

Recovering Short Generators of Principal Ideals in Cyclotomic Rings

Ronald Cramer* Léo Ducas† Chris Peikert‡ Oded Regev§

April 6, 2015¶

Abstract

A handful of recent cryptographic proposals rely on the conjectured hardness of the following problem in the ring of integers of a cyclotomic number field: given a basis of an ideal that is guaranteed to have a “rather short” generator, find such a generator. In the past year, Bernstein and Campbell-Groves-Shepherd have sketched potential attacks against this problem. Most notably, the latter authors claimed a *quantum polynomial-time* algorithm (alternatively, replacing the quantum component with an algorithm of Biasse and Fieker would yield a *classical subexponential-time* algorithm). A key claim of Campbell *et al.* is that one step of their algorithm—namely, decoding the *log-unit* lattice of the ring to recover a short generator from an arbitrary one—is efficient (whereas the standard approach takes exponential time). However, very few convincing details were provided to substantiate this claim, and as a result it has met with some skepticism.

In this work, we remedy the situation by giving a rigorous theoretical and practical confirmation that the log-unit lattice is indeed efficiently decodable, in cyclotomics of prime-power index. The proof consists of two main technical contributions: the first is a geometrical analysis, using tools from analytic number theory, of the canonical generators of the group of cyclotomic units. The second shows that for a wide class of typical distributions of the short generator, a standard lattice-decoding algorithm can recover it, given any generator.

1 Introduction

In recent years, *lattices* have emerged as an attractive foundation for cryptography. The most efficient (and potentially practical) lattice-based cryptosystems are related to *ideal lattices*, which correspond to ideals in

*Cryptology Group, CWI, Amsterdam, The Netherlands & Mathematical Institute, Leiden University, The Netherlands. Email: cramer@cwil.nl, cramer@math.leidenuniv.nl

†Cryptology Group, CWI, Amsterdam, The Netherlands. Supported by an NWO Free Competition Grant. Email: ducas@cwil.nl

‡School of Computer Science, College of Computing, Georgia Institute of Technology. This material is based upon work supported by the National Science Foundation under CAREER Award CCF-1054495, by DARPA under agreement number FA8750-11-C-0096, and by the Alfred P. Sloan Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation, DARPA or the U.S. Government, or the Sloan Foundation. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

§Courant Institute of Mathematical Sciences, New York University. Supported by the Simons Collaboration on Algorithms and Geometry and by the National Science Foundation (NSF) under Grant No. CCF-1320188. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

¶This is a slight editorial revision of [CDPR15].

certain families of rings, e.g., $\mathbb{Z}[X]/(X^{2^k} + 1)$. Representative works include [HPS98, Mic02, LMPR08, Gen09, LPR10].

In recent years, several cryptographic constructions have relied directly on *principal ideals* that have “*relatively short*” generators, which serve as secret keys.¹ These include a variant of Gentry’s original fully homomorphic encryption scheme [Gen09] due to Smart and Vercauteren [SV10], the closely related Soliloquy encryption scheme [CGS14], and candidate cryptographic multilinear maps [GGH13, LSS14]. Breaking these systems is no harder than solving the following problem, which we call the *Short Generator of a Principal Ideal Problem* (SG-PIP): given some \mathbb{Z} -basis of an ideal that is guaranteed to have a “short” generator g , find a sufficiently short one (not necessarily g itself).

In the past year, warnings about SG-PIP in certain rings have been issued by Bernstein [Ber14a] and Campbell, Groves, and Shepherd [CGS14], who sketched potential attacks. The basic structure of the attacks, which appears to be folklore in computational number theory, consists of two main parts:

- First, given a \mathbb{Z} -basis of the principal ideal, find some arbitrary (not necessarily short) generator of the ideal. For this task, which is known as the *Principal Ideal Problem* (PIP), the state of the art is an algorithm of Biassé and Fieker [BF14, Bia14], whose runtime has only a subexponential $2^{n^{2/3+\epsilon}}$ dependence on n , the degree of the ring (over \mathbb{Z}). In addition, building on the recent work of Eisenträger *et al.* [EHKS14], polynomial-time *quantum* algorithms for PIP have recently been described in two independent works [CGS14, BS15], the latter of which provides a fully rigorous treatment. (In this work, we will not be further concerned with algorithms for PIP.)
- Second, transform the generator found in the previous phase into a *short* generator, thereby recovering the secret key, or its functional equivalent. The standard approach casts this task as a *closest vector problem* (CVP) on the Dirichlet “log-unit” lattice. In general, the fastest known algorithm for CVP (even allowing quantum) runs in exponential $2^{\Omega(n)}$ time [MV10], or in less time but with much weaker guarantees on the solution quality (e.g., [LLL82, Bab85, Sch87]). Note that it is not obvious *a priori* whether this approach yields a sufficiently short generator; much depends on the geometry of the log-unit lattice and the quality of the CVP solution.

Recently, a more refined view of the second phase has emerged, centered around the kinds of rings and short generators typically suggested for cryptographic applications. In [Ber14a], Bernstein suggests an approach that may yield slightly subexponential runtimes in *cyclotomic* rings of *highly smooth* index (e.g., $m = 3 \cdot 5 \cdot 7 \cdot 11$). In addition, several researchers have noted that the CVP instances arising in the second phase have some implicit structure. For example, [CGS14, Ber14b] observe that the existence of a “rather short” generator (by choice of the secret key) implies that the target point is “somewhat close” to the log-unit lattice; CVP with such a distance guarantee is more commonly known as *bounded-distance decoding* (BDD). Previously, Garg, Gentry and Halevi [GGH13] gave an improved variant of the Gentry-Szydlo algorithm [GS02] which shows that in cyclotomic rings having power-of-two index, BDD on the log-unit lattice is efficiently solvable to within sub-polynomial $n^{-\log \log n}$ distance. However, this threshold is much too small to handle the BDD instances arising in cryptosystems.

Notably, the work of Campbell, Groves and Shepherd [CGS14] contains a surprising claim: that in cyclotomic rings having power-of-two index, the second phase described above is *easy*, simply by decoding the log-unit lattice using a standard algorithm such as LLL [LLL82]. The stated reason is that the secret generator corresponds to a vector that is short relative to the determinant of the log-unit lattice. Unfortunately, this observation is not sufficient to justify their claim: in addition, one needs to show how to efficiently compute a basis

¹A principal ideal in a commutative ring R is of the form $gR = \{g \cdot r : r \in R\}$ for some $g \in R$, called a *generator* of the ideal.

of good enough geometric “quality” for decoding to succeed. Nevertheless, experiments in cryptographically relevant choices of dimension have shown that decoding is indeed practically efficient [She14, Sch15].

Contribution. In this work, we give a rigorous theoretical and practical confirmation of the above-described claim from [CGS14], and generalize to all cyclotomics of prime-power index. By using a particular efficiently computable basis, we prove that the second phase succeeds with high probability for typical distributions of the secret short generator.

In more detail, we provide two main technical contributions. First, we use standard tools from analytical number theory, such as bounds on *Dirichlet L-series*, to elucidate the geometry of a canonical set of generators for the group of *cyclotomic units*. (The cyclotomic units correspond either to the log-unit lattice itself, or to a sublattice whose index is conjectured to be quite small, which is sufficient for our purposes.) Then we show that for a wide class of typical distributions of the secret generator—e.g., Gaussian-like distributions—a standard, efficient lattice-decoding algorithm [Bab85] recovers the secret short generator, given any generator of the ideal. Somewhat counterintuitively, the *variance* of the distribution is essentially irrelevant, because it is implicitly normalized by working with logarithmic embeddings. Finally, to complement these results, we also give concrete numerical data demonstrating that this part of the attack succeeds for all practical choices of dimension.

