# CVP Is NP-Complete for Principal Cyclotomic Ideals

Jiaqi Liu  Yansong Feng  Yanbin Pan

State Key Laboratory of Mathematical Sciences,
Academy of Mathematics and Systems Science, Beijing, China

{ljqi,fengyansong,panyanbin}@amss.ac.cn

**Abstract**

We prove that exact Euclidean decision-CVP is NP-complete on the coefficient lattices of nonzero principal ideals in the power-of-two cyclotomic rings $R_d := \mathbb{Z}[y]/(y^d + 1)$. Our deterministic reduction from Exact Cover by 3-Sets (X3C) produces a target and a squared threshold $\Delta$ such that the closest squared distance is exactly $\Delta$ in YES instances and at least $\Delta+4$ in NO instances. This also implies NP-hardness of exact search-CVP under polynomial-time Turing reductions.

We also transfer the resulting principal-ideal CVP instances to full-rank principal ideals of the cyclic quotient ring $\mathbb{Z}[X]/(X^D - 1)$, where $D := 2d$. Their coefficient lattices are invariant under cyclic coordinate shifts. The lift preserves principality and multiplies corresponding squared distances by eight. Thus, on principal cyclic ideal lattices, exact decision-CVP is NP-complete and exact search-CVP is NP-hard.

We also obtain uniformly computable fixed cyclotomic and cyclic families in which only the target and threshold depend on the X3C collection. Consequently, a polynomial-time solution to exact decision-CVPP on either family would imply $\mathsf{NP} \subseteq \mathsf{P}/\mathrm{poly}$ and collapse the polynomial hierarchy to $\Sigma_2^{\mathsf{P}}$. To our knowledge, the cyclic results answer Micciancio's questions of whether exact decision-CVP is NP-hard on cyclic lattices and on a fixed family of cyclic lattices, even when restricted to full-rank principal cyclic ideals.

Finally, under the coefficient embedding, we prove that exact decision-module-SIVP is NP-complete on free rank-two modules over the same cyclotomic rings.

# 1 Introduction

A *lattice* is the set of all integer combinations of linearly independent vectors. Three basic lattice problems are the *shortest vector problem* (SVP), the *closest vector problem* (CVP), and the *shortest independent vectors problem* (SIVP). SVP asks for a shortest nonzero lattice vector. CVP asks for a lattice point closest to a given target. In dimension $n$, SIVP asks for $n$ linearly independent lattice vectors whose largest norm is as small as possible. These problems are central in the study of lattices and in lattice-based cryptography [MG02].

Cryptographic schemes often use lattices with extra algebraic structure. An *ideal* in a polynomial quotient ring is closed under multiplication by every ring element, and a *principal ideal* is generated by a single element. In coefficient coordinates, ring multiplication becomes polynomial convolution. For example, ideals in $\mathbb{Z}[X]/(X^D - 1)$ give lattices closed under cyclic rotation. This structure gives shorter descriptions and faster arithmetic. It is used in lattice-based cryptography, including the NTRU cryptosystem, Ring-LWE-based schemes, compact one-way functions, collision-resistant hashing, and fully homomorphic encryption [HPS98, LPR10, Mic02, PR06, LM06, Gen09]. The same structure may also enable specialized algorithms.

Module lattices form an intermediate regime between ordinary and ideal lattices. A rank-$r$ module over the ring of integers of a degree-$d$ number field has $\mathbb{Z}$-rank $rd$. Rank one gives an ideal lattice, while taking the base field to be $\mathbb{Q}$ gives an ordinary lattice. Langlois and Stehlé introduced Module-SIS and Module-LWE and gave reductions from worst-case problems on module lattices to these average-case problems [LS15]. Fixed small module rank is now central in practice, as illustrated by the standardized ML-KEM and ML-DSA schemes [Nat24b, Nat24a].

For general lattices, much is known about hardness. Exact Euclidean decision-CVP is NP-complete [vEB81]. CVP is also NP-hard to approximate within almost-polynomial factors [DKS98]. Exact Euclidean SVP is NP-hard under randomized reductions [Ajt98]. Approximating Euclidean SVP within some constant factor is also NP-hard under randomized reductions [Mic98]. Hair and Sahai recently proved deterministic NP-hardness of exact SVP in the $\ell_p$-norm for every $p > 2$ [HS26b], and almost-polynomial-factor hardness for every finite $p \geq 2$ under deterministic subexponential-time reductions [HS26a]. These reductions can build arbitrary lattice bases and need not preserve the ideal or principal ideal structure. Earlier work gave reductions between problems on ideal lattices, but these reductions do not prove the NP-hardness of exact CVP on principal ideals [PR07]. This gap matters because cyclotomic structure can lead to special algorithms. Examples include recovering short generators of some principal ideals [CDPR16] and, under number-theoretic assumptions, finding mildly short vectors in cyclotomic ideal lattices in quantum polynomial time [CDW21].

Classical hardness results are also known for SIVP on general lattices [BS99, AC21]. For structured lattices, random self-reducibility was first established for approximate Ideal-SVP [dBDPMW20]. Under ERH, this was later extended to approximate SIVP on module lattices of any fixed rank [dBPTW26], yielding a polynomial-time worst-case-to-average-case reduction over a discretized invariant distribution with a polynomially bounded loss in the approximation factor. The exact worst-case complexity of SIVP on explicit fixed-rank module families is a separate question.

*CVP with preprocessing* (CVPP) asks what happens when the lattice is known before the target. The offline step may take unlimited time but store only polynomially many bits, while the online step must answer target-and-threshold queries in polynomial time. Micciancio showed that exact decision-CVPP has no polynomial-time solution on general lattices unless $\mathsf{NP} \subseteq \mathsf{P/poly}$ [Mic01]. For approximate CVPP, constant-factor approximation is hard [FM04]. Stronger inapproximability factors are known under quasipolynomial-time hardness assumptions [AKKV05, KPV12]. Separately, Micciancio asked whether CVP is NP-hard on cyclic lattices and whether it remains hard for a fixed family with one cyclic lattice for each input size [Mic07].

This leads to our main questions.

> *Is exact Euclidean CVP* NP*-hard for ideal and cyclic lattices, even for full-rank principal ideals? Does this hardness still hold for a uniformly computable family with one ideal per input size, fixed before the target is known?*

In the module setting, we additionally ask whether exact SIVP remains hard when the module rank is fixed at two.

### 1.1 Our results

Our answer to these questions is yes. Our main result concerns principal ideals in the power-of-two cyclotomic rings

$$R_d := \mathbb{Z}[y]/(y^d + 1),$$

where $d$ is a power of two. We write $K_d := \mathbb{Q}[y]/(y^d + 1)$ for the fraction field of $R_d$.

**Theorem 1.1** (Principal cyclotomic ideals, informal)**.** *For $d$ ranging over powers of two, exact Euclidean decision-CVP is* NP*-complete on the coefficient lattices of nonzero principal ideals in $R_d$.*

We use a deterministic polynomial-time many-one reduction from Exact Cover by 3-Sets (X3C). The reduction produces a nonzero principal ideal $(g) \subseteq R_d$, a target $T \in R_d$, and a squared threshold $\Delta \in \mathbb{Z}_{\geq 0}$. Its closest squared distance satisfies

$$\begin{aligned} \operatorname{dist}(T, (g))^2 &= \Delta && \text{in a YES instance,} \\ \operatorname{dist}(T, (g))^2 &\geq \Delta + 4 && \text{in a NO instance.} \end{aligned}$$

Moreover, the ideal elements at squared distance at most $\Delta$ from the target under the coefficient embedding are in bijection with the exact covers of the X3C instance.

Every nonzero ideal in $R_d$ has full $\mathbb{Z}$-rank. By Lemma 2.6, the theorem also holds under the canonical embedding after multiplying all squared distances and thresholds by $d$. This bijection also gives NP-hardness of exact search-CVP under polynomial-time Turing reductions.

The construction also lifts to full-rank principal ideals in $S_D := \mathbb{Z}[X]/(X^D - 1)$, where $D := 2d$, and proves NP-completeness of exact decision-CVP on their coefficient lattices, which are cyclic lattices. Exact search-CVP on these lattices is NP-hard under polynomial-time Turing reductions.

For each X3C universe size $m$, we also give one uniformly computable principal cyclotomic ideal and one principal cyclic ideal, both of dimension $\Theta(m^7)$. Only the target and squared threshold depend on the triple collection. Exact decision-CVP remains NP-complete on both fixed families. Unless $\mathsf{NP} \subseteq \mathsf{P/poly}$, exact decision-CVPP has no polynomial-time solution on either family. Consequently, the polynomial hierarchy would collapse to $\Sigma_2^{\mathsf{P}}$ if such a solution existed [KL80]. To our knowledge, these cyclic results answer Micciancio's two open questions about exact CVP.

We further transfer the principal cyclotomic CVP instances to exact module-SIVP on free rank-two modules over $R_d$. For every such module $\mathcal{M}$, the coefficient lattice $\mathcal{L}(\mathcal{M})$ has dimension $2d$, and

$$\lambda_2^{K_d}(\mathcal{M}) = \lambda_{2d}\big(\mathcal{L}(\mathcal{M})\big).$$

Thus exact decision-module-SIVP coincides with standard exact decision-SIVP on the underlying coefficient lattice.

**Theorem 1.2** (Rank-two cyclotomic modules, informal)**.** *For $d$ ranging over powers of two, exact decision-module-SIVP is* NP*-complete on free rank-two modules over $R_d$.*

### 1.2 Technical overview

The main challenge is to control every element of the ideal. In a reduction to a general lattice, we may choose the basis vectors separately. For a principal ideal $(g) \subseteq R_d$ generated by $g$, however, the natural basis vectors are the *negacyclic shifts*

$$g, yg, \ldots, y^{d-1}g,$$

and every element of $(g)$ has the form $gq$ for some $q \in R_d$. Thus we must encode Boolean vectors by suitable choices of $q$ and rule out every other $gq$.

