite factor is

$$\eta = \Delta_K^{-\frac{1}{2n}} \cdot \exp\left(\frac{\ln p}{n} \cdot \|v - e\|_1 + \|l - t\|_2\right),$$

where p is the smallest prime such that $p \equiv 1 \pmod m$. We can now predict the concrete Hermite factor of the quantum algorithms for Ideal-SVP.

Table 1. Asymptotic summary.

	Step 2 $\ v - e\ _1$	Step 4 $\ l - t\ _1$
Naive algorithms from [CDW17, CDPR16]	$0.039 \cdot n^{3/2}$	$0.32 \cdot \sqrt{n \ln n}$
HeuristicCVP with $n^{3/2}$ iterations	$0.032 \cdot n^{3/2}$	$0.117 \cdot \sqrt{n \ln n}$
Volumetric lower bound	$0.02927 \cdot n^{3/2}$	$0.1839 \cdot \sqrt{n}$
Halved volumetric lower bound (Remark 5)	$0.01463 \cdot n^{3/2}$	N/A




Fig. 5. Quality of Quantum Ideal-SVP vs. LLL and BKZ.

7.2 Comparison with Classical Algorithms

We now compare our prediction to the classical algorithms LLL and BKZ. For this comparison, we will consider the smallest possible value for $p = 2m + 1$ and the expected value $p = m \ln m$ derived from prime density. This comparison is provided in Fig. 5, using the root Hermite factor $\delta = \eta^{1/n}$.

We provide the reference root Hermite factors for LLL and BKZ with block-sizes $\beta \in \{80, 120, 160, 300\}$. The LLL algorithm is the cheapest lattice reduction available, and it should be noted that the quantum steps [EHKS14] 1 and 3

also make several quantum calls to LLL: the computational cost of the quantum algorithm is therefore bounded below by the cost of LLL.⁴ The cost of BKZ grows exponentially or even super-exponentially with β , depending on the choice of algorithm. Nevertheless, BKZ-80 remains a reasonably easy computation (say, about $8m$ core-minutes), while BKZ-120 is to be considered doable ($8m$ core-days). Running BKZ-160 is on the borderline of feasible: to this date, computational records almost correspond to one out of the $\approx 8m$ steps of such a lattice reduction [ADH+19, SG10]. Finally, BKZ-300 is roughly what is required to break the weakest lattice-based candidates to the NIST post-quantum standardization [ACD+18].

7.3 Conclusion

Our first conclusion is that the naive version of the quantum algorithm is not relevant for rings of ranks considered practical for use in cryptography, as it does not outperform classically feasible computation (BKZ-120) before prime conductor $m \approx 32000$. Nevertheless, our heuristic improvements allow to decrease this cross-over point down to $m \approx 6000$. Such a dimension is still one order of magnitude larger than what is used for NIST post-quantum standardization candidates, but is within the range of what is used by certain concrete Fully Homomorphic Encryption schemes, for example [HS15].

Finally, one may fear that further tricks could improve the heuristic CVP steps within [CDPR16, CDW17], and maybe reach the lower bound.⁵ The conclusion is somewhat reassuring for NIST candidates, as the cross-over point with BKZ-300 should not happen before ring rank $n \approx 6000$, even given a perfect CVP oracle for the Log-unit lattice and the Stickelberger lattice: NIST candidates use cyclotomic rings of rank at most $n = 1024$.

While the body of this article is focused on prime conductors m , we also considered the powers of 2 conductors, and found that both the experimental behavior and the numerical lower bounds were slightly worse in the powers of 2 case. This is reported in Appendix A.

7.4 Limitations

To avoid any over-interpretation of our results, we summarize here the limits of what can be concluded from the present work.

⁴ Unfortunately, while proved polynomial time, the algorithms of [EHKS14, BS16] have, to our knowledge not been the subject of refined complexity analysis. But already, one can note that the lower bound we suggest is far from tight, considering the overheads of running LLL quantumly rather than classically, and this, many times.

⁵ We recall that this bound is plausibly not tight, that is, even a perfect CVP oracle may not be able to reach it; see Remark 4.

Limitation of the lower bounds. We first remind that this lower-bound is only probabilistic, i.e., Proposition 1 states that the probability that a target falls closer to the lattice by a factor $\alpha < 1$ is at most α^d . That is, it may not be impossible to rerandomize the input to bruteforce a better solution, but it will raise the cost of the algorithm to exponential time.

