The Closest-Vector Problem over Cyclotomics and its Application to Homomorphic Encryption

Natalie Lang and Dana Dachman-Soled*

University of Maryland
{lang, danadach}@umd.edu

Abstract. Rounding error plays a key role in approximate homomorphic encryption (HE), as it determines the noise level and affects the achievable precision. While rounding is typically performed coordinate-wise in the coefficient basis, an alternative is to perform *lattice-based rounding in the canonical embedding*. These two approaches coincide for power-of-two (PoT) cyclotomics, but can differ for non-PoT cyclotomics. This difference opens the possibility of new parameter tradeoffs for approximate HE over general cyclotomic rings.

Motivated by this distinction, we study rounding error in the Closest-Vector Problem (CVP) over cyclotomic lattices of arbitrary order. We derive new worst- and average-case bounds, showing improved rounding for non-PoT cyclotomics over their PoT counterparts. To enable efficient rounding, we analyze Babai's nearest-plane algorithm as an approximation to exact nearest-point rounding, while retaining efficient ring arithmetic for broad families of cyclotomic orders.

These improved bounds, together with the additional flexibility afforded by moving beyond the sparse PoT orders, inform the choice of encryption parameters for approximate HE. Our concrete evaluation yields parameters with simultaneously smaller lattice dimension and ciphertext modulus than the corresponding PoT choices at fixed output precision and target security level. Our results highlight the potential benefits of non-PoT cyclotomics for approximate HE.

Keywords: Closest-Vector Problem · Cyclotomic Lattices · Approximate Homomorphic Encryption.

1 Introduction

Cyclotomic rings are a fundamental building block of modern lattice-based cryptography. Many efficient constructions, including public-key encryption, digital signatures, and homomorphic encryption (HE), are built over the ring of integers of a cyclotomic number field. Their rich algebraic structure enables compact

* Supported in part by NSF grant #CNS-2154705 and by Intel through the Intel Labs Crypto Frontiers Research Center.

representations and fast polynomial arithmetic via the number theoretic transform (NTT), making them the foundation of practical ring learning with errors (R-LWE)- and Module-LWE-based cryptosystems. Consequently, nearly all practical implementations use power-of-two cyclotomic rings because they admit particularly efficient arithmetic. However, from the perspective of the geometry of the underlying lattice, arbitrary non-power-of-two cyclotomic rings possess different properties that may offer advantages for *lattice-based rounding*—that is, closest-vector computations for arbitrary target points in $\mathbb{R}^n$—while still supporting fast arithmetic. We discuss efficient arithmetic for non-power-of-two cyclotomics in the sequel, and begin by studying the rounding properties of cyclotomic lattices.

Cyclotomic lattices. For each $m \in \mathbb{N}$, let $R = \mathbb{Z}[X]/\Phi_m(X)$ denote the ring of integers of the m-th cyclotomic number field, where $\Phi_m(X)$ is the m-th cyclotomic polynomial. The ring R has rank $n = \varphi(m)$, where $\varphi(\cdot)$ is Euler's totient function. We consider the lattice $\Lambda_m = \sigma(R)$, where σ denotes the canonical embedding of R into the complex numbers. Our goal is to characterize the error incurred when performing lattice-based rounding over Λ_m.

Lattice-based rounding and CVP. By *lattice-based rounding*, we refer to the problem of mapping an arbitrary target vector $\boldsymbol{t} \in \mathbb{R}^n$ to a nearby lattice point, that is, (approximately) solving the Closest-Vector Problem (CVP). Unlike the bounded-distance decoding problem that frequently arises in lattice-based cryptography, we make no assumption that the target lies close to the lattice. Instead, our objective is to understand the rounding error induced by Λ_m, both in the worst case and on average.

The geometry governing lattice-based rounding is captured by the Voronoi tessellation of the lattice. The Voronoi cell of a lattice point $\boldsymbol{v}$ consists of all points at least as close to $\boldsymbol{v}$ as to any other lattice point, and therefore describes exactly the region that rounds to $\boldsymbol{v}$ under exact CVP. For power-of-two cyclotomics, the Voronoi cell is simply a hypercube, whose rounding properties are well-understood but comparatively poor. For general composite values of m, however, the Voronoi cell has a much richer geometry, leaving open the possibility of substantially improved rounding properties. Our main results show that this is indeed the case: non-power-of-two cyclotomic lattices admit provably smaller worst- and average-case rounding error than their power-of-two counterparts.

Worst-case rounding error. The rounding error is the Euclidean distance between the target vector $\boldsymbol{t}$ and the lattice point $\boldsymbol{v}$ returned by the rounding algorithm. When exact CVP is solved, this is simply the distance from $\boldsymbol{t}$ to its nearest lattice point. The maximum rounding error over all target vectors is characterized by the covering radius $\rho(\Lambda_m)$. For power-of-two cyclotomics, the covering radius is exactly $n/2$. The exact covering radius is also known when $m = p^k$ is a prime power [8].

Average-case rounding error. While the covering radius characterizes worst-case behavior, many applications are governed by the average rounding error. We therefore study the case in which the target vector is drawn uniformly from a fundamental region of the lattice, corresponding to a uniformly random lattice coset. The resulting mean-squared error (MSE) measures the expected squared rounding error and serves as a natural metric for applications such as approximate HE, to be discussed later. Beyond the expectation, we also seek concentration bounds that show the rounding error is tightly concentrated around its mean with high probability.

Exact versus approximate CVP algorithms. For general lattices, exact CVP is NP-hard [14], and it remains NP-hard to approximate within almost-polynomial factors [11]. Moreover, this hardness can persist in certain preprocessing settings, where the lattice is fixed in advance and arbitrary preprocessing is allowed before the target is given [17, 22]. These results, however, do not rule out the possibility of efficient algorithms for highly structured lattice families. Table 1 summarizes the known CVP algorithms over Λ_m for several choices of m and their runtimes. As can be seen from the table, efficient exact algorithms are known only for certain families of cyclotomics. Furthermore, in our motivating applications, we require runtimes of $O(n \log n)$ in order to remain competitive with the power-of-two case. Thus, we also consider efficient approximate-CVP algorithms.

Babai's nearest-plane algorithm. Babai's nearest-plane algorithm [2] is one of the most well-known algorithms for approximate CVP. While for general lattices equipped with LLL-reduced bases it achieves only a worst-case approximation factor of $2^{n/2}$, it performs significantly better for cyclotomic lattices, which admit bases that are close to orthogonal. A straightforward implementation of Babai's algorithm runs in time $O(n^2)$, whereas the tensor-product structure of cyclotomic lattices allows this runtime to be reduced to $O\left(n \cdot \varphi\left(\mathrm{rad}(m)\right)\right)$, where $\mathrm{rad}(m)$ is the product of the distinct prime factors of m. In particular, this runtime is $O(n \log n)$ whenever $\varphi\left(\mathrm{rad}(m)\right) \in O(\log n)$.

Applications. Lattice-based rounding also has a direct motivation from HE. Ring-based HE schemes are typically built from the R-LWE assumption over quotient rings $R_q = R/qR$. Ciphertexts contain small error terms, and homomorphic operations increase these errors. Operations such as modulus switching and rescaling are used to control this growth, but they involve scaling operations that do not generally preserve the ring. Consequently, one must round the resulting point back to R, or equivalently to the embedded lattice Λ_m. This makes the rounding step a CVP problem over a cyclotomic lattice, and its error contributes directly to the scheme's noise budget.

In the approximate HE scheme of Cheon et al. (known as CKKS) [7], lattice-based rounding appears even more explicitly. The encoding procedure maps an arbitrary element of $\mathbb{R}^n$ to a nearby lattice point in Λ_m by performing a rounding step. Thus, improvements in the geometry of Λ_m or in the associated rounding algorithm directly translate into improvements in the encoding error.

Crucially, we observe that while rounding to the nearest point in R for power-of-two cyclotomics amounts to independently rounding each coordinate in the *polynomial representation* to the nearest integer, the same is no longer true for non-power-of-two cyclotomics, where exact rounding requires solving a nontrivial instance of CVP over the lattice Λ_m.

1.1 Efficient Arithmetic for non-power-of-two Cyclotomics

We now return to the computational cost of working over non-power-of-two cyclotomics. For power-of-two cyclotomics, ring multiplication can be performed in an evaluation domain in $O(n \log n)$ time using the usual power-of-two FFT/NTT. This efficiency is not restricted to the power-of-two case.

For arbitrary m, the fast cyclotomic transforms described in the Ring-LWE Toolkit of Lyubashevsky, Peikert, and Regev [25] exploit the tensor-product decomposition across the prime-power factors of m and run in $O(m \log m)$ time. This becomes $O(n \log n)$ for families with a bounded number of distinct prime divisors, such as $m = 2^s p^k q^\ell$. In particular, for $m = 2^s 3^k$, mixed-radix FFT techniques achieve $O(n \log n)$ complexity [26]. Such non-power-of-two cyclotomic families have also been considered in cryptographic constructions and implementations [12, 27]. Thus, broad families of non-power-of-two cyclotomics retain efficient arithmetic while offering the improved rounding geometry discussed above.

1.2 Our Contributions

Our contributions are threefold. First, we develop new worst-case and average-case analyses of lattice-based rounding for both general lattices and cyclotomic lattices. Second, we apply these results to approximate HE schemes instantiated over non-power-of-two cyclotomics, demonstrating how improved rounding leads to reduced error in key operations. Finally, we perform an end-to-end concrete parameter analysis, showing that the additional flexibility afforded by non-power-of-two cyclotomics enables reductions in both the lattice dimension and the modulus while maintaining a target security and the same output precision.

Analysis of lattice-based rounding error. We first develop a general analysis of lattice-based rounding error. For the worst-case analysis, we derive upper bounds on the covering radius of cyclotomic lattices for arbitrary cyclotomic indices. When expressed as a function of the dimension $n = \varphi(m)$, our upper bound is strictly smaller for non-power-of-two cyclotomics than the $n/2$ covering radius achieved by power-of-two cyclotomics. We define a reduction factor $\mathrm{Red}(m)$ capturing this improvement and derive a closed-form expression for it (Proposition 1). Defining $\mathrm{Gain}(m) = 1/\mathrm{Red}(m)$, we further prove that $\mathrm{Gain}(m) \in O(\sqrt{\log m})$ (see Lemma 2).

We then provide a new analysis of Babai's algorithm over general lattices, including a characterization of the distribution of its rounding error. We informally summarize below the main consequence of Lemma 4 and Theorem 2.

Theorem 1 (Informal). *Let the target vector $\boldsymbol{t}$ be sampled uniformly from a fundamental region of a lattice. Then, when expressed in the Gram–Schmidt basis, the rounding error produced by Babai's nearest-plane algorithm has independent coordinates, each uniformly distributed in* $[-0.5, 0.5)$.

This analysis yields an average-case characterization of the rounding error of Babai's algorithm over arbitrary lattices in terms of its MSE; see Corollary 2. We then specialize this result to cyclotomic lattices to obtain the explicit MSE in Corollary 3, and further establish the concentration bound of Proposition 2. Since the rounding error of Babai's algorithm upper bounds the rounding error of exact CVP, our analysis also yields corresponding bounds for exact lattice-based rounding.

Applications to approximate HE. We next apply our analysis to approximate HE over arbitrary cyclotomics. We identify two operations in which lattice-based rounding naturally arises: encoding and rescaling (or modulus switching). We propose replacing coordinate-wise rounding in the polynomial domain with lattice-based rounding in the canonical embedding. This rounding can be performed either exactly or approximately using Babai's nearest-plane algorithm, and our error analysis applies in both settings. Specifically, rescaling is typically modeled as introducing independent uniform error in $[-0.5, 0.5]$ in each coordinate with respect to the polynomial basis, thus leading to a uniform target vector over the corresponding fundamental region. This is precisely the setting in which our average-case and concentration bounds for lattice-based rounding using Babai's algorithm hold (see Informal Theorem 1).

Although lattice-based rounding is typically more expensive than coordinate-wise rounding, we observe that rounding algorithms that take time $O(n \log n)$ do not asymptotically increase the cost of the existing HE operations. In particular, in standard implementations, homomorphic multiplication and relinearization can be performed in the NTT domain, but the rounding step required during rescaling or modulus switching is performed in the polynomial domain. Consequently, each such operation requires a conversion from the NTT domain to the polynomial domain and back, incurring $O(n \log n)$ additional work. We instead propose to remain in the canonical embedding throughout: ring multiplication and relinearization take time $O(n)$, and reduction modulo q can be performed in $O(m \log m)$ time (and in $O(n \log n)$ time for m of certain forms discussed in Section 1.1). Moreover, as discussed above, Babai's nearest-plane algorithm runs in $O(n \log n)$ time for cyclotomic indices m satisfying $\phi(\mathrm{rad}(m)) \in O(\log n)$. Thus, replacing coordinate-wise rounding with lattice-based rounding does not asymptotically increase the cost of these operations for these cyclotomic indices.

To exploit the geometric advantages of arbitrary cyclotomics discussed above, we formulate the resulting error guarantees in terms of ℓ_2 error. This choice is particularly natural in applications such as machine learning, where convergence guarantees are typically expressed in terms of the ℓ_2 distance between the output model parameters and an optimal solution. We then analyze how the improved rounding guarantees translate into improved precision in approximate HE.

Concrete evaluation. Finally, we perform an end-to-end concrete comparison of approximate HE instantiated with power-of-two and non-power-of-two cyclotomics. The flexibility afforded by arbitrary cyclotomics allows us to select the lattice dimension more freely, enabling smaller dimensions while maintaining a target security level. Moreover, our improved noise analysis based on ℓ_2 error, combined with lattice-based rounding, enables us to reduce the ciphertext modulus relative to what can be achieved using ℓ_∞ error analysis and coordinate-based rounding, while maintaining the same target precision. Together, these improvements yield preferable security–precision tradeoffs for approximate HE.

For example, for the power circuit $f(x) = x^{16}$ at 10 bits of precision, Table 2 shows that replacing $m = 2^{15}$ by $m = 2^{12} \cdot 7$ reduces the dimension by 25% and the required modulus by about one bit, while preserving the target precision and security level.

1.3 Related Work

CVP and lattice-based rounding over cyclotomics. CVP over cyclotomic lattices Λ_m has been studied previously. Van Woerden [31], building on [28], gives efficient exact-CVP algorithms for certain cyclotomic orders, and Ducas and van Woerden [13] extend these results to broader classes; see Table 1. These works focus primarily on exact CVP for specific families of indices. In contrast, we consider CVP over general cyclotomics, characterize its worst- and average-case rounding error, and apply these results to approximate HE over non-power-of-two cyclotomics.

Cryptography over non-power-of-two cyclotomics. The use of non-power-of-two cyclotomics in lattice-based cryptography has been studied in several contexts. Lyubashevsky, Peikert, and Regev [24] introduced Ring-LWE over arbitrary cyclotomic rings and established its connection to hard problems on ideal lattices; their Ring-LWE Toolkit [25] further develops efficient arithmetic beyond the power-of-two setting. Ducas and Durmus [12] study Polynomial-LWE over arbitrary cyclotomics and conditions preserving security in the polynomial representation. More recently, Di Giusto and Marcolla [10] investigate non-power-of-two cyclotomics in HE, showing that the additional flexibility in parameter selection can yield more efficient instantiations. Our focus is complementary: we exploit the geometry of the corresponding cyclotomic lattices to improve lattice-based rounding while retaining efficient arithmetic.

Approximate HE and CKKS error analysis. The CKKS scheme [7] enables homomorphic computation over complex numbers by explicitly incorporating approximation error into its encoding and evaluation procedures. Subsequent work developed efficient RNS variants and analyzed RNS-based rescaling [6], reduced errors arising from encryption, key switching, and rescaling [18], and improved the precision of CKKS bootstrapping [19, 20]. In contrast, we focus on the rounding operation in encoding and rescaling, replacing coordinate-wise rounding with lattice-based rounding in the canonical embedding to exploit the geometry of

non-power-of-two cyclotomic lattices. This approach is also expected to remain compatible with RNS modulus decomposition: after lattice-based rounding, reduction modulo $q\Lambda_m$ is equivalent to reduction modulo qR, which can be implemented directly using the RNS decomposition of q.

