SVP Is NP-Hard for Some Rank-2 Cyclotomic Modules

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

Let q be an odd prime. Let ζ_q be a primitive qth root of unity, and put $K = \mathbb{Q}(\zeta_q)$, with ring of integers $\mathcal{O}_K = \mathbb{Z}[\zeta_q]$. As q ranges over odd primes, we prove that the decision version of the Shortest Vector Problem (SVP) in the canonical ℓ_2-norm is NP-complete on full-rank free submodules of $\mathcal{O}_K^2$ by giving a deterministic polynomial-time many-one reduction from Exact Cover by 3-Sets (X3C). The module rank is fixed at two. The main obstacle is closure under the action of $\mathcal{O}_K$. A module containing a nonzero vector also contains every scalar multiple of that vector by a nonzero element of $\mathcal{O}_K$, and some of these multiples may be shorter.

Three ideas overcome this obstacle. First, we map the Bennett–Peikert Reed–Solomon lattice to a principal cyclotomic ideal and use Wan's point-count estimates to prove that a coset of this ideal contains many binary coefficient representatives. Second, we construct a checker whose canonical squared norm encodes the X3C equations, with a controlled nonnegative error term. Third, the checker and a second module coordinate combine with a separation bound for ideal cosets to rule out short vectors in NO instances. Each constructed instance consists of an odd prime q, two integral generators whose 2×2 generator matrix has nonzero determinant, and an integer squared threshold. The construction also gives NP-hardness of search-SVP under deterministic polynomial-time Turing reductions with at most one oracle call.

1 Introduction

A lattice $\mathcal{L} \subseteq \mathbb{R}^m$ is the set of all integer linear combinations of linearly independent vectors. Given a basis of $\mathcal{L}$, the decision shortest vector problem (SVP) asks whether the lattice contains a nonzero vector of squared norm at most a given threshold. Its search version asks for a shortest nonzero lattice vector. In 1981, Van Emde Boas asked whether exact SVP is NP-hard [vEB81]. Later, Ajtai's breakthrough result established that exact Euclidean SVP is NP-hard under randomized reductions [Ajt98]. For Euclidean SVP on lattices of rank n, later randomized reductions established hardness of approximation within a factor of $1 + n^{-\varepsilon}$ for every fixed $\varepsilon > 0$ [CN98], within any fixed factor in $(1, \sqrt{2})$ [Mic01], and within every constant factor [Kho05, Mic12]. Stronger randomized-time assumptions yield superconstant and almost-polynomial factors [Kho05, HR12].

However, a deterministic NP-hardness proof for exact Euclidean SVP remained open for decades. Along one route, Hair and Sahai used probabilistically checkable proofs (PCPs) to obtain deterministic hardness for ℓ_p-norms with $p > 2$ [HS26b]. They also proved approximation hardness in the ℓ_p-norm for every fixed $1 \leq p < \infty$ under a subexponential-time assumption [HS26a]. In two recent preprints, Wan proved deterministic hardness for every constant approximation factor in the Euclidean norm [Wan26a], extending an earlier result for factors below $\sqrt{2}$, including the exact problem [Wan26c].

At an intuitive level, Micciancio's *local-density* program asks for an explicit lattice $\mathcal{L}$ and a coset $\mathbf{t} + \mathcal{L}$ containing many points of norm strictly below $\lambda_1(\mathcal{L})$ [Mic01, Mic12]. The minimum of $\mathcal{L}$ protects the soundness gap, while the abundance of short coset points allows the reduction to encode a solution to the input problem. For our deterministic reduction, these points must realize every prescribed assignment on a designated set of coordinates. Bennett and Peikert proposed a candidate by lifting Reed–Solomon codes [BP23]. Wan used finite-field point counts to prove that an explicit Reed–Solomon coset has the required representatives with prescribed coordinates [Wan26c]. Wan's result supplies the local-density statement needed for deterministic hardness results for general lattices in the ℓ_2-norm. It does not by itself show that the reduction remains sound when the output is required to have a prescribed cyclotomic module structure.

Under the canonical embedding of a number field, ideals yield ideal lattices [LPR13a]. Higher-rank modules similarly yield module lattices [LS15, FPMS22]. The ring action supports compact representations and fast arithmetic [HPS98, LPR13b], underlies Ring-LWE and Module-LWE [LPR13a, LS15], and enables specialized algorithms for module lattices [LPMSW19]. This algebraic structure is no longer just a theoretical design choice. In 2024, NIST standardized ML-KEM (FIPS 203), a module-lattice-based scheme derived from CRYSTALS–Kyber and based on Module-LWE [BDK+18, Nat24b]. It also standardized ML-DSA (FIPS 204), derived from CRYSTALS–Dilithium and relying on Module-LWE and a variant of Module-SIS [DKL+18, Nat24a].

The cryptographic importance of module lattices has motivated broad work on structured lattices. Prior work studies cyclic-lattice problems [Mic07], module unique-SVP and NTRU [FPMS22], probabilistic reductions for approximation problems from arbitrary modules to free modules [DMMPMT23], worst-case to average-case reductions for the approximate shortest independent vectors problem (SIVP) at fixed module rank under the Extended Riemann Hypothesis [dBPTW26], and the geometry of shortest vectors in random modules [GSVV25].

For the Closest Vector Problem (CVP), a recent preprint proves that exact Euclidean decision-CVP is NP-complete on coefficient lattices of nonzero principal ideals in power-of-two cyclotomic rings, with a lift to principal cyclic ideal lattices [LFP26]. To the best of our knowledge, no deterministic polynomial-time NP-hardness result for exact Euclidean SVP on full-rank free rank-two modules over rings of integers in cyclotomic fields was known before the first version of this work appeared.

This leads to the question addressed in this work:

Is Euclidean SVP NP*-hard for rank-two modules over cyclotomic rings?*

We answer this question affirmatively, even for full-rank free submodules of $\mathcal{O}_K^2$ with $K = \mathbb{Q}(\zeta_q)$ for odd primes q.

1.1 Our result

For every odd prime q, let $\zeta := \zeta_q$ be a primitive qth root of unity. Write $K := \mathbb{Q}(\zeta)$ and $\mathcal{O}_K := \mathbb{Z}[\zeta]$. For $\mathbf{y} = (y_1, y_2) \in K^2$, define the canonical product norm by

$$\|\mathbf{y}\|^2 := \sum_{\sigma : K \hookrightarrow \mathbb{C}} (|\sigma(y_1)|^2 + |\sigma(y_2)|^2).$$

We write $\|\cdot\|_2$ for the ℓ_2-norm of coefficient vectors. An $\mathcal{O}_K$-submodule $\mathcal{M} \subseteq \mathcal{O}_K^2$ is called *full rank* if it has finite index in $\mathcal{O}_K^2$.

We consider the following decision problem. An input consists of an odd prime q, two vectors $\mathbf{m}_1, \mathbf{m}_2 \in \mathcal{O}_K^2$ with $\det[\,\mathbf{m}_1\ \mathbf{m}_2\,] \neq 0$, and an integer $S \geq 0$. Each of the four ring coordinates of $\mathbf{m}_1$ and $\mathbf{m}_2$ is specified by its $q-1$ integer coefficients in the basis $1, \zeta, \ldots, \zeta^{q-2}$. For $\mathcal{M} := \mathcal{O}_K\mathbf{m}_1 + \mathcal{O}_K\mathbf{m}_2$, the question is whether

$$\lambda_1(\mathcal{M})^2 \leq S.$$

We call this problem *decision-*SVP *in the ℓ_2-norm on rank-two cyclotomic modules*. The corresponding search-SVP problem takes $q, \mathbf{m}_1, \mathbf{m}_2$ as input and asks for a shortest nonzero vector of $\mathcal{M}$. The determinant condition makes $\mathcal{M}$ a full-rank free submodule of $\mathcal{O}_K^2$ generated by two vectors.

Theorem 1.1. *As q ranges over odd primes, decision-*SVP *in the ℓ_2-norm on free rank-two cyclotomic modules is* NP*-complete. The hardness holds under deterministic polynomial-time many-one reductions.*

Our reduction is from Exact Cover by 3-Sets (X3C), which is NP-complete [GJ79]. Given an X3C instance, the reduction outputs a tuple $(q, \mathbf{m}_1, \mathbf{m}_2, S)$ of the form above. Here q is an odd prime, $\mathbf{m}_1, \mathbf{m}_2 \in \mathcal{O}_K^2$ generate the output module and have nonzero determinant, and S is an integer squared threshold. As a lattice, the output module has $\mathbb{Z}$-rank $2(q-1)$. The tuple has polynomial encoding length. Section 5.3 gives the choice of q and bounds the encoding length and running time.

This is a worst-case hardness result in which the cyclotomic ring varies with q. The theorem does not establish hardness for the power-of-two cyclotomic family, on average, or through a cryptographic security reduction. The same construction also gives NP-hardness of search-SVP.

Corollary 1.2. *As q ranges over odd primes, search-*SVP *in the ℓ_2-norm on free rank-two submodules of $\mathcal{O}_K^2$ is* NP*-hard. The hardness holds under deterministic polynomial-time Turing reductions with at most one oracle call.*

Subsequent work. Wan proved deterministic NP-completeness for exact decision-SVP on full-rank cyclic lattices in the coefficient ℓ_2-norm [Wan26b]. Martin gave a dimension-preserving deterministic polynomial-time reduction from general Euclidean SVP to ideal lattices over monogenic orders in totally real number fields under the canonical embedding [Mar26].

1.2 Proof overview

Our goal is to give a deterministic polynomial-time reduction from X3C to decision-SVP on rank-two modules over cyclotomic rings. The reduction directly constructs a rank-two $\mathcal{O}_K$-module $\mathcal{M} \subseteq \mathcal{O}_K^2$ and an integer squared threshold S such that

$$\lambda_1(\mathcal{M})^2 \le S \quad \Longleftrightarrow \quad \text{the X3C instance is positive.}$$

The main challenge is to prove the strict lower bound $\lambda_1(\mathcal{M})^2 > S$ for every negative instance despite closure under multiplication by $\mathcal{O}_K$.

X3C asks whether a given matrix $\mathbf{A} \in \{0,1\}^{M\times n}$ with exactly three ones in each column admits a vector $\boldsymbol{\xi} \in \{0,1\}^n$ satisfying

$$\mathbf{A}\boldsymbol{\xi} = \mathbf{1}_M, \tag{1}$$

where $\mathbf{1}_M$ is the all-one vector in $\{0,1\}^M$.

The proof of the reduction has three layers. First, we show that every $\boldsymbol{\xi}$ has an extension to a binary coefficient vector $\mathbf{x}$ whose evaluation at ζ lies in a fixed ideal coset in $\mathcal{O}_K$. Second, we construct an X3C checker whose squared norm encodes $\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M$. Third, we place the checker in a rank-two module. The reduction computes its output directly from $\mathbf{A}$ without finding $\mathbf{x}$.

The ideal coset. The reduction embeds $\boldsymbol{\xi} \in \{0,1\}^n$ in a longer binary coefficient vector $\mathbf{x} = (x_a)_{a\in\mathbb{F}_q} \in \{0,1\}^{\mathbb{F}_q}$ by choosing distinct field elements $\alpha_1, \dots, \alpha_n \in \mathbb{F}_q$ and requiring $x_{\alpha_j} = \xi_j$ for $1 \le j \le n$. Put $\pi := 1 - \zeta$. For a chosen integer k, define $\mathfrak{a} := \pi^k \mathcal{O}_K$. For $\mathbf{x} = (x_a)_{a\in\mathbb{F}_q} \in \mathbb{Z}^{\mathbb{F}_q}$, evaluation at ζ gives the corresponding ring element

$$[\mathbf{x}] := \sum_{a\in\mathbb{F}_q} x_a \zeta^a \in \mathcal{O}_K.$$

By Theorem 3.2, which uses the Bennett–Peikert minimum bound [BP23], evaluation maps a Reed–Solomon lattice to the ideal lattice $\mathfrak{a}$ and $\lambda_1(\mathfrak{a})^2 \ge 2kq$. Wan's point-count estimate [Wan26c], in the uniform form of Theorem 3.4, supplies many binary vectors whose corresponding ring elements lie in a fixed coset of $\mathfrak{a}$. The reduction chooses one coset $c + \mathfrak{a}$ with $c \in \mathcal{O}_K$. By Lemma 3.3 and Theorem 3.4, every $\boldsymbol{\xi} \in \{0,1\}^n$ has many extensions $\mathbf{x} \in \{0,1\}^{\mathbb{F}_q}$ such that

$$[\mathbf{x}] \in c + \mathfrak{a}, \qquad \sum_{a\in\mathbb{F}_q} x_a = h, \qquad x_{\alpha_j} = \xi_j \quad (1 \le j \le n),$$

where h is the prescribed Hamming weight. Thus $x_{\alpha_j} = \xi_j$ stores $\boldsymbol{\xi}$ in designated coordinates $\alpha_1, \dots, \alpha_n$ of $\mathbf{x}$, while $[\mathbf{x}]$ places the resulting ring element in the coset $c + \mathfrak{a}$. For the chosen parameters, $[\mathbf{x}]$ has squared canonical norm $H = h(q-h) < 2kq$ and is therefore shorter than every nonzero element of $\mathfrak{a}$.

