

CLUSTER ALGEBRAS OF FINITE TYPE AND POSITIVE SYMMETRIZABLE MATRICES

MICHAEL BAROT, CHRISTOF GEISS, AND ANDREI ZELEVINSKY

ABSTRACT. The paper is motivated by an analogy between cluster algebras and Kac-Moody algebras: both theories share the same classification of finite type objects by familiar Cartan-Killing types. However the underlying combinatorics beyond the two classifications is different: roughly speaking, Kac-Moody algebras are associated with (symmetrizable) Cartan matrices, while cluster algebras correspond to skew-symmetrizable matrices. We study an interplay between the two classes of matrices, in particular, establishing a new criterion for deciding whether a given skew-symmetrizable matrix gives rise to a cluster algebra of finite type.

1. INTRODUCTION

This paper is motivated by the theory of cluster algebras initiated in [2]. Here we deal exclusively with the combinatorial aspects of the theory, so no knowledge of algebraic properties of cluster algebras (including their definition) will be assumed or needed. The reader should just bear in mind an analogy between cluster algebras and Kac-Moody algebras. In both theories, there is an appropriate notion of finite type. (For Kac-Moody algebras, “finite type” just means being finite-dimensional, that is, a semisimple Lie algebra.) Cluster algebras of finite type were classified in [3], and the resulting classification turns out to be identical to the famous Cartan-Killing classification of semisimple Lie algebras. However the underlying combinatorics beyond the two classes of algebras is different: roughly speaking, Kac-Moody algebras correspond to (symmetrizable) Cartan matrices, while cluster algebras correspond to skew-symmetrizable matrices. In this paper, we study an interplay between the two classes of matrices. In particular, we establish a new criterion for deciding whether a given skew-symmetrizable matrix gives rise to a cluster algebra of finite type.

To state our main results, we need some terminology. In what follows, by a matrix we always mean a square integer matrix. A matrix A (resp. B) is symmetrizable (resp. skew-symmetrizable) if DA (resp. DB) is symmetric (resp. skew-symmetric) for some diagonal matrix D with positive diagonal entries. Thus, in a symmetrizable (resp. skew-symmetrizable) matrix, the two transpose entries have the same sign (resp. opposite signs). We say that a symmetrizable matrix is a *quasi-Cartan matrix* if all its diagonal entries are equal to 2 (note that a skew-symmetrizable matrix has all diagonal entries equal to 0). We say that a quasi-Cartan matrix A is *positive* if the symmetrized matrix DA is positive definite; by the Sylvester criterion, this means that the principal minors of A are all positive. For a skew-symmetrizable

Date: December 13, 2004; revised June 27, 2005.

2000 *Mathematics Subject Classification.* Primary: 05E15, Secondary: 05C50, 15A36, 17B67, Research supported by DGAPA grant IN101402-3 (C.G.) and NSF (DMS) grant 0200299 (A.Z.).

matrix B , we will refer to a quasi-Cartan matrix A with $|A_{ij}| = |B_{ij}|$ for all $i \neq j$ as a *quasi-Cartan companion* of B .

A quasi-Cartan matrix is a (generalized) *Cartan matrix* if its off-diagonal entries are non-positive. These are the matrices giving rise to Kac-Moody algebras, see [4]. As shown in [4], the Kac-Moody algebra associated to A is finite-dimensional if and only if A is positive. Thus, positive Cartan matrices form the combinatorial backbone of the Cartan-Killing classification: every such matrix can be transformed by a simultaneous permutation of rows and columns to a block-diagonal matrix whose diagonal blocks are of the familiar types $\mathbb{A}_n, \mathbb{B}_n, \mathbb{C}_n, \mathbb{D}_n, \mathbb{E}_6, \mathbb{E}_7, \mathbb{E}_8, \mathbb{F}_4$, and \mathbb{G}_2 , represented by Dynkin diagrams.

On the other hand, cluster algebras are associated with the mutation-equivalence classes of skew-symmetrizable matrices. Recall from [2] that, for each matrix index k , the *mutation* in direction k transforms a skew-symmetrizable matrix B into another skew-symmetrizable matrix $B' = \mu_k(B)$, whose entries are given by

$$(1.1) \quad B'_{ij} = \begin{cases} -B_{ij} & \text{if } i = k \text{ or } j = k; \\ B_{ij} + \operatorname{sgn}(B_{ik})[B_{ik}B_{kj}]_+ & \text{otherwise,} \end{cases}$$

where we use the notation $[x]_+ = \max(x, 0)$ and $\operatorname{sgn}(x) = x/|x|$, with the convention $\operatorname{sgn}(0) = 0$ (the formula (1.1) is easily seen to be equivalent to [2, (4.3)]). One easily checks that μ_k is involutive, implying that the repeated mutations in all directions give rise to the *mutation-equivalence* relation on skew-symmetrizable matrices. We are now ready to state the classification result from [3].

Theorem 1.1. *For a mutation-equivalence class \mathcal{S} of skew-symmetrizable matrices, the following are equivalent.*

- (1) *The cluster algebra associated to \mathcal{S} is of finite type.*
- (2) *\mathcal{S} contains a matrix B such that the Cartan matrix A with off-diagonal entries $A_{ij} = -|B_{ij}|$ is positive.*
- (3) *For every $B \in \mathcal{S}$ and all $i \neq j$, we have $|B_{ij}B_{ji}| \leq 3$.*

Furthermore, the Cartan-Killing type of the Cartan matrix A in (2) is uniquely determined by \mathcal{S} .

In this paper we address the following

Recognition Problem. Given a skew-symmetrizable matrix B , find an efficient way to determine whether the cluster algebra associated with (the mutation-equivalence class of) B is of finite type.

The problem makes sense since the mutations are hard to control, so each of the conditions (2) and (3) in Theorem 1.1 is hard to check in general. A nice solution of the problem was obtained by A. Seven in [6]. His answer is given in terms of “forbidden minors” of B . Our answer is very different and probably more practical: it is based on extending the criterion in (2) to *every* representative of a mutation class in question. To state it, we need a little bit more terminology.

To a skew-symmetrizable $n \times n$ matrix B we associate a directed graph $\Gamma(B)$ with vertices $1, \dots, n$ and directed edges $i \rightarrow j$ for all i, j with $B_{ij} > 0$. A *chordless cycle* in $\Gamma(B)$ is an induced subgraph isomorphic to a cycle (thus, the vertices of a

chordless cycle can be labeled by the elements of $\mathbb{Z}/p\mathbb{Z}$ for some $p \geq 3$ so that the edges between them are precisely $\{i, i + 1\}$ for $i \in \mathbb{Z}/p\mathbb{Z}$.

We are finally ready to state our main result.

Theorem 1.2. *The mutation-equivalence class of a skew-symmetrizable matrix B satisfies the equivalent conditions in Theorem 1.1 if and only if B satisfies:*

- (4) *Every chordless cycle in $\Gamma(B)$ is cyclically oriented, and B has a positive quasi-Cartan companion.*

Remark 1.3. We actually prove $(2) \Rightarrow (4) \Rightarrow (3)$, which gives a new proof of the implication $(2) \Rightarrow (3)$, more elementary and straightforward than the one in [3].

Condition (4) in Theorem 1.2 is not completely explicit since B can have many quasi-Cartan companions. Note that a quasi-Cartan companion of B is specified by choosing the signs of its off-diagonal matrix entries, with the only requirement that $\text{sgn}(A_{ij}) = \text{sgn}(A_{ji})$ for $i \neq j$. Thus, the number of choices for A is 2^N , where N is the number of edges in $\Gamma(B)$. The following two propositions allow us to considerably sharpen Theorem 1.2.