2 Preliminaries

We begin by introducing the notation used throughout the paper.

Numbers, vectors, and norms. $\mathbb{N}$ denotes the natural numbers, $\mathbb{Q}, \mathbb{R}$, and $\mathbb{C}$ are the fields of rational, real, and complex numbers, respectively. $\mathbb{Z}$ is the ring of integers, and for $a \in \mathbb{Z}_{>0}$ we let $\mathbb{Z}_a = \mathbb{Z}/a\mathbb{Z}$, and $[a] = \{0, 1, \ldots, a-1\}$. The radical of a positive integer m, denoted $\mathrm{rad}(m)$, is the product of all distinct primes dividing m. We denote $\lceil x \rfloor$ as rounding to the unique nearest integer, i.e., $x - \lceil x \rfloor \in [-1/2, 1/2)$.

For a vector $\boldsymbol{x}$ over $\mathbb{R}$ or $\mathbb{C}$, define the ℓ_2-norm as $\|\boldsymbol{x}\|_2 = \left(\sum_i |x_i|^2\right)^{1/2}$, and the ℓ_∞-norm as $\|\boldsymbol{x}\|_\infty = \max_i |x_i|$.

Gram–Schmidt orthogonalization. Viewing B as a matrix whose columns are linearly independent vectors $\{\boldsymbol{b}_j\}_{j\in[n]}$, the Gram–Schmidt orthogonalization $\widetilde{B} = \{\widetilde{\boldsymbol{b}}_j\}_{j\in[n]}$ is defined iteratively as follows. First, $\widetilde{\boldsymbol{b}}_0 = \boldsymbol{b}_0$. Then, for $j = 1, 2, \ldots, n-1$, the vector $\widetilde{\boldsymbol{b}}_j$ is the component of $\boldsymbol{b}_j$ orthogonal to the linear span of $\boldsymbol{b}_0, \ldots, \boldsymbol{b}_{j-1}$:

$$\widetilde{\boldsymbol{b}}_j = \boldsymbol{b}_j - \sum_{k\in[j]} \frac{\langle \boldsymbol{b}_j, \widetilde{\boldsymbol{b}}_k \rangle}{\|\widetilde{\boldsymbol{b}}_k\|_2^2} \widetilde{\boldsymbol{b}}_k.$$

2.1 Lattices and CVP

We define full-rank *lattices* and review some basic properties.

Definition 1. *A full-rank lattice is generated as the set of all integer linear combinations of n linearly independent basis vectors $B = \{\boldsymbol{b}_j\}_{j\in[n]}$:*

$$\Lambda = \mathcal{L}(B) = \left\{ \sum_{j\in[n]} z_j \boldsymbol{b}_j : z_j \in \mathbb{Z} \right\}.$$

We say that Λ has rank n and generator matrix B.

Two matrices B and B' generate the same lattice if (and only if) there exists a unimodular matrix U (that is, U has integer entries and determinant ± 1) so that $BU = B'$. The *determinant* or covolume of a lattice $\Lambda = \mathcal{L}(B)$ is $|\det(B)|$ and is independent of the choice of the basis B. The above assumes the basis is square. It can be generalized to $\sqrt{\det(B^T B)}$ for non-square basis. Geometrically, $\det(B)$ is the n-dimensional volume, or hypervolume, of one repeating cell in a tiling of $\mathrm{span}_{\mathbb{R}}(\Lambda)$ by lattice translates.

A *fundamental domain*, or fundamental region, for Λ is such a cell: a measurable set $F \subseteq \mathrm{span}_{\mathbb{R}}(\Lambda)$ whose translates $\boldsymbol{\lambda} + F,\ \boldsymbol{\lambda} \in \Lambda$, tessellate $\mathrm{span}_{\mathbb{R}}(\Lambda)$, up to boundary overlaps. Equivalently, F contains one representative of each coset modulo Λ, up to boundary points. Fundamental domains are therefore not unique, and different choices are useful for different purposes. We will use two main examples throughout this work. Given a basis B, the associated basis-dependent fundamental domain is the *fundamental parallelepiped*,

$$\mathcal{P}(B) := B[0,1)^n = \{B\boldsymbol{x} : \boldsymbol{x} \in [0,1)^n\}.$$

We will also use its symmetric version $\mathcal{P}_{1/2}(B) := B[-1/2, 1/2)^n$. These cells are adapted to basis coordinates, and hence to procedures such as coefficient-wise rounding. The second example is the basis-independent *Voronoi cell*,

$$\mathcal{V}(\Lambda) = \{\boldsymbol{x} \in \mathrm{span}_{\mathbb{R}}(\Lambda) : \|\boldsymbol{x}\|_2 \leq \|\boldsymbol{x} - \boldsymbol{\lambda}\|_2 \text{ for all } \boldsymbol{\lambda} \in \Lambda\}.$$

It consists of all points at least as close to $\mathbf{0}$ as to any other lattice point, and is the natural cell for exact nearest-neighbor rounding: a target is mapped to the lattice point whose Voronoi cell contains it, up to boundary ties.

Finally, the *covering radius* $\rho(\Lambda)$ is the minimal distance such that any point in $\mathrm{span}_{\mathbb{R}}(\Lambda)$ is at distance at most $\rho(\Lambda)$ to a lattice point. Note that, by the tessellation of the Voronoi cells,

$$\rho(\Lambda) = \max_{\boldsymbol{x} \in \mathcal{V}(\Lambda)} \|\boldsymbol{x}\|_2 .$$

Thus the exact covering radius is a property of the Voronoi cell of the lattice.

In this work, we focus on a classical lattice problem: CVP.

Definition 2. *Given a basis B of a lattice $\Lambda = \mathcal{L}(B)$ and an arbitrary target point $\boldsymbol{t} \in \mathrm{span}_{\mathbb{R}}(\Lambda)$, the goal is to find a closest lattice point of Λ to $\boldsymbol{t}$, i.e.,*

$$\boldsymbol{\lambda}^\star \in \arg\min_{\boldsymbol{\lambda} \in \Lambda} \|\boldsymbol{t} - \boldsymbol{\lambda}\|_2 .$$

Note that any such $\boldsymbol{\lambda}^\star$ is called a closest vector to $\boldsymbol{t}$. Equivalently, $\boldsymbol{\lambda}^\star$ is such that $\boldsymbol{t} - \boldsymbol{\lambda}^\star \in \mathcal{V}(\Lambda)$, where $\mathcal{V}(\Lambda)$ is the Voronoi cell, as it consists of all points that have $\mathbf{0}$ as a closest vector. Since $\boldsymbol{t}$ is only assumed to be in the linear span of the lattice, it is generally not in the lattice Λ itself. Therefore, the covering radius gives a tight upper bound on the distance between a target and its closest lattice point.

Next, we study CVP over lattices arising from cyclotomic integer rings, known to be highly structured.

2.2 The Cyclotomic Lattice

We now introduce the family of lattices that will be studied throughout the rest of the paper. Starting from the m-th cyclotomic field, its ring of integers carries a natural algebraic structure; after applying the canonical embedding, this

ring becomes a full-rank lattice in a Euclidean space. These embedded rings of integers are the cyclotomic lattices on which we study CVP. The main advantage of this family is that its geometry is far from arbitrary: using a particular choice of a basis, the corresponding generator matrix has a structured tensor-product form determined by the prime-power factorization of m. This structure will allow us to analyze cyclotomic lattices through their prime-power components.

Cyclotomics. First we recall some facts about the cyclotomic fields. For a positive integer $m \in \mathbb{N}$, an m-th root of unity is an element ζ_m such that $\zeta_m^m = 1$. If $\zeta_m^k \neq 1$ for any $k < m$, then ζ is called primitive. Adjoining a primitive m-th root of unity ζ_m (i.e., an element of multiplicative order m) to the rationals, the field extension $K = \mathbb{Q}(\zeta_m)$ is obtained, called the m-th *cyclotomic number field*.

The m-th *cyclotomic polynomial* is the minimal polynomial of ζ_m

$$\Phi_m(X) = \prod_{i \in \mathbb{Z}_m^*} (X - \omega_m^i) \in \mathbb{Z}[X] \tag{1}$$

where $\mathbb{Z}_m^* = \{i \in [m] : \gcd(i, m) = 1\}$ and ω_m^i is any primitive m-th root of unity in $\mathbb{C}$ (for instance, $\omega_m = e^{2\pi\sqrt{-1}/m}$). The degree of $\Phi_m(X)$ is $n = |\mathbb{Z}_m^*| = \varphi(m)$, where $\varphi(\cdot)$ denotes Euler's totient function.

By naturally identifying $\zeta_m \mapsto X$, we get that K is isomorphic to $\mathbb{Q}[X]/(\Phi_m(X))$, and can be viewed as a vector space over $\mathbb{Q}$ with the power basis $\left(\zeta_m^j\right)_{j \in [n]}$. The ring of integers $R \subset K$ is $R = \mathbb{Z}[\zeta_m] \cong \mathbb{Z}[X]/(\Phi_m(X))$, having the power basis as a $\mathbb{Z}$-basis.

Let $m = \prod_\ell m_\ell$ be the prime-power factorization of m, and write $K_{m_\ell} = \mathbb{Q}(\zeta_{m_\ell})$. Then the m-th cyclotomic field decomposes as the tensor product of its prime-power cyclotomic components:

$$K \cong \bigotimes_\ell K_{m_\ell} \cong \bigotimes_\ell \mathbb{Q}[X_\ell]/(\Phi_{m_\ell}(X_\ell)),$$

induced by the correspondence $X_s \mapsto X^{\prod_{\ell \neq s} m_\ell} = X^{m/m_s}$. Indeed, if X represents ζ_m, then X^{m/m_s} represents ζ_m^{m/m_s}, which is a primitive m_s-th root of unity. Thus X^{m/m_s} naturally plays the role of the variable X_s in the m_s cyclotomic component. Note that the same decomposition also holds integrally for the cyclotomic integer ring, i.e., $R \cong \bigotimes_\ell \mathbb{Z}[X_\ell]/(\Phi_{m_\ell}(X_\ell))$.

The Canonical Embedding. We now pass from the algebraic object R to the Euclidean lattice on which CVP will be studied. This is done through the canonical embedding $\sigma : K \to \mathbb{C}^{\mathbb{Z}_m^*}$. Formally, the canonical embedding of $a \in K$ is the vector $\sigma(a) = (\sigma_i(a))_{i \in \mathbb{Z}_m^*}$ where each $\sigma_i : K \to \mathbb{C}$ is the embedding determined by $\sigma_i(\zeta_m) = \omega_m^i$, where ω_m is the fixed primitive m-th root of unity used in Equation (1). The embeddings occur in conjugate pairs, satisfying $\sigma_{m-i}(a) = \overline{\sigma_i(a)}$ for every $a \in K$. Therefore, σ maps into a subspace $H \subset \mathbb{C}^{\mathbb{Z}_m^*}$,

$$H = \left\{ \boldsymbol{x} \in \mathbb{C}^{\mathbb{Z}_m^*} : x_i = \overline{x_{m-i}},\ \forall i \in \mathbb{Z}_m^* \right\},$$

in which multiplication and addition are both component-wise, and the canonical inner product on K is the standard Hermitian inner product on $\mathbb{C}^{\mathbb{Z}_m^*}$.

The Embedded Lattice. The canonical embedding identifies the ring of integers R with a full-rank lattice $\Lambda_m := \sigma(R)$ in a Euclidean space. Since the prime-power decomposition of $K = \mathbb{Q}(\zeta_m)$ also holds integrally for R, the canonical embedding is compatible with this decomposition. Thus, the embedded lattice can be identified with the tensor product of the embedded prime-power lattices: $\sigma(R) \cong \bigotimes_\ell \sigma_{m_\ell}(R_{m_\ell})$. As a result, much of the geometry can be understood from the prime-power case and then extended by tensoring the components.

Following Lyubashevsky, Peikert, and Regev [25], we adopt the useful *powerful basis* of $R \cong \mathbb{Z}[X]/(\Phi_m(X))$. For prime-power m, this is the usual power basis; for general m, it is obtained by tensoring the prime-power bases along the prime-power factorization of m. Thus, in general, it does not coincide with the standard power basis $1, X, X^2, \ldots, X^{\varphi(m)-1}$.

Definition 3. *The powerful basis $\boldsymbol{p}_m$ of R is as follows.*

- *If m is a prime power, then $\boldsymbol{p}_m$ is the power basis $\boldsymbol{p}_m = \left(\zeta_m^j\right)_{j \in [\varphi(m)]}$, viewed as an ordered basis with entries in R.*
- *Let $m = \prod_\ell m_\ell$, where the m_ℓ's are powers of distinct primes. Then $\boldsymbol{p}_m = \bigotimes_\ell \boldsymbol{p}_{m_\ell}$, where $\boldsymbol{p}_{m_\ell}$ is the powerful basis of R_{m_ℓ}.*

The advantage of the powerful basis is that the geometry of its induced lattice, generated by $\mathrm{CRT}_m := \sigma(\boldsymbol{p}_m^T)$, can be neatly characterized as follows [25].

Claim. The vectors comprising CRT_m are all short, i.e., the ℓ_2 norm of each is $\sqrt{\varphi(m)} = \sqrt{n}$. Its largest and smallest singular values are, respectively, $s_{\max}(m) = \sqrt{\widehat{m}}$ and $s_{\min}(m) = \sqrt{m/\operatorname{rad}(m)}$, where $\widehat{m} = m$ if m is odd and $\widehat{m} = m/2$ otherwise. Further, CRT_m is close to orthogonal: the ratio of $s_{\max}(m)$ to $\sqrt{\varphi(m)}$ is $O\left(\sqrt{\log \log m}\right)$.

Lemma 1. *Let $m = p^k$ be a prime power, and let $m' = m/p = p^{k-1}$. Then*

$$\mathrm{CRT}_m = Q_m \cdot \left(\sqrt{m'}\, D_p \otimes I_{[m']}\right) \cdot \left(U_p \otimes I_{[m']}\right),$$

where Q_m is unitary, D_p is a diagonal matrix with j-th diagonal entry $\sqrt{(p-1) - j/p - j}$ for $j \in [\varphi(p)]$, and U_p is an upper unitriangular matrix with its (i,j)-th entry given by $-1/(p-i-1)$ for $0 \le i < j < \varphi(p)$.

Corollary 1. *Let $m = \prod_\ell m_\ell$ be the prime-power factorization of m, where the m_ℓ's are powers of distinct primes. Then, $\mathrm{CRT}_m = \bigotimes_\ell \mathrm{CRT}_{m_\ell}$.*

Finally, the sparsity of the embedded ring of integers is determined by the absolute discriminant of K, i.e., $\Delta_K = \det(\sigma(R))^2$, known to be upper bounded by n^n, where equality holds exactly when m is a *power-of-two*. Larger discriminant means a sparser lattice with larger fundamental cells, where the sparsest is attained by power-of-two cyclotomics. This suggests that non-power-of-two cyclotomics may be favorable for lattice-based rounding, studied next.

3 Worst-Case Error Analysis

Recall that the covering radius is the largest possible CVP error over all targets. It is therefore a natural quantity to study for cyclotomic lattices: it measures how far an arbitrary point in H can lie from the nearest point of the embedded ring of integers. In this sense, the covering radius captures the worst-case geometry of $\sigma(R)$ as a lattice and provides a first measure of its rounding performance.

3.1 An Upper Bound on the Covering Radius

Computing the covering radius exactly amounts to identifying the farthest points of the Voronoi cell (deep holes). For some cyclotomic lattices, this cell can be determined exactly. When the cyclotomic index m has no odd prime divisors, i.e., m is a power-of-two, the canonically embedded powerful basis is orthogonal, and hence the Voronoi cell is a rectangular box; if $n = \varphi(m)$, then $\rho(\Lambda_m) = \frac{n}{2}$, since the farthest points from the origin are the corners of this box.

