Linear Descent for Rank-2 and Rank-4 Module-LIP in the Standard Cyclotomic Orbit

Anonymous authors

Anonymous submission

Abstract. Let

$$R = \mathbb{Z}[X]/(X^N + 1),$$

where $N \geq 4$ is a power of two. We study the following standard-orbit search problem in module ranks $r \in \{2, 4\}$: given a positive-definite Hermitian matrix $Q \in M_r(R)$ with $\det Q = 1$, promised that

$$B^* Q B = I_r \qquad \text{for some } B \in \mathrm{SL}_r(R),$$

recover and verify such a transition. Throughout, the source lattice is (R^r, I_r) and the target lies in its special-linear orbit.

In rank two, the symmetric square identifies a public rank-$3N/2$ lattice with the trace-zero Q-self-adjoint endomorphisms of R^2. At the standard point its complete shortest shell is explicit and consists of self-adjoint involutions. A target shortest vector is therefore a non-scalar Q-orthogonal automorphism. We then invoke the instance-preserving rank-two Module-LIP-to-Module-LAP reduction of Luo et al. [6] to reconstruct an equivalent Gram factor and hence a transition. A calibrated primal–dual lattice and the Bambury–Nguyen reduction lower the largest γ-SVP oracle rank to $3N/4 + 1$ for every fixed $\gamma < \sqrt{2}$.

In rank four, the Hermitian Hodge involution on $\bigwedge^2_{\mathbb{L}} \mathbb{L}^4$ defines a public rank-$3N$ fixed lattice. The standard fixed lattice is hypercubic. After recovering an orthonormal basis, the public action of $s = \zeta + \zeta^{-1}$ reduces the remaining matching problem to a polynomial-size enumeration of signed cycle intertwiners. For a matched exterior-square map, four wedge annihilators recover the image lines of the standard basis, and four rank-one Lenstra–Silverberg computations recover and verify the transition. The hypercubic Bambury–Nguyen algorithm restricts all γ-SVP calls to ranks at most $3N/2 + 1$ for $\gamma < \sqrt{2 - 2/(3N)}$.

Keywords: Module-LIP · linear descent · symmetric square · Hodge involution · exterior square · shortest vector problem

1 Introduction

The module lattice isomorphism problem asks to recover a module isometry between two promised-isometric Hermitian lattices carrying a public number-field action. In rank two over power-of-two cyclotomic rings, this is the structured isomorphism problem introduced with HAWK [3]. The public R-action makes tensor constructions over the cyclotomic field available, suggesting that one should move the hidden transition to a lower-dimensional representation before invoking lattice algorithms.

Dimension reduction alone is not a search reduction. A useful carrier must satisfy three requirements simultaneously. First, its integral lattice and Euclidean metric must be computable from the public form Q and transported by the hidden matrix. Second, the relevant output of the auxiliary lattice algorithm must be recognizable. Third, the recognized data must lift back to a matrix in $\mathrm{SL}_r(R)$ satisfying the original Hermitian equation. The last point is essential: an arbitrary isometry of a tensor lattice need not belong to the image of the tensor representation.

This paper treats only the standard free module and the standard unit form. Fix a power of two $N \geq 4$ and put

$$\mathbb{L} = \mathbb{Q}[X]/(X^N + 1), \qquad R = \mathbb{Z}[X]/(X^N + 1).$$

For $r \in \{2, 4\}$, the input is a positive-definite Hermitian $Q \in M_r(R)$ with $\det Q = 1$, together with the promise that

$$B^* Q B = I_r \qquad \text{for some } B \in \mathrm{SL}_r(R). \tag{1}$$

The task is to recover a verified matrix $C \in \mathrm{SL}_r(R)$ satisfying $C^* Q C = I_r$. Throughout, the source is the free standard lattice (R^r, I_r) and the target lies in its special-linear orbit. Fractional ideals appear only when a recovered line is intersected with R^r.

1.1 Results

Rank two. The identity

$$M^t J_2 M = (\det M) J_2, \qquad J_2 = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix},$$

turns the public form Q into a conjugate-linear symmetry J_Q of $\mathbb{L}^2$. Its symmetric square has a three-dimensional fixed space over the maximal real subfield. The natural map

$$\Phi : \mathrm{Sym}^2_{\mathbb{L}}(\mathbb{L}^2) \longrightarrow \mathrm{End}^0_{\mathbb{L}}(\mathbb{L}^2)$$

identifies this fixed space with the trace-zero Q-self-adjoint endomorphisms and converts the symmetric-square action of B into conjugation by B. Imposing integral matrix entries gives a public $\mathbb{Z}$-lattice $\mathcal{S}(Q)$ of rank $3N/2$.

At $Q = I_2$, the lattice has an explicit orthogonal basis. Its complete shortest shell is

$$\{\pm D\} \cup \{\pm E_j : 0 \leq j < N\},$$

where $D = \mathrm{diag}(1, -1)$ and $E_j = \left(\begin{smallmatrix} 0 & X^j \\ X^{-j} & 0 \end{smallmatrix}\right)$. The next distinct length is $\sqrt{2}$, and every shortest vector is an involution. Hence one γ-SVP output with $\gamma < \sqrt{2}$

on the target lattice is a non-scalar Q-orthogonal automorphism. The instance-preserving rank-two Module-LIP-to-Module-LAP reduction of Luo et al. [6, Thm. 3.1] then recovers an equivalent Gram factor in polynomial time; a final unitary determinant normalization yields a matrix in $\mathrm{SL}_2(R)$ satisfying $C^*QC = I_2$. A calibrated primal–dual version of the same carrier has minima product $1/2$; the Bambury–Nguyen block reduction [1] therefore restricts the largest oracle rank to $3N/4 + 1$.

Rank four. On $\bigwedge_{\mathbb{L}}^2 \mathbb{L}^4$, the wedge product and the Hermitian form define a conjugate-linear Hodge involution $\mathcal{J}_Q$. Its fixed space has dimension six over the maximal real subfield, and its integral intersection with $\bigwedge_R^2 R^4$ is a public $\mathbb{Z}$-lattice $\mathcal{W}(Q)$ of rank $3N$. For $Q = I_4$ this lattice is exactly three coefficient copies of R, hence hypercubic.

A shortest vector alone does not determine the exterior-square map. We first recover a complete orthonormal basis. Multiplication by $s = X + X^{-1}$ acts on the standard basis as three signed N-cycles. Since the hidden map is R-linear, it intertwines this public action. All signed basis bijections commuting with s can be enumerated in at most $6(4N)^3$ branches, one of which is the restriction of $\bigwedge^2 B$. Each branch extends to an $\mathbb{L}$-linear map on the full exterior square. Four wedge-annihilator computations then recover the image lines of the standard basis. Intersecting these lines with R^4 and applying Lenstra–Silverberg four times recovers the column scalars; the compound matrix, determinant, module, and Hermitian equations are checked exactly. The hypercubic algorithm of Bambury and Nguyen reduces the largest oracle rank to $3N/2 + 1$.

