hexagon:2611.00163v2

DECIDABILITY OF MUTATION-ACYCLICITY

QIYUE TANG

Abstract. Mutation-acyclicity is decidable for all finite integral skew-symmetrizable exchange matrices. Mutation equivalence up to relabeling is decidable when at least one input is mutation-acyclic.

1. Introduction

An exchange matrix is an integral $n \times n$ matrix $B = (b_{ij})$, with $n \geq 0$, that admits a *symmetrizer*: a diagonal matrix $D = \operatorname{diag}(d_1, \ldots, d_n)$ with positive integer entries such that $d_i b_{ij} = -d_j b_{ji}$ for all i, j. Such a matrix is called skew-symmetrizable. Its directed graph has an arrow $i \to j$ when $b_{ij} > 0$. If B itself is skew-symmetric, we take b_{ij} parallel arrows in this direction and obtain a quiver. We call B acyclic if its directed graph has no directed cycle.

For an integer x, put $[x]_+ = \max(x, 0)$. For $1 \leq k \leq n$, mutation μ_k at index k sends B to the matrix with entries

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

Applying μ_k twice negates the incident entries twice and subtracts the correction made to each other entry. To check that D remains a symmetrizer, use $d_i[b_{ik}]_+ = d_k[-b_{ki}]_+$ and $d_k[b_{kj}]_+ = d_j[-b_{jk}]_+$ and the identities with all signs reversed. They show that the corrections to opposite entries are negatives after multiplication by their row weights [16]. A matrix is *mutation-acyclic* if a finite sequence of mutations, possibly empty, takes it to an acyclic matrix. Otherwise it is *mutation-cyclic*. Two matrices are *mutation equivalent up to relabeling* if a sequence of mutations and a simultaneous permutation of rows and columns takes one to the other. Relabeling commutes with mutation after relabeling its index, so a single permutation at the end suffices.

Fomin and Zelevinsky introduced matrix mutation in their definition of cluster algebras [16]. A coefficient-free seed of size n consists of an exchange matrix B and an ordered transcendence basis $(x_1, \ldots, x_n)$ of a rational function field in n indeterminates over $\mathbb{Q}$. The basis is called its cluster. Mutation at k updates the matrix and replaces x_k by

$$\frac{\prod_{i=1}^n x_i^{[b_{ik}]_+} + \prod_{i=1}^n x_i^{[-b_{ik}]_+}}{x_k},$$

with empty products equal to one. The variables in all the resulting clusters generate the cluster algebra. The finite-type classification gives acyclic seeds of Dynkin type for the algebras with finitely many cluster variables [5]. Beyond finite type, an acyclic quiver has a hereditary path algebra, and its cluster category realizes mutations as exchanges of indecomposable summands [13, 12, 14]. The recognition problem asks whether an acyclic seed exists when the given seed has directed cycles.

Assem, Blais, Brüstle, and Samson [2] decided the first nontrivial case, that of skew-symmetric matrices on three vertices, by reducing arrow multiplicities under mutation. For

Date: November 7, 2026.

2020 *Mathematics Subject Classification.* Primary 13F60; Secondary 16G20.

Key words and phrases. exchange matrix, quiver mutation, mutation-acyclicity, quasitilted algebra, folding.

a directed triangle with multiplicities x, y, z, Beineke, Brüstle, and Hille [3, Theorem 1.2] showed that the triangle is mutation-cyclic precisely when $x, y, z \geq 2$ and $x^2 + y^2 + z^2 - xyz \leq 4$. Akagi [1, Theorem 1.2] gave the corresponding Markov-constant criterion for skew-symmetrizable matrices on three vertices. In arbitrary size, Buan, Marsh, and Reiten [12, Corollary 5.3] showed that full subquivers inherit mutation-acyclicity. A full subquiver retains all arrows between the selected vertices, so finding a mutation-cyclic full subquiver gives an obstruction. Neville [6] extended the use of the Markov invariant to mutation-acyclic quivers of arbitrary size.

Soukup explicitly posed algorithmic recognition for arbitrary quivers [7, Section 2.4.5]. For four vertices, Armstrong-Williams, Hirst, Jackson, and Lee [8, Theorem 1.1 and Corollaries 5.2–5.3] proved that a mutation-acyclic quiver of maximum arrow multiplicity h has a path to an acyclic quiver along which the multiplicities stay at most $h^4 + 4h^3 + 6h^2$. This bounds a finite search.

The two-input problem asks whether given matrices lie in the same mutation class. Fomin, Williams, and Zelevinsky posed the unrestricted problem for quivers and for skew-symmetrizable matrices [4, Problems 2.6.14 and 2.8.2]. For two acyclic quivers, the description of acyclic seeds by Caldero and Keller [14, Corollary 4] reduces comparison to mutations at sources and sinks. Liu and Pan [21, Theorem 3] give the matrix version used here. Under a one-sided promise, recognition of the other input lets this comparison apply. Armstrong-Williams, Hirst, Jackson, and Lee combined recognition and source–sink comparison for four vertices [8, Corollary 5.7].

Theorem 1.1. *Mutation-acyclicity is decidable for integral skew-symmetrizable matrices of arbitrary finite size.*

Theorem 1.2. *There is an algorithm that halts on every pair of integral skew-symmetrizable matrices. It reports when both are mutation-cyclic and otherwise decides whether they are mutation equivalent up to relabeling.*

For skew-symmetric inputs, a dimension bound gives a finite list of algebras whose relation-extension quivers can realize the input. A spectral test and finite-dimensional module witnesses then give a search that terminates on mutation-cyclic inputs. A finite folding reduction handles general skew-symmetrizable matrices.

Mutations act separately on the connected components of the underlying undirected graph. Every mutation term across two components vanishes, and involutivity prevents a connected component from splitting. Thus mutation-acyclicity is equivalent to mutation-acyclicity of every component. The empty matrix is acyclic.