We achieve our goal by reducing from X3C [GJ79]. An X3C instance consists of a universe of size $m$ and a collection of $n$ three-element subsets. It asks whether some subcollection covers every element exactly once. We express this condition as an integer linear system in Boolean variables:

$$\mathbf{A}\boldsymbol{\xi} = \mathbf{b}, \qquad \mathbf{A} \in \{0,1\}^{m \times n}, \quad \mathbf{b} \in \{0,1\}^m, \quad \boldsymbol{\xi} \in \{0,1\}^n. \tag{1}$$

Here $\mathbf{A}$ is the incidence matrix. Its entry $A_{ij}$ equals one exactly when the $i$-th universe element belongs to the $j$-th set. We take $\mathbf{b} := \mathbf{1}_m$. A Boolean solution $\boldsymbol{\xi}$ specifies an exact cover. We encode $\boldsymbol{\xi}$ as a polynomial with Boolean coefficients:

$$a_{\boldsymbol{\xi}}(x) := \sum_{k=1}^{n} \xi_k x^{\alpha_k}.$$

**Encoding the Boolean equations.** We choose the exponents $\alpha_1, \ldots, \alpha_n$ so that all nonzero ordered differences $\alpha_k - \alpha_j$ are pairwise distinct. A classical Sidon construction provides such exponents of size $O(n^2)$ [ET41]. We choose pairwise well-separated coefficient positions $\rho_1, \ldots, \rho_m$ for the $m$ rows and set

$$h(x) := \sum_{i,j:A_{ij}=1} x^{\rho_i - \alpha_j}.$$

The product $x^{\rho_i - \alpha_j} x^{\alpha_k}$ contributes at coefficient position

$$\rho_i - \alpha_j + \alpha_k.$$

When $k = j$, all contributions for row $i$ occur at $\rho_i$ and sum to $(\mathbf{A}\boldsymbol{\xi})_i$. When $k \neq j$, these products occupy pairwise distinct positions, all separate from the row positions $\rho_{i'}$.

To obtain a gap in coefficient $\ell_2$-distance, we define the target polynomial $t$ by assigning coefficient $2b_i$ at $x^{\rho_i}$ and coefficient one at every position $\rho_i - \alpha_j + \alpha_k$ with $A_{ij} = 1$ and $k \neq j$. The positions $\rho_i$ contribute $4\|\mathbf{A}\boldsymbol{\xi} - \mathbf{b}\|_2^2$. At every other specified position the difference is $2\xi_k - 1 \in \{-1, 1\}$. We denote the number of these positions by $\nu$, so their total squared contribution is $\nu$, independently of $\boldsymbol{\xi}$. This gives

$$\|2h(x)a_{\boldsymbol{\xi}}(x) - t(x)\|_2^2 = \nu + 4\|\mathbf{A}\boldsymbol{\xi} - \mathbf{b}\|_2^2.$$

Thus every Boolean vector satisfying (1) has squared distance $\nu$, whereas every Boolean vector violating at least one constraint in (1) has squared distance at least $\nu + 4$.

**Controlling every ideal element.** The identity above applies only to the polynomials $a_{\boldsymbol{\xi}}$ encoding Boolean vectors. To handle every $q \in R_d$, we place the construction in the even coordinates of a larger cyclotomic ring by replacing $x$ with $y^2$. The degree-$< d$ representative of $q$ then decomposes uniquely as

$$q(y) = a(y^2) + yb(y^2).$$

Here $a$ and $b$ are the integer polynomials formed by the even and odd coefficients of $q$, respectively. Define $u(x) := \sum_{k=1}^{n} x^{\alpha_k}$. For a sufficiently large integer $P$, we set

$$\begin{aligned} g(y) &:= 2h(y^2) + 2Py, \\ T(y) &:= t(y^2) + Py\,u(y^2), \\ \Delta &:= P^2 n + \nu. \end{aligned}$$

Separating $gq - T$ into its even- and odd-degree coefficients gives

$$gq - T = (2ha + 2Pxb - t)(y^2) + y(2hb + 2Pa - Pu)(y^2).$$

The terms involving $P$ in the two components above have squared norm $4P^2(\|b\|_2^2 + \|a - u/2\|_2^2)$. Over integral coefficients, this quantity is minimized exactly when $b = 0$ and $a = a_{\boldsymbol{\xi}}$ for some Boolean $\boldsymbol{\xi} \in \{0,1\}^n$. Every other integral pair increases this squared norm by at least $4P^2$. We choose $P$ large enough that the remaining terms cannot offset this increase, making every other $q$ have squared distance greater than $\Delta + 4$.

For a Boolean vector $\boldsymbol{\xi} \in \{0,1\}^n$, the two components give

$$\|g a_{\boldsymbol{\xi}}(y^2) - T\|_2^2 = \Delta + 4\|\mathbf{A}\boldsymbol{\xi} - \mathbf{b}\|_2^2.$$

Thus the elements of $(g)$ at squared distance at most $\Delta$ from $T$ are exactly those encoding solutions, and every other element has squared distance at least $\Delta + 4$.

Full rank follows because $R_d$ is an integral domain and the generator is nonzero. Fourier orthogonality shows that squared distances under the canonical embedding are $d$ times those under the coefficient embedding.

**Cyclic and fixed-family extensions.** For the cyclic case, we convert each constructed principal cyclotomic ideal into a principal cyclic ideal of twice the dimension, while multiplying the squared distances of corresponding elements by eight. For the fixed families, we use the matrix of all $\binom{m}{3}$ triples. A mask $\boldsymbol{\mu}$ indicates which triples belong to the input collection, and we replace $u(x)$ by $u_{\boldsymbol{\mu}}(x) := \sum_k \mu_k x^{\alpha_k}$. With this replacement, the odd-coordinate penalty is minimized only if $\boldsymbol{\xi} \le \boldsymbol{\mu}$. Hence absent triples cannot be selected. The resulting principal cyclotomic and cyclic ideals depend only on $m$. The details appear in Section 3.3 and Section 4, respectively.

**The module-SIVP extension.** We start with an exact-CVP instance $(g, T, \Delta)$ on a principal cyclotomic ideal. Here $g \in R_d \setminus \{0\}$ generates the ideal $(g)$, $T \in R_d$ is the target, and $\Delta \in \mathbb{Z}_{\ge 0}$ is the squared-distance threshold. Choose an integer $Q$ with $Q^2 > \max\{\|T\|_2^2, \|g\|_2^2\}$ and form

$$\mathcal{M}_{Q,g,T} := R_d(Q, T) + R_d(0, g) \subseteq K_d^2.$$

Thus $\mathcal{M}_{Q,g,T}$ is a free rank-two $R_d$-module, while its coefficient lattice $\mathcal{L}(\mathcal{M}_{Q,g,T})$ is a rank-two cyclotomic module lattice. The submodule $R_d(0, g)$ spans a one-dimensional $K_d$-subspace, so any two module vectors linearly independent over $K_d$ must include a vector outside this subspace. The choice of $Q$ forces a shortest such vector to have the form $\pm y^j(Q, T) + b(0, g)$ for some $b \in R_d$ and $0 \le j < d$. Minimizing the norm of its second coordinate then gives

$$\left(\lambda_2^{K_d}(\mathcal{M}_{Q,g,T})\right)^2 = Q^2 + \operatorname{dist}(T, (g))^2.$$

Thus the squared-distance gap transfers directly to exact module-SIVP.

**Organization.** Section 2 reviews lattices and ideal lattices. Section 3 proves NP-completeness of exact decision-CVP on principal cyclotomic and principal cyclic ideal lattices. Section 4 proves that exact decision-CVPP has no polynomial-time solution on the resulting fixed principal-ideal families unless $\mathsf{NP} \subseteq \mathsf{P/poly}$. Section 5 proves NP-completeness of exact module-SIVP on free rank-two cyclotomic modules. Section 6 summarizes the results.

# 2 Preliminaries

**Notation.** For $r \geq 1$, write $[r] := \{1, \ldots, r\}$. We use bold lowercase Roman or Greek letters for vectors and bold uppercase letters for matrices. For a vector $\mathbf{v}$, its $k$-th coordinate is denoted by $v_k$. For a matrix $\mathbf{A}$, its entry in row $i$ and column $j$ is denoted by $A_{ij}$. Vector inequalities are coordinatewise. For $1 \leq p \leq \infty$, $\|\mathbf{v}\|_p$ is the usual $\ell_p$-norm, while $\|\mathbf{v}\|_0 := |\{i : v_i \neq 0\}|$ is the *support size*. We write $\mathbf{0}_r$ and $\mathbf{1}_r$ for the length-$r$ all-zero and all-ones vectors. For scalars, finite sets, and binary strings, $|\cdot|$ denotes absolute value, cardinality, and bit length, respectively.

## 2.1 Lattices

For integers $m \geq n \geq 1$, let $\mathbf{B} \in \mathbb{R}^{m \times n}$ have linearly independent columns. The *lattice generated by* $\mathbf{B}$ is defined by

$$\mathcal{L}(\mathbf{B}) := \mathbf{B}\mathbb{Z}^n = \{\mathbf{Bz} : \mathbf{z} \in \mathbb{Z}^n\}.$$

The columns of $\mathbf{B}$ form a *basis* of $\mathcal{L}(\mathbf{B})$. If $m = n$, then $\mathcal{L}(\mathbf{B})$ is called *full rank*.

For a lattice $\mathcal{L} \subseteq \mathbb{R}^m$ and $\mathbf{t} \in \mathbb{R}^m$, define

$$\operatorname{dist}(\mathbf{t}, \mathcal{L}) := \min_{\mathbf{v} \in \mathcal{L}} \|\mathbf{t} - \mathbf{v}\|_2.$$

**Definition 2.1** (Exact CVP)**.** Given a rational lattice basis $\mathbf{B}$, a rational target $\mathbf{t}$, and a nonnegative rational squared threshold $\Delta$, exact decision-CVP asks whether

$$\operatorname{dist}\big(\mathbf{t}, \mathcal{L}(\mathbf{B})\big)^2 \leq \Delta.$$

Given $\mathbf{B}$ and $\mathbf{t}$, exact search-CVP asks for a vector $\mathbf{v} \in \mathcal{L}(\mathbf{B})$ satisfying

$$\|\mathbf{t} - \mathbf{v}\|_2 = \operatorname{dist}\big(\mathbf{t}, \mathcal{L}(\mathbf{B})\big).$$

The exact shortest vector problem (SVP) is the corresponding problem with $\mathbf{t} = \mathbf{0}_m$, except that the zero lattice vector is excluded.