Proposition 1.4. *To be positive, a quasi-Cartan companion A of a skew-symmetrizable matrix B must satisfy the following sign condition: for every chordless cycle Z in $\Gamma(B)$, the product $\prod_{\{i,j\} \in Z} (-A_{ij})$ over all edges of Z must be negative.*

Proposition 1.5. *If every chordless cycle in $\Gamma(B)$ is cyclically oriented, then B has a quasi-Cartan companion (not necessarily positive) satisfying the sign condition in Proposition 1.4; furthermore, such a quasi-Cartan companion is unique up to simultaneous sign changes in rows and columns.*

Proposition 1.4 (resp. Proposition 1.5) is a consequence of Proposition 2.6 (resp. Corollary 5.2) below. In view of these results, in checking condition (4), it is enough to test positivity of just one quasi-Cartan companion of B since simultaneous sign changes in rows and columns do not affect positivity. The following example provides an illustration.

Example 1.6. Let $B(n)$ be the $n \times n$ skew-symmetric matrix with the above-diagonal entries given by

$$(1.2) \quad B_{ij} = \begin{cases} -1 & \text{if } j - i = 1; \\ 1 & \text{if } j - i = 2; \\ 0 & \text{if } j - i > 2. \end{cases}$$

The graph $\Gamma(B(n))$ has $n - 2$ chordless cycles: they are formed by all triples of consecutive indices. An immediate check shows that all of them are cyclically oriented. Now let $A(n)$ be the quasi-Cartan companion of $B(n)$ such that $A_{ij} = B_{ij}$ for $j > i$. An immediate check shows that $A(n)$ satisfies the condition in Proposition 1.4. Let $d_n = \det(A(n))$, with the convention $d_0 = 1$. It is not hard to show that the generating function of this sequence is given by

$$(1.3) \quad \sum_{n \geq 0} d_n x^n = \frac{(1+x)(1+x+x^2)(1+x^2)(1+x^3)}{1-x^{12}},$$

implying that $d_{n+12} = d_n$ for $n \geq 0$. Since the numerator in (1.3) is a polynomial of degree 8, we see that $d_9 = d_{10} = d_{11} = 0$. The values of d_n for $1 \leq n \leq 8$ are given by the following table:

n	1	2	3	4	5	6	7	8
$\det(A(n))$	2	3	4	4	4	3	2	1
Cartan-Killing type	\mathbb{A}_1	\mathbb{A}_2	\mathbb{A}_3	\mathbb{D}_4	\mathbb{D}_5	\mathbb{E}_6	\mathbb{E}_7	\mathbb{E}_8

TABLE 1. Determinants and Cartan-Killing types of the $A(n)$

By the Sylvester criterion, $A(n)$ is positive if and only if $n \leq 8$. Applying Theorem 1.2 and Proposition 1.4, we conclude that the cluster algebra associated to $B(n)$ is of finite type precisely when $n \leq 8$. The corresponding Cartan-Killing types are given in the last line of Table 1 (their determination is left as an exercise to the reader).

The paper is organized as follows. The proof of Theorem 1.2 is carried out in the next three sections. In Section 2, which can be viewed as a “symmetrizable analogue” of [3, Section 9], we establish some needed properties of positive quasi-Cartan matrices. In Section 3, we show that, under some conditions, the mutation-equivalence of skew-symmetrizable matrices can be extended to a natural equivalence of properly chosen quasi-Cartan companions. The results of these two sections are put together in Section 4, where the proof of Theorem 1.2 is completed.

The concluding Section 5 is purely graph-theoretic. We call a graph Γ *cyclically orientable* if it admits an orientation in which any chordless cycle is cyclically oriented. (For example, the full graph on four vertices is *not* cyclically orientable.) The main result of Section 5 is Theorem 5.1 which gives several properties of graphs that are equivalent to cyclical orientability. As a consequence, we obtain a graph-theoretic statement (Corollary 5.2) which implies Proposition 1.5. The paper concludes with Remark 5.7 discussing a possible strategy for checking condition (4).

Remark 1.7. The preliminary version of Theorem 1.2 due to the first two authors was stated in terms of *integral quadratic forms* rather than quasi-Cartan matrices. (Recall that the quadratic form represented by a symmetric matrix C with positive diagonal entries is called *integral* if $2C_{ij}/C_{ii} \in \mathbb{Z}$ for all i, j ; in our present terminology, this means that $C = DA$ is the symmetrized version of a quasi-Cartan matrix A .) The classification of positive definite integral quadratic forms (up to natural equivalence) is well-known, and the answer is once again given by the Cartan-Killing classification, cf. Proposition 2.9. Our results imply that this classification is in agreement with the one given by Theorem 1.1: if A is a positive quasi-Cartan companion of B , then the Cartan-Killing type of the cluster algebra associated to B is the same as the Cartan-Killing type of the quadratic form represented by the symmetrized matrix DA .

2. PROPERTIES OF POSITIVE QUASI-CARTAN MATRICES

This section is a “symmetrizable analogue” of [3, Section 9]. We start with a simple observation.

Lemma 2.1. *Let A be a positive quasi-Cartan matrix. Then*

- (a) $0 \leq A_{ij}A_{ji} \leq 3$ for any $i \neq j$.
- (b) $A_{ik}A_{kj}A_{ji} \geq 0$ for any pairwise different i, j, k .

Proof. (a) is immediate from the positivity of the principal minor of A on the rows and columns i and j .

(b) Let $A_{ik}A_{kj}A_{ji} \neq 0$. Since A is symmetrizable, we have $A_{ki}A_{jk}A_{ij} = A_{ik}A_{kj}A_{ji}$. The positivity condition for the principal 3×3 minor of A on the rows and columns i, j, k can now be rewritten as

$$(2.1) \quad A_{ik}A_{kj}A_{ji} > A_{ij}A_{ji} + A_{ik}A_{ki} + A_{jk}A_{kj} - 4 \geq -1,$$

implying our claim. \square

We now introduce the *diagram* of a quasi-Cartan matrix, which is a symmetrizable analogue of [3, Definition 7.3].

Definition 2.2. The *diagram* $\Gamma(A)$ of a $n \times n$ quasi-Cartan matrix A is a (undirected) graph with vertices $\{1, 2, \dots, n\}$ and edges $\{i, j\}$ for each $i \neq j$ with $A_{ij} \neq 0$, where every edge $\{i, j\}$ is assigned the *weight* $A_{ij}A_{ji}$ and the *sign* $\varepsilon_{ij} = -\operatorname{sgn}(A_{ij}) = -\operatorname{sgn}(A_{ji})$.

In drawing the diagrams, all unspecified weights will be assumed to be equal to 1. With some abuse of notation, we denote by the same symbol $\Gamma(A)$ the underlying edge-weighted graph of the diagram, obtained by forgetting signs of the edges, and also the underlying graph obtained by forgetting both signs and edge weights. Note that an edge-weighted graph Γ which is of the form $\Gamma(A)$ must satisfy the following condition (see [4, Exercise 2.1]) :

$$(2.2) \quad \text{The product of edge weights along every cycle of } \Gamma \text{ is a perfect square.}$$

Definition 2.3. An edge-weighted graph will be called *positive* if some sign assignment to the edges makes it into the diagram $\Gamma(A)$ of some positive quasi-Cartan matrix A .


The following proposition is an analogue of [3, Proposition 9.3].

Proposition 2.4. *Positive edge-weighted trees are precisely Dynkin diagrams. Each of them becomes the diagram of a positive quasi-Cartan matrix under an arbitrary assignment of signs.*

Proof. Suppose $\Gamma(A)$ is a tree. We can assume without loss of generality that the signs of all edges are equal to 1, i.e., A is a generalized Cartan matrix (this can be achieved by replacing A if necessary by a positive quasi-Cartan matrix of the form $A' = EAE$, where E is a diagonal matrix with entries $E_{ii} = \pm 1$). Our statement now follows from the Cartan-Killing classification of positive generalized Cartan matrices, see, e.g., [4, Theorem 4.8] or Theorems 1 and 4 in [1, VI,4]. \square


In analogy with [3, Definition 9.1], by a *subdiagram* of $\Gamma(A)$ we mean a diagram of the form $\Gamma(A')$, where A' is a principal submatrix of A . Thus, $\Gamma(A')$ is obtained from $\Gamma(A)$ by taking an induced subgraph on a subset of vertices and keeping all the edge weights and signs the same as in $\Gamma(A)$. Since the positivity of A implies that of A' , we obtain the following corollary.