The X3C checker. The second step constructs $U, V \in \mathcal{O}_K$ so that the squared canonical norm of $U[\mathbf{x}] - V$ encodes the residual $\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M$ of Equation (1). By Theorem 4.3, every weight-h binary vector $\mathbf{x}$ with $\xi_j = x_{\alpha_j}$ satisfies

$$\|U[\mathbf{x}] - V\|^2 = Q^2\big(B_0 + 2\kappa\|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2 + \delta(\mathbf{x})\big). \tag{2}$$

Here B_0 and the positive parameters Q, κ are independent of $\mathbf{x}$, and $\delta(\mathbf{x}) \ge 0$ is the total error. We call a representative $\mathbf{x}$ *admissible* if it satisfies the correlation conditions in Section 4.3, which ensure $\delta(\mathbf{x}) < \kappa/8$. The uniform count in Theorem 3.4 and the union bound in the proof of Theorem 4.4

show that every $\boldsymbol{\xi} \in \{0,1\}^n$ has an admissible representative $\mathbf{x}$ with $[\mathbf{x}] \in c + \mathfrak{a}$. Consequently, an exact cover gives an admissible representative whose squared checker norm is at most $Q^2(B_0 + \kappa/8)$. For a negative instance, every weight-h binary representative has squared checker norm at least $Q^2(B_0 + 2\kappa)$.

The rank-two module. The third step places the value $U[\mathbf{x}] - V$ in a module. Choose a positive integer Γ to control the second module coordinate. Define the rank-two $\mathcal{O}_K$-submodule of $\mathcal{O}_K^2$

$$\mathcal{M} := \mathcal{O}_K(U\pi^k, 0) + \mathcal{O}_K(Uc - V, -Q\Gamma) \subset \mathcal{O}_K^2.$$

Every vector in $\mathcal{M}$ has the unique form

$$\mathbf{v}(y, w) := (Uy - Vw, -Q\Gamma w), \qquad y, w \in \mathcal{O}_K, \quad y - wc \in \mathfrak{a}. \tag{3}$$

When $w = 1$ and $y = [\mathbf{x}]$, the first coordinate is the checker value $U[\mathbf{x}] - V$. The second coordinate contributes $Q^2\Gamma^2\|w\|^2$ to the squared norm. The integer threshold in Equation (23) satisfies $S = Q^2(B_0 + F + \kappa)$, where $F = \Gamma^2(q-1)$. If $\boldsymbol{\xi}$ is an exact cover, choose an admissible representative $\mathbf{x}$. Taking $w = 1$ in Equation (3) then gives $\|\mathbf{v}([\mathbf{x}], 1)\|^2 < S$.

Soundness under multiplication. Suppose that the X3C instance is negative, and consider a nonzero vector $\mathbf{v}(y, w) \in \mathcal{M}$. The main issue is possible cancellation between Uy and Vw. Since $y \in wc + \mathfrak{a}$, one has $\|y\| \geq \operatorname{dist}(wc, \mathfrak{a})$. For an explicit $\nu = o(1)$, Lemma 5.1 gives

$$\|\mathbf{v}(y, w)\|^2 \geq Q^2(1 - \nu)(\|y\|^2 + \Gamma^2\|w\|^2). \tag{4}$$

This inequality lets us use lower bounds for $\|y\|$ and $\|w\|$ separately and transfers the geometric bounds from the ideal cosets to the module. The argument considers four cases. The comparisons with S below are verified in the proof of Theorem 5.2.

(i) If $w = 0$, then $0 \neq y \in \mathfrak{a}$. Lemma 4.2 gives $\|Uy\| \geq Q\|y\|$, so the minimum bound for $\mathfrak{a}$ puts the squared norm of every nonzero module vector with $w = 0$ above S.

(ii) If w is a root of unity, write $y = wu$ with $u \in c + \mathfrak{a}$. Then $\mathbf{v}(y, w) = w\mathbf{v}(u, 1)$, and multiplication by w preserves the canonical norm. Let $\mathbf{x} \in \mathbb{Z}^{\mathbb{F}_q}$ be the unique integral coefficient representative of u with coefficient sum h from Lemma 3.3. If $\mathbf{x}$ is binary, the checker bound for a negative instance and the second coordinate put the squared norm of the module vector above S. If $\mathbf{x}$ is not binary, Lemma 3.3 gives $\|u\|^2 \geq H + 2q$. Applying Equation (4) to u and 1 again puts the squared norm of the module vector above S.

(iii) If w is not a root of unity and $0 < \|w\|^2 \leq 3q$, Theorem 3.7 gives a lower bound for $\operatorname{dist}(wc, \mathfrak{a})$, and part (ii) of Lemma 3.5 gives $\|w\|^2 \geq 2q - 4$. Combining these bounds with Equation (4) puts the squared norm of the module vector above S.

(iv) If w is not a root of unity and $\|w\|^2 > 3q$, the second coordinate $-Q\Gamma w$ alone puts the squared norm of the module vector above S.

Organization. Section 2 recalls Euclidean lattices, cyclotomic rings, X3C, and Sidon sequences, and records the power-sum estimates used to count binary representatives in ideal cosets. Section 3 maps the lifted Reed–Solomon lattice to a principal cyclotomic ideal, proves the fixed-coordinate count, and establishes the norm-separation bounds used for soundness. Section 4 constructs the X3C checker and proves the existence of admissible representatives. Section 5 gives the main reduction. It assembles the rank-two module, proves completeness and soundness, and bounds the encoding length and running time.

2 Preliminaries

Notation. For a positive integer r, write $[r] := \{1, \dots, r\}$. For a prime q, let $\mathbb{F}_q$ denote the finite field with q elements, and let $\mathbb{F}_q^\times := \mathbb{F}_q \setminus \{0\}$ denote its multiplicative group. When an element of $\mathbb{F}_q$ occurs as an exponent or in a binomial coefficient, we use its unique integer representative in $\{0, \dots, q-1\}$. Vectors and matrices are represented in bold, e.g., $\mathbf{v}$ for a vector and $\mathbf{A}$ for a matrix, while their scalar coordinates are not. We use $\|\cdot\|_2$ for the ℓ_2-norm and $\mathbf{1}_r$ for the vector of ones in $\mathbb{R}^r$. The notation $\mathfrak{a} \triangleleft A$ means that $\mathfrak{a}$ is an ideal of the ring A.

2.1 Euclidean lattices

A Euclidean lattice of rank n is a discrete additive subgroup $\mathcal{L} \subseteq \mathbb{R}^m$ of the form

$$\mathcal{L} = \mathcal{L}(\mathbf{B}) := \mathbf{B}\mathbb{Z}^n,$$

where $\mathbf{B} \in \mathbb{R}^{m \times n}$ has linearly independent columns. These columns form a basis. The first minimum of $\mathcal{L}$ is the length of a shortest nonzero lattice vector:

$$\lambda_1(\mathcal{L}) = \min_{\mathbf{0} \neq \mathbf{v} \in \mathcal{L}} \|\mathbf{v}\|_2.$$

Definition 2.1 (SVP)**.** The *decision*-SVP problem takes as input a basis $\mathbf{B} \in \mathbb{Q}^{m \times n}$ with linearly independent columns and a rational squared threshold $\Delta \geq 0$, and asks whether

$$\lambda_1(\mathcal{L}(\mathbf{B}))^2 \leq \Delta.$$

The problem belongs to NP: a YES certificate is a nonzero integer coefficient vector for a short lattice point. Standard cofactor bounds give a certificate of polynomial bit length, and exact rational arithmetic checks both membership and the squared norm. See, e.g., [MG02]. The corresponding search-SVP problem takes a full-column-rank basis $\mathbf{B}$ and outputs a shortest nonzero vector of $\mathcal{L}(\mathbf{B})$.

2.2 Cyclotomic fields and module lattices

Cyclotomic field. In this paper, we only consider odd primes q. For a fixed q, let $\zeta := \zeta_q \in \mathbb{C}$ be a primitive qth root of unity. We suppress the subscript q after fixing the prime and put

$$K := \mathbb{Q}(\zeta), \qquad \mathcal{O}_K := \mathbb{Z}[\zeta], \qquad \pi := 1 - \zeta.$$

Here $\mathcal{O}_K$ is the ring of integers of K. The elements $1, \zeta, \dots, \zeta^{q-2}$ form an integral basis of K, so $[K : \mathbb{Q}] = q - 1$. The prime q is totally ramified in $\mathcal{O}_K$, and hence $(q) = (\pi)^{q-1}$ [Was97].

Canonical embedding. For $v \in K$, write $\overline{v}$ for its complex conjugate. In particular, $\overline{\zeta} = \zeta^{-1}$. For $a \in \mathbb{F}_q^\times$, let $\sigma_a : K \hookrightarrow \mathbb{C}$ be the embedding defined by $\sigma_a(\zeta) = \zeta^a$. These are all the embeddings of K. The canonical embedding sends v to $(\sigma_a(v))_{a \in \mathbb{F}_q^\times}$. For each integer $d \geq 1$ and vectors $\mathbf{y} = (y_1, \dots, y_d), \mathbf{y}' = (y'_1, \dots, y'_d) \in K^d$, define the real inner product

$$\langle \mathbf{y}, \mathbf{y}' \rangle := \mathrm{Re} \sum_{i=1}^{d} \sum_{a \in \mathbb{F}_q^\times} \sigma_a(y_i)\overline{\sigma_a(y'_i)}$$

and its induced product norm

$$\|\mathbf{y}\|^2 := \langle \mathbf{y}, \mathbf{y} \rangle = \sum_{i=1}^{d} \sum_{a \in \mathbb{F}_q^\times} |\sigma_a(y_i)|^2.$$

For $d = 1$, this is the canonical norm on K:

$$\|v\|^2 = \sum_{a \in \mathbb{F}_q^\times} |\sigma_a(v)|^2.$$

The roots of unity in K are exactly $\pm\zeta^j$ for $0 \le j < q$ [Was97]. If ε is a root of unity, then $\|\varepsilon\mathbf{y} - \varepsilon\mathbf{y}'\| = \|\mathbf{y} - \mathbf{y}'\|$ for all $\mathbf{y}, \mathbf{y}' \in K^d$. Thus multiplication by ε is an isometry. In particular, it preserves distances.

Module lattices. For an integer $s \ge 1$, an *$\mathcal{O}_K$-module lattice* in K^s is a finitely generated $\mathcal{O}_K$-submodule $\mathcal{M} \subseteq K^s$, equipped with the componentwise canonical embedding and the product norm. It has *full rank* if its K-span is K^s. When $s = 1$, a nonzero ideal $\mathfrak{a} \subseteq \mathcal{O}_K$ gives a module lattice in K called an *ideal lattice*. Both module and ideal lattices are discrete additive subgroups of their real spans.

In this paper, we only use free module lattices of full rank. If $\mathbf{m}_1, \ldots, \mathbf{m}_s \in K^s$ are linearly independent over K, then

$$\mathcal{M} = \mathcal{O}_K\mathbf{m}_1 + \cdots + \mathcal{O}_K\mathbf{m}_s$$

is free of module rank s, has full rank, and has $\mathbb{Z}$-rank $s(q-1)$. The square matrix with columns $\mathbf{m}_1, \ldots, \mathbf{m}_s$ has nonzero determinant precisely when these vectors are linearly independent over K. If $\mathcal{M} \subseteq \mathcal{O}_K^s$, then full rank is equivalent to finite index in $\mathcal{O}_K^s$. Under the canonical embedding, the product norm is the Euclidean ℓ_2-norm. We specify the lattice by its module generators in power-basis coordinates.

The first minimum of an ideal lattice or module lattice is the length of its shortest nonzero vector in the canonical or product norm. For $v \in K$ and an ideal lattice $\mathfrak{a}$, define

$$\operatorname{dist}(v, \mathfrak{a}) := \min_{y \in \mathfrak{a}} \|v - y\|.$$

Coefficient representatives. For $\mathbf{z} = (z_a)_{a \in \mathbb{F}_q} \in \mathbb{Q}^{\mathbb{F}_q}$, write

$$[\mathbf{z}] := \sum_{a \in \mathbb{F}_q} z_a \zeta^a \in K.$$

If $\mathbf{z} \in \mathbb{Z}^{\mathbb{F}_q}$, then $[\mathbf{z}] \in \mathcal{O}_K$, and we call $\mathbf{z}$ an *integral coefficient representative* of $[\mathbf{z}]$. Since $\Phi_q(X) = 1 + X + \cdots + X^{q-1}$ is the minimal polynomial of ζ, two integral coefficient vectors represent the same element precisely when they differ by an integer multiple of the all-one vector $\mathbf{1}_q$.

