

Introduction to Cluster Algebras

Chapters 1–7

(preliminary version)

SERGEY FOMIN

LAUREN WILLIAMS

ANDREI ZELEVINSKY

Preface

This is a preliminary draft of Chapters 1–3 of our forthcoming textbook *Introduction to cluster algebras*, joint with Andrei Zelevinsky (1953–2013). Other chapters have been posted as

- [arXiv:1707.07190](https://arxiv.org/abs/1707.07190) (Chapters 4–5),
- [arXiv:2008.09189](https://arxiv.org/abs/2008.09189) (Chapter 6), and
- [arXiv:2106.02160](https://arxiv.org/abs/2106.02160) (Chapter 7).

We expect to post additional chapters in the not so distant future.

This book grew from the ten lectures given by Andrei at the NSF CBMS conference on Cluster Algebras and Applications at North Carolina State University in June 2006. The material of his lectures is much expanded but we still follow the original plan aimed at giving an accessible introduction to the subject for a general mathematical audience.

Since its inception in [23], the theory of cluster algebras has been actively developed in many directions. We do not attempt to give a comprehensive treatment of the many connections and applications of this young theory. Our choice of topics reflects our personal taste; much of the book is based on the work done by Andrei and ourselves.

Comments and suggestions are welcome.

Sergey Fomin
Lauren Williams

Partially supported by the NSF grants DMS-1049513, DMS-1361789, and DMS-1664722.

2020 *Mathematics Subject Classification*. Primary 13F60.

© 2016–2021 by Sergey Fomin, Lauren Williams, and Andrei Zelevinsky

Contents

Chapter 1. Total positivity	1
§1.1. Totally positive matrices	1
§1.2. The Grassmannian of 2-planes in m -space	4
§1.3. The basic affine space	8
§1.4. The general linear group	13
Chapter 2. Mutations of quivers and matrices	19
§2.1. Quiver mutation	19
§2.2. Triangulations of polygons	21
§2.3. Wiring diagrams	22
§2.4. Double wiring diagrams	24
§2.5. Urban renewal	25
§2.6. Mutation equivalence	27
§2.7. Matrix mutation	30
§2.8. Invariants of matrix mutations	33
Chapter 3. Clusters and seeds	35
§3.1. Basic definitions	35
§3.2. Examples of rank 1 and 2	39
§3.3. Laurent phenomenon	44
§3.4. Connections to number theory	50
§3.5. Y -patterns	54
§3.6. Tropical semifields	62
Bibliography	67

Total positivity

Total positivity, along with G. Lusztig’s theory of canonical bases, was one of the main motivations for the development of cluster algebras. In this chapter, we present the basic notions of total positivity, focusing on three important examples (to be re-examined again in the future): square matrices, Grassmannians, and basic affine spaces. As our main goal here is to provide motivation rather than develop a rigorous theory, the exposition is somewhat informal. Additional details and references can be found in [18].

1.1. Totally positive matrices

An $n \times n$ matrix with real entries is called *totally positive* (TP for short) if all its minors—that is, determinants of square submatrices—are positive. A real matrix is called *totally nonnegative* (or TNN) if all its minors are nonnegative. The first systematic study of these classes of matrices was conducted in the 1930s by F. Gantmacher and M. Krein [28], following the pioneering work of I. Schoenberg [45]. In particular, they showed that the eigenvalues of an $n \times n$ totally positive matrix are real, positive, and distinct.

Total positivity is a remarkably widespread phenomenon: TP and TNN matrices play an important role, *inter alia*, in classical mechanics, probability, discrete potential theory, asymptotic representation theory, algebraic and enumerative combinatorics, and the theory of integrable systems.

The *Binet-Cauchy Theorem* implies that TP (resp., TNN) matrices in $G = \mathrm{SL}_n$ are closed under multiplication; they thus form a multiplicative semigroup, denoted by $G_{>0}$ (resp., $G_{\geq 0}$). The following “splitting lemma” due to C. Cryer [12, 13] shows that the study of $G_{\geq 0}$ can be reduced to the investigation of its subsemigroup of upper-triangular unipotent TNN matrices, i.e. upper-triangular TNN matrices with 1’s on the diagonal:

Lemma 1.1.1. *A matrix $z \in \mathrm{SL}_n$ is totally nonnegative if and only if it has a Gaussian decomposition*

$$z = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ * & 1 & 0 & \cdots & 0 \\ * & * & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ * & * & * & \cdots & 1 \end{bmatrix} \begin{bmatrix} * & 0 & 0 & \cdots & 0 \\ 0 & * & 0 & \cdots & 0 \\ 0 & 0 & * & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & * \end{bmatrix} \begin{bmatrix} 1 & * & * & \cdots & * \\ 0 & 1 & * & \cdots & * \\ 0 & 0 & 1 & \cdots & * \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{bmatrix}$$

in which all three factors (lower-triangular unipotent, diagonal, and upper-triangular unipotent) are totally nonnegative.

There is also a counterpart of this statement for totally positive matrices.

The *Loewner-Whitney Theorem* [36, 52] identifies the infinitesimal generators of $G_{\geq 0}$ as the *Chevalley generators* of the corresponding Lie algebra. In pedestrian terms, each $n \times n$ TNN matrix can be written as a product of matrices of the form $x_i(t)$, $y_i(t)$, and $z_i(t)$, where each parameter t is positive, the matrices $x_i(t)$ are defined by

$$x_i(t) = \begin{bmatrix} 1 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 1 & t & \cdots & 0 \\ 0 & \cdots & 0 & 1 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & 0 & \cdots & 1 \end{bmatrix},$$

where the t is in row i of $x_i(t)$, $y_i(t)$ is the transpose of $x_i(t)$, and $z_i(t)$ is the diagonal matrix with diagonal entries $(1, \dots, 1, t, t^{-1}, 1, \dots, 1)$ where t and t^{-1} are in positions i and $i+1$. This led G. Lusztig [37] to the idea of extending the notion of total positivity to other semisimple Lie groups G , by defining the set $G_{\geq 0}$ of TNN elements in G as the semigroup generated by the Chevalley generators. Lusztig proved that $G_{\geq 0}$ is a semialgebraic subset of G , and described it by inequalities of the form $\Delta(x) \geq 0$ where Δ lies in the appropriate *dual canonical basis*; see [38, Section 5]. A simpler description in terms of *generalized minors* [20] was given in [22].

A yet more general (if informal) concept is one of a *totally positive* (or *totally nonnegative*) (sub)variety of a given complex algebraic variety Z . Vaguely, the idea is this: suppose that Z comes equipped with a family $\mathbf{\Delta}$ of “important” regular functions on Z . The corresponding TP (resp., TNN) variety $Z_{>0}$ (resp., $Z_{\geq 0}$) is the set of points at which all of these functions take positive (resp., nonnegative) real values:

$$Z_{>0} = \{z \in Z : \Delta(z) > 0 \text{ for all } \Delta \in \mathbf{\Delta}\}.$$

If Z is the affine space of $n \times n$ matrices (or $Z = \mathrm{GL}_n(\mathbb{C})$, or $Z = \mathrm{SL}_n(\mathbb{C})$), and Δ is the set of all minors, then we recover the classical notions. One can restrict this construction to matrices lying in a given stratum of a Bruhat decomposition, or in a given *double Bruhat cell* [20, 37]. Another important example is the *totally positive (resp., nonnegative) Grassmannian* consisting of the points in a usual Grassmann manifold where all Plücker coordinates can be chosen to be positive (resp., nonnegative). This construction can be extended to arbitrary partial flag manifolds, and more generally to homogeneous spaces G/P associated to semisimple complex Lie groups.

We note that in each of the examples alluded to above, the notion of positivity depends on a particular choice of a coordinate system: a basis in a vector space allows us to view linear transformations as matrices, determines a system of Plücker coordinates, etc.

Among many questions which one may ask about totally positive/nonnegative varieties $Z_{>0}$ and $Z_{\geq 0}$, let us restrict our attention to the problem of *efficient TP testing*: how many inequalities (and which ones) does one need to check in order to ascertain that a given point in Z is totally positive? In particular, are there efficient ways for testing a very large matrix for total positivity? (It is not hard to see that an $n \times n$ matrix has altogether $\binom{2n}{n} - 1$ minors, a number which grows exponentially in n .) Examples 1.1.2 and 1.1.3 provide a glimpse into the tricks used to construct efficient TP criteria.

Example 1.1.2. A 2×2 matrix $z = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ has five minors: the matrix entries a, b, c, d and the determinant $\Delta = \det(z) = ad - bc$. Now the identity

$$(1.1.1) \quad ad = \Delta + bc$$

shows that we do not have to check all five minors: if a, b, c , and Δ are positive, then so is $d = (\Delta + bc)/a$. (Alternatively, test the minors d, b, c, Δ .)

Example 1.1.3. Now let $n = 3$. To keep it simple (cf. also Lemma 1.1.1), let us consider the subgroup of unipotent upper triangular matrices

$$z = \begin{bmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{bmatrix} \in \mathrm{SL}_3.$$

Since some of the entries of z are equal to 0, we modify the definition of total positivity by requiring that $\Delta(z) > 0$ for each minor Δ which does not identically vanish on the subgroup. This leaves us with four minors to check for positivity: the matrix entries a, b, c , and the 2×2 minor $P = ac - b$. Again we can reduce the number of needed checks from 4 to 3 using the identity

$$(1.1.2) \quad ac = P + b.$$

Thus each of the sets $\{a, b, P\}$ and $\{b, c, P\}$ provides an efficient TP test.

We note that in each of the above examples, the number of checks involved in each positivity test was equal to the dimension of the variety at hand. It seems implausible that one could do better.

1.2. The Grassmannian of 2-planes in m -space

Before developing efficient total positivity tests for square matrices, we shall discuss the somewhat simpler case of Grassmannians of 2-planes.

Recall that the complex *Grassmann manifold* (the *Grassmannian* for short), denoted $\mathrm{Gr}_{k,m} = \mathrm{Gr}_{k,m}(\mathbb{C})$, is the variety of all k -dimensional subspaces in an m -dimensional complex vector space. Let us fix a basis in this space, thereby identifying it with \mathbb{C}^m . Now any $k \times m$ matrix z of rank k defines a point $[z] \in \mathrm{Gr}_{k,m}$, the *row span* of z .

Given a k -element subset $J \subset \{1, \dots, m\}$, the *Plücker coordinate* $P_J(z)$ (evaluated at a matrix z as above) is, by definition, the $k \times k$ minor of z determined by the column set J . The collection $(P_J(z))_{|J|=k}$ only depends on the row span $[z]$ (up to common rescaling), and in fact provides an embedding of $\mathrm{Gr}_{k,m}$ into the complex projective space of dimension $\binom{m}{k} - 1$, called the *Plücker embedding*.

The Plücker coordinates P_J generate the *Plücker ring* $R_{k,m}$, the homogeneous coordinate ring of $\mathrm{Gr}_{k,m}$ with respect to the Plücker embedding. The ideal of relations that they satisfy is generated by the quadratic *Grassmann-Plücker relations*.

The Plücker coordinates are used to define the totally positive points of the Grassmannian, as follows.

Definition 1.2.1. The *totally positive Grassmannian* $\mathrm{Gr}_{k,m}^+$ is the subset of $\mathrm{Gr}_{k,m}$ consisting of points whose Plücker coordinates can be chosen so that all of them are positive real numbers. (Recall that Plücker coordinates are defined up to a common rescaling.)

In simple terms, an element $[z] \in \mathrm{Gr}_{k,m}$ defined by a full-rank $k \times m$ matrix z (without loss of generality, z can be assumed to have real entries) is TP if all the maximal (i.e., $k \times k$) minors $P_J(z)$ are of the same sign. We can begin by checking one particular value $P_J(z)$, and if it happens to be negative, replace the top row of z by its negative. Thus the problem at hand can be restated as follows: find an efficient method for checking whether all maximal minors of a given $k \times m$ matrix are positive. A brute force test requires $\binom{m}{k}$ checks. Can this number be reduced?

In this section, we systematically study the case $k = 2$ (Grassmannians of 2-planes). Chapter 8 will generalize the treatment below to arbitrary Grassmannians $\mathrm{Gr}_{k,m}$.

In the case of the Grassmannian $\mathrm{Gr}_{2,m}$, there are $\binom{m}{2}$ Plücker coordinates $P_{ij} = P_{\{i,j\}}$ labeled by pairs of integers $1 \leq i < j \leq m$. It turns out however

that in order to verify that all the 2×2 minors $P_{ij}(z)$ of a given $2 \times m$ matrix z are positive, it suffices to check the positivity of only $2m - 3$ special minors. (Note that $2m - 3 = 2(m - 2) + 1$ is the dimension of the affine cone over $\text{Gr}_{2,m}$.)

Exercise 1.2.2. Show that the 2×2 minors of a $2 \times m$ matrix (equivalently, the Plücker coordinates P_{ij}) satisfy the *three-term Grassmann-Plücker relations*

$$(1.2.1) \quad P_{ik} P_{jl} = P_{ij} P_{kl} + P_{il} P_{jk} \quad (1 \leq i < j < k < l \leq m).$$

We are going to construct a family of “optimal” tests for total positivity in $\text{Gr}_{2,m}$ using triangulations of an m -gon. Consider a convex m -gon \mathbf{P}_m with its vertices labeled clockwise. We associate the Plücker coordinate P_{ij} with the chord (i.e., a side or a diagonal of \mathbf{P}_m) whose endpoints are i and j .

Now, let T be a *triangulation* of \mathbf{P}_m by pairwise noncrossing diagonals. (Common endpoints are allowed.) We view T as a maximal collection of noncrossing chords; as such, it consists of m sides and $m - 3$ diagonals, giving rise to a collection $\tilde{\mathbf{x}}(T)$ of $2m - 3$ Plücker coordinates, which we call an *extended cluster*. The Plücker coordinates corresponding to the sides of \mathbf{P}_m are called *frozen variables*. They are present in every extended cluster $\tilde{\mathbf{x}}(T)$, hence the term “frozen;” an alternative terminology is *coefficient variables*. The remaining $m - 3$ Plücker coordinates corresponding to diagonals of \mathbf{P}_m are called *cluster variables*; they form a *cluster*. Thus each extended cluster consists of $m - 3$ cluster variables and m frozen variables. We note that these $2m - 3$ quantities are algebraically independent. See Figure 1.1.




Figure 1.1. A triangulation T of an octagon \mathbf{P}_8 . The extended cluster $\tilde{\mathbf{x}}(T)$ consists of the cluster variables $P_{17}, P_{24}, P_{27}, P_{46}, P_{47}$ and the frozen variables $P_{12}, P_{23}, \dots, P_{78}, P_{18}$.

Theorem 1.2.3. *Each Plücker coordinate P_{ij} can be written as a subtraction-free rational expression in the elements of a given extended cluster $\tilde{\mathbf{x}}(T)$. Thus, if the $2m - 3$ Plücker coordinates $P_{ij} \in \tilde{\mathbf{x}}(T)$ evaluate positively at a given $2 \times m$ matrix z , then all 2×2 minors of z are positive.*

To clarify, a *subtraction-free* expression is a formula involving variables and positive integers that uses only the operations of addition, multiplication, and division.

Proof of Theorem 1.2.3. Let us visualize the three-term relations (1.2.1) using the m -gon \mathbf{P}_m . Take four vertices $i < j < k < l$ of \mathbf{P}_m , cf. Figure 1.2. Then the relation (1.2.1) is reminiscent of the classical Ptolemy Theorem which asserts that for an inscribed quadrilateral, the products of the lengths of two pairs of opposite sides add up to the product of the two diagonals.




Figure 1.2. Three-term Grassmann-Plücker (or Ptolemy) relation (1.2.1).

Now Theorem 1.2.3 is an immediate consequence of the following three facts:

- (1) Every Plücker coordinate appears as an element of an extended cluster $\tilde{\mathbf{x}}(T)$ for some triangulation T of the polygon \mathbf{P}_m .
- (2) Any two triangulations of \mathbf{P}_m can be transformed into each other by a sequence of *flips*. Each flip removes a diagonal of a triangulation to create a quadrilateral, then replaces it with the other diagonal of the same quadrilateral. See Figure 1.3.
- (3) Each flip, say one involving the diagonals ik and jl , acts on extended clusters by exchanging the Plücker coordinates P_{ik} and P_{jl} . This exchange can be viewed as a subtraction-free transformation determined by the corresponding three-term relation (1.2.1). \square

Remark 1.2.4. In fact something stronger than Theorem 1.2.3 holds: every Plücker coordinate can be written as a *Laurent polynomial* with *positive coefficients* in the Plücker coordinates from $\tilde{\mathbf{x}}(T)$. This is an instance of very general phenomena of Laurentness and positivity in cluster algebras, which will be discussed later in the book.




Figure 1.3. A flip of a diagonal in a quadrilateral.

The combinatorics of flips is captured by the graph whose vertices are labeled by the triangulations of the polygon \mathbf{P}_m and whose edges correspond to flips. This is a regular graph: all its vertices have degree $m - 3$. Moreover, this graph is the 1-skeleton of an $(m - 3)$ -dimensional convex polytope (discovered by J. Stasheff [50]) called the *associahedron*. See Figure 1.4.




Figure 1.4. The 3-dimensional associahedron.

In the forthcoming terminology of cluster algebras, this graph is an example of an *exchange graph*. Its vertices correspond to extended clusters (all of which have the same cardinality) while its edges correspond to *exchange relations* (1.2.1): adjacent extended clusters are related to each other by interchanging the cluster variables appearing on the left-hand side of an exchange relation.

Recall that the Plücker ring $R_{2,m}$ is generated by the Plücker coordinates P_{ij} subject to the three-term relations (1.2.1). The combinatorics of extended clusters provides an elegant way to construct a linear basis for this ring. Define a *cluster monomial* to be a monomial in (i.e., a product of a multiset of) cluster and/or frozen variables, all of which belong to one extended cluster; in other words, the corresponding collection of arcs does not contain two diagonals which cross each other. These monomials appeared already in the classical 19th century literature on invariant theory; in particular, it is known [35, 51] that the set of all cluster monomials form a linear basis of $R_{2,m}$. Later on we shall discuss a far-reaching generalization of this result in the context of cluster algebras.

As mentioned above, the ideas we have discussed for the Grassmannian $\text{Gr}_{2,m}$ can be generalized to a beautiful theory that works for arbitrary Grassmannians $\text{Gr}_{k,m}$, cf. Chapter 8. In Chapter 10, we shall describe a generalization of the $\text{Gr}_{2,m}$ example in a different direction, which involves the combinatorics of flips for triangulations of a Riemann surface with boundary and punctures. That construction has an intrinsic interpretation in hyperbolic geometry where an analogue of the Ptolemy relation holds for the exponentiated hyperbolic distances between horocycles drawn around vertices of a polygon with geodesic sides and cusps at the vertices (the “Penner coordinates” on the corresponding *decorated Teichmüller space* [41]).

1.3. The basic affine space

We next turn our attention to total positivity criteria for square matrices. In view of Lemma 1.1.1, it makes sense to study lower- and upper-triangular matrices first, and then proceed to the whole group SL_n . We choose a related but different strategy: first study matrices for which a certain subset of “flag minors” are positive, then treat the general case.

Definition 1.3.1. The *flag minor* P_J labeled by a nonempty proper subset $J \subsetneq \{1, \dots, n\}$ is a function on the special linear group $G = \text{SL}_n$ defined by

$$P_J : z = (z_{ij}) \mapsto \det(z_{ij} \mid i \leq |J|, j \in J).$$

Since an n -element set has $2^n - 2$ proper nonempty subsets, a matrix in SL_n has $2^n - 2$ flag minors; each of them occupies the first several rows and an arbitrary subset of columns (of the same cardinality).

Let $U \subset G$ be the subgroup of unipotent lower-triangular matrices, i.e. the lower-triangular matrices with 1’s on the diagonal. The group U acts on G by multiplication on the left. It consequently acts on $\mathbb{C}[G]$, the ring of polynomials in the matrix entries of a generic matrix of determinant 1. It is easy to see that each flag minor P_J is an invariant of this action: for any $z \in G$ and $y \in U$, we have $P_J(yz) = P_J(z)$. Similarly to the case of Plücker

coordinates in a Plücker ring, we have (thanks to the appropriate versions of the First and Second Fundamental Theorems of invariant theory):

- (1) the flag minors generate the ring $\mathbb{C}[G]^U$ of U -invariant polynomials in the matrix entries, and
- (2) the ideal of relations among the flag minors is generated by certain quadratic *generalized Plücker relations*.

The ring $\mathbb{C}[G]^U$ plays an important role in the representation theory of the semisimple Lie group G : it naturally carries all irreducible representations of G , each with multiplicity 1. (We will not rely on this in what follows.) This ring is the coordinate ring of the *basic affine space*, the (Geometric Invariant Theory) quotient $U \backslash G$. This space is also known as the *base affine space*, the *fundamental affine space*, and the *principal affine space*. In this section, this space plays the role analogous to the role that the Grassmannians played in Section 1.2. As before, we can state everything we need in an elementary fashion, in terms of matrices and their minors.

Definition 1.3.2. An element $z \in G$ is *flag totally positive* (FTP) if all flag minors P_J take positive values at z .

We would like to detect flag total positivity in an efficient way, by testing as few of the $2^n - 2$ flag minors as possible. It turns out that the optimal test of this kind probes only $\frac{(n-1)(n+2)}{2}$ flag minors. We note that $\frac{(n-1)(n+2)}{2} = n^2 - 1 - \binom{n}{2}$ is the dimension of the basic affine space.

Following [3], we construct a family of tests for flag total positivity labeled by combinatorial objects called *wiring diagrams* which play the same role as triangulations did in Section 1.2.

This notion is best explained by an example such as the one in Figure 1.5. A wiring diagram consists of a family of n piecewise-straight lines, considered up to isotopy, which can be viewed as graphs of n continuous piecewise-linear functions defined on the same interval. The lines are labeled $1, \dots, n$ as shown in Figure 1.5. The key requirement is that each pair of lines intersects exactly once.




Figure 1.5. A wiring diagram.

The notion of a wiring diagram is closely related to that of a *reduced word* for an element of maximal length in the symmetric group \mathcal{S}_n . See [3, Section 2.3] for detailed explanations.

We then assign to each *chamber* of a wiring diagram (i.e., a connected component of its complement inside the vertical strip) a subset of $[1, n] = \{1, \dots, n\}$ indicating which lines pass *below* that chamber; cf. Figure 1.6.




Figure 1.6. Chamber minors P_1, P_2, P_3, P_{12} , and P_{23} .

Thus every chamber is naturally associated with a flag minor P_J , called a *chamber minor*, that occupies the columns specified by the set J in the chamber, and the rows $1, 2, \dots, |J|$. The total number of chamber minors is always $\frac{(n-1)(n+2)}{2}$.

The chamber minors associated to a wiring diagram make up an *extended cluster*. Such an extended cluster will always contain the $2n - 2$ flag minors associated to the unbounded chambers:

$$P_1, P_{1,2}, \dots, P_{1,2,\dots,n-1} \text{ and } P_n, P_{n-1,n}, \dots, P_{2,3,\dots,n};$$

these are the *frozen variables*. The $\binom{n-1}{2}$ chamber minors associated with bounded chambers are the *cluster variables*; they form the corresponding *cluster*.

Theorem 1.3.3 ([3]). *Every flag minor can be written as a subtraction-free rational expression in the chamber minors of any given wiring diagram. Thus, if these $\frac{(n-1)(n+2)}{2}$ chamber minors evaluate positively at a matrix $z \in \mathrm{SL}_n$, then z is FTP.*

Proof. Theorem 1.3.3 is implied by the following three facts:

- (1) Each flag minor appears as a chamber minor in some wiring diagram.
- (2) Any two wiring diagrams can be transformed into each other by a sequence of local *braid moves* of the form shown in Figure 1.7.
- (3) Under each braid move, the corresponding collections of chamber minors are obtained from each other by exchanging the minors Y and Z (cf. Figure 1.7), and these minors satisfy the identity

$$(1.3.1) \quad YZ = AC + BD .$$

Statement (1) is easily checked by direct inspection. Statement (2) is a theorem of G. Ringel [44]; it is also equivalent to a well known property of reduced words in the symmetric group (Tits' Lemma): any two such words are related by braid moves. Finally, formula (1.3.1) is one of the aforementioned generalized Plücker relations; see Exercise 1.3.4 below. \square




Figure 1.7. A braid move.

Exercise 1.3.4. Prove (1.3.1). Hint: use Proposition 1.3.6.

Proposition 1.3.6 and Proposition 1.3.5 below are used to prove (1.3.1), which in turn completes the proof of Theorem 1.3.3.

Proposition 1.3.5 (Muir's Law of extensible minors [39]). *Suppose that (I) is a polynomial identity involving minors $\Delta_{A,B}$ of a generic matrix, which is homogeneous in that every term in (I) is a product of the same number of determinants. Let R and C be finite sets of positive integers such that R (resp., C) is disjoint from every row set A (resp., every column set B) appearing in a determinant in (I). Then one can get a new identity (I') from I by replacing each minor $\Delta_{A,B}$ by $\Delta_{A \cup R, B \cup C}$.*

Proposition 1.3.6. *Suppose (I) is a polynomial identity involving flag minors P_B of a generic matrix, such that every term in (I) is a product of the same number of flag minors. Let C be a finite set of positive integers such that C is disjoint from every column set B appearing in a flag minor P_B in (I). Then one can get a new identity (I'') from (I) by replacing each term P_B by $P_{B \cup C}$.*

Proof. Let $b = |B|$ and $c = |C|$. Recall that the flag minor P_B is equal to $\Delta_{\{1,2,\dots,b\},B}$. Note that if the identity (I) is true, then the identity (I') is true, where (I') is obtained from (I) by replacing each term $\Delta_{\{1,2,\dots,b\},B}$ by the term $\Delta_{\{c+1,c+2,\dots,c+b\},B}$. But now if we apply Muir's Law with $R = \{1, 2, \dots, c\}$ and C, then we get a new identity (I'') from (I') by replacing each term $\Delta_{\{c+1,c+2,\dots,c+b\},B}$ by $\Delta_{\{1,2,\dots,c+b\},B \cup C} = P_{B \cup C}$. \square

Remark 1.3.7. Just as in the case of $\text{Gr}_{2,m}$ (cf. Remark 1.2.4), something much stronger than Theorem 1.3.3 is true: each flag minor can be written as a *Laurent polynomial with positive coefficients* in the chamber minors of a given wiring diagram.

Before moving on to our final example (total positivity for square matrices), let us pause to make a few observations based on our study of total positivity in the Grassmannian $\text{Gr}_{2,m}$, flag total positivity in $G = \text{SL}_n$, and the related rings $R_{2,m}$ and $\mathbb{C}[G]^U$. In both cases, we observed the following common features:

- a family of distinguished generators of the ring (Plücker coordinates and flag minors, respectively);
- a finite subset of *frozen* generators;
- a grouping of the generators into overlapping *extended clusters* all of which have the same size; each extended cluster contains the frozen generators;
- combinatorial data accompanying each extended cluster (triangulations and wiring diagrams, respectively);
- *exchange relations* that can be written using those data; these relations lead to subtraction-free birational maps relating extended clusters to each other;
- a “mutation rule” for producing new combinatorial data from the current one (flips of triangulations and braid moves, respectively).

In the case of the Grassmannian $\text{Gr}_{2,m}$, we defined a graph whose vertices are indexed by the set of triangulations of \mathbf{P}_m , and whose edges correspond to flips. This graph is *regular*, i.e., all its vertices have the same degree. Indeed, in any triangulation, we can flip any of the participating diagonals. Put another way, given an extended cluster and a cluster variable within it, there is a unique way to construct a new extended cluster by replacing that cluster variable by another one.

We could construct an analogous graph related to the flag TP elements of SL_n , with vertices corresponding to wiring diagrams, and edges corresponding to braid moves. However, this graph is not regular for $n \geq 4$. The framework of cluster algebras will rectify this issue by providing a recipe for constructing the missing cluster variables and clusters.

Another important property that we observed for the Grassmannian $\text{Gr}_{2,m}$ concerned cluster monomials. Recall that a cluster monomial is a product of cluster and frozen (=coefficient) variables (not necessarily distinct) all of which belong to one extended cluster. Cluster monomials form a linear basis for the Plücker ring $R_{2,m}$. Unfortunately the analogous statement does not hold for the ring $\mathbb{C}[G]^U$. However, the cluster monomials are still linearly independent, and hence can be included in an additive basis for the ring. Explicit constructions of such additive bases that possess “nice” properties (e.g., various versions of positivity) remain at the center of current research on cluster algebras.

1.4. The general linear group

We now turn to total positivity criteria for the general linear group, or equivalently, for square matrices of arbitrary size. It turns out that to test whether a given $n \times n$ matrix is TP it suffices to check the positivity of only n^2 special minors.

For an $n \times n$ matrix z , let $\Delta_{I,J}(z)$ denote the minor of z determined by the row set I and the column set J ; here I and J are nonempty subsets of $[1, n] = \{1, \dots, n\}$ of the same cardinality. Thus z is TP if and only if $\Delta_{I,J}(z) > 0$ for all such I and J .

Following [20, 21] we construct a family of “optimal” tests for total positivity, labeled by combinatorial objects called *double wiring diagrams*. They generalize the wiring diagrams we saw in the previous section. A double wiring diagram is basically a superposition of two ordinary wiring diagrams, each colored in its own color (‘thin’ or ‘thick’); see Figure 1.8.




Figure 1.8. A double wiring diagram.

The lines in a double wiring diagram are numbered separately within each color. (Note the difference in the numbering schemes for the two colors.) We then assign to every *chamber* of a diagram a pair of subsets of $[1, n]$: each subset indicates which lines of the corresponding color pass below that chamber; see Figure 1.9. Thus every chamber is naturally associated with a minor $\Delta_{I,J}$ (again called a *chamber minor*) that occupies the rows and columns of an $n \times n$ matrix specified by the sets I and J written into that chamber. The total number of chamber minors is always n^2 .




Figure 1.9. This double wiring diagram has $3^2 = 9$ chamber minors: $\Delta_{3,1}$, $\Delta_{3,2}$, $\Delta_{1,2}$, $\Delta_{1,3}$, $\Delta_{23,12}$, $\Delta_{13,12}$, $\Delta_{13,23}$, $\Delta_{12,23}$, and $\Delta_{123,123}$.

Theorem 1.4.1 ([20]). *Every minor of a square matrix can be written as a subtraction-free rational expression in the chamber minors of a given double wiring diagram. Thus, if these n^2 chamber minors evaluate positively at a given $n \times n$ matrix z , then z is totally positive.*

By now the reader can guess the strategy for proving this theorem.

Proof. Theorem 1.4.1 is a consequence of the following facts:

- (1) Every minor is a chamber minor for some double wiring diagram.
- (2) Any two double wiring diagrams are related to each other via a sequence of *local moves* of three different kinds, shown in Figure 1.10.
- (3) Under each local move, the corresponding collections of chamber minors transform by exchanging the minors Y and Z , and these minors satisfy the identity

$$(1.4.1) \quad YZ = AC + BD .$$




Figure 1.10. Local “moves.”

Statements (1) and (2) can be easily derived from their counterparts for ordinary wiring diagrams, which appeared in the proof of Theorem 1.3.3. Each instance of the relation (1.4.1) is a well known determinantal identity; the reader may enjoy finding its proof on her/his own. The identities corresponding to the top two cases in Figure 1.10 are nothing but the three-term relation (1.3.1) which we discussed earlier; the third one is sometimes called the “Lewis Carroll identity,” due to the role it plays in C. L. Dodgson’s

condensation method [14, pp. 170–180]. All of these identities were proved by P. Desnanot as early as in 1819, see [39, pp. 140–142]. \square

Exercise 1.4.2. A minor $\Delta_{I,J}$ is called *solid* if both I and J consist of several consecutive indices. It is easy to see that an $n \times n$ matrix z has n^2 solid minors $\Delta_{I,J}$ such that $I \cup J$ contains 1 (see Figure 1.11). Show that z is TP if and only if all these n^2 minors are positive.




Figure 1.11. Solid minors $\Delta_{I,J}$ with $1 \in I \cup J$.

Remark 1.4.3. Similarly to Remarks 1.2.4 and 1.3.7, Theorem 1.4.1 can be strengthened as follows: every minor of a square matrix can be written as a *Laurent polynomial with positive coefficients* in the chamber minors of a given double wiring diagram.

The algebraic/combinatorial construction described above possesses the same key features that we identified in the two previous settings. The collections of chamber minors associated to double wiring diagrams are again examples of *extended clusters* in the future cluster algebra setup. As noted above, all these extended clusters have the same cardinality n^2 . Each of them contains the $2n - 1$ minors of the form $\Delta_{[1,p],[n-p+1,n]}$ and/or $\Delta_{[n-p+1,n],[1,p]}$ (for $p \in [1, n]$), which correspond to unbounded chambers. These $2n - 1$ minors are the *frozen variables*. Removing them from an extended cluster, we obtain a *cluster* consisting of $(n - 1)^2$ *cluster variables*, the chamber minors associated with the bounded chambers. A “mutation” of one of the three kinds depicted in Figure 1.10 replaces a single cluster variable in a cluster by a new one; the product of these two cluster variables (minors) appears on the left-hand side of the corresponding *exchange relation* (1.4.1).

For $n = 3$, there are 34 clusters corresponding to double wiring diagrams. They are shown in Figure 1.12 as vertices of a graph whose edges correspond to local moves. Looking closely at this graph, we see that it is not regular: of the 34 vertices, 18 have degree 4, and 16 have degree 3. Thus, for each of the 16 clusters corresponding to vertices of degree 3, there is one minor that cannot be exchanged from this cluster to form another cluster.

This “irregularity” can be repaired using two additional polynomials in the matrix entries, see Exercise 1.4.4.

Exercise 1.4.4. For a 3×3 matrix $z = (z_{ij})$, let

$$(1.4.2) \quad K(z) = z_{33}\Delta_{12,12}(z) - \det(z),$$

$$(1.4.3) \quad L(z) = z_{11}\Delta_{23,23}(z) - \det(z).$$

Use K and L to add 16 more clusters to the graph in Figure 1.12. The resulting graph will be regular of degree 4. As an example, the cluster $\{e, f, g, A\}$ at the bottom of Figure 1.12 will be joined to the new cluster $\{e, f, g, K\}$ by an edge corresponding to a new exchange relation

$$(1.4.4) \quad \Delta_{23,23}K = \Delta_{12,23}\Delta_{23,12}z_{33} + \det(z)z_{23}z_{32}.$$

This construction will yield 16 additional TP tests for 3×3 matrices.

The theory of cluster algebras, to be developed in subsequent chapters, will unify the three examples we have treated here, and will provide a systematic way to produce the “missing” clusters and exchange relations, thereby generating a large class of new total positivity tests.




Figure 1.12. Total positivity tests for a 3×3 matrix $z = (z_{ij})$. Each test checks 9 minors. The frozen minors $z_{31}, z_{13}, \Delta_{23,12}(z), \Delta_{12,23}(z), \det(z)$ are common to all tests. The remaining 4 minors form a cluster shown near the corresponding vertex. To illustrate, the test derived from Figure 1.9 involves the cluster $bfCE = \{z_{32}, z_{12}, \Delta_{13,12}, \Delta_{13,23}\}$. The edges of the graph correspond to local moves.

Mutations of quivers and matrices

In this chapter we discuss mutations of quivers and of skew-symmetrizable matrices. These constructions lie at the heart of the combinatorial framework underlying the general theory of cluster algebras.

Quivers (or more generally, skew-symmetrizable matrices) are the combinatorial data which accompany (extended) clusters and determine exchange relations between them. The notion of mutation generalizes many examples of “local” transformations of combinatorial objects, including those discussed in Chapter 1: flips in triangulations, braid moves in wiring diagrams, etc.

In some guise, quiver mutation appeared in the work of theoretical physicists (cf. [9, 47] and the discussion in [15, Section 6]) several years before its discovery by mathematicians [23]. However, the systematic study of the combinatorics of mutations has only begun with the advent of cluster algebras.

2.1. Quiver mutation

Definition 2.1.1. A *quiver* is a finite oriented graph, consisting of vertices and directed edges (called *arrows*). We allow multiple edges, but we disallow loops (i.e., an arrow may not connect a vertex to itself) and oriented 2-cycles (i.e., no arrows of opposite orientation may connect the same pair of vertices). A quiver does not have to be connected.

In what follows, we will need a slightly richer notion, with some vertices in a quiver designated as *frozen*. The remaining vertices are called *mutable*. We will always assume that there are no arrows between pairs of frozen vertices. (Such arrows would make no difference in the future construction of a cluster algebra associated with a quiver.)

The terminology in Definition 2.1.1 anticipates the role that quiver mutations play in the forthcoming definition of a cluster algebra; wherein the vertices of a quiver are labeled by the elements of an extended cluster, so that the frozen vertices correspond to frozen variables, and the mutable vertices to the cluster variables. In this chapter, all of this remains in the background.

Definition 2.1.2. Let k be a mutable vertex in a quiver Q . The *quiver mutation* μ_k transforms Q into a new quiver $Q' = \mu_k(Q)$ via a sequence of three steps:

- (1) For each oriented two-arrow path $i \rightarrow k \rightarrow j$, add a new arrow $i \rightarrow j$ (unless both i and j are frozen, in which case do nothing).
- (2) Reverse the direction of all arrows incident to the vertex k .
- (3) Repeatedly remove oriented 2-cycles until unable to do so.

An example is given in Figure 2.1.




Figure 2.1. A quiver mutation μ_k . Vertices q and r are frozen. Step 1 adds arrows $a \rightarrow b$, $a \rightarrow q$, and two arrows $r \rightarrow b$. Step 2 reverses five arrows connecting k to a, b, q, r . Step 3 removes the arrows $a \rightarrow b$ and $b \rightarrow a$.

Remark 2.1.3. If a vertex k of a quiver is a sink or a source, then mutation at k reverses the orientations of all arrows incident to k , and does nothing else. This operation was first considered in the context of quiver representation theory (the reflection functors of Bernstein-Gelfand-Ponomarev [5]).

We next formulate some simple but important properties of quiver mutation.

Exercise 2.1.4. Verify the following properties of quiver mutation:

- (1) Mutation is an involution: $\mu_k(\mu_k(Q)) = Q$.
- (2) Mutation commutes with the simultaneous reversal of orientations of all arrows of a quiver.
- (3) Let k and ℓ be two mutable vertices which have no arrows between them (in either direction). Then mutations at k and ℓ commute with each other: $\mu_\ell(\mu_k(Q)) = \mu_k(\mu_\ell(Q))$.

In particular, mutations in different connected components of a quiver do not interact with each other.

2.2. Triangulations of polygons

Triangulations of polygons were discussed in Section 1.2 in the context of studying total positivity in the Grassmannian of 2-planes in m -space.

We now associate a quiver to each triangulation of a convex m -gon \mathbf{P}_m and explain how flips of such triangulations correspond to quiver mutations.

Definition 2.2.1. Let T be a triangulation of the polygon \mathbf{P}_m by pairwise noncrossing diagonals. The quiver $Q(T)$ associated to T is defined as follows. The frozen vertices of $Q(T)$ are labeled by the sides of \mathbf{P}_m , and the mutable vertices of $Q(T)$ are labeled by the diagonals of T . If two diagonals, or a diagonal and a boundary segment, belong to the same triangle, we connect the corresponding vertices in $Q(T)$ by an arrow whose orientation is determined by the clockwise orientation of the boundary of the triangle. See Figure 2.2.




Figure 2.2. The quiver $Q(T)$ associated to a triangulation T of an octagon.

Exercise 2.2.2. Let T be a triangulation of \mathbf{P}_m as above. Let T' be the triangulation obtained from T by flipping a diagonal γ . Verify that the quiver $Q(T')$ is obtained from $Q(T)$ by mutating at the vertex labeled by γ .

The construction of Definition 2.2.1 can be generalized to triangulations of more general oriented surfaces with boundary and punctures; this will be discussed in Chapter 10. Another generalization [16] was developed in the study of cluster structures arising in *higher Teichmüller theory*; a very special case is described in the exercise below.

Exercise 2.2.3. To each triangulation T of a convex polygon by noncrossing diagonals, let us associate a quiver $Q_3(T)$ as follows. Place two mutable vertices of $Q_3(T)$ on each diagonal of T , two frozen vertices on each side of the polygon, and one mutable vertex in the interior of each triangle of T . For a triangle in T , let $A_1, A_2, B_1, B_2, C_1, C_2$ be the vertices of $Q_3(T)$ lying on the boundary of the triangle, listed clockwise so that A_1 and A_2 (resp., B_1 and B_2 , C_1 and C_2) lie on the same side of the triangle. Let K denote the vertex of $Q_3(T)$ lying inside the triangle. Draw the arrows $A_1 \rightarrow K \rightarrow B_2 \rightarrow C_1 \rightarrow K \rightarrow A_2 \rightarrow B_1 \rightarrow K \rightarrow C_2 \rightarrow A_1$. Doing so for each triangle of T , and removing the arrows between frozen vertices, yields the quiver $Q_3(T)$. Show that if T and T' are connected by a flip, then $Q_3(T)$ and $Q_3(T')$ are connected by a sequence of mutations. See Figure 2.3.




Figure 2.3. The quiver $Q_3(T)$ associated to a triangulation T of a quadrilateral.

2.3. Wiring diagrams

Wiring diagrams were introduced in Section 1.3 in the context of studying total positivity in basic affine spaces. We also explained how to label each chamber of a wiring diagram by a subset of $[1, n]$, cf. Figure 1.6.

We now associate a quiver to each wiring diagram and demonstrate that braid moves in wiring diagrams translate into quiver mutations.

The *left end* (resp., *right end*) of a chamber is a crossing point of two wires located at the very left (resp., right) of the chamber. Each bounded chamber has two ends; an unbounded chamber has one.

Definition 2.3.1. The quiver $Q(D)$ associated to a wiring diagram D is defined as follows. The vertices of $Q(D)$ are labeled by the chambers of D . The

bounded chambers correspond to mutable vertices; the unbounded chambers correspond to frozen vertices. Let c and c' be two chambers, at least one of which is bounded. Then there is an arrow $c \rightarrow c'$ in $Q(D)$ if and only if one of the following conditions is met:

- (i) the right end of c coincides with the left end of c' ;
- (ii) the left end of c lies directly above c' , and the right end of c' lies directly below c ;
- (iii) the left end of c lies directly below c' , and the right end of c' lies directly above c .

An example is shown in Figure 2.4.




Figure 2.4. A quiver associated with a wiring diagram. All vertices but **2** are frozen (the latter corresponds to the only bounded chamber). Consequently the quiver does not include the arrow $\mathbf{12} \rightarrow \mathbf{23}$ because $\mathbf{12}$ and $\mathbf{23}$ are both frozen/ unbounded.

The somewhat technical construction of Definition 2.3.1 is justified by the fact that braid moves on wiring diagrams translate into mutations of associated quivers:

Proposition 2.3.2. *Let D and D' be wiring diagrams related by a braid move at chamber Y (cf. Figure 1.7). Then $Q(D') = \mu_Y(Q(D))$.*

We leave the proof of Proposition 2.3.2 as an exercise for the reader.

Exercise 2.3.3. Draw the wiring diagrams corresponding to the quivers in Figure 2.5. Verify that these wiring diagrams are related by a braid move, and that the quivers are related by a quiver mutation.

Remark 2.3.4. We note that the wiring diagrams introduced in Section 1.3 can be identified with reduced decompositions of the longest permutation w_0 of the symmetric group. One can extend the notion of wiring diagram to the setting of decompositions (not necessarily reduced) of an arbitrary element of the symmetric group. The quiver $Q(D)$ and the correspondence between braid moves and mutations also make sense in this setting.




Figure 2.5. Quivers for two wiring diagrams related by a braid move.

2.4. Double wiring diagrams

We next extend the constructions of Section 2.3 to the double wiring diagrams discussed in Section 1.4.

Recall that each chamber of a double wiring diagram D is labeled by a pair of subsets of $[1, n]$, cf. Figure 1.9. Similarly to the case of ordinary wiring diagrams, each chamber of D has either one or two “ends,” and each end is either “thick” or “thin” (formed by two thick lines or two thin lines).

Definition 2.4.1. The quiver $Q(D)$ associated with a double wiring diagram D is defined as follows. The vertices of $Q(D)$ are labeled by the chambers of D . The bounded chambers correspond to mutable vertices; the unbounded chambers correspond to frozen vertices. Let c and c' be two chambers, at least one of which is bounded. Then there is an arrow $c \rightarrow c'$ in $Q(D)$ if and only if one of the following conditions is met (cf. Figure 2.6):

- (i) the right (resp., left) end of c is thick (resp., thin), and coincides with the left (resp., right) end of c' ;
- (ii) the left end of c' is thin, the right end of c' is thick, and the entire chamber c' lies directly above or directly below c ;
- (iii) the left end of c is thick, the right end of c is thin, and the entire chamber c lies directly above or directly below c' ;
- (iv) the left (resp., right) end of c' is above c and the right (resp., left) end of c is below c' and both of these ends are thin (resp., thick);
- (v) the left (resp., right) end of c is above c' and the right (resp., left) end of c' is below c and both of these ends are thick (resp., thin).

Remark 2.4.2. One can check that the quiver $Q(D)$ defined as above depends only on the isotopy type of the double wiring diagram D .

As before, local moves translate into quiver mutations:

Proposition 2.4.3. *Suppose that double wiring diagrams D and D' are related by a local move (cf. Figure 1.10) at chamber Y . Then $Q(D') = \mu_Y(Q(D))$.*




Figure 2.6. A double wiring diagram D and the corresponding quiver $Q(D)$.

The verification of this statement, which generalizes Proposition 2.3.2, is left as an exercise for the reader.

Remark 2.4.4. While the constructions of the quiver associated to a (double) wiring diagram given in Definition 2.3.1 and Definition 2.4.1 appear a bit complicated, we will explain in Chapter 8 how these constructions can be viewed as a special case of the quiver associated to a planar bicolored graph in a disk. Section 2.5 below will give a first introduction to quivers associated to planar bipartite graphs.

2.5. Urban renewal

Urban renewal [33] is an operation on bipartite graphs which arises in several different contexts including statistical mechanics (spider moves in dimer models [30]), gauge theory (Seiberg duality action on brane tilings [27]), and total positivity (square moves in dual graphs of Postnikov diagrams [42]).

We will give urban renewal a more thorough treatment in Chapter 7. In Chapter 8, it will play an important role in the study of cluster structures on Grassmannians.

Definition 2.5.1. Let G be a connected planar bipartite graph, properly embedded in a disk, and considered up to isotopy. More precisely, we require the following:

- each vertex in G is colored either white or black and lies either in the interior of the disk or on its boundary;
- each edge in G connects two vertices of different colors, and is represented by a simple curve whose interior is disjoint from the other edges and from the boundary;

- each boundary vertex has degree 1, and each internal vertex has degree at least 2.
- the closure of each *face* (i.e, a connected component of the complement of G) is simply connected.

To such a bipartite planar graph G , we associate a quiver $Q(G)$ as follows. The vertices of $Q(G)$ are labeled by the faces of G . A vertex of $Q(G)$ is frozen if the corresponding face is incident to the boundary of the disk, and is mutable otherwise. For each edge e in G , we introduce an arrow connecting the (distinct) faces separated by e ; this arrow is oriented so that it “sees” the white endpoint of e to the left and the black endpoint to the right as it crosses over e , see Figure 2.7. We then remove oriented 2-cycles from the resulting quiver, one by one, to get $Q(G)$.



Figure 2.7. Constructing a quiver associated to a bipartite graph on a surface.

We assume that the quiver $Q(G)$ is connected. One simple observation is that if G has a vertex v of degree 2, and we construct \tilde{G} from G by deleting v and contracting its two incident edges, then the associated quiver does not change: $Q(G) = Q(\tilde{G})$.

Definition 2.5.2. *Urban renewal* is a local transformation of a bipartite graph G as above that takes a quadrilateral face, each of whose vertices has degree at least 3, and adds or removes four “legs” as shown in Figure 2.8.




Figure 2.8. Urban renewal.

Exercise 2.5.3. Verify that if bipartite graphs as above are related via urban renewal, then the corresponding quivers are related by a mutation.

Definition 2.5.2 can be generalized to bipartite graphs properly embedded into an oriented surface.

2.6. Mutation equivalence

Definition 2.6.1. Two quivers Q and Q' are called *mutation equivalent* if Q can be transformed into a quiver isomorphic to Q' by a sequence of mutations. (Equivalently, Q' can be transformed into a quiver isomorphic to Q .) The *mutation equivalence class* $[Q]$ of a quiver Q is the set of all quivers (up to isomorphism) which are mutation equivalent to Q .

Definition 2.6.2. Two quivers Q and Q' are said to have the same *type* if their mutable parts are mutation equivalent. Here the *mutable part* of the quiver refers to the mutable vertices together with all arrows which connect two mutable vertices. When the mutation equivalence class of the mutable part has a name (e.g. type ADE) then we will use that name to describe the type.

Example 2.6.3. Consider the quiver Q at the left of Figure 2.9; this is an orientation of the type A_3 Dynkin diagram. The mutation equivalence class $[Q]$ of Q consists of quivers isomorphic to one of those shown in Figure 2.9. In particular, we consider an oriented 3-cycle to be a quiver of type A_3 .



Figure 2.9. The isomorphism classes of the quivers of type A_3 . All three vertices are mutable.

Example 2.6.4. The *Markov quiver* is a quiver Q of the form shown in Figure 2.10. Mutating Q at any of its 3 vertices produces a quiver isomorphic to Q , so $[Q]$ consists of a single element (up to isomorphism).



Figure 2.10. The Markov quiver. All three vertices are mutable.

Exercise 2.6.5. Show that all orientations of a tree (with no frozen vertices) are mutation equivalent to each other via mutations at sinks and sources.

Exercise 2.6.6. Which orientations of an n -cycle are mutation equivalent?

An $a \times b$ *grid quiver* is an orientation of an $a \times b$ grid in which each 4-cycle is oriented either clockwise or counterclockwise; see Figure 2.11. All vertices are mutable.

Exercise 2.6.7. Show that a grid quiver is mutation equivalent to the corresponding *triangulated grid quiver* (see Figure 2.11).




Figure 2.11. The 3×4 grid quiver, and the corresponding triangulated grid quiver.

Exercise 2.6.8. Verify that in each row of Figure 2.12, the quiver on the left is mutation equivalent to any orientation of the Dynkin diagram on the right.




Figure 2.12. Quivers mutation equivalent to orientations of Dynkin diagrams of types D_4, D_5, E_6, E_7, E_8 .

A *triangular grid quiver* with k vertices on each side is a quiver with $\binom{k+1}{2}$ vertices and $3\binom{k}{2}$ arrows that has the form shown in Figure 2.13. All vertices are mutable.




Figure 2.13. A triangular grid quiver.

Exercise 2.6.9. Show that the triangular grid quiver with three vertices on each side (see Figure 2.13) is mutation equivalent to an orientation of a tree. (However, this is no longer true for a triangular grid quiver with four vertices on each side.)

Exercise 2.6.10. Show that the $k \times (2k + 1)$ grid quiver is mutation equivalent to the triangular grid quiver with $2k$ vertices on each side.

Definition 2.6.11. A quiver Q is said to have *finite mutation type* if the mutation equivalence class $[Q]$ of Q is finite.

Quivers with no frozen vertices of finite mutation type can be completely classified in explicit combinatorial terms. This classification will be described in Chapter 10.

We conclude this section by stating, without proof, an innocent-looking but rather nontrivial result about quiver mutation.

A quiver is called *acyclic* if it has no oriented cycles.

Theorem 2.6.12 (see [7]). *Let Q and Q' be acyclic quivers mutation equivalent to each other. Then Q can be transformed into a quiver isomorphic to Q' via a sequence of mutations at sources and sinks. Consequently (cf. Remark 2.1.3), all acyclic quivers in a given mutation equivalence class have the same underlying undirected graph.*

Corollary 2.6.13. *An acyclic quiver which is mutation equivalent to an orientation of a tree is itself an orientation of the same tree. In particular, orientations of non-isomorphic trees are not mutation equivalent.*

The proof of Theorem 2.6.12 (and hence Corollary 2.6.13) uses *cluster categories*; it would be very interesting to find a purely combinatorial proof of either of these results.

In general, it can be very hard to determine whether two quivers are mutation equivalent to each other.

Problem 2.6.14. Design a (reasonably efficient) algorithm for deciding whether two quivers are mutation equivalent or not.

2.7. Matrix mutation

In this section, we extend the notion of mutation from quivers to a certain class of matrices. We begin by explaining how matrices can be viewed as generalizations of quivers.

Definition 2.7.1. Let Q be a quiver (as in Definition 2.1.1) with m vertices, n of them mutable. Let us label the vertices of Q by the indices $1, \dots, m$ so that the mutable vertices are labeled $1, \dots, n$. The *extended exchange matrix* of Q is the $m \times n$ matrix $\tilde{B}(Q) = (b_{ij})$ defined by

$$b_{ij} = \begin{cases} \ell & \text{if there are } \ell \text{ arrows from vertex } i \text{ to vertex } j \text{ in } Q; \\ -\ell & \text{if there are } \ell \text{ arrows from vertex } j \text{ to vertex } i \text{ in } Q; \\ 0 & \text{otherwise.} \end{cases}$$

The *exchange matrix* $B(Q)$ is the $n \times n$ skew-symmetric submatrix of $\tilde{B}(Q)$ occupying the first n rows:

$$B(Q) = (b_{ij})_{i,j \in [1,n]}.$$

To illustrate, consider the Markov quiver Q shown in Figure 2.10. Then

$$\tilde{B}(Q) = B(Q) = \pm \begin{bmatrix} 0 & 2 & -2 \\ -2 & 0 & 2 \\ 2 & -2 & 0 \end{bmatrix},$$

where the sign depends on the labeling of the vertices.

Remark 2.7.2. While the definition of $\tilde{B}(Q)$ depends on the choice of labeling of the vertices of Q by the integers $1, \dots, m$, we often consider extended exchange matrices up to a simultaneous relabeling of rows and columns $1, 2, \dots, n$, and a relabeling of the rows $n + 1, n + 2, \dots, m$.

The proof of the following lemma is straightforward.

Lemma 2.7.3. *Let k be a mutable vertex of a quiver Q . The extended exchange matrix $\tilde{B}(\mu_k(Q)) = (b'_{ij})$ of the mutated quiver $\mu_k(Q)$ is given by*

$$(2.7.1) \quad b'_{ij} = \begin{cases} -b_{ij} & \text{if } i = k \text{ or } j = k; \\ b_{ij} + b_{ik}b_{kj} & \text{if } b_{ik} > 0 \text{ and } b_{kj} > 0; \\ b_{ij} - b_{ik}b_{kj} & \text{if } b_{ik} < 0 \text{ and } b_{kj} < 0; \\ b_{ij} & \text{otherwise.} \end{cases}$$

We next move from skew-symmetric matrices to a more general class of matrices.

Definition 2.7.4. An $n \times n$ matrix $B = (b_{ij})$ with integer entries is called *skew-symmetrizable* if $d_i b_{ij} = -d_j b_{ji}$ for some positive integers d_1, \dots, d_n . In other words, a matrix is skew-symmetrizable if it differs from a skew-symmetric matrix by a rescaling of its rows by positive scalars.

An $m \times n$ integer matrix, with $m \geq n$, whose top $n \times n$ submatrix is skew-symmetrizable is called an *extended skew-symmetrizable* matrix.

Exercise 2.7.5. Show that the class of matrices B described in Definition 2.7.4 would not change if instead of rescaling the rows of B , we allow to rescale its columns; alternatively, we could allow conjugation of B by a diagonal matrix with positive real diagonal entries.

We are now ready to define the notion of matrix mutation.

Definition 2.7.6. Let $\tilde{B} = (b_{ij})$ be an $m \times n$ extended skew-symmetrizable integer matrix. For $k \in [1, n]$, the *matrix mutation μ_k in direction k* transforms \tilde{B} into the $m \times n$ matrix $\mu_k(\tilde{B}) = (b'_{ij})$ whose entries are given by (2.7.1).

By Lemma 2.7.3, matrix mutation generalizes quiver mutation.

Exercise 2.7.7. Under the conventions of Definitions 2.7.4 and 2.7.6, verify that

- (1) the mutated matrix $\mu_k(\tilde{B})$ is again extended skew-symmetrizable, with the same choice of d_1, \dots, d_n ;
- (2) $\mu_k(\mu_k(\tilde{B})) = \tilde{B}$;
- (3) $\mu_k(-\tilde{B}) = -\mu_k(\tilde{B})$;
- (4) $\mu_k(B^T) = (\mu_k(B))^T$, where B^T denotes the transpose of B ;
- (5) if $b_{ij} = b_{ji} = 0$, then $\mu_i(\mu_j(\tilde{B})) = \mu_j(\mu_i(\tilde{B}))$.

For $b \in \mathbb{R}$, let $\text{sgn}(b)$ be 1, 0, or -1 , depending on whether b is positive, zero, or negative.

Definition 2.7.8. Let B be a skew-symmetrizable matrix. The skew-symmetric matrix $S(B) = (s_{ij})$ defined by

$$(2.7.2) \quad s_{ij} = \text{sgn}(b_{ij}) \sqrt{|b_{ij} b_{ji}|}$$

is called the *skew-symmetrization* of B . Note that $S(B)$ has real (not necessarily integer) entries. Exercise 2.7.9 shows that skew-symmetrization commutes with mutation (extended verbatim to matrices with real entries).

Exercise 2.7.9. For any skew-symmetrizable matrix B and any k , we have

$$(2.7.3) \quad S(\mu_k(B)) = \mu_k(S(B)).$$

Definition 2.7.10. The *diagram* of a skew-symmetrizable $n \times n$ matrix $B = (b_{ij})$ is the weighted directed graph $\Gamma(B)$ with the vertices $1, \dots, n$ such that there is a directed edge from i to j if and only if $b_{ij} > 0$, and this edge is assigned the weight $|b_{ij}b_{ji}|$. In particular, if $b_{ij} \in \{-1, 0, 1\}$ for all i and j , then $\Gamma(B)$ is a quiver whose exchange matrix is B .

To illustrate Definition 2.7.10, consider $B = \begin{bmatrix} 0 & 2 & -2 \\ -1 & 0 & 2 \\ 1 & -2 & 0 \end{bmatrix}$. Then $\Gamma(B)$ is an oriented cycle with edge weights 2, 4, and 2.

More generally, we use the term *diagram* in the rest of this chapter to mean a finite directed graph Γ (no loops, multiple edges, or 2-cycles allowed) whose edges are assigned positive real weights.

We note that the diagram $\Gamma(B)$ does *not* determine B : for instance, the matrix $(-B^T)$ has the same diagram as B . Here is another example:

$$\Gamma\left(\begin{bmatrix} 0 & 1 \\ -4 & 0 \end{bmatrix}\right) = \Gamma\left(\begin{bmatrix} 0 & 2 \\ -2 & 0 \end{bmatrix}\right).$$

Note that the diagram $\Gamma(B)$ and the skew-symmetric matrix $S(B)$ encode the same information about B : having an edge $i \rightarrow j$ in $\Gamma(B)$ supplied with weight c is the same as saying that $s_{ij} = \sqrt{c}$ and $s_{ji} = -\sqrt{c}$.

Proposition 2.7.11. *For a skew-symmetrizable matrix B , the diagram $\Gamma' = \Gamma(\mu_k(B))$ is uniquely determined by the diagram $\Gamma = \Gamma(B)$ and an index k .*

Proof. By Exercise 2.7.9, $S(\mu_k(B)) = \mu_k(S(B))$. It remains to translate this property into the language of diagrams. \square

In the situation of Proposition 2.7.11, we write $\Gamma' = \mu_k(\Gamma)$, and call the transformation μ_k a *diagram mutation in direction k* . A detailed description of diagram mutation can be found in [24, Proposition 8.1]. Two diagrams Γ and Γ' related by a sequence of mutations are called *mutation equivalent*, and we write $\Gamma \sim \Gamma'$.

Remark 2.7.12. While the entries of B are integers, the entries of $S(B)$ may be irrational, as the weights of $\Gamma(B)$ may not be perfect squares. On the other hand, one can deduce from the skew-symmetrizability of B that the product of weights over the edges of any cycle in the underlying graph of $\Gamma(B)$ is a perfect square.

Lemma 2.7.13. *If the diagram $\Gamma(B)$ of an $n \times n$ skew-symmetrizable matrix B is connected, then the skew-symmetrizing vector (d_1, \dots, d_n) is unique up to rescaling.*

Proof. Let (d_1, \dots, d_n) and (d'_1, \dots, d'_n) be two skew-symmetrizing vectors. We have $d_i b_{ij} = -d_j b_{ji}$ and $d'_i b_{ij} = -d'_j b_{ji}$ for all i and j . So if b_{ij} is nonzero,

then $\frac{b_{ij}}{b_{ji}} = \frac{-d_j}{d_i} = \frac{-d'_j}{d'_i}$ and hence $\frac{d_j}{d'_j} = \frac{d_i}{d'_i}$. Since $\Gamma(B)$ is connected, there exists an ordering $\ell_1, \ell_2, \dots, \ell_n$ of its vertices such that every vertex ℓ_j with $2 \leq j \leq n$ is connected by an edge in $\Gamma(B)$ to a vertex ℓ_i with $i < j$; in other words, $b_{\ell_i \ell_j} \neq 0$. It follows that $\frac{d_{\ell_1}}{d'_{\ell_1}} = \frac{d_{\ell_2}}{d'_{\ell_2}} = \dots = \frac{d_{\ell_n}}{d'_{\ell_n}}$, as desired. \square

2.8. Invariants of matrix mutations

The following notion is a straightforward extension of Definition 2.6.1.

Definition 2.8.1. Two skew-symmetrizable matrices B and B' are *mutation equivalent* if one can get from B to B' by a sequence of mutations, possibly followed by simultaneous renumbering of rows and columns. The *mutation equivalence class* $[B]$ of B is the set of all matrices mutation equivalent to B . These notions generalize to extended skew-symmetrizable matrices in an obvious way.

It is natural to extend Problem 2.6.14 to the setting of matrix mutations:

Problem 2.8.2. Find an effective way to determine whether two given $n \times n$ skew-symmetrizable matrices are mutation equivalent.

Problem 2.8.2 remains wide open, even in the case of skew-symmetric matrices (or equivalently quivers). For $n = 2$, the question is trivial, since mutation simply negates the entries of the matrix. For $n = 3$, there is an explicit algorithm for determining whether two skew-symmetric matrices are mutation equivalent, see [2].

Problem 2.8.2 is closely related to the problem of identifying explicit nontrivial invariants of matrix (or quiver) mutation. Unfortunately, very few invariants of this kind are known at present.

Theorem 2.8.3 ([4, Lemma 3.2]). *Mutations preserve the rank of a matrix.*

Proof. Let \tilde{B} be an $m \times n$ extended skew-symmetrizable integer matrix. Fix an index $k \in [1, n]$ and a sign $\varepsilon \in \{1, -1\}$. The rule (2.7.1) describing the matrix mutation in direction k can be rewritten as follows:

$$(2.8.1) \quad b'_{ij} = \begin{cases} -b_{ij} & \text{if } i = k \text{ or } j = k; \\ b_{ij} + \max(0, -\varepsilon b_{ik}) b_{kj} + b_{ik} \max(0, \varepsilon b_{kj}) & \text{otherwise.} \end{cases}$$

(To verify this, examine the four possible sign patterns for b_{ik} and b_{kj} .) Next observe that (2.8.1) can be restated as

$$(2.8.2) \quad \begin{aligned} \mu_k(\tilde{B}) &= J_{m,k} \tilde{B} J_{n,k} + J_{m,k} \tilde{B} F_k + E_k \tilde{B} J_{n,k} \\ &= (J_{m,k} + E_k) \tilde{B} (J_{n,k} + F_k) \end{aligned}$$

where

- $J_{m,k}$ (respectively, $J_{n,k}$) denotes the diagonal matrix of size $m \times m$ (respectively, $n \times n$) whose diagonal entries are all 1, except for the (k, k) entry, which is -1 ;
- $E_k = (e_{ij})$ is the $m \times m$ matrix with $e_{ik} = \max(0, -\varepsilon b_{ik})$, and all other entries equal to 0;
- $F_k = (f_{ij})$ is the $n \times n$ matrix with $f_{kj} = \max(0, \varepsilon b_{kj})$, and all other entries equal to 0.

(Here we used that $E_k \tilde{B} F_k = 0$ because $b_{ii} = 0$ for all i .) Since

$$(2.8.3) \quad \det(J_{m,k} + E_k) = \det(J_{n,k} + F_k) = -1,$$

it follows that $\text{rank}(\mu_k(\tilde{B})) = \text{rank}(\tilde{B})$. □

Theorem 2.8.4. *The determinant of a skew-symmetrizable matrix is invariant under mutation.*

Proof. This follows from (2.8.2) and (2.8.3) (taking $m = n$ and $B = \tilde{B}$). □

Another invariant of matrix mutations is the greatest common divisor of the matrix elements of B . A finer invariant is the greatest common divisor of the matrix elements of the i th row (or column) of B , for a fixed index i [48].

Remark 2.8.5. For skew-symmetric matrices (equivalently, quivers with no frozen vertices), formulas (2.8.2) and (2.8.3) allow us to interpret mutation as a transformation of a skew-symmetric bilinear form over the integers under a particular unimodular change of basis. One can then use the general theory of invariants of such transformations (the *skew Smith normal form*, see [40, Section IV.3]) to identify some invariants of quiver mutation. Unfortunately this approach does not yield much beyond the facts established above.

Clusters and seeds

This chapter introduces cluster algebras of *geometric type*. A more general construction of cluster algebras over an arbitrary semifield will be discussed in Chapter 12.

3.1. Basic definitions

Let us recall the three motivating examples discussed in Chapter 1: Grassmannians of 2-planes, affine base spaces, and general linear groups. In each of these examples, we manipulated two kinds of data:

- combinatorial data (triangulations, wiring diagrams) and
- algebraic data (Plücker coordinates, chamber minors).

Accordingly, transformations applied to these data occurred on two levels:

- on the “primary” level, we saw the combinatorial data evolve via local moves (flips in triangulations, braid moves in wiring diagrams); as shown in Chapter 2, a unifying description of this dynamics can be given using the language of quiver mutations;
- on the “secondary” level, we saw the algebraic data evolve in a way that was “driven” by the combinatorial dynamics, with subtraction-free birational transformations, called exchange relations, encoded by the current combinatorial data.

An attempt to write the exchange relations in terms of the quiver at hand naturally leads to the axiomatic setup of cluster algebras of geometric type, which we now describe.

Let m and n be two positive integers such that $m \geq n$. As an *ambient field* for a cluster algebra, we take a field \mathcal{F} isomorphic to the field of rational functions over \mathbb{C} (alternatively, over \mathbb{Q}) in m independent variables.

Definition 3.1.1. A *labeled seed* of geometric type in \mathcal{F} is a pair $(\tilde{\mathbf{x}}, \tilde{B})$ where

- $\tilde{\mathbf{x}} = (x_1, \dots, x_m)$ is an m -tuple of elements of \mathcal{F} forming a *free generating set*; that is, x_1, \dots, x_m are algebraically independent, and $\mathcal{F} = \mathbb{C}(x_1, \dots, x_m)$;
- $\tilde{B} = (b_{ij})$ is an $m \times n$ extended skew-symmetrizable integer matrix, see Definition 2.7.4.

We shall use the following terminology:

- $\tilde{\mathbf{x}}$ is the (labeled) *extended cluster* of the labeled seed $(\tilde{\mathbf{x}}, \tilde{B})$;
- the n -tuple $\mathbf{x} = (x_1, \dots, x_n)$ is the (labeled) *cluster* of this seed;
- the elements x_1, \dots, x_n are its *cluster variables*;
- the remaining elements x_{n+1}, \dots, x_m of $\tilde{\mathbf{x}}$ are the *frozen variables* (or *coefficient variables*);
- the matrix \tilde{B} is the *extended exchange matrix* of the seed;
- its top $n \times n$ submatrix B is the *exchange matrix*.

Definition 3.1.2. Let $(\tilde{\mathbf{x}}, \tilde{B})$ be a labeled seed as above. Take an index $k \in \{1, \dots, n\}$. The *seed mutation* μ_k in direction k transforms $(\tilde{\mathbf{x}}, \tilde{B})$ into the new labeled seed $\mu_k(\tilde{\mathbf{x}}, \tilde{B}) = (\tilde{\mathbf{x}}', \tilde{B}')$ defined as follows:

- $\tilde{B}' = \mu_k(\tilde{B})$ (cf. Definition 2.7.6).
- the extended cluster $\tilde{\mathbf{x}}' = (x'_1, \dots, x'_m)$ is given by $x'_j = x_j$ for $j \neq k$, whereas $x'_k \in \mathcal{F}$ is determined by the *exchange relation*

$$(3.1.1) \quad x_k x'_k = \prod_{b_{ik} > 0} x_i^{b_{ik}} + \prod_{b_{ik} < 0} x_i^{-b_{ik}}.$$

We note that if the indexing set for one of the two monomials above is the empty set, then by convention we set the corresponding product equal to 1.