Corollary 2.5. *None of the following edge-weighted graphs can appear as subdiagrams of the diagram of a positive quasi-Cartan matrix:*



Our next result is an analogue of [3, Proposition 9.7].

Proposition 2.6. *Positive edge-weighted cycles are precisely those of the following three types:*



Furthermore, a sign assignment makes each of these cycles the diagram of a positive quasi-Cartan matrix if and only if the product of signs along all the edges is equal to -1 .

Proof. First suppose that A is a positive quasi-Cartan matrix whose diagram is a cycle not of type (a). Then at least one edge of $\Gamma(A)$ has weight $a > 1$. By Lemma 2.1(a), the weight a is either 2 or 3. It follows from (2.2) that at least two edges of $\Gamma(A)$ have weight a . If $a = 3$ then the cycle must be a triangle since otherwise it would have a subdiagram $\tilde{\mathbb{G}}_2$, a contradiction to Corollary 2.5. We are left with the case of a 3×3 matrix A such that $A_{12}A_{21} = A_{23}A_{32} = 3$, and $A_{13}A_{31} = 1$. In view of (2.1), such a matrix A cannot be positive.

So, we can assume that all the edge weights are equal to 1 or 2, with at least two edges of weight 2. Then $\Gamma(A)$ must be of one of the types (b) or (c), since otherwise it would have a subdiagram $\tilde{\mathbb{C}}_n$ for some $n \geq 2$, again in contradiction to Corollary 2.5.

To finish the proof, it remains to show that, in each of the cases (a), (b) and (c), the positivity of a quasi-Cartan matrix A is equivalent to the condition that the product of signs along all the edges is equal to -1 . Since in each case, all the proper subdiagrams of $\Gamma(A)$ are Dynkin diagrams, we only need to show that the above sign condition is equivalent to $\det(A) > 0$. Using simultaneous permutations and

changes of signs of rows and columns, we can assume that $A_{ij} < 0$ for $|j - i| = 1$; and we need to show that $\det(A) > 0$ precisely when A_{1n} and A_{n1} are positive.

Dealing with case (a) first, let $A_n(x)$ be the $n \times n$ quasi-Cartan matrix with non-zero off-diagonal entries $A_{ij} = -1$ for $|j - i| = 1$, and $A_{1n} = A_{n1} = x$. By a standard calculation,

$$\det(A_n(x)) = n + 1 + 2x - x^2(n - 1);$$

therefore, $\det(A_n(1)) = 4$, and $\det(A_n(-1)) = 0$, as required. The cases (b) and (c) are similar (and simpler). \square

Remark 2.7. The sign condition in Proposition 2.6 implies Proposition 1.4.

Definition 2.8. We say that quasi-Cartan matrices A and A' are *equivalent* and write $A' \sim A$ if A and A' have the same symmetrizer D (that is, D is a diagonal matrix with positive diagonal entries such that $C = DA$ and $C' = DA'$ are symmetric), and the symmetrized matrices satisfy $C' = E^T C E$ for some integer matrix E with determinant ± 1 .

Note that this definition does not depend on the choice of a symmetrizer D .

We conclude this section by showing that the equivalence classes of positive quasi-Cartan matrices are classified by Cartan-Killing types. Although this result is well-known to experts, we were unable to find an adequate reference, so will outline the proofs.

Let A be a $n \times n$ quasi-Cartan matrix. For each $i = 1, \dots, n$, define an automorphism s_i of the lattice \mathbb{Z}^n by setting $s_i(e_j) = e_j - A_{ij}e_i$, where $\{e_1, \dots, e_n\}$ is the standard basis in \mathbb{Z}^n . Let $W(A) \subset GL_n(\mathbb{Z})$ be the group generated by s_1, \dots, s_n .

Proposition 2.9. *The following conditions on a quasi-Cartan matrix A are equivalent:*

- (1) A is positive.
- (2) The group $W(A)$ is finite.
- (3) There exist a (reduced) root system Φ and a linearly independent subset $\{\beta_1, \dots, \beta_n\} \subset \Phi$ such that $A_{ij} = \langle \beta_i^\vee, \beta_j \rangle$, where β^\vee is the coroot dual to a root β .
- (4) A is equivalent to a positive Cartan matrix A^0 .

Under these conditions, if $\Phi_0 \subset \Phi$ is the smallest root subsystem of Φ that contains the set $\{\beta_1, \dots, \beta_n\}$ in (3), then the Cartan-Killing type of Φ_0 is the same as the Cartan-Killing type of the matrix A^0 in (4), and it characterizes A up to equivalence. Furthermore, $W(A)$ is naturally identified with the Weyl group of Φ_0 .

Proof. The implication (4) \implies (1) is trivial. To prove (1) \implies (2), suppose that A is positive, and let $(\alpha | \beta)$ be the positive definite (symmetric) scalar product in \mathbb{Z}^n given by $(e_i | e_j) = C_{ij}$, where $C = (C_{ij}) = DA$ is the positive definite symmetric matrix corresponding to A . The standard check shows that, with respect to this scalar product, s_i is the orthogonal reflection in the orthogonal complement to e_i . Thus W is a discrete subgroup of the compact orthogonal group $O_n(\mathbb{R})$, and so is finite, as claimed.

The implication (2) \implies (3) follows from a well-known classification of finite crystallographic reflection groups given in [1]: namely, every such group is the Weyl group

of a reduced root system Φ . To be more specific, take $\Phi = W(A) \{e_1, \dots, e_n\} \subset \mathbb{Z}^n$. If $W(A)$ is finite then so is Φ ; the fact that Φ satisfies the rest of the axioms of the root systems in [1] is checked easily. Then (3) holds with $\beta_i = e_i$.

It remains to prove (3) \implies (4). Without loss of generality, we can assume that Φ is the smallest root system containing $\{\beta_1, \dots, \beta_n\}$, that is, $\Phi = W\{\beta_1, \dots, \beta_n\}$, where $W = W(A)$ is the group generated by the reflections corresponding to the roots β_1, \dots, β_n . By this assumption, Φ has rank n , and $\{\beta_1, \dots, \beta_n\}$ is a \mathbb{Z} -basis of the root lattice of Φ .

Let $(\alpha | \beta)$ be a W -invariant positive definite scalar product in the root lattice. By the definition, $A_{ij} = 2(\beta_i | \beta_j) / (\beta_i | \beta_i)$. Thus, the symmetric matrix $C = DA$ corresponding to A has entries $C_{ij} = (\beta_i | \beta_j)$, with the diagonal entries of D given by $D_{ii} = (\beta_i | \beta_i) / 2$.

Let $\{\alpha_1, \dots, \alpha_n\}$ be a system of simple roots in Φ , and A^0 the corresponding (positive) Cartan matrix, so that $A_{ij}^0 = \langle \alpha_i^\vee, \alpha_j \rangle$. Thus, the symmetric matrix $C^0 = D^0 A^0$ corresponding to A^0 has entries $C_{ij}^0 = (\alpha_i | \alpha_j)$, with the diagonal entries of D^0 given by $D_{ii}^0 = (\alpha_i | \alpha_i) / 2$. It follows that $C = E^T C^0 E$, where E is the transition matrix from the basis $\{\alpha_1, \dots, \alpha_n\}$ to the basis $\{\beta_1, \dots, \beta_n\}$ (since both families are \mathbb{Z} -bases of the root lattice, the matrix E is invertible over the integers). To finish the proof it remains to show that the simple roots $\alpha_1, \dots, \alpha_n$ can be chosen in such a way that $D^0 = D$. In other words, we need to show the following:

(2.3) One can reorder the roots β_1, \dots, β_n so that $(\beta_i | \beta_i) = (\alpha_i | \alpha_i)$ for all i .

We can assume without loss of generality that the root system Φ is irreducible. If Φ is simply-laced, i.e., all roots are of the same length, (2.3) becomes tautological. So it is enough to check (2.3) for Φ of one of the types $\mathbb{B}_n, \mathbb{C}_n, \mathbb{F}_4$ and \mathbb{G}_2 . In each case, there are two different root lengths (“long” and “short”). Our assumption that Φ is the smallest root system containing $\{\beta_1, \dots, \beta_n\}$ implies that among the roots β_i there is a long one and a short one. This finishes the story for \mathbb{G}_2 .

