hexagon:2611.00158v2

UNDECIDABILITY OF QUIVER MUTATION EQUIVALENCE

QIYUE TANG

Abstract. We prove that quiver mutation equivalence up to relabeling is undecidable even for a fixed source quiver. We also prove that it is strongly NP-hard.

1. Introduction

Fomin and Zelevinsky introduced matrix mutation as part of the definition of cluster algebras [10, Definition 4.2]. Mutation of a seed replaces one cluster variable and updates its exchange matrix. For a skew-symmetric exchange matrix, this update is an operation on a quiver, a finite directed multigraph without loops or directed two-cycles. Iterating it gives the mutation class of the quiver. Given two quivers, the question is whether their mutation classes coincide. A matrix is skew-symmetrizable if multiplying its rows by positive constants makes it skew-symmetric. A quiver is acyclic if it has no directed cycles.

Early classifications gave ways to describe particular classes. Fomin and Zelevinsky classified cluster algebras with finitely many cluster variables by finite Cartan type [11, Theorem 1.4]. The finiteness of the mutation class of an exchange matrix is a weaker property. Fomin, Shapiro and Thurston associated exchange matrices to triangulations of marked surfaces, with flips of arcs corresponding to mutations [3]. Felikson, Shapiro and Tumarkin classified the connected skew-symmetric quivers of finite mutation type on at least three vertices [4], obtaining the surface classes and eleven exceptional classes. When the class of the source is finite, enumerating mutations and recording the matrices already visited terminates and decides membership.

A different method works in small order, where the class need not be finite. Assem, Blais, Brüstle and Samson used descent in the arrow multiplicities to distinguish the three-vertex quivers whose classes contain a quiver without directed cycles from those whose classes do not. In the latter case they classified the classes, up to permutation and transposition of matrices, by minimal sorted triples of multiplicities [1]. Seven studied minimal representatives of the diagrams associated to skew-symmetrizable matrices of order three [15].

The general decision problem appears in the 2016 version of Fomin, Williams and Zelevinsky's *Introduction to Cluster Algebras* [8, Problem 2.6.12]. Fomin, Igusa and Lee restated it in their work on universal quivers [6, Remark 2.3]. Their construction puts every quiver of a given order into a representative of one mutation class as a full subquiver, obtained by retaining a subset of the vertices and all arrows between them. Fomin asked for a decision algorithm for two specified quivers in his OPAC 2022 lecture [5, slide 7]. The later version of the book records the quiver problem and its matrix extension [9, Problems 2.6.14 and 2.8.2].

Soukup obtained complexity lower bounds for a related question from the same lecture, in which the target is specified by a property of its arrows. An ice quiver has a specified subset of frozen vertices, at which mutation is disallowed. Soukup proved that deciding whether an ice quiver can be mutated so that there are no arrows between its frozen vertices

Date: November 6, 2026.

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

Key words and phrases. quiver mutation, matrix mutation, undecidability, strong NP-hardness.

is strongly NP-hard [16]. Fomin and Neville constructed arbitrarily long mutation cycles at every fixed order at least four. They showed that steepest descent cannot in general produce shortest mutation paths [7]. More recently, Armstrong-Williams, Hirst, Jackson and Lee gave a finite search for deciding whether a four-vertex quiver has an acyclic representative in its class. They also decide equivalence of two four-vertex quivers when at least one has such a representative [2, Corollaries 5.2 and 5.7].

We formulate the problem in matrix coordinates. For vertices i, j of a labeled quiver, let b_{ij} be the number of arrows $i \to j$ minus the number of arrows $j \to i$. This gives an integer skew-symmetric matrix $B = (b_{ij})$, and every such matrix specifies a quiver by this convention. Write $[x]_+ = \max(x, 0)$ for real x. Mutation of B at vertex k is the matrix $\mu_k B$ with entries

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

Two matrices B, C of order n are mutation equivalent if

$$C = P^T \mu_{k_t} \cdots \mu_{k_1}(B) P$$

for a permutation matrix P of order n, an integer $t \geq 0$, and indices $k_1, \ldots, k_t$ in $\{1, \ldots, n\}$, applied in that order. The empty word acts as the identity. Mutation is involutive and equivariant under relabeling by (1). Thus allowing relabelings between mutations gives the same equivalence. In constructions that temporarily designate vertices as frozen, we still use (1) on the entire square matrix and allow mutations only at the other vertices. In particular entries between frozen vertices are retained.

A language is computably enumerable, or recursively enumerable, if an algorithm halts and accepts exactly its members. It is computably enumerable complete if, in addition, every computably enumerable language reduces to it by a total computable map preserving membership. Here strong NP-hardness means NP-hardness under polynomial-time many-one reductions even when integer magnitudes are written in unary. The reduction below establishes it with a fixed polynomial bound in the matrix order.

Theorem 1.1. *Mutation equivalence of two integer skew-symmetric matrices is recursively enumerable complete under computable many-one reductions, with all indices mutable and with a simultaneous final row and column permutation allowed. There is a fixed all-even skew-symmetric matrix whose equivalence to an input all-even skew-symmetric matrix of the same order is already recursively enumerable complete.*

Theorem 1.2. *Mutation equivalence up to simultaneous permutation is NP-hard on pairs of even skew-symmetric integer matrices of order N whose entries have magnitude $O(N^3)$. The reduction is polynomial also when the matrix entries are given in unary.*

Theorem 1.1 answers Fomin's question: the general two-quiver problem is undecidable and hence lies outside NP. Computable completeness extends to skew-symmetrizable matrices. Indeed, one can check skew-symmetrizability by testing for a zero diagonal and opposite signs at every nonzero pair b_{ij}, b_{ji}. A pair with exactly one zero entry is rejected. On each connected component of the graph with edges $\{i, j\}$ for $b_{ij} \neq 0$, set one diagonal scale to one and propagate the ratios $-b_{ij}/b_{ji}$ from the scale at i to the scale at j; check that all prescribed ratios agree. On valid inputs, enumerate mutation words and endpoint permutations. The skew-symmetric subclass supplies the hardness.

The proof of Theorem 1.1 first realizes a fixed group word problem as simultaneous reachability of four rows under reflections. A kernel vector and auxiliary rows force a successful mutation word to implement those reflections. For Theorem 1.2, the controls select candidate triples, and the change in the data rows detects whether each element is covered exactly once.