Exercise 3.1.3. Consider each of the three settings that we discussed in Sections 1.2, 1.3, and 1.4. Construct a seed $(\tilde{\mathbf{x}}, \tilde{B}(Q))$ where Q is a quiver associated with a particular triangulation, wiring diagram, or double wiring diagram (see Definitions 2.2.1, 2.3.1, and 2.4.1, respectively), and $\tilde{\mathbf{x}}$ is the extended cluster consisting of the corresponding Plücker coordinates or chamber minors. Verify that applying the recipe (3.1.1) to these data recovers the appropriate exchange relations (1.2.1), (1.3.1), and (1.4.1), respectively.

To keep track of the various labeled seeds one can obtain by mutation from a given one, we introduce the following combinatorial setup.




Figure 3.1. The n -regular trees \mathbb{T}_n for $n = 1, 2, 3$.

Definition 3.1.4. Let \mathbb{T}_n denote the n -regular tree whose edges are labeled by the numbers $1, \dots, n$, so that the n edges incident to each vertex receive different labels. See Figure 3.1.

We shall write $t \xrightarrow{k} t'$ to indicate that vertices $t, t' \in \mathbb{T}_n$ are joined by an edge labeled by k .

Definition 3.1.5. A *seed pattern* is defined by assigning a labeled seed $(\tilde{x}(t), \tilde{B}(t))$ to every vertex $t \in \mathbb{T}_n$, so that the seeds assigned to the endpoints of any edge $t \xrightarrow{k} t'$ are obtained from each other by the seed mutation in direction k . A seed pattern is uniquely determined by any one of its seeds. See Figure 3.2.




Figure 3.2. Clusters in a seed pattern.

Now everything is in place for defining cluster algebras.

Definition 3.1.6. Let $(\tilde{\mathbf{x}}(t), \tilde{B}(t))_{t \in \mathbb{T}_n}$ be a seed pattern as above, and let

$$\mathcal{X} = \bigcup_{t \in \mathbb{T}_n} \mathbf{x}(t)$$

be the set of all cluster variables appearing in its seeds. We let the *ground ring* be $R = \mathbb{C}[x_{n+1}, \dots, x_m]$, the polynomial ring generated by the frozen variables. (A common alternative is to take $R = \mathbb{C}[x_{n+1}^{\pm 1}, \dots, x_m^{\pm 1}]$, the ring of Laurent polynomials in the frozen variables, see Section 6.6 for an example. Sometimes the scalars are restricted to \mathbb{Q} , or even to \mathbb{Z} .)

The *cluster algebra* \mathcal{A} (of geometric type, over R) associated with the given seed pattern is the R -subalgebra of the ambient field \mathcal{F} generated by all cluster variables: $\mathcal{A} = R[\mathcal{X}]$. To be more precise, a cluster algebra is the R -subalgebra \mathcal{A} as above together with a fixed seed pattern in it.

A common way to describe a cluster algebra is to pick an *initial (labeled) seed* $(\tilde{\mathbf{x}}_o, \tilde{B}_o)$ in \mathcal{F} and build a seed pattern from it. The corresponding cluster algebra, denoted $\mathcal{A}(\tilde{\mathbf{x}}_o, \tilde{B}_o)$, is generated over the ground ring R by all cluster variables appearing in the seeds mutation equivalent to $(\tilde{\mathbf{x}}_o, \tilde{B}_o)$.

Remark 3.1.7 (cf. Exercise 3.1.3). It can be shown that applying this construction in each of the three settings discussed in Chapter 1, one obtains cluster algebras naturally identified with the Plücker ring $R_{2,m}$ (cf. Section 6.7), the ring of invariants $\mathbb{C}[\mathrm{SL}_k]^U$ (cf. Section 6.5), and the polynomial ring $\mathbb{C}[z_{11}, z_{12}, \dots, z_{kk}]$ (cf. Section 6.6), respectively.

Remark 3.1.8. It is often more natural to work with (*unlabeled*) seeds, which differ from the labeled ones in that we identify two seeds $(\tilde{\mathbf{x}}, \tilde{B})$ and $(\tilde{\mathbf{x}}', \tilde{B}')$ in which \mathbf{x}' is a permutation of \mathbf{x} , and \tilde{B}' is obtained from \tilde{B} by the corresponding permutation of rows and columns. Note that ignoring the labeling does not affect the resulting cluster algebra in a meaningful way.

Remark 3.1.9. Many questions arising in cluster algebra theory and its applications do not really concern cluster algebras as such. These are questions which are not about commutative rings carrying a cluster structure; rather, they are about seed patterns and the birational transformations that relate extended clusters to each other. For those questions, the choice of the ground ring is immaterial: the formulas remain the same regardless.

Remark 3.1.10. Since any free generating collection of m elements in \mathcal{F} can be mapped to any other such collection by an automorphism of \mathcal{F} , the choice of the initial extended cluster $\tilde{\mathbf{x}}_o$ is largely inconsequential: the cluster algebra $\mathcal{A}(\tilde{\mathbf{x}}_o, \tilde{B}_o)$ is determined, up to an isomorphism preserving all the matrices $\tilde{B}(t)$, by the initial extended exchange matrix \tilde{B}_o , and indeed by its mutation equivalence class. Also, replacing \tilde{B}_o by $-\tilde{B}_o$ yields essentially the same cluster algebra (all matrices $\tilde{B}(t)$ change their sign).

Remark 3.1.11. The same commutative ring (or two isomorphic rings) can carry very different cluster structures. One can construct two seed patterns whose sets of exchange matrices are disjoint from each other, yet the two rings generated by their respective sets of cluster variables are isomorphic. We will give concrete examples later on.

Remark 3.1.12. We will soon encounter many examples in which different vertices of the tree \mathbb{T}_n correspond to identical labeled or unlabeled seeds. In spite of that, the set \mathcal{X} of cluster variables will typically be infinite. Note that this does not preclude a cluster algebra \mathcal{A} from being finitely generated (which is often the case). We shall also see in Section 6.8 (see Example 6.8.13) that even when \mathcal{X} is finite, the exchange relations (3.1.1) do not always generate the defining ideal of \mathcal{A} , i.e. the ideal of all relations satisfied by the cluster variables \mathcal{X} .

3.2. Examples of rank 1 and 2

The *rank* of a cluster algebra (or its underlying seed pattern) is the cardinality of each of its clusters (denoted above by n). In this section, we look at some examples of cluster algebras of small rank.

Rank 1. This case is very simple. The tree \mathbb{T}_1 has two vertices, so we only have two seeds, and two clusters (x_1) and (x'_1) . The extended exchange matrix \tilde{B}_\circ can be any $m \times 1$ matrix whose top entry is 0. The single exchange relation has the form $x_1 x'_1 = M_1 + M_2$ where M_1 and M_2 are monomials in the frozen variables x_2, \dots, x_m which do not share a common factor x_i . The cluster algebra is generated by $x_1, x'_1, x_2, \dots, x_m$, subject to this relation, and lies inside the ambient field $\mathcal{F} = \mathbb{C}(x_1, x_2, \dots, x_m)$.

Simple as they might be, cluster algebras of rank 1 do arise “in nature,” cf. Examples 1.1.2 and 1.1.3. We give two more examples here.

Example 3.2.1 (cf. Section 1.3). Let $U \subset G = \mathrm{SL}_3(\mathbb{C})$ be the subgroup of unipotent lower triangular 3×3 matrices. The ring $\mathbb{C}[G]^U$ is generated by the six flag minors P_J , for J a nonempty proper subset of $\{1, 2, 3\}$. This ring has the structure of a cluster algebra of rank 1 in which

- the ambient field is $\mathbb{C}(P_1, P_2, P_3, P_{12}, P_{23})$;
- the frozen variables are P_1, P_3, P_{12}, P_{23} ;
- the cluster variables are P_2 and P_{13} ;
- the single exchange relation is $P_2 P_{13} = P_1 P_{23} + P_3 P_{12}$.

The two seeds of this cluster algebra correspond to the two wiring diagrams with 3 strands. Their respective sets of chamber minors are the two extended clusters $\{P_2, P_1, P_3, P_{12}, P_{23}\}$ (cf. Figure 2.4) and $\{P_{13}, P_1, P_3, P_{12}, P_{23}\}$.

Example 3.2.2 (cf. Example 1.1.3). The coordinate ring of the subgroup U^+ of unipotent upper-triangular 3×3 matrices

$$\begin{bmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{bmatrix} \in \mathrm{SL}_3(\mathbb{C})$$

is $\mathbb{C}[a, b, c]$. This ring has the structure of a cluster algebra of rank 1 in which

- the ambient field is $\mathcal{F} = \mathbb{C}(a, b, c)$;
- the frozen variables are b and $P = ac - b$;
- the cluster variables are a and c ;
- the single exchange relation is $ac = P + b$.

We will discuss other examples in Chapter 6.

Rank 2. Any 2×2 skew-symmetrizable matrix looks like this:

$$(3.2.1) \quad \pm \begin{bmatrix} 0 & b \\ -c & 0 \end{bmatrix},$$

for some integers b and c which are either both positive, or both equal to 0. Applying a mutation μ_1 or μ_2 to a matrix of the form (3.2.1) simply changes its sign.

Example 3.2.3. In the case $b = c = 0$, the two mutations commute, because each μ_k changes the sign of the entries in column k of the extended exchange matrix while leaving the other column untouched; as the two matrix columns do not affect each other, the story reduces to the rank 1 case. We get four cluster variables x_1, x_2, x'_1, x'_2 , four clusters (x_1, x_2) , (x'_1, x_2) , (x_1, x'_2) , and (x'_1, x'_2) , and two exchange relations of the form $x_1 x'_1 = M_1 + M_2$ and $x_2 x'_2 = M_3 + M_4$, where M_1, M_2, M_3, M_4 are monomials in the frozen variables.

For the rest of this section, we assume that $b > 0$ and $c > 0$. We denote the cluster variables in our cluster algebra \mathcal{A} of rank 2 by

$$\dots, z_{-2}, z_{-1}, z_0, z_1, z_2, \dots,$$

so that the seed pattern looks like this:

$$\dots \xrightarrow{1} \begin{pmatrix} z_1 & z_0 \\ 0 & -b \\ c & 0 \end{pmatrix} \xrightarrow{2} \begin{pmatrix} z_1 & z_2 \\ 0 & b \\ -c & 0 \end{pmatrix} \xrightarrow{1} \begin{pmatrix} z_3 & z_2 \\ 0 & -b \\ c & 0 \end{pmatrix} \xrightarrow{2} \begin{pmatrix} z_3 & z_4 \\ 0 & b \\ -c & 0 \end{pmatrix} \xrightarrow{1} \dots$$

where we placed each cluster on top of the corresponding exchange matrix. (The extended exchange matrix may have additional rows.)

We denote by $\mathcal{A} = \mathcal{A}(b, c)$ a cluster algebra of rank 2 which has exchange matrices $\pm \begin{bmatrix} 0 & b \\ -c & 0 \end{bmatrix}$ and no frozen variables. (Cluster algebras without frozen variables are generally said to have *trivial coefficients*.) The exchange relations in $\mathcal{A}(b, c)$ are, in the notation introduced above:

$$(3.2.2) \quad z_{k-1} z_{k+1} = \begin{cases} z_k^c + 1 & \text{if } k \text{ is even;} \\ z_k^b + 1 & \text{if } k \text{ is odd.} \end{cases}$$

Example 3.2.4. The cluster variables in the cluster algebra $\mathcal{A}(1, 1)$ with trivial coefficients satisfy the recurrence

$$(3.2.3) \quad z_{k-1} z_{k+1} = z_k + 1.$$

Expressing everything in terms of the initial cluster (z_1, z_2) , we get:

$$(3.2.4) \quad z_3 = \frac{z_2 + 1}{z_1}, \quad z_4 = \frac{z_1 + z_2 + 1}{z_1 z_2}, \quad z_5 = \frac{z_1 + 1}{z_2}, \quad z_6 = z_1, \quad z_7 = z_2, \dots,$$

so the sequence is 5-periodic! Thus in this case, we have only 5 distinct cluster variables. In the seed pattern, we will have:

$$\dots \xrightarrow{2} \begin{pmatrix} z_1 & z_2 \\ 0 & 1 \\ -1 & 0 \end{pmatrix} \xrightarrow{1} \begin{pmatrix} z_3 & z_2 \\ 0 & -1 \\ 1 & 0 \end{pmatrix} \xrightarrow{2} \dots \xrightarrow{1} \begin{pmatrix} z_7 & z_6 \\ 0 & -1 \\ 1 & 0 \end{pmatrix} \xrightarrow{2} \dots$$

Note that even though the labeled seeds containing the clusters (z_1, z_2) and (z_7, z_6) are different, the corresponding unlabeled seeds coincide. Just switch z_6 and z_7 , and interchange the rows and the columns in the associated exchange matrix. Thus, this exchange pattern has 5 distinct (unlabeled) seeds.

Remark 3.2.5. The recurrence (3.2.3) arises in different mathematical contexts such as dilogarithm identities (cf., e.g., bibliographical pointers in [19, Section 1.1]), the Napier-Gauss *Pentagramma Mirificum* (cf. [8] and [11, Section 12.7]) and Coxeter’s frieze patterns [10].

Example 3.2.6. We now keep the same exchange matrices but introduce a single frozen variable y . Consider a seed pattern which looks like this:

$$\dots \begin{matrix} z_1 & z_2 & & z_3 & z_2 & & z_3 & z_4 & & z_5 & z_4 & & z_5 & z_6 & & z_7 & z_6 \\ \begin{bmatrix} 0 & 1 \\ -1 & 0 \\ p & q \end{bmatrix} & \xrightarrow{1} & \begin{bmatrix} 0 & -1 \\ 1 & 0 \\ -p & p+q \end{bmatrix} & \xrightarrow{2} & \begin{bmatrix} 0 & 1 \\ -1 & 0 \\ q & -p-q \end{bmatrix} & \xrightarrow{1} & \begin{bmatrix} 0 & -1 \\ 1 & 0 \\ -q & -p \end{bmatrix} & \xrightarrow{2} & \begin{bmatrix} 0 & 1 \\ -1 & 0 \\ -q & p \end{bmatrix} & \xrightarrow{1} & \begin{bmatrix} 0 & -1 \\ 1 & 0 \\ q & p \end{bmatrix} & \dots, \end{matrix}$$

where p and q are nonnegative integers. Relabeling the rows and columns to keep the 2×2 exchange matrices invariant, we get

$$\dots \begin{matrix} z_1 & z_2 & & z_2 & z_3 & & z_3 & z_4 & & z_4 & z_5 & & z_5 & z_6 & & z_6 & z_7 \\ \begin{bmatrix} 0 & 1 \\ -1 & 0 \\ p & q \end{bmatrix} & - & \begin{bmatrix} 0 & 1 \\ -1 & 0 \\ p+q & -p \end{bmatrix} & - & \begin{bmatrix} 0 & 1 \\ -1 & 0 \\ q & -p-q \end{bmatrix} & - & \begin{bmatrix} 0 & 1 \\ -1 & 0 \\ -p & -q \end{bmatrix} & - & \begin{bmatrix} 0 & 1 \\ -1 & 0 \\ -q & p \end{bmatrix} & - & \begin{bmatrix} 0 & 1 \\ -1 & 0 \\ p & q \end{bmatrix} & \dots, \end{matrix}$$

so the sequence of extended exchange matrices remains 5-periodic. We then compute the cluster variables:

$$z_3 = \frac{z_2 + y^p}{z_1}, \quad z_4 = \frac{y^{p+q}z_1 + z_2 + y^p}{z_1z_2}, \quad z_5 = \frac{y^qz_1 + 1}{z_2}, \quad z_6 = z_1, \quad z_7 = z_2;$$

the 5-periodicity persists! Just as in the case of trivial coefficients, there are five distinct cluster variables overall, and five distinct unlabeled seeds.

The above computations were based on the assumption that both entries in the third row of the initial extended exchange matrix are nonnegative. In fact, this condition is not required for 5-periodicity. Note that we could start with an initial seed containing the cluster (z_i, z_{i+1}) , for any $i \in \{1, 2, 3, 4, 5\}$, together with the associated extended exchange matrix in the relabeled sequence above, and we would get the same 5-periodic behaviour. Since any row vector in \mathbb{Z}^2 has the form (p, q) , $(p + q, -p)$, $(q, -p - q)$, $(-p, -q)$, or $(-q, p)$, for some $p, q \geq 0$ (see Figure 3.3), we conclude that any seed pattern with extended exchange matrices of the form $\pm \begin{bmatrix} 0 & 1 \\ * & * \end{bmatrix}$ has exactly five seeds.

As we shall later see, the general case of a seed pattern with exchange matrices $\pm \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$ and an arbitrary number of frozen variables exhibits the same qualitative behaviour: there will still be five cluster variables and five seeds.




Figure 3.3. Five types of “frozen rows” in extended exchange matrices with top rows $(0, 1)$ and $(-1, 0)$. Within each of the five cones, the points are parameterized by $p, q \geq 0$.

Example 3.2.7. The cluster variables in the cluster algebra $\mathcal{A}(1, 2)$ satisfy the recurrence

$$(3.2.5) \quad z_{k-1}z_{k+1} = \begin{cases} z_k^2 + 1 & \text{if } k \text{ is even;} \\ z_k + 1 & \text{if } k \text{ is odd.} \end{cases}$$

Expressing everything in terms of the initial cluster (z_1, z_2) , we get:

$$z_3 = \frac{z_2^2 + 1}{z_1}, \quad z_4 = \frac{z_2^2 + z_1 + 1}{z_1z_2}, \quad z_5 = \frac{z_1^2 + z_2^2 + 2z_1 + 1}{z_1z_2^2}, \quad z_6 = \frac{z_1 + 1}{z_2},$$

and then $z_7 = z_1$ and $z_8 = z_2$, so the sequence is 6-periodic! Thus in this case, we have only 6 distinct cluster variables, and 6 distinct seeds.

Exercise 3.2.8. Compute the cluster variables for the cluster algebra with the initial extended exchange matrix $\begin{bmatrix} 0 & 1 \\ -2 & 0 \\ p & q \end{bmatrix}$.

Exercise 3.2.9. Compute the cluster variables of the cluster algebra $\mathcal{A}(1, 3)$. (Start by evaluating them in the specialization $z_1 = z_2 = 1$; notice that all the numbers will be integers.)

Example 3.2.10. Consider the cluster algebra $\mathcal{A}(1, 4)$. Setting $z_1 = z_2 = 1$ and applying the recurrence (3.2.2), we see that the cluster variables z_3, z_4, \dots specialize to the following values:

$$2, 3, 41, 14, 937, 67, 21506, 321, 493697, 1538, 11333521, 7369, 260177282, \dots$$

This does not look like a periodic sequence, so the (unspecialized) sequence of cluster variables is not periodic either!

The good news is that all these numbers are integers. Why does this happen? To understand this, let us recursively compute the cluster variables z_3, z_4, \dots in terms of z_1 and z_2 :

$$\begin{aligned} z_3 &= \frac{z_2^4 + 1}{z_1}, \\ z_4 &= \frac{z_3 + 1}{z_2} = \frac{z_2^4 + z_1 + 1}{z_1 z_2}, \\ z_5 &= \frac{z_4^4 + 1}{z_3} \\ &= \frac{z_2^{12} + 4z_1 z_2^8 + 3z_2^8 + 6z_1^2 z_2^4 + 8z_1 z_2^4 + z_1^4 + 3z_2^4 + 4z_1^3 + 6z_1^2 + 4z_1 + 1}{z_1^3 z_2^4}, \\ z_6 &= \frac{z_5 + 1}{z_4} = \frac{z_2^8 + 3z_1 z_2^4 + 2z_2^4 + z_1^3 + 3z_1^2 + 3z_1 + 1}{z_1^2 z_2^3}, \text{ etc.} \end{aligned}$$

Now we see what is going on: the evaluations of these expressions at $z_1 = z_2 = 1$ are integers because they are *Laurent polynomials* in z_1 and z_2 , i.e., their denominators are monomials. (This is by no means to be expected: for example, the computation of z_6 involves dividing by $z_4 = \frac{z_2^4 + z_1 + 1}{z_1 z_2}$.)

3.3. Laurent phenomenon

The examples of Laurentness that we have seen before are special cases of the following general phenomenon.

Theorem 3.3.1. *In a cluster algebra of geometric type, each cluster variable can be expressed as a Laurent polynomial with integer coefficients in the elements of any extended cluster.*

The rest of this section is devoted to the proof of Theorem 3.3.1. First, we state a simple auxiliary lemma which can be obtained by direct inspection of the exchange relations (3.1.1).

Lemma 3.3.2. *Let \tilde{B}_\circ be an $m \times n$ extended exchange matrix. Let \tilde{B}'_\circ be the matrix obtained from \tilde{B}_\circ by deleting the rows labeled by a subset $I \subset \{n+1, \dots, m\}$. Then the formulas expressing the cluster variables in a cluster algebra $\mathcal{A}(\tilde{\mathbf{x}}'_\circ, \tilde{B}'_\circ)$ in terms of the initial extended cluster $\tilde{\mathbf{x}}'_\circ$ can be obtained from their counterparts for $\mathcal{A}(\tilde{\mathbf{x}}_\circ, \tilde{B}_\circ)$ by specializing the frozen variables x_i ($i \in I$) to 1, and relabeling the remaining variables accordingly.*

Remark 3.3.3. A specialization of the kind described in Lemma 3.3.2 sends Laurent polynomials to Laurent polynomials. This means that if we add extra frozen variables to the initial seed and establish Laurentness of an arbitrary cluster variable in this modified setting, it would then imply the Laurentness of the cluster variable's counterpart in the original setting.

Let us set up the notation needed for the proof of Theorem 3.3.1:

- $t_\circ \in \mathbb{T}_n$ is an (arbitrarily chosen) initial vertex;
- $(\tilde{\mathbf{x}}_\circ, \tilde{B}_\circ)$ is the initial seed;
- $\tilde{\mathbf{x}}_\circ = (x_1, \dots, x_m)$ is the initial extended cluster;
- $\tilde{B}_\circ = (b_{ij}^0)$ is the initial $m \times n$ extended exchange matrix;
- $t \in \mathbb{T}_n$ is an arbitrary vertex;
- $x \in \mathbf{x}(t)$ is a cluster variable at t ;
- t_1 and t_2 are the first two vertices on the unique path in \mathbb{T}_n connecting t_\circ to t , obtained via mutations in direction j then k , so that locally we have $t_\circ \xrightarrow{j} t_1 \xrightarrow{k} t_2$;
- d is the length of this path, i.e., the distance in \mathbb{T}_n between t_\circ and t ;
- $\tilde{\mathbf{x}}(t_1) = (\tilde{\mathbf{x}}(t_\circ) - \{x_j\}) \cup \{x'_j\}$;
- $\tilde{\mathbf{x}}(t_2) = (\tilde{\mathbf{x}}(t_1) - \{x_k\}) \cup \{x'_k\}$.

We will prove the Laurentness of x , viewed as a function of \mathbf{x}_\circ , by induction on d . (More precisely, the statement we prove by induction concerns

arbitrary seeds at distance d from each other in arbitrary cluster algebras of geometric type.) The base cases $d = 1$ and $d = 2$ are trivial.

There are two possibilities to consider.

Case 1: $b_{jk}^0 = b_{kj}^0 = 0$. Let t_3 be the vertex in \mathbb{T}_n connected to t_o by an edge labeled k . Since μ_j and μ_k commute at t_o (cf. Exercise 2.7.7(5)), each of the two seeds at t_1 and t_3 , respectively, lies at distance $d - 1$ from a seed containing x , and $\tilde{\mathbf{x}}(t_3) = (\tilde{\mathbf{x}}(t_o) - \{x_k\}) \cup \{x'_k\}$.

By the induction assumption, the cluster variable x is expressed as a Laurent polynomial in terms of the extended cluster $\tilde{\mathbf{x}}(t_1) = (x_1, \dots, x'_j, \dots, x_m)$. Also, $x'_j = \frac{M_1 + M_2}{x_j}$, where M_1 and M_2 are monomials in x_1, \dots, x_m . Substituting this into the aforementioned Laurent polynomial, we obtain a formula expressing x in terms of $\tilde{\mathbf{x}}_o$. Another such formula is obtained by taking the Laurent polynomial expression for x in terms of $\tilde{\mathbf{x}}(t_3)$, and substituting $x'_k = \frac{M_3 + M_4}{x_k}$, with M_3 and M_4 some monomials in x_1, \dots, x_m . Removing common factors, we obtain (necessarily identical) expressions for x as a ratio of coprime polynomials in x_1, \dots, x_m , with monic denominator.

Note that in the first computation, all non-monomial factors that can potentially remain in the denominator must come from $M_1 + M_2$; in the second one, they can only come from $M_3 + M_4$. If $M_1 + M_2$ and $M_3 + M_4$ were coprime to each other, the Laurentness of x would follow. This coprimality however does not hold in general. (For example, if columns j and k of \tilde{B}_o are equal to each other, then $M_1 + M_2 = M_3 + M_4$.) We can however use a trick based on Lemma 3.3.2, cf. Remark 3.3.3. Let us introduce a new frozen variable x_{m+1} and extend the matrix \tilde{B}_o by an extra row in which the $(m + 1, j)$ -entry is 1, and all other entries are 0. Now $M_1 + M_2$ has become a binomial which has degree 1 in the variable x_{m+1} . Note that the supports of M_1 and M_2 are disjoint, so $M_1 + M_2$ does not have a monomial factor, and if it had a nontrivial factorization, that factorization would give a nontrivial factorization of $1 + x_{m+1}$ after specializing the other cluster variables to 1. Therefore $M_1 + M_2$ is irreducible; moreover it cannot divide $M_3 + M_4$ as the latter does not depend on x_{m+1} . So $M_1 + M_2$ and $M_3 + M_4$ are coprime to each other, and we are done with Case 1.

Case 2: $b_{jk}^0 b_{kj}^0 < 0$. This case is much harder. The general shape of the proof remains the same: we use induction on d together with a coprimality argument assisted by the introduction of additional frozen variables. One new aspect of the proof is that we need to separately consider the case $d = 3$ since the induction step relies on it.

Without loss of generality we assume that $b_{jk}^0 < 0$ and $b_{kj}^0 > 0$. Otherwise, change the signs of all extended exchange matrices; this will not affect the formulas relating extended clusters to each other, see Exercise 2.7.7(3).

We denote by $t_3 \in \mathbb{T}_n$ the vertex connected to t_2 by an edge labeled j , and introduce notation

$$\tilde{\mathbf{x}}(t_3) = (\tilde{\mathbf{x}}(t_2) - \{x'_j\}) \cup \{x''_j\} = (x_1, \dots, x''_j, \dots, x'_k, \dots, x_m).$$

(Whether $j < k$ or $k < j$ is immaterial.) See Figure 3.4.

To be clear, we are *not* assuming that t_3 lies on the unique path in T_n connecting t_0 to t .




Figure 3.4. Cluster variables obtained via successive mutations μ_j, μ_k, μ_j .

Note that among cluster variables obtained by at most three mutations from the initial seed, those of the form x''_j are the only ones whose Laurentness is not obvious. Therefore to establish Case 2 when $d = 3$, it is enough to prove Lemma 3.3.4 below.

Lemma 3.3.4. *The cluster variable x''_j is a Laurent polynomial in $\tilde{\mathbf{x}}_0$.*

Proof. Let $\mu_j(\tilde{B}_0) = \tilde{B}(t_1) = (b_{ij})$ and $\mu_k(\tilde{B}(t_1)) = \tilde{B}(t_2) = (b'_{ij})$ be the extended exchange matrices at t_1 and t_2 , respectively. Our assumption $b^0_{jk} < 0$ implies that $b_{jk} > 0$ and $b'_{kj} > 0$.

We will view each of the cluster variables x'_j, x'_k, x''_j as a rational function in the elements x_1, \dots, x_m of the initial extended cluster $\tilde{\mathbf{x}}_0$. The notation $P \sim Q$ will mean that P and Q differ by a monomial factor, i.e., $P = QM$ where M is a Laurent monomial in x_1, \dots, x_m . Given a polynomial P , the notation $Q \equiv R \pmod{P}$ will mean that $Q - R = PS$ for some Laurent polynomial S . We also define

$$P_j = P_j(x_1, \dots, x_m) = \prod_i x_i^{b_{ij}} + 1.$$

The relevant instances of the exchange relation (3.1.1) imply that

$$(3.3.1) \quad x'_j \sim x_j^{-1} \left(\prod_i x_i^{b_{ij}} + 1 \right) \sim P_j,$$

$$(3.3.2) \quad x'_k = x_k^{-1} \left((x'_j)^{b_{jk}} \prod_{\substack{b_{ik} > 0 \\ i \neq j}} x_i^{b_{ik}} + \prod_{b_{ik} < 0} x_i^{-b_{ik}} \right) \equiv x_k^{-1} \prod_{b_{ik} < 0} x_i^{-b_{ik}} \pmod{P_j},$$

$$(3.3.3) \quad x''_j \sim (x'_j)^{-1} \left((x'_k)^{b'_{kj}} \prod_{i \neq k} x_i^{b'_{ij}} + 1 \right).$$

To establish that x''_j is a Laurent polynomial in x_1, \dots, x_m , we need to show that the second factor in (3.3.3) is divisible by P_j .

Working modulo P_j we obtain

$$\begin{aligned} (x'_k)^{b'_{kj}} \prod_{i \neq k} x_i^{b'_{ij}} + 1 &\equiv \left(x_k^{-1} \prod_{b_{ik} < 0} x_i^{-b_{ik}} \right)^{b'_{kj}} \prod_{i \neq k} x_i^{b'_{ij}} + 1 \\ &= x_k^{b_{kj}} \prod_{b_{ik} < 0} x_i^{b_{ik} b_{kj}} \prod_{i \neq k} x_i^{b'_{ij}} + 1 = \prod_i x_i^{b_{ij}} + 1 \equiv 0, \end{aligned}$$

as desired. In the last line, we used the fact that $b_{kj} < 0$ and consequently

$$b'_{ij} = \begin{cases} b_{ij} - b_{ik} b_{kj} & \text{if } b_{ik} < 0; \\ b_{ij} & \text{if } b_{ik} \geq 0 \text{ and } i \neq k. \quad \square \end{cases}$$

After proving the following technical lemma, we will be ready to complete the inductive proof of Case 2.

Lemma 3.3.5. *Suppose that distinct indices $q, r \in \{n + 1, \dots, m\}$ are such that $b_{qj}^0 = 1$ and $b_{rk}^0 = 1$, and moreover all other entries in rows q and r of \tilde{B}_\circ are equal to 0. Then x'_j is coprime to both x'_k and x''_j .*

Here “coprime” means that those cluster variables, viewed as Laurent polynomials in \mathbf{x}_\circ , have no common non-monomial factor.

Proof. Let us denote $b_{jk}^0 = -b$ and $b_{kj}^0 = c$. Recall that $b_{kj}^0 > 0$, so $b, c > 0$. The local structure of the extended exchange matrices at t_\circ, t_1 , and t_2 at the intersections of rows j, k, q, r and columns j, k is as follows:

$$\begin{array}{c|ccc|ccc|ccc} & & j & k & & & j & k & & & & j & k \\ \hline & & \vdots & \vdots & & & \vdots & \vdots & & & & \vdots & \vdots \\ j & \cdots & 0 & -b & \cdots & j & \cdots & 0 & b & \cdots & j & \cdots & 0 & -b & \cdots \\ k & \cdots & c & 0 & \cdots & k & \cdots & -c & 0 & \cdots & k & \cdots & c & 0 & \cdots \\ & & \vdots & \vdots & & & \vdots & \vdots & & & & \vdots & \vdots \\ q & \cdots & 1 & 0 & \cdots & q & \cdots & -1 & 0 & \cdots & q & \cdots & -1 & 0 & \cdots \\ r & \cdots & 0 & 1 & \cdots & r & \cdots & 0 & 1 & \cdots & r & \cdots & 0 & -1 & \cdots \\ & & \tilde{B}_\circ & & & & \mu_j(\tilde{B}) & & & & & \mu_k(\mu_j(\tilde{B})) & & & \end{array}$$

We then have

$$\begin{aligned} x'_j &= x_j^{-1} (x_k^c x_q M_1 + M_2) \\ x'_k &= x_k^{-1} ((x'_j)^b x_r M_3 + M_4) \\ x''_j &= (x'_j)^{-1} (x_q M_5 + (x'_k)^c M_6), \end{aligned}$$

where M_1, \dots, M_6 are monomials in the x_i 's, with $i \notin \{j, k, q, r\}$. We see that x'_j is linear in x_q and hence irreducible (as a Laurent polynomial in $\tilde{\mathbf{x}}$), i.e., it cannot be written as a product of two non-monomial factors. Since x'_j does not depend on x_r , we conclude that x'_k is linear in x_r and hence irreducible, and moreover coprime with x'_j .

It remains to show that x'_j and x''_j are coprime. Note that x'_k and x''_j can be regarded as polynomials in x_r ; we denote by $x'_k(0)$ and $x''_j(0)$ their specializations at $x_r = 0$. If we show that $x''_j(0)$ is coprime to $x'_j = x'_j(0)$, then we'll be done. To this end, note that $x'_k(0) = x_k^{-1}M_4$. Therefore

$$x''_j(0) = x_j \frac{x_q M_5 + (x_k^{-1} M_4)^c M_6}{x_k^c x_q M_1 + M_2}.$$

Here both the numerator and denominator are linear in x_q , and therefore the denominator (essentially, x'_j) cannot divide the numerator more than once. Also, x'_j is irreducible. This means that, in order for $x''_j(0)$ and x'_j to fail to be coprime, we would need the denominator of $x''_j(0)$ to divide the numerator at least twice. Hence $x''_j(0)$ and x'_j are coprime, as desired. \square

We are now ready to complete the proof of Case 2 of Theorem 3.3.1. We begin by augmenting the initial extended exchange matrix \tilde{B}_o by two additional rows corresponding to two new frozen variables x_q and x_r . We set

$$b_{qi}^0 = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j; \end{cases} \quad b_{ri}^0 = \begin{cases} 1 & \text{if } i = k, \\ 0 & \text{if } i \neq k, \end{cases}$$

so as to satisfy the conditions of Lemma 3.3.5.

By the induction assumption, the cluster variable x is expressed as a Laurent polynomial in terms of each of the extended clusters $\tilde{\mathbf{x}}(t_1)$ and $\tilde{\mathbf{x}}(t_3)$. The only elements of these clusters which do not appear in $\tilde{\mathbf{x}}_o$ are x'_j , x'_k , and x''_j , so

$$x = \frac{\text{Laurent polynomial in } \tilde{\mathbf{x}}_o}{(x'_j)^a} = \frac{\text{Laurent polynomial in } \tilde{\mathbf{x}}_o}{(x'_k)^b (x''_j)^c},$$

for $a, b, c \in \mathbb{Z}$. By Lemma 3.3.4, x'_j , x'_k , and x''_j are Laurent polynomials in $\tilde{\mathbf{x}}_o$. By Lemma 3.3.5, x'_j is coprime to both x'_k and x''_j . The theorem now follows by the same argument (based on Lemma 3.3.2) that we used in Case 1. \square

Theorem 3.3.1 can be sharpened as follows.

Theorem 3.3.6. *In a cluster algebra of geometric type, frozen variables do not appear in the denominators of the Laurent polynomials expressing cluster variables in terms of an initial extended cluster.*

Stated in our standard notation, Theorem 3.3.6 asserts that each cluster variable is a Laurent polynomial in the initial cluster variables x_1, \dots, x_n , with coefficients in $\mathbb{Z}[x_{n+1}, \dots, x_m]$.

Proof. We borrow the notation from the proof of Theorem 3.3.1 above. Let x be a cluster variable from a distant seed, and x_r a frozen variable ($n < r \leq m$). We will think of x as a Laurent polynomial $x(x_r)$ whose coefficients are integral Laurent polynomials in the variables x_i , with $i \neq r$. We want to show that x is in fact a polynomial in x_r ; Theorem 3.3.6 will then follow by varying r .

We will make use of the following trivial lemma.

Lemma 3.3.7. *Let P and Q be two polynomials (in any number of variables) with coefficients in a domain S , and with nonzero constant terms a and b , respectively. If the ratio P/Q is a Laurent polynomial over S , then it is in fact a polynomial over S with the constant term a/b .*

Our proof of Theorem 3.3.6 proceeds by induction on d , the smallest distance in \mathbb{T}_n between the initial seed and a seed containing x . We will inductively prove the following strengthening of the desired statement:

$x(x_r)$ is a polynomial in x_r whose constant term $x(0)$ can be written as a subtraction-free rational expression in the elements of $\tilde{\mathbf{x}}_\circ - \{x_r\}$; in particular, $x(0) \neq 0$.

If $d = 0$, then $x \in \tilde{\mathbf{x}}_\circ$, and there is nothing to prove. If $d > 0$, then x appears on the left-hand side of an exchange relation (3.1.1) $xx' = M_1 + M_2$, where M_1 and M_2 denote the monomials in the exchange relation, and where x' and the cluster variables in M_1 and M_2 come from a seed located at distance $d - 1$ from $\tilde{\mathbf{x}}_\circ$. By definition of the exchange relation, the frozen variable x_r will appear in at most one of M_1 and M_2 . Therefore if we express x' , M_1 , and M_2 in terms of $\tilde{\mathbf{x}}_\circ$, viewing them as Laurent polynomials in x_r (whose coefficients are integral Laurent polynomials in the variables x_i with $i \neq r$), the inductive assumption implies that $M_1 + M_2$ is a polynomial in x_r with nonzero constant term. It now follows from Lemma 3.3.7 that $x = \frac{M_1 + M_2}{x'}$ is a polynomial in x_r with nonzero constant term. \square

3.4. Connections to number theory

Example 3.4.1 (*Markov triples*). Consider the cluster algebra defined by the Markov quiver given in Figure 2.10. Since the quiver is invariant under mutations, exchange relations for any cluster (x_1, x_2, x_3) will look the same:

$$\begin{aligned}x'_1 x_1 &= x_2^2 + x_3^2, \\x'_2 x_2 &= x_1^2 + x_3^2, \\x'_3 x_3 &= x_1^2 + x_2^2.\end{aligned}$$

If we start with the triple $(1, 1, 1)$ and mutate in all possible directions, we will get an infinite set of triples in \mathbb{Z}^3 , including those shown in Figure 3.5.




Figure 3.5. Markov triples.

We next observe that all these triples satisfy the diophantine *Markov equation*

$$x_1^2 + x_2^2 + x_3^2 = 3x_1x_2x_3.$$

To see this, verify that each mutation in our cluster algebra transforms a solution of this equation (a *Markov triple*) into another solution. (This is an instance of *Vieta jumping*, which replaces one root of a quadratic equation by another root.) Starting with the solution $(1, 1, 1)$, we get the tree of Markov triples above. In fact, every Markov triple appears in this tree. The celebrated (still open) Uniqueness Conjecture asserts that the maximal elements of Markov triples are all distinct. See [1] for a detailed account.

Returning to the cluster-algebraic interpretation of this example, we note that more generally, the quantity

$$\frac{x_1^2 + x_2^2 + x_3^2}{x_1 x_2 x_3}$$

is invariant under mutations in this seed pattern. This is closely related to *integrability* of the Markov recurrence; cf. Chapter 11.

Example 3.4.2 (*Fermat numbers*). Sometime in the 1640's, Pierre Fermat conjectured that for every positive integer n , the number $F_n = 2^{2^n} + 1$ is prime. This was disproved in 1732 by Leonhard Euler, who discovered that

$$F_5 = 2^{32} + 1 = 641 \cdot 6700417.$$

Curiously, this factorization can be obtained by observing cluster mutations.

Consider the rank 2 cluster algebra with the initial seed $(\tilde{\mathbf{x}}, \tilde{B})$ where

$$\tilde{\mathbf{x}} = (x_1, x_2, x_3), \quad \tilde{B} = \begin{bmatrix} 0 & 4 \\ -1 & 0 \\ 1 & -3 \end{bmatrix}.$$

The mutation μ_1 produces a new extended cluster $\tilde{\mathbf{x}}' = (x'_1, x_2, x_3)$ where

$$x'_1 x_1 = x_2 + x_3.$$

Taking the specialization

$$(x_1, x_2, x_3) = (3, -1, 16),$$

we see that the mutated extended cluster specializes to

$$(x'_1, x_2, x_3) = (5, -1, 16).$$

Applying the sharp version of the Laurent phenomenon (Theorems 3.3.1 and 3.3.6) to the initial extended cluster $\tilde{\mathbf{x}}$, we see that every cluster variable specializes to an integer (possibly) divided by a power of 3; applying the same result to $\tilde{\mathbf{x}}'$, we conclude that every cluster variable is an integer (possibly) divided by a power of 5. Thus, every cluster variable specializes to an integer! Now let us see which integers we get. Alternately applying the mutations μ_1 and μ_2 , we obtain the following sequence, with specialized cluster variables written on top of the extended exchange matrices:

$$\begin{array}{c} \begin{array}{cc|c} 3 & -1 & \\ \hline 0 & 4 & \mu_1 \rightarrow \end{array} \quad \begin{array}{cc|c} 5 & -1 & \\ \hline 0 & -4 & \mu_2 \rightarrow \end{array} \quad \begin{array}{cc|c} 5 & -641 & \\ \hline 0 & 4 & \mu_1 \rightarrow \end{array} \quad \begin{array}{cc|c} -128 & -641 & \\ \hline 0 & -4 & \mu_2 \rightarrow \end{array} \quad \begin{array}{cc|c} -128 & \frac{-F_5}{641} & \\ \hline 0 & 4 & \end{array} \\ -1 & 0 & 1 & 0 & -1 & 0 & 1 & 0 & -1 & 0 & 0 \\ 1 & -3 & -1 & 1 & 0 & -1 & 0 & -1 & 0 & 0 & 1 \end{array}$$

This shows that $F_5/641$ is an integer, reproducing Euler's discovery.

Example 3.4.3. The *Somos-4 sequence* z_0, z_1, z_2, \dots is defined by the initial conditions $z_0 = z_1 = z_2 = z_3 = 1$ and the recurrence

$$z_{m+2}z_{m-2} = z_{m+1}z_{m-1} + z_m^2.$$


This sequence is named after M. Somos who discovered it (and its various generalizations) sometime in the 1980s; see, e.g., [6, 31] and references therein.

The first several terms of the Somos-4 sequence are

$$1, 1, 1, 1, 2, 3, 7, 23, 59, 314, 1529, 8209, 83313, 620297, 7869898, \dots$$

—all integers! An explanation of the integrality of this sequence can be given using cluster algebras.

Consider the quiver



(all four vertices are mutable). It is easy to check that mutating at the vertex labeled 1 produces a quiver that differs from the original one by clockwise rotation by $\pi/2$. It follows that subsequent quiver mutations at $2, 3, 4, 1, 2, 3, 4, 1, \dots$ will generate a sequence of cluster variables satisfying the Somos-4 recurrence above. In view of the Laurent phenomenon, the initial conditions $z_0 = z_1 = z_2 = z_3 = 1$ will result in a sequence of integers.

Remark 3.4.4. An alternative approach to establishing integrality of the Somos-4 and other related sequences is based on explicit combinatorial interpretations of their terms. In particular, the numbers z_m defined above can be shown to count *perfect matchings* in certain planar bipartite graphs, see [49] and Figure 3.6.

Remark 3.4.5. Somos sequences and their various generalizations are intimately related to the arithmetic of elliptic curves; see, e.g., [32] and references therein. Here is a typical result, stated here without proof (this version is due to D. Speyer). Consider the elliptic curve

$$y^2 = 1 - 8x + 12x^2 - 4x^3,$$

and let x_m be the x -coordinate of the point $P + mQ$ where $P = (0, 1)$, $Q = (1, -1)$, and we are using the standard group law of the elliptic curve. Then

$$x_m = \frac{z_{m-1}z_{m+1}}{z_m^2}$$

where (z_m) is the Somos-4 sequence above.




Figure 3.6. The number of perfect matchings in each of these bipartite graphs is 2, 3, 7, 23, respectively (cf. the Somos-4 sequence).

Exercise 3.4.6. Use the Laurent phenomenon to show that the sequence z_0, z_1, z_2, \dots defined by the initial conditions $z_0 = z_1 = z_2 = 1$ and the recurrence

$$z_{m+3}z_m = z_{m+2}z_{m+1} + 1$$

consists entirely of integers.

Exercise 3.4.7. Show that all elements of the sequence z_0, z_1, z_2, \dots defined by the generalized Somos-4 recurrence

$$z_{m+2}z_{m-2} = az_{m+1}z_{m-1} + bz_m^2$$

are Laurent polynomials in z_0, z_1, z_2, z_3 , with coefficients in $\mathbb{Z}[a, b]$.

Example 3.4.8. The *Somos-5* sequence

$$1, 1, 1, 1, 1, 2, 3, 5, 11, 37, 83, 274, 1217, 6161, 22833, 165713, \dots$$

is defined by the recurrence relation

$$(3.4.1) \quad z_m z_{m+5} = z_{m+1} z_{m+4} + z_{m+2} z_{m+3} \quad (m = 1, 2, \dots)$$


with the initial conditions $z_1 = \dots = z_5 = 1$. *A priori*, one expects the numbers z_m to be rational—but in fact, all of them are integers. Once again, this is a consequence of a stronger statement: viewed as a function of z_1, \dots, z_5 , every z_m is a Laurent polynomial with integer coefficients. To prove this, we need to find a cluster algebra with an initial cluster (z_1, \dots, z_5) (no frozen variables) which has all relations (3.4.1) among its exchange relations, so that all the z_m are among its cluster variables.

Exercise 3.4.9. Establish the integrality of all terms of the Somos-5 sequence by examining the sequence of mutations

$$\mu_1, \mu_2, \mu_3, \mu_4, \mu_5, \mu_1, \mu_2, \mu_3, \mu_4, \mu_5, \dots,$$

in the cluster algebra whose initial exchange matrix (and the corresponding quiver) are shown below.

$$B = \tilde{B} = \begin{bmatrix} 0 & -1 & 1 & 1 & -1 \\ 1 & 0 & -2 & 0 & 1 \\ -1 & 2 & 0 & -2 & 1 \\ -1 & 0 & 2 & 0 & -1 \\ 1 & -1 & -1 & 1 & 0 \end{bmatrix}$$



Example 3.4.10 (see [26]). Fix a positive integer n , and let (a_1, \dots, a_{n-1}) be a *palindromic* integer vector, that is, $a_i = a_{n-i}$ for $i = 1, \dots, n-1$. Consider the sequence z_1, z_2, \dots given by the recurrence

$$(3.4.2) \quad z_m z_{m+n} = \prod_{i=1}^{n-1} z_{m+i}^{[a_i]_+} + \prod_{i=1}^{n-1} z_{m+i}^{[-a_i]_+} \quad (m = 1, 2, \dots)$$

with indeterminates z_1, \dots, z_n as the initial terms; here we use the notation

$$(3.4.3) \quad [a]_+ = \max(a, 0).$$

(The recurrence (3.4.1) is a special case with $n = 5$ and $(a_1, \dots, a_{n-1}) = (1, -1, -1, 1)$.) Then all the terms z_m are integer Laurent polynomials in z_1, \dots, z_n .

To show this, we find an $n \times n$ skew-symmetric integer matrix $B = (b_{ij})$ such that $\mu_1(B)$ is obtained from B by the cyclic permutation of its rows and columns, and such that its first column is given by $b_{i1} = a_{i-1}$ for $i = 2, \dots, n$. If we can do so, then the sequence of mutations

$$\mu_1, \mu_2, \mu_3, \dots, \mu_n, \mu_1, \mu_2, \dots,$$

will produce the recurrence (3.4.2). It is not hard to see that setting

$$-b_{ji} = b_{ij} = a_{i-j} + \sum_{k=1}^{j-1} ([-a_{i-k}]_+ [a_{j-k}]_+ - [a_{i-k}]_+ [-a_{j-k}]_+)$$

for $1 \leq j < i \leq n$ produces a matrix B with the requisite properties.

3.5. Y -patterns

We keep the notational conventions used in Section 3.3.

One of our goals is to show that many structural properties of a seed pattern are determined by (the mutation class formed by) its $n \times n$ exchange matrices $B(t)$, and do not depend on the bottom parts of the matrices $\tilde{B}(t)$.

Theorem 3.5.1 below concerns certain Laurent monomials in the elements of a given seed; each of these Laurent monomials is simply a ratio of the two terms appearing on the right-hand side of an exchange relation (3.1.1). Surprisingly, the evolution of these ratios is completely controlled by the matrices $B(t)$. That is, the laws governing this evolution do not depend on the bottom parts of the matrices $\tilde{B}(t)$.

Theorem 3.5.1. *Let $(\tilde{\mathbf{x}}, \tilde{B})$ and $(\tilde{\mathbf{x}}', \tilde{B}')$ be two labeled seeds related by mutation at k , with extended clusters*

$$\tilde{\mathbf{x}} = (x_1, \dots, x_m), \quad \tilde{\mathbf{x}}' = (x'_1, \dots, x'_m)$$

and $m \times n$ extended exchange matrices

$$\tilde{B} = (b_{ij}), \quad \tilde{B}' = (b'_{ij}).$$

Define the n -tuples $\hat{\mathbf{y}} = (\hat{y}_1, \dots, \hat{y}_n)$ and $\hat{\mathbf{y}}' = (\hat{y}'_1, \dots, \hat{y}'_n)$ by

$$(3.5.1) \quad \hat{y}_j = \prod_{i=1}^m x_i^{b_{ij}}, \quad \hat{y}'_j = \prod_{i=1}^m (x'_i)^{b'_{ij}}.$$

Then

$$(3.5.2) \quad \hat{y}'_j = \begin{cases} \hat{y}_k^{-1} & \text{if } j = k; \\ \hat{y}_j (\hat{y}_k + 1)^{-b_{kj}} & \text{if } j \neq k \text{ and } b_{kj} \leq 0; \\ \hat{y}_j (\hat{y}_k^{-1} + 1)^{-b_{kj}} & \text{if } j \neq k \text{ and } b_{kj} \geq 0. \end{cases}$$

Proof. We check (3.5.2) case by case. The case $j = k$ is easy:

$$\hat{y}'_k = \prod_i (x'_i)^{b'_{ik}} = \prod_{i \neq k} x_i^{b'_{ik}} = \prod_{i \neq k} x_i^{-b_{ik}} = \hat{y}_k^{-1}.$$

If $j \neq k$ and $b_{kj} \leq 0$, then

$$\begin{aligned} \hat{y}'_j &= (x'_k)^{b'_{kj}} \prod_{i \neq k} x_i^{b'_{ij}} = (x'_k)^{-b_{kj}} \prod_{i \neq k} x_i^{b_{ij}} \prod_{b_{ik} < 0} x_i^{-b_{ik} b_{kj}} \\ &= x_k^{b_{kj}} \left(\prod_{b_{ik} > 0} x_i^{b_{ik}} + \prod_{b_{ik} < 0} x_i^{-b_{ik}} \right)^{-b_{kj}} \prod_{i \neq k} x_i^{b_{ij}} \prod_{b_{ik} < 0} x_i^{-b_{ik} b_{kj}} \\ &= \hat{y}_j (\hat{y}_k + 1)^{-b_{kj}}. \end{aligned}$$

If $j \neq k$ and $b_{kj} \geq 0$, then we can check (3.5.2) directly as before; alternatively, it follows from the previous case, using the fact that mutation is an involution and switching the roles of the two seeds.

Theorem 3.5.1 suggests the following definitions.

Definition 3.5.2. A Y -seed of rank n in a field \mathcal{F} is a pair (Y, B) where

- Y is an n -tuple of elements of \mathcal{F} ;
- B is a skew-symmetrizable $n \times n$ integer matrix.

We say that two Y -seeds (Y, B) and (Y', B') of rank n are related by a Y -seed mutation μ_k in direction k (here $1 \leq k \leq n$) if

- the matrices $B = (b_{ij})$ and $B' = (b'_{ij})$ are related via mutation at k ;
- the n -tuple $Y' = (Y'_1, \dots, Y'_n)$ is obtained from $Y = (Y_1, \dots, Y_n)$ by

$$(3.5.3) \quad Y'_j = \begin{cases} Y_k^{-1} & \text{if } j = k; \\ Y_j (Y_k + 1)^{-b_{kj}} & \text{if } j \neq k \text{ and } b_{kj} \leq 0; \\ Y_j (Y_k^{-1} + 1)^{-b_{kj}} & \text{if } j \neq k \text{ and } b_{kj} \geq 0. \end{cases}$$

It is easy to check that mutating (Y', B') at k recovers (Y, B) .

A Y -pattern of rank n is a collection of Y -seeds $(Y(t), B(t))_{t \in \mathbb{T}_n}$ labeled by the vertices of the n -regular tree \mathbb{T}_n , such that for any edge $t \xrightarrow{k} t'$ in \mathbb{T}_n , the Y -seeds $(Y(t), B(t))$ and $(Y(t'), B(t'))$ are related to each other by the Y -seed mutation in direction k .

Remark 3.5.3. In Definition 3.5.2, we do not require the elements Y_i to be algebraically independent, one reason being that this condition does not always hold for the monomials \hat{y}_j in Theorem 3.5.1. Consequently, one can not *a priori* guarantee that the mutation process can propagate to all vertices in \mathbb{T}_n (what if $Y_k = 0$ in (3.5.3)?). To ensure the existence of a Y -pattern with a given initial seed (Y, B) , one can for example require all elements of Y to be given by subtraction-free expressions in some set of variables, or alternatively take positive values under a particular specialization of these variables. As each of these conditions reproduces under mutations of Y -seeds, the mutation process can then proceed without hindrance.

Example 3.5.4 (Y -pattern of type A_2). Consider the Y -pattern of rank 2

$$\dots \xrightarrow{2} (Y(0), B(0)) \xrightarrow{1} (Y(1), B(1)) \xrightarrow{2} (Y(2), B(2)) \xrightarrow{1} \dots$$

with the exchange matrices

$$(3.5.4) \quad B(t) = (-1)^t \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}.$$

Note that the corresponding quivers are orientations of the type A_2 Dynkin diagram.

The rule (3.5.3) of *Y*-seed mutation gives the following recurrence for the *Y*-seeds $Y(t) = (Y_{1;t}, Y_{2;t})$. For t even, we have

$$Y(t+1) = \mu_1(Y(t)), \quad Y_{1;t+1} = Y_{1;t}^{-1}, \quad Y_{2;t+1} = Y_{2;t} (Y_{1;t}^{-1} + 1)^{-1},$$

whereas for t odd, we have

$$Y(t+1) = \mu_2(Y(t)), \quad Y_{1;t+1} = Y_{1;t} (Y_{2;t}^{-1} + 1)^{-1}, \quad Y_{2;t+1} = Y_{2;t}^{-1}.$$

We then recursively obtain the pairs $Y(t) = (Y_{1;t}, Y_{2;t})$ listed in Figure 3.7.

t	$Y_{1;t}$	$Y_{2;t}$
0	y_1	y_2
1	y_1^{-1}	$y_1 y_2 (y_1 + 1)^{-1}$
2	$y_2 (y_1 y_2 + y_1 + 1)^{-1}$	$(y_1 + 1) y_1^{-1} y_2^{-1}$
3	$(y_1 y_2 + y_1 + 1) y_2^{-1}$	$y_1^{-1} (y_2 + 1)^{-1}$
4	y_2^{-1}	$y_1 (y_2 + 1)$
5	y_2	y_1
6	$y_1 y_2 (y_1 + 1)^{-1}$	y_1^{-1}
7	$(y_1 + 1) y_1^{-1} y_2^{-1}$	$y_2 (y_1 y_2 + y_1 + 1)^{-1}$
8	$y_1^{-1} (y_2 + 1)^{-1}$	$(y_1 y_2 + y_1 + 1) y_2^{-1}$
9	$y_1 (y_2 + 1)$	y_2^{-1}
10	y_1	y_2

Figure 3.7. The *Y*-seeds $(Y(t), B(t)) = ((Y_{1;t}, Y_{2;t}), B(t))$ in type A_2 . The exchange matrices $B(t)$ are given by (3.5.4). The initial *Y*-seed is $(Y(0), B(0))$, with $Y(0) = (Y_{1;0}, Y_{2;0}) = (y_1, y_2)$. This sequence of *Y*-seeds is 10-periodic: $Y(t+10) = Y(t)$.

Using the terminology introduced in Definition 3.5.2, we can state the following direct corollary of Theorem 3.5.1.

Corollary 3.5.5. *Let $(\tilde{\mathbf{x}}(t), \tilde{B}(t))_{t \in \mathbb{T}_n}$ be a seed pattern in \mathcal{F} , with*

$$\tilde{\mathbf{x}}(t) = (x_{1;t}, \dots, x_{m;t}), \quad \tilde{B}(t) = (b_{ij}^t).$$

Let $B(t) = (b_{ij}^t)_{i,j \leq n}$ denote the exchange matrix at a vertex $t \in \mathbb{T}_n$, and let $\hat{\mathbf{y}}(t) = (\hat{y}_{1;t}, \dots, \hat{y}_{n;t})$ be the n -tuple of elements in \mathcal{F} given by

$$(3.5.5) \quad \hat{y}_{k;t} = \prod_{i=1}^m x_{i;t}^{b_{ik}^t}.$$

*Then $(\hat{\mathbf{y}}(t), B(t))_{t \in \mathbb{T}_n}$ is a *Y*-pattern in \mathcal{F} .*

Remark 3.5.6. The rules governing the evolution of Y -seeds may seem simpler than the corresponding rules of seed mutation:

- Y -seed mutations are driven by the $n \times n$ matrices B whereas ordinary seed mutations require the extended $m \times n$ matrices \tilde{B} ;
- in the Y -seed setting, there are no frozen variables;
- each recurrence (3.5.3) only involves two variables Y_j and Y_k whereas the exchange relation (3.1.1) potentially involves all cluster variables of the current seed.

On the other hand,

- a seed mutation only changes one cluster variable whereas a Y -seed mutation may potentially change all the variables Y_1, \dots, Y_n ;
- consequently, we end up getting “more” Y -variables than cluster variables (if the number of seeds is finite, then this is a precise statement);
- the Y -pattern recurrences do not, generally speaking, exhibit the Laurent phenomenon.

Remark 3.5.7. In various examples, including many cluster algebras arising as coordinate rings of algebraic varieties, the cluster algebra under investigation has a distinguished (multi-)grading, and its exchange relations are all (multi-)homogeneous. It follows that the rational expressions $\hat{y}_{k;t}$ defined by (3.5.5) have (multi-)degree 0. It is not surprising, then, that Y -patterns naturally arise in the study of configurations (of points, lines, flags, etc.) in projective spaces. See Examples 3.5.8 and 3.5.10 below.

Example 3.5.8 (*Configurations of points on the projective line*). The *cross-ratio* is a quantity associated with an ordered quadruple of collinear points, particularly points on the projective line \mathbb{P}^1 (say over \mathbb{C}). For our purposes, it will be convenient to use the following version of the cross-ratio. Let $P_1, P_2, P_3, P_4 \in \mathbb{P}^1$ be four distinct points on the projective line, with projective coordinates $(a_1 : b_1)$, $(a_2 : b_2)$, $(a_3 : b_3)$, and $(a_4 : b_4)$, respectively. We then define

$$(3.5.6) \quad Y(P_1, P_2, P_3, P_4) = \frac{P_{14} P_{23}}{P_{12} P_{34}},$$

where we use the notation

$$(3.5.7) \quad P_{ij} = \det \begin{pmatrix} a_i & a_j \\ b_i & b_j \end{pmatrix} = a_i b_j - a_j b_i.$$

This quantity is related to the conventional cross-ratio via the formula

$$Y(P_1, P_2, P_3, P_4) = -(P_1, P_3; P_4, P_2).$$

The symmetric group \mathcal{S}_4 acts on quadruples of collinear points by permuting the points in a quadruple. The permutations in \mathcal{S}_4 which preserve the cross-ratio form a subgroup isomorphic to the Klein four-group. There are therefore six different versions of the cross-ratio. All of them are uniquely determined by any one of them.

The cross-ratio is essentially the only projective invariant of a quadruple of collinear points. More generally (see, e.g., [43, Section 7.4]), any rational function of an ordered m -tuple of points on the projective line which is invariant under projective transformations can be expressed in terms of cross-ratios associated to various quadruples of points. In fact, one only needs cross-ratios associated with $m - 3$ quadruples to get all $\binom{m}{4}$ of them. One way to make this explicit is by using the machinery of Y -patterns. A configuration of m distinct ordered points $P_1, \dots, P_m \in \mathbb{P}^1$ with projective coordinates $(a_1 : b_1), \dots, (a_m : b_m)$ can be encoded by a $2 \times m$ matrix

$$z = \begin{bmatrix} a_1 & a_2 & \dots & a_m \\ b_1 & b_2 & \dots & b_m \end{bmatrix}.$$

Recall that the Plücker coordinates $P_{ij} = P_{ij}(z)$ are defined by the formula (3.5.7), for $1 \leq i < j \leq m$.

We now associate a Y -seed to an arbitrary triangulation T of a convex m -gon \mathbf{P}_m (cf. Sections 1.2 and 2.2) by $m - 3$ pairwise noncrossing diagonals. Recall that \mathbf{P}_m has m vertices labeled $1, \dots, m$, in clockwise order. We label the diagonals of T by the numbers $1, \dots, m - 3$, and define the exchange matrix B_T to be the $(m - 3) \times (m - 3)$ matrix associated to the mutable part of the quiver $Q(T)$, see Definition 2.2.1. (Ignore the frozen vertices associated with the sides of the polygon.) Consider a diagonal of T labeled d . This diagonal triangulates a quadrilateral with vertices labeled i, j, k, ℓ in clockwise order, connecting vertices i and k , cf. Figure 1.2. Define

$$Y_d = Y(P_i, P_j, P_k, P_\ell),$$

cf. (3.5.6). Note that since $Y(P_i, P_j, P_k, P_\ell) = Y(P_k, P_\ell, P_i, P_j)$, there is no ambiguity in this definition. Finally, define the Y -seed associated with T to be the pair (Y_T, B_T) , where $Y_T = (Y_1, \dots, Y_{m-3})$.

It is now an exercise to verify that these Y -seeds transform under flips by the Y -seed mutation rule (3.5.3). Note that this example is nothing but the application of the construction in Theorem 3.5.1 to the seed pattern associated to the $\text{Gr}_{2,m}$ example (cf. Section 2.2).

Exercise 3.5.9. Given six points P_1, \dots, P_6 on the projective line, express the cross-ratios for the quadruples $\{P_i, P_4, P_5, P_6\}$ in terms of the cross-ratios for the quadruples $\{P_i, P_1, P_2, P_3\}$.

Example 3.5.10 (*The pentagram map*). The *pentagram map*, introduced in [46], is a transformation of generic projective polygons (i.e., cyclically ordered tuples of points on the projective plane \mathbb{P}^2) defined by the following construction: given a polygon A as input, draw all of its “shortest” diagonals, and output the “smaller” polygon A' which they cut out. See Figure 3.8.




Figure 3.8. The pentagram map.

As shown in [29], the pentagram map is related to Y -seed mutation. To explain the connection, one needs to describe the pentagram map in properly chosen coordinates. We shall view a polygon with n vertices as an n -periodic sequence $A = (A_i)_{i \in \mathbb{Z}}$ of points in \mathbb{P}^2 . Given two polygons related by the pentagram map, it is convenient to index the points of one of them by the integers \mathbb{Z} and the points of the other by the half-integers $\mathbb{Z} + \frac{1}{2}$, as shown at the right in Figure 3.8.

Recall the definition (3.5.6) of the projective invariant $Y(P_1, P_2, P_3, P_4)$ (a negative cross-ratio) associated with a quadruple of distinct collinear points P_1, P_2, P_3, P_4 . One can associate a similar invariant to a quadruple of distinct concurrent lines L_1, L_2, L_3, L_4 in \mathbb{P}^2 passing through a point Q : any line L not passing through Q intersects these lines in four distinct points P_1, P_2, P_3, P_4 , and the number $Y(L_1, L_2, L_3, L_4) := Y(P_1, P_2, P_3, P_4)$ does not depend on the choice of the line L .

Definition 3.5.11. Let A be a polygon with n vertices indexed as above by either \mathbb{Z} or $\mathbb{Z} + \frac{1}{2}$. The y -parameters of A are the numbers $y_j(A)$ (for $1 \leq j \leq 2n$) defined by

$$(3.5.8) \quad \begin{aligned} y_{2k}(A) &= Y(\overleftarrow{A_k A_{k-1}}, \overleftarrow{A_k A_{k+2}}, \overleftarrow{A_k A_{k+1}}, \overleftarrow{A_k A_{k-2}})^{-1}, \\ y_{2k+1}(A) &= Y(\overleftarrow{A_k}, \overleftarrow{A_{k+2} A_{k+3} \cap L}, \overleftarrow{A_{k+1}}, \overleftarrow{A_{k-2} A_{k-1} \cap L}), \end{aligned}$$

where $L = \overleftrightarrow{A_k A_{k+1}}$. (Here $\overleftrightarrow{A_i A_j}$ denotes the line passing through A_i and A_j .) See Figure 3.9.




Figure 3.9. The y -parameters of a polygon.

We next define the $2n \times 2n$ exchange matrix $B = (b_{ij})$ by

$$(3.5.9) \quad b_{ij} = \begin{cases} (-1)^j & \text{if } i - j \equiv \pm 1 \pmod{2n}; \\ (-1)^{j+1} & \text{if } i - j \equiv \pm 3 \pmod{2n}; \\ 0 & \text{otherwise.} \end{cases}$$

We set $Y(A) = (y_1(A), \dots, y_{2n}(A))$. Thus $(Y(A), B)$ is a Y -seed of rank $2n$.

The following result, obtained in [29], is included without proof.

Proposition 3.5.12. *Let A be an n -gon indexed by \mathbb{Z} , and let A' be the n -gon (indexed by $\mathbb{Z} + \frac{1}{2}$) obtained from A via the pentagram map. Then applying the composition of Y -seed mutations*

$$\mu_{\text{even}} = \mu_2 \circ \mu_4 \circ \cdots \circ \mu_{2n}$$

to the Y -seed $(Y(A), B)$ (cf. (3.5.8)–(3.5.9)) produces the Y -seed $(Y(A'), -B)$.

Similarly, let A' be an n -gon indexed by $\mathbb{Z} + \frac{1}{2}$. Then applying

$$\mu_{\text{odd}} = \mu_1 \circ \mu_3 \circ \cdots \circ \mu_{2n-1}$$

to the Y -seed $(Y(A'), -B)$ produces the Y -seed $(Y(A''), B)$ associated with the n -gon A'' obtained from A' via the pentagram map.

(Note that the individual mutations in each of μ_{even} and μ_{odd} commute.)

Y -patterns have arisen in many other mathematical contexts. An incomplete list includes:

- Thurston’s *shear coordinates* in Teichmüller spaces and their generalizations (see, e.g., [17] and references therein);
- recursively defined sequences of points on *elliptic curves*, and associated Somos-like sequences, cf. Example 3.4.3 and Remark 3.4.5;
- *wall-crossing formulas* for motivic Donaldson-Thomas invariants introduced by M. Kontsevich and Y. Soibelman (see, e.g., [34]), and related wall-crossing phenomena for BPS states in theoretical physics;
- Fock-Goncharov varieties [16], including moduli spaces of point configurations in basic affine spaces;
- Zamolodchikov’s Y -systems [25, 53] in the theory of the Thermodynamic Bethe Ansatz.

We will return to some of the aforementioned applications in the subsequent chapters.

3.6. Tropical semifields

In this section, we re-examine the combinatorics of matrix mutations, relating it to the concept of Y -seeds and their mutations discussed in Section 3.5. We begin by introducing the notion of semifield, and in particular, the tropical semifield, which will give us an important alternative way to encode the bottom part of an extended exchange matrix \tilde{B} .

Definition 3.6.1. A *semifield* is an abelian group P , written multiplicatively, endowed with an operation of “auxiliary addition” \oplus which is required to be commutative and associative, and satisfy the distributive law with respect to the multiplication in P .

We emphasize that (P, \oplus) does not have to be a group, just a semigroup. Since every element of P has a multiplicative inverse, it does not contain an additive identity (or “zero”) element (unless P is trivial).

Definition 3.6.2. Let $\text{Trop}(q_1, \dots, q_\ell)$ denote the multiplicative group of Laurent monomials in the variables q_1, \dots, q_ℓ . We equip $\text{Trop}(q_1, \dots, q_\ell)$ with the binary operation of *tropical addition* \oplus defined by

$$(3.6.1) \quad \prod_{i=1}^{\ell} q_i^{a_i} \oplus \prod_{i=1}^{\ell} q_i^{b_i} = \prod_{i=1}^{\ell} q_i^{\min(a_i, b_i)} .$$

Lemma 3.6.3. *Tropical addition is commutative and associative, and it satisfies the distributive law with respect to the ordinary multiplication:*

$$(p \oplus q)r = pr \oplus qr.$$

Thus $\text{Trop}(q_1, \dots, q_\ell)$ is a semifield, which we call the *tropical semifield* generated by q_1, \dots, q_ℓ .

Remark 3.6.4. The above terminology differs from the one used in *tropical geometry* by what is essentially a notational convention: replacing Laurent monomials by the corresponding vectors of exponents, one gets a semifield in which multiplication is the ordinary addition, and auxiliary addition amounts to taking the minimum.

The formalism of the tropical semifield and its auxiliary addition allows us to restate the rules of matrix mutation in the following way.