More generally, for prime powers $m = p^k$, the cyclotomic lattice decomposes into orthogonal copies of the prime case $m = p$, which is isometric to the classical dual root lattice A^*_{p-1}, of known geometry. This gives the exact value [8]

$$\rho(\Lambda_{p^k}) = \frac{n}{p-1}\sqrt{\frac{p^2-1}{12}}. \tag{2}$$

For general composite m, however, the tensor-product structure is more complicated, leaving the Voronoi cell difficult to characterize. Instead, we derive a basis-dependent upper bound that is explicit for all cyclotomics. The bound follows from the following standard consequence of Babai's nearest-plane algorithm (discussed in detail in Section 4.2).

Claim. Let $B = \{\boldsymbol{b}_j\}_{j\in[n]}$ be the generator matrix of the lattice $\Lambda = \mathcal{L}(B)$ and let $\widetilde{B} = \{\widetilde{\boldsymbol{b}}_j\}_{j\in[n]}$ be the Gram–Schmidt orthogonalization of B. Then,

$$\rho(B) \le \frac{1}{2}\sqrt{\sum_{j\in[n]} \left\|\widetilde{\boldsymbol{b}}_j\right\|_2^2}. \tag{3}$$

3.2 The Covering Radius of Cyclotomics

In particular, for cyclotomic lattices with the canonically embedded powerful basis, B is CRT_m. Then, by Claim 2.2, it holds that $\sum_{j\in[n]} \left\|\boldsymbol{b}_j\right\|_2^2 = n^2$ and further that $\sum_{j\in[n]} \left\|\widetilde{\boldsymbol{b}}_j\right\|_2^2 < \sum_{j\in[n]} \left\|\boldsymbol{b}_j\right\|_2^2$ when m has an odd prime divisor. Incorporating both into Equation (3) yields

$$4\rho^2(B) \le \sum_{j\in[n]} \left\|\widetilde{\boldsymbol{b}}_j\right\|_2^2 \le n^2,$$

with the second inequality being tight exactly when m is a power-of-two.

We conclude that the maximum distance of a target to its nearest lattice point is smaller for non-power-of-two cyclotomics than for their power-of-two counterparts. The next proposition explicitly quantifies the relative reduction factor $\sqrt{\sum_{j\in[n]} \|\widetilde{\boldsymbol{b}}_j\|_2^2/n^2}$; its proof is given in Appendix A.1.

Proposition 1. *Let the matrix $B = \mathrm{CRT}_m$ be the canonical embedding of the powerful basis of $R = \mathbb{Z}[\zeta_m]$, and let $\widetilde{B} = \{\widetilde{\boldsymbol{b}}_j\}_{j\in[n]}$ be its Gram–Schmidt orthogonalization. Further, let the prime-power factorization of m be $\prod_\ell p_\ell^{k_\ell}$ and denote $n = \varphi(m) = \prod_\ell p_\ell^{k_\ell-1}(p_\ell - 1)$. Then, the sum of squared Gram–Schmidt lengths is*

$$S^2(m) := \sum_{j\in[n]} \|\widetilde{\boldsymbol{b}}_j\|_2^2 = \prod_\ell p_\ell^{2k_\ell-1}\,(p_\ell - H_{p_\ell})\,, \tag{4}$$

yielding the following rounding-error relative red*uction factor*

$$\mathrm{Red}(m) := S(m)/n = \prod_{prime\ p|m} \sqrt{p\,(p-H_p)/(p-1)^2}; \tag{5}$$

where $H_p := \sum_{k=1}^{p} \frac{1}{k}$ is the p-th harmonic number for a prime p dividing m.

Proposition 1 identifies the multiplicative factor by which non-power-of-two cyclotomics reduce the covering-radius upper bound relative to the power-of-two case. This factor, denoted by $\mathrm{Red}(m)$ and defined in Equation (5), depends only on the distinct prime divisors of m. Indeed, $\mathrm{Red}(m) = 1$ in the power-of-two case, whereas every odd prime divisor of m contributes a multiplicative factor strictly smaller than 1.

Another useful comparison comes from specializing our bound $\rho(B) \le \frac{1}{2}S(m)$, with $S^2(m)$ given in Equation (4), to the prime-power case $m = p^k$. In this case the exact covering radius is known from Equation (2), and the ratio is

$$\frac{\frac{1}{2}S(p^k)}{\rho(\Lambda_{p^k})} = \sqrt{\frac{3p\,(p-H_p)}{p^2-1}} \underset{p\to\infty}{\longrightarrow} \sqrt{3},$$

since $H_p \sim \log p$. Thus, for prime-power cyclotomics, the basis-dependent covering-radius bound is asymptotically within a factor $\sqrt{3}$ of the exact covering radius.

To complete the general picture, Figure 1a depicts the single-prime factors of $\mathrm{Red}(m)$, while Figure 1b depicts the value of $\mathrm{Red}(m)$ when m has the first k primes $2, 3, 5, \ldots, p_k$ as prime divisors. Finally, the asymptotic behavior of the improvement captured by the gain factor $\mathrm{Gain}(m) := 1/\mathrm{Red}(m)$ is given in the following lemma (see Appendix A.2 for the proof).

Lemma 2. *Let* $\mathrm{Red}(m) = \prod_{prime\ p|m} \sqrt{p\,(p-H_p)/(p-1)^2}$ *where* $H_p := \sum_{k=1}^{p} \frac{1}{k}$ *is the p-th harmonic number for a prime p dividing m. Then*

$$\mathrm{Gain}(m) := 1/\mathrm{Red}(m) = O\left(\sqrt{\log m}\right).$$

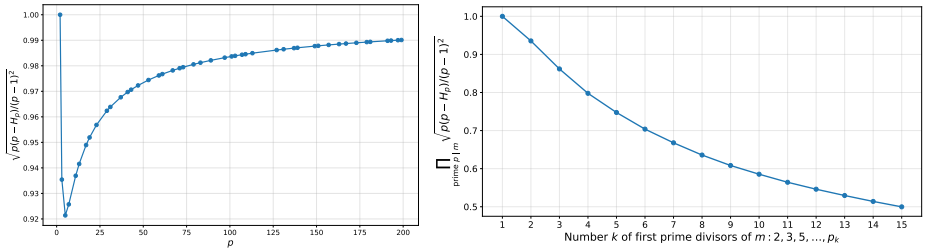

(a) $\mathrm{Red}(m)$ as a function of a single-prime factor p.

(b) $\mathrm{Red}(m)$ when m has the first k primes $2, 3, 5, \ldots, p_k$ as prime divisors.

Fig. 1: The behavior of the rounding-error relative reduction factor $\mathrm{Red}(m)$.

According to Figure 1, an odd prime divisor contributes the greatest relative reduction when it is small, and its contribution approaches 1 as the prime grows, while the cumulative reduction decreases steadily as more consecutive small prime divisors are included. Lemma 2 shows that the improvement remains controlled asymptotically, with $\mathrm{Gain}(m) = O\left(\sqrt{\log m}\right)$. This complements Claim 2.2, where the same departure from the power-of-two case appears through the bound $s_{\max}(m)/\sqrt{\varphi(m)} = O\left(\sqrt{\log\log m}\right)$. We now turn from geometric bounds to algorithms.

4 Algorithms

4.1 Exact Solution to CVP over Cyclotomics

For completeness, Table 1 summarizes known CVP algorithms for specific cyclotomic families, with references for the algorithmic details. The first row covers indices with at most one odd prime divisor, and the second covers indices with two odd prime divisors. For arbitrary m, CVP can be reduced to n/b independent b-dimensional instances. This yields both the generic exact blockwise algorithm detailed below and the $O(nb)$-time implementation of Babai's nearest-plane algorithm analyzed in Section 4.2.

The general composite case. Let $n = \varphi(m)$ and $b = \varphi\left(\mathrm{rad}(m)\right)$. By Corollary 1 and Lemma 1, the cyclotomic lattice Λ_m is isometric to the lattice generated by $I_{n/b} \otimes B_{\mathrm{blk}}$. Thus, exact CVP over Λ_m reduces to n/b independent b-dimensional instances over the lattice generated by B_{blk}.

For each block target $\boldsymbol{t}_i$, first reduce $B_{\mathrm{blk}}^{-1}\boldsymbol{t}_i$ modulo $\mathbb{Z}^b$ by setting $\boldsymbol{r}_i := B_{\mathrm{blk}}^{-1}\boldsymbol{t}_i - \left\lceil B_{\mathrm{blk}}^{-1}\boldsymbol{t}_i \right\rfloor \in [-1/2, 1/2)^b$. It then remains to solve $\min_{\boldsymbol{z}_i \in \mathbb{Z}^b} \left\| B_{\mathrm{blk}}\left(\boldsymbol{r}_i - \boldsymbol{z}_i\right)\right\|_2$. The following lemma bounds the resulting search space.

Lemma 3. *Consider CVP over the lattice $\mathcal{L}\left(B_{\mathrm{blk}}\right)$, with reduced target $B_{\mathrm{blk}}\boldsymbol{r}_i$ for $\boldsymbol{r}_i \in [-1/2, 1/2)^b$. It suffices to enumerate the vectors $\boldsymbol{z}_i \in \mathbb{Z}^b$ satisfying*

$$\left\|\boldsymbol{z}_i\right\|_2 \leq \frac{b}{2}\mathrm{Red}(m) + \frac{\sqrt{b}}{2},$$

Table 1: CVP algorithm complexities for cyclotomics. Here $p \geq q$ are distinct odd primes, $n = \varphi(m)$, $b = \varphi(\mathrm{rad}(m))$, and $c > 0$ is an absolute constant.

Type	**m's structure**	**Cost**	**Reference**
Exact	$m = 2^s p^k$ $s, k \geq 0, \quad s + k \geq 1$	$O(n)$	Reduction to the dual root lattice A^*_{p-1} [28].
Exact	$m = 2^s p^k q^\ell$ $s \geq 0, \quad k, \ell \geq 1$	$O\left(np^{q-1}\right)$	Ducas–van Woerden gluing algorithm [13].
Exact	Arbitrary m	$O\left(n\,(cb)^{b/2}\right)$	Blockwise exhaustive search; see Lemma 3.
Inexact	Arbitrary m	$O(nb)$	Babai's nearest-plane algorithm; see Section 4.2.

where $\mathrm{Red}(m) \leq 1$ *is given in Proposition 1. The number of candidates per block is at most* $(cb)^{b/2}$ *for some absolute constant* $c > 0$*. Consequently, exact CVP over* Λ_m *can be solved in* $O\left(n\,(cb)^{b/2}\right)$ *time.*

The proof is given in Appendix A.3. The algorithm is therefore useful when $b = \varphi(\mathrm{rad}(m))$ is small. When m is square-free, $b = n$, and the resulting complexity bound becomes super-exponential in the lattice rank.

4.2 Inexact Solution to CVP over Cyclotomics

In regimes where exact CVP is too costly, one may instead relax the problem and use an approximate solution, such as Babai's nearest-plane algorithm. While other alternatives exist, we focus on Babai's algorithm for two reasons. First, it can be especially fast for cyclotomics: on a block of dimension b, its cost is polynomial in b. Second, and more importantly for our purposes, its error geometry is explicit for the powerful basis considered here.

Furthermore, as we will explore in Section 5, it is possible to derive average-case and concentration bounds on the error of Babai's algorithm in the general case, i.e., for arbitrary lattices. For every instance, Babai's algorithm incurs error at least as large as the minimum achievable CVP error, namely the distance between the target and its nearest lattice point. Consequently, upper bounds on Babai's error also provide upper bounds on the average-case behavior and concentration of the exact CVP error.

Babai's Nearest-Plane Algorithm Babai's nearest-plane algorithm is a fast approximate algorithm for CVP [2]. Given an ordered basis, it processes the basis vectors from last to first, and at each step rounds the coefficient of the current residual in the corresponding Gram–Schmidt direction. As a result, at iteration i, the algorithm sets the coefficient of the updated residual in the direction $\widetilde{\boldsymbol{b}}_i$ to

lie in $[-1/2, 1/2)$. Later iterations only affect earlier Gram–Schmidt directions, and therefore do not change this coefficient. We next define the tessellating Babai region associated with the algorithm and then state the algorithm formally.

Definition 4. *Let $B = \{\boldsymbol{b}_j\}_{j\in[n]}$ be a basis and let $\widetilde{B} = \{\widetilde{\boldsymbol{b}}_j\}_{j\in[n]}$ be its Gram–Schmidt orthogonalization. The symmetric Babai fundamental domain of B, or Babai cell, is defined as*

$$\mathcal{P}_{1/2}(\widetilde{B}) := \widetilde{B}[-1/2, 1/2)^n = \left\{ \sum_{i\in[n]} c_i \widetilde{\boldsymbol{b}}_i : c_i \in [-1/2, 1/2) \right\}.$$

Algorithm 1 Babai Nearest-Plane Algorithm

Require: A basis $B = \{\boldsymbol{b}_j\}_{j\in[n]}$ of a lattice Λ, its Gram–Schmidt orthogonalization $\widetilde{B} = \{\widetilde{\boldsymbol{b}}_j\}_{j\in[n]}$, and a target $\boldsymbol{t} \in \mathrm{span}(\Lambda)$.
Ensure: $(\boldsymbol{v}, \boldsymbol{e})$, where $\boldsymbol{v} \in \Lambda$, $\boldsymbol{t} = \boldsymbol{v} + \boldsymbol{e}$, and $\boldsymbol{e} \in \mathcal{P}_{1/2}(\widetilde{B})$.
$\boldsymbol{e} \leftarrow \boldsymbol{t}$
$\boldsymbol{v} \leftarrow \boldsymbol{0}$
for $i = n-1, \ldots, 0$ **do**
  $k \leftarrow \left\lceil \dfrac{\langle \boldsymbol{e} \rangle \, \widetilde{\boldsymbol{b}}_i}{\|\widetilde{\boldsymbol{b}}_i\|_2^2} \right\rfloor$
  $\boldsymbol{e} \leftarrow \boldsymbol{e} - k\boldsymbol{b}_i$
  $\boldsymbol{v} \leftarrow \boldsymbol{v} + k\boldsymbol{b}_i$
end for
return $(\boldsymbol{v}, \boldsymbol{e})$

Complexity. Assuming the Gram–Schmidt orthogonalization is precomputed, Algorithm 1 performs n rounding steps. A direct implementation costs $O(n^2)$ arithmetic operations, since each step requires computing one Gram–Schmidt coefficient and updating an n-dimensional residual vector.

For cyclotomics, however, the runtime can potentially be faster by exploiting the block-diagonal structure arising from the tensor-product decomposition, as used by the algorithm solving CVP exactly for arbitrary cyclotomic index m in Section 4.1. That is, the cyclotomic lattice decomposes into n/b independent b-dimensional instances, where $b = \varphi(\mathrm{rad}(m))$. Applying Babai's algorithm independently on each block costs $O(b^2)$ arithmetic operations per block. Therefore the total cost is $O(b^2) \cdot n/b = O(nb)$.

Correctness. First note that Algorithm 1 never returns points that are too far away from the input target since each coordinate i of the output is within $\pm\frac{1}{2}\|\widetilde{\boldsymbol{b}}_i\|_2^2$ of that of $\boldsymbol{t}$. Namely, $\boldsymbol{e} \in \mathcal{P}_{1/2}(\widetilde{B})$, with associated norm $\|\boldsymbol{e}\|_2^2 \le \frac{1}{4}\sum_{i\in[n]} \|\widetilde{\boldsymbol{b}}_i\|_2^2$, as stated in Claim 3.1. Consequently, the quality of Babai's algorithm is basis-dependent. For an arbitrary basis, the Gram–Schmidt vectors

may be poorly conditioned, and the resulting Babai fundamental domain may be much larger than the Voronoi cell, which describes the exact-CVP error region. A standard way to obtain a worst-case approximation guarantee is therefore to first apply Lenstra–Lenstra–Lovász (LLL) reduction [21]. On an LLL-reduced basis, Babai's nearest-plane algorithm gives an exponential-factor approximation to CVP, commonly stated as a $2^{O(n)}$-approximation.