2. A fixed reflection group with undecidable membership

Fix an effective enumeration of Turing machines and let $\mathcal{K}$ contain the positive integers n for which machine $n-1$ halts on empty input. This set is computably enumerable complete with integers encoded in binary: on a proposed input to any semidecision procedure, encode the machine that runs that procedure on the input and halts exactly on acceptance. Consider the group

$$H = \langle a, b \mid [b, a^n b a^{-n}] = 1 \quad (n \in \mathcal{K}) \rangle, \qquad [g, h] = ghg^{-1}h^{-1}.$$

To identify these relations, let R have generators b_i, $i \in \mathbb{Z}$, with defining relations $[b_i, b_j] = 1$ whenever $|i - j| \in \mathcal{K}$. The shift $b_i \mapsto b_{i+1}$ defines an automorphism, and its semidirect product with the infinite cyclic group generated by a has the displayed presentation for H. Indeed, eliminating $b_i = a^i b_0 a^{-i}$ reduces the defining relations to those displayed, with $b = b_0$. If $n \notin \mathcal{K}$, mapping every b_i other than b_0, b_n to the identity gives a retraction onto the free group on b_0, b_n: every defining commutation relation has at most one retained generator. The commutator of the two free generators is a nonempty reduced word. Therefore

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

By Higman's embedding theorem [13], the finitely generated recursively presented group H embeds in a finitely presented group G. Fix such an embedding. The images of a, b are fixed finite words in the generators of G, so substitution gives a computable reduction to the word problem of G. That word problem is recursively enumerable by enumerating products of conjugates of the finitely many relators and their inverses.

For each generator a_i of G, adjoin a generator t_i with $t_i^2 = 1$ and $t_i a_i t_i = a_i^{-1}$. Each step is a free product with amalgamation along $\langle a_i \rangle$: the second factor is $\langle a_i \rangle \rtimes C_2$, with C_2 cyclic of order two, for the inversion action on the cyclic group $\langle a_i \rangle$. The normal form for a free product with amalgamation [14, Chapter I, Section 1] gives an embedding at each step. The resulting finitely presented group G_0 is generated by involutions $t_i, t_i a_i$, since $a_i = t_i(t_i a_i)$. It has a recursively enumerable complete word problem. We may take its free product with additional copies of C_2 so that it has at least two generators in the chosen involution presentation. Write this fixed presentation as $G_0 = \langle \mathcal{A} \mid R_1, \ldots, R_r \rangle$, where q^2 is among the defining relators for each $q \in \mathcal{A}$.

Let F be the free group on $\mathcal{A} \cup \{x\}$, where x is a new generator. Let θ be the automorphism of F sending each basis generator to its inverse. Define

$$\mathcal{T} = \mathcal{A} \cup \{x\} \cup \{R_j, qR_j : q \in \mathcal{A} \cup \{x\},\ 1 \le j \le r\}, \qquad E = F / \langle\langle \theta(g) g : g \in \mathcal{T} \rangle\rangle.$$

Here double angle brackets denote normal closure. Write $\pi : F \to E$ for the quotient map. The normal closure is θ-invariant since $\theta(\theta(g)g) = g\theta(g)$ is conjugate to $\theta(g)g$. Denote the induced automorphism of E by $\bar{\theta}$. The relator from R_j gives $\theta(R_j) = R_j^{-1}$ in E, and that from qR_j gives $q^{-1}R_j^{-1}qR_j = 1$. Thus the images of all R_j are central. Moreover the natural map $F \to G_0 * \langle x \rangle$ factors through E, since $\theta(R_j)$ and R_j have the same image in G_0 and $q^{-1}q = 1$. For a word u on $\mathcal{A}$, if $u = 1$ in G_0 its image in E is a product of central relators and their inverses. If $u \neq 1$ in G_0, then $[x, u] \neq 1$ in the free product by its reduced normal form. We obtain

$$u = 1 \text{ in } G_0 \quad \Longleftrightarrow \quad [x, u] = 1 \text{ in } E. \tag{2}$$

In $F \times F$, let M be generated by (g, g^{-1}) for $g \in \mathcal{T}$. The basis generators supply all $(a, \theta(a))$, $a \in F$, forming the twisted diagonal. For each $g \in \mathcal{T}$,

$$(g, \theta(g))^{-1}(g, g^{-1}) = (1, (g\theta(g))^{-1}).$$

Since $g\theta(g) = g(\theta(g)g)g^{-1}$, conjugation by the twisted diagonal generates $\{1\} \times \ker \pi$ from these products. Every generator of M satisfies the condition on the right below, and

multiplying (a, b) satisfying the condition by $(a, \theta(a))^{-1}$ gives an element of $\{1\} \times \ker \pi$.

$$M = \{(a, b) \in F \times F : \pi(b) = \bar{\theta}(\pi(a))\}. \tag{3}$$

The factor swap normalizes M since it sends each specified generator to its inverse.

We realize F inside the level-four congruence subgroup of $\mathrm{SL}_2(\mathbb{Z})$, consisting of the matrices congruent to the identity modulo 4. 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. In their fractional linear action on the real projective line, $U^n z = z + 4n$ and $V^n z = z/(4nz + 1)$. For $n \neq 0$, U^n sends $|z| < 1$ into $|z| > 1$. On the latter set, $|V^n z| \leq (4|n| - 1/|z|)^{-1} < 1$, with $V^n(\infty) = 1/(4n)$. To see faithfulness, a reduced alternating word can be conjugated, if necessary, to one beginning and ending with a power of the same letter: for a word beginning in U and ending in V, conjugate by a power of U that does not cancel the first syllable, and likewise with the letters exchanged. Such a word maps one of the two disjoint sets into the other. The projective action is therefore faithful as well. Assign distinct $U^i V U^{-i}$, for positive integers i, to the basis generators of F. A reduced word in these conjugates, after combining consecutive powers of the same conjugate, expands to nonzero V-powers separated by nonzero U-powers. This defines a faithful embedding and its image does not contain $-I$.

Use the coordinate order (a, b, c, d) on the entries of a 2×2 matrix and write

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

The bilinear form with matrix J has quadratic form $2 \det X$. Represent $(g, h) \in F \times F$ by $X \mapsto gXh^{-1}$ and the factor swap by $s(X) = -\operatorname{adj}(X)$. For 2×2 matrices, $\operatorname{adj}(X) = \operatorname{tr}(X)I - X$, where the trace is the sum of the diagonal entries. The adjugate reverses multiplication, and $\operatorname{adj}(g) = g^{-1}$ for $g \in F$. Thus $s(g\, s(X) h^{-1}) = hXg^{-1}$. If an element (g, h) acts trivially, taking $X = I$ gives $g = h$, which then commutes with every X and is scalar. The only possibility is $g = h = I$ since $-I$ is absent. The factor $F \times F$ acts by determinant-one maps on the four-dimensional matrix space, whereas s fixes the traceless subspace and negates I, so has determinant -1. Thus the semidirect product representation is faithful.

