

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 range over primes congruent to 3 modulo 4. Let ζ_q be a primitive q th root of unity, and put $K = \mathbb{Q}(\zeta_q)$, with ring of integers $\mathcal{O}_K = \mathbb{Z}[\zeta_q]$. We prove that the decision version of the Shortest Vector Problem (SVP) in the ℓ_2 -norm is NP-complete on full-rank free submodules of $\mathcal{O}_K^2$ by a deterministic polynomial-time many-one reduction from Exact Cover by 3-Sets (X3C). The module rank is fixed at two. As a $\mathbb{Z}$ -lattice, the module has rank $2(q-1)$, which grows with q . 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, a checker based on a quadratic Gauss sum turns the X3C equations into a canonical squared norm. Third, the checker and a second module coordinate combine with a separation bound for ideal cosets to rule out every unintended vector created by the $\mathcal{O}_K$ -action. Each constructed instance consists of a prime $q \equiv 3 \pmod{4}$, 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 polynomial-time Turing reductions.

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, later randomized reductions gave hardness within $1 + n^{-\varepsilon}$ for every fixed $\varepsilon > 0$, within every constant below $\sqrt{2}$, and within every constant factor [CN98, Mic01, 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 $p > 2$ and, under a subexponential-time assumption, approximation hardness for every finite p [HS26b, 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 [Wan26b].

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 derandomization, these points must realize every prescribed assignment on a designated set of coordinates. Bennett and Peikert proposed a transparent 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 [Wan26b]. 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 [MMPMT23], 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]. To the best of our knowledge, none of these results establishes NP-hardness of exact Euclidean SVP for rank-two modules under deterministic polynomial-time reductions. Most closely related to the present question, 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].

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 $\mathcal{O}_K^2$.

1.1 Our result

For every prime $q \equiv 3 \pmod{4}$, let $\zeta := \zeta_q$ be a primitive q th root of unity. Write $K := \mathbb{Q}(\zeta)$ and $\mathcal{O}_K := \mathbb{Z}[\zeta]$. For $\mathbf{x} = (x_1, x_2) \in K^2$, define

$$\|\mathbf{x}\|^2 := \sum_{\sigma: K \hookrightarrow \mathbb{C}} (|\sigma(x_1)|^2 + |\sigma(x_2)|^2)$$

for the canonical product norm. 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 a prime $q \equiv 3 \pmod{4}$, 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 $1, \zeta, \dots, \zeta^{q-2}$. 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 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 primes congruent to 3 modulo 4, decision-SVP in the ℓ_2 -norm on 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 $q \equiv 3 \pmod{4}$ is 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. The tuple has polynomial encoding length. The generator construction and the denominator-clearing argument appear in Section 5.3. Section 5.4 chooses q and computes an integer basis whose Euclidean lattice is isometric to the canonically embedded output module. This basis is used by the NP verifier and the search reduction.

The output module is a full-rank free submodule of the rank-two module $\mathcal{O}_K^2$. As a $\mathbb{Z}$ -lattice, the output module has rank $2(q-1)$. The same construction also gives the search consequence stated below. This is a worst-case hardness result in which the cyclotomic ring varies with q . The theorem does not establish hardness for a single fixed cyclotomic ring, for the power-of-two cyclotomic family, on average, or through a cryptographic security reduction.

Corollary 1.2. *As q ranges over primes congruent to 3 modulo 4, search-SVP in the ℓ_2 -norm on full-rank free rank-two $\mathcal{O}_K$ -module lattices is NP-hard. The hardness holds under polynomial-time Turing reductions.*

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. Before clearing denominators, the reduction constructs an intermediate rank-two $\mathcal{O}_K$ -module $\mathcal{M} \subset K^2$ and a squared threshold S such that

$$\lambda_1(\mathcal{M})^2 \leq S \iff \text{the X3C instance is positive.}$$

A common integer scaling later produces the integral output module in $\mathcal{O}_K^2$ and an integer squared threshold without changing this equivalence. 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 $\mathbf{A} \in \{0, 1\}^{M \times n}$ admits a vector $\boldsymbol{\xi} \in \{0, 1\}^n$ satisfying

$$\mathbf{A}\boldsymbol{\xi} = \mathbf{1}_M, \quad (1)$$

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

The proof of the reduction has three layers. First, it extends $\boldsymbol{\xi}$ to a binary coefficient vector $\mathbf{x}$ and evaluates $\mathbf{x}$ at ζ to obtain an element of one fixed ideal coset in $\mathcal{O}_K$. Second, it transforms this element into an X3C checker expression whose squared norm measures $\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M$. Third, it places the checker expression in a rank-two module. Writing a module vector as an $\mathcal{O}_K$ -linear combination of the two generators introduces a coefficient $w \in \mathcal{O}_K$ for the second generator. Completeness uses the intended choice $w = 1$, whereas soundness must control every possible $w \in \mathcal{O}_K$.

The ideal coset. The first step embeds $\boldsymbol{\xi} \in \{0, 1\}^n$ into $\mathbf{x} = (x_a)_{a \in \mathbb{F}_q} \in \{0, 1\}^{\mathbb{F}_q}$. The reduction chooses distinct field elements $\alpha_1, \dots, \alpha_n \in \mathbb{F}_q$ and stores ξ_j in the coordinate x_{α_j} . Let $\zeta := \zeta_q$ be a primitive q th root of unity and 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

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

By [Theorem 3.2](#), which uses the Bennett–Peikert minimum bound [\[BP23\]](#), evaluation maps the Reed–Solomon lattice to $\mathfrak{a}$ and $\lambda_1(\mathfrak{a})^2 \geq 2kq$. Wan’s point-count estimate [\[Wan26b\]](#), in the uniform form of [Theorem 3.4](#), supplies many binary vectors whose ring elements lie in a fixed coset. The reduction chooses one coset $c + \mathfrak{a}$. 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

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

where h is the prescribed Hamming weight. Thus $x_{\alpha_j} = \xi_j$ stores $\boldsymbol{\xi}$ in designated coordinates of $\mathbf{x}$, while $v_{\mathbf{x}}$ places the resulting ring element in the fixed ideal coset. This encoding does not yet test [Equation \(1\)](#).

The X3C checker. The second step encodes [Equation \(1\)](#) in the canonical norm. The reduction constructs $U, V \in K$. By [Theorem 4.4](#), every weight- h binary vector $\mathbf{x}$, with $\xi_j = x_{\alpha_j}$, satisfies

$$\|Uv_{\mathbf{x}} - V\|^2 = B_0 + 2\|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2 + \mathcal{E}(\mathbf{x}). \quad (2)$$

Here B_0 depends only on the constructed instance, and $\mathcal{E}(\mathbf{x}) \geq 0$ collects the additional terms not determined by $\boldsymbol{\xi}$. Thus $Uv_{\mathbf{x}} - V$ is the *checker value* whose squared norm encodes the residual $\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M$. We call a representative $\mathbf{x}$ *clean* if it satisfies the correlation conditions in [Section 4.3](#), and these conditions imply $\mathcal{E}(\mathbf{x}) = 0$. The uniform coset count in [Theorem 3.4](#) and the union bound in the proof of [Theorem 4.5](#) show that every $\boldsymbol{\xi} \in \{0, 1\}^n$ has a clean coefficient vector $\mathbf{x}$ with $v_{\mathbf{x}} \in c + \mathfrak{a}$. Consequently, an exact cover gives squared checker norm B_0 , whereas every binary coefficient vector in a negative instance has squared checker norm at least $B_0 + 2$.

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

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

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

$$\mathbf{v}(z, w) := (U(\pi^k z + wc) - Vw, -\Gamma w), \quad z, w \in \mathcal{O}_K. \quad (3)$$

When $w = 1$ and $\pi^k z + c = v_{\mathbf{x}}$, the first coordinate is the checker value $Uv_{\mathbf{x}} - V$. For arbitrary w , the element $\pi^k z + wc$ lies in $wc + \mathfrak{a}$, and the first coordinate is $U(\pi^k z + wc) - Vw$. The second coordinate contributes $\Gamma^2 \|w\|^2$ to the squared norm. Set $S := B_0 + \Gamma^2(q - 1)$. If ξ is an exact cover, choose a clean coefficient vector $\mathbf{x}$ and write $v_{\mathbf{x}} = \pi^k z + c$. Taking $w = 1$ in Equation (3) then gives $\|\mathbf{v}(z, 1)\|^2 = S$.

Soundness under multiplication. It remains to control all vectors in this module when the X3C instance is negative. The main issue is possible cancellation between $U(\pi^k z + wc)$ and Vw . Set $y := \pi^k z + wc \in wc + \mathfrak{a}$. Then $\|y\| \geq \text{dist}(wc, \mathfrak{a})$. The stability estimate in Lemma 5.1 gives an explicit $\nu = o(1)$ and the following lower bound for the squared norm of the module vector.

$$\|Uy - Vw\|^2 + \Gamma^2 \|w\|^2 \geq (1 - \nu)(\|y\|^2 + \Gamma^2 \|w\|^2). \quad (4)$$

This inequality transfers the geometric bounds from the ideal cosets to the module. The argument considers four cases.

- (i) If $w = 0$, Part (iii) of Lemma 4.3 gives $\|Uy\| \geq \|y\|$. This preserves the lower bound for the minimum of $\mathfrak{a}$ and puts every nonzero vector above S .
- (ii) If w is a root of unity, write $y = wu$ with $u \in c + \mathfrak{a}$. Multiplication by w is an isometry. Let $\mathbf{x}$ be the integral coefficient representative of u from Lemma 3.6. If $\mathbf{x}$ is binary, the checker identity gives a squared first-coordinate norm at least $B_0 + 2$ in a negative instance. If $\mathbf{x}$ is not binary, Lemma 3.6 gives $\|u\|^2 \geq H + 2q$. Applying Equation (4) to u and 1 therefore puts the module vector above S .
- (iii) If w is not a root of unity and $0 < \|w\|^2 \leq 3q$, Theorem 3.8 separates wc from $\mathfrak{a}$, while Lemma 3.9 gives $\|w\|^2 \geq 2q - 4$. These bounds and Equation (4) put the module vector above S .
- (iv) If w is not a root of unity and $\|w\|^2 > 3q$, the coordinate $-\Gamma w$ alone puts the module vector above S .