2. Recognizing mutation-acyclic quivers

Proposition 2.1. *There is an algorithm that decides whether a finite integral skew-symmetric exchange matrix is mutation-acyclic.*

It suffices to treat connected quivers. Enumerating finite mutation sequences already gives a procedure that halts with a positive answer on every mutation-acyclic input. We construct a second procedure that halts with a negative answer on every other input.

Throughout this section $K = \overline{\mathbb{Q}}$, the computable algebraic closure of the rational numbers. Algebras are finite-dimensional associative unital K-algebras, and modules are finite-dimensional right modules unless specified otherwise. For an algebra R, write $\operatorname{rad} R$ for its Jacobson radical, the ideal annihilating every simple right module, and $DR = \operatorname{Hom}_K(R, K)$ for the linear dual of its regular bimodule, with $(afb)(x) = f(bxa)$ for $a, b, x \in R$ and $f \in DR$. The trace tr below is the sum of the diagonal entries of an endomorphism in any basis. An algebra is basic if its semisimple quotient is a product of copies of K. For modules M, N, the space $\operatorname{Ext}^q_R(M, N)$ is the degree-q right derived functor of $\operatorname{Hom}_R(M, N)$, computed from a projective resolution of M. For a basic algebra with simple right modules $S_1, \ldots, S_n$, we take the ordinary quiver to have $\dim_K \operatorname{Ext}^1_R(S_i, S_j)$

arrows from i to j. The relation extension of R is the algebra on $R \oplus \operatorname{Ext}^2_R(DR, R)$ with multiplication, for $r, s \in R$ and $\xi, \eta \in \operatorname{Ext}^2_R(DR, R)$,

$$(r, \xi)(s, \eta) = (rs, r\eta + \xi s).$$

The left action on $\operatorname{Ext}^2_R(DR, R)$ is induced by left multiplication on R, and the right action is induced contravariantly by the left action on DR.

2.1. **A finite list of candidates.** Let Q be a connected input quiver on $n \geq 1$ vertices and let h be its maximum arrow multiplicity, with $h = 0$ when it has no arrows. If Q is mutation-acyclic, the exchange theorem for cluster categories realizes Q as the ordinary quiver of a cluster-tilted algebra of a connected hereditary path algebra [12, 13]. The relation extension theorem of Assem, Brüstle and Schiffler [9, Theorem 3.4] then gives a basic tilted algebra R whose relation extension has quiver Q. Here a tilted algebra is the endomorphism algebra of a tilting module over a finite-dimensional hereditary algebra. Its ordinary quiver is acyclic, and is a subquiver of the relation-extension quiver [9, Theorem 2.6]. The algebra R is connected. Indeed, a nontrivial product has relation extension the product of the relation extensions of its factors, since extensions between modules supported on distinct factors vanish. The tilted algebra furnished by the theorem has global dimension at most two. It is a quotient of the path algebra of its ordinary quiver. In a topological ordering, a path of length ℓ uses $\ell + 1$ increasing vertices and has at most h^ℓ choices of arrows for any given vertex sequence. Consequently

$$\dim_K R \;\leq\; \sum_{\ell=0}^{n-1} \binom{n}{\ell+1} h^\ell, \tag{1}$$

where the term for $\ell = 0$ counts the n stationary paths. Write $HH^q(R)$ for the degree-q Hochschild cohomology with coefficients in the regular bimodule. A path algebra has a projective bimodule resolution of length one. For clarity, let P be such an algebra, V the subalgebra spanned by the stationary paths, and W the V-bimodule spanned by the arrows. The resolution is

$$0 \longrightarrow P \otimes_V W \otimes_V P \longrightarrow P \otimes_V P \longrightarrow P \longrightarrow 0.$$

The right map multiplies the factors; the preceding map sends $p \otimes w \otimes q$ to $pw \otimes q - p \otimes wq$, for $p, q \in P$ and $w \in W$. For each path, the tensors in the middle term represent its cuts at vertices. The tensors to the left represent cuts at arrows, and their images are differences of consecutive vertex cuts. These differences are independent and span the kernel of multiplication on that path. Both tensor terms are projective bimodules, since the stationary idempotents decompose them into sums of tensor products of principal left and right projectives. Thus $HH^2(P) = 0$. The tilted algebra R is derived equivalent to a hereditary path algebra, and Hochschild cohomology is preserved under this equivalence [19].

Lemma 2.2. *For each positive integer d, one can compute a complete list, up to K-algebra isomorphism, of d-dimensional algebras R with $HH^2(R) = 0$.*

Proof. Fix the first vector e_1 of the standard basis of K^d as the unit. The associative laws with that unit form an affine algebraic set, defined by the equations for the structure constants of a bilinear multiplication. The group of invertible linear maps fixing e_1 acts by change of basis, and its orbits are the isomorphism classes.

For each law R, the normalized Hochschild cochains in degree q are the K-multilinear maps $R^q \to R$ vanishing when an argument is the unit. Denote the cochain space by C^q and its differential by $\delta^q : C^q \to C^{q+1}$. For a linear map u with $u(1) = 0$ and a normalized bilinear map v, the first two differentials are

$$(\delta^1 u)(a, b) = au(b) - u(ab) + u(a)b,$$
$$(\delta^2 v)(a, b, c) = av(b, c) - v(ab, c) + v(a, bc) - v(a, b)c \qquad (a, b, c \in R).$$

Indeed, a perturbation of the product by ϵv over $K[\epsilon]/(\epsilon^2)$, preserves the unit exactly when v is normalized. Expanding the two associations of abc and equating the coefficients of ϵ gives the normalized cocycles $\ker \delta^2$. Transporting the product by the basis change $1 + \epsilon u$ gives the coboundaries $\operatorname{im} \delta^1$, since the inverse change is $1 - \epsilon u$ and substitution gives $\delta^1 u$ up to sign. In characteristic zero the stabilizer is smooth and the orbit map has differential of rank equal to the orbit dimension. If $HH^2(R) = 0$, the orbit dimension thus equals the dimension of the linearized space. The local dimension of the algebraic set lies between the orbit dimension and the dimension of the Zariski tangent space, which is bounded above by the dimension of the linearized space. Hence the point is smooth and belongs to a unique irreducible component. The orbit is locally closed and has full dimension in that component, so it contains a neighborhood of the point. Translating this neighborhood by changes of basis shows that the orbit is open. Two different open orbits cannot meet the same irreducible component. There are only finitely many components, so there are finitely many isomorphism classes satisfying $HH^2(R) = 0$.

