

Cluster algebras and triangulated surfaces. Part I: Cluster complexes

by

SERGEY FOMIN

*University of Michigan
Ann Arbor, MI, U.S.A.*

MICHAEL SHAPIRO

*Michigan State University
East Lansing, MI, U.S.A.*

DYLAN THURSTON

*Columbia University
New York, NY, U.S.A.*

Contents

1. Introduction	84
2. Bordered surfaces with marked points and ideal triangulations . .	87
3. Arc complexes	91
4. Signed adjacency matrices and their mutations	95
5. Cluster algebras associated with triangulated surfaces	101
6. Types and diagrams	103
7. Tagged arc complexes	107
8. Denominators and intersection numbers	113
9. Proofs of main results	115
9.1. Basic properties of tagged arc complexes	116
9.2. Strata of tagged triangulations	117
9.3. Mutations and flips in the tagged case	118
9.4. Cycles in the graph of flips	119
9.5. Tropical exchange relations for intersection numbers	123
9.6. Proofs of Theorems 5.6, 7.11 and 8.6	126
10. On the topology of cluster complexes	128
11. Polynomial vs. exponential growth	131
12. Finite mutation classes	132
13. Block decompositions	136
14. Recovering topology from combinatorial data	139
Acknowledgments	143
References	144

This work was partially supported by NSF grants DMS-0245385 and DMS-0555880 (S. F.), DMS-0401178 and PHY-0555346 (M. S.), DMS-0071550 (D. T.), BSF grant 2002375 (M. S.), and an A. P. Sloan Research Fellowship (D. T.).

1. Introduction

Cluster algebras are a class of commutative rings endowed with an additional combinatorial structure, which involves a set of distinguished generators (cluster variables) grouped into overlapping subsets (clusters) of the same cardinality. The original motivation for cluster algebras came from representation theory, specifically the study of dual canonical bases and total positivity phenomena in semisimple Lie groups. (See [18] for an introduction to cluster algebras in the historical order of development.) In subsequent years, constructions involving cluster algebras were discovered in various mathematical contexts, including a topological/geometric one that is the main focus of this paper.

Shortly after the introduction of cluster algebras [16], cluster-algebraic structures in Teichmüller theory were defined and studied by M. Gekhtman, M. Shapiro and A. Vainshteyn [24] and, in a more general setting, by V. Fock and A. Goncharov [10], [12], making use of foundational work by R. Penner [33], [35] and V. Fock [9]. These cluster algebras were built using the following crucial observation: in hyperbolic geometry, an analogue of the classical Ptolemy theorem holds for the *lambda-lengths* of geodesic arcs triangulating a bordered and/or punctured surface with geodesic sides and cusps at the vertices (the arcs connect horocycles drawn around cusps and punctures); these lambda-lengths are the *Penner coordinates* on the corresponding *decorated Teichmüller space* [33], [34]. The constructions based on this observation made it possible to transport notions, results and insights from the theory of cluster algebras (cluster monomials, Laurent phenomenon, finite-type classification, canonical bases, etc.) into geometric settings.

This paper and its sequel [14] exploit the aforementioned connection in the opposite direction: we use ideas and techniques of combinatorial topology and hyperbolic geometry for the benefit of cluster algebra theory. We undertake a systematic study of the wide class of cluster algebras whose underlying combinatorial structures are determined and governed by the topology of bordered surfaces. We demonstrate that the pivotal role in the construction of a cluster algebra associated with a bordered surface is played by its *tagged arc complex*, a certain simplicial complex generalizing the classical notion of an arc complex.

Recall that a cluster algebra is built around a combinatorial scaffolding formed by *exchange matrices*, related to each other by *matrix mutations*. The basic observation in [10], [12], [24] is that this kind of structure arises when one considers *signed adjacency matrices* associated with triangulations of an oriented surface, with vertices at a fixed set of *marked points*. More specifically, matrix mutations arise as transformations of signed adjacency matrices that correspond to *flips* of triangulations.

Here is a sketch of the main construction studied in this paper, a class of cluster algebras associated with 2-dimensional surfaces, aimed at a reader familiar with the basics

of cluster algebra theory (see, e.g., [18]). Let \mathbf{S} be a connected oriented (2-dimensional) Riemann surface with boundary. Fix a finite set \mathbf{M} of *marked points* on \mathbf{S} that includes at least one marked point on each boundary component, plus possibly some interior points. For technical reasons, we need to exclude the possibility that \mathbf{S} is closed, with exactly two marked points. (This restriction will be removed in the sequel [14] to this paper.) Consider a triangulation T of \mathbf{S} whose vertices are located at marked points in \mathbf{M} and whose edges are simple pairwise non-intersecting curves (*arcs*), considered up to isotopy, connecting some of the marked points. The triangulation T defines a skew-symmetric matrix $B(T)$ defined by looking at the (signed) adjacencies between the edges of T . This matrix serves as an *exchange matrix* of a cluster algebra \mathcal{A} —actually, a whole class of such algebras, depending on the choice of *coefficients*. This class of cluster algebras is in fact independent of the choice of the triangulation T ; it only depends on the underlying surface (\mathbf{S}, \mathbf{M}) with marked points. This is because any two triangulations are connected by *flips*, and a flip corresponds to a *mutation* of the corresponding matrices $B(T)$.

The construction is more complicated than it may seem at first glance, because not every arc in any triangulation can be flipped, whereas the cluster algebra axioms do not allow exceptions: every cluster variable in every cluster is exchangeable. Consequently, not every seed in \mathcal{A} is naturally associated with a triangulation, and not every cluster variable is naturally labeled by an arc.

In this paper, we resolve these problems, and proceed further to establish the basic properties of cluster algebras associated with oriented bordered surfaces with marked points. In particular, we obtain the following results:

- *The seeds in \mathcal{A} are uniquely determined by their clusters* [Theorem 5.6].
- *The exchange graph and the cluster complex of \mathcal{A} are uniquely determined by the surface (\mathbf{S}, \mathbf{M}) ; i.e., they do not depend on the choice of coefficients* [Theorem 5.6].
- *The seeds containing a particular cluster variable form a connected subgraph of the exchange graph* [Theorem 5.6].
- *There are finitely many different exchange matrices appearing in the seeds of \mathcal{A}* [Corollary 12.3].
- *The class of exchange matrices of cluster algebras of topological origin (i.e., integer matrices that can be realized as signed adjacency matrices of some triangulated bordered surface) can be given a concrete combinatorial description* [Theorem 13.3].
- *The exponents appearing in denominators of Laurent expansions of cluster variables can be identified as certain intersection numbers* [Theorem 8.6].
- *The cluster complex of \mathcal{A} has a concrete combinatorial description in terms of “tagged arcs”* [Theorem 7.11].

- *The cluster complex of \mathcal{A} is a flag complex, i.e., it is the clique complex for its own 1-skeleton* [Theorem 5.6].
- *The cluster complex of \mathcal{A} is either contractible, or homotopy equivalent to a sphere* [Theorem 10.1].
- *The exchange graphs of such cluster algebras can be completely classified according to their growth (polynomial vs. exponential)* [§11].

The results above in particular establish, for the class of cluster algebras under consideration, several conjectures which are expected to hold for *any* cluster algebra ([18, Conjecture 4.14 (1)–(3)], [20, Conjecture 7.4] and Conjecture 5.5 of this paper).

Our main constructions depend only on the combinatorics of triangulated surfaces, and could in principle be formulated in entirely combinatorial language, with no reference to topology or geometry. In the sequel [14] to this paper, we are going to take a complementary, geometric approach to the construction and study of cluster algebras associated with bordered surfaces; this approach is based on an interpretation of cluster variables as certain generalized lambda-lengths.

Our results can be extended, with modifications, to cluster algebras whose exchange matrices are not necessarily skew-symmetric, such as the finite and affine types B and C . However, in this first paper we confine ourselves to the skew-symmetric case.

The paper is organized as follows. Requisite background on triangulated surfaces and cluster algebras is presented in §§2–4 and in §§5–6, respectively. In particular, we introduce, in §5, the class of cluster algebras \mathcal{A} whose exchange matrices are associated with a bordered surface with marked points (\mathbf{S}, \mathbf{M}) as explained above. The structural properties of any such algebra \mathcal{A} are stated in Theorem 5.6. In §7, we introduce our main new combinatorial construction of a *tagged arc complex* $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$. This simplicial complex is a pseudomanifold (Theorem 7.9). Our main result (Theorem 7.11) asserts that $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$ is in fact isomorphic to the cluster complex $\Delta(\mathbf{S}, \mathbf{M})$ of \mathcal{A} , provided (\mathbf{S}, \mathbf{M}) is not a closed surface with at most two marked points. Thus, the cluster variables in \mathcal{A} are naturally labeled by the tagged arcs in (\mathbf{S}, \mathbf{M}) , and they form a cluster if and only if the corresponding tagged arcs form a maximal simplex in $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$.

In §8, we introduce the notion of a (generalized) intersection number for tagged arcs, and show that these numbers satisfy the *tropical* version of the exchange relation recurrences. This leads to an interpretation of these intersection numbers as exponents in the denominators of Laurent expansions of cluster variables (see Theorem 8.6).

The proofs of Theorems 7.11 and 5.6 presented in §9 are based on a generalization to the tagged setup of a classical result from combinatorial topology (see Theorem 3.10) asserting that the fundamental group of the graph of flips is generated by cycles of length 4 and 5 (commuting flips plus the pentagon relation). In the tagged case, this property

fails for closed surfaces with two marked points, which explains our need to exclude these surfaces in the present paper.

In §10, we determine, up to homotopy, the topology of cluster complexes (or tagged arc complexes) associated with bordered surfaces. Specifically, we show that the tagged arc complex is homotopy equivalent to a sphere for polygons with at most one puncture and for closed surfaces (with punctures); in all other cases, the cluster complex is contractible.

In §11, we resolve the dichotomy of polynomial vs. exponential growth for the exchange graphs of the cluster algebras in question.

The mutation equivalence class formed by exchange matrices of a cluster algebra associated with a triangulated surface is finite. This leads to the discussion, in §12, of the general problem of recognizing (or classifying) finite mutation classes. Some examples of this kind (e.g., finite and affine types E_6 through E_8 , Grassmannians $\text{Gr}_{3,9}$ and $\text{Gr}_{4,8}$) do *not* come from triangulated surfaces.

The main result of §13 is a purely combinatorial description of the class of matrices B that arise as exchange matrices of cluster algebras associated with bordered surfaces. In other words, we characterize those matrices B for which there exists a bordered surface with marked points (\mathbf{S}, \mathbf{M}) such that $B = B(T)$ for some tagged (equivalently, ordinary—see Proposition 12.3) triangulation T of (\mathbf{S}, \mathbf{M}) .

The problem of recovering the topology of \mathbf{S} from the matrix $B(T)$ associated with a triangulation T of (\mathbf{S}, \mathbf{M}) is discussed in §14. For closed surfaces with punctures, all topological information can be recovered from $B(T)$; in general, it cannot. In Theorem 14.3, we compute the rank of $B(T)$ for an arbitrary (ideal) triangulation T . Specifically, the corank of $B(T)$ equals the number of punctures plus the number of boundary components with an even number of marked points.

2. Bordered surfaces with marked points and ideal triangulations

This section reviews some standard background from combinatorial topology of surfaces (cf., e.g., [6], [11], [27] and [29]), setting up terminology and notation to be used throughout the paper.

Definition 2.1. (Bordered surfaces with marked points) Let \mathbf{S} be a connected oriented 2-dimensional Riemann surface with boundary. Fix a finite set \mathbf{M} of *marked points* in the closure of \mathbf{S} . Marked points in the interior of \mathbf{S} are called *punctures*. In the terminology of [11] and [13], (\mathbf{S}, \mathbf{M}) is a *ciliated surface*.⁽¹⁾

⁽¹⁾ There is a slight difference between our setup and the one employed by V. Fock and A. Goncharov [11]: we use punctures where they use boundary components without marked points.

We will work with triangulations of \mathbf{S} whose vertices are located at the marked points in \mathbf{M} . We exclude situations in which such triangulations cannot be constructed or there is only one triangulation. Specifically, we assume that \mathbf{M} is non-empty, that there is at least one marked point on each connected component of the boundary of \mathbf{S} , and that (\mathbf{S}, \mathbf{M}) is none of the following:

- a sphere with one or two punctures;
- an unpunctured or once-punctured monogon;
- an unpunctured digon;
- an unpunctured triangle.

(In this paper, an m -gon is a disk with m marked points on the boundary.) It will also be convenient to exclude, for technical reasons,

- a sphere with three punctures.

A *bordered surface with marked points* is an object (\mathbf{S}, \mathbf{M}) satisfying the aforementioned constraints. All of our constructions depend on the choice of such an object.

Up to homeomorphism, (\mathbf{S}, \mathbf{M}) is defined by the following data:

- the genus g of the original Riemann surface;
- the number b of boundary components;
- the integer partition with b parts describing the number of marked points on each boundary component;
- the number p of punctures.

Definition 2.2. (Arcs) A (simple) *arc* γ in (\mathbf{S}, \mathbf{M}) is a curve in \mathbf{S} such that

- the endpoints of γ are marked points in \mathbf{M} ;
- γ does not intersect itself, except that its endpoints may coincide;
- except for the endpoints, γ is disjoint from \mathbf{M} and from the boundary of \mathbf{S} ;
- γ does not cut out an unpunctured monogon or an unpunctured digon. (In other words, γ is not contractible into \mathbf{M} or onto the boundary of \mathbf{S} .)

Each arc γ is considered up to isotopy inside the class of such curves.

The set $\mathbf{A}^\circ(\mathbf{S}, \mathbf{M})$ of all arcs in (\mathbf{S}, \mathbf{M}) is typically infinite; the cases where it is finite can be easily classified.

PROPOSITION 2.3. (See, e.g., [11, §2]) *The set $\mathbf{A}^\circ(\mathbf{S}, \mathbf{M})$ of all arcs in (\mathbf{S}, \mathbf{M}) is finite if and only if (\mathbf{S}, \mathbf{M}) is an unpunctured or once-punctured polygon.*

An arc whose endpoints coincide is called a *loop*.

Definition 2.4. (Compatibility of arcs) Two arcs are called *compatible* if they do not intersect in the interior of \mathbf{S} ; more precisely, there are curves in their respective isotopy classes which do not intersect in the interior of \mathbf{S} .

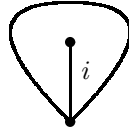


Figure 1. Self-folded ideal triangle.

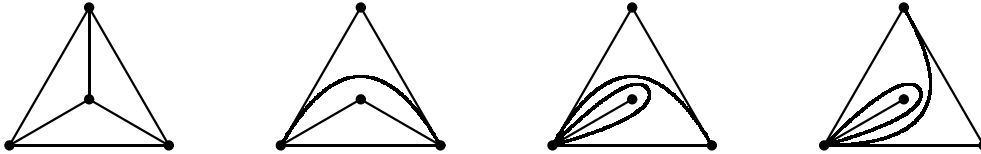


Figure 2. Ideal triangulations of a once-punctured triangle.

For instance, each arc is compatible with itself.

PROPOSITION 2.5. ([21]) *Any collection of pairwise compatible arcs can be realized by curves in their respective isotopy classes which do not intersect in the interior of \mathbf{S} .*

Definition 2.6. (Ideal triangulations; see, e.g., [29, §5.5]) A maximal collection of distinct pairwise compatible arcs is called an *ideal triangulation*. (Recall that we excluded all cases in which (\mathbf{S}, \mathbf{M}) cannot be triangulated.) The arcs of a triangulation cut the surface \mathbf{S} into *ideal triangles*. The three sides of an ideal triangle do not have to be distinct, i.e., we allow *self-folded* triangles. We also allow for a possibility that two triangles share more than one side.

Figure 1 shows a triangle folded along the arc i . Each side of an ideal triangle is either an arc or a segment of a boundary component between two marked points.

Example 2.7. Figure 2 shows four different kinds of ideal triangulations of a once-punctured triangle. In this case, there are 10 ideal triangulations altogether, obtained from these four by 120° and 240° rotations.

Remark 2.8. An alternative approach to working with triangulated surfaces is to start with a finite collection of disjoint oriented triangles, take a partial matching of their sides (i.e., a set of disjoint pairs of sides), then glue the triangles along the matched sides, making sure that orientations match.

Remark 2.9. Yet another model is that of *trivalent fat graphs* dual to (ideal) triangulations; see, e.g., [6], [10] and [12]. Note that by allowing self-folded triangles, we are allowing fat graphs with loops.

(\mathbf{S}, \mathbf{M})	g	b	p	c	n	Type
$(n+3)$ -gon	0	1	0	$n+3$	n	A_n
n -gon, 1 puncture	0	1	1	n	n	D_n
annulus, n_1+n_2 marked points	0	2	0	n_1+n_2	n_1+n_2	$\tilde{A}(n_1, n_2)$
$(n-3)$ -gon, 2 punctures	0	1	2	$n-3$	n	\tilde{D}_{n-1}
torus, 1 puncture	1	0	1	0	3	

Table 1. Examples of bordered surfaces with marked points.

The number n of arcs in an ideal triangulation is an invariant of (\mathbf{S}, \mathbf{M}) called the *rank* of (\mathbf{S}, \mathbf{M}) . (The terminology comes from the fact that n is the rank of any cluster algebra in a class, to be introduced later, associated with the surface (\mathbf{S}, \mathbf{M}) .) An easy count utilizing the Euler characteristic (see, e.g., [11, §2]) gives the following formula for the rank.

PROPOSITION 2.10. *Each ideal triangulation consists of*

$$n = 6g + 3b + 3p + c - 6$$

arcs, where g is the genus of \mathbf{S} , b is the number of boundary components, p is the number of punctures, and c is the number of marked points on the boundary.

COROLLARY 2.11. *For each positive integer n , there are finitely many (homeomorphism classes of) bordered surfaces (\mathbf{S}, \mathbf{M}) with marked points whose rank is equal to n .*

Example 2.12. By combining Proposition 2.10 with the restrictions in Definition 2.1, we can list all choices of (\mathbf{S}, \mathbf{M}) where the rank n is small (the nomenclature of “types” used below will be explained later):

- $n=1$: unpunctured square (type A_1);
- $n=2$: unpunctured pentagon (type A_2);
 - once-punctured digon (type $A_1 \times A_1$);
 - annulus with one marked point on each boundary component (type $\tilde{A}(1, 1)$);
- $n=3$: unpunctured hexagon (type A_3);
 - once-punctured triangle (type A_3);
 - annulus with one marked point on one boundary component and two marked points on another (type $\tilde{A}(2, 1)$);
 - once-punctured torus.

Table 1 shows the values of g , b , p , m and n for several examples of bordered surfaces with marked points.

LEMMA 2.13. *A bordered surface with marked points satisfying the restrictions in Definition 2.1 has a triangulation without self-folded triangles.*

Proof. Induction on n . Let \mathbf{S}' be the closed oriented surface obtained by gluing a disk to each boundary component of \mathbf{S} . If the genus of \mathbf{S}' is positive, then there is a loop γ in (\mathbf{S}, \mathbf{M}) that goes around a handle in \mathbf{S}' . Cut (\mathbf{S}, \mathbf{M}) open along γ , and proceed by induction.

It remains to treat the case of a sphere with holes and punctures. In the absence of boundary, connect the punctures cyclically by non-intersecting arcs; the polygons on both sides of the resulting closed curve can be triangulated without self-folded triangles. If there is boundary, say with components labeled $1, \dots, b$, then draw non-intersecting arcs connecting components 1 and 2, 2 and 3, etc., and cut them open to reduce the claim to the case of one boundary component (a disk with punctures). It follows from the restrictions in Definition 2.1 that there will be at least three marked points altogether, including at least one marked point on the boundary of the disk. Direct inspection shows that such a punctured disk can be triangulated without self-folded triangles. \square

3. Arc complexes

Definition 3.1. (Arc complex; cf. [25] and [26]) The *arc complex* $\Delta^\circ(\mathbf{S}, \mathbf{M})$ is the *clique complex* for the compatibility relation. That is, $\Delta^\circ(\mathbf{S}, \mathbf{M})$ is the simplicial complex on the ground set $\mathbf{A}^\circ(\mathbf{S}, \mathbf{M})$ of all arcs in (\mathbf{S}, \mathbf{M}) whose simplices are collections of distinct mutually compatible arcs, and whose maximal simplices are the ideal triangulations. As they are all of the same cardinality (see Proposition 2.10), the arc complex is *pure*.

Our terminology differs from that of [36] and [37], where the term “arc complex” is used for the quotient of $\Delta^\circ(\mathbf{S}, \mathbf{M})$ by the *pure mapping class group* of (\mathbf{S}, \mathbf{M}) .

Example 3.2. Figure 3 shows the arc complex of a hexagon, i.e., an unpunctured disk with six marked points on the boundary. The exterior region is a triangle, which should be filled in. It is then apparent that this arc complex is topologically a 2-dimensional sphere. Moreover, this complex can be realized as a boundary of a convex polytope, the (polar) dual for a 3-dimensional *associahedron*. See, e.g., [15, §3.1] and references therein.

Example 3.3. Figure 4 shows the arc complex of a once-punctured triangle, i.e., a disk with one puncture and three marked points on the boundary.

Example 3.4. Figure 5 shows the arc complex for an annulus with two marked points on one boundary component, and one on another. Each maximal simplex is labeled by the corresponding ideal triangulation.