Now let Φ be of type \mathbb{B}_n , with the standard choice of simple roots $\{\alpha_1, \dots, \alpha_n\} = \{e_1 - e_2, \dots, e_{n-1} - e_n, e_n\}$, where e_1, \dots, e_n is the standard basis in the root lattice \mathbb{Z}^n . The long (resp. short) roots form a subsystem of type \mathbb{D}_n (resp. \mathbb{A}_1^n) consisting of the roots $\pm e_i \pm e_j$ ($1 \leq i < j \leq n$) (resp. $\pm e_i$ ($1 \leq i \leq n$)). To prove (2.3), we observe the following: every subset of Φ containing at most $n - 2$ long roots, is contained in a root subsystem of type $\mathbb{B}_k \times \mathbb{B}_{n-k}$ for some $k = 1, \dots, n - 1$ (to see this, represent each long root $\pm e_i \pm e_j$ by an edge joining i and j , and use the obvious fact that every graph with n vertices and at most $n - 2$ edges is disconnected). Therefore, the set $\{\beta_1, \dots, \beta_n\}$ must consist of $n - 1$ long roots and a short one, proving (2.3). The type \mathbb{C}_n is treated in the same way (with long and short roots interchanged).

Finally, let Φ be of type \mathbb{F}_4 , with the simple roots $\{\alpha_1, \dots, \alpha_4\} = \{e_2 - e_3, e_3 - e_4, e_4, \frac{1}{2}(e_1 - e_2 - e_3 - e_4)\}$ (see [1]). In view of the symmetry between the long and short roots, to prove (2.3) it suffices to show that $\{\beta_1, \dots, \beta_4\}$ cannot consist of one long root and three short ones. Suppose on the contrary that, say $\beta_1 = \pm e_1 \pm e_2$, and each of the roots β_2, β_3 and β_4 is of the form $\pm e_i$ or $\frac{1}{2}(\pm e_1 \pm e_2 \pm e_3 \pm e_4)$. An easy inspection shows that $W\{\beta_1, \dots, \beta_4\}$ is then contained in the root subsystem of Φ consisting of all short roots and the eight long roots $\{\pm e_1 \pm e_2, \pm e_3 \pm e_4\}$, again

in contradiction with our assumption that Φ is the smallest root system containing $\{\beta_1, \dots, \beta_4\}$.

This concludes the proof of (2.3) and hence of the equivalence of conditions (1) – (4). It is well known that positive Cartan matrices of different types are not equivalent to each other. The rest of the statements in Proposition 2.9 have been already established in the course of the above argument. \square

Remark 2.10. In the case where a quasi-Cartan matrix is assumed to be *symmetric*, a proof of Proposition 2.9 was essentially given in [5, 1.2] (using the language of unit quadratic forms).

3. MUTATIONS AND QUASI-CARTAN COMPANIONS

In this section, we show that, under some conditions, the mutation-equivalence of skew-symmetrizable matrices can be extended to an equivalence (in the sense of Definition 2.8) of properly chosen quasi-Cartan companions.

Definition 3.1. For a matrix index k , we say that a quasi-Cartan companion A of a skew-symmetrizable matrix B is *k-compatible* with B if the signs of its entries satisfy

$$(3.1) \quad \text{if } B_{ik} > 0 \text{ and } B_{kj} > 0 \text{ for some } i, j \text{ then } \text{sgn}(A_{ik}A_{kj}A_{ji}) = \text{sgn}(B_{ji}).$$

Proposition 3.2. *Let A be a k -compatible quasi-Cartan companion of a skew-symmetrizable matrix B . Then there exists a k -compatible quasi-Cartan companion A' of $B' = \mu_k(B)$ such that $A' \sim A$. Explicitly, the signs of off-diagonal matrix entries of A' can be chosen as follows:*

$$(3.2) \quad \text{sgn}(A'_{ij}) = \text{sgn}(A'_{ji}) = \begin{cases} \text{sgn}(B_{ik}) \text{sgn}(A_{ik}) & \text{if } i \neq k, j = k; \\ \text{sgn}(B_{ij}B'_{ij}) \text{sgn}(A_{ij}) & \text{if } i \neq k, j \neq k. \end{cases}$$

Proof. We fix k and introduce the following three $n \times n$ matrices:

- J is the diagonal matrix with $J_{kk} = -1$ and $J_{ii} = 1$ for $i \neq k$.
- E is the matrix with all the entries outside the k -th column equal to 0, and the k -th column entries given by

$$E_{ik} = \begin{cases} A_{ik} & \text{if } B_{ik} > 0; \\ 0 & \text{otherwise.} \end{cases}$$

- F is the matrix with all the entries outside the k -th row equal to 0, and the k -th row entries given by

$$F_{kj} = \begin{cases} A_{kj} & \text{if } B_{kj} < 0; \\ 0 & \text{otherwise.} \end{cases}$$

We now set

$$(3.3) \quad A' = (J - E)A(J - F),$$

and claim that A' satisfies all the required properties. First, a direct calculation shows that the entries of $A' = JAJ - EAJ - JAF + EAF$ are given by

$$(3.4) \quad A'_{ij} = \begin{cases} 2 & \text{if } i = j = k; \\ \operatorname{sgn}(B_{ik})A_{ik} & \text{if } i \neq k = j; \\ -\operatorname{sgn}(B_{kj})A_{kj} & \text{if } i = k \neq j; \\ A_{ij} - \operatorname{sgn}(A_{ik}A_{kj})[B_{ik}B_{kj}]_+ & \text{if } i \neq k, j \neq k. \end{cases}$$

Comparing (3.3) with (1.1) and using (3.1), it is easy to see that A' is a quasi-Cartan companion of B' , and that it satisfies (3.2). Furthermore, the claim that A' is k -compatible with B' is a direct consequence of (3.2). Finally, to show that $A' \sim A$, we note that

$$D(J - E) = (D(J - F))^T = (J - F)^T D.$$

Therefore,

$$C' = DA' = D(J - E)A(J - F) = (J - F)^T DA(J - F) = (J - F)^T C(J - F),$$

and we are done. \square

Corollary 3.3. *Suppose a skew-symmetrizable matrix B satisfies condition (4) in Theorem 1.2, and let A be a positive quasi-Cartan companion of B . Then $\mu_k(B)$ has a positive quasi-Cartan companion for any matrix index k .*

Proof. In view of Lemma 2.1 (b), A is k -compatible with B . By Proposition 3.2, the matrix $\mu_k(B)$ has a quasi-Cartan companion which is equivalent to A and so is positive. \square

In the rest of this section, we use Proposition 3.2 to obtain more information on positive edge-weighted graphs, see Definition 2.3. First a piece of notation. Let Γ be an edge-weighted graph, $i \neq j$ two vertices of Γ , and t a positive integer. We denote by $\Gamma[i, j; t]$ the edge-weighted graph obtained from Γ by adjoining t new vertices i_1, \dots, i_t and $t+1$ new edges $\{i, i_1\}, \{i_1, i_2\}, \dots, \{i_{t-1}, i_t\}, \{i_t, j\}$, all of weight one.

Lemma 3.4 (Chain Contraction). *If the graph $\Gamma[i, j; t]$ is positive then so is $\Gamma[i, j; 1]$.*

As a preparation for the proof, we recall several facts from [3]. First recall from [3, Definition 7.3] that every skew-symmetrizable matrix B has the *diagram* $\Gamma(B)$ which is a directed edge-weighted graph. As an edge-weighted graph, it is completely analogous to its symmetrizable counterpart in Definition 2.2: the vertices correspond to matrix indices, with an edge $\{i, j\}$ for each $i \neq j$ with $B_{ij} \neq 0$, supplied with the *weight* $|B_{ij}B_{ji}|$. Note that the edge weights satisfy (2.2). The only difference between the two kinds of diagrams is in the way of encoding the signs of matrix entries: while for a quasi-Cartan matrix A , an edge $\{i, j\}$ in $\Gamma(A)$ is assigned the sign $\varepsilon_{ij} = -\operatorname{sgn}(A_{ij}) = -\operatorname{sgn}(A_{ji})$, in $\Gamma(B)$ we have an arrow $i \rightarrow j$ whenever $B_{ij} > 0$ (and so $B_{ji} < 0$). As before, in drawing the diagram $\Gamma(B)$, all unspecified weights will be equal to 1.