The condition $HH^2(R) = 0$ is the polynomial rank test

$$\operatorname{rank} \delta^1 + \operatorname{rank} \delta^2 = \dim_K C^2.$$

Each possible rank is expressed using minors. Isomorphism to a specified algebra is a polynomial condition with existentially quantified entries of an invertible matrix fixing e_1. We use quantifier elimination for algebraically closed fields in its effective form. To eliminate an existential quantifier, write the quantifier-free formula as a disjunction of conjunctions of polynomial equations and inequations. In each conjunction take the product of the inequation polynomials, with empty product one. In the quantified variable, branch on the vanishing of leading coefficients and compute the greatest common divisor of the nonzero equation polynomials by the Euclidean algorithm. If all equations vanish identically, the conjunction has a solution exactly when the product of the inequation polynomials is nonzero as a polynomial. Otherwise take the square-free part of the greatest common divisor by dividing it by its greatest common divisor with its derivative. There is a solution exactly when this square-free part fails to divide the product of inequation polynomials. In particular a constant greatest common divisor gives no solution. Branching makes these conditions polynomial conditions in the remaining variables. Negation handles universal quantifiers. Given a list of pairwise nonisomorphic representatives, effective quantifier elimination for algebraically closed fields decides whether there is a law satisfying the rank test and isomorphic to none of those representatives. If the answer is yes, enumerate tuples of algebraic numbers, testing that same condition on each tuple, until a new representative is found. If the answer is no, return the list. Finiteness of the isomorphism classes makes this process halt. □

Apply Lemma 2.2 to every dimension in the range (1). Retain the basic connected algebras of global dimension at most two with n simple modules, acyclic ordinary quiver, and relation-extension quiver isomorphic to Q.

For a finite-dimensional algebra R, let L_x be left multiplication by $x \in R$ on its underlying vector space. The Jacobson radical is the kernel of the bilinear form $(x, y) \mapsto \operatorname{tr} L_{xy}$. Indeed, the radical pairs to zero because xy is nilpotent when x is in the radical. The form on the semisimple quotient is a sum of matrix trace pairings, weighted by the multiplicities of the simple left modules in a composition series of the left regular module. Each multiplicity is positive since the regular module surjects onto the semisimple quotient. The weighted trace pairings are nondegenerate in characteristic zero. This gives the radical by linear algebra. The algebra is basic exactly when its semisimple quotient is commutative. In that case the quotient is a product of as many copies of K as its dimension. The corresponding one-dimensional simple modules are obtained by solving the unital character equations, one solution for each factor. Enumerate algebraic-number tuples satisfying these equations until the number of distinct solutions equals the dimension of the quotient.

For right modules M, N, the ordinary bar resolution of M in degree $q \geq 0$ is $M \otimes_K R^{\otimes_K q} \otimes_K R$, with the free right action on the last factor. The augmentation sends $m \otimes r_1$ to mr_1 for $m \in M$ and $r_1 \in R$. In degree $q \geq 1$ its differential sends $m \otimes r_1 \otimes \cdots \otimes r_{q+1}$, with $m \in M$ and $r_i \in R$, to

$$mr_1 \otimes r_2 \otimes \cdots \otimes r_{q+1} + \sum_{i=1}^{q} (-1)^i m \otimes r_1 \otimes \cdots \otimes (r_i r_{i+1}) \otimes \cdots \otimes r_{q+1}.$$

As a complex of vector spaces it contracts by appending the unit, with sign $(-1)^{q+1}$ in degree q and sign $+1$ from M to degree zero. Substitution cancels all terms except the identity, proving exactness. Each term is free as a right module. Applying $\mathrm{Hom}_R(-, N)$ gives a cochain complex whose cohomology in any fixed degree is computed from just the incoming and outgoing differentials in that degree. These are finite matrices with entries polynomial in the algebra and module structure constants. This computes the ordinary quiver using $M = S_i$, $N = S_j$, and $q = 1$. Connectivity and acyclicity are graph tests. With $S = R/\operatorname{rad} R$ considered as a right module, the global dimension is at most two exactly when $\mathrm{Ext}^3_R(S, S) = 0$. In a minimal projective resolution, applying $\mathrm{Hom}_R(-, S)$ makes the differentials zero, and each nonzero projective term has nonzero top. Thus the vanishing test forces the third projective term of every simple to vanish. The bound on the projective dimensions of the simples extends to all modules by composition series. For the relation extension, compute $\mathrm{Ext}^2_R(DR, R)$ as a kernel modulo an image in the bar cochain complex. On these cochains the left action applies left multiplication to the output, and the right action precomposes with left multiplication on the DR input. These operations commute with the differential and give the two actions on the quotient. Use them to form the multiplication table of the relation extension. Compute its quiver by the same radical and Ext^1 calculations and compare to Q over all vertex bijections.