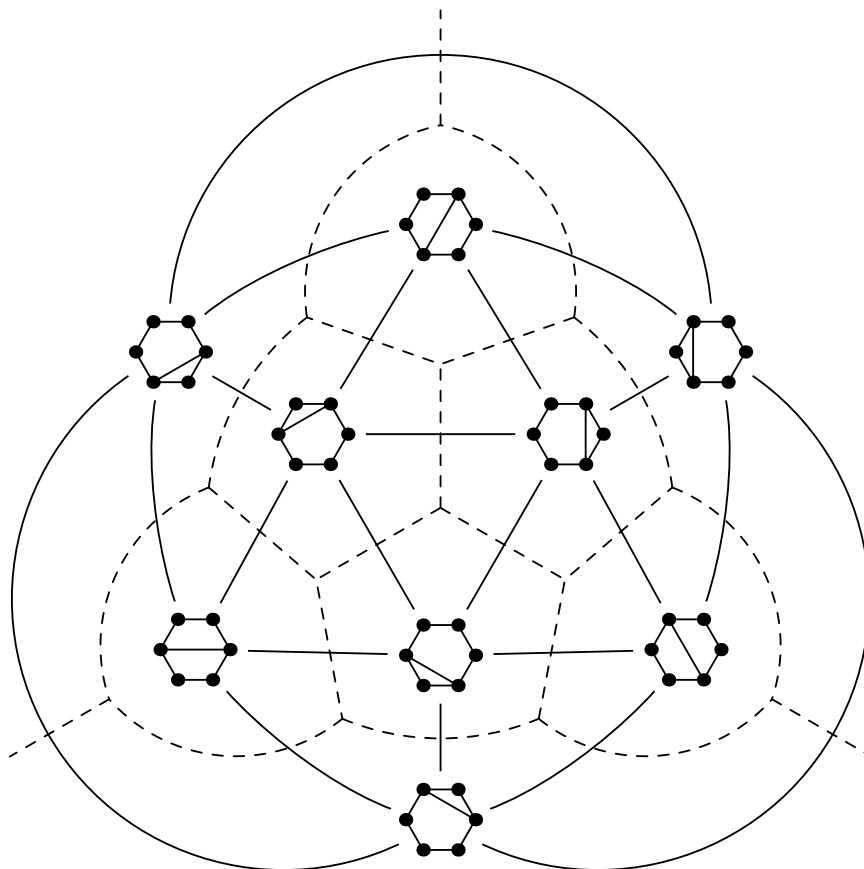


Figure 3. The arc complex of a hexagon (solid lines).

Definition 3.5. A *flip* is a transformation of an ideal triangulation T that removes an arc γ and replaces it with a (unique) different arc $\bar{\gamma}$ that, together with the remaining arcs, forms a new ideal triangulation T' .

To illustrate, adjacent triangulations in Figure 2 are related to each other by flips.

There is at most one way to flip an arc γ in an ideal triangulation T . If γ is a “folded” side of a self-folded ideal triangle (the interior edge in Figure 1), then γ cannot be flipped. Otherwise, removing γ creates a quadrilateral face on (\mathbf{S}, \mathbf{M}) , and we let $\bar{\gamma}$ be the other “diagonal” of that quadrilateral. Another way to look at flips is the following.

LEMMA 3.6. *The arc complex is a pseudomanifold with boundary, i.e., each maximal simplex is of the same dimension and each simplex of codimension 1 is contained in at most two maximal simplices.*

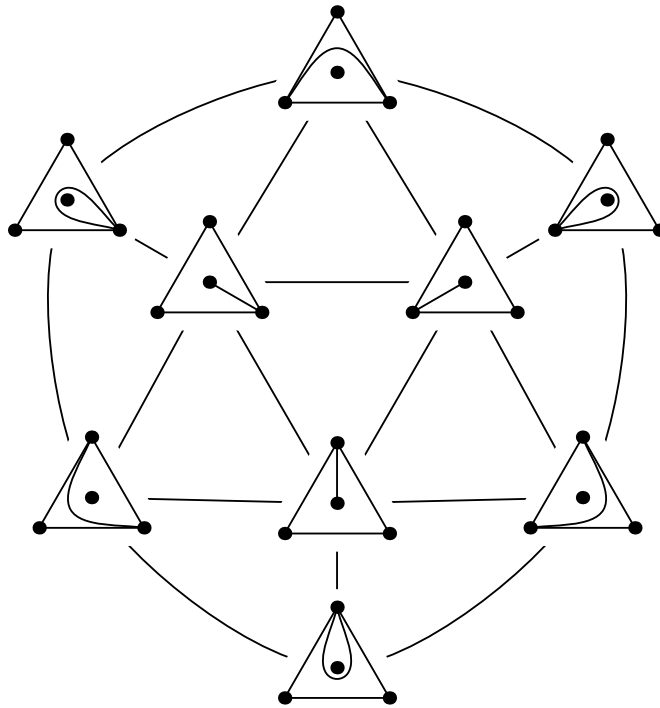


Figure 4. The arc complex of a once-punctured triangle.

The boundary of $\Delta^\circ(\mathbf{S}, \mathbf{M})$ consists of collections of arcs which contain a loop bounding a punctured monogon but do not contain the arc enclosed.

THEOREM 3.7. ([25], [27]) *The arc complex is contractible except when (\mathbf{S}, \mathbf{M}) is a polygon, i.e., a disk with no punctures.*

The *dual graph* of a pseudomanifold has maximal simplices as its vertices, with edges connecting maximal simplices sharing a codimension-1 face.

Let $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$ denote the dual graph of $\Delta^\circ(\mathbf{S}, \mathbf{M})$. Thus, the vertices of $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$ are (labeled by) the ideal triangulations of (\mathbf{S}, \mathbf{M}) , and the edges correspond to the flips. Examples are shown in Figure 3 (in dashed lines) and in Figures 6, 7 and 8.

In Figure 8, some readers may recognize the “two-layer brick wall” of [16, Figure 4]. This coincidence will be explained later; see Example 6.9.

PROPOSITION 3.8. ([26], [27], [32]) *The graph $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$ is connected.*

In other words, any two ideal triangulations are related by a sequence of flips. Combining this statement with Lemma 2.13, we obtain the following corollary, which can also be proved directly.

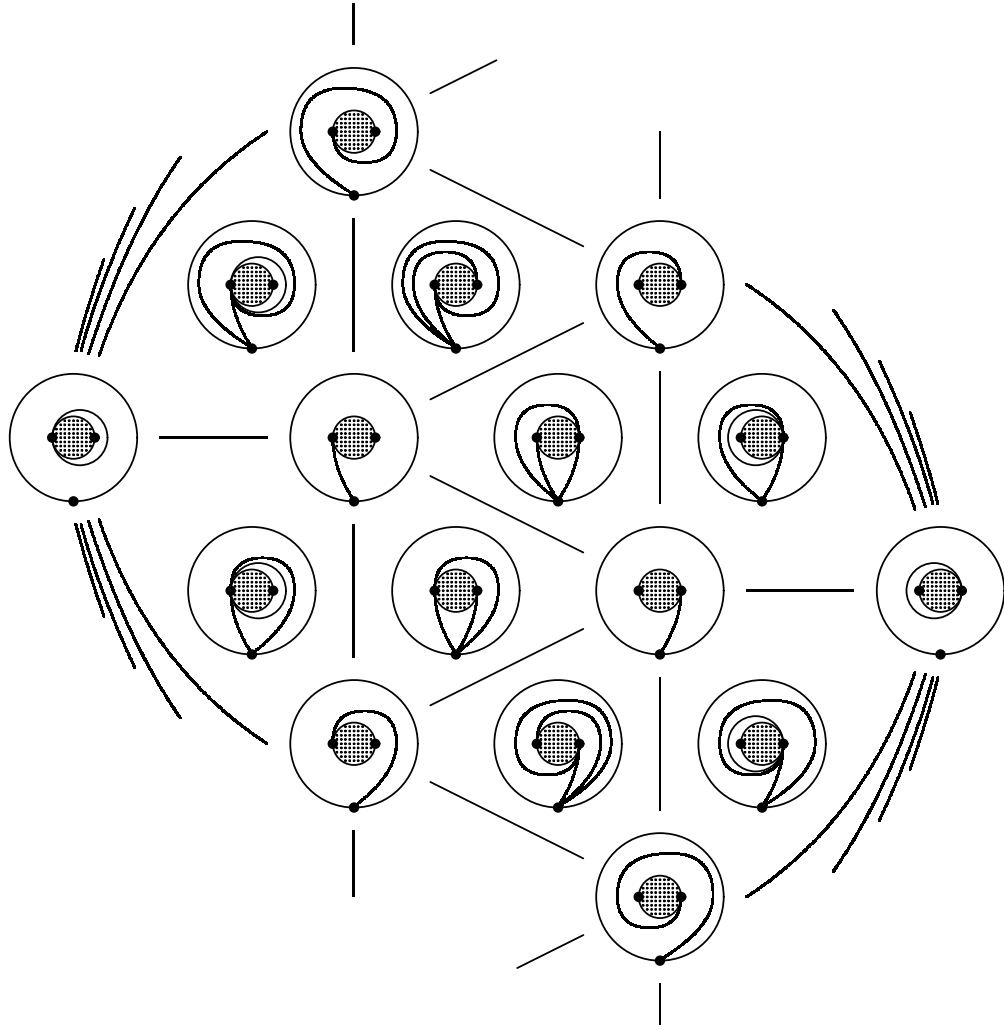


Figure 5. The arc complex $\Delta^\circ(\mathbf{S}, \mathbf{M})$ for an annulus of type $\tilde{A}(2, 1)$.

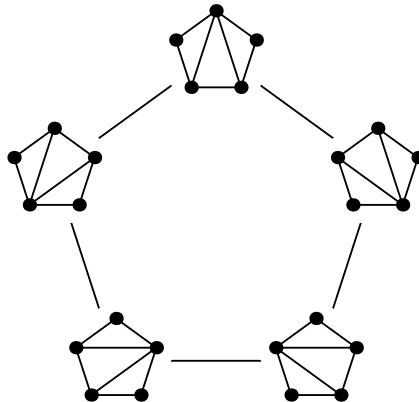


Figure 6. The graph $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$ for an unpunctured pentagon.

COROLLARY 3.9. *Any ideal triangulation can be transformed into a triangulation without self-folded triangles by a sequence of flips.*

A more difficult result is Theorem 3.10 below. It appeared explicitly in [6, Theorem 1.1]; see also [42, Appendix B] and [7, Theorem 2.7]. This result can also be extracted from [26, Theorems 1.3 and 2.1]; cf. the comments in [27, pp. 190–191].

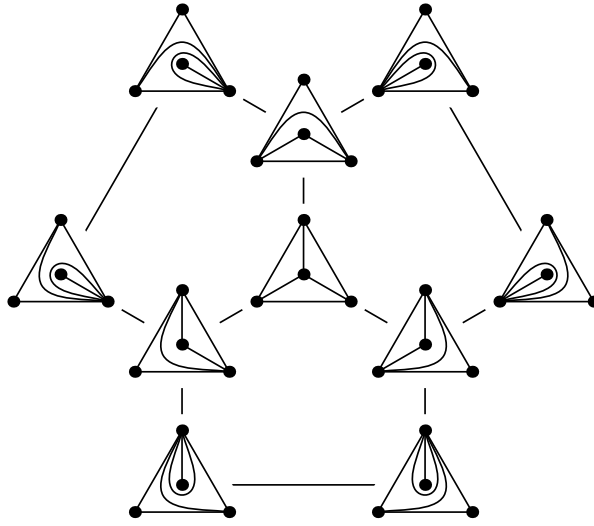
THEOREM 3.10. *The fundamental group of $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$ is generated by cycles of length 4 and 5, pinned down to a basepoint. More specifically, take the cycles of length 4 corresponding to pairs of commuting flips, and the cycles of length 5 obtained by alternately flipping two arcs whose removal would create a pentagonal face (see Figure 6).*

To clarify, by a “pentagonal face” in Theorem 3.10 we mean any face that looks like a pentagon to an observer located inside it. One example of such a face is obtained by taking a once-punctured triangle and connecting one vertex to the puncture.

The graph $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$ (or the complex $\Delta^\circ(\mathbf{S}, \mathbf{M})$) is finite, i.e., there are finitely many ideal triangulations of (\mathbf{S}, \mathbf{M}) , if and only if the total number of arcs in (\mathbf{S}, \mathbf{M}) is finite, that is, in the cases listed in Proposition 2.3.

4. Signed adjacency matrices and their mutations

In this section, we introduce and discuss the notion of a signed adjacency matrix, a skew-symmetric matrix derived from an ideal triangulation of a bordered surface. This construction is an extension of the ones proposed in [11, §2] and [24, §3], which use different vocabulary and more restrictive setup.

Figure 7. The graph $\mathbf{E}^s(\mathbf{S}, \mathbf{M})$ for a once-punctured triangle.

Definition 4.1. We associate with each ideal triangulation T the (generalized) *signed adjacency matrix* $B=B(T)$ that reflects the combinatorics of T . The rows and columns of $B(T)$ are naturally labeled by the arcs in T . For notational convenience, we arbitrarily label these arcs by the numbers $1, \dots, n$, so that the rows and columns of $B(T)$ are numbered from 1 to n as customary, with the understanding that this numbering of rows and columns is temporary rather than intrinsic. For an arc (labeled) i , let $\pi_T(i)$ denote (the label of) the arc defined as follows: if there is a self-folded ideal triangle in T folded along i (see Figure 1), then $\pi_T(i)$ is its remaining side (the enclosing loop); if there is no such triangle, set $\pi_T(i)=i$.

For each ideal triangle Δ in T which is not self-folded, define the $n \times n$ integer matrix $B^\Delta=(b_{ij}^\Delta)$ by setting

$$b_{ij}^\Delta = \begin{cases} 1, & \text{if } \Delta \text{ has sides labeled } \pi_T(i) \text{ and } \pi_T(j), \\ & \text{with } \pi_T(j) \text{ following } \pi_T(i) \text{ in the clockwise order,} \\ -1, & \text{if the same holds, with the counter-clockwise order,} \\ 0, & \text{otherwise.} \end{cases} \quad (4.1)$$

The matrix $B=B(T)=(b_{ij})$ is then defined by

$$B = \sum_{\Delta} B^\Delta,$$

where the summation runs over all ideal triangles Δ in T which are not self-folded. The $n \times n$ matrix B is skew-symmetric, and all its entries b_{ij} are equal to 0, 1, -1 , 2 or -2 .

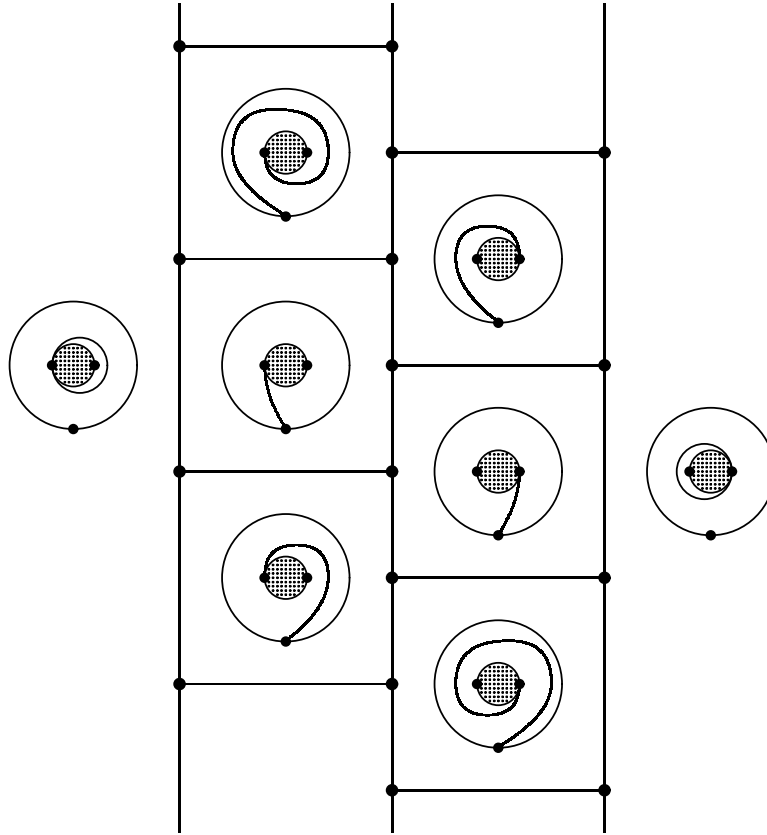


Figure 8. The graph $E^s(\mathbf{S}, \mathbf{M})$ for an annulus of type $\tilde{A}(2, 1)$. Each vertex corresponds to the triangulation made up of the three arcs labelling the adjacent regions.

For triangulations without self-folded triangles, there is no need to introduce the map π_T , and the above definition of $B(T)$ coincides with the ones given in [11] and [24].

Remark 4.2. A more concrete version of Definition 4.1 that avoids the use of the map π_T is based on representing an arbitrary ideal triangulation as a result of gluing together a number of “puzzle pieces” which “conceal” self-folded triangles inside them. Specifically, any ideal triangulation, with a single exception discussed below, can be obtained as follows:

- Take several puzzle pieces of the form shown in Figure 9 (a triangle, a triangulated once-punctured digon, or a triangulated twice-punctured monogon).
- Take a partial matching of the exposed (outer) sides of these puzzle pieces, never matching two sides of the same puzzle piece. In order to obtain a connected surface, we need to ensure that any two puzzle pieces in the collection are connected via matched pairs.

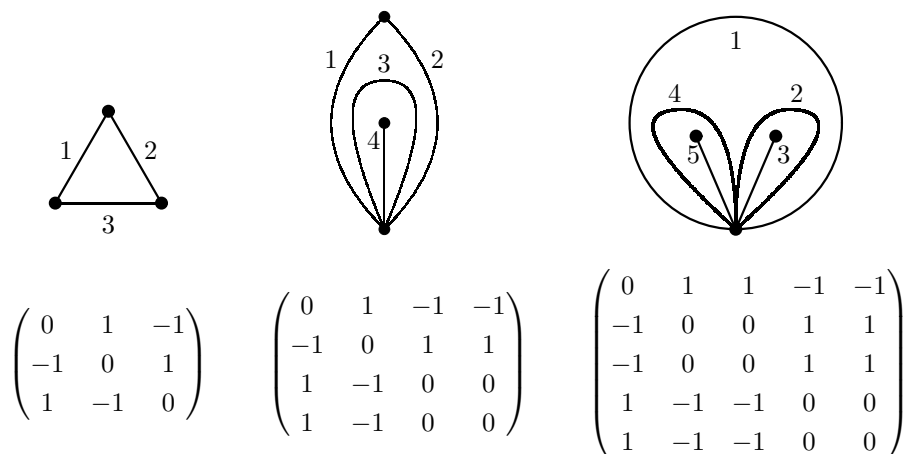


Figure 9. The three puzzle pieces and signed adjacencies within each.

- Glue the puzzle pieces along the matched sides, making sure that the orientations match.

The only exception to this construction is a triangulation of a 4-punctured sphere obtained by gluing three self-folded triangles to respective sides of an ordinary triangle. (Recall that we do not allow a once-punctured monogon or a thrice-punctured sphere, or else there would be more exceptions.)

Each entry b_{ij} of the matrix $B(T)$ is then obtained by adding up the $(0, \pm 1)$ -values describing signed adjacencies of the arcs i and j inside each puzzle piece. These are in turn determined by the matrices in Figure 9.

Example 4.3. (Twice-punctured monogon; annulus of type $\tilde{A}(2, 2)$) Figure 10 shows two (labeled) ideal triangulations of a twice-punctured monogon, and the corresponding signed adjacency matrices $B(T)$.

Figure 11 shows two ideal triangulations of an annulus with two marked points on each boundary component. The corresponding matrices $B(T)$ are exactly the same as the ones in Figure 10.

Example 4.4. (Annulus of type $\tilde{A}(1, 1)$) Let (\mathbf{S}, \mathbf{M}) be an annulus with one marked point on each boundary component and no punctures. Then, for any (ideal) triangulation T , the matrix $B(T)$ has the form

$$B(T) = \begin{pmatrix} 0 & 2 \\ -2 & 0 \end{pmatrix},$$

for an appropriate labeling of the arcs. See Figure 12.

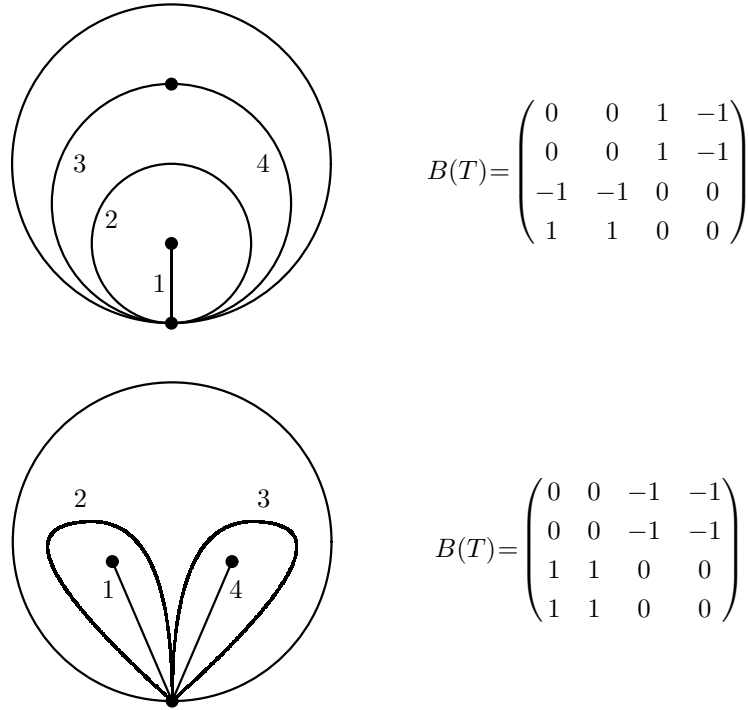


Figure 10. Ideal triangulations of a twice-punctured monogon.

Example 4.5. (Once-punctured triangle; unpunctured hexagon) Let (\mathbf{S}, \mathbf{M}) be a disk with three marked points on the boundary and one puncture. Then each matrix $B = B(T)$ associated with an ideal triangulation T of (\mathbf{S}, \mathbf{M}) has one of the following four forms:

$$\begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}, \quad \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 1 \\ 0 & -1 & 0 \end{pmatrix}, \quad \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix}, \quad \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 1 \\ 0 & -1 & 0 \end{pmatrix},$$

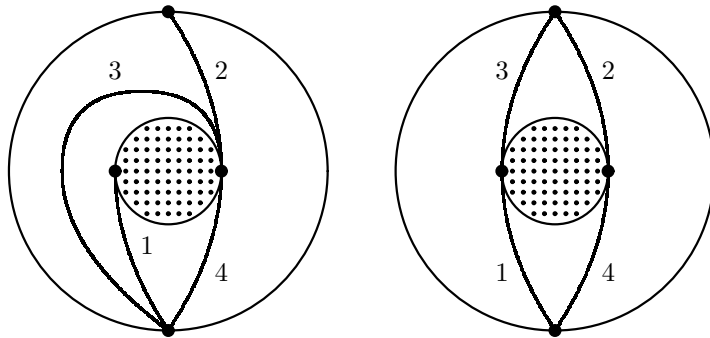
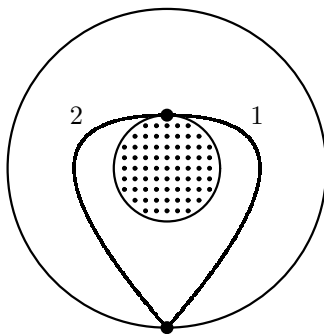
up to simultaneous relabeling of rows and columns. Cf. Figure 2.