In the cyclotomic setting considered here, we do not apply LLL as preprocessing. The usual role of LLL is to make an arbitrary basis sufficiently reduced for Babai's algorithm to be effective. For the natural powerful bases of cyclotomic lattices, the relevant Gram–Schmidt geometry is already explicit from the CRT_m tensor structure, so Babai's error region can be analyzed directly. To illustrate this, the following example considers $m = 2^s 3^k$, where the block structure is two-dimensional, and contrasts Babai's result with the true CVP solution.

4.3 Example: CVP over the Cyclotomic $m = 2^s 3^k$

We now focus on CVP over cyclotomics where $m = 2^s 3^k$, and compare the performance of Algorithm 1 with exact CVP. According to Table 1, one may use the specialized $m = 2^s p^k$ algorithm or the generic blockwise algorithm. We use the latter to illustrate the cyclotomic block decomposition.

Here $n = \varphi(m) = 2^s 3^{k-1}$ and $b = \varphi(\mathrm{rad}(m)) = \varphi(6) = 2$, so the block-diagonal form consists of $n/2 = 2^{s-1}3^{k-1}$ identical 2D blocks. By Lemma 1, up to a unitary factor that preserves Euclidean distances, each block has columns

$$\boldsymbol{b}_{\mathsf{even}} = \sqrt{3^{k-1}2^{s-1}}\sqrt{2}\begin{bmatrix}1\\0\end{bmatrix}, \qquad \boldsymbol{b}_{\mathsf{odd}} = \sqrt{3^{k-1}2^{s-1}}\sqrt{2}\begin{bmatrix}-\frac{1}{2}\\ \frac{\sqrt{3}}{2}\end{bmatrix}. \tag{6}$$

Let B_{blk} denote the 2×2 matrix with columns $\boldsymbol{b}_{\mathsf{even}}$ and $\boldsymbol{b}_{\mathsf{odd}}$. According to Equation (8) in the proof of Lemma 3, we account for all integer vectors $\boldsymbol{z}_i = (z_0, z_1) \in \mathbb{Z}^2$ such that

$$\|\boldsymbol{z}_i - \boldsymbol{r}_i\|_2 \le \frac{\rho(B_{\mathrm{blk}})}{s_{\min}(B_{\mathrm{blk}})} \le \sqrt{\frac{7}{8}}, \quad \text{where} \quad \boldsymbol{r}_i \in [-1/2, 1/2)^2.$$

Thus, every coordinate of $\boldsymbol{z}_i$ lies in $\{\pm 1, 0\}$, requiring testing only 9 candidates.

Babai need not agree with exact CVP, even on this two-dimensional block. For the reduced target $\boldsymbol{r}_i = (-0.3, 0.4)$, the exact closest vector is $\boldsymbol{b}_{\mathsf{odd}}$, with squared error $0.54 \cdot 2^{s-1}3^{k-1}$, whereas Babai returns either 0 or $-\boldsymbol{b}_{\mathsf{even}}$, depending on tie-breaking, with squared error $0.74 \cdot 2^{s-1}3^{k-1}$.

Figure 2 illustrates this behavior through a Monte Carlo experiment over reduced targets in the symmetric fundamental parallelepiped generated by the unscaled vectors in (6). Babai differs from exact CVP precisely on the regions decoded to $(0, \pm 1)$, with disagreement probability approximately $1/12$, consistent with [3, Corollary 1].

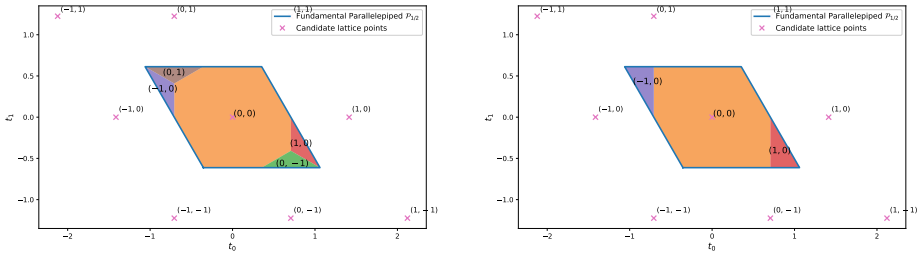

(a) Exact CVP solution.

(b) Babai's algorithm solution.

Fig. 2: Exact CVP vs. Babai decoding for the 2D block arising when $m = 2^s 3^k$.

5 Average-Case Error Analysis

Babai's nearest-plane algorithm is a fast approximate alternative for CVP, which is useful because no efficient exact algorithm is known for arbitrary cyclotomic lattices. Moreover, Babai's algorithm is amenable to an average-case analysis and to concentration bounds. We note that any average-case/concentration bounds achieved by Babai immediately imply upper bounds on the same quantities for exact algorithms. Before quantifying its performance for cyclotomic lattices, we first prove a general statement that holds for an arbitrary lattice basis.

5.1 The Probability Distribution of the Error

For Babai's algorithm, it is enough to study the distribution of the error when the target vector is reduced modulo the lattice: shifting the target by a lattice vector shifts Babai's output by the same lattice vector, and therefore leaves the error unchanged. This is formalized by the following lemma.

Lemma 4. *Let $\bar{\boldsymbol{t}} \in \Lambda + \boldsymbol{t}$, where $\boldsymbol{t} \in \mathcal{P}_{1/2}(B)$. Then the targets $\bar{\boldsymbol{t}}$ and $\boldsymbol{t}$ have the same error with respect to Babai's nearest-plane algorithm.*

Henceforth, we assume that the representative of the target in the symmetric fundamental parallelepiped $\mathcal{P}_{1/2}(B) = B[-1/2, 1/2)^n$ is uniformly distributed. Under this model, we show below that the Babai error is uniformly distributed over the Babai cell $\mathcal{P}_{1/2}(\widetilde{B})$, where $\widetilde{B}$ is the Gram–Schmidt orthogonalization of the lattice generator matrix B. The following theorem is proved in Appendix A.4 and sketched below.

Theorem 2. *Let a target $\boldsymbol{t}$ be uniformly distributed over $\mathcal{P}_{1/2}(B)$. Then the error produced by Babai's algorithm is uniformly distributed over $\mathcal{P}_{1/2}(\widetilde{B})$.*

Proof Sketch. Write the random target in basis coordinates, $\boldsymbol{t} = B\boldsymbol{x}$. Since $\boldsymbol{t}$ is uniform over $\mathcal{P}_{1/2}(B)$, the coefficients of $\boldsymbol{x}$ are independent and uniform on $[-1/2, 1/2)$. At step i, Babai has already processed the indices $j > i$, so the coefficient of the running error in the $\widetilde{\boldsymbol{b}}_i$-direction has the form $x_i + f_i\left(\{x_j\}_{j>i}\right)$,

where $f_i\left(\{x_j\}_{j>i}\right)$ denotes a function that depends only on $\{x_j\}_{j>i}$. After rounding (Step 4 in Algorithm 1), the remaining coefficient of the final error in this direction is $x_i + f_i\left(\{x_j\}_{j>i}\right)$ reduced modulo 1 back to $[-1/2, 1/2)$, and is therefore still uniform on $[-1/2, 1/2)$. Proceeding backward for all i shows that the Gram–Schmidt error coefficients are independent and uniform, hence the error is uniform over $\mathcal{P}_{1/2}(\widetilde{B})$.

5.2 The Expectation and Concentration of the Error

Having characterized the distribution of the error produced by Babai's algorithm, we now quantify its average magnitude through the MSE. The MSE is a natural distortion measure in this setting, capturing the average squared Euclidean norm of the rounding error, namely the expected squared distance between the target and the lattice point returned by the algorithm.

As immediate consequences of Theorem 2, we derive the MSE for an arbitrary lattice basis and then specialize the resulting expression to cyclotomic lattices.

Corollary 2. *Let a target $\boldsymbol{t}$ be uniformly distributed over $\mathcal{P}_{1/2}(B)$. Then the MSE of Babai's nearest-plane algorithm is*

$$\mathbb{E}\Big[\|\boldsymbol{e}\|_2^2\Big] = \frac{1}{12}\sum_{i\in[n]} \|\widetilde{\boldsymbol{b}}_i\|_2^2.$$

Proof. The error coefficients in the Gram–Schmidt directions are i.i.d. uniform over $[-1/2, 1/2)$ with second moment $1/12$. The corollary then follows from the orthogonality of the Gram–Schmidt vectors and linearity of expectation.

In particular, for cyclotomic lattices we have $B = \mathrm{CRT}_m$, for which the sum of the squared Gram–Schmidt lengths $\sum_{j\in[n]} \|\widetilde{\boldsymbol{b}}_j\|_2^2 := S^2(m) = n^2\mathrm{Red}^2(m)$ where $n = \varphi(m)$, was derived in Proposition 1.

Corollary 3. *For an integer $m \geq 2$, let CRT_m be the basis of the cyclotomic lattice of index m and let $n = \varphi(m)$. Consider a target $\boldsymbol{t}$ uniformly distributed over $\mathcal{P}_{1/2}(\mathrm{CRT}_m)$. Then the MSE of Babai's nearest-plane algorithm is*

$$\mathbb{E}\Big[\|\boldsymbol{e}\|_2^2\Big] = \frac{1}{12}n^2\mathrm{Red}^2(m) = \frac{n^2}{12}\prod_{prime\ p|m} \frac{p\,(p-H_p)}{(p-1)^2}, \qquad H_p := \sum_{k=1}^{p}\frac{1}{k}.$$

Finally, the following concentration bound turns the MSE estimate into a high-probability statement, showing that the actual squared error is very unlikely to exceed its mean by more than a small multiplicative factor. The proof can be found in Appendix A.5.

Proposition 2. *Let $m \geq 2$, let $n = \varphi(m)$, and let $B = \mathrm{CRT}_m$. Assume a target $\boldsymbol{t} \sim \mathrm{Unif}\big(\mathcal{P}_{1/2}(B)\big)$, and let $\boldsymbol{e}$ be the error produced by Babai's nearest-plane algorithm. Then, for every $\varepsilon > 0$,*

$$\Pr\Big(\|\boldsymbol{e}\|_2^2 \geq (1+\varepsilon)\mathbb{E}\Big[\|\boldsymbol{e}\|_2^2\Big]\Big) \leq \exp\left(-n\frac{\mathrm{Red}^2(m)\varepsilon^2}{\frac{8}{5}+\frac{4}{3}\varepsilon}\right),\ \mathrm{Red}^2(m) := \prod_{prime\ p|m}\frac{p\,(p-H_p)}{(p-1)^2}.$$

Proposition 2 shows that Babai's squared error is highly concentrated around its MSE, with failure probability decaying exponentially in $n\mathrm{Red}^2(m)\varepsilon^2$. Using Lemma 2, which gives $1/\mathrm{Red}(m) = O\left(\sqrt{\log m}\right)$, we have that the exponent is at least of order $\varepsilon^2 n/\log n$. Hence, for fixed $\varepsilon > 0$, the failure probability is negligible in the lattice rank n.

6 Applications

Homomorphic encryption (HE) schemes are typically designed over power-of-two cyclotomics, for which coefficient-wise rounding is exact nearest-point rounding. The previous sections studied the case of non-power-of-two cyclotomics, where rounding becomes a CVP instance over $\sigma(R)$, with different error behavior and algorithmic cost.

6.1 Homomorphic Encryption over Non-power-of-two Cyclotomics

Changing the cyclotomic index also affects the underlying R-LWE error distribution and the representation used for efficient ring operations. We discuss the dual-to-primal Gaussian scaling and the tradeoffs between the NTT and canonical-embedding domains below.

R-LWE error distribution and sampling. The R-LWE error over the dual $R^\vee$ is modeled as a spherical Gaussian in the canonical embedding [23]. To pass to the primal ring R, one can use the fact that $\widehat{m}R^\vee \subseteq R$ where $\widehat{m} = m$ if m is odd and $\widehat{m} = m/2$ otherwise [12, 25], and replace the Gaussian parameter σ with $\widehat{m}\sigma$. Furthermore, unlike in the power-of-two case, for non-power-of-two cyclotomics, sampling independent equal-variance coefficients in the polynomial basis does not produce a spherical Gaussian in the canonical embedding. For this purpose, one may use the sampling method of Ducas and Durmus [12], which samples in an auxiliary polynomial ring and then reduces modulo Φ_m.

Homomorphic evaluation in canonical coordinates. Homomorphic evaluation requires efficient multiplication in $R_q = \mathbb{Z}_q[X]/\Phi_m(X)$. Since coefficient-domain multiplication is expensive, implementations usually work in an NTT evaluation domain when q is a prime with $q \equiv 1 \pmod{m}$. In this domain, ring multiplication is coordinate-wise and modular arithmetic is native. By contrast, rounding is not naturally performed in this representation, so rescaling typically requires switching to the coefficient representation for rounding and then returning to the NTT domain.

Our perspective is to *formulate homomorphic evaluation directly in the canonical embedding*, which we call the complex CRT domain. In this representation, multiplication is again coordinate-wise, now over complex coordinates, and rescaling can be followed immediately by rounding to the embedded lattice $\sigma(R)$. This domain is also not tied to the existence of an NTT prime $q \equiv 1 \pmod{m}$, and it is the natural setting for canonical-norm noise analysis.

The price is that modular reduction is less efficient in this representation. In the canonical embedding, reduction modulo q corresponds to reduction modulo the sublattice $\sigma(qR) = q\sigma(R)$. In general, this can be implemented by returning to coefficient representation, reducing modulo q, and transforming back.

Cost of switching between domains. The CRT and NTT representations are analogous: both are evaluation domains. The complex CRT evaluates a ring element at the complex primitive m-th roots of unity, while the NTT evaluates it modulo q at suitable roots of unity in the finite field or residue ring used by the scheme. Thus both representation changes are Fourier-type transforms.

For power-of-two cyclotomics, this CRT/NTT conversion is the usual power-of-two FFT and costs $O(n \log n)$. The same asymptotic cost holds for $m = 2^s 3^k$ [26], where the transform is a mixed-radix FFT with radix-2 and radix-3 stages. More generally, the fast transform algorithms of [25] apply to arbitrary cyclotomic indices by decomposing the transform into prime-power factors and handling prime-length base cases using Rader, Winograd, or Bluestein-type methods. This gives $O(m \log m)$ time and $O(\log m)$ parallel depth.

Rounding in R versus rounding in $\sigma(R)$. Although rounding operations in ring-based HE are often written algebraically as rounding back to R, applying the canonical embedding turns the same operation into rounding back to the lattice $\sigma(R)$. Indeed, choosing a rounded value $r \in R$ for a scaled element x in the real span of R is equivalent to choosing the lattice point $\sigma(r) \in \sigma(R)$ for the embedded target $\sigma(x)$. For power-of-two cyclotomics, coefficient-wise rounding agrees with CVP in the canonical embedding.

For non-power-of-two cyclotomics, these two notions no longer coincide; using exact CVP or Babai's nearest-plane algorithm corresponds to rounding with respect to the Euclidean geometry of the canonical embedding. To see this concretely, consider the non-power-of-two cyclotomic index $m = 2^s 3^k$ studied in Section 4.3. After reduction modulo $\mathcal{P}_{1/2}(B_{\mathrm{blk}})$, all basis coefficients lie in $[-1/2, 1/2)$. Hence coefficient-wise rounding maps the entire cell to $\mathbf{0}$. In contrast, neither exact CVP nor Babai's nearest-plane algorithm maps the entire cell to $\mathbf{0}$, as shown in Figures 2a and 2b, respectively.

Cost of rounding. Replacing coefficient-wise rounding by lattice-based rounding has an algorithmic cost. In coefficient representation, coordinate-wise rounding in the power-of-two case costs $O(n)$. In an NTT-based evaluation pipeline, however, performing such a rounding step requires leaving the NTT domain and then returning to it, so the full round-trip costs $O(n \log n)$.

As discussed in Section 4.2, for non-power-of-two cyclotomics, applying Babai's nearest-plane algorithm while exploiting the block structure of cyclotomics costs

$$O(nb) \leq O(n \log n) \quad \text{if} \quad b = \varphi\left(\mathrm{rad}(m)\right) \leq \log n.$$

Since $n = \varphi(m) \leq m$, this cost does not asymptotically dominate the $O(m \log m)$ cost of the domain changes required for efficient ring multiplication.