The following identity converts the canonical inner product and norm into calculations with these q coefficients.

Lemma 2.2. *For all rational coefficient vectors $\mathbf{x}, \mathbf{y} \in \mathbb{Q}^{\mathbb{F}_q}$,*

$$\langle [\mathbf{x}], [\mathbf{y}] \rangle = q \sum_a x_a y_a - \left(\sum_a x_a \right) \left(\sum_a y_a \right).$$

In particular, for every $\mathbf{z} \in \mathbb{Z}^{\mathbb{F}_q}$,

$$\|[\mathbf{z}]\|^2 = q\sum_a z_a^2 - \left(\sum_a z_a\right)^2 = \sum_{a<b}(z_a - z_b)^2.$$

Here $a < b$ *uses the integer representatives* $0, \ldots, q-1$.

Proof. For $a, b \in \mathbb{F}_q$, a finite geometric sum gives

$$\sum_{t\in\mathbb{F}_q^\times} \zeta^{t(a-b)} = \begin{cases} q-1, & a = b, \\ -1, & a \neq b. \end{cases}$$

Hence, for rational coefficient vectors $\mathbf{x}$ and $\mathbf{y}$,

$$\langle[\mathbf{x}], [\mathbf{y}]\rangle = q\sum_a x_a y_a - \sum_{a,b} x_a y_b = q\sum_a x_a y_a - \left(\sum_a x_a\right)\left(\sum_b y_b\right).$$

Taking $\mathbf{x} = \mathbf{y} = \mathbf{z}$ gives the norm identity. Expanding $\sum_{a<b}(z_a - z_b)^2$ gives its final form. The inner-product formula is unchanged under replacing $\mathbf{x}$ by $\mathbf{x} + r\mathbf{1}_q$ and $\mathbf{y}$ by $\mathbf{y} + s\mathbf{1}_q$, for $r, s \in \mathbb{Q}$. Thus the identities do not depend on the choice of coefficient representatives. □

2.3 X3C and Sidon sequences

An Exact Cover by 3-Sets (X3C) instance consists of a universe $[M]$ and three-element subsets $C_1, \ldots, C_n \subseteq [M]$. Its incidence matrix $\mathbf{A} \in \{0,1\}^{M\times n}$ is defined by $A_{ij} = 1$ when $i \in C_j$ and $A_{ij} = 0$ otherwise. A *selection vector* $\boldsymbol{\xi} \in \{0,1\}^n$ specifies a choice of sets, with $\xi_j = 1$ exactly when C_j is selected. We call $\boldsymbol{\xi}$ an *X3C witness* if

$$\mathbf{A}\boldsymbol{\xi} = \mathbf{1}_M.$$

This equality says that every universe element belongs to exactly one selected set. The instance is positive exactly when it has an X3C witness, which specifies an exact cover of the universe. X3C is NP-complete [GJ79].

An instance with $M = 0$ is positive. An instance whose incidence matrix has a zero row is negative. Henceforth, assume $M, n \geq 1$ and that every row of $\mathbf{A}$ is nonzero. Then $M \leq 3n$.

To index the n selection bits in the reduction, we choose distinct nonnegative integers whose nonzero ordered differences are all distinct. Such a sequence is called a *Sidon sequence*. The following lemma gives a deterministic construction with largest entry $O(n^2)$.

Lemma 2.3. *For every* $n \geq 1$, *one can deterministically compute integers*

$$0 \leq \alpha_1 < \cdots < \alpha_n =: \alpha_\star, \qquad \alpha_\star = O(n^2),$$

such that all nonzero ordered differences $\alpha_\ell - \alpha_j$, $j \neq \ell$, *are distinct. The computation takes time polynomial in* n.

Proof. Use the deterministic quadratic-residue Sidon construction of Erdős and Turán [ET41]. For an explicit construction and the verification of the directed-difference property, see also Liu–Feng–Pan [LFP26]. □

2.4 Power-sum estimates

Let q be prime, let $2 \le k < m$, and let $\mathbf{b} = (b_1, \dots, b_{k-1}) \in \mathbb{F}_q^{k-1}$. Define

$$\mathcal{X}_{k,m}(\mathbf{b}) := \left\{ \mathbf{z} = (z_1, \dots, z_m) \in \mathbb{F}_q^m : \sum_{r=1}^{m} z_r^j = b_j \quad (1 \le j < k) \right\}.$$

The following proposition records the point-count bounds used below.

Proposition 2.4 ([Wan26c]). *Let q be prime and suppose that $2 \le k \le m-2 < q-2$. For every $\mathbf{b} \in \mathbb{F}_q^{k-1}$, the following estimates hold.*

(i)

$$\left| |\mathcal{X}_{k,m}(\mathbf{b})| - q^{m-k+1} \right| \le \frac{1}{2}(2k)^m q^{(m-k+2)/2}.$$

(ii) For distinct $i, \ell \in [m]$,

$$\left| \left| \{ \mathbf{z} \in \mathcal{X}_{k,m}(\mathbf{b}) : z_i = z_\ell \} \right| - q^{m-k} \right| \le \frac{1}{2}(2k)^{m-1} q^{(m-k+2)/2}.$$

(iii) For $i \in [m]$ and $a \in \mathbb{F}_q$,

$$\left| \left| \{ \mathbf{z} \in \mathcal{X}_{k,m}(\mathbf{b}) : z_i = a \} \right| - q^{m-k} \right| \le \frac{1}{2}(2k)^{m-1} q^{(m-k+1)/2}.$$

Proof. Wan [Wan26c, Propositions 5.2 and 6.2] proves part (i) and proves part (ii) for $z_1 = z_2$, respectively. The statements in that paper use one particular vector of prescribed power sums $\mathbf{b}$. The displayed estimates and their proofs depend only on q, k, and m, so they hold for every $\mathbf{b}$.

For distinct $i, \ell \in [m]$, a permutation of the m coordinates sends the condition $z_i = z_\ell$ to $z_1 = z_2$ and leaves each equation $\sum_{r=1}^m z_r^j = b_j$ unchanged. Therefore Wan's estimate for $z_1 = z_2$ proves part (ii) for every pair $i \ne \ell$.

For part (iii), fix $z_i = a$ and delete that coordinate. The remaining $m-1$ variables satisfy $\sum_{r \ne i} z_r^j = b_j - a^j$ for $1 \le j < k$. Since $k < m - 1 < q$, Wan's Proposition 5.2 [Wan26c] applies to these $m-1$ variables and proves part (iii). □

3 Reed–Solomon lattices and ideal cosets

Following the local-density approach in earlier SVP hardness reductions [Mic01, Mic12, BP23], we use an ideal lattice obtained from a Construction A Reed–Solomon lattice. We require one fixed ideal coset containing many short ring elements represented by weight-h binary coefficient vectors $\mathbf{x} \in \{0,1\}^{\mathbb{F}_q}$, with $[\mathbf{x}] = \sum_{a \in \mathbb{F}_q} x_a \zeta^a \in \mathcal{O}_K$. The first step of the reduction stores an X3C selection vector $\boldsymbol{\xi} \in \{0,1\}^n$ in designated coordinates of $\mathbf{x}$. Using Wan's point-count estimates [Wan26c], we show that for every $\boldsymbol{\xi}$ there are many such vectors $\mathbf{x}$ with $[\mathbf{x}]$ in the same fixed coset. At the same time, the ideal is designed to have a minimum larger than the norms of these coset elements, so that its nonzero elements cannot create unintended short vectors. This section constructs the ideal, proves the required count, and establishes the additional norm bounds for nonbinary representatives and scalar multiples.

3.1 From lattice to ideal

We first transfer an ℓ_2-minimum bound for the Construction A Reed–Solomon lattice to a principal ideal of $\mathcal{O}_K$.

Fix an ordering $\mathbb{F}_q = \{a_1, \dots, a_q\}$ and index the coordinates of $\mathbb{Z}^q$ accordingly. For $1 \le k < q$, define

$$\mathbf{H}_q(k) := \begin{pmatrix} 1 & 1 & \cdots & 1 \\ a_1 & a_2 & \cdots & a_q \\ a_1^2 & a_2^2 & \cdots & a_q^2 \\ \vdots & \vdots & \ddots & \vdots \\ a_1^{k-1} & a_2^{k-1} & \cdots & a_q^{k-1} \end{pmatrix} = (a^j)_{0 \le j < k,\ a \in \mathbb{F}_q} \in \mathbb{F}_q^{k \times q}.$$

The matrix $\mathbf{H}_q(k)$ defines the lattice

$$\mathcal{L}_{q,k} := \left\{ \mathbf{z} \in \mathbb{Z}^{\mathbb{F}_q} : \mathbf{H}_q(k)\mathbf{z} = \mathbf{0} \text{ in } \mathbb{F}_q^k \right\},$$

which is also known as the Construction A lattice of the Reed–Solomon code of dimension $q-k$ over $\mathbb{F}_q$. The following lemma gives a lower bound on the ℓ_2-minimum of $\mathcal{L}_{q,k}$.

Lemma 3.1 ([BP23, Theorem 14])**.** *Let q be prime and $1 \le k \le q/2$. Then the first minimum of $\mathcal{L}_{q,k}$ with respect to the ℓ_2-norm $\|\cdot\|_2$ satisfies*

$$\lambda_1(\mathcal{L}_{q,k})^2 \ge 2k.$$

Put $\mathfrak{a}_{q,k} := \pi^k \mathcal{O}_K$ and write $\mathbf{1}_q^\perp := \{\mathbf{z} \in \mathbb{R}^q : \mathbf{1}_q^{\mathsf{T}}\mathbf{z} = 0\}$. Let

$$\varphi : \mathbb{Z}[X]/(X^q - 1) \longrightarrow \mathcal{O}_K, \qquad X \longmapsto \zeta,$$

be the surjective ring homomorphism given by evaluation at ζ. Its kernel is generated by $\Phi_q(X) = 1 + X + \cdots + X^{q-1}$. We use the unique degree-$< q$ representative of each element $f \in \mathbb{Z}[X]/(X^q-1)$ and write $\bar{f} \in \mathbb{F}_q[X]$ for its reduction modulo q.

Theorem 3.2. *For $1 \le k \le q/2$, the map $\mathbf{z} \mapsto [\mathbf{z}]$ restricts to an additive-group isomorphism*

$$\mathcal{L}_{q,k} \cap \mathbf{1}_q^\perp \ \longrightarrow\ \mathfrak{a}_{q,k}.$$

Consequently, with respect to the canonical norm,

$$\lambda_1(\mathfrak{a}_{q,k})^2 \ge 2kq.$$

Proof. For $\mathbf{z} \in \mathbb{Z}^{\mathbb{F}_q}$, let $f := \sum_{a \in \mathbb{F}_q} z_a X^a \in \mathbb{Z}[X]/(X^q-1)$, so that $\varphi(f) = [\mathbf{z}]$. The coefficient of Y^j in $f(1+Y)$ is $\sum_a z_a \binom{a}{j}$. Thus

$$(X-1)^k \mid \bar{f} \quad \Longleftrightarrow \quad \sum_a z_a \binom{a}{j} = 0 \text{ in } \mathbb{F}_q \quad (0 \le j < k).$$

Since $k < q$, the functions $a \mapsto a^j$ and $a \mapsto \binom{a}{j}$, for $0 \le j < k$, are related by an invertible triangular change of basis over $\mathbb{F}_q$. Hence the defining equations of $\mathcal{L}_{q,k}$ are equivalent to $(X-1)^k \mid \bar{f}$. The identity $X^q - 1 = (X-1)^q$ in $\mathbb{F}_q[X]$ therefore gives

$$\mathbf{z} \in \mathcal{L}_{q,k} \quad \Longleftrightarrow \quad f \in (q, (X-1)^k) \lhd \mathbb{Z}[X]/(X^q-1).$$

The image of $(q,(X-1)^k)$ under φ is the ideal (q,π^k). Total ramification gives $(q)=(\pi^{q-1})$, so this image is $(\pi^k)=\mathfrak{a}_{q,k}$ since $k\le q-1$.