For an unpunctured hexagon, we obtain exactly the same matrices.

Example 4.6. (Once-punctured torus) Let (\mathbf{S}, \mathbf{M}) be a torus (without boundary) with one puncture. Then for any (ideal) triangulation T , the matrix $B(T)$ has the form

$$B(T) = \begin{pmatrix} 0 & 2 & -2 \\ -2 & 0 & 2 \\ 2 & -2 & 0 \end{pmatrix},$$

for a suitable labeling of the arcs.

Figure 11. Ideal triangulations of an annulus of type $\tilde{A}(2,2)$.Figure 12. Ideal triangulation of an annulus of type $\tilde{A}(1,1)$.

In §13, we discuss the “inverse problem” of recovering information about the triangulated surface from the matrix $B(T)$, and the related problem of direct combinatorial description of the class of matrices $B(T)$ that arise through this general construction.

Definition 4.7. ([16, Definition 4.2]) Let $B=(b_{ij})$ be an $n \times n$ integer matrix. We say that an $n \times n$ matrix $\bar{B}=(\bar{b}_{ij})$ is obtained from B by *matrix mutation* in direction k , and write $\bar{B}=\mu_k(B)$, if the entries of \bar{B} are given by

$$\bar{b}_{ij} = \begin{cases} -b_{ij}, & \text{if } i = k \text{ or } j = k, \\ b_{ij} + \frac{1}{2}(|b_{ik}|b_{kj} + b_{ik}|b_{kj}|), & \text{otherwise.} \end{cases} \quad (4.2)$$

Matrix mutations are involutive, i.e., $\mu_k(\mu_k(B))=B$.

If B is a skew-symmetric matrix, then so is $\mu_k(B)$.

Proposition 4.8 below shows that flips correspond to mutations of the matrix $B(T)$. This proposition can be verified by direct inspection. In restricted generality, it appeared in [10], [12] and [24].

PROPOSITION 4.8. *Suppose that an ideal triangulation \bar{T} is obtained from T by a flip replacing an arc labeled k . (The labeling of all other arcs remains unchanged.) Then $B(\bar{T}) = \mu_k(B(T))$.*

Proof. Let us use the puzzle-piece decomposition of \bar{T} described in Remark 4.2. Each flip occurs either inside a puzzle piece, or involves an arc shared by two puzzle pieces. In either case, the only entries of the signed adjacency matrix affected by the flip are the ones located in the rows and columns associated with the arcs in the puzzle piece(s) involved. The proof therefore reduces to a routine case-by-case verification. \square

Definition 4.9. Two matrices are called *mutation-equivalent* if they can be transformed into each other by a sequence of mutations. Let $\mathcal{B}(\mathbf{S}, \mathbf{M})$ denote the mutation equivalence class of a matrix $B(T)$ associated with a triangulation T of (\mathbf{S}, \mathbf{M}) .

PROPOSITION 4.10. *The mutation equivalence class $\mathcal{B}(\mathbf{S}, \mathbf{M})$ depends only on (\mathbf{S}, \mathbf{M}) but not on the choice of an ideal triangulation t .*

Proof. This follows from Propositions 3.8 and 4.8. \square

5. Cluster algebras associated with triangulated surfaces

We begin this section by a swift review of cluster algebras, closely following [17]. This is not meant as an introduction to the subject, but rather to fix notation and specify the generality we will be working in. For entry points to the theory, see, e.g., [18] or [15].

Let $(\mathbf{P}, \oplus, \cdot)$ be a *semifield*, i.e., (\mathbf{P}, \cdot) is an abelian group, (\mathbf{P}, \oplus) is a commutative semigroup, and the *auxiliary addition* \oplus is distributive with respect to the multiplication.

Let \mathcal{F} be a field isomorphic to the field of rational functions in n independent variables, with coefficients in $\mathbb{Z}\mathbf{P}$, the group ring of the multiplicative group of \mathbf{P} .

Definition 5.1. (Seeds and their mutations) A *seed* in \mathcal{F} is a triple $\Sigma = (\mathbf{x}, \mathbf{p}, B)$, where

- the *cluster* \mathbf{x} is a set of n algebraically independent elements in \mathcal{F} (called *cluster variables*) that generate \mathcal{F} over the field of fractions of $\mathbb{Z}\mathbf{P}$;
- the *coefficient tuple* $\mathbf{p} = (p_x^\pm)_{x \in \mathbf{x}}$ is a $2n$ -tuple of elements of \mathbf{P} satisfying the *normalization condition* $p_x^+ \oplus p_x^- = 1$ for all $x \in \mathbf{x}$;
- the *exchange matrix* $B = (b_{xy})_{x, y \in \mathbf{x}}$ is a skew-symmetrizable⁽²⁾ $n \times n$ integer matrix with rows and columns indexed by \mathbf{x} .

⁽²⁾ A square matrix is *skew-symmetrizable* if it becomes skew-symmetric upon rescaling of its columns by appropriate positive factors.

For such a seed $\Sigma=(\mathbf{x}, \mathbf{p}, B)$, and for any cluster variable $z \in \mathbf{x}$, the *seed mutation* μ_z transforms Σ into a new seed $\bar{\Sigma}=(\bar{\mathbf{x}}, \bar{\mathbf{p}}, \bar{B})$ defined as follows:

- $\bar{\mathbf{x}}=(\mathbf{x} \setminus \{z\}) \cup \{\bar{z}\}$, where $\bar{z} \in \mathcal{F}$ is determined by the *exchange relation*

$$z\bar{z} = p_z^+ \prod_{\substack{x \in \mathbf{x} \\ b_{xz} > 0}} x^{b_{xz}} + p_z^- \prod_{\substack{x \in \mathbf{x} \\ b_{xz} < 0}} x^{-b_{xz}}; \quad (5.1)$$

- the $2n$ -tuple $\bar{\mathbf{p}}=(\bar{p}_x^\pm)_{x \in \bar{\mathbf{x}}}$ is uniquely determined by the normalization conditions $\bar{p}_x^+ \oplus \bar{p}_x^- = 1$ together with

$$\bar{p}_x^+ / \bar{p}_x^- = \begin{cases} p_z^- / p_z^+, & \text{if } x = \bar{z}, \\ (p_z^+)^{b_{zx}} p_x^+ / p_x^-, & \text{if } b_{zx} \geq 0, \\ (p_z^-)^{b_{zx}} p_x^+ / p_x^-, & \text{if } b_{zx} \leq 0; \end{cases} \quad (5.2)$$

- the matrix \bar{B} is obtained from B by applying the matrix mutation in direction z and then relabelling one row and one column by replacing z by \bar{z} .

It is easy to check that $\mu_z(\bar{\Sigma}) = \Sigma$.

Definition 5.2. (Normalized skew-symmetrizable cluster algebra) Let \mathcal{S} be a mutation equivalence class of seeds in \mathcal{F} . Let \mathcal{X} denote the set of all cluster variables appearing in the seeds of \mathcal{S} . Let \mathcal{R} be any subring with unit in $\mathbb{Z}\mathbf{P}$ containing all elements $p_x^\pm \in \mathbf{p}$, for all seeds $\Sigma=(\mathbf{x}, \mathbf{p}, B) \in \mathcal{S}$. The (normalized) *cluster algebra* $\mathcal{A}=\mathcal{A}(\mathcal{S})$ is the \mathcal{R} -subalgebra of \mathcal{F} generated by \mathcal{X} .

The exchange matrices B appearing in the seeds of \mathcal{S} form a mutation equivalence class, denoted by $\mathcal{B}(\mathcal{A})$.

Conjecture 5.3. ([17, p. 70], [18, Conjecture 4.14 (2)]) Each seed in a cluster algebra (i.e., each seed in a given mutation equivalence class) is uniquely determined by its cluster. Two such seeds are related by a mutation if and only if the corresponding clusters share all elements but one.

Definition 5.4. (Exchange graph, cluster complex) Let $\mathcal{A}=\mathcal{A}(\mathcal{S})$ be the cluster algebra defined by \mathcal{S} , a mutation equivalence class of seeds. The *exchange graph* of $\mathcal{A}(\mathcal{S})$ is the n -regular graph whose vertices are labeled by the seeds in \mathcal{S} , and whose edges correspond to mutations.

Assuming that Conjecture 5.3 holds for a cluster algebra \mathcal{A} , its underlying combinatorics is governed by its *cluster complex*. This is the (possibly infinite) simplicial complex on the ground set \mathcal{X} whose maximal simplices are the clusters. The exchange graph is then the dual graph of the cluster complex: the vertices of the exchange graph can be identified with clusters, with edges corresponding to pairs of clusters whose intersection has cardinality $n-1$.

Two cluster variables are *compatible* if they appear together in some cluster.

Conjecture 5.5. The cluster complex is always a “flag complex.” That is, the cluster complex is the clique complex for the compatibility relation on the set of all cluster variables.

In other words, for any collection of pairwise compatible cluster variables, there is a cluster containing all of them.

We note that Conjecture 5.5 would follow from [20, Conjecture 7.4] and [18, Conjecture 4.14 (3)].

Our main object of study is the class of cluster algebras whose exchange matrices can be described by triangulated surfaces. More precisely, a cluster algebra \mathcal{A} belongs to this class if and only if $\mathcal{B}(\mathcal{A}) = \mathcal{B}(\mathbf{S}, \mathbf{M})$, for some bordered surface (\mathbf{S}, \mathbf{M}) with marked points. (See Definition 4.9 and Proposition 4.10.) As we will see in §6, this class in particular includes all cluster algebras of finite and affine types A and D .

The following result in particular establishes [18, Conjecture 4.14, parts (1)–(3)] and Conjecture 5.5 for the class of cluster algebras under consideration.

THEOREM 5.6. *Let (\mathbf{S}, \mathbf{M}) be a bordered surface with marked points. Assume that (\mathbf{S}, \mathbf{M}) is not a closed surface with exactly two punctures. Let \mathcal{A} be a cluster algebra with $\mathcal{B}(\mathcal{A}) = \mathcal{B}(\mathbf{S}, \mathbf{M})$. Then the following hold:*

- *Each seed in \mathcal{A} is uniquely determined by its cluster.*
- *The cluster complex of \mathcal{A} is uniquely determined by (\mathbf{S}, \mathbf{M}) (i.e., it is independent of the choice of coefficients p_z^\pm).*
- *The seeds containing a particular cluster variable form a connected subgraph of the exchange graph.*
- *The cluster complex of \mathcal{A} is the clique complex for the compatibility relation on the set of cluster variables.*

The proof of Theorem 5.6 is given in §9. The cluster complex of any cluster algebra with $\mathcal{B}(\mathcal{A}) = \mathcal{B}(\mathbf{S}, \mathbf{M})$ is denoted by $\Delta(\mathbf{S}, \mathbf{M})$. A concrete description of $\Delta(\mathbf{S}, \mathbf{M})$ in terms of the combinatorial topology of (\mathbf{S}, \mathbf{M}) is provided in Theorem 7.11.

We expect Theorem 5.6 to hold in full generality, including the case of closed surfaces with two punctures. The proof presented in this paper does not however work in this special case. A different proof covering all cases should appear in [14].

6. Types and diagrams

We begin this section by reviewing the encoding of exchange matrices by oriented graphs (quivers), the finite-type classification theorem [17], and the related nomenclature of

cluster algebras of finite and affine type. We then discuss four examples (finite and affine types A and D) that naturally fit into this paper's framework, i.e., they can be associated with bordered surfaces with marked points. (These examples correspond to the first four line entries in Table 1.) In each case, we briefly mention appearances of cluster algebras of the corresponding type as (homogeneous) coordinate rings of various classical algebraic varieties such as Grassmannians, double Bruhat cells, etc.; cf. [1], [17] and [40].

Definition 6.1. (Type of a cluster algebra) We say that two cluster algebras \mathcal{A} and \mathcal{A}' are of the same *type* if the corresponding mutation equivalence classes $\mathcal{B}(\mathcal{A})$ and $\mathcal{B}(\mathcal{A}')$ coincide, up to simultaneous relabeling of rows and columns.

A cluster algebra is of *finite type* if it has finitely many seeds. This terminology is consistent with Definition 6.1 due to the following result.

THEOREM 6.2. ([17]) *Whether a cluster algebra \mathcal{A} is of finite or infinite type depends only on the mutation equivalence class $\mathcal{B}(\mathcal{A})$.*

All cluster algebras of finite type have been classified [17]; this classification is completely parallel to the Cartan–Killing classification of, say, finite crystallographic root systems. As shown in [17], Conjecture 5.3 holds for any cluster algebra of finite type; its cluster complex is the *generalized associahedron* [19] of the corresponding type.

In this paper, we restrict our attention to cluster algebras whose exchange matrices are *skew-symmetric*. To state a more concrete version of Theorem 6.2 in this restricted generality, we will need a construction that associates a mutation equivalence class to an oriented graph (or *quiver*).

Definition 6.3. (Encoding matrices by graphs) Let Γ be a finite directed graph with no loops and no cycles of length 2; multiple edges are allowed. Let $B(\Gamma)=(b_{ij})$ denote the skew-symmetric matrix whose rows and columns are labeled by the vertices of Γ , and whose entry b_{ij} is equal to the number of edges going from i to j minus the number of edges going from j to i .

LEMMA 6.4. ([17]) *Let \mathbf{T} be an unoriented graph obtained from a finite forest (a disjoint union of trees) by replacing some of its edges by multiple ones. If Γ_1 and Γ_2 are two orientations of \mathbf{T} , then $B(\Gamma_1)$ and $B(\Gamma_2)$ are mutation-equivalent.*

It follows that there is a well-defined mutation equivalence class $\mathcal{B}(\mathbf{T})$ (associated with any such graph \mathbf{T} (i.e., with any finite forest with multiple edges)). This in particular applies in a situation where \mathbf{T} is a simply-laced Coxeter–Dynkin diagram of type A_n ($n \geq 1$), D_n ($n \geq 4$), or E_n ($n = 6, 7, 8$), or a disjoint union of such diagrams. See Figure 13.

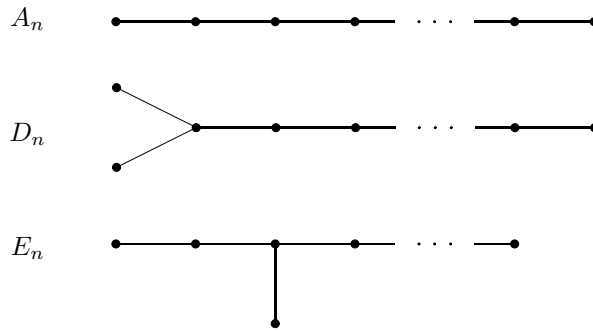


Figure 13. Simply-laced Coxeter–Dynkin diagrams.

THEOREM 6.5. ([17]) *A cluster algebra \mathcal{A} of rank n with skew-symmetric exchange matrices is of finite type if and only if the mutation equivalence class $\mathcal{B}(\mathcal{A})$ is of the form $\mathcal{B}(\mathbf{T})$, where \mathbf{T} is a disjoint union of simply-laced Coxeter–Dynkin diagrams.*

In the situation described in Theorem 6.5, the type of \mathbf{T} (in the Cartan–Killing nomenclature) is uniquely determined by the cluster algebra \mathcal{A} , and is called the (*cluster*) *type* of \mathcal{A} (or of the mutation equivalence class $\mathcal{B}(\mathcal{A})$, or of any matrix in that class). This terminology has already been used in §2.

Example 6.6. (Type A_n : A polygon) Let (\mathbf{S}, \mathbf{M}) be an unpunctured $(n+3)$ -gon, $n \geq 1$. Then $\mathcal{B}(\mathbf{S}, \mathbf{M})$ is a mutation equivalence class of type A_n . This can be verified by looking at a triangulation of the form shown in Figure 14 on the left.

As shown in [17, §12.2], a cluster algebra of this type can be constructed by taking (the \mathbb{Z} -form of) the homogeneous coordinate ring of the Grassmannian $\text{Gr}_{2,n+3}$ of 2-dimensional subspaces of an $(n+3)$ -dimensional complex vector space, with respect to its Plücker embedding. Another example is the coordinate ring of the affine base space $\text{SL}_4(\mathbb{C})/N$ (type A_3). (Here and below, N denotes the subgroup of unipotent upper-triangular matrices.)

Example 6.7. (Type D_n : A once-punctured polygon) Let (\mathbf{S}, \mathbf{M}) be a once-punctured n -gon, $n \geq 4$. Then $\mathcal{B}(\mathbf{S}, \mathbf{M})$ is a mutation equivalence class of type D_n . This can be verified by looking at a triangulation of the form shown in Figure 14 on the right.

Examples of cluster algebras of type D include the (homogeneous) coordinate rings of the Grassmannian $\text{Gr}_{3,6}$ (type D_4), the special linear group $\text{SL}_3(\mathbb{C})$ viewed as a surface in \mathbb{C}^9 (type D_4), the affine base space $\text{SL}_5(\mathbb{C})/N$ (type D_6) and (the affine cone over) the Schubert divisor in the Grassmannian $\text{Gr}_{2,n+2}$ (type D_n) [17, Example 12.15].

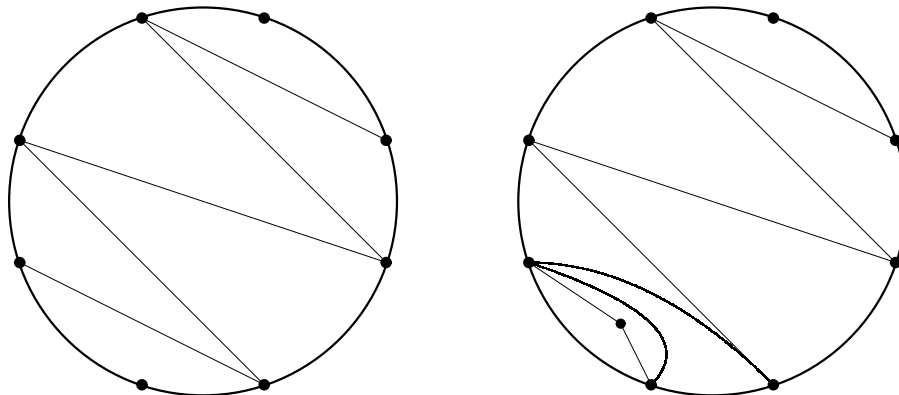


Figure 14. Triangulations for types A and D .

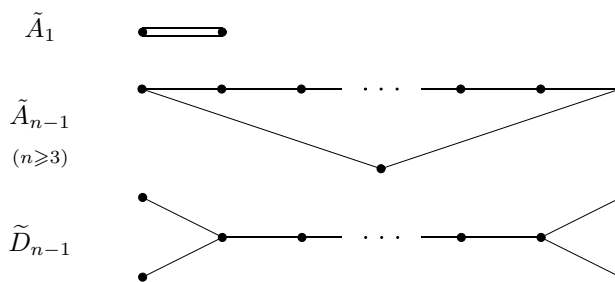
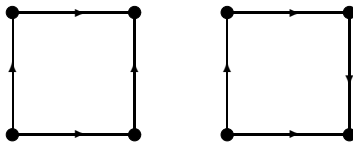


Figure 15. Some affine Coxeter–Dynkin diagrams.

By analogy with finite type, one can use (orientations of) affine Coxeter–Dynkin diagrams to define mutation equivalence classes of “affine types.” Some examples are shown in Figure 15. In the case of affine type \tilde{A}_{n-1} , $n \geq 3$, the underlying graph (an n -cycle) is not a forest, so Lemma 6.4 does not apply. In fact, the mutation equivalence class in this case does depend on the orientation of the edges in the diagram; see Lemma 6.8 below.

LEMMA 6.8. *Let Γ and Γ' be two n -cycles, $n \geq 3$, whose edges have been oriented, so that in Γ (resp. Γ'), there are n_1 (resp. n'_1) edges of one direction and $n_2 = n - n_1$ (resp. $n'_2 = n - n'_1$) edges of the opposite direction. Then the matrices $B(\Gamma)$ and $B(\Gamma')$ are mutation-equivalent (up to simultaneous relabeling of rows and columns) if and only if the unordered pairs $\{n_1, n_2\}$ and $\{n'_1, n'_2\}$ coincide. If $n_1 = 0$ or $n_1 = n$, then $B(\Gamma)$ has type D_n .*

The “only if” part of the lemma is proved in §14.

Figure 16. Diagrams of types $\tilde{A}(2, 2)$ and $\tilde{A}(3, 1)$.

Proof (“if” direction). Suppose that $\{n_1, n_2\} = \{n'_1, n'_2\}$. It is then easy to check that Γ can be transformed into Γ' by *shape-preserving mutations* (see [17, §9.2]), i.e., mutations performed at sinks or sources of the diagram. \square

In the situation described in Lemma 6.8, with $n_1 \geq n_2 > 0$, we say that the mutation equivalence class of $B(\Gamma)$ (or a cluster algebra with this exchange matrix) is of type $\tilde{A}(n_1, n_2)$. See Figure 16.

Example 6.9. (Type $\tilde{A}(n_1, n_2)$: An annulus) Let (\mathbf{S}, \mathbf{M}) be an unpunctured annulus with n_1 marked points on one boundary component and n_2 on the other. Then $\mathcal{B}(\mathbf{S}, \mathbf{M})$ is a mutation equivalence class of type $\tilde{A}(n_1, n_2)$. This can be verified by looking at a triangulation of the form shown in Figure 17 on the left.

In the special case of type $\tilde{A}(2, 1)$, we recover the example illustrated in Figures 5 and 8.

Type $\tilde{A}(2, 2)$ can also be obtained by taking a twice-punctured monogon; see Example 4.3.

Coordinate rings of certain double Bruhat cells and Schubert varieties/cells in (affine cones over) Grassmannians have natural cluster algebra structures of type $\tilde{A}(n_1, n_2)$. For instance, cluster type $\tilde{A}(6, 3)$ can be realized by coordinate rings of hypersurfaces in $\text{Gr}_{4,8}$ and/or SL_4 (embedding SL_4 as the hypersurface $\det=1$ in \mathbb{C}^{16}).

Example 6.10. (Type \tilde{D}_{n-1} : A twice-punctured polygon) Let (\mathbf{S}, \mathbf{M}) be a twice-punctured $(n-3)$ -gon, $n \geq 5$. Then $\mathcal{B}(\mathbf{S}, \mathbf{M})$ is a mutation equivalence class of type \tilde{D}_{n-1} . This can be verified by looking at a triangulation of the form shown in Figure 17, on the right.