**Definition 2.2** (Successive minima)**.** Let $\mathcal{L} \subseteq \mathbb{R}^m$ be a lattice of rank $n$, and let $\mathcal{B}_2^m := \{\mathbf{x} \in \mathbb{R}^m : \|\mathbf{x}\|_2 \leq 1\}$ be the closed Euclidean unit ball. For $1 \leq k \leq n$, the $k$-th successive minimum $\lambda_k(\mathcal{L})$ is the smallest radius $r$ such that $r\mathcal{B}_2^m$ contains $k$ linearly independent vectors of $\mathcal{L}$. Equivalently,

$$\lambda_k(\mathcal{L}) := \min\left\{r \geq 0 : \dim_{\mathbb{R}} \operatorname{span}_{\mathbb{R}}\big(\mathcal{L} \cap r\mathcal{B}_2^m\big) \geq k\right\}.$$

**Definition 2.3** (Exact SIVP)**.** Given a rational lattice basis $\mathbf{B} \in \mathbb{Q}^{m \times n}$ and a nonnegative rational squared threshold $\Delta$, exact decision-SIVP asks whether

$$\lambda_n\big(\mathcal{L}(\mathbf{B})\big)^2 \leq \Delta.$$

We use the following formulation of exact decision-CVPP, following Micciancio [Mic01].

**Definition 2.4** (Exact decision-CVPP)**.** An exact decision-CVPP scheme consists of a *preprocessing function* $\mathsf{Pre}$, a *uniform polynomial-time decoder* $\mathsf{Dec}$, and a polynomial $p$ with the following properties.

- Given a rational lattice basis $\mathbf{B}$, the *offline stage* outputs a string $\pi_{\mathbf{B}} := \mathsf{Pre}(\mathbf{B})$ satisfying $|\pi_{\mathbf{B}}| \le p(|\langle \mathbf{B} \rangle|)$, where $\langle \mathbf{B} \rangle$ is the standard binary encoding of $\mathbf{B}$. No running-time bound is imposed on $\mathsf{Pre}$.
- Given $\pi_{\mathbf{B}}$, a rational target $\mathbf{t}$, and a nonnegative rational squared threshold $\Delta$, the *online stage* satisfies
$$\mathsf{Dec}(\pi_{\mathbf{B}}, \mathbf{t}, \Delta) = 1 \quad \Longleftrightarrow \quad \operatorname{dist}(\mathbf{t}, \mathcal{L}(\mathbf{B}))^2 \le \Delta.$$
Its running time is polynomial in the total length of these inputs.

We say that exact decision-CVPP is solvable in polynomial time on a lattice family with one specified basis for each lattice if a single such scheme works for all these bases.

## 2.2 Algebraic Number Theory and Ideal Lattices

Let $K$ be a number field of degree $n$, and let $R := \mathcal{O}_K$ be its ring of integers.

**Coefficient embedding.** Fix an algebraic integer $\theta \in R$ such that $K = \mathbb{Q}(\theta)$, and let $F(z) \in \mathbb{Z}[z]$ be the monic minimal polynomial of $\theta$. Then $K \cong \mathbb{Q}[z]/(F(z))$. Every element $a \in K$ has a unique representation
$$a = a_0 + a_1\theta + \cdots + a_{n-1}\theta^{n-1}, \qquad a_0, \ldots, a_{n-1} \in \mathbb{Q}.$$
Its *coefficient embedding* with respect to the *power basis* $1, \theta, \ldots, \theta^{n-1}$ is
$$\operatorname{coeff}(a) := (a_0, a_1, \ldots, a_{n-1})^\top \in \mathbb{Q}^n.$$
If $I \subseteq R$ is a nonzero ideal, then
$$\mathcal{L}(I) := \{\operatorname{coeff}(a) : a \in I\}$$
is a full-rank lattice in $\mathbb{R}^n$. For $T \in K$, define
$$\operatorname{dist}(T, I) := \operatorname{dist}\big(\operatorname{coeff}(T), \mathcal{L}(I)\big) = \min_{a \in I} \|\operatorname{coeff}(T - a)\|_2.$$

**Canonical embedding.** Let $\sigma_1, \ldots, \sigma_n : K \longrightarrow \mathbb{C}$ be all embeddings of $K$ into $\mathbb{C}$. The *canonical embedding* is defined by
$$\sigma(a) := (\sigma_1(a), \ldots, \sigma_n(a)) \in \mathbb{C}^n.$$
Because the complex embeddings occur in conjugate pairs, $\sigma(K)$ lies in a real subspace of $\mathbb{C}^n$ of dimension $n$. If $I \subseteq R$ is a nonzero ideal, then
$$\sigma(I) := \{\sigma(a) : a \in I\}$$
is a full-rank lattice in this real subspace. For $T \in K$, its distance from $I$ under the canonical embedding is
$$\operatorname{dist}\big(\sigma(T), \sigma(I)\big) := \min_{a \in I} \|\sigma(T - a)\|_2.$$

In general, the coefficient and canonical embeddings induce different Euclidean lengths. We now specialize to power-of-two cyclotomic fields.

**The power-of-two cyclotomic case.** Let $d$ be a power of two and put $\zeta := e^{\pi i/d}$, a primitive $2d$-th root of unity. We specialize to

$$K_d := \mathbb{Q}(\zeta), \qquad R_d := \mathcal{O}_{K_d} = \mathbb{Z}[\zeta] \cong \mathbb{Z}[y]/(y^d + 1).$$

In a quotient by a monic polynomial of degree $r$, we identify each element with its unique representative of degree less than $r$. Every $f \in K_d$ has a unique representation

$$f = f_0 + f_1 y + \cdots + f_{d-1} y^{d-1}, \qquad f_0, \ldots, f_{d-1} \in \mathbb{Q},$$

with coefficient vector

$$\operatorname{coeff}(f) = (f_0, \ldots, f_{d-1})^\top.$$

The $d$ canonical embeddings of $K_d$ send $y$ to $\zeta^{2k+1}$ for $k = 0, \ldots, d-1$. Accordingly, we write

$$\sigma_d(f) := \big(f(\zeta), f(\zeta^3), \ldots, f(\zeta^{2d-1})\big).$$

For $1 \le p \le \infty$, we use $\|f\|_p := \|\operatorname{coeff}(f)\|_p$ for the *coefficient $\ell_p$-norm* of $f \in K_d$.

We will use *Young's convolution inequality* in the following form. For every $r \ge 1$ and $f, g \in \mathbb{R}[z]/(z^r + 1)$, with coefficient norms taken on their degree-$< r$ representatives,

$$\|fg\|_2 \le \|f\|_1 \|g\|_2.$$

Indeed, multiplication by a monomial is a signed coordinate permutation, so the inequality follows from the triangle inequality.

For $g \in R_d$, write $(g) := gR_d = \{gq : q \in R_d\}$ for the principal ideal generated by $g$. We write $I \lhd R$ to mean that $I$ is an ideal of $R$. For $0 \ne g \in R_d$, let

$$\mathbf{B}_g := \big[\operatorname{coeff}(g) \;\; \operatorname{coeff}(yg) \;\; \cdots \;\; \operatorname{coeff}(y^{d-1} g)\big].$$

Since $R_d$ is a domain, these columns form a $\mathbb{Z}$-basis of $\mathcal{L}((g))$. We refer to $\mathbf{B}_g$ as its *multiplication basis*.

We now formulate exact CVP on principal cyclotomic ideal lattices.

**Definition 2.5** (Principal cyclotomic CVP)**.** Given a power of two $d$, a nonzero generator $g \in R_d$, a target $T \in K_d$, and a nonnegative rational squared threshold $\Delta$, principal cyclotomic exact decision-CVP asks whether

$$\operatorname{dist}(T, (g))^2 \le \Delta.$$

Given $d$, $g$, and $T$, principal cyclotomic exact search-CVP asks for $\operatorname{coeff}(v)$, where $v \in (g)$ satisfies

$$\|T - v\|_2 = \operatorname{dist}(T, (g)).$$

The inputs $g$ and $T$ are represented by $\operatorname{coeff}(g)$ and $\operatorname{coeff}(T)$, respectively. Under this representation, the corresponding standard decision-CVP instance is $(\mathbf{B}_g, \operatorname{coeff}(T), \Delta)$.

In $K_d$, the canonical and coefficient $\ell_2$-norms differ only by a fixed factor $\sqrt{d}$.

**Lemma 2.6.** *For every $f \in K_d$,*

$$\|\sigma_d(f)\|_2^2 = d \|\operatorname{coeff}(f)\|_2^2.$$

*Proof.* For $0 \le r, s < d$, Fourier orthogonality gives

$$\sum_{k=0}^{d-1} \zeta^{(2k+1)(r-s)} = \begin{cases} d, & r = s, \\ 0, & r \neq s. \end{cases}$$

Expanding the squared canonical norm therefore yields

$$\|\sigma_d(f)\|_2^2 = d \sum_{r=0}^{d-1} |f_r|^2 = d\|\operatorname{coeff}(f)\|_2^2.$$

□

**Remark 2.7.** Lemma 2.6 implies that, for every nonzero ideal $I \subseteq R_d$ and every target $T \in K_d$,

$$\operatorname{dist}\big(\sigma_d(T), \sigma_d(I)\big)^2 = d \operatorname{dist}(T, I)^2.$$

Thus exact SVP and CVP under the coefficient and canonical embeddings are equivalent after multiplying each squared threshold by $d$.

**Rank-two cyclotomic module lattices.** For $\mathbf{v} = (v_1, v_2) \in K_d^2$, define

$$\operatorname{coeff}_2(\mathbf{v}) := \big(\operatorname{coeff}(v_1)^\top, \operatorname{coeff}(v_2)^\top\big)^\top \in \mathbb{Q}^{2d}, \qquad \|\mathbf{v}\|_2 := \|\operatorname{coeff}_2(\mathbf{v})\|_2.$$

A free rank-two $R_d$-module in $K_d^2$ has the form

$$\mathcal{M} = R_d\mathbf{b}_1 + R_d\mathbf{b}_2,$$

where $\mathbf{b}_1, \mathbf{b}_2 \in K_d^2$ are linearly independent over $K_d$. Its *coefficient lattice*

$$\mathcal{L}(\mathcal{M}) := \{\operatorname{coeff}_2(\mathbf{v}) : \mathbf{v} \in \mathcal{M}\}$$

is a full-rank lattice in $\mathbb{R}^{2d}$, called a *rank-two cyclotomic module lattice*.

