COMPLEXITY OF MUTATION-ACYCLICITY AND MATRIX MUTATION EQUIVALENCE

QIYUE TANG

Abstract. Mutation-acyclicity of integral skew-symmetrizable matrices belongs to NP. Mutation equivalence belongs to promise-NP when at least one input is mutation-acyclic. For pairs of mutation-cyclic matrices, mutation equivalence is RE-complete and strongly NP-hard. Both results hold even for connected skew-symmetric matrices.

hexagon:2611.00180v2

1. Introduction

Fomin and Zelevinsky introduced cluster algebras to study total positivity and canonical bases [1]. Their construction starts from a seed: a set of algebraically independent generators, called a cluster, together with an integral exchange matrix. An exchange replaces one generator by a rational expression determined by the other generators and updates the matrix by mutation. Iterating these operations produces the cluster variables and the associated cluster algebra. The matrix mutation class records the combinatorial part of this construction. In the skew-symmetric case it is a class of quivers, which are directed multigraphs with no loops or directed two-cycles.

Fomin and Zelevinsky [7] classified the cluster algebras with finitely many cluster variables by the finite crystallographic root systems. Each such algebra has an acyclic seed of Dynkin type. An exchange matrix is called *mutation-acyclic* if some sequence of mutations makes its directed support acyclic, and *mutation-cyclic* otherwise. Mutation-acyclicity also includes representation-infinite hereditary types, so it asks for more than finite-type recognition. For finite type, Dias and Castonguay [8] gave a polynomial recognition algorithm using the quasi-Cartan criterion.

Representation theory explains why recognizing an acyclic representative is useful. A finite acyclic quiver has a finite-dimensional hereditary path algebra. Buan, Marsh, Reineke, Reiten, and Todorov [24] introduced its cluster category and related cluster-tilting objects to tilting modules over hereditary algebras. Caldero and Keller [9] established the correspondence between clusters and cluster-tilting objects in the acyclic skew-symmetric case. Their description of the acyclic seeds implies that two mutation-equivalent acyclic quivers are connected by source and sink mutations. For valued quivers, Hubery's support-tilting complex [21] and Rupel's exchange matrices [22] supply the hereditary species formulation used here.

The recognition problem first becomes nontrivial on three vertices. Assem, Blais, Brüstle, and Samson [10] gave a descent algorithm and canonical representatives for mutation-cyclic classes, up to relabeling and transposition. Beineke, Brüstle, and Hille [30] expressed the distinction using the Markov constant: a cyclic triangle with positive integral arrow multiplicities a, b, c is mutation-cyclic precisely when

$$a, b, c \geq 2, \qquad a^2 + b^2 + c^2 - abc \leq 4.$$

Date: November 7, 2026.

2020 *Mathematics Subject Classification.* Primary 13F60; Secondary 68Q17, 03D25.

Key words and phrases. exchange matrix, mutation-acyclicity, mutation equivalence, polynomial certificate, undecidability.

Seven [11] obtained minimal representatives for the diagrams of skew-symmetrizable matrices on three vertices, and Akagi [12] gave the corresponding Markov-constant criterion. These results use the arithmetic of a three-vertex orbit to obtain a terminating descent.

Finiteness of the entire matrix class gives another decision procedure. Felikson, Shapiro, and Tumarkin [13] classified the connected skew-symmetric quivers of finite mutation type: besides the surface families and the two-vertex cases there are eleven exceptional classes. When a mutation class is finite, a search recording every visited matrix terminates. Infinite classes require a different stopping argument. Full subquivers provide useful obstructions: Buan, Marsh, and Reiten [31, Corollary 5.3] proved that every full subquiver of a mutation-acyclic quiver is mutation-acyclic. This is the obstruction used by the guard triangles in our reductions.

The general two-input question was recorded by Fomin, Williams, and Zelevinsky [2, Problems 2.6.14 and 2.8.2], first for quivers and then for skew-symmetrizable matrices. Fomin, Igusa, and Lee [14, Remark 2.3] explained the algorithmic issue: enumerating mutation words recognizes equivalent pairs, but gives no stopping rule for inequivalent pairs. Their universal-quiver theorem realizes every quiver of a prescribed order as a full subquiver of some member of one larger mutation class. Fomin's 2022 OPAC lecture [15, slide 7] again posed the decision problem with all vertices mutable.

The geometry of the mutation graph also obstructs a direct extension of three-vertex descent. Fomin and Neville [16] constructed arbitrarily long mutation cycles at every fixed order at least four. Their cycles cannot be paved by short mutation cycles; they also show that a steepest-descent rule need not find a shortest mutation path. To obtain a complexity upper bound for recognition, one needs either an effective search bound or a short certificate whose verification avoids writing a mutation path.

Armstrong-Williams, Hirst, Jackson, and Lee [3] obtained an effective search bound in order four. For maximum initial arrow multiplicity $h \geq 1$, their decision procedure searches $O(h^{24})$ states. Combined with the acyclic-seed comparison theorem, it also decides equivalence when one of the four-vertex inputs is mutation-acyclic. Their alternating affine families show that, in every order at least four, mutation-acyclicity cannot be expressed by a fixed finite Boolean combination of polynomial sign conditions in the arrow multiplicities. The height search is pseudopolynomial: binary input length measures the logarithms of the multiplicities. The qualitative general decidability result for integral skew-symmetrizable matrices was stated in [4]. Theorems 1.1 and 1.2 give binary certificate bounds and fixed-order numerical bounds.

For lower bounds, Soukup [5] studied targets specified by properties of their arrows. He proved NP-hardness for reaching an ice quiver with a prescribed arrow multiplicity, and strong NP-hardness for removing all arrows between frozen vertices. Here a frozen vertex is one at which mutation is prohibited. His discussion of undecidability separates these hardness results from the existence of a decision algorithm for two specified quivers. The all-mutable two-matrix undecidability construction and a strong NP-hardness reduction appeared in [6]. We include their proofs and show that their outputs can be required to be connected and mutation-cyclic.

1.1. **Inputs and complexity.** An *exchange matrix* is an integral $n \times n$ matrix $B = (b_{ij})$, $n \geq 0$, for which a diagonal matrix $\Delta = \operatorname{diag}(f_1, \ldots, f_n)$ with positive integral entries satisfies $f_i b_{ij} = -f_j b_{ji}$. The matrix is called *skew-symmetrizable*; it is *skew-symmetric* when $B^T = -B$. Its directed support has an arrow $i \to j$ when $b_{ij} < 0$. We call it *acyclic* when this support has no directed cycle. This property is a graph test in polynomial time. Our recognition problem concerns existence of an acyclic representative under mutation.

For an integer x put $[x]_+ = \max(x, 0)$. For $1 \le k \le n$, mutation at k is the matrix $\mu_k(B)$ defined by

$$(\mu_k(B))_{ij} = \begin{cases} -b_{ij}, & i = k \text{ or } j = k, \\ b_{ij} + [b_{ik}]_+[b_{kj}]_+ - [-b_{ik}]_+[-b_{kj}]_+, & i, j \neq k. \end{cases} \tag{1}$$

Mutation is an involution and preserves Δ. Two matrices of the same order are *mutation equivalent*, denoted $B \sim B'$, if a finite mutation sequence followed by a simultaneous permutation of rows and columns takes B to B'. We include the empty sequence. Relabeling mutation indices moves every intermediate permutation to the end. The empty exchange matrix is acyclic.

Entries are encoded in binary, with their signs; the order and all entries form the input. An assertion of polynomial time always refers to the total binary length unless unary encoding is specified. A language belongs to NP if membership is equivalent to the existence of a polynomial-length certificate accepted by a deterministic polynomial-time verifier. For a promise problem the verifier must accept precisely the positive instances among inputs satisfying the promise. We write P for deterministic polynomial time, coNP for complements of languages in NP, and E for deterministic time $2^{O(L)}$, where L is the binary input length. A set is *recursively enumerable* (RE, or semidecidable) if an algorithm halts and accepts exactly its members. Completeness for RE is under computable many-one reductions. Strong NP-hardness below is established by a polynomial-time reduction with matrix entries bounded by a polynomial in the order, so that the reduction remains polynomial for unary entries.

Theorem 1.1. *Mutation-acyclicity of integral skew-symmetrizable exchange matrices belongs to* NP*. For $n \ge 1$, put $h = \max(2, \max_{i,j} |b_{ij}|)$. There are effective universal polynomials P, Q such that a positive instance has an integral Euler/real-root certificate with entries of absolute value at most $(2nh)^{P(n)}$, and certificate enumeration decides the problem in time at most $(2nh)^{Q(n)}$.*

Theorem 1.2. *Mutation equivalence of integral skew-symmetrizable matrices, under the promise that at least one input is mutation-acyclic, belongs to promise-*NP*. The ordinary language of equivalent pairs whose two members are mutation-acyclic belongs to* NP*. For each fixed order there is a pseudopolynomial decision algorithm on the promised domain.*

Theorem 1.3. *For pairs of connected mutation-cyclic integral skew-symmetric matrices, mutation equivalence is* RE*-complete under promise-preserving computable many-one reductions. The reduction has a fixed order and a fixed source matrix. There is also a separate strong* NP*-hardness reduction with variable order N and entries of magnitude $O(N^3)$. These statements hold as well for the larger class of skew-symmetrizable inputs.*

The first two theorems use actual real-root bases and their Euler pairings as certificates. Their size is bounded after a common Coxeter translation: the full Euler matrix moments control products of the two exterior spectral coordinates, and a positive chamber bounds the root coordinates. An exact polynomial-time real-root recognition algorithm authenticates the columns. The last theorem encodes membership in a fixed integral reflection subgroup. A doubled Cartan construction controls arbitrary mutation words, and radical-vector arguments replace frozen vertices by mutable guards. The proofs of Theorems 1.1 and 1.2 are given in Sections 6 and 7; Propositions 8.6 and 9.7 prove Theorem 1.3.

Corollary 1.4. *The complement of mutation-acyclicity belongs to* coNP*. At every fixed order, mutation-acyclicity and the promised problem of Theorem 1.2 belong to* P *for unary entries and to* E *for binary entries. For the promise in Theorem 1.3, positive instances are semidecidable and negative instances are not semidecidable; the problem is undecidable and lies outside promise-*P*, promise-*NP*, and promise-*coNP*.*

Proof. The complement assertion follows from Theorem 1.1. For fixed order, the running-time bounds in Theorems 1.1 and 1.2 are polynomial in the maximum absolute input entry. They are therefore polynomial in unary length and exponential with linear exponent in binary length. Enumerating mutation words and final permutations semidecides equivalence. If the negative promised instances were also semidecidable, dovetailing the two procedures would decide the promised problem, contradicting Theorem 1.3. Every promise-NP or promise-coNP problem is decidable on its promised domain by finite certificate enumeration. □

2. Recognizing real roots in polynomial time

Let $n \geq 1$ and let $A = (a_{ij})_{1 \leq i,j \leq n}$ be an integral symmetrizable generalized Cartan matrix: $a_{ii} = 2$, $a_{ij} \leq 0$ for $i \neq j$, $a_{ij} = 0$ if and only if $a_{ji} = 0$, and some positive diagonal matrix makes A symmetric. The simple reflection s_i acts on column vectors in $\mathbb{Z}^n$ by

$$s_i x = x - (Ax)_i e_i,$$

where e_i is the ith standard basis vector. The Weyl group $W(A)$ is the group generated by these integral matrices. Its real roots are the vectors we_i, with $w \in W(A)$ and $1 \leq i \leq n$. Vector inequalities are coordinatewise. A positive real root is a nonzero real root with nonnegative coordinates; $x > 0$ means that every coordinate is strictly positive. All matrix entries and vector coordinates in this section are encoded in binary. A set subscript on a matrix denotes its principal submatrix on that index set, and a set subscript on a vector denotes the corresponding subvector. For a matrix M, the symbols $M_{i,*}$ and M_{*i} denote its ith row and column, respectively.

Theorem 2.1. *Given an integral symmetrizable generalized Cartan matrix A and $x \in \mathbb{Z}^n$, one can decide in deterministic polynomial time, in their total binary length, whether x is a real root of A.*

Ordinary root descent need not take polynomially many steps in binary coordinates. We compress its long affine portions. The compression uses the particular translation lattices of affine Weyl groups; the lattices and the arithmetic operations needed below are described explicitly.

2.1. **Height and a uniform spectral gap.** On each connected component, compute a primitive positive integral symmetrizer $d = (d_1, \ldots, d_n)$, so that $d_i a_{ij} = d_j a_{ji}$. Propagation of these rational ratios on a spanning tree, clearing denominators, and division by their common divisor compute d in polynomial time and with polynomial bit length. Inconsistent ratios reject a purported symmetrizable input. Set

$$q(x) = \tfrac{1}{2} x^T \operatorname{diag}(d) A x, \qquad w_i = \lfloor \sqrt{d_i} \rfloor, \qquad H(x) = \sum_i w_i x_i \quad (x \geq 0).$$

Integer square roots can be computed in polynomial time. In particular,

$$\tfrac{1}{2}\sqrt{d_i} \leq w_i \leq \sqrt{d_i}, \qquad H(x) \leq \sum_i \sqrt{d_i} x_i \leq 2H(x). \tag{2}$$

The invariant q is preserved by every s_i.

We use the standard positivity and descent properties of real roots [17, Chapters 3 and 5]: every real root has coordinates of one sign; a positive real root other than e_i remains positive after any reflection s_i with $(Ax)_i > 0$. Such a reflection decreases H by the positive integer $w_i(Ax)_i$. A positive real root has $q(x) = d_j > 0$ for some j, so at least one such pairing is positive. Consequently ordinary descent terminates at a simple root.

For estimates only, write

$$S = \operatorname{diag}(\sqrt{d_i}), \qquad \overline{D} = SAS^{-1} = 2I - \overline{M}, \qquad y = Sx.$$

The symmetric matrix $\overline{M}$ has zero diagonal and entries $\overline{M}_{ij} = \sqrt{a_{ij}a_{ji}} \geq 0$ for $i \neq j$. All algorithmic comparisons below use integers or rational numbers. Write $\rho(M)$ for the spectral

radius of a square matrix M, the maximum absolute value of its complex eigenvalues. For nonnegative M, the Perron–Frobenius theorem makes $\rho(M)$ a nonnegative eigenvalue.

Lemma 2.2. *Put $\gamma = 1/1024$. Every connected matrix $\overline{M}$ of the preceding form with Perron eigenvalue greater than 2 has Perron eigenvalue at least $2+\gamma$.*

Proof. An edge with $a_{ij}a_{ji} \geq 5$ gives a two-vertex principal matrix of Perron eigenvalue at least $\sqrt{5} > 2+\gamma$. Otherwise the edge weights belong to $\{1, \sqrt{2}, \sqrt{3}, 2\}$, the canonical Coxeter weights for labels $3, 4, 6, \infty$. For these matrices McMullen's classification of the 38 minimal hyperbolic Coxeter diagrams gives [18, Theorem 1.5 and Section 6]

$$\rho(\overline{M})^2 \geq \lambda + \lambda^{-1} + 2,$$

where λ is Lehmer's number. Indeed every indefinite diagram dominates a minimal hyperbolic one, the Perron eigenvalue is monotone under this domination, and the bicolored eigenvalue is related to the adjacency eigenvalue by $\beta + \beta^{-1} + 2 = \rho^2$. For a rational verification of the needed bound, Lehmer's polynomial

$$P(z) = z^{10} + z^9 - z^7 - z^6 - z^5 - z^4 - z^3 + z + 1$$

has $P(7/6) = -13481789/60466176 < 0$, whereas $P(z) > 0$ for large z. Its largest real root therefore exceeds $7/6$. Since $z + z^{-1}$ is increasing for $z > 1$, the preceding lower bound exceeds $169/42 > (2+\gamma)^2$. □

We also use the finite and affine crystallographic diagram classification [17, Chapters 4 and 6]. In a connected positive-semidefinite normalized diagram, the matrix is either positive definite or affine of corank one. An affine normalized null vector can be scaled so that all its coordinates lie in $[1, 12]$. Here is sufficient data to verify the uniform constant. The simply laced affine marks have maximum 6 and minimum 1. For the remaining untwisted models one may take the following right null labels δ and symmetrizers d; the normalized labels are $\sqrt{d_i}\delta_i$:

type	δ	d	$\max_i \sqrt{d_i}\delta_i / \min_i \sqrt{d_i}\delta_i$
$B_m^{(1)}$	$(1, 1, 2, \ldots, 2)$	$(2, 2, \ldots, 2, 1)$	2
$C_m^{(1)}$	$(1, 2, \ldots, 2, 1)$	$(2, 1, \ldots, 1, 2)$	$\sqrt{2}$
$F_4^{(1)}$	$(1, 2, 3, 4, 2)$	$(2, 2, 2, 1, 1)$	3
$G_2^{(1)}$	$(1, 3, 2)$	$(3, 1, 3)$	2

The twisted presentations have the same normalized diagrams, up to permutation. Rank-two affine diagrams have normalized off-diagonal entry -2 and equal normalized null labels. Every finite normalized diagram is obtained by deleting an extending vertex from one of these affine diagrams.