Examples of cluster algebras of this type include coordinate rings of double Bruhat cells, Schubert varieties, etc. For instance, cluster type \tilde{D}_8 can be realized by the coordinate ring of a hypersurface $P=1$ in $\text{Gr}_{3,9}$, where P is a Plücker coordinate. On the combinatorial level, the natural cluster structure on $\text{Gr}_{3,9}$ can be represented by a 2×5 grid graph on 10 vertices, with all 4-cycles properly oriented (see [40]); removing (or “freezing”) a vertex in the middle yields the aforementioned hypersurface.

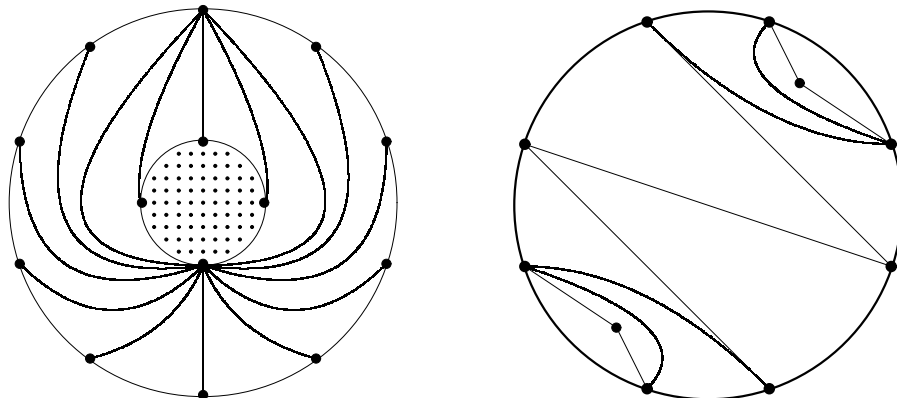
Figure 17. Triangulations for types $\tilde{A}(10, 4)$ and \tilde{D}_{10} .

Figure 18. Representing an arc bounding a punctured monogon by a tagged arc.

7. Tagged arc complexes

In view of Proposition 4.10, associated with each marked surface (\mathbf{S}, \mathbf{M}) there is a class of cluster algebras \mathcal{A} whose exchange matrices include the sign adjacency matrices of triangulations of (\mathbf{S}, \mathbf{M}) . Proposition 4.8 suggests (and we will later prove) that the arc complex $\Delta^\circ(\mathbf{S}, \mathbf{M})$ is a subcomplex of the cluster complex of \mathcal{A} , and its dual graph $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$ is a subgraph of the exchange graph of \mathcal{A} . (In general, one does not get the full picture because some arcs in some triangulations are not flippable.) We will now present a combinatorial construction for the entire cluster complex and exchange graph.

Definition 7.1. (Tagged arcs) Each arc γ in (\mathbf{S}, \mathbf{M}) has two *ends* obtained by arbitrarily cutting γ into three pieces, and then throwing out the middle one. We think of the two ends as locations near the endpoints to be used for labeling (“tagging”) an arc. A *tagged arc* is an arc in which each end has been tagged in one of two ways, *plain* or *notched*, so that the following conditions are satisfied:

- the arc does not cut out a once-punctured monogon;
- an endpoint lying on the boundary is tagged plain;
- both ends of a loop are tagged in the same way.

In the figures, the plain tags are omitted while the notched tags are represented by

(\mathbf{S}, \mathbf{M})	Boundary to boundary (not a loop)	Loop at boundary	Loop at puncture	Boundary to puncture	Puncture to puncture (not a loop)
$(n+3)$ -gon	$n(n+3)/2$	0	0	0	0
n -gon, 1 puncture	n^2-2n	0	0	$2n$	0
annulus $\tilde{A}(n_1, n_2); n_1, n_2 \geq 2$	∞	n_1+n_2	0	0	0
annulus $\tilde{A}(n_1, 1); n_1 \geq 2$	∞	n_1	0	0	0
annulus $\tilde{A}(1, 1)$	∞	0	0	0	0
$(n-3)$ -gon, 2 punctures; $n \geq 5$	∞	$n-3$	0	∞	4
monogon, 2 punctures	0	0	0	∞	4
torus, 1 puncture	0	0	∞	0	0

Table 2. Tagged arcs of different type.

the \bowtie symbol.

Definition 7.2. Each ordinary (“plain”) arc γ can be represented by a tagged arc $\tau(\gamma)$, as follows. If γ does not cut out a once-punctured monogon, then $\tau(\gamma)$ is γ with both ends tagged plain. Otherwise, let γ be a loop, based at a marked point a , cutting out a punctured monogon with the sole puncture b inside it; see Figure 18. Note that there is only one possible monogon since (\mathbf{S}, \mathbf{M}) is not the 3-punctured sphere (by assumption) so γ cannot bound a punctured monogon on both sides. Let β be the unique arc connecting a and b which is compatible with γ . Then $\tau(\gamma)$ is obtained by tagging β plain at a and notched at b .

The set of all tagged arcs in (\mathbf{S}, \mathbf{M}) is denoted by $\mathbf{A}^{\bowtie}(\mathbf{S}, \mathbf{M})$. As we just explained, there is a canonical map τ from the set of plain arcs $\mathbf{A}^\circ(\mathbf{S}, \mathbf{M})$ into $\mathbf{A}^{\bowtie}(\mathbf{S}, \mathbf{M})$.

Remark 7.3. Each tagged arc $\gamma \in \mathbf{A}^{\bowtie}(\mathbf{S}, \mathbf{M})$ belongs to one of the following classes:

- γ is a plain arc that connects different marked points on the boundary;
- γ is a plain arc that connects a marked point on the boundary to itself and does not enclose, on either side, a once-punctured monogon;
- γ connects a puncture to itself and does not enclose, on either side, a once-punctured monogon; γ may be either plain or with both ends notched;
- γ connects a marked point on the boundary to a puncture (two tagged arcs per untagged isotopy class);
- γ connects two punctures at different locations (four tagged arcs per untagged isotopy class).

Table 2 shows, for each of several examples of bordered surfaces with marked points, the cardinalities of these five classes of tagged arcs.

We next adapt the concept of compatibility of arcs (see Definition 2.4) to the tagged setup.

Definition 7.4. (Compatibility of tagged arcs) Two tagged arcs $\alpha, \beta \in \mathbf{A}^{\bowtie}(\mathbf{S}, \mathbf{M})$ are called *compatible* if and only if the following conditions are satisfied:

- the untagged versions of α and β are compatible;
- if the untagged versions of α and β are different, and α and β share an endpoint a , then the ends of α and β connected to a must be tagged in the same way;
- if the untagged versions of α and β coincide, then at least one end of α must be tagged in the same way as the corresponding end of β .

Remark 7.5. The following is a complete list of possible compatible pairs $\{\alpha, \beta\}$ of tagged arcs:

- α and β have endpoints at four different locations, and do not intersect;
- α and β connect from different locations to the same boundary point, and do not intersect in the interior of \mathbf{S} ;
- α and β connect from different locations to the same puncture, are tagged in the same way at it, and do not intersect in the interior of \mathbf{S} ;
- α and β are not loops, have the same endpoints, are tagged in the same way at their respective ends, and do not intersect in the interior of \mathbf{S} ;
- α and β are two different loops based at the same marked point, are tagged in the same way at all four ends, and do not intersect in the interior of \mathbf{S} ;
- α and β are not loops, their untagged versions coincide, they are tagged in the same way at one end and differently at the other;
- α and β coincide.

Remark 7.6. If two plain arcs α and β are compatible, then the tagged arcs $\tau(\alpha)$ and $\tau(\beta)$ representing them (in the sense of Definition 7.2) are compatible. The converse is *false*: in a once-punctured digon, the two loops are not compatible even though the tagged arcs representing them are compatible. On the other hand, the two notions coincide for arcs that do not cut out a once-punctured monogon.

Definition 7.7. (Tagged arc complex) As before, let (\mathbf{S}, \mathbf{M}) be a bordered surface with marked points. The *tagged arc complex* $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$ is the clique complex for the compatibility relation on the set of tagged arcs $\mathbf{A}^{\bowtie}(\mathbf{S}, \mathbf{M})$. That is, $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$ is the simplicial complex on the ground set $\mathbf{A}^{\bowtie}(\mathbf{S}, \mathbf{M})$ whose simplices are collections of mutually compatible tagged arcs. In the absence of punctures, the complex $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$ obviously coincides with the ordinary arc complex $\Delta^{\circ}(\mathbf{S}, \mathbf{M})$ of Definition 3.1.

Example 7.8. Figure 19 shows the tagged arc complex of a punctured triangle. In view of Remark 7.6, the untagged arc complex $\Delta^{\circ}(\mathbf{S}, \mathbf{M})$ can be viewed as a subcomplex of the tagged arc complex $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$, but $\Delta^{\circ}(\mathbf{S}, \mathbf{M})$ is *not* an induced subcomplex of $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$, as one can see by comparing Figures 4 and 19.

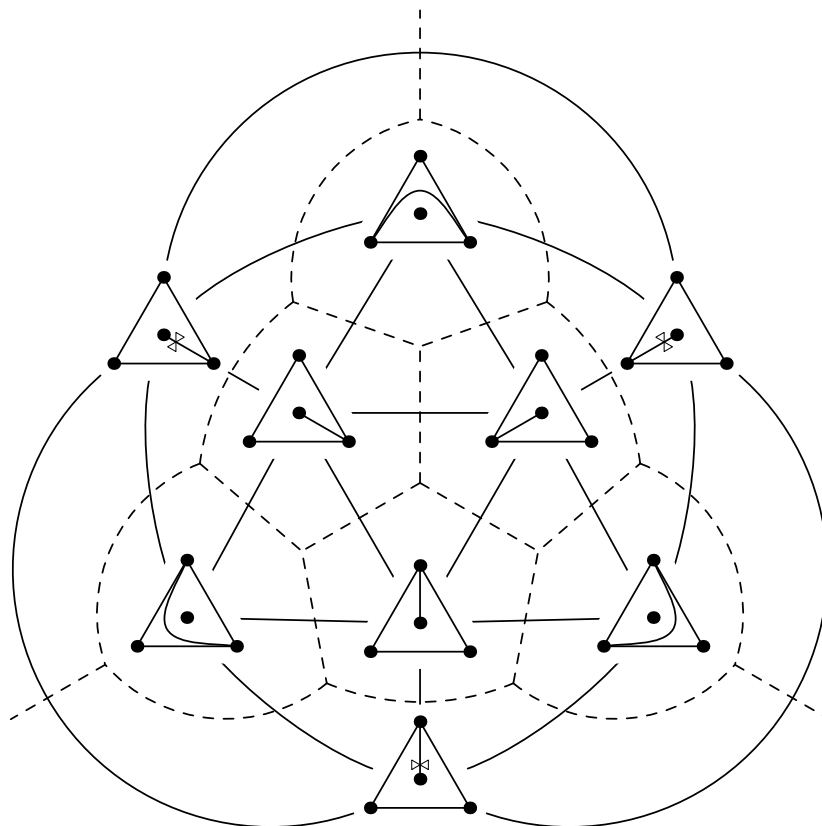


Figure 19. The tagged arc complex of a once-punctured triangle.

A maximal (by inclusion) collection of pairwise compatible tagged arcs in (\mathbf{S}, \mathbf{M}) is called a *tagged triangulation*.

THEOREM 7.9. *The tagged arc complex $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$ is a pseudomanifold. That is, each maximal simplex in $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$ (a tagged triangulation) has the same cardinality (equal to the rank of (\mathbf{S}, \mathbf{M})), and each simplex of codimension 1 is contained in precisely two maximal simplices.*

Theorem 7.9 is proved in §9.1.

We denote by $\mathbf{E}^{\bowtie}(\mathbf{S}, \mathbf{M})$ the dual graph of $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$. We will soon see that $\mathbf{E}^{\bowtie}(\mathbf{S}, \mathbf{M})$ almost always coincides with the exchange graph $\mathbf{E}(\mathbf{S}, \mathbf{M})$ of a cluster algebra \mathcal{A} with $\mathcal{B}(\mathcal{A}) = \mathcal{B}(\mathbf{S}, \mathbf{M})$ (see Definition 5.4). Again, if there are no punctures, then $\mathbf{E}^{\bowtie}(\mathbf{S}, \mathbf{M})$ is just the dual graph $\mathbf{E}^{\circ}(\mathbf{S}, \mathbf{M})$ of the ordinary arc complex. By analogy with the ordinary case, the edges of $\mathbf{E}^{\bowtie}(\mathbf{S}, \mathbf{M})$ are called *flips*. Thus, each flip replaces a tagged arc in a tagged triangulation by a (uniquely defined, different) new tagged arc.

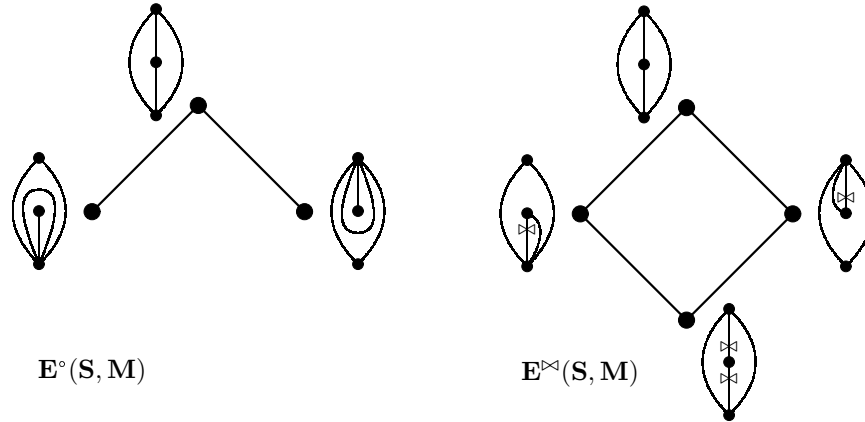


Figure 20. The graphs $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$ and $\mathbf{E}^\diamond(\mathbf{S}, \mathbf{M})$ for a once-punctured digon.

See Figure 19 (dashed lines) and Figure 20.

PROPOSITION 7.10. *If (\mathbf{S}, \mathbf{M}) is not a closed surface with exactly one puncture then the graph $\mathbf{E}^\diamond(\mathbf{S}, \mathbf{M})$ is connected, i.e., any two tagged triangulations are connected by a sequence of flips. Consequently, $\Delta^\diamond(\mathbf{S}, \mathbf{M})$ is connected.*

If (\mathbf{S}, \mathbf{M}) is a closed surface with one puncture, then $\mathbf{E}^\diamond(\mathbf{S}, \mathbf{M})$ and $\Delta^\diamond(\mathbf{S}, \mathbf{M})$ each have two isomorphic components: one in which all ends of all arcs are plain and one in which they are all notched.

The following is our main result, proved in §9.

THEOREM 7.11. *Assume that (\mathbf{S}, \mathbf{M}) is not a closed surface with one or two punctures. Let \mathcal{A} be a cluster algebra with $\mathcal{B}(\mathcal{A}) = \mathcal{B}(\mathbf{S}, \mathbf{M})$. Then the cluster complex $\Delta(\mathbf{S}, \mathbf{M})$ of \mathcal{A} is isomorphic to the tagged arc complex $\Delta^\diamond(\mathbf{S}, \mathbf{M})$, and the exchange graph of \mathcal{A} is isomorphic to $\mathbf{E}^\diamond(\mathbf{S}, \mathbf{M})$.*

If (\mathbf{S}, \mathbf{M}) is a closed surface with exactly one puncture, then $\Delta(\mathbf{S}, \mathbf{M})$ is isomorphic to a connected component of $\Delta^\diamond(\mathbf{S}, \mathbf{M})$, and the exchange graph of \mathcal{A} is isomorphic to a connected component of $\mathbf{E}^\diamond(\mathbf{S}, \mathbf{M})$.

The case of a two-punctured closed surface will be treated later [14].

Thus, the cluster variables in \mathcal{A} can be labeled by the tagged arcs in (\mathbf{S}, \mathbf{M}) so that two cluster variables are compatible if and only if the corresponding tagged arcs are compatible. If (\mathbf{S}, \mathbf{M}) is a closed surface with exactly one puncture, we only use the plain arcs.

By Theorem 7.11, the tagged arc complex of (\mathbf{S}, \mathbf{M}) is identified with the cluster complex of a cluster algebra \mathcal{A} whose seeds are labeled by the tagged triangulations

of (\mathbf{S}, \mathbf{M}) . The exchange matrix at each seed can be directly described in terms of the corresponding tagged triangulation; the exact combinatorial recipe is given in §9.3.

8. Denominators and intersection numbers

One of the most fundamental results of cluster algebra theory is the following.

THEOREM 8.1. (Laurent phenomenon for cluster algebras [16]) *Any cluster variable, when expressed as a rational function in the elements of a given cluster, is a Laurent polynomial.*

It is a natural idea—already used in many papers, starting with [17]—to use the *denominators* of these Laurent polynomials to distinguish the cluster variables from each other. Let us recall this setup, largely following [20, §7].

Definition 8.2. (Denominator vectors) Let \mathbf{x}_0 be a cluster from some seed in a cluster algebra \mathcal{A} . By Theorem 8.1, any non-zero element $z \in \mathcal{A}$ can be uniquely written as

$$z = \frac{P(\mathbf{x}_0)}{\prod_{x \in \mathbf{x}_0} x^{d(x|z)}}, \quad (8.1)$$

where $\mathbf{d}_{\mathbf{x}_0}(z) = (d(x|z) : x \in \mathbf{x}_0)$ is an n -tuple of integers, and $P(\mathbf{x}_0)$ is a polynomial not divisible by any cluster variable $x \in \mathbf{x}_0$. The vector $\mathbf{d}_{\mathbf{x}_0}(z)$ is called the *denominator vector* of z with respect to the cluster \mathbf{x}_0 .

According to [20, Conjecture 7.4 (3)], the number $d(x|z)$ depends only on x and z , but not on the choice of cluster \mathbf{x}_0 containing x . Since we cannot assume this to hold *a priori*, we will sometimes use the notation $d_{\mathbf{x}_0}(x|z)$ to reflect the possible dependence on \mathbf{x}_0 .

As noted in [20], the denominator vectors of cluster variables (and in fact their entire Newton polytopes) do not depend on the choice of coefficients. This is because the exchange relations (5.1) imply the following explicit recurrence for the denominator vectors.

LEMMA 8.3. ([20, §7]) *The components $d(x|z) = d_{\mathbf{x}_0}(x|z)$ of denominator vectors of cluster variables with respect to a fixed initial cluster \mathbf{x}_0 are uniquely determined by the initial conditions*

$$d(x|x) = -1 \quad \text{and} \quad d(x|z) = 0 \quad \text{for } z \in \mathbf{x}_0 \setminus \{x\}, \quad (8.2)$$

together with the recurrence relations

$$d(x|\bar{z}) = -d(x|z) + \max\left(\sum_{u \in \mathbf{x}} [b_{uz}]_+ d(x|u), \sum_{u \in \mathbf{x}} [-b_{uz}]_+ d(x|u)\right), \quad (8.3)$$

for each pair of adjacent seeds $\Sigma = (\mathbf{x}, \mathbf{p}, B)$ and $\mu_z(\Sigma) = (\bar{\mathbf{x}}, \bar{\mathbf{p}}, \bar{B})$, with $B = (b_{uv})_{u,v \in \mathbf{x}}$ and $\bar{\mathbf{x}} = (\mathbf{x} \setminus \{z\}) \cup \{\bar{z}\}$ (cf. (5.1)). Here we use the notation $[a]_+ = \max(a, 0)$.

The relation (8.3) can be viewed as a *tropical* version of the exchange relation (5.1). That is, (8.3) can be obtained from (5.1) by replacing ordinary multiplication and addition by their tropical counterparts, addition and maximum, respectively.

We will show that for the class of cluster algebras under our consideration, the denominators $d(x|z)$ can be interpreted as certain modified intersection numbers. The existence of such an interpretation was also suggested in [24].

Definition 8.4. (Intersection pairing) Let α and β be two tagged arcs in (\mathbf{S}, \mathbf{M}) . The *intersection number* $(\alpha|\beta)$ is defined as follows. Let α_0 and β_0 be two non-self-intersecting curves homotopic relative to \mathbf{M} to the untagged versions of α and β , respectively, and intersecting transversally in a minimum number of points in $\mathbf{S} \setminus \mathbf{M}$. Then

$$(\alpha|\beta) \stackrel{\text{def}}{=} A + B + C + D,$$

where

- A is the number of intersection points of α_0 and β_0 in $\mathbf{S} \setminus \mathbf{M}$;
- $B = 0$ unless α_0 is a loop based at a marked point a , in which case B is computed as follows: assume that β_0 intersects α_0 (in the interior of $\mathbf{S} \setminus \mathbf{M}$) at the points b_1, \dots, b_m (numbered along β_0 in this order); then B is the *negative* of the number of segments $\gamma_i = [b_i, b_{i+1}] \subset \beta_0$ such that γ_i together with the segments $[a, b_i] \subset \alpha_0$ and $[a, b_{i+1}] \subset \alpha_0$ form three sides of a contractible triangle disjoint from the punctures (see Figure 21);
- $C = 0$ unless $\alpha_0 = \beta_0$, in which case $C = -1$;
- D is the number of ends of β that are incident to an endpoint of α and, at that endpoint, carry a tag different from the one α does.

Example 8.5. To illustrate Definition 8.4, let α be a plain loop based at a puncture a , and let β be an arc connecting a with $b \neq a$, notched at a and not having other common points with α (see Figure 22). Then $(\alpha|\beta) = 1$ while $(\beta|\alpha) = 2$.

From now on, we assume that (\mathbf{S}, \mathbf{M}) is not a closed surface with exactly two punctures, and that \mathcal{A} is a cluster algebra with $\mathcal{B}(\mathcal{A}) = \mathcal{B}(\mathbf{S}, \mathbf{M})$. By Theorem 7.11, the cluster variables in \mathcal{A} can be labeled by the tagged arcs in (\mathbf{S}, \mathbf{M}) . Let $x[\beta]$ denote the cluster variable labeled by β .

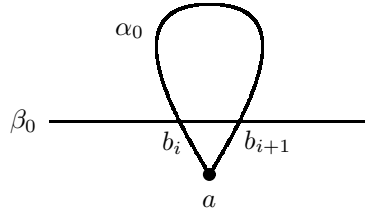


Figure 21. Computing B in Definition 8.4.