Let \tilde{B} be an $m \times n$ extended exchange matrix. As before, x_{n+1}, \dots, x_m are the frozen variables. We encode the bottom $(m-n) \times n$ submatrix of \tilde{B} by the *coefficient tuple* $\mathbf{y} = (y_1, \dots, y_n) \in \text{Trop}(x_{n+1}, \dots, x_m)^n$ defined by

$$(3.6.2) \quad y_j = \prod_{i=n+1}^m x_i^{b_{ij}} \quad (j \in \{1, \dots, n\}).$$

Thus the matrix \tilde{B} contains the same information as its top $n \times n$ submatrix B together with the coefficient tuple \mathbf{y} .

Proposition 3.6.5. *Let $\tilde{B} = (b_{ij})$ and \tilde{B}' be two extended skew-symmetrizable matrices related by a mutation μ_k , and let $\mathbf{y} = (y_1, \dots, y_n)$ and $\mathbf{y}' = (y'_1, \dots, y'_n)$ be the corresponding coefficient tuples (cf. (3.6.2)). Then*

$$(3.6.3) \quad y'_j = \begin{cases} y_k^{-1} & \text{if } j = k; \\ y_j(y_k \oplus 1)^{-b_{kj}} & \text{if } j \neq k \text{ and } b_{kj} \leq 0; \\ y_j(y_k^{-1} \oplus 1)^{-b_{kj}} & \text{if } j \neq k \text{ and } b_{kj} \geq 0. \end{cases}$$

Comparing (3.6.3) with (3.5.3), we can informally say that the coefficient tuple \mathbf{y} undergoes a “tropical Y -seed mutation” at k .

Proposition 3.6.5 can be proved by translating the rules of matrix mutation into the language of the tropical semifield. We outline a different proof which explains the connection between the formulas (3.5.2)–(3.5.3) and (3.6.3), and introduces some notions that will be useful in the sequel.

Definition 3.6.6. Let $\mathbb{Q}_{\text{sf}}(x_1, \dots, x_m)$ denote the set of nonzero rational functions in x_1, \dots, x_m which can be written as subtraction-free rational expressions in these variables, with positive rational coefficients. Thus, each element of $\mathbb{Q}_{\text{sf}}(x_1, \dots, x_m)$ can be written in the form $\frac{P(x_1, \dots, x_m)}{Q(x_1, \dots, x_m)}$, where P and Q are polynomials with positive coefficients. The set $\mathbb{Q}_{\text{sf}}(x_1, \dots, x_m)$ is a semifield with respect to the ordinary operations of addition and multiplication. We call it the *universal semifield* generated by x_1, \dots, x_m .

This terminology is justified by the following easy lemma, whose proof we omit (see [3, Lemma 2.1.6]). Informally speaking, this lemma says that the generators x_1, \dots, x_m of the semifield $\mathbb{Q}_{\text{sf}}(x_1, \dots, x_m)$ do not satisfy any relations, save for those which are implied by the axioms of a semifield.

Lemma 3.6.7. *For any semifield \mathcal{S} , any map $f : \{x_1, \dots, x_m\} \rightarrow \mathcal{S}$ extends uniquely to a semifield homomorphism $\mathbb{Q}_{\text{sf}}(x_1, \dots, x_m) \rightarrow \mathcal{S}$.*

Proof of Proposition 3.6.5. Let $\tilde{\mathbf{x}} = (x_1, \dots, x_m)$ be a collection of indeterminates. Define the semifield homomorphism

$$f : \mathbb{Q}_{\text{sf}}(x_1, \dots, x_m) \rightarrow \text{Trop}(x_{n+1}, \dots, x_m)$$

by setting (cf. Lemma 3.6.7)

$$f(x_i) = \begin{cases} 1 & \text{if } i \leq n; \\ x_i & \text{if } i > n. \end{cases}$$

Applying mutation at k to the seed $(\tilde{\mathbf{x}}, \tilde{B})$, we get a new seed $(\tilde{\mathbf{x}}', \tilde{B}')$, in which the only new cluster variable x'_k satisfies an exchange relation of the form $x_k x'_k = M_1 + M_2$. The two monomials M_1 and M_2 are coprime, and in particular do not share a frozen variable x_i . Applying the semifield homomorphism f , we obtain $1 \cdot f(x'_k) = f(M_1) \oplus f(M_2) = 1$, so $f(x'_k) = 1$.

Now let $\hat{\mathbf{y}}$ and $\hat{\mathbf{y}}'$ (resp., \mathbf{y} and \mathbf{y}') be defined by (3.5.1) (resp., (3.6.2)). Since all cluster variables in $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{x}}'$ are sent to 1 by f , we conclude that $f(\hat{\mathbf{y}}) = \mathbf{y}$ and $f(\hat{\mathbf{y}}') = \mathbf{y}'$. Therefore applying f to (3.5.2) yields (3.6.3). \square

Let $(\tilde{\mathbf{x}}, \tilde{B})$ be a labeled seed as before. Since the extended exchange matrix \tilde{B} contains the same information as the exchange matrix B together with the coefficient tuple \mathbf{y} defined by (3.6.2), we can identify the seed $(\tilde{\mathbf{x}}, \tilde{B})$ with the triple $(\mathbf{x}, \mathbf{y}, B)$. Abusing notation, we will also refer to such triples as (labeled) seeds:

Definition 3.6.8. Let \mathcal{F} be a field of rational functions (say over \mathbb{C}) in some m variables which include the *frozen variables* x_{n+1}, \dots, x_m . A labeled seed (of geometric type) of rank n is a triple $\Sigma = (\mathbf{x}, \mathbf{y}, B)$ consisting of

- a *cluster* \mathbf{x} , an n -tuple of elements of \mathcal{F} such that the *extended cluster* $\mathbf{x} \cup \{x_{n+1}, \dots, x_m\}$ freely generates \mathcal{F} ;
- an *exchange matrix* B , a skew-symmetrizable integer matrix;
- a *coefficient tuple* \mathbf{y} , an n -tuple of Laurent monomials in the tropical semifield $\text{Trop}(x_{n+1}, \dots, x_m)$.

We can now restate the rules of seed mutation in this language.

Proposition 3.6.9. *Let $(\mathbf{x}, \mathbf{y}, B)$, with $B = (b_{ij})$ and $\mathbf{y} = (y_1, \dots, y_n)$, and $(\mathbf{x}', \mathbf{y}', B')$, with $\mathbf{y}' = (y'_1, \dots, y'_n)$, be two labeled seeds related by a mutation μ_k . Then $(\mathbf{x}', \mathbf{y}', B')$ is obtained from $(\mathbf{x}, \mathbf{y}, B)$ as follows:*

- $B' = \mu_k(B)$;
- \mathbf{y}' is given by the “tropical Y -seed mutation rule” (3.6.3);
- $\mathbf{x}' = (\mathbf{x} - \{x_k\}) \cup \{x'_k\}$, where x'_k is defined by the exchange relation

$$(3.6.4) \quad x_k x'_k = \frac{y_k}{y_k \oplus 1} \prod_{b_{ik} > 0} x_i^{b_{ik}} + \frac{1}{y_k \oplus 1} \prod_{b_{ik} < 0} x_i^{-b_{ik}};$$

Proof. The only statement requiring proof is (3.6.4), which can be easily seen to be a rewriting of (3.1.1). \square

We will use Equation (3.6.4) in Chapter 12 to define cluster algebras over an arbitrary semifield.

We can now re-define the notion of a labeled seed pattern.

Definition 3.6.10. A *labeled seed pattern* of rank n is obtained by assigning a triple $\Sigma(t) = (\mathbf{x}(t), \mathbf{y}(t), B(t))$ as above to every vertex t in the n -regular tree \mathbb{T}_n , and requiring that the triples assigned to adjacent vertices of the tree are related by the corresponding mutation, as described in Proposition 3.6.9.

The advantage of the latest version of the definition of seed pattern is that it enables us to perform calculations for arbitrary extensions of a given exchange matrix to an extended exchange matrix; see (3.6.2) and the surrounding discussion.

Example 3.6.11 (*Type A_2*). Consider the seed pattern of rank 2

$$\dots \overset{2}{\Sigma(0)} \overset{1}{\Sigma(1)} \overset{2}{\Sigma(2)} \overset{1}{\Sigma(3)} \overset{2}{\Sigma(4)} \overset{1}{\Sigma(5)} \overset{2}{\Sigma(6)} \dots$$

formed by the seeds $\Sigma(t) = (\mathbf{x}(t), \mathbf{y}(t), B(t))$, for $t \in \mathbb{T}_2 \cong \mathbb{Z}$, with the exchange matrices

$$(3.6.5) \quad B(t) = (-1)^t \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

(cf. Example 3.5.4). Note that we do not specify the bottom part of the initial exchange matrix, nor even the number of frozen variables. Still, we can express all the seeds in terms of the initial one using the language of the tropical semifield, and following the recipe formulated in Proposition 3.6.9. The results of the computation are shown in Figure 3.10.

t	$\mathbf{y}(t)$		$\mathbf{x}(t)$	
0	y_1	y_2	x_1	x_2
1	$\frac{1}{y_1}$	$\frac{y_1 y_2}{y_1 \oplus 1}$	$\frac{y_1 + x_2}{x_1(y_1 \oplus 1)}$	x_2
2	$\frac{y_2}{y_1 y_2 \oplus y_1 \oplus 1}$	$\frac{y_1 \oplus 1}{y_1 y_2}$	$\frac{y_1 + x_2}{x_1(y_1 \oplus 1)}$	$\frac{x_1 y_1 y_2 + y_1 + x_2}{(y_1 y_2 \oplus y_1 \oplus 1)x_1 x_2}$
3	$\frac{y_1 y_2 \oplus y_1 \oplus 1}{y_2}$	$\frac{1}{y_1(y_2 \oplus 1)}$	$\frac{x_1 y_2 + 1}{x_2(y_2 \oplus 1)}$	$\frac{x_1 y_1 y_2 + y_1 + x_2}{(y_1 y_2 \oplus y_1 \oplus 1)x_1 x_2}$
4	$\frac{1}{y_2}$	$y_1(y_2 \oplus 1)$	$\frac{x_1 y_2 + 1}{x_2(y_2 \oplus 1)}$	x_1
5	y_2	y_1	x_2	x_1

Figure 3.10. Seeds in type A_2 . The exchange matrices are given by (3.6.5). We denote the initial cluster by $\mathbf{x}_\circ = (x_1, x_2)$, and the initial coefficient tuple by $\mathbf{y}_\circ = (y_1, y_2)$. The formulas for the coefficient tuples $\mathbf{y}(t)$ are the tropical versions of the formulas in Figure 3.7. Note that the labeled seed $\Sigma(5)$ is obtained from $\Sigma(0)$ by interchanging the indices 1 and 2; the sequence then continues by obvious periodicity (so that $\Sigma(10)$ is identical to $\Sigma(0)$, etc.).

Bibliography

- [1] AIGNER, M. *Markov's theorem and 100 years of the uniqueness conjecture*. Springer, Cham, 2013. A mathematical journey from irrational numbers to perfect matchings.
- [2] ASSEM, I., BLAIS, M., BRÜSTLE, T., AND SAMSON, A. Mutation classes of skew-symmetric 3×3 -matrices. *Comm. Algebra* 36, 4 (2008), 1209–1220.
- [3] BERENSTEIN, A., FOMIN, S., AND ZELEVINSKY, A. Parametrizations of canonical bases and totally positive matrices. *Adv. Math.* 122, 1 (1996), 49–149.
- [4] BERENSTEIN, A., FOMIN, S., AND ZELEVINSKY, A. Cluster algebras. III. Upper bounds and double Bruhat cells. *Duke Math. J.* 126, 1 (2005), 1–52.
- [5] BERNŠTEĪN, I. N., GEL'FAND, I. M., AND PONOMAREV, V. A. Coxeter functors, and Gabriel's theorem. *Uspehi Mat. Nauk* 28, 2(170) (1973), 19–33.
- [6] BOUSQUET-MÉLOU, M., PROPP, J., AND WEST, J. Perfect matchings for the three-term Gale-Robinson sequences. *Electron. J. Combin.* 16, 1 (2009), Research Paper 125, 37.
- [7] CALDERO, P., AND KELLER, B. From triangulated categories to cluster algebras. II. *Ann. Sci. École Norm. Sup. (4)* 39, 6 (2006), 983–1009.
- [8] CAYLEY, A. On Gauss's pentagramma mirificum. *Phil. Mag. (4)* XLII, 280 (1871), 311–312.
- [9] CECOTTI, S., AND VAFA, C. On classification of $N = 2$ supersymmetric theories. *Comm. Math. Phys.* 158, 3 (1993), 569–644.
- [10] COXETER, H. S. M. Frieze patterns. *Acta Arith.* 18 (1971), 297–310.
- [11] COXETER, H. S. M. *Non-Euclidean geometry*, sixth ed. MAA Spectrum. Mathematical Association of America, Washington, DC, 1998.
- [12] CRYER, C. W. The LU -factorization of totally positive matrices. *Linear Algebra and Appl.* 7 (1973), 83–92.
- [13] CRYER, C. W. Some properties of totally positive matrices. *Linear Algebra and Appl.* 15, 1 (1976), 1–25.
- [14] DODGSON, C. L. *The mathematical pamphlets of Charles Lutwidge Dodgson and related pieces*, vol. 2 of *The Pamphlets of Lewis Carroll*. Published by the Lewis Carroll Society of North America, Silver Spring, MD; and distributed by the University Press of Virginia, Charlottesville, VA, 1994. Compiled, with introductory essays, notes and annotations by Francine F. Abeles.

- [15] FENG, B., HANANY, A., HE, Y.-H., AND URANGA, A. M. Toric duality as Seiberg duality and brane diamonds. *J. High Energy Phys.*, 12 (2001), Paper 35, 29.
- [16] FOCK, V., AND GONCHAROV, A. Moduli spaces of local systems and higher Teichmüller theory. *Publ. Math. Inst. Hautes Études Sci.*, 103 (2006), 1–211.
- [17] FOCK, V. V., AND GONCHAROV, A. B. Dual Teichmüller and lamination spaces. In *Handbook of Teichmüller theory. Vol. I*, vol. 11 of *IRMA Lect. Math. Theor. Phys.* Eur. Math. Soc., Zürich, 2007, pp. 647–684.
- [18] FOMIN, S. Total positivity and cluster algebras. In *Proceedings of the International Congress of Mathematicians. Volume II* (2010), Hindustan Book Agency, New Delhi, pp. 125–145.
- [19] FOMIN, S., AND READING, N. Root systems and generalized associahedra. In *Geometric combinatorics*, vol. 13 of *IAS/Park City Math. Ser.* Amer. Math. Soc., Providence, RI, 2007, pp. 63–131.
- [20] FOMIN, S., AND ZELEVINSKY, A. Double Bruhat cells and total positivity. *J. Amer. Math. Soc.* 12, 2 (1999), 335–380.
- [21] FOMIN, S., AND ZELEVINSKY, A. Total positivity: tests and parametrizations. *Math. Intelligencer* 22, 1 (2000), 23–33.
- [22] FOMIN, S., AND ZELEVINSKY, A. Totally nonnegative and oscillatory elements in semisimple groups. *Proc. Amer. Math. Soc.* 128, 12 (2000), 3749–3759.
- [23] FOMIN, S., AND ZELEVINSKY, A. Cluster algebras. I. Foundations. *J. Amer. Math. Soc.* 15, 2 (2002), 497–529 (electronic).
- [24] FOMIN, S., AND ZELEVINSKY, A. Cluster algebras. II. Finite type classification. *Invent. Math.* 154, 1 (2003), 63–121.
- [25] FOMIN, S., AND ZELEVINSKY, A. Y -systems and generalized associahedra. *Ann. of Math. (2)* 158, 3 (2003), 977–1018.
- [26] FORDY, A. P., AND MARSH, R. J. Cluster mutation-periodic quivers and associated Laurent sequences. *J. Algebraic Combin.* 34, 1 (2011), 19–66.
- [27] FRANCO, S., HANANY, A., VEGH, D., WECHT, B., AND KENNAWAY, K. D. Brane dimers and quiver gauge theories. *J. High Energy Phys.*, 1 (2006), 096, 48 pp. (electronic).
- [28] GANTMACHER, F. P., AND KREIN, M. G. *Oscillation matrices and kernels and small vibrations of mechanical systems*, revised ed. AMS Chelsea Publishing, Providence, RI, 2002. Translation based on the 1941 Russian original, Edited and with a preface by Alex Eremenko.
- [29] GLICK, M. The pentagram map and Y -patterns. *Adv. Math.* 227, 2 (2011), 1019–1045.
- [30] GONCHAROV, A. B., AND KENYON, R. Dimers and cluster integrable systems. *Ann. Sci. Éc. Norm. Supér. (4)* 46, 5 (2013), 747–813.
- [31] HONE, A. N. W. Laurent polynomials and superintegrable maps. *SIGMA Symmetry Integrability Geom. Methods Appl.* 3 (2007), Paper 022, 18.
- [32] HONE, A. N. W., AND SWART, C. Integrality and the Laurent phenomenon for Somos 4 and Somos 5 sequences. *Math. Proc. Cambridge Philos. Soc.* 145, 1 (2008), 65–85.
- [33] KENYON, R. W., PROPP, J. G., AND WILSON, D. B. Trees and matchings. *Electron. J. Combin.* 7 (2000), Research Paper 25, 34 pp. (electronic).
- [34] KONTSEVICH, M., AND SOIBELMAN, Y. Motivic Donaldson-Thomas invariants: summary of results. In *Mirror symmetry and tropical geometry*, vol. 527 of *Contemp. Math.* Amer. Math. Soc., Providence, RI, 2010, pp. 55–89.

- [35] KUNG, J. P. S., AND ROTA, G.-C. The invariant theory of binary forms. *Bull. Amer. Math. Soc. (N.S.)* 10, 1 (1984), 27–85.
- [36] LOEWNER, C. On totally positive matrices. *Math. Z.* 63 (1955), 338–340.
- [37] LUSZTIG, G. Total positivity in reductive groups. In *Lie theory and geometry*, vol. 123 of *Progr. Math.* Birkhäuser Boston, Boston, MA, 1994, pp. 531–568.
- [38] LUSZTIG, G. Introduction to total positivity. In *Positivity in Lie theory: open problems*, vol. 26 of *de Gruyter Exp. Math.* de Gruyter, Berlin, 1998, pp. 133–145.
- [39] MUIR, T. *A treatise on the theory of determinants*. Revised and enlarged by William H. Metzler. Dover Publications, Inc., New York, 1960.
- [40] NEWMAN, M. *Integral matrices*. Academic Press, New York-London, 1972. Pure and Applied Mathematics, Vol. 45.
- [41] PENNER, R. C. The decorated Teichmüller space of punctured surfaces. *Comm. Math. Phys.* 113, 2 (1987), 299–339.
- [42] POSTNIKOV, A. Total positivity, grassmannians, and networks, [arXiv:math/0609764](https://arxiv.org/abs/math/0609764).
- [43] RICHTER-GEBERT, J. *Perspectives on projective geometry*. Springer, Heidelberg, 2011.
- [44] RINGEL, G. Teilungen der Ebene durch Geraden oder topologische Geraden. *Math. Z.* 64 (1955), 79–102 (1956).
- [45] SCHOENBERG, I. Über variationsvermindernde lineare Transformationen. *Math. Z.* 32, 1 (1930), 321–328.
- [46] SCHWARTZ, R. The pentagram map. *Experiment. Math.* 1, 1 (1992), 71–81.
- [47] SEIBERG, N. Electric-magnetic duality in supersymmetric non-abelian gauge theories. *Nuclear Phys. B* 435, 1-2 (1995), 129–146.
- [48] SEVEN, A. I. Mutation classes of skew-symmetrizable 3×3 matrices. *Proc. Amer. Math. Soc.* 141, 5 (2013), 1493–1504.
- [49] SPEYER, D. E. Perfect matchings and the octahedron recurrence. *J. Algebraic Combin.* 25, 3 (2007), 309–348.
- [50] STASHEFF, J. D. Homotopy associativity of H -spaces. I, II. *Trans. Amer. Math. Soc.* 108 (1963), 275–292; *ibid.* 108 (1963), 293–312.
- [51] STURMFELS, B. *Algorithms in invariant theory*, 2nd ed. Texts and Monographs in Symbolic Computation. Springer, 2008.
- [52] WHITNEY, A. M. A reduction theorem for totally positive matrices. *J. Analyse Math.* 2 (1952), 88–92.
- [53] ZAMOŁODCHIKOV, A. B. On the thermodynamic Bethe ansatz equations for reflectionless ADE scattering theories. *Phys. Lett. B* 253, 3-4 (1991), 391–394.

Introduction to Cluster Algebras

Chapters 4–5

(preliminary version)

SERGEY FOMIN

LAUREN WILLIAMS

ANDREI ZELEVINSKY

Preface

This is a preliminary draft of Chapters 4–5 of our forthcoming textbook *Introduction to cluster algebras*, joint with Andrei Zelevinsky (1953–2013). Other chapters have been posted as

- [arXiv:1608.05735](https://arxiv.org/abs/1608.05735) (Chapters 1–3),
- [arXiv:2008.09189](https://arxiv.org/abs/2008.09189) (Chapter 6), and
- [arXiv:2106.02160](https://arxiv.org/abs/2106.02160) (Chapter 7).

We expect to post additional chapters in the not so distant future.

We thank Gregg Musiker for explanations concerning the Sage package [24].

We are grateful to Colin Defant, Chris Fraser, Sergei Gelfand, Felix Gotti, Amal Mattoo, Hanna Mularczyk, Emmanuel Tsukerman, and Raluca Vlad for a number of comments on the earlier versions of these chapters.

Our work was partially supported by the NSF grants DMS-1361789, DMS-1664722 and DMS-1600447.

Comments and suggestions are welcome.

Sergey Fomin
Lauren Williams

Contents

Chapter 4. New patterns from old	1
§4.1. Restrictions and embeddings of quivers and matrices	1
§4.2. Seed subpatterns and cluster subalgebras	4
§4.3. Changing the coefficients	6
§4.4. Folding	11
Chapter 5. Finite type classification	19
§5.1. Finite type classification in rank 2	20
§5.2. Cartan matrices and Dynkin diagrams	22
§5.3. Seed patterns of type A_n	28
§5.4. Seed patterns of type D_n	35
§5.5. Seed patterns of types B_n and C_n	45
§5.6. Seed patterns of types E_6, E_7, E_8	49
§5.7. Seed patterns of types F_4 and G_2	52
§5.8. Decomposable types	53
§5.9. Enumeration of clusters and cluster variables	53
§5.10. 2-finite exchange matrices	56
§5.11. Quasi-Cartan companions	62
Bibliography	67

New patterns from old

This chapter provides several methods for obtaining new seed patterns (or new cluster algebras) from existing ones.

4.1. Restrictions and embeddings of quivers and matrices

We begin by discussing some purely combinatorial constructions involving mutations of quivers or matrices—but not clusters or seeds.

Definition 4.1.1. Let \tilde{B} be an $m \times n$ extended skew-symmetrizable matrix. For a subset $I \subset [1, m]$, consider the matrix \tilde{B}_I obtained from \tilde{B} by restricting to the row set I and to the column set $I \cap [1, n]$. It is easy to see that \tilde{B}_I is again an extended skew-symmetrizable matrix. We say that \tilde{B}_I is *obtained from \tilde{B} by restriction to I* . More generally, we say that a matrix \tilde{B}' is obtained from \tilde{B} by restriction if \tilde{B}' can be identified with a matrix \tilde{B}_I as above. (Note that we will use the convention that the rows and columns of \tilde{B}_I are labeled by I rather than by $\{1, 2, \dots, |I|\}$.)

If $\tilde{B} = \tilde{B}(Q)$ is the extended exchange matrix of a quiver Q , then $\tilde{B}_I = \tilde{B}(Q_I)$ is the extended exchange matrix of the quiver Q_I , where Q_I is obtained from Q by taking the subset of vertices of Q indexed by I along with all the arrows in Q that connect the vertices in I . Such a quiver Q_I is called a *full* (or *induced*) *subquiver* of Q . The vertices in Q_I inherit the property of being frozen or mutable from the ambient quiver Q .

The following property is easy to check.

Lemma 4.1.2. *Mutation of matrices/quivers commutes with restriction. More precisely, if \tilde{B}_I is the restriction of an extended skew-symmetrizable matrix \tilde{B} to a subset I , and $k \in I$ is mutable, then $\mu_k(\tilde{B}_I) = (\mu_k(\tilde{B}))_I$.*

Definition 4.1.3. We say that a property of extended skew-symmetrizable matrices is *hereditary* if it is preserved under restriction: for any matrix \tilde{B} which has this property, the same holds true for all its submatrices \tilde{B}_I . For quivers, a property is hereditary if it is inherited by the full subquivers of any quiver which has that property.

We are interested in hereditary properties that are preserved under mutations. The first example of this kind concerns the notion of finite mutation type, cf. Definition 2.6.11; this definition can be generalized in a straightforward manner to extended exchange matrices. Using Lemma 4.1.2, we obtain the following.

Proposition 4.1.4. *Finite mutation type is a hereditary property.*

In other words, an extended exchange matrix obtained by restriction from an extended exchange matrix of finite mutation type will again have finite mutation type.

We will show later that the property of being mutation-equivalent to an orientation of a (possibly disconnected, simply laced) Dynkin diagram is hereditary, see Remark 5.10.9. See also Theorem 10.4.1 for a version of this statement that includes extended Dynkin diagrams.

Example 4.1.5. Recall that an acyclic quiver is one containing no oriented cycles. A quiver with no frozen vertices is called *mutation-acyclic* if it is mutation equivalent to an acyclic quiver. It was shown in [3], using the machinery of quiver representations, that the property of being mutation-acyclic is hereditary. It would be interesting to find an elementary proof.

In light of Example 4.1.5, it is natural to consider the unoriented analogue of the notion of mutation-acyclicity.

Remark 4.1.6. A quiver is called *arborizable* if it is mutation-equivalent to an orientation of a forest (i.e., an undirected simple graph with no cycles).

Unfortunately, arborizability is not a hereditary property. A counterexample is given in Figure 4.1.

Lemma 4.1.7. *For mutation classes \mathbf{Q} and \mathbf{Q}' , the following are equivalent:*

- (i) *there exist $\tilde{B} \in \mathbf{Q}$ and $\tilde{B}' \in \mathbf{Q}'$ such that \tilde{B} is obtained from \tilde{B}' by restriction;*
- (ii) *for any $\tilde{B} \in \mathbf{Q}$, there exists $\tilde{B}' \in \mathbf{Q}'$ such that \tilde{B} is obtained from \tilde{B}' by restriction.*

Proof. The equivalence of (i) and (ii) follows from Lemma 4.1.2. □

The notion of restriction descends to a partial order on the set of mutation classes of extended exchange matrices, as follows.




Figure 4.1. The quiver Q shown on the left is arborizable since it is mutation equivalent to the quiver Q' on the right. On the other hand, the full subquiver Q_I on the vertex set $I = \{1, 2, 3\}$ is not arborizable: Q_I is of finite mutation type, and its mutation class $[Q_I]$ consists of two quivers (up to isomorphism) neither of which is an orientation of a tree.

Definition 4.1.8. Let \mathbf{Q} and \mathbf{Q}' be mutation classes. We say that \mathbf{Q} is *embeddable* into \mathbf{Q}' , and write $\mathbf{Q} \leq \mathbf{Q}'$, if either of the two equivalent conditions (i)–(ii) in Lemma 4.1.7 holds. In particular, when \mathbf{Q} and \mathbf{Q}' are mutation classes of quivers, we say that \mathbf{Q} is *embeddable* into \mathbf{Q}' , if there exist quivers $Q \in \mathbf{Q}$ and $Q' \in \mathbf{Q}'$ such that Q is a full subquiver of Q' .

Example 4.1.9. In Figure 2.12, the mutation class of each quiver (including any orientation of each of the Dynkin diagrams shown) is embeddable into the mutation class of every quiver appearing in the rows below. We revisit (and extend) this example in Remark 5.2.13.

Recall that $[\tilde{B}]$ denotes the mutation class of an extended exchange matrix \tilde{B} .

Remark 4.1.10. Fix a mutation class \mathbf{R} . Then embeddability into \mathbf{R} is a hereditary property. More precisely, if $[\tilde{B}] \leq \mathbf{R}$, then $[\tilde{B}_I] \leq \mathbf{R}$ for any matrix \tilde{B}_I obtained by restriction from \tilde{B} .

Remark 4.1.11. The equivalent conditions (i)–(ii) in Lemma 4.1.7 do not imply the condition

(iii) for any $\tilde{B}' \in \mathbf{Q}'$, there exists $\tilde{B} \in \mathbf{Q}$ such that \tilde{B} is obtained from \tilde{B}' by restriction.

To see this, consider the quivers Q and Q' shown in Figure 4.1. Let $\mathbf{Q} = [Q_I]$ with $I = \{1, 2, 3\}$ and $\mathbf{Q}' = [Q']$. Then (i)–(ii) hold while (iii) fails. (Note that that the mutation sequence relating the two quivers in Figure 4.1 uses mutations at vertices outside of I .)

Problem 4.1.12. Let T and T' be finite trees, and let \mathbf{T} and \mathbf{T}' denote the mutation classes containing their respective orientations. Is it true that \mathbf{T} is embeddable into \mathbf{T}' if and only if T can be obtained from T' by contracting some edges?

Remark 4.1.13. Given mutation classes \mathbf{Q} and \mathbf{R} , the problem of deciding whether \mathbf{Q} is embeddable into \mathbf{R} is generally very hard (unless \mathbf{R} is finite). For example, for any positive integer k , there is no known algorithm to determine whether the two-vertex *Kronecker quiver*

$$(4.1.1) \quad \bullet \begin{array}{c} \xrightarrow{\quad} \\ \cdots \\ \xrightarrow{\quad} \end{array} \bullet$$

with k arrows is embeddable into the mutation class of a given quiver R . That is, for any k , there is no known general method to determine whether R can be mutated to a quiver containing two vertices connected by k arrows.

4.2. Seed subpatterns and cluster subalgebras

Definition 4.2.1. Let $(\tilde{\mathbf{x}}, \tilde{B})$ be a seed of rank n , and let $x_i \in \tilde{\mathbf{x}}$ be a cluster variable. *Freezing* at the index i (or, of the variable x_i) is a transformation of the seed that reclassifies i and x_i as frozen, and accordingly removes the i th column from the exchange matrix \tilde{B} . (In addition, this would typically require a change of indexing, provided we want to keep using the smaller indices $1, \dots, n-1$ for the mutable variables.) More generally, we can freeze any subset of cluster variables. The order of freezing does not matter.

In the quiver case, freezing at a subset of mutable vertices amounts to reclassifying all these vertices, and the corresponding cluster variables, as frozen; and then removing all arrows connecting frozen vertices to each other.

Lemma 4.2.2. *Freezing commutes with seed mutation.*

Proof. This property is straightforward from the definitions. \square

It is natural to try to extend the operation of restriction (to a full subquiver, or more generally to a submatrix of the exchange matrix) to the level of seeds. Note however that the naive notion of restriction does *not*, generally speaking, commute with seed mutation—because an exchange relation for a cluster variable associated with a vertex in a subquiver may well involve variables coming from outside the subquiver. This observation explains the additional constraints appearing in Definition 4.2.3 below.

Definition 4.2.3. Let $(\tilde{\mathbf{x}}, \tilde{B})$ be a seed, and let $I \sqcup J$ be a partition of $[1, m]$ such that $b_{jk} = 0$ for any $j \in J$ and $k \in I \cap [1, n]$. (In other words, none of the variables x_j , for $j \in J$, appears on the right-hand side of an exchange relation $x_k x'_k = \dots$, for $k \in I$.) We then define the *restricted seed* $(\tilde{\mathbf{x}}_I, \tilde{B}_I)$ with the extended cluster $\tilde{\mathbf{x}}_I = (x_i)_{i \in I}$ and the extended exchange (sub)matrix \tilde{B}_I , cf. Definition 4.1.1. In the case where \tilde{B} comes from a quiver Q , this is equivalent to requiring that there are no arrows between $I \cap [1, n]$ and J .

Example 4.2.4. Let Q and $Q' = \mu_k(Q)$ be the quivers shown in Figure 2.1 on the left and on the right, respectively. The subset $I = \{a, q, k\}$ does not satisfy the conditions in Definition 4.2.3 (neither for Q , nor for Q') since the vertex $k \in I$ is connected by arrows to vertices not in I . On the other hand, if we first freeze k , then we can restrict to I in Q' (but not in Q).

Lemma 4.2.5. *Passing to a restricted seed commutes with seed mutation.*

It is easy to see that repeated applications of freezing and restriction produce an outcome that can be achieved by a single application of freezing followed by restriction.

Definition 4.2.6. Let Σ be a seed. Freeze some subset of cluster variables, as in Definition 4.2.1, to obtain a seed Σ' . After that, apply the construction in Definition 4.2.3 to the seed Σ' to obtain a restricted seed Σ'' . We then say that the seed pattern defined by Σ'' is a *seed subpattern* of the seed pattern defined by Σ ; and the cluster algebra associated to Σ'' is a *cluster subalgebra* of the cluster algebra associated to Σ .

Put simply, a cluster subpattern can be viewed as a part of the original pattern in which we are only allowed to exchange cluster variables labeled by a particular subset of indices, and in which we discard (some of) the coefficient variables that do not appear at all in the resulting exchange relations.

Note that if passing to a seed subpattern (resp., cluster subalgebra) involves freezing at least one cluster variable, then its rank is smaller than the rank of the original pattern (resp., cluster algebra).

Example 4.2.7. Let \mathbf{P}_{n+3} be a convex polygon whose vertices are labeled $\{1, 2, \dots, n+3\}$ in clockwise order. As explained in Section 2.2, each triangulation T of \mathbf{P}_{n+3} gives rise to a seed $(\tilde{\mathbf{x}}(T), Q(T))$ whose mutable variables are the Plücker coordinates P_{ij} labeled by the diagonals of T , and whose frozen variables are the Plücker coordinates labeled by the sides of \mathbf{P}_{n+3} . We thus obtain a seed pattern and a cluster algebra $R_{2,n+3}$ of rank n .

Let $S = \{s_1 < \dots < s_\ell\}$ be a subset of $\{1, 2, \dots, n\}$, with $\ell \geq 4$, and let \mathbf{P}_S be the convex polygon on the vertex set S . Let T be a triangulation of \mathbf{P}_{n+3} that includes the sides of the polygon \mathbf{P}_S , and consequently contains a triangulation T_S of \mathbf{P}_S . Take the seed $(\tilde{\mathbf{x}}(T), Q(T))$, and freeze all the cluster variables in $\tilde{\mathbf{x}}(T)$ corresponding to the sides of \mathbf{P}_S . We can now restrict from $\tilde{\mathbf{x}}(T)$ to the subset $\tilde{\mathbf{x}}(T_S)$ consisting of the elements labeled by diagonals and sides of \mathbf{P}_S . The resulting seed $(\tilde{\mathbf{x}}(T_S), Q(T_S))$ defines a rank $\ell - 3$ cluster subalgebra (isomorphic to $R_{2,\ell}$) of the cluster algebra $R_{2,n+3}$.

The following properties of seed subpatterns and cluster subalgebras follow immediately from the definitions.

Lemma 4.2.8.

- Let Σ be a seed pattern. A seed subpattern of a seed subpattern of Σ is again a seed subpattern of Σ .
- If a cluster algebra has finitely many cluster variables, then so does any of its cluster subalgebras.
- A connected component of a quiver gives rise to a seed subpattern.

4.3. Changing the coefficients

In this section we will establish a connection between seed patterns utilizing the same exchange matrices but different coefficient tuples. For a more thorough and comprehensive treatment, see [14].

We begin by a brief discussion of a couple of very simple special instances of coefficient change.

Definition 4.3.1. Let $(\tilde{\mathbf{x}}, \tilde{B})$ be a seed with an m -element extended cluster $\tilde{\mathbf{x}}$, and let $x_i \in \tilde{\mathbf{x}}$ be a coefficient (or frozen) variable. *Trivialization* at the index i (or, of the variable x_i) is a transformation of the seed that removes x_i from $\tilde{\mathbf{x}}$, and accordingly removes i from the set of indices, and the i th row from the exchange matrix \tilde{B} . (As in the case of freezing, cf. Definition 4.2.1, a renumbering may be required if we want to use the indices $1, \dots, m-1$ after the trivialization.) More generally, we can trivialize any subset of coefficient variables; the order of operations does not matter.

In the quiver case, trivialization amounts to a removal of a subset of frozen vertices, together with all arrows incident to them; and the removal of the corresponding coefficient variables.

Lemma 4.3.2. *Trivialization of coefficients commutes with seed mutation.*

Proof. The key observation is that trivialization of a coefficient variable x_i can be interpreted as setting $x_i = 1$. \square

It is also easy to see that trivializing coefficients commutes with taking a seed subpattern.

Remark 4.3.3. Here is another simple way of transforming a seed pattern into a new one: introduce (any number of) additional “dummy” coefficient variables that do not appear in any exchange relations. That is, enlarge the extended exchange matrices by adding rows consisting entirely of zeroes; in the quiver case, just add isolated frozen vertices. This transformation changes the associated cluster algebra by tensoring it with the polynomial ring generated by the dummy variables.

We next describe a large class of “rescaling” transformations of seed patterns. Roughly, the idea is to multiply each cluster variable by a Laurent

monomial in the coefficient variables, and then rewrite the exchange relations in terms of the “rescaled” variables. (We use quotation marks since we are not simply rescaling our variables, but also sending them to a different ambient field of rational functions.) The key property of this construction, formalized in Theorem 4.3.4 below, is that the resulting “rescaled” seeds are again related by (the same) mutations, yielding a seed pattern.

Theorem 4.3.4. *Let n, m, \bar{m} be positive integers, with $n \leq m$ and $n \leq \bar{m}$. Let \mathcal{F} and $\bar{\mathcal{F}}$ be the fields of rational functions in the variables x_1, \dots, x_m and $\bar{x}_1, \dots, \bar{x}_{\bar{m}}$, respectively. Let*

$$\varphi : \mathbb{Q}_{\text{sf}}(x_1, \dots, x_m) \rightarrow \text{Trop}(\bar{x}_{n+1}, \dots, \bar{x}_{\bar{m}})$$

be a semifield homomorphism (determined by an arbitrary choice of Laurent monomials $\varphi(x_i) \in \text{Trop}(\bar{x}_{n+1}, \dots, \bar{x}_{\bar{m}})$). Define the semifield map

$$\psi : \mathbb{Q}_{\text{sf}}(x_1, \dots, x_m) \rightarrow \mathbb{Q}_{\text{sf}}(\bar{x}_1, \dots, \bar{x}_{\bar{m}})$$

by setting

$$(4.3.1) \quad \psi(x_i) = \begin{cases} \bar{x}_i \varphi(x_i) & \text{if } i \leq n; \\ \varphi(x_i) & \text{if } i > n. \end{cases}$$

Let $(\mathbf{x}(t), \mathbf{y}(t), B(t))_{t \in \mathbb{T}_n}$ be a seed pattern in \mathcal{F} (cf. Definition 3.6.10), with

$$\begin{aligned} \mathbf{x}(t) &= (x_{1;t}, \dots, x_{n;t}) \in \mathcal{F}^n, \\ \mathbf{y}(t) &= (y_{1;t}, \dots, y_{n;t}) \in \text{Trop}(x_{n+1}, \dots, x_m)^n, \\ B(t) &= (b_{ij}^t), \end{aligned}$$

with the initial cluster $\mathbf{x}(t_0) = (x_1, \dots, x_n)$, and with the frozen variables x_{n+1}, \dots, x_m . Define $\bar{\mathbf{x}}(t) = (\bar{x}_{1;t}, \dots, \bar{x}_{n;t})$ and $\bar{\mathbf{y}}(t) = (\bar{y}_{1;t}, \dots, \bar{y}_{n;t})$ by

$$(4.3.2) \quad \bar{x}_{i;t} = \frac{\psi(x_{i;t})}{\varphi(x_{i;t})},$$

$$(4.3.3) \quad \bar{y}_{k;t} = \varphi(\hat{y}_{k;t}) = \varphi(y_{k;t}) \prod_{i=1}^n \varphi(x_{i;t})^{b_{ik}^t}.$$

Then $(\bar{\mathbf{x}}(t), \bar{\mathbf{y}}(t), B(t))_{t \in \mathbb{T}_n}$ is a seed pattern in $\bar{\mathcal{F}}$, with the same exchange matrices $B(t)$ and with the frozen variables $\bar{x}_{n+1}, \dots, \bar{x}_{\bar{m}}$.

Proof. By Corollary 3.5.5, the elements $\hat{y}_{i;t}$ satisfy the Y-pattern recurrences. It follows that their images under the semifield homomorphism φ , cf. (4.3.3), satisfy (3.6.3), the tropical version of these recurrences. It remains to check that the elements $\bar{x}_{i;t}$ and $\bar{y}_{j;t}$ satisfy the exchange relations

$$(4.3.4) \quad \bar{x}_{k;t} \bar{x}_{k;t'} = \frac{\bar{y}_{k;t}}{\bar{y}_{k;t} \oplus 1} \prod_{b_{ik}^t > 0} \bar{x}_{i;t}^{b_{ik}^t} + \frac{1}{\bar{y}_{k;t} \oplus 1} \prod_{b_{ik}^t < 0} \bar{x}_{i;t}^{-b_{ik}^t},$$

for $t \xrightarrow{k} t'$ (cf. (3.6.4)).

First note that by (4.3.3) and the distributive property for the tropical semifield (Lemma 3.6.3), we have that

$$(4.3.5) \quad \left(\bar{y}_{k;t} \oplus 1\right) \prod_{b_{ik}^t < 0} \varphi(x_{i;t})^{-b_{ik}^t} = \varphi(y_{k;t}) \prod_{b_{ik}^t > 0} \varphi(x_{i;t})^{b_{ik}^t} \oplus \prod_{b_{ik}^t < 0} \varphi(x_{i;t})^{-b_{ik}^t}.$$

We shall deduce (4.3.4) from the exchange relation

$$(4.3.6) \quad x_{k;t} x_{k;t'} = \frac{1}{y_{k;t} \oplus 1} \left(y_{k;t} \prod_{b_{ik}^t > 0} x_{i;t}^{b_{ik}^t} + \prod_{b_{ik}^t < 0} x_{i;t}^{-b_{ik}^t} \right).$$

Applying the semifield homomorphism ψ (resp., φ) to both sides of (4.3.6), dividing respective images by each other, using the fact that ψ and φ agree on the frozen variables x_{n+1}, \dots, x_m , and using (4.3.2) and (4.3.5), we get:

$$\begin{aligned} \bar{x}_{k;t} \bar{x}_{k;t'} &= \frac{\psi(x_{k;t}) \psi(x_{k;t'})}{\varphi(x_{k;t}) \varphi(x_{k;t'})} \\ &= \frac{\varphi(y_{k;t} \oplus 1)}{\psi(y_{k;t} \oplus 1)} \cdot \frac{\psi(y_{k;t}) \prod_{b_{ik}^t > 0} \psi(x_{i;t})^{b_{ik}^t} + \prod_{b_{ik}^t < 0} \psi(x_{i;t})^{-b_{ik}^t}}{\varphi(y_{k;t}) \prod_{b_{ik}^t > 0} \varphi(x_{i;t})^{b_{ik}^t} \oplus \prod_{b_{ik}^t < 0} \varphi(x_{i;t})^{-b_{ik}^t}} \\ &= \frac{1}{\bar{y}_{k;t} \oplus 1} \cdot \frac{\varphi(y_{k;t}) \prod_{b_{ik}^t > 0} \psi(x_{i;t})^{b_{ik}^t} + \prod_{b_{ik}^t < 0} \psi(x_{i;t})^{-b_{ik}^t}}{\prod_{b_{ik}^t < 0} \varphi(x_{i;t})^{-b_{ik}^t}} \\ &= \frac{1}{\bar{y}_{k;t} \oplus 1} \left(\varphi(y_{k;t}) \prod_{b_{ik}^t > 0} \psi(x_{i;t})^{b_{ik}^t} \prod_{b_{ik}^t < 0} \varphi(x_{i;t})^{b_{ik}^t} + \prod_{b_{ik}^t < 0} \bar{x}_{i;t}^{-b_{ik}^t} \right) \\ &= \frac{1}{\bar{y}_{k;t} \oplus 1} \left(\bar{y}_{k;t} \prod_{b_{ik}^t > 0} \bar{x}_{i;t}^{b_{ik}^t} + \prod_{b_{ik}^t < 0} \bar{x}_{i;t}^{-b_{ik}^t} \right). \quad \square \end{aligned}$$

We next restate Theorem 4.3.4 in terms of extended exchange matrices.

Proposition 4.3.5. *Keep the assumptions and notation of Theorem 4.3.4. Let $\tilde{B}_0 = \tilde{B}(t_0)$ be the initial extended exchange matrix of the original exchange pattern. Then the new initial seed $(\bar{\mathbf{x}}(t_0), \bar{\mathbf{y}}(t_0), B(t_0))$ has the extended exchange matrix $\bar{\tilde{B}}_0 = \Psi \tilde{B}_0$ where $\Psi = (\psi_{ij})$ is the $\bar{m} \times m$ matrix whose entries ψ_{ij} are defined by*

$$(4.3.7) \quad \psi(x_j) = \prod_{i=1}^{\bar{m}} \bar{x}_i^{\psi_{ij}}.$$

(Note that (4.3.1) implies that $\psi(x_j)$ is a Laurent monomial in $\bar{x}_1, \dots, \bar{x}_{\bar{m}}$.)

Proof. We use the notation $\overline{\tilde{B}}_\circ = (\bar{b}_{ik})$ and $\tilde{B}_\circ = (b_{ik})$. Recall that $\bar{b}_{ik} = b_{ik}$ for $i \leq n$. We have:

$$\begin{aligned} \prod_{i \leq \bar{m}} \bar{x}_i^{\bar{b}_{ik}} &= \bar{y}_{k;t_\circ} \prod_{j \leq n} \bar{x}_j^{\bar{b}_{jk}} = \varphi(y_{k;t_\circ}) \prod_{j \leq n} \varphi(x_j)^{b_{jk}} \prod_{j \leq n} \bar{x}_j^{b_{jk}} \\ &= \prod_{j \leq m} \varphi(x_j)^{b_{jk}} \prod_{j \leq n} \bar{x}_j^{b_{jk}} = \prod_{j \leq m} \psi(x_j)^{b_{jk}} = \prod_{i \leq \bar{m}} \prod_{j \leq m} \bar{x}_i^{\psi_{ij} b_{jk}}, \end{aligned}$$

establishing that $\overline{\tilde{B}}_\circ = \Psi \tilde{B}_\circ$. (Here we used (3.6.2) and (4.3.7).) \square

The following corollary will be particularly useful in the sequel; see the overview of Chapter 5 which follows Theorem 5.2.12.

Corollary 4.3.6. *Consider two seed patterns*

$$\begin{aligned} (\Sigma(t))_{t \in \mathbb{T}_n} &= (\mathbf{x}(t), \mathbf{y}(t), B(t))_{t \in \mathbb{T}_n}, \\ (\bar{\Sigma}(t))_{t \in \mathbb{T}_n} &= (\bar{\mathbf{x}}(t), \bar{\mathbf{y}}(t), \bar{B}(t))_{t \in \mathbb{T}_n} \end{aligned}$$

with the same exchange matrices $B(t) = \bar{B}(t)$, for $t \in \mathbb{T}_n$. Suppose that all rows of the initial extended exchange matrix \tilde{B}_\circ for the second seed pattern lie in the \mathbb{Z} -span of the rows of the initial extended exchange matrix \tilde{B}_\circ for the first seed pattern. If two labeled (resp., unlabeled) seeds $\Sigma(t) = \Sigma(t')$ coincide in the first pattern, then the corresponding seeds $\bar{\Sigma}(t) = \bar{\Sigma}(t')$ in the second pattern coincide as well.

Proof. We use the notation of Proposition 4.3.5. In particular, we have initial clusters $\mathbf{x}(t_\circ) = (x_1, \dots, x_n)$ and $\bar{\mathbf{x}}(t_\circ) = (\bar{x}_1, \dots, \bar{x}_n)$, with frozen variables $x_{n+1}, \dots, x_{\bar{m}}$ and $\bar{x}_{n+1}, \dots, \bar{x}_{\bar{m}}$ respectively. We also have exchange matrices \tilde{B}_\circ and $\bar{\tilde{B}}_\circ$, which are $m \times n$ and $\bar{m} \times n$ matrices whose top $n \times n$ submatrices equal $B(t_\circ)$. Recall from the discussion surrounding (3.6.2) that an extended exchange matrix contains the same information as its top $n \times n$ submatrix together with the coefficient tuple \mathbf{y} .

The \mathbb{Z} -span condition in our hypothesis means that $\overline{\tilde{B}}_\circ = \Psi \tilde{B}_\circ$ where $\Psi = (\psi_{ij})$ is an integer $\bar{m} \times m$ matrix. Since the top $n \times n$ submatrices of \tilde{B}_\circ and $\bar{\tilde{B}}_\circ$ coincide, we may assume that Ψ is a block matrix of the form

$$\Psi = \begin{bmatrix} I & 0 \\ \Psi_1 & \Psi_2 \end{bmatrix}$$

where I is the $n \times n$ identity matrix.

We then define the maps ψ and φ by

$$\psi(x_j) = \prod_{i=1}^{\bar{m}} \bar{x}_i^{\psi_{ij}}, \quad \varphi(x_j) = \prod_{i=n+1}^{\bar{m}} \bar{x}_i^{\psi_{ij}},$$

for $1 \leq j \leq m$ (cf. (4.3.7)), so as to agree with (4.3.1) and (4.3.7).

We claim that initial seeds $(\bar{\mathbf{x}}(t_0), \bar{\mathbf{y}}(t_0), \bar{B}(t_0))$ and $(\mathbf{x}(t_0), \mathbf{y}(t_0), B(t_0))$ are related via formulas (4.3.2)–(4.3.3); once we know the claim, it follows from Theorem 4.3.4 that *each* seed $(\bar{\mathbf{x}}(t), \bar{\mathbf{y}}(t), \bar{B}(t))$ of the second seed pattern is related to the corresponding seed $(\mathbf{x}(t), \mathbf{y}(t), B(t))$ of the first seed pattern via (4.3.2)–(4.3.3). And this implies the corollary.

To verify that our initial seeds satisfy (4.3.2), note that for $1 \leq j \leq n$, we have $\frac{\psi(x_j)}{\varphi(x_j)} = \prod_{i=1}^n \bar{x}_i^{\psi_{ij}}$, which is equal to \bar{x}_j , since Ψ restricts to the identity matrix on the first n rows and columns. To verify that our initial seeds also satisfy (4.3.3), we need to use (3.6.2) to relate the y -variables to the extended cluster variables. We find that the left-hand side of (4.3.3) becomes $\bar{y}_k = \prod_{\ell=n+1}^{\bar{m}} \bar{x}_\ell^{\bar{b}_{\ell k}}$, while the right-hand side becomes $\varphi(y_k) \prod_{i=1}^n \varphi(x_i)^{b_{ik}} = \varphi(\prod_{i=n+1}^m x_i^{b_{ik}}) \prod_{i=1}^n \varphi(x_i)^{b_{ik}} = \prod_{i=1}^m \prod_{\ell=n+1}^{\bar{m}} \bar{x}_\ell^{\psi_{\ell i} b_{ik}}$. We can see that the two sides are equal by using the fact that $\bar{B}_0 = \Psi \tilde{B}_0$, or equivalently, that $\bar{b}_{\ell k} = \sum_{i=1}^m \psi_{\ell i} b_{ik}$. \square

Remark 4.3.7. The \mathbb{Z} -span condition in Corollary 4.3.6 is in particular satisfied if the initial extended exchange matrix $\tilde{B}_0 = \tilde{B}(t_0)$ has *full \mathbb{Z} -rank*, i.e., the \mathbb{Z} -span of its rows is the entire lattice \mathbb{Z}^n of integer row-vectors.

This condition is also satisfied if the second seed pattern has trivial coefficients, i.e., if it has no frozen variables.

Exercise 4.3.8. Use Corollary 4.3.6 to show that any seed pattern with exchange matrices

$$(4.3.8) \quad B(t) = (-1)^t \begin{bmatrix} 0 & 1 \\ -c & 0 \end{bmatrix},$$

with $c \in \{1, 2, 3\}$, has finitely many distinct seeds. To this end, consider the seed with the initial extended exchange matrix

$$\tilde{B}_0 = \begin{bmatrix} 0 & 1 \\ -c & 0 \\ 1 & 0 \end{bmatrix}.$$

Note that the rows of \tilde{B}_0 span \mathbb{Z}^2 over \mathbb{Z} , so by Remark 4.3.7, it suffices to check that the seed pattern defined by \tilde{B}_0 has finitely many seeds. (In fact, it has five seeds for $c = 1$, six seeds for $c = 2$, and eight seeds for $c = 3$.)

Remark 4.3.9. In Corollary 4.3.6, the \mathbb{Z} -span condition can be replaced by one involving a \mathbb{Q} -span. The proof remains essentially the same but requires allowing rational powers of the variables. This can be handled in two alternative ways: algebraically, by working in the semifield of “subtraction-free Puiseux expressions;” or analytically, by always choosing the branch of a fractional power that takes positive values at positive arguments.

4.4. Folding

Folding is a procedure that, under certain conditions, produces new seed patterns from existing ones. The basic idea is to exploit symmetries of a quiver to construct a quotient object (a folded extended exchange matrix), then design an equivariant mutation dynamics that would drive the algebraic dynamics of “folded seeds.”

In this text, we discuss folding in a somewhat limited generality; see [6] for a more elaborate treatment. Our main application of folding will occur in Chapter 5, where we use it to construct cluster algebras of finite types *BCFG* from those of “simply-laced” types *ADE*. Folding has also been used in [7] to classify the cluster algebras of finite mutation type; see Section 10.2.

We begin with a motivating example. Consider the quiver

$$1 \longleftarrow 2 \longrightarrow 3$$

of type A_3 , with three mutable vertices. We notice the $\mathbb{Z}/2\mathbb{Z}$ symmetry of the quiver, and place only *two* distinct variables at its vertices, as follows:

$$x_0 \longleftarrow x_1 \longrightarrow x_0.$$

The exchange relations then become:

$$\begin{aligned} x_0 x'_0 &= x_1 + 1, \\ x_1 x'_1 &= x_0^2 + 1. \end{aligned}$$

If we want to preserve the symmetry, we can now mutate either at vertex 2, or simultaneously at 1 and 3. Continuing in this fashion, we recover the seed pattern from Example 3.2.7.

Definition 4.4.1. Let Q be a labeled quiver, as in Definition 2.7.1. More explicitly, we assume that Q has m vertices labeled $1, \dots, m$; the vertices labeled $1, \dots, n$ are mutable; the vertices labeled $n+1, \dots, m$ are frozen. Let G be a group acting on the vertex set of Q , or equivalently on $\{1, \dots, m\}$. (For all practical purposes, it is safe to assume that the group G is finite.) The notation $i \sim i'$ will mean that i and i' lie in the same G -orbit. We say that the quiver Q (or the corresponding $m \times n$ extended exchange matrix $\tilde{B} = \tilde{B}(Q) = (b_{ij})$) is *G-admissible* if

- (1) for any $i \sim i'$, index i is mutable (i.e., $i \leq n$) if and only if i' is;
- (2) for any indices i and j , and any $g \in G$, we have $b_{ij} = b_{g(i),g(j)}$;
- (3) for mutable indices $i \sim i'$, we have $b_{ii'} = 0$;
- (4) for any $i \sim i'$, and any mutable j , we have $b_{ij} b_{i'j} \geq 0$.

Assume that Q is G -admissible. We call a G -orbit *mutable* (resp., *frozen*) if it consists of mutable (resp., frozen) vertices, cf. condition (1) above. Let $\tilde{B}^G = \tilde{B}(Q)^G = (b_{IJ}^G)$ be the matrix whose rows (resp., columns) are labeled by the G -orbits (resp., mutable G -orbits), and whose entries are given by

$$(4.4.1) \quad b_{IJ}^G = \sum_{i \in I} b_{ij},$$

where j is an arbitrary index in J . (By condition (2), the right-hand side of (4.4.1) does not depend on the choice of j .) We then say that \tilde{B}^G is obtained from \tilde{B} (or from the quiver Q) by *folding* with respect to the given G -action.

An example of folding is shown in Figure 4.2.




Figure 4.2. The quiver Q shown on the left is G -admissible with respect to the action of the group $G = \mathbb{Z}/2\mathbb{Z}$ wherein the generator of G acts on the vertices of Q by a 180° rotation. All 5 vertices are mutable.

Remark 4.4.2. Condition (3) can be restated as saying that each G -orbit I is *totally disconnected*; that is, there is no arrow between two vertices in I . Condition (4) means that there is no oriented path of length 2 through a mutable vertex that connects two vertices belonging to the same G -orbit. These conditions are dictated by the following considerations. If i and i' are in the same G -orbit I , then in the folded seed (to be defined below in this section), the same variable x_I will be associated with both i and i' . We do not want x_I to appear on the right-hand side of the exchange relation for x_I (hence (3)), nor do we want x_I to appear in both monomials on the right-hand side of the exchange relation for some variable x_J (hence (4)).

Lemma 4.4.3. *Let Q be a G -admissible quiver. Then $\tilde{B}(Q)^G$ is an extended skew-symmetrizable matrix.*

Proof. As above, we use the notation $\tilde{B} = \tilde{B}(Q) = (b_{ij})$. Formula (4.4.1) implies that, for I and J mutable,

$$(4.4.2) \quad |J| b_{IJ}^G = \sum_{i \in I} \sum_{j \in J} b_{ij} = - \sum_{j \in J} \sum_{i \in I} b_{ji} = -|I| b_{JI}^G,$$

so the square matrix $(|J| b_{IJ}^G)$ is skew-symmetric, and hence $\tilde{B}(Q)^G = (b_{IJ}^G)$ is skew-symmetrizable. \square

Lemma 4.4.4. *Let Q be a G -admissible quiver, with $\tilde{B}(Q) = (b_{ij})$. Let I and J be two G -orbits, with J mutable. Then either all entries b_{ij} , for $i \in I$ and $j \in J$, are nonnegative, or all are nonpositive. Consequently the following are equivalent:*

- $b_{IJ}^G > 0$;
- there exist $i \in I$ and $j \in J$ such that $b_{ij} > 0$;
- for every $i \in I$, there exists $j \in J$ such that $b_{ij} > 0$.

Similar equivalences hold with all inequality signs reversed.

Proof. Let $i, i' \in I$ and $j, j' \in J$. Let $g \in G$ be such that $g(j') = j$. Using conditions (2) and (4) of Definition 4.4.1, we get $b_{ij}b_{i'j'} = b_{ij}b_{g(i'),j} \geq 0$, as desired. The equivalence statements in the lemma follow by (4.4.1) together with condition (2). \square

For Q a G -admissible quiver, condition (3) of Definition 4.4.1 ensures (cf. Exercise 2.1.4(3)) that the result of mutating Q at the set of all vertices in a mutable G -orbit K does not depend on the order of mutations. We will denote this composition of mutations by

$$(4.4.3) \quad \mu_K = \prod_{k \in K} \mu_k,$$

making sure to only use this notation when it is well defined.

Lemma 4.4.5. *Let Q be a G -admissible quiver, with $\tilde{B} = \tilde{B}(Q)$. Let K be a mutable G -orbit such that $\mu_K(Q)$ is also G -admissible. Then*

$$(\mu_K(\tilde{B}))^G = \mu_K(\tilde{B}^G).$$

Note the abuse of notation in the last equality: on the left, μ_K is a composition of mutations defined by (4.4.3); on the right, μ_K is a single mutation at the mutable index K of the folded matrix \tilde{B}^G .

Proof. Using the definition of matrix mutation (2.7.1) in combination with Lemma 4.4.4, we obtain:

$$(4.4.4) \quad \mu_K(\tilde{B}^G)_{IJ} = \begin{cases} -b_{IJ}^G & \text{if } K \in \{I, J\}; \\ b_{IJ}^G + b_{IK}^G b_{KJ}^G & \text{if } K \notin \{I, J\} \text{ and } b_{ik} > 0 \text{ and } b_{kj} > 0 \\ & \text{for some } i \in I, j \in J, k \in K; \\ b_{IJ}^G - b_{IK}^G b_{KJ}^G & \text{if } K \notin \{I, J\} \text{ and } b_{ik} < 0 \text{ and } b_{kj} < 0 \\ & \text{for some } i \in I, j \in J, k \in K; \\ b_{IJ}^G & \text{otherwise.} \end{cases}$$

On the other hand, mutating \tilde{B} at each $k \in K$ (recall that K is totally disconnected), we get:

$$(4.4.5) \quad \mu_K(\tilde{B})_{ij} = \begin{cases} -b_{ij} & \text{if } i \in K \text{ or } j \in K \text{ (i.e., } K \in \{I, J\}); \\ b_{ij} + \sum_{k \in K} b_{ik}b_{kj} & \text{if } i \notin K, j \notin K \text{ (i.e., } K \notin \{I, J\}), \text{ and} \\ & b_{ik} > 0 \text{ and } b_{kj} > 0 \text{ for some } k \in K; \\ b_{ij} - \sum_{k \in K} b_{ik}b_{kj} & \text{if } i \notin K, j \notin K \text{ (i.e., } K \notin \{I, J\}), \text{ and} \\ & b_{ik} < 0 \text{ and } b_{kj} < 0 \text{ for some } k \in K; \\ b_{ij} & \text{otherwise.} \end{cases}$$

By (4.4.1) (recall that $\mu_K(\tilde{B})$ is G -admissible), we have, for any $j \in J$:

$$((\mu_K(\tilde{B}))^G)_{IJ} = \sum_{i \in I} (\mu_K(\tilde{B}))_{ij}.$$

We note furthermore that all terms in the last sum will fall into the same case in (4.4.5). It is now easy to check that $(\mu_K(\tilde{B}^G))_{IJ} = ((\mu_K(\tilde{B}))^G)_{IJ}$. For example, in the second case of (4.4.5), we have:

$$\begin{aligned} ((\mu_K(\tilde{B}))^G)_{IJ} &= \sum_{i \in I} (b_{ij} + \sum_{k \in K} b_{ik}b_{kj}) \\ &= b_{IJ}^G + \sum_{k \in K} b_{kj} \sum_{i \in I} b_{ik} \\ &= b_{IJ}^G + b_{IK}^G b_{KJ}^G \\ &= \mu_K(\tilde{B}^G)_{IJ}. \end{aligned} \quad \square$$

The ‘‘mutation commutes with folding’’ statement in Lemma 4.4.5 comes with a caveat: it requires admissibility of *both* quivers Q and $\mu_K(Q)$ with respect to the group action at hand. Unfortunately, admissibility does not propagate via mutations: Q may be G -admissible while $\mu_K(Q)$ is not.

Example 4.4.6. Let Q be an oriented 6-cycle with six mutable vertices labeled 0 through 5 in clockwise order. Let the generator of $G = \mathbb{Z}/2\mathbb{Z}$ act by sending each vertex i to the vertex $(i+3) \bmod 6$. Then Q is G -admissible but for any mutable $\mathbb{Z}/2\mathbb{Z}$ -orbit K , the quiver $\mu_K(Q)$ is not.

We next proceed to the folding of seeds. This will require, in addition to an action of a group G , a choice of a semifield homomorphism that ‘‘bundles together’’ the variables associated with the vertices in the same G -orbit.

Definition 4.4.7. Let G be a group acting on the set of indices $\{1, \dots, m\}$ so that every $g \in G$ maps the subset $\{1, \dots, n\}$ to itself. Let m^G denote the number of orbits of this action. Let \mathcal{F} (resp., \mathcal{F}^G) be a field isomorphic to the

field of rational functions in m (resp., m^G) independent variables, and let \mathcal{F}_{sf} (resp., $\mathcal{F}_{\text{sf}}^G$) denote the corresponding semifields of subtraction-free rational expressions. Let $\psi : \mathcal{F}_{\text{sf}} \rightarrow \mathcal{F}_{\text{sf}}^G$ be a surjective semifield homomorphism.

Let Q be a quiver as above. A seed $\Sigma = (\tilde{\mathbf{x}}, \tilde{B}(Q))$ in \mathcal{F} , with the extended cluster $\tilde{\mathbf{x}} = (x_i)$, is called (G, ψ) -admissible if

- Q is a G -admissible quiver;
- for any $i \sim i'$, we have $\psi(x_i) = \psi(x_{i'})$.

In this situation, we define a new “folded” seed $\Sigma^G = (\tilde{\mathbf{x}}^G, \tilde{B}^G)$ in $\mathcal{F}_{\text{sf}}^G \subset \mathcal{F}^G$ whose extended exchange matrix \tilde{B}^G is given by (4.4.1), and whose extended cluster $\tilde{\mathbf{x}}^G = (x_I)$ has m^G elements x_I indexed by the G -orbits and defined by $x_I = \psi(x_i)$, for $i \in I$. Note that since ψ is a surjective homomorphism, the elements x_I generate \mathcal{F}^G , hence are algebraically independent.

We can now extend Lemma 4.4.5 to the folding of seeds.

Lemma 4.4.8. *Let $\Sigma = (\tilde{\mathbf{x}}, \tilde{B}(Q))$ be a (G, ψ) -admissible seed as above. Let K be a mutable G -orbit. If the quiver $\mu_K(Q)$ is G -admissible, then the seed $\mu_K(\Sigma)$ is (G, ψ) -admissible, and moreover $(\mu_K(\Sigma))^G = \mu_K(\Sigma^G)$.*

Proof. Let \mathcal{O} denote the set of G -orbits. As above, we use the notation

$$\begin{aligned} \Sigma &= (\tilde{\mathbf{x}}, \tilde{B}), & \tilde{\mathbf{x}} &= (x_i), & \tilde{B} &= \tilde{B}(Q) = (b_{ij}), \\ \Sigma^G &= (\tilde{\mathbf{x}}^G, \tilde{B}^G), & \tilde{\mathbf{x}}^G &= (x_I)_{I \in \mathcal{O}}, & \tilde{B}^G &= (b_{IJ}^G) \end{aligned}$$

By Lemma 4.4.5, all we need to show is that the extended clusters in $\mu_K(\Sigma^G)$ and in $(\mu_K(\Sigma))^G$ are the same. The extended cluster in $\mu_K(\Sigma^G)$ is obtained from $\tilde{\mathbf{x}}^G$ by replacing x_K by the element x'_K defined by

$$(4.4.6) \quad x_K x'_K = \prod_{b_{IK}^G > 0} x_I^{b_{IK}^G} + \prod_{b_{IK}^G < 0} x_I^{-b_{IK}^G}.$$

The extended cluster in $\mu_K(\Sigma)$ is obtained from $\tilde{\mathbf{x}}$ by replacing each x_k , for $k \in K$, by the element x'_k defined by

$$(4.4.7) \quad x_k x'_k = \prod_{b_{ik} > 0} x_i^{b_{ik}} + \prod_{b_{ik} < 0} x_i^{-b_{ik}}.$$

The extended cluster in $(\mu_K(\Sigma))^G$ is then obtained by applying the homomorphism ψ ; as we know, ψ sends each x_i to $x_{[i]}$ where $[i]$ denotes the G -orbit containing i . As a result, the cluster contains the variables x_I for $I \in \mathcal{O} - \{K\}$, together with $\psi(x'_k)$; here k is an arbitrary element of K . We need to show that $\psi(x'_k) = x'_K$ where x'_K is defined by (4.4.6).

Applying ψ to (4.4.7), we get

$$(4.4.8) \quad x_K \psi(x'_k) = \prod_{b_{ik} > 0} x_{[i]}^{b_{ik}} + \prod_{b_{ik} < 0} x_{[i]}^{-b_{ik}}.$$

Note that in the first monomial on the right-hand side of (4.4.8), the exponent of a given x_I will be nonzero if and only if there exists $i \in I$ such that $b_{ik} > 0$, in which case the exponent will be $\sum_{i \in I} b_{ik} = b_{IK}^G$. But by Lemma 4.4.4, the existence of such an index $i \in I$ is equivalent to the inequality $b_{IK}^G > 0$. It follows that the first monomials in the right-hand sides of (4.4.8) and (4.4.6) agree. A similar argument shows that the second monomials agree, and we are done. \square

Definition 4.4.9. Let G be a group acting on the vertex set of a quiver Q . We say that Q is *globally foldable* with respect to G if Q is G -admissible and moreover for any sequence of mutable G -orbits J_1, \dots, J_k , the quiver $(\mu_{J_k} \circ \dots \circ \mu_{J_1})(Q)$ is G -admissible.

Exercise 4.4.10. For the quiver Q and the action of the group $G = \mathbb{Z}/2\mathbb{Z}$ described in Figure 4.2, show that Q is globally foldable with respect to G .

Lemma 4.4.8 implies that if Q is globally foldable, then we can fold all the seeds in the corresponding seed pattern.