1.2 Technique overview

Both descents follow the same proof pattern. A determinant-one tensor identity produces a public conjugate-linear involution whose fixed space is a real form of a smaller representation. The hidden transition carries the standard fixed lattice isometrically to the target fixed lattice. The standard cyclotomic coefficient basis then makes the source geometry explicit.

The decoding mechanisms are different. In rank two, the shortest vector is already a nontrivial automorphism of the public Hermitian form. We do not repeat the eigensubmodule reconstruction: the existing instance-preserving rank-two Module-LIP-to-Module-LAP reduction converts this automorphism into an equivalent Gram factor for the same public instance.

In rank four, a basis of the fixed lattice still determines only an ordinary integral isometry. The action of the real subring removes this ambiguity and yields a short list of S-linear candidates. For a matched candidate $T = \bigwedge^2 B$, the three-dimensional subspace $T(e_i \wedge \mathbb{L}^4)$ has wedge annihilator $\mathbb{L}Be_i$. These four lines recover the columns projectively; four norm-constrained generators recover the missing scalars.

1.3 Related work and organization

The rank-two use of a nontrivial module-lattice automorphism follows the symplectic-automorphism line of work of Luo, Jiang, Pan, and Wang [7]. Its final factor-recovery step is exactly the instance-preserving Module-LIP-to-Module-LAP reduction of Luo, Jiang, Jin, Pan, and Wang [6, Thm. 3.1]; we cite that result rather than repeating its eigensubmodule analysis. The rank-one norm-constrained generator algorithm used in rank four is due to Lenstra and Silverberg [5]. The reductions that replace full-rank SVP calls by calls in essentially half the rank are used as black boxes from Bambury and Nguyen [1]; they are not new contributions of this paper.

Section 2 fixes notation and states the external lattice algorithms. Section 3 gives the symmetric-square carrier, extracts a nontrivial automorphism, and invokes the cited factor-recovery theorem to obtain the direct and half-rank reductions. Section 4 gives the Hodge carrier, signed-cycle completion, the four-line decoder, and the corresponding oracle-rank bounds. Section 5 records exact public constructions used by both algorithms.

2 Preliminaries

2.1 Cyclotomic and Hermitian notation

Let $N \geq 4$ be a power of two, let ζ denote the image of X, and put

$$\mathbb{L} = \mathbb{Q}(\zeta) = \mathbb{Q}[X]/(X^N + 1), \qquad R = \mathbb{Z}[\zeta] = \mathbb{Z}[X]/(X^N + 1).$$

Complex conjugation is the involution $a \mapsto \bar{a}$ determined by $\bar{\zeta} = \zeta^{-1}$. Its fixed field and fixed ring are

$$\mathbb{L}^+ = \{a \in \mathbb{L} : \bar{a} = a\}, \qquad S = R \cap \mathbb{L}^+.$$

Write

$$d = [\mathbb{L}^+ : \mathbb{Q}] = N/2.$$

The roots of unity in $\mathbb{L}$ are $\mu(\mathbb{L}) = \left\{\zeta^j : 0 \leq j < 2N\right\}$.

Vectors are columns. For a vector x and a matrix A over $\mathbb{L}$, write $x^* = \bar{x}^{\,t}$ and $A^* = \bar{A}^{\,t}$. A Hermitian matrix $Q = Q^*$ defines

$$h_Q(x, y) = x^* Q y, \qquad X^{\dagger_Q} = Q^{-1} X^* Q.$$

Thus $X^{\dagger_Q} = X$ means that X is self-adjoint for h_Q. The matrix Q is positive definite if each complex embedding gives a positive-definite complex Hermitian matrix.

For

$$f = \sum_{i=0}^{N-1} a_i \zeta^i \in R, \qquad \mathrm{vec}(f) = (a_0, \ldots, a_{N-1})^t,$$

the coefficient inner product satisfies

$$\langle f, g\rangle_{\mathrm{coeff}} = \mathrm{vec}(f)^t \,\mathrm{vec}(g) = \frac{1}{N}\,\mathrm{Tr}_{\mathbb{L}/\mathbb{Q}}(\bar{f} g). \tag{2}$$

Hence $1, \zeta, \ldots, \zeta^{N-1}$ is an orthonormal $\mathbb{Z}$-basis of R. Put

$$s = \zeta + \zeta^{-1}, \qquad s_i = \zeta^i + \zeta^{-i} \quad (1 \le i < d).$$

Then

$$S = \mathbb{Z}\cdot 1 \oplus \bigoplus_{i=1}^{d-1} \mathbb{Z}\cdot s_i = \mathbb{Z}[s]. \tag{3}$$

The basis $1, s_1, \ldots, s_{d-1}$ is orthogonal for the form $d^{-1}\,\mathrm{Tr}_{\mathbb{L}^+/\mathbb{Q}}(ab)$. Its squared lengths are $1, 2, \ldots, 2$. For the recurrence, set $s_0 = 2$ and $s_d = 0$. Then

$$s s_i = s_{i+1} + s_{i-1} \qquad (1 \le i < d),$$

which also proves $S = \mathbb{Z}[s]$.

2.2 The standard-orbit search problem

For $r \in \{2, 4\}$, a *standard-orbit instance* is a matrix $Q \in M_r(R)$ satisfying

$$Q = Q^* > 0, \qquad \det Q = 1,$$

together with the promise that there is $B \in \mathrm{SL}_r(R)$ such that $B^* Q B = I_r$. The output is any verified $C \in \mathrm{SL}_r(R)$ satisfying $C^* Q C = I_r$.

The matrix Q may be given directly, or reconstructed from the public R-action and the rational trace Gram matrix. Indeed, if $\omega_1, \ldots, \omega_N$ is a $\mathbb{Z}$-basis of R and $\eta_1, \ldots, \eta_N$ is its conjugate trace-dual basis,

$$\mathrm{Tr}_{\mathbb{L}/\mathbb{Q}}(\bar{\omega}_i \eta_j) = \delta_{ij},$$

then the normalized trace form $g_Q(x, y) = N^{-1}\,\mathrm{Tr}_{\mathbb{L}/\mathbb{Q}}\, h_Q(x, y)$ determines

$$h_Q(x, y) = N \sum_{j=1}^{N} g_Q(\omega_j x, y)\eta_j. \tag{4}$$

Thus every matrix, adjoint, module, and Hermitian equation used below is an exact public test.

2.3 Rank-one recovery

A fractional R-ideal is a nonzero finitely generated R-submodule of $\mathbb{L}$ spanning $\mathbb{L}$ over $\mathbb{Q}$.

Theorem 1 (Lenstra–Silverberg). *We use the following specialization of [5, Thm. 1.3 and Prop. 7.4]. Let $\mathfrak{I}$ be a fractional R-ideal and let $w \in \mathbb{L}^+$ be totally positive. There is a deterministic polynomial-time algorithm that decides whether there exists $v \in \mathbb{L}^\times$ such that*

$$\mathfrak{I} = Rv, \qquad v\bar{v} = w,$$

and returns one when it exists. If v is one solution, then all solutions are $v\mu(\mathbb{L})$.