Organization. Section 2 recalls Euclidean lattices, prime 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 clean completions. Section 5 gives the main reduction. It assembles the rank-two module, proves completeness and soundness, clears denominators, and establishes the claimed input representations and complexity bounds.

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. 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 $\mathbb{R}^r$.

Hasse derivatives. We will use Hasse derivatives in [Section 3](#) to recognize when a polynomial has a zero of prescribed order at $X = 1$. Let F be a field, let $f(X) = \sum_{i=0}^d c_i X^i \in F[X]$, and let $j \geq 0$ be an integer. The j th *Hasse derivative* of f is

$$D^{[j]}f(X) := \sum_{i=j}^d \binom{i}{j} c_i X^{i-j},$$

where each integer binomial coefficient is read as an element of F . We also use the binomial polynomials $\binom{X}{0} := 1$ and $\binom{X}{j} := X(X-1)\cdots(X-j+1)/j! \in F[X]$ for $j \geq 1$ whenever $j! \neq 0$ in F . For $b \in F$ and an integer $k \geq 1$, we say that b is a zero of order at least k when $(X-b)^k$ divides $f(X)$. The following standard characterization will be used [[LN97](#)]:

$$(X-b)^k \mid f(X) \iff D^{[j]}f(b) = 0 \quad (0 \leq j < k). \quad (5)$$

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:

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

For computational problems, the basis $\mathbf{B} \in \mathbb{Q}^{m \times n}$ is given explicitly. Exact decision-SVP also receives a nonnegative rational squared threshold Δ 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 Prime cyclotomic fields and module lattices

Cyclotomic field. Fix an odd prime q for the remainder of the construction, and let $\zeta := \zeta_q \in \mathbb{C}$ be a primitive q th root of unity. We suppress the subscript q after fixing the prime and put

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

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](#)]. Whenever $a \in \mathbb{F}_q$ occurs as an exponent of ζ , we identify it with its unique integer representative in $\{0, \dots, q-1\}$. A representation $v = \sum_{a \in \mathbb{F}_q} c_a \zeta^a$ with $c_a \in \mathbb{Q}$ is called a *q-coordinate expression* for v . Its q coefficients are indexed by $\mathbb{F}_q$ and need not be unique, because $\sum_{a \in \mathbb{F}_q} \zeta^a = 0$. By contrast, the expression of v in the power basis $1, \zeta, \dots, \zeta^{q-2}$ is unique.

Canonical embedding. For $v \in K$, write $\bar{v}$ for its complex conjugate. In particular, $\bar{\zeta} = \zeta^{-1}$. 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 field trace is

$$\mathrm{Tr}_{K/\mathbb{Q}}(v) := \sum_{a \in \mathbb{F}_q^\times} \sigma_a(v).$$

The canonical embedding sends v to $(\sigma_a(v))_{a \in \mathbb{F}_q^\times}$. For each integer $s \geq 1$, we equip K^s with the real inner product

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

and its induced product norm

$$\|\mathbf{x}\|^2 := \langle \mathbf{x}, \mathbf{x} \rangle = \sum_{i=1}^s \mathrm{Tr}_{K/\mathbb{Q}}(x_i \bar{x}_i).$$

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

$$\|v\|^2 = \sum_{a \in \mathbb{F}_q^\times} |\sigma_a(v)|^2 = \mathrm{Tr}_{K/\mathbb{Q}}(v \bar{v}).$$

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

Module lattices. For an integer $s \geq 1$, 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 $\mathbf{m}_1, \dots, \mathbf{m}_s \in K^s$ are linearly independent over K , then

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

is free of module rank s , has full rank, and has $\mathbb{Z}$ -rank $s(q-1)$. The square matrix with columns $\mathbf{m}_1, \dots, \mathbf{m}_s$ 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 $\mathcal{O}_K^s$.

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

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

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

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

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

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

Lemma 2.1. For all rational coefficient vectors $\mathbf{x}, \mathbf{y} \in \mathbb{Q}^q$,

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

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

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

Proof. For $a, b \in \mathbb{F}_q$, additive-character orthogonality 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. $\square$

2.3 X3C and Sidon sequences

An Exact Cover by 3-Sets (X3C) instance consists of a universe $[M]$ and three-element sets $C_1, \dots, C_n \subseteq [M]$. 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. X3C is NP-complete [GJ79].

A universe element is *isolated* if it belongs to none of the sets C_j . An instance with an isolated element is negative. After handling this case separately, every row of $\mathbf{A}$ is nonzero and $M \leq 3n$. Adjoining a disjoint three-element block together with its unique set preserves the answer, so we may also assume $n \geq 1$.

To index the n selection bits, 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 2.2. For every $n \geq 1$, one can deterministically compute integers

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

such that all nonzero ordered differences $\alpha_\ell - \alpha_j$, $j \neq \ell$, are distinct.

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]. $\square$

2.4 Power-sum estimates

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

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

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

Proposition 2.3 ([Wan26b]). *Let q be prime and suppose that $2 \leq k \leq m - 2 < q - 2$. For every $\mathbf{b} \in \mathbb{F}_q^{k-1}$, the following estimates hold.*

(i)

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

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

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

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

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

All three estimates are uniform in $\mathbf{b}$, and the last two are also uniform in the displayed choices of indices and field elements.

Proof. Wan [Wan26b, 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 uniformly 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 $i \neq \ell$.

For part (iii), fix $z_i = a$ and delete that coordinate. The remaining $m - 1$ variables satisfy $\sum_{r \neq i} z_r^j = b_j - a^j$ for $1 \leq j < k$. Applying Wan's first point-count estimate to these $m - 1$ variables proves part (iii). $\square$

3 Reed–Solomon lattices and ideal cosets

The first step of the reduction stores an X3C selection vector $\boldsymbol{\xi} \in \{0, 1\}^n$ in designated coordinates of a longer vector $\mathbf{x} \in \{0, 1\}^{\mathbb{F}_q}$. Evaluation at ζ maps $\mathbf{x}$ to $[\mathbf{x}] = \sum_{a \in \mathbb{F}_q} x_a \zeta^a \in \mathcal{O}_K$. The checker in Section 4 operates on fixed-weight binary vectors and reads $x_{\alpha_j} = \xi_j$ as the selection bits. Hence every $\boldsymbol{\xi}$ must admit a weight- h extension $\mathbf{x}$ whose evaluation lies in one fixed ideal coset. This coset is independent of $\boldsymbol{\xi}$. The ideal is designed to have a large minimum, so that a nonzero ideal element cannot produce an unintended short vector. This section obtains such an ideal from a lifted Reed–Solomon lattice [BP23], proves that the required binary extensions exist even after their designated coordinates are fixed, and establishes the additional norm bounds needed for nonbinary representatives and scalar multiples.

3.1 From lattice to ideal

We first transfer an ℓ_2 -minimum bound for the lifted 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 $1 \leq k < q$, define

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

When $\mathbf{H}_q(k)$ is applied to $\mathbf{z} \in \mathbb{Z}^q$, we first reduce the coordinates of $\mathbf{z}$ modulo q . The matrix $\mathbf{H}_q(k)$ defines the lifted Reed–Solomon lattice

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

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

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

Write $\mathbf{1}_q^\perp := \{\mathbf{z} \in \mathbb{R}^q : \mathbf{1}_q^\top \mathbf{z} = 0\}$. To express the lattice equations and evaluation at ζ in the same ring, associate $\mathbf{z} \in \mathbb{Z}^q$ with

$$f_{\mathbf{z}} := \sum_{a \in \mathbb{F}_q} z_a X^a \in \mathbb{Z}[X]/(X^q - 1).$$

Every element of this quotient has a unique degree- $< q$ representative, which we use to define its coefficients and its reduction $\bar{f} \in \mathbb{F}_q[X]$ modulo q . Let

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

be the ring homomorphism given by evaluation at ζ . Its kernel is generated by $\Phi_q(X) = 1 + X + \cdots + X^{q-1}$, and $\varphi(f_{\mathbf{z}}) = [\mathbf{z}]$.

Theorem 3.2. *For $1 \leq k \leq q/2$,*

$$\{[\mathbf{z}] : \mathbf{z} \in \mathcal{L}_{q,k}\} = \mathfrak{a}_{q,k} := \pi^k \mathcal{O}_K.$$

Moreover, 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,

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

Proof. Since $k < q$, the polynomials $1, X, \dots, X^{k-1}$ and $\binom{X}{0}, \binom{X}{1}, \dots, \binom{X}{k-1}$ 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

$$D^{[j]} \bar{f}_{\mathbf{z}}(1) = 0 \quad (0 \leq j < k),$$

because $D^{[j]} \bar{f}_{\mathbf{z}}(1) = \sum_a z_a \binom{a}{j}$ in $\mathbb{F}_q$. The multiplicity criterion in Equation (5) and the identity $X^q - 1 = (X - 1)^q$ in $\mathbb{F}_q[X]$ therefore give

$$\{f_{\mathbf{z}} : \mathbf{z} \in \mathcal{L}_{q,k}\} = (q, (X - 1)^k) \triangleleft \mathbb{Z}[X]/(X^q - 1).$$

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

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

$$\varphi^{-1}(\mathfrak{a}_{q,k}) = (q, (X - 1)^k) \quad \text{in } \mathbb{Z}[X]/(X^q - 1). \quad (6)$$

To prove surjectivity of the restricted map, fix $v \in \mathfrak{a}_{q,k}$. By Equation (6), choose $f \in (q, (X - 1)^k)$ with $\varphi(f) = v$. Then $f(1) \in q\mathbb{Z}$, and

$$g := f - \frac{f(1)}{q} \Phi_q$$

remains in $(q, (X - 1)^k)$, satisfies $\varphi(g) = v$, and has coefficient sum zero. Thus $g = f_{\mathbf{z}}$ for some $\mathbf{z} \in \mathcal{L}_{q,k} \cap \mathbf{1}_q^\perp$.