Discussion. Combining our results with known algorithms for PIP (which are the bottleneck in the full attack) [BF14, Bia14, CGS14, BS15], one obtains quantum polynomial-time, or classical $2^{n^{2/3+\epsilon}}$ -time, key-recovery algorithms for the cryptographic constructions of [SV10, GGH13, LSS14, CGS14].² An important open problem is to obtain even faster (classical) PIP algorithms, perhaps also using the guarantee that a short generator exists.

The overall attack is quite specialized to the specific combination of *principal* ideals having “*rather short*” generators, in *cyclotomic* number fields. In particular, it does not seem to apply to the approximate Shortest Vector Problem (SVP) on arbitrary ideals. The conjectured hardness of approx-SVP is the foundation of the ring-LWE problem [LPR10], which in turn is the heart of many ideal-lattice-based cryptosystems. The attack fails on approx-SVP because most ideals in cyclotomic rings are not principal, and moreover, most principal ideals do not have short generators (as compared with their shortest nonzero elements). An interesting and important question is whether these barriers can be overcome to attack SVP on arbitrary ideals, or ring-LWE. In a complementary direction, another interesting question is whether the attack on SG-PIP can be extended to other families of *non-cyclotomic* rings, such as those suggested in [Ber14a]. For this it may suffice to find (by analysis, computation, or both) a suitably good basis of the log-unit lattice, or of a sublattice of not too large index.

Acknowledgments. We thank Dan Bernstein, Jean-François Biasse, Sorina Ionica, Dimitar Jetchev, and Paul Kirchner for many insightful conversations on topics related to this work. We also especially thank Dan Shepherd [She14] for explaining many additional details about the claims made in [CGS14], and for sharing other helpful observations.

²Strictly speaking, these runtimes depend on number-theoretic conjectures regarding the class numbers $h^+(m)$; see Section 2.4 for details.

2 Preliminaries

We denote column vectors by lower-case bold letters (e.g., \mathbf{x}) and matrices by upper-case bold letters (e.g., \mathbf{X}). We often adopt the nonstandard, but very useful, convention of indexing rows and columns by particular finite sets (not necessarily $\{1, \dots, n\}$), and identify a matrix with its indexed set of column vectors. The canonical scalar product over \mathbb{R}^n and over \mathbb{C}^n is denoted $\langle \cdot, \cdot \rangle$, and $\|\cdot\|$ denotes the Euclidean norm. For a complex number $z \in \mathbb{C}$, \bar{z} denotes its complex conjugate, and $|z| = \sqrt{z \cdot \bar{z}}$ denotes its magnitude.

2.1 Lattices and BDD

A *lattice* \mathcal{L} is a discrete additive subgroup of \mathbb{R}^n for some positive integer n . The *minimum distance* of \mathcal{L} is $\lambda_1(\mathcal{L}) := \min_{\mathbf{v} \in \mathcal{L} \setminus \{0\}} \|\mathbf{v}\|$, the length of a shortest nonzero lattice vector. Every lattice is generated as the integer linear combinations of some (non-unique) \mathbb{R} -linearly independent *basis* vectors $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_k\}$, as $\mathcal{L} = \mathcal{L}(\mathbf{B}) := \{\sum_{j=1}^k \mathbb{Z} \cdot \mathbf{b}_j\}$, where $k \leq n$ is called the *rank* of the lattice.

Letting span denote the \mathbb{R} -linear span of a set, the *dual basis* $\mathbf{B}^\vee = \{\mathbf{b}_1^\vee, \dots, \mathbf{b}_k^\vee\} \subset \text{span}(\mathbf{B})$ and dual lattice $\mathcal{L} = \mathcal{L}(\mathbf{B}^\vee)$ are defined to satisfy $\langle \mathbf{b}_j^\vee, \mathbf{b}_{j'} \rangle = \delta_{j,j'}$ for all j, j' , where the Kronecker delta $\delta_{j,j'} = 1$ if $j = j'$, and is 0 otherwise. In other words, $\mathbf{B}^t \cdot \mathbf{B}^\vee = (\mathbf{B}^\vee)^t \cdot \mathbf{B}$ is the identity matrix.

In this work we deal with a computational problem on lattices called *bounded-distance decoding* (BDD): given a lattice basis $\mathbf{B} \subset \mathbb{R}^n$ of $\mathcal{L} = \mathcal{L}(\mathbf{B})$ and a target point $\mathbf{t} \in \text{span}(\mathcal{L})$ with the guarantee that $\min_{\mathbf{v} \in \mathcal{L}} \|\mathbf{v} - \mathbf{t}\| \leq r$ for some known $r < \lambda_1(\mathcal{L})/2$, find the unique $\mathbf{v} \in \mathcal{L}$ closest to \mathbf{t} (i.e., such that $\|\mathbf{v} - \mathbf{t}\| \leq r$). In fact, in our context \mathbf{B} and r will be fixed in advance, and \mathbf{t} is the only input that may vary.

A standard approach to solve BDD (and related problems) is the “round-off” algorithm of [Bab85], which simply returns $\mathbf{B} \cdot \lfloor (\mathbf{B}^\vee)^t \cdot \mathbf{t} \rfloor$, where the rounding function $\lfloor c \rfloor := \lfloor c + \frac{1}{2} \rfloor \in \mathbb{Z}$ is applied to each coordinate independently. (Notice that $(\mathbf{B}^\vee)^t \cdot \mathbf{t}$ is the coefficient vector of \mathbf{t} with respect to basis \mathbf{B} .) We recall the following standard fact about this algorithm, and include a brief proof for completeness.

Claim 2.1. *Let $\mathcal{L} \subset \mathbb{R}^n$ be a lattice with basis \mathbf{B} , and let $\mathbf{t} = \mathbf{v} + \mathbf{e} \in \mathbb{R}^n$ for some $\mathbf{v} \in \mathcal{L}$, $\mathbf{e} \in \mathbb{R}^n$. If $\langle \mathbf{b}_j^\vee, \mathbf{e} \rangle \in [-\frac{1}{2}, \frac{1}{2})$ for all j , then on input \mathbf{t} and basis \mathbf{B} , the round-off algorithm outputs \mathbf{v} .*

Proof. Because $\mathbf{v} = \mathbf{B}\mathbf{z}$ for some integer vector \mathbf{z} , we have $(\mathbf{B}^\vee)^t \cdot \mathbf{t} = \mathbf{z} + (\mathbf{B}^\vee)^t \cdot \mathbf{e}$, so by hypothesis on the $\langle \mathbf{b}_j^\vee, \mathbf{e} \rangle$, we have $\lfloor (\mathbf{B}^\vee)^t \cdot \mathbf{t} \rfloor = \mathbf{z}$. The claim follows. \square

2.2 Circulant Matrices

We recall some standard facts about *circulant* matrices for a finite abelian group (G, \cdot) , and their relationship with the *characters* of the group. See, e.g., see [Lan02] for further details and proofs.

Definition 2.2 (Circulant matrix). For a vector $\mathbf{a} = (a_g)_{g \in G}$ indexed by G , the G -*circulant* matrix associated with \mathbf{a} is the G -by- G matrix whose (i, j) th entry is $a_{ij^{-1}}$.

Note that the transpose of any G -circulant matrix (associated with $(a_g)_{g \in G}$) is also a G -circulant matrix (associated with $(a_{g^{-1}})_{g \in G}$).