Corollary 4.4.11. *Let Q be a quiver which is globally foldable with respect to a group G acting on the set of its vertices. Let $\Sigma = (\tilde{\mathbf{x}}, \tilde{B}(Q))$ be a seed in the field \mathcal{F} of rational functions freely generated by an extended cluster $\tilde{\mathbf{x}} = (x_i)$. Let $\tilde{\mathbf{x}}^G = (x_I)$ be a collection of formal variables labeled by the G -orbits I , and let \mathcal{F}^G denote the field of rational functions in these variables. Define the surjective homomorphism*

$$\begin{aligned} \psi : \mathcal{F}_{\text{sf}} &\rightarrow \mathcal{F}_{\text{sf}}^G \\ x_i &\mapsto x_I \quad (i \in I) \end{aligned}$$

of the corresponding semifields of subtraction-free rational expressions, so that Σ is a (G, ψ) -admissible seed. Then for any mutable G -orbits J_1, \dots, J_k , the seed $(\mu_{J_k} \circ \dots \circ \mu_{J_1})(\Sigma)$ is (G, ψ) -admissible, and moreover the folded seeds $((\mu_{J_k} \circ \dots \circ \mu_{J_1})(\Sigma))^G$ form a seed pattern in \mathcal{F}^G , with the initial extended exchange matrix $(\tilde{B}(Q))^G$.

In general, it may be very difficult to determine whether a quiver is globally foldable. Fortunately, the cases that we will need in Chapter 5 for the purposes of finite type classification will turn out to be easy to handle. One of these cases is discussed in Exercise 4.4.12 below. We will revisit this exercise in Section 5.7.

Exercise 4.4.12. Let Q be the quiver at the left of Figure 4.3, with all its vertices mutable. Let the generator of the group $G = \mathbb{Z}/2\mathbb{Z}$ act on the vertices of Q by fixing the vertices 1 and 2, exchanging the vertices 3 and 4, and exchanging the vertices 5 and 6. Then Q is G -admissible; the folded skew-symmetrizable matrix $B(Q)^G$ is shown at the right in Figure 4.3. (The rows and columns of this matrix are labeled by the orbits $\{1\}, \{2\}, \{3, 4\}, \{5, 6\}$.) Show that Q is globally foldable, by cataloguing all quivers (up to G -equivariant isomorphism) which can be obtained from Q by iterating the transformations μ_J associated with G -orbits J .




Figure 4.3. Folding a type E_6 quiver.

Remark 4.4.13. There exist skew-symmetrizable square matrices which cannot be obtained by folding a globally foldable quiver, see [22, Section 14]. Consequently the technique of folding (in the current state of the art) is not powerful enough to reduce the study of general seed patterns, and the associated cluster algebras, to the quiver case.

Finite type classification

A seed pattern (or the corresponding cluster algebra) is said to be of *finite type* if it has finitely many different seeds. To rephrase, a seed gives rise to a pattern (or cluster algebra) of finite type if the process of iterated mutation produces finitely many distinct seeds.

If a seed pattern has finite type, then it obviously has finitely many distinct cluster variables. In fact the converse is also true, see Proposition 5.9.4.

The main result of this chapter is the classification, originally obtained in [11], of seed patterns (equivalently cluster algebras) of finite type. It turns out that the property of a cluster algebra with an initial seed $(\mathbf{x}, \mathbf{y}, B)$ to be of finite type depends only on (the mutation class of) the exchange matrix B but not on the choice of a coefficient tuple \mathbf{y} . Even more remarkably, such mutation classes are in one-to-one correspondence with Cartan matrices of finite type, or equivalently with finite crystallographic root systems.

To show that a cluster algebra of finite type must come from a Cartan matrix of finite type, we follow the approach of [11]. In particular we show that a cluster algebra has finite type if and only if all its exchange matrices $B = (b_{ij})$ have the property that $|b_{ij}b_{ji}| \leq 3$ for any pair of indices i, j .

The fact that cluster algebras coming from finite-type Cartan matrices are of finite type is established by a string of case-by-case arguments. For each type, we construct a particular seed pattern that has finitely many seeds, and whose initial extended exchange matrix has full \mathbb{Z} -rank. A generalization to arbitrary coefficients is then obtained via Remark 4.3.7. We note that the original proof [11] of this direction of the classification theorem used a different (root-theoretic) strategy relying on some rather delicate properties of *generalized associahedra*, cf. Section 9.3.

We are not aware of a simple argument that would directly derive the classification of cluster algebras of finite type from one of the instances of the classical Cartan-Killing classification (see Theorem 5.2.6).

5.1. Finite type classification in rank 2

Any seed pattern of rank 1 has two seeds, so is of finite type. In this section, we determine which seed patterns of rank 2 are of finite type.

Recall that the seeds in any seed pattern of rank 2 can be labeled by integers $t \in \mathbb{Z}$. Without loss of generality, we may assume that the exchange matrices $B(t)$ are given by

$$(5.1.1) \quad B(t) = (-1)^t \begin{bmatrix} 0 & b \\ -c & 0 \end{bmatrix}$$

where the integers b and c are either both positive, or both equal to 0.

Theorem 5.1.1. *A seed pattern of rank 2, with exchange matrices given by (5.1.1), is of finite type if and only if $bc \leq 3$.*

The statement of Theorem 5.1.1 is very similar to the classification of finite crystallographic reflection groups of rank 2, which we recall in Proposition 5.1.2 below. (The proof of Theorem 5.1.1 will not depend on this proposition.)

Proposition 5.1.2. *Consider the subgroup $W \subset \mathrm{GL}_2$ generated by the reflections*

$$(5.1.2) \quad s_1 = \begin{bmatrix} -1 & b \\ 0 & 1 \end{bmatrix}, \quad s_2 = \begin{bmatrix} 1 & 0 \\ c & -1 \end{bmatrix}.$$

(As above, $b, c \in \mathbb{Z}$ are either both positive, or both equal to 0.) Then W is a finite group if and only if $bc \leq 3$.

Proof. Since $s_1^2 = s_2^2 = 1$, the group W is finite if and only if the element

$$s_1 s_2 = \begin{bmatrix} bc - 1 & -b \\ c & -1 \end{bmatrix}$$

is of finite order. An eigenvalue λ of $s_1 s_2$ satisfies the characteristic equation

$$(5.1.3) \quad \lambda^2 - (bc - 2)\lambda + 1 = 0.$$

If $bc \in \{1, 2, 3\}$, then the roots of (5.1.3) are two distinct roots of unity, so $s_1 s_2$ has finite order (specifically, order 3, 4, or 6, respectively). If $b = c = 0$, then $(s_1 s_2)^2 = 1$ by inspection. If $bc > 4$, then the eigenvalues are real and not equal to ± 1 , so W is infinite. For $bc = 4$, one checks that

$$(s_1 s_2)^k = \begin{bmatrix} 2k + 1 & -kb \\ kc & -2k + 1 \end{bmatrix},$$

implying that W is infinite in this case as well. □

Proof of Theorem 5.1.1. The case $b = c = 0$ is trivial. The cases $bc \in \{1, 2, 3\}$ are handled by the calculation from Exercise 4.3.8. An alternative (more conceptual) explanation of why the rank 2 cluster algebras with $bc \leq 3$ are of finite type will be given later in this chapter.

Recall from Exercise 4.3.8 that any seed pattern with exchange matrices

$$(5.1.4) \quad B(t) = (-1)^t \begin{bmatrix} 0 & 1 \\ -c & 0 \end{bmatrix},$$

with $c \in \{1, 2, 3\}$, has finitely many distinct seeds.

Now suppose that $bc \geq 4$. Let $(\mathbf{x}(t), \mathbf{y}(t), B(t))_{t \in \mathbb{Z}}$ be a seed pattern in an ambient field \mathcal{F} whose exchange matrices are given by (5.1.1). We denote the cluster variables in this pattern by z_t ($t \in \mathbb{Z}$), so that

$$\dots, \mathbf{x}(0) = (z_1, z_2), \quad \mathbf{x}(1) = (z_3, z_2), \quad \mathbf{x}(2) = (z_3, z_4), \quad \mathbf{x}(3) = (z_5, z_4), \dots$$

Our goal is to show that the set $\{z_t : t \in \mathbb{Z}\} \subset \mathcal{F}$ is infinite. (In fact, all cluster variables z_t turn out to be distinct.)

Let u be a formal variable, and consider the semifield $U = \{u^r : r \in \mathbb{R}\}$ of formal monomials in u with real exponents, with the operations defined by

$$\begin{aligned} u^r \oplus u^s &= u^{\max(r,s)}, \\ u^r \cdot u^s &= u^{r+s}. \end{aligned}$$

We shall prove that the set $\{z_t\}$ is infinite by constructing a semifield homomorphism $\psi : \mathcal{F} \rightarrow U$ such that the image $\{\psi(z_t) : t \in \mathbb{Z}\} \subset U$ is infinite. We consider the cases $bc > 4$ and $bc = 4$ separately.

Case 1: $bc > 4$. In this case, there is a real number $\lambda > 1$ satisfying (5.1.3). Let ψ be the map uniquely defined by setting the image of every frozen variable to be 1, and setting

$$(5.1.5) \quad \psi(z_1) = u^c, \quad \psi(z_2) = u^{\lambda+1}.$$

The exchange relations imply that the images $\psi(z_t)$ satisfy

$$(5.1.6) \quad \psi(z_{t-1}) \psi(z_{t+1}) = \begin{cases} \psi(z_t)^c \oplus 1 & \text{if } t \text{ is even;} \\ \psi(z_t)^b \oplus 1 & \text{if } t \text{ is odd} \end{cases}$$

(cf. (3.2.2)). We claim that, for $k = 0, 1, 2, \dots$, one has

$$(5.1.7) \quad \psi(z_{2k+1}) = u^{\lambda^k c}, \quad \psi(z_{2k+2}) = u^{\lambda^k(\lambda+1)}.$$

The base case $k = 0$ holds in view of (5.1.5). Induction step:

$$\begin{aligned} \psi(z_{2k+3}) &= \frac{\psi(z_{2k+2})^c \oplus 1}{\psi(z_{2k+1})} = u^{\lambda^k(\lambda+1)c - \lambda^k c} = u^{\lambda^{k+1}c}, \\ \psi(z_{2k+4}) &= \frac{\psi(z_{2k+3})^b \oplus 1}{\psi(z_{2k+2})} = u^{\lambda^{k+1}cb - \lambda^k(\lambda+1)} = u^{\lambda^k(\lambda cb - \lambda - 1)} = u^{\lambda^{k+1}(\lambda+1)}, \end{aligned}$$

where the last equality relies on (5.1.3). Since $\lambda > 1$, formulas (5.1.7) imply that the image set $\{\psi(z_t)\} \subset U$ is infinite.

Case 2: $bc = 4$. While the general logic of the argument remains the same, it has to be adjusted since in this case, the only root of the equation (5.1.3) is $\lambda = 1$. So we replace (5.1.5) by

$$\psi(z_1) = u, \quad \psi(z_2) = u^b,$$

and verify by induction that $\psi(z_{2k-1}) = u^{2k-1}$ and $\psi(z_{2k+2}) = u^{(k+1)b}$, implying the claim. Details are left to the reader. \square

Theorem 5.1.1 suggests the following definition.

Definition 5.1.3. A skew-symmetrizable matrix $B = (b_{ij})$ is called *2-finite* if and only if for any matrix B' mutation equivalent to B and any indices i and j , we have $|b'_{ij}b'_{ji}| \leq 3$.

Corollary 5.1.4. *In a seed pattern of finite type, every exchange matrix is 2-finite.*

Proof. This is an immediate consequence of Theorem 5.1.1. If an exchange matrix B is mutation equivalent to $B' = (b'_{ij})$ such that $|b'_{ij}b'_{ji}| \geq 4$ for some i and j , then “freezing” all the cluster variables in the corresponding seed except for x_i and x_j , and alternately applying mutations μ_i and μ_j to the corresponding seed, we obtain infinitely many distinct cluster variables, and infinitely many distinct seeds. \square

We will eventually show that the converse to Corollary 5.1.4 holds as well, see Theorem 5.10.1.

5.2. Cartan matrices and Dynkin diagrams

The classification of cluster algebras of finite type turns out to be completely parallel to the famous Cartan-Killing classification of semisimple Lie algebras, finite crystallographic root systems, etc. The latter classification can be found in many books, e.g., [15, 17]. In this section, we quickly review it, using the language of *Cartan matrices* and *Dynkin diagrams*. We then explain the connection between (symmetrizable) Cartan matrices and (skew-symmetrizable) exchange matrices.

Definition 5.2.1. A square integer matrix $A = (a_{ij})$ is called a *symmetrizable generalized Cartan matrix* if it satisfies the following conditions:

- all diagonal entries of A are equal to 2;
- all off-diagonal entries of A are non-positive;
- there exists a diagonal matrix D with positive diagonal entries such that the matrix DA is symmetric.

We call A *positive* if DA is positive definite; this is equivalent to the positivity of all principal minors $\Delta_{I,I}(A)$. In particular, any such matrix satisfies

$$\Delta_{\{i,j\},\{i,j\}}(A) = \det \begin{pmatrix} 2 & a_{ij} \\ a_{ji} & 2 \end{pmatrix} = 4 - a_{ij}a_{ji} > 0,$$

or equivalently

$$(5.2.1) \quad a_{ij}a_{ji} \leq 3 \text{ for } i \neq j.$$

Positive symmetrizable generalized Cartan matrices are often referred to simply as *Cartan matrices*, or *Cartan matrices of finite type*.

Example 5.2.2. In view of (5.2.1), a 2×2 matrix

$$A = \begin{bmatrix} 2 & -b \\ -c & 2 \end{bmatrix}$$

is a Cartan matrix of finite type if and only if one of the following holds:

- $b = c = 0$;
- $b = c = 1$;
- $b = 1, c = 2$ or $b = 2, c = 1$;
- $b = 1, c = 3$ or $b = 3, c = 1$.

Note that this matches the classifications in Theorem 5.1.1 and Proposition 5.1.2. The latter match has a well-known explanation: there is a canonical correspondence between Cartan matrices of finite type and finite Weyl groups (or finite crystallographic root systems). The relationship between these objects and cluster algebras of finite type is much more subtle.

Remark 5.2.3. A Cartan matrix encodes essential information about the geometry of a root system (or the corresponding Weyl group). The classification of finite crystallographic root systems (resp., associated reflection groups) can be reduced to classifying Cartan matrices of finite type. This standard material can be found in many books, see, e.g., [16].

Definition 5.2.4. The *Coxeter graph* of an $n \times n$ Cartan matrix A is a simple graph with vertices $1, \dots, n$ in which vertices i and $j \neq i$ are joined by an edge whenever $a_{ij} \neq 0$. If $a_{ij} \in \{0, -1\}$ for all $i \neq j$, then A is uniquely determined by its Coxeter graph. (Such matrices are called *simply-laced*.) If A is not simply-laced but of finite type then, in view of (5.2.1), one needs a little additional information to specify A . This is done by replacing the Coxeter graph of A with its *Dynkin diagram* in which, instead of being connected by a single edge, each pair of vertices i and j with $a_{ij}a_{ji} > 1$ is shown as follows:

$$\begin{array}{l} i \bullet \rightleftarrows j \quad \text{if } a_{ij} = -1 \text{ and } a_{ji} = -2; \\ i \bullet \rightleftarrows j \quad \text{if } a_{ij} = -1 \text{ and } a_{ji} = -3. \end{array}$$

Note that our usage of the terms *Coxeter graph* and *Dynkin diagram* is a bit non-standard: Coxeter graphs are usually defined as edge-labeled graphs, and Dynkin diagrams are often assumed to be connected. (We make no such requirement.) Also, as in [11], we use the conventions of [17] (as opposed to those in [2]) in going between Dynkin diagrams and Cartan matrices.

A couple of examples are shown in Figure 5.1. The notation B_3 and C_3 is explained in Figure 5.2.

$$\begin{array}{cc}
 B_3 & \begin{array}{c} \bullet \longleftarrow \bullet \text{---} \bullet \\ \left[\begin{array}{ccc} 2 & -2 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 2 \end{array} \right] \end{array} \\
 C_3 & \begin{array}{c} \bullet \longrightarrow \bullet \text{---} \bullet \\ \left[\begin{array}{ccc} 2 & -1 & 0 \\ -2 & 2 & -1 \\ 0 & -1 & 2 \end{array} \right] \end{array}
 \end{array}$$

Figure 5.1. Dynkin diagrams and Cartan matrices of types B_3 and C_3 .

Remark 5.2.5. It is important to stress that the meaning of double and triple arrows in a Dynkin diagram is very different from the meaning of multiple arrows in a quiver. A double arrow

$$1 \rightleftarrows 2$$

in a Dynkin diagram corresponds to the submatrix

$$\begin{bmatrix} 2 & -1 \\ -2 & 2 \end{bmatrix}$$

of the associated Cartan matrix. Meanwhile, a double arrow

$$1 \rightrightarrows 2$$

in a quiver corresponds to the submatrix

$$\begin{bmatrix} 0 & 2 \\ -2 & 0 \end{bmatrix}$$

of the associated exchange matrix.

A Cartan matrix is called *indecomposable* if its Dynkin diagram is connected. By a simultaneous permutation of rows and columns, any Cartan matrix A can be transformed into a block-diagonal matrix with indecomposable blocks. This corresponds to decomposing the Dynkin diagram of A

into connected components. The *type* of A (i.e., its equivalence class with respect to simultaneous permutations of rows and columns) is determined by specifying the multiplicity of each type of connected Dynkin diagram in this decomposition. Thus to classify Cartan matrices, one needs to produce the list of all connected Dynkin diagrams. The celebrated *Cartan-Killing classification* asserts that this list is given as follows.

Theorem 5.2.6. *Figure 5.2 gives a complete list of connected Dynkin diagrams corresponding to indecomposable Cartan matrices of finite type.*




Figure 5.2. Dynkin diagrams of indecomposable Cartan matrices. The subscript n indicates the number of nodes in the diagram.

We do not include the proof of Theorem 5.2.6 in this book, nor do we rely on this theorem anywhere in our proofs. We will, however, make extensive use of the standard nomenclature of Dynkin diagrams presented in Figure 5.2. The notation $X_n \sqcup Y_{n'}$ will denote the disjoint union of two Dynkin diagrams X_n and $Y_{n'}$.

The relationship between Cartan matrices and skew-symmetrizable matrices is based on the following definition [11].

Definition 5.2.7. Let $B = (b_{ij})$ be a skew-symmetrizable integer matrix. Its *Cartan counterpart* is the symmetrizable generalized Cartan matrix

$$(5.2.2) \quad A = A(B) = (a_{ij})$$

of the same size, defined by

$$(5.2.3) \quad a_{ij} = \begin{cases} 2 & \text{if } i = j; \\ -|b_{ij}| & \text{if } i \neq j. \end{cases}$$

The main result of this chapter is the following classification of cluster algebras of finite type [11].

Theorem 5.2.8. *A seed pattern (or the corresponding cluster algebra) is of finite type if and only if it contains an exchange matrix B whose Cartan counterpart $A(B)$ (see Definition 5.2.7) is a Cartan matrix of finite type.*

The proof of Theorem 5.2.8 spans Sections 5.3–5.10. An overview of the proof is given below in this section.

One important feature of Theorem 5.2.8 is that the finite type property depends solely on the exchange matrix B but not on the coefficient tuple \mathbf{y} . In other words, the top $n \times n$ submatrix B of an extended exchange matrix \tilde{B} determines whether the seed pattern at hand has finitely many seeds. The bottom part of \tilde{B} has no effect on this property.

Definition 5.2.9. Let X_n be a Dynkin diagram on n vertices. A seed pattern of rank n (or the corresponding cluster algebra) is said to be of (*Cartan-Killing*) *type X_n* if one of its exchange matrices B has Cartan counterpart of type X_n .

Example 5.2.10. The matrices

$$\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ -2 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ -3 & 0 \end{bmatrix}$$

define seed patterns (and cluster algebras) of types $A_1 \sqcup A_1$, A_2 , B_2 , and G_2 , respectively.

Remark 5.2.11. Suppose X_n is simply laced, i.e., is one of the types A_n , D_n , E_6 , E_7 , E_8 . Then a seed pattern is of type X_n if one of its exchange matrices B corresponds to a quiver that is an orientation of a Dynkin diagram of type X_n . We note that by Exercise 2.6.5, all orientations of a tree are mutation equivalent to each other, so if one of them is present in the pattern, then all of them are.

A priori, Definition 5.2.9 allows for the possibility that a given seed pattern is simultaneously of two different types. However, the following companion to Theorem 5.2.8 shows that this cannot happen.

Theorem 5.2.12. *Let B' and B'' be skew-symmetrizable square matrices whose Cartan counterparts $A(B')$ and $A(B'')$ are Cartan matrices of finite type. Then the following are equivalent:*

- (1) *the Cartan matrices $A(B')$ and $A(B'')$ have the same type;*
- (2) *B' and B'' are mutation equivalent.*

This theorem will be proved in Section 5.9.

By Theorem 5.2.8, a seed pattern of finite type must contain exchange matrices whose Cartan counterparts are Cartan matrices of finite type. By Theorem 5.2.12, all these matrices have the same type. Consequently the (Cartan-Killing) type of a seed pattern (resp., cluster algebra) of finite type is unambiguously defined.

Remark 5.2.13. Mutation classes of finite type can be partially ordered via embeddability (cf. Definition 4.1.8), which in turn can be interpreted in the language of cluster subalgebras (cf. Definition 4.2.6). For example, the inclusion of Dynkin diagrams $A_n \subset D_{n+1}$ (see Figure 5.2) yields an embedding $\mathbf{A}_n \leq \mathbf{D}_{n+1}$ of mutation classes of type A_n and type D_{n+1} quivers; consequently there is a cluster subalgebra of type A_n within each cluster algebra of type D_{n+1} . Similarly, we have embeddings $\mathbf{D}_n \leq \mathbf{E}_{n+1}$ for $5 \leq n \leq 7$, and $\mathbf{E}_6 \leq \mathbf{E}_7 \leq \mathbf{E}_8$, cf. Example 4.1.9 and Figure 2.12.

Going beyond the quiver case, we have the embeddings $\mathbf{A}_n \leq \mathbf{B}_{n+1}$, $\mathbf{A}_n \leq \mathbf{C}_{n+1}$, $\mathbf{B}_3 \leq \mathbf{F}_4$, $\mathbf{C}_3 \leq \mathbf{F}_4$, and the corresponding inclusions for cluster algebras of finite type.

We conclude this section by an overview of the remainder of Chapter 5.

Sections 5.3–5.8 are dedicated to showing that any seed pattern that has an exchange matrix whose Cartan counterpart is of one of the types A_n, B_n, \dots, G_2 has finitely many seeds. This is done case by case. The idea is to explicitly construct, for each (indecomposable) type, a particular seed pattern whose exchange matrices have full \mathbb{Z} -rank, and show that this pattern has finitely many seeds. Then an argument based on Corollary 4.3.6 and Remark 4.3.7 will imply the same for *all* cluster algebras of the corresponding type.

In Sections 5.3–5.4 we exhibit seed patterns of types A_n and D_n possessing the requisite properties. In type A_n , we utilize the construction involving the homogeneous coordinate ring of the Grassmannian $\text{Gr}_{2,n+3}$, cf. Section 1.2. Type D_n is treated using a similar construction, admittedly much more technical than in the type A_n case. In Section 5.5, we handle the types B_n and C_n using the technique of folding introduced in Section 4.4.

The exceptional types are treated in a different way. In Section 5.6, we use a computer check to verify that the cluster algebras of type E_8 are of finite type. This implies the same for the types E_6 and E_7 . Types F_4 and G_2 are then handled in Section 5.7, via folding of E_6 and D_4 , respectively.

In Section 5.8, we discuss decomposable types, and complete the first part of the proof of Theorem 5.2.8.

We note that Sections 5.3–5.6 contain many incidental examples of cluster algebras of finite type, in addition to those used in the proof of the classification theorem.

Theorem 5.2.12 is proved in Section 5.9. Also in this section, we complete and summarize the enumeration of the seeds (equivalently, clusters) and the cluster variables for all finite types.

The proof of Theorem 5.2.8 is completed in Section 5.10, by demonstrating that any seed pattern of finite type comes from a Cartan matrix of finite type. This is done by exploiting the fact that every exchange matrix appearing in such a seed pattern is 2-finite, see Corollary 5.1.4.

Theorem 5.2.8 provides a characterization of seed patterns of finite type which, while conceptually satisfying, is not particularly useful in practice. In Section 5.11, we discuss an alternative criterion for recognizing whether a seed pattern is of finite type. This criterion is formulated directly in terms of the given exchange matrix B (as opposed to its mutation class).

5.3. Seed patterns of type A_n

The main result of this section is Theorem 5.3.2, asserting that seed patterns of type A_n are of finite type. The proof uses the fact that seed patterns of type A_n are governed by the combinatorics of triangulated polygons. A case in point is the cluster algebra structure in the Plücker ring $R_{2,n+3}$ (the homogeneous coordinate ring of $\text{Gr}_{2,n+3}$) discussed in Section 1.2. In this example, we verify that a cluster variable indexed by a diagonal in a triangulation T only depends on the diagonal and not on the choice of T , or a sequence of mutation steps relating T to an initial triangulation. After verifying that the cluster structure in $R_{2,n+3}$ is of finite type, we check that one of its exchange matrices has full \mathbb{Z} -rank, and the general case follows.

Let T be a triangulation of a convex $(n+3)$ -gon \mathbf{P}_{n+3} by n noncrossing diagonals labeled $1, \dots, n$. We define the $n \times n$ matrix $B(T) = (b_{ij}(T))$ by

$$(5.3.1) \quad b_{ij}(T) = \begin{cases} 1 & \text{if } i \text{ and } j \text{ label two sides of a triangle in } T, \\ & \text{with } j \text{ following } i \text{ in the clockwise order;} \\ -1 & \text{if } i \text{ and } j \text{ label two sides of a triangle in } T, \\ & \text{with } i \text{ following } j \text{ in the clockwise order;} \\ 0 & \text{if } i \text{ and } j \text{ do not belong to the same triangle in } T. \end{cases}$$

The skew-symmetric matrix $B(T)$ corresponds to the mutable part of the quiver $Q(T)$ described in Definition 2.2.1 and Figure 2.2.

The following easy lemma provides an alternative definition of the notion of a seed pattern of type A_n .

Lemma 5.3.1. *A seed pattern has type A_n if and only if one (equivalently, any) of its exchange matrices can be identified with the exchange matrix $B(T)$ corresponding to a triangulation T of a convex $(n+3)$ -gon, cf. (5.3.1).*




Figure 5.3. A triangulation T_0 of the polygon \mathbf{P}_{n+3} (here $n = 5$). The mutable part of the quiver $Q(T_0)$ (see Definition 2.2.1) is the *equioriented Dynkin quiver* of type A_n .

Theorem 5.3.2. *Seed patterns of type A_n are of finite type.*

Remark 5.3.3. Definition 5.2.9 imposes no restrictions on the bottom part of the extended exchange matrices, nor on the number of frozen variables. In light of Lemma 5.3.1, Theorem 5.3.2 asserts that as long as the mutable part of a quiver comes from a triangulated polygon, the total number of seeds generated by the quiver is finite.

Proof. We start by showing that a particular seed pattern of type A_n , the one associated with the Plücker ring $R_{2,n+3}$, is of finite type.

As in Section 1.2, we label the vertices of the polygon \mathbf{P}_{n+3} clockwise by the numbers $1, \dots, n+3$. For a triangulation T of \mathbf{P}_{n+3} as in Definition 2.2.1, we use the labels $1, \dots, n$ for the diagonals of T (in arbitrary fashion), and use the labels $n+1, \dots, 2n+3$ for the sides of \mathbf{P}_{n+3} , as follows:

- the side with vertices ℓ and $\ell+1$ is labeled $n+\ell$, for $\ell = 1, \dots, n+2$;
- the side with vertices 1 and $n+3$ is labeled $2n+3$.

We define the matrix $\tilde{B}(T) = (b_{ij}(T))$ by the formula (5.3.1), this time with $i \in \{1, \dots, 2n+3\}$ and $j \in \{1, \dots, n\}$. Thus $\tilde{B}(T)$ is the extended exchange matrix for the quiver $Q(T)$ from Definition 2.2.1. By Exercise 2.2.2, flips of triangulations translate into mutations of associated quivers.

We now reformulate the construction described in Section 1.2. Let $V = \mathbb{C}^2$ be a 2-dimensional complex vector space, and let $\langle \cdot, \cdot \rangle$ denote the standard skew-symmetric nondegenerate bilinear form on V . Simply put, $\langle u, v \rangle$ is the determinant of the 2×2 matrix with columns $u, v \in V$. Let \mathcal{K} be the field of rational functions on V^{n+3} written in terms of $2n + 6$ variables, the coordinates of $n + 3$ vectors v_1, \dots, v_{n+3} . The special linear group naturally acts on V , hence on V^{n+3} and on \mathcal{K} . Let $\mathcal{F} = \mathcal{K}^{\mathrm{SL}_2}$ be the subfield of SL_2 -invariant rational functions. All the action will take place in \mathcal{F} .

We associate the Plücker coordinate $P_{ij} = \langle v_i, v_j \rangle \in \mathcal{F}$ with the line segment connecting vertices i and j . For each triangulation T of \mathbf{P}_{n+3} , we let $\tilde{\mathbf{x}}(T)$ be the collection of the $2n + 3$ Plücker coordinates P_{ij} associated with the sides and diagonals of T , as in Section 1.2. By Lemma 5.3.4 below, the elements of $\tilde{\mathbf{x}}(T)$ are algebraically independent. We view the Plücker coordinates associated with the sides of \mathbf{P}_{n+3} as frozen variables. As observed earlier, the Grassmann-Plücker relations (1.2.1) satisfied by the elements P_{ij} can be interpreted as exchange relations encoded by the matrices $\tilde{B}(T)$.

We set $\Sigma(T) = (\tilde{\mathbf{x}}(T), \tilde{B}(T))$. Therefore the seeds obtained from an initial seed $\Sigma(T_\circ)$ form a seed pattern of type A_n , and the initial cluster variables consist of the Plücker coordinates labeling the triangulation T_\circ . Because the Plücker coordinates satisfy the exchange relations, which correspond to flips of triangulations, this gives a canonical identification of cluster variables and clusters with diagonals and triangulations of \mathbf{P}_{n+3} . (A priori, a seed and the cluster variables in it could depend not only on the triangulation they correspond to, but the sequence of mutations we used to arrive at that triangulation from the initial seed.) Therefore the number of distinct seeds in this pattern is finite.

We complete the proof of Theorem 5.3.2 using an argument based on Corollary 4.3.6 and Remark 4.3.7. All we need to do is to check that for some triangulation T_\circ , the matrix $\tilde{B}(T_\circ)$ has full \mathbb{Z} -rank. Taking T_\circ as in Figure 5.3, we obtain the matrix

$$\tilde{B}(T_\circ) = \begin{array}{c} \left[\begin{array}{cccccc} 0 & -1 & 0 & \cdots & 0 & 0 \\ 1 & 0 & -1 & \cdots & 0 & 0 \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 0 & -1 \\ 0 & 0 & 0 & \cdots & 1 & 0 \end{array} \right], \\ \hline \left[\begin{array}{cccccc} -1 & 0 & 0 & \cdots & 0 & 0 \\ 1 & 0 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \end{array} \right] \end{array},$$

where the line is drawn under the n th row. The matrix is easily seen to have full \mathbb{Z} -rank. \square

Lemma 5.3.4. *For any triangulation T of the polygon \mathbf{P}_{n+3} , the elements of $\tilde{\mathbf{x}}(T)$ are algebraically independent. Thus $(\tilde{\mathbf{x}}(T), \tilde{B}(T))$ is a seed in \mathcal{F} .*

Proof. One way to establish this is to observe that $\tilde{\mathbf{x}}(T)$ generates the field \mathcal{F} (since each Plücker coordinate P_{ij} is a rational function in $\tilde{\mathbf{x}}(T)$), and combine this with the fact that the transcendence degree of \mathcal{F} over \mathbb{C} (equivalently, the dimension of the affine cone over $\text{Gr}_{2,n+3}$) is $2n + 3$. \square

Remark 5.3.5. An alternative proof of Theorem 5.3.2 can be based on the description of the fundamental group of the graph whose vertices are the triangulations of the polygon \mathbf{P}_{n+3} , and whose edges correspond to the flips. (We view this graph as a 1-dimensional simplicial complex, the 1-skeleton of the n -dimensional associahedron.) The fundamental group of this graph is generated by 4-cycles and 5-cycles (the boundaries of 2-dimensional faces of the associahedron) pinned down to a basepoint. For each of these cycles, the corresponding sequence of 4 or 5 mutations in a seed pattern of type A_n brings us back to the original seed; this follows from the analysis of the type A_2 case in Section 5.1. Consequently, the seeds in such a pattern can be labeled by the triangulations of \mathbf{P}_{n+3} , implying the claim of finite type.

Corollary 5.3.6. *Cluster variables in a seed pattern of type A_n can be labeled by the diagonals of a convex $(n + 3)$ -gon \mathbf{P}_{n+3} so that*

- *clusters correspond to triangulations of the polygon \mathbf{P}_{n+3} by noncrossing diagonals,*
- *mutations correspond to flips, and*
- *exchange matrices are given by (5.3.1).*

Cluster variables labeled by different diagonals are distinct, so there are altogether $\frac{n(n+3)}{2}$ cluster variables and $\frac{1}{n+2} \binom{2n+2}{n+1}$ seeds (and as many clusters).

Proof. It is well known that the number of triangulations a convex $(n + 3)$ -gon has is equal to the Catalan number $\frac{1}{n+2} \binom{2n+2}{n+1}$, see e.g., [26, Exercise 6.19a]. So the only claim remaining to be proved is that all these cluster variables are distinct in any seed pattern of type A_n . Let x and x' be two cluster variables labeled by distinct diagonals d and d' . If d and d' do not cross each other, then there is a cluster containing x and x' , so x and x' are algebraically independent and therefore distinct. If d and d' do cross, then there is an exchange relation of the form $xx' = M_1 + M_2$ where M_1 and M_2 are monomials in the elements of some extended cluster $\tilde{\mathbf{x}}$ containing x . Now the equality $x = x'$ would imply $x^2 = M_1 + M_2$, contradicting the condition that the elements of $\tilde{\mathbf{x}}$ are algebraically independent. \square

In the rest of this section, we examine several seed patterns (or cluster algebras) of type A_n which naturally arise in various mathematical contexts.

Exercise 5.3.7. A *frieze pattern* [4, 5] is a table of the form

$$n + 2 \text{ rows} \left\{ \begin{array}{cccccccccc} \cdots & 1 & & 1 & & 1 & & 1 & & 1 & & 1 & & 1 & & 1 & \cdots \\ \cdots & & * & & * & & * & & * & & * & & * & & * & \cdots \\ \cdots & * & & * & & * & & * & & * & & * & & * & & * & \cdots \\ \cdots & & * & & * & & * & & * & & * & & * & & * & \cdots \\ \cdots & * & & * & & * & & * & & * & & * & & * & & * & \cdots \\ \cdots & & 1 & & 1 & & 1 & & 1 & & 1 & & 1 & & 1 & \cdots \end{array} \right.$$

with (say) positive integer entries such that every quadruple

$$\begin{array}{ccc} & B & \\ A & & C \\ & D & \end{array}$$

satisfies $AC - BD = 1$. Identify the entries in a frieze pattern with cluster variables in a seed pattern of type A_n . How many distinct entries does a frieze pattern have? What kind(s) of periodicity does it possess?

Example 5.3.8. Let us discuss, somewhat informally, the example of a seed pattern associated with the basic affine space for SL_4 . Choose the initial seed for this pattern as shown in Figure 5.4. (This seed has already appeared in Figure 2.5; it corresponds to a particular choice of a wiring diagram.) The variables P_2 , P_3 , and P_{23} are mutable; the remaining six variables are frozen. The mutable part of the initial quiver is an oriented 3-cycle; as such, it is easily identified as the mutable part of a quiver associated with a particular triangulation of a hexagon. Thus we are dealing here with a seed of type A_3 .




Figure 5.4. A seed of flag minors in $\mathbb{C}[SL_4]^U$.

A matrix in SL_4 has $2^4 - 2 = 14$ nontrivial flag minors: the 6 frozen variables

$$P_1, P_{12}, P_{123}, P_4, P_{34}, P_{234}$$

(recall that they correspond to the unbounded chambers in a wiring diagram) and 8 additional flag minors

$$P_2, P_3, P_{13}, P_{14}, P_{23}, P_{24}, P_{124}, P_{134},$$

all of which can be obtained by the mutation process from our initial seed. Note however that a seed pattern of type A_3 should have 9 cluster variables—so one of them is still missing!

Examining the initial seed shown in Figure 5.4, we see that it can be mutated in three possible ways. The corresponding exchange relations are:

$$\begin{aligned} P_2 P_{13} &= P_{12} P_3 + P_1 P_{23}, \\ P_3 P_{24} &= P_4 P_{23} + P_{34} P_2, \\ P_{23} \Omega &= P_{123} P_{34} P_2 + P_{12} P_{234} P_3. \end{aligned}$$

The first two relations correspond to the two braid moves that can be applied to the initial wiring diagram D_o . The third relation is different in nature: it produces a seed that does not correspond to any wiring diagram, as the new cluster variable Ω is not a flag minor. Ostensibly, Ω is a rational expression (indeed, a Laurent polynomial) in the flag minors. One can check that in fact

$$(5.3.2) \quad \Omega = \frac{P_{123} P_{34} P_2 + P_{12} P_{234} P_3}{P_{23}} = -P_1 P_{234} + P_2 P_{134}$$

—so Ω is not merely a rational function on the basic affine space but a *regular* function. It follows that the corresponding cluster algebra is precisely the invariant ring $\mathbb{C}[\mathrm{SL}_4]^U$ (recall that the latter is generated by the flag minors). It turns out that this phenomenon holds for any special linear group SL_k , resulting in a cluster algebra structure in $\mathbb{C}[\mathrm{SL}_k]^U$.

In the case under consideration, we get 14 distinct extended clusters, in agreement with Corollary 5.3.6. See Figure 5.5.

Example 5.3.9. We conclude this section by presenting a family of seed patterns of type A_n introduced in [28]. They correspond to cluster structures in particular *double Bruhat cells* for the special linear groups $\mathrm{SL}_{n+1}(\mathbb{C})$, more specifically in the cells associated with *pairs of Coxeter elements* in the associated symmetric group \mathcal{S}_{n+1} . This construction can be extended to arbitrary simply connected semisimple complex Lie groups, see [28].

Let $L_n \subset \mathrm{SL}_{n+1}(\mathbb{C})$ be the subvariety of tridiagonal matrices

$$(5.3.3) \quad z = \begin{bmatrix} v_1 & q_1 & 0 & \cdots & 0 \\ 1 & v_2 & q_2 & \ddots & \vdots \\ 0 & 1 & v_3 & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & q_n \\ 0 & \cdots & 0 & 1 & v_{n+1} \end{bmatrix}$$

of determinant 1. For $i, j \in \{1, \dots, n+3\}$ satisfying $i+2 \leq j$, consider the *solid principal minor* (cf. Exercise 1.4.2)

$$U_{ij} = \Delta_{[i,j-2],[i,j-2]} \in \mathbb{C}[L_n],$$

the determinant of the submatrix with rows and columns $i, i+1, \dots, j-2$. For example, $U_{i,i+2} = v_i$ and $U_{1,n+3} = \det(z) = 1$. By convention, $U_{i,i+1} = 1$.




Figure 5.5. Clusters in $\mathbb{C}[\mathrm{SL}_4]^U$. The 14 clusters for this seed pattern are shown as the vertices of a graph; the edges of the graph correspond to seed mutations. Note that there is one additional vertex at infinity, so the graph should be viewed as drawn on a sphere rather than a plane. The regions are labeled by cluster variables. Each cluster consists of three elements labeling the regions adjacent to the corresponding vertex. The 6 frozen variables are not shown. This graph is isomorphic to the 1-skeleton of the three-dimensional associahedron, shown in Figure 1.4.

Exercise 5.3.10. Prove that these functions satisfy the relations

$$(5.3.4) \quad U_{ik} U_{j\ell} = q_{j-1} q_j \cdots q_{k-2} U_{ij} U_{k\ell} + U_{il} U_{jk},$$

for $1 \leq i < j < k < \ell \leq n + 3$. Then show that these relations are the exchange relations in a particular seed pattern of type A_n . More precisely, show that there is a seed pattern of type A_n , with the frozen variables q_1, \dots, q_n , in which the cluster variables U_{ij} associated to the diagonals $[i, j]$ of the convex polygon \mathbf{P}_{n+3} satisfy the exchange relations (5.3.4).

What is required to show that the relations (5.3.4) are the exchange relations in a seed pattern of type A_n ? Let us extend each $n \times n$ exchange matrix $B(T)$ associated to a triangulation T of \mathbf{P}_{n+3} (see (5.3.1)) to the $2n \times n$ matrix $\tilde{B}(T)$ whose columns encode the relations (5.3.4) corresponding to the n possible flips from T . One then needs to verify that whenever triangulations T and T' are related by a flip, the associated matrices $\tilde{B}(T)$ and $\tilde{B}(T')$ are related by the corresponding mutation.

One of the clusters in this seed pattern consists of the n leading principal minors $U_{13}, \dots, U_{1, n+2}$. The exchange relations from this cluster are the relations (5.3.4) with $(i, j, k, \ell) = (1, k-1, k, k+1)$. They can be rewritten as follows:

$$(5.3.5) \quad U_{1, k+1} = v_{k-1} U_{1, k} - q_{k-2} U_{1, k-1} \quad (k = 3, \dots, n+2).$$

These relations play important roles in the classical theory of orthogonal polynomials in one variable [27], in the study of a generalized Toda lattice [21], and in other mathematical contexts.

5.4. Seed patterns of type D_n

In this section we show that seed patterns of type D_n are of finite type. The proof of this theorem, while substantially more technical than the proof of Theorem 5.3.2, follows the same general plan. It relies on two main ingredients. The first ingredient is a combinatorial construction (“tagged arcs” in a punctured disk) that enables us to explicitly describe the combinatorics of mutations in type D_n and introduce the relevant nomenclature. The second ingredient is an algebraic construction of a particular seed pattern of type D_n . In this pattern, each cluster variable associated with a tagged arc has an intrinsic definition independent of the mutation path; this will imply that the number of seeds is finite. The “full \mathbb{Z} -rank” argument will then allow us to generalize the finiteness statement to arbitrary coefficients.

While type D_n Dynkin diagrams are usually defined for $n \geq 4$, in this section we will allow for the possibility of $n = 3$ (in which case one recovers type A_3).

Exercise 5.4.1. Show that a seed pattern is of type D_n if and only if one of its exchange matrices corresponds to a quiver which is an oriented n -cycle.




Figure 5.6. Dynkin diagram of type D_n and an oriented n -cycle.

Theorem 5.4.2. *Seed patterns of type D_n are of finite type.*

As in Remark 5.3.3, the key point of Theorem 5.4.2 is that a pattern of type D_n has finitely many seeds regardless of the number of frozen variables and of the entries in the bottom parts of extended exchange matrices.

Definition 5.4.3. Let \mathbf{P}_n^\bullet be a convex n -gon ($n \geq 3$) with a distinguished point p (a *puncture*) in its interior. We label the vertices of \mathbf{P}_n^\bullet clockwise from 1 to n . These vertices and the puncture p are the *marked points* of \mathbf{P}_n^\bullet .

An *arc* in \mathbf{P}_n^\bullet is a non-selfintersecting curve γ in \mathbf{P}_n^\bullet such that

- the endpoints of γ are two different marked points;
- except for its endpoints, γ is disjoint from the boundary of \mathbf{P}_n^\bullet and from the puncture p ;
- γ does not cut out an unpunctured digon.

We consider each arc up to isotopy within the class of such curves.

The arcs incident to the puncture are called *radii*.

Definition 5.4.4. A *tagged arc* in \mathbf{P}_n^\bullet is either an ordinary non-radius arc, or a radius that has been labeled (“tagged”) in one of two possible ways, *plain* or *notched*. Two tagged arcs are called *compatible* with one another if their untagged versions do not cross each other, with the following modification: the plain and notched versions of the same radius are compatible, but the the plain and notched versions of two different radii are not.

A *tagged triangulation* is a maximal (by inclusion) collection of pairwise compatible tagged arcs. See Figure 5.7.

Lemma 5.4.5. *Any tagged triangulation T of \mathbf{P}_n^\bullet consists of n tagged arcs. Any tagged arc in a tagged triangulation T can be replaced in a unique way by a tagged arc not belonging to T , to form a new tagged triangulation T' .*

Proof. It is easy to see that tagged triangulations come in three flavors:

- (1) a triangulation in the usual (topological) sense, with every radius plain;
- (2) a triangulation in the usual sense, with every radius notched;
- (3) the plain and notched versions of the same radius inside a punctured digon. Outside of the digon, it is an ordinary triangulation.

In each of these cases, the total number of tagged arcs is n . □

In the situation described in Lemma 5.4.5, we say that the tagged triangulations T and T' are related by a *flip*.

Exercise 5.4.6. Verify that any two tagged triangulations of \mathbf{P}_n^\bullet are connected via a sequence of flips.

We next define an extended exchange matrix $\tilde{B}(T)$ associated with a tagged triangulation T of the punctured polygon \mathbf{P}_n^\bullet . The construction is similar to the one in type A_n . In the cases (1) and (2) above, the rule (5.3.1) is used; in the case (3), some adjustments are needed around the radii.

Figure 5.8 illustrates the recipe used to define the matrix $\tilde{B}(T)$, or equivalently the corresponding quiver. The vertices of the quiver corresponding to boundary segments (i.e., the sides of the polygon) are frozen; the ones corresponding to tagged arcs are mutable. Details are left to the reader.




Figure 5.7. Tagged arcs in a once-punctured triangle \mathbf{P}_3^* . Solid lines indicate which arcs are compatible. The vertices of the dashed graph correspond to tagged triangulations; its edges correspond to flips. Note that there is one additional vertex at infinity, so the graph should be viewed as drawn on a sphere rather than a plane.

Exercise 5.4.7. Verify that if two tagged triangulations T and T' are related by a flip, then the extended skew-symmetric matrices $\tilde{B}(T)$ and $\tilde{B}(T')$ are related by the corresponding matrix mutation.

Unfortunately, the matrices $\tilde{B}(T)$ have rank $< n$, so exhibiting a seed pattern with these matrices and finitely many seeds would not yield a proof of Theorem 5.4.2, cf. Remark 4.3.7. We will overcome this obstacle by introducing additional frozen variables besides those labeled by boundary segments.

We now prepare the algebraic ingredients for the proof of Theorem 5.4.2. Similarly to the type A_n case, the idea is to interpret tagged arcs in a once-punctured polygon as a family of rational functions. These rational functions satisfy the type D_n exchange relations, which in turn are associated to




Figure 5.8. Quivers associated with tagged triangulations of a once-punctured convex polygon \mathbf{P}_n^* .

flips of tagged arcs. The astute reader will notice that the algebraic construction we associate to the type D_n case contains the algebraic construction associated to the type A_{n-1} case (Plücker coordinates of an $(n+2)$ -tuple of vectors in \mathbb{C}^2), reflective of the fact that the type A_{n-1} Dynkin diagram is contained inside the type D_n Dynkin diagram.

As in Section 5.3, we set $V = \mathbb{C}^2$, and let $\langle u, v \rangle$ be the determinant of the 2×2 matrix with columns $u, v \in V$. Let \mathcal{K} be the field of rational functions on $V^n \times V \times V \times \mathbb{C}^2 \cong \mathbb{C}^{2n+6}$, written in terms of $2n+6$ variables: the coordinates of $n+2$ vectors

$$v_1, \dots, v_n, a, \bar{a} \in V,$$

plus two additional variables λ and $\bar{\lambda}$.

Let $A \in \text{End}(V)$ denote the linear operator defined by

$$(5.4.1) \quad Av = \frac{\bar{\lambda} \langle v, a \rangle \bar{a} - \lambda \langle v, \bar{a} \rangle a}{\langle \bar{a}, a \rangle},$$

so that a (resp. \bar{a}) is an eigenvector for A with eigenvalue λ (resp. $\bar{\lambda}$). Let

$$a^{\bowtie} = \frac{\bar{\lambda} - \lambda}{\langle \bar{a}, a \rangle} \bar{a};$$

thus a^{\bowtie} is also an eigenvector for A , with eigenvalue $\bar{\lambda}$. We choose this normalization because of the following property, which is immediate from the definitions.

Lemma 5.4.8. *For $v \in V$, we have $\langle v, Av \rangle = \langle v, a \rangle \langle v, a^{\bowtie} \rangle$.*

We next describe a seed pattern inside \mathcal{K} , and in fact inside its subfield \mathcal{F} of $\mathrm{SL}_2(\mathbb{C}) \times \mathbb{C}^*$ -invariant rational functions. (The group SL_2 acts in the standard way on each of the vectors $v_1, \dots, v_n, a, \bar{a}$, and fixes λ and $\bar{\lambda}$. The group \mathbb{C}^* acts by rescaling the vector \bar{a} . The subfield \mathcal{F} can be thought of as a field of rational functions on a $(2n+2)$ -dimensional variety, and indeed, extended clusters in our seed pattern will have size $2n + 2$.) Informally speaking, we are going to think of the vectors v_1, \dots, v_n as located at the corresponding vertices of \mathbf{P}_n^\bullet , and we shall associate both eigenvectors a and a^\boxtimes with the puncture p . We make a *cut* running from the puncture p to the boundary segment (the side of the polygon) that connects the vertices 1 and n . We will think of crossing the cut as picking up an application of A .

Definition 5.4.9. We associate an element $P_\gamma \in \mathcal{K}$ to each tagged arc or boundary segment γ in \mathbf{P}_n^\bullet , as follows (see Figure 5.9):

$$P_\gamma = \begin{cases} \langle v_i, v_j \rangle & \text{if } \gamma \text{ doesn't cross the cut, and has endpoints } i \text{ and } j > i; \\ \langle v_j, Av_i \rangle & \text{if } \gamma \text{ crosses the cut, and has endpoints } i \text{ and } j > i; \\ \langle v_i, a \rangle & \text{if } \gamma \text{ is a plain radius with endpoints } p \text{ and } i; \\ \langle v_i, a^\boxtimes \rangle & \text{if } \gamma \text{ is a notched radius with endpoints } p \text{ and } i. \end{cases}$$




Figure 5.9. Elements of the field \mathcal{K} associated with tagged arcs in \mathbf{P}_n^\bullet . The boundary segment crossed by the cut corresponds to $\langle v_n, Av_1 \rangle$.

We can now verify that the P_γ satisfy a family of relations that topologists know as “skein relations” for tagged arcs.

For example, Figure 5.10 illustrates the equation

$$(5.4.2) \quad \langle v_i, v_l \rangle \langle v_k, Av_i \rangle = \langle v_k, v_l \rangle \langle v_i, a \rangle \langle v_i, a^\bowtie \rangle + \langle v_i, v_k \rangle \langle v_l, Av_i \rangle,$$

which follows from the Grassmann-Plücker relation (1.2.1) combined with Lemma 5.4.8.




Figure 5.10. Pictorial representation of the relation (5.4.2) (or (5.4.7)).

More generally, we have the following relations. The reader may want to draw a figure corresponding to each relation to get some intuition about the form of the relations.

Proposition 5.4.10. *The elements P_γ described in Definition 5.4.9 satisfy the following relations:*

$$(5.4.3) \quad \langle v_i, v_k \rangle \langle v_j, v_l \rangle = \langle v_i, v_j \rangle \langle v_k, v_l \rangle + \langle v_i, v_l \rangle \langle v_j, v_k \rangle,$$

$$(5.4.4) \quad \langle v_j, v_l \rangle \langle v_k, Av_i \rangle = \langle v_j, v_k \rangle \langle v_l, Av_i \rangle + \langle v_j, Av_i \rangle \langle v_k, v_l \rangle,$$

$$(5.4.5) \quad \langle v_k, Av_i \rangle \langle v_l, Av_j \rangle = \lambda \bar{\lambda} \langle v_i, v_j \rangle \langle v_k, v_l \rangle + \langle v_k, Av_j \rangle \langle v_l, Av_i \rangle,$$

$$(5.4.6) \quad \langle v_i, v_k \rangle \langle v_l, Av_j \rangle = \langle v_l, Av_i \rangle \langle v_j, v_k \rangle + \langle v_l, Av_k \rangle \langle v_i, v_j \rangle,$$

$$(5.4.7) \quad \langle v_i, v_l \rangle \langle v_k, Av_i \rangle = \langle v_i, v_k \rangle \langle v_l, Av_i \rangle + \langle v_k, v_l \rangle \langle v_i, a \rangle \langle v_i, a^\bowtie \rangle,$$

$$(5.4.8) \quad \langle v_j, Av_i \rangle \langle v_l, Av_j \rangle = \lambda \bar{\lambda} \langle v_i, v_j \rangle \langle v_j, v_l \rangle + \langle v_j, a^\bowtie \rangle \langle v_j, a \rangle \langle v_l, Av_i \rangle,$$

$$(5.4.9) \quad \langle v_i, v_l \rangle \langle v_l, Av_j \rangle = \langle v_l, Av_i \rangle \langle v_j, v_l \rangle + \langle v_l, a^\bowtie \rangle \langle v_l, a \rangle \langle v_i, v_j \rangle$$

$$(5.4.10) \quad \langle v_i, v_k \rangle \langle v_j, a \rangle = \langle v_i, v_j \rangle \langle v_k, a \rangle + \langle v_i, a \rangle \langle v_j, v_k \rangle,$$

$$(5.4.11) \quad \langle v_j, Av_i \rangle \langle v_k, a \rangle = \bar{\lambda} \langle v_j, v_k \rangle \langle v_i, a \rangle + \langle v_j, a \rangle \langle v_k, Av_i \rangle,$$

$$(5.4.12) \quad \langle v_k, Av_j \rangle \langle v_i, a \rangle = \langle v_k, Av_i \rangle \langle v_j, a \rangle + \lambda \langle v_k, a \rangle \langle v_i, v_j \rangle,$$

$$(5.4.13) \quad \langle v_i, a^\bowtie \rangle \langle v_j, a \rangle = \bar{\lambda} \langle v_i, v_j \rangle + \langle v_j, Av_i \rangle.$$

where $1 \leq i < j < k < l \leq n$. In addition, they satisfy the relations obtained from (5.4.10)–(5.4.13) by interchanging λ with $\bar{\lambda}$ and a with a^\bowtie throughout.

Proof. Each of these relations follows from a suitable instance of the Grassmann-Plücker relation (1.2.1), using Lemma 5.4.8 and the identities

$$\begin{aligned}\langle Av, Av' \rangle &= \det(A) \langle v, v' \rangle = \lambda \bar{\lambda} \langle v, v' \rangle, \\ \langle Av, a \rangle &= \bar{\lambda} \langle v, a \rangle, \\ \langle Av, a^{\boxtimes} \rangle &= \lambda \langle v, a^{\boxtimes} \rangle.\end{aligned}$$

For example, (5.4.12) can be obtained from the identity

$$\langle v_k, Av_j \rangle \langle Av_i, a \rangle = \langle v_k, Av_i \rangle \langle Av_j, a \rangle + \langle v_k, a \rangle \langle Av_i, Av_j \rangle,$$

while (5.4.13) follows (using Lemma 5.4.8) from

$$\langle v_i, Av_i \rangle \langle v_j, a \rangle = \langle v_i, v_j \rangle \langle Av_i, a \rangle + \langle v_i, a \rangle \langle v_j, Av_i \rangle. \quad \square$$

For a tagged triangulation T of the punctured polygon \mathbf{P}_n^\bullet , let $\tilde{\mathbf{x}}(T)$ be the $(2n+2)$ -tuple consisting of the elements P_γ labeled by the tagged arcs and boundary segments in T , together with λ and $\bar{\lambda}$. We view the elements $P_\gamma \in \tilde{\mathbf{x}}(T)$ labeled by tagged arcs as cluster variables, and those labeled by the boundary segments as frozen variables; λ and $\bar{\lambda}$ are frozen variables as well.

Let γ be a tagged arc in a tagged triangulation T , and let γ' be the tagged arc that replaces γ when the latter is flipped. One can check that in every such instance, exactly one of the relations in Proposition 5.4.10 has the product $P_\gamma P_{\gamma'}$ on the left-hand side; the corresponding right-hand side is always a sum of two monomials in the elements of $\tilde{\mathbf{x}}(T)$. We let $\tilde{B}^\bullet(T)$ denote the matrix encoding these relations for all tagged arcs in T . The matrix $\tilde{B}^\bullet(T)$ can be seen to be an extension of the matrix $\tilde{B}(T)$ by two extra rows corresponding to λ and $\bar{\lambda}$. (Put differently, setting $\lambda = \bar{\lambda} = 1$ produces relations encoded by $\tilde{B}(T)$.)

Example 5.4.11. Figure 5.11 shows a triangulation T_\circ of a punctured pentagon, and the associated matrix $\tilde{B}^\bullet(T_\circ)$. Columns 1 and 5 of $\tilde{B}^\bullet(T_\circ)$ encode the exchange relations

$$\begin{aligned}\langle v_5, Av_2 \rangle \langle v_1, a \rangle &= \langle v_5, Av_1 \rangle \langle v_2, a \rangle + \lambda \langle v_5, a \rangle \langle v_1, v_2 \rangle = x_{10}x_2 + \lambda x_5x_6 \\ \langle v_4, Av_1 \rangle \langle v_5, a \rangle &= \bar{\lambda} \langle v_4, v_5 \rangle \langle v_1, a \rangle + \langle v_4, a \rangle \langle v_5, Av_1 \rangle = \bar{\lambda} x_9x_1 + x_4x_{10}.\end{aligned}$$

The matrix $\tilde{B}^\bullet(T_\circ)$ has full \mathbb{Z} -rank. However, the submatrix of $\tilde{B}(T_\circ)$ consisting of the first ten rows does not have full rank (each row sum is 0).

This example can be straightforwardly generalized to $n \neq 5$.

Proposition 5.4.12. *For any tagged triangulation T of \mathbf{P}_n^\bullet , the elements of $\tilde{\mathbf{x}}(T)$ are algebraically independent, so $(\tilde{\mathbf{x}}(T), \tilde{B}^\bullet(T))$ is a seed in the field generated by $\tilde{\mathbf{x}}(T)$. The seeds associated to tagged triangulations related by a flip are related to each other by the corresponding mutation.*




Figure 5.11. The (tagged) triangulation T_\circ of a punctured pentagon ($n = 5$), and the corresponding matrix $\tilde{B}^\bullet(T_\circ)$. The columns of $\tilde{B}^\bullet(T_\circ)$ correspond to x_1, \dots, x_5 . The rows correspond to $x_1, \dots, x_{10}, \lambda, \bar{\lambda}$, in this order.

Proof. It is a straightforward but tedious exercise to verify that the matrix $\tilde{B}^\bullet(T)$ undergoes a mutation when a tagged arc in T is flipped to produce a new tagged triangulation T' . We already know that the elements of $\tilde{\mathbf{x}}(T)$ and $\tilde{\mathbf{x}}(T')$ satisfy the corresponding exchange relation. It remains to show algebraic independence. Since each exchange is a birational transformation, it suffices to prove algebraic independence for one particular choice of T .

Consider the triangulation T_\circ made up of n plain radii, cf. Figure 5.11 and Example 5.4.11. Then $\lambda, \bar{\lambda}$, and $\langle v_n, Av_1 \rangle$ are the only elements of $\tilde{\mathbf{x}}(T_\circ)$ which involve $\lambda, \bar{\lambda}$, or \bar{a} . The remaining $2n - 1$ elements are 2×2 minors of the matrix with columns v_1, v_2, \dots, v_n, a ; they form an extended cluster in the corresponding seed pattern of type A_{n-2} . Hence they are algebraically independent, and the claim follows. \square

Proof of Theorem 5.4.2. By Proposition 5.4.12, the seeds $(\tilde{\mathbf{x}}(T), \tilde{B}^\bullet(T))$ form a seed pattern. This pattern has finitely many seeds because a once-punctured polygon has finitely many tagged triangulations. As verified in Example 5.4.11, the extended exchange matrix $\tilde{B}^\bullet(T_\circ)$ has full \mathbb{Z} -rank. The theorem follows by Corollary 4.3.6 and Remark 4.3.7. \square

Corollary 5.4.13. *Cluster variables in a seed pattern of type D_n can be labeled by the tagged arcs in a once-punctured convex n -gon \mathbf{P}_n^\bullet so that clusters correspond to tagged triangulations, and mutations correspond to flips. Cluster variables labeled by different tagged arcs are distinct. There are altogether n^2 cluster variables and $\frac{3n-2}{n} \binom{2n-2}{n-1}$ seeds (or clusters).*

Most of the work in the proof of the corollary concerns the enumeration of seeds. Let a_n (resp., d_n) denote the number of seeds in a pattern of

type A_n (resp., D_n), including $a_0 = 1$, $d_2 = 4$, and $d_3 = 14$ by convention. We know from Corollary 5.3.6 that $a_n = \frac{1}{n+2} \binom{2n+2}{n+1}$.

Lemma 5.4.14. *The number of tagged triangulations T containing a given radius γ is equal to a_{n-1} .*

Proof. Assume without loss of generality that γ is a plain radius connecting the puncture p to a boundary point i . If T contains the notched counterpart of γ , then the remaining $n-1$ arcs of T form a triangulation of the $(n+1)$ -gon obtained from \mathbf{P}_n^\bullet by cutting it open along γ . The number of such triangulations is a_{n-2} . If T does not contain the notched counterpart of γ , then it is an ordinary triangulation of \mathbf{P}_n^\bullet cut open along γ , in which we are not allowed to use the arc connecting the two boundary points obtained from i . The number of such triangulations is $a_{n-1} - a_{n-2}$. \square

Lemma 5.4.15. *The numbers d_n satisfy the recurrence*

$$d_n = \sum_{k=0}^{n-3} a_k d_{n-1-k} + 2a_{n-1}.$$

Proof. Consider two ways of counting the triples (T, γ, i) in which T is a tagged triangulation of \mathbf{P}_n^\bullet , γ is a tagged arc in T , and i is an endpoint of γ . Selecting T first, then γ and then i , we see that the number of such triples is equal to $d_n \cdot n \cdot 2$. Selecting i first, then γ and T together, treating the cases $i \neq p$ and $i = p$ separately, and using Lemma 5.4.14, we obtain:

$$2nd_n = n \left(2 \sum_{k=0}^{n-3} a_k d_{n-1-k} + 2a_{n-1} \right) + 2na_{n-1},$$

as desired. (The factor of 2 before the sum accounts for the two ways of cutting up the polygon: leaving the puncture to the left or to the right of γ , as we move away from i .) \square

Proof of Corollary 5.4.13. The statements in the first sentence of the corollary have already been established. The claim of distinctness can be verified in the special case $n = 4$ by direct calculation; the general case then follows by restriction.

The cluster variables are labeled by $\frac{n^2-3n+2}{2}$ ordinary arcs not crossing the cut, $\frac{n^2-n-2}{2}$ ordinary arcs crossing the cut, and n radii of each of the two flavors, bringing the total to n^2 .

The formula $d_n = \frac{3n-2}{n} \binom{2n-2}{n-1}$ can now be proved by induction using the recurrence in Lemma 5.4.15. We leave this step to the reader. \square

We conclude this section by examining a couple of examples of cluster algebras of type D_n that occur in the settings discussed in Chapter 1.

Example 5.4.16. The ring of polynomials in 9 variables z_{ij} ($i, j \in \{1, 2, 3\}$), viewed as matrix entries of a 3×3 matrix

$$z = \begin{bmatrix} z_{11} & z_{12} & z_{13} \\ z_{21} & z_{22} & z_{23} \\ z_{31} & z_{32} & z_{33} \end{bmatrix} \in \text{Mat}_{3,3}(\mathbb{C}) \cong \mathbb{C}^9,$$

carries a natural cluster algebra structure of type D_4 , see, e.g., [12]. The seed pattern giving rise to this cluster structure contains, among others, the seeds associated with double wiring diagrams with 3 wires of each kind.

Let us use the initial seed associated with the diagram D in Figure 2.6. Thus the initial cluster is

$$(5.4.14) \quad \mathbf{x} = (x_1, x_2, x_3, x_4) = (\Delta_{1,2}, \Delta_{3,2}, \Delta_{13,12}, \Delta_{13,23}),$$

the 5 coefficient variables are

$$(5.4.15) \quad (x_5, \dots, x_9) = (\Delta_{123,123}, \Delta_{12,23}, \Delta_{1,3}, \Delta_{3,1}, \Delta_{23,12}),$$

and the exchange relations (encoded by the quiver shown in Figure 2.6) are:

$$\begin{aligned} \Delta_{1,2} \Delta_{3,3} &= \Delta_{1,3} \Delta_{3,2} + \Delta_{13,23}, \\ \Delta_{3,2} \Delta_{1,1} &= \Delta_{13,12} + \Delta_{3,1} \Delta_{1,2}, \\ \Delta_{13,12} \Delta_{23,23} &= \Delta_{23,12} \Delta_{13,23} + \Delta_{123,123} \Delta_{3,2}, \\ \Delta_{13,23} \Delta_{12,12} &= \Delta_{123,123} \Delta_{1,2} + \Delta_{12,23} \Delta_{13,12}. \end{aligned}$$

The mutable part of $Q(D)$ is an oriented 4-cycle, so we are indeed dealing with a pattern of type D_4 . By Corollary 5.4.13, it has 50 seeds, which include 34 seeds associated with double wiring diagrams, see Figure 1.12. There are 16 cluster variables, labeled by the 16 tagged arcs in a once-punctured quadrilateral. They include the 14 minors of z not listed in (5.4.15). As suggested by Exercise 1.4.4, the remaining two cluster variables are the polynomials $K(z)$ and $L(z)$ given by (1.4.2)–(1.4.3).

Example 5.4.17. The coordinate ring $\mathbb{C}[\text{SL}_5]^U$ of the basic affine space for SL_5 has a natural cluster algebra structure of type D_6 , to be discussed in detail in Example 6.5.6. This cluster algebra has 36 cluster variables and 8 frozen variables. This set of 44 generators includes $2^5 - 2 = 30$ flag minors plus 14 non-minor elements. The total number of clusters is 672.

Remark 5.4.18. For any irreducible representation of a semisimple algebraic group G , Lusztig [23] introduced the concept of a canonical (resp., dual canonical) basis; these correspond to the lower and upper global bases of Kashiwara [19]. While we will not define the dual canonical bases here, we note that they are strongly connected to the theory of cluster algebras.




Figure 5.12. The triangulation of a once-punctured square representing the initial cluster (5.4.14). Note that the frozen variables are not shown; the frozen variables of the initial cluster at the right do not correspond to the edges of the square.

In particular, it was conjectured in [10, p. 498] that every cluster monomial in $\mathbb{C}[\mathrm{SL}_k]^U$ (i.e., a monomial in the elements of some extended cluster) belongs to the dual canonical basis; this conjecture was proved in [18]. For $k \leq 5$, cluster monomials make up the entire dual canonical basis in $\mathbb{C}[\mathrm{SL}_k]^U$.

Example 5.4.19. The homogeneous coordinate ring of a Schubert divisor in the Grassmannian $\mathrm{Gr}_{2,n+2}$ has a structure of a cluster algebra of type D_n ; see Example 6.3.5 for details.

5.5. Seed patterns of types B_n and C_n

By Definition 5.2.9, a seed pattern of rank $n \geq 2$ (or the associated cluster algebra) is of type B_n if one of its exchange matrices is

$$(5.5.1) \quad B = \begin{bmatrix} 0 & -2 & 0 & 0 & \cdots & 0 & 0 \\ 1 & 0 & -1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & -1 & \cdots & 0 & 0 \\ 0 & 0 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & -1 \\ 0 & 0 & 0 & 0 & \cdots & 1 & 0 \end{bmatrix}$$

(up to simultaneous permutation of rows and columns).

Similarly, a seed pattern (or cluster algebra) of rank $n \geq 3$ is of type C_n if one of its exchange matrices has the form

$$(5.5.2) \quad B = \begin{bmatrix} 0 & -1 & 0 & 0 & \cdots & 0 & 0 \\ 2 & 0 & -1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & -1 & \cdots & 0 & 0 \\ 0 & 0 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & -1 \\ 0 & 0 & 0 & 0 & \cdots & 1 & 0 \end{bmatrix}.$$

(We continue to use the conventions of [17], cf. Definition 5.2.4.)

Theorem 5.5.1. *Seed patterns of type B_n are of finite type.*

Theorem 5.5.2. *Seed patterns of type C_n are of finite type.*

Note that type B_2 is already covered by Theorem 5.1.1, with $bc = 2$.

We will prove Theorems 5.5.1–5.5.2 using the technique of folding introduced in Section 4.4. Specifically, we will obtain seed patterns of type C_n by folding seed patterns of type A_{2n-1} , and we will get type B_n from type D_{n+1} .

The following statement follows easily from the definitions.