Figure 22. An example with $(\alpha|\beta) \neq (\beta|\alpha)$.

THEOREM 8.6. *For any pair of tagged arcs α and β , the denominator component $d(x[\alpha]|x[\beta])$ is equal to the intersection number $(\alpha|\beta)$. In particular, $d(x[\alpha]|x[\beta])$ does not depend on the choice of a cluster containing $x[\alpha]$.*

In the special case of ordinary (rather than tagged) arcs, Theorem 8.6 rectifies a similar albeit incomplete statement made in [24, Theorem 3.4].

Theorem 8.6 establishes [20, Conjecture 7.4] for the class of cluster algebras considered in this paper.

Theorem 8.6 can be restated as saying that for any tagged arcs $\alpha_1, \dots, \alpha_n$ forming a tagged triangulation of (\mathbf{S}, \mathbf{M}) , and any tagged arc β , we have

$$x[\beta] = \frac{P(x[\alpha_1], \dots, x[\alpha_n])}{\prod_{i=1}^n x[\alpha_i]^{(\alpha_i|\beta)}}, \tag{8.4}$$

where P is a polynomial not divisible by any variable $x[\alpha_i]$.

Theorem 8.6 is proved in §9 by verifying that the intersection numbers $(\alpha|\beta)$ satisfy the recurrence (8.3).

9. Proofs of main results

In this section, we prove Theorems 5.6, 7.9, 7.11, 8.6 and Proposition 7.10.

9.1. Basic properties of tagged arc complexes

Definition 9.1. (Signature of a partial tagged triangulation) Let C be a simplex in $\Delta^\times(\mathbf{S}, \mathbf{M})$, i.e., a collection of pairwise compatible tagged arcs. For a puncture a , let us temporarily denote by $S_C(a) \subset \{\text{plain}, \text{notched}\}$ the set of tags at a that appear in C . The *signature* δ_C of C is the $(0, \pm 1)$ -valued function on the set of all punctures in (\mathbf{S}, \mathbf{M}) defined by

$$\delta_C(a) = \begin{cases} 1, & \text{if } S_C(a) = \{\text{plain}\}, \\ -1, & \text{if } S_C(a) = \{\text{notched}\}, \\ 0, & \text{if } S_C(a) = \{\text{plain}, \text{notched}\} \text{ or } S_C(a) = \emptyset. \end{cases} \quad (9.1)$$

Note that if $S_C(a) = \{\text{plain}, \text{notched}\}$, then there are precisely two arcs in C incident to a : the untagged versions of these two arcs coincide and they are tagged in the same way at the end different from a .

Definition 9.2. Let ε be a ± 1 -valued function on the set of punctures in (\mathbf{S}, \mathbf{M}) . For γ an (ordinary) arc, define $\tau(\gamma, \varepsilon)$ to be the tagged arc obtained from $\tau(\gamma)$ (see Definition 7.2) by changing the tags at all punctures a with $\varepsilon(a) = -1$ from plain to notched and vice versa. For a collection C of compatible ordinary arcs, define

$$\tau(C, \varepsilon) = \{\tau(\gamma, \varepsilon) : \gamma \in C\}.$$

It is clear that the tagged arcs in $\tau(C, \varepsilon)$ are pairwise compatible.

Conversely, for a collection C of compatible tagged arcs, let C° be the collection of ordinary arcs obtained as follows:

- replace all (notched) tags at the punctures a with $\delta_C(a) = -1$ by plain ones;
- for each puncture b with $\delta_C(b) = 0$, replace the arc β notched at b (if any) by a loop γ enclosing b and β ; that is, γ is based at the endpoint $a \neq b$ of β , and is obtained by closely wrapping around β , as shown in Figure 18.

It is straightforward to check that the arcs in C° are pairwise compatible and distinct.

LEMMA 9.3. *Let C be a collection of compatible tagged arcs. Let ε be any ± 1 -valued function on the set of punctures in (\mathbf{S}, \mathbf{M}) such that $\varepsilon(a)\delta_C(a) \geq 0$ for all punctures a . Then $\tau(C^\circ, \varepsilon) = C$.*

Proof. This is clear from the definitions. □

Proof of Theorem 7.9. We need to show that the tagged arc complex $\Delta^\times(\mathbf{S}, \mathbf{M})$ is a pseudomanifold. First, let us prove that $\Delta^\times(\mathbf{S}, \mathbf{M})$ is pure of dimension $n-1$. (As before, n denotes the rank of (\mathbf{S}, \mathbf{M}) , i.e., the cardinality of each maximal simplex in $\Delta^\circ(\mathbf{S}, \mathbf{M})$.)

For any collection C of k pairwise compatible tagged arcs, C° is a collection of k pairwise compatible ordinary arcs. It follows that any simplex in $\Delta^\bowtie(\mathbf{S}, \mathbf{M})$ has cardinality at most n .

Suppose that C is a simplex in $\Delta^\bowtie(\mathbf{S}, \mathbf{M})$ of cardinality less than n . We must show that C is not maximal. Since C° is not maximal, there exists an arc $\gamma \notin C^\circ$ compatible with all arcs in C° . Let ε be any ± 1 -valued function with $\varepsilon(a)\delta_C(a) \geq 0$ for all punctures a . Then, by Lemma 9.3, $\tau(\gamma, \varepsilon)$ is compatible with C .

It remains to demonstrate that each codimension-1 simplex C in $\Delta^\bowtie(\mathbf{S}, \mathbf{M})$ is contained in exactly two maximal simplices (tagged triangulations). Note that if a loop starting at a puncture a enclosing a monogon with a single puncture b appears in C° , the enclosed arc from a to b also appears in C° . Thus for any codimension-1 simplex C , the ordinary collection C° is not on the boundary of $\Delta^\circ(\mathbf{S}, \mathbf{M})$ and so can be completed by two different arcs γ and γ' . The tagged arcs $\tau(\gamma, \varepsilon)$ and $\tau(\gamma', \varepsilon)$ give two different ways to complete C to a maximal simplex. To show that these are the only two completions, reverse the construction: Let $\tilde{C} \supset C$ be a maximal simplex in $\Delta^\bowtie(\mathbf{S}, \mathbf{M})$. Then $\tilde{C} \setminus C^\circ$ consists of a single ordinary arc compatible with C° , so it must be one of the arcs γ and γ' . Furthermore, we have $\varepsilon(d)\delta_{\tilde{C}}(d) \geq 0$ for all punctures d , so, by Lemma 9.3, the sole arc in $\tilde{C} \setminus C$ is either $\tau(\gamma, \varepsilon)$ or $\tau(\gamma', \varepsilon)$. \square

9.2. Strata of tagged triangulations

We next discuss the relationship between the graphs $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$ and $\mathbf{E}^\bowtie(\mathbf{S}, \mathbf{M})$ of ordinary and tagged flips, making use of the following natural concepts.

Definition 9.4. (Open and closed strata of tagged triangulations) For a $(0, \pm 1)$ -function δ on the set of punctures, the *open stratum* Ω_δ is the set of all tagged triangulations T with signature $\delta_T = \delta$. For a ± 1 -function ε on the set of punctures, the *closed stratum* $\bar{\Omega}_\varepsilon$ consists of all tagged triangulations T such that $\delta_T(a) = \varepsilon(a)$ whenever $\delta_T(a) = \pm 1$. In other words, $\bar{\Omega}_\varepsilon$ is the union of all open strata Ω_δ satisfying $\delta(a)\varepsilon(a) \geq 0$ for all punctures a .

The ‘‘positive’’ closed stratum $\bar{\Omega}_1$ associated with the function $\varepsilon = \mathbf{1}$ defined by setting $\varepsilon(a) = 1$ for all punctures a is identified with the set of all ideal triangulations of (\mathbf{S}, \mathbf{M}) by the map $T \mapsto \tau(T) = \tau(T, \mathbf{1})$. Accordingly, it makes sense to talk about the *signature* of an ideal triangulation, defined as a $(0, 1)$ -function on the set of punctures which takes values 0 and 1 depending on whether a puncture is located inside a self-folded triangle or not.

As before, let p denote the number of punctures in (\mathbf{S}, \mathbf{M}) . For each ± 1 -function ε on the set of punctures, the induced subgraph of $\mathbf{E}^\bowtie(\mathbf{S}, \mathbf{M})$ whose vertex set is the

closed stratum $\bar{\Omega}_\varepsilon$ is canonically isomorphic to the graph $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$ via the construction of Definition 9.2. Thus, $\mathbf{E}^{\bowtie}(\mathbf{S}, \mathbf{M})$ can be obtained by gluing together 2^p copies of $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$.

Proof of Proposition 7.10. (Connectedness of the dual graph $\mathbf{E}^{\bowtie}(\mathbf{S}, \mathbf{M})$.) If (\mathbf{S}, \mathbf{M}) is a closed surface with one puncture, then the tagged arcs forming a simplex in $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$ have either all plain ends or all notched ends (by Definition 7.1), so there are two connected components. Switching between plain and notched tags gives an isomorphism between the two components.

Otherwise, in view of Proposition 3.8, it suffices to show that any tagged triangulation can be transformed by a sequence of flips into an ideal triangulation—more precisely, into a tagged triangulation with no notches. First, we pick some closed stratum $\bar{\Omega}_\varepsilon$ and use Corollary 3.9 within that stratum to deduce that any tagged triangulation can be transformed by flips into a tagged triangulation T without punctures a with $\delta_T(a)=0$. We then transform the punctures a with $\delta_T(a)=-1$, one by one, into punctures with $\delta_T(a)=1$ by first placing a puncture inside a once-punctured digon and then applying the (tagged) flips shown in Figure 20 (moving from the bottom to the top). Note that we have excluded all cases in which placing a puncture inside a digon is not possible, either in Definition 2.1 or in the case when (\mathbf{S}, \mathbf{M}) is a closed surface with one puncture. \square

LEMMA 9.5. *If \mathbf{S} is not closed, then each open stratum Ω_δ is non-empty. If (\mathbf{S}, \mathbf{M}) is closed, then Ω_δ is non-empty unless $\delta(a)=0$ for all a , in which case Ω_δ is empty.*

Proof. It is enough to prove the lemma for the open strata contained in the distinguished “positive” closed stratum $\bar{\Omega}_1$. That is, it suffices to show that for any $(0, 1)$ -vector δ , except for the zero vector in the closed case, there exists an ideal triangulation with signature δ .

If there is some boundary, let a be a marked point on the boundary. Otherwise, let a be a puncture with $\delta(a) \neq 0$. In either case, connect a with each puncture b such that $\delta(b)=0$, so that these arcs do not intersect pairwise. Wrap a loop based at a around each of these arcs, forming a self-folded triangle. Triangulate the rest without such triangles if necessary (see Lemma 2.13). Each self-folded triangle that we cut out reduces the number of punctures p by 1, possibly increases the number of boundary components, and adds one to the number of marked points on the boundary; this never goes from a surface allowing triangulations to one with no triangulations.

To see that the stratum with $\delta(a)=0$ for all punctures a is empty in the closed case, note that for a self-folded triangle in an ideal triangulation, the exterior vertex a must have $\delta(a)=\pm 1$. \square

9.3. Mutations and flips in the tagged case

We next define adjacency matrices of tagged triangulations, and verify that, in analogy with Proposition 4.8, flips (i.e., edges in $\mathbf{E}^{\triangleright}(\mathbf{S}, \mathbf{M})$) correspond to mutations of these matrices.

Definition 9.6. (Signed adjacency matrix of a tagged triangulation) Let T be a tagged triangulation of (\mathbf{S}, \mathbf{M}) , with its tagged arcs labeled 1 through n . We define the (generalized) signed adjacency matrix $B(T) = (b_{ij})$ as follows. Let T° be the ideal triangulation associated with T ; for each tagged arc in T , the corresponding ordinary arc in T° retains the same label. We then set $B(T) = B(T^\circ)$.

Definition 9.6 can also be restated in the spirit of Remark 4.2.

LEMMA 9.7. Let T and \bar{T} be labeled tagged triangulations related by a flip of a tagged arc labeled k . (The rest of the labeling stays put.) Then $B(\bar{T}) = \mu_k(B(T))$.

Proof. There is always a closed stratum $\bar{\Omega}_\varepsilon$ containing both T and \bar{T} . Within that stratum Proposition 4.8 applies. \square

9.4. Cycles in the graph of flips

Our next goal is to prove, using Theorem 3.10, that all cycles in the graph of “tagged flips” are generated by cycles of length 4 and 5 (barring the exceptional cases of closed surfaces with two punctures). Our proof is based on a connectivity result (Lemma 9.13 below) which in turn relies on the following more general version of the notion of arc complex.

Definition 9.8. Let (\mathbf{S}, \mathbf{M}) be a surface as in Definition 2.1, without the condition that there be at least one marked point on each boundary component. An *arc* in (\mathbf{S}, \mathbf{M}) is as in Definition 2.2, and two arcs are *compatible* if they do not intersect in the interior of \mathbf{S} (as in Definition 2.4). The *arc complex* $\Delta^\circ(\mathbf{S}, \mathbf{M})$ is as in Definition 3.1 the clique complex for the compatibility relation, while a (*generalized*) *ideal triangulation* of (\mathbf{S}, \mathbf{M}) is a maximal collection of disjoint compatible arcs. The triangulation is generalized because if there are boundary components without marked points (such a component is called a *hole*), then some of the complementary regions will not be triangles but rather annuli with one marked point on one boundary component.

Inside these generalized ideal triangulations, there are two kinds of (generalized) flips: the flips inside a quadrilateral that we saw before, and the flips inside a digon with a hole removed, as in Figure 23. The graph $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$ is defined as before with this new notion of flips.

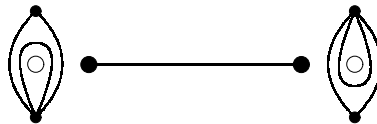
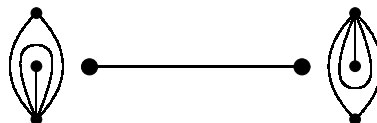


Figure 23. A flip inside a digon with a hole.

Figure 24. A \diamond -flip.

PROPOSITION 9.9. ([27]) *For (\mathbf{S}, \mathbf{M}) as in Definition 9.8, the graph $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$ is connected.*

In fact, Hatcher proves a much more general result.

THEOREM 9.10. ([27]) *For (\mathbf{S}, \mathbf{M}) as in Definition 9.8, the arc complex $\Delta^\circ(\mathbf{S}, \mathbf{M})$ is contractible unless (\mathbf{S}, \mathbf{M}) is an unpunctured polygon or a polygon with one hole removed; in these exceptional cases $\Delta^\circ(\mathbf{S}, \mathbf{M})$ is a sphere [26].*

Since $\Delta^\circ(\mathbf{S}, \mathbf{M})$ is still a pseudomanifold, this implies the proposition.

Definition 9.11. (\diamond -flips) A \diamond -flip is a transformation of an ideal triangulation that is either an ordinary flip, or a combination of two flips occurring inside a once-punctured digon (see Figure 24). This notion is naturally extended to tagged triangulations by replicating the above definition inside each closed stratum. That is, two tagged triangulations T_1 and T_2 are related by a \diamond -flip if, first, they belong to the same closed stratum $\bar{\Omega}_\epsilon$ and, second, the ideal triangulations T_1° and T_2° associated to them are related by a \diamond -flip.

Comparing Figures 23 and 24, we see that a \diamond -flip corresponds to a generalized flip for a surface in which the puncture inside the digon has been replaced by a hole.

Definition 9.12. (\diamond -paths) A \diamond -path is a sequence of (ideal or tagged) triangulations in which any two consecutive triangulations are related by a \diamond -flip. A path Π in $\mathbf{E}^{\times}(\mathbf{S}, \mathbf{M})$ is a realization of a \diamond -path Π° if Π can be obtained from Π° by refining each \diamond -flip which is not an ordinary flip by a pair of ordinary flips.

LEMMA 9.13. *Assume that (\mathbf{S}, \mathbf{M}) has more than one puncture. Let b be a puncture. Then any two ideal triangulations T with $\delta_T(b)=0$ (i.e., b lies inside a self-folded triangle) are \diamond -connected. Consequently, the intersection of each closed stratum of tagged triangulations with the set $\{T:\delta_T(b)=0\}$ is \diamond -connected.*

Proof. This follows from Proposition 9.9. \square

Definition 9.14. (R2-cycles) Let T be a tagged triangulation, and let $B(T)=(b_{ij})$ be the corresponding (generalized) signed adjacency matrix. Consider two arcs in T labeled i and j such that $b_{ij}=0$ or $b_{ij}=\pm 1$. It is easy to check that alternately flipping the two arcs labeled i and j (while retaining the labeling of the remaining arcs) brings us back to the original tagged triangulation T after 4 flips (if $b_{ij}=0$) or after 5 flips (if $b_{ij}=\pm 1$). The corresponding cycle in the exchange graph $\mathbf{E}^{\heartsuit}(\mathbf{S}, \mathbf{M})$ is called an R2-cycle.

Remark 9.15. In [17, §2.1], the R2-cycles were called *geodesic loops* (with respect to the canonical discrete connection on the exchange graph). We avoid using this terminology here because of the misleading connotations it brings as we speak of curves on surfaces.

The terminology in Definition 9.14 reflects the fact that, in effect, we are performing flips in a (possibly disconnected) bordered surface of rank $n=2$, specifically a pentagon, or a disjoint union of two quadrilaterals, or a once-punctured digon.

Definition 9.16. (R2-homotopy) Two paths in $\mathbf{E}^{\heartsuit}(\mathbf{S}, \mathbf{M})$ are called R2-homotopic to each other if they can be obtained from each other by a sequence of transformations of the following two kinds:

- removing two consecutive edges on a path which trace the same edge in $\mathbf{E}^{\heartsuit}(\mathbf{S}, \mathbf{M})$ in opposite directions;
- replacing a fragment of a path (i.e., several consecutive edges of it) contained in an R2-cycle by the complement of that fragment inside this R2-cycle.

We note that any two realizations of the same \diamond -path (see Definition 9.12) are R2-homotopic. Accordingly, two \diamond -paths are called R2-homotopic if (any of) their respective realizations in $\mathbf{E}^{\heartsuit}(\mathbf{S}, \mathbf{M})$ are R2-homotopic to each other.

THEOREM 9.17. *Assume that (\mathbf{S}, \mathbf{M}) is not a closed surface with exactly two punctures. Then any cycle in $\mathbf{E}^{\heartsuit}(\mathbf{S}, \mathbf{M})$ is R2-homotopic to a point.*

Proof. By Theorem 3.10, the claim holds for any cycle contained in any closed stratum $\bar{\Omega}_\varepsilon$, since we can identify it with $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$ via the map $\tau(\cdot, \varepsilon)$ (see Definition 9.4). To rephrase, any two paths in $\mathbf{E}^{\heartsuit}(\mathbf{S}, \mathbf{M})$ that have common endpoints and are entirely contained in $\bar{\Omega}_\varepsilon$ are R2-homotopic to each other. The idea of the proof is to successively use such R2-homotopies to contract a given cycle Π .

First, we fix a linear order on the set of closed strata. This will determine the order in which the portions of Π contained in these strata will be deformed. Specifically, let $\varepsilon_1, \dots, \varepsilon_{2^p}$ be any linear ordering on the 2^p -element set of ± 1 -functions on the set of

punctures that satisfies

$$\sum_p \varepsilon_i(p) < \sum_p \varepsilon_j(p) \implies i < j.$$

The only purpose of this condition is to ensure that, for any i with $1 \leq i < 2^p$, we have

$$\bar{\Omega}_{\varepsilon_i} \cap \left(\bigcup_{j>i} \bar{\Omega}_{\varepsilon_j} \right) = \bigcup_{j>i} (\bar{\Omega}_{\varepsilon_i} \cap \bar{\Omega}_{\varepsilon_j}) = \bigcup_{\substack{j>i \\ \varepsilon_j \succ \varepsilon_i}} (\bar{\Omega}_{\varepsilon_i} \cap \bar{\Omega}_{\varepsilon_j}), \quad (9.2)$$

where the notation $\varepsilon_j \succ \varepsilon_i$ means that the ± 1 -functions ε_i and ε_j coincide at all punctures except for a single puncture a for which $\varepsilon_j(a) > \varepsilon_i(a)$.

Starting with the given cycle Π , we are going to successively perform R2-homotopies of the fragments of a cycle contained inside the closed strata $\bar{\Omega}_{\varepsilon_1}, \bar{\Omega}_{\varepsilon_2}, \dots$ (in this order). For this plan to work, we need to make sure that, after each deformation occurring inside a closed stratum $\bar{\Omega}_{\varepsilon_i}$, the deformed fragment is contained in the union of closed strata $\bigcup_{j>i} \bar{\Omega}_{\varepsilon_j}$.

Let Θ be a fragment of interest to us; that is, Θ is a path in $\mathbf{E}^{\bowtie}(\mathbf{S}, \mathbf{M})$ contained entirely in $\bar{\Omega}_{\varepsilon_i}$ which begins and ends at some tagged triangulations T_1 and T_2 located at the “boundary” of this closed stratum. More precisely, we may assume, taking (9.2) into account, that

- $T_1 \in \bar{\Omega}_{\varepsilon_j}$ for some $j > i$ such that $\varepsilon_j \succ \varepsilon_i$; say, $\varepsilon_j(a) > \varepsilon_i(a)$;
- $T_2 \in \bar{\Omega}_{\varepsilon_k}$ for some $k > i$ such that $\varepsilon_k \succ \varepsilon_i$; say, $\varepsilon_k(b) > \varepsilon_i(b)$.

In particular, $\delta_{T_1}(a) = \delta_{T_2}(b) = 0$.

Suppose that $j = k$ (so that $a = b$). Then both T_1 and T_2 lie in $\bar{\Omega}_{\varepsilon_i} \cap \bar{\Omega}_{\varepsilon_j}$, which is precisely the set of tagged triangulations T in $\bar{\Omega}_{\varepsilon_i}$ for which $\delta_T(a) = 0$. By Lemma 9.13, there exists a \diamond -path inside $\bar{\Omega}_{\varepsilon_i} \cap \bar{\Omega}_{\varepsilon_j}$ that connects T_1 and T_2 . By Theorem 3.10, this \diamond -path is R2-homotopic to Θ . The latter is thus R2-homotopic to a path in $\mathbf{E}^{\bowtie}(\mathbf{S}, \mathbf{M})$ contained entirely in $\bar{\Omega}_{\varepsilon_j}$, and we are done.