For $g \in F$, let v_g be the column vector of its matrix entries. The 2×2 adjugate identity $\operatorname{adj}(Y) = \operatorname{tr}(Y)I - Y$ gives

$$r_g(X) := -g \operatorname{adj}(X) g = X - \operatorname{tr}(g^{-1}X) g, \qquad r_g = I - v_g(v_g^{\mathsf{T}} J), \quad v_g^{\mathsf{T}} J v_g = 2.$$

Indeed, $g \operatorname{adj}(g^{-1}X) = g \operatorname{adj}(X) g$, and the trace pairing $\operatorname{tr}(g^{-1}X)$ equals $v_g^{\mathsf{T}} J$ applied to the entry column of X. Taking $X = g$ gives $v_g^{\mathsf{T}} J v_g = 2$. The group K_0 generated by r_1 and the r_g for $g \in \mathcal{T}$ is the image of $M \rtimes \langle s \rangle$, since $r_1 = s$ and $r_g s$ represents (g, g^{-1}). For $w = [x, u]$, write L_w for the matrix of left multiplication by w. Faithfulness, (3) and (2) give

$$L_w \in K_0 \iff (w, 1) \in M \iff u = 1 \text{ in } G_0. \tag{4}$$

3. A simultaneous orbit of four rows

List the matrices 1 and $g \in \mathcal{T}$ as $g_1, \ldots, g_m$, allowing repetitions. Write $v_i = v_{g_i}$, and let T_0 have these columns. They span $\mathbb{Q}^4$: the list includes I and at least three conjugates

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

for distinct positive i. Subtracting I gives three traceless matrices whose coordinates are independent by the Vandermonde determinant on $1, i, i^2$. Set $A = T_0^{\mathsf{T}} J T_0$. Since T_0 is surjective and J invertible, $\ker A = \ker T_0$, so A has rank four. Its entries are $A_{ij} = \operatorname{tr}(g_i^{-1} g_j)$. Every A_{ij} is congruent to 2 modulo 16, since the determinant condition for a matrix $I + 4Z$ gives $\operatorname{tr}(I + 4Z) = 2 - 16 \det Z$.

Choose an odd integer $L > 2 + \max_{i \neq j} |A_{ij}|/2$. Index a symmetric matrix D by $i+, i-$ for $1 \leq i \leq m$. Set the diagonal to 2, set $D_{i+,i-} = D_{i-,i+} = 0$, and for $i \neq j$ put

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

The off-diagonal entries joining distinct blocks $\{i+, i-\}$ are negative even integers of magnitude at least 2. Let W be the free product of m copies of $C_2 \times C_2$, with generators s_{i+}, s_{i-} in factor i. It acts on column vectors by $R_k = I - e_k D_{k,*}$, where e_k is the coordinate column and $D_{k,*}$ is row k. Since $D_{kk} = 2$, $R_k^2 = I$ and $R_k^{\mathsf{T}} D R_k = D$ by expansion. For mates k, j, $D_{kj} = D_{jk} = 0$, so the cross terms in $R_k R_j$ and $R_j R_k$ vanish and the matrices commute. For $h \in W$, write R_h for the resulting column action. The action by right multiplication on rows has the coordinate formula

$$x'_k = -x_k, \qquad x'_j = x_j - D_{kj} x_k \quad (j \neq k). \tag{5}$$

Let ι swap $i+$ with $i-$ for every i, both on vectors and on the generators of W. The formula for R_k gives $\iota R_k \iota^{-1} = R_{\iota(k)}$. In the eigenspaces (z, z) and $(z, -z)$, $z \in \mathbb{Q}^m$, with the two groups of coordinates ordered by sign, the operator D acts by

$$(2 + 2L) I - 2L \mathbf{1}\mathbf{1}^{\mathsf{T}}, \qquad A,$$

where $\mathbf{1}$ is the all-ones column of length m. The Gram matrices in these coordinates are twice the displayed operator matrices. The first operator block is nonsingular, with eigenvalues $2+2L$ and $2-2L(m-1)$. Thus the radical of the form with matrix D is the copy of $\ker A$ in the antisymmetric part. The quotient decomposes into orthogonal nondegenerate symmetric and antisymmetric subspaces of dimensions m and four, respectively.

The quotient action is faithful. To prove this on the dual row space, begin with the row $-\mathbf{1}_{2m}^{\mathsf{T}} D$, all of whose entries are $2L(m-1) - 2 > 0$, where $\mathbf{1}_{2m}$ denotes the all-ones column of length $2m$. Each syllable of a reduced free product word fires a nonempty subset of one block by (5). After the first syllable its selected coordinates are negative, the unselected coordinate, if present, is positive, and every coordinate outside the block is at least the sum $R > 0$ of the absolute values of the negative coordinates: each outside coordinate has gained at least $2R$. Suppose the next syllable uses a different block, with fired coordinates totaling $T \geq R$. Each old negative entry becomes at least $-R + 2T \geq T$, and each other coordinate outside the new block gains at least $2T$ from a positive starting value. The selected entries become negative with total absolute value T, and any unselected mate remains positive. This inductive invariant shows that the final row differs from the initial row. The isomorphism $x + \ker D \mapsto x^{\mathsf{T}} D$ from the column quotient to the row space commutes with ι. Since $R_h^{\mathsf{T}} D R_h = D$, right multiplication of this row by R_h gives $(R_h^{-1} x)^{\mathsf{T}} D$. This proves faithfulness on the column quotient as well.

Write S for the four-dimensional antisymmetric subspace of the quotient. An element preserving S also preserves $S^{\perp}$, so it commutes with ι on the quotient. Faithfulness implies that it is fixed by the automorphism ι of W. Uniqueness of reduced words in the free product shows that each syllable must be fixed by the swap within its block. Thus the stabilizer of S is generated by $c_i = s_{i+} s_{i-}$, $1 \leq i \leq m$. On an antisymmetric vector $(z, -z)$, c_i replaces z by $z - e_i (Az)_i$. The map $z \mapsto T_0 z$ identifies S with $\mathbb{Q}^4$ since $\ker A = \ker T_0$. In this identification c_i acts as the reflection with normal v_i.

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

$$f(z)_{i+} = z^{\mathsf{T}} J v_i, \qquad f(z)_{i-} = -z^{\mathsf{T}} J v_i.$$

For $z = T_0 y$ the row $f(z)$ equals $(y, -y)^{\mathsf{T}} D$, where $(y, -y)$ denotes a column. Since T_0 is surjective and J invertible, f is injective with image the antisymmetric part of the row space of D. Integral columns give integral rows. Formula (5) gives $f(z)c_i = f(r_{g_i} z)$. For a product h of the c_i this becomes $f(z)h = f(\rho(h)^{-1} z)$, where $\rho(h)$ denotes the action on S in the identification with $\mathbb{Q}^4$. This follows by applying the generators in row order and using that each reflection is an involution.