Furthermore, $\Phi_q(X)\equiv(X-1)^{q-1} \pmod q$, so $\ker(\varphi)=(\Phi_q)$ is contained in $(q,(X-1)^k)$. Since $\varphi((q,(X-1)^k))=\mathfrak{a}_{q,k}$, it follows that

$$\varphi^{-1}(\mathfrak{a}_{q,k})=(q,(X-1)^k) \quad \text{in } \mathbb{Z}[X]/(X^q-1). \tag{5}$$

To prove surjectivity of the restricted map, fix $v\in\mathfrak{a}_{q,k}$ and choose $f\in(q,(X-1)^k)$ with $\varphi(f)=v$. Since $f(1)\in q\mathbb{Z}$, replace f by $f-(f(1)/q)\Phi_q$. The new f remains in $(q,(X-1)^k)$, still satisfies $\varphi(f)=v$, and has coefficient sum zero. Its coefficient vector $\mathbf{z}$ therefore lies in $\mathcal{L}_{q,k}\cap\mathbf{1}_q^\perp$ and satisfies $[\mathbf{z}]=v$.

For injectivity, if $\mathbf{z}\in\mathcal{L}_{q,k}\cap\mathbf{1}_q^\perp$ satisfies $[\mathbf{z}]=0$, then $\mathbf{z}\in\mathbb{Z}\mathbf{1}_q$ by the description of $\ker(\varphi)$. Its coefficient sum is zero, so $\mathbf{z}=\mathbf{0}$.

Finally, for every nonzero $\mathbf{z}\in\mathcal{L}_{q,k}\cap\mathbf{1}_q^\perp$, Lemmas 3.1 and 2.2 give

$$\|[\mathbf{z}]\|^2=q\|\mathbf{z}\|_2^2\ge 2kq.$$

The isomorphism therefore gives the claimed minimum bound. □

3.2 Binary representatives in ideal cosets

For $\mathbf{x}\in\mathbb{Z}^{\mathbb{F}_q}$, its *syndrome* $\mathbf{H}_q(k)\mathbf{x}\in\mathbb{F}_q^k$ has jth coordinate $\sum_{a\in\mathbb{F}_q}x_a a^j$ modulo q for $0\le j<k$. When $\mathbf{x}$ is binary, these coordinates are the power sums of the positions a for which $x_a=1$. Lemma 3.3 identifies the vectors with fixed syndrome and coefficient sum with one coset of $\mathfrak{a}_{q,k}$.

Lemma 3.3. *Assume $1\le k\le q/2$. Fix $h\in\mathbb{Z}$. Let $\mathbf{u}=(u_0,\dots,u_{k-1})\in\mathbb{F}_q^k$ satisfy $u_0=h \bmod q$. There exists $\mathbf{z}_0\in\mathbb{Z}^{\mathbb{F}_q}$ such that*

$$\mathbf{H}_q(k)\mathbf{z}_0=\mathbf{u}, \qquad \mathbf{1}_q^{\mathsf{T}}\mathbf{z}_0=h.$$

Moreover, evaluation at ζ induces a bijection

$$\{\mathbf{z}\in\mathbb{Z}^{\mathbb{F}_q}:\mathbf{H}_q(k)\mathbf{z}=\mathbf{u},\ \mathbf{1}_q^{\mathsf{T}}\mathbf{z}=h\}\ \longrightarrow\ [\mathbf{z}_0]+\mathfrak{a}_{q,k}, \qquad \mathbf{z}\longmapsto[\mathbf{z}].$$

Put $H:=h(q-h)$. For each $v\in[\mathbf{z}_0]+\mathfrak{a}_{q,k}$, let $\mathbf{z}$ be its unique representative with syndrome $\mathbf{u}$ and coefficient sum h. If $\mathbf{z}\in\{0,1\}^{\mathbb{F}_q}$, then it has weight h and $\|v\|^2=H$. Otherwise, $\|v\|^2\ge H+2q$.

Proof. A Vandermonde $k\times k$ minor of $\mathbf{H}_q(k)$ has determinant $\prod_{1\le i<j\le k}(a_j-a_i)\neq 0$ in $\mathbb{F}_q$. Hence the matrix has rank k. Choose $\mathbf{x}_0\in\mathbb{Z}^{\mathbb{F}_q}$ whose reduction modulo q satisfies $\mathbf{H}_q(k)\mathbf{x}_0=\mathbf{u}$ over $\mathbb{F}_q$. Since $h-\mathbf{1}_q^{\mathsf{T}}\mathbf{x}_0\in q\mathbb{Z}$, adding this integer to the coordinate indexed by 0 gives a vector $\mathbf{z}_0$ with coefficient sum h and the same syndrome $\mathbf{u}$.

Every other vector with syndrome $\mathbf{u}$ and coefficient sum h differs from $\mathbf{z}_0$ by an element of $\mathcal{L}_{q,k}\cap\mathbf{1}_q^\perp$. The restricted isomorphism in Theorem 3.2 now gives the stated bijection.

By Lemma 2.2,

$$\|v\|^2=q\sum_a z_a^2-h^2=H+q\sum_a z_a(z_a-1).$$

For every integer z_a, the quantity $z_a(z_a-1)$ is a nonnegative even integer. The sum is zero exactly when every $z_a\in\{0,1\}$, and is otherwise at least two. This proves the stated claim. □

We next count binary vectors with a prescribed syndrome and fixed coordinate values. We use the following parameters to specialize Proposition 2.4. Let C_0 and k_0 be sufficiently large absolute integer constants. For $k \geq k_0$ and an odd prime $q \geq (2k)^{C_0}$, put

$$h := \left\lfloor \frac{21k}{20} \right\rfloor. \tag{6}$$

Set

$$H := h(q-h), \qquad L := 2kq. \tag{7}$$

By Lemma 3.3, every weight-h binary vector $\mathbf{x}$ satisfies $\|[\mathbf{x}]\|^2 = H$. By Theorem 3.2, $\lambda_1(\mathfrak{a}_{q,k})^2 \geq L$. Thus $H/L = 21/40 + o(1) < 1$ as $k \to \infty$, uniformly for these q. In particular, the ring elements represented by weight-h binary vectors are shorter than a shortest nonzero vector of $\mathfrak{a}_{q,k}$. The count below guarantees many such elements in every coset described in Lemma 3.3.

For a prescribed syndrome $\mathbf{u} = (u_0, \ldots, u_{k-1})$, let $P \subseteq \mathbb{F}_q$ be the *position set* indexing the coordinates fixed in advance, and let $\boldsymbol{\eta} \in \{0,1\}^P$ specify their values. Define

$$\mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta}) := \{\mathbf{x} \in \{0,1\}^{\mathbb{F}_q} : \mathbf{H}_q(k)\mathbf{x} = \mathbf{u},\ \mathbf{1}_q^{\mathsf{T}}\mathbf{x} = h,\ \mathbf{x}|_P = \boldsymbol{\eta}\}.$$

The first row of $\mathbf{H}_q(k)$ forces $u_0 = h \bmod q$ in $\mathbb{F}_q$.

Theorem 3.4. *For every $\mathbf{u} \in \mathbb{F}_q^k$ with $u_0 = h \bmod q$, every position set $P \subseteq \mathbb{F}_q$ with $|P| \leq \sqrt{k}$, and every $\boldsymbol{\eta} \in \{0,1\}^P$, put $t := \sum_{a \in P} \eta_a$. Then*

$$|\mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta})| = (1 + o(1))\frac{q^{h-t-k+1}}{(h-t)!}. \tag{8}$$

As $k \to \infty$, this estimate is uniform in q, $\mathbf{u}$, P, and $\boldsymbol{\eta}$. Its relative error is at most $q^{-\Omega(k)} + 3k^2/q < 1/4$.

Proof. Fix $\mathbf{u}$, P, and $\boldsymbol{\eta}$, and put $m_0 := h - t$. By the parameter choices above,

$$m_0 \geq 21k/20 - \sqrt{k} - 1 \geq k + 2, \qquad m_0 < q - 2.$$

The fixed coordinates contain exactly t ones. Subtract their contribution by defining

$$u'_j := u_j - \sum_{a \in P} \eta_a a^j \qquad (1 \leq j < k). \tag{9}$$

For each $\mathbf{x} \in \mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta})$, ordering the positions of its m_0 ones outside P gives exactly $m_0!$ tuples $(z_1, \ldots, z_{m_0}) \in (\mathbb{F}_q \setminus P)^{m_0}$ with pairwise distinct entries satisfying

$$\sum_{i=1}^{m_0} z_i^j = u'_j \qquad (1 \leq j < k). \tag{10}$$

Conversely, each such tuple and the prescribed values on P determine a unique vector in $\mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta})$, so the number of these tuples is $m_0!\,|\mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta})|$.

Let A count the tuples satisfying Equation (10) without requiring their entries to be pairwise distinct or to lie outside P. In Proposition 2.4, set $m = m_0$ and divide the error bound in part (i) by q^{m_0-k+1} and those in parts (ii) and (iii) by q^{m_0-k}. The three ratios are at most

$$(2k)^{m_0} q^{-(m_0-k-2)/2} \leq q^{-\Omega(k)}.$$

This bound follows from $q \geq (2k)^{C_0}$, $m_0 \geq 21k/20 - \sqrt{k} - 1$, and the choices of C_0 and k_0. Part (i) of Proposition 2.4 gives $|A/q^{m_0-k+1} - 1| \leq q^{-\Omega(k)}$. We bound the excluded tuples in two cases.
Case I (collision). For fixed $i < \ell$, part (ii) of Proposition 2.4 shows that the number of solutions with $z_i = z_\ell$ is at most $2q^{m_0-k}$.
Case II (entry in P). For fixed i and $a \in P$, part (iii) of Proposition 2.4 bounds the number of solutions with $z_i = a$ by $2q^{m_0-k}$.

The union bound over $\binom{m_0}{2}$ equalities $z_i = z_\ell$ and $m_0|P|$ equalities $z_i = a$ with $a \in P$ shows that the nonnegative difference satisfies

$$0 \leq A - m_0!\,|\mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta})| \leq 2\left(\binom{m_0}{2} + m_0|P|\right)q^{m_0-k} \leq 3k^2 q^{m_0-k}.$$

Together with the estimate for A, this gives

$$\left|\frac{m_0!\,|\mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta})|}{q^{m_0-k+1}} - 1\right| \leq q^{-\Omega(k)} + \frac{3k^2}{q} < \frac{1}{4}.$$

Substituting $m_0 = h - t$ proves the theorem. □

3.3 Norm separation

Theorem 3.4 supplies the binary representatives used by the checker in Section 4. Lemma 3.3 gives the norm gap for nonbinary representatives. We now choose a coset center whose multiples by short nonzero elements of $\mathcal{O}_K$ remain separated from the ideal.

We first record the norm gap between roots of unity and all other nonzero elements of $\mathcal{O}_K$, together with a sparse representative for short elements. The sparse representative will be used in the polynomial argument for separating short scalar multiples from the ideal.

Lemma 3.5. *The following statements hold.*

(i) Every $v \in \mathcal{O}_K$ has an integral coefficient representative $\mathbf{z}$ satisfying $\|\mathbf{z}\|_2^2 \leq 2\|v\|^2/q$, and hence having at most $2\|v\|^2/q$ nonzero coefficients.

(ii) For every nonzero $w \in \mathcal{O}_K$, either $w = \pm\zeta^j$ for some j, in which case $\|w\|^2 = q - 1$, or $\|w\|^2 \geq 2q - 4$.

Proof. For (i), let $\mathbf{z} = (z_a)_{a\in\mathbb{F}_q} \in \mathbb{Z}^{\mathbb{F}_q}$ satisfy $[\mathbf{z}] = v$. Choose $b \in \mathbb{F}_q$ minimizing $\sum_a (z_a - z_b)^2$. By Lemma 2.2,

$$\sum_a (z_a - z_b)^2 \leq \frac{1}{q}\sum_{r\in\mathbb{F}_q}\sum_{a\in\mathbb{F}_q}(z_a - z_r)^2 = \frac{2\|v\|^2}{q}.$$

Subtracting z_b from every coordinate gives another integral coefficient representative of v. Every nonzero integer coefficient has square at least one, so the resulting representative has at most $2\|v\|^2/q$ nonzero coefficients.

For (ii), if $\|w\|^2 \geq q^2/4$, then $\|w\|^2 \geq 2q - 4$. Otherwise, part (i) gives an integral coefficient representative with exactly $s < q/2$ nonzero coefficients. If $s \geq 2$, each of the $s(q-s)$ pairs formed by one nonzero coordinate and one zero coordinate contributes at least one to the pairwise norm formula in Lemma 2.2, so

$$\|w\|^2 \geq s(q-s) \geq 2(q-2).$$

If $s = 1$, then $w = t\zeta^j$ for a nonzero integer t. The cases $t = \pm 1$ are precisely the roots of unity, whereas $|t| \geq 2$ gives squared norm at least $4(q-1)$. □

Lemma 3.6. *If*

$$F(X) = \sum_{i=1}^{s} c_i X^{a_i} \in \mathbb{F}_q[X]$$

is nonzero, where every $c_i \neq 0$ and the exponents $a_i \in \{0, \dots, q-1\}$ are distinct, then its order of vanishing at $X = 1$ is at most $s - 1$.