For injectivity, suppose that $\mathbf{z} \in \mathcal{L}_{q,k} \cap \mathbf{1}_q^\perp$ and $[\mathbf{z}] = 0$. The degree- $< q$ representative of $f_{\mathbf{z}}$ vanishes at ζ , so it equals $\ell \Phi_q$ for some $\ell \in \mathbb{Z}$. Its coefficient sum is zero, whereas $\Phi_q(1) = q$. Thus $\ell = 0$ and $\mathbf{z} = \mathbf{0}$.

Finally, every nonzero $v \in \mathfrak{a}_{q,k}$ has a unique nonzero $\mathbf{z} \in \mathcal{L}_{q,k} \cap \mathbf{1}_q^\perp$ with $v = [\mathbf{z}]$. By Lemmas 2.1 and 3.1,

$$\|v\|^2 = q \|\mathbf{z}\|_2^2 \geq 2kq,$$

which proves the minimum bound. $\square$

3.2 Binary representatives in ideal cosets

For $\mathbf{x} \in \mathbb{Z}^q$, its *syndrome* $\mathbf{H}_q(k)\mathbf{x} \in \mathbb{F}_q^k$ has the coordinate indexed by j equal to $\sum_{a \in \mathbb{F}_q} x_a a^j$ modulo q for $0 \leq j < k$. 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 $1 \leq k \leq q/2$. Let $\mathbf{u} = (u_0, \dots, u_{k-1}) \in \mathbb{F}_q^k$, and let $s \in \mathbb{Z}$ have residue u_0 modulo q . There exists $\mathbf{x}_{\mathbf{u},s} \in \mathbb{Z}^q$ such that*

$$\mathbf{H}_q(k)\mathbf{x}_{\mathbf{u},s} = \mathbf{u}, \quad \mathbf{1}_q^\top \mathbf{x}_{\mathbf{u},s} = s.$$

Moreover, evaluation at ζ induces a bijection

$$\{\mathbf{x} \in \mathbb{Z}^q : \mathbf{H}_q(k)\mathbf{x} = \mathbf{u}, \mathbf{1}_q^\top \mathbf{x} = s\} \longrightarrow [\mathbf{x}_{\mathbf{u},s}] + \mathfrak{a}_{q,k}, \quad \mathbf{x} \longmapsto [\mathbf{x}].$$

Thus every element of this ideal coset has a unique representative with syndrome $\mathbf{u}$ and coefficient sum s .

Proof. A Vandermonde $k \times k$ minor of $\mathbf{H}_q(k)$ has determinant $\prod_{1 \leq i < j \leq k} (a_j - a_i) \neq 0$ in $\mathbb{F}_q$. Hence the matrix has rank k . Choose $\mathbf{x}_0 \in \mathbb{Z}^q$ whose reduction modulo q satisfies $\mathbf{H}_q(k)\mathbf{x}_0 = \mathbf{u}$ over $\mathbb{F}_q$. Its coefficient sum is congruent to u_0 , and hence to s modulo q . The vector $\mathbf{1}_q$ belongs to $\mathcal{L}_{q,k}$: its zeroth power sum is $q = 0$ in $\mathbb{F}_q$, and $\sum_{a \in \mathbb{F}_q} a^j = 0$ for $1 \leq j < k < q - 1$. Moreover, $[\mathbf{1}_q] = 0$. Therefore

adding a suitable integer multiple of $\mathbf{1}_q$ to $\mathbf{x}_0$ gives a vector $\mathbf{x}_{\mathbf{u},s}$ with the required coefficient sum without changing its syndrome or its image at ζ .

Every other vector with syndrome $\mathbf{u}$ and coefficient sum s differs from $\mathbf{x}_{\mathbf{u},s}$ 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. $\square$

We use the following parameters to specialize [Proposition 2.3](#). For every sufficiently large prime q , put

$$k := \left\lfloor \frac{q^{1/1000}}{2} \right\rfloor, \quad h := \left\lfloor \frac{201k}{200} \right\rfloor, \quad T := \lfloor q^{1/2000} \rfloor. \quad (7)$$

These integers can be computed without real-number approximations. The value k is the largest nonnegative integer satisfying $(2k)^{1000} \leq q$, and T is the largest nonnegative integer satisfying $T^{2000} \leq q$. Fixed-degree integer-root algorithms compute k and T in time polynomial in $\log q$, after which integer division gives h . Here k is the number of rows of $\mathbf{H}_q(k)$. For each vector $\mathbf{x} \in \{0, 1\}^{\mathbb{F}_q}$ considered below, h is its weight, and T is the maximum number of its coordinates whose values may be fixed in advance. The relevant scale relations are $T = o(k)$, $h = (201/200)k + O(1)$, and $h = o(q)$. Set

$$H := h(q - h), \quad L := 2kq. \quad (8)$$

By [Lemma 2.1](#), every $\mathbf{x} \in \{0, 1\}^{\mathbb{F}_q}$ of weight h satisfies $\|\mathbf{x}\|^2 = H$. By [Theorem 3.2](#), every nonzero element of $\mathfrak{a}_{q,k}$ has canonical squared norm at least L .

For a prescribed syndrome $\mathbf{u} = (u_0, \dots, u_{k-1})$, 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

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

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

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

$$|\mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta})| = \Theta \left(\frac{q^{h-t-k+1}}{(h-t)!} \right). \quad (9)$$

Both the threshold for q and the implied constants are uniform in $\mathbf{u}$, P , and $\boldsymbol{\eta}$.

Proof. Fix $\mathbf{u}$, P , and $\boldsymbol{\eta}$, and put $m_0 := h - t$. By [Equation \(7\)](#), all sufficiently large q satisfy

$$k + 2 \leq m_0 < q - 2.$$

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

$$u'_j := u_j - \sum_{a \in P} \eta_a a^j \quad (1 \leq j < k). \quad (10)$$

Every $\mathbf{x} \in \mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta})$ has exactly m_0 coordinates outside P equal to one. Ordering their positions gives a tuple $(z_1, \dots, z_{m_0}) \in (\mathbb{F}_q \setminus P)^{m_0}$ such that $z_i \neq z_\ell$ for every $i \neq \ell$ and

$$\sum_{i=1}^{m_0} z_i^j = u'_j \quad (1 \leq j < k). \quad (11)$$

Conversely, every pairwise distinct tuple in $(\mathbb{F}_q \setminus P)^{m_0}$ satisfying [Equation \(11\)](#) determines a unique vector by setting $x_{z_i} = 1$, setting $\mathbf{x}|_P = \boldsymbol{\eta}$, and setting all remaining coordinates to zero. The equation for $j = 0$ follows from $u_0 = h$. Hence the number of these ordered tuples is $m_0! |\mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta})|$.

Let A count the tuples satisfying Equation (11) without requiring their entries to be pairwise distinct or to lie outside P . After the shift in Equation (10), part (i) of Proposition 2.3 gives $A = q^{m_0-k+1}(1 + o(1))$. We bound the excluded tuples in two cases.

Case I (collision). For fixed $i < \ell$, part (ii) of Proposition 2.3 shows that the number of solutions with $z_i = z_\ell$ is $O(q^{m_0-k})$.

Case II (entry in P). For fixed i and $a \in P$, substituting $z_i = a$ in Equation (11) gives $\sum_{r \neq i} z_r^j = u'_j - a^j$ for $1 \leq j < k$. Part (iii) of Proposition 2.3 shows that the number of solutions with $z_i = a$ is $O(q^{m_0-k})$.

The explicit error bounds in Proposition 2.3 are uniform in the prescribed power-sum vector $(u'_j)_{1 \leq j < k}$. For the parameters in Equation (7), each relative error is at most

$$(2k)^{m_0} q^{-(m_0-k-2)/2} = q^{-\Omega(k)} = o(1).$$

Thus the estimate for A , the bound for $z_i = z_\ell$, and the bound for $z_i = a$ are uniform over all shifted power-sum vectors $(u'_j)_{1 \leq j < k}$ and all choices of i , ℓ , and a .

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

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

After division by the main term q^{m_0-k+1} , this error is $O((m_0^2 + m_0|P|)/q) = o(1)$ by the scale relations following Equation (7). Together with the estimate for A , this gives

$$m_0! |\mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta})| = q^{m_0-k+1}(1 + o(1)).$$

Substituting $m_0 = h - t$ proves Equation (9). □

Corollary 3.5. *Let $\mathbf{u}, P, \boldsymbol{\eta}$, and t be as in Theorem 3.4. For $r \in \{1, 2\}$, let $Q \subseteq \mathbb{F}_q \setminus P$ have size r and satisfy $|P| + r \leq T$. Put $P' := P \cup Q$ and extend $\boldsymbol{\eta}$ to $\boldsymbol{\eta}' \in \{0, 1\}^{P'}$ by setting $\eta'_a = 1$ for every $a \in Q$. Then*

$$\frac{|\mathcal{F}(\mathbf{u}; P', \boldsymbol{\eta}')|}{|\mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta})|} = O((h/q)^r). \quad (12)$$

The implied constant is uniform in $\mathbf{u}$, P , $\boldsymbol{\eta}$, and Q .

Proof. For all sufficiently large q , $t + r \leq |P| + r \leq T < h$, so $h - t - r \geq 0$. The numerator in Equation (12) counts exactly the vectors in $\mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta})$ that equal one at every coordinate in Q . Applying Theorem 3.4 to the numerator and denominator gives

$$\frac{|\mathcal{F}(\mathbf{u}; P', \boldsymbol{\eta}')|}{|\mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta})|} = O\left(q^{-r} \frac{(h-t)!}{(h-t-r)!}\right) = O((h/q)^r).$$

□

3.3 Norm separation

The preceding count supplies the intended binary representatives. To use their ideal coset in the module construction, we must rule out shorter representatives with nonbinary coefficients and short points arising from nonzero scalar multiples.