The module analogue of the second successive minimum is defined as

$$\lambda_2^{K_d}(\mathcal{M}) := \min\left\{r \ge 0 : \dim_{K_d} \operatorname{span}_{K_d}\{\mathbf{v} \in \mathcal{M} : \|\mathbf{v}\|_2 \le r\} \ge 2\right\}.$$

For every nonzero $\mathbf{v} \in \mathcal{M}$, the $d$ vectors $y^j\mathbf{v}$ $(0 \le j < d)$ have equal norm and are $\mathbb{R}$-linearly independent. Hence

$$\lambda_2^{K_d}(\mathcal{M}) = \lambda_{d+1}\big(\mathcal{L}(\mathcal{M})\big) = \lambda_{2d}\big(\mathcal{L}(\mathcal{M})\big).$$

For $d = 1$, this is the ordinary second successive minimum $\lambda_2$.

**Definition 2.8** (Cyclotomic module-SIVP)**.** Given a power of two $d$, an $R_d$-basis of a free rank-two module $\mathcal{M} \subseteq K_d^2$, and a nonnegative integer squared threshold $\Gamma$, exact decision-module-SIVP asks whether

$$\big(\lambda_2^{K_d}(\mathcal{M})\big)^2 \le \Gamma.$$

The input basis vectors are represented by their coefficient vectors under $\operatorname{coeff}_2$. For $\mathbf{v} = (v_1, v_2) \in K_d^2$, applying Lemma 2.6 componentwise gives

$$\|\big(\sigma_d(v_1), \sigma_d(v_2)\big)\|_2^2 = d\|\mathbf{v}\|_2^2.$$

Thus the coefficient norm and the componentwise canonical norm differ by the fixed factor $\sqrt{d}$.

### 2.3 Exact Cover by 3-Sets (X3C).

Fix a positive multiple $m$ of 3, and let $\mathcal{C} := \{C_1, \ldots, C_n\}$ be a collection of three-element subsets of $[m]$ (called *triples*). The *incidence matrix* $\mathbf{H} \in \{0,1\}^{m \times n}$ of the instance is defined by

$$H_{ij} := \begin{cases} 1 & \text{if } i \in C_j, \\ 0 & \text{otherwise.} \end{cases}$$

An *exact cover* is a vector $\boldsymbol{\xi} \in \{0,1\}^n$ satisfying the binary linear system

$$\mathbf{H}\boldsymbol{\xi} = \mathbf{1}_m,$$

where $\mathbf{1}_m \in \{0,1\}^m$ denotes the all-ones vector. X3C asks whether such an exact cover exists. Equivalently, the selected subfamily $\mathcal{C}' := \{C_j : \xi_j = 1\} \subseteq \mathcal{C}$ is pairwise disjoint and covers $[m]$ exactly once:

$$\bigcup_{C_j \in \mathcal{C}'} C_j = [m], \qquad C_j \cap C_{j'} = \varnothing \ \text{ for all distinct } C_j, C_{j'} \in \mathcal{C}'.$$

It is well known that X3C is NP-complete [GJ79].

## 3 Reduction from X3C to CVP on principal ideal lattices

We use the following masked generalization of X3C. For the general masked system below, $M$ denotes the number of equations. In the X3C specialization, $M = m$. Let $M, n \geq 1$, and consider the *masked Boolean system*

$$\mathbf{A}\boldsymbol{\xi} = \mathbf{b}, \qquad \boldsymbol{\xi} \in \{0,1\}^n, \qquad \boldsymbol{\xi} \leq \boldsymbol{\mu}, \tag{2}$$

where $\mathbf{A} \in \{0,1\}^{M \times n}$, $\mathbf{b} \in \{0,1\}^M$, and $\boldsymbol{\mu} \in \{0,1\}^n$.

In this paper, we call $\boldsymbol{\mu}$ the *mask*. The constraint $\boldsymbol{\xi} \leq \boldsymbol{\mu}$ permits $\xi_k = 1$ only when $\mu_k = 1$. The feasibility problem for masked Boolean systems is NP-complete, since X3C is the special case obtained by setting

$$(\mathbf{A}, \mathbf{b}, \boldsymbol{\mu}) := (\mathbf{H}, \mathbf{1}_m, \mathbf{1}_n).$$

We reduce this feasibility problem to exact decision-CVP on principal ideal lattices. Roughly speaking, for each Boolean vector $\boldsymbol{\xi} \leq \boldsymbol{\mu}$, we form a sparse polynomial $a_{\boldsymbol{\xi}}$ and construct a principal ideal $(g)$, a target $T_{\boldsymbol{\mu}}$, and a threshold $\Delta_{\boldsymbol{\mu}}$ with

$$\|g(y)a_{\boldsymbol{\xi}}(y^2) - T_{\boldsymbol{\mu}}(y)\|_2^2 = \Delta_{\boldsymbol{\mu}} + 4\|\mathbf{A}\boldsymbol{\xi} - \mathbf{b}\|_2^2.$$

We then add coefficients at odd powers of $y$. If $q$ is not of the form $a_{\boldsymbol{\xi}}(y^2)$ for any $\boldsymbol{\xi} \in \{0,1\}^n$ satisfying $\boldsymbol{\xi} \leq \boldsymbol{\mu}$, the resulting ideal element $gq$ has squared distance from $T_{\boldsymbol{\mu}}$ greater than $\Delta_{\boldsymbol{\mu}} + 4$. Thus the closest squared distance is $\Delta_{\boldsymbol{\mu}}$ for a YES instance and at least $\Delta_{\boldsymbol{\mu}} + 4$ for a NO instance.

### 3.1 Encoding the Boolean equations

We first choose coefficient positions $\alpha_1, \ldots, \alpha_n$ whose nonzero ordered differences are pairwise distinct. Such positions are obtained from the classical Erdős–Turán Sidon-sequence construction [ET41].

**Lemma 3.1.** *There exists a deterministic polynomial-time algorithm that, on input $n \geq 1$, outputs integers*

$$0 = \alpha_1 < \cdots < \alpha_n, \qquad \alpha_n = \Theta(n^2) \quad (n \geq 2),$$

*such that the ordered differences $\alpha_k - \alpha_j$ $(j, k \in [n],\ j \neq k)$ are pairwise distinct.*

*Proof.* For $n = 1$, take $\alpha_1 = 0$. Otherwise we can always choose an odd prime $p \in [n, 2n)$ by Bertrand's postulate. Let $r_j$ be the least residue of $(j-1)^2$ modulo $p$, and set

$$\alpha_j := 2p(j-1) + r_j.$$

This is the classical Erdős–Turán construction, which ensures that all nonzero ordered differences $\alpha_k - \alpha_j$ are pairwise distinct [ET41]. The construction runs in polynomial time. Finally, $p < 2n$ and $r_n < p$ give $\alpha_n = O(n^2)$, while $p \geq n$ gives

$$\alpha_n \geq 2p(n-1) \geq 2n(n-1).$$

Hence $\alpha_n = \Theta(n^2)$ for $n \geq 2$. □

Set

$$c := \alpha_n, \quad L := 2c + 1, \quad \rho_i := c + (i-1)L \quad (i \in [M]).$$

We call $\rho_i$ the *row coordinate* of row $i$. For $A_{ij} = 1$ and $k \neq j$, the position $\rho_i - \alpha_j + \alpha_k$ is a *cross-term coordinate*. Set

$$W := \sum_{i \in [M]} \sum_{j \in [n]} A_{ij}, \qquad \nu := (n-1)W.$$

Let $N$ be the least power of two greater than $\rho_M + c$, and work in $R_N^- := \mathbb{Z}[x]/(x^N + 1)$. We use degree-$< N$ representatives and their coefficient norms. Define

$$h(x) := \sum_{A_{ij}=1} x^{\rho_i - \alpha_j},$$

$$t(x) := 2\sum_{i=1}^{M} b_i x^{\rho_i} + \sum_{\substack{A_{ij}=1 \\ k \neq j}} x^{\rho_i - \alpha_j + \alpha_k},$$

$$a_{\boldsymbol{\xi}}(x) := \sum_{k=1}^{n} \xi_k x^{\alpha_k}.$$

**Lemma 3.2.** *For every Boolean vector $\boldsymbol{\xi} \in \{0,1\}^n$,*

$$\|2h(x)a_{\boldsymbol{\xi}}(x) - t(x)\|_2^2 = \nu + 4\|\mathbf{A}\boldsymbol{\xi} - \mathbf{b}\|_2^2.$$

*Proof.* Note that

$$2h(x)a_{\boldsymbol{\xi}}(x) = 2\sum_{A_{ij}=1} \sum_{k=1}^{n} \xi_k x^{\rho_i - \alpha_j + \alpha_k}.$$

Splitting the inner sum according to $k = j$ and $k \neq j$ yields

$$2h(x)a_{\boldsymbol{\xi}}(x) = 2\sum_{A_{ij}=1} \xi_j x^{\rho_i} + 2\sum_{\substack{A_{ij}=1 \\ k \neq j}} \xi_k x^{\rho_i - \alpha_j + \alpha_k}. \tag{3}$$

Subtracting $t(x)$ from (3) gives

$$2h(x)a_{\boldsymbol{\xi}}(x) - t(x) = 2\sum_{i=1}^{M} \Big( \sum_{j : A_{ij}=1} \xi_j - b_i \Big) x^{\rho_i} + \sum_{\substack{A_{ij}=1 \\ k \neq j}} (2\xi_k - 1) x^{\rho_i - \alpha_j + \alpha_k}. \tag{4}$$

Because $(\mathbf{A}\boldsymbol{\xi})_i = \sum_{j:A_{ij}=1} \xi_j$, the coefficient of $x^{\rho_i}$ in (4) is precisely $2((\mathbf{A}\boldsymbol{\xi})_i - b_i)$.

To compute the $\ell_2$-norm of (4), we must verify that all its monomials have distinct exponents. The row coordinates $\rho_i$ are pairwise distinct because $L > 0$.

*(i) Cross-term coordinates do not equal row coordinates.* A row coordinate is some $\rho_{i'} = c + (i'-1)L$. A cross-term coordinate is $\rho_i - \alpha_j + \alpha_k$ with $A_{ij} = 1$ and $k \neq j$. Equality of the two would require $\rho_{i'} - \rho_i = \alpha_k - \alpha_j$. If $i' = i$ this forces $\alpha_j = \alpha_k$, hence $j = k$, contradicting $k \neq j$. If $i' \neq i$ then $|\rho_{i'} - \rho_i| \geq L = 2c + 1$, whereas $|\alpha_k - \alpha_j| \leq c$ because $0 \leq \alpha_j, \alpha_k \leq c$. Thus equality is impossible.