Definition 2.3 (Character group). A *character* is a group morphism $\chi: G \rightarrow \{u \in \mathbb{C} : |u| = 1\}$, i.e., $\chi(g \cdot h) = \chi(g) \cdot \chi(h)$ for all $g, h \in G$. The *character group* (\hat{G}, \cdot) is the set of characters of G , with the group operation being the usual multiplication of functions, i.e., $(\chi \cdot \psi)(g) = \chi(g) \cdot \psi(g)$.

A basic fact is that $|\hat{G}| = |G|$. Notice that for a character $\chi \in \hat{G}$, we have $\overline{\chi(g)} = \chi(g)^{-1} = \chi(g^{-1})$. We identify χ with the vector $(\chi(g))_{g \in G}$. Then all characters χ have Euclidean norm $\|\chi\| = \sqrt{|G|}$, because

$$\langle \chi, \chi \rangle = \sum_{g \in G} \chi(g) \cdot \overline{\chi(g)} = \sum_{g \in G} 1 = |G|.$$

Moreover, distinct characters χ, ψ are orthogonal:

$$\langle \chi, \psi \rangle = \sum_{g \in G} \chi(g) \cdot \overline{\psi(g)} = \sum_{g \in G} (\chi \cdot \psi^{-1})(g) = 0.$$

Therefore, the complex G -by- \hat{G} matrix

$$\mathbf{P}_G := |G|^{-1/2} \cdot (\chi(g))_{g \in G, \chi \in \hat{G}}$$

is unitary, i.e., $\mathbf{P}_G^{-1} = \mathbf{P}_G^*$, the conjugate transpose of \mathbf{P}_G .

Lemma 2.4. *A complex matrix \mathbf{A} is G -circulant if and only if the \hat{G} -by- \hat{G} matrix $\mathbf{P}_G^{-1} \cdot \mathbf{A} \cdot \mathbf{P}_G$ is diagonal; equivalently, the columns of \mathbf{P}_G are the eigenvectors of \mathbf{A} . If \mathbf{A} is the G -circulant matrix associated with $\mathbf{a} = (a_g)_{g \in G}$, its eigenvalue corresponding to $\chi \in \hat{G}$ is $\lambda_\chi = \langle \mathbf{a}, \chi \rangle = \sum_{g \in G} a_g \cdot \overline{\chi(g)}$.*

It follows that every row and column of \mathbf{A} has squared Euclidean norm

$$\|\mathbf{a}\|^2 = \|\mathbf{P}_G^* \cdot \mathbf{a}\|^2 = |G|^{-1} \cdot \sum_{\chi \in \hat{G}} |\lambda_\chi|^2.$$

It also follows that \mathbf{A}^{-1} (when defined) is G -circulant, with eigenvalue λ_χ^{-1} for eigenvector χ .

Proof. Suppose that \mathbf{A} is G -circulant, and let $\chi \in \hat{G}$ be a character of G . Then

$$(\mathbf{A} \cdot \chi)_g = \sum_{h \in G} a_{gh^{-1}} \cdot \chi(h) = \left(\sum_{k \in G} a_k \cdot \overline{\chi(k)} \right) \cdot \chi(g),$$

where in the final equality we have substituted $k = gh^{-1}$ and used $\chi(h) = \overline{\chi(k)} \cdot \chi(g)$. So $\mathbf{A} \cdot \chi = \lambda_\chi \cdot \chi$.

For the other direction, it suffices by linearity to show that $\mathbf{A}_\chi = \mathbf{P}_G \cdot \mathbf{D}_\chi \cdot \mathbf{P}_G^{-1}$ is G -circulant for every $\chi \in \hat{G}$, where \mathbf{D}_χ is the diagonal \hat{G} -by- \hat{G} matrix with 1 in its (χ, χ) th entry and zeros elsewhere. Indeed, by definition of \mathbf{P}_G and because $\mathbf{P}_G^{-1} = \mathbf{P}_G^*$, the (i, j) th entry of \mathbf{A}_χ is simply $|G|^{-1} \cdot \chi(i) \cdot \overline{\chi(j)} = |G|^{-1} \cdot \chi(ij^{-1})$, which depends only on ij^{-1} as required. \square

2.3 Dirichlet Characters and L -Series

A *Dirichlet character* χ is a character of \mathbb{Z}_k^* for some positive integer k . Note that if $k|\ell$ then χ induces a character of \mathbb{Z}_ℓ^* via the natural morphism $\mathbb{Z}_\ell^* \rightarrow \mathbb{Z}_k^*$, so we can equivalently view χ as being defined modulo either k or ℓ . The *conductor* f_χ of χ is the smallest positive f such that χ is induced by a Dirichlet character modulo f . The character is said to be *even* if $\chi(-1) = 1$; note that the even Dirichlet characters correspond with the characters of $\mathbb{Z}_k^*/\{\pm 1\}$. We often implicitly extend χ to a completely multiplicative function from \mathbb{Z} to \mathbb{C} , by letting $\chi(a) = 0$ if $\gcd(a, k) > 1$.

Definition 2.5 (Dirichlet L -Series). For a Dirichlet character χ , the Dirichlet L -function $L(\cdot, \chi)$ is defined as the formal series

$$L(s, \chi) = \sum_{k \geq 1} \frac{\chi(k)}{k^s}.$$

For any Dirichlet character χ , the series $L(s, \chi)$ is absolutely convergent for all $s \in \mathbb{C}$ with $\Re(s) > 1$. It is also known that $L(1, \chi)$ converges and is nonzero for any nontrivial Dirichlet character (i.e., $\chi \neq 1$). We have the following asymptotic bounds on its value.

Theorem 2.6. *There exists a real constant $C > 0$ such that, for any character χ of conductor $f > 1$,*

$$\frac{1}{\ell(f)} \leq |L(1, \chi)| \leq \ell(f) \quad \text{where } \ell(f) = C \ln f.$$

The above result can be traced back to Landau [Lan27], and improving the constant C is an active field of research [Lou15]. Additionally, under the Generalized Riemann Hypothesis, the above bound can be improved to $\ell(f) = C \ln \ln f$ (see [YXK13]).

2.4 Cyclotomic Number Fields and the Log-Unit Lattice

Cyclotomic number fields. Let L be a field. An element $\zeta \in L$ is a root of unity if $\zeta^m = 1$ for some positive integer m . The order of a root of unity $\zeta \in L$ is the order of the finite multiplicative subgroup of L^* generated by ζ . A primitive m th root of unity in L is a root of unity $\zeta \in L$ of order m . Note that if $\zeta \in L$ is a primitive m th root of unity, then the polynomial $X^m - 1 \in L[X]$ factors as $\prod_{i=0}^{m-1} (X - \zeta^i)$ over $L[X]$. Also note that the complete set of primitive m th roots in L consists of the powers ζ^j for $j \in \mathbb{Z}_m^*$.

An algebraic number field K is an extension field of the rationals \mathbb{Q} such that its dimension $[K : \mathbb{Q}]$ as a \mathbb{Q} -vector space (i.e., its degree) is finite. If $\Omega \supset K$ is an extension field such that Ω is algebraically closed over \mathbb{Q} , then there are exactly $[K : \mathbb{Q}]$ field embeddings of K into Ω .³ An algebraic number field is Galois if the order of its automorphism group equals its degree.⁴ A number field K is cyclotomic if $K = \mathbb{Q}(\zeta)$ for some root of unity $\zeta \in K$. Its degree is $\varphi(m)$, where $\varphi(\cdot)$ is the Euler totient function and m is the order of ζ , and its ring of integers R is monogenic, i.e., $R = \mathbb{Z}[\zeta]$. We let U denote the cyclic (multiplicative) subgroup of m th roots of unity, which is generated by ζ .