Lemma 3.6. Let $\mathbf{u} \in \mathbb{F}_q^k$ satisfy $u_0 = h$, and let $\mathbf{x}' \in \mathbb{Z}^q$ have syndrome $\mathbf{u}$ and coefficient sum h . Put $c' := [\mathbf{x}']$. Every $v \in c' + \mathfrak{a}_{q,k}$ has a unique integral coefficient representative $\mathbf{z} \in \mathbb{Z}^q$ satisfying

$$[\mathbf{z}] = v, \quad \mathbf{H}_q(k)\mathbf{z} = \mathbf{u}, \quad \sum_a z_a = h.$$

If $\mathbf{z} \in \{0, 1\}^{\mathbb{F}_q}$, then it has weight h and $\|v\|^2 = H$. Otherwise, $\|v\|^2 \geq H + 2q$.

Proof. The existence and uniqueness of $\mathbf{z}$ follow directly from Lemma 3.3. By Lemma 2.1,

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

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. $\square$

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

$$m := \left\lfloor \frac{k}{2} \right\rfloor, \quad c := h + \pi^m. \quad (13)$$

The scalar summand h fixes the coefficient sum at h , since $h + (1 - X)^m$ evaluates to h at $X = 1$. The term π^m , with $m = \lfloor k/2 \rfloor$, balances the two orders of vanishing, m and $k - m$, that arise from the π^k -divisibility condition. Thus each of the two sparse-polynomial arguments below retains $k/2 + O(1)$ orders of vanishing. Let $\mathbf{c} \in \mathbb{Z}^q$ be the coefficient vector of $h + (1 - X)^m$. Then $[\mathbf{c}] = c$, and its syndrome

$$\mathbf{u}_c := \mathbf{H}_q(k)\mathbf{c} \in \mathbb{F}_q^k$$

satisfies $(u_c)_0 = h$. Thus Lemma 3.3 identifies the vectors of syndrome $\mathbf{u}_c$ and coefficient sum h with the coset $c + \mathfrak{a}_{q,k}$.

The proof of Theorem 3.8 combines a sparse integral coefficient representative with a bound on the multiplicity of a sparse polynomial at $X = 1$. The next lemma records both facts.

Lemma 3.7. *The following statements hold.*

(i) *If $v \in \mathcal{O}_K$ and $\|v\|^2 < q^2/4$, then v has an integral coefficient representative with at most $2\|v\|^2/q$ nonzero coefficients.*

(ii) *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. For (i), choose a coefficient vector $\mathbf{z} = (z_a)_{a \in \mathbb{F}_q} \in \mathbb{Z}^q$ with $[\mathbf{z}] = v$. If no value occurs among the coordinates z_a more than $q/2$ times, then at least $q^2/4$ unordered pairs have unequal coordinates, contradicting the last expression in Lemma 2.1. Subtract the most frequent coordinate value from every z_a . This does not change $[\mathbf{z}] = v$. Let s be the number of nonzero coefficients in the resulting representative. The preceding argument gives $s < q/2$. Each pair formed by a nonzero coordinate and a zero coordinate gives

$$\|v\|^2 \geq s(q - s) \geq \frac{sq}{2},$$

which proves (i).

For (ii), multiplicity at least s would make the first s Hasse derivatives vanish at one. The resulting homogeneous system has determinant

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

which is nonzero in $\mathbb{F}_q$: the exponents are distinct, and $s \leq q$ makes every factorial in the denominator nonzero. Hence every c_i would vanish, a contradiction. $\square$

With $b := 1/100$, the next theorem separates every short nonzero multiple of c from $\mathfrak{a}_{q,k}$.

Theorem 3.8. *For all sufficiently large primes q , every nonzero $w \in \mathcal{O}_K$ with $\|w\|^2 \leq 3q$ satisfies*

$$\text{dist}(wc, \mathfrak{a}_{q,k})^2 > bL.$$

Proof. We prove the theorem by contradiction. Suppose that there exist a nonzero $w \in \mathcal{O}_K$ with $\|w\|^2 \leq 3q$ and an element $e \in wc + \mathfrak{a}_{q,k}$ such that $\|e\|^2 \leq bL = 2bkq$. For sufficiently large q , both $2bkq$ and $3q$ are below $q^2/4$. By part (i) of Lemma 3.7, choose integral coefficient representatives $\mathbf{z}_e$ and $\mathbf{z}_w$ of e and w with at most $4bk$ and 6 nonzero coefficients, respectively. Put $f_e := f_{\mathbf{z}_e}$ and $f_w := f_{\mathbf{z}_w}$. Since $\varphi(h + (1 - X)^m) = c$ and $e - wc \in \mathfrak{a}_{q,k}$, Equation (6) 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

$$\bar{f}_e - h\bar{f}_w \equiv \bar{f}_w(1 - X)^m \pmod{(X - 1)^k}. \quad (14)$$

In particular, $\bar{f}_e - h\bar{f}_w$ is divisible by $(X - 1)^m$. For sufficiently large k , this polynomial has at most $4bk + 6 \leq m$ nonzero coefficients. Part (ii) of Lemma 3.7 therefore forces $\bar{f}_e - h\bar{f}_w = 0$ in $\mathbb{F}_q[X]$.

It follows from Equation (14) that $\bar{f}_w$ has a zero of order at least $k - m$ at $X = 1$. This polynomial is nonzero. Otherwise, $\mathbf{z}_w = q\mathbf{z}'$ for some $\mathbf{z}' \in \mathbb{Z}^q$. Since $[\mathbf{z}'] = w/q \neq 0$, the vector $\mathbf{z}'$ is not constant. At least $q - 1$ pairs of its coordinates therefore differ, and Lemma 2.1 gives

$$\|w\|^2 = q^2 \|[\mathbf{z}']\|^2 \geq q^2(q - 1) > 3q.$$

Since $\bar{f}_w$ has at most six nonzero coefficients, part (ii) of Lemma 3.7 bounds its order of vanishing at $X = 1$ by 5. This is impossible because $k - m \rightarrow \infty$. $\square$

We finally record the norm gap between roots of unity and all other nonzero elements of $\mathcal{O}_K$.

Lemma 3.9. *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. If $\|w\|^2 \geq q^2/4$, then $\|w\|^2 \geq 2q - 4$. Otherwise, part (i) of Lemma 3.7 gives an integral coefficient representative with exactly $s < q/2$ nonzero coefficients. If $s \geq 2$, every pair formed by one nonzero coordinate and one zero coordinate contributes to

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

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 $4(q - 1)$. $\square$

4 The X3C checker

Section 3 supplies weight- h binary representatives in fixed ideal cosets. For such a representative $\mathbf{x}$, put $v_{\mathbf{x}} := [\mathbf{x}]$ and let $\xi_j := x_{\alpha_j}$. This section constructs $U, V \in K$ so that the squared norm of the checker value $Uv_{\mathbf{x}} - V$ records the residual $\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M$, up to a baseline independent of $\mathbf{x}$ and a nonnegative error term. In Section 4.3, we define clean completions for which this error term vanishes. The final corollary records the conclusions used in Section 5.

4.1 Separated offsets

For a selection vector $\boldsymbol{\xi} \in \{0, 1\}^n$ and each $i \in [M]$, define

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

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

$$2\|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2 = 2 \underbrace{\sum_{i=1}^M (s_i^2 - s_i)}_{\text{quadratic}} - 2 \underbrace{\sum_{i=1}^M s_i}_{\text{linear}} + 2M. \quad (15)$$

Define the *incidence set*

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

Every set C_j contains three elements. We assume here that $M, n \geq 1$ and that every universe element lies in at least one set. These conditions are enforced by the preprocessing in Section 5.4. Under these assumptions,

$$W := |\mathcal{I}| = 3n, \quad M \leq W. \quad (16)$$

Let $0 \leq \alpha_1 < \dots < \alpha_n = \alpha_*$ be the Sidon positions from Lemma 2.2. The coordinate indexed by α_j will store the selection bit for C_j , so 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

$$\tau := 2\alpha_* + 1, \quad \beta_i := \tau i \quad (i \in [M]), \quad d_{ij} := \beta_i - \alpha_j \quad ((i, j) \in \mathcal{I}). \quad (17)$$

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

$$q > 4\beta_M = 4\tau M. \quad (18)$$

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

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

In the norm expansion, each witness offset carries a product $\xi_j \xi_\ell$ from the quadratic term in Equation (15). Every correlation at a bad offset contributes instead to $\mathcal{E}(\mathbf{x})$. To preserve the products at witness offsets while eliminating the correlations at bad offsets, the two offset sets must be disjoint. The next lemma proves this separation and bounds the number of bad offsets.

Lemma 4.1. *Under the bound in Equation (18), the $2W$ residues d_{ij} and $-d_{ij}$, for $(i, j) \in \mathcal{I}$, are nonzero and pairwise distinct. Moreover,*

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

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

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

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

Case 1: $i = r$.

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

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,

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

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

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

Together with $q > 4\beta_M$, these bounds place the same-row differences, cross-row differences, and sums strictly inside $(-q/2, q/2)$, so reduction modulo q creates no new equalities. The difference bounds separate residues of the same sign, and the sum bound separates residues of opposite signs. 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, every cross-row difference and every sum in $\mathcal{D}_{\text{bad}}$ has a representative in $(-q/2, q/2)$ with absolute value greater than $\alpha_\star$. Hence $0 \notin \mathcal{D}_{\text{bad}}$ and $\mathcal{D}_{\text{bad}} \cap \mathcal{D}_{\text{wit}} = \emptyset$. The two sets in the definition of $\mathcal{D}_{\text{bad}}$ contain at most $O(W^2) = O(n^2)$ offsets. $\square$

4.2 Checker elements and norm formula

Assume throughout this subsection that $q \equiv 3 \pmod{4}$. Fix $\mathbf{x} \in \{0, 1\}^{\mathbb{F}_q}$ of weight h and put

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

Define $\xi_j := x_{\alpha_j}$ and let $s_i := \sum_{j \in \mathcal{R}_i} \xi_j$ as in Section 4.1. By Lemma 2.1 and Equation (8), $\|v_{\mathbf{x}}\|^2 = H$.

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

The preceding decomposition separates the target into quadratic and linear parts. We first define p , whose shifted monomials encode the quadratic part through cyclic correlations. A Gauss sum is then used to define U , so that the quadratic contribution enters $\|Uv_{\mathbf{x}}\|^2$ and multiplication by U does not decrease the canonical norm. Next, D encodes the linear part, and V combines both parts in $\|Uv_{\mathbf{x}} - V\|^2$.