Now suppose that $j \neq k$ (so that $a \neq b$). It follows from Lemma 9.5 (here we use the condition in the theorem) that there exists a tagged triangulation $T_0 \in \bar{\Omega}_{\varepsilon_i} \cap \bar{\Omega}_{\varepsilon_j} \cap \bar{\Omega}_{\varepsilon_k}$. By Lemma 9.13, there exist \diamond -paths inside $\bar{\Omega}_{\varepsilon_i} \cap \bar{\Omega}_{\varepsilon_j}$ and $\bar{\Omega}_{\varepsilon_i} \cap \bar{\Omega}_{\varepsilon_k}$, respectively, that connect T_0 with T_1 and T_2 , respectively. Again by Theorem 3.10, Θ is R2-homotopic to the concatenation of these paths, which can in turn be realized as a path in $\mathbf{E}^{\bowtie}(\mathbf{S}, \mathbf{M})$ contained entirely in $\bar{\Omega}_{\varepsilon_j} \cup \bar{\Omega}_{\varepsilon_k}$. The theorem is proved. \square

Remark 9.18. The above proof can be viewed as a sequence of repeated applications of the Mayer–Vietoris theorem.

Remark 9.19. The statement of Theorem 9.17 is *false* for closed surfaces with two punctures, for the following reasons. Let (\mathbf{S}, \mathbf{M}) be such a surface. Then the set of all tagged triangulations is a disjoint union of the eight open strata determined by the signatures at the two punctures (recall that the signature $(0, 0)$ is impossible):

$$\begin{array}{ccc}
 \Omega_{-+} \longleftrightarrow \Omega_{0+} \longleftrightarrow \Omega_{++} & & \\
 \updownarrow & & \updownarrow \\
 \Omega_{-0} & & \Omega_{+0} \\
 \updownarrow & & \updownarrow \\
 \Omega_{--} \longleftrightarrow \Omega_{0-} \longleftrightarrow \Omega_{+-} . & &
 \end{array} \tag{9.3}$$

All edges in $\mathbf{E}^{\text{tri}}(\mathbf{S}, \mathbf{M})$ connect tagged triangulations that either belong to the same open stratum or else to a pair of strata connected by an arrow in the diagram (9.3). Consider a cycle Π that goes through these eight strata clockwise. Since each R2-cycle is of length 4 or 5, it is impossible to represent Π as a product of R2-cycles.

9.5. Tropical exchange relations for intersection numbers

In this section, we return to the study of the intersection numbers $(\alpha|\beta)$ introduced in §8. We will demonstrate that these numbers satisfy the “tropical recurrence” (8.3).

LEMMA 9.20. *Let $T=(\beta_1, \dots, \beta_n)$ be a tagged triangulation of (\mathbf{S}, \mathbf{M}) , with the signed adjacency matrix $B(T)=(b_{ij})$. Let \bar{T} be the tagged triangulation obtained from T by a flip replacing β_k by a tagged arc $\bar{\beta}_k$. Then, for any tagged arc α in (\mathbf{S}, \mathbf{M}) , we have*

$$(\alpha|\bar{\beta}_k) = -(\alpha|\beta_k) + \max\left(\sum_{i=1}^n [b_{ik}]_+ (\alpha|\beta_i), \sum_{i=1}^n [-b_{ik}]_+ (\alpha|\beta_i)\right). \tag{9.4}$$

Proof. Without loss of generality, we may assume that the tagged triangulations T and \bar{T} belong to the positive closed stratum (see Definition 9.4). The proof then proceeds case by case, for each combinatorially different type of a flip $T \rightarrow \bar{T}$. (Note that the roles of T and \bar{T} can be switched without affecting the statement (9.4) to be proved.) These types can be catalogued using the puzzle-piece decomposition of (the ideal triangulation corresponding to) T , similarly to the proof of Proposition 4.8.

The two most basic types of flips are (cf. Figure 25):

(F1) flipping a diagonal of a quadrilateral Q formed by two triangles (none of them self-folded) sharing precisely one edge;

(F2) flipping a tagged arc inside a once-punctured digon.

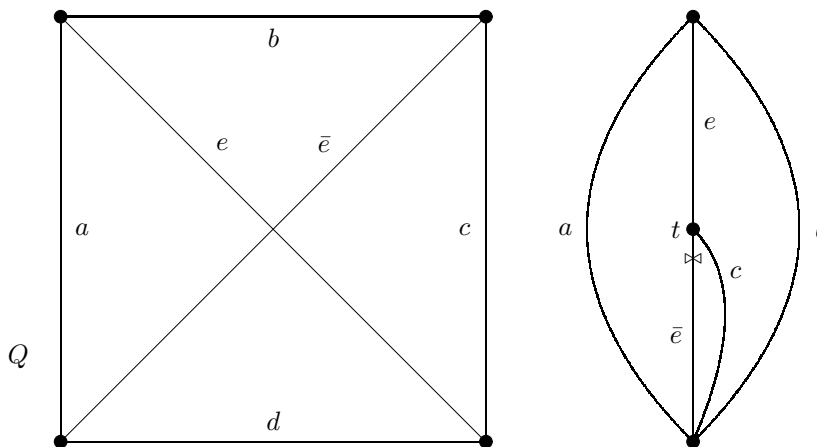


Figure 25. Flips of types F1 and F2.

We will prove (9.4) for the flips of type F1, sketch the proof for type F2, and omit the tedious verification of the remaining cases.

The only arcs $\beta \in T$ contributing to (9.4) are $\beta_k, \bar{\beta}_k$ and those β_i for which $b_{ik} \neq 0$. In the case of a flip of type F1, these arcs are precisely the sides and the diagonals of the quadrilateral Q . Let a, \dots, e, \bar{e} denote the corresponding intersection numbers, as shown in Figure 25 on the left. Then (9.4) takes the form

$$e + \bar{e} = \max(a + c, b + d). \quad (9.5)$$

To prove this, we first get rid of the tagging. If α is notched at some vertex of Q , then removing the notch would subtract one from each of the sums $a + c, b + d$ and $e + \bar{e}$. We may therefore assume that α is a plain arc. We may also assume that α is none of the six arcs shown in Figure 25, since otherwise (9.5) can be checked directly. Thus, we need not worry about the summands C and D in the definition of intersection numbers, and only take careful account of the summands A and B .

First, let us assume that α is not a loop, so that $B = 0$ throughout. Each segment $\sigma \subset \alpha$ cutting across Q makes contributions to the three sums $a + c, b + d$ and $e + \bar{e}$; we denote these contributions by q_{ac}, q_{bd} and $q_{e\bar{e}}$, respectively. Then the following cases may occur:

- (F1-1) σ cuts through two consecutive sides of Q , with $q_{ac} = q_{bd} = q_{e\bar{e}} = 1$;
- (F1-2) σ cuts through two opposite sides of Q , with $q_{e\bar{e}} = 2, q_{ac} = 2$ and $q_{bd} = 0$;
- (F1-3) σ cuts through two opposite sides of Q , with $q_{e\bar{e}} = 2, q_{ac} = 0$ and $q_{bd} = 2$;
- (F1-4) σ starts at a vertex of Q , with $q_{e\bar{e}} = 1, q_{ac} = 1$ and $q_{bd} = 0$;
- (F1-5) σ starts at a vertex of Q , with $q_{e\bar{e}} = 1, q_{ac} = 0$ and $q_{bd} = 1$.

It is easy to check (keeping in mind that different segments σ do not intersect each other, and that α is not a loop) that we are left with the following possibilities:

- all σ 's are of types F1-1, F1-2 or F1-4, so that $q_{e\bar{e}}=q_{ac} \geq q_{bd}$ in each case;
- all σ 's are of types F1-1, F1-3 or F1-5, so that $q_{e\bar{e}}=q_{bd} \geq q_{ac}$ in each case.

We then conclude, by additivity, that in these cases one has $e+\bar{e}=a+c \geq b+d$ or $e+\bar{e}=b+d \geq a+c$, respectively, yielding (9.5).

If α is a loop, we need to make some adjustments to the above proof. Namely, we need to account for the summands B in Definition 8.4, and also allow for the additional possibility that

- all σ 's but two are of type F1-1, one is of type F1-4 and one of type F1-5.

In the latter case, the total B -contribution to $e+\bar{e}$ is -1 , whereas the B -contributions to $a+c$ and $b+d$ are both 0. Combined with the A -contributions of the two segments of types F1-4 and F1-5, this results in equal quantities (namely -1) for all three sums; so in this case we actually get $e+\bar{e}=a+c=b+d$. In all other cases (that is, unless F1-4 and F1-5 occur simultaneously), the B -contributions associated with pairs of segments of α should be considered together with the corresponding A -contributions; each time, the same argument as above works, i.e., all inequalities will go the same way, consistently favouring $a+c$ over $b+d$ (or vice versa).

We note that essentially the same proof works in the case where two opposite sides of Q (or both pairs of opposite sides) are identified; keep in mind that in that case, the corresponding entries of the matrix $B(T)$ are doubled, making up for the gluing of intersection points.

For a flip of type F2, the formula (9.4) takes the form

$$e+\bar{e}=\max(a,b). \quad (9.6)$$

The argument is similar to the one given above in the F1 case, with the additional complication due to the presence of a puncture t inside the domain in which the flip occurs—in this case, inside a digon. As a result, we should consider the possibility that α may have t as an endpoint, or even be a loop based at t . As above, the proof examines several possible ways in which segments of α may cut through the digon. Each time the contributions to $e+\bar{e}$ turn out to be equal to the maximum of their counterparts for a and b . As before, the maximum is determined consistently for all segments (or pairs of segments, in the case of B -contributions, which are pooled together with the A -contributions from the same segments); that is, either all contributions to a are larger than the respective contributions to b , or the opposite happens. Either way, (9.6) follows.

Other types of flips are examined in a similar fashion. One of them is illustrated in Figure 26. This flip (dubbed F3) occurs at an arc shared by a triangle and a punctured

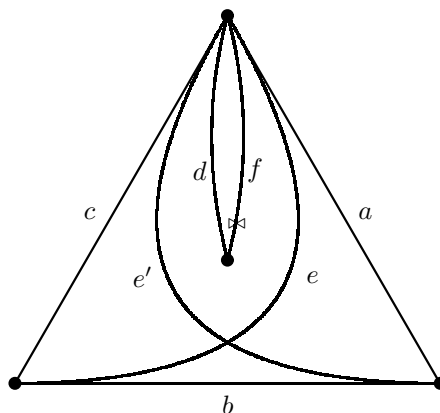


Figure 26. A flip of type F3.

digon. The formula (9.4) takes the form

$$e + \bar{e} = \max(a + c, b + d + f). \quad (9.7)$$

We omit the consideration of the remaining types of flips, namely flipping an arc shared by two once-punctured digons, and flips involving a twice-punctured monogon. \square

9.6. Proofs of Theorems 5.6, 7.11 and 8.6

The proofs are based on the following result.

PROPOSITION 9.21. ([17, Proposition 2.3 and Lemma 2.4]) *Let Δ be a simplicial complex on a (possibly infinite) ground set Ψ satisfying the following conditions:*

- Δ is an $(n-1)$ -dimensional pseudomanifold with a connected dual graph \mathbf{E} ;
- the dual graph of the link of every non-maximal simplex in Δ is connected;
- the fundamental group of \mathbf{E} is generated by cycles (pinned down to a fixed base-point) that arise as links of codimension-2 simplices.

Suppose that $\{B(C) = (b_{\beta\gamma})_{\beta, \gamma \in C}\}$ is a collection of skew-symmetrizable matrices labeled by the maximal simplices of Δ such that

- for each edge (C, \bar{C}) in \mathbf{E} with $\bar{C} = (C \setminus \{\gamma\}) \cup \{\bar{\gamma}\}$, the matrix $B(\bar{C})$ is obtained from $B(C)$ by a matrix mutation in direction (of the label of) γ (we assume that the labels for the elements of $C \cap \bar{C}$ are the same in C and in \bar{C});
- for each codimension-2 simplex $D \in \Delta$ whose link is a finite cycle (say of length ℓ), and for any maximal simplex $C = D \cup \{\beta, \gamma\}$, the number $|b_{\beta\gamma} b_{\gamma\beta}|$ is equal to 0, 1, 2 or 3, and furthermore ℓ is equal to 4, 5, 6 or 8, respectively.

Finally, let \mathcal{A} be a normalized cluster algebra with exchange matrix $B=B(C_0)$ for some $C_0 \in \mathbf{E}$. Then there is a (canonically defined) surjection $C \mapsto \Sigma(C)$ from \mathbf{E} to the set of seeds in \mathcal{A} and a (canonical) surjection $\beta \mapsto x[\beta]$ from Ψ to the set of cluster variables of \mathcal{A} such that

- for each edge (C, \bar{C}) in \mathbf{E} with $\bar{C}=(C \setminus \{\gamma\}) \cup \{\bar{\gamma}\}$, the seeds $\Sigma(C)$ and $\Sigma(\bar{C})$ are related by a mutation in direction γ ;
- the cluster variables in each seed $\Sigma(C)$ are precisely $\{x[\beta]: \beta \in C\}$.

We are going to use Proposition 9.21 for $\Delta=\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$, so that Ψ is the set of tagged arcs in (\mathbf{S}, \mathbf{M}) , and $\mathbf{E}=\mathbf{E}^{\bowtie}(\mathbf{S}, \mathbf{M})$ is the graph of tagged triangulations. The restrictions on Δ stated in Proposition 9.21 are satisfied by Theorem 7.9, Proposition 7.10 and Theorem 9.17, as long as (\mathbf{S}, \mathbf{M}) is not a surface with exactly two punctures. The connectedness of links follows by noting that each link is in fact a tagged arc complex of another (possibly disconnected) surface.

The matrices $B(C)$ associated with the maximal simplices of Δ are the signed adjacency matrices of tagged triangulations. The restrictions on them stated in Proposition 9.21 are satisfied by Lemma 9.7 and by direct examination of all possible R2-cycles (which are links of codimension-2 simplices). The only values of $|b_{\beta\gamma} b_{\gamma\beta}|$ that actually come up in this context are 0 and 1, corresponding to 4-cycles and 5-cycles in $\mathbf{E}^{\bowtie}(\mathbf{S}, \mathbf{M})$.

The conclusion of Proposition 9.21 then tells us that in a cluster algebra \mathcal{A} associated with (\mathbf{S}, \mathbf{M}) , there is a well-defined cluster variable $x[\beta]$ corresponding to each tagged arc β , and a well-defined seed $\Sigma(T)$ corresponding to each tagged triangulation T ; the cluster variables participating in $\Sigma(T)$ are precisely the $x[\beta]$'s for $\beta \in T$. In short, $\mathbf{E}^{\bowtie}(\mathbf{S}, \mathbf{M})$ and $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$ cover the exchange graph and cluster complex of \mathcal{A} .

To complete the proof of the main claims in Theorems 5.6 and 7.11, we need to show that this covering is in fact an isomorphism, that is, the aforementioned surjections are actually bijections. In practical terms, we need to show that all cluster variables $x[\beta]$ are distinct.

But first, let us prove Theorem 8.6. Now that there is a well-defined cluster variable $x[\beta]$ for each tagged arc β , the statement of Theorem 8.6 follows directly from Lemmas 8.3 and 9.20.

The fact that all cluster variables $x[\beta]$ are distinct is then immediate from Theorem 8.6, as the latter implies that for any $\gamma \neq \beta$, the expression for $x[\gamma]$ in terms of a seed containing $x[\beta]$ may not be equal to $x[\beta]$, since $(\gamma|\beta) > -1 = (\beta|\beta)$.

It remains to verify that the seeds containing a given cluster variable form a connected subgraph of $\mathbf{E}^{\bowtie}(\mathbf{S}, \mathbf{M})$. This statement is a special case of the ‘‘connectedness of links’’ property discussed earlier in the proof.

Remark 9.22. The assertion that all cluster variables $x[\beta]$ are distinct can be given a different proof that avoids the use of denominators and intersection numbers but instead relies on hyperbolic geometry. This argument will be given in [14].

10. On the topology of cluster complexes

In this section we determine (up to homotopy) the topology of the cluster complex associated with a general bordered surface with marked points—or, equivalently (see Theorem 7.11), the topology of the corresponding tagged arc complex $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$.

As before, (\mathbf{S}, \mathbf{M}) is a bordered surface with marked points (as described in Definition 2.1), and $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$ is its tagged arc complex (see Definition 7.7).

THEOREM 10.1. *The cluster complex of a cluster algebra \mathcal{A} associated with (\mathbf{S}, \mathbf{M}) is either contractible or homotopy equivalent to a sphere provided (\mathbf{S}, \mathbf{M}) is not a closed surface with exactly two punctures. Specifically:*

- *If \mathcal{A} is of finite type, the cluster complex is homeomorphic to an $(n-1)$ -dimensional sphere S^{n-1} .*
- *If (\mathbf{S}, \mathbf{M}) is a closed surface with $p \geq 2$ punctures, the cluster complex is homotopy equivalent to S^{p-1} .*
- *In all other cases, the cluster complex is contractible.*

Theorem 10.1 is a consequence of the following version for tagged arc complexes.

THEOREM 10.2. *The tagged arc complex $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$ is either contractible or homotopy equivalent to a sphere. Specifically:*

- *If (\mathbf{S}, \mathbf{M}) is a polygon or a once-punctured polygon, then $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$ is homeomorphic to S^{n-1} .*
- *If (\mathbf{S}, \mathbf{M}) is a closed surface with p punctures, then $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$ is homotopy equivalent to S^{p-1} .*
- *In all other cases, $\Delta^{\bowtie}(\mathbf{S}, \mathbf{M})$ is contractible.*

The proof of Theorem 10.2 will use the following classical tool from combinatorial topology (see, e.g., [2, Theorem 10.7] and references therein).

Definition 10.3. The *nerve* $\mathcal{N}(X)$ of a collection of sets $X = (X_i)_{i \in I}$ is the simplicial complex whose ground set is I and whose simplices are finite subsets $\{i_1, \dots, i_k\} \subset I$ such that $X_{i_1} \cap \dots \cap X_{i_k} \neq \emptyset$.

LEMMA 10.4. (The nerve lemma) *Let Δ be a triangulable topological space and $X=(X_i)_{i \in I}$ be a finite family of closed subsets of Δ such that $\Delta=\bigcup_{i \in I} X_i$. If each intersection $X_{i_1} \cap \dots \cap X_{i_k}$ is either empty or contractible, then Δ is homotopy equivalent to the nerve $\mathcal{N}(X)$.*

Proof of Theorem 10.2. If (\mathbf{S}, \mathbf{M}) has no punctures and is not a polygon, then the tagged arc complex $\Delta^\bowtie(\mathbf{S}, \mathbf{M})$ coincides with the ordinary untagged arc complex $\Delta^\circ(\mathbf{S}, \mathbf{M})$, which is contractible by Theorem 3.7. If (\mathbf{S}, \mathbf{M}) is an unpunctured or a once-punctured polygon, then $\Delta^\bowtie(\mathbf{S}, \mathbf{M})$ is the cluster complex of a cluster algebra of finite type A_n or D_n , which is homeomorphic to S^{n-1} [17, Theorem 1.13], [19, Corollary 1.11].

From now on, assume that $p \geq 1$ and that (\mathbf{S}, \mathbf{M}) is not a once-punctured polygon. Consider the map $\varphi: \Delta^\bowtie(\mathbf{S}, \mathbf{M}) \rightarrow \mathbb{R}^p$ which is linear on simplices and defined on vertices of $\Delta^\bowtie(\mathbf{S}, \mathbf{M})$ (tagged arcs γ) by

$$\varphi(\gamma; i) = \begin{cases} 0, & \text{if } \gamma \text{ does not meet puncture } i, \\ 1, & \text{if } \gamma \text{ is tagged plain at } i, \\ -1, & \text{if } \gamma \text{ is notched at } i. \end{cases} \tag{10.1}$$

For a ± 1 -valued function ε on the set of punctures in (\mathbf{S}, \mathbf{M}) , define

$$X_\varepsilon = X_\varepsilon(\mathbf{S}, \mathbf{M}) \stackrel{\text{def}}{=} \{x \in \Delta^\bowtie(\mathbf{S}, \mathbf{M}) : \varphi(x; i) \varepsilon(i) \geq 0 \text{ for all punctures } i\}.$$

The sets X_ε are closed and cover $\Delta^\bowtie(\mathbf{S}, \mathbf{M})$. Similarly, for a $(0, \pm 1)$ -valued function δ on the set of punctures, define

$$X_\delta \stackrel{\text{def}}{=} \{x \in \Delta^\bowtie(\mathbf{S}, \mathbf{M}) : \delta(i) = 0 \Rightarrow \varphi(x; i) = 0 \text{ and } \varphi(x; i) \delta(i) \geq 0 \text{ for all } i\}.$$

The intersection of any collection of X_ε 's is an X_δ .

For δ as above, let $(\mathbf{S}, \mathbf{M})_\delta$ be the surface with marked points obtained from (\mathbf{S}, \mathbf{M}) by replacing each puncture i with $\delta(i)=0$ by a *hole*, a boundary component with no marked points (cf. Definition 9.8).

LEMMA 10.5. *For a $(0, \pm 1)$ -valued function δ , X_δ is homeomorphic to the arc complex $\Delta^\circ(\mathbf{S}, \mathbf{M})_\delta$.*

Proof. By symmetry, we may assume that δ only has values of 0 and +1. Define the map $\tilde{\tau}: \Delta^\circ(\mathbf{S}, \mathbf{M}) \rightarrow \Delta^\bowtie(\mathbf{S}, \mathbf{M})$, a variant of the map τ of Definition 7.2, as follows. If γ does not cut out a once-punctured monogon, then $\tilde{\tau}(\gamma)$ is γ with both ends tagged plain, as before. Otherwise, let β be the plain arc enclosed by γ , and let β^\bowtie be the same arc with the encircled end notched and the other end plain; then $\tilde{\tau}(\gamma) = \frac{1}{2}(\beta + \beta^\bowtie)$.

Extend $\tilde{\tau}$ by linearity over simplices in $\Delta^\circ(\mathbf{S}, \mathbf{M})$. Note that the vertices of any simplex in $\Delta^\circ(\mathbf{S}, \mathbf{M})$ map inside a single simplex of $\Delta^\bowtie(\mathbf{S}, \mathbf{M})$, so $\tilde{\tau}$ is well defined.