Lemma 2.3. *If F is a positive-definite connected normalized Cartan matrix of order $m \leq n$, then*

$$0 \leq (F^{-1})_{ij} \leq 144n, \qquad \sum_i (F^{-1})_{ij} \leq 144n^2.$$

If D_J is an affine normalized Cartan matrix with null labels $1 \leq \delta_i \leq 12$, then

$$z^T D_J z = \sum_{\{i,j\} \text{ an edge}} (\overline{M}_{ij}\delta_i\delta_j) \left(\frac{z_i}{\delta_i} - \frac{z_j}{\delta_j} \right)^2. \tag{3}$$

Proof. Scale an affine extension of F by its null labels. Its Cartan matrix becomes a graph Laplacian with conductance $\overline{M}_{ij}\delta_i\delta_j \geq 1$ on each edge. Deleting the extending vertex gives the grounded Laplacian $L = \operatorname{diag}(\delta_i) F \operatorname{diag}(\delta_i)$. The maximum principle gives $L^{-1} \geq 0$. Its ith diagonal entry is the effective resistance to the grounded vertex: the unit flow on a simple path of at most m edges has energy at most m, so $(L^{-1})_{ii} \leq m$. The Cauchy–Schwarz inequality for the positive definite matrix L^{-1} bounds every entry by m. Multiplication by

the null labels gives the claimed bounds for F^{-1}. Expanding the Laplacian quadratic form gives (3). □

2.2. **A large affine block when ordinary descent is slow.** Fix the rational constants

$$t = \frac{1}{100000n^4}, \qquad \varepsilon_0 = \frac{\gamma t^2}{16}. \tag{4}$$

The following lemma gives a block that can be found by integral tests.

Lemma 2.4. *Let $x \geq 0$ be nonzero with $q(x) > 0$, and put $H = H(x)$. If*

$$\max_i w_i(Ax)_i < \varepsilon_0 H,$$

form $U = \{i : w_i x_i \geq tH\}$. All components of A_U are finite or affine. An affine component J maximizing $\sum_{j\in J} w_j x_j$ satisfies, with $\overline{H} = \sum_i \sqrt{d_i}x_i$,

$$\sum_{j\in J} w_j x_j \geq \frac{H}{2n}, \qquad \sqrt{d_j}x_j \geq \frac{\overline{H}}{64n^2} \quad (j \in J).$$

Every vertex $i \notin J$ adjacent to J satisfies $\sqrt{d_i}x_i < 2t\overline{H}$.

Proof. For analysis put $v = Sx/\overline{H}$, $p = \overline{D}v$, and $\varepsilon = \max_i p_i$. Then $\sum_i v_i = 1$, $v^T\overline{D}v > 0$, and (2) gives $0 < \varepsilon < 2\varepsilon_0$. Since every $p_i \leq \varepsilon$ and the total sum $\sum_i v_i p_i$ is nonnegative, every partial sum is at least $-\varepsilon$; in particular $v_i p_i \geq -\varepsilon$. On U we have $v_i \geq t/2$; off U we have $v_i < 2t$.

Suppose a component of $\overline{M}_U$ has Perron eigenvalue $\rho > 2$, and normalize its positive Perron vector u to coordinate sum 1. By Lemma 2.2,

$$u^T p \leq u^T \overline{D}_U v_U = (2-\rho)u^T v_U \leq -\gamma t/2.$$

But $p_i \geq -2\varepsilon/t$ on U, a contradiction to $\varepsilon < \gamma t^2/4$. Thus $\overline{D}_U$ is positive semidefinite.

Write

$$C_{\mathrm{cross}} = \sum_{i\in U, j\notin U} \overline{M}_{ij} v_i v_j, \qquad c_i = \sum_{j\notin U} \overline{M}_{ij} v_j \quad (i \in U).$$

Expanding the outside pairings gives

$$\sum_{j\notin U} v_j p_j = 2\sum_{j\notin U} v_j^2 - \sum_{j,k\notin U} \overline{M}_{jk} v_j v_k - C_{\mathrm{cross}}.$$

The left side is at least $-\varepsilon$, while the double sum is nonnegative. Since $v_j < 2t$ off U,

$$C_{\mathrm{cross}} \leq 2\sum_{j\notin U} v_j^2 + \varepsilon \leq 8nt^2 + \varepsilon.$$

Also $C_{\mathrm{cross}} = \sum_{i\in U} v_i c_i$ and $v_i \geq t/2$ on U, so $\sum_{i\in U} c_i \leq 2C_{\mathrm{cross}}/t$. On a finite component F, one has $\overline{D}_F v_F = p_F + c_F$. Lemma 2.3, its nonnegative inverse, and $p_i \leq \varepsilon$ show that the total mass in all finite components is at most

$$144n^2(n\varepsilon + 2C_{\mathrm{cross}}/t) \leq 2736n^3 t.$$

The mass outside U is at most $2nt$. Their sum is less than $1/4$. By (2), their fraction of integral height is less than $1/2$. Hence some affine component has integral-height mass at least $H/(2n)$, and the maximizing J has normalized true mass at least $1/(4n)$.

Its internal affine energy obeys

$$E_J = v_J^T \overline{D}_J v_J = \sum_{j\in J} v_j p_j + \sum_{j\in J, i\notin U} \overline{M}_{ji} v_j v_i \leq \varepsilon + C_{\mathrm{cross}} \leq 10nt^2.$$

Choose affine null labels in $[1, 12]$. Along a path of at most n edges, (3) and Cauchy–Schwarz give

$$|v_i/\delta_i - v_j/\delta_j| \leq \sqrt{nE_J} \leq 4nt.$$

The δ-weighted average of these ratios is at least $1/(48n^2)$. Therefore $v_j \geq 1/(48n^2) - 4nt \geq 1/(64n^2)$. An adjacent outside vertex cannot belong to another component of U, so it lies outside U and has coordinate less than $2t$. □

Lemma 2.5. *Let x be a positive real root and let A_J be affine, with $x_j > 0$ for every $j \in J$. It is impossible that $(Ax)_j \leq 0$ for all $j \in J$ and*

$$2\sqrt{d_i}x_i < \min_{j\in J} \sqrt{d_j}x_j$$

for every outside vertex i adjacent to J.

Proof. Follow ordinary positive-pairing descent to a simple root. Until the first reflection in J or one of its neighbors, all these coordinates are unchanged, and all pairings in J remain nonpositive. For a neighbor i, in normalized coordinates the pairing is at most

$$2y_i - \sum_{j\in J} \overline{M}_{ij} y_j < 0,$$

because every nonzero $\overline{M}_{ij}$ is at least 1. No first reflection at any such vertex is possible. The positive coordinates on the affine diagram J therefore remain throughout descent. Since an affine diagram has at least two vertices, its endpoint cannot be simple. □

2.3. **Computing an affine fold.** Let J be an affine component, keeping all coordinates outside J fixed. Its type and an extending vertex can be found in polynomial time by the finite and affine diagram patterns: paths and cycles of type A, paths with the indicated end labels in types B, C, the two branching patterns of type D, and the finitely many exceptional patterns. Enumerate the possible endpoints and branching vertices, then traverse the paths between them and compare their lengths and edge labels with these patterns. Cycles are matched by choosing a starting vertex and direction; the exceptional patterns have bounded order. These enumerations are polynomial in $|J|$. Checking all vertices as candidates for the extending vertex also takes polynomial time. Exact rational null-space computation supplies the primitive positive right null vector δ_J.

Choose an untwisted affine model A^{mod} with the same Coxeter labels as A_J. This includes the rank-two model $\left(\begin{smallmatrix} 2 & -2 \\ -2 & 2 \end{smallmatrix}\right)$. There is a positive rational diagonal matrix $R = \mathrm{diag}(R_j)_{j\in J}$ with

$$A_J = R^{-1} A^{\mathrm{mod}} R. \tag{5}$$

Indeed, equality of the edge products determines each successive ratio R_i/R_j from the corresponding entries of the two Cartan matrices. Affine diagrams other than the unit A cycle are trees; on that cycle all entries are -1. Thus edge propagation constructs R, whose entries have polynomial bit length. The identity can be checked exactly. It also covers twisted Cartan orientations. The same diagonal matrix transports the actual root action: for $i \in J$,

$$R(I - e_i(A_J)_{i,*})R^{-1} = I - e_i(A^{\mathrm{mod}})_{i,*}.$$

In particular the orbit of e_s in the original coordinates is sent to the orbit of $R_s e_s$ in the model coordinates.

Number the model extending vertex by 0. Write $(\ ,\ )$ for the Euclidean inner product in the finite root space and $\|u\|^2 = (u,u)$. Let α_i for $i \neq 0$ be its finite simple roots, and let θ be the highest finite root. For a finite root β, put $\beta^\vee = 2\beta/(\beta,\beta)$ and let $s_\beta(u) = u - (u, \beta^\vee)\beta$ be its reflection. Write $d_i^{\mathrm{mod}} = (\alpha_i, \alpha_i)/2$. The model right null vector has $\delta_0^{\mathrm{mod}} = 1$ and $\theta = \sum_{i\neq 0} \delta_i^{\mathrm{mod}} \alpha_i$. Put $\lambda_i = d_i^{\mathrm{mod}} \delta_i^{\mathrm{mod}}$, including $d_0^{\mathrm{mod}} = (\theta,\theta)/2$. The vector λ is the model's left null vector. For the actual pairings, set

$$p = (Ax)_J, \qquad p^{\mathrm{mod}} = Rp, \qquad C = -\lambda^T p^{\mathrm{mod}} \geq 0. \tag{6}$$

The inequality follows from $x \geq 0$, the null-vector identity on J, and the nonpositive outside attachments. Moreover $C > 0$ precisely when a positive outside coordinate is adjacent to J.

Assume $C > 0$. Define a point r in the finite Euclidean space by

$$(r, \alpha_i^\vee) = -p_i^{\mathrm{mod}}/C \quad (i \neq 0). \tag{7}$$

These equations are solved by rational linear algebra in the usual classical coordinate models, or in a fixed root basis for exceptional types. Under s_i, the pairing vector transforms as

$$p^{\mathrm{mod}} \longmapsto p^{\mathrm{mod}} - (A^{\mathrm{mod}})_{*i} p_i^{\mathrm{mod}}.$$

Thus the finite generators act on r by ordinary finite reflections. For the extending generator, the identity

$$(r, \theta) = 1 + d_0^{\mathrm{mod}} p_0^{\mathrm{mod}}/C$$

shows that it acts by reflection in $(r, \theta) = 1$. The closed fundamental alcove is

$$\mathcal{A} = \{r : (r, \alpha_i^\vee) \geq 0 \ (i \neq 0), \ (r, \theta) \leq 1\}.$$

Each orbit has exactly one point in $\mathcal{A}$. The affine Weyl group is $Q^\vee \rtimes W_0$, where W_0 is the finite Weyl group generated by the s_i with $i \neq 0$, and $Q^\vee$ is the lattice spanned over $\mathbb{Z}$ by its finite coroots [19, Section 2.2]. In particular $s_0 s_\theta$ is translation by $\theta^\vee$.

For specificity, the classical coordinate models below use an orthonormal coordinate basis $f_1, f_2, \ldots$, distinct from the input root-coordinate vectors e_i:

type	finite simple roots	θ
A_m	$f_i - f_{i+1}$ $(1 \leq i \leq m)$	$f_1 - f_{m+1}$
B_m	$f_i - f_{i+1}$ $(i < m)$, f_m	$f_1 + f_2$
C_m	$f_i - f_{i+1}$ $(i < m)$, $2f_m$	$2f_1$
D_m	$f_i - f_{i+1}$ $(i < m)$, $f_{m-1} + f_m$	$f_1 + f_2$

For A_m the ambient space is the coordinate-sum-zero subspace. The usual ranges $m \geq 2$ in B, C and $m \geq 4$ in D suffice; the rank-one case is A_1. The coroot lattices are respectively the integral sum-zero lattice, the even-sum integer lattice, $\mathbb{Z}^m$, and the even-sum integer lattice. For the exceptional models, one may work in the simple-root basis with rational Gram matrix $\mathrm{diag}(d^{\mathrm{mod}})A^{\mathrm{mod}}$ on the finite nodes. All dimensions and matrices in those exceptional computations are fixed.

Lemma 2.6. *For the affine block J at positive level C in (6), in polynomial binary time one can construct and evaluate an actual element of the subgroup generated by s_j, $j \in J$, which takes p^{mod} to nonpositive coordinates. Its full action includes all outside attachments of J.*

Proof. Find a closest coroot-lattice point z to r, then move $u = r - z$ to the finite dominant chamber by a finite Weyl element. The comparison with $z + \alpha^\vee$ gives

$$(u, \alpha^\vee) \leq \|\alpha^\vee\|^2/2$$

for every finite root α. Taking $\alpha = \theta$ gives $(u, \theta) \leq 1$, so the resulting dominant point belongs to $\mathcal{A}$.

The closest-point calculation for these particular lattices is as follows. For A_m, use coordinates in $\mathbb{R}^{m+1}$ with sum zero; the lattice is the integral sum-zero lattice. Write $\{u\} = u - \lfloor u \rfloor$ for the fractional part of a real number u. Take the fractional parts of r_i. Their sum k is an integer. Subtract 1 from the k largest fractional parts, and sort the resulting coordinates decreasingly. This minimizes squared distance, since changing a floor to a ceiling has extra cost $1 - 2\{r_i\}$, and imposes exactly the required sum. The sorted residual satisfies $u_1 \geq \cdots \geq u_{m+1} \geq u_1 - 1$.

For D_m the lattice is $\{z \in \mathbb{Z}^m : \sum z_i \text{ is even}\}$. Round independently to nearest integers, resolving ties by a fixed rule. If their sum is odd, change one rounded coordinate by one unit toward its residual, choosing a coordinate with largest absolute residual. For a zero residual either direction is allowed. The additional cost $1 - 2|r_i - z_i|$ is minimal. Any parity correction has at least this cost, so the result is closest. An even signed permutation places the residual in $u_1 \geq \cdots \geq u_{m-1} \geq |u_m|$; the root comparison gives $u_1 + u_2 \leq 1$. For B_m,

the coroot lattice is this same even-sum lattice, while arbitrary signed permutations give the finite chamber. For C_m, the coroot lattice is $\mathbb{Z}^m$; independent nearest rounding and arbitrary signed permutations apply, and the highest-root inequality is $\max_i |u_i| \le 1/2$. These finite factors have simple-generator words of length $O(m^2)$.

For E_6, E_7, E_8, F_4, G_2, the dimension is at most 8. Round the coordinates of r in a fixed coroot basis, obtaining z_0. Its residual lies in a fixed bounded parallelepiped, of radius R_0. A closest point differs from z_0 by a lattice vector of length at most $2R_0$. Enumerate this fixed finite set, compare exact rational squared distances, and then enumerate the fixed finite Weyl group to find a dominant representative. All constants depend only on these five fixed root systems.

It remains to evaluate the element just computed on the full input space. Conjugates of $s_0 s_\theta$ by W_0 translate by the orbit of $\theta^\vee$. In the classical systems they give a lattice basis using words of length $O(m)$: adjacent differences for A, adjacent differences and the last sum for D and B, and coordinate vectors for C. This also follows by moving the indicated roots along their standard coordinate chains. The exceptional words have fixed length. Rational linear algebra expresses z in this basis with integer coefficients of polynomial bit length. The translation therefore uses at most m binary powers, followed by the polynomial-length finite factor.

Modulo the model null vector, the finite simple roots are α_i and the extending simple root is $-\theta$. The quotient action of an affine Weyl element is its finite linear part. A translation therefore acts trivially on this quotient and fixes the null vector. By diagonal conjugacy the same holds in the original affine root span. Thus $T - I$ has image in $\mathbb{Q}\delta_J$ and annihilates δ_J, giving $(T - I)^2 = 0$ on that span. On the full space it fixes outside coordinates and has the block form

$$T - I = \begin{pmatrix} N & B \\ 0 & 0 \end{pmatrix}, \qquad N^2 = 0.$$

Consequently $(T - I)^3 = 0$ on the full space, and, for every signed integer k,

$$T^k = I + k(T - I) + \frac{k(k-1)}{2}(T - I)^2. \tag{8}$$

Compute the base words with the actual full matrices s_j, then use (8). There are polynomially many matrix factors, their base words have polynomial length, and all exponents have polynomial bit length. This evaluates the full action in polynomial bit complexity. □

Lemma 2.7. *For a positive real root with $C > 0$, the endpoint of Lemma 2.6 is coordinatewise no larger than x. For a block J chosen in Lemma 2.4, its integral height decreases by at least $H(x)/(256n^2)$.*

Proof. Repeatedly reflect at a positive actual pairing in J. Each step stays positive and lowers integral height, so the process terminates. It cannot terminate at a simple root in J, since some positive outside coordinate stays fixed. Its terminal J-pairings are nonpositive. This legal descent is coordinatewise decreasing. Closed-alcove uniqueness gives the same p^{mod} endpoint as the explicit fold. The two full vector endpoints have the same outside coordinates, so their difference is a scalar $s\delta_J$. For the invariant q,

$$q(x + s\delta_J) - q(x) = s\delta_J^T \operatorname{diag}(d_J)(Ax)_J.$$

The coefficient is strictly negative when $C > 0$. Equality of the two invariants forces $s = 0$. Thus the endpoints agree.