2.2. **A test on Euler forms.** For each remaining algebra R, define its Euler matrix E_R by

$$(E_R)_{ij} = \sum_{q=0}^{2} (-1)^q \dim_K \mathrm{Ext}^q_R(S_i, S_j).$$

It represents the Euler pairing on the Grothendieck group of finite-dimensional modules in the simple basis. Its pairing between the classes of the indecomposable projectives and the simple classes is the identity matrix, so E_R is invertible. Write $\rho(-E_R^{-1}E_R^{\mathsf{T}})$ for the maximum modulus of the complex eigenvalues of the indicated Coxeter matrix. Retain R only if

$$\rho(-E_R^{-1}E_R^{\mathsf{T}}) > 1 \quad \text{or} \quad E_R + E_R^{\mathsf{T}} \text{ is positive semidefinite of corank at most one.} \tag{2}$$

The roots of the Coxeter characteristic polynomial are algebraic numbers. For each root λ, compare $\lambda\overline{\lambda}$ with one by exact real algebraic arithmetic. Rational congruence diagonalization gives the signs and number of zero entries of the symmetrized form. Derived equivalence induces an integral change of basis preserving the Euler pairing. Under $E_R \mapsto T^{\mathsf{T}} E_R T$ for an invertible integral change-of-basis matrix T, the symmetrized form changes by congruence and the Coxeter matrix changes by conjugation by T.

For a connected acyclic path algebra, order the vertices so that each arrow goes from an earlier vertex to a later one. Let N be the strictly upper triangular matrix of arrow multiplicities. The Euler matrix is $I - N$, where I is the identity matrix, since the path algebra is hereditary and the first extensions of simples count arrows. If $2I - N - N^{\mathsf{T}}$ is positive semidefinite, its corank is at most one. For more than one vertex, $N + N^{\mathsf{T}}$ is irreducible and nonnegative by connectedness. If 2 is an eigenvalue, it is the largest eigenvalue and is simple by Perron–Frobenius. The one-vertex form is positive definite.

Otherwise $N + N^{\mathsf{T}}$ has an eigenvalue greater than 2. For real $t \geq 1$, the spectral radius of the nonnegative matrix

$$\frac{tN + N^{\mathsf{T}}}{t+1}$$

is continuous in t. At $t = 1$ it exceeds one, and as t tends to infinity it tends to zero, the spectral radius of the nilpotent matrix N. At some $t > 1$ the radius is one. By Perron–Frobenius, one is then an eigenvalue. Hence $t(I - N) + (I - N^{\mathsf{T}})$ is singular, so t is an eigenvalue of $-(I - N)^{-1}(I - N^{\mathsf{T}})$. Derived invariance now gives (2) for every algebra derived equivalent to a connected acyclic path algebra.

A quasitilted algebra is the endomorphism algebra of a tilting object in a hereditary Ext-finite abelian K-category. Happel's classification [17, Theorem 3.1] implies that a connected quasitilted algebra over an algebraically closed field is derived equivalent to a connected hereditary path algebra or to a canonical algebra. A canonical algebra has an integer number $t \geq 2$ of directed arms from a common source to a common sink. Arm i has $p_i \geq 2$ edges, and the $t - 2$ minimal relations are linear combinations of the full arm paths. Write E for its Euler matrix. Its diagonal entries are one, arrows contribute -1 in their source–target positions, and the relations contribute $t-2$ in the common source–sink position. In a minimal projective resolution of a simple, the outgoing arrows index the projectives in degree one, and the minimal relations starting at its vertex index those in degree two [9, Lemma 2.2].

For the Coxeter polynomial, let z be an indeterminate and consider $zE + E^{\mathsf{T}}$. The block of internal vertices of each arm is tridiagonal with diagonal $z + 1$, upper diagonal $-z$, and lower diagonal -1. The determinant D_r of an r-vertex block satisfies $D_r = (z + 1)D_{r-1} - zD_{r-2}$ for $r \geq 2$, by expansion at an end, with $D_0 = 1$ and $D_1 = z + 1$. Induction gives $D_r = 1 + z + \cdots + z^r$. Put $f_i = D_{p_i-1}$. In the inverse of the internal block of arm i, the two end diagonal entries are D_{p_i-2}/f_i, and the first-to-last and last-to-first entries are z^{p_i-2}/f_i and $1/f_i$. The diagonal cofactors are the smaller tridiagonal determinants, and the off-diagonal cofactors are products of the respective adjacent entries with the cofactor sign. These formulas also apply when the internal block has size one.

Eliminate the internal blocks over the field of rational functions in z. Each arm subtracts $(f_i - 1)/f_i$ from both endpoint diagonal entries and subtracts z^{p_i}/f_i and $1/f_i$ from the upper and lower off-diagonal entries, respectively. This uses $zD_{p_i-2} = f_i - 1$ and the two end entries of the inverse just computed. The original off-diagonal entries are $z(t-2)$ and $t - 2$. Using $z^{p_i} = (z - 1)f_i + 1$, the resulting endpoint block is

$$\begin{pmatrix} d & -d + 1 - z \\ -d + z - 1 & d \end{pmatrix}, \qquad d = z + 1 - t + \sum_{i=1}^{t} 1/f_i.$$

Its determinant is $(z - 1)^2$. Since E is unit upper triangular in a topological ordering, the Coxeter polynomial is

$$\det(zE + E^{\mathsf{T}}) = (z - 1)^2 \prod_{i=1}^{t} (1 + z + \cdots + z^{p_i-1}).$$

In particular its spectral radius is one [20, Proposition 2.4]. Let $x = (x_u)_u$ be a real vector indexed by the vertices of the canonical quiver, with entries a and b at its source and sink. The Euler quadratic form $x^{\mathsf{T}}Ex$ is

$$\frac{1}{2} \sum_{\text{arm edges } u \to v} (x_u - x_v)^2 - \frac{t-2}{2}(a - b)^2.$$