A cyclotomic number field is Galois. If $K = \mathbb{Q}(\zeta)$ is a cyclotomic number field with $\zeta \in K$ an m th primitive root of unity then each automorphism is characterized by the assignment $\zeta \mapsto \zeta^j$ for some $j \in \mathbb{Z}_m^*$. As a consequence, if L is an extension field of a cyclotomic field K , then K is situated uniquely in L . For concreteness, we situate cyclotomic number fields in the complex numbers \mathbb{C} . Let m be a positive integer and define $\omega = \omega_m = \exp(2\pi i/m) \in \mathbb{C}$. Then ω is a primitive m th root of unity and $K = \mathbb{Q}(\omega)$ is the m th cyclotomic number field. The embeddings of K into the complex numbers (i.e., the automorphisms of K) are denoted σ_j for $j \in \mathbb{Z}_m^*$, where σ_j sends ω to ω^j .

Logarithmic embedding. The embeddings σ_i of K , being complex, come in conjugate pairs, i.e., $\sigma_j(x) = \overline{\sigma_{-j}(x)}$. We will mainly be concerned with their *magnitudes*, so we identify the pairs by indexing over the

³These embeddings are merely ring morphisms $\psi : K \rightarrow \Omega$. Each such ψ is automatically injective because K is a field. Also note that any such ψ fixes \mathbb{Q} pointwise.

⁴An automorphism of a field L is a ring isomorphism $\psi : L \rightarrow L$. The automorphisms of L form a group with functional composition as the group operation.

multiplicative quotient group $G := \mathbb{Z}_m^*/\{\pm 1\}$. We then have the *logarithmic embedding*, defined as

$$\begin{aligned} \text{Log}: K &\rightarrow \mathbb{R}^{\varphi(m)/2} \\ a &\mapsto (\log|\sigma_i(a)|)_{i \in G}. \end{aligned}$$

The logarithmic embedding defines a group morphism, mapping the multiplicative group K^* to an additive subgroup of $\mathbb{R}^{\varphi(m)/2}$. The kernel of Log restricted to R^* is $\{\pm 1\} \cdot U$. The Dirichlet Unit Theorem (see [Sam70, Chapter 4.4, Theorem 1]) implies that $\Lambda = \text{Log}(R^*)$, the image of the multiplicative unit group of R under the logarithmic embedding, is a full-rank lattice in the linear subspace of $\mathbb{R}^{\varphi(m)/2}$ orthogonal to the all-1s vector $\mathbf{1}$. We refer to Λ as the *log-unit lattice*.

Cyclotomic units. Let A be the multiplicative subgroup of K^* generated by $\pm\zeta$ and

$$z_j := \zeta^j - 1, \quad j \in \mathbb{Z}_m \setminus \{0\}.$$

Notice that $z_j = -\zeta^j \cdot z_{-j}$, so z_j and z_{-j} are equivalent modulo $\pm U$; in particular, $\text{Log}(z_j) = \text{Log}(z_{-j})$. The group of *cyclotomic units*, denoted C , is defined by

$$C = A \cap R^*.$$

The z_j given above are not necessarily units in R , and thus do not generate C . However, a closely related generating set, which we call the *canonical generators*, is given by the following lemma. Recall that $G = \mathbb{Z}_m^*/\{\pm 1\}$, and identify it with some canonical set of representatives in \mathbb{Z}_m^* .

Lemma 2.7 (Lemma 8.1 of [Was97]). *Let m be a prime power, and define $b_j := z_j/z_1 = (\zeta^j - 1)/(\zeta - 1)$. The group C of cyclotomic units is generated by $\pm\zeta$ and b_j for $j \in G \setminus \{1\}$.*

Notice that $\text{Log } C$ is a sublattice of Λ . As shown below, the index of Λ over $\text{Log } C$ is finite. In fact, it is $h^+(m)$, the *class number* of the real subfield $K^+ = \mathbb{Q}(\zeta + \bar{\zeta})$, defined as the index of the subgroup of principal fractional ideals in the multiplicative group of all fractional ideals (in K^+). The proof of this theorem is left as Exercise 8.5 in [Was97]. For completeness, we sketch the solution in Appendix B.

Theorem 2.8. *For a prime power $m > 2$, the index of the log-unit lattice Λ over $\text{Log } C$ is*

$$[\Lambda : \text{Log } C] = h^+(m).$$

Some facts and conjectures concerning h^+ . For our purposes, we need $h^+(m)$ not to be very big. For all power-of-two m up to $m = 256$, and also for $m = 512$ under GRH, it is known that $h^+(m) = 1$ (see [Mil14a]). Whether $h^+(m) = 1$ for all power-of-two m is known as Weber's class number problem, and is presented in the literature as a reasonable conjecture.

In the case of odd primes, it also appears that h^+ is quite small. Computations of Schoof [Sch03] and Miller [Mil14b] show that $h^+(p) \leq 11$ for all primes $p \leq 241$. For powers of odd primes it has been conjectured (with support of the Cohen-Lenstra heuristic) that, for all but finitely many pairs (p, ℓ) where p is a prime, $h^+(p^{\ell+1}) = h^+(p^\ell)$ [BPR04]. A direct consequence is that $h^+(p^\ell)$ is bounded for a fixed p and increasing ℓ .

3 Geometry of the Canonical Generators

Throughout this section, let the cyclotomic index m be a prime power. Our goal here is to show that the canonical generators of the cyclotomic units, under the logarithmic embedding, are geometrically well-suited for bounded-distance decoding.

Recalling that $G = \mathbb{Z}_m^*/\{\pm 1\}$ is identified with some set of canonical representatives in \mathbb{Z}_m^* and that $\text{Log}(b_j) = \text{Log}(b_{-j})$, define

$$\mathbf{b}_j = \text{Log}(b_j), \quad j \in G \setminus \{1\},$$

to be the log-embeddings of the canonical generators $b_j = (\zeta^j - 1)/(\zeta - 1)$ defined in Lemma 2.7. By Lemma 2.7, these \mathbf{b}_j form a basis of the sublattice $\text{Log } C$, which by Theorem 2.8 has index $h^+(m)$ in Λ .

In order to apply the round-off algorithm and Claim 2.1 with this basis, we bound the norms $\|\mathbf{b}_j^\vee\|$ of the dual basis vectors. The remainder of this section is dedicated to proving the following theorem.

Theorem 3.1. *Let $m = p^k$ for a prime p , and let $\{\mathbf{b}_j^\vee\}_{j \in G \setminus \{1\}}$ denote the basis dual to $\{\mathbf{b}_j\}_{j \in G \setminus \{1\}}$. Then all $\|\mathbf{b}_j^\vee\|$ are equal, and*

$$\|\mathbf{b}_j^\vee\|^2 \leq 2k|G|^{-1} \cdot \ell(m)^2 = O(m^{-1} \cdot \log^3 m).$$

To prove the theorem we start by relating the basis vectors \mathbf{b}_j to a certain G -circulant matrix. Recalling that $z_j = \zeta^j - 1$ is the numerator of b_j , define $\mathbf{z}_j := \text{Log}(z_j) = \mathbf{b}_j + \mathbf{z}_1$. Collect these vectors into a square G -by- G matrix \mathbf{Z} whose j th column is \mathbf{z}_{j-1} , and notice that its (i, j) th entry $\log|\omega^{i \cdot j^{-1}} - 1|$ is determined by $ij^{-1} \in G$ alone, so \mathbf{Z} is the G -circulant matrix associated with \mathbf{z}_1 . For each eigenvector $\chi \in \hat{G}$ of \mathbf{Z} , let $\lambda_\chi := \langle \mathbf{z}_1, \chi \rangle$ denote the corresponding eigenvalue.