The next proposition reproduces [3, Proposition 8.1].

Proposition 3.5. *For a skew-symmetrizable matrix B , the diagram $\Gamma' = \Gamma(\mu_k(B))$ is uniquely determined by the diagram $\Gamma = \Gamma(B)$ and a matrix index k . Specifically, Γ' is obtained from Γ as follows:*

- *The orientations of all edges incident to k are reversed, their weights intact.*
- *For any vertices i and j which are connected in Γ via a two-edge oriented path going through k (refer to Figure 1 for the rest of notation), the direction of the edge (i, j) in Γ' and its weight c' are uniquely determined by the rule*

$$(3.5) \quad \pm\sqrt{c} \pm \sqrt{c'} = \sqrt{ab},$$

where the sign before \sqrt{c} (resp., before $\sqrt{c'}$) is “+” if i, j, k form an oriented cycle in Γ (resp., in Γ'), and is “-” otherwise. Here either c or c' can be equal to 0.

- *The rest of the edges and their weights in Γ remain unchanged.*




FIGURE 1. Diagram mutation

Proof of Lemma 3.4. It suffices to show the following:


$$(3.6) \quad \text{If } t \geq 2 \text{ then the positivity of } \Gamma[i, j; t] \text{ implies that of } \Gamma[i, j; t - 1].$$

Choose any skew-symmetrizable matrix B such that $\Gamma(B) = \Gamma[i, j; t]$ (as edge-weighted graphs), and the edges $\{i, i_1\}, \{i_1, i_2\}, \dots, \{i_{t-1}, i_t\}, \{i_t, j\}$ are oriented as $i \rightarrow i_1 \rightarrow \dots \rightarrow i_t \rightarrow j$. Since $\Gamma[i, j; t]$ is assumed to be positive, B has a positive quasi-Cartan companion A . Clearly, our choice of orientations implies that A is i_t -compatible with B . By Proposition 3.2, the diagram $\Gamma' = \Gamma(\mu_{i_t}(B))$ is positive (as an edge-weighted graph). By (3.5), Γ' is obtained from $\Gamma[i, j; t]$ by adjoining the edge $\{i_{t-1}, j\}$. Thus, Γ' has $\Gamma[i, j; t - 1]$ as an induced subgraph, and so $\Gamma[i, j; t - 1]$ is positive. This concludes the proof of (3.6) and Lemma 3.4. \square

We conclude this section with one more application of Lemma 3.4, which allows us to construct a family of non-positive edge-weighted graphs.

Lemma 3.6. *Let Γ be an edge-weighted graph obtained from a cycle Z with all edge weights 1 by adjoining a vertex k joined with an even number of vertices of Z with edges of weights not exceeding 2. Then Γ is non-positive, with the only exception, where k is joined with just two adjacent vertices of Z , both edges having weight 1.*

Proof. First assume that at least one edge through k has weight 2. In view of (2.2), then every edge through k must have weight 2. If two of these edges connect k with non-adjacent vertices of Z , then they form a subdiagram $\tilde{\mathbb{C}}_2$, so Γ cannot be positive. This leaves us with the case where k is joined with just two adjacent vertices of Z , both edges having weight 2. Then the graph Γ is of the form $\Delta[i, j; t]$, where Δ is a triangle with edge weights 2, 2 and 1, the edge $\{i, j\}$ being of weight 1.



By Lemma 3.4, the non-positivity of Γ follows from that of a 4-vertex graph $\Delta[i, j; 1]$ obtained from Δ by adjoining one more vertex $\ell = i_1$ and two edges $\{i, \ell\}$ and $\{j, \ell\}$, both of weight 1, as shown in the picture above. Arguing as in the proof of Lemma 3.4, choose a skew-symmetrizable matrix B such that $\Gamma(B) = \Delta[i, j; 1]$ (as edge-weighted graphs), and both triangles $\{i, j, k\}$ and $\{i, j, \ell\}$ are cyclically oriented. Performing the mutation μ_ℓ destroys the edge $\{i, j\}$, and so transforms $\Delta[i, j; 1]$ into a 4-cycle with edge weights (in cyclical order) 2, 2, 1, 1. By Proposition 2.6, the latter cycle is non-positive, hence by Corollary 3.3, so are $\Delta[i, j; 1]$ and Γ .

It remains to consider the case where all the $2p$ edges through k have weight 1. We start with the case $p = 1$, that is, k is joined with two non-adjacent vertices of Z . Then Γ has exactly three chordless cycles, with each edge belonging to exactly two of these cycles. This makes it impossible to attach the signs to edges to satisfy the sign condition in Proposition 2.6, so Γ cannot be positive.

Next let $p = 2$, that is, k is joined with four vertices of Z . Applying Lemma 3.4 in the same way as above, we can assume that Z has no other vertices, so Γ is made up of four triangles. Orienting all the edges so that these triangles become all cyclically oriented, and performing the mutation μ_k transforms Γ into the graph $\tilde{\mathbb{D}}_4$. Thus, Γ is non-positive in this case as well.

If $p = 3$, arguing as above we can assume that Γ is built from a 6-cycle Z with all its vertices joined with k . Orienting the edges of Γ so that these triangles become all cyclically oriented, and performing mutations at two opposite vertices of Z destroys four out of six edges through k , leaving us with the case $p = 1$ which we already dealt with.

Finally, if $p \geq 4$ then Γ contains a subdiagram $\tilde{\mathbb{D}}_4$, and so is again non-positive by Corollary 2.5. \square

4. PROOF OF THEOREM 1.2

The following result is the key ingredient of our proof of Theorem 1.2.

Lemma 4.1. *Mutations of skew-symmetrizable matrices preserve property (4) in Theorem 1.2.*

Proof of Theorem 1.2. To deduce Theorem 1.2 from Lemma 4.1, we show that the conditions (2) and (3) in Theorem 1.1 satisfy the implications (2) \implies (4) \implies (3). The implication (4) \implies (3) is immediate from Lemma 4.1 and Lemma 2.1. As for (2) \implies (4), it is again immediate from Lemma 4.1 once we observe that a matrix B in condition (2) of Theorem 1.1 satisfies (4) (by virtue of the Cartan-Killing classification, the graph $\Gamma(B)$ is a Dynkin diagram, hence a tree, so there are no (chordless) cycles, and the corresponding condition in (4) is vacuous). \square

Turning to the proof of Lemma 4.1, we will deduce it from a sharper statement (Lemma 4.4 below). To state it, we need some more terminology.

The following definition is motivated by Propositions 2.4 and 2.6, and Lemma 3.6.

Definition 4.2. An edge-weighted graph Γ will be called *quasi-finite* if it satisfies the following conditions:

- any induced subgraph of Γ which is a tree, is a Dynkin diagram.
- any induced subgraph of Γ which is a chordless cycle, is of one of the types (a), (b), (c) in Proposition 2.6.
- none of the non-positive edge-weight graphs in Lemma 3.6 appears as an induced subgraph of Γ .

We say that a skew-symmetrizable matrix B and its diagram $\Gamma(B)$ are quasi-finite if so is the underlying edge-weighted graph.

Remark 4.3. Any skew-symmetrizable matrix B which has a positive quasi-Cartan companion is quasi-finite as follows from Propositions 2.4 and 2.6, and Lemma 3.6.

We will say that a skew-symmetrizable matrix B (and its diagram $\Gamma(B)$) is *cyclically oriented* if every chordless cycle in $\Gamma(B)$ is cyclically oriented. We will deduce Lemma 4.1 from the following statement.