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

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

By Lemma 4.1, the displayed q -coordinate expression for $p = \bar{p}$ has $2W$ terms with pairwise distinct exponents. Same-row differences produce witness offsets, while the other differences and sums produce bad offsets. The resulting additional correlations are collected in $\mathcal{E}(\mathbf{x})$ below.

The element U . Let χ be the quadratic character of $\mathbb{F}_q$. Thus χ is 1 on nonzero squares, -1 on nonsquares, and 0 at zero. Define the quadratic Gauss sum

$$\vartheta := \sum_{a \in \mathbb{F}_q} \chi(a) \zeta^a \in \mathcal{O}_K.$$

Lemma 4.2. *The Gauss sum ϑ satisfies*

$$\bar{\vartheta} = -\vartheta, \quad \vartheta \bar{\vartheta} = q, \quad \vartheta^2 = -q.$$

Proof. Since $q \equiv 3 \pmod{4}$, one has $\chi(-1) = -1$. Conjugation followed by $a \mapsto -a$ therefore gives $\bar{\vartheta} = -\vartheta$. Since $\chi(0) = 0$, writing $a = tb$ in $\vartheta \bar{\vartheta}$ gives

$$\vartheta \bar{\vartheta} = \sum_{t \in \mathbb{F}_q^\times} \chi(t) \sum_{b \in \mathbb{F}_q^\times} \zeta^{(t-1)b} = q.$$

The inner sum is $q - 1$ for $t = 1$ and -1 otherwise. Combining $\vartheta \bar{\vartheta} = q$ with $\bar{\vartheta} = -\vartheta$ gives $\vartheta^2 = -q$. $\square$

Using Lemma 4.2, define

$$U := 1 + \frac{\vartheta p}{q} \in K. \tag{19}$$

Lemma 4.3. *The element U has the following properties.*

(i) $U\bar{U} = 1 + p^2/q$.

(ii) For every embedding $\sigma : K \hookrightarrow \mathbb{C}$, $|\sigma(U)|^2 = 1 + \sigma(p)^2/q \geq 1$. In particular, $U \neq 0$.

(iii) For every $y \in K$, one has $\|Uy\| \geq \|y\|$.

Proof. Since $p = \bar{p}$, Lemma 4.2 gives $U\bar{U} = 1 + p^2/q$, proving (i). The value $\sigma(p)$ is real for every embedding σ , so applying σ to (i) gives (ii). Summing the lower bound in (ii) over all embeddings proves (iii). $\square$

Linear term. Define

$$D := -3 \sum_{j=1}^n \zeta^{\alpha_j} \in \mathcal{O}_K.$$

Each selected set contributes to three row sums, which motivates the coefficient -3 . Part (ii) of Lemma 4.3 gives $U \neq 0$, so define

$$V := -\frac{D}{q\bar{U}} = -\frac{D}{q - \vartheta p} \in K. \tag{20}$$

Norm formula. We call $Uv_{\mathbf{x}} - V \in K$ the *checker value*. The correlations in $q^{-1}\|pv_{\mathbf{x}}\|^2$ that are not determined by the counts s_i are collected in the following term:

$$\begin{aligned} \mathcal{E}(\mathbf{x}) := & 2 \sum_{i=1}^M \sum_{\substack{j, \ell \in \mathcal{R}_i \\ j \neq \ell}} (\mathcal{C}_{\mathbf{x}}(\alpha_{\ell} - \alpha_j) - \xi_j \xi_{\ell}) \\ & + 2 \sum_{\substack{i, t \in [M] \\ i \neq t}} \sum_{\substack{j \in \mathcal{R}_i \\ \ell \in \mathcal{R}_t}} \mathcal{C}_{\mathbf{x}}(\beta_i - \beta_t - \alpha_j + \alpha_{\ell}) \\ & + 2 \sum_{i, t \in [M]} \sum_{\substack{j \in \mathcal{R}_i \\ \ell \in \mathcal{R}_t}} \mathcal{C}_{\mathbf{x}}(\beta_i + \beta_t - \alpha_j - \alpha_{\ell}). \end{aligned} \quad (21)$$

The three lines record, respectively, the additional same-row correlations, the cross-row difference correlations, and the sum-offset correlations.

Theorem 4.4. Assume that $q \equiv 3 \pmod{4}$. 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

$$\|Uv_{\mathbf{x}} - V\|^2 = B_0 + 2\|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2 + \mathcal{E}(\mathbf{x}), \quad (22)$$

where

$$B_0 := H + \|V\|^2 + 2Wh - 2M + \frac{2hW - 4h^2W^2}{q} \quad (23)$$

is a rational quantity determined by the constructed checker and independent of $\mathbf{x}$. Moreover, $\mathcal{E}(\mathbf{x}) \geq 0$ for every binary $\mathbf{x}$.

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 $pv_{\mathbf{x}} = \sum_a y_a \zeta^a$. Each of the $2W$ shifted copies of $\mathbf{x}$ has coefficient sum h , so $\sum_a y_a = 2Wh$. For $r, s \in \mathbb{F}_q$, cyclic reindexing gives

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

For $e = (i, j)$ and $f = (t, \ell)$ in $\mathcal{I}$, write $d_e := d_{ij}$ and $d_f := d_{t\ell}$. Expanding $\sum_a y_a^2$ and using the preceding identities gives

$$\sum_{a \in \mathbb{F}_q} y_a^2 = 2Wh + 2 \sum_{\substack{e, f \in \mathcal{I} \\ e \neq f}} \mathcal{C}_{\mathbf{x}}(d_e - d_f) + 2 \sum_{e, f \in \mathcal{I}} \mathcal{C}_{\mathbf{x}}(d_e + d_f).$$

The diagonal products contribute $2Wh$, and the remaining correlations fall into three classes.

Case 1: $d_e - d_f$ with $e \neq f$ and $i = t$. Here $e = (i, j)$ and $f = (i, \ell)$, so $j \neq \ell$ and $d_{ij} - d_{i\ell} = \alpha_{\ell} - \alpha_j$. Since $\sum_{\substack{j, \ell \in \mathcal{R}_i \\ j \neq \ell}} \xi_j \xi_{\ell} = s_i^2 - s_i$, this class contributes $2 \sum_i (s_i^2 - s_i)$ together with the first line of Equation (21).

Case 2: $d_e - d_f$ with $i \neq t$. Here $d_{ij} - d_{t\ell} = \beta_i - \beta_t - \alpha_j + \alpha_{\ell}$. These terms form the second line of Equation (21).

Case 3: $d_e + d_f$. For all $(i, j), (t, \ell) \in \mathcal{I}$, one has $d_{ij} + d_{t\ell} = \beta_i + \beta_t - \alpha_j - \alpha_\ell$. These terms form the third line of Equation (21).

Combining the diagonal term and the three cases yields

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

For $j \neq \ell$, the difference $\mathcal{C}_{\mathbf{x}}(\alpha_\ell - \alpha_j) - \xi_j \xi_\ell$ counts the ordered pairs $(a, a + \alpha_\ell - \alpha_j) \neq (\alpha_j, \alpha_\ell)$ for which $x_a = x_{a+\alpha_\ell-\alpha_j} = 1$. The remaining two lines of Equation (21) are sums of correlation counts. Thus $\mathcal{E}(\mathbf{x}) \geq 0$.

Using Lemma 2.1 and $\sum_a y_a = 2Wh$, we obtain

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

Linear term. The coefficient sum of D is $-W$. Since each selected set contains three universe elements, $3 \sum_j \xi_j = \sum_i s_i$. Hence Lemma 2.1 gives

$$\langle v_{\mathbf{x}}, D \rangle = -3q \sum_j \xi_j + hW = -q \sum_i s_i + hW.$$

Combining the terms. Part (i) of Lemma 4.3, the definition of V , the equality $\|v_{\mathbf{x}}\|^2 = H$, and the two preceding identities give

$$\begin{aligned} \|Uv_{\mathbf{x}} - V\|^2 &= H + \frac{1}{q} \|pv_{\mathbf{x}}\|^2 + \frac{2}{q} \langle v_{\mathbf{x}}, D \rangle + \|V\|^2 \\ &= H + \|V\|^2 + 2Wh + \frac{2hW - 4h^2W^2}{q} + 2 \sum_i (s_i^2 - 2s_i) + \mathcal{E}(\mathbf{x}) \\ &= B_0 + 2 \sum_i (s_i - 1)^2 + \mathcal{E}(\mathbf{x}). \end{aligned}$$

Since $s_i = (\mathbf{A}\boldsymbol{\xi})_i$, the last line is Equation (22). The formula for B_0 is independent of $\mathbf{x}$, and $B_0 \in \mathbb{Q}$ because $\|V\|^2 = \text{Tr}_{K/\mathbb{Q}}(V\bar{V})$. $\square$

4.3 Clean completions

Put $P := \{\alpha_1, \dots, \alpha_n\}$. For each selection vector $\boldsymbol{\xi} \in \{0, 1\}^n$, define $\boldsymbol{\eta}^{\boldsymbol{\xi}} \in \{0, 1\}^P$ by $\eta_{\alpha_j}^{\boldsymbol{\xi}} := \xi_j$. For $\mathbf{u} \in \mathbb{F}_q^k$ with $u_0 = h$, a *completion* of $\boldsymbol{\xi}$ with syndrome $\mathbf{u}$ is a vector $\mathbf{x} \in \mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta}^{\boldsymbol{\xi}})$. Thus $\mathbf{x}$ is binary, has weight h and syndrome $\mathbf{u}$, and satisfies $x_{\alpha_j} = \xi_j$ for every $j \in [n]$.

A completion $\mathbf{x}$ is *clean* if

$$\mathcal{C}_{\mathbf{x}}(\alpha_\ell - \alpha_j) = \xi_j \xi_\ell \quad (j, \ell \in [n], j \neq \ell), \quad \mathcal{C}_{\mathbf{x}}(d) = 0 \quad (d \in \mathcal{D}_{\text{bad}}). \quad (24)$$

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 at a bad offset. Hence every term in Equation (21) vanishes and $\mathcal{E}(\mathbf{x}) = 0$.