In settings where solving exact CVP is no more expensive than running Babai's algorithm, see Table 1, one can compute the exact nearest lattice point while still using the Babai analysis as a convenient upper bound on the error. Indeed, since exact CVP minimizes the Euclidean error for every target, the MSE and concentration bounds derived for Babai remain valid for exact CVP, although they are generally not tight.

6.2 Homomorphic Encryption Schemes

We focus on ring-based HE schemes built from the R-LWE assumption, such as BGV, BFV, and CKKS [5, 15, 7]. In BGV and BFV, noise growth is controlled by modulus switching, while in CKKS it is controlled by rescaling. Both operations involve division by a scaling factor and therefore require mapping the result back to R; in canonical coordinates, this is a lattice-rounding step over $\sigma(R)$, which we perform using exact CVP or Babai's nearest-plane algorithm. In this work, we focus on CKKS, which supports arithmetic on approximate numbers and in which rounding appears not only in rescaling but also in encoding.

CKKS over the Canonical Embedding At a high level, CKKS encodes complex vectors as elements of a cyclotomic ring of integers R, encrypts them using an R-LWE-based public key, and supports homomorphic addition and multiplication over a modulus chain $q_L \to q_{L-1} \to \cdots \to q_0$. We refer the reader to [7, Section 3.4] for a detailed description of the scheme. In the following, we record only the decryption-level behavior, in canonical-embedding coordinates, of the operations most relevant to the introduction and propagation of rounding noise: encoding, P-scaled encryption (a variant of the ordinary plain encryption), rescaling, and multiplication. We first introduce the auxiliary notation.

Coordinates and decryption convention. We use both ring coordinates and canonical-embedding coordinates. For a ring element $a \in R$, $\sigma(a) \in H$ denotes its canonical embedding. For a ciphertext vector $\boldsymbol{c} = (c_1, c_2) \in R_{q_\ell}^2$, $\sigma(\boldsymbol{c}) = (\sigma(c_1), \sigma(c_2)) \in H^2$. We say that $\boldsymbol{c}$ encrypts a ring plaintext $\mu \in R$ with noise e, with respect to a decryption vector $\boldsymbol{sk} \in R^2$, if $\langle \boldsymbol{c}, \boldsymbol{sk} \rangle = \mu + e \pmod{q_\ell}$. CKKS is an approximate scheme: the decrypted value $\mu + e$ is interpreted as an approximation to μ.

Lattice-based rounding. Consider the cyclotomic lattice $\Lambda_m = \sigma(R)$ and a target $\boldsymbol{t} \in \mathrm{span}_{\mathbb{R}}(\Lambda_m) = H$. We denote nearest-neighbor rounding over $\sigma(R)$ as

$$\lceil \boldsymbol{t} \rfloor_{\sigma(R)} \in \operatorname*{arg\,min}_{\lambda \in \sigma(R)} \|\boldsymbol{t} - \lambda\|_2 .$$

The notation $\lceil \cdot \rfloor_{\sigma(R)^2}$ denotes the corresponding product-lattice rounding in H^2, applied componentwise to ciphertext vectors.

– Encoding. CKKS encodes a message vector $\boldsymbol{z} \in \mathbb{C}^{n/2}$ at scale Δ (here we do not assume it is a Gaussian integer vector) by first mapping (expanding) it

to the canonical-embedding space H via π^{-1}, then scaling by Δ, and finally rounding to the lattice $\sigma(R)$. The encoded plaintext $\mu \in R$ is then

$$\sigma(\mu) = \left\lceil \Delta \cdot \pi^{-1}(\boldsymbol{z}) \right\rfloor_{\sigma(R)} := \Delta \cdot \pi^{-1}(\boldsymbol{z}) + \boldsymbol{e}_{\mathsf{ecd}},$$

where $\boldsymbol{e}_{\mathsf{ecd}} \in H$ is the encoding rounding-error vector. Evidently, encoding already contains a CVP instance over the lattice $\sigma(R)$.

– Plain encryption. For a plaintext $\mu \in R$, encryption outputs a fresh ciphertext $\boldsymbol{c} \in R_{q_L}^2$ satisfying $\langle \boldsymbol{c}, \boldsymbol{sk} \rangle = \mu + e_{\mathsf{clean}} \pmod{q_L}$, where e_{clean} is the fresh encryption noise.
We use a P-scaled encryption variant, common in CKKS implementations [4], which can reduce the initial noise. We first isolate its rescaling step.

– Rescaling. For a ciphertext $\boldsymbol{c} \in R_{q_\ell}^2$ at level ℓ, and for a lower level $\ell' < \ell$, rescaling is defined in canonical coordinates by

$$\sigma(\boldsymbol{c}') := \mathsf{RS}_{\ell \to \ell'}(\sigma(\boldsymbol{c})) = \left\lceil \frac{q_{\ell'}}{q_\ell} \sigma(\boldsymbol{c}) \right\rfloor_{\sigma(R)^2} \pmod{q_{\ell'} \sigma(R)^2}.$$

Here, reduction modulo $q_{\ell'}$ in the canonical embedding corresponds to reduction modulo the sublattice $q_{\ell'}\sigma(R)$. The embedded rounding error introduced by rescaling is

$$\boldsymbol{e}_{\mathsf{scale}} = \sigma(\boldsymbol{c}') - \frac{q_{\ell'}}{q_\ell} \sigma(\boldsymbol{c}) \in H^2.$$

Let $\boldsymbol{\tau}_{\mathsf{scale}}$ be the corresponding element in the real span of the product ring, defined by $\sigma(\boldsymbol{\tau}_{\mathsf{scale}}) = \boldsymbol{e}_{\mathsf{scale}}$. The decryption-level rounding error is $e_{\mathsf{scale}} = \langle \boldsymbol{\tau}_{\mathsf{scale}}, \boldsymbol{sk} \rangle$. Therefore, if $\langle \boldsymbol{c}, \boldsymbol{sk} \rangle = \mu + e \pmod{q_\ell}$, then

$$\langle \boldsymbol{c}', \boldsymbol{sk} \rangle = \frac{q_{\ell'}}{q_\ell} \mu + \left(\frac{q_{\ell'}}{q_\ell} e + e_{\mathsf{scale}} \right) \pmod{q_{\ell'}}. \tag{7}$$

Thus rescaling multiplies both the plaintext and the existing noise by $q_{\ell'}/q_\ell$, and adds the new rounding-noise term e_{scale}.

– P-scaled encryption. The P-scaled encryption variant first encrypts the scaled plaintext $P\mu$ modulo the larger modulus $P \cdot q_L$ and outputs $\boldsymbol{c} \in R_{P \cdot q_L}^2$ such that

$$\langle \boldsymbol{c}, \boldsymbol{sk} \rangle = P\mu + e_{\mathsf{clean}} \pmod{Pq_L}.$$

It then immediately rescales by P, returning to modulus q_L, obtaining $\sigma(\boldsymbol{c}') := \mathsf{RS}_{Pq_L \to q_L}(\sigma(\boldsymbol{c}))$. By the rescaling relation in Equation (7),

$$\langle \boldsymbol{c}', \boldsymbol{sk} \rangle = \mu + \left(P^{-1} e_{\mathsf{clean}} + e_{\mathsf{scale}} \right) \pmod{q_L}.$$

For large P, the clean encryption error becomes negligible compared with the CVP rounding error introduced by the immediate rescaling.

– Multiplication. Let $\boldsymbol{c}_i \in R_{q_\ell}^2$ encrypt μ_i with noise e_i, namely

$$\langle \boldsymbol{c}_i, \boldsymbol{sk} \rangle = \mu_i + e_i \pmod{q_\ell}, \qquad i = 1, 2.$$

Multiplication followed by relinearization and key switching outputs a ciphertext $\boldsymbol{c}_{\mathsf{mult}} \in R_{q_\ell}^2$ encrypting $\mu_1\mu_2$ with total multiplication noise e_{Mult}:

$$\langle \boldsymbol{c}_{\mathsf{mult}}, \boldsymbol{sk} \rangle = \mu_1\mu_2 + e_{\mathsf{Mult}} \pmod{q_\ell}.$$

6.3 Noise Estimates of Homomorphic Operations

As discussed in the previous section, CKKS is an approximate scheme: after homomorphic evaluation, the accumulated decryption error must remain small compared with the scale of the encoded message for the output to be meaningful. In this section, we estimate the noise terms introduced by encoding, encryption, rescaling, and multiplication.

Conventionally, CKKS noise bounds are stated using canonical ℓ_p-norms,

$$\|e\|_p^{\mathrm{can}} = \|\sigma(e)\|_p ,$$

which are well matched to the coordinate-wise multiplication structure of the canonical embedding. The most common choice is the canonical ℓ_∞-norm, which gives a worst-case coordinate-wise error bound. Since the error terms in the previous section were already formulated in canonical coordinates, passing to canonical norms is immediate.

In this work, we use the canonical ℓ_2-norm instead. This choice is natural because Babai's MSE analysis is Euclidean and because the ℓ_2-norm is the relevant metric in, e.g., CKKS-based machine learning on real- or complex-valued data, a central application of CKKS [29, 30]. In such settings, encrypted training yields model parameters that deviate from those obtained by the corresponding plaintext computation, and convergence guarantees are naturally expressed in squared ℓ_2-distance [32].

Rounding and sampling. We use the average-case rounding-error model underlying the MSE analysis of Babai's algorithm and apply its concentration bound to each CVP instance, with exact CVP incurring no larger error for any target. Let $n = \varphi(m)$. For every $\varepsilon > 0$, define

$$B_{\mathsf{rnd}} := \sqrt{1+\varepsilon}\,\frac{n\mathrm{Red}(m)}{\sqrt{12}}, \qquad \delta_{\mathsf{rnd}} := \exp\left(-n\frac{\varepsilon^2\mathrm{Red}(m)^2}{\frac{8}{5}+\frac{4}{3}\varepsilon}\right);$$

$$\text{where} \quad \mathrm{Red}(m) := \prod_{\text{prime } p\mid m}\sqrt{\frac{p\,(p-H_p)}{(p-1)^2}}, \qquad H_p := \sum_{k=1}^{p}\frac{1}{k}.$$

Thus, under this model, the canonical ℓ_2-norm of each rounding-error vector is at most B_{rnd}, except with probability at most δ_{rnd}. By default, we set $\varepsilon = 0.1$.

We assume that the secret key, ephemeral polynomial, and all error polynomials are sampled so that their canonical embeddings follow centered spherical Gaussian distributions with standard deviation $\widehat{m}\sigma$, where

$$\widehat{m} := \begin{cases} m, & m \text{ odd}, \\ m/2, & m \text{ even}, \end{cases} \qquad \sigma > 0.$$

For every such sampled element $a \in R$, $\|a\|_\infty^{\mathrm{can}} \le B_{\mathsf{G},\infty} := 6\widehat{m}\sigma$ except with probability at most $\delta_{\mathsf{G},\infty} := ne^{-18}$, by the standard Gaussian tail bound; see [16, Corollary 10]. The canonical ℓ_2-bound used in our analysis is stated next and proved in Appendix A.6.

Lemma 5. *Let $a \in R$ be such that its canonical embedding is a centered spherical Gaussian with standard deviation $\widehat{m}\sigma$. Then*

$$\|a\|_2^{\text{can}} \le B_{\mathsf{G},2}, \qquad B_{\mathsf{G},2} := 2\widehat{m}\sigma\sqrt{n},$$

except with probability at most $\delta_{\mathsf{G},2} := e^{-n/2}$.

Let $\boldsymbol{sk} = (1, s)$ denote the decryption vector. The following lemmas give high-probability upper bounds on the error terms introduced in the previous section; for simplicity, we omit the resulting combined failure probabilities. The bounds are obtained by adapting the proof steps of [7, Lemmas 1–3] to the ℓ_2-norm, using $\|\boldsymbol{x} \odot \boldsymbol{y}\|_2 \le \|\boldsymbol{x}\|_2 \|\boldsymbol{y}\|_\infty$ for any complex vectors $\boldsymbol{x}, \boldsymbol{y} \in H$, where $\odot$ denotes entry-wise multiplication.

Lemma 6 (Encoding). *Let $\mu \in R$ be the encoding of $\boldsymbol{z} \in \mathbb{C}^{n/2}$ at scale Δ, so that $\sigma(\mu) = \Delta\pi^{-1}(\boldsymbol{z}) + \boldsymbol{e}_{\mathsf{ecd}}$. Then $\|\boldsymbol{e}_{\mathsf{ecd}}\|_2 \le B_{\mathsf{rnd}}$.*

Lemma 7 (Plain encryption). *Let $\boldsymbol{c} \in R_{q_L}^2$ be a plain encryption of $\mu \in R$, satisfying $\langle \boldsymbol{c}, \boldsymbol{sk} \rangle = \mu + e_{\mathsf{clean}} \pmod{q_L}$. Then*

$$\|e_{\mathsf{clean}}\|_2^{\text{can}} \le B_{\mathsf{clean}}, \qquad B_{\mathsf{clean}} := 2B_{\mathsf{G},\infty}B_{\mathsf{G},2} + B_{\mathsf{G},2}.$$

Lemma 8 (Rescaling). *Let $\boldsymbol{\tau}_{\mathsf{scale}} = (\tau_0, \tau_1)$ be the rounding error introduced during rescaling, and let $e_{\mathsf{scale}} := \langle \boldsymbol{\tau}_{\mathsf{scale}}, \boldsymbol{sk} \rangle = \tau_0 + \tau_1 s$. Then*

$$\|e_{\mathsf{scale}}\|_2^{\text{can}} \le B_{\mathsf{scale}}, \qquad B_{\mathsf{scale}} := B_{\mathsf{rnd}}\left(1 + B_{\mathsf{G},\infty}\right).$$

The preceding bounds immediately yield the total encryption error for the P-scaled variant.

Corollary 4 (Encoding and P-scaled encryption). *Let $\mu \in R$ be the encoding of $\boldsymbol{z} \in \mathbb{C}^{n/2}$ at scale Δ, so that $\sigma(\mu) = \Delta\pi^{-1}(\boldsymbol{z}) + \boldsymbol{e}_{\mathsf{ecd}}$, and let $\boldsymbol{c} \in R_{q_L}^2$ be its P-scaled encryption, satisfying*

$$\langle \boldsymbol{c}, \boldsymbol{sk} \rangle = \mu + e_{\mathsf{enc}}^{(P)} \pmod{q_L}, \qquad e_{\mathsf{enc}}^{(P)} = P^{-1}e_{\mathsf{clean}} + e_{\mathsf{scale}}.$$

Then the total error relative to the expanded scaled message $\Delta\pi^{-1}(\boldsymbol{z})$ satisfies

$$\left\| \boldsymbol{e}_{\mathsf{ecd}} + \sigma\left(e_{\mathsf{enc}}^{(P)}\right) \right\|_2 \le B_{\mathsf{fresh}}, \qquad B_{\mathsf{fresh}} := B_{\mathsf{rnd}} + \frac{B_{\mathsf{clean}}}{P} + B_{\mathsf{scale}}.$$

As shown in the next section, for the HE parameters considered below, P is sufficiently large that B_{clean}/P is negligible. Hence, $B_{\mathsf{fresh}} \approx B_{\mathsf{rnd}} + B_{\mathsf{scale}}$.