In the two decoders below, the source line is free. Given a fractional ideal $\mathfrak{b}$ and totally positive $q_0, q_1 \in \mathbb{L}^+$, applying the algorithm to

$$\mathfrak{I} = \mathfrak{b}, \qquad w = q_0/q_1$$

recovers all scalars c, up to roots of unity, satisfying

$$\mathfrak{b} = Rc, \qquad q_1 c\bar{c} = q_0.$$

2.4 Symmetric and exterior powers

For an $\mathbb{L}$-space V, write $v \odot w$ for the class of $(v \otimes w + w \otimes v)/2$ in $\mathrm{Sym}^2_{\mathbb{L}}(V)$. A linear or conjugate-linear map $A : V \to V$ induces the diagonal action $\mathrm{Sym}^2(A)$; if $A^2 = -I$, then $-(A \otimes A)$ induces a conjugate-linear involution on $\mathrm{Sym}^2_{\mathbb{L}}(V)$.

The second exterior power is denoted $\bigwedge^2_{\mathbb{L}} V$. In the standard basis of $\mathbb{L}^4$, write $e_{ij} = e_i \wedge e_j$ and use the ordered basis

$$\mathcal{E}_2 = (e_{12}, e_{13}, e_{14}, e_{23}, e_{24}, e_{34}). \tag{5}$$

For $M \in M_4(\mathbb{L})$, the second compound matrix $C_2(M)$ is the matrix of $\bigwedge^2 M$ in $\mathcal{E}_2$. It satisfies

$$C_2(AB) = C_2(A)C_2(B), \qquad C_2(M^*) = C_2(M)^*.$$

The Hermitian form induced by Q on $\bigwedge^2 \mathbb{L}^4$ has matrix

$$Q_2 = C_2(Q).$$

We shall repeatedly use the following elementary real-form fact.

Lemma 2 (Fixed space of a conjugate-linear involution). *Let U be an n-dimensional $\mathbb{L}$-space and let $\tau : U \to U$ be a conjugate-linear involution. Then* $\mathrm{Fix}(\tau)$ *has dimension n over $\mathbb{L}^+$, and scalar extension gives*

$$\mathbb{L} \otimes_{\mathbb{L}^+} \mathrm{Fix}(\tau) \simeq U.$$

Proof. Choose $\vartheta \in \mathbb{L}^\times$ with $\bar{\vartheta} = -\vartheta$. Every $u \in U$ has the unique decomposition

$$u = \frac{u + \tau(u)}{2} + \vartheta \frac{u - \tau(u)}{2\vartheta},$$

and both displayed coefficients lie in $\mathrm{Fix}(\tau)$. □

2.5 Euclidean lattices and half-rank reductions

For a Euclidean lattice Λ, let $\Lambda^\vee$ be its Euclidean dual, let $\lambda_1(\Lambda)$ be its first minimum, and let $\mathrm{Min}(\Lambda)$ be its complete shortest shell. A γ-SVP oracle returns a nonzero $x \in \Lambda$ satisfying $\|x\| \le \gamma\lambda_1(\Lambda)$. A lattice is *hypercubic* if it has an orthonormal $\mathbb{Z}$-basis after a common scaling.

We use the following two consequences of Bambury and Nguyen. They are stated in the form needed later.

Theorem 3 (Primal–dual shell extraction). *We use the following specialization of [1, Thm. 1 and Eq. (3)]. Let Λ be an integral rank-m Euclidean lattice. Suppose*

$$\lambda_1(\Lambda) = a, \qquad \lambda_1(\Lambda^\vee) = a^\vee,$$

and suppose that, in both the primal and dual lattices, every vector outside the shortest shell has norm at least Δ times the first minimum, for some $\Delta > 1$. Let $1 \le \gamma < \Delta$. If

$$\gamma^2 a a^\vee < 1 - \varepsilon \qquad \text{for some } \varepsilon \ge 1/\operatorname{poly}(m), \tag{6}$$

then a vector in $\mathrm{Min}(\Lambda)$ or in $\mathrm{Min}(\Lambda^\vee)$ is recovered with polynomially many γ-SVP calls, each in rank at most

$$\lfloor m/2 \rfloor + 1.$$

Proof (Justification of the stated interface). Bambury and Nguyen's block reduction uses two reduction calls and two exact shell-detection calls per iteration. Their inequality (3) permits the reduction calls to use approximation factor γ under Equation (6). A shell-detection sublattice has the same ambient length spectrum: if it contains a global shortest vector, a γ-SVP output is still in the shortest shell because $\gamma < \Delta$; if it does not, every nonzero vector has norm at least Δa (or $\Delta a^\vee$). Thus the exact shell tests can also be implemented with the same γ-SVP oracle. The rank bound is the maximum of the four block ranks in their theorem. $\square$

Theorem 4 (Hypercubic basis recovery). *We use [1, Thm. 4] in the following form. Let Λ be a rank-m hypercubic lattice. For every*

$$1 \le \gamma < \sqrt{2 - \frac{2}{m}},$$

an orthonormal lattice basis, up to the common scale, is recovered with polynomially many γ-SVP calls, each in rank at most $\lfloor m/2 \rfloor + 1$.

All Gram matrices supplied to these two algorithms below are integral. At the standard point this is explicit; at the target point it follows from the promised integral isometry carrying an integral source basis to an integral target basis. No rational-Gram rescaling is needed.

3 Rank two: symmetric-square descent

Throughout this section, $Q \in M_2(R)$ is a standard-orbit input and $B \in \mathrm{SL}_2(R)$ denotes a hidden matrix satisfying $B^*QB = I_2$.

3.1 The public symmetry and its matrix model

Put

$$J_2 = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \qquad \tau_2(v) = J_2\bar{v}.$$

The elementary identity

$$M^t J_2 M = (\det M) J_2 \qquad (M \in M_2(\mathbb{L})) \tag{7}$$

will be used repeatedly. Define the public conjugate-linear map

$$J_Q(v) = J_2\bar{Q}\,\bar{v}. \tag{8}$$

Proposition 5 (Conjugated symplectic symmetry). *The map J_Q satisfies $J_Q^2 = -I_2$. Moreover,*

$$\boxed{J_Q = B\tau_2 B^{-1}.} \tag{9}$$

Consequently

$$\Sigma_Q = -\left.(J_Q \otimes J_Q)\right|_{\mathrm{Sym}^2_{\mathbb{L}}(\mathbb{L}^2)}$$

is a conjugate-linear involution, and

$$W_2(Q) := \mathrm{Fix}(\Sigma_Q)$$

has dimension three over $\mathbb{L}^+$.