Proof. Multiplicity at least s would make the first s coefficients of $F(1+Y)$ vanish. The resulting homogeneous system has determinant

$$\det\left(\binom{a_i}{j}\right)_{0\le j<s,\ 1\le i\le s} = \frac{\prod_{1\le i<\ell\le s}(a_\ell - a_i)}{\prod_{j=0}^{s-1} j!},$$

which is nonzero in $\mathbb{F}_q$: the exponents are distinct, and $s \le q$ makes every factorial in the denominator nonzero. Hence every c_i would vanish. This contradicts $F \neq 0$. □

We next choose c so that wc remains separated from $\mathfrak{a}_{q,k}$ for every nonzero $w \in \mathcal{O}_K$ with $\|w\|^2 \le 3q$. Define $c \in \mathcal{O}_K$ by

$$c := h + \pi^{\lfloor k/2 \rfloor}. \tag{11}$$

Let $\mathbf{c} \in \mathbb{Z}^{\mathbb{F}_q}$ be the coefficient vector of $h + (1-X)^{\lfloor k/2 \rfloor}$. Then $[\mathbf{c}] = c$, and evaluating the polynomial at $X = 1$ gives $\sum_{a\in\mathbb{F}_q} c_a = h$. Thus Lemma 3.3 and Theorem 3.4 identify and count binary representatives in $c + \mathfrak{a}_{q,k}$ with prescribed coordinate values.

The next theorem separates every short nonzero multiple of c from $\mathfrak{a}_{q,k}$.

Theorem 3.7. *For the parameter choices in Section 3.2, every nonzero $w \in \mathcal{O}_K$ with $\|w\|^2 \le 3q$ satisfies*

$$\operatorname{dist}(wc, \mathfrak{a}_{q,k})^2 > L/10.$$

Proof. We prove the theorem by contradiction. Suppose that there exist a nonzero $w \in \mathcal{O}_K$ with $\|w\|^2 \le 3q$ and an element $e \in wc + \mathfrak{a}_{q,k}$ such that $\|e\|^2 \le L/10 = kq/5$. By part (i) of Lemma 3.5, choose $f_e, f_w \in \mathbb{Z}[X]/(X^q - 1)$ with $\varphi(f_e) = e$ and $\varphi(f_w) = w$, whose coefficient vectors have squared ℓ_2-norm at most $2k/5$ and 6, respectively. In particular, f_e and f_w have at most $2k/5$ and 6 nonzero coefficients. Put $m := \lfloor k/2 \rfloor$.

Since $\varphi(h+(1-X)^m) = c$ and $e - wc \in \mathfrak{a}_{q,k}$, Equation (5) gives $f_e - f_w(h+(1-X)^m) \in (q, (X-1)^k)$. Reducing modulo q and using $X^q - 1 = (X-1)^q$ in $\mathbb{F}_q[X]$ gives

$$\bar{f}_e - h\bar{f}_w \equiv \bar{f}_w(1-X)^m \pmod{(X-1)^k}. \tag{12}$$

By Equation (12), $(X-1)^m$ divides $\bar{f}_e - h\bar{f}_w$. This polynomial has at most $2k/5 + 6 \le m$ nonzero coefficients by the choice of k_0, so Lemma 3.6 gives $\bar{f}_e - h\bar{f}_w = 0$.

Substituting this equality into Equation (12) gives $(X-1)^{k-m} \mid \bar{f}_w$. Since $\bar{f}_w$ has at most six nonzero coefficients and $k - m > 5$, Lemma 3.6 gives $\bar{f}_w = 0$. Since $w \neq 0$, the coefficient vector of f_w is nonzero. All its coefficients are divisible by q, so its squared ℓ_2-norm is at least q^2. This contradicts the upper bound of 6 for the squared norm. □

4 The X3C checker

Section 3 supplies weight-h binary representatives in fixed ideal cosets. A representative $\mathbf{x}$ encodes the X3C selection vector $\boldsymbol{\xi}$ through prescribed coordinates. This section constructs $U, V \in \mathcal{O}_K$ so that $\|U[\mathbf{x}] - V\|^2$ records the residual $\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M$, up to nonnegative error terms. In Section 4.3, we define admissible representatives for which one of the error terms vanishes. The final corollary records the conclusions used in Section 5.

4.1 Separated offsets

Consider the X3C instance with universe $[M]$, three-element subsets $C_1, \dots, C_n$, and incidence matrix $\mathbf{A} \in \{0,1\}^{M \times n}$. A selection vector $\boldsymbol{\xi} \in \{0,1\}^n$ specifies which subsets are chosen. For each $i \in [M]$, define

$$\mathcal{R}_i := \{j \in [n] : i \in C_j\} \qquad \text{and} \qquad s_i := \sum_{j \in \mathcal{R}_i} \xi_j = (\mathbf{A}\boldsymbol{\xi})_i.$$

Thus s_i is the number of selected sets containing i, and the sets selected by $\boldsymbol{\xi}$ form an exact cover exactly when $s_i = 1$ for every i. The quantity encoded by the checker is

$$2\|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2 = \underbrace{2\sum_{i=1}^{M}(s_i^2 - s_i)}_{\text{quadratic}} \underbrace{- 2\sum_{i=1}^{M} s_i}_{\text{linear}} + 2M. \tag{13}$$

Since $\xi_j^2 = \xi_j$, we have $s_i^2 - s_i = \sum_{j,\ell \in \mathcal{R}_i,\ j \neq \ell} \xi_j \xi_\ell$, so the first sum in Equation (13) is quadratic in $\boldsymbol{\xi}$.

Define the *incidence set*

$$\mathcal{I} := \{(i,j) \in [M] \times [n] : j \in \mathcal{R}_i\}.$$

Every set C_j contains three elements. Under the assumptions in Section 2.3,

$$W := |\mathcal{I}| = 3n, \qquad M \leq W. \tag{14}$$

Let $0 \leq \alpha_1 < \cdots < \alpha_n = \alpha_\star$ be the Sidon positions from Lemma 2.3. We encode $\boldsymbol{\xi}$ in a vector $\mathbf{x} \in \{0,1\}^{\mathbb{F}_q}$ by requiring $x_{\alpha_j} = \xi_j$ for $j \in [n]$, so the coordinate indexed by α_j stores the selection bit ξ_j for C_j. We call α_j a *witness position*. For $a, b \in \mathbb{F}_q$, the *offset* from a to b is the cyclic difference $b - a \in \mathbb{F}_q$. We separate the offsets associated with different rows of $\mathbf{A}$. Define

$$\tau := 2\alpha_\star + 1, \qquad \beta_i := \tau i \quad (i \in [M]), \qquad d_{ij} := \beta_i - \alpha_j \quad ((i,j) \in \mathcal{I}). \tag{15}$$

We regard α_j, β_i, and d_{ij} as integers and reduce them modulo q when using them as coordinate indices or offsets. Here β_i is the position assigned to row i of $\mathbf{A}$, and d_{ij} is the offset from the witness position α_j to β_i. We choose q large enough that

$$q > 4\beta_M = 4\tau M. \tag{16}$$

We use two types of offsets. The *witness offsets* are the cyclic differences between distinct witness positions:

$$\mathcal{D}_{\text{wit}} := \{\alpha_\ell - \alpha_j : j, \ell \in [n],\ j \neq \ell\}.$$

By the Sidon property and $q > 4\beta_M > 2\alpha_\star$, these differences remain distinct modulo q. The *bad offsets* are the cross-row differences and the sums of the offsets d_{ij}:

$$\mathcal{D}_{\text{bad}} := \{d_{ij} - d_{r\ell} : (i,j),(r,\ell) \in \mathcal{I},\ i \neq r\} \cup \{d_{ij} + d_{r\ell} : (i,j),(r,\ell) \in \mathcal{I}\}.$$

The expansion of $\|U[\mathbf{x}] - V\|^2$ in the proof of Theorem 4.3 uses witness offsets to represent the products $\xi_j \xi_\ell$ in the quadratic term of Equation (13). For $d \in \mathbb{F}_q$, the *correlation* of $\mathbf{x}$ at offset d counts all ordered pairs $(a, a+d)$ with $x_a = x_{a+d} = 1$, as defined in Section 4.2. The position pair (α_j, α_ℓ) contributes $\xi_j \xi_\ell = x_{\alpha_j} x_{\alpha_\ell}$ to the correlation of $\mathbf{x}$ at $d = \alpha_\ell - \alpha_j$. Correlations at bad offsets contribute to a nonnegative error term. To preserve the products at witness offsets while eliminating the correlations at bad offsets, the two offset sets are designed to be disjoint. The next lemma proves this separation and bounds the number of bad offsets.

Lemma 4.1. *Under the bound in Equation (16), the $2W$ residues d_{ij} and $-d_{ij}$, for $(i,j) \in \mathcal{I}$, are nonzero and pairwise distinct. Moreover,*

$$0 \notin \mathcal{D}_{\text{bad}}, \qquad \mathcal{D}_{\text{bad}} \cap \mathcal{D}_{\text{wit}} = \varnothing, \qquad |\mathcal{D}_{\text{bad}}| = O(W^2) = O(n^2).$$

Proof. For every $(i,j) \in \mathcal{I}$,

$$\alpha_\star + 1 = \beta_1 - \alpha_\star \le d_{ij} \le \beta_M < q/4,$$

so d_{ij} is nonzero modulo q. Fix distinct $(i,j),(r,\ell) \in \mathcal{I}$.

Case 1: $i = r$.

$$d_{ij} - d_{i\ell} = \alpha_\ell - \alpha_j \neq 0, \qquad |d_{ij} - d_{i\ell}| \le \alpha_\star.$$

Case 2: $i \neq r$.

$$d_{ij} - d_{r\ell} = \beta_i - \beta_r + \alpha_\ell - \alpha_j = \tau(i-r) + \alpha_\ell - \alpha_j.$$

Consequently,

$$\alpha_\star + 1 \le |d_{ij} - d_{r\ell}| \le \tau(M-1) + \alpha_\star = \beta_M - \alpha_\star - 1.$$

For all $(i,j),(r,\ell) \in \mathcal{I}$, one also has

$$2\alpha_\star + 2 = 2\beta_1 - 2\alpha_\star \le d_{ij} + d_{r\ell} \le 2\beta_M.$$

Since $q > 4\beta_M$, the bounds above give $0 < |d_{ij} - d_{r\ell}| < q/2$ whenever $(i,j) \neq (r,\ell)$, so $d_{ij} \not\equiv d_{r\ell} \pmod q$. For all $(i,j),(r,\ell) \in \mathcal{I}$, we also have $0 < d_{ij} + d_{r\ell} < q/2$, so $d_{ij} \not\equiv -d_{r\ell} \pmod q$. Thus the $2W$ residues d_{ij} and $-d_{ij}$ are nonzero and pairwise distinct.

Every witness offset is represented in $[-\alpha_\star, \alpha_\star]$. By contrast, the integer differences $d_{ij} - d_{r\ell}$ with $i \neq r$ and the integer sums $d_{ij} + d_{r\ell}$ lie in $(-q/2, q/2)$, with absolute value greater than $\alpha_\star$. Hence $0 \notin \mathcal{D}_{\text{bad}}$ and $\mathcal{D}_{\text{bad}} \cap \mathcal{D}_{\text{wit}} = \varnothing$. The two sets in the definition of $\mathcal{D}_{\text{bad}}$ contain at most $O(W^2) = O(n^2)$ offsets. □

4.2 Checker elements and norm formula

Fix $\mathbf{x} \in \{0,1\}^{\mathbb{F}_q}$ of weight h and define $\xi_j := x_{\alpha_j}$. Recall $s_i := \sum_{j \in \mathcal{R}_i} \xi_j$. By Lemma 2.2 and Equation (7), $\|[\mathbf{x}]\|^2 = H = h(q-h)$.

For $d \in \mathbb{F}_q$, define the *cyclic correlation* of $\mathbf{x}$ at offset d by

$$\mathcal{C}_{\mathbf{x}}(d) := \sum_{a \in \mathbb{F}_q} x_a x_{a+d}.$$

The indices lie in $\mathbb{F}_q$, and $a + d$ is computed modulo q. Thus $\mathcal{C}_{\mathbf{x}}(d)$ counts the ordered pairs $(a, a+d) \in \mathbb{F}_q^2$ for which $x_a = x_{a+d} = 1$.

Equation (13) separates the target into quadratic and linear parts. We first define $g \in \mathcal{O}_K$ so that $\|g[\mathbf{x}]\|^2$ represents the quadratic part through cyclic correlations. Next, V encodes the linear part, and the checker combines both parts in $\|U[\mathbf{x}] - V\|^2$.