Moreover, it should be noted that these lower bounds apply only to algorithms that are slight variations of [CDPR16, CDW17]. It has been proved that ideas beyond this framework make it asymptotically possible to go below those lower bounds [PMHS19], but at the cost of a sub-exponential running time, together with an exponential amount of precomputation.

Limitation of the cryptanalytic impact. On the other hand, we also remind the reader that we have made several working assumptions for the sake of simplicity, putting ourselves in the most favorable set-up. In particular, if one were to need not 1 but 2 ideals to generate the class group, this would asymptotically double the constant for Step 2.

Most importantly, we also recall that this work only studies the concreteness of the first obstacle discussed in our introduction, while the second obstacle remains unsolved. That is, these results concern only Ideal-SVP, and it remains unclear how they could be generalized to Ring-SIS, Ring-LWE, or NTRU.

A The Power of 2 Case

In this section we compare the power of 2 case to the prime case. The experimental behavior and lower bounds for step 2 and step 4 are given in Fig. 6. We see that the asymptotic lower bounds for the power of 2 case is similar to the prime case, yet for both step 2 and 4, the experimental behavior is slightly worse for the power of 2 case.

We also need to account for the inverse root discriminant, which is also a factor in final Hermite factor η given by Formula (5). A quick calculation shows that this factor is a similar function of the rank n in both cases. Indeed, when m is prime, the inverse root discriminant $|\Delta_K|^{-1/2n}$ appearing in the formula for the root Hermite factor (5) is given by

$$|\Delta_K|^{-1/2n} = m^{-(n-1)/2n} \sim 1/\sqrt{m} \sim 1/\sqrt{n}.$$

On the other hand for $m = 2^k$ we have

$$|\Delta_K|^{-1/2n} = 2^{-n(k-1)/2n} = 2^{(1-k)/2} = \sqrt{2/m} = 1/\sqrt{n}.$$

In conclusion, we expect that the quantum algorithm for Ideal-SVP at hand provides vectors slightly longer for power of 2 conductors than for prime conductors.

B Estimation of the Regulator

In this appendix we prove Theorem 3, which states that for any prime power $m = p^k$, we have $(\text{Vol}(\Lambda)/h^+)^{\frac{1}{n/2-1}} \sim \sqrt{m}/2$. First, we recall that the volume of

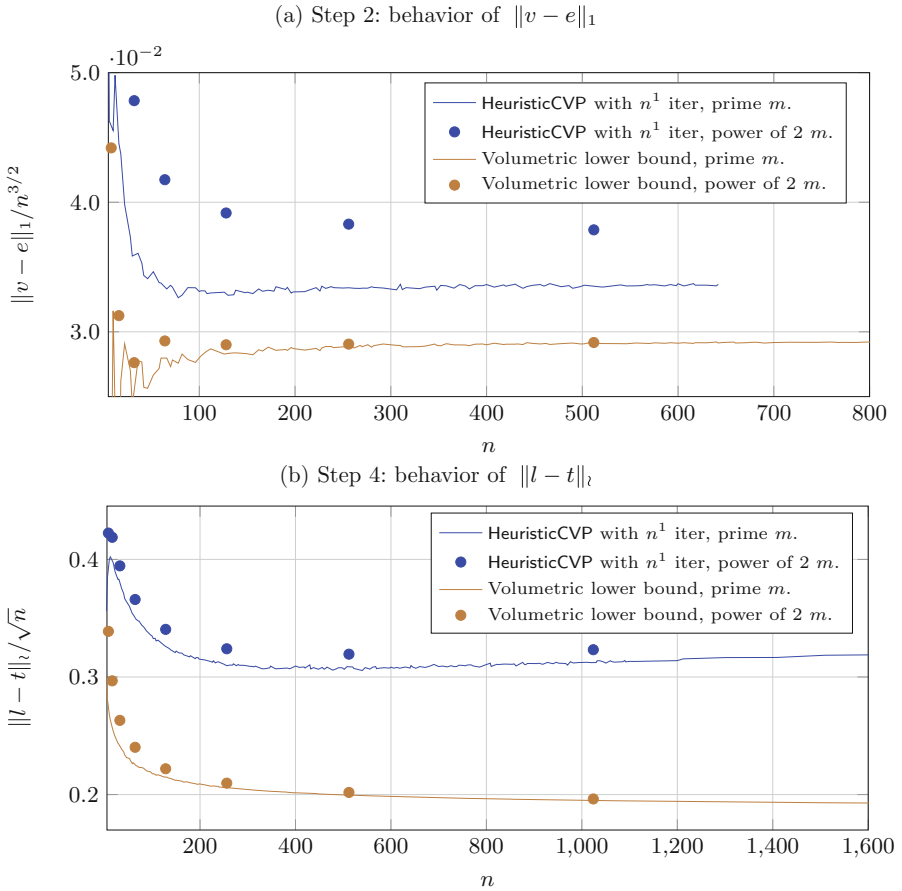


Fig. 6. Comparison of the prime conductors and power of 2 conductor.