Lemma 4.4. *Suppose that two skew-symmetrizable matrices B and $B' = \mu_k(B)$ are both quasi-finite. Then B is cyclically oriented if and only if so is B' .*

Assuming Lemma 4.4, we can prove Lemma 4.1 as follows.

Proof of Lemma 4.1. Suppose that a skew-symmetrizable matrix B satisfies (4), and let A be a positive quasi-Cartan companion of B . We need to show that $B' = \mu_k(B)$ satisfies (4) for any k . Note that B is quasi-finite by Remark 4.3 and cyclically oriented by the definition of (4). Note also that A is k -compatible with B : to check (3.1), combine the fact that B is cyclically oriented with Lemma 2.1 (b). Applying Proposition 3.2, we see that $B' = \mu_k(B)$ has a positive quasi-Cartan companion A' and therefore B' is quasi-finite, by Remark 4.3. By Lemma 4.4, B' is cyclically oriented, and so satisfies (4), as desired. \square

To finish the proof of Theorem 1.2, it remains to prove Lemma 4.4. Using Proposition 3.5, in our proof of Lemma 4.4 we can forget about skew-symmetrizable matrices and just work with diagrams and their mutations given by the rules in Proposition 3.5. So in the following argument by a diagram we will just mean a directed edge-weighted graph Γ satisfying (2.2).

Proof of Lemma 4.4. Let Γ be a diagram which is quasi-finite but not cyclically oriented, and let Γ' be obtained from Γ by mutation in some direction k (here we think of k as a vertex of Γ). We need to show that Γ' is either not quasi-finite, or not cyclically oriented. The proof will split into several cases.


Case 1. k belongs to some chordless cycle in Γ which is not cyclically oriented. Without loss of generality we can assume that this cycle is the whole diagram Γ . Let $\{i, k\}$ and $\{j, k\}$ be the two edges through k in Γ .

Case 1.1. The edges $\{i, k\}$ and $\{j, k\}$ are oriented both towards k or both away from k . Then Γ' is obtained from Γ by just reversing the orientations of these edges, and so is not cyclically oriented.

Case 1.2. The edges $\{i, k\}$ and $\{j, k\}$ are oriented as $i \rightarrow k \rightarrow j$.

Case 1.2.1. The cycle Γ has at least four vertices. Then in Γ' we have $j \rightarrow k \rightarrow i \rightarrow j$, and so, removing the vertex k from Γ' leaves us with a cycle which is again not cyclically oriented.

Case 1.2.2. The cycle Γ is a triangle with vertices i, j, k and weighted edges oriented as follows:



Applying (3.5), we see that the weight c' of the edge $\{i, j\}$ in Γ' satisfies

$$c' = (\sqrt{ab} + \sqrt{c})^2 = ab + 2\sqrt{abc} + c \geq 4,$$

so in this case Γ' is not quasi-finite.

Case 2. k does not belong to any non-oriented chordless cycle in Γ . Since Γ is not cyclically oriented, it has at least one non-oriented chordless cycle Z . Without loss of generality we can assume that Γ is just the union of Z and $\{k\}$. In order for the mutation μ_k to change Z , there must be at least two edges through k in Γ . Furthermore, if we list these edges in the cyclical order along Z , then the incoming and outgoing edges must alternate, in particular, there must be an even number of them; otherwise, k would belong to a non-oriented cycle. Since Γ is assumed to be quasi-finite, the cycle Z is of one of the three types (a), (b), or (c) in Proposition 2.6.

Case 2.1. Z is of type (b), i.e., it is a non-oriented triangle with edge weights 2, 2, and 1. Let k be joined with vertices i and j of Z , with orientations $i \rightarrow k \rightarrow j$. Since the triangle $\{i, j, k\}$ in Γ must be oriented, the edge $\{i, j\}$ is oriented as $j \rightarrow i$. The mutation μ_k reverses the orientations of the edges $\{i, k\}$ and $\{j, k\}$, and changes the weight and orientation of the edge $\{i, j\}$ in accordance with (3.5) (in particular, the latter edge can be absent in Γ'). An easy inspection shows that the only chance for Γ' to be cyclically oriented, while Γ is not cyclically oriented, is to have the edge $\{i, j\}$ to be of opposite orientations in Γ and Γ' . By (3.5), this can only happen if $\{i, j\}$ has weight 1 in Γ , while each of the remaining four edges has weight 2. But then Γ has a subdiagram of type $\widetilde{\mathbb{C}}_2$ (see Corollary 2.5), so cannot be quasi-finite.

Case 2.2. Z is of type (c), i.e., it is a non-oriented 4-vertex cycle with edge weights (in cyclic order) 1, 2, 1, 2.




FIGURE 2. Different cases to be considered

Case 2.2.1. k is joined precisely with the two adjacent vertices i and j of Z . The same argument as in Case 2.1 shows that, the only chance to have Γ' cyclically

oriented and Γ not cyclically oriented appears if $\{i, j\}$ has weight 1 in Γ , while both $\{i, k\}$ and $\{j, k\}$ have weight 2. But then Γ has a subdiagram of type $\widetilde{\mathbb{C}}_2$, so cannot be quasi-finite.

Case 2.2.2. k is joined precisely with the two opposite vertices i and j of Z . This case is even easier than the previous one: in view of Proposition 2.6, to satisfy (2.2), the edges $\{i, k\}$ and $\{j, k\}$ must have weights 1 and 2, which again forces Γ to have a subdiagram of type $\widetilde{\mathbb{C}}_2$.

Case 2.2.3. k is joined with all four vertices of Z . The only balance of edge weights that makes Γ quasi-finite is the one shown in Figure 2. Furthermore, since we are still in Case 2, all four triangles containing k must be cyclically oriented. Performing the mutation μ_k , we obtain a diagram, which contains a subdiagram $\widetilde{\mathbb{B}}_3$ and so is not quasi-finite.

Case 2.3. Z is of type (a), i.e., has all edge weights equal to 1. According to Lemma 3.6, if Γ is quasi-finite then k is joined with just two adjacent vertices i and j of Z , both edges $\{i, k\}$ and $\{j, k\}$ having weight 1. As before, the triangle $\{i, j, k\}$ must be cyclically oriented. Applying the mutation μ_k destroys the edge $\{i, j\}$, and so transforms Γ into a chordless cycle. If this cycle is cyclically oriented, then so is the original cycle Z , contradicting our assumption. This completes the proofs of Lemma 4.4 and Theorem 1.2. \square

5. CYCLICALLY ORIENTABLE GRAPHS

This section is purely graph-theoretic. We call a graph Γ *cyclically orientable* if it admits an orientation in which any chordless cycle is cyclically oriented. For example, the full graph on four vertices is *not* cyclically orientable. We will give several properties of graphs that are equivalent to cyclical orientability. This requires some notation.

For a given graph Γ , we denote by $\text{Ver}(\Gamma)$ its set of vertices, by $\text{Edg}(\Gamma)$ its set of edges, by $\text{Con}(\Gamma)$ its set of connected components, and by $\text{Cyc}(\Gamma)$ its set of chordless cycles; if no confusion can arise, we drop the dependence of Γ . For a finite set X , let \mathbb{F}_2^X denote the vector space of functions $f : X \rightarrow \mathbb{F}_2$ with the values in the 2-element field \mathbb{F}_2 . To any incidence relation $I \subset X \times Y$ between finite sets X and Y we associate a linear map $\rho = \rho_I : \mathbb{F}_2^X \rightarrow \mathbb{F}_2^Y$ given by

$$(5.1) \quad (\rho f)(y) = \sum_{(x,y) \in I} f(x) .$$

In particular, there is a natural sequence of linear maps:

$$(5.2) \quad 0 \rightarrow \mathbb{F}_2^{\text{Con}} \rightarrow \mathbb{F}_2^{\text{Ver}} \rightarrow \mathbb{F}_2^{\text{Edg}} \rightarrow \mathbb{F}_2^{\text{Cyc}} \rightarrow 0 ,$$

where all the maps are of the form (5.1) associated to the obvious incidence relations.