*(ii) Cross-term coordinates for a fixed row are distinct.* Fix $i$ and suppose

$$\rho_i - \alpha_j + \alpha_k = \rho_i - \alpha_{j'} + \alpha_{k'} \quad \text{with } k \neq j,\ k' \neq j'.$$

Cancelling $\rho_i$ gives $\alpha_k - \alpha_j = \alpha_{k'} - \alpha_{j'}$. By Lemma 3.1 all nonzero ordered differences $\alpha_k - \alpha_j$ are distinct, so $(j, k) = (j', k')$.

*(iii) Cross-term coordinates for different rows are distinct.* All cross-term coordinates for row $i$ lie in $[\rho_i - c, \rho_i + c] = [(i-1)L, iL - 1]$. These intervals are disjoint for distinct rows.

Since $N$ is a power of two larger than $\rho_M + c$, all exponents lie in $[0, N)$ and no wrap-around modulo $x^N + 1$ occurs. Consequently, the two sums in (4) have disjoint supports. The second sum contains exactly $\nu = (n-1)W$ monomials, one for each triple $(i, j, k)$ with $A_{ij} = 1$ and $k \neq j$. Because $\xi_k \in \{0, 1\}$, each coefficient satisfies $(2\xi_k - 1)^2 = 1$. Therefore

$$\|2h(x)a_{\boldsymbol{\xi}}(x) - t(x)\|_2^2 = 4\|\mathbf{A}\boldsymbol{\xi} - \mathbf{b}\|_2^2 + \nu.$$

□

### 3.2 Controlling every ideal element

Lemma 3.2 applies only to the polynomials $a_{\boldsymbol{\xi}}$ encoding Boolean vectors. We now add coefficients at odd powers of $y$ so that every $q \in R_d$ not of the form $a_{\boldsymbol{\xi}}(y^2)$ with $\boldsymbol{\xi} \leq \boldsymbol{\mu}$ satisfies $\|gq - T_{\boldsymbol{\mu}}\|_2^2 > \Delta_{\boldsymbol{\mu}} + 4$.

**Lemma 3.3.** *The polynomial $h(x) = \sum_{A_{ij}=1} x^{\rho_i - \alpha_j}$ has exactly $W$ distinct monomials. Equivalently, the exponents $\rho_i - \alpha_j$ with $A_{ij} = 1$ are pairwise distinct.*

*Proof.* Suppose $(i, j) \neq (i', j')$, $A_{ij} = A_{i'j'} = 1$, and

$$\rho_i - \alpha_j = \rho_{i'} - \alpha_{j'}.$$

*Case 1:* $i = i'$. Then $\alpha_j = \alpha_{j'}$, so strict monotonicity of the $\alpha_j$'s gives $j = j'$, a contradiction.
*Case 2:* $i \neq i'$. The spacing of the row coordinates gives

$$|\rho_{i'} - \rho_i| \geq L = 2c + 1, \qquad |\alpha_{j'} - \alpha_j| \leq c,$$

which contradicts $\rho_{i'} - \rho_i = \alpha_{j'} - \alpha_j$.

Finally, $0 \leq \rho_i - \alpha_j \leq \rho_M < N$, so no wrap-around occurs in $R_N^-$. Hence all $W$ exponents are distinct. □

Since these $W$ monomials have unit coefficients, $\|h\|_1 = W$. Set

$$s := \lceil \sqrt{n+4} \rceil, \qquad P := 1 + s\big(sW + \lceil \|t\|_2 \rceil + \lceil \sqrt{\nu + 4} \rceil\big).$$

Let $d := 2N$ and identify $R_N^-$ with the even-coordinate subring of $R_d = \mathbb{Z}[y]/(y^d + 1)$ via $x \mapsto y^2$. For a mask $\boldsymbol{\mu} \in \{0,1\}^n$, put

$$u_{\boldsymbol{\mu}}(x) := \sum_{k=1}^{n} \mu_k x^{\alpha_k}$$

and define

$$g(y) := 2h(y^2) + 2Py, \tag{5}$$

$$T_{\boldsymbol{\mu}}(y) := t(y^2) + Py\, u_{\boldsymbol{\mu}}(y^2), \tag{6}$$

$$\Delta_{\boldsymbol{\mu}} := P^2 \|\boldsymbol{\mu}\|_0 + \nu. \tag{7}$$

Consider the principal ideal generated by $g$. Every $q \in R_d$ has a unique form $q = a(y^2) + yb(y^2)$ with $a, b \in R_N^-$. Direct multiplication gives

$$gq - T_{\boldsymbol{\mu}} = (2ha + 2Pxb - t)(y^2) + y(2hb + 2Pa - Pu_{\boldsymbol{\mu}})(y^2), \tag{8}$$

where both $2ha + 2Pxb - t$ and $2hb + 2Pa - Pu_{\boldsymbol{\mu}}$ belong to $R_N^-$.

**Lemma 3.4.** *If $\boldsymbol{\xi} \in \{0,1\}^n$ and $\boldsymbol{\xi} \le \boldsymbol{\mu}$, then*

$$\|g(y)a_{\boldsymbol{\xi}}(y^2) - T_{\boldsymbol{\mu}}(y)\|_2^2 = \Delta_{\boldsymbol{\mu}} + 4\|\mathbf{A}\boldsymbol{\xi} - \mathbf{b}\|_2^2. \tag{9}$$

*Every $q \in R_d$ not of the form $a_{\boldsymbol{\xi}}(y^2)$ for some $\boldsymbol{\xi} \in \{0,1\}^n$ with $\boldsymbol{\xi} \le \boldsymbol{\mu}$ satisfies*

$$\|g(y)q(y) - T_{\boldsymbol{\mu}}(y)\|_2^2 > \Delta_{\boldsymbol{\mu}} + 4.$$

*Proof.* Let $\omega := \|\boldsymbol{\mu}\|_0$ be the Hamming weight of $\boldsymbol{\mu}$.

*Case 1: $q = a_{\boldsymbol{\xi}}(y^2)$ for a Boolean vector $\boldsymbol{\xi} \le \boldsymbol{\mu}$.* Equation (8) gives

$$g(y)q(y) - T_{\boldsymbol{\mu}}(y) = \underbrace{2h(y^2)a_{\boldsymbol{\xi}}(y^2) - t(y^2)}_{\text{even block}} + \underbrace{Py\big(2a_{\boldsymbol{\xi}}(y^2) - u_{\boldsymbol{\mu}}(y^2)\big)}_{\text{odd block}}. \tag{10}$$

**The even block.** The even block in (10) is supported on the even coordinates, and the substitution $x = y^2$ preserves its coefficient $\ell_2$-norm. Hence Lemma 3.2 gives

$$\|2h(y^2)a_{\boldsymbol{\xi}}(y^2) - t(y^2)\|_2^2 = \|2h(x)a_{\boldsymbol{\xi}}(x) - t(x)\|_2^2 = \nu + 4\|\mathbf{A}\boldsymbol{\xi} - \mathbf{b}\|_2^2.$$

**The odd block.** The odd block in (10) is

$$Py\big(2a_{\boldsymbol{\xi}}(y^2) - u_{\boldsymbol{\mu}}(y^2)\big) = P\sum_{k=1}^{n}(2\xi_k - \mu_k)y^{2\alpha_k+1}.$$

For every $k$, the conditions $\boldsymbol{\xi} \le \boldsymbol{\mu}$ and $\boldsymbol{\xi}, \boldsymbol{\mu} \in \{0,1\}^n$ imply $(2\xi_k - \mu_k)^2 = \mu_k$. Moreover, $2\alpha_k + 1 \le 2c + 1 \le N < d$, so the corresponding odd-coordinate positions are distinct and do not wrap around in $R_d$. Hence the odd block in (10) has squared norm

$$P^2\sum_{k=1}^{n}(2\xi_k - \mu_k)^2 = P^2\|\boldsymbol{\mu}\|_0 = P^2\omega.$$

Since the even and odd coordinates are disjoint, adding the two contributions proves (9).

*Case 2: all remaining $q \in R_d$.* Write $q(y) = a(y^2) + yb(y^2)$ with $a, b \in R_N^-$, and set

$$\Phi(a,b) := \|b\|_2^2 + \Big\|a - \frac{u_{\boldsymbol{\mu}}}{2}\Big\|_2^2.$$

For $0 \le r < N$, write $a_r$ and $b_r$ for the coefficients of $x^r$ in $a$ and $b$, respectively. For each $k$ with $\mu_k = 1$, integrality gives $(a_{\alpha_k} - 1/2)^2 \ge 1/4$, with equality exactly when $a_{\alpha_k} \in \{0,1\}$. Each coefficient $a_r$ with $r \notin \{\alpha_k : \mu_k = 1\}$ contributes $a_r^2$ to $\Phi(a,b)$, and each coefficient $b_r$ contributes $b_r^2$. Consequently,

$$\Phi(a,b) = \frac{\omega}{4}$$

exactly when $b = 0$ and $a(x) = a_{\boldsymbol{\xi}}(x)$ for some Boolean vector $\boldsymbol{\xi} \le \boldsymbol{\mu}$. For every other pair, $\Phi(a,b)$ is larger by at least 1. Hence

$$\Phi(a,b) \ge \frac{\omega + 4}{4}.$$

Identify a polynomial in $R_d$ with the ordered pair of its even- and odd-degree coefficient blocks, equipped with the product Euclidean norm. Equation (8) becomes

$$gq - T_{\boldsymbol{\mu}} = 2P\big((xb, a) - (0, u_{\boldsymbol{\mu}}/2)\big) + 2(ha, hb) - (t, 0).$$

Multiplication by $x$ is an isometry in $R_N^-$, so $\|xb\|_2 = \|b\|_2$. Young's inequality and $\|h\|_1 = W$ give the first bound below, while the triangle inequality gives the second:

$$\|(ha, hb)\|_2 \le W\|(a,b)\|_2, \qquad \|(a,b)\|_2 \le \sqrt{\Phi(a,b)} + \frac{\sqrt{\omega}}{2}.$$

The reverse triangle inequality therefore yields

$$\|gq - T_{\boldsymbol{\mu}}\|_2 \ge 2(P - W)\sqrt{\Phi(a,b)} - W\sqrt{\omega} - \|t\|_2.$$