The next theorem combines Theorem 3.4 with the fixed-coordinate bound in Corollary 3.5 to prove that a clean completion always exists.

Theorem 4.5. Suppose that q is a sufficiently large prime, $n+2 \leq T$, and Equation (18) holds. For every syndrome $\mathbf{u} \in \mathbb{F}_q^k$ with $u_0 = h$, and every $\boldsymbol{\xi} \in \{0,1\}^n$, there is a clean completion $\mathbf{x} \in \{0,1\}^{\mathbb{F}_q}$ such that

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

Proof. Write $\mathcal{F} := \mathcal{F}(\mathbf{u}; P, \boldsymbol{\eta}^\xi)$. By Theorem 3.4, one has $|\mathcal{F}| > 0$. For a set $Q \subseteq \mathbb{F}_q \setminus P$ of one or two additional coordinates, put

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

Applying Corollary 3.5 gives

$$\frac{|\mathcal{F}_Q|}{|\mathcal{F}|} = O((h/q)^{|Q|}).$$

Case 1: $d = \alpha_\ell - \alpha_j \in \mathcal{D}_{\text{wit}}$. By Lemma 2.2 and Equation (18), (α_j, α_ℓ) is the unique ordered pair in P at offset d . This ordered pair contributes $\xi_j \xi_\ell$ to $\mathcal{C}_{\mathbf{x}}(d)$. For every $a \neq \alpha_j$, the first equality in Equation (24) requires $x_a x_{a+d} = 0$, and at least one of $a, a+d$ lies outside P .

Case 2: $d \in \mathcal{D}_{\text{bad}}$. The second equality in Equation (24) requires $x_a x_{a+d} = 0$ for every $a \in \mathbb{F}_q$. By Lemma 4.1, $d \notin \mathcal{D}_{\text{wit}}$ and $d \neq 0$. Thus no ordered pair in P has offset d , so at least one of $a, a+d$ lies outside P .

Every required zero product from the two cases has at least one endpoint outside P . If exactly one endpoint lies outside P , no completion violates the condition when the endpoint in P is fixed to zero. When the endpoint in P is fixed to one, the violating completions lie in $\mathcal{F}_Q$ with $|Q| = 1$ and occupy an $O(h/q)$ fraction of $\mathcal{F}$. If both endpoints lie outside P , the corresponding fraction is $O((h/q)^2)$.

The definition of $\mathcal{D}_{\text{wit}}$ and Lemma 4.1 give $|\mathcal{D}_{\text{wit}} \cup \mathcal{D}_{\text{bad}}| = O(n^2)$. For each of these offsets, at most $2n$ ordered pairs have exactly one endpoint outside P , and at most q have both endpoints outside P . A union bound shows that the number of completions violating at least one required zero-product condition is at most

$$O\left(\frac{n^3 h}{q} + \frac{n^2 h^2}{q}\right) |\mathcal{F}|.$$

Since $n \leq T = O(q^{1/2000})$ and $h = O(q^{1/1000})$, the relative factor is $o(1)$. Therefore some completion in $\mathcal{F}$ is clean. $\square$

We summarize the checker bounds used in Section 5.

Corollary 4.6. For every sufficiently large prime $q \equiv 3 \pmod{4}$ satisfying $n+2 \leq T$ and the bound in Equation (18), the following statements hold.

(i) Every $\mathbf{x} \in \{0,1\}^{\mathbb{F}_q}$ of weight h , with $\xi_j := x_{\alpha_j}$, satisfies

$$\|Uv_{\mathbf{x}} - V\|^2 \geq B_0 + 2\|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2.$$

(ii) For the center c and syndrome $\mathbf{u}_c$ defined in Section 3, every $\boldsymbol{\xi} \in \{0,1\}^n$ has a clean completion $\mathbf{x} \in \mathcal{F}(\mathbf{u}_c; P, \boldsymbol{\eta}^\xi)$ such that

$$v_{\mathbf{x}} \in c + \mathbf{a}_{q,k}, \quad \|Uv_{\mathbf{x}} - V\|^2 = B_0 + 2\|\mathbf{A}\boldsymbol{\xi} - \mathbf{1}_M\|_2^2. \quad (25)$$

Proof. Part (i) follows directly from Theorem 4.4. For part (ii), choose $\mathbf{x}$ by Theorem 4.5 with $\mathbf{u} = \mathbf{u}_c$. Cleanliness gives $\mathcal{E}(\mathbf{x}) = 0$, so Equation (22) gives the norm equality in Equation (25). The vectors $\mathbf{x}$ and $\mathbf{c}$ have syndrome $\mathbf{u}_c$ and coefficient sum h . Hence Lemma 3.3 gives $v_{\mathbf{x}} = [\mathbf{x}] \in c + \mathbf{a}_{q,k}$. $\square$

5 The rank-two reduction

Let $q \equiv 3 \pmod{4}$ be prime. This section assembles the checker and the ideal coset into a rank-two module and chooses the scale Γ for its second coordinate. Let $U, V \in K$ be the checker elements constructed in [Section 4.2](#) for the given X3C instance. 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 \(13\)](#). Until [Section 5.4](#), assume that $M, n \geq 1$, that every universe element occurs in at least one set, that $n + 2 \leq T$, and that [Equation \(18\)](#) holds. That subsection enforces these conditions when choosing q and handles isolated elements separately. 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. Before clearing denominators, define the generators and their span by

$$\mathbf{m}_1 := (U\pi^k, 0), \quad \mathbf{m}_2 := (Uc - V, -\Gamma),$$

$$\mathcal{M} := \mathcal{O}_K \mathbf{m}_1 + \mathcal{O}_K \mathbf{m}_2 \subset K^2.$$

[Lemma 4.3](#) gives $U \neq 0$, and the choice of Γ below is positive. Hence every vector in $\mathcal{M}$ has the unique form

$$\mathbf{v}(z, w) := (U(\pi^k z + wc) - Vw, -\Gamma w), \quad z, w \in \mathcal{O}_K. \quad (26)$$

When $w = 1$ and $\pi^k z + c = v_{\mathbf{x}}$ for some $\mathbf{x} \in \{0, 1\}^{\mathbb{F}_q}$, the first coordinate is the checker value $Uv_{\mathbf{x}} - V$. Part (ii) of [Corollary 4.6](#) guarantees the existence of such a clean completion for the completeness branch $w = 1$. The reduction does not compute this completion. Part (i) gives a lower bound for every binary coefficient vector. The second coordinate contributes $\Gamma^2 \|w\|^2$ to the squared norm. For soundness, we consider separately $w = 0$, roots of unity, and elements that are not roots of unity.

5.1 Gap parameters

Recall from [Equation \(8\)](#) that every $\mathbf{x} \in \{0, 1\}^{\mathbb{F}_q}$ of weight h satisfies $\|v_{\mathbf{x}}\|^2 = H = h(q - h)$, and that L is the lower bound on the squared norm of a nonzero element of $\mathfrak{a}$. By [Equation \(16\)](#), $W = 3n$. The constant $b = 1/100$ is used in [Theorem 3.8](#), and B_0 is defined in [Equation \(23\)](#). Let Γ be the nearest positive integer to $\sqrt{99L/(200(q-1))}$, with ties rounded upward. Then Γ can be computed exactly in time polynomial in $\log q$.

$$\Gamma = \sqrt{\frac{99L}{200(q-1)}} + O(1) = \Theta(\sqrt{k}). \quad (27)$$

Squaring the estimate in [Equation \(27\)](#) gives $\Gamma^2 = 99k/100 + O(\sqrt{k})$. Since $L = 2kq$, the error contributes $O(q\sqrt{k}) = o(L)$ after multiplication by $q - 1$. Define

$$F := \Gamma^2(q - 1) = \frac{198}{400}L + o(L) = \Theta(L), \quad S := B_0 + F, \quad \nu := \frac{W}{q\Gamma}. \quad (28)$$

A root of unity w contributes $\Gamma^2 \|w\|^2 = F$ through the second coordinate, so S is the completeness threshold. The parameter ν bounds the cancellation caused by Vw .

Lemma 5.1. *Under the standing assumptions of this section, for every sufficiently large prime $q \equiv 3 \pmod{4}$, one has $B_0 > 0$, $0 < \nu < 1$, and*

$$(1 - \nu)(H + 2q + F) > S, \quad S < L,$$

$$(1 - \nu)(bL + \Gamma^2(2q - 4)) > S, \quad 3F > S.$$

Moreover, every $y, w \in K$ satisfy the stability estimate

$$\|Uy - Vw\|^2 + \Gamma^2\|w\|^2 \geq (1 - \nu)(\|y\|^2 + \Gamma^2\|w\|^2). \quad (29)$$

Proof. The definitions of h and L , together with $k = o(q)$, give

$$h = \frac{201}{200}k + O(1), \quad H = \frac{201}{400}L + o(L).$$

The estimate for Γ^2 preceding Equation (28) also gives $\Gamma^2(2q - 4) = (99/100)L + o(L)$. For every embedding $\sigma : K \hookrightarrow \mathbb{C}$, Part (ii) of Lemma 4.3 gives

$$|\sigma(q\bar{U})| = q|\sigma(U)| \geq q.$$

The triangle inequality gives $|\sigma(D)| \leq W$. Since $V = -D/(q\bar{U})$, it follows that

$$\max_{\sigma} |\sigma(V)| \leq \frac{W}{q}, \quad \|V\|^2 \leq \frac{(q-1)W^2}{q^2} < \frac{W^2}{q}.$$

Since $W = 3n = O(T)$ and $M \leq W$, the checker formula in Equation (23) and the scale relations following Equation (7) give

$$|B_0 - H| = o(q), \quad \nu L = O(W\sqrt{k}) = o(q).$$

In particular, $B_0 > 0$ and $0 < \nu < 1$ for large q . Since $q = o(L)$, the same estimates give $B_0 = (201/400)L + o(L)$ and $S = (399/400)L + o(L)$. Therefore,

$$\begin{aligned} (1 - \nu)(H + 2q + F) - S &= 2q + o(q) > 0, & L - S &= \frac{1}{400}L + o(L) > 0, \\ (1 - \nu)(bL + \Gamma^2(2q - 4)) - S &= \frac{1}{400}L + o(L) > 0, & 3F - S &= \frac{39}{80}L + o(L) > 0. \end{aligned}$$

For the first line, expanding $S = B_0 + F$ cancels the F terms. The remaining error is $o(q)$ because $B_0 - H = o(q)$ and $\nu(H + 2q + F) = O(\nu L) = o(q)$. This proves every assertion of the lemma except the stability estimate.