Lemma 3.2. *For all $j \in G \setminus \{1\}$ we have*

$$\|\mathbf{b}_j^\vee\|^2 = |G|^{-1} \cdot \sum_{\chi \in \hat{G} \setminus \{1\}} |\lambda_\chi|^{-2}. \quad (1)$$

Proof. Let \mathbf{z}_j^\vee denote the vectors dual to the \mathbf{z}_j , i.e., the columns of \mathbf{Z}^{-t} . (As shown below in the proof of Theorem 3.1, \mathbf{Z}^{-1} is indeed well defined because all eigenvalues λ_χ of \mathbf{Z} are nonzero.)

We first claim that \mathbf{b}_j^\vee is simply the projection of \mathbf{z}_j^\vee orthogonal to $\mathbf{1}$, i.e., $\mathbf{b}_j^\vee = \mathbf{z}_j^\vee - |G|^{-1} \cdot \langle \mathbf{z}_j^\vee, \mathbf{1} \rangle \cdot \mathbf{1}$. Indeed, these vectors are all in $\text{span}(\mathbf{b}_{j'})_{j'}$, the space orthogonal to $\mathbf{1}$, and moreover, for all $j, j' \in G \setminus \{1\}$ they satisfy

$$\langle \mathbf{z}_j^\vee - |G|^{-1} \cdot \langle \mathbf{z}_j^\vee, \mathbf{1} \rangle \cdot \mathbf{1}, \mathbf{b}_{j'} \rangle = \langle \mathbf{z}_j^\vee, \mathbf{b}_{j'} \rangle = \langle \mathbf{z}_j^\vee, \mathbf{z}_{j'} - \mathbf{z}_1 \rangle = \delta_{j,j'} - 0.$$

Now,

$$\|\mathbf{b}_j^\vee\|^2 = \|\mathbf{z}_j^\vee\|^2 - |G|^{-1} \cdot \langle \mathbf{z}_j^\vee, \mathbf{1} \rangle^2.$$

Recall by Lemma 2.4 that \mathbf{Z}^{-t} is the G -circulant matrix associated with \mathbf{z}_1^\vee , which has eigenvalue $\lambda_\chi^{-1} = \langle \mathbf{z}_1^\vee, \chi \rangle$ for eigenvector $\chi \in \hat{G}$. By the remarks following Lemma 2.4, $\|\mathbf{z}_j^\vee\|^2 = |G|^{-1} \cdot \sum_{\chi \in \hat{G}} |\lambda_\chi|^{-2}$. The lemma follows by noting that $\langle \mathbf{z}_j^\vee, \mathbf{1} \rangle = \langle \mathbf{z}_1^\vee, \mathbf{1} \rangle = \lambda_1^{-1}$. \square

We now provide an upper bound on the right-hand side of Equation (1). Our proof is similar to the proof that the cyclotomic units have finite index in the full group of units [Was97, Theorem 8.2].

Theorem 3.3 ([Was97, Lemma 4.8 and Theorem 4.9]). Let χ be an even Dirichlet character of conductor $f > 1$, and let $\omega_f = \exp(2\pi i/f) \in \mathbb{C}$. Then

$$\left| \sum_{a \in \mathbb{Z}_f^*} \overline{\chi(a)} \cdot \log|1 - \omega_f^a| \right| = \sqrt{f} \cdot |L(1, \chi)|.$$

For completeness, we briefly explain how the finite sum on the left hand side gives rise to an L -series, and refer to [Was97] for the details. Using the Taylor expansion

$$\log|1 - x| = - \sum_{k \geq 1} x^k/k,$$

one gets a sum over finitely many a and infinitely many k of terms $\overline{\chi(a)} \cdot \omega_f^{ak}/k$. For a fixed k , the sum over a can easily be rewritten as $\tau(\chi) \cdot \chi(k)/k$, where $\tau(\chi)$ is a Gauss sum (see [Was97, Lemma 4.7]), which makes the Dirichlet L -function apparent.

Corollary 3.4. Suppose $f > 1$ divides a prime power m . For any even Dirichlet character of conductor f ,

$$\left| \sum_{a \in \mathbb{Z}_m^*} \overline{\chi(a)} \cdot \log|1 - \omega_m^a| \right| = \sqrt{f} \cdot |L(1, \chi)|.$$

Proof. Let $\phi: \mathbb{Z}_m^* \rightarrow \mathbb{Z}_f^*$ be the map given by reduction modulo f . We have

$$\begin{aligned} \sum_{a \in \mathbb{Z}_m^*} \overline{\chi(a)} \cdot \log|1 - \omega_m^a| &= \sum_{a \in \mathbb{Z}_f^*} \overline{\chi(a)} \sum_{\substack{b \in \mathbb{Z}_m^* \\ \phi(b)=a}} \log|1 - \omega_m^b| \\ &= \sum_{a \in \mathbb{Z}_f^*} \overline{\chi(a)} \cdot \log \left| \prod_{\substack{b \in \mathbb{Z}_m^* \\ \phi(b)=a}} (1 - \omega_m^b) \right| \\ &= \sum_{a \in \mathbb{Z}_f^*} \overline{\chi(a)} \cdot \log|1 - \omega_f^a|, \end{aligned}$$

where in the last equality we have used the identity $\prod_{i \in \mathbb{Z}_n} (1 - \omega_n^i Y) = 1 - Y^n$ and $\omega_m^n = \omega_f$ with $n = m/f$. The claim follows by applying Theorem 3.3. \square

We are now ready to complete the proof of the main theorem.

Proof of Theorem 3.1. Recall that the characters $\chi \in \hat{G}$ correspond to the even characters of \mathbb{Z}_m^* , because $\chi(\pm 1) = 1$. Also recall that by Lemma 2.4, the eigenvalues are

$$\lambda_\chi = \langle \mathbf{z}_1, \chi \rangle = \sum_{a \in G} \overline{\chi(a)} \cdot \log|1 - \omega_m^a| = \frac{1}{2} \sum_{a \in \mathbb{Z}_m^*} \overline{\chi(\pm a)} \cdot \log|1 - \omega_m^a|,$$

where the second equality holds because $|1 - \omega_m^{-a}| = |1 - \omega_m^a|$. Therefore, using Corollary 3.4 and Theorem 2.6, we have

$$|\lambda_\chi| = \frac{1}{2} \sqrt{f_\chi} \cdot |L(1, \chi)| \geq \frac{\sqrt{f_\chi}}{2\ell(f_\chi)} \geq \frac{\sqrt{f_\chi}}{2\ell(m)}.$$

Hence, by Lemma 3.2,

$$\|\mathbf{b}_j^\vee\|^2 = |G|^{-1} \cdot \sum_{\chi \in \widehat{G} \setminus \{1\}} |\lambda_\chi|^{-2} \leq 4|G|^{-1} \cdot \ell(m)^2 \sum_{\chi \in \widehat{G} \setminus \{1\}} f_\chi^{-1} \leq 2k|G|^{-1} \cdot \ell(m)^2,$$

where the last inequality follows from Claim 3.5 below. \square

Claim 3.5. *Let $m = p^k$ for a prime p . Then, for $G = \mathbb{Z}_m^*/\{\pm 1\}$,*

$$\sum_{\chi \in \widehat{G} \setminus \{1\}} f_\chi^{-1} \leq \frac{k}{2}.$$

Proof. Notice that there are at most f Dirichlet characters of conductor f , at most half of which are even (when $f > 1$), so

$$\sum_{\chi \in \widehat{G} \setminus \{1\}} f_\chi^{-1} \leq \sum_{\ell=1}^k \frac{p^\ell}{2} \cdot \frac{1}{p^\ell} = \frac{k}{2}. \quad \square$$

4 Algorithmic Implications

The following is our main algorithmic result, showing that under mild restrictions on the distribution of the short generator, we can recover it from any generator that differs from it by a unit in C .