Since $P \ge 1 + s^2 W > W$, the lower bound on $\Phi(a,b)$ gives

$$\|gq - T_{\boldsymbol{\mu}}\|_2 \ge P\sqrt{\omega + 4} - W\big(\sqrt{\omega + 4} + \sqrt{\omega}\big) - \|t\|_2.$$

Because $0 \le \omega \le n$ and $s = \lceil \sqrt{n+4} \rceil$,

$$\sqrt{\omega + 4} - \sqrt{\omega} \ge \frac{2}{s}, \qquad \sqrt{\omega + 4} + \sqrt{\omega} \le 2s.$$

The definition of $P$ also gives the strict inequality

$$\frac{2P}{s} = \frac{2}{s} + 2sW + 2\lceil \|t\|_2 \rceil + 2\lceil \sqrt{\nu + 4} \rceil > 2sW + \|t\|_2 + \sqrt{\nu + 4}.$$

Combining these estimates gives

$$\begin{aligned} \|gq - T_{\boldsymbol{\mu}}\|_2 - P\sqrt{\omega} &\ge P\big(\sqrt{\omega + 4} - \sqrt{\omega}\big) - W\big(\sqrt{\omega + 4} + \sqrt{\omega}\big) - \|t\|_2 \\ &> \sqrt{\nu + 4}. \end{aligned}$$

Squaring and discarding the nonnegative cross term now gives

$$\|gq - T_{\boldsymbol{\mu}}\|_2^2 > \big(P\sqrt{\omega} + \sqrt{\nu + 4}\big)^2 \ge P^2\omega + \nu + 4 = \Delta_{\boldsymbol{\mu}} + 4.$$

This proves the second assertion. □

The preceding lemma yields the reduction for masked Boolean systems. Its X3C specialization gives the main hardness result, while the fact that the ideal is independent of the mask will be used for the fixed-family construction in Section 4. For a full-rank principal ideal $I = (g) \lhd R_d$, we use $(I, T, \Delta)$ as shorthand for the standard CVP input $(\mathbf{B}_g, \mathrm{coeff}(T), \Delta)$.

**Theorem 3.5.** *There is a polynomial-time reduction from the masked Boolean system (2) to exact decision-CVP that maps $(\mathbf{A}, \mathbf{b}, \boldsymbol{\mu})$ to $(I, T_{\boldsymbol{\mu}}, \Delta_{\boldsymbol{\mu}})$, where $g \in R_d \setminus \{0\}$ and $I := (g) \lhd R_d$ is a full-rank principal ideal with a multiplication basis depending only on $(\mathbf{A}, \mathbf{b})$. Here $d = O((M+1)n^2)$ is a power of two, and $T_{\boldsymbol{\mu}}$ and $\Delta_{\boldsymbol{\mu}}$ are integral.*

*Proof.* Equation (9) shows that the element $ga_{\boldsymbol{\xi}}(y^2)$ associated with a Boolean vector $\boldsymbol{\xi} \le \boldsymbol{\mu}$ has squared distance $\Delta_{\boldsymbol{\mu}}$ from $T_{\boldsymbol{\mu}}$ precisely when $\mathbf{A}\boldsymbol{\xi} = \mathbf{b}$, i.e., when $\boldsymbol{\xi}$ is a feasible solution of the masked Boolean system (2). If $\mathbf{A}\boldsymbol{\xi} \neq \mathbf{b}$, then the integrality of $\mathbf{A}\boldsymbol{\xi} - \mathbf{b}$ implies $\|\mathbf{A}\boldsymbol{\xi} - \mathbf{b}\|_2^2 \ge 1$, so the squared distance is at least $\Delta_{\boldsymbol{\mu}} + 4$. The second assertion of Lemma 3.4 gives the same lower bound for every remaining $q \in R_d$. Consequently,

$$\begin{aligned} \mathrm{dist}(T_{\boldsymbol{\mu}}, I)^2 &= \Delta_{\boldsymbol{\mu}} && \text{in a YES instance,} \\ \mathrm{dist}(T_{\boldsymbol{\mu}}, I)^2 &\ge \Delta_{\boldsymbol{\mu}} + 4 && \text{in a NO instance.} \end{aligned}$$

The coefficient of $y$ in $g$ is $2P \neq 0$, so $g \neq 0$. Since $R_d$ is an integral domain, multiplication by $g$ is injective. Hence $(g)$ has full $\mathbb{Z}$-rank. Moreover, the elements $v \in (g)$ satisfying $\|v - T_{\boldsymbol{\mu}}\|_2^2 \le \Delta_{\boldsymbol{\mu}}$ are in bijection with the solutions of (2).

Finally, $c = O(n^2)$ and $\rho_M + c = O((M+1)n^2)$, so $d = O((M+1)n^2)$. We also have $W \le Mn$, $\nu \le Mn^2$, and $\|t\|_2^2 = 4\|\mathbf{b}\|_0 + \nu$. Hence every constructed quantity has polynomial encoding length, and the reduction runs in polynomial time. The generator $g \in R_d \setminus \{0\}$ and its multiplication basis depend only on $(\mathbf{A}, \mathbf{b})$, whereas $T_{\boldsymbol{\mu}}$ and $\Delta_{\boldsymbol{\mu}}$ also depend on the mask $\boldsymbol{\mu}$. □

We now apply Theorem 3.5 with $(\mathbf{A}, \mathbf{b}, \boldsymbol{\mu}) := (\mathbf{H}, \mathbf{1}_m, \mathbf{1}_n)$.

**Theorem 3.6.** *Exact decision-CVP is* NP*-complete on the coefficient lattices of full-rank principal ideals in the power-of-two cyclotomic rings $R_d$. The hard instances have squared-distance gap $\Delta$ versus $\Delta + 4$.*

*Proof.* Theorem 3.5 gives a deterministic reduction from X3C and the claimed gap. For membership in NP, a YES certificate is $\mathrm{coeff}(v)$ for an element $v \in (g)$ satisfying $\|v - T\|_2^2 \le \Delta$. Write $\mathrm{coeff}(v) = (v_0, \ldots, v_{d-1})^\top$ and $\mathrm{coeff}(T) = (T_0, \ldots, T_{d-1})^\top$. Each coefficient satisfies $|v_r| \le |T_r| + \lceil\sqrt{\Delta}\rceil$, and membership in $(g)$ is verified by solving the nonsingular multiplication system $\mathbf{B}_g\mathbf{z} = \mathrm{coeff}(v)$ and checking $\mathbf{z} \in \mathbb{Z}^d$. The verifier also checks $\|v - T\|_2^2 \le \Delta$.

□

The same characterization of the elements satisfying the distance threshold shows that exact search-CVP on these ideals is NP-hard under polynomial-time Turing reductions. Indeed, one call to an exact search-CVP oracle returns a closest lattice vector $\mathrm{coeff}(v)$ for some $v \in (g)$. If $\|v - T\|_2^2 \le \Delta$, solving $\mathbf{B}_g\mathbf{z} = \mathrm{coeff}(v)$ recovers $q$ through $\mathbf{z} = \mathrm{coeff}(q)$. The coefficients of $q$ at $y^{2\alpha_k}$ give the exact cover. Otherwise the X3C instance is negative.

By Lemma 2.6, both results hold under the canonical embedding. Squared thresholds and gaps are multiplied by $d$, and the closest elements are unchanged.

### 3.3 Cyclic extension

Let $S_D := \mathbb{Z}[X]/(X^D - 1)$, identified with $\mathbb{Z}^D$ by coefficient vectors. Multiplication by $X$ rotates the coordinates, so every full-rank ideal of $S_D$ is a *cyclic lattice*.

The following lemma lifts the principal cyclotomic instances to principal cyclic ideals in $S_{2d}$. Here $I$ is the original principal ideal, $T$ is its target, and $H$ is the squared-distance threshold. Under the lift, squared distances are multiplied by 8, and points of $I$ within squared distance $H$ of $T$ correspond bijectively to points of the lifted ideal within squared distance $8H$ of the lifted target.

**Lemma 3.7.** *Let $d$ be a power of two, $I = (g) \lhd R_d$ with $g \neq 0$, and $T \in R_d$. Fix $H \geq 0$, choose an integer $B$ with $B^2 > 4H$, and set $D := 2d$. Let $\iota_d : R_d \to S_D$ copy the degree-$< d$ representative, and define*

$$G := (1 - X^d)\iota_d(g) + B(1 + X^d), \qquad \widehat{T} := 2(1 - X^d)\iota_d(T).$$

*Then $J := (G)$ is full rank, and*

$$\|2(1 - X^d)\iota_d(v) - \widehat{T}\|_2^2 = 8\|v - T\|_2^2 \qquad (v \in I).$$

*Every $w \in J$ satisfying $\|w - \widehat{T}\|_2^2 \leq 8H$ arises uniquely in this way.*

*Proof.* For an arbitrary $Q \in S_D$, write $Q = p + X^d q$, where $p, q \in \mathbb{Z}[X]$ have degree less than $d$. Put $r := p + q$ and $z := p - q$. We view these degree-$< d$ polynomials as elements of $R_d$ in products with $g$. Equivalently,

$$Q = \frac{1}{2}\big((1 + X^d)r + (1 - X^d)z\big). \tag{11}$$

If $F := GQ - \widehat{T} = F_0 + X^d F_1$, direct expansion in $R_d$ gives

$$\begin{aligned} F_0 &= Br + gz - 2T, \qquad F_1 = Br - (gz - 2T), \\ \|F\|_2^2 &= \|F_0\|_2^2 + \|F_1\|_2^2 = 2\|gz - 2T\|_2^2 + 2B^2\|r\|_2^2. \end{aligned} \tag{12}$$

If $r \neq 0$, then the integer coefficient vector of $r$ has squared norm at least one, so (12) is at least $2B^2 > 8H$. Hence $\|GQ - \widehat{T}\|_2^2 \leq 8H$ implies $r = 0$. Then $q = -p$ and $z = 2p$. Substituting into (11) gives

$$Q = \frac{1}{2}(1 - X^d)z = (1 - X^d)p.$$

Using $(1 - X^d)(1 + X^d) = 0$ and $(1 - X^d)^2 = 2(1 - X^d)$, we first get $GQ = 2(1 - X^d)\iota_d(g)p$. By the definition of $\iota_d(gp)$, the difference $\iota_d(g)p - \iota_d(gp)$ is divisible by $X^d + 1$. Since $(1 - X^d)(X^d + 1) = 1 - X^{2d} = 0$ in $S_D$, we obtain

$$GQ = 2(1 - X^d)\iota_d(gp), \qquad \|GQ - \widehat{T}\|_2^2 = 8\|gp - T\|_2^2.$$

Thus $v = gp$ maps to $2(1 - X^d)\iota_d(v)$. The coefficients of degrees $0, \ldots, d - 1$ in this image equal those of $2v$, so $v$ is uniquely determined.