The map $\tilde{\tau}$ maps $\Delta^\circ(\mathbf{S}, \mathbf{M})_\delta \subset \Delta^\circ(\mathbf{S}, \mathbf{M})$ inside $X_\delta \subset \Delta^\bowtie(\mathbf{S}, \mathbf{M})$, since for every vertex γ of $\Delta^\circ(\mathbf{S}, \mathbf{M})_\delta$ and every puncture i with $\delta(i)=0$, we have $\varphi(\tilde{\tau}(\gamma); i)=0$. Conversely, we can define an inverse map from X_δ to $\Delta^\circ(\mathbf{S}, \mathbf{M})_\delta$: the arcs incident to a puncture i with $\delta(i)=0$ must come in equal weights plain and notched, and we map such a mixture to a loop enclosing the puncture. \square

LEMMA 10.6. *X_δ is contractible if it is non-empty. X_δ is empty only if \mathbf{S} is a closed surface and δ is zero everywhere.*

Proof. The first statement follows from Lemma 10.5 and Theorem 9.10, since we have already ruled out the non-contractible cases in that theorem. The second statement is analogous to Lemma 9.5. \square

To complete the proof of Theorem 10.2, consider the covering of $\Delta^\bowtie(\mathbf{S}, \mathbf{M})$ by the sets X_ε . By Lemma 10.6, this covering satisfies the hypotheses of Lemma 10.4. If (\mathbf{S}, \mathbf{M}) is not a closed surface with punctures, then every intersection of a subset of the X_ε 's is non-empty. Thus $\mathcal{N}(X)$ and $\Delta^\bowtie(\mathbf{S}, \mathbf{M})$ are contractible. Otherwise, we use the following lemma.

LEMMA 10.7. *Let \mathcal{N}_p be the simplicial complex defined as follows:*

- *The ground set of \mathcal{N}_p is the set $\{1, -1\}^p$ of all ± 1 -vectors $\varepsilon=(\varepsilon_1, \dots, \varepsilon_p)$.*
- *A subset $J \subset \{1, -1\}^p$ is a simplex in \mathcal{N}_p if and only if there exist an index $k \in \{1, \dots, p\}$ and a sign $\sigma \in \{1, -1\}$ such that $\varepsilon_k \sigma \geq 0$ for all $\varepsilon \in J$.*

Then \mathcal{N}_p is homotopy equivalent to the $(p-1)$ -dimensional sphere S^{p-1} .

Proof. Let Δ be the boundary of the cube $[-1, 1]^p \subset \mathbb{R}^p$. For $\varepsilon=(\varepsilon_1, \dots, \varepsilon_p) \in \{1, -1\}^p$, let

$$X_\varepsilon = \Delta \cap \{x = (x_1, \dots, x_p) : x_k \varepsilon_k \geq 0 \text{ for all } k\}.$$

Then $\mathcal{N}(X) = \mathcal{N}_p$ and $\Delta \cong S^{p-1}$, so applying Lemma 10.4 establishes the claim. \square

If (\mathbf{S}, \mathbf{M}) is a closed surface with p punctures, then in view of Lemma 10.6 the nerve of the family $(X_\varepsilon(\mathbf{S}, \mathbf{M}))$ coincides with \mathcal{N}_p . Therefore by Lemma 10.7, $\mathcal{N}(X)$ and thus $\Delta^\bowtie(\mathbf{S}, \mathbf{M})$ are homotopy equivalent to the sphere S^{p-1} . \square

Remark 10.8. The mapping class group of (\mathbf{S}, \mathbf{M}) naturally acts on the cluster complex, or the tagged arc complex $\Delta^\bowtie(\mathbf{S}, \mathbf{M})$. The quotient is a finite simplicial complex that generalizes the ‘‘arc complexes’’ studied by R. Penner in [37].

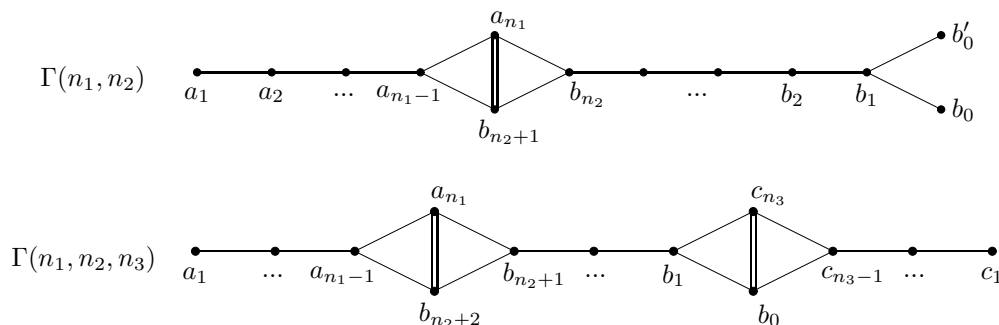


Figure 27. Diagrams for the cluster algebras of quadratic and cubic growth. All triangles are oriented. Orientations of the remaining edges are of no importance.

11. Polynomial vs. exponential growth

A cluster algebra (or the corresponding exchange graph) has *polynomial growth* if the number of distinct seeds which can be obtained from a fixed initial seed by at most n mutations is bounded from above by a polynomial function of n . A cluster algebra has *exponential growth* if the number of such seeds is bounded from below by an exponentially growing function of n . (There is always an exponential upper bound.)

PROPOSITION 11.1. *Let \mathcal{A} be a cluster algebra associated with a connected bordered surface with marked points. Then \mathcal{A} has polynomial growth if and only if it is defined by a diagram/quiver in the following list:*

- finite type A_n (finite);
- finite type D_n (finite);
- affine type $\tilde{A}(n_1, n_2)$ (linear growth);
- affine type \tilde{D}_{n-1} (linear growth);
- diagram $\Gamma(n_1, n_2)$, $n_1, n_2 \in \mathbb{Z}_{>0}$, shown in Figure 27 (quadratic growth);
- diagram $\Gamma(n_1, n_2, n_3)$, $n_1, n_2, n_3 \in \mathbb{Z}_{>0}$, shown in Figure 27 (cubic growth).

Otherwise \mathcal{A} has exponential growth.

For the class of cluster algebras considered in this paper, the growth rate coincides with the growth rate of the graph $\mathbf{E}^\times(\mathbf{S}, \mathbf{M})$ of tagged flips (by Theorem 7.11). The counterpart of Proposition 11.1 formulated in terms of the surface is given below.

PROPOSITION 11.2. *Let (\mathbf{S}, \mathbf{M}) be a connected bordered surface with marked points, which has b boundary components and p punctures. Then the graph of tagged flips $\mathbf{E}^\times(\mathbf{S}, \mathbf{M})$ has polynomial growth if and only if \mathbf{S} is a sphere and $b+p \leq 3$. In all other cases, $\mathbf{E}^\times(\mathbf{S}, \mathbf{M})$ has exponential growth.*

Proposition 11.2 implies Proposition 11.1: the possibilities for the numbers of punctures and boundary components are

- one boundary component (a disk) with $n+3$ marked points: type A_n ;
- one boundary component with n marked points, one puncture (a once-punctured disk): type D_n ;
- two boundary components (an annulus) with n_1 and n_2 marked points, respectively: type $\tilde{A}(n_1, n_2)$;
- one boundary component with $n-3$ marked points and two punctures (a twice-punctured disk): type \tilde{D}_{n-1} ;
- two boundary components with n_1 and n_2 marked points and one puncture (a once-punctured annulus): type $\Gamma(n_1, n_2)$;
- three boundary components with n_1, n_2 and n_3 marked points (a pair of pants): type $\Gamma(n_1, n_2, n_3)$.

Proof of Proposition 11.2. Let a *feature* of a surface be either a puncture or a boundary component (with a positive number of marked points). The graph $\mathbf{E}^{\triangleright}(\mathbf{S}, \mathbf{M})$ is covered by finitely many copies of the ordinary graph of ideal triangulations $\mathbf{E}^\circ(\mathbf{S}, \mathbf{M})$, which in turn is quasi-isometric to the mapping class group $\mathcal{MCG}(\mathbf{S}, \mathbf{M})$ (see, e.g., [29]). If (\mathbf{S}, \mathbf{M}) is a sphere with at most three features, then $\mathcal{MCG}(\mathbf{S}, \mathbf{M})$ is generated by Dehn twists along the boundary components, which commute with each other; hence $\mathcal{MCG}(\mathbf{S}, \mathbf{M})$ has polynomial growth. Otherwise, we can find two simple closed curves C_1 and C_2 which intersect non-trivially on \mathbf{S} , as follows. If (\mathbf{S}, \mathbf{M}) has more than three features, pick features p_1 through p_4 and pick a separating curve C_1 which has p_1 and p_2 on one side, and p_3 and p_4 on the other, and a separating curve C_2 which has p_2 and p_3 on one side, and p_4 and p_1 on the other. If \mathbf{S} has genus $g > 0$, take C_1 and C_2 which have algebraic intersection number 1 in $H_1(\mathbf{S})$. In either case, the Dehn twists along C_1 and C_2 will not commute. Therefore, by the work of Ivanov [28] and McCarthy [31] on the Tits alternative for mapping class groups, $\mathcal{MCG}(\mathbf{S}, \mathbf{M})$ contains a free group on two generators and hence has exponential growth. \square

It is straightforward to verify the precise growth rate.

Remark 11.3. Curiously, the mutation classes defined by the diagrams in Figure 27 (and thus corresponding to a once-punctured annulus and a pair of pants, respectively) can also be obtained by reorienting some of the edges in the “minimal 2-infinite diagrams” of types $D_n^{(1)}(m, r)$ and $D_n^{(1)}(m, r, s)$, in Seven’s nomenclature [41, §8].

12. Finite mutation classes

It is an important problem of cluster combinatorics to describe all skew-symmetrizable $n \times n$ matrices B whose mutation equivalence class is finite; see, e.g., [3] or [20, §4]. This is obviously the case for $n \leq 2$, and also whenever B defines a cluster algebra of finite type. But there are many more examples, as we will see, even if one restricts (as we do in this paper) to the skew-symmetric case.

A. Buan and I. Reiten used representation theory of quivers to obtain the following non-trivial result.

THEOREM 12.1. ([3, Theorem 3.6]) *Let Γ be a finite directed graph with $n \geq 3$ vertices and no oriented cycles. Then the mutation equivalence class of the corresponding $n \times n$ skew-symmetric matrix $B = B(\Gamma)$ (see Definition 6.3) is finite if and only if Γ is an orientation of a simply-laced Dynkin or extended Dynkin diagram.*

Our main theorems imply the following statement.

COROLLARY 12.2. *The mutation class of any matrix $B(T)$ associated with a triangulation T of a bordered surface with marked points is finite.*

Corollary 12.2 is immediate from the fact that all matrices in such a mutation class have entries ≤ 2 . This is of course a very crude argument that hardly tells us anything about the mutation class in question. The following simple observations shed some light on this issue.

PROPOSITION 12.3. *The mutation class $\mathcal{B}(\mathbf{S}, \mathbf{M})$ consists of the signed adjacency matrices of all ideal triangulations of (\mathbf{S}, \mathbf{M}) (up to simultaneous permutations of rows and columns).*

Proposition 12.3 is immediate from Definition 9.6. It can be used to give an alternative proof of the mutation-finiteness property: Even though the number of ideal triangulations is typically infinite (cf. Proposition 2.3), the number of orbits of the mapping class group action is always finite. Since the matrix $B(T)$ is invariant under this action, the claim follows.

Yet another, completely combinatorial proof of mutation-finiteness can be given using block decompositions discussed in §13.

A skew-symmetric integer matrix B is *acyclic* if the rows and columns of B can be simultaneously permuted so that all entries above the diagonal are non-negative and all entries below the diagonal are non-positive. Equivalently, $B = B(\Gamma)$, where Γ is an acyclic directed graph (as in Theorem 12.1).

COROLLARY 12.4. *The mutation class $\mathcal{B}(\mathbf{S}, \mathbf{M})$ contains an acyclic matrix if and only if $\mathcal{B}(\mathbf{S}, \mathbf{M})$ is of finite or affine type.*

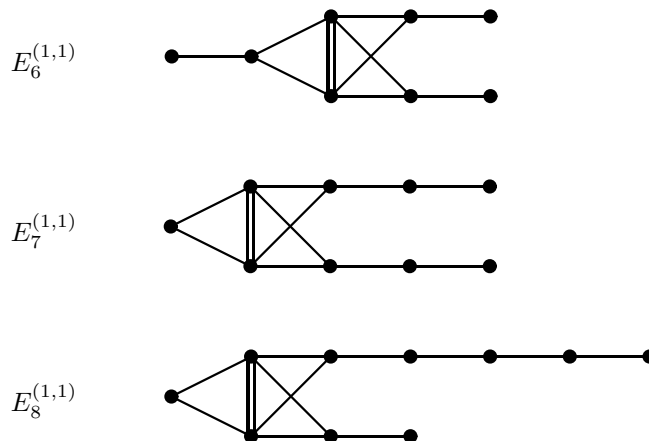


Figure 28. Extended affine diagrams $E_6^{(1,1)}$, $E_7^{(1,1)}$ and $E_8^{(1,1)}$. All triangles are oriented.

Proof. This follows from Corollary 12.2 and Theorem 12.1. \square

Problem 12.5. (Recognizing finite mutation classes) Find an effective criterion/algorithm for determining whether a given skew-symmetric (or, more generally, skew-symmetrizable) integer matrix defines a finite mutation class.

Theorem 12.1 and Corollary 12.2 provide two sources of finite mutation classes: the simply-laced Dynkin or extended Dynkin diagrams, and the signed adjacency matrices of triangulations. We next discuss three additional examples of this kind.

Definition 12.6. A directed graph (quiver) Γ (or the corresponding skew-symmetric matrix $B=B(\Gamma)$) is of *extended affine type* $E_6^{(1,1)}$, $E_7^{(1,1)}$ or $E_8^{(1,1)}$ if Γ is an orientation of the corresponding graph in Figure 28 in which all triangles are oriented. More broadly, this nomenclature applies to the mutation equivalence class of Γ or B , or any quiver therein.

The quivers for the extended affine exceptional types shown in Figure 28 are orientations of the Dynkin diagrams for extended affine root systems first described by Saito [39, Table 1]. The connection between cluster combinatorics and extended affine root systems was first noticed by Geiss, Leclerc and Schröer [22].

Neither of these three mutation classes is associated with triangulated surfaces. This can be shown using block decompositions of §13.

PROPOSITION 12.7. *The mutation equivalence class of a quiver (or matrix) of extended affine type $E_6^{(1,1)}$, $E_7^{(1,1)}$ or $E_8^{(1,1)}$ is finite.*

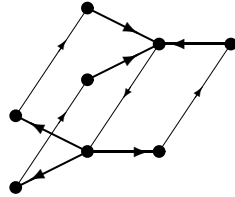


Figure 29. $A_2 \times D_4$ diagram.

Proof. This statement was verified by exhaustive computer search with the help of Java applets for matrix mutations written by B. Keller [30] and L. Williams. \square

Remark 12.8. Dynkin diagrams for extended affine root systems of types A and D do not generally define a finite mutation class (for any orientation). On the other hand, various subdiagrams of these diagrams do come from surfaces. One nice example is a cluster algebra associated with a 4-punctured sphere, which has extended affine type $D_4^{(1,1)}$.

Remark 12.9. The extended affine types $E_6^{(1,1)}$, $E_7^{(1,1)}$ and $E_8^{(1,1)}$ can be alternatively described using a construction of a “direct product” of Dynkin quivers (discovered independently by F. Chapoton [5], who also posed a version of Problem 12.10 [4]). The definition of this construction should be clear from the example of the quiver $A_2 \times D_4$ shown in Figure 29; this quiver is mutation-equivalent to the extended affine quiver of type $E_6^{(1,1)}$. Similarly, $A_3 \times A_3$ yields the extended affine type $E_7^{(1,1)}$ (the cluster type of $\text{Gr}_{4,8}$), while $A_2 \times A_5$ gives the extended affine type $E_8^{(1,1)}$ (the cluster type of $\text{Gr}_{3,9}$).

Problem 12.10. Are there finitely many finite mutation classes of indecomposable skew-symmetric $n \times n$ matrices, $n \geq 3$, which are not associated with triangulated surfaces? Are there any at all, besides the finite exceptional types E_6 , E_7 and E_8 , the corresponding affine types \tilde{E}_6 , \tilde{E}_7 and \tilde{E}_8 , and the extended affine types $E_6^{(1,1)}$, $E_7^{(1,1)}$ and $E_8^{(1,1)}$?

Note added in proof. Two additional finite mutation classes have been recently discovered in [8].

As an application, we classify the Grassmannians of “finite mutation type.” Recall that a cluster algebra structure on the homogeneous coordinate ring $\mathbb{C}[X]$ of a Grassmann manifold $X = \text{Gr}_{k,k+\ell}$ (with respect to its Plücker embedding) was described by J. Scott [40], who showed, in the notation of Remark 12.9, that $\mathbb{C}[X]$ has cluster type $A_{k-1} \times A_{\ell-1}$.

PROPOSITION 12.11. *Let \mathcal{A} be the canonical cluster algebra structure on the homogeneous coordinate ring $\mathbb{C}[X]$ of a Grassmannian $\text{Gr}_{k,k+\ell}$. The mutation equivalence class $\mathcal{B}(\mathcal{A})$ is finite if and only if $(k-2)(\ell-2) \leq 4$.*

Grassmannian X	Cluster type of $\mathbb{C}[X]$
$\text{Gr}_{2,n+3} \cong \text{Gr}_{n+1,n+3}$	A_n
$\text{Gr}_{3,6}$	D_4
$\text{Gr}_{3,7} \cong \text{Gr}_{4,7}$	E_6
$\text{Gr}_{3,8} \cong \text{Gr}_{5,8}$	E_8
$\text{Gr}_{3,9} \cong \text{Gr}_{6,9}$	$E_8^{(1,1)}$
$\text{Gr}_{4,8}$	$E_7^{(1,1)}$

Table 3. Grassmannians of finite mutation type.

A more concrete version of Proposition 12.11 is presented in Table 3.

Proof. The “only if” part is proved by checking that in the special cases $\{k, \ell\} = \{3, 7\}$ and $\{k, \ell\} = \{4, 5\}$, the mutation class $\mathcal{B}(\mathcal{A})$ is infinite. This is done by identifying, with the help of a computer search, a sequence of mutations producing a graph Γ' some of whose edges have multiplicity ≥ 3 . (Any such graph with at least three vertices is easily seen to generate an infinite mutation class.)

For the “if” part, the corresponding Grassmannians $\text{Gr}_{k,k+\ell}$ are listed in Table 3 on the left. It remains to show that in each of these cases, the quiver $A_{k-1} \times A_{\ell-1}$ has cluster type given in the right column. This is done by exhibiting an appropriate sequence of mutations (which can be found by hand or with the help of a computer). \square

We note that the strict inequality $(k-2)(\ell-2) < 4$ corresponds precisely to the situations where $\mathbb{C}[\text{Gr}_{k,k+\ell}]$ has finite cluster type.

Problem 12.12. Provide a concrete combinatorial description for cluster complexes of extended affine types $E_6^{(1,1)}$, $E_7^{(1,1)}$ and $E_8^{(1,1)}$.

Remark 12.13. It is not hard to show that the coordinate ring $\mathbb{C}[\text{SL}_4]$ (viewing SL_4 as the hypersurface $\det=1$ in \mathbb{C}^{16} ; cf. Example 6.9) has cluster type $E_7^{(1,1)}$, with respect to its usual cluster algebra structure closely related to the one described in [1] for an open double Bruhat cell. Similarly, the coordinate ring $\mathbb{C}[\text{SL}_6/N]$ (here N is the subgroup of unipotent upper-triangular matrices in SL_6) has cluster type $E_8^{(1,1)}$. Thus, a solution of Problem 12.12 is likely to lead to a better understanding of the cluster combinatorics underlying the special linear group $\text{SL}_4(\mathbb{C})$ and the affine base space $\text{SL}_6(\mathbb{C})$, and the corresponding dual canonical bases. The cluster algebras $\mathbb{C}[\text{SL}_n]$, $n \geq 5$, and $\mathbb{C}[\text{SL}_n/N]$, $n \geq 7$, are of infinite mutation type (“wild”), and therefore are unlikely to be understood in sufficient detail.

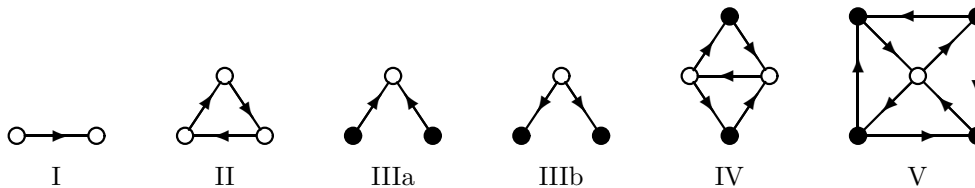


Figure 30. Blocks of types I–V.

13. Block decompositions

In view of Proposition 12.3, the following classes of skew-symmetric integer matrices coincide:

- exchange matrices of cluster algebras associated with bordered surfaces with marked points;
- signed adjacency matrices of ideal triangulations;
- signed adjacency matrices of tagged triangulations.

In this section, we describe, in concrete combinatorial terms, the set of matrices satisfying any of these equivalent descriptions.

Definition 13.1. A *block* is a directed graph isomorphic to one of the graphs shown in Figure 30. Depending on which graph it is, we call it a block of *type* I, II, III, IV or V. The vertices marked by unfilled circles in Figure 30 are called *outlets*. A directed graph Γ is called *block-decomposable* if it can be obtained from a collection of disjoint blocks by the following procedure. Take a partial matching of the combined set of outlets; matching an outlet to itself or to another outlet from the same block is not allowed. Identify (or “glue”) the vertices within each pair of the matching. We require that the resulting graph Γ' be connected. If Γ' contains a pair of edges connecting the same pair of vertices but going in opposite directions, then remove each such pair of edges. The result is a block-decomposable graph Γ .

By design, a block-decomposable graph has no loops, and all edge multiplicities are 1 or 2.

Example 13.2. Here are some examples illustrating Definition 13.1. Two blocks of type I can be glued to create any of the following block-decomposable graphs:

- two isolated vertices;
- a two-vertex graph with two edges of the same orientation;
- an arbitrary orientation of a type A_3 Coxeter–Dynkin diagram.