Theorem 5.1. *Each of the following conditions on a finite graph Γ is equivalent to cyclical orientability:*

(5.3) *The edges of Γ can be linearly ordered so that different chordless cycles in Γ will have different maximal edges.*

(5.4) $|\text{Cyc}| = |\text{Edg}| - |\text{Ver}| + |\text{Con}|$.

(5.5) *The sequence (5.2) is exact.*

Using (5.5), we obtain the following corollary implying Proposition 1.5.

Corollary 5.2. *In a cyclically orientable graph, one can attach the signs ± 1 to all edges in such a way that the product of these signs along every chordless cycle is equal to -1 . Furthermore, such an attachment is unique up to simultaneous changes of signs for all edges incident to a given vertex.*

As a preparation for the proof of Theorem 5.1, we show that certain weaker versions of (5.3)-(5.5) hold for arbitrary finite graphs.

Proposition 5.3. *The following properties hold for an arbitrary finite graph Γ :*

(5.6) *The edges of Γ can be linearly ordered so that an edge is maximal in some chordless cycle if and only if it is one of the last $|\text{Edg}| - |\text{Ver}| + |\text{Con}|$ edges on the list.*

(5.7) $|\text{Cyc}| \geq |\text{Edg}| - |\text{Ver}| + |\text{Con}|$.

(5.8) *The sequence $0 \rightarrow \mathbb{F}_2^{\text{Con}} \rightarrow \mathbb{F}_2^{\text{Ver}} \rightarrow \mathbb{F}_2^{\text{Edg}} \rightarrow \mathbb{F}_2^{\text{Cyc}}$ is exact.*

Proof. To construct a linear ordering of edges satisfying (5.6), take the first $|\text{Ver}| - |\text{Con}|$ edges (in an arbitrary order) so that they form a spanning forest in Γ (the union of spanning trees in all connected components). Then order the remaining $|\text{Edg}| - |\text{Ver}| + |\text{Con}|$ edges so that every new edge connects two vertices at the minimal possible distance in the graph formed by the preceding edges. By virtue of this construction, each of the last $|\text{Edg}| - |\text{Ver}| + |\text{Con}|$ edges becomes the maximal edge in at least one chordless cycle of Γ .

The bound (5.7) follows at once from (5.6).

Finally, the only statement in (5.8) which may be not immediately clear is the inclusion

$$\text{Ker}(\mathbb{F}_2^{\text{Edg}} \rightarrow \mathbb{F}_2^{\text{Cyc}}) \subset \text{Im}(\mathbb{F}_2^{\text{Ver}} \rightarrow \mathbb{F}_2^{\text{Edg}}).$$

To prove this inclusion, first consider the case where Γ is a forest. In this case, we need to show that the map $\mathbb{F}_2^{\text{Ver}} \rightarrow \mathbb{F}_2^{\text{Edg}}$ is surjective, which is well-known (and follows easily by induction on the number of edges). Returning to the general case, choose a linear ordering of the edges satisfying (5.6). It is clear that a function $g \in \text{Ker}(\mathbb{F}_2^{\text{Edg}} \rightarrow \mathbb{F}_2^{\text{Cyc}})$ is uniquely determined by its restriction to the first $|\text{Ver}| - |\text{Con}|$ edges. Since these edges form a forest, there is a function $f \in \mathbb{F}_2^{\text{Ver}}$ whose image in $\mathbb{F}_2^{\text{Edg}}$ agrees with g on them; therefore, the image of f is g , as desired. \square


Proof of Theorem 5.1, (5.3) \iff (5.4) \iff (5.5). The implication (5.5) \implies (5.4) is clear, while (5.4) \implies (5.3) is immediate from (5.6). Finally, to show (5.3) \implies (5.5), we note that, in view of (5.8), the exactness of the sequence (5.2) simply means that

its right-most map $\mathbb{F}_2^{\text{Edg}} \rightarrow \mathbb{F}_2^{\text{Cyc}}$ is surjective. The latter condition follows easily from (5.3), and we are done. \square

To complete the proof of Theorem 5.1, we will show that the cyclical orientability of Γ is equivalent to (5.3). First some preparation. Let i and j be two non-adjacent vertices of a finite graph Γ' , and let Γ be obtained from Γ' by adjoining the edge $\{i, j\}$. By a *chain* connecting i and j in Γ' we will mean a sequence of $t \geq 3$ distinct vertices $(i = i_1, i_2, \dots, i_t = j)$ such that the only edges between them are $\{i_p, i_{p+1}\}$ for $p = 1, \dots, t - 1$. It is then clear that the chordless cycles in Γ containing the edge $\{i, j\}$ are in a bijection with the chains connecting i and j in Γ' . In particular, i and j belong to a *unique* chordless cycle in Γ if and only if they are connected by a *unique* chain in Γ' .

Lemma 5.4. *Suppose Γ' is a cyclically orientable graph, and i and j are two non-adjacent vertices of Γ' connected by a unique chain $(i = i_1, i_2, \dots, i_t = j)$. Then there exists an orientation of Γ' making all chordless cycles cyclically oriented and having the chain (i_1, i_2, \dots, i_t) linearly oriented: $i_1 \rightarrow i_2 \rightarrow \dots \rightarrow i_t$.*

Proof. Let Γ'' be the induced subgraph of Γ' on the set of vertices $\text{Ver}(\Gamma') - \{i_1, \dots, i_t\}$. For $p = 1, \dots, t - 1$, let Γ_p denote the union of all connected components of Γ'' joined by an edge with i_p and i_{p+1} . We claim that the subgraphs $\Gamma_1, \dots, \Gamma_{t-1}$ are pairwise disjoint, and so are disconnected from each other in Γ'' . Indeed, if $\Gamma_p \cap \Gamma_q \neq \emptyset$ for some $p < q$ then i_p and i_{q+1} would be connected by a chain having all interior vertices in Γ'' ; but this contradicts the uniqueness of a chain connecting i and j . Clearly, if an edge $\{i_p, i_{p+1}\}$ belongs to some chordless cycle Z in Γ' , then the rest of the vertices of Z belong to Γ_p .



Now take any orientation of Γ' making all the chordless cycles cyclically oriented. If all the edges $\{i_p, i_{p+1}\}$ are oriented as $i_p \rightarrow i_{p+1}$, there is nothing to prove. Otherwise, modify the orientation as follows: reverse every directed edge $i_{p+1} \rightarrow i_p$ together with all the edges inside Γ_p and all the edges connecting i_p and i_{p+1} with Γ_p . Clearly, the new orientation still has all the chordless cycles in Γ' cyclically oriented, so we are done. \square

Proof of Theorem 5.1, (5.3) \implies cyclical orientability. Thus, we assume that Γ is a finite graph whose edges are ordered in accordance to (5.3). To show that Γ is cyclically orientable, we proceed by induction on the number of edges N of Γ . If $N = 0$, there is nothing to prove; so we assume that Γ has at least one edge. Let $e = \{i, j\}$ be the maximal edge in the chosen ordering, and let Γ' be the graph obtained from Γ by removing e . Since e belongs to at most one chordless cycle in Γ , it is easy to see that the chordless cycles in Γ' are precisely the chordless cycles in Γ that do not contain e . It follows that Γ' satisfies (5.3). By induction, we can assume that Γ' is cyclically orientable. Choose an orientation of Γ' making all the chordless cycles cyclically oriented. If Γ has no chordless cycle containing e then i and j belong to different connected components of Γ' , and so e can be oriented either way, making

all chordless cycles in Γ cyclically oriented. This leaves us with the case where Γ has a unique chordless cycle containing e . Then i and j are connected by a unique chain in Γ' . By Lemma 5.4, we can assume that this chain is linearly oriented in a chosen orientation of Γ' . Therefore, orienting e so that the only chordless cycle containing e becomes cyclically oriented, yields an orientation of Γ making all chordless cycles cyclically oriented, thus showing that Γ is cyclically orientable. \square

It remains to show that a cyclically orientable graph satisfies (5.3). Again we need some preparation. We start with two lemmas.