For the stability estimate, the bound $\max_{\sigma} |\sigma(V)| \leq W/q$ gives $\|Vw\| \leq (W/q)\|w\| = \nu\Gamma\|w\|$. Hence

$$2|\langle Uy, Vw \rangle| \leq \nu\|Uy\|^2 + \nu^{-1}\|Vw\|^2 \leq \nu\|Uy\|^2 + \nu\Gamma^2\|w\|^2.$$

Using this inequality together with $\|Uy\| \geq \|y\|$ from Part (iii) of Lemma 4.3, we obtain

$$\begin{aligned} \|Uy - Vw\|^2 + \Gamma^2\|w\|^2 &\geq (1 - \nu)\|Uy\|^2 + \|Vw\|^2 + (1 - \nu)\Gamma^2\|w\|^2 \\ &\geq (1 - \nu)(\|y\|^2 + \Gamma^2\|w\|^2), \end{aligned}$$

which is Equation (29). □

5.2 Completeness and soundness

Theorem 5.2. *For every sufficiently large prime $q \equiv 3 \pmod{4}$, the following implications hold for every X3C instance such that each universe element occurs in at least one set, $n + 2 \leq T$, and Equation (18) holds:*

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

Proof. Completeness. Suppose $\mathbf{A}\boldsymbol{\xi} = \mathbf{1}_M$. Part (ii) of [Corollary 4.6](#) gives a clean $\mathbf{x} \in \{0, 1\}^{\mathbb{F}_q}$ such that $v_{\mathbf{x}} \in c + \mathfrak{a}$ and $\|Uv_{\mathbf{x}} - V\|^2 = B_0$. Hence $v_{\mathbf{x}} = \pi^k z + c$ for some $z \in \mathcal{O}_K$. Taking $w = 1$ in [Equation \(26\)](#) gives

$$\|\mathbf{v}(z, 1)\|^2 = B_0 + \Gamma^2 \|1\|^2 = B_0 + \Gamma^2(q - 1) = S.$$

Soundness. Assume that the X3C instance is negative, and consider any nonzero $\mathbf{v}(z, w) \in \mathcal{M}$.

Case 1: $w = 0$. Then $z \neq 0$. [Lemma 4.3](#) and [Theorem 3.2](#) give

$$\|\mathbf{v}(z, 0)\|^2 = \|U\pi^k z\|^2 \geq \|\pi^k z\|^2 \geq L > S,$$

where the final inequality follows from [Lemma 5.1](#).

Case 2: w is a root of unity. Write $w = \pm \zeta^a$ for some $a \in \mathbb{F}_q$. Since $w\mathfrak{a} = \mathfrak{a}$, the element $u := w^{-1}\pi^k z + c$ belongs to $c + \mathfrak{a}$, and $\pi^k z + wc = wu$. The first coordinate is $w(Uu - V)$. Since multiplication by w is an isometry and $\|w\|^2 = q - 1$, both subcases below satisfy $\|\mathbf{v}(z, w)\|^2 = \|Uu - V\|^2 + F$. By [Lemma 3.6](#), u has a unique integral coefficient representative $\mathbf{x} \in \mathbb{Z}^q$ 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 [Corollary 4.6](#) yields

$$\|\mathbf{v}(z, w)\|^2 = \|Uu - V\|^2 + \Gamma^2(q - 1) \geq B_0 + 2 + F = S + 2 > S.$$

- (b) If $\mathbf{x} \notin \{0, 1\}^{\mathbb{F}_q}$, then [Lemma 3.6](#) gives $\|u\|^2 \geq H + 2q$. Apply [Equation \(29\)](#) to the arguments u and 1, and then use the inequality $(1 - \nu)(H + 2q + F) > S$ from [Lemma 5.1](#):

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

Case 3: w is not a root of unity and $0 < \|w\|^2 \leq 3q$. Put $y := \pi^k z + wc$. Since $\pi^k z \in \mathfrak{a}$, the coset-separation bound in [Theorem 3.8](#) gives

$$\|y\|^2 \geq \text{dist}(wc, \mathfrak{a})^2 > bL.$$

Because w is not a root of unity, [Lemma 3.9](#) also gives $\|w\|^2 \geq 2q - 4$. The stability estimate and the inequality $(1 - \nu)(bL + \Gamma^2(2q - 4)) > S$ from [Lemma 5.1](#) imply

$$\|\mathbf{v}(z, w)\|^2 > (1 - \nu)(bL + \Gamma^2(2q - 4)) > S.$$

Case 4: w is not a root of unity and $\|w\|^2 > 3q$. The second coordinate alone gives

$$\|\mathbf{v}(z, w)\|^2 \geq \Gamma^2 \|w\|^2 > 3q\Gamma^2 > 3F > S,$$

using the last inequality in [Lemma 5.1](#). These cases exhaust every nonzero module vector. $\square$

5.3 Integral module and encoding length

The generators of $\mathcal{M}$ may have denominators. We multiply both generators by a common positive integer. This scales every squared norm and the threshold by the same square and preserves the gap.

A common denominator. Put $r := q - \vartheta p = q\bar{U} \in \mathcal{O}_K \setminus \{0\}$, the denominator of V in Equation (20). Let

$$N_r := \prod_{a \in \mathbb{F}_q^\times} \sigma_a(r) \in \mathbb{Z} \setminus \{0\}, \quad \rho := q|N_r|.$$

Every σ_a permutes the factors defining N_r , so N_r is rational. It is also an algebraic integer and therefore lies in $\mathbb{Z}$. Since σ_1 is the identity embedding, $N_r/r = \prod_{a \in \mathbb{F}_q^\times \setminus \{1\}} \sigma_a(r) \in \mathcal{O}_K$. Using Equations (19) and (20),

$$\rho U = |N_r|(q + \vartheta p) \in \mathcal{O}_K, \quad \rho V = -qD \frac{|N_r|}{r} \in \mathcal{O}_K. \quad (30)$$

Define

$$\widehat{\mathbf{m}}_i := \rho \mathbf{m}_i \in \mathcal{O}_K^2 \quad (i \in \{1, 2\}), \quad \widehat{\mathcal{M}} := \mathcal{O}_K \widehat{\mathbf{m}}_1 + \mathcal{O}_K \widehat{\mathbf{m}}_2, \quad \widehat{S} := \rho^2 S. \quad (31)$$

Its generator matrix

$$\widehat{\mathbf{B}} = \begin{pmatrix} \rho U \pi^k & \rho(Uc - V) \\ 0 & -\rho \Gamma \end{pmatrix} \in \mathcal{O}_K^{2 \times 2},$$

has determinant $-\rho^2 \Gamma U \pi^k \neq 0$. By the module-lattice facts in Section 2, $\widehat{\mathcal{M}}$ is therefore a full-rank free submodule of $\mathcal{O}_K^2$. Since $\widehat{\mathcal{M}} = \rho \mathcal{M}$ and $\lambda_1(\widehat{\mathcal{M}}) = \rho \lambda_1(\mathcal{M})$, Theorem 5.2 yields

$$\begin{aligned} \text{positive X3C} &\implies \lambda_1(\widehat{\mathcal{M}})^2 \leq \widehat{S}, \\ \text{negative X3C} &\implies \lambda_1(\widehat{\mathcal{M}})^2 > \widehat{S}. \end{aligned} \quad (32)$$

Integral threshold. Since $\rho V \in \mathcal{O}_K$,

$$\rho^2 \|V\|^2 = \|\rho V\|^2 = \text{Tr}_{K/\mathbb{Q}}((\rho V) \overline{\rho V})$$

is an integer. Every denominator q in Equation (23) is cleared because $\rho^2/q = qN_r^2 \in \mathbb{Z}$. Hence $\widehat{S} \in \mathbb{Z}$. Since $\rho > 0$ and $S > 0$, we have $\widehat{S} = \rho^2 S > 0$.

Encoding length. For the module $\widehat{\mathcal{M}}$ constructed above, we bound the power-basis coefficients of its two generators and the bit length of $\widehat{S}$ polynomially in q and the X3C input length. By Equation (31), the generator bounds amount to controlling $\rho U \pi^k$, $\rho(Uc - V)$, and $\rho \Gamma$. In view of Equation (30), the only new quantities requiring estimates are N_r and the power-basis coefficients of N_r/r . Write

$$r = \sum_{i=0}^{q-2} r_i \zeta^i, \quad B_r := \max \left\{ 1, \max_i |r_i| \right\}.$$

Since $q \equiv 3 \pmod{4}$, one has $\chi(-1) = -1$. Using $\zeta^{q-1} = -(1 + \zeta + \dots + \zeta^{q-2})$, we obtain

$$\vartheta = \sum_{a=1}^{q-2} \chi(a) \zeta^a - \zeta^{q-1} = 1 + \sum_{a=1}^{q-2} (1 + \chi(a)) \zeta^a.$$

Thus the power-basis coefficients of ϑ lie in $\{0, 1, 2\}$ and are efficiently computable from $a^{(q-1)/2} \bmod q$. The power-basis coefficient bounds for p , D , and ϑ are therefore 1, 3, and 2, respectively. In q -coordinates, p has $2W$ nonzero coefficients and every coefficient of ϑ has absolute value at most one. Each coefficient of ϑp is therefore a sum of at most $2W$ coefficients of ϑ . Hence the coefficients of ϑp are bounded by $2W$ in q -coordinates and by $4W$ in the power basis. Consequently,

$$B_r \leq q + 4W.$$

Let $\mathbf{M}_r \in \mathbb{Z}^{(q-1) \times (q-1)}$ be the matrix of multiplication by r in the power basis $1, \zeta, \dots, \zeta^{q-2}$. For $0 \leq j \leq q-2$, the column indexed by j is the coefficient vector of $r\zeta^j$. First use $\zeta^q = 1$ to reduce its exponents to $0, \dots, q-1$. The resulting powers are distinct, and the only power outside the power basis is ζ^{q-1} . Replacing this term with

$$\zeta^{q-1} = -(1 + \zeta + \dots + \zeta^{q-2})$$

shows that every entry of $\mathbf{M}_r$ is a sum of at most two terms of the form $\pm r_i$. Hence every entry has absolute value at most $2B_r$, and every row has ℓ_2 -norm at most $2\sqrt{q-1} B_r$.