Over $\mathbb{Q}$, the Chinese remainder theorem gives

$$S_D \otimes_{\mathbb{Z}} \mathbb{Q} \cong \mathbb{Q}[X]/(X^d - 1) \times \mathbb{Q}[X]/(X^d + 1), \qquad G \longmapsto (2B, 2g).$$

Since $d$ is a power of two, $X^d + 1$ is irreducible over $\mathbb{Q}$, so $\mathbb{Q}[X]/(X^d + 1)$ is a field. Both components are units: $2B$ is a nonzero scalar in the first factor, and $2g \neq 0$ in the second. Hence multiplication by $G$ is invertible over $\mathbb{Q}$, and $J$ has full rank. □

**Theorem 3.8.** *Exact decision-CVP for tuples $(D, G, \widehat{T}, \widehat{\Delta})$, where $D$ is a power of two, $G, \widehat{T} \in S_D$, $(G) \lhd S_D$ is full rank, and $\widehat{\Delta} \in \mathbb{Z}_{\geq 0}$, is* NP*-complete. The lattice basis is computed from the $D$ cyclic shifts of $G$. The reduction has squared-distance gap $\widehat{\Delta}$ versus $\widehat{\Delta} + 32$.*

*Proof.* Apply Lemma 3.7 to the instance from Theorem 3.6 with

$$H := \Delta + 4, \qquad B := 2\lceil\sqrt{H}\rceil + 1, \qquad \widehat{\Delta} := 8\Delta.$$

The closest squared distance equals $\widehat{\Delta}$ in a YES instance and is at least $\widehat{\Delta} + 32$ in a NO instance. All lifted data have polynomial encoding length. Membership in NP follows as in Theorem 3.6, using the cyclic shifts of $G$ as a basis. □

**Remark 3.9.** Theorem 3.8 settles one of Micciancio's open questions: exact Euclidean decision-CVP is NP-hard on cyclic lattices [Mic07]. The lift also gives NP-hardness of exact search-CVP on the same cyclic ideals under polynomial-time Turing reductions. Given a closest $\widehat{v} \in J$, if $\|\widehat{v} - \widehat{T}\|_2^2 > \widehat{\Delta}$, the X3C instance is negative. Suppose instead that $\|\widehat{v} - \widehat{T}\|_2^2 \leq \widehat{\Delta}$. Lemma 3.7 then gives a unique $v \in I$ with $\widehat{v} = 2(1 - X^d)\iota_d(v)$. The coefficients of $\widehat{v}$ in degrees $0, \ldots, d-1$ are twice the corresponding coefficients of $v$. Thus they determine $\mathrm{coeff}(v)$. Solving $\mathbf{B}_g \mathbf{z} = \mathrm{coeff}(v)$ now yields $\mathbf{z} = \mathrm{coeff}(q)$. The coefficients of $q$ at $y^{2\alpha_k}$ give the exact cover.

## 4 Fixed principal-ideal families and preprocessing

The closest vector problem with preprocessing (CVPP) [Mic01] allows all lattice-dependent work to be performed before the target is known. We show that exact decision-CVPP on the fixed principal-ideal family below has no polynomial-time solution unless $\mathsf{NP} \subseteq \mathsf{P/poly}$.

For a fixed lattice family, the preprocessing string is common to all target queries and can serve as a nonuniform advice string depending only on the input length. We therefore encode every X3C instance of a given size using only the target and threshold.

**The X3C encoding.** For $m \geq 3$, let $k_m := \binom{m}{3}$, and let $\mathbf{U}_m \in \{0,1\}^{m \times k_m}$ be the matrix whose columns are the incidence vectors of all three-element subsets of $[m]$, listed in lexicographic order. A mask $\boldsymbol{\mu} \in \{0,1\}^{k_m}$ specifies which of these triples belong to the input collection. The collection has an exact cover precisely when

$$\mathbf{U}_m \boldsymbol{\xi} = \mathbf{1}_m \quad \text{for some} \quad \boldsymbol{\xi} \in \{0,1\}^{k_m} \text{ with } \boldsymbol{\xi} \leq \boldsymbol{\mu}.$$

Thus $(\mathbf{U}_m, \mathbf{1}_m)$ depends only on $m$. The mask $\boldsymbol{\mu}$ encodes the input collection.

**Theorem 4.1.** *There exists a polynomial-time reduction from X3C to exact decision-CVP such that every X3C instance on a universe of size $m$ is mapped to an instance $(I_m, T_{\boldsymbol{\mu}}, \Delta_{\boldsymbol{\mu}})$, where $I_m := (g_m) \lhd R_{d_m}$ is a full-rank principal ideal and its multiplication basis $\mathbf{B}_{g_m}$ depends only on $m$.*

*Proof.* Apply Theorem 3.5 to the fixed system $(\mathbf{U}_m, \mathbf{1}_m)$. The reduction produces the ideal $I_m$ and its multiplication basis independently of the mask. The target $T_{\boldsymbol{\mu}}$ and threshold $\Delta_{\boldsymbol{\mu}}$ encode the collection. In the notation of Section 3, set $M := m$ and $n := k_m$. Thus $k_m = \Theta(m^3)$, and the explicit Sidon construction gives $c = \alpha_{k_m} = \Theta(k_m^2)$. Since $\rho_m + c = \Theta(mc)$, rounding up to a power of two gives $d_m = \Theta(mk_m^2) = \Theta(m^7)$. The closest squared distance is $\Delta_{\boldsymbol{\mu}}$ for a YES instance and at least $\Delta_{\boldsymbol{\mu}} + 4$ for a NO instance. □

Together with the certificate argument of Theorem 3.6, Theorem 4.1 also shows that exact decision-CVP on this fixed family is NP-complete. Its main consequence is the following preprocessing lower bound.

**Corollary 4.2.** *If exact decision-CVPP is solvable in polynomial time on the family $\{\mathcal{L}(I_m)\}_{m\geq 3}$, each of which is represented by the corresponding multiplication basis from Theorem 4.1, then*

$$\mathsf{NP} \subseteq \mathsf{P/poly}.$$

*Proof.* Let $(\mathsf{Pre}, \mathsf{Dec}, p)$ be an exact decision-CVPP scheme for $\{\mathcal{L}(I_m)\}_{m\geq 3}$, with each lattice represented by its multiplication basis $\mathbf{B}_{g_m}$. Define the language $L \subseteq \{0,1\}^*$ as follows. For every $m \geq 3$, its length-$k_m$ strings are precisely the masks $\boldsymbol{\mu} \in \{0,1\}^{k_m}$ for which the collection encoded by $\boldsymbol{\mu}$ admits an exact cover. No strings of any other length belong to $L$. An X3C collection on $[m]$ can be converted in polynomial time to its length-$k_m$ mask, and vice versa. Hence $L$ is NP-complete. Since $k_m = \binom{m}{3}$ is strictly increasing, the input length determines $m$. For length $k_m$, use as advice $\pi_m := \mathsf{Pre}(\mathbf{B}_{g_m})$. The basis has encoding length $m^{O(1)}$, hence $k_m^{O(1)}$. Thus $|\pi_m| \leq k_m^{O(1)}$, as required by Definition 2.4. Given $\boldsymbol{\mu}$, Theorem 4.1 computes $T_{\boldsymbol{\mu}}$ and $\Delta_{\boldsymbol{\mu}}$ in polynomial time. Set $\mathbf{t}_{\boldsymbol{\mu}} := \mathrm{coeff}(T_{\boldsymbol{\mu}})$. Definition 2.4 and the correctness of the reduction give

$$\mathsf{Dec}(\pi_m, \mathbf{t}_{\boldsymbol{\mu}}, \Delta_{\boldsymbol{\mu}}) = 1 \quad \Longleftrightarrow \quad \mathrm{dist}(\mathbf{t}_{\boldsymbol{\mu}}, \mathcal{L}(I_m))^2 \leq \Delta_{\boldsymbol{\mu}} \quad \Longleftrightarrow \quad \boldsymbol{\mu} \in L.$$

Therefore $L \in \mathsf{P/poly}$, implying $\mathsf{NP} \subseteq \mathsf{P/poly}$. □

By the Karp–Lipton theorem [KL80], such a preprocessing scheme would collapse the polynomial hierarchy to $\Sigma_2^{\mathsf{P}}$.

**Remark 4.3.** Analogues of Theorem 4.1 and Corollary 4.2 hold after lifting the fixed family to principal cyclic ideal lattices. Let $P_m, \nu_m$ be the values from Section 3 for $(\mathbf{U}_m, \mathbf{1}_m)$, and set $H_m := P_m^2 k_m + \nu_m + 4$ and $B_m := 2\lceil\sqrt{H_m}\rceil + 1$. We have $\Delta_{\boldsymbol{\mu}} + 4 \leq H_m$. Applying Lemma 3.7 with $H = H_m$ and $B = B_m$ yields one principal cyclic ideal for each $m$. Its basis consists of the cyclic shifts of its generator. Only the lifted target and threshold depend on the mask. The threshold is $8\Delta_{\boldsymbol{\mu}}$, and the squared-distance gap is 32.

A polynomial-time solution to exact decision-CVPP on this fixed cyclic family would imply $\mathsf{NP} \subseteq \mathsf{P/poly}$. Exact decision-CVP on the same family remains NP-complete, thereby answering the exact decision version of Micciancio's question of whether CVP is NP-hard even for a fixed family of cyclic lattices [Mic07].

## 5 Exact module-SIVP on free rank-two cyclotomic modules

This section reduces exact CVP on principal cyclotomic ideals to exact module-SIVP on free rank-two modules over the power-of-two cyclotomic rings $R_d$.

The following lemma gives the embedding used in the reduction and identifies the second $K_d$-linear successive minimum.