Lemma 5.5. *Let Γ be a cyclically orientable graph, and suppose two chordless cycles Z and Z' share a common edge $\{i, j\}$. Then i and j are the only common vertices of Z and Z' , and $\{i, j\}$ is the only edge joining a vertex of Z with a vertex of Z' .*

Proof. Let i, j, i_1, \dots, i_p (resp. i, j, i'_1, \dots, i'_q) be all the vertices of Z (resp. Z') written in the cyclical order. Without loss of generality, we can assume that $i_1 \neq i'_1$ (this can be always achieved by replacing $\{i, j\}$ with another edge if necessary).

To show that i and j are the only common vertices of Z and Z' , we suppose that $i_k = i'_\ell$ for some k and ℓ , and choose the smallest possible k with this property. Thus, the sequence of vertices $j, i_1, \dots, i_k, i'_{\ell-1}, \dots, i'_1, j$ forms a cycle of length $k + \ell$. Now choose $r \in \{1, \dots, k\}$ and $s \in \{1, \dots, \ell - 1\}$ so that i_r and i'_s are joined by an edge, and $r + s$ is smallest possible with this property. Then the vertices $j, i_1, \dots, i_r, i'_s, i'_{s-1}, \dots, i'_1$ form a chordless cycle, say Z'' . To obtain the desired contradiction, it remains to observe that no orientation of edges of Γ can make all three chordless cycles Z, Z' , and Z'' cyclically oriented: indeed, if say j is a head of $\{i, j\}$ then j must be a tail of both $\{j, i_1\}$ and $\{j, i'_1\}$, and so Z'' will not be cyclically oriented. Having proven that the vertices $i_1, \dots, i_k, i'_1, \dots, i'_\ell$ are all distinct, the same argument as above shows that there are no edges of the form $\{i_r, i'_s\}$, and we are done. \square

Lemma 5.6. *Suppose that in a cyclically orientable graph Γ , there is at least one edge through a vertex j . Then there is an edge through j that belongs to at most one chordless cycle.*

Proof. Fix an orientation of Γ making all chordless cycles cyclically oriented. Arguing by contradiction, suppose that every edge through j belongs to at least two chordless cycles. Start with an arbitrary edge $\{j, i_1\}$, let Z_1 be a chordless cycle containing this edge, and let $\{j, i_2\}$ be the second edge through j that belongs to Z_1 . Then pick a chordless cycle $Z_2 \neq Z_1$ that contains the edge $\{j, i_2\}$, and let $\{j, i_3\}$ be the second edge through j that belongs to Z_2 . Continuing in the same way, we construct a sequence of chordless cycles Z_1, Z_2, \dots and a sequence of vertices i_1, i_2, \dots such that each Z_k contains edges $\{j, i_k\}$ and $\{j, i_{k+1}\}$. Renaming the vertices if necessary, we can assume without loss of generality that the vertices i_1, \dots, i_p and the chordless cycles Z_1, \dots, Z_p are pairwise distinct, but $i_{p+1} = i_1$ and $Z_{p+1} = Z_1$. In view of Lemma 5.5, we have $p \geq 3$. Also since all the chordless cycles Z_1, \dots, Z_p must be cyclically oriented, p is even, and so $p \geq 4$.



Now let $i_1, \dots, i_2, \dots, i_3, \dots, i_p, \dots, i_1$ be the sequence of vertices in $Z_1 \cup \dots \cup Z_p - \{j\}$ written in the natural cyclical order, so that each interval i_k, \dots, i_{k+1} is a chain $Z_k - \{j\}$. Choose a shortest possible interval (in the cyclical order) j_1, \dots, j_q in this sequence consisting of $q \geq 3$ pairwise distinct vertices and such that $\{j_1, j_q\}$ is also an edge. (It is easy to see that such intervals exist because every three consecutive vertices in our sequence are distinct.) The vertices j_1, \dots, j_q obviously form a chordless cycle. Clearly, this chordless cycle is not a part of any $Z_k - \{j\}$. Thus, it must contain some i_k as an interior vertex. But this contradicts the cyclical orientability since the two edges through i_k belong to the chordless cycles Z_k and Z_{k-1} (with the convention that $Z_0 = Z_p$) and so either both are oriented towards i_k , or both are oriented away from i_k . This contradiction proves the result. \square

Proof of Theorem 5.1, cyclical orientability \implies (5.3). Let Γ be a cyclically orientable graph having at least one edge. By Lemma 5.6, Γ has an edge e belonging to at most one chordless cycle. Let Γ' be obtained from Γ by removing the edge e . While proving the reverse implication, we have already noticed that the chordless cycles in Γ' are precisely the chordless cycles in Γ that do not contain e . It follows that an orientation of Γ making all chordless cycles cyclically oriented, restricts to the orientation of Γ' with the same property. Thus, Γ' is cyclically orientable. By induction on the number of edges, there is a linear ordering of the edges of Γ' satisfying (5.3). Adding e as the maximal element gives a desired linear ordering of edges for Γ , finishing the proof of Theorem 5.1. \square

Remark 5.7. Theorem 5.1 suggests the following way to check whether a given skew-symmetrizable matrix satisfies condition (4) in Theorem 1.2. Start by checking whether the edges of (the underlying graph of) the diagram $\Gamma(B)$ can be ordered in accordance with (5.3) (we leave aside the question of a practical implementation of such a check). If (5.3) cannot be satisfied, then B is *not* of finite type. Otherwise, go over the list of edges ordered in accordance with (5.3), and attach the signs to them in the following way. If a current edge $\{i, j\}$ is not maximal in any chordless cycle, choose the sign ε_{ij} arbitrarily; otherwise, $\{i, j\}$ is maximal in a unique chordless cycle Z , and we determine ε_{ij} from the condition that the product of signs along the edges of Z is equal to -1 . Finally, define a quasi-Cartan companion A of B by setting $A_{ij} = -\varepsilon_{ij}|B_{ij}|$ for all $i \neq j$. Then B satisfies (4) and so is of finite type if and only if A is positive (the latter condition can be checked for example by using the Sylvester criterion).

ACKNOWLEDGMENTS

The first solution of the recognition problem for cluster algebras of finite type was given by A. Seven in [6] as a part of his Ph.D. thesis written under the supervision of A. Zelevinsky at Northeastern University. M. Barot and C. Geiss started working on an alternative solution after attending A. Seven's talk at the Sixth International Joint AMS-SMM Meeting in May 2004. They joined forces with A. Zelevinsky after he attended their presentation of a preliminary version of Theorem 1.2 at the XI International Conference on Representations of Algebras, Mexico, August 2004.

The authors thank Alex Postnikov for stimulating discussions and help in proving the results in Section 5.

REFERENCES

- [1] N. Bourbaki, *Groupes and algèbres de Lie, Ch. 4,5 et 6. Hermann, Paris 1968.*
- [2] S. Fomin, A. Zelevinsky, *Cluster algebras I: Foundations. J. Amer. Math. Soc.* **15** (2002), no. 2, 497-529.
- [3] S. Fomin, A. Zelevinsky, *Cluster algebras II: Finite type classification. Invent. Math.* **154** (2003), no 1. 63-121.
- [4] V. Kac, *Infinite dimensional Lie algebras*, 3rd edition, Cambridge University Press, 1990.
- [5] C. M. Ringel, *Tame algebras and integral quadratic forms*, Lecture Notes in Mathematics, **1099**, Springer-Verlag, Berlin, 1984.
- [6] A. Seven, *Recognizing cluster algebras of finite type*, *math.CO/0406545*.

INSTITUTO DE MATEMÁTICAS, UNAM, CIUDAD UNIVERSITARIA, 04510 MEXICO D.F., MEXICO

E-mail address: `barot@matem.unam.mx`

INSTITUTO DE MATEMÁTICAS, UNAM, CIUDAD UNIVERSITARIA, 04510 MEXICO D.F., MEXICO

E-mail address: `christof@matem.unam.mx`

DEPARTMENT OF MATHEMATICS, NORTHEASTERN UNIVERSITY, BOSTON, MA 02115, USA

E-mail address: `andrei@neu.edu`