Indeed, each internal vertex occurs in two edge squares, and each endpoint occurs in t edge squares. Expanding the formula gives a unit square at each vertex, the negative arrow terms, and the term $(t - 2)ab$ from the $t - 2$ canonical relations between the endpoints. On arm i, subtract $a + j(b - a)/p_i$ at distance j from the source, for $0 \leq j \leq p_i$. The

squared differences then split into a positive definite form on the internal deviations with zero endpoint values and the contribution $(a-b)^2/p_i$. The mixed term vanishes because differences of deviations sum to zero along the arm. Hence the coefficient on $(a-b)^2$ in the full form is

$$\frac{1}{2}\left(2-t+\sum_{i=1}^{t}\frac{1}{p_i}\right).$$

If this coefficient is zero, the form has corank two. If it is negative, the form is indefinite. Neither case satisfies (2). The canonical algebras with positive coefficient are the domestic ones, which are derived equivalent to affine hereditary path algebras [20, proof of Proposition 2.8]. It follows that every connected quasitilted algebra satisfying (2) is derived equivalent to a hereditary path algebra.

2.3. **Finite-dimensional module witnesses.** Happel, Reiten and Smalø characterize quasitilted algebras among the algebras of global dimension at most two [18, Theorem II.2.3]. In this range, R is quasitilted if and only if every indecomposable module has projective dimension at most one or injective dimension at most one. Failure is therefore witnessed by an indecomposable module M such that

$$\operatorname{Ext}^2_R(M,S)\neq 0 \qquad \text{and} \qquad \operatorname{Ext}^2_R(S,M)\neq 0,. \tag{3}$$

The first condition in (3) detects a nonzero second term in a minimal projective resolution, since every simple occurs in S. In a minimal injective resolution, each differential vanishes on the socle, the sum of the simple submodules of its source. Applying $\operatorname{Hom}_R(S,-)$ therefore makes the differentials zero. Every nonzero finite-dimensional injective has a simple submodule, so the second condition detects a nonzero term in degree two.

For a fixed dimension $m\geq 1$, an R-action on K^m is specified by $m\times m$ matrices for a basis of R, subject to its multiplication and unit equations. The common centralizer consists of module endomorphisms. A nontrivial idempotent gives a decomposition into its image and kernel, and a direct-sum decomposition gives such an idempotent by projection. Thus indecomposability is expressed by the absence of these idempotents. Both conditions are first-order over K. The nonvanishing tests (3) are also first-order. In degree two of each bar cochain complex, the ranks of the incoming and outgoing maps sum to less than the dimension of the middle space exactly when the cohomology is nonzero. Quantifier elimination thus decides whether a witness of dimension m exists. Testing $m=1,2,\dots$ gives a procedure that halts exactly when R is not quasitilted.

Proof of Proposition 2.1. Use the finite candidate list after all the preceding tests. For each candidate run the module witness search, advancing the searches in parallel. If each candidate has a witness, return a negative answer. This includes the case of an empty list.

When Q is mutation-acyclic, a tilted algebra from the relation extension theorem belongs to the candidate list. It passes (2) by its hereditary derived type. It is quasitilted, so the negative procedure never returns an answer in this case.

Conversely, suppose a candidate is quasitilted. The spectral test forces it to be derived equivalent to a hereditary path algebra, hence to be iterated tilted. By Barot, Fernández, Platzeck, Pratti and Trepode [11, Theorem 1.1], an iterated tilted algebra of global dimension at most two has the same relation-extension quiver as the ordinary quiver of a cluster-tilted algebra. The cluster-tilting graph of the hereditary cluster category is connected [13, Proposition 3.5], and exchange of a cluster-tilting summand mutates the ordinary quiver [12, Theorem 1.3]. Therefore Q is mutation-acyclic. If Q is not mutation-acyclic, every candidate has a module witness and the negative procedure eventually returns an answer. Alternate steps of the two searches until one returns an answer. The positive search halts on mutation-acyclic inputs, and the negative search halts on the remaining inputs. □

3. Reduction to skew-symmetric matrices

Let $C = (c_{uv})$ be a skew-symmetric integer matrix and let g be a permutation of its vertices preserving C. We call the action admissible if there are no arrows within a vertex orbit and all nonzero entries of each block between two orbits have the same sign. Write $E_1, \dots, E_n$ for the orbits. The fold is the matrix with entries

$$b_{ij} = \sum_{u \in E_i} c_{uv} \qquad (v \in E_j).$$

The sum is independent of v by invariance, and $|E_j| b_{ij} = -|E_i| b_{ji}$ by summing over the whole block and using skew-symmetry. Thus the fold is symmetrized by a diagonal matrix with entries proportional to $1/|E_i|$, scaled to positive integers. An orbit mutation mutates once at each vertex of one orbit. These mutations commute because there are no arrows within the orbit. We use Demonet's theorem, as stated in [15, Theorem 2.23], that an admissible action on an acyclic quiver remains admissible after every orbit mutation sequence. For the fold-mutation identity [15, Theorem 2.24], fix an orbit E_k to mutate and $v \in E_j$ with $j \neq k$. For $i \neq k$, sum the two product terms of the mutation formula over $u \in E_i$ and $w \in E_k$. The factors at w are unchanged by mutations at other vertices of E_k, since there are no arrows inside E_k. The common signs in the blocks give

$$\sum_{w \in E_k} \sum_{u \in E_i} [c_{uw}]_+ [c_{wv}]_+ = [b_{ik}]_+ [b_{kj}]_+.$$

The same computation with all entries negated gives $[-b_{ik}]_+ [-b_{kj}]_+$. Entries incident with the mutated orbit change sign. Thus the folded update is exactly matrix mutation.

Every action by a finite group of quiver automorphisms on an acyclic quiver is admissible, and its fold is acyclic. Indeed, a directed path between vertices of the same orbit can be concatenated with its successive translates under an element taking the first endpoint to the second. This would give a directed closed walk, hence a directed cycle. An arrow inside an orbit or arrows in opposite directions between two orbits would give such a path, translating one arrow to meet the other if necessary. The same translation argument lifts a directed cycle in the fold to a path between vertices of the same orbit.