Lemma 5.5.3. *Let Q be a quiver globally foldable with respect to an action of a group G . Let \overline{Q} be a quiver constructed from Q by adding some new frozen vertices together with some arrows connecting them to the mutable vertices in Q . We extend the action of G from Q to \overline{Q} by making G fix every newly added vertex. Assume that the new arrows are compatible with the action of G on \overline{Q} , i.e. this action satisfies condition (2) of Definition 4.4.1. Then the quiver \overline{Q} is globally foldable with respect to G .*

Corollary 5.5.4. *Let Q be a quiver without frozen vertices. Suppose that Q is globally foldable with respect to an action of a group G . If every seed pattern with the initial exchange matrix $B(Q)$ is of finite type (regardless of the choice of an initial extended exchange matrix \tilde{B} containing $B(Q)$), then every seed pattern with the initial exchange matrix $B(Q)^G$ is of finite type.*

Proof. Any extended exchange matrix \tilde{B} that extends $B(Q)^G$ can be obtained from an extended exchange matrix that extends $B(Q)$ via the folding procedure described in Lemma 5.5.3. The claim then follows from Corollary 4.4.11. \square

Proof of Theorem 5.5.2. Our proof strategy is as follows. We will construct a type A_{2n-1} quiver Q_0 with a group G acting on its vertices, so that Q_0 is globally foldable with respect to G , and $B(Q_0)^G$ is the $n \times n$ exchange matrix of type C_n from Equation (5.5.2). Theorem 5.5.2 will then follow from Theorem 5.3.2 and Corollary 5.5.4.

The combinatorial model for a seed pattern of type A_{2n-1} presented in Section 5.3 uses triangulations of a convex $(2n+2)$ -gon \mathbf{P}_{2n+2} , with vertices numbered $1, 2, \dots, 2n+2$ in clockwise order. Consider the centrally symmetric triangulation T_0 (see Figure 5.13) consisting of the following diagonals:

- a “diameter” d_1 connecting vertices 1 and $n+2$;
- diagonals d_2, d_3, \dots, d_n connecting $n+2$ with $2, 3, \dots, n$;
- diagonals $d_{2'}, d_{3'}, \dots, d_{n'}$ connecting 1 with $n+3, n+4, \dots, 2n+1$.

Let Q_0 denote the mutable part of the quiver associated to the triangulation T_0 (see Section 2.2); Q_0 is an orientation of the type A_{2n-1} Dynkin diagram with vertices labeled $n', \dots, 3', 2', 1, 2, 3, \dots, n$ in order, and arrows directed towards the central vertex 1, see Figure 5.13. The group $G = \mathbb{Z}/2\mathbb{Z}$ acts on the vertices of Q_0 by exchanging i' and i for $2 \leq i \leq n$, and fixing the vertex 1. It is easy to see that Q_0 is G -admissible, and moreover $B(Q_0)^G$ is the exchange matrix of type C_n from Equation (5.5.2).




Figure 5.13. A centrally symmetric triangulation T_0 and its quiver $Q_0 = Q(T_0)$. The folded matrix $B(Q(T_0))^G$ has type C_n ; here $n = 3$.

It remains to show that Q_0 is globally foldable. Besides acting on the set $\{n', \dots, 3', 2', 1, 2, 3, \dots, n\}$, the group $G = \mathbb{Z}/2\mathbb{Z}$ naturally acts—by central symmetry—on the set of diagonals of the polygon \mathbf{P}_{2n+2} . This action has two kinds of orbits: (1) the “diameters” of \mathbf{P}_{2n+2} fixed by the G -action, and (2) pairs of centrally symmetric diagonals.

A transformation $\mu_{J_k} \circ \dots \circ \mu_{J_1}$ associated with a sequence J_1, \dots, J_k of G -orbits in $\{n', \dots, 3', 2', 1, 2, 3, \dots, n\}$ corresponds to a sequence of flips associated with diameters or pairs of centrally symmetric diagonals. Such a sequence transforms T_0 into another centrally symmetric triangulation T in which the diameter retains the label 1 and each pair of centrally symmetric diagonals have labels i and i' , for some i . It is easy to see that the quiver associated to T is G -admissible. Thus Q_0 is globally foldable. \square

Proof of Theorem 5.5.1. We will follow the strategy used in the proof of Theorem 5.5.2. Namely, we will construct a type D_{n+1} quiver Q_0 with a group G acting on its vertices, so that Q_0 is globally foldable with respect to G , and $B(Q_0)^G$ is the exchange matrix of type B_n from (5.5.1). Theorem 5.5.1 will then follow from Theorem 5.4.2 and Corollary 5.5.4.

The combinatorial model for a seed pattern of type D_{n+1} ($n \geq 3$) uses tagged triangulations of a convex $(n+1)$ -gon \mathbf{P}_{n+1}^\bullet with a puncture p in

its interior. The vertices of \mathbf{P}_{n+1}^\bullet are numbered $1, 2, \dots, n+1$ in clockwise order. Consider the tagged triangulation T_0 formed by:

- two radii d_1 and d_2 (tagged plain and notched) connecting the vertex 1 with the puncture p ;
- plain arcs d_3, d_4, \dots, d_{n+1} connecting the vertex 1 with vertices $2, 3, \dots, n$, respectively, as shown in Figure 5.14.




Figure 5.14. A tagged triangulation T_0 with quiver Q_0 such that $B(Q_0)^G$ is the type B_n exchange matrix; here $n = 7$.

The corresponding quiver $Q_0 = Q(T_0)$ is an orientation of the type D_{n+1} Dynkin diagram. The group $G = \mathbb{Z}/2\mathbb{Z}$ acts on the vertices of Q_0 by exchanging the vertices 1 and 2, and fixing all the other vertices. It is easy to see that Q_0 is G -admissible, and moreover $B(Q_0)^G$ is the exchange matrix of type B_n from Equation (5.5.1).

It remains to show that Q_0 is globally foldable. In addition to the action on the set $\{1, 2, \dots, n+1\}$ described above, the group $G = \mathbb{Z}/2\mathbb{Z}$ acts on the set of tagged arcs in the punctured polygon \mathbf{P}_{n+1}^\bullet by fixing every non-radius arc, and toggling the tags of the radii. A transformation $\mu_{J_k} \circ \dots \circ \mu_{J_1}$ associated with a sequence J_1, \dots, J_k of G -orbits in $\{1, 2, \dots, n+1\}$ corresponds to a sequence of flips associated with non-radius arcs or pairs of “parallel” radii with different tagging. (The latter step replaces a pair of parallel radii inside a punctured digon by another such pair of radii – the pair incident to the other vertex of the digon.) Such a sequence transforms T_0 into another G -invariant tagged triangulation T of the punctured polygon \mathbf{P}_{n+1}^\bullet in which the labels 1 and 2 are assigned to a pair of parallel radii. It is easy to see that the quiver associated to T is G -admissible (cf., e.g., the right picture in Figure 5.8), and so Q_0 is globally foldable. \square

Lemma 5.5.5. *The cluster variables labeled by different G -orbits in a cluster algebra of type B_n (respectively, C_n) are not equal to each other.*

Proof. If two cluster variables appear in the same cluster, then they are necessarily distinct. We now consider two cluster variables which do not appear in the same cluster.

In type C , consider two pairs of centrally symmetric diagonals, or a diameter and a pair of centrally symmetric diagonals, or two diameters. By the observation in the first paragraph, we may assume that at least of the diagonals/diameters cross over each other. In all three cases, there is an octagon that contains the two G -orbits. Freezing all cluster variables except those corresponding to the diagonals of this octagon yields a seed subpattern of type C_3 . In type C_3 , one can check directly that each of 12 distinct G -orbits corresponds to a different cluster variable, say by verifying that the Laurent expansions expressing these cluster variables in terms of a particular initial cluster have different denominators.

In type B , consider two pairs of radii in \mathbf{P}_{n+1}^\bullet , or a pair of radii and an arc which is not a radius, or two arcs which are not radii. In all three cases, the two G -orbits lie inside a certain punctured quadrilateral, reducing the problem to the treatment of a seed subpattern of type B_3 . In type B_3 , there are 12 different G -orbits; using the same method as in type C_3 , we can verify that they correspond to 12 distinct cluster variables. \square

Exercise 5.5.6. By enumerating G -orbits, verify that a cluster algebra of type B_n or C_n has $n^2 + n$ cluster variables. By enumerating G -invariant tagged triangulations and centrally symmetric triangulations, verify that a cluster algebra of type B_n or C_n has $\binom{2n}{n}$ clusters (or seeds).

Proposition 5.5.7. *The same seed pattern cannot be simultaneously of type B_n and of type C_n for $n \geq 3$.*

Proof. By Lemma 2.7.13, if the diagram of an exchange matrix is connected, then its skew-symmetrizing vector is unique up to rescaling. By Exercise 2.7.7, the skew-symmetrizing vector is preserved under mutation. It remains to note that the skew-symmetrizing vectors for the exchange matrices of types B_n and C_n are $(1, 2, 2, \dots, 2)$ and $(2, 1, 1, \dots, 1)$, respectively (up to rescaling and permuting the entries). \square

Examples of coordinate rings having natural cluster algebra structures of types B_n and C_n will be given in Section 6.3.

5.6. Seed patterns of types E_6, E_7, E_8

In this section, we describe a computer-assisted proof of the statement that the cluster algebras (or seed patterns) of exceptional types E_6, E_7, E_8 are of

finite type. The proof utilizes one of several software packages for computing with cluster algebras, freely available online. Our personal favorites are the Java applet [20] and the Sage package [24], cf. the links at [8]. Among other things, both the applet and the Sage package allow one to compute seeds and Laurent expansions of cluster variables obtained by applying a sequence of mutations to a given initial seed.

Theorem 5.6.1. *Seed patterns of types E_6 , E_7 , and E_8 are of finite type.*

Proof. It suffices to verify that a seed pattern of type E_8 is of finite type. A pattern of type E_6 or E_7 can be viewed as a subpattern of a pattern of type E_8 , so if the latter has finitely many seeds, then so does the former.

The main part of the proof is a verification that a cluster algebra of type E_8 with trivial coefficients has finitely many seeds. (The case of general coefficients will follow easily, see below.) With the Sage package [24], this is done as follows. The Sage command

```
S24 = ClusterSeed([[0,1],[1,2],[2,3],[4,5],[5,6],[6,7],
                  [0,4],[1,5],[2,6],[3,7],[5,0],[6,1],[7,2]]);
```

defines a seed S24 with the quiver shown in Figure 5.15. Recall that by Exercise 2.6.8, this quiver is mutation equivalent to any orientation of a Dynkin diagram of type E_8 . Next, the Sage command

```
VC = S24.variable_class(ignore_bipartite_belt=True);
```

performs an exhaustive depth-first search to find all seeds that can be obtained from S24 using $\leq N$ mutations, for $N = 0, 1, 2, 3, \dots$; the cluster variables appearing in these seeds are recorded in the list VC. As soon as the calculation stops, the finiteness of the seed pattern is thereby established. Then the command

```
len(VC);
```

produces the output

```
128
```


which is the total number of cluster variables in the pattern. These 128 cluster variables, or more precisely their Laurent expansions in terms of the chosen initial seed, can be displayed by executing the commands

```
for k in range(128): print(VC[k]); print("...");
```

To get a better idea of what happens in the course of the depth-first search, one can run the command

```
SC=S24.mutation_class(show_depth=True, return_paths=True);
```

its output will show how many seeds have been obtained after each stage. These data are recorded in Figure 5.15. We see that no new seeds are found for $N = 14$. The total number of seeds is 25080.



0	1	2	3	4	5	6	7	8	9	10	11	12	≥ 13
1	9	50	196	614	1582	3525	6863	11626	17098	21706	24220	24974	25080

Figure 5.15. Top: the triangulated grid quiver of type E_8 . Bottom: the table showing, for each $N \geq 0$, the number of distinct seeds that can be obtained from the seed with this quiver using $\leq N$ mutations.

The proof of Theorem 5.6.1 for the case of general coefficients can now be completed using a standard argument based on Corollary 4.3.6 and Remark 4.3.7. One only needs to check that the 8×8 exchange matrix

$$B = \begin{bmatrix} 0 & 1 & 0 & 0 & 1 & -1 & 0 & 0 \\ -1 & 0 & 1 & 0 & 0 & 1 & -1 & 0 \\ 0 & -1 & 0 & 1 & 0 & 0 & 1 & -1 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 & 1 \\ -1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & -1 & 0 & 0 & -1 & 0 & 1 & 0 \\ 0 & 1 & -1 & 0 & 0 & -1 & 0 & 1 \\ 0 & 0 & 1 & -1 & 0 & 0 & -1 & 0 \end{bmatrix}.$$

associated to the quiver in Figure 5.15 has full \mathbb{Z} -rank. \square

Remark 5.6.2. The reader may be wondering: why we chose as the initial quiver the triangulated grid quiver in Figure 5.15, rather than an orientation of a Dynkin diagram of type E_8 ? The answer is that a straightforward implementation of the latter strategy appears to be computationally infeasible. A seed pattern of type E_8 has 128 cluster variables. The formulas expressing them as Laurent polynomials in the 8 initial cluster variables are recursively computed in the process of the depth-first search, and then used to compare the seeds to each other. When the triangulated grid quiver is chosen as the initial one, these Laurent polynomials turn out to be quite manageable. As a result, the entire calculation took less than one hour on a MacBook Pro laptop computer (manufactured in 2013), with a 2.6 GHz processor and 8 GB RAM. By comparison, a similar calculation using, as the initial quiver, the Dynkin diagram of type E_8 with an alternating orientation (i.e., each vertex is either a source or a sink) did not terminate within a few days. The explanation likely lies in the fact that the Laurent polynomials expressing some of the 128 cluster variables in terms of the initial ones are extremely cumbersome in this case. To get an idea of the size of these Laurent polynomials (which are known to have positive coefficients), one can

specialize the 8 initial variables to 1, and compute the remaining 120 cluster variables recursively. When the initial quiver is the alternating Dynkin quiver of type E_8 , the largest of these specializations turns out to be equal to 2820839; for the triangulated grid quiver, the corresponding value is 107.

Remark 5.6.3. One naturally arising cluster algebra of type E_8 is the homogeneous coordinate ring of the Grassmannian $\text{Gr}_{3,8}$ of 3 planes in 8-space. Another closely related example is the coordinate ring of the affine space of 3×5 matrices. (See Chapter 6 and Chapter 8 for more details.) Each of these constructions can in principle be used, with or without a computer, to verify that seed patterns of type E_8 are of finite type.

5.7. Seed patterns of types F_4 and G_2

We will now use folding to take care of types F_4 and G_2 .

By Exercise 4.4.12, one can realize a type F_4 exchange matrix as $B(Q)^G$, where Q is an orientation of the type E_6 Dynkin diagram, $G = \mathbb{Z}/2\mathbb{Z}$, and Q is globally foldable. Now Corollary 5.5.4, together with the fact that type E_6 cluster algebras are of finite type, implies that the same is true in type F_4 .

An alternative approach is to use a computer to check directly that a cluster algebra of type F_4 (with no frozen variables) has finitely many seeds (there are 105 of them). We can then use our standard \tilde{B} argument based on Remark 4.3.7, together with the fact that the matrix \tilde{B}^G from Figure 4.3 has full \mathbb{Z} -rank, to complete the proof.

We now turn to cluster algebras of type G_2 . While it follows from the results of Section 5.1 that cluster algebras of type G_2 are of finite type, this result can also be obtained via folding of the type D_4 quiver shown in Figure 5.16. It is easy to see that this quiver is G -foldable, with respect to the natural action of $G = \mathbb{Z}/3\mathbb{Z}$. Since every seed pattern of type D_4 is of finite type, Corollary 5.5.4 implies the same for the type G_2 .




Figure 5.16. The generator of the group $G = \mathbb{Z}/3\mathbb{Z}$ acts on the vertices of the quiver shown on the left by sending $1 \mapsto 2 \mapsto 3 \mapsto 1$ and $4 \mapsto 4$. All vertices are mutable. The rows and columns of the matrix \tilde{B}^G are indexed by the G -orbits $\{1, 2, 3\}$ and $\{4\}$.

5.8. Decomposable types

We refer to the disjoint union of two Dynkin diagrams of types X_n and $Y_{n'}$ as a Dynkin diagram of type $X_n \sqcup Y_{n'}$; and similarly for disjoint unions of three or more Dynkin diagrams.

We have now defined cluster algebras of types A_n , B_n , C_n , D_n , E_6 , E_7 , E_8 , F_4 , and G_2 , corresponding to the indecomposable Cartan matrices of finite type, or equivalently, the connected Dynkin diagrams. A seed pattern (or cluster algebra) of rank $n + n'$ is said to be of type $X_n \sqcup Y_{n'}$ if one of its exchange matrices is (up to a simultaneous permutation of rows and columns) a block-diagonal matrix with blocks whose Cartan counterparts are of types X_n and $Y_{n'}$. We have the following simple lemma.

Lemma 5.8.1. *If \mathcal{A} is a cluster algebra of decomposable type $X_n \sqcup Y_{n'}$, then the number of cluster variables is the sum of those of the cluster algebras of types X_n and $Y_{n'}$, and the number of clusters is the product.*

Proof. Suppose that \mathcal{A} is of type $X_n \sqcup Y_{n'}$. Then (after a simultaneous permutation of rows and columns) one of its exchange matrices B is block-diagonal with blocks whose Cartan counterparts are of types X_n and $Y_{n'}$. Then the labeled seeds of \mathcal{A} are all pairs of the form $(\mathbf{x}^1 \sqcup \mathbf{x}^2, B^1 \sqcup B^2)$, where each (\mathbf{x}^i, B^i) is a labeled seed of the cluster algebra associated to the i th block of B , $\mathbf{x}^1 \sqcup \mathbf{x}^2$ denotes the concatenation of the clusters \mathbf{x}^1 and \mathbf{x}^2 , and $B^1 \sqcup B^2$ denotes the block-diagonal matrix with blocks B^1 and B^2 . The statement of the lemma follows. \square

We are now ready to complete the proof of the “if” direction of Theorem 5.2.8, restated below for the convenience of the reader.

Corollary 5.8.2. *If the Cartan counterpart of an exchange matrix of a seed pattern (or a cluster algebra) is a Cartan matrix of finite type, then the seed pattern is of finite type.*

Proof. Putting together Theorems 5.3.2, 5.4.2, 5.5.1, 5.5.2, and the results of Sections 5.6–5.7, we conclude that if the Cartan counterpart $A = A(B)$ of one of the exchange matrices B associated to a seed pattern is an indecomposable Cartan matrix of finite type, then the seed pattern has finite type. In the general (decomposable) case, the same conclusion follows from Lemma 5.8.1. \square

5.9. Enumeration of clusters and cluster variables

In this section we present formulas for the number of cluster variables and clusters in each cluster algebra of finite type. These enumerative invariants

provide a way to distinguish cluster algebras of different types, leading to a proof of Theorem 5.2.12.

One important property of finite type cluster algebras is that the underlying combinatorics does not depend on the choice of coefficient tuple. (Conjecturally this holds for arbitrary cluster algebras, see Section 9.1.) We have seen that in cluster algebras of type A_n , seeds are in bijection with triangulations of an $(n+3)$ -gon, regardless of the choice of coefficient tuple. Similarly, in cluster algebras of type D_n , seeds are in bijection with tagged triangulations of a punctured n -gon, for any choice of coefficients. Seeds of cluster algebras of types C_n and B_n are in bijection with folded triangulations and tagged triangulations, respectively. For cluster algebras of exceptional types other than E_7 , the exchange matrices all have full \mathbb{Z} -rank. It then follows from Corollary 4.3.6 and Remark 4.3.7 that the combinatorics of seeds is independent of the choice of coefficients: if an exchange matrix $B(t)$ has full \mathbb{Z} -rank, and $\tilde{B}(t)$ is obtained from $B(t)$ by appending some additional rows, then the rows of $\tilde{B}(t)$ lie in the \mathbb{Z} -span of the rows of $B(t)$, and also the rows of $B(t)$ lie in the \mathbb{Z} -span of the rows of $\tilde{B}(t)$. Therefore Corollary 4.3.6 implies that the seeds of the seed patterns associated with $B(t)$ and with $\tilde{B}(t)$ are in bijection with each other.

In type E_7 , one needs to add one additional row to the exchange matrix B in order to obtain an exchange matrix \tilde{B} of full \mathbb{Z} -rank. One can then check (for example by computer) that the seeds of the seed patterns associated with B and \tilde{B} , respectively, are in bijection. The same argument as before implies that the underlying combinatorics of any cluster algebra of type E_7 does not depend on the choice of coefficient tuple.

Proposition 5.9.1. *Let X_n be a connected Dynkin diagram. The numbers of seeds and cluster variables in a seed pattern of type X_n are given by the values in the corresponding column of the table in Figure 5.17. Alternatively, let Φ be an irreducible finite crystallographic root system of type X_n . Then*

$$(5.9.1) \quad \#\text{seeds}(X_n) = \prod_{i=1}^n \frac{e_i + h + 1}{e_i + 1},$$

$$(5.9.2) \quad \#\text{clvar}(X_n) = \frac{n(h+2)}{2},$$

where e_1, \dots, e_n are the exponents of Φ , and h is the corresponding Coxeter number.

Proof. The values in Figure 5.17 can be verified case by case. The types A_n , B_n/C_n , and D_n were worked out in Corollary 5.3.6, Exercise 5.5.6, and Corollary 5.4.13, respectively. Exceptional types can be handled using the software packages discussed in Section 5.6. It is then straightforward

X_n	A_n	B_n, C_n	D_n	E_6	E_7	E_8	F_4	G_2
$\#\text{seeds}(X_n)$	$\frac{1}{n+2} \binom{2n+2}{n+1}$	$\binom{2n}{n}$	$\frac{3n-2}{n} \binom{2n-2}{n-1}$	833	4160	25080	105	8
$\#\text{clvar}(X_n)$	$\frac{n(n+3)}{2}$	$n(n+1)$	n^2	42	70	128	28	8

Figure 5.17. Enumeration of seeds and cluster variables

to check that the formulas (5.9.1)–(5.9.2) match the values shown in Figure 5.17. (See, e.g., [2, 9] for the values of the exponents and Coxeter numbers for all types.) \square

Remark 5.9.2. Since the number $\frac{1}{n+2} \binom{2n+2}{n+1}$ of seeds in type A_n is a Catalan number, the numbers in (5.9.1) can be regarded as generalizations of the Catalan numbers to arbitrary Dynkin diagrams.

Remark 5.9.3. The number of cluster variables is alternatively given by

$$\#\text{clvar}(X_n) = \frac{\#\text{roots}(X_n)}{2} + n,$$

where $\#\text{roots}(X_n)$ denotes the number of roots in the root system Φ of type X_n . Thus cluster variables are equinumerous to the roots which are either positive or negative simple (i.e., the negatives of simple roots). A natural labeling of the cluster variables by these “almost positive” roots was described and studied in [13].

The reader is referred to [9] for a detailed discussion of cluster combinatorics of finite type, and further references.

Recall from Lemma 5.8.1 that if \mathcal{A} is a cluster algebra of decomposable type $X_n \sqcup Y_{n'}$, then the number of cluster variables is the sum of those for the cluster algebras of types X_n and $Y_{n'}$, and the number of seeds is the product. Therefore Proposition 5.9.1 allows us to compute the number of cluster variables and seeds for any cluster algebra of finite type.

Proof of Theorem 5.2.12. The implication (1) \Rightarrow (2) is easy to establish. Suppose that the Cartan counterparts $A(B')$ and $A(B'')$ are Cartan matrices of the same finite type. In the case of simply laced types ADE , this means that the corresponding quivers are (possibly different) orientations of isomorphic Dynkin diagrams. By Exercise 2.6.5, these two quivers are related to each other by a sequence of mutations at sources and sinks, and consequently B' and B'' are mutation equivalent. The remaining cases $BCFG$ are treated in a similar fashion, using an appropriate analogue of Exercise 2.6.5.

Let us prove the implication (2) \Rightarrow (1). We first observe that it suffices to establish this result in the indecomposable case, since mutations transform the connected components of a quiver (or their analogues for skew-symmetrizable matrices) independently of each other.

Let B' and B'' be mutation equivalent exchange matrices of types X' and X'' , respectively, where X' and X'' are connected Dynkin diagrams. We need to show that X' and X'' are of the same type. This is done as follows. By Proposition 5.9.1, the number of cluster variables in a seed pattern associated with such an exchange matrix is uniquely determined by the type of the corresponding Dynkin diagram (i.e., it does not depend on the choice of coefficient tuple). Moreover no two connected Dynkin types of the same rank produce the same number of cluster variables—with the exception of the pairs (B_n, C_n) . For B_n versus C_n , the claim follows from Proposition 5.5.7. \square

We conclude this section by an elementary observation that shows that the problem of enumerating cluster variables does not make sense outside of finite type.

Proposition 5.9.4. *A seed pattern is of finite type if and only if it has finitely many cluster variables.*

Proof. One direction is obvious: if a seed pattern has finitely many distinct seeds, then it has finitely many cluster variables. Conversely, if a seed pattern has finitely many cluster variables, then the only way it could possibly have infinitely many distinct seeds is if there were infinitely many distinct extended exchange matrices. By the Pigeonhole principle, one cluster would have to appear with infinitely many different extended exchange matrices \tilde{B} , which implies that one of its cluster variables x_j would appear in infinitely many exchange relations, leading to infinitely many different x'_j 's. \square

5.10. 2-finite exchange matrices

In this section, we complete the proof of Theorem 5.2.8, closely following [11]. The notion of a 2-finite (skew-symmetrizable) matrix introduced in Definition 5.1.3 plays a key role.

We establish Theorem 5.2.8 by including it in the following statement.

Theorem 5.10.1. *For a seed $\Sigma = (\mathbf{x}, \mathbf{y}, B)$, the following are equivalent:*

- (1) *there exists a matrix B' mutation equivalent to B such that its Cartan counterpart $A(B')$ is a Cartan matrix of finite type;*
- (2) *the seed pattern (or the cluster algebra) defined by Σ is of finite type;*
- (3) *the exchange matrix B is 2-finite.*

The implication $\boxed{(1) \Rightarrow (2)}$ in Theorem 5.10.1 is nothing but Corollary 5.8.2. To the best of our knowledge, the only known proof of the reverse implication $(2) \Rightarrow (1)$ goes through property (3).

The implication $\boxed{(2) \Rightarrow (3)}$ is precisely Corollary 5.1.4.

All that remains in order to complete the proof of Theorem 5.10.1 (hence Theorem 5.2.8) is to prove the implication $\boxed{(3) \Rightarrow (1)}$. We reformulate the latter below as a standalone statement.

Proposition 5.10.2. *Let $B = (b_{ij})$ be a 2-finite skew-symmetrizable integer matrix. Then there exists a matrix B' mutation equivalent to B such that its Cartan counterpart $A(B')$ is a Cartan matrix of finite type.*

In the rest of this section, we outline the proof of Proposition 5.10.2 given in [11, Sections 7–8]. The proof is purely combinatorial and rather technical. The missing details (all of them relatively minor) can be found in *loc. cit.*

The proof of Proposition 5.10.2 makes heavy use of the notion of diagram from Definition 2.7.10. Recall from Proposition 2.7.11 that mutation is well-defined for diagrams, and we write $\Gamma \sim \Gamma'$ to denote that diagrams Γ and Γ' are mutation equivalent.

A diagram Γ is called *2-finite* if every diagram $\Gamma' \sim \Gamma$ has all edge weights equal to 1, 2 or 3; otherwise we refer to Γ as *2-infinite*. Thus a matrix B is 2-finite if and only if its diagram $\Gamma(B)$ is 2-finite. Note that a diagram is 2-finite if and only if so are all its connected components.

We now restate Proposition 5.10.2 in the language of diagrams.

Proposition 5.10.3. *Any 2-finite diagram is mutation equivalent to an orientation of a Dynkin diagram (where the weight of each edge is understood as its multiplicity in the corresponding Dynkin diagram).*

We note that all orientations of a given Dynkin diagram are mutation equivalent to each other, as they are related by source-or-sink mutations as in the proof of Theorem 5.2.12. This is true more generally for any diagram whose underlying graph is a tree.

A *subdiagram* of a diagram Γ is a diagram $\Gamma' \subset \Gamma$ obtained by taking an induced directed subgraph of Γ and keeping all its edge weights intact.

The proof of Proposition 5.10.3 repeatedly makes use of the following obvious property: any subdiagram of a 2-finite diagram is 2-finite. Equivalently, any diagram that has a 2-infinite subdiagram is 2-infinite. Thus, in order to show that a given diagram is 2-infinite, it suffices to exhibit a sequence of mutations that creates an edge of weight 4 or larger, or a subdiagram which is already known to be 2-infinite. The strategy is to catalogue

enough 2-infinite subdiagrams to be able to show that any diagram avoiding them has to be mutation equivalent to an orientation of a Dynkin diagram.

We first examine two special classes of diagrams, those whose underlying graphs are trees or cycles, respectively. We refer to them as *tree diagrams* and *cycle diagrams*, respectively.

Proposition 5.10.4. *Any 2-finite tree diagram is an orientation of a connected Dynkin diagram.*

For the proof we consider a class of diagrams defined as follows. A diagram Γ is called an *extended Dynkin tree diagram* if

- Γ is a tree diagram with edge weights ≤ 3 ;
- Γ is not on the Dynkin diagram list;
- any proper subdiagram of Γ is a Dynkin diagram (possibly disconnected).

To prove the proposition, it is enough to show that any orientation of any extended Dynkin tree diagram is 2-infinite. Direct inspection shows that Figure 5.18 provides a complete list of such diagrams (as discussed above, the choice of an orientation for a tree diagram is immaterial). We note that all these diagrams are associated with untwisted affine Lie algebras and can be found in the tables in [2] or in [17, Chapter 4, Table Aff 1]. The only diagram from those tables that is missing in Figure 5.18 is $A_n^{(1)}$, which is an $(n + 1)$ -cycle; it will appear later in our discussion of cycle diagrams.

In showing that an extended Dynkin tree diagram is 2-infinite, we can arbitrarily choose its orientation. We start with the three infinite series $B_n^{(1)}$, $C_n^{(1)}$, and $D_n^{(1)}$, each time orienting all the edges left to right. Let us denote the diagram in question by $X_n^{(1)}$, and let n_\circ be the minimal value of n . So if $X = D$ (resp., B , C), then n_\circ equals 4 (resp., 3, 2). If $n > n_\circ$, then mutating at the second vertex from the left, and subsequently removing this vertex (together with all incident edges) leaves us with a subdiagram of type $X_{n-1}^{(1)}$. Using induction on n , it suffices to check the base cases $D_4^{(1)}$, $B_3^{(1)}$ and $C_2^{(1)}$. It is not hard to check that each of these three diagrams is 2-infinite. The same applies to extended Dynkin trees of types $F_4^{(1)}$ and $G_2^{(1)}$.

The remaining three cases $E_6^{(1)}$, $E_7^{(1)}$ and $E_8^{(1)}$ can be treated in a similar manner (with or without a computer) but we prefer another approach. To describe it, we will need to introduce some notation.

Definition 5.10.5. For $p, q, r \in \mathbb{Z}_{\geq 0}$, we denote by $T_{p,q,r}$ the tree diagram (with all edge weights equal to 1) on $p + q + r + 1$ vertices obtained by connecting an endpoint of each of the three chains A_p , A_q and A_r to a single extra vertex (see Figure 5.19).




Figure 5.18. Extended Dynkin tree diagrams. Each tree $X_n^{(1)}$ has $n+1$ vertices. All unspecified edge weights are equal to 1.




Figure 5.19. The tree diagram $T_{5,4,2}$.

Definition 5.10.6. For $p, q, r \in \mathbb{Z}_{>0}$ and $s \in \mathbb{Z}_{\geq 0}$, let $S_{p,q,r}^s$ denote the diagram (with all edge weights equal to 1) on $p+q+r+s$ vertices obtained by attaching three branches A_{p-1} , A_{q-1} , and A_{r-1} to three consecutive vertices on a *cyclically oriented* $(s+3)$ -cycle (see Figure 5.20).




Figure 5.20. The diagram $S_{4,3,2}^5$.

Note that in both definitions, the choice of orientations for the edges where they are not shown is immaterial: different choices lead to mutation-equivalent diagrams. For $T_{p,q,r}$, this follows from Exercise 2.6.5; for $S_{p,q,r}^s$, one needs a slight generalization of this result, see [11, Proposition 9.2].

Exercise 5.10.7. Show that the diagram $S_{p,q,r}^s$ is mutation equivalent to $T_{p+r-1,q,s}$.

With the help of Exercise 5.10.7, the proof of Proposition 5.10.4 can now be completed, using the observations that

$$\begin{aligned} E_6^{(1)} &= T_{2,2,2} \sim S_{2,2,1}^2 \supset D_5^{(1)}; \\ E_7^{(1)} &= T_{3,1,3} \sim S_{3,1,1}^3 \supset E_6^{(1)}; \\ E_8^{(1)} &= T_{2,1,5} \sim S_{2,1,1}^5 \supset E_7^{(1)}. \end{aligned}$$

Turning to the cycle diagrams, we have the following classification.

Exercise 5.10.8. Let Γ be a 2-finite diagram whose underlying graph is an n -cycle (with some orientation of edges). Show that Γ is cyclically oriented, and moreover it must be one of the following (see Figure 5.21):

- (a) an n -cycle with all weights equal to 1 (in this case, $\Gamma \sim D_n$);
- (b) a 3-cycle with edge weights 2, 2, 1 (in this case, $\Gamma \sim B_3$);
- (c) a 4-cycle with edge weights 2, 1, 2, 1 (in this case, $\Gamma \sim F_4$).

Proof of Proposition 5.10.3. We proceed by induction on n , the number of vertices in Γ . If $n \leq 3$, then Γ is either a tree or a cycle, and the theorem follows by Proposition 5.10.4 and Exercise 5.10.8. So let us assume that the statement is already known for some $n \geq 3$; we need to show that it holds for a diagram Γ on $n + 1$ vertices. Pick a vertex $v \in \Gamma$ such that the subdiagram $\Gamma' = \Gamma - \{v\}$ is connected. Since Γ' is 2-finite, it is mutation equivalent to some Dynkin diagram X_n . Furthermore, we may assume that




Figure 5.21. 2-finite cycles.

Γ' is (isomorphic to) our favorite representative of the mutation equivalence class of X_n . For each X_n , we will choose a representative that is most convenient for the purposes of this proof, and use the classifications of 2-finite tree and cycle diagrams obtained above to achieve the desired goal.

Case 1: Γ' is an orientation of a Dynkin diagram with no branching point, i.e., is of one of the types $A_n, B_n, C_n, F_4,$ or G_2 . Let us orient the edges of Γ' so that they all point in the same direction. If v is adjacent to exactly one vertex of Γ' , then Γ is a tree, and we are done by Proposition 5.10.4. If v is adjacent to more than two vertices of Γ' , then Γ has a cycle subdiagram whose edges are not cyclically oriented, contradicting Exercise 5.10.8. Thus we may assume that v is adjacent to precisely two vertices v_1 and v_2 of Γ' , see Figure 5.22. Then Γ has precisely one cycle \mathcal{C} , which must be of one of the types (a)–(c) shown in Figure 5.21.

Suppose that \mathcal{C} is an oriented cycle with unit edge weights. If Γ has an edge of weight ≥ 2 , then it contains a subdiagram of type $B_m^{(1)}$ or $G_2^{(1)}$, unless \mathcal{C} is a 3-cycle, in which case $\mu_v(\Gamma)$ is a tree, and we are done by Proposition 5.10.4. If all edges in Γ are of weight 1, then it is one of the diagrams $S_{p,q,r}^s$ in Exercise 5.10.7 (with $q = 0$). Hence Γ is mutation equivalent to a tree, and we are done.

Suppose that \mathcal{C} is as in Figure 5.21(b). If one of the edges (v, v_1) and (v, v_2) has weight 1, then μ_v removes the edge (v_1, v_2) , resulting in a tree, and we are done again. So assume that both (v, v_1) and (v, v_2) have weight 2. If at least one edge outside \mathcal{C} has weight ≥ 2 , then $\Gamma \supset C_m^{(1)}$ or $\Gamma \supset G_2^{(1)}$. It remains to consider the case shown in Figure 5.22. A direct check shows that $\mu_\ell \circ \cdots \circ \mu_2 \circ \mu_1 \circ \mu_{v_2} \circ \mu_v(\Gamma) = B_{n+1}$, and we are done.




Figure 5.22. Second subcase in Case 1.

Suppose that \mathcal{C} is as in Figure 5.21(c). It suffices to show that any diagram \mathcal{C}' obtained from \mathcal{C} by adjoining a single vertex adjacent to one of its vertices is 2-infinite. Indeed, if this extra edge has weight 1 (resp., 2, 3), then \mathcal{C}' has a 2-infinite subdiagram of type $B_3^{(1)}$ (resp., $C_2^{(1)}$, $G_2^{(1)}$).

Case 2: $\Gamma' \sim D_n$ ($n \geq 4$). We may assume that Γ' is an oriented n -cycle with unit edge weights.

If v is adjacent to two non-adjacent vertices of Γ' (and possibly others), then Γ contains an improperly oriented cycle, contradicting Exercise 5.10.8.

Suppose v is adjacent to a single vertex $v_1 \in \Gamma'$. If the edge (v, v_1) has weight ≥ 2 , then Γ has a subdiagram $B_3^{(1)}$ or $G_2^{(1)}$. If (v, v_1) has weight 1, then by Exercise 5.10.7, Γ is mutation equivalent to a tree, and we are done.

Suppose that v is adjacent to exactly two vertices v_1 and v_2 which are adjacent to each other. Then the triangle (v, v_1, v_2) is either an oriented 3-cycle with unit edge weights or the diagram in Figure 5.21(b). In the former case, $\mu_v(\Gamma)$ is an oriented $(n+1)$ -cycle, so $\Gamma \sim D_{n+1}$. In the latter case, $\mu_v(\Gamma)$ contains an improperly oriented (hence 2-infinite) cycle.

Case 3: $\Gamma' \sim E_n = T_{1,2,n-4}$, for $n \in \{6, 7, 8\}$. By Exercise 5.10.7, we may assume that $\Gamma' = S_{1,2,1}^{n-4}$, i.e., Γ' consists of an oriented $(n-1)$ -cycle \mathcal{C} with unit edge weights, and an extra edge of weight 1 connecting a vertex in \mathcal{C} to a vertex $v_1 \notin \mathcal{C}$.

There are several subcases to examine, depending on how v connects to \mathcal{C} . It is routine (if tedious) to check that in each of these subcases, Γ must be equivalent to an orientation of a Dynkin diagram (e.g. because it is equivalent to a tree, or to one of the diagrams treated in Cases 1 and 2 above), or else Γ is not 2-finite. Details can be found in [11].

This concludes the proof of Proposition 5.10.3. As a consequence, we obtain Proposition 5.10.2, Theorem 5.10.1, and Theorem 5.2.8. \square

Remark 5.10.9. The property of being 2-finite is clearly hereditary. Therefore the other two equivalent properties of exchange matrices appearing in Theorem 5.10.1 are hereditary as well.

5.11. Quasi-Cartan companions

An unpleasant feature of the finite type classification (see Theorem 5.10.1) is that it does not provide an effective way to verify whether a given exchange matrix B defines a seed pattern of finite type: both condition (3) (2-finiteness) and condition (1) (being mutation equivalent to a skew-symmetrizable version of a Cartan matrix) impose a restriction on *all* matrices in the mutation class of B . An alternative criterion formulated directly in terms of the matrix B (rather than its mutation class) was given in [1]. We reproduce this result below while omitting the technical part of the proof.

Remark 5.11.1. A different finite type recognition criterion was given in [25], by explicitly listing all minimal obstructions to finite type. More precisely, [25] provides a list of all (up to isomorphism) *minimal 2-infinite diagrams*, i.e., all diagrams which are not 2-finite but whose proper subdiagrams are all 2-finite. Then B is of finite type if and only if $\Gamma(B)$ does not contain a subdiagram on this list. Unfortunately, the list is rather long: it includes 10 infinite series and a large number of exceptional diagrams of size ≤ 9 .

Definition 5.11.2. A *quasi-Cartan matrix* is a symmetrizable (square) matrix $A = (a_{ij})$ with integer entries such that $a_{ii} = 2$ for all i . Note that in such a matrix, the entries a_{ij} and a_{ji} always have the same sign. Unlike for generalized Cartan matrices, they are allowed to be positive.

A quasi-Cartan matrix A is *positive* if the symmetrized matrix is positive definite, or equivalently if the principal minors of A are all positive.

A quasi-Cartan matrix A is called a *quasi-Cartan companion* of a skew-symmetrizable integer matrix B if $|a_{ij}| = |b_{ij}|$ for all $i \neq j$. Thus B can have several quasi-Cartan companions one of which is the Cartan counterpart of B given by Definition 5.2.7. (To be precise, the number of quasi-Cartan companions of B is 2^e where e is the number of edges in the diagram of B .)

A *chordless cycle* in the diagram $\Gamma(B)$ is an induced subgraph isomorphic to a cycle (with arbitrary orientation).

Theorem 5.11.3. *For a skew-symmetrizable integer matrix B , each of the conditions (1)–(3) in Theorem 5.10.1 is equivalent to*

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

The key ingredient in the proof of Theorem 5.11.3 given in [1] is the following lemma, whose proof we omit.

Lemma 5.11.4 ([1, Lemma 4.1]). *Property (4) in Theorem 5.11.3 is preserved under mutations of skew-symmetrizable integer matrices.*

Proof of Theorem 5.11.3 modulo Lemma 5.11.4. We deduce the implications (1) \Rightarrow (4) \Rightarrow (3) from Lemma 5.11.4. To prove that (1) \Rightarrow (4), it is enough to observe that if $\Gamma(B)$ is a Dynkin diagram, then B satisfies (4). (Indeed, the Cartan counterpart of B is positive, and $\Gamma(B)$ has no cycles.) To prove (4) \Rightarrow (3), note that any positive quasi-Cartan matrix $A = (a_{ij})$ satisfies $|a_{ij}a_{ji}| \leq 3$ for all $i \neq j$ because of the positivity of the principal 2×2 minor of A occupying the rows and columns i and j . \square

Remark 5.11.5. As explained above, one can use Lemma 5.11.4 to establish the implications (1) \Rightarrow (4) \Rightarrow (3). In combination with the arguments

given in Section 5.10, this yields a self-contained combinatorial proof of the equivalence (3) \Leftrightarrow (1).

A skew-symmetrizable integer matrix B can have many quasi-Cartan companions A , corresponding to different choices for the signs of its off-diagonal matrix entries. Note that the positivity property of A is preserved by simultaneous sign changes in rows and columns. It turns out that for the purposes of checking (4) for a given matrix B , there is a *unique*, up to these transformations, sign pattern for A that needs to be checked for positivity. More precisely, we have the following result, cf. [1, Propositions 1.4–1.5].

Proposition 5.11.6. *Let B be a skew-symmetrizable integer matrix such that each chordless cycle in $\Gamma(B)$ is cyclically oriented. Then B has a quasi-Cartan companion $A = (a_{ij})$ such that the sign condition*


$$(5.11.1) \quad \prod_{\{i,j\} \in Z} (-a_{ij}) < 0$$

(product over all edges $\{i, j\}$ in Z) is satisfied for every chordless cycle Z . In fact, A is unique up to simultaneous sign changes in rows and columns. Moreover B satisfies conditions (1)–(3) in Theorem 5.10.1 if and only if A is positive.

Remark 5.11.7. Several characterizations of positive quasi-Cartan matrices have been given in [1, Proposition 2.9]. In particular, these matrices are, up to certain equivalence, also classified by Cartan-Killing types. More precisely, any positive quasi-Cartan matrix corresponding to a root system Φ has the entries $a_{ij} = \langle \beta_i^\vee, \beta_j \rangle$, where $\{\beta_1, \dots, \beta_n\} \subset \Phi$ is a \mathbb{Z} -basis of the root lattice generated by Φ , and β^\vee is the coroot dual to a root $\beta \in \Phi$.

We conclude this section by an example illustrating the use of Proposition 5.11.6 for checking whether a particular skew-symmetric matrix is an exchange matrix of a seed pattern of finite type.

Example 5.11.8. Let $Q(n)$ be the following quiver with vertices $1, 2, \dots, n$:



The quiver $Q(n)$ is the diagram of its $n \times n$ exchange matrix

$$B(n) = B(Q(n)) = \begin{bmatrix} 0 & -1 & 1 & 0 & \cdots & 0 & 0 \\ 1 & 0 & -1 & 1 & \cdots & 0 & 0 \\ -1 & 1 & 0 & -1 & \cdots & 0 & 0 \\ 0 & -1 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & -1 \\ 0 & 0 & 0 & 0 & \cdots & 1 & 0 \end{bmatrix}$$

(that is, $\Gamma(B(n)) = Q(n)$). This quiver has $n - 2$ chordless cycles, the 3-cycles with vertices $\{i, i + 1, i + 2\}$, for $i = 1, \dots, n - 2$. All of them are cyclically oriented. Now let $A(n)$ be the quasi-Cartan companion of $B(n)$ such that $a_{ij} = b_{ij}$ for $i < j$. One immediately checks that $A(n)$ satisfies the sign condition (5.11.1). Let $\delta_n = \det(A(n))$. By Sylvester's criterion, $A(n)$ is positive if and only if all the numbers $\delta_1, \dots, \delta_n$ are positive.

It is not hard to compute the generating function of the sequence (δ_n) , with the convention $\delta_0 = 1$:

$$(5.11.2) \quad \sum_{n \geq 0} \delta_n x^n = \frac{(1+x)(1+x+x^2)(1+x^2)(1+x^3)}{1-x^{12}}.$$

We see that $\delta_{n+12} = \delta_n$ for $n \geq 0$. Since the numerator in (5.11.2) is a polynomial of degree 8, we conclude that $\delta_9 = \delta_{10} = \delta_{11} = 0$. The fact that $\delta_9 = 0$ implies that $A(n)$ is not positive (hence $B(n)$ is not 2-finite) for $n \geq 9$.

The values of δ_n for $1 \leq n \leq 8$ are given in Figure 5.23; cf. Exercise 2.6.8. As all of them are positive, we conclude that $A(n)$ is positive (and so $B(n)$ is 2-finite) if and only if $n \leq 8$. The corresponding Cartan-Killing types are shown in Figure 5.23; we leave the verification to the reader.

n	1	2	3	4	5	6	7	8
$\delta_n = \det(A(n))$	2	3	4	4	4	3	2	1
Cartan-Killing type	A_1	A_2	A_3	D_4	D_5	E_6	E_7	E_8

Figure 5.23. Determinants and Cartan-Killing types of the matrices $A(n)$.

Bibliography

- [1] BAROT, M., GEISS, C., AND ZELEVINSKY, A. Cluster algebras of finite type and positive symmetrizable matrices. *J. London Math. Soc. (2)* 73, 3 (2006), 545–564.
- [2] BOURBAKI, N. *Éléments de mathématique. Fasc. XXXIV. Groupes et algèbres de Lie. Chapitre IV: Groupes de Coxeter et systèmes de Tits. Chapitre V: Groupes engendrés par des réflexions. Chapitre VI: systèmes de racines.* Actualités Scientifiques et Industrielles, No. 1337. Hermann, Paris, 1968.
- [3] BUAN, A. B., MARSH, R. J., AND REITEN, I. Cluster mutation via quiver representations. *Comment. Math. Helv.* 83, 1 (2008), 143–177.
- [4] CONWAY, J. H., AND COXETER, H. S. M. Triangulated polygons and frieze patterns. *Math. Gaz.* 57, 400 (1973), 87–94.
- [5] COXETER, H. S. M. Frieze patterns. *Acta Arith.* 18 (1971), 297–310.
- [6] DUPONT, G. An approach to non-simply laced cluster algebras. *J. Algebra* 320, 4 (2008), 1626–1661.
- [7] FELIKSON, A., SHAPIRO, M., AND TUMARKIN, P. Cluster algebras of finite mutation type via unfoldings. *Int. Math. Res. Not. IMRN*, 8 (2012), 1768–1804.
- [8] FOMIN, S. Cluster algebras portal, <http://www.math.lsa.umich.edu/~fomin/cluster.html>.
- [9] FOMIN, S., AND READING, N. Root systems and generalized associahedra. In *Geometric combinatorics*, vol. 13 of *IAS/Park City Math. Ser.* Amer. Math. Soc., Providence, RI, 2007, pp. 63–131.
- [10] FOMIN, S., AND ZELEVINSKY, A. Cluster algebras. I. Foundations. *J. Amer. Math. Soc.* 15, 2 (2002), 497–529 (electronic).
- [11] FOMIN, S., AND ZELEVINSKY, A. Cluster algebras. II. Finite type classification. *Invent. Math.* 154, 1 (2003), 63–121.
- [12] FOMIN, S., AND ZELEVINSKY, A. Cluster algebras: notes for the CDM-03 conference. In *Current developments in mathematics, 2003*. Int. Press, Somerville, MA, 2003, pp. 1–34.
- [13] FOMIN, S., AND ZELEVINSKY, A. Y-systems and generalized associahedra. *Ann. of Math. (2)* 158, 3 (2003), 977–1018.
- [14] FRASER, C. Quasi-homomorphisms of cluster algebras. *Adv. in Appl. Math.* 81 (2016), 40–77.

- [15] FULTON, W., AND HARRIS, J. *Representation theory*, vol. 129 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991.
- [16] HUMPHREYS, J. E. *Reflection groups and Coxeter groups*, vol. 29 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1990.
- [17] KAC, V. G. *Infinite-dimensional Lie algebras*, third ed. Cambridge University Press, Cambridge, 1990.
- [18] KANG, S.-J., KASHIWARA, M., KIM, M., AND OH, S.-J. Monoidal categorification of cluster algebras. *J. Amer. Math. Soc.* 31, 2 (2018), 349–426.
- [19] KASHIWARA, M. On crystal bases of the Q -analogue of universal enveloping algebras. *Duke Math. J.* 63, 2 (1991), 465–516.
- [20] KELLER, B. Quiver mutation in Java, <https://webusers.imj-prg.fr/~bernhard.keller/quivermutation/>.
- [21] KOSTANT, B. The solution to a generalized Toda lattice and representation theory. *Adv. in Math.* 34, 3 (1979), 195–338.
- [22] LABARDINI-FRAGOSO, D., AND ZELEVINSKY, A. Strongly primitive species with potentials I: mutations. *Bol. Soc. Mat. Mex. (3)* 22, 1 (2016), 47–115.
- [23] LUSZTIG, G. *Introduction to quantum groups*, vol. 110 of *Progress in Mathematics*. Birkhäuser Boston, Inc., Boston, MA, 1993.
- [24] MUSIKER, G., AND STUMP, C. A compendium on the cluster algebra and quiver package in Sage. *Sem. Lothar. Comb.* 65 (2011). Article B65d.
- [25] SEVEN, A. I. Recognizing cluster algebras of finite type. *Electron. J. Combin.* 14, 1 (2007), Research Paper 3, 35 pp. (electronic).
- [26] STANLEY, R. P. *Enumerative combinatorics. Vol. 2*, vol. 62 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999. With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin.
- [27] SZEGŐ, G. *Orthogonal polynomials*, fourth ed. American Mathematical Society, Providence, R.I., 1975. American Mathematical Society, Colloquium Publications, Vol. XXIII.
- [28] YANG, S.-W., AND ZELEVINSKY, A. Cluster algebras of finite type via Coxeter elements and principal minors. *Transform. Groups* 13, 3-4 (2008), 855–895.

Introduction to Cluster Algebras

Chapter 6

(preliminary version)

SERGEY FOMIN

LAUREN WILLIAMS

ANDREI ZELEVINSKY

Preface

This is a preliminary draft of Chapter 6 of our forthcoming textbook *Introduction to cluster algebras*, joint with Andrei Zelevinsky (1953–2013).

Other chapters have been posted as

- [arXiv:1608.05735](#) (Chapters 1–3),
- [arXiv:1707.07190](#) (Chapters 4–5), and
- [arXiv:2106.02160](#) (Chapter 7).

We expect to post additional chapters in the not so distant future.

We thank Bernard Leclerc and Karen Smith for their invaluable advice. The algebraic version of the argument in the proof of Proposition 6.4.1 was suggested by Karen.

We are grateful to Colin Defant, Chris Fraser, Anne Larsen, Amal Mattoo, Hanna Mularczyk, and Raluca Vlad for a number of comments on the earlier versions of this chapter, and for assistance with creating figures.

Our work was partially supported by the NSF grants DMS-1664722 and DMS-1854512.

Comments and suggestions are welcome.

Sergey Fomin
Lauren Williams

Contents

Chapter 6. Cluster structures in commutative rings	1
§6.1. Introductory examples	2
§6.2. Cluster algebras and coordinate rings	4
§6.3. Examples of cluster structures of classical types	5
§6.4. Starfish lemma	10
§6.5. Cluster structure in the ring $\mathbb{C}[\mathrm{SL}_k]^U$	14
§6.6. Cluster structure in the rings $\mathbb{C}[\mathrm{Mat}_{k \times k}]$ and $\mathbb{C}[\mathrm{SL}_k]$	20
§6.7. The cluster structure in the ring $\mathbb{C}[\widehat{\mathrm{Gr}}_{a,b}]$	22
§6.8. Defining cluster algebras by generators and relations	29
Bibliography	37

Cluster structures in commutative rings

Cluster algebras are commutative rings endowed with a particular kind of combinatorial structure (a *cluster structure*, as we call it). In this chapter, we study the problem of identifying a cluster structure in a given commutative ring, or equivalently the problem of verifying that certain additional data make a given ring a cluster algebra.

Sections 6.1–6.3 provide several examples of cluster structures in coordinate rings of affine algebraic varieties. General techniques used to verify that a given commutative ring is a cluster algebra are introduced in Section 6.4. These techniques are applied in Sections 6.5–6.7 to treat several important classes of cluster algebras: the basic affine spaces for SL_k (Section 6.5), the coordinate rings of $\mathrm{Mat}_{k \times k}$ and SL_k (Section 6.6), and the homogeneous coordinate rings of Grassmannians, also called Plücker rings (Section 6.7). An in-depth study of the latter topic will be given later in Chapter 8, following the development of the required combinatorial tools in Chapter 7. The problem of defining cluster algebras by generators and relations is discussed in Section 6.8.

Section 6.1 is based on [16, Section 12.1] and [20, Section 2]. Sections 6.2 and 6.3 follow [16, Sections 11.1 and 12]. Section 6.4 follows [12, Section 3], which in turn extends the ideas used in the proofs of [3, Theorem 2.10] and subsequently [38, Proposition 7]. The constructions presented in Sections 6.5 and 6.6 predate the general definition of cluster algebras; they essentially go back to [2] and [13], respectively. Our development of cluster structures in Grassmannians (Section 6.7) is different from the original sources [38] and [24, Section 3.3]. The material in Section 6.8 is mostly new.

6.1. Introductory examples

As a warm-up, we discuss several simple examples of cluster structures in commutative rings.

Example 6.1.1. Let $V = \mathbb{C}^{2k}$, with $k \geq 3$, be an even-dimensional vector space with coordinates (x_1, \dots, x_{2k}) . Consider the nondegenerate quadratic form Q on V given by

$$(6.1.1) \quad Q(x_1, \dots, x_{2k}) = \sum_{i=1}^k (-1)^{i-1} x_i x_{2k+1-i}.$$

Let

$$\mathcal{C} = \{v \in V \mid Q(v) = 0\}$$

be the isotropic cone and $\mathbb{P}(\mathcal{C})$ the corresponding smooth quadric in $\mathbb{P}(V)$.

The homogeneous coordinate ring of the quadric (or equivalently the coordinate ring of \mathcal{C}) is the quotient

$$(6.1.2) \quad \mathcal{A} = \mathbb{C}[x_1, \dots, x_{2k}] / \langle Q(x_1, \dots, x_{2k}) \rangle.$$

To see that \mathcal{A} is a cluster algebra, we define, for $1 \leq s \leq k-3$, the functions

$$p_s = \sum_{i=1}^{s+1} (-1)^{s+1-i} x_i x_{2k+1-i}.$$

Then \mathcal{A} has cluster variables $\{x_2, x_3, \dots, x_{k-1}\} \cup \{x_{k+2}, x_{k+3}, \dots, x_{2k-1}\}$ and frozen variables $\{x_1, x_k, x_{k+1}, x_{2k}\} \cup \{p_s \mid 1 \leq s \leq k-3\}$. It has 2^{k-2} clusters defined by choosing, for each $i \in \{2, \dots, k-1\}$, precisely one of x_i and x_{2k+1-i} . The exchange relations are (here $2 \leq i \leq k-1$):

$$x_i x_{2k+1-i} = \begin{cases} p_{i-1} + p_{i-2} & \text{if } 3 \leq i \leq k-2; \\ p_1 + x_1 x_{2k} & \text{if } i = 2 \text{ and } k \neq 3; \\ x_k x_{k+1} + x_1 x_{2k} & \text{if } i = 2 \text{ and } k = 3; \\ x_k x_{k+1} + p_{k-3} & \text{if } i = k-1 \text{ and } k \neq 3. \end{cases}$$

This cluster algebra is of finite type $A_1^{k-2} = A_1 \times A_1 \times \dots \times A_1$. Figure 6.1 shows a seed of \mathcal{A} in the case $k = 5$.

The quadric $\mathbb{P}(\mathcal{C})$ is a homogeneous space G/P (a ‘‘partial flag variety’’) for the special orthogonal group attached to Q . The fact that the coordinate ring $\mathbb{C}[\mathcal{C}]$ is a cluster algebra is a special case of a more general phenomenon. The (multi-)homogeneous coordinate ring of any type A partial flag variety carries a natural cluster algebra structure [19]. For generalizations to other semisimple Lie groups G and parabolic subgroups $P \subset G$, see the survey [22].




Figure 6.1. A seed for the cluster structure on the ring (6.1.2), for $k = 5$. Here $p_1 = -x_1x_{10} + x_2x_9$, $p_2 = x_1x_{10} - x_2x_9 + x_3x_8 = x_4x_7 - x_5x_6$.

Example 6.1.2. Let $\mathcal{A} = \mathbb{C}[a_1, \dots, a_{n+1}, b_1, \dots, b_{n+1}]$ be the coordinate ring of the affine space of $2 \times (n+1)$ matrices

$$(6.1.3) \quad \begin{bmatrix} a_1 & a_2 & \cdots & a_{n+1} \\ b_1 & b_2 & \cdots & b_{n+1} \end{bmatrix}.$$

We will show that \mathcal{A} carries several pairwise non-isomorphic cluster algebra structures.

First, we can identify \mathcal{A} with a cluster algebra of type A_n as follows. The cluster variables and frozen variables are the $2(n+1)$ matrix entries $a_1, \dots, a_{n+1}, b_1, \dots, b_{n+1}$ together with the $\binom{n+1}{2}$ minors (Plücker coordinates) $P_{ij} = a_i b_j - a_j b_i$. The exchange relations are:

$$\begin{aligned} a_i b_j &= P_{ij} + a_j b_i \quad (1 \leq i < j \leq n+1), \\ a_j P_{ik} &= a_i P_{jk} + a_k P_{ij} \quad (1 \leq i < j < k \leq n+1), \\ b_j P_{ik} &= b_i P_{jk} + b_k P_{ij} \quad (1 \leq i < j < k \leq n+1), \\ P_{ik} P_{j\ell} &= P_{ij} P_{k\ell} + P_{i\ell} P_{jk} \quad (1 \leq i < j < k < \ell \leq n+1). \end{aligned}$$

By adding a column $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ at the beginning and a column $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ at the end of the $2 \times (n+1)$ matrix, we obtain a full rank $2 \times (n+3)$ matrix, which we can view as an element of the Grassmannian $\text{Gr}_{2, n+3}$. Under this identification, matrix entries of the $2 \times (n+1)$ matrix are equal to Plücker coordinates of the corresponding element of $\text{Gr}_{2, n+3}$: $a_i = P_{i+1, n+3}$ and $b_i = P_{1, i+1}$. Note also that $P_{1, n+3} = 1$. We can thus identify \mathcal{A} with the quotient $R_{2, n+3} / \langle P_{1, n+3} - 1 \rangle$ of the Plücker ring $R_{2, n+3}$. Our cluster variables and frozen variables for \mathcal{A} are inherited from the cluster structure on $R_{2, n+3}$. The cluster algebra \mathcal{A} has rank n , with $n+2$ frozen variables. In the case $n = 1$ we recover Example 1.1.2.

On the other hand, subdividing a $2 \times (n+1)$ matrix (6.1.3) into a $2 \times i$ matrix and a $2 \times (n+1-i)$ matrix, we can make \mathcal{A} into a cluster algebra of type $A_{i-1} \times A_{n-i}$. The cluster and frozen variables would include all matrix entries as well as every 2×2 minor contained in one of the two distinguished submatrices. The total number of frozen variables in this cluster algebra is $(i+1) + (n+2-i) = n+3$. For each i , this gives a cluster algebra structure of rank $(i-1) + (n-i) = n-1$, with $n+3$ frozen variables.

More generally, we can partition a $2 \times (n + 1)$ matrix (6.1.3) into k matrices of sizes $2 \times i_1, \dots, 2 \times i_k$, where $i_1 + \dots + i_k = n + 1$. The cluster variables and frozen variables would include all matrix entries plus every 2×2 minor contained in one of the k distinguished submatrices. This gives a cluster structure of type $A_{i_1-1} \times \dots \times A_{i_k-1}$ on the ring \mathcal{A} . This cluster algebra has rank $(i_1 - 1) + (i_2 - 1) + \dots + (i_k - 1) = n - k + 1$, and has $(i_1 + 1) + (i_2 + 1) + \dots + (i_k + 1) = n + k + 1$ frozen variables.

6.2. Cluster algebras and coordinate rings

Suppose a collection of regular functions on an algebraic variety X satisfies relations which can be interpreted as exchange relations for a seed pattern. Then—subject to conditions articulated below—the coordinate ring of X can be naturally identified with the corresponding cluster algebra:

Proposition 6.2.1. *Let \mathcal{A} be a cluster algebra (of geometric type, over \mathbb{C}) of rank n , with frozen variables x_{n+1}, \dots, x_m . Let \mathcal{X} denote the set of cluster variables in \mathcal{A} . Let X be a rational affine irreducible algebraic variety of dimension m . Suppose we are given a family of nonzero regular functions*

$$\{\varphi_z : z \in \mathcal{X}\} \cup \{\varphi_{n+1}, \dots, \varphi_m\} \subset \mathbb{C}[X]$$

satisfying the following conditions:

(6.2.1) the functions φ_z ($z \in \mathcal{X}$) and φ_i ($n + 1 \leq i \leq m$) generate $\mathbb{C}[X]$;

(6.2.2) replacing each cluster variable z by φ_z , and each frozen variable x_i by φ_i makes every exchange relation (3.1.1) into an identity in $\mathbb{C}[X]$.

Then there is a unique \mathbb{C} -algebra isomorphism $\varphi : \mathcal{A} \rightarrow \mathbb{C}[X]$ such that $\varphi(z) = \varphi_z$ for all $z \in \mathcal{X}$ and $\varphi(x_i) = \varphi_i$ for $i \in \{n + 1, \dots, m\}$.

Remark 6.2.2. We briefly comment on the general assumptions on the variety X made above. Irreducibility implies that the ring of regular functions $\mathbb{C}[X]$ is a domain, so its fraction field is well defined (and coincides with the field $\mathbb{C}(X)$ of rational functions on X). Rationality of X means that $\mathbb{C}(X)$ is isomorphic to the field of rational functions over \mathbb{C} in $\dim(X)$ independent variables. In a typical application, X contains an open subset isomorphic to an affine space, so this condition is satisfied.

Proof. The key assertion to be proved is that each cluster in \mathcal{A} gives rise to a transcendence basis of the field of rational functions $\mathbb{C}(X)$. Pick a seed in \mathcal{A} ; let \mathbf{x} (resp., $\tilde{\mathbf{x}}$) be the corresponding cluster (resp., extended cluster). Every cluster variable $z \in \mathcal{X}$ is expressed as a rational function in $\tilde{\mathbf{x}}$ by iterating the exchange relations away from the chosen seed. By (6.2.2), we

can apply the same procedure to express all functions φ_z and φ_i inside the field $\mathbb{C}(X)$ as rational functions in the set

$$\Phi = \{\varphi_x : x \in \mathbf{x}\} \cup \{\varphi_{n+1}, \dots, \varphi_m\}.$$

Since X is rational and $|\Phi| = m = \dim(X)$, we conclude from (6.2.1) that Φ is a transcendence basis of the field of rational functions $\mathbb{C}(X)$, and that the correspondence

$$z \mapsto \varphi_z \quad (z \in \mathcal{X}), \quad x_i \mapsto \varphi_i \quad (n < i \leq m)$$

extends uniquely to an isomorphism of fields $\mathcal{F} \rightarrow \mathbb{C}(X)$, and hence yields an isomorphism of algebras $\mathcal{A} \rightarrow \mathbb{C}[X]$. \square

6.3. Examples of cluster structures of classical types

Informally speaking, Proposition 6.2.1 tells us that in order to identify a coordinate ring of a rational algebraic variety as a cluster algebra, it suffices to find elements of that ring that satisfy the requisite exchange relations. In reality, this approach is only practical for cluster algebras of finite type. In this section, we present four examples of coordinate rings endowed with cluster structures of types A_n , B_n , C_n , and D_n , respectively. All four rings are closely related to each other; the first two of them are actually identical (as commutative rings) even though the cluster structures are different.

Our first example, the homogeneous coordinate ring of a Grassmannian of 2-planes, has already been thoroughly examined in Sections 1.2 and 5.3.

Example 6.3.1 (*Type A_n*). Let $X = \widehat{\text{Gr}}_{2,n+3}$ be the affine cone over the Grassmannian $\text{Gr}_{2,n+3}$ of 2-dimensional subspaces in \mathbb{C}^{n+3} taken in its Plücker embedding. Equivalently, X can be viewed as the variety of nonzero decomposable bivectors:

$$X \cong \{u \wedge v \neq 0 \mid u, v \in \mathbb{C}^{n+3}\}.$$

This is a $(2n+3)$ -dimensional affine algebraic variety. Its coordinate ring is the Plücker ring $R_{2,n+3} = \mathbb{C}[X]$. This ring is generated by the standard Plücker coordinates $P_{ab} \in \mathbb{C}[X]$, for $1 \leq a < b \leq n+3$.

Alternatively, we can view the Plücker ring $R_{2,n+3}$ as the ring of SL_2 -invariant polynomial functions on the space of $(n+3)$ -tuples of vectors in \mathbb{C}^2 . Representing these vectors as columns of a $2 \times (n+3)$ matrix $z = (z_{ab})$, one identifies the Plücker coordinates with the 2×2 minors of z :

$$P_{ab} = z_{1a}z_{2b} - z_{1b}z_{2a} \quad (1 \leq a < b \leq n+3).$$

In Section 5.3, we constructed a seed pattern of type A_n in the field of rational functions $\mathbb{C}(X)$. The seeds in this pattern are labeled by triangulations of a convex $(n+3)$ -gon \mathbf{P}_{n+3} by pairwise noncrossing diagonals.

Each cluster consists of the Plücker coordinates P_{ij} corresponding to the diagonals in a given triangulation. The frozen variables are the Plücker coordinates associated with the sides of \mathbf{P}_{n+3} . The exchange relations of the seed pattern are exactly the Grassmann-Plücker relations (1.2.1). Thus, we can view this example as an instance of Proposition 6.2.1.

As a cluster algebra of type A_n , the Plücker ring $R_{2,n+3}$ is generated by the cluster and coefficient variables, which are precisely the Plücker coordinates P_{ij} . It is moreover well known (and not hard to see) that the ideal of relations among the Plücker coordinates is generated by the exchange relations, or more precisely by the polynomials

$$P_{ik}P_{jl} - P_{ij}P_{kl} - P_{il}P_{jk} \quad (1 \leq i < j < k < l \leq n+3)$$

(cf. (1.2.1)). We will see in Section 6.8 that this phenomenon does not hold in general: even when a cluster algebra is of finite type, some of the relations among its generators may not lie in the ideal generated by the exchange relations.

While the type A cluster structure on a Plücker ring $R_{2,m}$ is perhaps the most natural one, we can also endow this ring with a type B cluster structure, as we now explain.

Example 6.3.2 (*Type B_n*). The two-element group $\mathbb{Z}/2\mathbb{Z}$ acts on the set of tagged arcs and boundary segments in the punctured polygon \mathbf{P}_{n+1}^\bullet (see Definition 5.4.3) by switching the tagging on radii, and leaving everything else intact. Let us associate an element P_γ of the Plücker ring $R_{2,n+2}$ to every $\mathbb{Z}/2\mathbb{Z}$ -orbit γ as follows (cf. Definition 5.4.9):

$$P_\gamma = \begin{cases} P_{ab} & \text{if } \gamma \text{ doesn't cross the cut, and has endpoints } a \text{ and } b > a; \\ P_{a\bar{b}} & \text{if } \gamma \text{ crosses the cut, and has endpoints } a \text{ and } b > a; \\ P_{a,n+2} & \text{if } \gamma \text{ is an orbit of radii with endpoints } p \text{ and } a, \end{cases}$$

where we use the notation

$$(6.3.1) \quad P_{a\bar{b}} = P_{a,n+2}P_{b,n+2} - P_{ab}.$$

The cluster variables and frozen variables are the elements P_γ , where γ ranges over orbits of tagged arcs and boundary segments, respectively.