the log-unit lattice is related to the so-called *regulator* R of K by the formula⁶

$$\text{Vol}(\Lambda) = \frac{R\sqrt{n/2}}{2^{n/2-1}}.$$

Therefore $\text{Vol}(\Lambda)^{\frac{1}{n/2-1}} \sim R^{\frac{1}{n/2-1}}/2$, and it remains to estimate Rh^+ . Let Δ_{K^+} denote the discriminant of K^+ , the maximal real subfield of K . We have that $|\Delta_{K^+}| = |\Delta_K/p|^{1/2}$ when m is a power of $p \neq 2$ (for $p = 2$, the following results should adjust for the fact that $|\Delta_{K^+}| = |\Delta_K/4|^{1/2}$). From [Was12, p. 42],

⁶ The denominator $2^{n/2-1}$ may not be standard in the literature, and is due to our definition of the logarithmic embedding. Indeed since the field at hand is totally complex, we only use one embedding from each pair of conjugate embeddings.

we get

$$Rh^+ = |\Delta_K/p|^{1/4} \prod_{\chi \neq 1 \text{ even}} L(1, \chi),$$

where the product is over all non-trivial even Dirichlet characters modulo m . We have

$$\begin{aligned} \log \left(\prod_{\chi \neq 1 \text{ even}} L(1, \chi) \right) &= - \sum_{\chi} \sum_q \log \left(1 - \frac{\chi(q)}{q} \right) \\ &= \sum_{\chi} \sum_q \sum_{i=1}^{\infty} \frac{\chi(q^i)}{iq^i} \\ &= \sum_q \sum_{i=1}^{\infty} \frac{1}{iq^i} \sum_{\chi} \chi(q^i). \end{aligned}$$

Since

$$\sum_{\chi} \chi(a) = \begin{cases} n/2 - 1 & \text{if } a \equiv \pm 1 \pmod{m}, \\ -1 & \text{otherwise,} \end{cases}$$

we deduce that

$$\log \left(\prod_{\chi \neq 1 \text{ even}} L(1, \chi) \right) = \lim_{x \rightarrow \infty} \left(\frac{n-2}{2} \sum_{\substack{q^i \leq x \\ q^i \equiv \pm 1 \pmod{m}}} \frac{1}{iq^i} - \sum_{\substack{q^i \leq x \\ q^i \not\equiv \pm 1 \pmod{m}}} \frac{1}{iq^i} \right).$$

Let us first deal with the terms where $i = 1$. From [Pom77], for any a such that $(a, m) = 1$, we have

$$\sum_{\substack{q \leq x \\ q \equiv a \pmod{m}}} \frac{1}{q} = \frac{\log \log(x)}{n} + \frac{1}{P(m, a)} + O \left(\frac{\log(m)}{n} \right),$$

where $P(m, a)$ is the first prime q such that $q \equiv a \pmod{m}$. We get

$$\begin{aligned} &\lim_{x \rightarrow \infty} \left(\frac{n-2}{2} \sum_{\substack{q \leq x \\ q \equiv \pm 1 \pmod{m}}} \frac{1}{q} - \sum_{\substack{q \leq x \\ q \not\equiv \pm 1 \pmod{m}}} \frac{1}{q} \right) \\ &= \frac{n-2}{2P(m, 1)} + \frac{n-2}{2P(m, -1)} - \sum_{\substack{a \in \{2, \dots, m-2\} \\ (a, m) = 1}} \frac{1}{P(m, a)} + O(\log(m)) \\ &= O(\log(m)). \end{aligned}$$

For the terms where $i \geq 2$, we have from [AC49] that

$$\sum_{i \geq 2} \sum_{\substack{q^i \leq x \\ q^i \equiv \pm 1 \pmod{m}}} \frac{1}{iq^i} = O(1/m).$$

The proof in [AC49] is given for m prime, but is easily adapted to powers of primes. We deduce that

$$\log \left(\prod_{\chi \neq 1 \text{ even}} L(1, \chi) \right) = O(\log(m)).$$

We get the estimate

$$(Rh^+)^{\frac{1}{n/2-1}} = p^{\frac{p^{k-1}(pk-k-1)-1}{2(n-2)}} e^{O(\frac{\log(m)}{n})} = m^{\frac{1}{2}+o(1)},$$

from which we conclude that $(\text{Vol}(A)/h^+)^{\frac{1}{n/2-1}} \sim \sqrt{m}/2$.