At this endpoint all J-pairings are nonpositive. If every normalized J-coordinate still exceeded $4t\overline{H}$, all these coordinates would be positive, and every adjacent outside coordinate, fixed during the fold, would be less than half their minimum. This satisfies every hypothesis of Lemma 2.5, a contradiction. Some coordinate drops from at least $\overline{H}/(64n^2)$ to at most $4t\overline{H}$. Its contribution to H drops by at least half this difference, which is at least $H/(256n^2)$ by (4). □

2.4. **The affine terminal test.** First note that the nonzero support of a positive real root is connected. Indeed, at a zero coordinate the pairing is nonpositive, so a legal positive-pairing descent never introduces a new support vertex. If the support had two components, each reflection would act separately inside one of them. Its restriction there is invertible and cannot annihilate an entire nonzero component. Descent could not reach a simple root.

If $C = 0$ in (6), every neighbor of J outside J has coordinate zero. A positive coordinate elsewhere would make the support disconnected. Hence the only case requiring a terminal test has x supported entirely in J. The following orbit test treats twisted as well as untwisted Cartan matrices.

A positive real root supported on J is a real root of A_J. Indeed, the equality

$$2q(x) = \sum_{j \in J} d_j x_j (Ax)_j > 0$$

supplies a positive pairing at a vertex in J. Ordinary descent at such vertices stays positive and never adds support outside J, so it reaches a simple root using only reflections in J. Reversing this word proves the assertion.

Lemma 2.8. *For an integral affine Cartan matrix A_J and an integral vector x supported on J, real-root membership can be decided in polynomial binary time.*

Proof. Use the model identification (5), with extending vertex 0, and let $\delta = \delta_J$ be the primitive integral right null vector in the original coordinates. On $\mathbb{Q}^J$ put

$$\pi(v) = v - (v_0/\delta_0)\delta.$$

Let W_0 be generated by the remaining simple reflections. For every original simple root e_s, enumerate the W_0-orbit of $\pi(e_s)$ by breadth-first search, retaining an actual vector $y = we_s$ for each projected vector. This enumeration has polynomial size. Under rational diagonal conjugacy these projected vectors are scalar multiples of finite roots; the classical finite root systems have $O(|J|^2)$ roots, and the exceptional ones have bounded size. Moreover W_0 fixes coordinate 0, so a projected vector uniquely determines the retained actual vector for each fixed s.

Construct a polynomial generating set $T_1, \dots, T_b$ for the translation subgroup by conjugating $s_0 s_\theta$ over the finite orbit of $\theta^\vee$. Its size is $O(|J|^2)$ in the classical types and bounded in the exceptional ones. Breadth-first search on this finite root orbit supplies polynomial-length representatives. These translations generate: the subgroup they generate is W_0-stable, and together with W_0 contains both s_0 and all finite simple generators.

For every retained y satisfying $\pi(y) = \pi(x)$, compute $x - y = k\delta$. Check directly that k is an integer. Since each T_i acts trivially modulo the null line, write

$$T_i y = y + c_i(y)\delta, \qquad c_i(y) \in \mathbb{Z}.$$

The coefficients are integral because δ is primitive and $T_i y - y$ is integral. Translations commute and fix δ, so the shifts attainable from y are exactly

$$\gcd(c_1(y), \dots, c_b(y))\mathbb{Z}.$$

Accept if and only if the target k is divisible by this gcd for at least one y. A zero gcd permits only $k = 0$. The decomposition $W(A_J) = Q^\vee \rtimes W_0$ proves both directions of this test. All orbit vectors, matrices, gcds, and comparisons have polynomial bit length, and their number is polynomial. □

2.5. **The decision algorithm.**

Proof of Theorem 2.1. Reject zero and mixed-sign vectors. Change a global negative sign to positive; this preserves real-root membership because $-e_i = s_i e_i$. Reject disconnected nonzero support and $q(x) \le 0$. A connected support lies in one component of A, so work with that component. Accept if $x = e_i$ for some i.

At each subsequent iteration compute $H = H(x)$. If $\max_i w_i(Ax)_i \geq \varepsilon_0 H$, reflect at a maximizing index. Reject unless the endpoint is nonnegative, and otherwise continue. This step lowers height by at least $\varepsilon_0 H$. Otherwise find J by the integral tests in Lemma 2.4, and compute the rational level C. For $C = 0$, reject a positive coordinate outside J and otherwise apply Lemma 2.8. For $C > 0$, evaluate the fold of Lemma 2.6. Reject unless its endpoint is nonnegative and has height at most $H - H/(256n^2)$; otherwise continue. At the beginning of every iteration repeat the simple-vector and support tests.

All continuing transformations are actual Weyl elements. The only acceptances are simple roots and vectors accepted by the exact affine orbit test, so the algorithm is sound. Positivity and ordinary descent justify the first branch on a true root. Lemmas 2.4 and 2.7 justify the second branch and its height test. The support argument and Lemma 2.8 justify the zero-level branch. Thus no true root is rejected.

Every continuing iteration reduces height by a factor at most

$$1 - \min\{\varepsilon_0, 1/(256n^2)\}.$$

Since ε_0 is a fixed constant times n^{-8}, there are $O(n^8 \log(2 + H(x)))$ iterations. The initial symmetrizer and height have polynomial binary length. Retained nonnegative vectors have no coordinate exceeding the initial height. Exact inertia, type matching, rational linear algebra, sorting, finite root-orbit enumeration, and integer gcd computations all have polynomial complexity here. The compressed translations use the degree-two formula (8); their number and the bit lengths of their matrices and exponents are polynomial in the current input length. The total running time is therefore polynomial in the original binary length. □

3. Euler certificates for mutation-acyclicity

For an integral matrix M, write $H(M) = \max(\{1\} \cup \{|M_{ij}| : i, j\})$. Bounds of the form $(2nH)^{\mathrm{poly}(n)}$ mean that one effective polynomial, independent of the matrices and of n, bounds the exponent. Different occurrences may use different such polynomials.

A positive symmetrizer of an exchange matrix B is a diagonal integral matrix $\Delta = \mathrm{diag}(f_1, \ldots, f_n)$, with $f_i > 0$, such that ΔB is skew-symmetric. It is primitive on a connected component when the corresponding f_i have greatest common divisor one. A connected component here and below means a connected component of the graph with an edge ij when $b_{ij} \neq 0$.

The primitive symmetrizer is computable in polynomial time. To see also its height bound, let $h = \max(2, H(B))$ and fix a spanning tree in a connected component. For each vertex i, direct the tree edges toward i and multiply $|b_{uv}|$ over these directed edges $u \to v$. The resulting product is a positive integer at most h^{n-1}. The products for adjacent vertices have ratio $|b_{ji}|/|b_{ij}|$, so they satisfy the symmetrizer relations. Divide by their common greatest common divisor and check every non-tree edge. This gives $1 \leq f_i \leq h^{n-1}$. Isolated vertices have $f_i = 1$.

A *hereditary Euler matrix* will mean an upper triangular integral matrix E with positive diagonal $D_0 = \mathrm{diag}(d_1, \ldots, d_n)$, nonpositive off-diagonal entries, and $d_i \mid E_{ij}$ and $d_j \mid E_{ij}$ for all i, j. Thus

$$E = D_0(I - N), \qquad A = D_0^{-1}(E + E^T), \qquad C = -E^{-1}E^T, \tag{9}$$

where N is strictly upper triangular and nonnegative integral, $D_0^{-1}N^T D_0$ is integral, and A is a symmetrizable generalized Cartan matrix. We use its integral simple reflections and real roots as defined in Section 2; the bilinear Euler form on its lattice is $x^T E y$. The Coxeter product is $C = s_1 \cdots s_n$. Indeed, set $y = s_1 \cdots s_n x$, applying the reflections from right to left. At step i its defining reflection formula gives

$$d_i y_i + \sum_{j>i} E_{ij} y_j = -d_i x_i - \sum_{j<i} E_{ji} x_j.$$

Thus $Ey = -E^T x$. Direct multiplication also gives $C^T EC = E$.

We recall the categorical objects used to interpret the certificate. A species over a finite field k assigns to each vertex a finite extension k_i/k and to each arrow a finite-dimensional bimodule over its endpoint fields. Its tensor algebra over $\prod_i k_i$, formed from the arrow bimodules, is a finite-dimensional hereditary algebra when the orientation is acyclic. We denote this algebra by H and use finite-dimensional left H-modules. The ith coordinate of a module's dimension vector is its dimension over k_i at vertex i. Its Euler form is the difference of the k-dimensions of Hom and first Ext. The field degree $[k_i : k]$ is the corresponding diagonal Euler entry.

An exceptional module X over a hereditary algebra H has a division endomorphism ring and $\operatorname{Ext}^1_H(X, X) = 0$. An exceptional sequence $(X_1, \dots, X_n)$ has $\operatorname{Hom}_H(X_j, X_i) = \operatorname{Ext}^1_H(X_j, X_i) = 0$ for $j > i$. A tilting module is a direct sum of n distinct indecomposable modules with vanishing mutual first extension groups. Dimension vectors throughout are coordinates in the basis of simple modules; the diagonal Euler entries are the dimensions of their division endomorphism rings over the ground field. A module is sincere when its support contains every vertex. A support-tilting module is tilting over the full hereditary species on its support. Its support-tilting pair includes the missing vertices as formal negative projectives; mutation replaces one member of such a pair. The support-tilting complex has these pairs as maximal faces, with adjacency given by one replacement.

Write $D^b(H)$ for the bounded derived category of finite-dimensional H-modules and $[1]$ for its cohomological shift. The Grothendieck group identifies module classes with their dimension vectors and satisfies $[X[1]] = -[X]$. Let τ_D denote the derived Auslander–Reiten translation. The cluster category is the orbit category $\mathcal{C}_H = D^b(H)/(\tau_D^{-1}[1])$, equipped with its triangulated structure. Its translation induced by τ_D is denoted by τ_C. A cluster-tilting object is a direct sum of n pairwise nonisomorphic indecomposables with vanishing first extension groups. Mutation replaces one summand by its other complement. The two exchange triangles encode this replacement and its reverse; their middle-term multiplicities, with the endomorphism division-field degrees, determine the exchange matrix.

Proposition 3.1. *The following statements hold for hereditary species over finite fields.*

(1) *Every hereditary Euler matrix is realized by such a species. A complete integral basis of real roots whose Euler matrix is upper triangular is, up to the signs of its members, the dimension sequence of an actual exceptional sequence.*

(2) *The support-tilting complex is connected. Its exchange matrices mutate by ordinary matrix mutation. For a sincere tilting module with dimension matrix V, put $F = V^T EV$ and let Δ be the diagonal matrix of its endomorphism division-field degrees. Its exchange matrix in our convention is*

$$B = (F^{-T} - F^{-1})\Delta. \tag{10}$$

The multiset of these degrees is the multiset of the simple division-field degrees.

(3) *Every cluster-tilting object in the cluster category of a connected representation-infinite hereditary species becomes a tilting module after a common power of the cluster translation. In finite type it becomes a tilting module over a hereditary species obtained by source and sink tilts. These tilts preserve the underlying valued graph and its division degrees.*

Here is the finite-field realization in the first assertion. Fix $k = \mathbb{F}_q$ and choose all extension fields inside one algebraic closure of k. At vertex i take $k_i = \mathbb{F}_{q^{d_i}}$. For $i < j$ take a k_j–k_i-bimodule M_{ji} consisting of $-E_{ij}/\operatorname{lcm}(d_i, d_j)$ copies of $\mathbb{F}_{q^{\operatorname{lcm}(d_i,d_j)}}$, with the two actions given by the common subfield embeddings. This multiplicity is a nonnegative integer by the Euler divisibility conditions. Let H be the tensor algebra of these bimodules over $\prod_i k_i$. Writing ε_i for its vertex idempotents and S_i for its vertex simples, the length-one projective

resolution is

$$0 \longrightarrow \bigoplus_{j>i} H\varepsilon_j \otimes_{k_j} M_{ji} \longrightarrow H\varepsilon_i \longrightarrow S_i \longrightarrow 0.$$

The first map concatenates an arrow with a path. Increasing vertex indices make the tensor algebra finite dimensional and the displayed map injective. Applying $\mathrm{Hom}_H(-, S_j)$ gives $\dim_k \mathrm{End}_H(S_i) = d_i$ and $\dim_k \mathrm{Ext}^1_H(S_i, S_j) = -E_{ij}$ for $i < j$; the reverse extension spaces vanish. Thus its Euler matrix is E. This is the canonical finite-field species used in [22]. The realization by the single extension field of degree $-E_{ij}$ in [20, Lemma 4.1] gives the same bimodule, since it is a vector space over the degree-$\mathrm{lcm}(d_i, d_j)$ common subfield.

The lifting assertion is Proposition 4.6 of [20]; a complete exceptional lattice sequence is an integral basis with the stated Euler orthogonality, and requiring all its members to be real roots gives precisely their complete real exceptional sequences. The second assertion uses the connectivity and division-ring results of [21, Theorem 19, Proposition 24, Corollary 25] and the exchange construction of [22, §§3.2,6, Theorem 6.7, Proposition 7.1]. The third is the hereditary-category argument in [24, Propositions 1.5–1.7, Theorem 3.3]. These statements use the hereditary Artin setting over arbitrary ground fields. The matrix order in (10) follows from Euler duality. The left and right Euler duals of the ordered tilting basis have column matrices L, R characterized by $L^T F = \Delta$ and $FR = \Delta$. Thus $L = F^{-T}\Delta$ and $R = F^{-1}\Delta$, and our exchange matrix is $L - R$.

Triangle equivalences preserve the labeled exchange matrix: exchange triangles are carried to exchange triangles with the same middle-term multiplicities, and the endomorphism division rings are preserved. For an exceptional module the extra orbit morphisms in its cluster endomorphism algebra form a nilpotent ideal. Its radical quotient is therefore its original division endomorphism ring. This also shows that a common cluster translation preserves each labeled degree when the translated summands are again represented by modules.

Proposition 3.2. *An integral skew-symmetrizable exchange matrix B is mutation-acyclic if and only if, after a simultaneous permutation of its labels, there exist a hereditary Euler matrix E and an integral matrix V satisfying:*

1. *the diagonal entries of E have the same multiset as the entries of the componentwise primitive symmetrizer Δ of B;*
2. *$\det V = \pm 1$, and each column of V is a positive real root for $D_0^{-1}(E + E^T)$;*
3. *$F = V^T E V$ is upper triangular, is nonnegative above the diagonal, and has diagonal Δ;*
4. *identity (10) holds.*

Proof. Suppose first that these conditions hold. Realize E by a hereditary species using Proposition 3.1. The real-root basis is a complete real exceptional sequence, and hence lifts to exceptional modules $X_1, \dots, X_n$ with the given positive dimension vectors. For $i < j$, its forward Euler pairing is nonnegative and its reverse Hom and Ext groups vanish. We show that $\mathrm{Ext}^1_H(X_i, X_j) = 0$.

If $\mathrm{Hom}_H(X_i, X_j) = 0$, nonnegativity of the Euler pairing gives this immediately. Otherwise every nonzero map $X_i \to X_j$ is monic or epic by [23, Lemma 4.1], since the reverse first extension vanishes. In the monic case, apply $\mathrm{Hom}_H(-, X_j)$ to $0 \to X_i \to X_j \to Z \to 0$. The exact segment

$$0 = \mathrm{Ext}^1_H(X_j, X_j) \longrightarrow \mathrm{Ext}^1_H(X_i, X_j) \longrightarrow \mathrm{Ext}^2_H(Z, X_j) = 0$$

forces the asserted vanishing. In the epic case write $0 \to Z \to X_i \to X_j \to 0$ and apply $\mathrm{Hom}_H(X_i, -)$. Its exact segment is

$$0 = \mathrm{Ext}^1_H(X_i, X_i) \longrightarrow \mathrm{Ext}^1_H(X_i, X_j) \longrightarrow \mathrm{Ext}^2_H(X_i, Z) = 0.$$

Thus $\bigoplus_i X_i$ is tilting. Its Hom Euler matrix is F, and its exchange matrix is B by (10). Connectivity of the support-tilting complex shows that B is mutation-equivalent to an acyclic matrix.