Theorem 4.1. *Let D be a distribution over $\mathbb{Q}(\zeta)$ with the property that for any tuple of vectors $\mathbf{v}_1, \dots, \mathbf{v}_{\varphi(m)/2-1} \in \mathbb{R}^{\varphi(m)/2}$ of Euclidean norm 1 that are orthogonal to the all-1 vector $\mathbf{1}$, the probability that $|\langle \text{Log}(g), \mathbf{v}_i \rangle| < cm^{1/2} \log^{-3/2} m$ holds for all i is at least some $\alpha > 0$, where g is chosen from D and c is a universal constant. There is an efficient algorithm that given $g' = g \cdot u$, where g is chosen from D and $u \in C$ is a cyclotomic unit, outputs an element of the form $\pm \zeta^j g$ with probability at least α .*

Proof. The algorithm applies the roundoff algorithm from Claim 2.1 to $\text{Log}(g') = \text{Log}(g) + \text{Log}(u)$, using the vectors \mathbf{b}_j (defined and analyzed in Section 3) as the basis. By the assumption on D and Theorem 3.1, with probability at least α the output is $\text{Log}(u) \in \text{Log}(C)$. We next find integer coefficients a_j such that $\text{Log}(u) = \sum a_j \mathbf{b}_j$, and compute $u' = \prod b_j^{a_j}$. Since $\text{Log}(u') = \text{Log}(u)$ it follows that u' must be of the form $\pm \zeta^j u$ for some sign and some j . Therefore, g'/u' is the desired element. \square

In the next section we show that the condition on D in the theorem is satisfied by several natural distributions.

One possible concern with the above algorithm is that it expects as input $g \cdot u$ for a *cyclotomic unit* $u \in C$, whereas the first phase of the attack described in the introduction, i.e., a PIP algorithm, is only guaranteed to output $g \cdot u$ for an *arbitrary unit* $u \in R^*$. There are several reasons why this should not be an issue. First, as mentioned in Section 2, in some cases, e.g., for power-of-2 cyclotomic, it is conjectured that $C = R^*$. More generally, the index of C in R^* , which we recall is h^+ , the class number of the totally real subfield, is often small. In such a case, if we have a list of coset representatives of C in R^* , we can enumerate over all of them and use the algorithm above to recover g , increasing the running time only by a factor of h^+ . In order to obtain such a list of representatives, we can use an algorithm for computing the unit group, either

classical [BF14] or quantum [EHKS14]. These algorithms are no slower than the known PIP algorithms and moreover, need only be applied once for a given cyclotomic field (as opposed to once for each public key). Alternatively, by running the PIP algorithm multiple times on a basis of a principal ideal with a *known* short generator chosen using the secret key generation algorithm, we can recover a list of representatives for all the cosets that show up as output of the PIP algorithm with non-negligible probability; we can then enumerate over that list.

5 Tail Bounds

In this section we show that the condition on D in Theorem 4.1 is satisfied by two natural distributions: the continuous Gaussian and a wide enough discrete Gaussian (over any lattice). This section is independent of the other sections in this paper, and we avoid the use of notation from algebraic number theory. Instead, we identify elements of K with vectors in $\mathbb{R}^{\varphi(m)}$ by taking the real and the imaginary part of their $\varphi(m)/2$ complex embeddings, i.e., a is mapped to $(\Re(\sigma_j(a)), \Im(\sigma_j(a)))_{j \in G}$. As a result, all random variables appearing here are real. The results in this section should be easy to extend to other distributions.

We start with Lemma 5.2, a tail bound on the sum of subexponential random variables. The proof is based on a standard Bernstein argument, and follows the proof in [Ver12] apart from some minor modifications for convenience.

Definition 5.1. For $\alpha, \beta > 0$, we say that a random variable X is (α, β) -subexponential if

$$\mathbb{E}[\cosh(\alpha X)] \leq \beta,$$

where recall that $\cosh(x) := (e^x + e^{-x})/2$.

Lemma 5.2 (Tail bound). *Let X_1, \dots, X_n be independent centered (i.e., expectation zero) (α, β) -subexponential random variables. Then, for any $\mathbf{a} = (a_1, \dots, a_n) \in \mathbb{R}^n$ and every $t \geq 0$,*

$$\Pr\left[\left|\sum a_i X_i\right| \geq t\right] \leq 2 \exp\left(-\min\left(\frac{\alpha^2 t^2}{8\beta \|\mathbf{a}\|_2^2}, \frac{\alpha t}{2\|\mathbf{a}\|_\infty}\right)\right).$$

Proof. By scaling, we can assume without loss of generality that $\alpha = 1$. Next, we use the inequality

$$e^{\delta x} - \delta x - 1 \leq (e^{\delta x} - \delta x - 1) + (e^{-\delta x} + \delta x - 1) = 2(\cosh(\delta x) - 1) \leq 2\delta^2(\cosh(x) - 1)$$

which holds for all $-1 \leq \delta \leq 1$ and all $x \in \mathbb{R}$, where the second inequality follows from the Taylor expansion. By applying this inequality to a $(1, \beta)$ -subexponential centered random variable X , and taking expectations we see that for all $-1 \leq \delta \leq 1$,

$$\begin{aligned} \mathbb{E}[\exp(\delta X)] &\leq 1 + 2\delta^2 \mathbb{E}[\cosh(X) - 1] \\ &\leq 1 + 2\delta^2(\beta - 1) \leq \exp(2\delta^2\beta). \end{aligned} \tag{2}$$

Using Markov's inequality, we can bound the upper tail probability for any $\lambda > 0$ as

$$\begin{aligned} \Pr\left[\sum a_i X_i \geq t\right] &= \Pr\left[\exp\left(\lambda \sum a_i X_i\right) \geq \exp(\lambda t)\right] \\ &\leq \exp(-\lambda t) \cdot \mathbb{E}\left[\exp\left(\lambda \sum a_i X_i\right)\right] \\ &= \exp(-\lambda t) \cdot \prod \mathbb{E}[\exp(\lambda a_i X_i)] \\ &\leq \exp(-\lambda t + 2\beta \lambda^2 \|\mathbf{a}\|_2^2), \end{aligned}$$

where in the second inequality we used (2) and assumed that $\lambda \|\mathbf{a}\|_\infty \leq 1$. Taking $\lambda = \min(t/(4\beta\|\mathbf{a}\|_2^2), 1/\|\mathbf{a}\|_\infty)$ this bound becomes at most

$$\exp\left(-\min\left(\frac{t^2}{8\beta\|\mathbf{a}\|_2^2}, \frac{t}{2\|\mathbf{a}\|_\infty}\right)\right).$$

We complete the proof by applying the same argument with $-\mathbf{a}$. \square

The next claim follows immediately from Definition 5.1.

Claim 5.3. *If Y is a non-negative random variable such that both $\mathbb{E}[Y]$ and $\mathbb{E}[Y^{-1}]$ are finite, then $\log Y$ is a $(1, \beta)$ -subexponential random variable for some $\beta > 0$.*

The following is an immediate corollary of the tail bound. It shows that the condition in Theorem 4.1 holds with overwhelming probability for a continuous Gaussian distribution of any radius that is spherical in the embedding basis. Notice that the parameter r plays no role in the conclusion of the statement.