References

- [AC49] Ankeny, N.C., Chowla, S.: The class number of the cyclotomic field. Proc. Natl. Acad. Sci. **35**(9), 529–532 (1949)
- [ACD+18] Albrecht, M.R., et al.: Estimate all the {LWE, NTRU} schemes! (2018). https://doi.org/10.1007/978-3-319-98113-0_19
- [ADH+19] Albrecht, M.R., Ducas, L., Herold, G., Kirshanova, E., Postlethwaite, E.W., Stevens, M.: The general sieve kernel and new records in lattice reduction. Cryptology ePrint Archive, Report 2019/089 (2019). <https://eprint.iacr.org/2019/089>, https://doi.org/10.1007/978-3-030-17656-3_25
- [Ajt99] Ajtai, M.: Generating hard instances of the short basis problem. In: Wiederemann, J., van Emde Boas, P., Nielsen, M. (eds.) ICALP 1999. LNCS, vol. 1644, pp. 1–9. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48523-6_1
- [Bab86] Babai, L.: On Lovász' lattice reduction and the nearest lattice point problem. Combinatorica **6**(1), 1–13 (1986). Preliminary version in STACS 1985
- [BCLvV17] Bernstein, D.J., Chuengsatiansup, C., Lange, T., van Vredendaal, C.: NTRU prime: reducing attack surface at low cost. In: Adams, C., Camenisch, J. (eds.) SAC 2017. LNCS, vol. 10719, pp. 235–260. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-72565-9_12
- [BS16] Biasse, J.-F., Song, F.: 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, pp. 893–902. SIAM (2016)
- [CDPR16] Cramer, R., Ducas, L., Peikert, C., Regev, O.: Recovering short generators of principal ideals in cyclotomic rings. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9666, pp. 559–585. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_20
- [CDW17] Cramer, R., Ducas, L., Wesolowski, B.: Short Stickelberger class relations and application to ideal-SVP. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10210, pp. 324–348. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56620-7_12

- [CGS14] Campbell, P., Groves, M., Shepherd, D.: Soliloquy: a cautionary tale. In: ETSI 2nd Quantum-Safe Crypto Workshop (2014). http://docbox.etsi.org/Workshop/2014/201410_CRYPTOS07_Systems_and_Attacks/S07_Groves_Annex.pdf
- [CN11] Chen, Y., Nguyen, P.Q.: BKZ 2.0: better lattice security estimates. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 1–20. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25385-0_1
- [DB15] Dadush, D., Bonifas, N.: Short paths on the Voronoi graph and closest vector problem with preprocessing. In: Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 295–314. Society for Industrial and Applied Mathematics (2015)
- [DLdW19] Doulgerakis, E., Laarhoven, T., de Weger, B.: Finding closest lattice vectors using approximate Voronoi cells. In: Ding, J., Steinwandt, R. (eds.) PQCrypto 2019. LNCS, vol. 11505, pp. 3–22. Springer, Cham (2019). <https://eprint.iacr.org/2016/888>
- [Duc17] Ducas, L.: Advances on quantum cryptanalysis of ideal lattices. *Nieuw Archief voor Wiskunde* **5**, 184–189 (2017)
- [EHKS14] Eisenträger, K., Hallgren, S., Kitaev, A., Song, F.: A quantum algorithm for computing the unit group of an arbitrary degree number field. In: STOC, pp. 293–302. ACM (2014)
- [HPS98] Hoffstein, J., Pipher, J., Silverman, J.H.: NTRU: a ring-based public key cryptosystem. In: Buhler, J.P. (ed.) ANTS 1998. LNCS, vol. 1423, pp. 267–288. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0054868>
- [HS15] Halevi, S., Shoup, V.: Bootstrapping for `HElib`. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 641–670. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46800-5_25
- [JW18] Jetchev, D., Wesolowski, B.P.C.: Horizontal isogeny graphs of ordinary abelian varieties and the discrete logarithm problem. *Acta Arithmetica* (2018, in press)
- [Laa16] Laarhoven, T.: Finding closest lattice vectors using approximate Voronoi cells. In: Published at SAC 2016 (2016). <https://eprint.iacr.org/2016/888/20161219:141310>
- [Lan27] Landau, E.: Über Dirichletsche Reihen mit komplexen Charakteren. *Journal für die reine und angewandte Mathematik* **157**, 26–32 (1927)
- [Len75] Lenstra Jr., H.W.: Euclid’s algorithm in cyclotomic fields. *J. London Math. Soc.* **10**, 457–465 (1975)
- [Lep74] Lepistö, T.: On the growth of the first factor of the class number of the prime cyclotomic field. *Ann. Acad. Sci. Fenn. Ser.* **577**(Ser. A I), 1–21 (1974)
- [LLL82] Lenstra, A.K., Lenstra Jr., H.W., Lovász, L.: Factoring polynomials with rational coefficients. *Math. Ann.* **261**(4), 515–534 (1982)
- [LPR10] Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 1–23. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5_1
- [LPR13] Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. *J. ACM* **60**(6), 43:1–43:35 (2013). Preliminary version in Eurocrypt 2010