For $w = [x, u]$, the four rows $f(e_a)$, $1 \le a \le 4$, and the four rows $f(L_w e_a)$, where e_a are the coordinate columns of $\mathbb{Q}^4$, each span the same antisymmetric subspace. A word sending the former frame to the latter must preserve that subspace. By the column-to-row isomorphism, its stabilizer is again generated by the c_i. Equality of the four rows is equivalent to $\rho(h)^{-1} = L_w$. Conversely an expression for L_w in the generating reflections supplies the reversed expression in the c_i for the row action. Consequently

$$(f(e_a))_{a=1}^4 \text{ and } (f(L_w e_a))_{a=1}^4 \text{ are in the same } W\text{-orbit} \iff u = 1 \text{ in } G_0. \tag{6}$$

4. Radical transport and mutation words

We use some intermediate skew-symmetric matrices with specified frozen indices, at which mutation is forbidden. An index outside this set is called active. The principal block is the submatrix on active indices. Mutations are computed on the full square matrix, including the block between frozen indices. A row from a frozen index, restricted to the principal columns, is called a frozen row.

For a skew-symmetric principal block $B = (b_{ij})$ and a column v in its kernel, set

$$S_k = \sum_j [b_{kj}]_+ v_j = \sum_j [-b_{kj}]_+ v_j.$$

The equality follows from $Bv = 0$, with j ranging over the active indices. Transport v across mutation at k by replacing v_k with $S_k - v_k$. To check the kernel equation, write $E_k = I$ except that its k-th column has entry -1 at k and entry $[b_{ik}]_+$ at $i \ne k$. Write $B' = \mu_k B$. In $E_k B E_k^{\mathsf{T}}$ the incident entries are negated. Off the k-th row and column the correction is

$$[b_{ik}]_+ b_{kj} + b_{ik}[-b_{kj}]_+ = [b_{ik}]_+ [b_{kj}]_+ - [-b_{ik}]_+ [-b_{kj}]_+.$$

Thus $B' = E_k B E_k^{\mathsf{T}}$, and $E_k^2 = I$ since the k-th diagonal entry is -1. Thus the transported column is $E_k^{\mathsf{T}} v$, since $\sum_i [b_{ik}]_+ v_i = \sum_i [-b_{ki}]_+ v_i = S_k$. The transport is involutive because the positive and negative budgets at k exchange under mutation. For a frozen row q, with updated row q', the two sums in S_k give

$$\sum_{j \ne k} (q'_j - q_j) v_j = [q_k]_+ S_k - [-q_k]_+ S_k = q_k S_k.$$

The new k-term is $-q_k(S_k - v_k)$, so

$$q' v' = q v. \tag{7}$$

In particular the equation holds for a frozen row incident to only one connected component, using its restriction to that component.

If two indices have zero exchange entry, either mutation preserves the entire row and column of the other index. The two mutations therefore commute: the two corrections to any entry outside the selected rows and columns depend on unchanged data and add in either order. The transported coordinate updates also commute, since each budget is independent of the other's coordinate.

Lemma 4.1. *Suppose every entry of the principal block is even, $v \ge 0$, and $S_k \ge 2v_k$ initially at every active index. Every word has a representative with the same full matrix and transported endpoint column, along which all coordinates of v are nondecreasing. The representative is obtained by zero-entry commutations and cancellations of consecutive*

equal letters. If $v > 0$ initially and all the initial budget inequalities are strict, every step of the representative is a strict increase.

Proof. Evenness is preserved by the mutation formula. At a nonnegative column, suppose mutation at k strictly decreases its coordinate and mutation at $j \neq k$ would not increase its coordinate. Then $S_k < 2v_k$, so $v_k > 0$. If $v_j = 0$, then $S_j = 0$, so j is not joined to k, whose coordinate is positive. If $v_j > 0$, a nonzero edge of magnitude $c \geq 2$ would give $S_k \geq cv_j$ and $S_j \geq cv_k$, using the sum with the corresponding edge sign in each row. Their product would contradict $S_k < 2v_k$, $S_j \leq 2v_j$. Thus $b_{kj} = 0$.

In a word with a strict decrease, take the first such step k. The inverse of the preceding step is a nonincrease. If its letter differs from k, the preceding argument permits commuting k back one step, with its budget and coordinate unchanged. Continue left until the letter cancels against a preceding k. If no such letter were met, there would be a strict decrease at the initial column. All states tested before the decreasing step belong to the original nondecreasing prefix. Repeating the shortening operation terminates with the asserted representative.

Under the strict initial hypothesis no equality $S_k = 2v_k$ can arise on a nondecreasing path. Inductively, the selected step is strict, so its inverse is a strict decrease. Each neighbor must then be strictly increasing by the first part of the proof. At every nonneighbor, the row, budget and coordinate are unchanged. The selected direction itself strictly decreases in reverse. Thus there are no flat directions after the step, completing the induction. □

5. Realizing the row orbit by mutations

Use the index set of D from Section 3, and choose a skew-symmetric matrix $E = (E_{st})$ on that index set with $|E_{st}| = -D_{st}$ for $s \neq t$. Replace each s by two active vertices $s0, s1$, and define the $4m \times 4m$ principal block by

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

The two vertices replacing s have opposite columns and no edge between them. Vertices over distinct blocks $\{i+, i-\}$ of D are adjacent, so the principal graph is connected. The all-ones column lies in the kernel because each pair has opposite columns. Each budget at $s0$ or $s1$ counts one incidence of magnitude $-D_{st}$ for every $t \neq s$. The two indices in any other block contribute $2L$ in total. Thus each budget is $2L(m-1) > 2$.

For each s, add one frozen row $4(e_{s0} - e_{s1})^{\mathsf{T}}$, where $e_{s\epsilon}$ denotes the active coordinate column at $s\epsilon$. Call these the marker rows. Add four frozen rows obtained by duplicating each entry of $4f(e_a)$, $1 \leq a \leq 4$, on the corresponding pair. The target has the same principal block and marker rows and replaces $4f(e_a)$ by $4f(L_w e_a)$ in the four duplicated rows. In this section compare the principal block and frozen rows and allow an arbitrary permutation of the active vertices, with every frozen index fixed.

Proposition 5.1. *The specified target is reachable if and only if the two frames in (6) are in the same W-orbit.*

Proof. Suppose a mutation word reaches the target up to permutation. Pull back the target data, including the marker supports, by the endpoint permutation. Normalize the word by Lemma 4.1, transporting the all-ones column. Every remaining step strictly increases its coordinate. At the endpoint, (7) forces the two transported coordinates in each pair to agree, since the corresponding marker has initial charge zero. The pair's two budgets also agree by the opposite principal rows.