Over $\mathbb{C}$, the eigenvalues of multiplication by r are the conjugates $\sigma_a(r)$ for $a \in \mathbb{F}_q^\times$. Consequently, the determinant identity and Hadamard's determinant bound give

$$\det(\mathbf{M}_r) = \prod_{a \in \mathbb{F}_q^\times} \sigma_a(r) = N_r, \quad |N_r| \leq \left(2\sqrt{q-1} B_r\right)^{q-1}. \quad (33)$$

We also need to control the coefficients of N_r/r . Its multiplication matrix is $N_r \mathbf{M}_r^{-1}$. For $1 \leq i, j \leq q-1$, let $\mathbf{M}_r^{(j,i)}$ be the matrix obtained from $\mathbf{M}_r$ by deleting row j and column i . The cofactor formula for the inverse gives

$$(N_r \mathbf{M}_r^{-1})_{ij} = (-1)^{i+j} \det(\mathbf{M}_r^{(j,i)}).$$

Every row of this $(q-2) \times (q-2)$ matrix has ℓ_2 -norm at most $2\sqrt{q-2} B_r$. A second application of Hadamard's bound therefore shows that

$$\left| (N_r \mathbf{M}_r^{-1})_{ij} \right| \leq \left(2\sqrt{q-2} B_r\right)^{q-2}. \quad (34)$$

Applying $N_r \mathbf{M}_r^{-1}$ to the coefficient vector of 1 shows that its first column is the power-basis coefficient vector of N_r/r . Thus [Equations \(33\) and \(34\)](#) show that N_r and every coefficient of N_r/r have bit length $O(q(\log B_r + \log q))$. The matrix $\mathbf{M}_r$ is computable from the coefficients of r . Standard polynomial-time exact integer linear algebra computes its determinant and the first column of its adjugate, and hence computes N_r and the coefficients of N_r/r within these bounds.

The bound $B_r \leq q + 4W$ and [Equation \(33\)](#) show that $\rho = q|N_r|$ has polynomial bit length. The coefficients of π^k and $c = h + \pi^{\lfloor k/2 \rfloor}$ have $O(k)$ bit length, and Γ has $O(\log k)$ bit length by [Equation \(27\)](#). Addition and multiplication in the power basis use $q-1$ integer coefficients and increase their bit lengths by only a polynomial amount. Since the next subsection chooses q polynomially bounded in the X3C input length, the two generators in [Equation \(31\)](#) are computable in deterministic polynomial time and have polynomial encoding length.

The bit-length bound for ρ and the coefficient bounds for ρV , together with [Equations \(23\) and \(28\)](#), show that $\hat{S}$ has bit length polynomial in q and the X3C input length. It is also computable in deterministic polynomial time. Compute $\|\rho V\|^2$ from the power-basis coefficients of ρV using [Lemma 2.1](#), and then compute $\hat{S} = \rho^2 S$ from [Equations \(23\) and \(28\)](#) using exact integer arithmetic.

5.4 Prime choice and Euclidean input

For an X3C instance with an isolated universe element, the reduction outputs the fixed module lattice over $\mathbb{Z}[\zeta_3]$ generated by $(2, 0)$ and $(0, 2)$, with squared threshold 1. Its squared minimum is $8 > 1$, so this is a NO instance. For every remaining instance, adjoin a disjoint three-element block together with its unique set, as in [Section 2](#). This answer-preserving preprocessing ensures $M, n \geq 1$. We retain the notation M, n , and $\mathbf{A}$ for the resulting instance, in which every universe element occurs in at least one set.

Let ℓ be the bit length of this preprocessed X3C instance, and recall from Equation (7) that T is the number of coordinates that may be prescribed. By the prime number theorem for arithmetic progressions [IK04, Equation (17.2)], choose an absolute integer X_0 exceeding the threshold implicit in Theorem 5.2 and large enough that every interval $[Y, 2Y]$ with $Y \geq X_0$ contains a prime congruent to 3 modulo 4. Since $n = O(\ell)$ for this explicit encoding, Lemma 2.2 and Equations (16) and (17) give $4\tau M = O(n^3) = O(\ell^3)$. Since $T = \lfloor q^{1/2000} \rfloor$, an absolute constant C_0 can therefore be chosen so that every $q \geq (\ell + 2)^{C_0}$ satisfies $n + 2 \leq T$ and the inequality in Equation (18). Put

$$X := \max \left\{ X_0, \left\lceil (\ell + 2)^{C_0} \right\rceil \right\}.$$

By the defining property of X_0 , the interval $[X, 2X]$ contains a prime congruent to 3 modulo 4. The reduction scans the integers in this interval that are congruent to 3 modulo 4 and applies deterministic primality testing [AKS04]. It finds a prime $q \equiv 3 \pmod{4}$ with $q = \ell^{O(1)}$ in deterministic polynomial time.

It remains to represent the canonical norm by the standard coordinate ℓ_2 -norm. Let $d := q - 1$. Using the integral basis of $\mathcal{O}_K$ from Section 2, consider the vectors

$$\zeta^t \widehat{\mathbf{m}}_i, \quad 0 \leq t < d, \quad i \in \{1, 2\},$$

which form a $\mathbb{Z}$ -basis of $\widehat{\mathcal{M}}$. Let $\mathbf{C}_{\widehat{\mathcal{M}}} \in \mathbb{Z}^{2d \times 2d}$ be the matrix whose columns are their power-basis coefficient vectors.

To realize the canonical inner product, let $\mathbf{E}_q \in \mathbb{Z}^{\binom{q}{2} \times d}$ be an oriented incidence matrix of the complete graph on q vertices with one vertex column deleted. Its Gram matrix is

$$\mathbf{E}_q^\top \mathbf{E}_q = q\mathbf{I}_d - \mathbf{1}_d \mathbf{1}_d^\top.$$

Thus Lemma 2.1 gives the following integer basis matrix:

$$\mathbf{B} := (\mathbf{I}_2 \otimes \mathbf{E}_q) \mathbf{C}_{\widehat{\mathcal{M}}} \in \mathbb{Z}^{2\binom{q}{2} \times 2d}, \quad (35)$$

The lattice $\mathcal{L}(\mathbf{B})$ is isometric to $\widehat{\mathcal{M}}$. The Gram identity above shows that $\mathbf{I}_2 \otimes \mathbf{E}_q$ has full column rank, while $\det(\widehat{\mathbf{B}}) \neq 0$ shows that $\mathbf{C}_{\widehat{\mathcal{M}}}$ has full rank. Hence $\mathbf{B}$ has full column rank. Its dimensions and encoding length are polynomial in ℓ .

5.5 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 $\mathbb{Z}$ -lattice, each output has rank $2(q - 1)$. For NP-hardness, consider the two branches in Section 5.4. If the input has an isolated universe element, then both the input and the fixed output are negative. Otherwise, the reduction maps the preprocessed X3C instance to $(q, \widehat{\mathbf{m}}_1, \widehat{\mathbf{m}}_2, \widehat{S})$. On this branch, the construction is deterministic and runs in polynomial time. The two module generators lie in $\mathcal{O}_K^2$ and have nonzero determinant, and Equation (32) proves the exact YES-NO equivalence. For the polynomial-time computable integer matrix $\mathbf{B}$ in Equation (35), the lattice $\mathcal{L}(\mathbf{B})$ is isometric to the canonically embedded module $\widehat{\mathcal{M}}$.

For membership in NP, first parse the encoded string and check the coefficient format, the nonnegative integer threshold, primality, $q \equiv 3 \pmod{4}$, and the nonzero determinant condition. Reject the input if any check fails. For a valid input, apply the coefficient-matrix construction in Section 5.4 to the input generators $\mathbf{m}_1, \mathbf{m}_2$. This produces in polynomial time an integer matrix $\mathbf{B}$ such that $\mathcal{L}(\mathbf{B})$ is isometric to the module lattice defined by the input generators. The nonzero K -determinant guarantees that $\mathbf{B}$ has full column rank. The standard verifier for decision-SVP from Section 2 then applies. This proves Theorem 1.1. $\square$

Proof of Corollary 1.2. Given an X3C instance, first check for an isolated universe element. If one exists, return NO without calling the oracle. Otherwise, apply the answer-preserving preprocessing above, carry out the reduction, and construct the integer basis $\mathbf{B}$ from Section 5.4. The lattice $\mathcal{L}(\mathbf{B})$ is isometric to the full-rank free module lattice $\widehat{\mathcal{M}} \subseteq \mathcal{O}_K^2$. A search-SVP oracle therefore returns a shortest nonzero vector $\mathbf{y} \in \mathcal{L}(\mathbf{B})$. By Equation (32), the X3C instance is positive exactly when $\|\mathbf{y}\|_2^2 \leq \widehat{S}$. Thus at most one oracle call decides X3C in polynomial time, which proves the claimed Turing hardness. $\square$

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 [Wan26b], and then developed the deterministic reduction proved in this paper. 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, Tancrede 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 $(1 + 1/\dim^\epsilon)$ 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, Tancrede 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.
- [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. *arXiv: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 $2^{\log^{1-\epsilon} n}$. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1716–1727. ACM, 2026.
- [IK04] Henryk Iwaniec and Emmanuel Kowalski. *Analytic Number Theory*, volume 53 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, 2004.
- [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.
- [LN97] Rudolf Lidl and Harald Niederreiter. *Finite Fields*. Cambridge University Press, 2 edition, 1997.
- [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.
- [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.
- [MMPMT23] 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.
- [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. 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, 2 edition, 1997.