- [Mic07] Micciancio, D.: Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Comput. Complex.* **16**(4), 365–411 (2007). Preliminary version in FOCS 2002
- [MV10] Micciancio, D., Voulgaris, P.: A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. In: STOC, pp. 351–358 (2010)
- [NS06] Nguyen, P.Q., Stehlé, D.: LLL on the average. In: Hess, F., Pauli, S., Pohst, M. (eds.) ANTS 2006. LNCS, vol. 4076, pp. 238–256. Springer, Heidelberg (2006). https://doi.org/10.1007/11792086_18
- [PMHS19] Pellet-Mary, A., Hanrot, G., Stehlé, D.: Approx-SVP in ideal lattices with pre-processing. *Cryptology ePrint Archive*, Report 2019/215 (2019). <https://eprint.iacr.org/2019/215>. To appear at EUROCRYPT 2019. https://doi.org/10.1007/978-3-030-17656-3_24
- [Pom77] Pomerance, C.: On the distribution of amicable numbers. *J. Reine Angew. Math.* **293**(294), 217–222 (1977)
- [PRSD17] Peikert, C., Regev, O., Stephens-Davidowitz, N.: Pseudorandomness of ring-LWE for any ring and modulus. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 461–473. ACM (2017)
- [Reg09] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. *J. ACM* **56**(6), 1–40 (2009). Preliminary version in STOC 2005
- [Sch98] Schoof, R.: Minus class groups of the fields of the l -th roots of unity. *Math. Comput. Am. Math. Soc.* **67**(223), 1225–1245 (1998)
- [Sch03] Schoof, R.: Class numbers of real cyclotomic fields of prime conductor. *Math. Comput.* **72**(242), 913–937 (2003)
- [Sch10] Schoof, R.: Catalan’s Conjecture. Springer, London (2010). <https://doi.org/10.1007/978-1-84800-185-5>
- [SD19] Stephens-Davidowitz, N.: A time-distance trade-off for GDD with preprocessing—instantiating the DLW heuristic. *arXiv* (2019). <https://arxiv.org/abs/1902.08340>
- [SG10] Schneider, M., Gama, N.: Darmstadt SVP challenges (2010). <https://www.latticechallenge.org/svp-challenge/index.php>. Accessed 02 Feb 2019
- [Sin80] Sinnott, W.: On the Stickelberger ideal and the circular units of an abelian field. *Inventiones Math.* **62**, 181–234 (1980)
- [SS11] Stehlé, D., Steinfeld, R.: Making NTRU as secure as worst-case problems over ideal lattices. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 27–47. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20465-4_4
- [SSTX09] Stehlé, D., Steinfeld, R., Tanaka, K., Xagawa, K.: Efficient public key encryption based on ideal lattices. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 617–635. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10366-7_36
- [SV10] Smart, N.P., Vercauteren, F.: Fully homomorphic encryption with relatively small key and ciphertext sizes. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 420–443. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13013-7_25
- [Was12] Washington, L.C.: *Introduction to Cyclotomic Fields*, vol. 83. Springer, New York (2012)
- [Wes18] Wesolowski, B.P.C.: *Arithmetic and geometric structures in cryptography*. Ph.D. thesis, EPFL (2018)