Conversely, start with an acyclic representative of B and realize its Euler lattice by a hereditary species. Follow a mutation sequence to obtain a cluster-tilting object with exchange matrix B. By Proposition 3.1, after componentwise common cluster translations, and finite-type source and sink tilts when needed, it is an actual tilting module over a hereditary species. Order its summands in the Hom order of [23, Corollary 4.2, Lemma 4.3]. Their dimension columns form an integral basis, their Gram matrix F is upper triangular nonnegative, and (10) holds. The division-degree invariance gives the diagonal multiset condition. Choosing the initial species with primitive simple degrees gives exactly the primitive symmetrizer of each connected component. Reordering the summands only permutes the input labels. □

4. Bounds for hereditary representatives

The criterion becomes an NP certificate once E and V can be chosen with polynomial binary length. We first bound E using B alone.

Lemma 4.1. *Let the matrices E, V, F, Δ satisfy Proposition 3.2, with $n \geq 1$. Put $h = \max(2, H(B))$. Then*

$$H(E) + H(F) + H\big(D_0^{-1}(E + E^T)\big) \leq (2nh)^{\mathrm{poly}(n)}.$$

Proof. For $n = 1$ the assertion follows from the primitive degree $d_1 = 1$. Assume $n \geq 2$. Let $U = I - \mathrm{strictupper}(B)$, where the displayed operator keeps precisely the entries strictly above the diagonal and sets all other entries to zero. The upper triangular part of (10) gives

$$U = F^{-1}\Delta, \qquad F = \Delta U^{-1}. \tag{11}$$

The finite geometric expansion of U^{-1} gives, for $i < j$,

$$|(U^{-1})_{ij}| \leq h(1 + h)^{j-i-1}.$$

Indeed, a term corresponds to an increasing path from i to j; summing over the possible intermediate vertices gives the displayed binomial sum. The primitive symmetrizer bound therefore bounds F as required.

The trace $\mathrm{tr}(M)$ of a square matrix is the sum of its diagonal entries. Congruence gives $-F^{-1}F^T = V^{-1}CV$, so the same trace is computed from F. Using (11) yields

$$\mathrm{tr}\, C + n = \sum_{i<j} \frac{b_{ij}F_{ij}}{f_j} \leq \binom{n}{2} \frac{\max_i f_i}{\min_i f_i} h^2(1 + h)^{n-2}. \tag{12}$$

For this estimate only, put $\overline{E} = D_0^{-1/2} E D_0^{-1/2} = I - \overline{N}$. Then $\overline{N}$ is strictly upper triangular and nonnegative, and

$$\mathrm{tr}\, C + n = \mathrm{tr}\big((I - \overline{N})^{-1}\overline{N}^T\big) \geq \sum_{i<j} \overline{N}_{ij}^2.$$

Here expand the inverse as $I + \overline{N} + \cdots + \overline{N}^{n-1}$; all summands are nonnegative and the term $\mathrm{tr}(\overline{N}\overline{N}^T)$ is the displayed sum. Since $\overline{N}_{ij} = \sqrt{d_i/d_j}\, N_{ij}$ and the diagonal multisets agree,

$$|N_{ij}| \leq \frac{\max_i f_i}{\min_i f_i} \sqrt{\binom{n}{2}}\, h(1 + h)^{(n-2)/2}.$$

Now $1 \leq d_i \leq h^{n-1}$. Multiplication by d_i bounds E, and $A = D_0^{-1}(E + E^T)$ has the same asserted form of height bound. □

5. Simultaneous normalization of real-root bases

We bound a real-root basis by applying one common Coxeter translation. For a real vector x, put $\|x\|_\infty = \max_i |x_i|$, $\|x\|_1 = \sum_i |x_i|$, and $\|x\|_2 = (\sum_i x_i^2)^{1/2}$.

Theorem 5.1. *Let E be a connected hereditary Euler matrix, with A, C as in (9), and let V be an integral basis consisting of real roots of either sign. Put $F = V^T E V$ and $H = \max(2, H(E), H(F))$. If A is of infinite type, there is an integer k such that*

$$H(C^k V) \le (2nH)^{\mathrm{poly}(n)}.$$

In finite type every real root already has coordinates bounded by the same form of estimate.

5.1. **Algebraic estimates.** We will use the following elementary bound on a single algebraic number field. The coefficient height of a polynomial is the largest absolute value of its coefficients. Suppose ρ is an algebraic integer of degree $e \le n$, its minimal polynomial has coefficient height at most 2^b, and p, q are integral polynomials of degree at most d and coefficient height at most 2^a, where $d \ge 0$ and $a, b \ge 0$. If $p(\rho)q(\rho) \ne 0$, then

$$U^{-e} \le \left|\frac{p(\rho)}{q(\rho)}\right| \le U^e, \qquad U = (d+1)2^a(1+2^b)^d. \tag{13}$$

Every conjugate of ρ has absolute value at most $1 + 2^b$. The field norm of $p(\rho)$, the product of its values under the e embeddings of $\mathbb{Q}(\rho)$ into $\mathbb{C}$, is a nonzero integer. Each factor has absolute value at most U. Thus $U^{-(e-1)} \le |p(\rho)| \le U$, and division gives (13). Clearing rational denominators adds their binary lengths to a.

For an integral n by n matrix with height at most H, characteristic polynomials, cofactors, and products of at most $\mathrm{poly}(n)$ matrices of height at most $(2nH)^{\mathrm{poly}(n)}$ have degrees and coefficient logarithms bounded by $\mathrm{poly}(n)\log(2nH)$. The same holds for an irreducible factor of a characteristic polynomial: its roots are bounded by a Cauchy bound, and its coefficients are elementary symmetric functions of at most n such roots. Consequently (13) makes every nonzero rational function arising from these data, with the stated polynomial degree and coefficient bounds, lie between $(2nH)^{-\mathrm{poly}(n)}$ and $(2nH)^{\mathrm{poly}(n)}$ in absolute value. We apply this only to the explicit cofactors, pairings, projector coefficients and matrix products below.

5.2. **Exterior positivity in indefinite type.** Assume A is connected and neither finite nor affine. Ringel's theorem [25] gives a simple eigenvalue $\rho > 1$ of C, strictly larger in modulus than its other eigenvalues. Since $C^T E C = E$, its spectrum is reciprocal and ρ^{-1} is likewise simple and strictly dominant for C^{-1}. The matrix C is integral unimodular, and $C = s_1 \cdots s_n$.

For the positivity and spectral-gap arguments, conjugate by $D_0^{1/2}$:

$$\overline{E} = I - \overline{N}, \qquad \overline{C} = D_0^{1/2} C D_0^{-1/2}.$$

The matrix $\overline{N} + \overline{N}^T$ is symmetric with off-diagonal entries $\sqrt{a_{ij}a_{ji}}$. For $\lambda > 0$, the equation $\overline{C}z = \lambda z$ is equivalent to

$$\frac{\lambda\overline{N} + \overline{N}^T}{\lambda + 1} z = z.$$

Levinger's theorem [27] makes the Perron root of this nonnegative matrix nonincreasing for $\lambda \ge 1$. It is analytic there and is strictly decreasing: a constant interval would make it constant throughout, whereas its value at $\lambda = 1$ exceeds one and its limit at infinity is zero. Thus there is exactly one $\lambda > 1$ at which its Perron root equals one. Its positive eigenvector is a $\overline{C}$ eigenvector. The spectral radius theorem, and the fact that the matrix at $\lambda = \rho$ has eigenvalue one, identify this λ with ρ. Transpose symmetry gives the corresponding positive

eigenvector for ρ^{-1}. Returning to the original coordinates, choose strictly positive vectors r_+, r_- with

$$Cr_+ = \rho r_+, \quad Cr_- = \rho^{-1} r_-, \quad \sum_i (r_+)_i = \sum_i (r_-)_i = 1. \tag{14}$$

The same uniqueness argument shows that the only positive real eigenvalues with nonzero nonnegative eigenvectors are ρ and ρ^{-1}.

The normalized-adjacency gap of Lemma 2.2, obtained from [18, Theorem 1.5 and §6], gives $\rho(\overline{N} + \overline{N}^T) \geq 2 + g$ with $g = 1/1024$. Perron monotonicity gives

$$1 = \rho\left(\frac{\rho\overline{N} + \overline{N}^T}{\rho + 1}\right) \geq \frac{\rho(\overline{N} + \overline{N}^T)}{\rho + 1}, \qquad \rho \geq 1 + g.$$

Here $\rho(M)$ denotes the spectral radius of a matrix M; the unadorned ρ is the eigenvalue in (14). If $a = \max(1, \max_{i \neq j} |a_{ij}|)$, the edge attaining a has normalized weight at least $\sqrt{a}$. Hence $\rho \geq \sqrt{a} - 1$ as well. Together these inequalities imply

$$\log \rho \geq c(1 + \log a) \tag{15}$$

for an absolute $c > 0$.

Set $D = E + E^T$ and $\gamma = r_+^T D r_-$. This scalar is nonzero. Indeed, C preserves the nondegenerate Euler form, so a simple eigenspace pairs nontrivially only with its reciprocal eigenspace, and $D = E(I - C)$. It is negative. To check the sign after conjugation, put $u_\pm = D_0^{1/2} r_\pm$, and set $P = u_-^T \overline{N} u_+$, $Q = u_-^T \overline{N}^T u_+$, $R = u_-^T u_+ > 0$. The Perron equation gives $\rho P + Q = (\rho + 1)R$. The derivative of the Perron root at ρ is $(P - Q)/((\rho + 1)^2 R) \leq 0$ by Levinger's theorem. Therefore

$$\gamma = 2R - P - Q = \frac{(\rho - 1)(P - Q)}{\rho + 1} < 0.$$

Lemma 5.2. *If $0 \neq x \in \mathbb{Z}^n$ and $C^m x \geq 0$ for every integer $m \geq 0$, the ρ-eigenprojection of x is a strictly positive multiple of r_+. The analogous assertion for C^{-1} holds with r_-.*

Proof. Let W be the real span of the forward orbit, and let K be the closed convex cone generated by that orbit. It lies in the nonnegative orthant, satisfies $K \cap (-K) = \{0\}$, spans W, and is preserved by $C|_W$. Write r for the spectral radius of this restriction. Every forward iterate of x is a nonzero integral vector, so $r \geq 1$: if $r < 1$, all these iterates would tend to zero.

We construct a nonnegative eigenvector for r. Choose u in the interior of K relative to W, and write $y \leq_K z$ when $z - y \in K$. For $t > r$ the convergent resolvent series gives

$$v_t = (tI - C|_W)^{-1} u = \sum_{m \geq 0} t^{-m-1} C^m u \in K, \qquad v_t \geq_K u/t.$$

If v_t had a bounded subsequence as $t \downarrow r$, its limit v would satisfy $v \geq_K u/r$ and $Cv = rv - u$. In particular v is an interior point of K. Choose $0 < \varepsilon < r$ with $u \geq_K \varepsilon v$. Then $0 \leq_K Cv \leq_K (r - \varepsilon)v$. The norm

$$\|z\|_v = \inf\{b \geq 0 : -bv \leq_K z \leq_K bv\} \quad (z \in W)$$

is finite because v is interior and is a norm because K is closed and contains no line. Positivity of $C|_W$ now gives the induced operator-norm bound $\|C|_W\|_v \leq r - \varepsilon$, contradicting its spectral radius r. Thus $\|v_t\|_2 \to \infty$. A convergent subsequence of $v_t/\|v_t\|_2$ has a nonzero limit $z \in K$ and satisfies $Cz = rz$. Exterior positivity therefore identifies r with ρ. The ρ-projection of x cannot vanish, because x generates W cyclically. Finally $\rho^{-m} C^m x$ tends to that projection; it is nonnegative and nonzero, and hence is a strictly positive multiple of r_+. Apply the argument to the integral unimodular matrix C^{-1} for the second assertion. □

5.3. **A chamber bound for all real roots.** Let P_ρ and $P_{\rho^{-1}}$ be the linear projections onto the respective eigenspaces along the other generalized eigenspaces of C. For a vector x, define its exterior coefficients $a(x)$ and $b(x)$ by $P_\rho x = a(x)r_+$ and $P_{\rho^{-1}}x = b(x)r_-$. The remaining generalized eigenspaces contribute neither to these projections nor to pairing with r_+ or r_- under D.

Lemma 5.3. *There is a constant $K_E \le (2nH(E))^{\mathrm{poly}(n)}$ such that every real root x satisfies*

$$\|x\|_\infty \le K_E(|a(x)| + |b(x)|).$$

Proof. Define the small inversion roots

$$\beta_i = s_n \cdots s_{i+1} e_i, \qquad \zeta_i = s_1 \cdots s_{i-1} e_i \quad (1 \le i \le n).$$

Empty products are the identity. They have height at most $(2na)^{O(n)}$ and satisfy $C\beta_i = -\zeta_i$. Speyer's reduced-powers theorem [26, Theorem 1] says that every positive power of the Coxeter word is reduced, meaning that its length is minimal among words in the simple reflections for the same element. The inversion roots of an element are the positive real roots it sends to negative real roots. The inversion sets give the disjoint partition of all positive real roots into

$$\begin{gathered}\{C^{-j}\beta_i : j \ge 0,\ 1 \le i \le n\}, \qquad \{C^j\zeta_i : j \ge 0,\ 1 \le i \le n\},\\ \{x \text{ a real root} : x \ge 0,\ C^j x \ge 0 \text{ for all } j \in \mathbb{Z}\}.\end{gathered} \tag{16}$$

For every integer $m \ge 1$, the inversion roots of C^m are exactly $C^{-j}\beta_i$ with $0 \le j < m$, and those of C^{-m} are $C^j\zeta_i$ with $0 \le j < m$. They are distinct; disjointness of the two unions follows from length additivity for C^{2m}.

Lemma 5.2 gives $a(x), b(x) > 0$ in the third set, $a(\beta_i) < 0 < b(\beta_i)$, and $b(\zeta_i) < 0 < a(\zeta_i)$. For every integer $m \ge 0$, both $C^{-m}\beta_i$ and $-C^{m+1}\beta_i$ are positive real roots, which gives the two asserted signs for β_i.

All these coefficients for the small roots lie in the single field $\mathbb{Q}(\rho)$. The eigenvectors in (14) are normalized cofactor columns of $\rho I - C$ and $\rho^{-1}I - C$. Write $\chi_C(t) = \det(tI - C)$, let a prime denote differentiation in t, and let adj denote the transpose of the cofactor matrix. The projectors are

$$P_\rho = \frac{\mathrm{adj}(\rho I - C)}{\chi_C'(\rho)}, \qquad P_{\rho^{-1}} = \frac{\mathrm{adj}(\rho^{-1} I - C)}{\chi_C'(\rho^{-1})}.$$

Their denominators are nonzero by simplicity. Clearing powers of ρ makes their entries rational functions of polynomial degree and coefficient logarithms $\mathrm{poly}(n)\log(2na)$. The same holds for the exterior coefficients of the small roots and their nonzero ratios by (13).

For $1 \le \eta \le 2$, consider the covector

$$h_\eta(x) = -(r_+ + \eta r_-)^T Dx = -\gamma(b(x) + \eta a(x)).$$

It is positive on the third set in (16). On the first two sets a nonpositive value requires, respectively,

$$\rho^{2j} \le \frac{-\eta a(\beta_i)}{b(\beta_i)}, \qquad \rho^{2j} \le \frac{-b(\zeta_i)}{\eta a(\zeta_i)}.$$

The logarithms of the right sides are at most $\mathrm{poly}(n)\log(2na)$. By (15), every such index j is at most $\mathrm{poly}(n)\log(2n)$, hence at most a fixed polynomial in n. There are consequently only polynomially many positive roots on which any h_η can be nonpositive. Each of these roots forbids at most one value of η by a zero pairing. Choose η from a rational grid in $[1, 2]$ with one more point than the number of these roots. Its binary length is polynomial in n. Then $h = h_\eta$ has no zero real-root pairing and is negative on only polynomially many positive roots.

If $h(e_i) < 0$, replace h by $h \circ s_i$. A simple reflection permutes the positive roots other than e_i and negates e_i, so this decreases the number of negative positive-root pairings by

exactly one. Repeating terminates after polynomially many reflections, with a covector h' whose coordinates are all strictly positive. Write $h' = h \circ w^{-1}$ for the resulting Weyl element. Both w, w^{-1} have heights at most $(2na)^{\mathrm{poly}(n)}$. The coordinates of h' are linear combinations of those of h with these integer coefficients, so (13) gives $\min_i h'(e_i) \geq (2na)^{-\mathrm{poly}(n)}$.