Proof. Since $Q = Q^*$ and $\det Q = 1$, applying Equation (7) to Q gives $\bar{Q}J_2Q = J_2$. Hence

$$J_Q^2 = J_2\bar{Q}J_2Q = J_2^2 = -I_2.$$

From $B^*QB = I_2$ we obtain $\bar{Q} = B^{-t}\bar{B}^{-1}$. Applying Equation (7) to B^t gives $J_2B^{-t} = BJ_2$, and therefore

$$J_Q(v) = J_2B^{-t}\bar{B}^{-1}\bar{v} = BJ_2\overline{B^{-1}v} = B\tau_2(B^{-1}v).$$

The assertion about Σ_Q follows by taking the diagonal tensor square; the dimension statement is Lemma 2. $\square$

For $u \in \mathrm{Sym}^2_{\mathbb{L}}(\mathbb{L}^2)$, let $\kappa(u)$ be the corresponding symmetric matrix, characterized by

$$\kappa(v \odot w) = \frac{vw^t + wv^t}{2}.$$

Define

$$\Phi(u) = \kappa(u)J_2. \tag{10}$$

Thus, if $u = xe_1^2 + 2y(e_1 \odot e_2) + ze_2^2$, then

$$\Phi(u) = \begin{pmatrix} -y & x \\ -z & y \end{pmatrix}.$$

Proposition 6 (Symmetric square as the adjoint representation). *The map*

$$\Phi : \mathrm{Sym}^2_{\mathbb{L}}(\mathbb{L}^2) \longrightarrow \{X \in M_2(\mathbb{L}) : \operatorname{tr} X = 0\}$$

is an $\mathbb{L}$-linear isomorphism. It satisfies

$$\Phi(\Sigma_Q u) = \Phi(u)^{\dagger_Q}, \tag{11}$$

$$\Phi(\mathrm{Sym}^2(C)u) = C\Phi(u)C^{-1} \qquad (C \in \mathrm{SL}_2(\mathbb{L})). \tag{12}$$

Hence

$$\Phi(W_2(Q)) = \left\{X \in M_2(\mathbb{L}) : X^{\dagger_Q} = X,\ \operatorname{tr} X = 0\right\}. \tag{13}$$

Proof. If $S^t = S$, then $\operatorname{tr}(SJ_2) = 0$; conversely, if $\operatorname{tr} X = 0$, then $-XJ_2$ is symmetric. This proves that Φ is an isomorphism.

For $C \in \mathrm{SL}_2(\mathbb{L})$,

$$\kappa(\mathrm{Sym}^2(C)u) = C\kappa(u)C^t, \qquad C^t J_2 = J_2 C^{-1},$$

which gives Equation (12). Let $S = \kappa(u)$. The semilinear matrix of J_Q is $J_2\bar{Q}$, so

$$\kappa((J_Q \otimes J_Q)u) = J_2\bar{Q}\,\bar{S}\,(J_2\bar{Q})^t.$$

Using $(J_2\bar{Q})^t = -QJ_2$ and the minus sign in Σ_Q gives

$$\Phi(\Sigma_Q u) = -J_2\bar{Q}\,\bar{S}Q.$$

The identity $QJ_2\bar{Q} = J_2$ implies $Q^{-1}J_2 = J_2\bar{Q}$, whence

$$\Phi(u)^{\dagger_Q} = Q^{-1}(SJ_2)^*Q = -Q^{-1}J_2\bar{S}Q = -J_2\bar{Q}\,\bar{S}Q.$$

This proves Equation (11) and then Equation (13). □

3.2 The integral carrier and its standard shell

We use the matrix model from this point onward. Define

$$\boxed{\mathcal{S}(Q) = \left\{X \in M_2(R) : X^{\dagger_Q} = X,\ \operatorname{tr} X = 0\right\}.} \tag{14}$$

For $X, Y \in \mathcal{S}(Q)$, put

$$\boxed{\langle X, Y\rangle_2 = \frac{1}{4d}\operatorname{Tr}_{\mathbb{L}/\mathbb{Q}} \operatorname{tr}(XY).} \tag{15}$$

At every embedding, a Q-self-adjoint trace-zero matrix has real eigenvalues $\lambda, -\lambda$, so this form is positive definite.

Proposition 7 (Integral transport). *The map*

$$\mathrm{Ad}_B : X \longmapsto BXB^{-1}$$

is an isometry from $\mathcal{S}(I_2)$ onto $\mathcal{S}(Q)$. In particular, $\mathcal{S}(Q)$ has $\mathbb{Z}$-rank $3d = 3N/2$.

Proof. Because $B, B^{-1} \in M_2(R)$, conjugation preserves integral matrices in both directions. From $B^*QB = I_2$ one checks directly that $X^* = X$ is equivalent to $(BXB^{-1})^{\dagger_Q} = BXB^{-1}$, and trace is invariant under conjugation. Finally,

$$\operatorname{tr}\big((BXB^{-1})(BYB^{-1})\big) = \operatorname{tr}(XY),$$

so the metric is preserved. The source rank is computed below. □

At the standard point, every element has the unique form

$$X(a, z) = \begin{pmatrix} a & z \\ \bar{z} & -a \end{pmatrix}, \qquad a \in S, \quad z \in R. \tag{16}$$

It satisfies

$$X(a, z)^2 = (a^2 + z\bar{z})I_2, \qquad \|X(a, z)\|_2^2 = \frac{1}{d} \operatorname{Tr}_{\mathbb{L}^+/\mathbb{Q}}(a^2 + z\bar{z}). \tag{17}$$

Define

$$D = \operatorname{diag}(1, -1), \qquad E_j = \begin{pmatrix} 0 & \zeta^j \\ \zeta^{-j} & 0 \end{pmatrix} \quad (0 \le j < N). \tag{18}$$

Proposition 8 (Standard rank-two geometry). *The lattice $\mathcal{S}(I_2)$ has the orthogonal $\mathbb{Z}$-basis*

$$D, \quad X(s_1, 0), \ldots, X(s_{d-1}, 0), \quad E_0, \ldots, E_{N-1}, \tag{19}$$

with squared lengths

$$1, \quad \underbrace{2, \ldots, 2}_{d-1}, \quad \underbrace{1, \ldots, 1}_{N}.$$

Its complete shortest shell is

$$\operatorname{Min}(\mathcal{S}(I_2)) = \{\pm D\} \cup \{\pm E_j : 0 \le j < N\}. \tag{20}$$

The next distinct length is $\sqrt{2}$, and every vector in Equation (20) *is an involution.*

Proof. The description Equation (16), the orthogonal bases of S and R from Equations (2) and (3), and Equation (17) give the displayed orthogonal basis and its lengths. A nonzero coefficient vector outside the displayed unit coordinate vectors has squared norm at least two. Hence Equation (20) is the complete shortest shell. Finally, $D^2 = E_j^2 = I_2$. □

3.3 Recovering the factor from a nontrivial automorphism

The auxiliary shortest vector already solves the rank-two automorphism problem for the same public form. We use the following specialization of the instance-preserving reduction of Luo et al.