For a nonempty word, let k be the first letter in reverse order and j its mate in the endpoint pair. Undoing k is a strict decrease, and the equal coordinates and budgets make mutation at j a strict decrease as well. They are nonadjacent. After undoing k, the

decrease at j persists. Every subsequent reversed letter before j is a strict decrease and hence has zero edge to j by the critical-pair argument in Lemma 4.1. Its mutation leaves the budget and coordinate at j unchanged. The letter j must occur, since otherwise j would still be a strict decrease at the initial state. Commuting it left next to k preserves strict decrease at each crossed step because each interchange is across a zero entry.

Mutating both members of a pair with opposite principal columns swaps the two principal labels. For two entries x, y on a two-step path, its correction is $[x]_+[y]_+ - [-x]_+[-y]_+$. Replacing (x, y) by $(-x, -y)$ negates this correction. In the principal block the two paths through the selected pair have these opposite entries, so their corrections cancel, and the incidences with the pair reverse. The selected marker also sees opposite entries on the two paths and hence changes only by negation on its support. All other marker rows remain fixed. The transported coordinates on the selected pair change by the same amount. For a row constant on each pair, with value d_s on pair s, the two mutations give

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

Indeed the incidences from the two selected columns to a fixed outside column are opposite and their magnitudes are $-D_{st}$, so their corrections sum to $([d_s]_+ - [-d_s]_+)|D_{st}| = -D_{st}d_s$. The resulting row is again constant on each pair. These calculations apply equally in the reverse word.

Remove the two reversed mutations and repeat until the word is empty, at which point the pulled-back principal and mixed data are those of the source. The endpoint of the remaining prefix has the same pair structure and the same named marker supports, with possible changes of signs. The degrees in each pair agree, and the remaining reverse steps are strict decreases. Thus the whole word decomposes into complete pair moves. Since every named support is preserved down to the source, the endpoint permutation preserves each individual pair and acts trivially on the four duplicated rows. The displayed action on their common pair values is exactly (5).

Conversely every simple row reflection is realized by its complete pair move, followed by swapping the two active labels to restore the principal matrix and the sign of its marker. Such swaps can be moved to the end, conjugating the intervening mutations by relabeling their indices. A word carrying one frame to the other therefore realizes the required endpoint. □

6. Restoring the entries between frozen indices

For each unordered pair of frozen indices f, g, add a separate four-vertex active component. Choose an order f, g for the pair. Its principal block and its only nonzero frozen rows are

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

Number the vertices $0, 1, 2, 3$ and mutate $0, 2, 1, 3$. The following table gives the six upper-triangular entries of the principal block in the order $01, 02, 03, 12, 13, 23$, the two frozen rows, and the cumulative change of b_{fg}. All entries follow from the mutation formula.

last mutation	principal entries	q_f	q_g	Δb_{fg}
0	$(-2, 0, 2, 2, -4, 2)$	$(-2, 2, 2, -2)$	$(0, 2, -2, 0)$	0
2	$(-2, 0, 2, -2, 0, -2)$	$(-2, 2, -2, 2)$	$(0, -2, 2, 0)$	4
1	$(2, -4, 2, 2, 0, -2)$	$(2, -2, -2, 2)$	$(0, 2, -2, 0)$	8
3	$(2, 0, -2, 2, 0, 2)$	$(2, -2, 2, -2)$	$(0, 2, -2, 0)$	8

Thus the word restores all data except the designated entry, which increases by 8, and its skew-symmetric counterpart. The inverse word decreases that entry by 8. The initial value of b_{fg} does not enter any correction at the four active vertices, so both identities hold at every initial value of that entry.

Use the same four-vertex data at both endpoints and set the frozen–frozen block to zero at both. In the main component all frozen incidences start as multiples of four and remain so. Every change of a frozen–frozen entry during a main-component word is therefore divisible by 16. The four-vertex words and their inverses can cancel all those changes independently.

Conversely, no active mutation can join distinct active components, and involutivity then prevents a component from splitting. The main component is the unique component with more than four vertices. Any successful endpoint permutation therefore preserves its vertex set. Mutations in other active components do not affect its principal or mixed data. Restricting a successful word to the main component now gives the endpoint of Proposition 5.1. This realizes (6) by an equivalence of full square matrices with specified frozen indices.

7. Making every index mutable

We use the following construction on two all-even skew-symmetric endpoints with specified frozen indices. At each frozen index f attach two oriented three-cycles of weight G_f, where G_f is an even integer at least two. Join only one vertex of each triangle to f by a single edge of magnitude G_f, with opposite signs for the two entries in row f. The triangles are otherwise disjoint from the old matrix and from one another. Their data are the same at both endpoints.

Proposition 7.1. *Suppose every successful ordinary endpoint permutation of the enlarged matrices preserves the sets of old active and old frozen indices. Then any ordinary equivalence gives an equivalence on the old matrices by mutating only old active indices, with the induced permutation on the old sets. Conversely such an equivalence lifts whenever its endpoint permutation extends to the attached triangles preserving their attachment data.*

Proof. Give every old index coordinate zero and every new index coordinate one. This column lies in the initial kernel. Each oriented triangle has row sums zero, and its contribution at its attachment index cancels that of the opposite triangle. The budgets are G_f on both triangles and at f, and zero on old active indices. Lemma 4.1 applies with all indices allowed.

At the enlarged target, sum the three equations of the kernel column over the vertices of either triangle at f. The internal terms cancel by the zero column sums. Indeed, for the triangle vertex set T and its attachment vertex h,

$$0 = \sum_{i\in T}(Bv)_i = \sum_{j\in T}\Big(\sum_{i\in T} b_{ij}\Big)v_j + b_{hf}v_f = \pm G_f v_f.$$

The same holds at a permuted target by pulling back the target triangles by the permutation, so all old frozen coordinates are zero there. Along a normalized successful word they stay zero by monotonicity.

Suppose the word has a first mutation at an old frozen index f. Before this step no mutation joins a triangle to the rest of the matrix except at its attachment index. A mutation inside a triangle cannot introduce an edge between two different old indices, since the triangle has just one old neighbor. Each triangle column is in the kernel of its own current principal block because the attachment coordinate is zero. Its attachment row from f has conserved charge by (7), so the two charges remain G_f and $-G_f$. On the triangle T with charge G_f, nonnegativity gives $\sum_{j\in T}[b_{fj}]_+ v_j \geq \sum_{j\in T} b_{fj}v_j = G_f$. Thus the positive incidence budget at f is at least G_f. Mutation at f would increase

its coordinate from zero, which is impossible. Deleting the triangle letters now gives the asserted old word. Conversely old active mutations leave all triangle data fixed, which gives the lift. □

To enforce the hypothesis of Proposition 7.1, use the gcd of each labelled row, taken to be zero if the row is zero. Mutation at its own index negates the row. At any other index k, the mutation negates its entry in column k and changes each remaining entry by an integer multiple of that entry. Thus the old gcd divides the new gcd, and the inverse mutation gives equality. It follows that any successful endpoint permutation must match row gcds.