Lemma 3.1. *Suppose C is connected and mutation-acyclic, and the action of g on C is admissible. If $\gcd(|E_1|, \dots, |E_n|) = 1$, then the fold of C is mutation-acyclic.*

Proof. Choose an acyclic representative Q of C. A seed with exchange quiver C is a seed of the coefficient-free cluster algebra of Q. Permuting its cluster variables by g induces a direct cluster automorphism f, that is, one preserving exchange quivers under the induced variable bijections. This follows by equivariance of the exchange relations under simultaneous permutation of the variables and the quiver. The order m of f is the order of g, since the variables of this seed form a transcendence basis.

An invariant acyclic seed. Over an algebraically closed field, the cluster category of a hereditary path algebra is the orbit of its bounded derived category by $F = \tau^{-1}[1]$, with τ the Auslander–Reiten translation and $[1]$ the derived shift. We use the correspondence between clusters and basic cluster-tilting objects, which identifies a mutation with exchange of one indecomposable summand [10, Section 3.1] [14, 12]. A basic object has pairwise nonisomorphic indecomposable summands. The direct automorphism f lifts to a derived equivalence and acts on the transjective Auslander–Reiten component, the component containing the initial shifted projectives [10, Lemma 3.1 and Corollary 3.2].

Consider the repetition quiver of the hereditary section formed by the initial shifted projectives. Index its vertices by (p, i), with $p \in \mathbb{Z}$ and i an initial vertex, so that $\tau(p, i) = (p - 1, i)$. For each arrow $i \to j$ of this section, the repetition quiver has arrows

$$(p, i) \longrightarrow (p, j), \qquad (p, j) \longrightarrow (p + 1, i) \qquad (p \in \mathbb{Z}).$$

In particular, the orientation used here is that of the Auslander–Reiten section. An integer vector $(h_i)_i$ defines a section $\{(h_i, i)\}_i$ whenever

$$(*) \qquad 0 \leq h_i - h_j \leq 1 \qquad \text{for every arrow } i \to j \text{ of the initial section.}$$

Each initial arrow contributes one arrow to the section, and along every arrow the quantity $p - h_i$ at (p, i) increases by 0 or 1. Thus paths between section vertices stay in the section. Such a section projects to an acyclic seed. To see this using source and sink exchanges, first subtract the minimum height from all heights, which is a common translate by a power of τ. If the resulting vector is nonzero, choose a sink in the initial section restricted to the vertices of maximal height. In the current section, an incident edge to a lower vertex is oriented into the chosen vertex: by $(*)$ it had the opposite orientation in the initial section. Edges at equal heights keep their initial orientation. Thus the chosen vertex is a sink in the current section. Lowering its height by one preserves $(*)$ and reverses its incident edges. Iteration reaches the zero vector. Read the moves in reverse. Each move raises a source by one, which replaces its object by its τ^{-1}-translate in the source exchange of a hereditary slice [10, Section 3.2]. Starting from the zero section, this proves that the sections project to acyclic seeds. A common τ-translate also preserves cluster-tilting objects and their quivers.

The inequalities $(*)$ are preserved by pointwise minimum: write them as $h_j \leq h_i$ and $h_i \leq h_j + 1$, and take the minimum on both sides of each inequality. An automorphism of the repetition quiver commuting with τ takes (p, i) to $(p + c_i, \sigma(i))$ for integers c_i and a permutation σ. The image of a section is again a section: it has one vertex from each translation orbit and an arrow over every adjacency of the initial quiver. By the repetition-quiver formula, such an arrow requires the two chosen heights to satisfy $(*)$ on that adjacency. The coordinate permutation and additions commute with pointwise minimum. For a finite-order translation-quiver automorphism, the pointwise minimum of the finitely many images of the zero section is an invariant section.

If Q is not Dynkin, the transjective component in the cluster category is the repetition quiver itself [10, Section 3.2]. The induced action of f has finite order, so the invariant section just constructed gives an f-invariant acyclic seed. If Q is Dynkin, its derived Auslander–Reiten quiver is the repetition quiver, and the projection onto the cluster Auslander–Reiten quiver is a connected regular covering with deck group generated by F [13, Propositions 1.3 and 1.6]. The rank-one case is immediate. In higher rank let a be the action of the derived lift on the repetition quiver. The derived lift commutes with the shift and the Auslander–Reiten translation, hence with F. In Dynkin type every indecomposable object of the cluster category corresponds to a cluster variable. Thus f^m fixes every vertex of the quotient, and a^m acts over the identity there. Its image at a chosen vertex equals that of F^k for some integer k. Unique lifting of walks from that vertex and connectedness of the covering give $a^m = F^k$ on all vertices. For each original orbit E_i, choose a lift v_i of a summand corresponding to a variable in E_i. Since $f^{|E_i|}$ fixes that variable, $a^{|E_i|}v_i = F^{k_i}v_i$ for some $k_i \in \mathbb{Z}$. Commutativity with F and freeness of the deck action give

$$k|E_i| = mk_i.$$

An integer linear combination of the numbers $|E_i|$ equals one. Multiplying that combination by k and using the displayed identities gives $m \mid k$. Hence $F^{-k/m}a$ is a finite-order lift of the same action, and its invariant section projects to an f-invariant acyclic seed in this case as well.

An orbit mutation path. Write S for the invariant acyclic seed and Q_S for its quiver. The permutation by f on S preserves Q_S. Extend it to an automorphism of the path algebra by numbering the parallel arrows in each ordered pair and sending each arrow to the arrow of the same number in the image pair. The resulting cluster action agrees with f on the transcendence basis of S, and therefore agrees on every cluster variable.