Theorem 9 (Rank-two factor recovery from an automorphism). *Let Q be a standard-orbit instance. Given a matrix $A \in \mathrm{GL}_2(R)$ such that*

$$A^*QA = Q, \qquad A \notin \mu(\mathbb{L})I_2, \tag{21}$$

there is a deterministic polynomial-time exact algorithm that returns $C \in \mathrm{SL}_2(R)$ satisfying $C^*QC = I_2$.

Proof. The promise gives a factorization $Q = U^*U$ with $U = B^{-1} \in \mathrm{SL}_2(R)$. By the instance-preserving Module-LIP-to-Module-LAP reduction of [6, Thm. 3.1], the public form Q and the nontrivial automorphism A determine, in deterministic polynomial time, an equivalent factor $\widetilde{U} \in \mathrm{GL}_2(R)$ satisfying

$$Q = \widetilde{U}^*\widetilde{U}.$$

It remains only to impose the determinant-one convention of the present paper. Put $\delta = \det \widetilde{U}$. Since $\det Q = 1$, $\bar{\delta}\delta = 1$. Thus every conjugate of the algebraic unit δ has absolute value one, and hence $\delta \in \mu(\mathbb{L})$. The unitary matrix

$$V = \mathrm{diag}(\delta^{-1}, 1)$$

satisfies $V^*V = I_2$. Therefore $U' = V\widetilde{U} \in \mathrm{SL}_2(R)$ still satisfies $Q = U'^*U'$. Returning $C = U'^{-1}$ gives $C \in \mathrm{SL}_2(R)$ and $C^*QC = I_2$. □

3.4 Direct and half-rank reductions

Theorem 10 (Direct rank-two reduction). *For every $1 \le \gamma < \sqrt{2}$, one γ-SVP call on the rank-$3N/2$ lattice $\mathcal{S}(Q)$, followed by deterministic polynomial-time exact computations, recovers a verified $C \in \mathrm{SL}_2(R)$ satisfying $C^*QC = I_2$.*

Proof. By Proposition 7, the complete target shortest shell is the image under Ad_B of Equation (20). The gap in Proposition 8 forces every γ-SVP output with $\gamma < \sqrt{2}$ to be a target shortest vector A. Hence $A^2 = I_2$, $A^{\dagger_Q} = A$, and $\operatorname{tr} A = 0$. The first two identities imply $A^*QA = Q$, while trace zero and $A \neq 0$ imply $A \notin \mu(\mathbb{L})I_2$. Apply Theorem 9. □

To reduce the oracle rank, define

$$\mathcal{H}(Q) = \left\{A \in M_2(R) : A^{\dagger_Q} = A\right\}, \tag{22}$$
$$\mathcal{T}(Q) = \left\{2A - \mathrm{tr}(A)I_2 : A \in \mathcal{H}(Q)\right\}, \tag{23}$$
$$\Gamma(Q) = 2\mathcal{T}(Q)^\vee, \tag{24}$$

where the dual is taken in the trace-zero self-adjoint space with the metric Equation (15). Conjugation by B transports $\mathcal{H}$, $\mathcal{T}$, and Γ isometrically.

Proposition 11 (Cyclotomic calibration). *The lattice $\Gamma(I_2)$ has the orthogonal basis*

$$E_0, \ldots, E_{N-1}, \quad X(s_1, 0), \ldots, X(s_{d-1}, 0), \quad 2D, \tag{25}$$

with squared lengths

$$\underbrace{1,\dots,1}_{N},\quad \underbrace{2,\dots,2}_{d-1},\quad 4.$$

Consequently

$$\lambda_1(\Gamma(I_2)) = 1, \qquad \lambda_1(\Gamma(I_2)^\vee) = \frac{1}{2}.$$

The primal shortest shell is $\{\pm E_j\}$*, the dual shortest shell is* $\{\pm D/2\}$*, and in both lattices the ratio to the next distinct length is* $\sqrt{2}$*.*

Proof. The trace-removal lattice $\mathcal{T}(I_2)$ has the orthogonal basis

$$D,\quad X(s_1,0),\dots,X(s_{d-1},0),\quad 2E_0,\dots,2E_{N-1},$$

with squared lengths $1,2,\dots,2,4,\dots,4$. Taking the Euclidean dual and multiplying by two gives Equation (25). The primal and dual minima and the two gaps follow immediately. □

Theorem 12 (Rank-two reduction with half-rank oracles). *Fix* $1 \le \gamma < \sqrt{2}$*. A verified transition* $C \in \mathrm{SL}_2(R)$ *satisfying* $C^*QC = I_2$ *is recovered with polynomially many* γ*-SVP calls, every one of rank at most*

$$\boxed{\frac{3N}{4} + 1.} \tag{26}$$

Proof. By transport, $\Gamma(Q)$ has the same primal and dual geometry as $\Gamma(I_2)$. In Theorem 3, take

$$m = 3d = 3N/2, \qquad a = 1, \qquad a^\vee = 1/2, \qquad \Delta = \sqrt{2}.$$

Since $\gamma^2/2 < 1$, condition Equation (6) holds with a positive constant ε. If the returned vector is primal, set A equal to that vector; if it is dual, set A equal to twice that vector. By Proposition 11 and transport, in either case A is a trace-zero Q-orthogonal involution and therefore satisfies Equation (21). Apply Theorem 9. The rank bound is $\lfloor m/2 \rfloor + 1 = 3N/4 + 1$. □

4 Rank four: Hodge descent in the exterior square

Throughout this section, $Q \in M_4(R)$ is a standard-orbit input and $B \in \mathrm{SL}_4(R)$ satisfies $B^*QB = I_4$.

4.1 The Hermitian Hodge involution

Let

$$\mathrm{vol} = e_1 \wedge e_2 \wedge e_3 \wedge e_4.$$

Algorithm 1 Rank-two recovery in the standard orbit

Require: $Q \in M_2(R)$ satisfying the promise Equation (1)
Ensure: A verified $C \in \mathrm{SL}_2(R)$ with $C^*QC = I_2$

```
Compute Γ(Q) from Equations (22) to (24).
Apply Theorem 3 with the parameters in the proof of Theorem 12, obtaining a
  shortest vector Z in the primal or dual lattice.
if Z is primal then
    A ← Z.
else
    A ← 2Z.
end if
Apply the factor-recovery algorithm of Theorem 9 to (Q, A).
return the resulting verified transition C.
```

The wedge product defines a symmetric $\mathbb{L}$-bilinear form

$$\beta : \bigwedge_{\mathbb{L}}^{2} \mathbb{L}^4 \times \bigwedge_{\mathbb{L}}^{2} \mathbb{L}^4 \longrightarrow \mathbb{L}, \qquad \xi \wedge \eta = \beta(\xi, \eta)\,\mathrm{vol}\,.$$