Lemma 9 (Multiplication). *For $i = 1, 2$, let $\boldsymbol{c}_i \in R_{q_\ell}^2$ encode and encrypt $\boldsymbol{z}_i \in \mathbb{C}^{n/2}$ at scale Δ, with total error $\boldsymbol{e}_i \in H$, so that*

$$\sigma\left(\langle \boldsymbol{c}_i, \boldsymbol{sk} \rangle\right) = \Delta\pi^{-1}(\boldsymbol{z}_i) + \boldsymbol{e}_i \pmod{q_\ell \sigma(R)}.$$

Suppose that $\left\|\Delta\pi^{-1}(\boldsymbol{z}_i)\right\|_2 \le \nu_i$, and $\|\boldsymbol{e}_i\|_2 \le B_i$. Multiplication followed by relinearization and key switching outputs a ciphertext $\boldsymbol{c}_{\mathsf{mult}} \in R_{q_\ell}^2$ satisfying

$$\sigma\left(\langle \boldsymbol{c}_{\mathsf{mult}}, \boldsymbol{sk} \rangle\right) = \Delta^2\pi^{-1}(\boldsymbol{z}_1 \odot \boldsymbol{z}_2) + \boldsymbol{e}_{\mathsf{Mult}} \pmod{q_\ell \sigma(R)},$$

where $\boldsymbol{e}_{\mathsf{Mult}} \in H$ is the total multiplication error. Then, defining $B_{\mathsf{ks}}^{(2)} := 2\sqrt{3}\,\widehat{m}\sigma n$,

$$\|\boldsymbol{e}_{\mathsf{Mult}}\|_2 \le B_{\mathsf{Mult}}, \qquad B_{\mathsf{Mult}} := \nu_1 B_2 + \nu_2 B_1 + B_1 B_2 + \frac{q_\ell}{P} B_{\mathsf{ks}}^{(2)} + B_{\mathsf{scale}}.$$

Relative-error propagation. Following the original CKKS analysis, we measure the error of a ciphertext relative to its corresponding expanded scaled message, before decoding. Using the notation of Lemma 9 above, define

$$\beta_i := \frac{B_i}{\nu_i}, \qquad \left\|\Delta\pi^{-1}\left(\boldsymbol{z}_i\right)\right\|_2 \le \nu_i, \qquad \left\|\boldsymbol{e}_i\right\|_2 \le B_i, \qquad i = 1, 2.$$

We assume throughout the computation that the expanded scaled messages satisfy $\left\|\Delta\pi^{-1}\left(\boldsymbol{z}\right)\right\|_\infty \le p$, and hence $\left\|\Delta\pi^{-1}\left(\boldsymbol{z}\right)\right\|_2 \le p\sqrt{n}$. By Corollary 4, the relative error of a fresh P-scaled encryption is therefore $\beta_0 := \frac{B_{\mathsf{fresh}}}{p\sqrt{n}}$.

By Lemma 9, multiplication followed by relinearization, key switching, and immediate rescaling by p yields relative error

$$\beta' = \frac{B_{\mathsf{Mult}}/p + B_{\mathsf{scale}}}{\nu_1\nu_2/p} = \beta_1 + \beta_2 + \beta_1\beta_2 + \frac{\frac{q_\ell}{P}B_{\mathsf{ks}}^{(2)} + (1+p)\,B_{\mathsf{scale}}}{\nu_1\nu_2}.$$

Let β_* bound the nonlinear and operation-dependent contribution at every multiplication; then $\beta' \le \beta_1 + \beta_2 + \beta_*$.

We now consider the homomorphic evaluation of the degree-d power polynomial $f(x) = x^d$, obtained by repeated squaring and immediate rescaling. Applying the relative-error bound of [7, Lemma 6] gives $\beta_d \le d\beta_0 + (d-1)\,\beta_*$. In particular, under the parameter condition $\beta_* \le \beta_0$, we obtain $\beta_d \le 2d\beta_0$.

6.4 Parameter Evaluation

Focusing on the homomorphic evaluation of a degree-d power polynomial analyzed above, we give a concrete end-to-end parameter evaluation for CKKS. It illustrates how non-power-of-two cyclotomics can reduce both the modulus and the lattice dimension: the broader range of available dimensions enables finer parameter selection, while the smaller rounding error reduces the required modulus, all while achieving the targeted output precision and security level.

We assume a fixed scale and rescaling base $\Delta = p > 1$, with $q_\ell = p^\ell q_0$, for $\ell = 1, \ldots, L$, and take the common choice $P = q_L$. The largest modulus entering the security estimate is therefore $Pq_L = q_L^2$. This remains true even without P-scaled encryption, since relinearization uses an evaluation key defined modulo Pq_L. We choose p to achieve a prescribed output precision and q_0 to prevent wraparound at decryption; the remaining moduli are accordingly determined.

Precision requirement. As discussed in the previous section, the relative error after evaluating the power polynomial $f(x) = x^d$ satisfies $\beta_d \le 2d\beta_0$, where $\beta_0 := B_{\mathsf{fresh}}/(p\sqrt{n})$. To obtain T bits of output precision, it suffices to require $\beta_d \le 2^{-T}$. Hence, it is sufficient to choose $p \ge 2^{T+1}dB_{\mathsf{fresh}}/\sqrt{n}$.

No-wraparound condition. We next choose q_0 so that decryption does not wrap around, as guaranteed by the following lemma, proved in Appendix A.7.

Lemma 10. *Suppose that the level-0 decrypted plaintext $\mu_{\mathsf{out}} \in R$ satisfies $\sigma(\mu_{\mathsf{out}}) = \boldsymbol{u}_{\mathsf{out}} + \boldsymbol{e}_{\mathsf{out}}$ where $\|\boldsymbol{u}_{\mathsf{out}}\|_2 \le \nu_{\mathsf{out}}$, and $\|\boldsymbol{e}_{\mathsf{out}}\|_2 \le B_{\mathsf{out}}$. Then centered lifting modulo $q_0 R$, with respect to the powerful basis, recovers μ_{out} without wraparound whenever $q_0 > 2\sqrt{\mathrm{rad}(m)/m}\,(\nu_{\mathsf{out}} + B_{\mathsf{out}})$.*

It remains to determine ν_{out} and B_{out}. By the relative error definition, we have $B_{\mathsf{out}} \le \beta_d \nu_{\mathsf{out}}$. Since $\Delta = p$ and the homomorphic evaluation rescales after each multiplication, the reference output of the power circuit on the input $p\pi^{-1}(\boldsymbol{z})$ is, in canonical coordinates, $\left(p\pi^{-1}(\boldsymbol{z})\right)^{\odot d}/p^{d-1}$. Under the assumption $\left\|p\pi^{-1}(\boldsymbol{z})\right\|_\infty \le p$, this output has canonical ℓ_∞-norm at most p, and hence canonical ℓ_2-norm at most $p\sqrt{n}$. We may therefore take $\nu_{\mathsf{out}} = p\sqrt{n}$. Consequently, $B_{\mathsf{out}} = \beta_d p\sqrt{n} \le 2d\beta_0 p\sqrt{n} = 2dB_{\mathsf{fresh}}$. Substituting this bound into Lemma 10, it is sufficient to choose $q_0 \ge \left\lceil 2\sqrt{\mathrm{rad}(m)/m}\,(p\sqrt{n} + 2dB_{\mathsf{fresh}}) \right\rceil$.

Numerical setting. We evaluate $f(x) = x^{16}$, which has multiplicative depth $L = 4$, and we require $T = 10$ bits of output precision. We compare the power-of-two cyclotomic dimensions $n = 2^{13}$ and $n = 2^{14}$ with non-power-of-two cyclotomic indices m satisfying $2^{13} \le n = \varphi(m) \le 2^{14}$. For each m, we choose p and q_0 at their lower bounds using $B_{\mathsf{fresh}} \approx B_{\mathsf{rnd}} + B_{\mathsf{scale}}$. We then set $q_L = p^L q_0,\ P = q_L$.

To estimate security, we supply the Albrecht Lattice Estimator [1] with dimension n, largest modulus $q_{\mathsf{sec}} = P q_L = q_L^2$, and Gaussian standard deviation $\sigma_{\mathsf{est}} = \frac{\widehat{m}}{n}\sigma_0$ for both the secret and error distributions, where $\sigma_0 = 3.19$ is the conventional power-basis standard deviation. For power-of-two cyclotomics, $\widehat{m} = n$, and hence $\sigma_{\mathsf{est}} = 3.19$. We report the BDD security estimate λ_{BDD} and the corresponding optimized BKZ block size β_{BKZ} with bit-cost $0.292\beta_{\mathsf{BKZ}}$.

For non-power-of-two cyclotomics, a spherical Gaussian in the canonical embedding generally induces a correlated, non-spherical Gaussian in the polynomial basis. Therefore, σ_{est} gives only an isotropic approximation. We additionally use the covariance-aware estimator of [9], which retains the full polynomial-basis covariance, and report the resulting block size $\beta_{\mathsf{BKZ,cov}}$ of bit-cost $0.265\beta_{\mathsf{BKZ,cov}}$.

Numerical results. Table 2 illustrates the tradeoff specific to non-power-of-two cyclotomics. The columns $\mathrm{Red}(m)$ and σ_{est} capture, respectively, the reduction in rounding error and the increase in the Gaussian parameter incurred when passing to the primal representation. In particular, a smaller $\mathrm{Red}(m)$ does not always yield a smaller modulus, since it may be offset by a larger value of $\widehat{m}/n$, and hence by a larger σ_{est}. Nevertheless, the choice $m = 2^{12} \cdot 7$ achieves $\lambda_{\mathsf{BDD}} = 133.84$ bits with both smaller dimension and smaller modulus than the next power-of-two choice: $n = 12288 < 16384$ and $\log_2 q_{\mathsf{sec}} = 316.91 < 317.36$. At fixed dimension, the table also isolates the effect of the cyclotomic factorization. For $n = 11520$ or $n = 12672$, different choices of m trade a smaller rounding factor $\mathrm{Red}(m)$ against a larger σ_{est}, and hence a larger modulus.

We repeat the same parameter evaluation for the deeper power circuit $f(x) = x^{1024}$, with multiplicative depth $L = 10$ and target output precision $T = 25$. We choose non-power-of-two cyclotomic indices m satisfying $2^{15} \le n = \varphi(m) \le 2^{16}$.

Table 2: CKKS parameter and security estimates for $f(x) = x^{16}$, with $T = 10$.

m	n	$\mathrm{Red}(m)$	σ_{est}	$\log_2 q_{\mathsf{sec}}$	λ_{BDD}	β_{BKZ}	$\beta_{\mathsf{BKZ,cov}}$	
2^{14}	8192	1.000	3.19	307.37	89.42	192	192.07	
$2^{10} \cdot 23$	11264	0.957	3.34	316.42	121.75	301	301.59	
$3^2 \cdot 5^2 \cdot 7 \cdot 17$	11520	0.757	7.41	330.01	120.03	295	296.81	
$7 \cdot 11 \cdot 13 \cdot 17$	11520	0.775	4.71	327.72	120.48	296	298.07	
$2^5 \cdot 7 \cdot 11 \cdot 13$	11520	0.817	4.43	323.60	122.06	302	303.37	
$2^{12} \cdot 7$	12288	0.926	3.72	316.91	133.84	342	342.93	128-bit thresh.
$2^2 \cdot 5 \cdot 11 \cdot 13^2$	12480	0.813	4.75	325.10	132.69	338	339.52	
$3 \cdot 5 \cdot 7 \cdot 13 \cdot 23$	12672	0.719	7.90	335.60	130.97	332	333.88	
$2^4 \cdot 7 \cdot 13 \cdot 23$	12672	0.834	4.22	325.68	134.49	344	345.61	
$3 \cdot 7^2 \cdot 11 \cdot 17$	13440	0.770	6.52	331.96	141.26	367	368.57	
2^{15}	16384	1.000	3.19	317.36	184.47	514	514.41	

Table 3 shows that the same tradeoff persists in this setting. In particular, $m = 2^{12} \cdot 23$ achieves $\lambda_{\mathsf{BDD}} = 130.72$ bits with both smaller dimension and smaller modulus than the next power-of-two choice.

Thus, for both depths $L = 4$ and $L = 10$, non-power-of-two cyclotomics achieve the target precision while exceeding the 128-bit security threshold, without requiring the full jump to the next power-of-two cyclotomic index m. This reduces both the lattice dimension $n = \varphi(m)$ and the required modulus, and hence the computational and storage costs of the scheme.

These savings require lattice-based rather than coefficient-wise rounding. For $\varphi(\mathrm{rad}(m)) = O(\log n)$, the blockwise Babai computation runs in $O(n \log n)$ time and does not asymptotically dominate the $O(m \log m)$ domain conversions [25].

Table 3: CKKS parameter and security estimates for $f(x) = x^{1024}$, with $T = 25$.

m	n	$\mathrm{Red}(m)$	σ_{est}	$\log_2 q_{\mathsf{sec}}$	λ_{BDD}	β_{BKZ}	$\beta_{\mathsf{BKZ,cov}}$	
2^{16}	32768	1.000	3.19	1179.78	95.00	202	202.31	
$2^{13} \cdot 11$	40960	0.937	3.51	1191.14	118.31	283	282.78	
$3^2 \cdot 5 \cdot 7 \cdot 17 \cdot 19$	41472	0.721	7.83	1219.09	117.18	279	279.25	
$3 \cdot 5^2 \cdot 7 \cdot 11 \cdot 19$	43200	0.712	8.10	1220.41	122.24	296	296.62	
$2^2 \cdot 7 \cdot 11 \cdot 13 \cdot 31$	43200	0.787	4.58	1206.85	123.65	301	300.93	
$2^{12} \cdot 23$	45056	0.957	3.34	1194.36	130.72	325	324.86	128-bit thresh.
$3 \cdot 5^2 \cdot 7 \cdot 13 \cdot 17$	46080	0.713	8.03	1222.33	131.03	326	325.88	
$7 \cdot 13 \cdot 17 \cdot 41$	46080	0.802	4.39	1209.23	132.23	330	330.30	
$2^3 \cdot 7 \cdot 11 \cdot 13 \cdot 17$	46080	0.775	4.71	1208.38	132.50	331	330.77	
$2^4 \cdot 3 \cdot 5 \cdot 7 \cdot 11 \cdot 13$	46080	0.704	8.32	1222.35	131.04	326	325.93	
$2^2 \cdot 29^3$	47096	0.962	3.30	1196.00	136.94	346	346.11	
2^{17}	65536	1.000	3.19	1201.78	197.26	551	550.76	

References

1. Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of learning with errors. Cryptology ePrint Archive, Report 2015/046 (2015), https://eprint.iacr.org/2015/046
2. Babai, L.: On lovász' lattice reduction and the nearest lattice point problem. Combinatorica **6**(1), 1–13 (1986)
3. Bollauf, M.F., Vaishampayan, V.A., Costa, S.I.: On the communication cost of determining an approximate nearest lattice point. In: 2017 IEEE International Symposium on Information Theory (ISIT). pp. 1838–1842. IEEE (2017)
4. Bossuat, J.P., Costache, A., Mouchet, C., Nürnberger, L., Troncoso-Pastoriza, J.R.: Accurate and composable noise estimates for CKKS with application to exact HE computation. CiC **2**(2), 8 (2025). https://doi.org/10.62056/a3n5tx4e-
5. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) fully homomorphic encryption without bootstrapping. In: Goldwasser, S. (ed.) ITCS 2012. pp. 309–325. ACM (Jan 2012). https://doi.org/10.1145/2090236.2090262
6. Cheon, J.H., Han, K., Kim, A., Kim, M., Song, Y.: A full RNS variant of approximate homomorphic encryption. In: Cid, C., Jacobson, Jr., M.J. (eds.) SAC 2018. LNCS, vol. 11349, pp. 347–368. Springer, Cham (Aug 2019). https://doi.org/10.1007/978-3-030-10970-7_16
7. Cheon, J.H., Kim, A., Kim, M., Song, Y.S.: Homomorphic encryption for arithmetic of approximate numbers. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part I. LNCS, vol. 10624, pp. 409–437. Springer, Cham (Dec 2017). https://doi.org/10.1007/978-3-319-70694-8_15
8. Conway, J.H., Sloane, N.J.A.: Sphere Packings, Lattices and Groups, Grundlehren der mathematischen Wissenschaften, vol. 290. Springer (1988). https://doi.org/10.1007/978-1-4757-2016-7, https://doi.org/10.1007/978-1-4757-2016-7
9. Dachman-Soled, D., Ducas, L., Gong, H., Rossi, M.: LWE with side information: Attacks and concrete security estimation. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part II. LNCS, vol. 12171, pp. 329–358. Springer, Cham (Aug 2020). https://doi.org/10.1007/978-3-030-56880-1_12
10. Di Giusto, A., Marcolla, C.: Breaking the power-of-two barrier: noise estimation for bgv in ntt-friendly rings. Designs, Codes and Cryptography **93**(3), 467–502 (2025)
11. Dinur, I., Kindler, G., Safra, S.: Approximating-CVP to within almost-polynomial factors is NP-hard. In: 39th FOCS. pp. 99–111. IEEE Computer Society Press (Nov 1998). https://doi.org/10.1109/SFCS.1998.743433
12. Ducas, L., Durmus, A.: Ring-LWE in polynomial rings. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 34–51. Springer, Berlin, Heidelberg (May 2012). https://doi.org/10.1007/978-3-642-30057-8_3
13. Ducas, L., van Woerden, W.P.J.: The closest vector problem in tensored root lattices of type A and in their duals. DCC **86**(1), 137–150 (2018). https://doi.org/10.1007/s10623-017-0332-x
14. van Emde Boas, P.: Another NP-complete partition problem and the complexity of computing short vectors in a lattice. Tech. Rep. 81-04, Department of Mathematics, University of Amsterdam (1981)
15. Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144 (2012), https://eprint.iacr.org/2012/144
16. Giusto, A.D., Marcolla, C.: Breaking the power-of-two barrier: noise estimation for BGV in NTT-friendly rings. DCC **93**(3), 467–502 (2025). https://doi.org/10.1007/s10623-024-01524-5