For every real root x, the root wx has one sign. Therefore

$$\|wx\|_1 \leq \frac{|h'(wx)|}{\min_i h'(e_i)} = \frac{|h(x)|}{\min_i h'(e_i)}.$$

Multiplication by w^{-1} and the expression for $h(x)$ give the asserted bound. If D_0 has a common scalar factor, divide E by it for this argument; neither the root system nor C changes. Primitive symmetrizers have entries at most a^{n-1} by the spanning-tree construction, so all constants have the stated dependence on E. □

5.4. **Normalization of the basis.**

Proof of Theorem 5.1. Assume first that A is indefinite. Write $a_i = a(v_i)$ and $b_i = b(v_i)$ for the columns of V, and put $C_F = -F^{-1}F^T = V^{-1}CV$. For $0 \leq j < n$,

$$V^T E C^j V = F C_F^j. \tag{17}$$

Each matrix on the right has height at most $(2nH)^{\mathrm{poly}(n)}$: use determinant bounds for F^{-1} and at most $n-1$ multiplications. The projector P_ρ is a polynomial in C of degree less than n with coefficients in $\mathbb{Q}(\rho)$ of the bounds established above. For example it is the polynomial $\chi_C(t)/((t-\rho)\chi_C'(\rho))$ evaluated at C. Thus every entry of $V^T E P_\rho V$, computed from (17), has absolute value at most $(2nH)^{\mathrm{poly}(n)}$.

Euler invariance pairs an eigenspace only with its reciprocal generalized eigenspace. Consequently

$$(V^T E P_\rho V)_{ij} = \kappa b_i a_j, \qquad \kappa = r_-^T E r_+ \neq 0.$$

The same single-field norm bound controls κ^{-1}, giving $|b_i a_j| \leq (2nH)^{\mathrm{poly}(n)}$ for every i, j. Both maxima $\max_i |a_i|$ and $\max_i |b_i|$ are nonzero because V is a basis. Choose an integer k such that

$$1 \leq \frac{\max_i |a_i|}{\max_i |b_i|} \rho^{2k} < \rho^2.$$

Every exterior coefficient of a column of $C^k V$ has absolute value at most

$$\rho \sqrt{(\max_i |a_i|)(\max_i |b_i|)} \leq (2nH)^{\mathrm{poly}(n)}.$$

Lemma 5.3 bounds the entire columns, proving the indefinite case.

Assume $D = E + E^T$ is positive semidefinite of corank one, and let $\delta \in \mathbb{Z}^n_{>0}$ be its primitive null vector. Kernel minors bound $\|\delta\|_\infty$ by $(2nH)^{\mathrm{poly}(n)}$. Fix any coordinate 0; this label is local to the present proof. For a real root v, define

$$\beta = v - \frac{v_0}{\delta_0}\delta.$$

Its coordinate 0 is zero and its denominators divide δ_0. The principal matrix D_J deleting coordinate 0 is positive definite integral. Its determinant is at least one and its greatest eigenvalue is at most $nH(D)$, so its least eigenvalue is at least $(nH(D))^{-(n-2)}$. Since $v^T D v = 2d_i \leq 2 \max d_i$ for a real root,

$$\|\beta\|_2^2 \leq 2 \max_i d_i \, (nH(D))^{n-2}.$$

Thus the projections of all real roots lie in a rational grid of coordinate height, denominator, and cardinality at most $(2nH)^{\mathrm{poly}(n)}$.

The induced action of C on the quotient by $\mathbb{R}\delta$ is invertible. Each projected simple-root orbit lies in the finite grid, so it is periodic. If m is the least common multiple of their n periods, then $m \le (2nH)^{\mathrm{poly}(n)}$ and

$$C^m = I + \delta\ell^T, \qquad \ell \in \mathbb{Z}^n, \quad \ell^T\delta = 0. \tag{18}$$

Integrality of ℓ follows from primitivity of δ. It is nonzero because the Coxeter element has infinite order by the reduced-powers theorem and the root representation is faithful. The Euler invariance in (18) gives

$$\ell f^T - f\ell^T = 0, \qquad f = E^T\delta,$$

so $\ell = \kappa E^T\delta$ for a nonzero rational κ. These quantities have height at most $(2nH)^{\mathrm{poly}(n)}$. Indeed, at each step of a simple-root orbit its projection stays in the bounded grid. Writing its null coefficient as q_t gives $q_{t+1} - q_t = (C\beta_t)_0/\delta_0$, which is uniformly bounded by that same form of estimate. Hence $H(C^t)$ grows at most linearly in $1+t$ with a bounded prefactor, and in particular C^m has the claimed height.

Within this affine argument, write the columns of V as $v_i = \beta_i + k_i\delta$ and put $f_i = \delta^T E v_i = \delta^T E\beta_i \in \mathbb{Z}$. At least one f_a is nonzero, since V is a basis and $\delta^T E \neq 0$. Because $E\delta = -E^T\delta$,

$$F_{ij} = \beta_i^T E\beta_j + k_i f_j - k_j f_i. \tag{19}$$

A common power C^{ms} sends k_i to $k_i + s\kappa f_i$. Each step $\kappa f_i = \ell^T v_i$ is integral. Choose $s \in \mathbb{Z}$ so that $0 \le k_a + s\kappa f_a < |\kappa f_a|$. The projections β_i do not change. Applying (19) with the second index a now bounds every other null coefficient, since $|f_a| \ge 1$ and all the other terms have height at most $(2nH)^{\mathrm{poly}(n)}$. This proves the affine case.

Finally, when D is positive definite, its determinant is a positive integer and its least eigenvalue is at least $(nH(D))^{-(n-1)}$. The bound $v^T Dv \le 2\max d_i$ directly bounds every real-root coordinate by $(2nH)^{\mathrm{poly}(n)}$. □

6. Polynomial certificates

We restore module dimension vectors after normalization.

Lemma 6.1. *Suppose a tilting module over a connected representation-infinite hereditary species has dimension matrix V and Gram matrix F. There is a tilting module with the same labeled exchange matrix and division degrees whose dimension matrix has height at most*

$$(2n\max(2, H(E), H(F)))^{\mathrm{poly}(n)}.$$

Proof. Choose k from Theorem 5.1. The derived translation τ_D^k sends each summand to an indecomposable object with signed Grothendieck class $C^k v_i$. Let P_i and I_i be the projective cover and the injective hull of the vertex simple S_i, respectively. Their dimension matrices are $E^{-T}D_0$ and $E^{-1}D_0$: their Euler pairings with the simple basis are the corresponding diagonal pairings. Triangular elimination with the simple reflections identifies their columns with β_i and ζ_i, respectively. For example, to compute β_i, the reflections with indices $i+1, i+2, \dots, n$ successively solve the lower triangular system $E^T v = d_i e_i$; the reflections with indices $i-1, i-2, \dots, 1$ solve $Ev = d_i e_i$ and give ζ_i. The derived translation satisfies $\tau_D P_i = I_i[-1]$ and $\tau_D^{-1} I_i = P_i[1]$. Along a derived translation orbit, a module leaves the subcategory of modules only by crossing a projective or injective boundary. The nested inversion sets of the reduced Coxeter powers show that a real-root class changes sign at most once in either direction. A boundary crossing changes the cohomological shift by one and the class sign; a second crossing in the same direction would therefore contradict these inversion sets. It follows that, for $k \ge 0$, the translated summand is a module or a module shifted by $[-1]$; for $k \le 0$ it is a module or a module shifted by $[1]$. Its underlying dimension is $|C^k v_i|$ coordinatewise.

The cluster category is the orbit category $D^b(H)/(\tau_D^{-1}[1])$. Its fundamental domain consists of modules and shifted projectives $P_i[1]$. A representative $M[-1]$ can be replaced by $\tau_D^{-1}M$ by one application of the orbit functor; this is a module unless M is injective, when it is a shifted projective. A representative $M[1]$ with M nonprojective can be replaced by $\tau_D M$; if M is projective it is already in the domain. These replacements multiply the dimension bound by at most $\max(1, H(C), H(C^{-1}))$, up to the factor n for matrix multiplication.

To count the excluded shifts, Speyer's inversion rays give $C^{-j}\beta_i \geq 0$ and $C^j\zeta_i \geq 0$ for every $j \geq 0$, and these roots are nonzero. Consequently $\tau_D^{-j}P_i$ and $\tau_D^j I_i$ remain modules for all $j \geq 0$: a first further crossing would change their root sign, contrary to these inequalities. In the cluster category, $\tau_C^{-1}(P_i[1]) = P_i$ and $\tau_C(P_i[1]) = I_i$. Thus every nonzero translate of $P_i[1]$ is a module, so each τ_C-orbit contains at most one shifted projective. Among the $n+1$ common shifts τ_C^j, $0 \leq j \leq n$, each summand therefore excludes at most one choice. Some common j makes every summand a module. At each step the fundamental-domain action either applies C to a module dimension, changes P_i to $P_i[1]$, or changes $P_i[1]$ to I_i. At most n steps preserve the asserted height bound.

The resulting object is cluster-tilting with n distinct module summands. Since Ext^1_H embeds as a direct summand in cluster Ext^1, it is an actual tilting module. Common triangle equivalences preserve its labeled exchange triangles and division degrees as explained after Proposition 3.1. □

Proof of Theorem 1.1. Accept the empty matrix directly. For positive order the qualitative criterion supplies a tilting realization. Its initial E, F have height at most $(2nh)^{\mathrm{poly}(n)}$ by Lemma 4.1. In finite type the root coordinates are already bounded by Theorem 5.1. In infinite type use Lemma 6.1. For disconnected matrices apply the construction separately to each hereditary component; a real root is supported on a single component, and componentwise triangle equivalences preserve the block exchange matrix. Reordering each final tilting basis gives the desired positive certificate. Compute its Gram matrix from the final E, V; it satisfies the criterion for the unchanged exchange matrix.

For verification, compute the primitive input symmetrizer and guess the input permutation, a diagonal ordering for E, and the bounded integral matrices E, V. Check Euler divisibility, diagonal multiset, positivity, $\det V = \pm 1$, triangular nonnegative $F = V^T EV$, and (10) by exact arithmetic. Authenticate each column as a real root using the deterministic polynomial algorithm of Theorem 2.1. Every check has polynomial binary complexity. The criterion proves soundness, while the bounds just proved give a polynomial-length certificate for every positive instance.

There are $O(n^2)$ bounded entries to enumerate in E, V. Including the permutations gives at most $(2nh)^{\mathrm{poly}(n)}$ candidates. The verification cost has the same form, proving the deterministic bound. □

7. Equivalence with a mutation-acyclic input

We bound source and sink reversals between the hereditary matrices in the certificates.

Lemma 7.1. *If two acyclic orientations of a fixed weighted graph on n vertices are connected by source and sink reversals, then they are connected by at most $n(n-1)$ source reversals.*

Proof. A sink reversal at v can be replaced by source reversals at all other vertices in a topological order ending at v. Indeed, reversing every vertex once in a topological order reverses each edge twice and returns to the initial orientation. Omitting its final v therefore has the same endpoint as the sink reversal.

Consider a legal source-only word. On the two endpoints of an edge its letters alternate, beginning at the initial source of that edge. Hence the total counts of adjacent vertices differ by at most one. If every vertex occurs, mark its first occurrence. A repeated occurrence of i before the first occurrence of j can only happen when i, j are nonadjacent: otherwise edge

alternation requires the first j before the second i. Thus commuting nonadjacent letters moves all marked first occurrences to the front. Their order is a legal source order containing every vertex once, so it returns to the initial orientation and can be deleted. The remaining source word is legal from the same starting orientation and has the same endpoint.

Continue deleting complete rounds until some vertex has count zero. On a connected component, every other count is at most its distance from a zero-count vertex, hence at most $n-1$. Its residual length is at most $n(n-1)$. For disconnected graphs perform the deletions independently on the components and concatenate their residual words; the sum of the component bounds is at most $n(n-1)$. □

Proof of Theorem 1.2. For each input guess and verify a certificate (E_i, V_i) from Theorem 1.1. Write D_i for the diagonal matrix of E_i. The hereditary acyclic exchange matrix is

$$S(E_i) = D_i^{-1}(E_i - E_i^T).$$

It is integral by Euler divisibility. Its convention can also be checked using the projective dimension matrix $V_P = E_i^{-T} D_i$: its Gram is $D_i E_i^{-T} D_i$, and substitution in (10) gives $D_i^{-1}(E_i - E_i^T)$. Guess a relabeling between the two hereditary vertex sets and a source word of length at most $n(n-1)$ taking $S(E_1)$ to the relabeled $S(E_2)$. Check each source and its matrix mutation. Such mutations just reverse incident arrows and retain their valuations, so intermediate bit lengths stay bounded. An accepted certificate gives mutation paths from each B_i to its hereditary matrix and a bridge between them. This proves soundness for the ordinary language in Theorem 1.2.

For completeness, suppose the two inputs are equivalent and mutation-acyclic. Choose a common acyclic representative and realize it by one hereditary species. The mutation paths to B_1, B_2 give two cluster-tilting objects in this same category. For an infinite-type connected component, independent common translation powers make the two objects tilting modules over the same hereditary species. The trace bound of Lemma 4.1 bounds this common E using either input, and normalization is applied independently to the two bases while preserving E and their exchange matrices. Thus the two certificates have an empty hereditary bridge.

For a finite-type component, Proposition 3.1 makes the objects tilting modules after possibly different source and sink tilts. Their hereditary matrices have the same valued graph and division degrees and are connected by source and sink reversals. The finite root bound makes both certificates small, and Lemma 7.1 supplies a bridge of length at most $n(n-1)$. Apply these arguments componentwise and guess the corresponding relabeling. The total witness has polynomial binary length.

Under the stated promise, equivalent inputs are both mutation-acyclic, so the ordinary language gives the promised decision problem. Finally, the certificate bounds, permutations, and bounded source words can all be enumerated in time $(2nh)^{\mathrm{poly}(n)}$, where h is the larger input height. At each fixed order this is polynomial in the numerical height. □

8. Mutation equivalence of mutation-cyclic matrices

We reduce membership in a fixed reflection subgroup to matrix mutation equivalence, developing the construction of [6] with connected mutation-cyclic endpoints.

During a construction, a frozen vertex means a vertex at which mutation is temporarily forbidden; its row, column and entries between frozen vertices still undergo matrix mutation. All frozen vertices will subsequently be made mutable. We use the arrow convention of the introduction, so $b_{ij} < 0$ represents arrows from i to j.

8.1. **Radical vectors and frozen vertices.**

We use two elementary identities on the full square matrix, including its entries between frozen vertices. For $v \in \ker B$ define

$$S_k(B, v) = \sum_j [b_{kj}]_+ v_j = \sum_j [-b_{kj}]_+ v_j.$$

The equality follows from $Bv = 0$. Transport v under mutation at k by

$$v'_k = S_k(B, v) - v_k, \qquad v'_i = v_i \quad (i \neq k). \tag{20}$$

For any row q updated against the principal block, the row formula is

$$q'_k = -q_k, \qquad q'_j = q_j + [q_k]_+[b_{kj}]_+ - [-q_k]_+[-b_{kj}]_+ \quad (j \neq k).$$

Its pairing with the transported vector satisfies

$$\begin{aligned} q'v' &= -q_k(S_k - v_k) + \sum_{j\neq k} q_j v_j + ([q_k]_+ - [-q_k]_+)S_k \\ &= qv. \end{aligned} \tag{21}$$

Applying this calculation to each row of B other than row k gives zero. Row k is negated, and its pairing remains zero because $b_{kk} = 0$. Thus $(\mu_k B)v' = 0$. The same calculation proves (21) for a frozen row with arbitrary pairing qv. Both identities allow arbitrary signs of q and v.

Lemma 8.1. *Let B have even entries and let $v \geq 0$ lie in its kernel. Suppose initially $S_k(B, v) \geq 2v_k$ for every allowed mutation. Any allowed mutation word can be replaced, using involution cancellations and commutations at zero exchange entries, by a word with the same full matrix endpoint and transported vector, along which every coordinate of the transported vector is nondecreasing. The transported vector is nonnegative along every original word as well. If initially all coordinates are positive and all these inequalities are strict, every step of a nondecreasing representative is a strict increase.*

Proof. Call k a strict descent when $S_k < 2v_k$, and a nonincrease when $S_k \leq 2v_k$. At a nonnegative state a strict descent has $v_k > 0$. If $j \neq k$ is a nonincrease and $b_{kj} \neq 0$, put $c = |b_{kj}| \geq 2$. If $v_j = 0$, its zero budget contradicts adjacency to $v_k > 0$. Otherwise

$$S_k \geq cv_j, \qquad S_j \geq cv_k, \qquad S_kS_j \geq 4v_kv_j > S_kS_j,$$

a contradiction. Thus a strict descent commutes with every distinct nonincrease.