In the basis $\mathcal{E}_2$ from Equation (5), its matrix is

$$\Omega = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \tag{27}$$

Thus, for coordinate columns $\mathbf{x} = [\xi]_{\mathcal{E}_2}$ and $\mathbf{y} = [\eta]_{\mathcal{E}_2}$,

$$\beta(\xi, \eta) = \mathbf{x}^t \Omega \mathbf{y}. \tag{28}$$

The exterior-power identity $\bigwedge^4 M = (\det M)I$ gives

$$\boxed{C_2(M)^t \Omega C_2(M) = (\det M)\Omega} \qquad (M \in M_4(\mathbb{L})). \tag{29}$$

Let $Q_2 = C_2(Q)$. The Hermitian Hodge operator is the unique conjugate-linear map $\mathcal{J}_Q$ satisfying

$$\beta(\mathcal{J}_Q \xi, \eta) = h_{Q,2}(\xi, \eta). \tag{30}$$

Since $\Omega^t = \Omega = \Omega^{-1}$, its coordinate rule is

$$\boxed{[\mathcal{J}_Q(\xi)]_{\mathcal{E}_2} = \Omega \overline{Q_2}\, \overline{[\xi]_{\mathcal{E}_2}}.} \tag{31}$$

At the standard point, $\mathcal{J}_0$ is the semilinear rule $\mathbf{x} \mapsto \Omega \bar{\mathbf{x}}$.

Proposition 13 (Hodge transport). *The map $\mathcal{J}_Q$ is a conjugate-linear involution. If $B_2 = \bigwedge^2 B$, then*

$$\boxed{\mathcal{J}_Q = B_2 \mathcal{J}_0 B_2^{-1}.} \tag{32}$$

Consequently

$$W_4(Q) := \mathrm{Fix}(\mathcal{J}_Q)$$

has dimension six over $\mathbb{L}^+$ and satisfies $W_4(Q) = B_2 W_4(I_4)$.

Proof. Since $Q_2 = Q_2^*$ and $\det Q = 1$, applying Equation (29) to Q gives $\overline{Q_2} \Omega Q_2 = \Omega$. Therefore

$$\Omega \overline{Q_2} \Omega Q_2 = I_6,$$

and Equation (31) implies $\mathcal{J}_Q^2 = I$.

Put $A = C_2(B)$. From $B^* Q B = I_4$ we obtain $Q_2 = (A^{-1})^* A^{-1}$ and $\overline{Q_2} = A^{-t} \bar{A}^{-1}$. Applying Equation (29) to B^t gives $\Omega A^{-t} = A \Omega$. Hence

$$\Omega \overline{Q_2} \bar{\mathbf{x}} = A \Omega \overline{A^{-1} \mathbf{x}},$$

which is Equation (32). The fixed-space statements follow from Lemma 2. □

4.2 The integral Hodge lattice and its standard geometry

Define

$$\boxed{\mathcal{W}(Q) = W_4(Q) \cap \bigwedge_R^2 R^4.} \tag{33}$$

For $\xi, \eta \in W_4(Q)$, put

$$\boxed{\langle \xi, \eta \rangle_4 = \frac{1}{2N} \mathrm{Tr}_{\mathbb{L}/\mathbb{Q}}\, h_{Q,2}(\xi, \eta).} \tag{34}$$

On the fixed space, Equation (30) gives $h_{Q,2}(\xi, \eta) = \beta(\xi, \eta) \in \mathbb{L}^+$; thus Equation (34) is a positive Euclidean inner product.

Proposition 14 (Integral Hodge transport). *The map $B_2 = \bigwedge^2 B$ is an S-linear Euclidean isometry*

$$B_2 : \mathcal{W}(I_4) \xrightarrow{\sim} \mathcal{W}(Q).$$

In particular, $\mathcal{W}(Q)$ has $\mathbb{Z}$-rank $3N$.

Proof. The matrix B_2 and its inverse have entries in R, so they transport the integral exterior modules in both directions. The fixed spaces are transported by Equation (32), and B_2 preserves the induced Hermitian form because $B^* Q B = I_4$. The source rank is computed below. □

Solving $\mathbf{x} = \Omega\bar{\mathbf{x}}$ at the standard point gives

$$[w(a,b,c)]_{\mathcal{E}_2} = (a,b,c,\bar{c},-\bar{b},\bar{a})^t, \qquad a,b,c \in R. \tag{35}$$

Moreover,

$$\|w(a,b,c)\|_4^2 = \frac{1}{d}\operatorname{Tr}_{\mathbb{L}^+/\mathbb{Q}}(a\bar{a}+b\bar{b}+c\bar{c}). \tag{36}$$

For $0 \le j < N$, define

$$u_{1,j} = w(\zeta^j,0,0), \qquad u_{2,j} = w(0,\zeta^j,0), \qquad u_{3,j} = w(0,0,\zeta^j), \tag{37}$$

and let $\mathcal{B}_0$ be their union.

Proposition 15 (Standard rank-four geometry). *The set $\mathcal{B}_0$ is an orthonormal $\mathbb{Z}$-basis of $\mathcal{W}(I_4)$. Consequently $\mathcal{W}(I_4)$ is hypercubic of rank $3N$, its shortest shell is*

$$\{\pm u_{k,j} : 1 \le k \le 3,\ 0 \le j < N\},$$

and its next distinct length is $\sqrt{2}$.

Proof. Equations Equations (35) and (36) identify $\mathcal{W}(I_4)$ isometrically with three orthogonal coefficient copies of R. The result follows from Equation (2). □

4.3 Recovering a transition from a matched exterior-square map

For a nonzero v in a four-dimensional $\mathbb{L}$-space V, put

$$U(v) = v \wedge V \subseteq \bigwedge_{\mathbb{L}}^2 V.$$

For $U \subseteq \bigwedge_{\mathbb{L}}^2 V$, put

$$\operatorname{Ann}_\wedge(U) = \{x \in V : x \wedge u = 0 \text{ for every } u \in U\}.$$

Lemma 16 (Wedge annihilator). *For every nonzero $v \in V$,*

$$\boxed{\operatorname{Ann}_\wedge(U(v)) = \mathbb{L}v.} \tag{38}$$

Moreover, for every isomorphism $C : V \to V$,

$$(\bigwedge^2 C)U(v) = U(Cv).$$

Proof. Every scalar multiple of v annihilates $v \wedge V$. If $x \notin \mathbb{L}v$, choose $y \notin \operatorname{Span}_{\mathbb{L}}(x,v)$. Then $x \wedge v \wedge y \neq 0$, so x is not in the annihilator. The transport identity follows from $(\bigwedge^2 C)(v \wedge y) = Cv \wedge Cy$. □

Proposition 17 (Four-line decoder). *Let*

$$T : \bigwedge_{\mathbb{L}}^{2} \mathbb{L}^4 \longrightarrow \bigwedge_{\mathbb{L}}^{2} \mathbb{L}^4$$

be an $\mathbb{L}$-linear isomorphism. There is a deterministic polynomial-time procedure that returns only verified solutions to the standard-orbit instance. If

$$T = \bigwedge^{2} B, \tag{39}$$

then the procedure returns a solution. It uses four calls to Theorem 1.