Quadratic term. Using the separated offsets d_{ij}, define

$$g := \sum_{(i,j) \in \mathcal{I}} (\zeta^{d_{ij}} + \zeta^{-d_{ij}}) \in \mathcal{O}_K.$$

We have $g = \overline{g}$. By Lemma 4.1, the $2W$ exponents in the sum defining g are distinct modulo q.

Linear term. Define

$$V := 6 \sum_{j=1}^{n} \zeta^{\alpha_j} \in \mathcal{O}_K.$$

The coefficient $6 = 2 \cdot 3$ comes from the factor 2 in Equation (13) and the three elements in each selected set.

Checker elements. To obtain the remainder bound in Theorem 4.3, put $Q := 128W^4h$ and define

$$U := Q + g^2 \in \mathcal{O}_K.$$

Since $g = \overline{g}$, we have $\langle [\mathbf{x}], g^2[\mathbf{x}] \rangle = \|g[\mathbf{x}]\|^2$. Thus the expansion of $\|U[\mathbf{x}] - V\|^2$ contains $2Q\|g[\mathbf{x}]\|^2$, which represents the quadratic part. The following lemma gives norm bounds for multiplication by U and V.

Lemma 4.2. *For every $y, w \in K$, one has $\|Uy\| \geq Q\|y\|$ and $\|Vw\| \leq 2W\|w\|$.*

Proof. For every embedding $\sigma : K \hookrightarrow \mathbb{C}$, $\sigma(g)$ is real, so $\sigma(U) = Q + \sigma(g)^2 \geq Q > 0$. Hence $|\sigma(Uy)| \geq Q|\sigma(y)|$. The triangle inequality gives $|\sigma(g)| \leq 2W$ and $|\sigma(V)| \leq 2W$, so $|\sigma(Vw)| \leq 2W|\sigma(w)|$. Summing squares over all embeddings proves the two inequalities. □

Norm formula. We call $U[\mathbf{x}] - V \in \mathcal{O}_K$ the *checker value*. Put $\kappa := 2q/Q$ and $\mathcal{R}(\mathbf{x}) := \|g^2[\mathbf{x}] - V\|^2/Q^2$. The next theorem relates the squared norm of the checker value to $\|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2$ and bounds the remainder $\mathcal{R}(\mathbf{x})$.

Theorem 4.3. *For every $\mathbf{x} \in \{0,1\}^{\mathbb{F}_q}$ of weight h, define $\boldsymbol{\xi} \in \{0,1\}^n$ by $\xi_j := x_{\alpha_j}$ for $j \in [n]$. Then there exists $\mathcal{E}(\mathbf{x}) \geq 0$ such that*

$$\frac{\|U[\mathbf{x}] - V\|^2}{Q^2} = B_0 + 2\kappa\|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2 + \kappa\mathcal{E}(\mathbf{x}) + \mathcal{R}(\mathbf{x}), \tag{17}$$

where

$$B_0 := H + \kappa\left(2Wh - 2M + \frac{2hW - 4h^2W^2}{q}\right) \in \mathbb{Q} \tag{18}$$

is independent of $\mathbf{x}$. Moreover, $0 \leq \mathcal{R}(\mathbf{x}) < \kappa/8$.

Proof. **Quadratic term.** For $a \in \mathbb{F}_q$, define

$$y_a := \sum_{(i,j)\in\mathcal{I}} (x_{a-d_{ij}} + x_{a+d_{ij}}).$$

Then $g[\mathbf{x}] = \sum_a y_a \zeta^a$. To compute $\|g[\mathbf{x}]\|^2$ by Lemma 2.2, we determine $\sum_a y_a$ and $\sum_a y_a^2$. For each $(i,j) \in \mathcal{I}$, the vectors $(x_{a-d_{ij}})_{a\in\mathbb{F}_q}$ and $(x_{a+d_{ij}})_{a\in\mathbb{F}_q}$ are coordinate permutations of $\mathbf{x}$, so both have coordinate sum h. Hence $\sum_a y_a = 2Wh$. For $r, s \in \mathbb{F}_q$, cyclic reindexing gives

$$\sum_{a\in\mathbb{F}_q} x_{a-r}x_{a-s} = \mathcal{C}_{\mathbf{x}}(r-s), \qquad \mathcal{C}_{\mathbf{x}}(-r) = \mathcal{C}_{\mathbf{x}}(r).$$

Expanding $\sum_a y_a^2$ and using the preceding identities gives

$$\sum_{a\in\mathbb{F}_q} y_a^2 = 2Wh + 2 \sum_{\substack{(i,j),(t,\ell)\in\mathcal{I} \\ (i,j)\neq(t,\ell)}} \mathcal{C}_{\mathbf{x}}(d_{ij} - d_{t\ell}) + 2 \sum_{(i,j),(t,\ell)\in\mathcal{I}} \mathcal{C}_{\mathbf{x}}(d_{ij} + d_{t\ell}).$$

The squared terms $x^2_{a-d_{ij}}$ and $x^2_{a+d_{ij}}$ contribute $2Wh$ after summing over $a \in \mathbb{F}_q$ and $(i,j) \in \mathcal{I}$. Both double sums are over ordered pairs in $\mathcal{I} \times \mathcal{I}$. The products $x_{a-d_{ij}}x_{a-d_{t\ell}}$ and $x_{a+d_{ij}}x_{a+d_{t\ell}}$ each contribute $\mathcal{C}_{\mathbf{x}}(d_{ij} - d_{t\ell})$ after summing over a. The sum over a of $x_{a-d_{ij}}x_{a+d_{t\ell}} + x_{a+d_{ij}}x_{a-d_{t\ell}}$ is $2\mathcal{C}_{\mathbf{x}}(d_{ij} + d_{t\ell})$. For the terms with $i = t$ and $j \neq \ell$, we have $d_{ij} - d_{i\ell} = \alpha_\ell - \alpha_j$. The ordered pair (α_j, α_ℓ) contributes $\xi_j\xi_\ell$ to $\mathcal{C}_{\mathbf{x}}(\alpha_\ell - \alpha_j)$. The difference $\mathcal{C}_{\mathbf{x}}(\alpha_\ell - \alpha_j) - \xi_j\xi_\ell$ counts all other ordered pairs $(a, a + \alpha_\ell - \alpha_j)$ with $x_a = x_{a+\alpha_\ell-\alpha_j} = 1$, so it is nonnegative. Since $\sum_{\substack{j,\ell\in\mathcal{R}_i \\ j\neq\ell}} \xi_j\xi_\ell = s_i^2 - s_i$, the products $\xi_j\xi_\ell$ contribute $2\sum_i(s_i^2 - s_i)$.

We collect the remaining terms in

$$\mathcal{E}(\mathbf{x}) := 2\sum_{i=1}^{M} \sum_{\substack{j,\ell\in\mathcal{R}_i \\ j\neq\ell}} (\mathcal{C}_{\mathbf{x}}(\alpha_\ell - \alpha_j) - \xi_j\xi_\ell) + 2 \sum_{\substack{(i,j),(t,\ell)\in\mathcal{I} \\ i\neq t}} \mathcal{C}_{\mathbf{x}}(d_{ij} - d_{t\ell}) + 2 \sum_{(i,j),(t,\ell)\in\mathcal{I}} \mathcal{C}_{\mathbf{x}}(d_{ij} + d_{t\ell}). \quad (19)$$

The differences in the first term and the correlation counts in the other two terms are nonnegative. Hence $\mathcal{E}(\mathbf{x}) \geq 0$, and the expansion becomes

$$\sum_a y_a^2 = 2Wh + 2\sum_{i=1}^{M}(s_i^2 - s_i) + \mathcal{E}(\mathbf{x}).$$

Using Lemma 2.2 and $\sum_a y_a = 2Wh$, we obtain

$$\|g[\mathbf{x}]\|^2 = q\left(2Wh + 2\sum_i(s_i^2 - s_i) + \mathcal{E}(\mathbf{x})\right) - 4W^2h^2.$$

Linear term. In the expression $V = 6\sum_{j=1}^{n} \zeta^{\alpha_j}$, the coefficient sum is $6n = 2W$. Since each selected set contains three universe elements, $3\sum_j \xi_j = \sum_i s_i$. Hence Lemma 2.2 gives

$$\langle [\mathbf{x}], V\rangle = 6q\sum_j \xi_j - 2hW = 2q\sum_i s_i - 2hW.$$

Combining the terms. Since $g = \overline{g}$, one has $\langle [\mathbf{x}], g^2[\mathbf{x}]\rangle = \|g[\mathbf{x}]\|^2$. The definitions of U and V, the equality $\|[\mathbf{x}]\|^2 = H$, and the two preceding identities give

$$\begin{aligned}
\frac{\|U[\mathbf{x}] - V\|^2}{Q^2} &= H + \frac{2}{Q}\|g[\mathbf{x}]\|^2 - \frac{2}{Q}\langle [\mathbf{x}], V\rangle + \mathcal{R}(\mathbf{x}) \\
&= B_0 + 2\kappa\sum_i (s_i - 1)^2 + \kappa\mathcal{E}(\mathbf{x}) + \mathcal{R}(\mathbf{x}) \\
&= B_0 + 2\kappa\|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2 + \kappa\mathcal{E}(\mathbf{x}) + \mathcal{R}(\mathbf{x}).
\end{aligned}$$

The last two equalities use $\sum_i(s_i^2 - 2s_i) = \|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2 - M$. The remaining terms independent of $\mathbf{x}$ give B_0 from Equation (18).

Finally, $H \leq hq$ and the embedding bounds in the proof of Lemma 4.2 give

$$0 \leq \mathcal{R}(\mathbf{x}) \leq \frac{\left(4W^2\sqrt{hq} + 2W\sqrt{q}\right)^2}{Q^2} < \frac{\kappa}{8}.$$

The last inequality uses $W \geq 3$, $h \geq 1$, and the definition of Q. □

4.3 Representatives with prescribed correlations

Let $\mathbf{x} \in \{0,1\}^{\mathbb{F}_q}$ have weight h, and put $\xi_j := x_{\alpha_j}$ for every $j \in [n]$. We call $\mathbf{x}$ an *admissible representative* if

$$\mathcal{C}_{\mathbf{x}}(\alpha_\ell - \alpha_j) = \xi_j \xi_\ell \quad (j, \ell \in [n],\ j \neq \ell), \qquad \mathcal{C}_{\mathbf{x}}(d) = 0 \quad (d \in \mathcal{D}_{\mathrm{bad}}). \tag{20}$$

The first equality allows only the ordered pair (α_j, α_ℓ) to contribute at the witness offset $\alpha_\ell - \alpha_j$. This ordered pair contributes $\xi_j \xi_\ell$. The second equality excludes every ordered pair $(a, a+d)$ with $d \in \mathcal{D}_{\mathrm{bad}}$ and $x_a = x_{a+d} = 1$. Hence every term in Equation (19) vanishes and $\mathcal{E}(\mathbf{x}) = 0$ whenever $\mathbf{x}$ is an admissible representative.

The next theorem uses Theorem 3.4 and a union bound to prove that an admissible representative always exists.

Theorem 4.4. *With the parameters from Section 3.2, assume that $n+2 \le \lfloor \sqrt{k} \rfloor$ and Equation (16) holds. For every syndrome $\mathbf{u} \in \mathbb{F}_q^k$ with $u_0 = h \bmod q$, and every $\boldsymbol{\xi} \in \{0,1\}^n$, there is an admissible representative $\mathbf{x} \in \{0,1\}^{\mathbb{F}_q}$ such that*

$$\mathbf{H}_q(k)\mathbf{x} = \mathbf{u}, \qquad \sum_a x_a = h, \qquad x_{\alpha_j} = \xi_j \quad (j \in [n]).$$

Proof. Put $P := \{\alpha_1, \dots, \alpha_n\}$. Let $\mathcal{F}$ be the set of weight-h vectors $\mathbf{x} \in \{0,1\}^{\mathbb{F}_q}$ with syndrome $\mathbf{u}$ and $x_{\alpha_j} = \xi_j$ for every $j \in [n]$. Applying Theorem 3.4 with position set P and values ξ_j at α_j gives $|\mathcal{F}| > 0$.

By Lemma 2.3 and Equation (16), the ordered pair (α_j, α_ℓ) is the unique ordered pair in P at the witness offset $d = \alpha_\ell - \alpha_j$. The contribution $x_{\alpha_j} x_{\alpha_\ell} = \xi_j \xi_\ell$ is already fixed. The first equality in Equation (20) therefore requires $x_a x_{a+d} = 0$ only for $a \neq \alpha_j$. For $d \in \mathcal{D}_{\mathrm{bad}}$, the second equality requires $x_a x_{a+d} = 0$ for every $a \in \mathbb{F}_q$. By Lemma 4.1, no ordered pair in P has a bad offset. Thus every pair whose product is required to vanish has at least one coordinate outside P. All these offsets are nonzero.