**Lemma 5.1.** *Let $g \in R_d \setminus \{0\}$, let $T \in R_d$, and choose an integer $Q > 0$ such that*

$$Q^2 > \max\{\|T\|_2^2, \|g\|_2^2\}.$$

*Set*

$$\mathcal{M}_{Q,g,T} := R_d(Q,T) + R_d(0,g) = \{(Qa, Ta + gb) : a, b \in R_d\}.$$

*Then $\mathcal{M}_{Q,g,T}$ is a free rank-two $R_d$-module with basis*

$$\begin{pmatrix} Q & 0 \\ T & g \end{pmatrix},$$

*and*

$$\left(\lambda_2^{K_d}(\mathcal{M}_{Q,g,T})\right)^2 = Q^2 + \operatorname{dist}(T,(g))^2.$$

*Proof.* The matrix in the statement has determinant $Qg \neq 0$ in $K_d$, so $(Q,T)$ and $(0,g)$ are $K_d$-linearly independent. Since these vectors span $\mathcal{M}_{Q,g,T}$ over $R_d$, they form an $R_d$-basis, and hence $\mathcal{M}_{Q,g,T}$ is free of rank two.

Let $V := K_d(0,g)$, and set

$$r^2 := Q^2 + \operatorname{dist}(T,(g))^2.$$

Every vector of $\mathcal{M}_{Q,g,T} \setminus V$ has the form $\mathbf{v}(a,b) := (Qa, Ta + gb)$ with $a \neq 0$. Since $a \in R_d \setminus \{0\}$, $\|a\|_2^2$ is a positive integer.

*Case 1:* $\|a\|_2^2 = 1$. Then $a = \pm y^j$ for some $0 \leq j < d$. Multiplication by $a$ is an isometry and $a^{-1} \in R_d$. Therefore

$$\|\mathbf{v}(a,b)\|_2^2 = Q^2 + \|T + g(ba^{-1})\|_2^2 \geq r^2.$$

*Case 2:* $\|a\|_2^2 \geq 2$. In this case,

$$\|\mathbf{v}(a,b)\|_2^2 \geq 2Q^2 > Q^2 + \operatorname{dist}(T,(g))^2 = r^2,$$

where $\operatorname{dist}(T,(g)) \leq \|T\|_2 < Q$. Thus every vector outside $V$ has norm at least $r$.

Let $c \in (g)$ be a closest element to $T$. Then $(Q, T-c) \in \mathcal{M}_{Q,g,T}$ has norm $r$ and lies outside $V$, while $\|(0,g)\|_2 < Q \leq r$. Since $(0,g)$ and $(Q,T-c)$ are linearly independent over $K_d$,

$$\lambda_2^{K_d}(\mathcal{M}_{Q,g,T}) = r.$$

□

We now apply Lemma 5.1 to the principal-ideal CVP instances from Theorem 3.6.

**Theorem 5.2.** *Exact decision-module-SIVP is* NP*-complete on free rank-two modules over the power-of-two cyclotomic rings $R_d$. Hardness holds even for lower-triangular $R_d$-bases of the form*

$$\begin{pmatrix} Q & 0 \\ T & g \end{pmatrix}, \qquad Q \in \mathbb{N}^*, \quad g, T \in R_d, \quad g \neq 0.$$

*Proof.* Apply Theorem 3.6 to obtain a principal-ideal CVP instance $(d, g, T, \Delta)$ satisfying

$$\begin{aligned} \operatorname{dist}(T, (g))^2 &= \Delta && \text{in a YES instance,} \\ \operatorname{dist}(T, (g))^2 &\geq \Delta + 4 && \text{in a NO instance.} \end{aligned}$$

Set

$$Q := 1 + \left\lceil \sqrt{\max\{\|T\|_2^2, \|g\|_2^2\}} \right\rceil, \qquad \mathcal{M} := \mathcal{M}_{Q,g,T}, \qquad \Gamma := Q^2 + \Delta.$$

Lemma 5.1 gives

$$\begin{aligned} \left(\lambda_2^{K_d}(\mathcal{M})\right)^2 &= \Gamma && \text{in a YES instance,} \\ \left(\lambda_2^{K_d}(\mathcal{M})\right)^2 &\geq \Gamma + 4 && \text{in a NO instance.} \end{aligned}$$

The basis of $\mathcal{M}$ and the threshold $\Gamma$ have polynomially bounded encoding length, so the construction is a polynomial-time reduction.

For membership in NP, a YES certificate consists of two vectors in $\mathcal{M}$ of squared norm at most $\Gamma$. The verifier checks membership of both coefficient vectors in $\mathcal{L}(\mathcal{M})$, their norm bounds, and $K_d$-linear independence by testing whether the corresponding $2 \times 2$ determinant is nonzero. After clearing the denominators of the input basis, the norm bound gives certificates of polynomial bit length. □

By the componentwise form of Lemma 2.6, the theorem also holds under the componentwise canonical embedding after multiplying squared thresholds by $d$.

## 6 Conclusion

In this work, we prove that exact Euclidean decision-CVP is NP-complete on the coefficient lattices of nonzero principal ideals in power-of-two cyclotomic rings. The construction also yields NP-hardness for exact search-CVP. By Lemma 2.6, both results also hold for the corresponding lattices under the canonical embedding after scaling squared distances and thresholds by $d$, the degree of the cyclotomic field. We extend both results to full-rank principal cyclic ideal lattices using a lift that scales squared distances exactly. We also construct uniformly computable fixed cyclotomic and cyclic principal-ideal families on which exact decision-CVP remains NP-complete. A polynomial-time solution to exact decision-CVPP on either family would imply $\mathsf{NP} \subseteq \mathsf{P/poly}$.

We also prove that exact decision-module-SIVP is NP-complete on free rank-two modules over the same cyclotomic rings.

**Use of generative AI.** The authors used GPT-5.6-sol (Codex) to assist with language editing, presentation, and preliminary verification of mathematical arguments during the drafting stage. All mathematical claims and proofs in this manuscript were checked by the authors, who take full responsibility for the content of the paper.

# References


[AC21] Divesh Aggarwal and Eldon Chung. A note on the concrete hardness of the shortest independent vector in lattices. *Information Processing Letters*, 167:106065, 2021.

[Ajt98] Miklós Ajtai. The shortest vector problem in $L_2$ is NP-hard for randomized reductions (extended abstract). In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC)*, pages 10–19. ACM, 1998.

[AKKV05] Mikhail Alekhnovich, Subhash A. Khot, Guy Kindler, and Nisheeth K. Vishnoi. Hardness of approximating the closest vector problem with pre-processing. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 216–225. IEEE Computer Society, 2005.

[BS99] Johannes Blömer and Jean-Pierre Seifert. On the complexity of computing short linearly independent vectors and short bases in a lattice. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC)*, pages 711–720, 1999.

[CDPR16] Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In *Advances in Cryptology—EUROCRYPT 2016, Part II*, volume 9666 of *Lecture Notes in Computer Science*, pages 559–585. Springer, 2016.

[CDW21] Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Mildly short vectors in cyclotomic ideal lattices in quantum polynomial time. *Journal of the ACM*, 68(2):1–26, 2021.

[dBDPMW20] Koen de Boer, Léo Ducas, Alice Pellet-Mary, and Benjamin Wesolowski. Random self-reducibility of ideal-SVP via arakelov random walks. In *Advances in Cryptology—CRYPTO 2020, Part II*, volume 12171 of *Lecture Notes in Computer Science*, pages 243–273. Springer, 2020.

[dBPTW26] Koen de Boer, Aurel Page, Radu Toma, and Benjamin Wesolowski. Average hardness of SIVP for module lattices of fixed rank. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC)*, pages 814–822, 2026.

[DKS98] Irit Dinur, Guy Kindler, and Shmuel Safra. Approximating-CVP to within almost-polynomial factors is NP-hard. In *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 99–111. IEEE Computer Society, 1998.

[ET41] Paul Erdős and Pál Turán. On a problem of Sidon in additive number theory, and on some related problems. *Journal of the London Mathematical Society*, 16(4):212–215, 1941.

[FM04] Uriel Feige and Daniele Micciancio. The inapproximability of lattice and coding problems with preprocessing. *Journal of Computer and System Sciences*, 69(1):45–67, 2004.

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 169–178. ACM, 2009.

[GJ79] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, San Francisco, 1979.

[HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *Algorithmic Number Theory*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer, 1998.

[HS26a] Isaac M. Hair and Amit Sahai. Deterministic hardness of approximation for SVP in all finite $\ell_p$ norms. In *Proceedings of the 67th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE Computer Society, 2026. To appear.

[HS26b] Isaac M. Hair and Amit Sahai. $\mathrm{SVP}_p$ Is Deterministically NP-Hard for All $p > 2$, Even to Approximate within a Factor of $2^{\log^{1-\varepsilon} n}$. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1716–1727. ACM, 2026.

[KL80] Richard M. Karp and Richard J. Lipton. Some connections between nonuniform and uniform complexity classes. In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing (STOC)*, pages 302–309, 1980.

[KPV12] Subhash A. Khot, Preyas Popat, and Nisheeth K. Vishnoi. $2^{\log^{1-\varepsilon} n}$ hardness for the closest vector problem with preprocessing. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC)*, pages 277–288. ACM, 2012.

[LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *Automata, Languages and Programming, Part II*, volume 4052 of *Lecture Notes in Computer Science*, pages 144–155. Springer, 2006.

[LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *Advances in Cryptology—EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23. Springer, 2010.

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

[MG02] Daniele Micciancio and Shafi Goldwasser. *Complexity of Lattice Problems: A Cryptographic Perspective*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, MA, 2002.

[Mic98] Daniele Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. In *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 92–98. IEEE Computer Society, 1998.

[Mic01] Daniele Micciancio. The hardness of the closest vector problem with preprocessing. *IEEE Transactions on Information Theory*, 47(3):1212–1215, 2001.

[Mic02] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 356–365. IEEE Computer Society, 2002.

[Mic07] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007.

[Nat24a] National Institute of Standards and Technology. Module-lattice-based digital signature standard. Technical Report FIPS 204, National Institute of Standards and Technology, 2024.

[Nat24b] National Institute of Standards and Technology. Module-lattice-based key-encapsulation mechanism standard. Technical Report FIPS 203, National Institute of Standards and Technology, 2024.

[PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Theory of Cryptography*, volume 3876 of *Lecture Notes in Computer Science*, pages 145–166. Springer, 2006.

[PR07] Chris Peikert and Alon Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*, pages 478–487, 2007.

[vEB81] Peter van Emde Boas. Another NP-complete partition problem and the complexity of computing short vectors in a lattice. Technical Report 81-04, Mathematisch Instituut, Universiteit van Amsterdam, 1981.