Lemma 5.4. *Let $X_1, \dots, X_n, X'_1, \dots, X'_n$ be i.i.d. $N(0, r)$ variables for some $r > 0$, and let $\hat{X}_i = (X_i^2 + X_i'^2)^{1/2}$. Then, for any vectors $\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(\ell)} \in \mathbb{R}^n$ of Euclidean norm 1 that are orthogonal to the all-1 vector, and every $t \geq C$ for some universal constant C ,*

$$\Pr\left[\exists j, \left|\sum_i a_i^{(j)} \log(\hat{X}_i)\right| \geq t\right] \leq 2\ell \exp(-t/2).$$

Proof. By union bound, it suffices to prove the lemma for the case $\ell = 1$, and we let $\mathbf{a} = \mathbf{a}^{(1)}$. Since $\sum a_i = 0$, we can assume without loss of generality that $r = 1$. Notice that \hat{X}_i has a chi distribution with 2 degrees of freedom (also known as a Rayleigh distribution) whose density function is given by $xe^{-x^2/2}$ for $x > 0$ and zero otherwise. In particular, it is easy to see that both $\mathbb{E}[\hat{X}_i]$ and $\mathbb{E}[\hat{X}_i^{-1}]$ are finite (both are $\sqrt{\pi/2}$). Therefore, by Claim 5.3, $\log \hat{X}_i$ is $(1, \beta)$ subexponential for some constant $\beta > 0$. From this it follows that $\hat{X}_i = \log \hat{X}_i - \mathbb{E}[\log \hat{X}_i]$ are centered $(1, \beta')$ subexponential random variables for some constant $\beta' > 0$. The result now follows by applying Lemma 5.2 to $\hat{X}_1, \dots, \hat{X}_n$, using the bound $\|\mathbf{a}\|_\infty \leq 1$, and the observation that $\sum_i a_i \mathbb{E}[\log \hat{X}_i] = 0$. \square

In the next lemma we show that small perturbations of the continuous Gaussian distribution still satisfy the condition in Theorem 4.1.

Lemma 5.5. *Let $X = (X_1, \dots, X_n, X'_1, \dots, X'_n)$ be i.i.d. $N(0, r)$ variables for some $r > 0$, and let $Y = (Y_1, \dots, Y_n, Y'_1, \dots, Y'_n)$ be a (not necessarily independent) random vector satisfying $\|Y\|_2 \leq u$ with probability 1 for some $u \leq r/(20\sqrt{n})$. Let $Z = X + Y$ and define $\hat{X}_i, \hat{Y}_i, \hat{Z}_i$ as before. Then for any vectors $\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(\ell)} \in \mathbb{R}^n$ of Euclidean norm 1 that are orthogonal to the all-1 vector, it holds with constant probability that for all j ,*

$$\left|\sum_i a_i^{(j)} \log(\hat{Z}_i)\right| \leq 1 + 10 \log \ell.$$

Proof. By Lemma 5.4 we have that with some constant probability close to 1,

$$\forall j, \left|\sum_i a_i^{(j)} \log(\hat{X}_i)\right| < 10 \log \ell. \quad (3)$$

Moreover, since $\hat{X}_i < r/(10\sqrt{n})$ implies that both $|X_i|$ and $|X'_i|$ are smaller than $r/(10\sqrt{n})$, we see that by independence of X_i, X'_i , the probability of the former event is at most c/n for some small constant c . As a result we have that with constant probability close to 1,

$$\forall i, \hat{X}_i > r/(10\sqrt{n}).$$

In the following we assume that these two conditions hold (which happens with constant probability close to 1 by union bound), and bound the effect of Y . Now let \mathbf{a} be one of the vectors in the statement of the lemma. Then,

$$\begin{aligned} \left| \sum_i a_i \log(\hat{Z}_i) \right| &\leq \left| \sum_i a_i \log(\hat{X}_i) \right| + \left| \sum_i a_i \log(\hat{Z}_i/\hat{X}_i) \right| \\ &\leq 10 \log \ell + \left| \sum_i a_i \log(\hat{Z}_i/\hat{X}_i) \right|, \end{aligned}$$

where we used Eq. (3). Notice that by the triangle inequality (for two-dimensional Euclidean space),

$$\hat{X}_i - \hat{Y}_i \leq \hat{Z}_i \leq \hat{X}_i + \hat{Y}_i.$$

Since $\hat{Y}_i \leq \|Y\|_2 \leq u \leq r/(20\sqrt{n}) \leq \hat{X}_i/2$, and using the inequality $|\log(1 + \delta)| \leq 2|\delta|$ valid for all $\delta \in [-1/2, 1/2]$,

$$\begin{aligned} \left| \sum_i a_i \log(\hat{Z}_i/\hat{X}_i) \right| &\leq \left(\sum_i (\log(\hat{Z}_i/\hat{X}_i))^2 \right)^{1/2} \\ &\leq \left(\sum_i (2\hat{Y}_i/\hat{X}_i)^2 \right)^{1/2} \\ &\leq 20\sqrt{n}/r \cdot \left(\sum_i \hat{Y}_i^2 \right)^{1/2} \\ &\leq 20\sqrt{nu}/r \leq 1, \end{aligned}$$

where the first inequality follows from Cauchy-Schwarz. □

Finally, we consider the spherical (in the embedding basis) discrete Gaussian distribution over an arbitrary lattice $L \subseteq \mathbb{R}^{2n}$. Such distributions show up often in cryptographic constructions (see, e.g., [LPR13]), and often that lattice is the (embedding of the) ring of integers R . For background on the discrete Gaussian distribution and the smoothing parameter, see, e.g., [MR04]. In order to apply Lemma 5.5 to this distribution, take X to be the continuous Gaussian D_r for some $r \geq 100n\eta_\varepsilon(L)$, and Y the discrete Gaussian $D_{L-X,s}$ over the coset $L - X$ of parameter $s = \eta_\varepsilon(L)$ for some negligible parameter ε . Using Banaszczyk's result [Ban93] we have that with all but exponentially small probability in n , $\|Y\|_2 \leq \sqrt{2n}\eta_\varepsilon(L) \leq r/(60\sqrt{n})$. Moreover, by the lemma below, the distribution of $Z = X + Y$ is within negligible statistical distance of the discrete Gaussian distribution $D_{L,r'}$ for $r' = (r^2 + \eta_\varepsilon(L)^2)^{1/2}$. We therefore see that the condition in Theorem 4.1 holds for the discrete Gaussian distribution $D_{L,r'}$ for any lattice L and any $r' > 200n\eta_\varepsilon(L)$.

Lemma 5.6 (Special case of [Pei10, Theorem 3.1]). *Let L be a lattice and $r, s > 0$ be such that $s \geq \eta_\varepsilon(L)$ for some $\varepsilon \leq 1/2$. Then if we choose \mathbf{x} from the continuous Gaussian D_r and then choose \mathbf{y} from the discrete Gaussian $D_{L-\mathbf{x},s}$ then $\mathbf{x} + \mathbf{y}$ is within statistical distance 8ε of the discrete Gaussian $D_{L,(r^2+s^2)^{1/2}}$.*

A Numeric Data

The previous sections established *asymptotic* bounds related to the log-embeddings of the cyclotomic units. Figure 1 gives concrete numeric data for several practical (and even impractical) choices of cyclotomic fields. This data confirms that the method works in practice.

$k (m = 2^k)$	6	7	8	9	≥ 10
$\ \mathbf{b}_j^\vee\ ^{-1}$	5.04	8.56	14.69	25.71	≥ 45.85

$k (m = 3^k)$	4	5	≥ 6
$\ \mathbf{b}_j^\vee\ ^{-1}$	5.72	13.65	≥ 34.04

$k (m = 5^k)$	3	4	≥ 5
$\ \mathbf{b}_j^\vee\ ^{-1}$	10.04	36.43	≥ 143

Figure 1: Lower bounds on the inverse lengths of the dual vectors \mathbf{b}_j^\vee defined in Section 3, for various cyclotomics of prime-power index. Larger values correspond to larger decoding distances for the log-embedding of the cyclotomic units.