We use Proposition 6.2.1 to show that this yields a cluster structure of type B_n in $R_{2,n+2}$. The only nontrivial task is to check condition (6.2.2),

which amounts to verifying the following six identities:

$$(6.3.2) \quad P_{ac} P_{bd} = P_{ab} P_{cd} + P_{ad} P_{bc} \quad (1 \leq a < b < c < d \leq n+1),$$

$$(6.3.3) \quad P_{a\bar{c}} P_{bd} = P_{a\bar{b}} P_{cd} + P_{a\bar{d}} P_{bc} \quad (1 \leq a < b < c < d \leq n+1),$$

$$(6.3.4) \quad P_{a\bar{c}} P_{b\bar{d}} = P_{ab} P_{cd} + P_{a\bar{d}} P_{b\bar{c}} \quad (1 \leq a < b < c < d \leq n+1),$$

$$(6.3.5) \quad P_{ac} P_{a\bar{b}} = P_{ab} P_{a\bar{c}} + P_{a,n+2}^2 P_{bc} \quad (1 \leq a < b < c \leq n+1),$$

$$(6.3.6) \quad P_{a\bar{b}} P_{b\bar{c}} = P_{ab} P_{bc} + P_{b,n+2}^2 P_{a\bar{c}} \quad (1 \leq a < b < c \leq n+1),$$

$$(6.3.7) \quad P_{a,n+2} P_{b,n+2} = P_{ab} + P_{a\bar{b}} \quad (1 \leq a < b \leq n+1).$$

While these identities can be directly deduced from the Grassmann-Plücker relations, we prefer another route, presented below in a slightly informal way.

Consider the algebraic model of a seed pattern of type D_{n+1} described in Section 5.4. (Note that we are replacing n by $n+1$.) Recall that it involves working with $n+1$ two-dimensional vectors v_1, \dots, v_{n+1} , two “special” vectors a and \bar{a} , and two scalars λ and $\bar{\lambda}$. To get a seed pattern of type B_n , we specialize the type D_{n+1} seed pattern as follows. The vectors v_1, \dots, v_{n+1} are kept without change. We take vectors a and b satisfying

$$(6.3.8) \quad \langle b, a \rangle = 1$$

(we shall later identify a with v_{n+2}), set

$$\bar{a} = a + \varepsilon b,$$

$$\lambda = 1,$$

$$\bar{\lambda} = 1 + \varepsilon,$$

and take the limit $\varepsilon \rightarrow 0$. We then have

$$a^{\boxtimes} = \frac{\bar{\lambda} - \lambda}{\langle \bar{a}, a \rangle} \bar{a} = \frac{\varepsilon}{\langle a + \varepsilon b, a \rangle} (a + \varepsilon b) \rightarrow a,$$

so in the limit we get $a^{\boxtimes} = a$ and $\bar{\lambda} = \lambda$. This yields the folding of our type D_{n+1} seed pattern into a type B_n pattern. It remains to check that the (folded) cluster variables of the type D_{n+1} pattern specialize to the cluster variables P_γ defined above. The only nontrivial case is the second one, wherein $P_\gamma = P_{a\bar{b}}$. The operator A defined in (5.4.1) specializes via

$$Av = \frac{\bar{\lambda} \langle v, a \rangle \bar{a} - \lambda \langle v, \bar{a} \rangle a}{\langle \bar{a}, a \rangle} \rightarrow \langle v, a \rangle a - \langle v, b \rangle a + \langle v, a \rangle b.$$

As a result, we get, using the Grassmann-Plücker relation and (6.3.8):

$$\langle w, Av \rangle \rightarrow \langle v, a \rangle \langle w, a \rangle - \langle v, b \rangle \langle w, a \rangle + \langle v, a \rangle \langle w, b \rangle = \langle v, a \rangle \langle w, a \rangle - \langle v, w \rangle,$$

So in particular $\langle v_j, Av_i \rangle \rightarrow \langle v_i, a \rangle \langle v_j, a \rangle - \langle v_i, v_j \rangle$, matching (6.3.1).

Remark 6.3.3. Examples 6.1.2 and 6.3.1-6.3.2 demonstrate that a given ring can carry different non-isomorphic cluster structures. A particularly striking example was given in [12, Figure 20]: the “mixed Plücker ring” $\mathbb{C}[V^3 \times (V^*)^4]^{\mathrm{SL}(V)}$, with $V \cong \mathbb{C}^3$, can carry a cluster structure of finite type D_6 or E_6 , or a cluster structure of an infinite mutation type (hence of infinite type).

Example 6.3.4 (*Type C_n*). Let SO_2 be the group of complex matrices

$$\begin{bmatrix} u & -v \\ v & u \end{bmatrix}$$

with $u^2 + v^2 = 1$. Consider the ring $R = \mathbb{C}[V^{n+1}]^{\mathrm{SO}_2}$ of SO_2 -invariant polynomial functions on the space of $(n+1)$ -tuples of vectors

$$(6.3.9) \quad (v_1, \dots, v_{n+1}) \in V^{n+1}, \quad V = \mathbb{C}^2,$$

or equivalently SO_2 -invariant polynomials in the entries of a $2 \times (n+1)$ matrix

$$z = \begin{bmatrix} z_{11} & \cdots & z_{1,n+1} \\ z_{21} & \cdots & z_{2,n+1} \end{bmatrix}.$$

This ring is generated by the Plücker coordinates

$$P_{ab} = \langle v_a, v_b \rangle = z_{1a}z_{2b} - z_{1b}z_{2a} \quad (1 \leq a < b \leq n+1)$$

together with the polynomials

$$P_{a\bar{b}} = \langle v_b, Mv_a \rangle = z_{1a}z_{1b} + z_{2a}z_{2b} \quad (1 \leq a \leq b \leq n+1),$$

where $M = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \in \mathrm{SO}_2$. The ring $R = \mathbb{C}[V^{n+1}]^{\mathrm{SO}_2}$ can also be viewed as the coordinate ring $\mathbb{C}[X]$ of the variety X of complex $(n+1) \times (n+1)$ matrices of rank ≤ 1 ; even more geometrically, X is the affine cone over the product of two copies of the projective space $\mathbb{C}\mathbb{P}^n$ taken in the Segre embedding. Specifically, the map

$$z = \begin{bmatrix} z_{11} & \cdots & z_{1,n+1} \\ z_{21} & \cdots & z_{2,n+1} \end{bmatrix} \mapsto ((z_{1a} - iz_{2a})(z_{1b} + iz_{2b}))_{a,b \in \{1, \dots, n+1\}} \in X$$

induces an algebra isomorphism $\mathbb{C}[X] \rightarrow \mathbb{C}[V^{n+1}]^{\mathrm{SO}_2}$. (Note that

$$(z_{1a} - iz_{2a})(z_{1b} + iz_{2b}) = P_{a\bar{b}} + iP_{ab}.)$$

To construct a cluster algebra structure of type C_n in this ring, let us associate an element $P_\gamma \in R$ to every orbit γ of the action of $\mathbb{Z}/2\mathbb{Z}$ on the set of diagonals and sides of a regular $(2n+2)$ -gon \mathbf{P}_{2n+2} . Specifically, we set

$$P_\gamma = \begin{cases} P_{ab} & \text{if } \gamma = \{(a, b), (a+n+1, b+n+1)\} \text{ for } a < b \leq n+1; \\ P_{a\bar{b}} & \text{if } \gamma = \{(a, b+n+1), (a+n+1, b)\} \text{ for } a \leq b \leq n+1. \end{cases}$$

where we use the notation (a, b) to denote a diagonal or side with endpoints a and b . The cluster variables and frozen variables are the elements P_γ , where γ ranges over orbits of diagonals and boundary segments, respectively.

The verification is similar to Example 6.3.2 above. The only substantive task is to check that the functions P_{ab} and $P_{a\bar{b}}$ satisfy the requisite exchange relations:

$$(6.3.10) \quad P_{ac} P_{bd} = P_{ab} P_{cd} + P_{ad} P_{bc} \quad (1 \leq a < b < c < d \leq n+1),$$

$$(6.3.11) \quad P_{a\bar{c}} P_{bd} = P_{a\bar{b}} P_{cd} + P_{a\bar{d}} P_{bc} \quad (1 \leq a < b < c < d \leq n+1),$$

$$(6.3.12) \quad P_{a\bar{c}} P_{b\bar{d}} = P_{ab} P_{cd} + P_{a\bar{d}} P_{b\bar{c}} \quad (1 \leq a < b < c < d \leq n+1),$$

$$(6.3.13) \quad P_{ac} P_{a\bar{b}} = P_{ab} P_{a\bar{c}} + P_{a\bar{a}} P_{bc} \quad (1 \leq a < b < c \leq n+1),$$

$$(6.3.14) \quad P_{a\bar{b}} P_{b\bar{c}} = P_{ab} P_{bc} + P_{b\bar{b}} P_{a\bar{c}} \quad (1 \leq a < b < c \leq n+1),$$

$$(6.3.15) \quad P_{a\bar{a}} P_{b\bar{b}} = P_{ab}^2 + P_{a\bar{b}}^2 \quad (1 \leq a < b \leq n+1).$$

This can either be done directly or via folding, this time going from a cluster structure of type A_{2n-1} in the Plücker ring $R_{2,2n+2}$ to the type C_n cluster structure in $\mathbb{C}[X]$, as follows. Starting with the $(n+1)$ -tuple (6.3.9), we build the $(2n+2)$ -tuple (v_1, \dots, v_{2n+2}) by setting $v_{n+1+a} = Mv_a$ for $a \in \{1, \dots, n+1\}$. In this specialization, using the fact that $M^2 = -1$, we have

$$\begin{aligned} \langle v_{n+1+a}, v_{n+1+b} \rangle &= \langle Mv_a, Mv_b \rangle = \langle v_a, v_b \rangle, \\ \langle v_b, v_{n+1+a} \rangle &= \langle v_b, Mv_a \rangle = \langle Mv_b, -v_a \rangle = \langle v_a, v_{n+1+b} \rangle. \end{aligned}$$

This shows that two Plücker coordinates corresponding to centrally symmetric diagonals in \mathbf{P}_{2n+2} are equal when evaluated at the $(2n+2)$ -tuple (v_1, \dots, v_{2n+2}) . Thus, the elements $P_\gamma \in \mathbb{C}[X]$ defined earlier come from Plücker coordinates in $R_{2,2n+2}$ via folding. The exchange relations in question are obtained by specializing the exchange relations in the Plücker ring.

Example 6.3.5 (*Type D_n*). Let $\widehat{\text{Gr}}_{2,n+2}$ be the affine cone over the Grassmannian $\text{Gr}_{2,n+2}$, taken in its Plücker embedding. Let X be the ‘‘Schubert’’ divisor in $\widehat{\text{Gr}}_{2,n+2}$ given by the equation $P_{n+1,n+2} = 0$; thus, we have

$$\mathbb{C}[X] = \mathbb{C}[\widehat{\text{Gr}}_{2,n+2}] / \langle P_{n+1,n+2} \rangle.$$

A cluster structure of type D_n in the coordinate ring $R = \mathbb{C}[X]$ can be obtained by associating an element $P_\gamma \in R$ to each tagged arc γ in the punctured polygon \mathbf{P}_n^\bullet , as follows (cf. Definition 5.4.9):

$$P_\gamma = \begin{cases} P_{ab} & \text{if } \gamma \text{ doesn't cross the cut} \\ & \text{and has endpoints } a \text{ and } b > a; \\ P_{a,n+1}P_{b,n+2} - P_{ab} & \text{if } \gamma \text{ crosses the cut} \\ & \text{and has endpoints } a \text{ and } b > a; \\ P_{a,n+1} & \text{if } \gamma \text{ is a plain radius with endpoint } a; \\ P_{a,n+2} & \text{if } \gamma \text{ is a notched radius with endpoint } a. \end{cases}$$

(We thus identify the two eigenvectors of (5.4.1) with v_{n+1} and v_{n+2} , respectively.) The verification is left to the reader; or see [16, Example 12.15].

Remark 6.3.6. While the seed pattern of type D_n described in Example 6.3.5 is much simpler than the one used in Section 5.4, we did not use it there because the corresponding exchange matrices do not have full rank.

6.4. Starfish lemma

In what follows, we denote by $\mathcal{A}(\tilde{\mathbf{x}}, \tilde{B})$ the cluster algebra defined by a seed $(\tilde{\mathbf{x}}, \tilde{B})$ in some ambient field of rational functions freely generated by $\tilde{\mathbf{x}}$.

Any cluster algebra, being a subring of a field, is an integral domain (and under our conventions, a \mathbb{C} -algebra). Conversely, given such a domain, one may be interested in identifying it as a cluster algebra.

For the remainder of this section, we let \mathcal{R} be an integral domain and a \mathbb{C} -algebra, and we denote by \mathcal{F} the quotient field of \mathcal{R} . The challenge is to find a seed $(\tilde{\mathbf{x}}, \tilde{B})$ in \mathcal{F} such that $\mathcal{A}(\tilde{\mathbf{x}}, \tilde{B}) = \mathcal{R}$. The difficulties here are two-fold. To prove the inclusion $\mathcal{A}(\tilde{\mathbf{x}}, \tilde{B}) \supset \mathcal{R}$, we need to demonstrate that (a subset of) cluster variables in this seed pattern, together with the frozen variables, generates \mathcal{R} . To prove the reverse inclusion $\mathcal{A}(\tilde{\mathbf{x}}, \tilde{B}) \subset \mathcal{R}$, we need to show that each cluster variable in the seed pattern generated by $(\tilde{\mathbf{x}}, \tilde{B})$ is an element of \mathcal{R} rather than merely a rational function in \mathcal{F} . In this section, we give sufficient conditions that guarantee the latter inclusion.

Recall that \mathcal{R} is *normal* if it is integrally closed in \mathcal{F} . This property is in particular satisfied if \mathcal{R} is *factorial* (or a *unique factorization domain*).

Recall that \mathcal{R} is called *Noetherian* if any ascending chain of ideals stabilizes. This is in particular satisfied if \mathcal{R} is *finitely generated* (over \mathbb{C}).

All rings of interest to us will be factorial and finitely generated, hence normal and Noetherian.

Let us call two elements $r, r' \in \mathcal{R}$ *coprime* if they are not contained in the same prime ideal of height 1. If \mathcal{R} is factorial, then such ideals are principal, and one recovers the usual definition of coprimality (r and r' are coprime if $\gcd(r, r')$ is a unit).

Proposition 6.4.1 (“Starfish lemma”). *Let \mathcal{R} be a \mathbb{C} -algebra and a normal Noetherian domain. Let $(\tilde{\mathbf{x}}, \tilde{B})$ be a seed of rank n in the fraction field \mathcal{F} with $\tilde{\mathbf{x}} = (x_1, \dots, x_m)$ for $n \leq m$ such that*

- (1) *all elements of $\tilde{\mathbf{x}}$ belong to \mathcal{R} ;*
- (2) *the cluster variables in $\tilde{\mathbf{x}}$ are pairwise coprime;*
- (3) *for each cluster variable $x_k \in \tilde{\mathbf{x}}$, the seed mutation μ_k replaces x_k with an element x'_k (cf. (3.1.1)) that lies in \mathcal{R} and is coprime to x_k .*

Then $\mathcal{A}(\tilde{\mathbf{x}}, \tilde{B}) \subset \mathcal{R}$.

We will give two proofs of Proposition 6.4.1, one using commutative algebra, and one using algebraic geometry. The commutative algebra proof relies on two lemmas.

For P a prime ideal in \mathcal{R} , let $\mathcal{R}_P = \mathcal{R}[(\mathcal{R} \setminus P)^{-1}]$ denote the localization of \mathcal{R} at $\mathcal{R} \setminus P$.

Lemma 6.4.2 ([32, Theorem 11.5]). *For a normal Noetherian domain \mathcal{R} , the natural inclusion $\mathcal{R} \subset \bigcap_{\text{ht } P=1} \mathcal{R}_P$ (intersection over prime ideals P of height 1) is an equality.*

Lemma 6.4.3. *Let P be a height 1 prime ideal in \mathcal{R} . Then at least one of the $n + 1$ products*

$$(x_1 \cdots x_n), (x'_1 x_2 \cdots x_n), \dots, (x_1 \cdots x_{n-1} x'_n)$$

does not belong to P .

Proof. Suppose that $x_1 \cdots x_n \in P$. Since P is prime, we have $x_k \in P$ for some $k \leq n$. Since $\text{ht } P = 1$, the coprimality assumption (6.4.1) implies that $x_j \notin P$ for $j \in \{1, \dots, n\} - \{k\}$. Similarly, (6.4.1) implies that $x'_k \notin P$. Again using that P is prime, we conclude that $x_1 \cdots x'_k \cdots x_n \notin P$. \square

Algebraic proof of Proposition 6.4.1. We need to prove that each cluster variable z from any seed mutation equivalent to $(\tilde{\mathbf{x}}, \tilde{B})$ belongs to \mathcal{R} . By Lemma 6.4.2, it suffices to show that $z \in \mathcal{R}_P$ for each prime ideal P of height 1. By Lemma 6.4.3, for any height 1 prime P in \mathcal{R} , there exists a cluster \mathbf{x}' such that $\prod_{x \in \mathbf{x}'} x \in \mathcal{R} \setminus P$. By the Laurent Phenomenon (Theorems 3.3.1 and 3.3.6), the cluster variable z can be expressed as a Laurent polynomial in the elements of \mathbf{x}' , with coefficients in $\mathbb{C}[x_{n+1}, \dots, x_m]$. Thus $z \in \mathcal{R}[(\mathcal{R} \setminus P)^{-1}] = \mathcal{R}_P$, as desired. \square

Geometric proof of Proposition 6.4.1. Our assumptions on the ring \mathcal{R} mean that it can be identified with the coordinate ring of an (irreducible) normal affine complex algebraic variety $X = \text{Spec}(\mathcal{R})$. Then the field of fractions of \mathcal{R} is $\text{Frac}(\mathcal{R}) = \mathbb{C}(X)$, the field of rational functions on X . We need to show that each cluster variable z from any seed mutation equivalent to $(\tilde{\mathbf{x}}, \tilde{B})$ belongs to \mathcal{R} . The key property that we need is the algebraic version of *Hartogs' continuation principle* for normal varieties (see, e.g., [7, Chapter 2, 7.1]) which asserts that a function on X that is regular outside a closed algebraic subset of codimension ≥ 2 is in fact regular everywhere on X .

Consider the subvariety

$$Y = \bigcup_{1 \leq i < j \leq n} \{x_i = x_j = 0\} \cup \bigcup_{1 \leq k \leq n} \{x_k = x'_k = 0\} \subset X.$$

The coprimeness conditions imposed on $(\tilde{\mathbf{x}}, \tilde{B})$ imply that $\text{codim}(Y) \geq 2$. By the algebraic Hartogs' principle mentioned above, it now suffices to show that z is regular on $X \setminus Y$.

The complement $X \setminus Y$ consists of the points $x \in X$ such that

- at most one of the cluster variables in $\tilde{\mathbf{x}}$ vanishes at x , and
- for each pair (x_k, x'_k) as above, either x_k or x'_k does not vanish at x .

Hence there is a seed $(Q', \tilde{\mathbf{x}}')$ (either the original seed $(\tilde{\mathbf{x}}, \tilde{B})$ or one of the adjacent seeds $\mu_k(\tilde{\mathbf{x}}, \tilde{B})$) none of whose cluster variables vanishes at x ; moreover $\tilde{\mathbf{x}}' \subset \mathbb{C}[X]$. Then the Laurent Phenomenon (Theorems 3.3.1 and 3.3.6) implies that our distant cluster variable z is regular at x , as desired. \square

Remark 6.4.4. The arguments given above actually establish a stronger statement: under the conditions of Proposition 6.4.1, the ring \mathcal{R} contains the *upper cluster algebra* associated with $\mathcal{A}(\tilde{\mathbf{x}}, \tilde{B})$ (see [3]), or more precisely the subalgebra of \mathcal{F} consisting of the elements which, when expressed in terms of any extended cluster, are Laurent polynomials in the cluster variables and ordinary polynomials in the coefficient variables.

Remark 6.4.5. The versions of the Starfish Lemma and the Laurent phenomenon given in [3] (implicit) and [26] are predicated on invertibility of coefficient variables (that is, the ground ring is the ring of Laurent polynomials in the coefficient variables) and concern the upper cluster algebra.

Corollary 6.4.6. *Let \mathcal{R} be a finitely generated factorial \mathbb{C} -algebra. Let $(\tilde{\mathbf{x}}, \tilde{B})$ be a seed in the quotient field of \mathcal{R} such that all cluster variables of $\tilde{\mathbf{x}}$ and all elements of clusters adjacent to $\tilde{\mathbf{x}}$ are irreducible elements of \mathcal{R} . Then $\mathcal{A}(\tilde{\mathbf{x}}, \tilde{B}) \subset \mathcal{R}$.*

Proof. The only conditions in Proposition 6.4.1 that we need to check are the ones concerning coprimality. Two elements of $\tilde{\mathbf{x}}$ cannot differ by a scalar factor since they are algebraically independent. Similarly, if x_k and x'_k were to differ by a scalar factor, then the exchange relation (3.1.1) would give an algebraic dependence in $\tilde{\mathbf{x}}$. \square

Suppose that a \mathbb{C} -algebra \mathcal{R} satisfies the conditions in the first sentence of Proposition 6.4.1 (or Corollary 6.4.6). In order to identify a cluster structure in \mathcal{R} , it suffices to exhibit a seed $(\tilde{\mathbf{x}}, \tilde{B})$ such that

- (i) all cluster variables of $\tilde{\mathbf{x}}$ and of the clusters adjacent to $\tilde{\mathbf{x}}$ are irreducible elements of \mathcal{R} ;
- (ii) the seed pattern generated by $(\tilde{\mathbf{x}}, \tilde{B})$ contains a generating set for \mathcal{R} .

This is however easier said than done.

Regarding condition (ii) above, let us make the following simple observation.

Proposition 6.4.7. *Let \mathcal{R} be a cluster algebra that is finitely generated (over \mathbb{C}) as a \mathbb{C} -algebra. Then \mathcal{R} is generated by a finite subset of cluster and coefficient variables.*

Proof. Let S be a finite generating set for \mathcal{R} , and let \mathcal{X} be the set of cluster and coefficient variables. Since each $s \in S$ can be written as a polynomial in the elements of a finite subset $\mathcal{X}_s \subset \mathcal{X}$, we conclude that the finite set $\bigcup_{s \in S} \mathcal{X}_s \subset \mathcal{X}$ generates \mathcal{R} . \square

We next review some general algebraic criteria that can be used to check that a given \mathbb{C} -algebra \mathcal{R} satisfies the conditions in Proposition 6.4.1 or Corollary 6.4.6.

The fact that \mathcal{R} is a domain will usually be immediate, e.g. when \mathcal{R} is given as a subring of a polynomial ring.

Perhaps the most famous result concerning finite generation is (the modern version of) Hilbert's Theorem, see, e.g., [35, Theorem 3.5]:

Theorem 6.4.8. *Let G be a reductive algebraic group acting on an affine algebraic variety X . Then the ring of invariants $\mathbb{C}[X]^G$ is finitely generated.*

For the purposes of applying Proposition 6.4.1, the following version is particularly useful.

Theorem 6.4.9 ([8, Proposition 3.1]). *Let G be a reductive algebraic group acting algebraically on a normal finitely generated \mathbb{C} -algebra A . Then A^G is a normal finitely generated \mathbb{C} -algebra.*

Even when a group is not reductive, the ring of its invariants may be finitely generated. The most important case to us is the following.

Theorem 6.4.10 ([8, Theorem 5.4]). *Let G be a reductive group acting rationally on a finitely generated \mathbb{C} -algebra A . Let U be a maximal unipotent group of G . Then the subalgebra A^U of U -invariant elements in A is finitely generated over \mathbb{C} .*

We conclude this section with a couple of factoriality criteria, see [35, Theorem 3.17], [42] and references therein.

Proposition 6.4.11. *Let G be a connected, simply connected semisimple complex Lie group. Then the ring of regular functions $\mathbb{C}[G]$ is factorial.*

Theorem 6.4.12. *Let G be a connected algebraic group acting on an affine algebraic variety X . If G has no nontrivial characters and $\mathbb{C}[X]$ is factorial, then so is $\mathbb{C}[X]^G$.*

Remark 6.4.13. The paper [23] provides factoriality criteria for cluster algebras. It also shows that a cluster algebra contains no nontrivial units, and all cluster variables are irreducible elements.

6.5. Cluster structure in the ring $\mathbb{C}[\mathrm{SL}_k]^U$

In this section, we identify a cluster algebra structure in the ring $\mathbb{C}[\mathrm{SL}_k]^U$, the coordinate ring of the basic affine space for the special linear group. This ring made its first appearance in Section 1.3. We start by reviewing the key features of this construction.

Let $V \cong \mathbb{C}^k$ be a k -dimensional complex vector space. After choosing a basis in V , we can identify SL_k with the special linear group $G = \mathrm{SL}(V)$ of complex matrices with determinant 1. The subgroup $U \subset G$ of unipotent lower-triangular matrices acts on G by left multiplication. This action induces the action of U on the coordinate ring $\mathbb{C}[G]$. We will show that the ring $\mathbb{C}[G]^U$ of U -invariant regular functions on G has a natural structure of a cluster algebra.

We note that U is not reductive, so Theorem 6.4.8 does not apply. Still, $\mathbb{C}[G]^U$ is finitely generated by Theorem 6.4.10. As mentioned in Section 1.3, this can be made explicit as follows. Recall that a *flag minor* P_J of a $k \times k$ matrix z (here $J \subset \{1, \dots, k\}$) is the determinant of the submatrix of z occupying the columns labeled by J and the rows labeled $1, 2, \dots, |J|$.

Theorem 6.5.1. *The ring of invariants $\mathbb{C}[\mathrm{SL}_k]^U$ is generated by the $2^k - 2$ flag minors P_J ; here J runs over nonempty proper subsets of $\{1, \dots, k\}$.*

The ideal of relations satisfied by the flag minors is generated by certain generalized Grassmann-Plücker relations (which we will not rely upon).

Theorem 6.5.1 is a consequence of the classical construction of irreducible representations of special linear groups, see, e.g., [18, 37]. This construction generalizes to an arbitrary connected, simply connected semisimple complex Lie group G and its maximal unipotent subgroup U . The role of flag minors is played by certain matrix elements in fundamental representations of G . See the end of this section for additional details.

Corollary 6.5.2. *The ring $\mathbb{C}[\mathrm{SL}_k]^U$ is factorial.*

Proof. This follows from Proposition 6.4.11 and Theorem 6.4.12. (The polynomial ring $\mathbb{C}[U]$ has no nontrivial units.) \square

Definition 6.5.3. Let D be a wiring diagram with k strands. Let \mathcal{F} be the field of rational functions in the chamber minors of D , cf. Section 1.3. We associate to D the pair $(\tilde{\mathbf{x}}(D), \tilde{B}(D))$, where

- $\tilde{\mathbf{x}}(D)$ consists of the chamber minors of D , listed so that the minors indexed by the bounded chambers precede the minors indexed by the unbounded ones;
- $\tilde{B}(D)$ is the signed adjacency matrix of the quiver $Q(D)$ from Definition 2.3.1.

Note that $\tilde{B}(D)$ is a $\frac{(k-1)(k+2)}{2} \times \frac{(k-1)(k-2)}{2}$ integer matrix whose rows are indexed by all chambers and whose columns are indexed by the bounded chambers. The minors corresponding to the bounded (resp., unbounded) chambers are the cluster variables (resp., frozen variables) of this seed.

Theorem 6.5.4. *Let $G = \mathrm{SL}_k(\mathbb{C})$.*

- (1) *For any wiring diagram with k strands, the pair $(\tilde{\mathbf{x}}(D), \tilde{B}(D))$ is a seed in the field of fractions for $\mathbb{C}[G]^U$.*
- (2) *All seeds $(\tilde{\mathbf{x}}(D), \tilde{B}(D))$ are mutation equivalent to each other.*
- (3) *The seed pattern containing the seeds $(\tilde{\mathbf{x}}(D), \tilde{B}(D))$ defines a cluster algebra structure in $\mathbb{C}[G]^U$. That is, $\mathcal{A}(\tilde{\mathbf{x}}(D), \tilde{B}(D)) = \mathbb{C}[G]^U$.*

Proof. Statement (2) is part of Exercise 3.1.3. We then conclude that any flag minor can be expressed as a rational function in the elements of a given extended cluster $\tilde{\mathbf{x}}(D)$. Since $|\tilde{\mathbf{x}}(D)| = \frac{(k-1)(k+2)}{2} = \dim(U \setminus G)$, it follows that the elements of $\tilde{\mathbf{x}}(D)$ are algebraically independent, proving statement (1).

It remains to prove statement (3). Since each flag minor appears in some extended cluster $\tilde{\mathbf{x}}(D)$, Theorem 6.5.1 implies that $\mathcal{A}(\tilde{\mathbf{x}}(D), \tilde{B}(D)) \supset \mathbb{C}[G]^U$.

We prove the inclusion $\mathcal{A}(\tilde{\mathbf{x}}(D), \tilde{B}(D)) \subset \mathbb{C}[G]^U$ using Proposition 6.4.1. Let us choose the seed associated to the wiring diagram D of the kind shown in Figure 6.2. The quiver is shown in Figure 6.3.

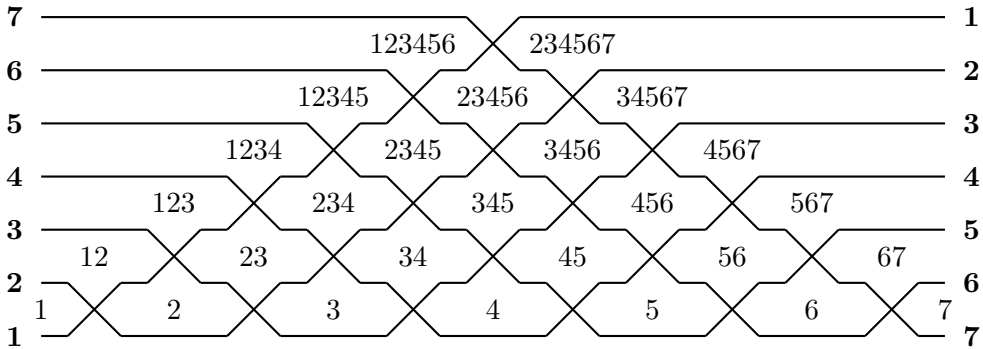


Figure 6.2. A special wiring diagram for $k = 7$, and its chamber minors.

All the elements of $\tilde{\mathbf{x}}(D)$ are flag minors, so they belong to $\mathbb{C}[G]^U$. Moreover they are irreducible polynomials, hence irreducible elements of $\mathbb{C}[G]^U$. This follows from the well-known fact that the determinant of a matrix of indeterminates is an irreducible polynomial, see [39, Theorem 3.2].

Let us compute the elements of the clusters adjacent to $\tilde{\mathbf{x}}(D)$. Note that the chamber minors of D are *solid*, i.e., they have column sets of the form

$$[a, d] = \{a, a + 1, \dots, d - 1, d\},$$




Figure 6.3. The seed corresponding to the wiring diagram in Figure 6.2.

for $1 \leq a \leq d \leq k$. Note that mutations at vertices in the bottom row of the quiver can be understood using the braid moves we studied earlier, *cf.* Figure 1.7. To understand mutations at the other vertices of the quiver, note that a typical chamber minor $P_{[b,c]} \in \tilde{\mathbf{x}}(D)$ is exchanged with the element $\Omega \in \text{Frac}(\mathbb{C}[G]^U)$ given by

$$(6.5.1) \quad \Omega = \frac{P_{Jabc} P_{Jcd} P_{Jb} + P_{Jab} P_{Jbcd} P_{Jc}}{P_{Jbc}},$$

where we used the shorthand

$$(6.5.2) \quad \begin{aligned} a &= b - 1, \\ d &= c + 1, \\ J &= [b + 1, c - 1], \\ Jbc &= J \cup \{b, c\} = [b, c], \end{aligned}$$

and similarly for Jb , Jc , etc. This can be seen by examining the quiver $Q(D)$ in the vicinity of the vertex associated with $P_{[b,c]} = P_{Jbc}$, see Figure 6.4.

We verify that $\Omega \in \mathbb{C}[G]^U$ by expressing Ω as a polynomial in flag minors:

Lemma 6.5.5. *The rational function Ω defined by (6.5.1) satisfies*

$$(6.5.3) \quad \Omega = -P_{Ja} P_{Jbcd} + P_{Jb} P_{Jacd}.$$

Proof. Follows from (5.3.2) by virtue of Muir's Law (Proposition 1.3.5). \square

It remains to prove that Ω is coprime to P_{Jbc} . Let f denote the specialization that sets the top $|J| + 2$ entries in column a to zero. Note that f leaves P_{Jbc} unchanged, and $f(P_{Ja}) = 0$. We know that P_{Jbc} is irreducible. If Ω were divisible by P_{Jbc} , then $f(\Omega) = \pm P_{Jb} \cdot z_{|J|+3,a} \cdot P_{Jcd}$ would also be divisible by P_{Jbc} . But $f(\Omega)$ is a product of three irreducible polynomials none of which is a scalar multiple of P_{Jbc} . \square

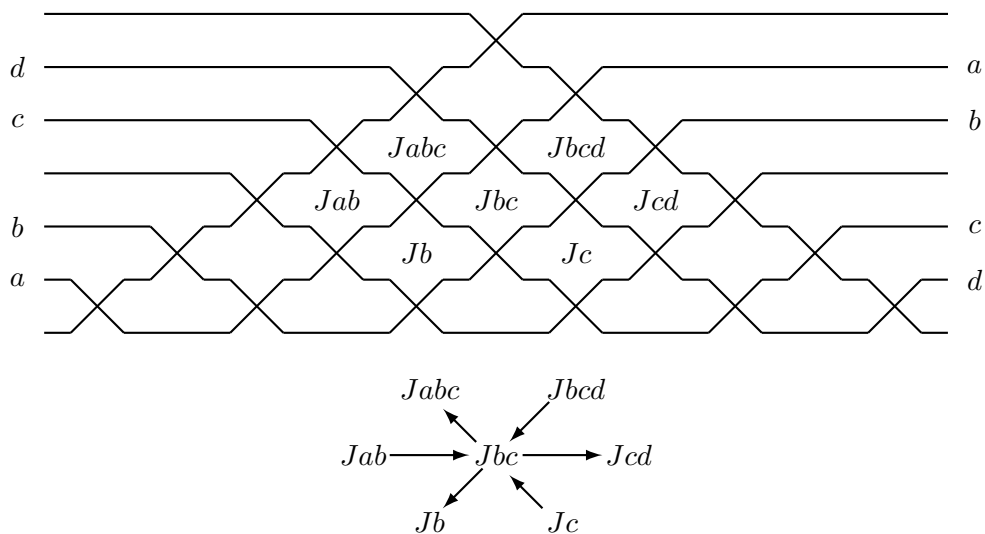


Figure 6.4. Chamber minors appearing in the exchange relation for a flag minor P_{Jbc} in a special wiring diagram, and part of the associated quiver. Only the arrows incident to the vertex Jbc are shown.

The special cases $G = \mathrm{SL}_3$ and $G = \mathrm{SL}_4$ of this construction have been presented in Examples 3.2.1 and 5.3.8, respectively. We now discuss the case $G = \mathrm{SL}_5$.

Example 6.5.6. Let $G = \mathrm{SL}_5$. In this case, the cluster structure in $\mathbb{C}[G]^U$ is of type D_6 , and accordingly has 36 cluster variables, cf. Figure 5.17. There are 8 coefficient variables, all of them flag minors.

To compute all cluster variables, one can use the following method, cf. [17, Proposition 11.1(1)]. Start with a seed coming from a wiring diagram. Apply mutations to obtain a seed whose quiver is a bipartite orientation of the Dynkin diagram type D_6 . (That is, each vertex is either a source or a sink.) Then repeatedly alternate between mutating at all sources and mutating at all sinks until all 36 cluster variables are computed. After each mutation, one needs to represent the new cluster variable as a regular function, not just a rational one. At the end of this process, one determines that each of the $22 = 2^5 - 2 - 8$ flag minors which is not a coefficient variable is a cluster variable. The remaining 14 cluster variables are described as follows. Define

$$\begin{aligned}
 g(a, b|c, d) &= -P_a P_{bcd} + P_b P_{acd}, \\
 g(a, b|c, d|J) &= -P_{Ja} P_{Jbcd} + P_{Jb} P_{Jacd}, \\
 h(a, b, c|d, e) &= -P_{ab} P_{cde} + P_{ac} P_{bde}, \\
 j(a, b|c, d, e) &= -P_a P_{bcde} + P_b P_{acde},
 \end{aligned}$$

where we use the shorthand $bcd = \{b, c, d\}$, $Ja = J \cup \{a\}$, etc. Note that $g(a, b|c, d|J)$ is precisely the regular function Ω from (6.5.1) and (6.5.3). With this notation, the 14 cluster variables in question turn out to be:

$$(6.5.4) \quad g(1, 2|3, 4), g(2, 3|4, 5), g(4, 5|1, 2), g(1, 3|4, 5), g(1, 2|3, 5),$$

$$(6.5.5) \quad g(1, 2|3, 4|5), g(2, 3|4, 5|1), g(4, 5|1, 2|3), g(1, 3|4, 5|2), g(1, 2|3, 5|4),$$

$$(6.5.6) \quad h(1, 2, 3|4, 5), h(5, 4, 3|2, 1),$$

$$(6.5.7) \quad j(1, 2|3, 4, 5), j(5, 4|3, 2, 1).$$

For $k \geq 6$, the cluster structure on $\mathbb{C}[\mathrm{SL}_k]^U$ that we described above is of infinite type. See Table 6.1.

Ring	Cluster type
$\mathbb{C}[\mathrm{SL}_3]^U$	A_1
$\mathbb{C}[\mathrm{SL}_4]^U$	A_3
$\mathbb{C}[\mathrm{SL}_5]^U$	D_6
$\mathbb{C}[\mathrm{SL}_k]^U$ for $k \geq 6$	infinite type

Table 6.1. The type of the standard cluster structure on $\mathbb{C}[\mathrm{SL}_k]^U$.

Remark 6.5.7. Let $U^+ \subset G = \mathrm{SL}_k$ be the subgroup of unipotent *upper-triangular* matrices. (Recall that we denoted by U the subgroup of unipotent *lower-triangular* matrices.) Let $\varphi : \mathbb{C}[G]^U \rightarrow \mathbb{C}[U^+]$ be the ring map defined by restricting U -invariant functions on G to the subgroup U^+ . Since every matrix entry of an element of U^+ can be written as a flag minor, φ is onto. The map φ can be used to transform a cluster structure on $\mathbb{C}[G]^U$ into a cluster structure on $\mathbb{C}[U^+]$. (This boils down to removing the coefficient variables corresponding to the leading principal minors $P_{1, \dots, j}$, as $P_{1, \dots, j}(u) = 1$ for $u \in U^+$.)

More generally, the coordinate ring of a maximal unipotent subgroup of any Kac-Moody group is a cluster algebra [21].

In this section we have so far discussed the basic affine space in the case of $G = \mathrm{SL}_k$. We now give a quick review of basic affine spaces and their significance for general semisimple Lie groups G . An excellent introduction is given in [5, Section 2.1]. Additional material can be found in [4, 27]. For general background on linear algebraic groups, see, e.g., [8, Section 3.3] or [34].

Let G be a simply connected semisimple complex algebraic group. Let U be a maximal unipotent subgroup of G . The variety $X = G/U$ is smooth and

quasi-affine (i.e., open in an affine variety). To be more specific, let $\mathcal{O}(X)$ denote the ring of regular functions on X . Then X embeds as an open subvariety into the affine (irreducible) variety $\overline{X} = \mathrm{Spec}(\mathcal{O}(X))$, the “affine completion” of X . The rings of regular functions on X and on \overline{X} coincide: $\mathcal{O}(X) = \mathcal{O}(\overline{X})$. Moreover these rings are naturally identified with the ring of invariants $\mathbb{C}[G]^U$. The variety \overline{X} is called the *basic affine space* for G . It is normal and usually singular.

In this section we proved that $\mathbb{C}[\mathrm{SL}_k]^U$ is a cluster algebra. More generally, for a simply-laced G , with U as above, there is a cluster algebra contained inside the coordinate ring $\mathbb{C}[G]^U$ [19]; it conjecturally coincides with the coordinate ring after a certain localization, see [19, Conjecture 10.4]. In this context, some *generalized minors* [14, Definition 1.4] (see also [31, Definition 6.2]) play the role of flag minors, and are used to define a collection of special seeds for the corresponding cluster algebra.

The significance of the basic affine space stems from the well known fact that its coordinate ring $\mathbb{C}[G]^U$ is a direct sum of all irreducible rational representations of G , each occurring with multiplicity 1. A more detailed description is as follows. The subgroup U is the unipotent radical of a Borel subgroup $B \subset G$. The action of the Cartan subgroup $H = B/U$ on X commutes with the natural G -action. The H -action induces a grading of $\mathbb{C}[G]^U$ by the weight lattice of G . The graded components are labeled by the dominant weights λ , and carry irreducible representations of G (of highest weight λ).

When $G = \mathrm{SL}_k$ is the special linear group, the irreducible representation with the highest weight $d_1\omega_1 + d_2\omega_2 + \cdots$ (here d_1, d_2, \dots are nonnegative integers, and $\omega_1, \omega_2, \dots$ are the fundamental weights of G in the standard order) appears as the space of polynomials in the flag minors which are homogeneous of degree d_1 in the flag minors P_1, P_2, P_3, \dots , of degree d_2 in $P_{12}, P_{13}, P_{23}, \dots$, and so on. To rephrase, these polynomials have degree $d_1 + d_2 + d_3 + \cdots$ with respect to the first row entries of a $k \times k$ matrix; degree $d_2 + d_3 + \cdots$ with respect to the second row entries; and so on.

A monomial in the flag minors is called a *cluster monomial* if all these minors belong to the same extended cluster. In the finite type cases where G is SL_3 , SL_4 , or SL_5 , the cluster monomials form a \mathbb{C} -basis of $\mathbb{C}[G]^U$. This is an instance of (the classical limit of) the *dual canonical basis* of G . Lusztig [30], also known as the *upper global basis* of M. Kashiwara [29]. (In the case $G = \mathrm{SL}_3$, this basis was introduced and studied in detail in [25].) This description of the basis served as the key original motivation for the introduction of cluster algebras in [15].

In infinite type, the picture turns out to be much more complicated. The fundamental result obtained in [28] asserts that in general, the dual canonical (or upper global) basis contains all cluster monomials. The rest of the basis still awaits an explicit description.

6.6. Cluster structure in the rings $\mathbb{C}[\text{Mat}_{k \times k}]$ and $\mathbb{C}[\text{SL}_k]$

The coordinate ring $\mathbb{C}[\text{Mat}_{k \times k}]$ of the space of $k \times k$ complex matrices is the polynomial ring $\mathbb{C}[z_{ij}]$; here $z = (z_{ij})$ denotes a generic $k \times k$ matrix. We will use double wiring diagrams to describe a cluster structure in this ring. An adaptation of this construction will then produce a cluster structure in the coordinate ring $\mathbb{C}[\text{SL}_k]$ of the special linear group.

Theorem 6.6.1. *The ring $\mathbb{C}[\text{Mat}_{k \times k}]$ has a cluster structure whose set of coefficient and cluster variables includes all minors of a $k \times k$ matrix.*

Proof. Recall from Section 2.4 and Exercise 3.1.3 that one can associate a seed $(\tilde{\mathbf{x}}, \tilde{B})$ to every double wiring diagram. All cluster and frozen variables in such a seed are minors of z . In particular, the set of frozen variables consists of all minors of the form $\Delta_{I,J}$ or $\Delta_{J,I}$ where $I = \{1, 2, \dots, i\}$ and $J = \{k - i + 1, k - i + 2, \dots, k\}$, with $i \in \{1, \dots, k\}$.

Recall that any two double wiring diagrams can be connected by local moves, cf. Figure 1.10. Since these moves correspond to mutations of the corresponding seeds, all such seeds define the same cluster algebra.

Let $(\tilde{\mathbf{x}}, \tilde{B})$ be a seed associated to a double wiring diagram. To show that $\mathbb{C}[\text{Mat}_{k \times k}] \subset \mathcal{A}(\tilde{\mathbf{x}}, \tilde{B})$, it suffices to show that each matrix entry z_{ij} lies in $\mathcal{A}(\tilde{\mathbf{x}}, \tilde{B})$. The latter statement follows from the fact that one can always construct a double wiring diagram whose extended cluster $\tilde{\mathbf{x}}$ contains z_{ij} .

The inclusion $\mathcal{A}(\tilde{\mathbf{x}}, \tilde{B}) \subset \mathbb{C}[\text{Mat}_{k \times k}]$ can be shown using either of the two arguments outlined below. These two arguments use two different seeds as well as two different versions of the Starfish lemma. The first argument relies on Proposition 6.4.1, which requires checking a coprimality condition on cluster variables; the second argument uses Corollary 6.4.6, which requires checking an irreducibility condition.

The polynomial ring $\mathbb{C}[z_{ij}]$ is factorial and therefore normal. In order to use Proposition 6.4.1, we first observe that $\tilde{\mathbf{x}} \subset \mathbb{C}[\text{Mat}_{k \times k}]$. Moreover, any two cluster variables in $\tilde{\mathbf{x}}$ are pairwise coprime, since the determinant is an irreducible polynomial. It remains to check property (3) of Proposition 6.4.1. Let us choose the seed $(\tilde{\mathbf{x}}, \tilde{B})$ shown in Figure 6.5. (The figure shows the example for $k = 4$ but the generalization to an arbitrary k is clear from the picture.) Note that the minors appearing in this seed are very simple: they are solid minors that “stick” to the left edge or the upper edge of the matrix. We now need to check that for each minor x_ℓ in the seed, the new cluster variable x'_ℓ obtained by an exchange with x_ℓ is a polynomial in the matrix entries that is moreover coprime to x_ℓ . Although some cluster variables x'_ℓ will no longer be minors (in particular, those resulting from mutations at degree 6 vertices in the quiver), one can adapt the argument from Section 6.5 to prove that they are nevertheless polynomials coprime to x_ℓ .




Figure 6.5. A double wiring diagram D whose extended cluster consists of solid minors. In the corresponding quiver, many mutable vertices will have degree 6. Mutation at such a vertex will result in a cluster variable which is not a minor.

An alternative approach relies on Corollary 6.4.6 to establish the inclusion $\mathcal{A}(\tilde{\mathbf{x}}, \tilde{B}) \subset \mathbb{C}[\text{Mat}_{k \times k}]$. Here we use a “grid seed” for $\mathbb{C}[\text{Mat}_{k \times k}]$, see Figures 6.6 and 6.7. This seed has the property that every mutable vertex in its quiver has degree three or four, and the corresponding exchange relation is a three-term Grassmann-Plücker relation. It follows that for each cluster variable x_ℓ the adjacent cluster variable x'_ℓ is a minor, and hence an irreducible polynomial. The claim $\mathcal{A}(\tilde{\mathbf{x}}, \tilde{B}) \subset \mathbb{C}[\text{Mat}_{k \times k}]$ follows. \square




Figure 6.6. A double wiring diagram for $\mathbb{C}[\text{Mat}_{4 \times 4}]$ whose associated quiver is the “grid quiver” shown in Figure 6.7.

Theorem 6.6.2. *The coordinate ring $\mathbb{C}[\text{SL}_k]$ of the special linear group has a cluster structure whose set of coefficient and cluster variables includes all minors of a $k \times k$ matrix, except for the determinant of the matrix.*

Proof. To adapt the above arguments to the case of $\mathbb{C}[\text{SL}_k]$, we use Proposition 6.4.11 to show that the ring $\mathbb{C}[\text{SL}_k]$ is factorial, and hence normal. The cluster variables, frozen variables, and clusters are the same as for $\mathbb{C}[\text{Mat}_{k \times k}]$, except that the $k \times k$ determinant of the entire matrix is no longer a frozen variable (as it is now equal to 1). \square




Figure 6.7. A grid seed for $\mathbb{C}[\text{Mat}_{4 \times 4}]$, cf. Figure 6.6. Every mutable vertex in the quiver has degree three or four.

Remark 6.6.3. One might expect that the cluster structures on $\mathbb{C}[\text{Mat}_{k \times k}]$ and $\mathbb{C}[\text{SL}_k]$ described in this section can be modified to yield a cluster structure in the coordinate ring of a general linear group GL_k . However this cannot be achieved without tweaking the basic definitions, because the inverse of the determinant $\det^{-1} \in \mathbb{C}[\text{GL}_k]$ is a regular function that does not lie in the cluster algebra. (The ground ring for a cluster algebra is the polynomial ring generated by the coefficient variables; it does not include their inverses.) As noted in Definition 3.1.6, a common alternative is to change the ground ring, adjoining the inverses of the coefficient variables (or “localizing at coefficients”). With this convention, the coordinate ring $\mathbb{C}[\text{GL}_k]$ becomes a cluster algebra.

Remark 6.6.4. The constructions presented above allow multiple generalizations and variations. In particular, one can replace SL_k by any connected, simply connected semisimple complex Lie group G and/or consider various subvarieties of G , such as those related to *double Bruhat cells*, see [3].

6.7. The cluster structure in the ring $\mathbb{C}[\widehat{\text{Gr}}_{a,b}]$

The Grassmannian $\text{Gr}_{a,b}$ of a -dimensional subspaces in \mathbb{C}^b can be embedded in the projective space of dimension $\binom{b}{a} - 1$ via the Plücker embedding; see, e.g., [8, Corollary 2.3]. Let $\widehat{\text{Gr}}_{a,b}$ denote the affine cone over $\text{Gr}_{a,b}$ taken in this embedding. The ring $\mathbb{C}[\widehat{\text{Gr}}_{a,b}]$ (the homogeneous coordinate ring of $\text{Gr}_{a,b}$) is generated by the Plücker coordinates P_J , where J ranges over all a -element subsets of $\{1, \dots, b\}$. These generators satisfy the quadratic *Grassmann-Plücker relations*.

Example 6.7.1 (cf. Section 1.2). The homogeneous coordinate ring $\mathbb{C}[\widehat{\text{Gr}}_{2,4}]$ is generated by the six Plücker coordinates $P_{12}, P_{13}, P_{14}, P_{23}, P_{24}, P_{34}$, which

are subject to the single Grassmann-Plücker relation

$$(6.7.1) \quad P_{13}P_{24} = P_{12}P_{34} + P_{14}P_{23}.$$

This ring carries the structure of a cluster algebra of rank 1 in which

- the ambient field is the field $\mathbb{C}(P_{12}, P_{13}, P_{14}, P_{23}, P_{34})$ of rational functions in five algebraically independent variables;
- the frozen variables are $P_{12}, P_{23}, P_{34}, P_{14}$;
- the cluster variables are P_{13} and P_{24} ;
- the single exchange relation is (6.7.1).

The ring $\mathbb{C}[\widehat{\text{Gr}}_{a,b}]$ is an archetypal object of classical invariant theory; see, e.g., [8, Chapter 2], [35, §9], [37, Section 11], and [41]. In invariant theory, this ring is typically given a somewhat different description:

Definition 6.7.2. Let $V \cong \mathbb{C}^a$ be an a -dimensional complex vector space equipped with a volume form. The special linear group $\text{SL}(V) \cong \text{SL}_a(\mathbb{C})$ naturally acts on the vector space V^b of b -tuples of vectors, hence on its coordinate (polynomial) ring. The *Plücker ring* $R_{a,b}$ is the ring

$$(6.7.2) \quad R_{a,b} = \mathbb{C}[V^b]^{\text{SL}(V)}$$

of $\text{SL}(V)$ -invariant polynomials on V^b . As a subring of a polynomial ring, $R_{a,b}$ is a domain.

In coordinate notation, the Plücker ring is described as follows. Consider a matrix $z = (z_{ij})$ of size $a \times b$ filled with indeterminates. The ring $R_{a,b}$ consists of polynomials in these ab variables that are invariant under the transformations $z \mapsto gz$, for $g \in \text{SL}_a(\mathbb{C})$. One example of such a polynomial is a *Plücker coordinate* P_J where J is an a -element subset of columns in z ; by definition, P_J is the $a \times a$ minor of z occupying the columns in J .

The First Fundamental Theorem of invariant theory, which goes back to A. Clebsch and H. Weyl, states the following.

Theorem 6.7.3. *The Plücker ring $R_{a,b}$ is generated by the $\binom{b}{a}$ Plücker coordinates P_J .*

Theorem 6.7.3 implies that for $a \leq b$, the Plücker ring $R_{a,b}$ is isomorphic to $\mathbb{C}[\widehat{\text{Gr}}_{a,b}]$, the homogeneous coordinate ring of the Grassmannian $\text{Gr}_{a,b}$. Therefore we can talk interchangeably about these two rings.

We note that the fact that the Plücker ring is finitely generated is a special case of Theorem 6.4.8.

Remark 6.7.4. The Second Fundamental Theorem, which we will not need, describes the ideal of relations among the generators P_J of the Plücker

ring $R_{a,b}$. As mentioned above, this ideal is generated by certain quadratic relations, the Grassmann-Plücker relations. The 3-term Grassmann-Plücker relations are among the exchange relations of the standard cluster structure on $R_{a,b}$ described below. When $3 \leq a \leq b - 3$, the Grassmann-Plücker relations include some longer quadratic relations which are not generated by the 3-term ones, cf. Example 6.8.6 below.

Corollary 6.7.5 ([40, Section 1.6b]). *The Plücker ring $R_{a,b}$ is factorial.*

Proof. This follows from Theorem 6.4.12 and (6.7.2). (Being semisimple, SL_a has no nontrivial characters.) \square

We next describe a cluster structure in the Plücker ring $R_{a,b}$ [38]. While canonical up to a ring automorphism, this structure will depend on the choice of a *cyclic ordering* of the b vectors.

The set of coefficient variables for this cluster structure in $R_{a,b}$ consists of the b Plücker coordinates P_J where J is a contiguous segment modulo b . For example, the coefficient variables for $R_{3,7}$ are the Plücker coordinates

$$P_{123}, P_{234}, P_{345}, P_{456}, P_{567}, P_{167}, P_{127}.$$

We will work with some distinguished seeds: the *rectangles seed* $\Sigma_{a,b}$, together with its cyclic shifts $\Sigma_{a,b}^i$ for $1 \leq i \leq b - 1$. To define the rectangles seed $\Sigma_{a,b}$, we first construct a quiver $Q_{a,b}$ whose vertices are labeled by the rectangles contained in an $a \times (b - a)$ rectangle R , including the empty rectangle \emptyset . The frozen vertices of $Q_{a,b}$ are labeled by the rectangles of sizes $a \times j$ (with $1 \leq j \leq b - a$), rectangles of sizes $i \times (b - a)$ (with $1 \leq i \leq a$), and the empty rectangle. The arrows from an $i \times j$ rectangle go to rectangles of sizes $i \times (j + 1)$, $(i + 1) \times j$, and $(i - 1) \times (j - 1)$ (assuming those rectangles have nonzero dimensions, fit inside R , and the arrow does not connect two frozen vertices). There is also an arrow from the frozen vertex labeled by \emptyset to the vertex labeled by the 1×1 rectangle. See Figure 6.8.

We map each rectangle r contained in the $a \times (b - a)$ rectangle R to an a -element subset of $\{1, 2, \dots, b\}$ (representing a Plücker coordinate), as follows. We justify r so that its upper left corner coincides with the upper left corner of R . There is a path of length b from the northeast corner of R to the southwest corner of R which cuts out the smaller rectangle r ; we label the steps of this path from 1 to b . We then map r to the set of labels $J(r)$ of the vertical steps on this path. This construction allows us to assign to each vertex of the quiver $Q_{a,b}$ a particular Plücker coordinate. We set

$$\tilde{\mathbf{x}}^{a,b} = \{P_{J(r)} \mid r \text{ is a rectangle contained in an } a \times (b - a) \text{ rectangle}\},$$

and then define the *rectangles seed* $\Sigma_{a,b} = (\tilde{\mathbf{x}}^{a,b}, \tilde{B}(Q_{a,b}))$. See Figure 6.9.




Figure 6.8. The quiver $Q_{3,7}$. The vertices are labeled by rectangles contained in a 3×4 rectangle, and arranged in a (triangulated) grid. The width and height of the rectangles increase from left to right and from top to bottom, respectively.




Figure 6.9. The rectangles seed $\Sigma_{3,7}$.

Remark 6.7.6. When $a = 2$, the rectangles seed $\Sigma_{2,b}$ coincides with the seed associated to the triangulation of the polygon \mathbf{P}_b that uses all of the diagonals incident to the vertex b , cf. Example 6.3.1 and Definition 2.2.1.

Given an a -element subset $J = \{j_1, j_2, \dots, j_a\} \subset \{1, 2, \dots, b\}$ and a positive integer i , we define

$$(J + i) \bmod b = \{j_1 + i, j_2 + i, \dots, j_a + i\},$$

where the sums are taken modulo b . We define a quiver $Q_{a,b}^i$ and seed $\Sigma_{a,b}^i$ by replacing each vertex label J in $Q_{a,b}$ by $(J + i) \bmod b$. (The quivers $Q_{a,b}$ and $Q_{a,b}^i$ are exactly the same; only their vertex labels are different.)

Exercise 6.7.7. Start from the seed $\Sigma_{a,b}$ and mutate at each of the mutable vertices of $Q_{a,b}$ exactly once, in the following order: mutate each row from

left to right, starting from the bottom row and ending at the top row. Show that at the end of this process, one obtains the seed $\Sigma_{a,b}^1$. For example, in Figure 6.9, mutating at 457, 347, 237, 467, 367, 267 (in this order) recovers the same quiver but with each label cyclically shifted.

The same mutation sequence transforms the seed $\Sigma_{a,b}^i$ into $\Sigma_{a,b}^{i+1}$. Therefore the rectangles seed and all its cyclic shifts are mutation equivalent.

Theorem 6.7.8. *The seed pattern defined by $\Sigma_{a,b}$ (or by any of its cyclic shifts, cf. Exercise 6.7.7) gives the Plücker ring $R_{a,b} = \mathbb{C}[\widehat{\text{Gr}}_{a,b}]$ the structure of a cluster algebra.*

Theorem 6.7.8 was first proved in [38], using results from [36]. Below in this section we give a different (and self-contained) proof.

Remark 6.7.9. There is a well known isomorphism $R_{a,b} \rightarrow R_{b-a,b}$ defined by $P_J \mapsto P_{J^c}$, where $J^c = \{1, 2, \dots, b\} \setminus J$. This isomorphism extends to an isomorphism between respective seed patterns in $R_{a,b}$ and $R_{b-a,b}$.

Lemma 6.7.10. *There is an injective ring homomorphism $R_{a-1,b-1} \rightarrow R_{a,b}$ which sends P_I to $P_{I \cup \{b\}}$.*

Proof. The fact that the correspondence $P_I \mapsto P_{I \cup \{b\}}$ extends to a ring homomorphism follows from Muir's law (Proposition 1.3.5). \square

We call the map $R_{a-1,b-1} \rightarrow R_{a,b}$ described above the *Muir embedding*.

Recall the notion of a seed subpattern from Definition 4.2.6.

Lemma 6.7.11. *The Muir embedding sends the seed pattern in $R_{a-1,b-1}$ defined by $\Sigma_{a-1,b-1}$ to a subpattern of the seed pattern in $R_{a,b}$ defined by $\Sigma_{a,b}$.*

Proof. Delete the bottom row of vertices in the rectangles quiver $Q_{a,b}$. Freeze the vertices of the new bottom row of the resulting quiver. Delete any arrows connecting two frozen vertices. Then remove the index b from every label. This will produce the rectangles quiver $Q_{a-1,b-1}$, with its standard labeling. \square

Remark 6.7.12. Similarly, the seed pattern in $R_{a,b-1}$ defined by $\Sigma_{a,b-1}$ is isomorphic to a seed subpattern of the seed pattern in $R_{a,b}$ defined by $\Sigma_{a,b}$. (This involves deleting the rightmost column of the quiver $Q_{a,b}$.)

Lemma 6.7.13. $R_{a,b} \subset \mathcal{A}(\Sigma_{a,b})$.

Proof. Since the Plücker coordinates P_J generate the Plücker ring $R_{a,b}$, it suffices to show that each P_J lies in the cluster algebra $\mathcal{A}(\Sigma_{a,b})$.

We will prove this claim by induction on a . The base case $a = 2$ holds from our earlier analysis of $\text{Gr}_{2,b}$. By induction, for any $(a - 1)$ -element

subset I of $\{1, 2, \dots, b-1\}$, there is a sequence of mutations that we can apply to $\Sigma_{a-1, b-1}$ to obtain the Plücker coordinate P_I . By Lemma 6.7.11, we can apply the same sequence of mutations to $\Sigma_{a,b}$ to obtain the Plücker coordinate $P_{I \cup \{b\}}$. Consequently all Plücker coordinates of the form $P_{I \cup \{b\}}$ belong to the cluster algebra $\mathcal{A}(\Sigma_{a,b})$. It then follows by Exercise 6.7.7 that all Plücker coordinates in $R_{a,b}$ lie in $\mathcal{A}(\Sigma_{a,b})$. \square

Proof of Theorem 6.7.8. In view of Lemma 6.7.13, it remains to show that $\mathcal{A}(\Sigma_{a,b}) \subset R_{a,b}$. By the Starfish lemma, all we need to establish is that mutating at any vertex P_J of the quiver $Q_{a,b}$ yields a cluster variable $(P_J)'$ which is coprime to P_J .

The mutable vertices of $Q_{a,b}$ all have degree 4 or 6. If we mutate at a degree 4 vertex of $Q_{a,b}$, then the corresponding exchange relation is a 3-term Grassmann-Plücker relation, and the resulting cluster variable is a Plücker coordinate. Since the determinant is an irreducible polynomial, the old and new cluster variables are coprime.

If we mutate at a degree 6 vertex, the resulting cluster variable is not a Plücker coordinate; however one can use an argument similar to that from Section 6.5 to prove that the old and new cluster variables are still coprime. In this case, our degree 6 vertex is labeled by some Plücker coordinate P_{ijkS} , where the subset $S \subset \{1, \dots, b\}$ of size $a-3$ is disjoint from $\{i, j, k\}$. The exchange relation has the form

$$P_{ijkS}P'_{ijkS} = P_{ikfS}P_{ijdS}P_{jkeS} + P_{ikdS}P_{ijeS}P_{jkfS},$$

where the subset $\{d, e, f\} \subset \{1, \dots, b\}$ is disjoint from $\{i, j, k\} \cup S$. One can then check that $P'_{ijkS} = P_{ikfS}P_{jdeS} - P_{jkdS}P_{iefS}$.

We need to show that P_{ijkS} and P'_{ijkS} are coprime. Since the determinant is an irreducible polynomial, the only way P_{ijkS} and P'_{ijkS} can fail to be coprime is if P_{ijkS} divides P'_{ijkS} . Let us show that this cannot happen. Let z be a generic $3 \times b$ matrix; let us augment it to an $a \times b$ matrix \hat{z} by adding new rows 4 through a , where the submatrix located in rows $4 \dots a$ and columns S is the identity, and all other entries in rows $4 \dots a$ are 0. Then $P_{ijkS}(\hat{z})$ divides $P'_{ijkS}(\hat{z})$ implies that $P_{ijk}(z)$ divides $P'_{ijk}(z)$. If we specialize $z_{1d} = z_{1e} = z_{2d} = z_{2e} = 0$, then P_{ijk} is unchanged whereas P_{jde} becomes 0, P_{jkd} becomes $z_{3,d} \Delta_{12,jk}(z)$, and P_{ief} becomes $-z_{3,e} \Delta_{12,if}$. Thus P'_{ijk} specializes to $z_{3,d} \Delta_{12,jk}(z) z_{3,e} \Delta_{12,if}$. Now if P_{ijk} divides P'_{ijk} then the same is true after specialization. But P_{ijk} does not divide $z_{3,d} \Delta_{12,jk}(z) z_{3,e} \Delta_{12,if}$. Thus P_{ijkS} and P'_{ijkS} are coprime, and we are done. \square

Typically, the cluster structure in a Plücker ring $R_{a,b}$ is of infinite type. The few exceptional cases where it has finite type are listed in Table 6.2.

Ring	Cluster type
$R_{2,b}$ and $R_{b-2,b}$	A_{b-3}
$R_{3,6}$	D_4
$R_{3,7}$ and $R_{4,7}$	E_6
$R_{3,8}$ and $R_{5,8}$	E_8
$R_{a,b}$ for other a, b	infinite type

Table 6.2. The type of the cluster structure of $R_{a,b}$.

Remark 6.7.14. It is natural to seek an explicit description for all cluster and coefficient variables for the cluster structure in $R_{a,b}$ described above. As we have seen, this set contains all Plücker coordinates. In the cases $a = 2$ and $a = b - 2$, there is nothing else; in all other cases, the list includes non-Plücker cluster variables. For the finite types listed in Table 6.2, the formulas for non-Plücker variables were given in [38]. Beyond finite type, the problem remains open. The case $a = 3$ was extensively studied in [12].

Remark 6.7.15. The above construction can be adapted to yield a cluster structure in the coordinate ring $\mathbb{C}[\text{Mat}_{a \times (b-a)}]$ of the affine space of $a \times (b-a)$ matrices. Append an identity matrix to the right of an $a \times (b-a)$ matrix z to obtain an $a \times b$ matrix z' . Up to SL_a action, the only restriction on z' is that its $a \times a$ minor occupying the last a columns is equal to 1. We can now identify the minors of z with the maximal minors of z' (the Plücker coordinates): a Plücker coordinate $P_J \in R_{a,b}$ corresponds to the minor $\varphi(P_J) = \Delta_{KL}(z) \in \mathbb{C}[\text{Mat}_{a \times (b-a)}]$ whose row and column sets are given by

$$K = (\{b-a+1, b-a+2, \dots, b\} \setminus J) - b + a,$$

$$L = J \cap \{1, 2, \dots, b-a\};$$

here the notation $S-c$ means $\{s-c \mid s \in S\}$. Given a seed for $R_{a,b}$, applying the map φ to all cluster and coefficient variables (except for the coefficient variable $P_{b-a+1, b-a+2, \dots, b}$) yields a seed for $\mathbb{C}[\text{Mat}_{a \times b}]$.

In the special case $b = 2a$, this identification shows that the cluster structures in the rings $\mathbb{C}[\text{Mat}_{a \times a}]$ and $\mathbb{C}[\text{SL}_a]$ introduced in Section 6.6 are very closely related to the cluster structure in the Plücker ring $R_{a,2a}$.

The constructions presented in Sections 6.5–6.7 can be generalized and modified to build cluster structures in many other rings naturally arising in the context of classical invariant theory as well as algebraic Lie theory. We already mentioned generalizations and extensions to other semisimple Lie

groups, their subgroups, parabolic quotients, and double Bruhat cells; the excellent survey [22] describes the state of the art circa 2013.

To keep our exposition within reasonable bounds, we did not discuss the constructions of cluster structures in the rings of SL_k invariants of collections of vectors, covectors, and/or matrices [6, 12, 11]. Likewise, we left out the treatment of the Fock-Goncharov configuration spaces [9, 10] and related topics of higher Teichmüller theory.

6.8. Defining cluster algebras by generators and relations

One traditional way of describing a commutative algebra \mathcal{A} (say over \mathbb{C}) is in terms of generators and relations. In this approach, \mathcal{A} is represented as a quotient of a \mathbb{C} -algebra $\mathbb{C}[\mathbf{z}] = \mathbb{C}[z_1, z_2, \dots]$ freely generated by a (finite or countable) set of “variables” $\mathbf{z} = \{z_1, z_2, \dots\}$ modulo an explicitly given ideal $I \subset \mathbb{C}[\mathbf{z}]$. In typical applications, the set \mathbf{z} is finite (so that $P = \mathbb{C}[\mathbf{z}]$ is a polynomial ring) and the ideal I is finitely generated: $I = \langle g_1, \dots, g_N \rangle$, where g_1, \dots, g_N are polynomials in the variables z_1, z_2, \dots . (By common abuses of terminology, we identify polynomials $f \in \mathbb{C}[\mathbf{z}]$ with the elements of $\mathcal{A} \cong \mathbb{C}[\mathbf{z}]/I$ they represent. We also conflate the polynomials $g \in I$ with the relations $g(z_1, z_2, \dots) = 0$ holding in \mathcal{A} .)

The definition of a cluster algebra (Definition 3.1.6) is set up differently: a cluster algebra \mathcal{A} is defined inside a field \mathcal{F} of rational functions in several variables as the algebra generated by certain (recursively determined) elements of \mathcal{F} , the cluster variables of \mathcal{A} . While the relations among these generators are not given explicitly, we do know some of them, namely the exchange relations (3.1.1).

It is natural to try to extract from this definition a traditional-style description of a cluster algebra as a quotient of a polynomial ring. This runs into two issues. First, the set of cluster variables is typically infinite. Second, the exchange relations do not, in general, generate the ideal of all relations among cluster variables. We will discuss these two issues one by one.

The following statement, provided here without proof, shows that some cluster algebras are *not* finitely generated:

Proposition 6.8.1 ([3, Theorem 1.26]). *A cluster algebra of rank 3 with trivial coefficients is finitely generated if and only if it has an acyclic seed.*

In the terminology of Example 4.1.5, a cluster algebra defined by a 3-vertex quiver Q with no frozen vertices is finitely generated if and only if Q is mutation-acyclic.

To illustrate Proposition 6.8.1, the cluster algebra defined by the Markov quiver (see Figure 2.10) is not finitely generated. The following result, combined with Proposition 6.8.1, provides many more examples.