Take the first descending step of a word. All preceding states are nonnegative and all preceding steps are nondecreasing. At the state before the descent, the inverse of the preceding step is a nonincrease. If its index differs from the descending index, the two mutations commute by the preceding paragraph. Moving the descent left preserves its descending status. It must meet its previous occurrence and cancel, because reaching the initial state would contradict the initial inequalities. These are identities on the entire square matrix and on the transported vector. Each cancellation shortens the word by two. Iteration gives a nondecreasing representative. Applying this argument to each prefix proves nonnegativity at the original intermediate states, including a proposed first negative state.

Under the strict initial assumptions, no flat direction occurs along a nondecreasing path. Induct on its length. After a strict ascent its inverse is a strict descent; hence all adjacent directions are strict ascents by the first paragraph. At nonadjacent directions both the coordinate and its budget are unchanged. This proves the last assertion. □

The gcd of the entries of a labelled row is another useful invariant. Mutation at that label negates the row. Mutation at a different label changes every other entry by an integer multiple of the incident pivot entry. The old row gcd therefore divides the new one, and involutivity gives the reverse divisibility. Thus each labelled row gcd is invariant; a final permutation only permutes these gcds.

Lemma 8.2. *Let X, Y be even skew-symmetric matrices on the same finite labelled vertex set, with a common nonempty set of designated frozen vertices, called ports. Suppose each active endpoint row has a nonzero active-principal entry of magnitude at most H, and each port has a fixed entry of magnitude 2 to an active vertex and all its entries are even. Assume endpoint permutations are required to fix the ports. The equivalence problem with these ports*

frozen reduces to ordinary all-mutable equivalence, with endpoint permutations preserving the active set and fixing every port. If the old support is connected, the constructed supports are connected. Every constructed endpoint is mutation-cyclic.

Proof. Choose distinct odd primes $p_f > H$, one for each port f, and multiply its row and column by $d_f = 2p_f$. Leave active scales equal to one. Congruence scaling by a positive diagonal matrix commutes with active mutation: both products in each correction acquire precisely the product of the two endpoint scales. The port row gcd becomes $2d_f = 4p_f$.

Attach to each port two disjoint oriented triangles, with every internal cyclic edge of weight $G_f = 4p_f^2$. Join one vertex of the first triangle to f with entry $+G_f$, and one vertex of the second with entry $-G_f$, using the row directed from the triangle to f. Add no other incidences. Every vertex is now mutable. An unchanged nonzero active-principal entry bounds each active row gcd by H. The exact row gcds are therefore at most H on active vertices, $4p_f$ at port f, and $4p_f^2$ on its six guards. These classes are disjoint and distinguish the ports individually. Thus every endpoint permutation preserves the active set, fixes each port, and preserves its six guards as a set.

Give all old vertices coordinate zero and all guards coordinate one. The resulting vector v lies in the kernel: each triangle has zero row sum, and its two attachments cancel at the port. Its initial budgets are zero on active vertices and G_f on ports and guards. Lemma 8.1 applies. At a permuted target, sum the three kernel equations on the pulled-back vertices of either target triangle. Internal terms cancel, giving $\pm G_f v_f = 0$. Consequently every endpoint port coordinate is zero.

Normalize a successful word to be coordinatewise nondecreasing. Suppose it has a first port mutation, at f. Before this event no active vertex has a guard neighbor; each guard triangle is attached only to its port, whose coordinate is zero. The restriction of v to that triangle belongs to its current principal kernel. By (21), its weighted attachment charge remains $+G_f$ or $-G_f$. Nonnegativity therefore makes both budgets at f at least G_f. Its first mutation raises its zero coordinate to a positive value, contradicting its zero endpoint coordinate. Thus the normalized word has no port mutations.

Guard moves now commute with active moves and change no old–old entry. Delete them and restrict to the old vertices. Conversely, an allowed old word lifts unchanged, with all guard data fixed. This proves the equivalence.

All attachments are nonzero, so connectivity is preserved. An equal-weight oriented triangle of weight $G \geq 2$ is mutation-cyclic by the rank-three criterion of [30]: its three weights are at least two and its Markov constant $3G^2 - G^3$ is at most four. Indeed it is four at $G = 2$ and nonpositive for integral $G \geq 3$. A full subquiver of a mutation-acyclic quiver is mutation-acyclic by [31, Corollary 5.3]. Each endpoint contains such a full guard triangle and is therefore mutation-cyclic. □

8.2. **A fixed hard reflection subgroup.** Higman's embedding theorem states that every finitely generated recursively presented group embeds in a finitely presented group [28]. We use it to fix a group with a many-one RE-complete word problem.

Let $\mathcal{K}$ be an RE-complete subset of the positive integers and set

$$H = \langle a, b \mid [b, a^n b a^{-n}] = 1 \ (n \in \mathcal{K}) \rangle,$$

where $[g, h] = ghg^{-1}h^{-1}$. To identify its word problem, let R have generators b_i, $i \in \mathbb{Z}$, with commutation relations exactly when $|i - j| \in \mathcal{K}$. Translation of indices acts on R, and the corresponding semidirect product by $\mathbb{Z}$ is H. For $n \notin \mathcal{K}$, killing all b_i except b_0, b_n retracts R onto the free group on these two generators. Hence

$$n \in \mathcal{K} \quad \Longleftrightarrow \quad [b, a^n b a^{-n}] = 1 \text{ in } H.$$

Higman's theorem embeds this fixed H in a fixed finitely presented group G. Fixed words for the images of a, b give a computable many-one reduction. The word problem of G is RE by enumeration of products of conjugates of its relators, so it is RE-complete.

For each generator a_i of G, adjoin t_i with $t_i^2 = 1$ and $t_i a_i t_i = a_i^{-1}$. This is an amalgam with the generalized dihedral group of the actual cyclic subgroup $\langle a_i \rangle$; the amalgam normal-form theorem embeds the preceding group [29]. No knowledge of the order of a_i is needed. The resulting fixed finitely presented group has involution generators $t_i, t_i a_i$. Write it as

$$G_0 = \langle A \mid R_1, \dots, R_r \rangle,$$

including q^2 among the relators for each $q \in A$. Its word problem remains RE-complete. Free involution factors can be added so that A has at least two elements.

Let F be the free group on $A \cup \{x\}$. Define its automorphism θ by inverting each basis generator, preserving the written order of factors. Put

$$\mathcal{T} = (A \cup \{x\}) \cup \{R_j, qR_j : q \in A \cup \{x\},\ 1 \le j \le r\}.$$

Let N be the normal closure of $\theta(g)g$ for $g \in \mathcal{T}$, and let $\pi : F \to E = F/N$ be the quotient map. This normal subgroup is θ-invariant since $\theta(\theta(g)g) = g\theta(g)$ is a conjugate of $\theta(g)g$.

Lemma 8.3. *For a word u on A, put $w = [x, u]$. Then $u = 1$ in G_0 if and only if $w = 1$ in E. Moreover the subgroup $M \le F \times F$ generated by (g, g^{-1}), $g \in \mathcal{T}$, is*

$$M = \{(c, d) : \pi(d) = \theta(\pi(c))\}.$$

Proof. The relations for R_j and qR_j give $\theta(R_j) = R_j^{-1}$ and $q^{-1}R_j^{-1}qR_j = 1$. Thus all relators R_j are central in E. If $u = 1$ in G_0, its image in E is a product of central relators and is central, so $[x, u] = 1$. Conversely E maps onto $G_0 * \langle x \rangle$: inversion of a basis letter of A acts trivially on its image, since that image is an involution. All defining relations of E therefore hold in this free product. If $u \neq 1$ in G_0, free-product normal form gives $[x, u] \neq 1$ there.

The basis pairs generate the twisted diagonal $(c, \theta(c))$. Comparing $(g, \theta(g))$ and (g, g^{-1}) supplies $(1, (g\theta(g))^{-1})$. Conjugation by the twisted diagonal supplies $1 \times N$, because $\theta(F) = F$. Conversely every specified generator satisfies the displayed quotient condition. These inclusions prove the formula for M. □

Embed F faithfully in $\mathrm{SL}_2(\mathbb{Z})$, with all matrices congruent to the identity modulo 4 and without $-I$. The matrices

$$U = \begin{pmatrix} 1 & 4 \\ 0 & 1 \end{pmatrix}, \qquad V = \begin{pmatrix} 1 & 0 \\ 4 & 1 \end{pmatrix}$$

generate a free group: nonzero powers of U send $|z| < 1$ into $|z| > 1$ on the real projective line, and nonzero powers of V send $|z| > 1$ into $|z| < 1$. This ping-pong argument gives projective faithfulness and excludes $-I$. Enumerate the $d = |A| + 1$ free basis generators of F and send the ith one to $U^i V U^{-i}$, for $1 \le i \le d$. These conjugates freely generate: after collecting adjacent equal conjugates, a nonempty reduced word retains nonzero V-powers separated by nonzero U-powers.

Identify $\mathrm{Mat}_2(\mathbb{Z})$ with $\mathbb{Z}^4$ by the coordinate order (a, b, c, d) and define

$$J = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \\ 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}.$$

Let (g, h) act by $X \mapsto gXh^{-1}$, and let the factor interchange s act by $X \mapsto -\operatorname{adj}(X)$, where $\operatorname{adj}\begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} d & -b \\ -c & a \end{pmatrix}$. This gives a faithful representation of $(F \times F) \rtimes \langle s \rangle$. If an even element acts trivially, evaluation at I gives $g = h$, and evaluation at every X makes this common matrix scalar. Only I occurs among these scalars. Every even action has determinant one, while $-\operatorname{adj}$ has determinant -1 on the four-dimensional space, so no odd element acts trivially.

If v_g is the column of entries of g, then

$$r_g = (g, g^{-1})s : X \longmapsto -g\,\mathrm{adj}(X)g = X - \mathrm{tr}(g^{-1}X)g, \qquad r_g = I - v_g(v_g^t J), \quad v_g^t J v_g = 2.$$

Here tr is the sum of the two diagonal entries. For every two-by-two matrix Y, one has $Y + \mathrm{adj}(Y) = \mathrm{tr}(Y)I$. Apply this identity to $g^{-1}X$, use $\mathrm{adj}(g^{-1}X) = \mathrm{adj}(X)g$, and multiply on the left by g to obtain the displayed formula. Let K_0 be generated by r_1 and r_g, $g \in \mathcal{T}$. Factor interchange normalizes M, sending its specified generators to their inverses. Thus K_0 represents $M \rtimes \langle s \rangle$. For left multiplication $L_w : X \mapsto wX$, Lemma 8.3 gives

$$L_w \in K_0 \quad \Longleftrightarrow \quad (w, 1) \in M \quad \Longleftrightarrow \quad u = 1 \text{ in } G_0. \tag{22}$$

Enumerate $\{1\} \cup \mathcal{T}$ as $g_1, \ldots, g_m$, with $g_1 = 1$, and write $v_i = v_{g_i}$ and $r_i = r_{g_i}$. Form the fixed matrix $T_0 = (v_1, \ldots, v_m) \in \mathbb{Z}^{4\times m}$. Its columns span $\mathbb{Q}^4$. Indeed $d \geq 3$, and

$$U^i V U^{-i} = \begin{pmatrix} 1 + 16i & -64i^2 \\ 4 & 1 - 16i \end{pmatrix}.$$

The traceless parts for three distinct indices are independent by the Vandermonde determinant for $1, i, i^2$; I supplies the scalar direction. Put $A = T_0^t J T_0$. Since T_0 is onto $\mathbb{Q}^4$ and J is nonsingular, $\ker A = \ker T_0$, so A has rank four. Its diagonal is 2. Each entry equals $\mathrm{tr}(g_i^{-1}g_j)$ and is 2 modulo 16: if $I + 4Z$ has determinant one, then $\mathrm{tr}(Z) = -4 \det Z$. In particular all entries of A are nonzero and even.

8.3. **A four-frame Weyl orbit.** Choose a fixed odd integer $L > \max_{i\neq j} |A_{ij}|/2 + 2$. Define a symmetric generalized Cartan matrix D, indexed by $i+, i-$ for $1 \leq i \leq m$, with diagonal 2, with $D_{i+,i-} = 0$, and, for $i \neq j$, with

$$D_{i+,j+} = D_{i-,j-} = -(L - A_{ij}/2), \qquad D_{i+,j-} = D_{i-,j+} = -(L + A_{ij}/2).$$

Every nonzero off-diagonal magnitude is even and at least 2. Let s_k be its simple reflection, acting on columns by $I - e_k D_{k,*}$, where e_k is the corresponding coordinate column and $D_{k,*}$ is row k of D. Let W be the free product of m Klein-four groups, with generators named s_{i+}, s_{i-} in its ith factor. The displayed reflections commute within each inner pair and therefore define a representation of W. The next lemma proves faithfulness on the quotient by the radical $\ker D$.

Let ι interchange all inner signs, and let $\mathbf{1}$ denote the all-ones column of length m. On symmetric and antisymmetric coordinate vectors the operator blocks of D are

$$D_+ = (2 + 2L)I - 2L\mathbf{1}\mathbf{1}^t, \qquad D_- = A.$$

The Gram matrices in the sum and difference bases are $2D_+$ and $2A$. The first is nonsingular, with eigenvalues $2 + 2L$ and $2 - 2L(m - 1)$. After quotienting by the radical of D, the symmetric and antisymmetric subspaces are therefore nondegenerate, orthogonal, and of dimensions m and four, respectively. Call the latter subspace S.

Lemma 8.4. *The representation of W on the quotient by the radical of D is faithful, and*

$$\mathrm{Stab}_W(S) = \langle c_i = s_{i+}s_{i-} : 1 \leq i \leq m \rangle.$$

On S, the generator c_i acts as r_i under the identification induced by T_0.

Proof. In dual row coordinates a simple reflection has the formula

$$x'_k = -x_k, \qquad x'_j = x_j - D_{kj}x_k \quad (j \neq k). \tag{23}$$

The row $x_0 = -\sum_{k=1}^{2m} D_{k,*}$ belongs to the row space of D and has every coordinate $2L(m - 1) - 2 > 0$. Write a reduced free-product word in syllables, each a nonempty subset of one inner pair. After a syllable the fired coordinates are negative. If their total absolute value is R, every coordinate outside that pair is at least R, and the unfired mate, if any, is positive. This holds after the first syllable because every coupling to another pair is at least 2. For the next syllable the fired positive coordinates have total $T \geq R$. Each formerly negative coordinate becomes at least $-R + 2T \geq T$; every other coordinate outside the

new pair remains positive and at least T. Exactly the newly fired coordinates are negative. Induction shows that no nonempty reduced word fixes x_0.

Thus the action on the row space is faithful. The identity $Ds_k = s_k^t D$ identifies it with the primal quotient representation, proving faithfulness there too. If a group element preserves S, it preserves its orthogonal complement and commutes with ι on the quotient. Faithfulness makes it fixed by the sign-interchange automorphism of W. Unique free-product normal form then forces each syllable to be the product of the two generators in its pair. This proves the stabilizer formula.

Before quotienting, c_i sends $(z, -z)$ to

$$(z - e_i(Az)_i, -z + e_i(Az)_i).$$

Applying T_0 gives $y \mapsto y - v_i(v_i^t Jy) = r_i y$ on $\mathbb{Q}^4$, as required. □

For $z \in \mathbb{Q}^4$ define the row $f(z)$ by

$$f(z)_{i+} = z^t Jv_i, \qquad f(z)_{i-} = -z^t Jv_i.$$

It is injective, has integral entries for integral z, and its rational image is the antisymmetric part of the row space of D. The row action satisfies $f(z)c_i = f(r_i z)$. More generally, if a word h in the c_i has primal restriction R_h, its row action is $f(z)h = f(R_h^{-1} z)$.

The four source rows $f(e_1), \dots, f(e_4)$ and target rows $f(L_w e_1), \dots, f(L_w e_4)$ are W-equivalent exactly when $L_w \in K_0$. Indeed any element carrying these frames preserves their span, hence lies in the stabilizer from Lemma 8.4, and equality on the four rows forces its inverse restriction to equal L_w. Conversely an expression in the reflections lifts, in reverse order, to a word in the c_i. Together with (22), this is a fixed four-frame orbit problem with an RE-complete set of targets.

8.4. **Compiling the frame orbit into mutation.** For distinct indices s, t of D, put $a_{st} = -D_{st}$, and choose a skew-symmetric matrix E with $|E_{st}| = a_{st}$. Its zero entries are those of D off the diagonal. Replace each index s by an outer pair $s0, s1$ and define the principal block

$$B_{s\epsilon,t\eta} = (-1)^{\epsilon+\eta} E_{st} \quad (\epsilon, \eta \in \{0, 1\}). \tag{24}$$

There are $4m$ principal vertices. Its all-ones vector lies in the kernel, because the two outer columns cancel. Each budget is $\sum_t a_{st} = 2L(m-1) > 2$. For each s, add a frozen marker row with entry 4 at $s0$, entry -4 at $s1$, and zero elsewhere. Add four frozen payload rows by duplicating each coordinate of $4f(e_a)$, $1 \le a \le 4$, equally on its outer pair; here e_a is the ath coordinate column of $\mathbb{Q}^4$. The targets have the same principal block and markers, with payload rows obtained from $4f(L_w e_a)$. Port labels are fixed at this stage.