17. Khot, S., Popat, P., Vishnoi, N.K.: $2^{\log^{1-\epsilon} n}$ hardness for the closest vector problem with preprocessing. In: Karloff, H.J., Pitassi, T. (eds.) 44th ACM STOC. pp. 277–288. ACM Press (May 2012). https://doi.org/10.1145/2213977.2214004
18. Kim, A., Papadimitriou, A., Polyakov, Y.: Approximate homomorphic encryption with reduced approximation error. In: Galbraith, S.D. (ed.) CT-RSA 2022. LNCS, vol. 13161, pp. 120–144. Springer, Cham (Mar 2022). https://doi.org/10.1007/978-3-030-95312-6_6
19. Lee, J.W., Lee, E., Lee, Y., Kim, Y.S., No, J.S.: High-precision bootstrapping of RNS-CKKS homomorphic encryption using optimal minimax polynomial approximation and inverse sine function. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021, Part I. LNCS, vol. 12696, pp. 618–647. Springer, Cham (Oct 2021). https://doi.org/10.1007/978-3-030-77870-5_22
20. Lee, Y., Lee, J.W., Kim, Y.S., Kim, Y., No, J.S., Kang, H.: High-precision bootstrapping for approximate homomorphic encryption by error variance minimization. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022, Part I. LNCS, vol. 13275, pp. 551–580. Springer, Cham (May / Jun 2022). https://doi.org/10.1007/978-3-031-06944-4_19
21. Lenstra, A.K., Lenstra, H.W.J., Lovász, L.: Factoring polynomials with rational coefficients. Mathematische Annalen **261**(4), 515–534 (1982). https://doi.org/10.1007/BF01457454
22. Liu, J., Feng, Y., Pan, Y.: Exact CVP is NP-complete for principal cyclotomic ideals. Cryptology ePrint Archive, Paper 2026/1793 (2026), https://eprint.iacr.org/2026/1793
23. 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, Berlin, Heidelberg (May / Jun 2010). https://doi.org/10.1007/978-3-642-13190-5_1
24. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. J. ACM **60**(6) (Nov 2013). https://doi.org/10.1145/2535925, https://doi.org/10.1145/2535925
25. Lyubashevsky, V., Peikert, C., Regev, O.: A toolkit for ring-LWE cryptography. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 35–54. Springer, Berlin, Heidelberg (May 2013). https://doi.org/10.1007/978-3-642-38348-9_3
26. Lyubashevsky, V., Seiler, G.: NTTRU: Truly fast NTRU using NTT. IACR TCHES **2019**(3), 180–201 (2019). https://doi.org/10.13154/tches.v2019.i3.180-201, https://tches.iacr.org/index.php/TCHES/article/view/8293
27. Mayer, C.M.: Implementing a toolkit for ring-LWE based cryptography in arbitrary cyclotomic number fields. Cryptology ePrint Archive, Report 2016/049 (2016), https://eprint.iacr.org/2016/049
28. McKilliam, R.G., Clarkson, I.V.L., Smith, W.D., Quinn, B.G.: A linear-time nearest point algorithm for the lattice a n. In: 2008 International Symposium on Information Theory and Its Applications. pp. 1–5. IEEE (2008)
29. Mihara, K., Yamaguchi, R., Mitsuishi, M., Maruyama, Y.: Neural network training with homomorphic encryption (2020). https://doi.org/10.48550/arXiv.2012.13552
30. Pirillo, A., Colombo, L.: ReBoot: Encrypted training of deep neural networks with CKKS bootstrapping (2025). https://doi.org/10.48550/arXiv.2506.19693
31. van Woerden, W.P.J.: The Closest Vector Problem in Cyclotomic Lattices. Master's thesis, Leiden University (2015), https://math.leidenuniv.nl/scripties/BachVanWoerden.pdf

32. Zhou, Y., Liang, M., Brugere, I., Dervovic, D., Guo, Y., Polychroniadou, A., Wu, M., Dachman-Soled, D.: Revisiting ML training under fully homomorphic encryption: Convergence guarantees, differential privacy, and efficient algorithms. CoRR **abs/2605.27782** (2026). https://doi.org/10.48550/ARXIV.2605.27782, https://doi.org/10.48550/arXiv.2605.27782

A Proofs

A.1 Proof of Proposition 1

Since both the powerful basis and the canonical embedding factor over the prime-power components of m, we first consider the prime-power case $m = p^k$, and let $m' = m/p = p^{k-1}$. For this case, Lemma 1 characterizes the spectrum of CRT_m in terms of Q_m, D_p, and U_p, which are unitary, diagonal, and upper unitriangular, respectively.

The unitary factor Q_m does not change Gram–Schmidt lengths. Moreover, multiplying a matrix on the right by an upper unitriangular matrix only adds linear combinations of previous columns to each column, and therefore also preserves the Gram–Schmidt lengths. Hence the squared Gram–Schmidt lengths are exactly the diagonal entries of $m' \left(D_p^2 \otimes I_{[m']}\right)$. Thus, each value $m'\left((p-1) - j/(p-j)\right)$ appears exactly $m' = p^{k-1}$ times, and therefore

$$\sum_{j \in [\varphi(p^k)]} \left\|\widetilde{\boldsymbol{b}}_j\right\|_2^2 = p^{2k-2}\left((p-1)^2 - \sum_{j=0}^{p-2} \frac{j}{p-j}\right) = p^{2k-1}\left(p - H_p\right),$$

for $H_p := \sum_{k=1}^{p} \frac{1}{k}$, where the last equality holds by setting $r = p - j$ and simplifying the remaining terms.

For a general m having prime-power factorization $\prod_\ell p_\ell^{k_\ell}$, we use the fact that the sums of squared Gram–Schmidt lengths multiply under tensor products, i.e.,

$$\sum_{j \in [\varphi(m)]} \left\|\widetilde{\boldsymbol{b}}_j\right\|_2^2 = \prod_\ell \sum_{j \in [\varphi(m_\ell)]} \left\|\widetilde{\boldsymbol{b}}_j\right\|_2^2 = \prod_\ell p_\ell^{2k_\ell - 1}\left(p_\ell - H_{p_\ell}\right).$$

□

A.2 Proof of Lemma 2

We have

$$-\log \mathrm{Red}(m) = \sum_{\text{prime } p \mid m} \left[-\frac{1}{2}\log\left(1 - \frac{H_p}{p}\right) + \log\left(1 - \frac{1}{p}\right)\right].$$

Using $\log(1-x) = -x + O(x^2)$, for all sufficiently large primes p, the summand satisfies

$$-\frac{1}{2}\log\left(1 - \frac{H_p}{p}\right) + \log\left(1 - \frac{1}{p}\right) = \frac{H_p}{2p} - \frac{1}{p} + O\left(\frac{H_p^2}{p^2}\right).$$

Since $H_p \le 1 + \log p$, this implies

$$-\frac{1}{2}\log\left(1-\frac{H_p}{p}\right)+\log\left(1-\frac{1}{p}\right) \le \frac{\log p}{2p}+O\left(\frac{(\log p)^2}{p^2}\right).$$

After absorbing finitely many small primes into the constant, we can write

$$-\log \mathrm{Red}(m) \le \frac{1}{2}\sum_{\text{prime } p\mid m}\frac{\log p}{p}+O\left(\sum_{\text{prime } p}\frac{(\log p)^2}{p^2}\right)=\frac{1}{2}\sum_{\text{prime } p\mid m}\frac{\log p}{p}+O(1),$$

where the last equality holds since $\sum_{\text{prime } p}(\log p)^2/p^2 < \infty$.

It remains to bound the prime-divisor sum. Set $y = \log m$. Then

$$\sum_{\text{prime } p\mid m}\frac{\log p}{p}=\sum_{\substack{\text{prime } p\mid m\\ p\le y}}\frac{\log p}{p}+\sum_{\substack{\text{prime } p\mid m\\ p>y}}\frac{\log p}{p}.$$

For the first term, by the Mertens estimate for prime sums,

$$\sum_{\substack{\text{prime } p\mid m\\ p\le y}}\frac{\log p}{p}\le\sum_{\text{prime } p\le y}\frac{\log p}{p}=\log y+O(1)=\log\log m+O(1).$$

For the second term, let $r = \#\{\text{prime } p \mid m :\ p > y\}$. Since the product of these r primes divides m, we have $y^r \le m$, and hence $r \le \log m/\log y$. Moreover, the function $x \mapsto (\log x)/x$ is decreasing for $x > e$, so for $p > y$, $\frac{\log p}{p} \le \frac{\log y}{y}$. Therefore

$$\sum_{\substack{\text{prime } p\mid m\\ p>y}}\frac{\log p}{p}\le r\cdot\frac{\log y}{y}\le\frac{\log m}{\log y}\cdot\frac{\log y}{y}=\frac{\log m}{y}=1.$$

Together, $\sum_{\text{prime } p\mid m}\log p/p \le \log\log m+O(1)$, thus $-\log\mathrm{Red}(m) \le \frac{1}{2}\log\log m+O(1)$, and exponentiating yields $1/\mathrm{Red}(m) = O\left(\sqrt{\log m}\right)$. □

A.3 Proof of Lemma 3

CVP is solved directly over the lattice generated by B_{blk} by narrowing down the search space using the minimum singular value and the covering radius, satisfying the chain

$$s_{\min}(B_{\mathrm{blk}})\,\|\boldsymbol{z}_i-\boldsymbol{r}_i\|_2 \le \|B_{\mathrm{blk}}(\boldsymbol{z}_i-\boldsymbol{r}_i)\|_2 \le \rho(B_{\mathrm{blk}}).$$

It is thus enough to test all integer vectors $\boldsymbol{z}_i \in \mathbb{Z}^b$ such that

$$\|\boldsymbol{z}_i\|_2 \le \|\boldsymbol{z}_i-\boldsymbol{r}_i\|_2+\|\boldsymbol{r}_i\|_2 \le \frac{\rho(B_{\mathrm{blk}})}{s_{\min}(B_{\mathrm{blk}})}+\frac{1}{2}\sqrt{b} \le \frac{1}{2}\frac{S(m)}{n}b+\frac{1}{2}\sqrt{b}. \tag{8}$$

The last inequality follows as the covering radius is upper bounded by $2\rho(B_{\mathrm{blk}}) \le S(m)/\sqrt{n/b}$ with $S(m)$ given in Proposition 1; while, by Claim 2.2, the minimum singular value of the block is $s_{\min}(B_{\mathrm{blk}}) = \sqrt{m/\mathrm{rad}(m)} = \sqrt{n/b}$. Using the standard volume bound for the number of integer points in a Euclidean ball, the number of candidates per block is at most

$$C_{\mathrm{blk}}(m) := V_b \left(\frac{1}{2} \frac{S(m)}{n} b + \sqrt{b} \right)^b, \quad V_b := \frac{\pi^{b/2}}{\Gamma\left(\frac{b}{2}+1\right)};$$

where V_b is the volume of the b-dimensional Euclidean unit ball and $\Gamma(\cdot)$ is the Gamma function.

To estimate $C_{\mathrm{blk}}(m)$, first note that $\frac{S(m)}{n} = \mathrm{Red}(m) \le 1$ and so $\left(\frac{1}{2}\frac{S(m)}{n} b + \sqrt{b}\right)^b = O(b)^b$. Further, V_b can be estimated using Stirling's formula in the form $\Gamma(x+1) = \sqrt{2\pi x}\left(\frac{x}{e}\right)^x \left(1 + O\left(\frac{1}{x}\right)\right)$, which gives $V_b = \exp(O(b))\, b^{-b/2}$. Therefore,

$$C_{\mathrm{blk}}(m) \le \exp\left(\frac{1}{2} b \log b + O(b) \right).$$

Since there are n/b independent blocks, the total number of candidates tested is at most $C_{\mathrm{blk}}(m) \cdot n/b \le O\left(n\,(cb)^{b/2}\right)$ for some absolute constant $c > 0$. □

A.4 Proof of Theorem 2

We use the following observations.

Observation 1. *Let $A^{n\times n}$ be an invertible matrix and define $\mathcal{S} \subset \mathbb{R}^n$. For $\boldsymbol{v} \in \mathcal{S}$, consider $\boldsymbol{w} := A\boldsymbol{v}$. Then*

$$\boldsymbol{w} \sim \mathrm{Unif}(A\mathcal{S}) \iff \boldsymbol{v} \sim \mathrm{Unif}(\mathcal{S}),$$

where $\mathrm{Unif}(\cdot)$ *stands for a uniform probability distribution.*

Proof. This follows from the standard change-of-variables formula, since the linear map $\boldsymbol{v} \mapsto A\boldsymbol{v}$ is invertible and scales volume by the constant factor $|\det A|$.

Observation 2. *Denote by x_i the i-th entry of the vector $\boldsymbol{x} \in \mathbb{R}^n$, and, for $i \in [n]$, consider $\mathcal{S}_i \subset \mathbb{R}$. Then*

$$\boldsymbol{x} \sim \mathrm{Unif}\left(\prod_{i\in[n]} \mathcal{S}_i \right) \iff x_0, \dots, x_{n-1} \text{ are independent and } x_i \sim \mathrm{Unif}(\mathcal{S}_i)\ \forall i.$$

Define $\mathcal{S} := [-0.5, 0.5)^n$, so that $\mathcal{P}_{1/2}(B) = B\mathcal{S}$, and set $\boldsymbol{x} := B^{-1}\boldsymbol{t} \in \mathcal{S}$. Since $\boldsymbol{t} \sim \mathrm{Unif}\big(\mathcal{P}_{1/2}(B)\big)$, Observations 1 and 2 imply that $x_i \overset{\text{i.i.d.}}{\sim} \mathrm{Unif}([-0.5, 0.5))$ for every $i \in [n]$.