B Proof of Theorem 2.8

Proof. First, Corollary 4.13 of [Was97] gives that $\mathbb{Z}[\zeta]^*$ is generated by $\mathbb{Z}[\zeta + \bar{\zeta}]^*$ and ζ , so it follows that

$$\Lambda = \text{Log } \mathbb{Z}[\zeta]^* = \text{Log } \mathbb{Z}[\zeta + \bar{\zeta}]^*,$$

since the kernel of Log is the group $\{\pm 1\} \cdot U$.

Next, recall that the group of *cyclotomic units* is defined as $C = A \cap R^*$. We define the group of *real cyclotomic units* as $C^+ = A \cap \mathbb{Z}[\zeta + \bar{\zeta}]^*$. The analogue of Lemma 2.7 for the real cyclotomic units, also included in Lemma 8.1 of [Was97], says that the group C^+ of real cyclotomic units is generated by -1 and $\zeta^{(1-j)/2} \cdot b_j$. So as above, we obtain that

$$\text{Log } C = \text{Log } C^+ .$$

The theorem then follows from the sequence of equalities

$$[\Lambda : \text{Log } C] = [\text{Log } \mathbb{Z}[\zeta + \bar{\zeta}]^* : \text{Log } C^+] = [\mathbb{Z}[\zeta + \bar{\zeta}]^* : C^+] = h^+ ,$$

where the second equality follows from $\ker(\text{Log}) \cap C^+ = \ker(\text{Log}) \cap \mathbb{Z}[\zeta + \bar{\zeta}]^* (= \{\pm 1\})$, and the third equality is Theorem 8.2 of [Was97]. \square

References

- [Bab85] L. Babai. On Lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986. Preliminary version in STACS 1985.
- [Ban93] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(4):625–635, 1993.
- [Ber14a] D. Bernstein. A subfield-logarithm attack against ideal lattices. <http://blog.cr.yp.to/20140213-ideal.html>, February 2014.
- [Ber14b] D. Bernstein, June 2014. Personal communication.
- [BF14] J.-F. Biasse and C. Fieker. Subexponential class group and unit group computation in large degree number fields. *LMS J. Comput. Math.*, 17(suppl. A):385–403, 2014.
- [Bia14] J.-F. Biasse. Subexponential time relations in the class group of large degree number fields. *Adv. Math. Commun.*, 8(4):407–425, 2014.
- [BPR04] J. Buhler, C. Pomerance, and L. Robertson. Heuristics for class numbers of prime-power real cyclotomic fields. *Fields Inst. Commun*, 41:149–157, 2004.
- [BS15] J.-F. Biasse and F. Song. A polynomial time quantum algorithm for computing class groups and solving the principal ideal problem in arbitrary degree number fields. <http://www.lix.polytechnique.fr/Labo/Jean-Francois.Biasse/>, 2015. In preparation.
- [CDPR15] R. Cramer, L. Ducas, C. Peikert, and O. Regev. Recovering short generators of principal ideals in cyclotomic rings. <https://www.cwi.nl/brief/2015/candidate-highly-versatile-crypto-systems-less-secure>, 2015. Posted 23 March.
- [CGS14] P. Campbell, M. Groves, and D. Shepherd. Soliloquy: A cautionary tale. ETSI 2nd Quantum-Safe Crypto Workshop, 2014. Available at http://docbox.etsi.org/Workshop/2014/201410_CRYPTOS/S07_Systems_and_Attacks/S07_Groves_Annex.pdf.
- [EHKS14] K. Eisenträger, S. Hallgren, A. Kitaev, and F. Song. A quantum algorithm for computing the unit group of an arbitrary degree number field. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 293–302. ACM, 2014.
- [Gen09] C. Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178. 2009.
- [GGH13] S. Garg, C. Gentry, and S. Halevi. Candidate multilinear maps from ideal lattices. In *EUROCRYPT*, pages 1–17. 2013.
- [GS02] C. Gentry and M. Szydło. Cryptanalysis of the revised NTRU signature scheme. In *EUROCRYPT*, pages 299–320. 2002.
- [HPS98] J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, pages 267–288. 1998.

- [Lan27] E. Landau. Über Dirichletsche Reihen mit komplexen Charakteren. *Journal für die reine und angewandte Mathematik*, 157:26–32, 1927.
- [Lan02] S. Lang. Algebra, volume 211 of graduate texts in mathematics, 2002.
- [LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, December 1982.
- [LMPR08] V. Lyubashevsky, D. Micciancio, C. Peikert, and A. Rosen. SWIFFT: A modest proposal for FFT hashing. In *FSE*, pages 54–72. 2008.
- [Lou15] S. Louboutin. An explicit lower bound on moduli of Dirichlet L -functions at $s = 1$. *J. Ramanujan Math. Soc.*, 30(1):101–113, 2015.
- [LPR10] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *Journal of the ACM*, 60(6):43:1–43:35, November 2013. Preliminary version in Eurocrypt 2010.
- [LPR13] V. Lyubashevsky, C. Peikert, and O. Regev. A toolkit for ring-LWE cryptography. In *EUROCRYPT*, pages 35–54. 2013.
- [LSS14] A. Langlois, D. Stehlé, and R. Steinfeld. Gghlite: More efficient multilinear maps from ideal lattices. In *Advances in Cryptology—EUROCRYPT 2014*, pages 239–256. Springer, 2014.
- [Mic02] D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007. Preliminary version in FOCS 2002.
- [Mil14a] J. C. Miller. Class numbers of totally real fields and applications to the Weber class number problem., 2014.
- [Mil14b] J. C. Miller. Real cyclotomic fields of prime conductor and their class numbers. *arXiv preprint arXiv:1407.2373*, 2014.
- [MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Preliminary version in FOCS 2004.
- [MV10] D. Micciancio and P. Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. In *STOC*, pages 351–358. 2010.
- [Pei10] C. Peikert. An efficient and parallel Gaussian sampler for lattices. In *CRYPTO*, pages 80–97. 2010.
- [Sam70] P. Samuel. *Algebraic Theory of Numbers*. Hermann, Paris, 1970.
- [Sch87] C.-P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.
- [Sch03] R. Schoof. Class numbers of real cyclotomic fields of prime conductor. *Mathematics of computation*, 72(242):913–937, 2003.
- [Sch15] J. Schank. LOGCVP, Pari implementation of CVP in $\text{Log } \mathbb{Z}[\zeta_{2^n}]^*$. <https://github.com/jschanck-si/logcvp>, March 2015.

- [She14] D. Shepherd, December 2014. Personal communication.
- [SV10] N. P. Smart and F. Vercauteren. Fully homomorphic encryption with relatively small key and ciphertext sizes. In *Public Key Cryptography*, pages 420–443. 2010.
- [Ver12] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. In *Compressed sensing*, pages 210–268. Cambridge Univ. Press, Cambridge, 2012. Available at <http://www-personal.umich.edu/~romanv/papers/non-asymptotic-rmt-plain.pdf>.
- [Was97] L. Washington. *Introduction to Cyclotomic Fields*. Graduate Texts in Mathematics. Springer New York, 1997. ISBN 9780387947624.
- [YXK13] L. Youness, L. Xiannan, and S. Kannan. Conditional bounds for the least quadratic non-residue and related problems. <http://arxiv.org/abs/1309.3595>, 2013.