For a set $J \subseteq \mathbb{F}_q \setminus P$ of one or two additional coordinates, put

$$\mathcal{F}_J := \{\mathbf{x} \in \mathcal{F} : x_a = 1 \text{ for every } a \in J\}.$$

Since $|P| + |J| \le n + 2 \le \sqrt{k}$, Theorem 3.4 applies to both $\mathcal{F}_J$ and $\mathcal{F}$. Let t be the weight of $\boldsymbol{\xi}$, so $t = \sum_{j=1}^n \xi_j$. By Theorem 3.4, each count lies between $3/4$ and $5/4$ times its main term. Dividing the upper bound for $|\mathcal{F}_J|$ by the lower bound for $|\mathcal{F}|$ gives

$$\frac{|\mathcal{F}_J|}{|\mathcal{F}|} \le \frac{5}{3} \frac{(h-t)!}{q^{|J|}(h-t-|J|)!} < 2(h/q)^{|J|}.$$

Consider one of the conditions $x_a x_{a+d} = 0$ above. If exactly one of $a, a+d$ lies outside P, no representative violates the condition when the coordinate in P is fixed to zero. When the coordinate in P is fixed to one, the violating representatives form $\mathcal{F}_J$ for $J = \{a, a+d\} \setminus P$, so $|\mathcal{F}_J|/|\mathcal{F}| \le 2h/q$. If both a and $a+d$ lie outside P, the violating representatives form $\mathcal{F}_{\{a,a+d\}}$ and occupy a fraction of at most $2(h/q)^2$.

The definition of $\mathcal{D}_{\mathrm{wit}}$ and Lemma 4.1 give $|\mathcal{D}_{\mathrm{wit}} \cup \mathcal{D}_{\mathrm{bad}}| = O(n^2)$. For each of these offsets, at most $2n$ ordered pairs $(a, a+d)$ have exactly one of $a, a+d$ outside P, and at most q have both $a, a+d$ outside P. A union bound shows that the proportion of representatives violating at least one of these conditions is at most

$$O(n^2)\left(2n \cdot \frac{2h}{q} + q \cdot 2\left(\frac{h}{q}\right)^2\right) = O\left(\frac{n^3 h + n^2 h^2}{q}\right).$$

Since $n \le \sqrt{k}$ and $h = O(k)$, the proportion of violating representatives is $O(k^3/q)$, uniformly in $\mathbf{u}$ and $\boldsymbol{\xi}$. By $q \ge (2k)^{C_0}$ and the choice of k_0, this proportion is less than one. Therefore some representative in $\mathcal{F}$ is admissible. □

We summarize the checker bounds used in Section 5.

Corollary 4.5. *Under the assumptions of Theorem 4.4, let $c \in \mathcal{O}_K$ have an integral coefficient representative $\mathbf{c} \in \mathbb{Z}^{\mathbb{F}_q}$ with $\sum_a c_a = h$. Then for every $\boldsymbol{\xi} \in \{0,1\}^n$, there exists an admissible representative $\mathbf{x} \in \{0,1\}^{\mathbb{F}_q}$ of weight h such that $[\mathbf{x}] \in c + \mathfrak{a}_{q,k}$, $x_{\alpha_j} = \xi_j$ for $j \in [n]$, and*

$$B_0 + 2\kappa\|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2 \le \frac{\|U[\mathbf{x}] - V\|^2}{Q^2} \le B_0 + 2\kappa\|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2 + \frac{\kappa}{8}. \tag{21}$$

Proof. Put $\mathbf{u}_c := \mathbf{H}_q(k)\mathbf{c} \in \mathbb{F}_q^k$ and choose $\mathbf{x}$ by Theorem 4.4 with $\mathbf{u} = \mathbf{u}_c$. The vectors $\mathbf{x}$ and $\mathbf{c}$ have syndrome $\mathbf{u}_c$ and coefficient sum h. Hence Lemma 3.3 gives $[\mathbf{x}] \in c + \mathfrak{a}_{q,k}$. Admissibility gives $\mathcal{E}(\mathbf{x}) = 0$, so Theorem 4.3 gives the norm bounds in Equation (21). □

5 The rank-two reduction

Let q be an odd prime with the parameters specified in Section 3.2. This section assembles the checker and the ideal coset into a rank-two module and chooses the scale Γ for its second coordinate. Let $U, V \in \mathcal{O}_K$ and Q be the checker elements and scale constructed in Section 4.2 for the given X3C instance, respectively. The reduction uses the ideal $\mathfrak{a} := \mathfrak{a}_{q,k} = \pi^k\mathcal{O}_K$ and the center $c \in \mathcal{O}_K$ from Equation (11). Assume that $n + 2 \le \lfloor\sqrt{k}\rfloor$ and Equation (16) holds. The center c fixes the ideal coset whose representatives encode X3C selections. The elements U and V test the X3C equations in the first coordinate. The integer Γ weights the second coordinate. Define the generators and their $\mathcal{O}_K$-span by

$$\mathbf{m}_1 := (U\pi^k, 0), \qquad \mathbf{m}_2 := (Uc - V, -Q\Gamma),$$
$$\mathcal{M} := \mathcal{O}_K\mathbf{m}_1 + \mathcal{O}_K\mathbf{m}_2 \subset \mathcal{O}_K^2.$$

Taking $y = 1$ in Lemma 4.2 gives $U \ne 0$, and the choice of Γ below is positive. The generator matrix has nonzero determinant $-Q\Gamma U\pi^k$, so $\mathcal{M}$ is a free module of rank two and $\mathbb{Z}$-rank $2(q-1)$. Hence every vector in $\mathcal{M}$ has the unique form

$$\mathbf{v}(y, w) := (Uy - Vw, -Q\Gamma w), \qquad y, w \in \mathcal{O}_K, \quad y - wc \in \mathfrak{a}. \tag{22}$$

The coefficients of $\mathbf{v}(y, w)$ in the module basis $\mathbf{m}_1, \mathbf{m}_2$ are $(y - wc)/\pi^k$ and w, respectively. When $w = 1$ and $y = [\mathbf{x}]$ for some $\mathbf{x} \in \{0,1\}^{\mathbb{F}_q}$ of weight h, the first coordinate is the checker value $U[\mathbf{x}] - V$. Corollary 4.5 guarantees the existence of such an admissible representative for the completeness proof with $w = 1$. Theorem 4.3 gives a lower bound for every weight-h binary coefficient vector. The second coordinate contributes $Q^2\Gamma^2\|w\|^2$ to the squared norm. For soundness, we consider separately $w = 0$, roots of unity, and elements that are not roots of unity.

5.1 Parameters

We choose Γ and S for the norm comparisons in the proof of Theorem 5.2. Recall $H = h(q-h)$, $L = 2kq$, $W = 3n$, and $Q = 128W^4h$, and let B_0 be as in Equation (18). Define

$$\Gamma := \lfloor\sqrt{9k/10}\rfloor, \quad F := \Gamma^2(q-1), \quad S := Q^2(B_0 + F + \kappa), \quad \nu := 2W/(Q\Gamma). \tag{23}$$

When w is a root of unity, the second coordinate contributes Q^2F to $\|\mathbf{v}(y,w)\|^2$. By Equation (18), both Q^2B_0 and $Q^2\kappa = 2qQ$ are integers, so $S \in \mathbb{Z}$.

The formula in Equation (22) also defines $\mathbf{v}(y,w)$ for arbitrary $y, w \in K$.

Lemma 5.1. *For the parameters in Equation (23), we have $0 < \nu < 1$. For every $y, w \in K$,*

$$\|\mathbf{v}(y,w)\|^2 \geq Q^2(1-\nu)(\|y\|^2 + \Gamma^2\|w\|^2). \tag{24}$$

Proof. Since $Q > 2W$ and $\Gamma \geq 1$, we have $0 < \nu < 1$. Lemma 4.2 gives $\|Uy\| \geq Q\|y\|$ and $\|Vw\| \leq 2W\|w\| = \nu Q\Gamma\|w\|$. Hence

$$2|\langle Uy, Vw\rangle| \leq 2\nu Q\Gamma\|Uy\|\|w\| \leq \nu(\|Uy\|^2 + Q^2\Gamma^2\|w\|^2).$$

Expanding $\|\mathbf{v}(y,w)\|^2$ and using this bound together with $\|Uy\| \geq Q\|y\|$ proves Equation (24). □

5.2 Completeness and soundness

Theorem 5.2. *With the parameters from Section 3.2 and Equation (23), assume $n+2 \leq \lfloor\sqrt{k}\rfloor$ and Equation (16) holds. Then the following implications hold:*

$$\begin{aligned} \exists \boldsymbol{\xi} \in \{0,1\}^n : \mathbf{A}\boldsymbol{\xi} = \mathbf{1}_M \quad &\Longrightarrow \quad \lambda_1(\mathcal{M})^2 < S, \\ \forall \boldsymbol{\xi} \in \{0,1\}^n : \mathbf{A}\boldsymbol{\xi} \neq \mathbf{1}_M \quad &\Longrightarrow \quad \lambda_1(\mathcal{M})^2 > S. \end{aligned}$$

Proof. **Completeness.** Let $\boldsymbol{\xi} \in \{0,1\}^n$ satisfy $\mathbf{A}\boldsymbol{\xi} = \mathbf{1}_M$. Corollary 4.5 gives an admissible representative $\mathbf{x} \in \{0,1\}^{\mathbb{F}_q}$ such that $[\mathbf{x}] \in c + \mathfrak{a}$ and $\|U[\mathbf{x}] - V\|^2 \leq Q^2(B_0 + \kappa/8)$. By Equation (22), the vector $\mathbf{v}([\mathbf{x}],1) \in \mathcal{M}$ has second coordinate $-Q\Gamma \neq 0$ and satisfies

$$\|\mathbf{v}([\mathbf{x}],1)\|^2 \leq Q^2\left(B_0 + F + \frac{\kappa}{8}\right) = S - \frac{7\kappa Q^2}{8} < S.$$

Soundness. We first record the parameter estimates used in the soundness proof. Since $h = \Theta(k)$ and $W \geq 3$, the definitions of Q, Γ, ν, and κ give

$$\Gamma^2 = \frac{9k}{10} + O(\sqrt{k}), \qquad \nu = O(k^{-3/2}), \qquad \kappa = O(q/k).$$

By Equation (18) and $M \leq W$, we have $|B_0 - H| = O(q/W^3 + h/W^2) = o(L)$. Since $h = 21k/20 + O(1)$ and $k = o(q)$,

$$H = \frac{21}{40}L + o(L), \qquad F = \frac{18}{40}L + o(L), \qquad \frac{S}{Q^2} = \frac{39}{40}L + o(L).$$

These estimates hold as $k \to \infty$, uniformly for $q \geq (2k)^{C_0}$, $1 \leq M \leq 3n$, and $n+2 \leq \lfloor\sqrt{k}\rfloor$. In particular, $S > 0$ by the choice of k_0.

Assume that the X3C instance is negative, and consider any nonzero $\mathbf{v}(y,w) \in \mathcal{M}$. We use the checker identity only when w is a root of unity and the coefficient representative of $w^{-1}y$ with coefficient sum h is binary. In the other cases, geometric norm bounds suffice to prove $\|\mathbf{v}(y,w)\|^2 > S$.

Case 1: $w = 0$. Then $0 \neq y \in \mathfrak{a}$. Lemma 4.2 and Theorem 3.2 give

$$\|\mathbf{v}(y,0)\|^2 = \|Uy\|^2 \geq Q^2\|y\|^2 \geq Q^2L > S,$$

since $S/Q^2 = (39/40)L + o(L) < L$.

Case 2: w is a root of unity. Since $y - wc \in \mathfrak{a}$ and $w\mathfrak{a} = \mathfrak{a}$, we have $u := w^{-1}y \in c + \mathfrak{a}$. Multiplication by w is an isometry, so the identities $\mathbf{v}(y, w) = w\mathbf{v}(u, 1)$ and $\|w\|^2 = q - 1$ give $\|\mathbf{v}(y, w)\|^2 = \|Uu - V\|^2 + Q^2F$. By Lemma 3.3, u has a unique integral coefficient representative $\mathbf{x} \in \mathbb{Z}^{\mathbb{F}_q}$ with coefficient sum h.