THEOREM 13.3. *An integer matrix B is a signed adjacency matrix of an ideal triangulation of a bordered surface with marked points if and only if B is encoded by a block-decomposable graph Γ (i.e., $B=B(\Gamma)$ in the sense of Definition 6.3).*

Proof. Let $B=B(T)$ be the signed adjacency matrix of an ideal triangulation T . As explained in Remark 4.2, with one exception the matrix $B(T)$ can be computed by taking a puzzle-piece decomposition of T and adding up the contributions of the puzzle pieces that make up T . It is routine to verify that this leads to a block decomposition of the graph Γ associated with B , as follows:

- the puzzle pieces shown in Figure 9 whose outer sides do not lie on the boundary of T (that is, are matched to outer sides of other puzzle pieces) contribute blocks of types II, IV and V, respectively;
- a triangle with two sides on the boundary does not contribute to the block decomposition;
- a triangle with one side on the boundary contributes a block of type I;
- a punctured-digon puzzle piece with both sides on the boundary yields two disjoint vertices (see Example 13.2);
- a punctured-digon puzzle piece with one side on the boundary contributes a block of type III;
- a twice-punctured monogon puzzle piece with the side on the boundary yields a graph of type $\tilde{A}(2, 2)$ which can be obtained by gluing four blocks of type I.

We also need to consider the case of the exceptional triangulation of a 4-punctured sphere that cannot be decomposed into puzzle pieces. The corresponding 6×6 matrix $B(T)$ is encoded by a 6-vertex graph Γ which is an orientation of the 1-skeleton of an octahedron. This graph Γ can be obtained by gluing four blocks of type II.

It remains to prove that for each block-decomposable graph Γ , we have $B(\Gamma)=B(T)$ for some tagged triangulation T . This is done by reversing the argument in the previous paragraph: each block in Figure 30 corresponds to (a piece of) a triangulated bordered surface in which the signed adjacencies are described by the arrows in the block, and whose exposed sides are precisely the outlets. Glue the pieces of the triangulated surface according to the matching of the outlets. Glue an additional triangle to the side corresponding to any unmatched outlet, leaving the other two sides of this new triangle exposed. \square

Remark 13.4. A block decomposition of a given graph Γ may not be unique. For example, the graph Γ of type $\tilde{A}(2, 2)$ (see Figure 16) can be obtained by gluing either four blocks of type I, or a block of type I and a block of type IV.

Remark 13.5. It can be proved by a direct case-by-case combinatorial argument that the class of block-decomposable graphs is closed under mutations. (This leads to an alternative proof of mutation finiteness.) It would be interesting to find a way to enlarge the set of blocks while keeping this invariance property. This would yield a

better understanding (or even a complete description) of the class of diagrams whose mutation equivalence class is finite; cf. Problem 12.5.

It can be shown by a tedious case-by-case analysis that a graph Γ obtained by arbitrarily orienting the edges of a Coxeter–Dynkin diagram of type E_n , $n=6, 7, 8$, or an affine extension thereof, is *not* block-decomposable. On the other hand, all these graphs have finite mutation classes (either provably or conjecturally); it would be nice to find a conceptual explanation of these phenomena.

14. Recovering topology from combinatorial data

Let B be a skew-symmetric matrix satisfying the criterion in Theorem 13.3, so that there exists a (non-unique) bordered surface (\mathbf{S}, \mathbf{M}) and a (non-unique) ideal (equivalently, tagged) triangulation T of (\mathbf{S}, \mathbf{M}) such that $B=B(T)$. How much information about (\mathbf{S}, \mathbf{M}) can we recover from B ?

Examples 4.3 and 4.5 show that the topology of \mathbf{S} is not in general determined by $B(T)$, and even if it is, the number of punctures is not uniquely determined.

Furthermore, there exist topologically inequivalent triangulations T_1 and T_2 of the same marked surface (\mathbf{S}, \mathbf{M}) for which $B(T_1)=B(T_2)$. Example: let (\mathbf{S}, \mathbf{M}) be a quadrilateral $abcd$ with a puncture t (type D_4). Triangulation T_1 is obtained by connecting t to a, b, c and d , while T_2 is obtained by connecting t to a and c , and connecting a with c by two arcs. Additional examples can be found by considering a twice-punctured digon (type \tilde{D}_4) or a sphere with four punctures (mentioned in the proof of Theorem 13.3).

On the positive side, we have the following result.

PROPOSITION 14.1. *Let T be an ideal triangulation of (\mathbf{S}, \mathbf{M}) such that $B(T)=B(\Gamma)$, where the graph Γ has a unique block decomposition. Then the topology of (\mathbf{S}, \mathbf{M}) is uniquely determined by $B(T)$. More precisely, for any $(\mathbf{S}', \mathbf{M}') \not\cong (\mathbf{S}, \mathbf{M})$ and any ideal triangulation T' of $(\mathbf{S}', \mathbf{M}')$, the matrix $B(T')$ is mutation-inequivalent to $B(T)$.*

Proof. It is enough to show that $B(T') \neq B(T)$. Suppose not. By Remark 4.2, every triangulation T has a puzzle-piece decomposition, except for one triangulation of a sphere with four punctures. By the proof of Theorem 13.3, every puzzle-piece decomposition gives a block decomposition, except for one triangulation of a twice-punctured monogon. In both exceptional cases the graph Γ associated with T has more than one block decomposition. As a result, if the hypotheses of the theorem are satisfied, the triangulations T and T' can be decomposed into puzzle pieces leading to block decompositions of Γ , which must coincide by hypothesis. By the argument at the end of the proof of Theorem 13.3,

the puzzle pieces of T (or T') and the way they are glued to each other are uniquely determined by the block decomposition of Γ . Thus $T=T'$ as desired. \square

Proof of Lemma 6.8. ($\tilde{A}(n_1, n_2) \neq \tilde{A}(n'_1, n'_2)$.) If $n_1=n_2=1$ or $\{n_1, n_2\}=\{1, 2\}$, then there is nothing to prove. Otherwise we may suppose that $n_1 \geq 3$ (possibly switching $\{n_1, n_2\}$ and $\{n'_1, n'_2\}$ and/or n_1 and n_2). Let Γ be the graph of type $\tilde{A}(n_1, n_2)$ obtained by orienting a cycle of length n_1+n_2 so that n_1 consecutive edges are oriented the same way. A straightforward inspection shows that Γ has a unique block decomposition, and the lemma follows by Proposition 14.1. \square

Suppose that $(\mathbf{S}, \mathbf{M}) \mapsto f((\mathbf{S}, \mathbf{M}))$ is a surface invariant that can be determined from the signed adjacency matrix $B(T)$ of an ideal triangulation of (\mathbf{S}, \mathbf{M}) ; that is, there exists a function $B \mapsto F(B)$ such that $f((\mathbf{S}, \mathbf{M}))=F(B(T))$. Then $F(B)$ must be invariant under matrix mutations. Consequently, the problem at hand is closely related to the following fundamental problem in cluster combinatorics.

Problem 14.2. (Mutation invariants) Describe functions on skew-symmetric (or skew-symmetrizable) matrices that are invariant under matrix mutations. Use such invariants to develop efficient algorithms that distinguish between mutation-inequivalent matrices.

Ideally, one would like to have a complete system of invariants for matrices of given order, so that if B and B' are not mutation-equivalent, then one of these invariants takes different values at B and B' .

Very few mutation invariants have been identified so far. The best known example is the (ordinary) *rank* of a matrix; see [23, Lemma 1.2] and [1, Lemma 3.2]. Note that for a skew-symmetric matrix B , $\text{rank}(B)$ is even [38, §21]. The following theorem is stated in terms of the *corank*, the dimension n minus the rank.

THEOREM 14.3. *For an arbitrary ideal triangulation T of a marked surface (\mathbf{S}, \mathbf{M}) , the corank of $B(T)$ equals the number of punctures plus the number of boundary components having an even number of marked points.*

The statement of Theorem 14.3 becomes more elegant if one follows Fock and Goncharov [11] and treats a puncture as a boundary component with no marked points; then the corank is just the number of boundary components with an even number of marked points (possibly none).

Proof. We argue by joint induction on the size n of the matrix and the number of boundary components. Since $\text{rank}(B)$ is invariant under mutations (flips), we are free to choose the triangulation T of (\mathbf{S}, \mathbf{M}) as we please.

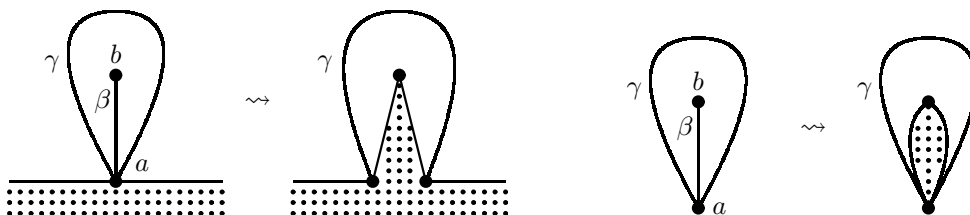


Figure 31. Making a cut inside a self-folded triangle.

The following lemma lets us trade one puncture for two marked points on the boundary without changing $\text{rank}(B)$.

LEMMA 14.4. *Let (\mathbf{S}, \mathbf{M}) be a surface with a puncture b and another marked point a . Form a surface $(\mathbf{S}', \mathbf{M}')$ by taking (\mathbf{S}, \mathbf{M}) and deleting the puncture b and either:*

- *if a is on the boundary of \mathbf{S} , add two marked points to the boundary component containing a ; or*
- *if a is a puncture, replace a by a boundary component with two marked points.*

Then for any triangulations T of (\mathbf{S}, \mathbf{M}) and T' of $(\mathbf{S}', \mathbf{M}')$, the rank of $B(T)$ equals the rank of $B(T')$.

Proof. Construct a triangulation T of (\mathbf{S}, \mathbf{M}) that includes an arc β connecting a to b and a loop γ encircling β and based at a , making a self-folded triangle (see Figure 31). By Definition 4.1, the two rows (resp. columns) corresponding to β and γ in $B(T)$ are identical. Thus the rank is unchanged (and the corank goes down by 1) if we delete the row and column corresponding to β from $B(T)$. The new matrix is the same as $B(T')$ for a triangulation T' of $(\mathbf{S}', \mathbf{M}')$ obtained by cutting T open along β . \square

If there is a puncture and at least one other marked point, we may remove all the punctures by repeated applications of Lemma 14.4. At each application, we leave $\text{rank}(B)$ unchanged and decrease n by 1, thus decreasing the corank by 1, as desired. Thus, we may henceforth assume that (\mathbf{S}, \mathbf{M}) has no punctures, or else is closed with a single puncture.

Suppose that there is more than one boundary component. If a boundary component has at least three marked points, we can reverse the transformation in Lemma 14.4 to temporarily trade two of them for a puncture in the interior (increasing n by 1 and not changing $\text{rank}(B)$), and then trade the puncture for two marked points on another boundary component, again without changing $\text{rank}(B)$. Thus we can move pairs of marked points from one boundary component to another without changing either n or the rank, or the number of even boundary components.

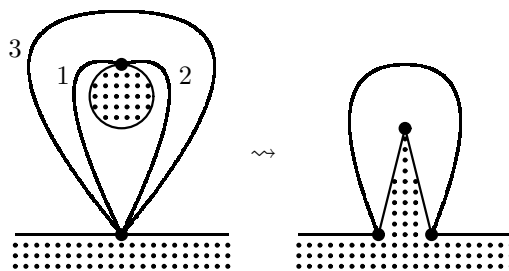


Figure 32. Removing a connected component with one marked point.

If a boundary component has two marked points, we can again reverse Lemma 14.4 to trade it for two punctures, and then trade them back for four marked points on another boundary component. As a result, we decrease n by 1, leave $\text{rank}(B)$ unchanged, and remove one component with an even number of punctures.

So we only need to treat the cases where at most one boundary component has more than one marked point. If there are at least two boundary components, then take a component with one marked point, and construct a triangulation T that includes the arcs shown in Figure 32 on the left.

The signed adjacency matrix $B(T)$ has the form

$$\begin{pmatrix} 0 & -2 & 1 & 0 & 0 & \dots \\ 2 & 0 & -1 & 0 & 0 & \dots \\ -1 & 1 & 0 & * & * & \dots \\ 0 & 0 & * & * & * & \dots \\ 0 & 0 & * & * & * & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

Adding the half-sum of columns 1 and 2 to column 3, and doing the same for the rows, we obtain the matrix

$$\begin{pmatrix} 0 & -2 & 0 & 0 & 0 & \dots \\ 2 & 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & * & * & \dots \\ 0 & 0 & * & * & * & \dots \\ 0 & 0 & * & * & * & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}.$$

This matrix is a block sum of two matrices. One is of dimension 2 and rank 2, while the other is $B(T')$, where T' is the triangulation shown on the right of Figure 32. Thus

we can remove this (odd) boundary component, decrease n by 2, and add two points to another boundary component without changing the corank, as desired.

We are left with the cases where (\mathbf{S}, \mathbf{M}) is a surface of genus $g \geq 0$ with either one boundary component or one puncture. If $g > 0$ and (\mathbf{S}, \mathbf{M}) is not the once-punctured torus, we can encircle one handle of \mathbf{S} with an arc based at a puncture and triangulate the enclosed torus with four additional arcs, forming a triangulation T . Again by row and column operations we can reduce the matrix to a block-diagonal form, with one block of dimension 4 and rank 4, and the other block of the form $B(T')$, where T' is a triangulation of a surface of genus $g-1$ and two more marked points on the boundary. In particular, the corank is unchanged.

In this way, we can reduce everything to the base cases: a once-punctured torus and an unpunctured c -gon, $c > 3$. In these cases, the theorem is verified by direct calculation: for the once-punctured torus, $\text{corank}(B) = 1$, while for the c -gon, $\text{corank}(B) = 0$ if c is odd and $\text{corank}(B) = 1$ if c is even. \square

COROLLARY 14.5. *If \mathbf{S} is closed (so that \mathbf{M} consists exclusively of punctures), then its genus g and the number of punctures p are determined by the size n and the rank r of the matrix $B(T)$ associated with any ideal triangulation T of (\mathbf{S}, \mathbf{M}) :*

$$g = \frac{3r - 2n + 6}{6} \quad \text{and} \quad p = n - r.$$

Proof. This is immediate from Theorem 14.3. \square

Acknowledgments

The authors thank Vladimir Fock, Alexander Goncharov, Nikolai Ivanov, Robert Penner, David Speyer, Lauren Williams and Andrei Zelevinsky for helpful advice and stimulating discussions. We thank Bernhard Keller and Lauren Williams for providing software that computes matrix mutations, and Nathan Reading for letting us borrow Figures 3 and 6 from [15].

We are grateful to the anonymous referee and to Daniel Labardini Fragoso for their thorough readings of the original submitted version, and for a number of valuable editorial suggestions whose implementation improved the quality of the paper.

S. F. acknowledges the hospitality of the Mittag-Leffler Institute in April–June, 2005; some of our main results were first presented there. Much of the work was completed while D. T. was a Benjamin Peirce Assistant Professor at Harvard University.

References

- [1] BERENSTEIN, A., FOMIN, S. & ZELEVINSKY, A., Cluster algebras. III. Upper bounds and double Bruhat cells. *Duke Math. J.*, 126 (2005), 1–52.
- [2] BJÖRNER, A., Topological methods, in *Handbook of Combinatorics*, Vol. 2, pp. 1819–1872. Elsevier, Amsterdam, 1995.
- [3] BUAN, A. B. & REITEN, I., Acyclic quivers of finite mutation type. *Int. Math. Res. Not.*, (2006), Art. ID 12804.
- [4] CHAPOTON, F., Lectures at the summer school “Cluster Algebras”. Luminy, May 2005.
- [5] — Rogers dilogarithm and y -systems, Lecture given at the workshop “Teichmüller Theory and Cluster Algebras”. Leicester, May 2006.
- [6] CHEKHOV, L. O. & PENNER, R. C., Introduction to Thurston’s quantum theory. *Uspekhi Mat. Nauk*, 58 (2003), 93–138 (Russian); English translation in *Russian Math. Surveys*, 58 (2003), 1141–1183.
- [7] — On quantizing Teichmüller and Thurston theories, in *Handbook of Teichmüller Theory*. Vol. I, IRMA Lect. Math. Theor. Phys., 11, pp. 579–645. Eur. Math. Soc., Zürich, 2007.
- [8] DERKSEN, H. & OWEN, T., New graphs of finite mutation type. Preprint, 2008. [arXiv:0804.0787](https://arxiv.org/abs/0804.0787).
- [9] FOCK, V. V., Dual Teichmüller spaces. Preprint, 1997. [arXiv:dg-ga/9702018](https://arxiv.org/abs/dg-ga/9702018).
- [10] FOCK, V. V. & GONCHAROV, A. B., Moduli spaces of local systems and higher Teichmüller theory. *Publ. Math. Inst. Hautes Études Sci.*, (2006), 1–211.
- [11] — Dual Teichmüller and lamination spaces, in *Handbook of Teichmüller Theory*. Vol. I, IRMA Lect. Math. Theor. Phys., 11, pp. 647–684. Eur. Math. Soc., Zürich, 2007.
- [12] — Cluster ensembles, quantization and the dilogarithm. Preprint, 2003. [arXiv:math.AG/0311245](https://arxiv.org/abs/math/0311245).
- [13] FOCK, V. V. & ROSLY, A. A., Poisson structure on moduli of flat connections on Riemann surfaces and the r -matrix, in *Moscow Seminar in Mathematical Physics*, Amer. Math. Soc. Transl. Ser. 2, 191, pp. 67–86. Amer. Math. Soc., Providence, RI, 1999.
- [14] FOMIN, S., & THURSTON, D., Cluster algebras and triangulated surfaces. Part II: Lambda lengths. Preprint, 2008.
- [15] FOMIN, S. & READING, N., Root systems and generalized associahedra, in *Geometric Combinatorics* (Park City, UT, 2003), IAS/Park City Math. Ser., 14, pp. 63–131. Amer. Math. Soc., Providence, RI, 2007.
- [16] FOMIN, S. & ZELEVINSKY, A., Cluster algebras. I. Foundations. *J. Amer. Math. Soc.*, 15 (2002), 497–529.
- [17] — Cluster algebras. II. Finite type classification. *Invent. Math.*, 154 (2003), 63–121.
- [18] — Cluster algebras: notes for the CDM-03 conference, in *Current Developments in Mathematics*, 2003, pp. 1–34. Int. Press, Somerville, MA, 2003.
- [19] — Y -systems and generalized associahedra. *Ann. of Math.*, 158 (2003), 977–1018.
- [20] — Cluster algebras. IV. Coefficients. *Compos. Math.*, 143 (2007), 112–164.
- [21] FREEDMAN, M., HASS, J. & SCOTT, P., Closed geodesics on surfaces. *Bull. London Math. Soc.*, 14 (1982), 385–391.
- [22] GEISS, C., LECLERC, B. & SCHRÖER, J., Semicanonical bases and preprojective algebras. *Ann. Sci. École Norm. Sup.*, 38 (2005), 193–253.
- [23] GEKHTMAN, M., SHAPIRO, M. & VAINSHTEIN, A., Cluster algebras and Poisson geometry. *Mosc. Math. J.*, 3 (2003), 899–934, 1199.
- [24] — Cluster algebras and Weil–Petersson forms. *Duke Math. J.*, 127 (2005), 291–311.
- [25] HARER, J. L., Stability of the homology of the mapping class groups of orientable surfaces. *Ann. of Math.*, 121 (1985), 215–249.

- [26] — The virtual cohomological dimension of the mapping class group of an orientable surface. *Invent. Math.*, 84 (1986), 157–176.
- [27] HATCHER, A., On triangulations of surfaces. *Topology Appl.*, 40 (1991), 189–194.
- [28] IVANOV, N. V., Algebraic properties of the Teichmüller modular group. *Dokl. Akad. Nauk SSSR*, 275 (1984), 786–789 (Russian); English translation in *Soviet Math. Dokl.*, 29 (1984), 288–291.
- [29] — Mapping class groups, in *Handbook of Geometric Topology*, pp. 523–633. North-Holland, Amsterdam, 2002.
- [30] KELLER, B., Quiver mutation in java.
<http://www.institut.math.jussieu.fr/quivermutation>.
- [31] MCCARTHY, J., A “Tits-alternative” for subgroups of surface mapping class groups. *Trans. Amer. Math. Soc.*, 291 (1985), 583–612.
- [32] MOSHER, L., Tiling the projective foliation space of a punctured surface. *Trans. Amer. Math. Soc.*, 306 (1988), 1–70.
- [33] PENNER, R. C., The decorated Teichmüller space of punctured surfaces. *Comm. Math. Phys.*, 113 (1987), 299–339.
- [34] — Universal constructions in Teichmüller theory. *Adv. Math.*, 98 (1993), 143–215.
- [35] — Decorated Teichmüller theory of bordered surfaces. *Comm. Anal. Geom.*, 12 (2004), 793–820.
- [36] — Probing mapping class groups using arcs, in *Problems on Mapping Class Groups and Related Topics*, Proc. Sympos. Pure Math., 74, pp. 97–114. Amer. Math. Soc., Providence, RI, 2006.
- [37] — The structure and singularities of quotient arc complexes. *J. Topol.*, 1 (2008), 527–550.
- [38] PRASOLOV, V. V., *Problems and Theorems in Linear Algebra*. Translations of Mathematical Monographs, 134. Amer. Math. Soc., Providence, RI, 1994.
- [39] SAITO, K., Extended affine root systems. I. Coxeter transformations. *Publ. Res. Inst. Math. Sci.*, 21 (1985), 75–179.
- [40] SCOTT, J. S., Grassmannians and cluster algebras. *Proc. London Math. Soc.*, 92 (2006), 345–380.
- [41] SEVEN, A. I., Recognizing cluster algebras of finite type. *Electron. J. Combin.*, 14 (2007), Research Paper 3, 35 pp.
- [42] TESCHNER, J., An analog of a modular functor from quantized Teichmüller theory, in *Handbook of Teichmüller Theory*. Vol. I, IRMA Lect. Math. Theor. Phys., 11, pp. 685–760. Eur. Math. Soc., Zürich, 2007

SERGEY FOMIN
Department of Mathematics
University of Michigan
Ann Arbor, MI 48109
U.S.A.
fomin@umich.edu

MICHAEL SHAPIRO
Department of Mathematics
Michigan State University
East Lansing, MI 48824
U.S.A.
mshapiro@math.msu.edu

DYLAN THURSTON
Department of Mathematics
Barnard College, Columbia University
New York, NY 10027
U.S.A.
dpt@math.columbia.edu

Received October 2, 2006