Lemma 8.5. *Allow an arbitrary final permutation of principal vertices. A principal mutation word carries the prescribed marker and payload rows to their targets if and only if the four source frames and four target frames are W-equivalent.*

Proof. Normalize a successful word using the all-ones principal radical and Lemma 8.1. Every remaining step is a strict ascent. By (21), the scalar product of a marker row with the transported vector remains zero. At the endpoint each pulled-back target marker thus forces the coordinates of its outer pair to agree. The two principal columns in that target pair are opposite, so their budgets agree as well.

Reverse the normalized word. Its first letter k is a strict descent; its target outer mate j is another strict descent, since its coordinate and budget agree with those of k. After undoing k, the direction j remains a strict descent because the two vertices are nonadjacent. Move backwards to the first occurrence of j. Every intervening backwards step is a descent. The critical-pair calculation in Lemma 8.1 makes it nonadjacent to j, permitting an actual zero-edge commutation. Interchanging these two letters preserves the descending status

of the letter crossed. If j did not occur, its descent would persist to the initial strict local minimum, which is impossible. Hence it can be commuted next to k.

Undo both outer mates. On the doubled principal matrix, this interchanges the two selected outer labels. It changes the sign of their marker, preserves all other marker rows, and preserves the supports of every marker. Balanced rows, meaning rows with equal values on each outer pair, remain balanced. The two transported coordinates stay equal. This leaves a shorter strictly descending reverse path of the same kind. Induction decomposes the normalized word into complete outer-pair moves.

To control the initially unspecified final permutation, begin this peeling with the pair partition specified by the pulled-back target markers. Each individual marker support persists through peeling to the source. Source marker supports are precisely the original individually labelled pairs. Thus the final permutation preserves every such pair, allowing only swaps within it. These swaps fix balanced payload rows.

For a balanced row with common value d_s on pair s, mutating both selected outer vertices gives

$$d'_s = -d_s, \qquad d'_t = d_t + a_{st}d_s \quad (t \neq s). \tag{25}$$

To verify this for either sign of d_s, the two opposite incidences contribute, respectively, $a_{st}[d_s]_+$ and $-a_{st}[-d_s]_+$, whose sum is $a_{st}d_s$. Formula (25) is exactly (23). Peeling therefore proves that any successful word induces a W-word on the payload frames.

Conversely, mutate both vertices of the outer pair and then swap their labels. The principal matrix and every marker return to their initial values, while balanced rows undergo the desired simple reflection. A word of such moves realizes any W-word. Mutation is equivariant under relabelling, so its intermediate swaps can be moved to a single final permutation by relabelling subsequent mutation indices. This proves both directions. □

Entries between ports can change during these moves. The following four-vertex component repairs them. For every unordered pair of ports f, g, add an active component with principal block and two nonzero port rows

$$P = \begin{pmatrix} 0 & 2 & 0 & -2 \\ -2 & 0 & 2 & 0 \\ 0 & -2 & 0 & 2 \\ 2 & 0 & -2 & 0 \end{pmatrix}, \qquad q_f = (2, -2, 2, -2), \qquad q_g = (0, 2, -2, 0).$$

All its other port rows are zero. Its four vertices are numbered $0, 1, 2, 3$. The word $(0, 2, 1, 3)$ restores P and both rows, adding 8 to b_{fg}. The row calculation is

last vertex	q_f	q_g	cumulative change of b_{fg}
0	$(-2, 2, 2, -2)$	$(0, 2, -2, 0)$	0
2	$(-2, 2, -2, 2)$	$(0, -2, 2, 0)$	4
1	$(2, -2, -2, 2)$	$(0, 2, -2, 0)$	8
3	$(2, -2, 2, -2)$	$(0, 2, -2, 0)$	8

For the principal block, list the six upper entries in the order $(01, 02, 03, 12, 13, 23)$. Their successive values are

initial	2	0	−2	2	0	2
0	−2	0	2	2	−4	2
2	−2	0	2	−2	0	−2
1	2	−4	2	2	0	−2
3	2	0	−2	2	0	2

Thus the principal block also returns to P. No other port–port entry changes, because every other port row vanishes on this component. The inverse word subtracts 8.

All main-core port entries are divisible by 4 along active words, so every main-core change between ports is divisible by 16. Independent translator loops therefore cancel every discrepancy. Take both endpoint port–port blocks to be zero, and prescribe each translator's

principal and mixed endpoint data to be the same. Active components are disconnected from each other while ports remain frozen. Their mutations commute on the entire square matrix. Mutation preserves the vertex sets of connected active components: it cannot join distinct components, and reversibility prevents a split. The main component has $4m > 4$ vertices, and each translator has four, so a final active permutation preserves the main component. Deleting translator letters in the converse leaves a main-core word covered by Lemma 8.5.

Let H be the maximum entry magnitude of the fixed active endpoint block. Every active row has a nonzero principal entry of magnitude at most H, which is unaffected by port scaling. Each port has an incidence of magnitude 2 to a translator and all old entries are even. Apply Lemma 8.2. It gives ordinary all-mutable source and target matrices X, Y_u with

$$X \sim Y_u \quad \Longleftrightarrow \quad u = 1 \text{ in } G_0. \tag{26}$$

All constructed endpoints are connected. The main core is connected because every pair of distinct inner blocks has nonzero couplings. Every marker attaches to its outer pair; every payload attaches because f is injective and the four source and target vectors are nonzero. Each translator attaches to its two ports, and each guard attaches to its port.

There are m reflection normals, $4m$ main vertices, $2m + 4$ ports, four translator vertices for each pair of ports, and six guards per port. Thus the common order is the fixed integer

$$n_* = 4m + 4\binom{2m+4}{2} + 7(2m+4) = 8m^2 + 46m + 52.$$

Only the four target payload rows vary with u. Products of fixed integral generator matrices have bit length $O(|u|)$, so the map from the written word u to Y_u is polynomial-time. The fixed presentation and fixed embedding words are chosen once; existence of these finite data suffices for the computable reduction from an RE-complete set.

Proposition 8.6. *The connected mutation-cyclic equivalence promise problem is RE-complete, already at a fixed order and with a fixed source matrix.*

Proof. Equation (26) gives RE-hardness, and Lemma 8.2 ensures the promise. Enumerating all finite mutation words and final permutations semidecides equivalence. If a decision procedure existed on the promised domain, composing it with the reduction would decide the word problem of G_0. A NO semidecider on that domain, run in parallel with the YES semidecider, would give such a decision procedure. Thus neither exists. Every promise-NP or promise-coNP problem has a decision procedure on its promise, by finite enumeration of polynomially bounded certificates, so these memberships are excluded as well. □

9. Strong NP-hardness

We use Exact Cover by 3-Sets: given a universe of size $3K$ and a list of three-element subsets, determine whether K of them partition the universe. This problem is NP-complete [32]. We shall construct connected mutation-cyclic endpoints of polynomial order whose entries have polynomial magnitude.

9.1. Growth outside commuting cubes.

For a skew-symmetric matrix, its mass is $\sum_{i<j} |b_{ij}|$. An independent set is a set of pairwise nonadjacent vertices. We require a mutation growth lemma with independent blocks.

Call a quiver a multipartite fork if its vertices have a partition

$$D \sqcup U \sqcup O_1 \sqcup \cdots \sqcup O_a \sqcup I_1 \sqcup \cdots \sqcup I_b$$

with D, $O = \bigcup_t O_t$, and $I = \bigcup_t I_t$ nonempty, with U possibly empty, and satisfying the following conditions. Its zero-edge equivalence classes are $D \cup U$ and the individual nonempty O_t, I_t; between different classes every edge has weight at least 2. On the complement of D,

arrows point forwards in the order $O_1, \dots, O_a, U, I_1, \dots, I_b$. The edges incident to D point $D \to O$ and $I \to D$. For $d \in D, o \in O, i \in I$, the edge $o \to i$ has weight strictly larger than each of the edges $d \to o$ and $i \to d$. The set D is called its return set.

Lemma 9.1. *Mutation at $k \notin D$ preserves the multipartite-fork class, weakly increases each edge magnitude, preserves every existing support edge, and strictly increases mass. Its new blocks are as follows, with empty blocks omitted:*

1. *If $k \in U$, the return and unused sets become $D \cup \{k\}$ and $U \setminus \{k\}$.*
2. *If $k \in O_t$, the return set is $\{k\}$, the unused set is $O_t \setminus \{k\}$, the out-blocks are $D, O_1, \dots, O_{t-1}$, and the in-blocks are $O_{t+1}, \dots, O_a, U, I_1, \dots, I_b$.*
3. *If $k \in I_t$, the return set is $\{k\}$, the unused set is $I_t \setminus \{k\}$, the out-blocks are $O_1, \dots, O_a, U, I_1, \dots, I_{t-1}$, and the in-blocks are $I_{t+1}, \dots, I_b, D$.*

The return set is intrinsic: it is the union of the vertices opposite the uniquely largest edges of directed triangles.

Proof. For $k \in U$, the two-step paths are $o \to k \to i$. Every corresponding cross weight increases by the product of two weights at least 2, proving both the old and new dominance inequalities. All other nonincident edges remain unchanged.

For $k \in O_t$, its in-neighbors are $D \cup O_1 \cup \cdots \cup O_{t-1}$ and its out-neighbors are $O_{t+1} \cup \cdots \cup O_a \cup U \cup I$. A changed edge is acyclic or initially missing, except for a cyclic triple $d \to k \to i \to d$. In the first cases the new weight is the old nonnegative weight plus the product, with the indicated forward orientation. In the cyclic case write its three weights as x, y, z, respectively. Here $x \geq 2$ and $y > \max(x, z)$, so mutation replaces the edge $i \to d$ by $d \to i$ of weight

$$xy - z \geq 2y - z > y > x, \qquad xy - z > z.$$

Thus every changed edge increases in magnitude and supplies the required dominance inequalities in the displayed new blocks. The old class $D \cup U$ separates into D, U when both are nonempty; the class containing k remains independent and becomes the new return-plus-unused class. All other orientations are unchanged or reversed at k, giving exactly the stated block order. The two new sides are nonempty, and there is a cyclic D, k, I triple, so mass increases strictly. Global reversal proves the I_t case.

The complement of D is acyclic, so every directed triangle contains exactly one vertex $d \in D$, together with $o \in O, i \in I$. Its uniquely largest edge is $o \to i$. Conversely each d lies in such a triangle. This proves the intrinsic description. □

A growth step means a mutation outside the current return set. A commuting block of a word is a subset of one zero class, with each vertex used once. Its letters commute as actual mutations on the full matrix, including unmutated rows.

Lemma 9.2. *Let P be a multipartite fork and let S be one of its O- or I-classes. Starting with $\mu_s(P)$ for $s \in S$, every finite mutation path, up to its first return to the exact labelled matrix P, has an endpoint represented by a nonempty growth word from P. Its first commuting block is a nonempty subset of S. The representing word is obtained by actual zero-edge commutations and involution cancellations.*

Proof. Initialize a block representation with $\{s\}$. Maintain the following invariant: every block was selected from one zero class disjoint from the preceding return set; the current last block is exactly the current return set. Lemma 9.1 maintains this when a mutation in an O- or I-class starts a new block, or a mutation in U enlarges the last block.

If the next mutation is d in the current return set, commute its previous occurrence to the end of the last block and cancel dd. The surviving letters of that block may be replayed in any order from its preceding state. They belong to one independent class disjoint from the preceding return set; the first is a growth step and the others are its unused twins. Deleting one of these independent letters changes no incidence of another before its turn.

If the last block becomes empty, remove it. If the whole word becomes empty, the actual endpoint is P.

Until this happens, the first block remains a nonempty subset of its original class S, although its individual letters may change. This proves the assertion by processing the path one letter at a time. A nonempty growth word cannot represent P, since its mass is strictly larger. □

We will apply the lemma at the partial-clone states described below.

9.2. **The exact-cover matrices.** Let the universe have $3K$ elements, with $K \geq 1$, and let m_0 be the number of listed triples. If an element occurs in no triple, output the fixed pair

$$M_2 = \begin{pmatrix} 0 & 2 & -2 \\ -2 & 0 & 2 \\ 2 & -2 & 0 \end{pmatrix}, \qquad M_4 = \begin{pmatrix} 0 & 4 & -4 \\ -4 & 0 & 4 \\ 4 & -4 & 0 \end{pmatrix}.$$

These matrices are connected and mutation-cyclic. Their row-gcd multisets are $\{2,2,2\}$ and $\{4,4,4\}$, so they are inequivalent. For the remaining inputs pad the triple list by empty sets until its length is

$$m = \max(m_0, 2K+1, 7), \qquad a = 4(K+3), \qquad M = a+1, \qquad G = 2M^2.$$

Write $\varepsilon_{ei} = 1$ if element e belongs to the ith set and zero otherwise, including the padded empty sets.

First form a matrix B with controls $u_1, \dots, u_m$, outputs p, q, and frozen data vertices r_e for the universe elements and r_0 for an additional neutral row. Its specified entries are

$$\begin{aligned} B_{u_i p} &= -a, & B_{u_i q} &= a, & B_{pq} &= -a^2 K, \\ B_{r_e u_i} &= -2M\varepsilon_{ei}, & B_{r_e p} &= 0, & B_{r_e q} &= 2aM, \\ B_{r_0 u_i} &= -2M, & B_{r_0 p} &= 0, & B_{r_0 q} &= 2aKM. \end{aligned}$$

All unspecified entries between distinct controls or data vertices are zero, and skew symmetry determines reverse entries.

For such a matrix define $E_p(B)$ by replacing p with four mutually nonadjacent clones $p_1, \dots, p_4$, each having half of p's entry to every other vertex. Entries between other vertices are retained. Let σ interchange p and q in B. Set

$$X_{\text{old}} = E_p(B), \qquad Y_{\text{old}} = E_p(\sigma B).$$

All entries are even integers. At each data vertex attach the same two triangles as in Lemma 8.2, now with internal weights and opposite attachment weights $G = 2M^2$, and make every vertex mutable. Call the resulting endpoints X, Y. There are $3K+1$ data vertices and

$$N = m + 5 + 7(3K+1) = m + 21K + 12$$

vertices in total. Every entry has magnitude at most

$$\max\{a^2K/2, 2aKM, 2M^2\} = O(N^3). \tag{27}$$

Lemma 9.3. *Every successful ordinary word between X and a permutation of Y induces a word between their old matrices using only controls and outputs. Its endpoint permutation preserves the control set, clone set, vertex q, and data set. Conversely such an old word lifts if its endpoint permutation extends to the attached guards.*

Proof. The endpoint row gcds on controls, clones, q, data, and guards are respectively

$$2, \qquad a/2, \qquad a, \qquad 2M, \qquad 2M^2.$$

For a control, the neutral incidence supplies $2M$ and the clone incidence supplies $a/2$; their gcd is 2 because a is divisible by 4 and $M = a+1$. A data row has an incidence $2M$ because its element occurs, or because it is neutral; all its other entries are multiples of $2M$. The output and guard gcds follow from their control or internal incidences. These five values

are distinct for $a \geq 16$. Row-gcd invariance therefore gives the stated vertex classes at any endpoint permutation.

Put zero radical coordinates on old vertices and one on guards. The same kernel calculation, peak normalization, and opposite-charge argument as in Lemma 8.2 prohibit the first data mutation. At the endpoint, triangle equations force all data coordinates to zero even when the data vertices and their guards are permuted. Deleting guard moves changes no old–old entry. Conversely an old word changes no guard data and lifts whenever the endpoint permutation carries attached triangles with their data vertices. □

We analyze the principal matrix on controls, clones and q. Mutations at distinct controls commute on the entire matrix. After mutating precisely a subset S of controls, the selected controls have entries $(a/2, -a)$ towards each clone and q, while the unselected controls have $(-a/2, a)$. Every clone has q-entry

$$-a^2(K - |S|)/2.$$

Call these the first control-cube states. There is a second control cube based at Y_{old}. The word cube refers to the 2^m states given by these commuting control involutions.

Lemma 9.4. *Every principal word to an allowed target permutation can be reduced on the entire old matrix to motion in the two control cubes and complete traversals of the commuting five-output cubes based at subsets S with $|S| = K$. In this reduced word the number of complete traversals is odd, and the residual control subset is empty.*

Proof. If $|S| < K$, the principal matrix is a multipartite fork: its return set is the unselected controls, its unused set is the selected controls, its out-side is the clone block, and its in-side is $\{q\}$. If $|S| > K$, interchange selected and unselected controls and interchange the two sides. The required dominant cross weight is $a^2|K - |S||/2 > a$, and all nonzero entries have magnitude at least 2. These statements use the negative-entry arrow convention.