Proof. For $1 \le i \le 4$, compute

$$\ell_i = \mathrm{Ann}_{\wedge}\big(T(U(e_i))\big). \tag{40}$$

Reject unless every ℓ_i is an $\mathbb{L}$-line. Intersect with the fixed target module and compute

$$R^4 \cap \ell_i = \mathfrak{b}_i v_i, \qquad q_i = v_i^* Q v_i \in \mathbb{L}^+. \tag{41}$$

Each q_i is totally positive. The source line is Re_i and has norm one. Apply Theorem 1 to

$$\mathfrak{I}_i = \mathfrak{b}_i, \qquad w_i = 1/q_i. \tag{42}$$

For every successful call, enumerate all root-of-unity branches c_i with

$$\mathfrak{b}_i = Rc_i, \qquad q_i c_i \bar{c}_i = 1.$$

There are at most $(2N)^4$ four-tuples. For each one, define $C(e_i) = c_i v_i$ and retain it only if

$$C(R^4) = R^4, \qquad C^* Q C = I_4, \qquad \det C = 1, \qquad C_2(C) = T. \tag{43}$$

Every returned matrix is therefore a verified solution compatible with the candidate exterior-square map.

If Equation (39) holds, then Lemma 16 gives $\ell_i = \mathbb{L} B e_i$. Since $B \in \mathrm{GL}_4(R)$,

$$R^4 \cap \ell_i = R B e_i.$$

Thus the restriction $e_i \mapsto Be_i$ occurs among the root-of-unity branches of the ith Lenstra–Silverberg call. Their assembly is B and passes Equation (43). $\square$

4.4 Completing a fixed-lattice isometry

A single shortest vector of the Hodge lattice does not determine $\bigwedge^2 B$. We first recover a complete orthonormal basis and then use the public action of S to identify the hidden basis matching.

Lemma 18 (Successive shortest-vector basis recovery). *Let Λ be a rank-m hypercubic lattice of shortest length one. For $1 \le \gamma < \sqrt{2}$, an orthonormal $\mathbb{Z}$-basis is recovered with at most m γ-SVP calls, in ranks $m, m-1, \ldots, 1$.*

Proof. In coordinates relative to an orthonormal lattice basis, every squared length is a sum of integer squares. Hence a γ-SVP output with $\gamma < \sqrt{2}$ is a signed basis vector. After recovering $x_1, \ldots, x_k$, compute exactly

$$\Lambda \cap \operatorname{Span}_{\mathbb{R}}(x_1, \ldots, x_k)^{\perp}.$$

It is the hypercubic lattice spanned by the remaining basis vectors. Iterating gives the result. □

Multiplication by $s = \zeta + \zeta^{-1}$ acts on the standard basis by

$$su_{k,j} = u_{k,j-1} + u_{k,j+1}, \tag{44}$$

where

$$u_{k,-1} = -u_{k,N-1}, \qquad u_{k,N} = -u_{k,0}.$$

Since $s \in S$, multiplication by s preserves every Hodge lattice $\mathcal{W}(Q)$ and is public. In the standard basis its signed support graph is the disjoint union of three N-cycles, each with edge-sign product -1.

Proposition 19 (Signed-cycle completion). *Let $\mathcal{B}_1$ be any orthonormal $\mathbb{Z}$-basis of $\mathcal{W}(Q)$. From the public matrix of multiplication by s in $\mathcal{B}_1$, one can enumerate at most*

$$6(4N)^3 \tag{45}$$

signed bijections $T_W : \mathcal{B}_0 \to \mathcal{B}_1$ whose $\mathbb{Z}$-linear extensions commute with s. Every enumerated map extends to an S-linear Euclidean lattice isometry

$$T_W : \mathcal{W}(I_4) \longrightarrow \mathcal{W}(Q),$$

and then uniquely to an $\mathbb{L}^+$-linear isometry $W_4(I_4) \to W_4(Q)$. The restriction of $\bigwedge^2 B$ occurs in the list.

Proof. The hidden images $B_2\mathcal{B}_0$ form an orthonormal $\mathbb{Z}$-basis of $\mathcal{W}(Q)$. Two orthonormal integral bases of the same Euclidean lattice differ by a signed permutation. Hence B_2 sends every element of $\mathcal{B}_0$ to a signed element of $\mathcal{B}_1$.

There are 3! matchings between the three source and target components. Between two N-cycles there are $2N$ underlying dihedral graph isomorphisms. Once the sign of one image vertex is chosen, the intertwining equation determines all remaining signs; consistency holds because both edge-sign products are -1. Thus there are at most $4N$ signed intertwiners per component and at most $6(4N)^3$ in total. Since $S = \mathbb{Z}[s]$, commuting with s is equivalent to S-linearity. The hidden map commutes with s and therefore appears in the enumeration. □

Each candidate T_W extends $\mathbb{L}^+$-linearly from the lattice to the fixed space. Choose $\vartheta \in \mathbb{L}^\times$ with $\bar{\vartheta} = -\vartheta$. By Lemma 2, every $\xi \in \bigwedge_{\mathbb{L}}^2 \mathbb{L}^4$ has a unique expression $\xi = w_0 + \vartheta w_1$ with $w_i \in W_4(I_4)$. Define

$$T(w_0 + \vartheta w_1) = T_W(w_0) + \vartheta T_W(w_1). \tag{46}$$

This is the unique $\mathbb{L}$-linear extension of T_W. It is then submitted to Proposition 17; invalid extensions are rejected by Equation (43).

4.5 Direct and half-rank reductions

Theorem 20 (Direct rank-four reduction). *For every $1 \leq \gamma < \sqrt{2}$, a verified $C \in \mathrm{SL}_4(R)$ satisfying $C^*QC = I_4$ is recovered using at most $3N$ γ-SVP calls in ranks at most $3N$, followed by polynomially many exact computations.*

Proof. By Propositions 14 and 15, the target Hodge lattice is hypercubic. Apply Lemma 18 to recover an orthonormal basis $\mathcal{B}_1$. Enumerate the maps of Proposition 19, extend each by Equation (46), and invoke Proposition 17. The hidden branch $T = \bigwedge^2 B$ occurs in the list and returns a solution; every accepted branch passes the original exact equations. □

Theorem 21 (Rank-four reduction with half-rank oracles). *For every*

$$1 \leq \gamma < \sqrt{2 - \frac{2}{3N}}, \tag{47}$$

a verified $C \in \mathrm{SL}_4(R)$ satisfying $C^*QC = I_4$ is recovered with polynomially many γ-SVP calls, every one of rank at most

$$\boxed{\frac{3N}{2} + 1.} \tag{48}$$

Proof. Apply the hypercubic basis-recovery theorem Theorem 4 to $\mathcal{W}(Q)$ with $m = 3N$. Since N is even, $\lfloor m/2 \rfloor + 1 = 3N/2 + 1$. The signed-cycle completion and the four-line decoder are unchanged from Theorem 20. □

5 Exact public computation

This section records how the auxiliary lattices and the decoding data are obtained from the standard exact input. All computations use the coefficient basis $1, \zeta, \ldots, \zeta^{N-1}$ of R and reduce to rational linear algebra, integer kernels, and ideal arithmetic in the fixed cyclotomic field. The required normal-form operations are polynomial time in the standard exact encoding [2,4].