Work over $\overline{\mathbb{F}_q}$, the algebraic closure of a finite field $\mathbb{F}_q$. The cluster labels and quivers in this category are obtained by the same seed mutation words, so the permutation of the cluster variables is independent of the field of the path algebra. Compose the path algebra automorphism with the scalar Frobenius $\lambda \mapsto \lambda^q$ and write H for the fixed algebra. Frobenius descent makes H a finite-dimensional hereditary $\mathbb{F}_q$-algebra whose scalar extension is the path algebra $\overline{\mathbb{F}_q}Q_S$ [22, Section 3.1]. It is indecomposable because Q_S is connected. The corresponding action on cluster-tilting objects agrees on indecomposable summands with the action on variables by f. Indeed, scalar Frobenius fixes the shifted projectives and hence every reachable indecomposable summand by equivariance of exchange and uniqueness of the other complement. Denote by M the basic cluster-tilting object corresponding to the original seed with quiver C. Its summands are permuted by the Frobenius action.

Frobenius descent in cluster categories gives the correspondence [22, Theorem 3.2 and Proposition 3.5]. The sum of a Frobenius orbit of indecomposable objects descends to one indecomposable object, and the Hom and extension spaces downstairs are the Frobenius-fixed spaces of their counterparts upstairs, with scalar extension recovering the latter spaces. In particular, a Frobenius-stable basic cluster-tilting object U descends to a basic cluster-tilting object in the category of H, denoted U^{desc}. Rigidity descends by the extension formula. For maximality, if an object downstairs has vanishing first extensions from U^{desc}, its scalar extension lies in the additive closure of U, and the orbit correspondence puts the original object in the additive closure of U^{desc}. The additive closure here consists of direct summands of finite direct sums of copies.

By [13, Proposition 3.5], the cluster-tilting graph of H is connected. Choose a path from the shifted regular object $H[1]$ to M^{desc}. Lift it inductively, starting at the object for S. At a stage with lifted object U, the summand exchanged downstairs comes from a Frobenius orbit of summands of U. The already lifted path consists of orbit mutations from Q_S, so Demonet's theorem ensures admissibility at this stage. The mutations at the vertices of the orbit commute and produce a basic cluster-tilting object U'. By equivariance under Frobenius, the new summands constitute one orbit. They are distinct by basicness. They also differ from every removed summand: the mutation producing any given replacement can be performed first, when all the other old summands are still present. Consequently U'^{desc} has the same unchanged summands as U^{desc} and one different summand. By the two-complement theorem [13, Theorem 5.1], it is the next object of the chosen path. At its end the lifted object is isomorphic to M, since scalar extension recovers the split object. We have obtained an orbit mutation path from Q_S to C, with the prescribed action at its endpoint. Folding gives a mutation path from the acyclic fold of Q_S to the fold of C, with the orbit labels transported along the path. Matching the specified labels at the endpoint relabels the full path and preserves its acyclic endpoint. □

Proposition 3.2. *Mutation-acyclicity of a skew-symmetrizable integer matrix reduces to a computable finite disjunction of skew-symmetric mutation-acyclicity queries.*

Proof. Treat a connected matrix $B = (b_{ij})$ of size n. Compute a positive integral diagonal symmetrizer $D = \mathrm{diag}(d_1, \dots, d_n)$, so that DB is skew-symmetric. Set the diagonal weight of one vertex to one and propagate $d_j/d_i = -b_{ij}/b_{ji}$ along a spanning tree of the undirected support. The input admits a symmetrizer, so the propagated rational weights agree with a common rescaling of that symmetrizer and satisfy all edges. Clear denominators. Put

$$L = \mathrm{lcm}(d_1, \dots, d_n), \qquad e_i = L/d_i, \qquad E_i = \mathbb{Z}/e_i\mathbb{Z}.$$

Here the sets E_i are taken disjointly. We have $\gcd(e_1, \dots, e_n) = 1$, since each prime dividing L attains its maximum exponent in some d_i. Let g act by simultaneous successor on all the E_i.

Enumerate all skew-symmetric matrices on $\bigsqcup_i E_i$ invariant under g, admissible for its action, and folding to B. In the (i, j) block every entry has absolute value at most

$|b_{ij}|$, by the common sign condition and the prescribed column sums. This gives a finite effective enumeration. Each such matrix is connected. Indeed, for a component X write $f_i = |X \cap E_i|$. A nonzero block has nonzero constant weighted degree at both ends, so X meets every E_i by connectedness of B. Counting the arrows of a nonzero block in X from its two ends gives $f_j b_{ij} = -f_i b_{ji}$, whereas $e_j b_{ij} = -e_i b_{ji}$. Hence the ratios f_i/e_i all coincide. Their common reduced denominator divides every e_i, so the ratio is a positive integer at most one. Thus X is the entire vertex set. Lemma 3.1 now gives mutation-acyclicity of B whenever one of the enumerated matrices is mutation-acyclic.

Conversely, suppose B has an acyclic representative $A = (a_{ij})$ reached by a mutation sequence. The same diagonal D symmetrizes every matrix along the sequence. Construct an acyclic lift of A on the prescribed E_i as follows. For each pair with $a_{ij} > 0$, simultaneous successor on $E_i \times E_j$ has orbit length $\operatorname{lcm}(e_i, e_j)$. During one period, a given column occurs $e_i/\gcd(e_i, e_j)$ times and a given row occurs $e_j/\gcd(e_i, e_j)$ times. The equation $e_j a_{ij} = -e_i a_{ji}$ implies that $e_i/\gcd(e_i, e_j)$ divides a_{ij}. Give every pair in one such orbit weight $a_{ij}\gcd(e_i, e_j)/e_i$ in the (i, j) block, put zero on the other pairs, and take the negative transpose in the opposite block. Use zero blocks when $a_{ij} = a_{ji} = 0$ and within each E_i. The resulting matrix folds to A, and every arrow projects to an arrow of the acyclic support of A, so the lift is acyclic. Lift the reverse mutation sequence $A \to B$ by orbit mutations. By stable admissibility and the fold-mutation formula, the endpoint is among the enumerated matrices and is mutation-acyclic. For a disconnected input, enumerate all block sums obtained by taking one lift per component. Such a block sum is mutation-acyclic exactly when each of its summands is. The empty block sum handles the empty input. □