For the endpoints of the preceding section let K_* be the largest absolute entry in their common active principal block. Each active row contains a nonzero principal entry, and therefore its gcd is positive and at most K_*. For each frozen index f choose a distinct odd prime $p_f > K_*$. Multiply the frozen row and column at f by $2p_f$, leaving the active scales equal to one. This positive diagonal scaling commutes with each active mutation: both terms of the correction at (i, j) receive the same factor, the product of the scales at i, j. It is injective, and the frozen indices are required to be fixed in the intermediate equivalence, so it preserves equivalence in both directions. The scaled row at f has gcd exactly $4p_f$ since the old entries were even and a four-vertex component supplies an incidence of magnitude 2 before scaling.

Apply Proposition 7.1 with $G_f = 4p_f^2$. The old active gcds are still at most K_*; the gcd at f is $4p_f$ and at each of its six triangle vertices it is $4p_f^2$. The three types are distinct, and the choice of primes makes the old frozen indices individually identifiable. The converse lift fixes the triangles pointwise. Consequently ordinary equivalence of the enlarged matrices is exactly the intermediate equivalence already constructed.

8. Fixed-source undecidability

Proof of Theorem 1.1. Given a word u in the generators of the fixed group G_0, construct the two endpoints above with $w = [x, u]$. The reduction is effective: the fixed parameters may be built into the algorithm, and the four variable rows are computed by multiplication of the fixed generator matrices. The source is fixed. If m is the number of reflection normals in Section 3, there are $4m$ vertices in the main component, $2m + 4$ frozen indices, four new active vertices for each unordered pair of frozen indices, and six new vertices for each formerly frozen index. Thus the final order is

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

The two matrices are equivalent exactly when $u = 1$ in G_0 by (6), Proposition 5.1, the correction of the frozen–frozen block, and Proposition 7.1. This proves recursively enumerable hardness for the fixed source. To semidecide equivalence, enumerate all finite mutation words and the finitely many final permutations, computing the matrix entries exactly and accepting on equality. □

9. Exact cover and strong NP-hardness

Proof of Theorem 1.2. We reduce from Exact Cover by 3-Sets, which is NP-complete [12, SP2]. The instance consists of a set E of size $3K$ and a list of m_0 three-element subsets of E. The question is whether K members of the list partition E. Send the empty universe to two zero matrices of order two. If some element appears in no triple, use the zero matrix and the matrix for a single edge of weight two on two vertices. In order two mutations and relabeling preserve the absolute value of the off-diagonal entry. For the remaining instances $K \geq 1$. Extend the original list by empty sets to obtain $E_1, \ldots, E_m$, where

$m = \max\{m_0, 2K+1, 7\}$. Write $\mathbf{1}_{\mathcal{P}}$ for the indicator of a condition $\mathcal{P}$, equal to one if it holds and zero otherwise. Set

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

9.1. **The two matrices.** The matrix B has control vertices $u_1, \dots, u_m$, two vertices p, q, and data vertices r_e for $e \in E$ and r_0 (the latter has a label distinct from all r_e). Its entries are specified by

$$\begin{aligned} b_{u_i p} &= -a, & b_{u_i q} &= a, & b_{pq} &= -a^2 K, \\ b_{r_e u_i} &= -2M\mathbf{1}_{e\in E_i}, & b_{r_e p} &= 0, & b_{r_e q} &= 2aM, \\ b_{r_0 u_i} &= -2M, & b_{r_0 p} &= 0, & b_{r_0 q} &= 2aKM. \end{aligned}$$

The reverse entries follow by skew symmetry, and all remaining entries are zero. To replace p by four copies, replace it with mutually nonadjacent vertices $p_1, \dots, p_4$ and give each of them half of each incidence of p. Call the clones and q the five outputs. Write X_0 for the result applied to B, and Y_0 for the result applied after exchanging p and q in B.

In each of X_0, Y_0, attach to each data vertex two separate directed three-cycles of edge weight G. Join the data vertex to just one vertex of each cycle, with entry G in the row of the data vertex for the first cycle and $-G$ for the second. There are no other edges incident to the new vertices. Call the resulting matrices $X = (x_{ij})$ and $Y = (y_{ij})$. Their order and maximum entry magnitude satisfy

$$N = m + 21K + 12, \qquad \max_{i,j}\{|x_{ij}|, |y_{ij}|\} \le \max\{a^2K/2, 2aKM, G\}.$$

All entries are even. These bounds give polynomial output size in either encoding.

The row gcd is invariant under mutation as shown in Section 7. At both endpoints the row greatest common divisors are

vertices	u_i	p_j	q	$r_e,\ r_0$	new cycle vertices
row greatest common divisor	2	$a/2$	a	$2M$	G.

In a control row, the clone and neutral-row incidences give $\gcd(a/2, 2M) = 2$, since $a/2$ is even and $2M = 4(a/2) + 2$. All entries are even, giving the exact gcd. Each data row has an entry of magnitude $2M$, and the displayed formulas show all its entries are multiples of $2M$. For the other classes a control incidence or an internal cycle incidence attains the indicated value, which divides all entries in the row. The five values are distinct. Thus the final permutation of any successful word preserves these five vertex classes.

9.2. **Eliminating data mutations.** Use Proposition 7.1 with the data vertices as the frozen indices and weight G at each attachment. The five row gcds verify its permutation hypothesis. Every successful word from X to Y therefore gives a word from X_0 to a class-preserving permutation of Y_0 using only controls and outputs. Conversely a word on controls and outputs taking X_0 exactly to Y_0 lifts with all attached triangles fixed.

9.3. **Growth away from the control and output cubes.** We use a multipartite version of the growth argument for forks [17, Section 2]. We apply it to words on the principal vertices, meaning the controls and the five outputs. The weight of an arrow is its positive matrix entry. Partition the vertices into classes having no edges within them, with all pairs in different classes adjacent by edges of weight at least two. One class is partitioned into D and U, where $D \neq \varnothing$. The other classes form two nonempty ordered lists, denoted $O_1, \dots, O_s$ and $I_1, \dots, I_t$. Outside D, arrows point forward in the block order

$$O_1, \dots, O_s, U, I_1, \dots, I_t,$$

omitting U when empty. The arrows incident with D point from D to each O_j and from each I_j to D. Require that for every $d \in D$, $o \in \bigcup_j O_j$, and $i \in \bigcup_j I_j$, the weight of $o \to i$ exceed both weights of $d \to o$ and $i \to d$.