Proposition 6.8.2 ([1, Theorem 1.2]). *Let $Q(a, b, c)$ denote the quiver with vertices 1, 2, 3 and $a + b + c$ arrows: a arrows $1 \rightarrow 2$, b arrows $2 \rightarrow 3$, and c arrows $3 \rightarrow 1$. (See Figure 6.10.) The following are equivalent:*

- the quiver $Q(a, b, c)$ is not mutation-acyclic;
- $a, b, c \geq 2$ and $\det \begin{pmatrix} 2 & a & c \\ a & 2 & b \\ c & b & 2 \end{pmatrix} \geq 0$.




Figure 6.10. The 3-vertex quiver $Q(a, b, c)$ with $a = 3$, $b = 2$, $c = 1$.

Various cluster algebras arising in Lie theory, such as the ones discussed in Sections 6.5, 6.6, and 6.7, are finitely generated, for the reasons given in Theorems 6.4.8, 6.4.9, and 6.4.10.

Another class of finitely generated cluster algebras is provided by the following result, stated here without proof.

Theorem 6.8.3 ([3, Corollary 1.21]). *Any cluster algebra defined by an acyclic quiver (with no frozen vertices) is finitely generated. In fact, it is generated by the cluster variables belonging to the initial acyclic seed together with the cluster variables obtained from this seed by a single mutation; the ideal of relations among these cluster variables is generated by the exchange relations out of the initial acyclic seed.*

Theorem 6.8.3 was extended in [33] to the much larger class of “locally acyclic” cluster algebras.

Remark 6.8.4. Any cluster algebra of finite type is finitely generated. This follows from the appropriate generalization of Theorem 6.8.3: a cluster algebra of finite type always has an acyclic seed of the kind described in Theorem 5.2.8. (In the quiver case, the quiver at such a seed is an orientation of the corresponding Dynkin diagram.) See [3, Remark 1.22].

Many cluster algebras of infinite type (and even infinite mutation type) are finitely generated. For example, any Plücker ring $R_{a,b}$ is finitely generated whereas its cluster structure is typically of infinite type, see Table 6.2.

When thinking about finite generation, it is helpful to keep in mind that by Proposition 6.4.7, a cluster algebra \mathcal{A} is finitely generated if and only if \mathcal{A} is generated by a finite subset of cluster and coefficient variables.

We next turn to the problem of describing the ideal of relations satisfied by a set of generators of a cluster algebra. Even when a cluster algebra is of finite type, this is a delicate issue, as Examples 6.8.5 and 6.8.6 below demonstrate.

Example 6.8.5. Let $\mathcal{A} = \mathcal{A}(1, 2)$ be the cluster algebra of type B_2 with trivial coefficients. In the notation of Example 3.2.7, \mathcal{A} is generated by the 6-periodic sequence of cluster variables z_1, z_2, \dots satisfying the exchange relations

$$(6.8.1) \quad z_1 z_3 - z_2^2 - 1 = 0,$$

$$(6.8.2) \quad z_2 z_4 - z_3 - 1 = 0,$$

$$(6.8.3) \quad z_3 z_5 - z_4^2 - 1 = 0,$$

$$(6.8.4) \quad z_4 z_6 - z_5 - 1 = 0,$$

$$(6.8.5) \quad z_5 z_1 - z_6^2 - 1 = 0,$$

$$(6.8.6) \quad z_6 z_2 - z_1 - 1 = 0.$$

It turns out that these relations do *not* generate the ideal of all relations satisfied by z_1, \dots, z_6 . It is not hard to check (using the formulas in Example 3.2.7) that these cluster variables also satisfy the relations

$$(6.8.7) \quad z_1 z_4 - z_2 - z_6 = 0,$$

$$(6.8.8) \quad z_3 z_6 - z_4 - z_2 = 0,$$

$$(6.8.9) \quad z_5 z_2 - z_6 - z_4 = 0,$$

none of which lies in the ideal generated by (6.8.1)–(6.8.6) inside the polynomial ring in six formal variables z_1, \dots, z_6 .

The last claim, like several others made below in this section, can be readily checked using any of the widely available software packages for commutative algebra.

Example 6.8.6. Consider the Plücker ring $R_{3,6} = \mathbb{C}[\widehat{\text{Gr}}(3, 6)]$, viewed as a cluster algebra of finite type D_4 , as explained in Section 6.7. The set of its cluster variables contains the Plücker coordinates P_{ijk} . This cluster algebra is graded, with $\deg(P_{ijk}) = 1$ for all i, j, k . All cluster variables are homogeneous elements, and all relations among them are homogeneous as well. These relations in particular include the Grassmann-Plücker relation

$$(6.8.10) \quad P_{135}P_{246} - P_{134}P_{256} - P_{136}P_{245} - P_{123}P_{456} = 0.$$

The relation (6.8.10) cannot be written as a polynomial combination of exchange relations, since all those relations have degree at least 2, and none of them involves the monomial $P_{135}P_{246}$. Thus the ideal of relations among the cluster variables of $R_{3,6}$ is *not* generated by the exchange relations.

This example can be extended to bigger Grassmannians using Muir’s Law (Proposition 1.3.5).

On the bright side, the rings discussed in Examples 6.8.5 and 6.8.6 do have “nice” explicit presentations. The ideal of relations among the Plücker

coordinates is generated by the classical (quadratic) Grassmann-Plücker relations. As to the cluster algebra $\mathcal{A}(1,2)$ from Example 6.8.5, the ideal of relations among the six cluster variables z_1, \dots, z_6 is generated by (the left-hand sides of) the relations (6.8.1)–(6.8.6) and (6.8.7)–(6.8.9).

In Theorem 6.8.10 below, we will provide a general description of a finitely generated cluster algebra \mathcal{A} in terms of generators and relations. This will require some preparation.

Definition 6.8.7. Recall that a cluster algebra \mathcal{A} of rank n is defined by a seed pattern whose seeds are labeled by the vertices of the n -regular tree \mathbb{T}_n , cf. Definition 3.1.4. Let T be a finite subtree of \mathbb{T}_n . For $i \in \{1, \dots, n\}$, let $T[i]$ denote the forest obtained from T by removing the edges labeled by i . We denote by \mathbf{z}_T the (finite) set of formal variables which includes

- one formal variable for each coefficient variable of \mathcal{A} , and
- one formal variable for each connected component of $T[i]$, for every $i \in \{1, \dots, n\}$.

The formal variable associated to a connected component C of $T[i]$ naturally corresponds to the unique cluster variable in \mathcal{A} that is indexed by i within each of the labeled seeds in C .

We denote by $\mathbb{C}[\mathbf{z}_T]$ the ring of polynomials in the set of variables \mathbf{z}_T . The *exchange ideal* $I_T \subset \mathbb{C}[\mathbf{z}_T]$ is the ideal generated by the exchange relations corresponding to the edges of T . More precisely, for each exchange relation $zz' = M_1 + M_2$ corresponding to an edge of T (here $z, z' \in \mathbf{z}_T$, and M_1, M_2 are monomials in the elements of \mathbf{z}_T), the exchange ideal I_T contains the polynomial $zz' - M_1 - M_2$.

Let \mathcal{Z}_T denote the set of all cluster and coefficient variables appearing in the seeds labeled by the vertices of T . Let $\mathcal{A}_T \subset \mathcal{A}$ be the subalgebra generated by \mathcal{Z}_T . We are especially interested in the cases where $\mathcal{A}_T = \mathcal{A}$, so that \mathcal{Z}_T generates the entire cluster algebra \mathcal{A} .

In what follows, we habitually use the same notation for a formal variable $z \in \mathbf{z}_T$ and the corresponding cluster variable $z \in \mathcal{Z}_T$. When this abuse of notation becomes dangerously confusing, we write $f(\mathbf{z}_T)$ and $f(\mathcal{Z}_T)$ to distinguish between a polynomial $f \in \mathbb{C}[\mathbf{z}_T]$ and its evaluation in $\mathcal{A}_T \subset \mathcal{A}$.

Remark 6.8.8. Let \mathcal{A} be a finitely generated cluster algebra. By Proposition 6.4.7, \mathcal{A} is generated by a finite subset \mathbf{z} of cluster and coefficient variables. Enlarging \mathbf{z} if necessary, we may furthermore assume that all cluster variables in \mathbf{z} come from clusters connected to each other by mutations that pass through clusters all of whose cluster variables belong to \mathbf{z} . It follows that we can find a finite tree T such that $\mathcal{A}_T = \mathcal{A}$.

Example 6.8.9. Let $\mathcal{A} = \mathcal{A}(1, 2)$, as in Example 6.8.5. We first consider the 3-vertex tree T corresponding to the following triple of clusters:

$$(6.8.11) \quad (z_1, z_2) \xrightarrow{1} (z_3, z_2) \xrightarrow{2} (z_3, z_4).$$

Then $\mathcal{Z}_T = \{z_1, z_2, z_3, z_4\}$, in the notation of Definition 6.8.7. The relations

$$(6.8.12) \quad z_6 = z_1 z_4 - z_2,$$

$$(6.8.13) \quad z_5 = z_4 z_6 - 1 = z_1 z_4^2 - z_2 z_4 - 1 = z_1 z_4^2 - z_3 - 2$$

(cf. (6.8.7) and (6.8.4)) imply that the cluster algebra \mathcal{A} is generated by \mathcal{Z}_T . This is an instance of Theorem 6.8.3: the four cluster variables in \mathcal{Z}_T come from the cluster (z_2, z_3) and the two clusters obtained from it by single mutations.

The four elements of \mathcal{Z}_T satisfy the exchange relations (6.8.1) and (6.8.2). The left-hand sides of these relations correspond to the generators of the exchange ideal $I_T \subset \mathbb{C}[\mathbf{z}_T]$. One can check that this exchange ideal contains all relations satisfied by the cluster variables z_1, z_2, z_3, z_4 , in agreement with Theorem 6.8.3. Consequently $\mathcal{A} \cong \mathbb{C}[\mathbf{z}_T]/I_T$.

Alternatively, consider the subtree T spanning the four clusters

$$(6.8.14) \quad (z_1, z_2) \xrightarrow{1} (z_3, z_2) \xrightarrow{2} (z_3, z_4) \xrightarrow{1} (z_5, z_4).$$

Again, the set $\mathcal{Z}_T = \{z_1, \dots, z_5\}$ generates \mathcal{A} . The three exchange relations associated with the edges of T are (6.8.1), (6.8.2), and (6.8.3). It turns out that the exchange ideal $I_T \subset \mathbb{C}[\mathbf{z}_T]$ generated by (the left-hand sides of) these relations does not contain some of the relations satisfied by the cluster variables z_1, \dots, z_5 . For example, $f(\mathbf{z}_T) = z_1 z_4^2 - z_3 - z_5 - 2 \notin I_T$ even though $f(\mathcal{Z}_T) = 0$ in \mathcal{A} , cf. (6.8.13). Thus for this choice of a tree T , we have $\mathcal{A} \not\cong \mathbb{C}[\mathbf{z}_T]/I_T$. (As the exchange ideal I_T is radical in this instance, the gap cannot be explained by the discrepancy between the ideal of an affine variety and an ideal coming from its set-theoretic description.)

To describe the relationship between the algebra \mathcal{A}_T and the quotient $\mathbb{C}[\mathbf{z}_T]/I_T$, we will need the following notation. Let M_T denote the product of all formal variables in \mathbf{z}_T that correspond to the (mutable) cluster variables. We denote by

$$(6.8.15) \quad J_T = (I_T : \langle M_T \rangle^\infty) = \{f \in \mathbb{C}[\mathbf{z}_T] : (M_T)^a f \in I_T \text{ for some } a\}$$

the saturation of the exchange ideal I_T by the principal ideal $\langle M_T \rangle$. In plain language, J_T consists of all polynomials that can be multiplied by a monomial so that the product lies in the exchange ideal I_T .

Theorem 6.8.10. *For a polynomial $f(\mathbf{z}_T)$, the following are equivalent:*

- $f(\mathcal{Z}_T) = 0$, i.e., f describes a relation among cluster variables in \mathcal{A}_T ;
- $f(\mathbf{z}_T)$ lies in the saturated ideal J_T .

We thus have the canonical isomorphism

$$(6.8.16) \quad \mathcal{A}_T \cong \mathbb{C}[\mathbf{z}_T]/J_T.$$

Informally, a polynomial in the cluster variables vanishes in the cluster algebra if and only if this polynomial can be multiplied by some monomial in cluster variables so that the product lies in the exchange ideal.

Remark 6.8.11. If the cluster algebra \mathcal{A} is finitely generated, with $\mathcal{A} = \mathcal{A}_T$ (cf. Remark 6.8.8), then (6.8.16) provides an implicit presentation of \mathcal{A} in terms of generators and relations. Furthermore, in each specific example, the saturated ideal J_T can be explicitly computed using existing efficient algorithms of computational commutative algebra.

Before proving Theorem 6.8.10, we illustrate it with a couple of examples.

Example 6.8.12. Continuing with Example 6.8.9, let $\mathcal{A} = \mathcal{A}(1, 2)$. Take the tree T shown in (6.8.14). We saw that the polynomial $f = z_1 z_4^2 - z_3 - z_5 - 2$ does not lie in the exchange ideal I_T , even though f describes an identity among the generators of \mathcal{A} . On the other hand,

$$\begin{aligned} z_3 f &= z_3(z_1 z_4^2 - z_3 - z_5 - 2) \\ &= z_4^2(z_1 z_3 - z_2^2 - 1) + (z_2 z_4 + z_3 + 1)(z_2 z_4 - z_3 - 1) - (z_3 z_5 - z_4^2 - 1) \in I_T, \end{aligned}$$

so f lies in the saturated ideal J_T .

Example 6.8.13. Continuing with Example 6.8.6, consider the Plücker ring $R_{3,6}$. Although the Grassmann-Plücker relation (6.8.10) does not lie in the ideal generated by the exchange relations, we *can* multiply (6.8.10) by a monomial (in fact, by a single variable) and get inside the ideal:

$$(6.8.17) \quad \begin{aligned} &P_{124}(P_{135}P_{246} - P_{134}P_{256} - P_{136}P_{245} - P_{123}P_{456}) \\ &= P_{246}(P_{124}P_{135} - P_{123}P_{145} - P_{125}P_{134}) \end{aligned}$$

$$(6.8.18) \quad - P_{134}(P_{124}P_{256} - P_{125}P_{246} + P_{126}P_{245})$$

$$(6.8.19) \quad - P_{245}(P_{124}P_{136} - P_{123}P_{146} - P_{126}P_{134})$$

$$(6.8.20) \quad - P_{123}(P_{124}P_{456} - P_{145}P_{246} + P_{146}P_{245}) \in I_T.$$

(Each of the four parenthetical expressions in (6.8.17)–(6.8.20) is a three-term Grassmann-Plücker relation, thus an instance of an exchange relation.)

Proof of Theorem 6.8.10. Going in one direction, let us verify that if $f(\mathbf{z}_T) \in J_T$, then $f(\mathcal{Z}_T) = 0$. Suppose that $M(\mathbf{z}_T)$ is a monomial such that $f(\mathbf{z}_T)M(\mathbf{z}_T) \in I_T$. Since every exchange relation holds when we substitute cluster variables into it, this implies that $f(\mathcal{Z}_T)M(\mathcal{Z}_T) = 0$ in \mathcal{A}_T . But \mathcal{A}_T is contained in a field \mathcal{F} , and $M(\mathcal{Z}_T)$ is a nonzero element of \mathcal{F} . Therefore $f(\mathcal{Z}_T) = 0$ as desired.

To prove the converse, we will need the following definitions. Fix a root vertex t_0 in the tree T . Let us linearly order the set \mathbf{z}_T so that

- the coefficient variables and the variables associated with the root cluster are smaller than the remaining variables;
- for each exchange relation $zz' = \dots$, we have $z' < z$ if the “cluster” containing z' is closer to t_0 in T than the “cluster” containing z .

(We put the word “cluster” in quotation marks since we are dealing with formal variables rather than the associated cluster variables.)

Let $f = f(\mathbf{z}_T)$ be a polynomial such that $f(\mathcal{Z}_T) = 0$. We need to show that there exists a monomial $M \in \mathbb{C}[\mathbf{z}_T]$ such that $fM \in I_T$. We will prove this by double induction: first, on the largest variable $z \in \mathbf{z}_T$ appearing in f , and for a given z , on the degree with which z appears. In other words, we will argue as follows. Let $z \in \mathbf{z}_T$ be the largest variable appearing in f , say with degree d . Then we can assume, while proving the claim above, that a similar statement holds for any polynomial that only involves the variables smaller than z , and perhaps also z in degrees $< d$.

We begin by writing

$$f = zg + h,$$

where $g, h \in \mathbb{C}[\mathbf{z}_T]$ are polynomials, with h not involving z . Let

$$E = zz' - M_1 - M_2 \in I_T$$

be the polynomial associated with the unique exchange relation among the variables in \mathbf{z}_T that corresponds to an edge in T and where $z' < z$. Thus $M_1, M_2 \in \mathbb{C}[\mathbf{z}_T]$ are monomials that only involve variables smaller than z .

Now set

$$f' = z'f - Eg.$$

Then $f'(\mathcal{Z}_T) = 0$ because $f(\mathcal{Z}_T) = E(\mathcal{Z}_T) = 0$. Moreover the calculation

$$\begin{aligned} f' &= z'f - Eg \\ &= zz'g + z'h - (zz' - M_1 - M_2)g \\ &= z'h + (M_1 + M_2)g \end{aligned}$$

shows that f' satisfies the conditions of the induction hypothesis. (Indeed, the polynomials z', h, M_1 and M_2 only involve variables $< z$ whereas

$\deg_z(g) = \deg_z(f) - 1$.) Hence there exists a monomial $M' = M'(\mathbf{z}_T)$ satisfying $f'M' \in I_T$. Now let $M = M'z'$ and conclude that

$$fM = M'z'f = M'f' + M'Eg \in I_T,$$

proving the claim. □

Bibliography

- [1] BEINEKE, A., BRÜSTLE, T., AND HILLE, L. Cluster-cyclic quivers with three vertices and the Markov equation. *Algebr. Represent. Theory* 14, 1 (2011), 97–112. With an appendix by Otto Kerner.
- [2] BERENSTEIN, A., FOMIN, S., AND ZELEVINSKY, A. Parametrizations of canonical bases and totally positive matrices. *Adv. Math.* 122, 1 (1996), 49–149.
- [3] BERENSTEIN, A., FOMIN, S., AND ZELEVINSKY, A. Cluster algebras. III. Upper bounds and double Bruhat cells. *Duke Math. J.* 126, 1 (2005), 1–52.
- [4] BEZRUKAVNIKOV, R., BRAVERMAN, A., AND POSITSSELSKII, L. Gluing of abelian categories and differential operators on the basic affine space. *J. Inst. Math. Jussieu* 1, 4 (2002), 543–557.
- [5] BRION, M. *Introduction to actions of algebraic groups*. Les cours du CIRM, 2010, <https://eudml.org/doc/116362>.
- [6] CARDE, K. Cluster algebras and classical invariant rings, Ph.D. thesis, University of Michigan, 2014, <https://deepblue.lib.umich.edu/handle/2027.42/108931>.
- [7] DANILOV, V. I. Algebraic varieties and schemes. In *Algebraic geometry, I*, vol. 23 of *Encyclopaedia Math. Sci.* Springer, Berlin, 1994, pp. 167–297.
- [8] DOLGACHEV, I. *Lectures on invariant theory*, vol. 296 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 2003.
- [9] FOCK, V., AND GONCHAROV, A. Moduli spaces of local systems and higher Teichmüller theory. *Publ. Math. Inst. Hautes Études Sci.*, 103 (2006), 1–211.
- [10] FOCK, V. V., AND GONCHAROV, A. B. Dual Teichmüller and lamination spaces. In *Handbook of Teichmüller theory. Vol. I*, vol. 11 of *IRMA Lect. Math. Theor. Phys.* Eur. Math. Soc., Zürich, 2007, pp. 647–684.
- [11] FOMIN, S., AND PLYAVSKYY, P. Webs on surfaces, rings of invariants, and clusters. *Proc. Natl. Acad. Sci. USA* 111, 27 (2014), 9680–9687.
- [12] FOMIN, S., AND PLYAVSKYY, P. Tensor diagrams and cluster algebras. *Adv. Math.* 300 (2016), 717–787.
- [13] FOMIN, S., AND ZELEVINSKY, A. Double Bruhat cells and total positivity. *J. Amer. Math. Soc.* 12, 2 (1999), 335–380.
- [14] FOMIN, S., AND ZELEVINSKY, A. Double Bruhat cells and total positivity. *J. Amer. Math. Soc.* 12, 2 (1999), 335–380.

- [15] FOMIN, S., AND ZELEVINSKY, A. Cluster algebras. I. Foundations. *J. Amer. Math. Soc.* 15, 2 (2002), 497–529 (electronic).
- [16] FOMIN, S., AND ZELEVINSKY, A. Cluster algebras. II. Finite type classification. *Invent. Math.* 154, 1 (2003), 63–121.
- [17] FOMIN, S., AND ZELEVINSKY, A. Cluster algebras. IV. Coefficients. *Compos. Math.* 143, 1 (2007), 112–164.
- [18] FULTON, W., AND HARRIS, J. *Representation theory*, vol. 129 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991.
- [19] GEISS, C., LECLERC, B., AND SCHRÖER, J. Partial flag varieties and preprojective algebras. *Ann. Inst. Fourier (Grenoble)* 58, 3 (2008), 825–876.
- [20] GEISS, C., LECLERC, B., AND SCHRÖER, J. Preprojective algebras and cluster algebras. In *Trends in representation theory of algebras and related topics*, EMS Ser. Congr. Rep. Eur. Math. Soc., Zürich, 2008, pp. 253–283.
- [21] GEISS, C., LECLERC, B., AND SCHRÖER, J. Kac-Moody groups and cluster algebras. *Adv. Math.* 228, 1 (2011), 329–433.
- [22] GEISS, C., LECLERC, B., AND SCHRÖER, J. Cluster algebras in algebraic Lie theory. *Transform. Groups* 18, 1 (2013), 149–178.
- [23] GEISS, C., LECLERC, B., AND SCHRÖER, J. Factorial cluster algebras. *Doc. Math.* 18 (2013), 249–274.
- [24] GEKHTMAN, M., SHAPIRO, M., AND VAINSHTEIN, A. Cluster algebras and Poisson geometry. *Mosc. Math. J.* 3 (2003), 899–934, 1199. {Dedicated to Vladimir Igorevich Arnold on the occasion of his 65th birthday}.
- [25] GEL’FAND, I. M., AND ZELEVINSKIĬ, A. V. A canonical basis in irreducible representations of \mathfrak{gl}_3 and its applications. In *Group-theoretic methods in physics, Vol. 2 (Russian) (Jūrmala, 1985)*. “Nauka”, Moscow, 1986, pp. 31–45.
- [26] GROSS, M., HACKING, P., AND KEEL, S. Birational geometry of cluster algebras. *Algebr. Geom.* 2, 2 (2015), 137–175.
- [27] GROSSHANS, F. Observable groups and Hilbert’s fourteenth problem. *Amer. J. Math.* 95 (1973), 229–253.
- [28] KANG, S.-J., KASHIWARA, M., KIM, M., AND OH, S.-J. Monoidal categorification of cluster algebras. *J. Amer. Math. Soc.* 31, 2 (2018), 349–426.
- [29] KASHIWARA, M. On crystal bases of the Q -analogue of universal enveloping algebras. *Duke Math. J.* 63, 2 (1991), 465–516.
- [30] LUSZTIG, G. Canonical bases arising from quantized enveloping algebras. *J. Amer. Math. Soc.* 3, 2 (1990), 447–498.
- [31] MARSH, R. J., AND RIETSCH, K. C. Parametrizations of flag varieties. *Represent. Theory* 8 (2004), 212–242 (electronic).
- [32] MATSUMURA, H. *Commutative ring theory*, second ed., vol. 8 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1989. Translated from the Japanese by M. Reid.
- [33] MULLER, G. Locally acyclic cluster algebras. *Adv. Math.* 233 (2013), 207–247.
- [34] MURNAGHAN, F. Linear algebraic groups. In *Harmonic analysis, the trace formula, and Shimura varieties*, vol. 4 of *Clay Math. Proc.* Amer. Math. Soc., Providence, RI, 2005, pp. 379–391.
- [35] POPOV, V., AND VINBERG, E. B. *Algebraic geometry. IV*, vol. 55 of *Encyclopaedia of Mathematical Sciences*. Springer-Verlag, Berlin, 1994. Linear algebraic groups. Invariant theory, A translation of it Algebraic geometry. 4 (Russian), Akad. Nauk SSSR Vsesoyuz. Inst. Nauchn. i Tekhn. Inform., Moscow, 1989 [MR1100483 (91k:14001)], Translation edited by A. N. Parshin and I. R. Shafarevich.

-
- [36] POSTNIKOV, A. Total positivity, grassmannians, and networks, [arXiv:math/0609764](https://arxiv.org/abs/math/0609764).
 - [37] PROCESI, C. *Lie groups. An approach through invariants and representations*. Universitext. Springer, New York, 2007.
 - [38] SCOTT, J. S. Grassmannians and cluster algebras. *Proc. London Math. Soc. (3)* 92, 2 (2006), 345–380.
 - [39] SERRE, D. *Matrices*, second ed., vol. 216 of *Graduate Texts in Mathematics*. Springer, New York, 2010.
 - [40] VINBERG, È. B., AND POPOV, V. L. A certain class of quasihomogeneous affine varieties. *Izv. Akad. Nauk SSSR Ser. Mat.* 36 (1972), 749–764.
 - [41] WEYL, H. *The classical groups: Their invariants and representations*. Princeton Landmarks in Mathematics. Princeton University Press, Princeton, NJ, 1997.
 - [42] WILLENBRING, J. F., AND ZUCKERMAN, G. J. Small semisimple subalgebras of semisimple Lie algebras. In *Harmonic analysis, group representations, automorphic forms and invariant theory*, vol. 12 of *Lect. Notes Ser. Inst. Math. Sci. Natl. Univ. Singap.* World Sci. Publ., Hackensack, NJ, 2007, pp. 403–429.

Introduction to Cluster Algebras

Chapter 7

(preliminary version)

SERGEY FOMIN

LAUREN WILLIAMS

ANDREI ZELEVINSKY

Preface

This is a preliminary draft of Chapter 7 of our forthcoming textbook *Introduction to cluster algebras*, joint with Andrei Zelevinsky (1953–2013).

Other chapters have been posted as

- [arXiv:1608.05735](#) (Chapters 1–3),
- [arXiv:1707.07190](#) (Chapters 4–5), and
- [arXiv:2008.09189](#) (Chapter 6).

We expect to post additional chapters in the not so distant future.

Anne Larsen and Raluca Vlad made a number of valuable suggestions that helped improve the quality of the manuscript. We are also grateful to Zenan Fu, Amal Mattoo, Hanna Mularczyk, and Ashley Wang for their comments on the earlier versions of this chapter, and for assistance with creating figures.

Our work was partially supported by the NSF grants DMS-1664722, DMS-1854512, and DMS-1854316, and by the Radcliffe Institute for Advanced Study.

Comments and suggestions are welcome.

Sergey Fomin
Lauren Williams

Contents

Chapter 7. Plabic graphs	1
§7.1. Plabic graphs and their quivers	3
§7.2. Triangulations and wiring diagrams via plabic graphs	7
§7.3. Trivalent plabic graphs	11
§7.4. Trips. Reduced plabic graphs	16
§7.5. Triple diagrams and normal plabic graphs	23
§7.6. Minimal triple diagrams	32
§7.7. From minimal triple diagrams to reduced plabic graphs	40
§7.8. The bad features criterion	44
§7.9. Affine permutations	47
§7.10. Bridge decompositions	51
§7.11. Edge labels of reduced plabic graphs	54
§7.12. Face labels of reduced plabic graphs	58
§7.13. Grassmann necklaces and weakly separated collections	62
Bibliography	67

Plabic graphs

In this chapter, we present the combinatorial machinery of *plabic graphs*, introduced and developed by A. Postnikov [21]. These are planar (unoriented) graphs with bicolored vertices satisfying some mild technical conditions. Plabic graphs can be transformed using certain *local moves*. A key observation is that each plabic graph gives rise to a quiver, so that local moves on plabic graphs translate into (a subclass of) quiver mutations.

Crucially, the combinatorics underlying several important classes of cluster structures that arise in applications fits into the plabic graphs framework. This in particular applies to the basic examples introduced in Chapter 1. More concretely, we show that the combinatorics of flips in triangulations of a convex polygon (resp., braid moves in wiring diagrams, either ordinary or double) can be entirely recast in the language of plabic graphs. In these and other examples, an important role is played by the subclass of *reduced* plabic graphs that are analogous to—and indeed generalize—reduced decompositions in symmetric groups.

D. Thurston’s *triple diagrams* [24] are closely related to plabic graphs. After making this connection precise and developing the machinery of triple diagrams, we use this machinery to establish the fundamental properties of reduced plabic graphs.

Plabic graphs and related combinatorics have arisen in the study of shallow water waves [16, 17] (via the KP equation) and in connection with scattering amplitudes in $\mathcal{N} = 4$ super Yang-Mills theory [2]. Constructions closely related to plabic graphs were studied by T. Kawamura [13] in the context of the topological theory of graph divides.

A reader interested exclusively in the combinatorics of plabic graphs can read this chapter independently of the previous ones. While we occasionally

refer to the combinatorial constructions introduced in Chapters 1–2, they are not relied upon in the development of the general theory of plabic graphs.

Cluster algebras as such do not appear in this chapter. On the other hand, reduced plabic graphs introduced herein will prominently feature in the upcoming study of cluster structures in Grassmannians and related varieties, see Chapter 8.

The structure of this chapter is as follows.

Section 7.1 introduces plabic graphs and their associated quivers.

In Section 7.2, we recast the combinatorics of triangulations of a polygon and (ordinary or double) wiring diagrams in the language of plabic graphs.

Section 7.3 discusses the version of the theory in which all internal vertices of plabic graphs are *trivalent*. (This version naturally arises in some applications, cf., e.g., [9].) As Section 7.3 is not strictly necessary for the sections that follow, it can be skipped if desired.

The important notions of a *trip*, a *trip permutation* and a *reduced plabic graph* are introduced in Section 7.4. Here we state (but do not prove) some key results, including the “fundamental theorem of reduced plabic graphs” which characterizes the move equivalence classes of reduced plabic graphs in terms of associated *decorated permutations*.

Section 7.5 introduces the basic notions of *triple diagrams*. We then show that triple diagrams are in bijection with *normal plabic graphs*.

In Section 7.6, we study *minimal* triple diagrams, largely following [24]. These diagrams can be viewed as counterparts of reduced plabic graphs.

In Section 7.7, we explain how to go between minimal triple diagrams and reduced plabic graphs. We then use this correspondence to prove the fundamental theorem of reduced plabic graphs.

In Section 7.8, we state and prove the *bad features criterion* that detects whether a plabic graph is reduced or not.

In Section 7.9, we describe a bijection between decorated permutations and a certain subclass of *affine permutations*.

In Section 7.10, a factorization algorithm for affine permutations is used to construct a family of reduced plabic graphs called *bridge decompositions*.

Section 7.11 discusses *edge labelings* of reduced plabic graphs and gives a particularly transparent *resonance criterion* for recognizing whether a plabic graph is reduced.

Section 7.12 introduces *face labelings* of reduced plabic graphs. In Section 7.13, we provide an intrinsic combinatorial characterization of collections of face labels that arise via this construction. Face labels will reappear in Chapter 8 in the study of cluster structures in Grassmannians.

7.1. Plabic graphs and their quivers

Definition 7.1.1. A *plabic* (planar bicolored) *graph* is a planar graph G embedded into a closed disk \mathbf{D} , such that:

- the embedding of G into \mathbf{D} is proper, i.e., the edges do not cross;
- each internal vertex is colored black or white;
- each internal vertex is connected by a path to some boundary vertex;
- the (uncolored) vertices lying on the boundary of \mathbf{D} are labeled $1, 2, \dots, b$ in clockwise order, for some positive integer b ;
- each of these b *boundary vertices* is incident to a single edge.

Loops and multiple edges are allowed.

We consider plabic graphs up to isotopy of the ambient disk \mathbf{D} fixing the disk's boundary. The *faces* of G are the connected components of the complement of G inside the ambient disk \mathbf{D} . A degree 1 internal vertex which is connected by an edge to a boundary vertex is called a *lollipop*.

Two examples of plabic graphs are shown in Figure 7.1. Many more examples appear throughout this chapter. In what follows, we will often omit the boundary of the ambient disk when drawing plabic graphs.




Figure 7.1. (a) A plabic graph G . (b) A plabic graph G' with a white lollipop.

Remark 7.1.2. Plabic graphs were introduced by A. Postnikov [21, Section 12], who used them to describe parametrizations of cells in totally non-negative Grassmannians. A closely related class of graphs was defined by T. Kawamura [13] in the context of the topological theory of *graph divides*. Our definition is very close to Postnikov's.

The key role in the theory of plabic graphs is played by a particular equivalence relation that is generated by a family of transformations called (*local*) *moves*.

Definition 7.1.3. We say that two plabic graphs G and G' are *move-equivalent*, and write $G \sim G'$, if G and G' can be related to each other via a sequence of the following *local moves*, denoted (M1), (M2), and (M3):

- (M1) (The *square move*) Change the colors of all vertices on the boundary of a quadrilateral face, provided these colors alternate and these vertices are trivalent. See Figure 7.2.




Figure 7.2. Move (M1) on plabic graphs.

- (M2) Remove a bivalent vertex (of any color) and merge the edges adjacent to it; or, conversely, insert a bivalent vertex in the middle of an edge. See Figure 7.3.




Figure 7.3. Move (M2) on plabic graphs.

- (M3) Contract an edge connecting two internal vertices of the same color; or split an internal vertex into two vertices of the same color joined by an edge. See Figure 7.4.




Figure 7.4. Move (M3) on plabic graphs. The number of “hanging” edges on either side can be any nonnegative integer.

We next explain how to associate a quiver to a plabic graph. This construction is closely related to the one presented in Definition 2.5.1.

Definition 7.1.4. The *quiver* $Q(G)$ associated to a plabic graph G is defined as follows. The vertices of $Q(G)$ are in one-to-one correspondence with the faces of G . A vertex of $Q(G)$ is declared mutable or frozen depending on whether the corresponding face is *internal* (i.e., disjoint from the boundary of \mathbf{D}) or not. The arrows of $Q(G)$ are constructed in the following way. Let e be an edge in G that connects a white vertex to a black vertex and separates two distinct faces, at least one of which is internal. For each such edge e , we introduce an arrow in $Q(G)$ connecting the faces separated by e ; this arrow is oriented so that when we move along the arrow in the direction

of its orientation, we see the white endpoint of e on our left and the black endpoint on our right. We then remove oriented 2-cycles from the resulting quiver, one by one, to get $Q(G)$. See Figure 7.5.




Figure 7.5. Two plabic graphs and their associated quivers. Shown on the left is the graph G from Figure 7.1(a). The quiver on the right has double arrows, corresponding to the instances where a pair of faces share two boundary segments disconnected from each other. The frozen vertex v at the top of the picture is isolated: the two arrows between v and an internal vertex located underneath v cancel each other.

Proposition 7.1.5. *Let G and G' be two plabic graphs related to each other by one of the local moves (M1), (M2), or (M3). Additionally, if G and G' are related via the square move (M1), then we require that*

(7.1.1) *among the four faces surrounding the square, the consecutive ones must be distinct, see Figure 7.6.*

Then the quivers $Q(G)$ and $Q(G')$ are mutation equivalent.




Figure 7.6. Restriction (7.1.1) allows the square move at A —but not at B , since face C is adjacent to two consecutive sides of B .

Proof. It is straightforward to check that a square move in a plabic graph translates into a quiver mutation at the vertex associated to that square face, provided that condition (7.1.1) is satisfied. It is also straightforward to check that the quiver associated with a plabic graph does not change under moves (M2) or (M3), see Figure 7.7. \square




Figure 7.7. Fragments of plabic graphs and their associated quivers. The first two plabic graphs are related by a square move (M1); their quivers are related by a single mutation. The second and the third graphs are related by moves of type (M3), and have isomorphic quivers.

Remark 7.1.6. Suppose condition (7.1.1) fails at a square face B , with B incident to another face C along two consecutive edges, as in Figure 7.6. Then the arrows transversal to these edges cancel each other, so they do not appear in the associated quiver. This leads to a discrepancy between the square move and the quiver mutation.

Remark 7.1.7. The key difference between the setting of this section *vs.* Section 2.5 is that here we do not require plabic graphs to be bipartite. This distinction is not particularly important, since we can always apply a sequence of moves (M2) to a plabic graph to make it bipartite.

In the bipartite setting, the square move (M1) corresponds to urban renewal, see Definition 2.5.2 as well as Definition 7.5.13 below.

Remark 7.1.8. In light of Proposition 7.1.5, one may choose to adjust the definition of the square move (M1)—hence the notion of move equivalence of plabic graphs—by forbidding square moves violating condition (7.1.1). (This convention was adopted in [9].) We note that for the important subclass of *reduced* plabic graphs (see Definition 7.4.2 below), condition (7.1.1) is automatically satisfied, so there is no need to worry about it.

Remark 7.1.9. Using Definition 7.1.4, we can associate a seed pattern—hence a cluster algebra—to any plabic graph G . By Proposition 7.1.5, this cluster algebra only depends on the move equivalence class of G , assuming that we adopt a restricted notion of move equivalence, cf. Remark 7.1.8. We will soon see that this family of cluster algebras includes all the main examples of cluster algebras (defined by quivers) introduced in the earlier chapters. This justifies the importance of the combinatorial study of plabic graphs, and in particular their classification up to move equivalence.

The quivers arising from plabic graphs are rather special. In particular, each quiver $Q(G)$ is planar.

On the other hand, the mutation class of any quiver without frozen vertices can be embedded into the mutation class of a quiver of a plabic graph:

Proposition 7.1.10 ([8]). *Let Q be a quiver whose vertices are all mutable. Then there exists a plabic graph G such that Q is a full subquiver (see Definition 4.1.1) of a quiver mutation-equivalent to $Q(G)$.*

7.2. Triangulations and wiring diagrams via plabic graphs

In this section, we explain how the machinery of local moves on plabic graphs unifies the combinatorial constructions of Chapter 2, including:

- flips in triangulations of a polygon (Section 2.2);
- braid moves for wiring diagrams (Section 2.3);
- their analogues for double wiring diagrams (Section 2.4).

As mentioned in Remark 7.1.7, the urban renewal transformations (Section 2.5) can also be interpreted in terms of local moves in plabic graphs.

Example 7.2.1 (Triangulations of a polygon, see [16, Algorithm 12.1]). Let T be a triangulation of a convex m -gon \mathbf{P}_m . The plabic graph $G(T)$ associated to T is constructed as follows:

- (1) Place a white vertex of $G(T)$ at each vertex of \mathbf{P}_m .
- (2) Place a black vertex of $G(T)$ in the interior of each triangle of T . Connect it by edges to the three white vertices of the triangle.
- (3) Embed \mathbf{P}_m into the interior of a disk \mathbf{D} .
- (4) Place m uncolored vertices of $G(T)$ on the boundary of \mathbf{D} .
- (5) Connect each white vertex of $G(T)$ to a boundary vertex. These edges must not cross.

We emphasize that the set of edges of $G(T)$ includes neither the sides of \mathbf{P}_m nor the diagonals of T . See Figure 7.8.




Figure 7.8. A triangulation T of an octagon, and the corresponding plabic graph $G(T)$, cf. Figure 2.2.

Exercise 7.2.2. Show that $Q(G(T)) = Q(T)$, i.e., the quiver associated to the plabic graph of a triangulation T coincides with the quiver $Q(T)$ associated to T , as in Definition 2.2.1.

Exercise 7.2.3. Show that if triangulations T and T' are related by a flip, then the plabic graphs $G(T)$ and $G(T')$ are move-equivalent to each other.

More concretely, flipping a diagonal in T translates into a square move at the corresponding quadrilateral face of $G(T)$, plus some (M3) moves to make each vertex of that face trivalent.

Example 7.2.4 (Wiring diagrams). Let D be a wiring diagram, as in Section 1.3. We associate a plabic graph $G(D)$ to D by replacing each crossing in D by a pair of trivalent vertices connected vertically, with a black vertex on top and a white vertex on the bottom. We then enclose the resulting graph in a disk.

This construction applies to a more general version of wiring diagrams. Let s_i denote the simple transposition in the symmetric group \mathcal{S}_n that exchanges i and $i + 1$. Given a sequence $\mathbf{w} = s_{i_1}s_{i_2}\dots s_{i_m}$ of simple transpositions, we associate to it a diagram $D(\mathbf{w})$ by concatenating m graphs; here the graph associated to s_j consists of n wires, of which $n - 2$ are horizontal, while the j th and $(j + 1)$ st wires cross over each other. See Figure 7.9.




Figure 7.9. Top: the wiring diagrams D_1 and D_2 associated to reduced expressions $s_2s_3s_2s_1s_2s_3$ and $s_3s_2s_3s_1s_2s_3$ for $w_0 = (4, 3, 2, 1) \in S_4$. These wiring diagrams (resp., reduced expressions) are related via a braid move. Middle: the plabic graphs $G(D_1)$ and $G(D_2)$. Bottom: the quivers $Q(G(D_1))$ and $Q(G(D_2))$, with isolated frozen vertices removed.

Exercise 7.2.5. Show that after removing isolated frozen vertices at the top and bottom, the quiver $Q(G(D))$ associated to the plabic graph of a wiring diagram D coincides with the quiver $Q(D)$ associated to D , as in Definition 2.3.1, up to a global reversal of arrows.

Remark 7.2.6. If we changed our convention in Example 7.2.4, swapping the colors of the black and white vertices, we'd recover precisely the quiver $Q(D)$ associated to the wiring diagram. However, we prefer the convention used in Example 7.2.4 because it will lead to a transparent algorithm for recovering the chamber minors, as shown in Figure 7.62. And as noted in Remark 3.1.10, the cluster algebra associated to a given quiver is the same as the cluster algebra associated to the opposite quiver.

Remark 7.2.7. If two wiring diagrams D and D' are related by a braid move, then the corresponding plabic graphs $G(D)$ and $G(D')$ are related by a square move plus some (M3) moves, see Figure 7.10.




Figure 7.10. A braid move on wiring diagrams, and a corresponding sequence of moves on plabic graphs. The first two (resp., the last two) plabic graphs are related by two (M3) moves; the two plabic graphs in the middle are related by an (M1) move.

Example 7.2.8 (Double wiring diagrams). Let D be a double wiring diagram, as in Section 1.4. The plabic graph $G(D)$ associated to D is defined by adjusting the construction of Example 7.2.4 in the following way: as before, we replace each crossing in the double wiring diagram by a pair of trivalent vertices connected vertically, and color one of these vertices white and the other black. If the crossing is thin, the top vertex gets colored white and the bottom one black; if the crossing is thick, the colors of the two vertices are reversed. See Figure 7.11.

As in Example 7.2.4, the above construction works for a more general version of double wiring diagrams. Given two sequences $\mathbf{w} = s_{i_1} s_{i_2} \dots s_{i_m}$ and $\bar{\mathbf{w}} = \bar{s}_{j_1} \bar{s}_{j_2} \dots \bar{s}_{j_\ell}$, choose an arbitrary shuffle of \mathbf{w} and $\bar{\mathbf{w}}$. Then we can

associate a (generalized) double wiring diagram to this shuffle, where thick crossings are associated to factors in \mathbf{w} , and thin crossings are associated to factors in $\overline{\mathbf{w}}$. So, e.g., the double wiring diagram in Figure 7.11 is associated to the shuffle $\overline{s_2}s_1s_2\overline{s_1}\overline{s_2}s_1$.




Figure 7.11. A double wiring diagram D , the corresponding plabic graph $G(D)$, and the quiver associated to $G(D)$. If one removes the bottom frozen vertex, one recovers the quiver from Figure 2.6 (up to a global reversal of arrows).

Exercise 7.2.9. Extend the statements of Exercise 7.2.5 and Remark 7.2.7 to the case of double wiring diagrams.

In addition to triangulations and (ordinary or double) wiring diagrams, plabic graphs can also be used to describe Fock-Goncharov cluster structures [7]:

Exercise 7.2.10. Construct a plabic graph whose associated quiver is the quiver shown in Figure 2.3. How does this construction generalize to a quiver $Q_3(T)$ associated to an arbitrary triangulation T of a convex polygon, cf. Exercise 2.2.3?

7.3. Trivalent plabic graphs

In Section 7.1, we introduced plabic graphs and described local moves that generate an equivalence relation on them. In this section, we focus on *trivalent plabic graphs*, i.e., those plabic graphs whose interior vertices are all trivalent. This will require working with an alternative set of moves that preserve the property of being trivalent.

Remark 7.3.1. Trivalent plabic graphs arise naturally in the studies of

- soliton solutions to the KP equation [16],
- sections of fine zonotopal tilings of 3-dimensional cyclic zonotopes [10],
- combinatorics of planar divides and associated links [9], and
- π -induced subdivisions for a projection π from the hypersimplex $\Delta_{k,n}$ to an n -gon [22].

We begin with the following simple observation.

Lemma 7.3.2. *Any plabic graph with no interior leaves (i.e., no degree 1 interior vertices) can be transformed by a sequence of moves of type (M2) and/or (M3) into a plabic graph all of whose interior vertices are trivalent.*

Proof. We can get rid of bivalent vertices using the moves (M2). If there are any vertices of degree ≥ 4 , split those vertices using (M3) until all internal vertices are trivalent. \square

The alternative set of moves for trivalent plabic graphs consists of the square move (M1) together with the *flip move* (M4) defined below.

Definition 7.3.3. The *flip move* (sometimes also called the *Whitehead move*) for trivalent plabic graphs is defined as follows:

(M4) Replace a fragment containing two trivalent vertices of the same color connected by an edge by another such fragment, see Figure 7.12.




Figure 7.12. The flip move (or Whitehead move) for trivalent graphs. The four vertices shown should either all be white, or all be black.

Remark 7.3.4. A flip move (M4) can be expressed as a composition of two moves of type (M3).

The main result of this section is the following.

Theorem 7.3.5. *Two trivalent plabic graphs G and G' are related via a sequence of local moves of types (M1), (M2), and (M3) if and only if they are related by a sequence of moves of types (M1) and (M4).*

The rest of this section is devoted to the proof of Theorem 7.3.5.

Lemma 7.3.6. *If two trivalent plabic graphs G and G' are connected by a sequence of moves of type (M2) or (M3), then they are connected by a sequence of moves of type (M3).*

Proof. We first note that in many cases, move (M2) can be thought of as an instance of move (M3): instead of using (M2) to add or remove a bivalent vertex that is adjacent via an edge e to a vertex of the same color, we can use (M3) to (un)contract e , to the same effect.

The only (M2) moves that are genuinely different from (M3) moves are the (M2) moves that add or remove a white (resp., black) vertex in the middle of a black-black or black-boundary (resp., white-white or white-boundary) edge. We call them *creative* or *destructive* (M2) moves, see the left and middle of Figure 7.13. We need to show that it is never necessary to use a creative or destructive (M2) move to connect G and G' as above.




Figure 7.13. Using a creative (M2) move and then an (M3) move to grow a white leaf.

Consider a shortest sequence of (M2)/(M3) moves connecting G and G' :

$$(7.3.1) \quad G = G_0 \sim G_1 \sim \cdots \sim G_k = G'.$$

We first claim that each G_i does not contain an internal leaf. Suppose otherwise. Let $G_{i-1} \sim G_i$ be the last step when we grow a leaf (as in Figure 7.13). That is, i is maximal such that G_i has a leaf v' attached via edge e to some vertex v , where v' and e were not present in G_{i-1} . Since G' has no leaves, there must be some $j > i$ such that G_j is obtained from G_{j-1} by contracting the edge adjacent to v' . By construction, each graph $G_i, G_{i+1}, \dots, G_{j-1}$ contains leaf v' along with a path connecting v' to a vertex of the same color; this path is obtained by subdividing the edge e by adding bivalent vertices (which must remain bivalent since no new leaves may be added). But then we can construct a shorter sequence of moves connecting G to G' by removing all steps used to create and contract this path.

Having established that none of the graphs appearing on the shortest path (7.3.1) have internal leaves, we will now demonstrate that this path never uses a creative or destructive (M2) move. Note that we cannot have only destructive (M2) moves because (M3) moves alone cannot create a vertex whose removal requires a destructive (M2) move. Thus, it is enough to show that we cannot have a creative (M2) move along (7.3.1). Suppose there is one, and that the last creative (M2) move adds a bivalent white vertex w along a black-black edge. Along (7.3.1), this bivalent vertex w might split into multiple bivalent white vertices (pairwise adjacent, in a row). However, all these vertices must get removed somewhere along (7.3.1), since G' has no bivalent vertices. We can then shorten (7.3.1) by removing all moves involving these bivalent vertices. The lemma is proved. \square

Lemma 7.3.7. *Let G and G' be two trivalent plabic graphs such that*

- *each of the graphs G and G' is connected;*
- *each of the graphs G and G' has f interior faces, b boundary vertices, and b boundary faces (the number of boundary vertices equals the number of boundary faces since the graphs are connected);*
- *in each of the graphs G and G' , all interior vertices have the same color, and this color is the same in both graphs.*

Then G and G' can be connected by a sequence of flip moves (M4).

Proof. The *dual graph* G_{dual} of a trivalent connected plabic graph G is obtained as follows. Place a vertex of G_{dual} in the interior of each face of G . For each edge e of G , introduce a (transversal) edge of G_{dual} connecting the vertices of G_{dual} located in the faces of G on both sides of e , see Figure 7.14. (This new edge may be a loop.) Under the conditions of the lemma, the dual graph G_{dual} is a generalized triangulation T of a (dual) b -gon. (We note that T may contain *self-folded* triangles – loops with an interior “pendant” edge – coming from the faces of G enclosed by a loop in G , as in Figure 7.14.) The triangulation T has $b + f$ vertices: the b vertices of the dual b -gon together with the f interior points (“punctures”).




Figure 7.14. A trivalent plabic graph G and its dual graph G_{dual} (in red); the latter contains a self-folded triangle. Here $b = 3$ and $f = 1$.

Figure 7.15 shows that a flip move in a trivalent plabic graph corresponds to a flip in the corresponding triangulation. The claim that G and G' are connected by flip moves can now be obtained from the well-known fact [11, 12] that any two triangulations of a b -gon with f interior points are connected by flips. \square




Figure 7.15. The flip move for trivalent graphs corresponds to a flip of the corresponding dual triangulations.

Definition 7.3.8. A *white component* W of a plabic graph G is obtained by taking a maximal (by inclusion) connected induced subgraph of G all of whose internal vertices are white, together with the half-edges extending from the (white) vertices of W towards black vertices outside W or towards boundary vertices of G . *Black components* of G are defined in the same way, with the roles of black and white vertices reversed.

Remark 7.3.9. Each black or white component C of a plabic graph G can itself be regarded as a (generalized) plabic graph. To this end, enclose C by a simple closed curve γ passing through the endpoints of the half-edges on the outer boundary of C . If the portion of G located inside γ is exactly C , then we get a usual plabic graph. It may however happen that C contains “holes,” i.e., some of the half-edges on the boundary of C may be entirely contained in the interior of the disk enclosed by γ . In that case, we need to draw simple closed curves through the endpoints of those half-edges, so that C becomes a generalized plabic graph inside a “swiss-cheese” shape (a disk with some smaller disks removed), as in Figure 7.16. The argument in the proof of Lemma 7.3.7 extends to this setting, so Lemma 7.3.7 also holds for black/white components of trivalent plabic graphs. We will use this generalization in the proof of Proposition 7.3.10 below.

Proposition 7.3.10. *Let two trivalent plabic graphs be related to each other by moves (M2) or (M3). Then they are related by a series of flip moves (M4).*

Proof. Let G and G' be the plabic graphs in question. Without loss of generality we may assume that G and G' are connected. By Lemma 7.3.6, G and G' are connected by moves (M3).

Each of the graphs G and G' breaks into disjoint (white or black) components. Each (M3) move only affects a single component. It follows that the white (resp., black) components W_1, \dots, W_ℓ (resp., B_1, \dots, B_m) of G




Figure 7.16. A generalized plabic graph inside a “swiss-cheese” shape (in this case, a disk with two smaller disks removed).

are in bijection with the components W'_1, \dots, W'_ℓ (resp., B'_1, \dots, B'_m) of G' , so that each W_i (resp., B_j) is related to W'_i (resp., B'_j) via (M3) moves.

Since an (M3) move preserves both the number of boundary vertices and the number of faces of a graph, both W_i and W'_i (respectively, B_j and B'_j) have the same number of boundary vertices and the same number of faces. It now follows from Lemma 7.3.7 (more precisely, from its extension to components of plabic graphs, see Remark 7.3.9) that each pair W_i and W'_i can be connected by flip moves, and similarly for B_j and B'_j . The proposition follows. \square

Proof of Theorem 7.3.5. The “if” direction immediately follows from Remark 7.3.4.

Suppose that G and G' are related via a sequence of (M1), (M2), and (M3) moves. Let k denote the number of square moves (M1) in the sequence. We then have a sequence of move-equivalences

$$G = G'_0 \sim G_1 \sim G'_1 \sim G_2 \sim G'_2 \sim \dots \sim G_k \sim G'_k \sim G_{k+1} = G',$$

where for all i ,

- G_i is related to G'_i by a single square move;
- G'_i is related to G_{i+1} by a sequence of (M2) and (M3) moves.

Since a square move only involves trivalent vertices, we may assume, applying extra (M2) and (M3) moves as needed, that all plabic graphs G_i and G'_i are trivalent. It then follows by Proposition 7.3.10 that for every i , the graphs G'_i and G_{i+1} are related by flip moves alone, and we are done. \square

Remark 7.3.11. The plabic graphs associated to wiring diagrams and double wiring diagrams as in Example 7.2.4 and Example 7.2.8 are trivalent, and consequently one can express the transformation corresponding to braid moves using square moves and flip moves, as shown in Figure 7.10.

On the other hand, the plabic graphs associated to triangulations of a polygon (see Example 7.2.1 and Figure 7.8) are *not* trivalent.

7.4. Trips. Reduced plabic graphs

Reduced plabic graphs are a subclass of plabic graphs that play a critically important role in the study of cluster structures and total positivity in Grassmannians. In this section, we introduce reduced plabic graphs and various related concepts. The main result of the section, Theorem 7.4.25, provides a criterion for move equivalence of reduced plabic graphs. The proof of this theorem will occupy Sections 7.5–7.7.

Definition 7.4.1. We say that T is a *collapsible tree* in a plabic graph G if

- T is a tree with at least one edge not incident to the boundary of G ;
- T is an induced subgraph of G ;
- T is attached to the rest of G at a single vertex v , the *root* of T , so that
 - if v is a boundary vertex, one can use local moves (M2)–(M3) to collapse T to a lollipop based at v .
 - if v is an internal vertex, one can use local moves (M2)–(M3) to collapse T onto the vertex v .

In other words, local moves can be used to replace the tree T by either a lollipop based at v , or by v . See Figure 7.17.




Figure 7.17. A collapsible tree T whose root is the internal vertex v .

Definition 7.4.2. A plabic graph G is *reduced* if no plabic graph $G' \sim G$ contains one of the following “forbidden configurations:”

- a *hollow digon*, i.e., two edges connecting a pair of distinct vertices, with no other edges in the region between them; or
- an internal leaf that is not a lollipop and does not belong to a collapsible tree.

See Figure 7.18 for a slightly modified (but equivalent) version.

Remark 7.4.3. For plabic graphs, the property of being reduced is, by definition, invariant under local moves.

Remark 7.4.4. As stated, this property is not readily testable, because the move equivalence class of a plabic graph is usually infinite. We will later obtain criteria (see Theorems 7.8.6 and 7.11.5) for testing this property.




Figure 7.18. A plabic graph is not reduced if and only if it is move-equivalent to a graph containing a hollow digon, as in (a,b,c), or a graph containing an internal leaf u adjacent to a trivalent vertex v of opposite color such that v is not a root of a collapsible tree, see (d,e).

Remark 7.4.5. The notion of a reduced plabic graph, as introduced above in Definition 7.4.2, may appear artificial, as the choice of forbidden configurations is not well motivated. We will soon present alternative definitions of reducedness (in slightly restricted generality) that are both more conceptual and more elegant; see Proposition 7.4.9 and Corollary 7.4.26. In particular, these alternative definitions do not involve the forbidden configurations in Figure 7.18(d,e).

Remark 7.4.6. The requirement in Figure 7.18(d,e) concerning the collapsible tree cannot be removed: as shown in Figure 7.19, dropping this requirement would make *any* plabic graph non-reduced.




Figure 7.19. Any plabic graph can be transformed via moves (M2)–(M3) into a graph containing a configuration as in Figure 7.18(d,e) in which vertex v is the root of a collapsible tree that contains vertex u .

Reduced plabic graphs can be interpreted as generalizations of reduced expressions in symmetric groups:

Exercise 7.4.7. Consider sequences (or “words”) $\mathbf{w} = s_{i_1} \cdots s_{i_m}$ of simple transpositions in a symmetric group, as in Example 7.2.4. We say that two such words are *braid-equivalent* if they can be related to each other using the braid relations $s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1}$ and $s_i s_j = s_j s_i$ for $|i - j| \geq 2$. A word \mathbf{w} is called a *reduced expression* if no word in its braid equivalence class has two consecutive equal entries: $\cdots s_i s_i \cdots$. Show that if \mathbf{w} fails to be a reduced expression, then $G(D(\mathbf{w}))$ fails to be a reduced plabic graph.

Remark 7.4.8. Conversely, if \mathbf{w} is reduced, then $G(D(\mathbf{w}))$ is reduced; this can be proved using Theorem 7.8.6 or Theorem 7.11.5.

Proposition 7.4.9. *Let G be a plabic graph that has no internal leaves, other than lollipops. Then the following are equivalent:*

- (i) G is not reduced;
- (ii) G can be transformed, via local moves that do not create internal leaves, into a plabic graph containing a hollow digon, cf. Figure 7.18(a,b,c).

We note that in the absence of internal leaves, the only forbidden configurations are the hollow digons.

The proof of Proposition 7.4.9 will require some technical preparations, see Definition 7.4.10 and Lemma 7.4.11 below.

Definition 7.4.10. We denote by \overline{G} the plabic graph obtained from G by repeatedly collapsing all collapsible trees. It is not hard to see that \overline{G} is uniquely defined.

Lemma 7.4.11. *Let G be a plabic graph such that \overline{G} has no internal leaves, other than lollipops. Let G' be a plabic graph obtained from G by a local move. Assume that this move is not a square move (M1) where one of the vertices of the square (in G) is the root of a collapsible tree. Then $\overline{G'}$ has no internal leaves, other than lollipops. Moreover, \overline{G} and $\overline{G'}$ are either equal to each other or related by a single local move.*

Proof. If the changes resulting from the local move $G \rightarrow G'$ occur within a collapsible tree, then the tree remains collapsible, so $\overline{G'} = \overline{G}$. Otherwise, the same move can be applied in \overline{G} , yielding $\overline{G'}$ (and not creating any leaves). \square

Proof of Proposition 7.4.9. The implication (ii) \Rightarrow (i) is immediate. Let us establish (i) \Rightarrow (ii). Assume that G is not reduced. Then there exists a sequence of plabic graphs

$$(7.4.1) \quad G = G_0, G_1, \dots, G_N$$

in which each pair (G_i, G_{i+1}) is related by a local move and moreover G_N contains a forbidden configuration from Figure 7.18.

Case 1: the sequence (7.4.1) does not include a square move $G_k \xrightarrow{(M1)} G_{k+1}$ where one of the vertices of the square in G_k is the root of a collapsible tree. Repeatedly applying Lemma 7.4.11, we conclude that the graphs $\overline{G_i}$ do not contain internal leaves, other than lollipops. Moreover, for each i , the plabic graphs $\overline{G_i}$ and $\overline{G_{i+1}}$ either coincide or are related via a single local move. It is furthermore easy to see that since G_N contains a forbidden configuration, then the same must be true for $\overline{G_N}$. That is, $\overline{G_N}$ contains a hollow digon. We conclude that $G = G_0 = \overline{G_0}$ is connected by local moves that do not create internal leaves to a plabic graph $\overline{G_N}$ containing a hollow digon.

Case 2: the sequence (7.4.1) includes a square move $G_k \xrightarrow{(M1)} G_{k+1}$ in which one of the vertices of the square in G_k is the root of a collapsible tree. Let (G_k, G_{k+1}) be the first such occurrence (i.e., the one with the smallest k). Repeatedly applying Lemma 7.4.11, we conclude that $G = \overline{G_0}$ is related to $\overline{G_k}$ via local moves that do not create internal leaves. Moreover, $\overline{G_k}$ contains

a square configuration in which one of the vertices of the square is bivalent. Removing this vertex using (M2) and contracting the resulting edge yields a bicolored hollow digon, and we are done. \square

The most fundamental result concerning reduced plabic graphs is their classification up to move equivalence (cf. Remark 7.4.3), to be given in Theorem 7.4.25 below. To state this result, we will need some preparation.

Definition 7.4.12. A *trip* τ in a plabic graph G is a directed walk along the edges of G that either begins and ends at boundary vertices (with all intermediate vertices internal), or is a closed walk entirely contained in the interior of the disk, which obeys the following *rules of the road*:

- at a black vertex, τ always makes the sharpest possible right turn;
- at a white vertex, τ always makes the sharpest possible left turn.

If the trip begins and ends at boundary vertices we call it a *one-way trip*; if it is a closed walk entirely contained in the interior of the disk, we call it a *roundtrip*.

The endpoints of a one-way trip may coincide with each other. For example, the trip corresponding to a lollipop rooted at a vertex i starts and ends at i .

Remark 7.4.13. Just as different countries have different rules regarding which side of the road one should drive on, different authors make conflicting choices for the rules of the road for plabic graphs. In this book, we consistently use the convention chosen in Definition 7.4.12.

Remark 7.4.14. The notion of a trip and the condition of being reduced have appeared in the study of dimer models in statistical mechanics, wherein trips have been called *zigzag paths* [14]. Reduced plabic graphs were called “marginally geometrically consistent” in [4, Section 3.4], and were said to “obey condition Z” in [3, Section 8].

Remark 7.4.15. For any edge e in G , there is a unique trip traversing e in each of the two directions. It may happen that the same trip traverses e twice (once in each direction).

Exercise 7.4.16. Show that one-way trips starting at different vertices terminate at different vertices.

Exercise 7.4.17. Let $G(D)$ be a plabic graph associated to some wiring diagram D , see Example 7.2.4 and Figure 7.9. Show that the trips starting at the left side of $G(D)$ follow the pattern determined by the strands of D , while the trips starting at the right side of $G(D)$ proceed horizontally to the left. Describe the trips in a plabic graph associated to a double wiring diagram.

Definition 7.4.18. Let G be a plabic graph with b boundary vertices. The *trip permutation* $\pi_G : \{1, \dots, b\} \rightarrow \{1, \dots, b\}$ is defined by setting $\pi_G(i) = j$ whenever the trip originating at i terminates at j . We will mostly use the one-line notation $\pi_G = (\pi_G(1), \dots, \pi_G(b))$ to represent these permutations.

To illustrate, in Figure 7.1(a), we have $\pi_G = (3, 4, 5, 1, 2)$.

Exercise 7.4.19. Show that move-equivalent plabic graphs have the same trip permutation.

The notion of a trip permutation can be further enhanced to construct finer invariants of local moves. For example, we can record, in addition to the trip permutation, the suitably defined *winding number* of each trip. These winding numbers do not change under local moves (with one subtle exception, cf. Figure 7.34 below). A more powerful invariant associates to any plabic graph a particular (transverse) link, see [9].

Definition 7.4.20. A *decorated permutation* $\tilde{\pi}$ on b letters is a permutation of the set $\{1, \dots, b\}$ together with a *decoration* of each fixed point by either an overline or an underline. In other words, for every i , we have

$$\tilde{\pi}(i) \in \{\bar{i}, \underline{i}\} \cup \{1, \dots, b\} \setminus \{i\}.$$

An example of a decorated permutation on 6 letters is $(3, 4, 5, 1, 2, \bar{6})$.

Exercise 7.4.21. Show that the number of decorated permutations on b letters is equal to $b! \sum_{k=0}^b \frac{1}{k!}$.

The following statement will be proved in Section 7.7.

Proposition 7.4.22. *Let G be a reduced plabic graph. If $\pi_G(i) = i$, then the connected component of G containing the boundary vertex i collapses to a lollipop.*

In Proposition 7.4.22, the requirement that G is reduced cannot be dropped, see Figure 7.20.




Figure 7.20. A non-reduced plabic graph G with $\pi_G(i) = i$, cf. Figure 7.18(d). The component containing i is not collapsible.

Definition 7.4.23. Let G be a reduced plabic graph with b boundary vertices. The *decorated trip permutation* associated with G is defined as follows:

$$\tilde{\pi}_G(i) = \begin{cases} \pi_G(i) & \text{if } \pi_G(i) \neq i; \\ \bar{i} & \text{if } G \text{ contains a tree collapsing to a white lollipop at } i; \\ \underline{i} & \text{if } G \text{ contains a tree collapsing to a black lollipop at } i. \end{cases}$$

Figure 7.21 shows two reduced plabic graphs with the same decorated trip permutation $(3, 4, 5, 1, 2, \bar{6})$.




Figure 7.21. Two reduced plabic graphs sharing the same decorated trip permutation $(3, 4, 5, 1, 2, \bar{6})$. Cf. Figure 7.1(b).

Exercise 7.4.19 can be strengthened as follows.

Exercise 7.4.24. The decorated trip permutation of a reduced plabic graph is invariant under local moves.

We will later show (see Corollary 7.10.4) that for each decorated permutation $\tilde{\pi}$ on b letters, there exists a reduced plabic graph whose decorated trip permutation is $\tilde{\pi}$.

Crucially, the move-equivalence class of a reduced plabic graph is completely determined by its decorated trip permutation:

Theorem 7.4.25 (Fundamental theorem of reduced plabic graphs). *Let G and G' be reduced plabic graphs. The following statements are equivalent:*

- (1) G and G' are move-equivalent;
- (2) G and G' have the same decorated trip permutation.

To illustrate, the two reduced plabic graphs shown in Figure 7.21 have the same decorated trip permutation and consequently are move-equivalent.

The statement (1) \Rightarrow (2) in Theorem 7.4.25 is easy, cf. Exercise 7.4.24. The converse implication (2) \Rightarrow (1) is much harder. In Section 7.7, we give a proof of this implication that utilizes D. Thurston's machinery of triple diagrams, which is presented in Sections 7.5–7.6.

A very intricate argument justifying the implication (2) \Rightarrow (1) was given in A. Postnikov's original preprint [21, Section 13]. Another proof of Theorem 7.4.25, involving some difficult results about *plabic tilings* (and relying on Theorem 7.13.4 below), was given by S. Oh and D. Speyer [20].

Corollary 7.4.26. *Let π be a permutation on b letters. Consider all plabic graphs G without internal leaves (other than lollipops) whose trip permutation is π ; in particular, G has b boundary vertices. Among all such plabic graphs G , the reduced ones are precisely those that have the smallest number of faces.*

Proof. Local moves do not change the number of faces. It follows by Theorem 7.4.25 that all reduced plabic graphs with a given decorated trip permutation have the same number of faces.

Changing the color of a lollipop transforms a reduced plabic graph into another reduced graph with the same number of faces and the same trip permutation (but with different decoration). Therefore all reduced plabic graphs G with $\pi(G) = \pi$ have the same number of faces.

It remains to show that if G is not reduced and has no internal leaves other than lollipops, then there exists a plabic graph G' with $\pi(G') = \pi$ and with fewer faces than G . We note that under our assumptions on G , Proposition 7.4.9 applies, so G can be transformed by local moves that do not create internal leaves into a plabic graph G'' containing a hollow digon. We now claim that G'' can be replaced by a plabic graph G''' (not move-equivalent to G'') such that G'' and G''' have the same trip permutation but G''' has fewer faces than G'' . The recipe for constructing G''' is as follows. If the vertices of the hollow digon in G'' are of the same color, then remove one of the sides of the digon (keeping its vertices) to get G''' . If, on the other hand, the vertices of the digon have different colors, then remove both sides of the digon; if one of the vertices was bivalent, then remove it as well. It is straightforward to check that in each case, the trip permutation does not change whereas the number of faces decreases by 1 or 2. \square

In Corollary 7.10.5, we will give a formula for the number of faces in a reduced plabic graph in terms of the associated decorated trip permutation.

Remark 7.4.27. In Corollary 7.4.26, the requirement that G has no internal leaves cannot be dropped. For example, the graph in Figure 7.20 has a single face but is not reduced.

Remark 7.4.28. Some authors call a plabic graph reduced if it has the smallest number of faces among all graphs with a given decorated trip permutation, cf. Corollary 7.4.26. If one adopts this definition, then the graph G in Figure 7.20 becomes reduced. This leads to a failure of Proposition 7.4.22, since the component of G attached to the boundary vertex i does not collapse to a lollipop.