An output move from a nonwinning cube state is an O- or I-move. Lemma 9.2 describes every excursion until its first return to its exact parent. At a proper nonempty S, every possible first output move creates all edges between S and its complement. The two resulting control zero classes have sizes less than m; every other zero class has size at most four. Thus the resulting matrix has no independent set of size m. Subsequent growth cannot lose support, whereas every permutation of a control-cube endpoint has an independent set of size m. Every such excursion must therefore return to its parent before reaching a cube endpoint. Its deletion is an identity on the full old matrix.

For the empty or full subset the first output move is handled by mass. Put $d = K$ or $m - K$, respectively, so $d \geq 1$. A single clone mutation gives principal mass

$$ma^3d/4 + ma + 2a^2d.$$

Mutation at q instead gives $2ma^3d - ma + 2a^2d$: its $4m$ control–clone edges have weight $a^3d/2 - a/2$, its m control–q edges have weight a, and the clone–q edges have total weight $2a^2d$. This mass exceeds the single-clone mass by $ma(7a^2d/4 - 2) > 0$. Each control-cube state has mass at most $3am + 2a^2(m - 1)$. The single-clone mass exceeds this bound: after subtracting and dividing by a, the difference is $m(a^2d/4 - 2 - 2a) + 2ad + 2a > 0$ for $a \geq 16, d \geq 1$. All subsequent growth retains this mass obstruction. The same parent-return deletion applies at these two extreme subsets.

At $|S| = K$ the four clones and q are mutually nonadjacent, hence form a commuting five-cube. After r distinct clones and e copies of q have been mutated, with $0 \leq r \leq 4$ and $e \in \{0, 1\}$, the cross-control magnitude is

$$|r - 4e|a^2/4.$$

It vanishes exactly at $(r, e) = (0, 0)$ and $(4, 1)$. Every intermediate principal state is a multipartite fork. If $e = 0$ and $1 \leq r \leq 4$, its return set consists of the r mutated clones together with q, its unused set consists of the unmutated clones, and its out- and in-sides

are the unselected and selected controls, respectively. If $e = 1$ and $0 \le r \le 3$, its return set consists of the unmutated clones together with q, its unused set consists of the mutated clones, and its out- and in-sides are the selected and unselected controls. These partitions follow by reversing the incidences at each selected output; the cross-control entry is $(r - 4e)a^2/4$. Its absolute value is at least $a^2/4 > a$, which gives the fork inequalities. All control mutations leaving this cube are O- or I-moves. After such a first move every zero class has size at most $\max\{K, m - K, 5\} < m$. Lemma 9.2 and support preservation again force return to the exact parent before a control-cube target is reached.

Delete these excursions in a finite word, and repeat. All deletions use zero-edge commutations and involution cancellations, so they are identities on all unmutated data rows and all data–data entries. Within each remaining cube, cancelling repeated commuting letters reduces its word to the selected subset. A complete five-output traversal takes a first-cube winning state to its counterpart in the second cube, and conversely. Between them only control-cube motion remains.

At an allowed target permutation, row gcds preserve controls, clones and q. After an even number of traversals, the control signs could match the target only if all controls were selected. Their clone–q magnitude would then be $a^2(m - K)/2$, different from the target value $a^2K/2$ because $m > 2K$. After an odd number, the signs require no residual selected controls. This proves the stated normal form and parity. □

9.3. **Data rows and exact covers.** A wrapped gate at a K-subset S means mutation at all its controls, then at the four clones and q, then at the same controls in reverse order. It interchanges the two principal control-cube bases.

Fold the four equal clone rows and columns to one output: its incidence to any other vertex is twice a clone incidence. A control mutation preserves this relation, because each correction to a clone incidence has one factor one half of the corresponding folded factor. The four clones remain nonadjacent. During their complete commuting block of mutations, the total correction between two outside vertices is four products of halves, hence the same correction as one mutation at the folded output. The incidences of all four clones are then negated and again equal. This proves the folding identity on the entire matrix, including data–data entries, for arbitrary signed data rows.

After each complete gate, interchange the names p, q in the folded description so that its principal base is again the initial one. The folded output represented by the four physical clones therefore alternates with the parity of the number of gates. Each gate mutates both folded outputs, so its physical realization is always the same four-clone block and the singleton output. The exchange of names is an identification of folded coordinates; the physical vertex sets retain their original sizes and labels.

For one such folded data row put

$$U_i = -B_{ru_i}, \qquad T = B_{rp} + B_{rq}, \qquad \ell_R = T - a\sum_{i\in R} U_i \quad (|R| = K).$$

The vector $(\ell_R)_{|R|=K}$ records all its exact-cover defects. For a constraint row its initial values are

$$\frac{\ell_R}{2aM} = 1 - \sum_{i\in R} \varepsilon_{ei} \in [-K, K]. \tag{28}$$

For the neutral row all defects are zero.

Lemma 9.5. *For arbitrary signed data rows, a normalized wrapped gate at S acts by*

$$\ell'_R = \ell_R + (a^2(K - |R \cap S|) - 2)\ell_S \quad (|R| = K). \tag{29}$$

In particular $\ell'_S = -\ell_S$. Each wrapped gate, with the exchange of p, q at its end, is an involution on the entire old matrix.

Proof. After the initial control mutations, the two row entries at the folded outputs are

$$P = B_{rp} - a\sum_{i\in S}[U_i]_+, \qquad Q = B_{rq} + a\sum_{i\in S}[-U_i]_+.$$

Thus $P + Q = \ell_S$. Define $g = [P]_+ - [-Q]_+$. The two output mutations negate P, Q. Their corrections to the data–control entries are

$$b_{ru_i} = \begin{cases} U_i + a([Q]_+ - [-P]_+), & i \in S, \\ -U_i + ag, & i \notin S. \end{cases}$$

For example, when $i \in S$, its incidences from the two outputs are $-a, a$, so the respective corrections are $-a[-P]_+$ and $a[Q]_+$. When $i \notin S$, these incidences are $a, -a$, giving ag. Since $[Q]_+ - [-P]_+ = P + Q - g = \ell_S - g$, undoing the selected controls and taking negatives of their entries gives

$$\begin{aligned} U_i' &= U_i + a\ell_S \mathbf{1}_{i\in S} - ag, \\ T' &= -\ell_S + a\sum_{i\in S}(U_i + a(\ell_S - g)) \\ &= T + (a^2K - 2)\ell_S - a^2Kg. \end{aligned}$$

Here $\mathbf{1}_{i\in S}$ is one for $i \in S$ and zero otherwise. For the second identity, the output sum is $-\ell_S$ before undoing controls; each selected control adds $a(U_i + a(\ell_S - g))$ to it. Subtracting $a\sum_{i\in R} U_i'$ from T' cancels the common g terms and gives (29).

For involutivity, the initial and final control words are mutual inverses. At their middle the outputs are mutually nonadjacent, so their simultaneous mutation is an involution commuting with output-name interchange. Name interchange also commutes with every control mutation, by mutation equivariance. Conjugating this commuting product by the control word proves the assertion on the entire matrix, including data–data entries. The same argument unfolds to four clones at the complete gate endpoints. □

Lemma 9.6. *Let $m > K \geq 1$. Let an integer vector have coordinates indexed by the K-subsets of an m-set, all initially in $[-K, K]$. For every index S, let s_S negate its coordinate and add c_{RS} times that coordinate to coordinate $R \neq S$, where each c_{RS} is an integer greater than $2K$. Any reduced word in these involutions whose first active letter acts on a nonzero coordinate leaves the box $[-K, K]$ and cannot return to it.*

Proof. The first active letter acts on an integer $t \neq 0$. Its new S-coordinate is $-t$. Every other coordinate has the sign of t and magnitude at least $(2K + 1)|t| - K > |t|$, and in particular exceeds K. Thus there is one exceptional coordinate, opposite in sign and strictly smaller in magnitude than every other coordinate.

A next different letter selects a nonexceptional coordinate t'. Its magnitude is larger than that of the previous exceptional coordinate. Adding $c_{RS}t'$ makes every other coordinate have the sign of t' and magnitude greater than $|t'|$: the old exceptional coordinate contributes at worst its smaller absolute value, while all others already have the sign of t'. The selected coordinate becomes $-t'$, the new exception. Induction preserves this sign pattern and strictly increases the exceptional magnitude at every subsequent letter. Hence at every nonempty stage some coordinate lies outside the box. □

Proposition 9.7. *The promised equivalence problem for connected mutation-cyclic integer skew-symmetric matrices is strongly NP-hard.*

Proof. For $R \neq S$ the coefficient in (29) is at least $a^2 - 2 > 2K$. Divide the constraint defects by $2aM$; they remain integers under every gate and initially lie in the box by (28). A target permutation preserves controls and data rows, so it merely permutes these rows and their subset indices. After the output normalization, every target row therefore still has defects in the same box.

By Lemmas 9.3 and 9.4, a successful word gives an odd word of normalized wrapped gates. Cancel adjacent identical gates using their full-matrix involutivity. The reduced word is nonempty. In a NO exact-cover instance its first gate S has a nonzero initial defect in some constraint row. Lemma 9.6 shows that this row cannot reach the target box under the reduced word, contradicting the endpoint condition.

Conversely let S be an exact cover. Every constraint defect at S is zero, as is the neutral defect. At the selected control state, p, q are opposite twins on the entire folded old matrix, including data rows, and have zero mutual entry. Their simultaneous mutations cancel each change outside the two outputs and interchange their roles. Four-cloning realizes this move by the five commuting output mutations. Thus the word

$$S,\ p_1, p_2, p_3, p_4, q,\ \mathrm{reverse}(S)$$

carries X exactly to Y and has length $2K + 5$. It changes no guard entry.

A K-subset involving a padded empty set cannot have all defects zero: K sets of sizes at most three, one empty, have fewer than $3K$ incidences, whereas zero defects require each of the $3K$ elements exactly once. Thus the reduction recognizes precisely the original exact covers.

The principal support is connected, every data vertex attaches through an occurring incidence or the neutral row, and every guard attaches to its data vertex. Each endpoint contains a full equal-weight guard triangle with $G \geq 2$, so it is mutation-cyclic. The missing-element preprocessing pair has the same promise. Finally (27) bounds both the order and all numerical entries by polynomials in the input size. The reduction is therefore polynomial for unary as well as binary multiplicities, proving strong NP-hardness. □

Propositions 8.6 and 9.7 prove Theorem 1.3. Since skew-symmetric matrices form a subclass of skew-symmetrizable matrices, the same lower bounds hold in the latter class.

References

[1] S. Fomin and A. Zelevinsky, *Cluster algebras I: Foundations*, J. Amer. Math. Soc. **15** (2002), 497–529.

[2] S. Fomin, L. Williams, and A. Zelevinsky, *Introduction to cluster algebras. Chapters 1–3*, arXiv:1608.05735v4; combined preliminary draft of Chapters 1–7 (2021), Problems 2.6.14 and 2.8.2, https://pages.uoregon.edu/arkadiy/intro.clust.alg.pdf.

[3] K. T. K. Armstrong-Williams, E. Hirst, B. Jackson, and K.-H. Lee, *Mutation-acyclicity in rank four is decidable but not semialgebraic*, arXiv:2610.36179v1, 2026.

[4] Q. Tang, *Decidability of mutation-acyclicity*, Hexagon preprint 2611.00163v1, 2026.

[5] D. Soukup, *Complexity of ice quiver mutation equivalence*, Ann. Comb. **29** (2025), 91–100, doi:10.1007/s00026-023-00668-w.

[6] Q. Tang, *Undecidability of quiver mutation equivalence*, Hexagon preprint 2611.00158v1, 2026.

[7] S. Fomin and A. Zelevinsky, *Cluster algebras II: Finite type classification*, Invent. Math. **154** (2003), 63–121, doi:10.1007/s00222-003-0302-y.

[8] E. S. Dias and D. Castonguay, *Polynomial recognition of cluster algebras of finite type*, arXiv:1507.03844v1, 2015.

[9] P. Caldero and B. Keller, *From triangulated categories to cluster algebras II*, Ann. Sci. Éc. Norm. Supér. (4) **39** (2006), 983–1009, doi:10.1016/j.ansens.2006.09.003.

[10] I. Assem, M. Blais, T. Brüstle, and A. Samson, *Mutation classes of skew-symmetric* 3×3*-matrices*, Comm. Algebra **36** (2008), 1209–1220, arXiv:math/0610627.

[11] A. I. Seven, *Mutation classes of skew-symmetrizable* 3×3 *matrices*, arXiv:1012.3318v2, 2011.

[12] R. Akagi, *Cluster-cyclic condition of skew-symmetrizable matrices of rank 3 via the Markov constant*, arXiv:2411.07083v2, 2025.

[13] A. Felikson, M. Shapiro, and P. Tumarkin, *Skew-symmetric cluster algebras of finite mutation type*, J. Eur. Math. Soc. **14** (2012), 1135–1180, doi:10.4171/JEMS/329.

[14] S. Fomin, K. Igusa, and K. Lee, *Universal quivers*, Algebr. Comb. **4** (2021), 683–702, doi:10.5802/alco.175.

[15] S. Fomin, *Quiver mutations*, Open Problems in Algebraic Combinatorics, Minneapolis, May 16, 2022, https://www.samuelfhopkins.com/OPAC/files/slides/fomin.pdf.

[16] S. Fomin and S. Neville, *Long mutation cycles*, Selecta Math. (N.S.) **31** (2025), article 103, doi:10.1007/s00029-025-01090-1.

[17] V. G. Kac, *Infinite-dimensional Lie algebras*, third edition, Cambridge University Press, Cambridge, 1990.
[18] C. T. McMullen, Coxeter groups, Salem numbers and the Hilbert metric, *Publications Mathématiques de l'IHÉS* **95** (2002), 151–183, doi:10.1007/s102400200001.
[19] N. Chapelier-Laget and C. Hohlweg, Shi arrangements and low elements in affine Coxeter groups, *Canadian Journal of Mathematics* **77** (2025), 683–713, doi:10.4153/S0008414X24000130.
[20] A. Hubery and H. Krause, *A categorification of non-crossing partitions*, J. Eur. Math. Soc. **18** (2016), 2273–2313, doi:10.4171/JEMS/641.
[21] A. Hubery, *The cluster complex of an hereditary Artin algebra*, Algebr. Represent. Theory **14** (2011), 1163–1185, arXiv:0812.1475.
[22] D. Rupel, *Quantum cluster characters for valued quivers*, Trans. Amer. Math. Soc. **367** (2015), 7061–7102, doi:10.1090/S0002-9947-2015-06251-5.
[23] D. Happel and C. M. Ringel, *Tilted algebras*, Trans. Amer. Math. Soc. **274** (1982), 399–443.
[24] A. B. Buan, R. J. Marsh, M. Reineke, I. Reiten, and G. Todorov, *Tilting theory and cluster combinatorics*, Adv. Math. **204** (2006), 572–618, arXiv:math/0402054.
[25] C. M. Ringel, *The spectral radius of the Coxeter transformations for a generalized Cartan matrix*, Math. Ann. **300** (1994), 331–339, doi:10.1007/BF01450490.
[26] D. E. Speyer, *Powers of Coxeter elements in infinite groups are reduced*, Proc. Amer. Math. Soc. **137** (2009), 1295–1302, arXiv:0710.3188.
[27] M. Fiedler, *Numerical range of matrices and Levinger's theorem*, Linear Algebra Appl. **220** (1995), 171–180, doi:10.1016/0024-3795(94)00073-M.
[28] G. Higman, Subgroups of finitely presented groups, *Proc. Roy. Soc. London Ser. A* **262** (1961), 455–475. doi:10.1098/rspa.1961.0132.
[29] R. C. Lyndon and P. E. Schupp, *Combinatorial Group Theory*, Ergebnisse der Mathematik und ihrer Grenzgebiete, vol. 89, Springer-Verlag, Berlin–New York, 1977.
[30] A. Beineke, T. Brüstle and L. Hille, Cluster-cyclic quivers with three vertices and the Markov equation, with an appendix by O. Kerner, *Algebr. Represent. Theory* **14** (2011), 97–112. doi:10.1007/s10468-009-9179-9.
[31] A. B. Buan, B. R. Marsh and I. Reiten, Cluster mutation via quiver representations, *Comment. Math. Helv.* **83** (2008), 143–177. doi:10.4171/CMH/121.
[32] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, San Francisco, 1979.

Zhili College, Tsinghua University, Beijing 100084, China
Email address: tangqy24@mails.tsinghua.edu.cn