Proof of Theorem 1.1. Construct the finite family in Proposition 3.2 and test its members by Proposition 2.1. Return yes exactly when one member is mutation-acyclic. □

4. Mutation equivalence with a mutation-acyclic input

Lemma 4.1. *For two acyclic exchange matrices, mutation equivalence up to relabeling is decidable.*

Proof. For a matrix $A = (a_{ij})$, a source is a vertex with no incoming arrow in the directed support $i \to j$ for $a_{ij} > 0$, and a sink is a vertex with no outgoing arrow. A mutation at a source or a sink reverses the arrows incident to that vertex and preserves every absolute entry. Indeed, the two-step terms in the matrix mutation formula vanish at such a vertex. A source becomes a sink and a sink becomes a source, so no directed cycle is created.

By the acyclic source–sink theorem of Liu and Pan [21, Theorem 3], if two acyclic sign-skew-symmetric matrices are mutation equivalent up to relabeling, they can be joined by source and sink mutations up to relabeling. Here sign-skew-symmetric means that for every pair of distinct indices the two opposite entries are both zero or have opposite signs, as they do for every skew-symmetrizable exchange matrix. In the skew-symmetric case the theorem is due to Caldero and Keller [14, Corollary 4]. The converse holds because source and sink mutations are mutations.

Matrices of different sizes are inequivalent. For equal sizes, search all matrices reachable from the first by source and sink mutations. If m is the number of unordered pairs carrying nonzero entries in the first matrix, there are at most 2^m states, because the absolute entries are fixed and each such pair has two possible orientations. Compare every visited matrix with every simultaneous row and column relabeling of the second matrix. The search finds a match exactly when the matrices are mutation equivalent up to relabeling. □

Proof of Theorem 1.2. Apply Theorem 1.1 to both matrices. If neither is mutation-acyclic, report that both are mutation-cyclic. If exactly one is mutation-acyclic, return inequivalence, since the existence of an acyclic representative is invariant under mutation and relabeling.

Suppose both are mutation-acyclic. For each matrix, enumerate all mutation words in order of length, including the empty word, and stop when the directed support of the resulting matrix is acyclic. Both searches terminate because both inputs are mutation-acyclic. By reversibility of mutation and equivariance under relabeling, the input matrices are mutation equivalent up to relabeling exactly when the two acyclic representatives are. Apply Lemma 4.1 to them. □

References

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

[2] I. Assem, M. Blais, T. Brüstle, and A. Samson, *Mutation classes of skew-symmetric* 3×3*-matrices*, arXiv:math/0610627, 2006.

[3] A. Beineke, T. Brüstle, and L. Hille, *Cluster-cyclic quivers with three vertices and the Markov equation*, Algebr. Represent. Theory **14** (2011), 97–112.

[4] S. Fomin, L. Williams, and A. Zelevinsky, *Introduction to cluster algebras. Chapters 1–3*, arXiv:1608.05735; problem numbers refer to the combined draft.

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

[6] S. Neville, *Mutation-acyclic quivers are totally proper*, arXiv:2409.17832v2, 2025.

[7] D. Soukup, *Computational complexity in combinatorics and algebra*, Ph.D. thesis, University of California, Los Angeles, 2024.

[8] 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.

[9] I. Assem, T. Brüstle, and R. Schiffler, *Cluster-tilted algebras as trivial extensions*, Bull. Lond. Math. Soc. **40** (2008), 151–162.

[10] I. Assem, R. Schiffler, and V. Shramchenko, *Cluster automorphisms*, Proc. Lond. Math. Soc. (3) **104** (2012), 1271–1302.

[11] M. Barot, E. Fernández, M. I. Platzeck, N. I. Pratti, and S. Trepode, *From iterated tilted algebras to cluster-tilted algebras*, Adv. Math. **223** (2010), 1468–1494.

[12] A. B. Buan, R. J. Marsh, and I. Reiten, *Cluster mutation via quiver representations*, Comment. Math. Helv. **83** (2008), 143–177.

[13] A. B. Buan, R. J. Marsh, M. Reineke, I. Reiten, and G. Todorov, *Tilting theory and cluster combinatorics*, Adv. Math. **204** (2006), 572–618.

[14] P. Caldero and B. Keller, *From triangulated categories to cluster algebras II*, Ann. Sci. Éc. Norm. Sup. (4) **39** (2006), 983–1009.

[15] G. Dupont, *An approach to non-simply laced cluster algebras*, J. Algebra **320** (2008), 1626–1661.

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

[17] D. Happel, *A characterization of hereditary categories with tilting object*, Invent. Math. **144** (2001), 381–398.

[18] D. Happel, I. Reiten, and S. O. Smalø, *Tilting in abelian categories and quasitilted algebras*, Mem. Amer. Math. Soc. **120** (1996), no. 575.

[19] B. Keller, *Derived invariance of higher structures on the Hochschild complex*, preprint, 2003.

[20] H. Lenzing and J. A. de la Peña, *Spectral analysis of finite dimensional algebras and singularities*, arXiv:0805.1018, 2008.

[21] S. Liu and J. Pan, *A note on mutation equivalence*, Mathematics **13** (2025), no. 3, 339.

[22] D. Yang, *Non-simply-laced clusters of finite type via Frobenius morphism*, arXiv:math/0608114, 2006.

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