A mutation outside D weakly increases all absolute entries, strictly increases their sum over unordered vertex pairs, and preserves the described form. If the mutated vertex belongs to U, its only directed two-step paths go from an outgoing class to an incoming class. Mutation increases each such weight by the product of the two incident weights. This product exceeds each factor, and the mutated vertex moves from U to D. If the mutated vertex k lies in O_j, its incoming vertices lie in $D, O_1, \dots, O_{j-1}$; its outgoing vertices lie in $O_{j+1}, \dots, O_s, U, I_1, \dots, I_t$. These two lists become the new ordered lists on the two sides, the new return set is $\{k\}$, and the new unused part of its class is $O_j \setminus \{k\}$. Along each directed two-step path through k, the endpoint edge increases by the product if it is missing or points forward. The product is at least twice the larger of the incident weights, so the new cross inequality is strict. An opposite edge occurs only on a path $d \to k \to i$ with $d \in D$, $i \in \bigcup_j I_j$. If the path weights are x, y and the opposite edge has weight z, then $y > x, z$ and $x \geq 2$. The new forward weight satisfies $xy - z > (x-1)y \geq y > \max\{x, z\}$. These inequalities give the new cross inequalities and strict growth, as there is at least one such path. Edges from the old D to the old U are created if U was nonempty. The case $k \in I_j$ follows by reversing all arrows and both block orders, which commutes with mutation.

Start at a matrix of this form and take a mutation in an O_j or I_j class. Until the path returns to its starting matrix, it can be represented by a nonempty word of the growth moves just described, with the first commuting block a subset of that same class. Maintain a list of blocks, each a nonempty subset of one ordered class at the state before the block. Its first mutation starts a new return set, and the other mutations add members of that class to it. The last block therefore equals the current return set. A mutation at an unused vertex extends the last block, and one in an ordered class starts a new block. If the next letter belongs to the return set, commute its occurrence in the last block to the end and cancel the two equal letters. All members of the block are mutually nonadjacent throughout it. A surviving subset can be replayed in any order: its first move starts a return set, and the others are unused vertices. If the block empties, its starting state and the preceding block are recovered. The first block stays in the initial class until it empties, at which point the principal matrix has returned to the starting matrix. Each cancellation and commutation is an identity on the entire matrix, including the data rows. In particular the excursion up to its first return acts as the identity on these rows.

A subset $S \subseteq \{1, \dots, m\}$ of control mutations commutes at X_0. On its principal matrix the controls stay independent. For $i \in S$ the entries from u_i to each clone and to q are $a/2, -a$, respectively; for $i \notin S$ they are $-a/2, a$. The clone-to-q entry is $-a^2(K - |S|)/2$, since each selected control adds $(a/2)a = a^2/2$ through the path from a clone to q. The principal states indexed by these subsets form the first control cube, in which mutating a control toggles its membership in the subset. The second has the signs reversed at corresponding principal matrices, beginning at Y_0. Mutating all the controls in a specified subset means mutating each once.

If $|S| \neq K$, the principal matrix has the growth form above. When $|S| < K$, take $D = \{u_i : i \notin S\}$, $U = \{u_i : i \in S\}$, the O-class $\{q\}$ and the I-class of four clones. When $|S| > K$, swap the two groups of controls and interchange the output classes. The cross weight $a^2|K - |S||/2 > a$ gives the strict inequalities. Mutation at any output begins an excursion. For S proper and nonempty the first move creates every edge between the controls in S and its complement. The independent classes then all have size less than m, and the growth representative preserves that bound since no absolute entry decreases. A control-cube state has an independent set of size m.

If S is empty or full, put $d = |K - |S||$. A single clone mutation gives total arrow weight $ma^3d/4 + ma + 2a^2d$. Indeed, the output edges contribute $2a^2d$, the clone incidences with the controls contribute $2ma$, and each path through the selected clone contributes $a^3d/4$ against the old q-control edge of weight a. Mutating q first instead contributes $a^3d/2$

against each old clone-control edge of weight $a/2$, giving $4m(a^3d/2-a/2)+ma+2a^2d$, which is larger. The maximum total weight in either control cube is at most $3am + 2a^2(m-1)$. Subtracting this bound from the first expression gives at least $m(a^3/4-2a-2a^2)+4a^2 > 0$. Thus in these cases too, an excursion must return to its starting matrix before reaching a control cube.

If $|S| = K$, all five outputs are nonadjacent and form a commuting cube. Mutating r of the four clones, with $0 \le r \le 4$, and mutating q exactly when $e = 1$, with $e \in \{0, 1\}$, creates a common cross-control weight $|r - 4e|a^2/4$. In the first control cube each mutated clone adds $a^2/4$ to the entry from a selected to an unselected control, and mutating q subtracts a^2. All output incidences with controls reverse precisely at the mutated outputs. The two ends $(r, e) = (0, 0), (4, 1)$ belong to opposite control cubes. At an intermediate state, the two control groups are the O- and I-classes, ordered in the direction of the cross-control arrows. The return set consists of the outputs pointing towards the O-class, and the other outputs form U. When $r - 4e > 0$, the return set consists of the mutated clones and unmutated q. When $r - 4e < 0$, it consists of the unmutated clones and mutated q. Both control classes are nonempty since $0 < K < m$; the return set is nonempty and cross dominance holds because $a^2/4 > a$. A control move begins a growth excursion with independent classes of size at most $\max\{K, m-K, 5\} < m$. It cannot reach a control cube before returning to the intermediate state. Reversing arrows gives all these assertions for the second control cube.

Track a successful word from its initial control cube. At a control-cube state, stay in that cube for a control mutation or enter its attached output cube at $|S| = K$ for an output mutation. In an output cube track the subset of outputs mutated so far. The endpoint has m independent controls, so it cannot occur in the interior of an output cube. Each move outside the tracked cubes begins an excursion that must return to its entrance parent before the endpoint. Erase it using the block argument. This gives a word in control cubes and output cubes. Within a cube, commuting and canceling letters replaces any segment by the symmetric difference of its two endpoint subsets, on the full matrix. An output segment returning to its entrance corner thus erases, and a segment crossing to the opposite control cube uses all five outputs exactly once. Wrap each such traversal by the same control subset S on both sides, so it runs between the two control bases. Between traversals the inserted subsets cancel against the actual control movements by symmetric difference. A single residual control subset remains at the end. With an even number of traversals, matching the target incidences would require this subset to be all controls, but its clone-to-q weight would be $a^2(m-K)/2$, which differs from the required $a^2K/2$. With an odd number the residual subset must be empty. Thus any successful word has a representative made of an odd number of wrapped traversals.