Algorithm 2 Rank-four recovery in the standard orbit

Require: $Q \in M_4(R)$ satisfying the promise Equation (1)
Ensure: A verified $C \in \mathrm{SL}_4(R)$ with $C^*QC = I_4$
Compute the Hodge lattice $\mathcal{W}(Q)$ from Equations (31) and (33).
Recover an orthonormal basis $\mathcal{B}_1$ using Theorem 4.
Compute multiplication by $s = \zeta + \zeta^{-1}$ in $\mathcal{B}_1$.
for each signed-cycle intertwiner T_W from Proposition 19 **do**
    Extend T_W to T by Equation (46).
    Run the four-line decoder of Proposition 17 on T.
    return the first verified transition.
end for
return FAIL if no branch passes.

Rank two. Write an unknown matrix $X \in M_2(R)$ in the coefficient basis. The conditions

$$X^*Q = QX, \qquad \operatorname{tr} X = 0$$

are rational homogeneous linear equations in its integer coefficients. Their integer kernel is exactly $\mathcal{S}(Q)$. Evaluating Equation (15) on an integer-kernel basis gives its exact Gram matrix. The same computation without the trace equation gives $\mathcal{H}(Q)$; applying the public map $A \mapsto 2A - \operatorname{tr}(A)I_2$ and taking Hermite normal form gives $\mathcal{T}(Q)$. If $t_1, \ldots, t_{3d}$ is a basis of $\mathcal{T}(Q)$ with Gram matrix G, then

$$t_i^\vee = \sum_j (G^{-1})_{ij} t_j$$

is the dual basis, and $2t_1^\vee, \ldots, 2t_{3d}^\vee$ is a basis of $\Gamma(Q)$. Under the promise, this is the integral isometric image of $\Gamma(I_2)$; its computed Gram matrix is therefore integral, as required by Theorem 3.

A primal shortest vector of $\Gamma(Q)$, or twice a dual shortest vector, is an explicit matrix $A \in M_2(R)$. The identities $A^2 = I_2$, $A^*QA = Q$, and $\operatorname{tr} A = 0$ are checked exactly before invoking the instance-preserving factor-recovery algorithm of Theorem 9. Its output factor is normalized to determinant one by the explicit unitary correction in that proof, and the resulting transition is finally checked against $C^*QC = I_2$.

Rank four. Compute $Q_2 = C_2(Q)$. Writing $\mathbf{x} \in R^6$ in the coefficient basis, the fixed-point equation

$$\mathbf{x} = \Omega \overline{Q_2}\, \bar{\mathbf{x}}$$

is a rational homogeneous linear system in $6N$ integer coefficients. Its integer kernel is exactly $\mathcal{W}(Q)$, and Equation (34) gives the exact Gram matrix. Multiplication by $s = \zeta + \zeta^{-1}$ on this kernel basis is obtained from the public coefficient action of R.

For a candidate extension T, each space $T(U(e_i))$ is the column span of three explicitly known bivectors. The equations $x \wedge u = 0$ for these generators are ordinary linear equations over $\mathbb{L}$ and yield the line ℓ_i in Equation (40). Intersecting

ℓ_i with R^4 gives $\mathfrak{b}_i v_i$; the ideal $\mathfrak{b}_i$ and norm value $1/q_i$ are the exact input to the ith Lenstra–Silverberg call. The final tests in Equation (43) are direct coefficient, determinant, compound, and Hermitian matrix equalities. Thus neither reduction requires a hidden basis, a numerical embedding, or a heuristic recognition step.

6 Conclusion

For the standard power-of-two cyclotomic orbit, the symmetric square in rank two and the exterior square in rank four give complete linear descents rather than dimension counts alone. In rank two, the public carrier has rank $3N/2$, and its transported shortest shell consists of nontrivial Q-orthogonal involutions. The existing instance-preserving rank-two Module-LIP-to-Module-LAP reduction then recovers an equivalent factor for the same public instance. In rank four, the Hodge fixed lattice has rank $3N$; hypercubic basis recovery, the public action of $\mathbb{Z}[\zeta + \zeta^{-1}]$, and four wedge annihilators recover the exterior-square transition and then the four columns of the original matrix.

The resulting reductions are exact and self-contained within the standard orbit: every accepted candidate is checked against R^r, determinant one, and the original equation $C^*QC = I_r$. Invoking the existing Bambury–Nguyen algorithms only at the shortest-vector stage yields largest oracle ranks

$$\frac{3N}{4} + 1 \quad \text{in rank two,} \qquad \frac{3N}{2} + 1 \quad \text{in rank four.}$$

These are oracle-dimension statements; no claim about a concrete SVP implementation or practical security level is made here.

References

1. Bambury, H., Nguyen, P.Q.: Improved provable reduction of NTRU and hypercubic lattices. In: Joye, M., Leander, G. (eds.) Post-Quantum Cryptography – PQCrypto 2024, Part I. LNCS, vol. 14771, pp. 343–370. Springer (2024). `https://doi.org/10.1007/978-3-031-62743-9_12`
2. Cohen, H.: A Course in Computational Algebraic Number Theory. Graduate Texts in Mathematics, vol. 138. Springer (1993). `https://doi.org/10.1007/978-3-662-02945-9`
3. Ducas, L., Postlethwaite, E.W., Pulles, L.N., van Woerden, W.: HAWK: Module LIP makes lattice signatures fast, compact and simple. In: Advances in Cryptology – ASIACRYPT 2022, Part IV. LNCS, vol. 13794, pp. 65–94. Springer (2022). `https://doi.org/10.1007/978-3-031-22972-5_3`
4. Kannan, R., Bachem, A.: Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix. SIAM Journal on Computing 8(4), 499–507 (1979). `https://doi.org/10.1137/0208040`
5. Lenstra, H.W., Silverberg, A.: Testing isomorphism of lattices over CM-orders. SIAM Journal on Computing 48(4), 1300–1334 (2019). `https://doi.org/10.1137/17M115390X`

6. Luo, H., Jiang, K., Jin, R., Pan, Y., Wang, A.: Commitment schemes based on Module-LIP. Cryptology ePrint Archive, Paper 2025/431 (2025). https://eprint.iacr.org/2025/431
7. Luo, H., Jiang, K., Pan, Y., Wang, A.: Cryptanalysis of rank-2 Module-LIP with symplectic automorphisms. In: Advances in Cryptology – ASIACRYPT 2024. LNCS, vol. 15487, pp. 359–385. Springer (2024). https://doi.org/10.1007/978-981-96-0894-2_12