(a) If $\mathbf{x} \in \{0, 1\}^{\mathbb{F}_q}$, the coefficient-sum condition shows that $\mathbf{x}$ has weight h. Define $\xi_j := x_{\alpha_j}$ for $j \in [n]$. Since the instance is negative, $\|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2 \geq 1$. Thus Theorem 4.3 yields

$$\|\mathbf{v}(y, w)\|^2 \geq Q^2(B_0 + 2\kappa + F) = S + \kappa Q^2 > S.$$

(b) If $\mathbf{x} \notin \{0, 1\}^{\mathbb{F}_q}$, then Lemma 3.3 gives $\|u\|^2 \geq H + 2q$. Dropping the negative terms in Equation (18) gives $B_0 - H \leq (q + 1)/(32W^3)$. Since $\kappa + \nu(H + 2q + F) = O(q/\sqrt{k}) = o(q)$,

$$\begin{aligned}(1 - \nu)(H + 2q + F) - S/Q^2 &= 2q - (B_0 - H) - \kappa - \nu(H + 2q + F)\\ &\geq 2q - \frac{q + 1}{32W^3} - o(q) > 0.\end{aligned}$$

The last inequality uses $W \geq 3$ and the choice of k_0. Applying Equation (24) to u and 1 therefore gives

$$\|\mathbf{v}(y, w)\|^2 \geq Q^2(1 - \nu)(H + 2q + F) > S.$$

Case 3: w is not a root of unity and $0 < \|w\|^2 \leq 3q$. Since $y - wc \in \mathfrak{a}$, the coset-separation bound in Theorem 3.7 gives

$$\|y\|^2 \geq \operatorname{dist}(wc, \mathfrak{a})^2 > L/10.$$

Because w is not a root of unity, part (ii) of Lemma 3.5 also gives $\|w\|^2 \geq 2q-4$. Hence Equation (24) yields

$$\|\mathbf{v}(y, w)\|^2 > Q^2(1 - \nu)(L/10 + \Gamma^2(2q - 4)) = Q^2(L + o(L)) > S.$$

Case 4: w is not a root of unity and $\|w\|^2 > 3q$. Since $F = (18/40)L + o(L)$ and $S/Q^2 < L$, we have $3F > L > S/Q^2$. The second coordinate alone gives

$$\|\mathbf{v}(y, w)\|^2 \geq Q^2\Gamma^2\|w\|^2 > 3Q^2F > S.$$

These cases exhaust every nonzero module vector. Thus $\lambda_1(\mathcal{M})^2 > S$. □

5.3 Encoding length and running time

We choose q and bound the encoding length and running time of the reduction.

Over $\mathbb{Z}[\zeta_3]$, the module generated by $(2, 0)$ and $(0, 2)$ has squared minimum 8. If $M = 0$, the reduction outputs this module with squared threshold 8 as a fixed YES instance. If $M > 0$ and $\mathbf{A}$ has a zero row, the reduction outputs this module with squared threshold 1 as a fixed NO instance. For all remaining instances, $M, n \geq 1$ and every row of $\mathbf{A}$ is nonzero.

Choosing k and q. Recall that k_0 and C_0 are the fixed absolute integer constants from Section 3.2. Let ℓ be the bit length of the given X3C instance, and put

$$k := \max\{k_0, (n + 2)^2\}, \qquad Y := (2k)^{C_0}.$$

Since $n = O(\ell)$ for this explicit encoding, $Y = O((n + 2)^{2C_0})$. By Bertrand's postulate [Erd32], the interval $[Y, 2Y]$ contains an odd prime. The reduction scans the odd integers in this interval and applies deterministic primality testing [AKS04]. It finds a prime $Y \leq q \leq 2Y$ with $q = \ell^{O(1)}$ in deterministic polynomial time. There are at most $Y + 1 = \ell^{O(1)}$ candidates. The choice of k gives $n + 2 \leq \lfloor\sqrt{k}\rfloor$. Lemma 2.3 and Equations (14) and (15) give $\beta_M = O(n^3) = O(k^{3/2})$, so the same sufficiently large k_0 ensures $q > 4\beta_M$.

Generator coefficients. The sum defining g has $2W$ nonzero coefficients, and the sum defining V has total absolute coefficient sum $2W$. Expanding $U = Q + g^2$ gives an integral coefficient representative whose total absolute coefficient sum is at most $Q + 4W^2$. The expansions of $(1-\zeta)^k$ and $h + (1-\zeta)^{\lfloor k/2 \rfloor}$ give representatives of π^k and c with total absolute coefficient sums at most 2^k and $h + 2^{\lfloor k/2 \rfloor}$, respectively. Hence $U\pi^k$ and $Uc - V$ have integral coefficient representatives whose coefficients are bounded in absolute value by $(Q + 4W^2)(h + 2^k) + 2W$. Converting $\sum_{a=0}^{q-1} c_a \zeta^a$ to the power basis gives coefficients $c_a - c_{q-1}$ for $0 \le a < q-1$. The maximum coefficient bit length therefore increases by at most one and is $O(k + \log q + \log n)$. The second coordinate $-Q\Gamma$ has $O(\log k + \log n)$ bit length, by Equation (23) and the definition of Q.

Running time. The output is computed from $\mathbf{A}$ and the explicit center c. By Equation (23), S has $O(\log q + \log k + \log n)$ bits. The generators and S are computed by an integer square-root computation for Γ and power-basis arithmetic on $q-1$ integer coefficients. Since k and q are polynomially bounded in ℓ, the coefficient bounds above give a deterministic polynomial-time construction with polynomial encoding length. Admissible representatives serve only as witnesses in the completeness proof and need not be found by the reduction.

5.4 Proof of the main theorem

Proof of Theorem 1.1. Every output of the reduction is a full-rank free $\mathcal{O}_K$-submodule of $\mathcal{O}_K^2$ and hence has module rank two. As a lattice, each output has $\mathbb{Z}$-rank $2(q-1)$. For NP-hardness, Theorem 5.2 gives the YES–NO equivalence for instances with $M, n \ge 1$ and no zero row in $\mathbf{A}$. Together with the fixed YES and NO outputs in Section 5.3, this proves the correctness of the reduction on all X3C instances. The reduction is deterministic, runs in polynomial time, and outputs $(q, \mathbf{m}_1, \mathbf{m}_2, S)$. The two module generators lie in $\mathcal{O}_K^2$ and have nonzero determinant.

A certificate for a positive instance is a nonzero $\mathbf{y} \in \mathcal{O}_K^2$ in power-basis coordinates with $\|\mathbf{y}\|^2 \le S$. By Lemma 2.2, every coefficient has absolute value at most $\sqrt{S}$, so the certificate has polynomial bit length. The verifier first checks that each ring coordinate of the input generators is given by $q-1$ integer coefficients and verifies the remaining input conditions in polynomial time. It then solves $\mathbf{y} = z_1\mathbf{m}_1 + z_2\mathbf{m}_2$ over K and checks that $z_1, z_2 \in \mathcal{O}_K$ to verify module membership. Exact power-basis arithmetic performs this check and computes $\|\mathbf{y}\|^2$ in polynomial time. Thus the problem belongs to NP, completing the proof of Theorem 1.1. □

Proof of Corollary 1.2. Given an X3C instance, apply the reduction from the proof of Theorem 1.1 to obtain $(q, \mathbf{m}_1, \mathbf{m}_2, S)$, and query a search-SVP oracle on $(q, \mathbf{m}_1, \mathbf{m}_2)$. It returns a shortest nonzero vector $\mathbf{y} \in \mathcal{M}$ in power-basis coordinates. These coordinates have polynomial bit length by Lemma 2.2 and $\|\mathbf{y}\|^2 \le \|\mathbf{m}_1\|^2$. Compute $\|\mathbf{y}\|^2$ using Lemma 2.2. By the correctness of the reduction, the X3C instance is positive exactly when $\|\mathbf{y}\|^2 \le S$. Thus at most one oracle call decides X3C in polynomial time, which proves the claimed Turing hardness. □

AI disclosure. GPT-5.6 Sol Ultra discovered the proof through an iterative conversation conducted in a single session. The discussion began by asking whether the locally dense lattices that Bennett and Peikert construct from Reed–Solomon codes for approximate SVP [BP23] could be adapted to module lattices. The model first found a randomized reduction by combining this construction with Wan's point-count estimates [Wan26c], and then developed the deterministic reduction proved in this paper. In the second version, the authors removed the restriction $q \equiv 3 \pmod 4$ and simplified the checker construction to avoid lengthy Gauss-sum calculations. The authors independently verified every mathematical claim and proof, checked every use of the cited references, simplified

and refined the argument, and wrote the final manuscript. The authors take full responsibility for the paper.

References

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

[AKS04] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. PRIMES is in P. *Annals of Mathematics*, 160(2):781–793, 2004.

[BDK+18] Joppe W. Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS–Kyber: A CCA-secure module-lattice-based KEM. In *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 353–367. IEEE, 2018.

[BP23] Huck Bennett and Chris Peikert. Hardness of the (approximate) shortest vector problem: A simple proof via Reed–Solomon codes. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2023)*, volume 275 of *LIPIcs*, pages 37:1–37:20, 2023.

[CN98] Jin-Yi Cai and Ajay Nerurkar. Approximating the SVP to within a factor $(1 + 1/\dim^{\varepsilon})$ is NP-hard under randomized reductions. In *Proceedings of the 13th Annual IEEE Conference on Computational Complexity (CCC)*, pages 46–55. IEEE Computer Society, 1998.

[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. ACM, 2026.

[DKL+18] Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-Dilithium: A lattice-based digital signature scheme. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2018(1):238–268, 2018.

[DMMPMT23] Gabrielle De Micheli, Daniele Micciancio, Alice Pellet-Mary, and Nam Tran. Reductions from module lattices to free module lattices, and application to dequantizing Module-LLL. In *Advances in Cryptology—CRYPTO 2023, Part V*, pages 836–865. Springer, 2023.

[Erd32] Paul Erdős. Beweis eines Satzes von Tschebyschef. *Acta Scientiarum Mathematicarum (Szeged)*, 5:194–198, 1932.

[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.

[FPMS22] Joël Felderhoff, Alice Pellet-Mary, and Damien Stehlé. On module unique-SVP and NTRU. In *Advances in Cryptology—ASIACRYPT 2022*, pages 709–740. Springer, 2022.

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

[GSVV25] Nihar Gargava, Vlad Serban, Maryna Viazovska, and Ilaria Viglino. Module lattices and their shortest vectors. *CoRR*, abs/2510.12893, 2025.

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

[HR12] Ishay Haviv and Oded Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. *Theory of Computing*, 8(23):513–531, 2012.

[HS26a] Isaac M. Hair and Amit Sahai. Deterministic hardness of approximation for SVP in all finite ℓ_p norms. *CoRR*, abs/2604.01451, 2026. Accepted to FOCS 2026.

[HS26b] Isaac M. Hair and Amit Sahai. 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.

[Kho05] Subhash Khot. Hardness of approximating the shortest vector problem in lattices. *Journal of the ACM*, 52(5):789–808, 2005.

[LFP26] Jiaqi Liu, Yansong Feng, and Yanbin Pan. Exact CVP is NP-complete for principal cyclotomic ideals. *CoRR*, abs/2608.23828, 2026.

[LPMSW19] Changmin Lee, Alice Pellet-Mary, Damien Stehlé, and Alexandre Wallet. An LLL algorithm for module lattices. In *Advances in Cryptology—ASIACRYPT 2019*, pages 59–90. Springer, 2019.

[LPR13a] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. *Journal of the ACM*, 60(6):43:1–43:35, 2013.

[LPR13b] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for Ring-LWE cryptography. In *Advances in Cryptology—EUROCRYPT 2013*, pages 35–54. Springer, 2013.

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

[Mar26] Daniel E. Martin. NP-hardness of ideal lattice problems. *CoRR*, abs/2609.15813, 2026.

[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, 2002.

[Mic01] Daniele Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. *SIAM Journal on Computing*, 30(6):2008–2035, 2001.

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

[Mic12] Daniele Micciancio. Inapproximability of the shortest vector problem: Toward a deterministic reduction. *Theory of Computing*, 8(22):487–512, 2012.

[Nat24a] National Institute of Standards and Technology. Module-Lattice-Based Digital Signature Standard. Federal Information Processing Standards Publication NIST FIPS 204, U.S. Department of Commerce, 2024.

[Nat24b] National Institute of Standards and Technology. Module-Lattice-Based Key-Encapsulation Mechanism Standard. Federal Information Processing Standards Publication NIST FIPS 203, U.S. Department of Commerce, 2024.

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

[Wan26a] Daqing Wan. Euclidean SVP is deterministically NP-hard to approximate within any constant factor. *CoRR*, abs/2608.12664, 2026.

[Wan26b] Daqing Wan. Euclidean SVP is NP-hard for cyclic lattices. *CoRR*, abs/2609.16711, 2026.

[Wan26c] Daqing Wan. NP-hardness of SVP in Euclidean space. *CoRR*, abs/2603.27398, 2026.

[Was97] Lawrence C. Washington. *Introduction to Cyclotomic Fields*, volume 83 of *Graduate Texts in Mathematics*. Springer, second edition, 1997.