9.4. **The exact cover condition.** Use the four-copy construction on the full matrix B. A mutation at an index other than p commutes with this construction: corrections involving a copy have one factor halved, and paths between copies through the mutated vertex have opposite signs and contribute zero. For mutation at p, the four copies are mutually nonadjacent, so their mutations commute. Their corrections to any pair outside the copies add up to the original correction at p, since each contributes one quarter. They reverse all half-incidences. Thus a wrapped traversal can be computed on the full matrix before copying.

Let σ exchange the labels p, q before copying. At a control subset S of size K, the outputs are nonadjacent. Mutating them both is an involution and their two mutations commute. Conjugation by σ exchanges their order and fixes every control mutation. The wrapped word is the conjugate of these two output mutations by the control word. Composing it with σ gives an involution returning to the first principal base. After any number of traversals the normalized matrix is the unexpanded endpoint with σ applied if

and only if the number is odd. This follows from equivariance and the commutation with σ at the two bases.

Write b_{ij} for the entries at the current normalized base. For any data vertex r, put

$$U_i = -b_{ru_i}, \qquad T = b_{rp} + b_{rq},$$
$$\ell_R = T - a\sum_{i\in R} U_i \qquad \big(R \subseteq \{1,\dots,m\},\ |R| = K\big).$$

After the first control subset S in the wrapped traversal, the two output entries of this row are

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

Here $P + Q = \ell_S$. For selected controls the output-to-control entries are $-a, a$ from p, q, respectively, so their combined correction in row r is $a([Q]_+ - [-P]_+)$. For unselected controls it is $a([P]_+ - [-Q]_+)$. Set $g = [P]_+ - [-Q]_+$, so the first correction is $a(\ell_S - g)$ and the second is ag. Undoing the control subset therefore gives

$$U_i' = U_i + a\ell_S \mathbf{1}_{i\in S} - ag.$$

The output entries after this undoing and exchange are $-Q + a\sum_{i\in S}[U_i']_+$ and $-P - a\sum_{i\in S}[-U_i']_+$. Their sum is $-\ell_S + a\sum_{i\in S} U_i'$. Therefore

$$\ell_R' = -\ell_S + a\sum_{i\in S} U_i' - a\sum_{i\in R} U_i'.$$

The terms in g cancel because $|S| = |R|$. Substituting $a\sum_{i\in S} U_i = T - \ell_S$ and $a\sum_{i\in R} U_i = T - \ell_R$ gives

$$\ell_R' = \ell_R + \big(a^2(K - |R\cap S|) - 2\big)\ell_S.$$

In particular $\ell_S' = -\ell_S$. This calculation applies for every sign of every row entry.

For each $e \in E$, the initial values of $\ell_R/(2aM)$ on row r_e are $1 - \sum_{i\in R} \mathbf{1}_{e\in E_i}$, all in $[-K, K]$. The row r_0 has every defect zero. The normalized values remain integral under the displayed action. For $R \neq S$, the coefficient of ℓ_S in the correction to ℓ_R is at least $a^2 - 2 > 2K$. If a vector of normalized defects starts in $[-K, K]$ and its coordinate at S is nonzero, the move at S makes every other coordinate have that coordinate's old sign and magnitude at least $a^2 - 2 - K > K$. Coordinate S now has the opposite sign and smaller magnitude. At the next distinct letter, its coordinate has the common sign, with magnitude larger than the exceptional magnitude. After the move it becomes the exceptional coordinate, and all other coordinates have the common sign and magnitude larger than the new exceptional magnitude. To see this, multiply all coordinates by their common sign off the exceptional index. If the old exceptional coordinate is $-t$ and the new letter has value $x > t$, then the old exceptional coordinate becomes at least $-t + (a^2 - 2)x > x$. Every other unchanged-sign coordinate gains at least $(a^2 - 2)x$. Thus any reduced word after its first active letter remains outside the box.

Canceling equal consecutive normalized traversals preserves the endpoint and odd parity. If there is no exact cover, the first remaining letter S has a nonzero defect on at least one constraint row. Indeed, a joint zero requires every element to occur exactly once in the chosen K sets, which also excludes empty padded sets. The reduced word sends that row outside the box. At a permitted endpoint, exchanging the output labels back to the first principal base gives the original data up to permutations of controls and data rows. Such permutations keep all normalized defects in $[-K, K]$. This proves the NO direction.

If there is an exact cover S, after its control mutations the columns of p and q are opposite on every other old vertex, since $P+Q = 0$ on all data rows and the principal incidences are opposite. The outputs are nonadjacent, so mutating both cancels their corrections to every other entry and exchanges their labels. The control mutations commute with that exchange, so undoing the control subset gives σB. After copying p this is Y_0, and the

attached triangles are unchanged. Thus mutating the controls of S, the four copies and q, and then the controls of S in reverse order takes X to Y. Its length is $2K + 5$. Together with the order and magnitude bounds this proves the theorem. □

References

[1] 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.

[2] 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.

[3] S. Fomin, M. Shapiro, and D. Thurston, *Cluster algebras and triangulated surfaces. Part I: Cluster complexes*, Acta Math. **201** (2008), 83–146. arXiv:math/0608367.

[4] 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.

[5] S. Fomin, *Quiver mutations*, Open Problems in Algebraic Combinatorics, University of Minnesota, Minneapolis, May 16, 2022.

[6] S. Fomin, K. Igusa, and K. Lee, *Universal quivers*, Algebr. Comb. **4** (2021), 683–702. arXiv:2003.01244.

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

[8] S. Fomin, L. Williams, and A. Zelevinsky, *Introduction to Cluster Algebras*, Chapters 1–3, arXiv:1608.05735v1, 2016.

[9] S. Fomin, L. Williams, and A. Zelevinsky, *Introduction to Cluster Algebras*, Chapters 1–3, arXiv:1608.05735v5, 2024.

[10] S. Fomin and A. Zelevinsky, *Cluster algebras I: Foundations*, J. Amer. Math. Soc. **15** (2002), 497–529. doi:10.1090/S0894-0347-01-00385-X.

[11] S. Fomin and A. Zelevinsky, *Cluster algebras II: Finite type classification*, Invent. Math. **154** (2003), 63–121. arXiv:math/0208229.

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

[13] G. Higman, *Subgroups of finitely presented groups*, Proc. Roy. Soc. London Ser. A **262** (1961), 455–475. doi:10.1098/rspa.1961.0132.

[14] J.-P. Serre, *Trees*, translated by J. Stillwell, Springer, Berlin, 1980.

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

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

[17] M. Warkentin, *Exchange Graphs via Quiver Mutation*, Ph.D. thesis, Technische Universität Chemnitz, 2014. Available from the German National Library.

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