At each iteration of the algorithm, we represent the error for target $\boldsymbol{t}$ at the beginning and end of iteration i as $\boldsymbol{e}_{\boldsymbol{t},\mathsf{b}}^{(i)}$ and $\boldsymbol{e}_{\boldsymbol{t},\mathsf{e}}^{(i)}$, respectively. The algorithm starts with $\boldsymbol{e}_{\boldsymbol{t},\mathsf{b}}^{(n-1)} = \boldsymbol{t}$ and, at each iteration, follows the update rule

$$\boldsymbol{e}_{\boldsymbol{t},\mathsf{b}}^{(i)} = \sum_{j=0}^{i} x_j \boldsymbol{b}_j + \sum_{j=i+1}^{n-1} \left(x_j - \left\lceil \frac{\left\langle \boldsymbol{e}_{\boldsymbol{t},\mathsf{b}}^{(j)}, \widetilde{\boldsymbol{b}}_j \right\rangle}{\left\| \widetilde{\boldsymbol{b}}_j \right\|_2^2} \right\rfloor \right) \boldsymbol{b}_j; \quad \boldsymbol{e}_{\boldsymbol{t},\mathsf{e}}^{(i)} = \boldsymbol{e}_{\boldsymbol{t},\mathsf{b}}^{(i)} - \left\lceil \frac{\left\langle \boldsymbol{e}_{\boldsymbol{t},\mathsf{b}}^{(i)}, \widetilde{\boldsymbol{b}}_i \right\rangle}{\left\| \widetilde{\boldsymbol{b}}_i \right\|_2^2} \right\rfloor \boldsymbol{b}_i.$$

Recall that the Gram–Schmidt process gives

$$\forall j \in [n],\ \widetilde{\boldsymbol{b}}_j = \boldsymbol{b}_j - \sum_{k \in [j]} \frac{\left\langle \boldsymbol{b}_j, \widetilde{\boldsymbol{b}}_k \right\rangle}{\left\| \widetilde{\boldsymbol{b}}_k \right\|_2^2} \widetilde{\boldsymbol{b}}_k \quad \text{and thus} \quad \boldsymbol{b}_j = \widetilde{\boldsymbol{b}}_j + \sum_{k \in [j]} \mu_{j,k} \widetilde{\boldsymbol{b}}_k,$$

implying that $\widetilde{\boldsymbol{b}}_i$ is orthogonal to $\{\boldsymbol{b}_j\}_{j \in [i]}$. It follows that the coefficient of $\boldsymbol{e}_{\boldsymbol{t},\mathsf{e}}^{(i)}$ in the direction of $\widetilde{\boldsymbol{b}}_i$, denoted $\tilde{d}_{\boldsymbol{t},\mathsf{e}}^{(i),i}$, will not change after the i-th iteration, and is given by

$$\begin{aligned}
\tilde{d}_{\boldsymbol{t},\mathsf{e}}^{(i),i} &:= \frac{\left\langle \boldsymbol{e}_{\boldsymbol{t},\mathsf{e}}^{(i)}, \widetilde{\boldsymbol{b}}_i \right\rangle}{\left\| \widetilde{\boldsymbol{b}}_i \right\|_2^2} = \frac{\left\langle \boldsymbol{e}_{\boldsymbol{t},\mathsf{b}}^{(i)}, \widetilde{\boldsymbol{b}}_i \right\rangle}{\left\| \widetilde{\boldsymbol{b}}_i \right\|_2^2} - \left\lceil \frac{\left\langle \boldsymbol{e}_{\boldsymbol{t},\mathsf{b}}^{(i)}, \widetilde{\boldsymbol{b}}_i \right\rangle}{\left\| \widetilde{\boldsymbol{b}}_i \right\|_2^2} \right\rfloor := s_{\boldsymbol{t},i} - \lceil s_{\boldsymbol{t},i} \rfloor, \text{ where} \\
s_{\boldsymbol{t},i} &:= \frac{\left\langle \boldsymbol{e}_{\boldsymbol{t},\mathsf{b}}^{(i)}, \widetilde{\boldsymbol{b}}_i \right\rangle}{\left\| \widetilde{\boldsymbol{b}}_i \right\|_2^2} = x_i + \sum_{j=i+1}^{n-1} \left(x_j - \left\lceil \frac{\left\langle \boldsymbol{e}_{\boldsymbol{t},\mathsf{b}}^{(j)}, \widetilde{\boldsymbol{b}}_j \right\rangle}{\left\| \widetilde{\boldsymbol{b}}_j \right\|_2^2} \right\rfloor \right) \frac{\left\langle \boldsymbol{b}_j, \widetilde{\boldsymbol{b}}_i \right\rangle}{\left\| \widetilde{\boldsymbol{b}}_i \right\|_2^2} \\
&= x_i + \sum_{j=i+1}^{n-1} \left(x_j - \lceil s_{\boldsymbol{t},j} \rfloor \right) \mu_{j,i}.
\end{aligned}$$

We prove by backward induction that for every $i \in [n]$, $s_{\boldsymbol{t},i}$ depends only on $\{x_j\}_{j \geq i}$. For the base case, $s_{\boldsymbol{t},n-1} = x_{n-1}$. For the inductive case, we assume that for every $j > i$, $s_{\boldsymbol{t},j}$ depends only on $\{x_{j'}\}_{j' \geq j}$ and show that $s_{\boldsymbol{t},i}$ depends only on $\{x_j\}_{j \geq i}$. Indeed, since $s_{\boldsymbol{t},i} = x_i + \sum_{j=i+1}^{n-1} (x_j - \lceil s_{\boldsymbol{t},j} \rfloor)\, \mu_{j,i}$, each of the summands depends on $s_{\boldsymbol{t},j}$ for $j > i$. By the inductive hypothesis, the term $\sum_{j=i+1}^{n-1} (x_j - \lceil s_{\boldsymbol{t},j} \rfloor)\, \mu_{j,i}$ depends only on $\{x_{j'}\}_{j' > i}$, and thus $s_{\boldsymbol{t},i}$ depends only on $\{x_j\}_{j \geq i}$. We can therefore write,

$$s_{\boldsymbol{t},i} = x_i + \sum_{j=i+1}^{n-1} \left(x_j - \lceil s_{\boldsymbol{t},j} \rfloor \right) \mu_{j,i} := x_i + f_i \left(\{x_j\}_{j>i} \right)$$

where $f_i \left(\{x_j\}_{j>i} \right)$ denotes a function that depends only on $\{x_j\}_{j>i}$.

Altogether, the coefficient of Babai's error in the direction of $\widetilde{\boldsymbol{b}}_i$ satisfies

$$\tilde{d}_{\boldsymbol{t},\mathsf{e}}^{(i),i} = s_{\boldsymbol{t},i} - \lceil s_{\boldsymbol{t},i} \rfloor = x_i + f_i \left(\{x_j\}_{j>i} \right) - \left\lceil x_i + f_i \left(\{x_j\}_{j>i} \right) \right\rfloor.$$

As for the distribution of $\tilde{d}^{(i),i}_{\boldsymbol{t},\mathsf{e}}$, let us consider the probability of the event $\left\{\tilde{d}^{(i),i}_{\boldsymbol{t},\mathsf{e}} \in I\right\}$ for $I \subseteq \mathbb{R}$:

$$\Pr\left(\tilde{d}^{(i),i}_{\boldsymbol{t},\mathsf{e}} \in I\right) = \mathbb{E}\Big[\Pr\Big(x_i + f_i\left(\{x_j\}_{j>i}\right) - \Big\lceil x_i + f_i\left(\{x_j\}_{j>i}\right)\Big\rfloor \in I \Big| \{x_j\}_{j>i}\Big)\Big].$$

Note that, given $\{x_j\}_{j>i}$, the value of $f_i\left(\{x_j\}_{j>i}\right)$ equals some fixed constant C, and further, $x_i \overset{\text{i.i.d.}}{\sim} \mathrm{Unif}([-0.5, 0.5))$ for every $i \in [n]$. Adding an arbitrary constant to a uniformly distributed point in $[-0.5, 0.5)$ and reducing the result modulo 1 back to $[-0.5, 0.5)$ preserves the uniform distribution. This is stated in the following claim.

Claim. Let $x \sim \mathrm{Unif}([-0.5, 0.5))$, let $C \in \mathbb{R}$ be an arbitrary constant, and define $y := x + C - \lceil x + C \rfloor$. Then $y \sim \mathrm{Unif}([-0.5, 0.5))$.

By the claim, $\Pr\left(\tilde{d}^{(i),i}_{\boldsymbol{t},\mathsf{e}} \in I \Big| \{x_j\}_{j>i}\right) = |I \cap [-0.5, 0.5)|$, and thus $\tilde{d}^{(i),i}_{\boldsymbol{t},\mathsf{e}} \sim \mathrm{Unif}([-0.5, 0.5))$.

It remains to show the independence of $\left\{\tilde{d}^{(i),i}_{\boldsymbol{t},\mathsf{e}}\right\}_{i\in[n]}$. Indeed,

$$\begin{aligned}\Pr\left(\tilde{d}^{(i),i}_{\boldsymbol{t},\mathsf{e}} \in I \Big| \left\{\tilde{d}^{(j),j}_{\boldsymbol{t},\mathsf{e}}\right\}_{j>i}\right) &= \mathbb{E}\left[\Pr\left(\tilde{d}^{(i),i}_{\boldsymbol{t},\mathsf{e}} \in I \Big| \{x_j\}_{j>i}\right) \Big| \left\{\tilde{d}^{(j),j}_{\boldsymbol{t},\mathsf{e}}\right\}_{j>i}\right] \\ &= |I \cap [-0.5, 0.5)| = \Pr\left(\tilde{d}^{(i),i}_{\boldsymbol{t},\mathsf{e}} \in I\right),\end{aligned}$$

where the first equality follows directly from the tower property of conditional expectation, as $\left\{\tilde{d}^{(j),j}_{\boldsymbol{t},\mathsf{e}}\right\}_{j>i}$ is a deterministic function of $\{x_j\}_{j>i}$.

We conclude that $\tilde{d}^{(i),i}_{\boldsymbol{t},\mathsf{e}} \overset{\text{i.i.d.}}{\sim} \mathrm{Unif}([-0.5, 0.5))$ for every $i \in [n]$, and thus $\boldsymbol{e}^{(0)}_{\boldsymbol{t},\mathsf{e}} = \sum_{i\in[n]} \tilde{d}^{(i),i}_{\boldsymbol{t},\mathsf{e}} \widetilde{\boldsymbol{b}}_i$ is uniformly distributed over $\mathcal{P}_{1/2}(\widetilde{B})$ by Observations 1 and 2. □

A.5 Proof of Proposition 2

We use Bernstein's inequality and basic properties of a uniform random variable, stated next.

Theorem 3 (Bernstein's inequality). *Let $Y_0, \ldots, Y_{n-1}$ be independent centered random variables satisfying $|Y_j| \le M$ almost surely. Then, for every $t > 0$,*

$$\Pr\left(\sum_{j\in[n]} Y_j \ge t\right) \le \exp\left(-\frac{t^2}{2\sum_{j\in[n]} \mathbb{E}\left[Y_j^2\right] + \frac{2}{3}Mt}\right).$$

Observation 3. *For $c \sim \mathrm{Unif}([-0.5, 0.5))$,*

$$\mathbb{E}\left[c^2\right] = \frac{1}{12}, \qquad \mathrm{Var}\left(c^2\right) = \frac{1}{180}, \qquad \left|c^2 - \frac{1}{12}\right| \le \frac{1}{6}.$$

Let $\boldsymbol{e}$ denote the error of Babai's nearest-plane algorithm, and let $n = \varphi(m)$. By the error characterization in Theorem 2,

$$\boldsymbol{e} = \sum_{j\in[n]} c_j \widetilde{\boldsymbol{b}}_j, \qquad c_j \overset{\text{i.i.d.}}{\sim} \mathrm{Unif}([-0.5, 0.5)),$$

where $\{\widetilde{\boldsymbol{b}}_j\}_{j\in[n]}$ are the Gram–Schmidt vectors of $B = \mathrm{CRT}_m$. By orthogonality and Corollary 2,

$$\|\boldsymbol{e}\|_2^2 - \mathbb{E}\Big[\|\boldsymbol{e}\|_2^2\Big] = \sum_{j\in[n]} \big\|\widetilde{\boldsymbol{b}}_j\big\|_2^2 \left(c_j^2 - \frac{1}{12}\right) := \sum_{j\in[n]} Y_j.$$

The random variables Y_j are independent and centered. Moreover,

$$|Y_j| \le \frac{1}{6}\big\|\widetilde{\boldsymbol{b}}_j\big\|_2^2 \le \frac{n}{6} =: M,$$

where $\big\|\widetilde{\boldsymbol{b}}_j\big\|_2^2 \le \|\boldsymbol{b}_j\|_2^2 = n$ by Claim 2.2. Furthermore,

$$\sum_{j\in[n]} \mathbb{E}\big[Y_j^2\big] = \frac{1}{180}\sum_{j\in[n]} \big\|\widetilde{\boldsymbol{b}}_j\big\|_2^4 \le \frac{n}{180}\sum_{j\in[n]} \big\|\widetilde{\boldsymbol{b}}_j\big\|_2^2 = \frac{n^3\mathrm{Red}^2(m)}{180},$$

where the last equality follows from Corollary 3, which gives $\sum_{j\in[n]} \big\|\widetilde{\boldsymbol{b}}_j\big\|_2^2 = n^2\mathrm{Red}^2(m)$. We conclude the proof by applying Bernstein's inequality with $t := \varepsilon\mathbb{E}\Big[\|\boldsymbol{e}\|_2^2\Big] = \varepsilon n^2\mathrm{Red}^2(m)/12$:

$$\begin{aligned}\Pr\Big(\|\boldsymbol{e}\|_2^2 \ge (1+\varepsilon)\,\mathbb{E}\Big[\|\boldsymbol{e}\|_2^2\Big]\Big) &\le \exp\left(-\frac{\varepsilon^2 n^4\mathrm{Red}^2(m)^2/144}{n^3\mathrm{Red}^2(m)/90 + \varepsilon n^3\mathrm{Red}^2(m)/108}\right)\\ &= \exp\left(-n\frac{\varepsilon^2\mathrm{Red}^2(m)}{\frac{8}{5}+\frac{4}{3}\varepsilon}\right).\end{aligned}$$

□

A.6 Proof of Lemma 5

Let $\boldsymbol{h}_1, \ldots, \boldsymbol{h}_n$ be an orthonormal basis of H. Since $\sigma(a)$ is a spherical Gaussian with standard deviation $\widehat{m}\sigma$, its coordinates in this basis are independent Gaussians with standard deviation $\widehat{m}\sigma$. Hence,

$$\sigma(a) = \widehat{m}\sigma\sum_{i=1}^{n} g_i \boldsymbol{h}_i, \qquad g_1, \ldots, g_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1).$$

By orthonormality, $\|\sigma(a)\|_2^2/(\widehat{m}\sigma)^2 = \sum_{i=1}^n g_i^2 \sim \chi_n^2$. The standard chi-square concentration bound gives

$$\Pr\Big(\chi_n^2 \ge n + 2\sqrt{nt} + 2t\Big) \le e^{-t}.$$

Taking $t = n/2$, we have $n + 2\sqrt{nt} + 2t = \left(2 + \sqrt{2}\right) n < 4n$. Therefore,

$$\Pr\left(\|a\|_2^{\mathrm{can}} > 2\widehat{m}\sigma\sqrt{n}\right) \leq e^{-n/2}.$$

□

A.7 Proof of Lemma 10

Since CRT_m is the canonical-embedding matrix with respect to the powerful basis $\boldsymbol{p}_m$, the vector $\mathrm{CRT}_m^{-1}\,\sigma\left(\mu_{\mathsf{out}}\right)$ is precisely the coefficient vector of μ_{out} in $\boldsymbol{p}_m$. Centered lifting modulo $q_0 R$ therefore succeeds whenever

$$\left\|\mathrm{CRT}_m^{-1}\,\sigma\left(\mu_{\mathsf{out}}\right)\right\|_\infty < \frac{q_0}{2}.$$

Let $\left\|A^{-1}\right\|_2$ denote the spectral norm of the matrix A^{-1}, given by $1/s_{\min}(A)$ where $s_{\min}(A)$ is the smallest singular value of A. Now,

$$\begin{aligned}\left\|\mathrm{CRT}_m^{-1}\,\sigma\left(\mu_{\mathsf{out}}\right)\right\|_\infty &\leq \left\|\mathrm{CRT}_m^{-1}\right\|_2 \cdot \left\|\sigma\left(\mu_{\mathsf{out}}\right)\right\|_2 \\ &\leq \frac{1}{s_{\min}\left(\mathrm{CRT}_m\right)}\left(\nu_{\mathsf{out}} + B_{\mathsf{out}}\right).\end{aligned}$$

Finally, by Claim 2.2, $s_{\min}\left(\mathrm{CRT}_m\right) = \sqrt{m/\mathrm{rad}(m)}$, which concludes the proof.
□