Another reason to treat this kind of plabic graph G as non-reduced will arise in the context of triple diagrams, cf. Section 7.7. While reduced plabic graphs should correspond to minimal triple diagrams (see Theorem 7.7.3), the triple diagram corresponding to G is not minimal.

7.5. Triple diagrams and normal plabic graphs

Triple diagrams (or triple crossing diagrams), introduced by D. Thurston in [24], are planar topological gadgets closely related to plabic graphs. Our treatment of triple diagrams in Sections 7.5–7.6 is largely based on [24].

Definition 7.5.1. Consider a collection \mathfrak{X} of oriented intervals and/or circles immersed into a disk \mathbf{D} . Their images are collectively called *strands*. The images of immersed intervals (resp., circles) are the *arcs* (resp., *closed strands*) of \mathfrak{X} . A *face* of \mathfrak{X} is a connected component of the complement of the union of the strands within \mathbf{D} . We call \mathfrak{X} a *triple diagram* if

- the preimage of each point in \mathbf{D} consists of either 0, 1, or 3 points; in the latter case of a *triple point*, the three local branches intersect transversally in the interior of \mathbf{D} ;
- the endpoints of arcs are distinct points located on the boundary $\partial\mathbf{D}$; each arc meets the boundary transversally;
- the union of the strands and the boundary of the disk is connected; this ensures that each face is homeomorphic to an open disk;
- the strand segments lying on the boundary of each face are oriented consistently (i.e., clockwise or counterclockwise); in particular, as we move along the boundary, the *sources* (endpoints where an arc runs away from $\partial\mathbf{D}$) alternate with the *targets* (where an arc runs into $\partial\mathbf{D}$).

Each triple diagram, say with b arcs, comes with a selection of b points on $\partial\mathbf{D}$ (called *boundary vertices*) labeled $1, \dots, b$ in clockwise order. There is one such boundary vertex within every other segment of $\partial\mathbf{D}$ between two consecutive arc endpoints. Specifically, we place boundary vertices so that, moving clockwise along the boundary, each boundary vertex follows (resp., precedes) a source (resp., a target). See Figure 7.22.




Figure 7.22. A triple diagram \mathfrak{X} with the strand permutation $\pi_{\mathfrak{X}} = (3, 4, 5, 1, 2)$.

Triple diagrams are considered up to smooth isotopy among such diagrams. This makes them essentially combinatorial objects: 6-valent/univalent directed graphs with some additional structure.

Remark 7.5.2. In [24], the definition of a triple diagram does not include the restriction appearing in Definition 7.5.1 that requires the union of the strands and the boundary $\partial\mathbf{D}$ to be connected. In the terminology of [24], all our triple diagrams are *connected*.

Remark 7.5.3. In order to ensure consistent orientations along the face boundaries, the orientations of strands must alternate between “in” and “out” around each triple point. Given a triple diagram with unoriented strands, we can satisfy this condition as follows: start anywhere and propagate out by assigning alternating orientations around vertices.

Definition 7.5.4. Let \mathfrak{X} be a triple diagram with b boundary vertices (hence b arcs). For each boundary vertex i , let s_i (resp., t_i) denote the source (resp., target) arc endpoint located next to i on the boundary of \mathbf{D} .

The *strand permutation* $\pi_{\mathfrak{X}}$ is defined by setting $\pi_{\mathfrak{X}}(i) = j$ whenever the arc originating at s_i ends up at t_j . Thus, the strand permutation describes the connectivity of the arcs. See Figure 7.22.

We will soon see (cf. Definition 7.6.8 below) that any permutation can arise as a strand permutation of a triple diagram.

Just as the local moves on plabic graphs preserve the (decorated) trip permutation (see Exercise 7.4.24), there is a notion of a local move on triple diagrams that keeps the strand permutation invariant.

Definition 7.5.5. We say that two triple diagrams \mathfrak{X} and \mathfrak{X}' are *move-equivalent* to each other, and write $\mathfrak{X} \sim \mathfrak{X}'$, if one can get between \mathfrak{X} and \mathfrak{X}' via a sequence of *swivel moves* shown in Figure 7.23. (These moves are called $2 \leftrightarrow 2$ moves in [24].)




Figure 7.23. The swivel move replaces one of these fragments of a triple diagram by the other fragment, then smooths out the strands. The orientation of each strand on the left should match the orientation of the strand on the right that has the same endpoints.

Remark 7.5.6. The connectivity of strands on both sides of Figure 7.23 is the same, so the strand permutation is invariant under the swivel move.

We will soon see that triple diagrams are essentially cryptomorphic to plabic graphs—or more precisely, to the subclass of normal plabic graphs defined below.

Definition 7.5.7. We say that a plabic graph G is *normal* if the coloring of its internal vertices is bipartite, all white vertices in G are trivalent, and each boundary vertex is adjacent to a black vertex. See Figure 7.24.




Figure 7.24. A normal plabic graph G . This graph was obtained from the one in Figure 7.1(a) by inserting several bivalent black vertices.

Remark 7.5.8. Since a normal plabic graph does not have white leaves, it cannot contain a tree that collapses to a white lollipop. Therefore, in the case of normal graphs, there is no need to decorate the trip permutation.

Definition 7.5.9. The triple diagram $\mathfrak{X}(G)$ associated to a normal plabic graph G is constructed as follows. To each trip in G —either a one-way trip or a roundtrip—we associate a strand in the ambient disk \mathbf{D} by slightly deforming the trip, as shown in Figure 7.25, so that the strand

- runs along each edge of the trip, keeping the edge on its left,
- makes a U-turn at each black internal leaf (a vertex of degree 1);
- ignores black vertices of degree 2,
- makes a right turn (as sharp as possible) at each other black vertex, and
- makes a left turn at each white vertex v along the trip, passing through v .

This collection of strands forms a triple diagram $\mathfrak{X}(G)$, see Figure 7.26 and Lemma 7.5.10.




Figure 7.25. Constructing a triple diagram from a normal plabic graph.




Figure 7.26. Left: A normal plabic graph G (cf. Figure 7.24) together with the associated triple diagram $\mathfrak{X} = \mathfrak{X}(G)$. Conversely, $G = G(\mathfrak{X})$. Right: The triple diagram $\mathfrak{X} = \mathfrak{X}(G)$. The trip permutation of G and the strand permutation of \mathfrak{X} are equal: $\pi_G = \pi_{\mathfrak{X}} = (3, 4, 5, 1, 2)$.

Lemma 7.5.10. *The diagram $\mathfrak{X}(G)$ associated to a normal plabic graph G as in Definition 7.5.9 is a triple diagram.*

Proof. Since the white vertices in G are trivalent, $\mathfrak{X}(G)$ has a triple point for every white vertex in G , and no other crossings. We need to check that $\mathfrak{X} = \mathfrak{X}(G)$ is connected, or more precisely, that the union of the strands and the boundary of the disk is connected, as required by Definition 7.5.1.

Let us ignore any component consisting of a single black vertex which is adjacent only to boundary vertices (e.g. a black lollipop), as the corresponding strands are clearly connected to the boundary of the disk. Consider any other strand S in \mathfrak{X} . By construction, S passes through at least one white vertex of the (bipartite) graph G , which is a triple point on S . It therefore suffices to show that every triple point in \mathfrak{X} is connected to the boundary within \mathfrak{X} (i.e., via strand segments of \mathfrak{X}).

Let u be a k -valent black vertex in G and let v_1, \dots, v_k be the white or boundary vertices adjacent to u . The strands of \mathfrak{X} that run along the k edges of G incident to u cyclically connect the triple points v_1, \dots, v_k to each other. (If the list v_1, \dots, v_k includes boundary vertices, then the corresponding strand segments are connected via the boundary.) We conclude that for any two-edge path $v - u - v'$ in G connecting two white or boundary vertices v and v' via a black vertex u , the triple (or nearby boundary) points v and v' are connected within \mathfrak{X} . It follows that for any path in the bipartite graph G connecting a white vertex v to the boundary, there is a path in \mathfrak{X} that connects the triple point v to the boundary.

It remains to note that by Definition 7.1.1, any white vertex v in G is connected by a path in G to some boundary vertex. \square

We now go in the opposite direction, from a triple diagram to a normal plabic graph.

Definition 7.5.11. The normal plabic graph $G = G(\mathfrak{X})$ associated to a triple diagram \mathfrak{X} is constructed as follows. Place a white vertex of G at each triple crossing in \mathfrak{X} . Treat each boundary vertex of \mathfrak{X} as a boundary vertex of G . For each region R of \mathfrak{X} whose boundary is oriented counterclockwise, place a black vertex in the interior of R and connect it to the white and boundary vertices lying on the boundary of R , so that each white (resp., boundary) vertex is trivalent (resp., univalent). The resulting plabic graph $G = G(\mathfrak{X})$ is normal by construction.

Proposition 7.5.12. *The maps $G \mapsto \mathfrak{X}(G)$ and $\mathfrak{X} \mapsto G(\mathfrak{X})$ described in Definitions 7.5.9 and 7.5.11 are mutually inverse bijections between normal plabic graphs and triple diagrams with the same number of boundary vertices. The trip permutation π_G of a normal graph G is equal to the strand permutation $\pi_{\mathfrak{X}}$ of the corresponding triple diagram $\mathfrak{X} = \mathfrak{X}(G)$.*

Proof. Starting from a normal graph G , let us decompose it into star-shaped subgraphs S_v each of which includes a black vertex v , all the edges incident to v , and the endpoints of those edges. Each of these stars will give rise to a fragment of the triple diagram $\mathfrak{X}(G)$ that “hugs” the edges of S_v and whose boundary is oriented counterclockwise (looking from S_v). Moreover, $\mathfrak{X}(G)$ is obtained by stitching these fragments together. Applying the map $\mathfrak{X} \mapsto G(\mathfrak{X})$ to $\mathfrak{X}(G)$ will recover the original graph G .

One similarly shows that if we start from a triple diagram \mathfrak{X} , construct the normal graph $G(\mathfrak{X})$, and then apply the map $G \mapsto \mathfrak{X}(G)$ to $G(\mathfrak{X})$, then we recover the original triple diagram \mathfrak{X} . The key property to keep in mind is that each face of \mathfrak{X} is homeomorphic to an open disk.

The strands of \mathfrak{X} run alongside the trips of G , implying that $\pi_{\mathfrak{X}} = \pi_G$. \square




Figure 7.27. The six reduced normal plabic graphs with three boundary vertices, shown together with the corresponding triple diagrams, cf. Definition 7.5.9. The associated trip (resp., strand) permutations are precisely the six permutations of $\{1, 2, 3\}$.

In Theorem 7.7.3, we will characterize triple diagrams that correspond, under the bijection of Proposition 7.5.12, to *reduced* normal plabic graphs.

The bijective correspondence between triple diagrams and normal plabic graphs can be used to translate the move equivalence of triple diagrams (under the swivel moves, see Definition 7.5.5) into a version of the move equivalence of plabic graphs (cf. Definition 7.1.3) formulated entirely within the setting of normal plabic graphs:

Definition 7.5.13. The (*normal*) *urban renewal* move is the local transformation of normal plabic graphs described in Figure 7.28. (This differs slightly from Definition 2.5.2. In this chapter, we will consistently use the new definition.) Unlike the square move (M1), the normal urban renewal move does not require the vertices of the square to be trivalent. They can even be bivalent, see Figure 7.29.




Figure 7.28. Normal urban renewal replaces one of these configurations by the other. In contrast to the square move of Figure 7.2, where each of the four vertices of the quadrilateral face must have exactly one incident edge leading outside the configuration, we allow each black vertex of the quadrilateral face to have 0 or more incident edges leading outside the configuration. Cf also Figure 7.29.




Figure 7.29. Special case of normal urban renewal: one of the black vertices is not incident to any edges leading outside the configuration.

Definition 7.5.14. The *normal flip move* is the local transformation shown in Figure 7.30. Ignoring the bivalent black vertices, this local move is the same as the (white) flip move for trivalent plabic graphs.




Figure 7.30. The normal flip move.

Lemma 7.5.15. *Let G and G' be normal plabic graphs related via a sequence of normal urban renewal moves and normal flip moves. Then $G \sim G'$.*

Proof. Each instance of normal urban renewal can be expressed as a square move (M1) together with (M2) and/or (M3) moves. Each normal flip move can be expressed as a combination of (M2) and (M3) moves. \square

Theorem 7.5.16. *Let G and G' be normal plabic graphs and let $\mathfrak{X} = \mathfrak{X}(G)$ and $\mathfrak{X}' = \mathfrak{X}(G')$ be the corresponding triple diagrams. Then the following are equivalent:*

- G and G' are related via a sequence of urban renewal moves and normal flip moves;
- \mathfrak{X} and \mathfrak{X}' are move-equivalent (i.e., related via swivel moves).

Proof. Figure 7.31 shows how one can translate back-and-forth between

- an arbitrary swivel move in a triple diagram and
- either an urban renewal move or the normal flip move in the corresponding normal plabic graph.

The latter choice depends on the orientations of the strands involved. \square



Figure 7.31. Depending on the orientations of the strands involved, a swivel move in a triple diagram may correspond to (a) an urban renewal move or (b) a normal flip move in the associated normal plabic graph.

We next generalize Definition 7.5.9 to arbitrary plabic graphs.

Definition 7.5.17. Let G be a plabic graph (not necessarily normal). The *generalized triple diagram* $\mathfrak{X}(G)$ associated to G is defined as follows. (To be precise, we define $\mathfrak{X}(G)$ up to move equivalence.) The recipe is essentially the same as in Definition 7.5.9, with the following additional rules dealing with non-trivalent white vertices:

- at a univalent white vertex in G , make a U-turn, see Figure 7.32;
- at a bivalent white vertex in G , go straight through, see Figure 7.32;
- at a white vertex v of degree ≥ 4 , replace v by a trivalent tree (with white vertices); then, at each of the vertices of the tree, apply the rule shown in Figure 7.25 on the right. See Figure 7.33.

Although the trivalent tree replacing v is not unique, all these trees are related to each other by flip moves, cf. Figure 7.12. Hence all triple diagrams constructed from them are move-equivalent to each other, cf. Figure 7.31(b) (remove the black vertex in the center).



Figure 7.32. Constructing a triple diagram around a white vertex of degree 1 or 2 in a general plabic graph.




Figure 7.33. Constructing a triple diagram around a high-degree white vertex in a general plabic graph. We first replace this vertex by a trivalent tree, then construct the corresponding fragment of the triple diagram following the rule shown in Figure 7.25.

The following statement is immediate from the definitions.

Lemma 7.5.18. *Let G be a plabic graph. If the union of the strands in $\mathfrak{X}(G)$ and the boundary $\partial\mathbf{D}$ is connected, then $\mathfrak{X}(G)$ is a triple diagram in the sense of Definition 7.5.1.*

The connectedness condition in Lemma 7.5.18 does not hold in general. To be concrete, if G contains a cycle C all of whose vertices are black, then the strands located inside C are disconnected from the rest of $\mathfrak{X}(G)$.

Remark 7.5.19. Unfortunately, the extension of the definition of the triple diagram $\mathfrak{X}(G)$ described in Definition 7.5.17 does not allow a straightforward generalization of Theorem 7.5.16: move-equivalent plabic graphs do not necessarily yield move-equivalent triple diagrams. The only problematic local move is the one shown in Figure 7.34: contracting an edge connecting a white internal leaf to a white vertex of degree ≥ 3 removes a triple point in the associated triple diagram, thereby altering its move equivalence class.

This last complication prompts the following definition.

Definition 7.5.20. Let G and G' be plabic graphs. We write $G \overset{\circ}{\sim} G'$ if G and G' can be related to each other via a sequence of local moves (M1)–(M3) that does not include an instance of the (M3) move shown in Figure 7.34, cf. Remark 7.5.19.




Figure 7.34. Contracting (or decontracting) an edge that connects a white internal leaf to a white vertex of degree ≥ 3 produces a move-equivalent triple diagram.

Lemma 7.5.21. *Let G and G' be plabic graphs such that $G \simeq G'$. Then the corresponding (generalized) triple diagrams $\mathfrak{X}(G)$ and $\mathfrak{X}(G')$ are move equivalent (i.e., related to each other via swivel moves).*

We note that $\mathfrak{X}(G)$ and $\mathfrak{X}(G')$ are defined up to move equivalence, so the statement that they are move-equivalent to each other makes sense.

Proof. It is straightforward to verify, case by case, that each of the local moves (M1)–(M3), with the exception of the move shown in Figure 7.34, either leaves the associated (generalized) triple diagram invariant or applies a swivel move to it (more precisely, to any of the possible diagrams obtained using the construction in Definition 7.5.17). To be specific:

- a square move (M1) translates into a swivel move, see Figure 7.31(a);
- both the move (M2) and a black (de)contraction move (M3) leave the triple diagram invariant (up to isotopy);
- a white (de)contraction move (M3), other than the instance shown in Figure 7.34, translates into a swivel move, see Figure 7.31(b) (remove the black vertex in the center). \square

Corollary 7.5.22. *Let G and G' be normal plabic graphs. The following are equivalent:*

- (1) $G \simeq G'$;
- (2) G and G' are related via a sequence of normal urban renewal moves and normal flip moves;
- (3) $\mathfrak{X}(G)$ and $\mathfrak{X}(G')$ are move-equivalent (in the sense of Definition 7.5.5).

Proof. The implication (2) \Rightarrow (1) is an easy enhancement of Lemma 7.5.15. The equivalence (2) \Leftrightarrow (3) was established in Theorem 7.5.16. The implication (1) \Rightarrow (3) was proved in Lemma 7.5.21. \square

We note the similarity between Corollary 7.5.22 and Theorem 7.3.5.

7.6. Minimal triple diagrams

Definition 7.6.1. A triple diagram is called *minimal* if it has no more triple points than any other triple diagram with the same strand permutation.

We will show in Section 7.7 that minimal triple diagrams are the natural counterparts of reduced normal plabic graphs.

Much of this section is devoted to the proof of the following key result.

Theorem 7.6.2. *Any two minimal triple diagrams with the same strand permutation are move-equivalent to each other.*

Lemma 7.6.3. *If a triple diagram \mathfrak{X} is minimal, then so is every triple diagram move-equivalent to \mathfrak{X} .*

Proof. It is easy to see that a swivel move preserves both the number of triple points and the strand permutation. The claim follows. \square

We next describe certain “bad features” (of a triple diagram) and show that they cannot occur in a minimal triple diagram.

Definition 7.6.4. A strand in a triple diagram that intersects itself forms a *monogon*. A pair of strands that intersect at two points x and y form either a *parallel digon* or *anti-parallel digon*, depending on whether their segments connecting x and y run in the same or opposite direction, see Figure 7.35. We use the term *badgon* to refer to either a monogon or a parallel digon.




Figure 7.35. A monogon, a parallel digon, and an anti-parallel digon. The actual picture will contain (potentially many) additional strands and intersections.

Lemma 7.6.5. *A triple diagram without badgons has no closed strands.*

Proof. Let \mathfrak{X} be a triple diagram without badgons. Since \mathfrak{X} does not contain monogons, no strand of \mathfrak{X} can intersect itself. Suppose that \mathfrak{X} contains a closed strand S . Let T be another strand of \mathfrak{X} intersecting S at points x and y ; such T exists since \mathfrak{X} must be connected to the boundary $\partial\mathbf{D}$. Then the segment of T between x and y together with one of the segments of S connecting x and y form a parallel digon, which is a contradiction. \square

Lemma 7.6.6. *A minimal triple diagram does not contain badgons. Therefore (cf. Lemma 7.6.5) it does not contain closed strands.*

Proof. Let \mathfrak{X} be a triple diagram containing a monogon, i.e., a strand S with a self-intersection at a triple point v . Construct the triple diagram \mathfrak{X}' by deforming \mathfrak{X} around v so that S “spins off” a closed strand while the triple point disappears, see Figure 7.36. (If the spun-off portion is disconnected from the rest of \mathfrak{X} , then remove it altogether.) The triple diagram \mathfrak{X}' has the same strand permutation as \mathfrak{X} but fewer triple points; thus \mathfrak{X} is not minimal.




Figure 7.36. In the presence of a monogon, we can reduce the number of triple points while keeping the same strand permutation. The triple diagram may contain additional strands intersecting the monogon, as well as additional points of self-intersection.

Now suppose that \mathfrak{X} does not contain monogons but does contain two strands S and T that form a parallel digon. Say, S and T contain segments \overline{S} and \overline{T} that run from a triple point x to a triple point y . Let U (resp., V) be the third strand passing through x (resp., y). We then deform \mathfrak{X} around both x and y by smoothing each of the two triple points: the strands U and V continue to go straight through, whereas the endpoints of \overline{S} (resp., \overline{T}) get connected to T (resp., S). Thus, the strands S and T swap their segments \overline{S} and \overline{T} with each other (with appropriate smoothings), the overall connectivity (i.e., the strand permutation) is preserved, and the triple points at x and y disappear. (If the diagram becomes disconnected from $\partial\mathbf{D}$, then remove the disconnected portion.) We then conclude that \mathfrak{X} was not minimal. \square

Definition 7.6.7. Let S be an arc in a triple diagram, i.e., a strand whose endpoints s and t lie on the boundary of the ambient disk \mathbf{D} . We call S *boundary-parallel* if it runs along a segment I of the boundary $\partial\mathbf{D}$ between s and t (in either direction), so that every other strand with an endpoint inside I runs directly to or from S , without any triple crossings in between. See Figure 7.37.




Figure 7.37. A boundary-parallel strand S in a triple diagram.

We next describe a particular way to construct, for any given permutation π , a triple diagram whose strand permutation is π .

Definition 7.6.8. Let π be a permutation of b letters $1, \dots, b$. A triple diagram in the disk \mathbf{D} is called *standard* (for π) if it can be constructed using the following recursive process. (The process involves some choices, so a standard diagram for π is not unique.)


We place b boundary vertices on the boundary $\partial\mathbf{D}$ and label them $1, \dots, b$ clockwise. Next to each boundary vertex v , we mark two endpoints of the future strands: a source endpoint that precedes v in the clockwise order and a target endpoint that follows v in this order. We know which source is to be matched to which target by the strand permutation π . Each such pair of endpoints divides the circle $\partial\mathbf{D}$ into two intervals. Let us partially order these $2b$ intervals by inclusion and select a *minimal interval* I with respect to this partial order.

We start constructing the triple diagram by running a boundary-parallel strand S along the interval I , introducing a triple crossing for each pair of strands that need to terminate in the interior of I , as shown in Figure 7.37. There will always be an even number (possibly zero) of strands to cross over, so the construction will proceed without a hitch.

Let \mathbf{D}' be the disk obtained from \mathbf{D} by removing the region between the boundary segment I and the strand S together with a small neighborhood of S ; so \mathbf{D}' is the shaded region in Figure 7.37. We accordingly remove S and its endpoints from the original pairing of the in- and out-endpoints, and swap each pair that S crossed over. This yields $2(b-1)$ endpoints on the boundary of \mathbf{D}' ; note that the in- and out-endpoints alternate, as before. We then determine the new pairing of these endpoints (thus, a new strand permutation, after an appropriate renumbering) and recursively continue the process in \mathbf{D}' until the desired (standard) triple diagram is constructed.


We shall keep in mind that a standard triple diagram is constructed by choosing a sequence of minimal intervals.

Exercise 7.6.9. For each of the three pairs of triple diagrams shown below, demonstrate that the two diagrams are move-equivalent to each other, i.e., related via a sequence of swivel moves.



(In each of the three cases, the central section can involve an arbitrary number of repetitions.)

Exercise 7.6.10. Use (7.6.3) to prove the move equivalence (7.6.4) below:



Lemma 7.6.11. Let \mathfrak{X} be a triple diagram such that no triple diagram move-equivalent to \mathfrak{X} contains a monogon. Then the following statements hold:

- (i) No triple diagram move-equivalent to \mathfrak{X} has a badgon or a closed strand.
- (ii) Let I be a minimal interval for the strand permutation associated with \mathfrak{X} . Then \mathfrak{X} is move-equivalent to a diagram \mathfrak{X}' in which the strand connecting the endpoints of I is boundary-parallel along I .

Proof. We will simultaneously prove statements (i) and (ii) by induction on the number of triple points in \mathfrak{X} . Thus, we assume that both (i) and (ii) hold for triple diagrams that have fewer triple points than \mathfrak{X} .

We first prove (i). Suppose that a triple diagram $\mathfrak{X}' \sim \mathfrak{X}$ contains (non-self-intersecting) strands S and U forming a parallel digon. The strand S cuts the disk \mathbf{D} into two regions. Let R be the region containing the digon, with a small neighbourhood of S removed. Since the boundaries of the faces of \mathfrak{X}' are consistently oriented, the same is true for the portion of \mathfrak{X}' contained inside R , so this portion can be viewed as a (smaller) triple diagram. Suppose that U bounds a minimal interval within S (viewed as a portion of the boundary of R). Then by the induction assumption, U can be moved to be boundary-parallel to S . Since S and U are co-oriented, we get the picture on the left-hand side of (7.6.4) (with U running horizontally at the bottom). Applying (7.6.4), we obtain a monogon, a contradiction.

If the subinterval of S cut out by U is not minimal, then there is a strand T that cuts across S twice, creating a minimal interval within S and forming a digon inside R . We may assume that this digon is anti-parallel (or else replace U by T and repeat). By the induction assumption, we can apply swivel moves inside R to make T boundary-parallel to S . We then apply (7.6.1) to remove the digon, as shown in Figure 7.38. Repeating this operation if necessary, we obtain a triple diagram in which U bounds a minimal interval within S ; we then argue as above to arrive at a contradiction.




Figure 7.38. Removing double intersections with S .

Thus, no triple diagram $\mathfrak{X}' \sim \mathfrak{X}$ contains badgons. By Lemma 7.6.5, we conclude that any such \mathfrak{X}' does not contain closed strands either. This completes the induction step for statement (i).

We now proceed to proving statement (ii). In addition to the induction assumption for (ii), we may assume that neither \mathfrak{X} nor any triple diagram move-equivalent to \mathfrak{X} contains a badgon or a closed strand.

Let S be the strand connecting the endpoints of the minimal interval I .

Step 1: Removing double intersections with S , see Figure 7.38. Let R be the region between the strand S and the interval I , with a small neighborhood of S removed. Suppose there is a strand that intersects S more than once. Among such strands, take one that cuts out a minimal interval along the boundary of R . Let T denote the segment of this strand contained in R . The portion of \mathfrak{X} contained inside R has fewer triple crossings than \mathfrak{X} , so by the induction assumption, we can make T boundary-parallel to S by applying swivel moves inside R . Now T and S form a (necessarily anti-parallel) digon, which we then remove using (7.6.1). We repeat this procedure until there are no strands left that intersect S more than once. Since the number of triple points along S decreases each time, the process terminates.

Step 2: Combing out the triple crossings. At this stage, no strand crosses S more than once. Since I is minimal, no strand has both ends at I . Since \mathfrak{X} contains no closed strands, every (non-self-intersecting) strand appearing between S and I must start or end at a point in I and cross S . Suppose that S is not boundary-parallel. Then there exists a strand T with an endpoint at I that passes through a triple point before hitting S . Among all such T , choose the one with the leftmost endpoint along I , cf. Figures 7.39 and 7.40 on the left. Let R be the part of the region between I and S that lies to the right of any strand T' located to the left of T . (By our choice of T , all such strands T' run directly from I to S , with no crossings in between.) As we have eliminated all double intersections with S , the interval corresponding to T (looking to the left) is minimal inside R . We can therefore use the induction assumption inside R to make T boundary-parallel.

What we do next depends on the orientation of T relative to S . If T is anti-parallel to S , as in Figure 7.39, then we apply (7.6.2) to make T run directly to S . If T is parallel to S , as in Figure 7.40, then we apply (7.6.3).

We repeat this step until S is boundary-parallel. \square

Lemma 7.6.12. *For a triple diagram \mathfrak{X} , the following are equivalent:*

- (a) *Any diagram \mathfrak{X}' move-equivalent to \mathfrak{X} does not contain a monogon.*
- (b) *\mathfrak{X} is move-equivalent to any standard triple diagram with the same strand permutation.*
- (c) *\mathfrak{X} is minimal.*

In particular, any standard triple diagram is minimal.




Figure 7.39. Combing out the triple crossings: the anti-parallel case.




Figure 7.40. Combing out the triple crossings: the parallel case.

Proof. The implication (c) \Rightarrow (a) follows from Lemmas 7.6.3 and 7.6.6. To prove the implication (a) \Rightarrow (b), choose a sequence of minimal intervals and repeatedly apply Lemma 7.6.11. We have now established (c) \Rightarrow (b), so any minimal triple diagram is move-equivalent to any standard triple diagram with the same strand permutation. It follows by Lemma 7.6.3 that any standard triple diagram is minimal, hence so is any diagram move-equivalent to a standard one. Thus (b) \Rightarrow (c) is proved. \square

Proof of Theorem 7.6.2. By Lemma 7.6.12, any two minimal triple diagrams with strand permutation π are move-equivalent to any standard diagram with strand permutation π , and therefore to each other. \square

Lemma 7.6.13. *Let \mathfrak{X} and \mathfrak{X}' be triple diagrams related by a swivel move. If \mathfrak{X} contains a badgon, then so does \mathfrak{X}' .*

Proof. We label the strands and the triple points involved in this swivel move by a, b, c, d , and x, y , as shown in Figure 7.41.




Figure 7.41. A swivel move relating \mathfrak{X} and \mathfrak{X}' .

If \mathfrak{X} contains a badgon that involves neither x nor y , then this badgon persists in \mathfrak{X}' .

Suppose \mathfrak{X} contains a monogon whose self-intersection point is (say) x . Thus, two of the strands $\{a, b, c\}$ coincide. If $a = c$ (resp., $b = c$), then the same monogon persists in \mathfrak{X}' because in Figure 7.41, strands a and c (resp., b and c) intersect in both \mathfrak{X} and \mathfrak{X}' .

If, on the other hand, $a = b$, then \mathfrak{X}' has a parallel digon, see Figure 7.42.




Figure 7.42. A monogon in \mathfrak{X} results in a parallel digon in \mathfrak{X}' .

From now on, we can assume that there is no monogon in \mathfrak{X} . Suppose \mathfrak{X} has a parallel digon whose two intersection points include x but not y . The sides of this parallel digon are either $\{a, b\}$ or $\{a, c\}$ or $\{b, c\}$. The last two cases are easy because such a parallel digon will persist in \mathfrak{X}' , since the strands a and c (resp., b and c) intersect in both \mathfrak{X} and \mathfrak{X}' .

Now suppose that our parallel digon has sides a and b , see Figure 7.43 on the left. (If the strands a and b go to the left and meet again there, then we get the same picture but with the roles of x and y interchanged.) Note that the end of strand a shown inside the digon must extend outside of it, but it cannot intersect a , as this would create a monogon. So strand a must intersect strand b again, see Figure 7.43 in the middle. Then, after the swivel move, we get a parallel digon as shown in Figure 7.43 on the right.

Finally, suppose there is a parallel digon in \mathfrak{X} whose two intersection points are x and y . We can assume it is oriented from x to y . The two arcs of the parallel digon should come from the following list:

- (aa) the arc along a from x to y ;
- (bb) the arc along b from x to y ;
- (cd) an arc leaving x along c , and returning to y along d (so $c = d$);
- (cb) an arc leaving x along c , and returning to y along b (so $c = b$);
- (bd) an arc leaving x along b , and returning to y along d (so $b = d$).

In case (bb), we get a closed strand; it will persist in \mathfrak{X}' and yield a badgon by Lemma 7.6.5. In cases (cb) and (bd), we get a monogon, contradicting our assumption. The remaining case is when the parallel digon has sides (aa) and (cd), in which case we get a monogon in \mathfrak{X}' . (The picture is like Figure 7.42, with the roles of \mathfrak{X} and \mathfrak{X}' swapped and some strands relabeled.)

□




Figure 7.43. Persistence of parallel digons under swivel moves.

Theorem 7.6.14. *A triple diagram is minimal if and only if it has no badgons.*

Proof. The “only if” direction is Lemma 7.6.6. The “if” direction follows from Lemma 7.6.13 and Lemma 7.6.12 (implication (a) \Rightarrow (c)). \square

Lemma 7.6.15. *Assume that a triple diagram \mathfrak{X} is not minimal. Then there exists a diagram \mathfrak{X}' move-equivalent to \mathfrak{X} that contains a hollow monogon.*

Proof. We will argue by induction on the number of faces in \mathfrak{X} . If this number is 1 or 2, then the claim is vacuously true.

By Lemma 7.6.12, there exists $\mathfrak{X}' \sim \mathfrak{X}$ such that \mathfrak{X}' has a monogon. Let M be the segment of a strand in \mathfrak{X}' that forms a monogon; we may assume M does not intersect itself except at its endpoints (or else replace M by its subsegment). If the monogon encircled by M is hollow, we are done. Otherwise, consider the disk \mathbf{D}_\circ obtained by removing a small neighborhood of M from the interior of the monogon. Let \mathfrak{X}'_\circ denote the portion of \mathfrak{X}' contained in \mathbf{D}_\circ ; this is a triple diagram with fewer faces than \mathfrak{X}' (or equivalently \mathfrak{X}).

The rest of the argument proceeds by showing that either we can apply local moves to \mathfrak{X}'_\circ to create a hollow monogon inside \mathbf{D}_\circ or we can apply moves to reduce the number of faces inside the monogon encircled by M (eventually producing a hollow monogon). If \mathfrak{X}'_\circ is not minimal, then the induction assumption applies, so we can transform \mathfrak{X}'_\circ (thus \mathfrak{X}' or \mathfrak{X}) into a move-equivalent triple diagram containing a hollow monogon. Therefore, we may assume that \mathfrak{X}'_\circ is minimal. Let M_\circ denote the interval obtained from the boundary of \mathbf{D}_\circ by removing a point located near the vertex of our monogon. Let $I \subset M_\circ$ be a minimal interval of the triple diagram \mathfrak{X}'_\circ . Since this triple diagram is minimal, we can, by virtue of Lemma 7.6.12 (or Lemma 7.6.11), apply local moves inside \mathbf{D}_\circ to transform \mathfrak{X}'_\circ into a triple diagram in which the strand T connecting the endpoints of I is boundary-parallel to I . Let us now look at the digon formed by T and the portion of M that runs along I . If this digon is anti-parallel, then we can push T outside the monogon as in Figure 7.38, reducing the number of faces enclosed by M . If, on the other hand, the digon is parallel, then we can use (7.6.4) to create a hollow monogon. \square

7.7. From minimal triple diagrams to reduced plabic graphs

In this section, we use the machinery of triple diagrams and normal plabic graphs to prove Proposition 7.4.22 and Theorem 7.4.25.

Lemma 7.7.1. *A normal plabic graph contains no collapsible trees.*

Proof. All internal leaves in a normal plabic graph G are black. Whatever (M2)–(M3) moves we apply to a tree, it will always have a black leaf (so it can't collapse to a white root or lollipop), and it will always have a white vertex of degree at least 3 (so it can't collapse to a black root or lollipop). \square

Lemma 7.7.2. *Let G be a non-reduced normal plabic graph. Then there exists a plabic graph $G' \simeq G$ (cf. Definition 7.5.20) containing one of the forbidden configurations shown in Figure 7.18.*

Proof. Suppose G has an internal (necessarily black) leaf u that is not a lollipop. Let v be the unique (white, trivalent) vertex adjacent to u . This gives us a forbidden configuration as in Figure 7.18(d). (Note that by Lemma 7.7.1, G has no collapsible trees.)

If, on the other hand, G has no such internal leaves, then by Proposition 7.4.9, G can be transformed, via local moves that do not create leaves, into a graph containing a hollow digon. \square

Recall that by Proposition 7.5.12, the map $G \rightarrow \mathfrak{X}(G)$ defined in Definition 7.5.9 gives a bijection between normal plabic graphs and triple diagrams with the same number of boundary vertices; moreover, this bijection preserves the associated (resp., trip or strand) permutation.

Theorem 7.7.3. *A normal plabic graph G is reduced if and only if the triple diagram $\mathfrak{X}(G)$ is minimal. Thus the map $G \mapsto \mathfrak{X}(G)$ restricts to a bijection between reduced normal plabic graphs and minimal triple diagrams.*

Proof. Suppose $\mathfrak{X}(G)$ is not minimal. By Lemma 7.6.15, there is a triple diagram $\mathfrak{X}' \sim \mathfrak{X}(G)$ such that \mathfrak{X}' has a hollow monogon. By Proposition 7.5.12, we have $\mathfrak{X}' = \mathfrak{X}(G')$ for some normal plabic graph G' . Moreover, by Corollary 7.5.22, G and G' are move equivalent. The hollow monogon in \mathfrak{X}' corresponds in the normal graph G' to one of the configurations shown in Figure 7.44: either a hollow digon or a black leaf adjacent to a white trivalent vertex. Either way, G' contains one of the forbidden configurations from Figure 7.18 (cf. Lemma 7.7.1), so G is not reduced.

Going in the other direction, suppose that a normal plabic graph G is not reduced. By Lemma 7.7.2, there exists $G' \simeq G$ containing a forbidden configuration. By Lemma 7.5.21, the triple diagram $\mathfrak{X}(G)$ is move-equivalent




Figure 7.44. A hollow monogon in a triple diagram yields a forbidden configuration in the corresponding normal plabic graph, cf. Figure 7.18.

to the (generalized) triple diagram $\mathfrak{X}(G')$. Since $\mathfrak{X}(G)$ is connected, so is $\mathfrak{X}(G')$. It follows by Lemma 7.5.18 that $\mathfrak{X}(G')$ is an honest triple diagram.

The remaining argument depends on the type of a forbidden configuration present in G' . Since G is normal and $G' \simeq G$, it follows that G' has no white leaves. If G' contains a digon whose vertices are of the same color, then $\mathfrak{X}(G')$ has a closed strand; hence $\mathfrak{X}(G')$ is not minimal (by Lemma 7.6.6) and neither is $\mathfrak{X}(G)$. Finally, if G' contains one of the configurations shown in Figure 7.44, then $\mathfrak{X}(G')$ contains a monogon, hence is not minimal. \square

Lemma 7.7.4. *Let G be a reduced plabic graph that does not contain a white lollipop, nor a tree that collapses to a white lollipop. Then G is move-equivalent to a normal plabic graph.*

See Figure 7.45.




Figure 7.45. A reduced plabic graph G and a normal plabic graph move-equivalent to G .

Proof. We use induction on the number of faces of G . Note that each local move keeps this number invariant. We begin by collapsing all collapsible trees (cf. Definition 7.4.1) and removing all bivalent vertices. The resulting plabic graph has no white leaves, since a white leaf would either be a lollipop or else be adjacent to a black vertex of degree ≥ 3 , which is impossible since the plabic graph at hand is reduced and has no collapsible trees.

We then split each white vertex of degree ≥ 4 into a tree made of trivalent white vertices. After that, we insert a bivalent black vertex in the middle of each edge with both endpoints white, as well as near each boundary vertex connected to a white vertex. By an abuse of notation, we keep calling our plabic graph G , even as it undergoes these and subsequent transformations.

It remains to contract the edges with both endpoints black to make the graph bipartite. This step may however be problematic if such an edge e is a loop connecting some black vertex v to itself. We will demonstrate that this in fact cannot happen.

If e encloses a face (i.e., there are no other vertices/edges of G inside e), then we can insert a black vertex into e and obtain a double edge, contradicting the fact that G is reduced. Now suppose that e encloses some nontrivial subgraph of G . Let \mathbf{D}' be the region bounded by e . We split v into two vertices v_1 and v_2 using an (M3) move, so that v_1 stays incident to e and to the edges located outside \mathbf{D}' whereas v_2 is incident to the edges inside \mathbf{D}' . See Figure 7.46.

We thus obtain a subgraph G' , viewed as a plabic graph in \mathbf{D}' with a single boundary vertex v_1 . The plabic graph G' must be reduced, or else G would not be. Moreover G' has fewer faces than G . Also, G' does not collapse to a white lollipop (even if v_2 is bivalent), because this would result (after the collapse) in a forbidden configuration in G . (We note that G has no collapsible trees.) Thus, the induction hypothesis applies, so we can transform G' via local moves into a normal graph G'' . Since G'' is reduced, its triple diagram $\mathfrak{X}(G'')$ must be minimal by Theorem 7.7.3. Given that $\mathfrak{X}(G'')$ only has one in- and one out-endpoint, this means that $\mathfrak{X}(G'')$ consists of a single strand connecting these endpoints to each other, with no triple points. In other words, the normal graph G'' is a black lollipop at v_1 . We then contract it into v_1 , creating a loop enclosing a face, and arrive at a contradiction with G being reduced. \square




Figure 7.46. Excluding a loop e based at a black vertex v .

Proof of Proposition 7.4.22. Let G be a reduced plabic graph such that $\pi_G(i) = i$. We need to show that the connected component of G containing the boundary vertex i collapses to a lollipop at i , cf. Definition 7.4.1.

Suppose otherwise. Without loss of generality, we can assume that G has no trees collapsing to other lippops either. Since G is reduced, Lemma 7.7.4 applies, so G is move-equivalent to a normal plabic graph G' . The trip permutations of G and G' coincide with each other (by Exercise 7.4.19) and with the strand permutation of the triple diagram $\mathfrak{X}(G')$ (by Proposition 7.5.12). Since G is reduced, so is G' ; hence $\mathfrak{X}(G')$ is minimal by Theorem 7.7.3.

Let d be the degree of the black vertex adjacent to the boundary vertex i in G' . If $d = 1$, then there is a sequence of local moves relating the component of G containing i to the black lollipop at i in G' . Since local moves preserve the number of internal faces, and a black lollipop has no

internal faces, no (M1) move appears in the sequence of moves. Therefore this component must be a (collapsible) tree and we are done. If $d = 2$ (see Figure 7.47 on the left), then $\pi_G(i) = i$ implies that $\mathfrak{X}(G')$ has a monogon, so it cannot be minimal, cf. Lemma 7.6.6. If $d \geq 3$, then we get a parallel digon (see Figure 7.47 on the right), again contradicting the minimality of $\mathfrak{X}(G')$.

□



Figure 7.47. The vicinity of i in G' .

Proof of Theorem 7.4.25. Let G and G' be reduced plabic graphs. If $G \sim G'$, then $\tilde{\pi}_G = \tilde{\pi}_{G'}$ by Exercise 7.4.24. We need to show the converse.

Let G and G' be reduced plabic graphs such that $\tilde{\pi}_G = \tilde{\pi}_{G'}$. If this decorated permutation has a fixed point at some vertex i , then by Proposition 7.4.22, after applying local moves if needed, both G and G' have a lollipop (of the same color) in position i . We can delete this lollipop in both graphs; the resulting graphs are still reduced, and their decorated trip permutations still coincide. So without loss of generality, we may assume that $\tilde{\pi}_G = \tilde{\pi}_{G'}$ has no fixed points and correspondingly G and G' have no trees collapsing to lollipops. Applying local moves as needed, we can furthermore assume, in light of Lemma 7.7.4, that both G and G' are normal. Since they are also reduced, Theorem 7.7.3 implies that the triple diagrams $\mathfrak{X}(G)$ and $\mathfrak{X}(G')$ are minimal. By Proposition 7.5.12, we moreover have $\pi_{\mathfrak{X}(G)} = \pi_G = \pi_{G'} = \pi_{\mathfrak{X}(G')}$. Invoking Theorem 7.6.2, we conclude that $\mathfrak{X}(G)$ and $\mathfrak{X}(G')$ are move-equivalent. Then by Corollary 7.5.22, G and G' are move-equivalent as well. □

Remark 7.7.5. As we have seen, A. Postnikov's theory of plabic graphs [21] is closely related to D. Thurston's theory of triple diagrams [24]. In particular, reduced plabic graphs are essentially minimal triple diagrams in disguise. While we have not discussed it here, there are some *reduction moves* that can be repeatedly applied to a non-reduced plabic graph (resp., a non-minimal triple diagram) in order to—eventually—make it reduced (resp., minimal). Here the two theories diverge: reduction moves for triple diagrams preserve the strand permutation, but reduction moves for plabic graphs do not preserve the trip permutation. In spite of that, reduction moves for plabic graphs fit into the theory of the totally nonnegative Grassmannian, as they are compatible with its cell decomposition, cf. [21, Section 12]. We will discuss this in the next chapter.

7.8. The bad features criterion

Here we provide an algorithm for deciding whether a plabic graph is reduced or not. We first explain (see Definition 7.8.2) how to transform an arbitrary plabic graph G into a normal plabic graph $N(G)$ move-equivalent to G (or conclude that G is not reduced). We then use a criterion based on Theorem 7.7.3 to determine whether $N(G)$ (hence G) is reduced or not.

Lemma 7.8.1. *Let G be a reduced plabic graph. Then G has no roundtrips. Also, G has no loops, i.e., edges whose endpoints coincide.*

Proof. We can assume that G does not contain white lollipops or trees that collapse to white lollipops. (Collapsing such trees and removing lollipops does not affect whether a graph is reduced or whether it has a roundtrip.) Now by Lemma 7.7.4, G is move-equivalent (up to the removal of some white lollipops) to a normal plabic graph G' . Since G' is reduced, $\mathfrak{X}(G')$ is minimal (see Theorem 7.7.3). Hence $\mathfrak{X}(G')$ has no closed strands (see Lemma 7.6.6), so G' has no roundtrips. Since roundtrips persist under local moves, G has no roundtrips either.

Suppose G has a loop e based at a black (resp., white) vertex. (Some edges and vertices might be enclosed by e .) Then the trip that traverses e clockwise (resp., counterclockwise) is a roundtrip, a contradiction. \square

Definition 7.8.2. Let G be a plabic graph. The following algorithm is similar to the procedure employed in the proof of Lemma 7.7.4. It either determines that G is not reduced, or outputs a normal plabic graph $N(G)$ which, up to the addition/removal of lollipops, is move-equivalent to G . In the latter scenario, the plabic graph $N(G)$ —hence the original graph G —may be either reduced or not.

At each stage of the algorithm, the plabic graph G undergoes some transformations that do not affect whether it is reduced.

Stage 1. Use moves (M2)–(M3) to collapse all collapsible trees in G .

Stage 2. Use moves (M2) to remove all bivalent vertices.

Stage 3. Remove all lollipops.

Stage 4. If G has an internal leaf u , then let v be the vertex adjacent to u . By construction, $\deg(v) \geq 3$. Since G has no collapsible trees, we conclude that u and v are of different color and G is not reduced.

If G has no internal leaves, then each internal vertex has degree ≥ 3 .

Stage 5. Use moves (M3) to contract all edges with both endpoints black. If at any point G has a loop, then it is not reduced, by Lemma 7.8.1.

Stage 6. Use moves (M3) to replace each white vertex by a trivalent tree with white vertices.

Stage 7. Use moves (M2) to insert a black vertex into every edge with no black endpoints. The resulting plabic graph $N(G)$ is normal.

Definition 7.8.3. If a trip passes through an edge e of a plabic graph twice (in the opposite directions), we call this an *essential self-intersection*.

If the edges e_1 and e_2 are such that there are two distinct trips each of which passes first through e_1 and then through e_2 , we call this a *bad double crossing*.

We use the term *bad features* to collectively refer to

- roundtrips (see Definition 7.4.12),
- essential self-intersections, and
- bad double crossings.

These notions are illustrated in Figure 7.48.




Figure 7.48. Plabic graph fragments representing “bad features:” (a) a roundtrip; (b) essential self-intersection; (c) bad double crossing.

Lemma 7.8.4. *A normal plabic graph G has a bad feature if and only if the associated triple diagram $\mathfrak{X}(G)$ has a badgon.*

Proof. Let G be a normal plabic graph. The strands in the triple diagram $\mathfrak{X} = \mathfrak{X}(G)$ closely follow the trips in G . Therefore \mathfrak{X} has a closed strand if and only if G has a roundtrip.

If G has an essential self-intersection (resp., a bad double crossing), then \mathfrak{X} has a monogon (resp., a parallel digon). To see that, take each edge e involved in a bad feature and consider the white end v of e . The strands corresponding to the trips involved in the bad feature will intersect at v ; thus v will be a vertex of the corresponding badgon. Cf. Figures 7.44 and 7.49.

Conversely, suppose that \mathfrak{X} has a monogon with self-intersection corresponding to the white vertex v of G . There are three strand segments of \mathfrak{X} that pass through v , each running along two distinct edges incident to v ; because we have a self-intersection, two of these strands segments are part of the same strand s . Since v is trivalent, the pigeonhole principle implies




Figure 7.49. A bad double crossing in G yields a parallel digon in $\mathfrak{X}(G)$.

that two of the four edges that s runs along must coincide. This yields an essential self-intersection in G . A similar argument shows that if \mathfrak{X} has a parallel digon, then G has a bad double crossing. \square

Corollary 7.8.5. *Let G be a normal plabic graph. Let $\mathfrak{X} = \mathfrak{X}(G)$ be the corresponding triple diagram. Then the following are equivalent:*

- G is reduced;
- \mathfrak{X} is minimal;
- G has no bad features;
- \mathfrak{X} has no badgons.

Proof. By Theorem 7.7.3, G is reduced if and only if \mathfrak{X} is minimal. By Theorem 7.6.14, \mathfrak{X} is minimal if and only if \mathfrak{X} has no badgons. By Lemma 7.8.4, \mathfrak{X} has no badgons if and only if G has no bad features. \square

Corollary 7.8.5 provides the following criterion, which is a version of [21, Theorem 13.2].

Theorem 7.8.6. *A normal plabic graph is reduced if and only if it does not contain any bad features.*

For example, any plabic graph containing one of the fragments shown in Figure 7.48 is necessarily not reduced.

Remark 7.8.7. For *any* plabic graph G , Theorem 7.8.6 can be used in conjunction with the procedure described in Definition 7.8.2 to determine whether G is reduced or not.

Remark 7.8.8. Recall from Remark 2.3.4 that a factorization (not necessarily reduced) of an element of a symmetric group into a product of simple reflections can be represented by (a version of) a wiring diagram. As plabic graphs can be viewed as generalizations of wiring diagrams (see Example 7.2.4), reduced plabic graphs may be viewed as a generalization of reduced expressions. In this context, the criterion of Theorem 7.8.6 corresponds to the condition that each pair of lines in the wiring diagram intersect at most once.

7.9. Affine permutations

By Theorem 7.4.25, move equivalence classes of reduced plabic graphs are labeled by decorated permutations. An alternative (sometimes more useful) labeling utilizes $((a, b)$ -bounded) *affine permutations*, introduced and studied in this section.

Definition 7.9.1. For a decorated permutation $\tilde{\pi}$ on b letters, we say that $i \in \{1, \dots, b\}$ is an *anti-excedance* of $\tilde{\pi}$ if either $\tilde{\pi}^{-1}(i) > i$ or if $\tilde{\pi}(i) = \bar{i}$. The number of anti-excedances of $\tilde{\pi}$ (which we usually denote by a) is equal to the number of values $i \in \{1, \dots, b\}$ such that $\tilde{\pi}(i) < i$ or $\tilde{\pi}(i) = \bar{i}$.

Example 7.9.2. The decorated permutation $\tilde{\pi} = (5, \underline{2}, \bar{3}, 6, 4, 1)$ on $b = 6$ letters (see Figure 7.59) has $a = 3$ anti-excedances, namely, 1, 4, and $\bar{3}$. Indeed, $\tilde{\pi}^{-1}(1) = 6 > 1$, and $\tilde{\pi}^{-1}(4) = 5 > 4$.

Definition 7.9.3. Let $\tilde{\pi}$ be a decorated permutation on b letters with a anti-excedances. The *affinization* of $\tilde{\pi}$ is the map $\tilde{\pi}_{\text{aff}} : \mathbb{Z} \rightarrow \mathbb{Z}$ constructed as follows. For $i \in \{1, \dots, b\}$, we set

$$\tilde{\pi}_{\text{aff}}(i) = \begin{cases} \tilde{\pi}(i) & \text{if } \tilde{\pi}(i) > i, \\ i & \text{if } \tilde{\pi}(i) = \underline{i}, \\ \tilde{\pi}(i) + b & \text{if } \tilde{\pi}(i) < i, \\ i + b & \text{if } \tilde{\pi}(i) = \bar{i}. \end{cases}$$

We then extend $\tilde{\pi}_{\text{aff}}$ to \mathbb{Z} so that it satisfies

$$(7.9.1) \quad \tilde{\pi}_{\text{aff}}(i + b) = \tilde{\pi}_{\text{aff}}(i) + b \quad (i \in \mathbb{Z}).$$

We note that

$$(7.9.2) \quad i \leq \tilde{\pi}_{\text{aff}}(i) \leq i + b \quad (i \in \mathbb{Z})$$

and

$$(7.9.3) \quad \sum_{i=1}^b (\tilde{\pi}_{\text{aff}}(i) - i) = b \cdot \#\{i \in \{1, \dots, b\} \mid \tilde{\pi}(i) < i \text{ or } \tilde{\pi}(i) = \bar{i}\} = ab.$$

Example 7.9.4. Continuing with $\tilde{\pi} = (5, \underline{2}, \bar{3}, 6, 4, 1)$ from Example 7.9.2, we get $\tilde{\pi}_{\text{aff}}(1) = 5$, $\tilde{\pi}_{\text{aff}}(2) = 2$, $\tilde{\pi}_{\text{aff}}(3) = 9$, $\tilde{\pi}_{\text{aff}}(4) = 6$, $\tilde{\pi}_{\text{aff}}(5) = 10$, $\tilde{\pi}_{\text{aff}}(6) = 7$, or more succinctly,

$$\tilde{\pi}_{\text{aff}} = (\dots, 5, 2, 9, 6, 10, 7, \dots) = (\dots \boxed{5 \ 2 \ 9 \ 6 \ 10 \ 7} \ 11 \ 8 \ 15 \ 12 \ 16 \ 13 \ \dots).$$

(The boxed terms are the values at $1, \dots, b$. They determine the rest of the sequence by virtue of (7.9.1).) In accordance with (7.9.3), we have

$$(5 + 2 + 9 + 6 + 10 + 7) - (1 + \dots + 6) = 39 - 21 = 18 = 3 \cdot 6 = ab.$$

With the above construction in mind, we introduce the following notion.

Definition 7.9.5. Let a and b be positive integers. An (a, b) -bounded affine permutation is a bijection $f : \mathbb{Z} \rightarrow \mathbb{Z}$ satisfying the following conditions:

- $f(i + b) = f(i) + b$ for all $i \in \mathbb{Z}$;
- $i \leq f(i) \leq i + b$ for all $i \in \mathbb{Z}$;
- $\sum_{i=1}^b (f(i) - i) = ab$.

Lemma 7.9.6. [15] *The correspondence $\tilde{\pi} \mapsto \tilde{\pi}_{\text{aff}}$ described in Definition 7.9.3 restricts to a bijection between decorated permutations on b letters with a anti-excedances and the (a, b) -bounded affine permutations.*

Proof. If $\tilde{\pi}$ is a decorated permutation on b letters with a anti-excedances, then (7.9.1)–(7.9.3) show that $\tilde{\pi}_{\text{aff}}$ is an (a, b) -bounded affine permutation.

Conversely, given an (a, b) -bounded affine permutation $f : \mathbb{Z} \rightarrow \mathbb{Z}$, we can define the decorated permutation $\tilde{\pi}$ on b letters by

$$\tilde{\pi}(i) = \begin{cases} \underline{i} & \text{if } f(i) = i; \\ \bar{i} & \text{if } f(i) = i + b; \\ f(i) & \text{if } f(i) \leq b \text{ and } f(i) \neq i; \\ f(i) - b & \text{if } f(i) > b \text{ and } f(i) \neq i + b. \end{cases}$$

We claim that $\tilde{\pi}$ has a anti-excedances. Using the inequality $i \leq f(i) \leq i + b$, we conclude that the anti-excedances of $\tilde{\pi}$ are in bijection with the values $i \in \{1, \dots, b\}$ such that $f(i) > b$. The claim follows from the observation that $ab = \sum_{i=1}^b (f(i) - i) = b \cdot \#\{i \in \{1, \dots, b\} \mid f(i) > b\}$. \square

Lemma 7.9.6 is illustrated in Figure 7.50 (the first two columns).

$\tilde{\pi}$	$\tilde{\pi}_{\text{aff}}$	$\ell(\tilde{\pi}_{\text{aff}})$
$\bar{1} \underline{2} \underline{3}$	$\dots \boxed{4 \ 2 \ 3} \ 7 \ 5 \ 6 \ \dots$	2
$\underline{1} \bar{2} \underline{3}$	$\dots \boxed{1 \ 5 \ 3} \ 4 \ 8 \ 6 \ \dots$	2
$\underline{1} \underline{2} \bar{3}$	$\dots \boxed{1 \ 2 \ 6} \ 4 \ 5 \ 9 \ \dots$	2
$2 \ 1 \underline{3}$	$\dots \boxed{2 \ 4 \ 3} \ 5 \ 7 \ 6 \ \dots$	1
$\underline{1} \ 3 \ 2$	$\dots \boxed{1 \ 3 \ 5} \ 4 \ 6 \ 8 \ \dots$	1
$3 \underline{2} \ 1$	$\dots \boxed{3 \ 2 \ 4} \ 6 \ 5 \ 7 \ \dots$	1
$2 \ 3 \ 1$	$\dots \boxed{2 \ 3 \ 4} \ 5 \ 6 \ 7 \ \dots$	0

Figure 7.50. Decorated permutations $\tilde{\pi}$ on $b=3$ letters with $a=1$ anti-excedance; the corresponding (a, b) -bounded affine permutations $\tilde{\pi}_{\text{aff}}$; and the lengths $\ell(\tilde{\pi}_{\text{aff}})$ of these affine permutations.

Recall from Exercise 7.4.21 that the number of decorated permutations on b letters is equal to $b! \sum_{k=0}^b \frac{1}{k!}$. The following result, stated here without proof, refines this count by taking into account the number of anti-excedances.

Proposition 7.9.7 ([25, Theorem 4.1], [21, Proposition 23.1]). *Let $D_{a,b}$ be the number of decorated permutations on b letters with a anti-excedances (or the number of (a, b) -bounded affine permutations, cf. Lemma 7.9.6). Then*

$$D_{a,b} = \sum_{i=0}^{a-1} (-1)^i \binom{b}{i} ((a-i)^i (a-i+1)^{b-i} - (a-i-1)^i (a-i)^{b-i}),$$

$$\sum_{0 \leq a \leq b} D_{a,b} x^a \frac{y^b}{b!} = e^{xy} \frac{x-1}{x - e^{y(x-1)}}.$$

Definition 7.9.8. Let $\tilde{\pi}_{\text{aff}}$ be an (a, b) -bounded affine permutation, an affinization of a decorated permutation $\tilde{\pi}$, cf. Lemma 7.9.6. We refer to a position $i \in \mathbb{Z}$ such that $\tilde{\pi}_{\text{aff}}(i) \equiv i \pmod{b}$ (in other words, $\tilde{\pi}_{\text{aff}}(i) \in \{i, i+b\}$; and if $1 \leq i \leq b$ then $\tilde{\pi}(i) \in \{\underline{i}, \bar{i}\}$) as a *fixed point* of $\tilde{\pi}_{\text{aff}}$. If every $i \in \mathbb{Z}$ is a fixed point of $\tilde{\pi}_{\text{aff}}$, then we say that $\tilde{\pi}_{\text{aff}}$ is *equivalent to the identity modulo b* (or that $\tilde{\pi}$ is a decoration of the identity).

Lemma 7.9.9. *If $\tilde{\pi}_{\text{aff}}$ is not equivalent to the identity modulo b , then there exist $i, j \in \mathbb{Z}$ such that*

$$(7.9.4) \quad 1 \leq i < j \leq b,$$

$$(7.9.5) \quad \tilde{\pi}_{\text{aff}}(i) < \tilde{\pi}_{\text{aff}}(j),$$

$$(7.9.6) \quad \text{every position } h \text{ such that } i < h < j \text{ is a fixed point of } \tilde{\pi}_{\text{aff}}, \text{ and}$$

$$(7.9.7) \quad \text{neither } i \text{ nor } j \text{ are fixed points of } \tilde{\pi}_{\text{aff}}.$$

Proof. Suppose such a pair (i, j) does not exist. Let $i_1 < \dots < i_m$ be the elements of $\{1, \dots, b\}$ that are not fixed points of $\tilde{\pi}_{\text{aff}}$. Then

$$i_1 < \dots < i_m < \tilde{\pi}_{\text{aff}}(i_m) \leq \dots \leq \tilde{\pi}_{\text{aff}}(i_1).$$

We conclude that none of the values $\tilde{\pi}_{\text{aff}}(i_j)$ is of the form i_ℓ and consequently is of the form $i_\ell + b$. In particular, $\tilde{\pi}_{\text{aff}}(i_j) = i_m + b$ for some $j \neq m$. This implies $\tilde{\pi}_{\text{aff}}(i_j) > i_j + b$, a contradiction. \square

Definition 7.9.10. An *inversion* of $\tilde{\pi}_{\text{aff}}$ is a pair of integers (i, j) such that $i < j$ and $\tilde{\pi}_{\text{aff}}(i) > \tilde{\pi}_{\text{aff}}(j)$. Two inversions (i, j) and (i', j') are *equivalent* if $i' - i = j' - j \in b\mathbb{Z}$. The *length* $\ell(\tilde{\pi}_{\text{aff}})$ of $\tilde{\pi}_{\text{aff}}$ is the number of equivalence classes of inversions. (These classes correspond to the *alignments* of $\tilde{\pi}$, as defined in [21].) This number is finite since for any inversion (i, j) , we have $i < j < i + b$. Indeed, if $j \geq i + b$, then $\tilde{\pi}_{\text{aff}}(j) \geq j \geq i + b \geq \tilde{\pi}_{\text{aff}}(i)$. See Figure 7.50.

Lemma 7.9.11. *Let $\tilde{\pi}_{\text{aff}}$ be an (a, b) -bounded affine permutation that is equivalent to the identity modulo b . Then $\ell(\tilde{\pi}_{\text{aff}}) = a(b - a)$.*

Proof. Let $I = \{i \in \{1, \dots, b\} \mid \tilde{\pi}_{\text{aff}}(i) = i + b\}$ and $\underline{I} = \{i \in \{1, \dots, b\} \mid \tilde{\pi}_{\text{aff}}(i) = i\}$. Then $|I| = a$ and $|\underline{I}| = b - a$. The equivalence classes of inversions of $\tilde{\pi}_{\text{aff}}$ are described by the following list of representatives:

$$\{(i, j) \in I \times \underline{I} \mid 1 \leq i < j \leq b\} \cup \{(i, j + b) \mid (i, j) \in I \times \underline{I}, 1 \leq j < i \leq b\}.$$

The cardinality $\ell(\tilde{\pi}_{\text{aff}})$ of this set is equal to $|I \times \underline{I}| = a(b - a)$. \square

We next describe an algorithm for factoring affine permutations.

Definition 7.9.12. Let $\tilde{\pi}_{\text{aff}}$ be an (a, b) -bounded affine permutation. If $\tilde{\pi}_{\text{aff}}$ is not equivalent to the identity modulo b , then by Lemma 7.9.9, there exist positions $i, j \in \mathbb{Z}$ satisfying (7.9.4)–(7.9.7). We then swap the values of $\tilde{\pi}_{\text{aff}}$ in positions i and j (and more generally, in positions $i + mb$ and $j + mb$, for all $m \in \mathbb{Z}$). We repeat this procedure until we obtain an affine permutation that is equivalent to the identity modulo b . The algorithm terminates because each swap increases the length of the affine permutation by 1; this number is bounded by Definition 7.9.10. See Figure 7.51.

(i, j)	1	2	3	4	5	6
(34)	4	6	5	7	8	9
(23)	4	6	7	5	8	9
(12)	4	7	6	5	8	9
(56)	7	4	6	5	8	9
(45)	7	4	6	5	9	8
(34)	7	4	6	9	5	8
(46)	7	4	9	6	5	8
(24)	7	4	9	8	5	6
	7	8	9	4	5	6

Figure 7.51. Factoring $\tilde{\pi}_{\text{aff}} = (4, 6, 5, 7, 8, 9)$. The entries corresponding to the fixed points of $\tilde{\pi}_{\text{aff}}$ are boxed. Here $\tilde{\pi} = (4, 6, 5, 1, 2, 3)$.

Remark 7.9.13. In view of (7.9.3), the affine permutation at hand remains (a, b) -bounded after each step of the algorithm in Definition 7.9.12. It then follows from Lemma 7.9.11 that among all (a, b) -bounded affine permutations $\tilde{\pi}_{\text{aff}}$, the ones that have the maximal possible length $\ell(\tilde{\pi}_{\text{aff}}) = a(b - a)$ are precisely the ones that are equivalent to the identity modulo b .

7.10. Bridge decompositions

Bridge decompositions [2, Section 3.2] provide a useful recursive construction of reduced plabic graphs with a given decorated trip permutation.

Definition 7.10.1. A *bridge* is a graph fragment shown in Figure 7.52 on the left. Let $\tilde{\pi}$ be a decorated permutation on b letters that has a anti-excedances, and let $\tilde{\pi}_{\text{aff}}$ be the corresponding affine permutation. To build a plabic graph associated to $\tilde{\pi}_{\text{aff}}$, we begin by introducing a white (resp., black) lollipop in each position i with $\tilde{\pi}(i) = \bar{i}$ (resp., $\tilde{\pi}(i) = \underline{i}$). If $\tilde{\pi}$ is a decoration of the identity, we are done. Otherwise, we generate a sequence of transpositions (i, j) following the algorithm in Definition 7.9.12, then attach successive bridges in the corresponding positions, as in Figure 7.52. The resulting graph is called a *bridge decomposition* of $\tilde{\pi}_{\text{aff}}$, or sometimes a BCFW bridge decomposition, due to its relation with the Britto-Cachazo-Feng-Witten recursion in quantum field theory, see [2].




Figure 7.52. *Left:* a single bridge. *Right:* the bridge decomposition associated to the factorization constructed in Figure 7.51. The resulting plabic graph has trip permutation $\tilde{\pi} = (4, 6, 5, 1, 2, 3)$. Moreover, we have $\tilde{\pi} = (24)(46)(34)(45)(56)(12)(23)(34)$, the product of the transpositions (i, j) generated by the algorithm (reading right to left).

Proposition 7.10.2. A bridge decomposition of an (a, b) -bounded affine permutation $\tilde{\pi}_{\text{aff}}$ uses $a(b - a) - \ell(\tilde{\pi}_{\text{aff}})$ bridges.

Proof. See Lemma 7.9.11 and Definitions 7.9.12 and 7.10.1. □

Theorem 7.10.3. Let $\tilde{\pi}$ be a decorated permutation on b letters that has a anti-excedances. Let $\tilde{\pi}_{\text{aff}}$ be the associated (a, b) -bounded affine permutation. Then any bridge decomposition of $\tilde{\pi}_{\text{aff}}$ is a reduced plabic graph with the decorated trip permutation $\tilde{\pi}$.

Proof. We use induction on the number of bridges $\beta = a(b - a) - \ell(\tilde{\pi}_{\text{aff}})$. If $\beta = 0$, then $\tilde{\pi}$ is a decoration of the identity (see Remark 7.9.13), so the bridge decomposition consists entirely of lollipops, and we are done.

Now suppose that $\tilde{\pi}$ is not a decoration of the identity. Proceeding as in Definition 7.9.12, we construct a sequence of transpositions $\sigma_1, \sigma_2, \dots, \sigma_\beta$,

where $\sigma_1 = (ij)$ satisfies (7.9.4)–(7.9.7). Let G be the plabic graph obtained by attaching bridges according to $\sigma_1, \dots, \sigma_\beta$ (from top to bottom). By the induction assumption, attaching bridges according to $\sigma_2, \dots, \sigma_\beta$ as in Definition 7.10.1 produces a reduced plabic graph G' with the trip permutation $\tilde{\pi}' = \sigma_\beta \cdots \sigma_2$. This graph has $\beta - 1$ bridges and is obtained by removing the topmost horizontal edge e from G and applying local moves (M2) to remove the endpoints of e .

Conversely, G is obtained from G' by attaching a bridge in position (i, j) at the top of G' . (To illustrate, in Figure 7.52 we have $(i, j) = (3, 4)$.) When we add this bridge to G' , the trips starting at i and j get their “tails” swapped: the trip T_i (resp., T_j) in G that begins at i (resp., at j) traverses e and continues along the trip that used to begin at j (resp., at i) in G' ; all other trips remain the same. Hence the trip permutation of G is $\tilde{\pi}'\sigma_1 = \tilde{\pi}$.

It remains to show that G is reduced. It is tempting to try to use for this purpose the “bad features” criterion of Theorem 7.8.6. However, this theorem requires the plabic graph to be normal, so we will need to replace G by a suitable normal graph $N(G)$. It will also be more convenient to utilize the triple diagram version of the criterion, cf. Corollary 7.8.5.

We begin by constructing the normal graph $N(G)$ using Definition 7.8.2. (Note that stages 1, 4 and 6 of the algorithm are not needed since every vertex in G has degree 2 or 3, or is a lollipop that gets removed.) Up to the addition/removal of lollipops, the triple diagram $\mathfrak{X}(N(G))$ is isotopic to the triple diagram $\mathfrak{X}(G)$ constructed directly from G as in Definition 7.5.17. The graph G is reduced if and only if $N(G)$ is reduced, which is equivalent to the triple diagram $\mathfrak{X} = \mathfrak{X}(N(G)) = \mathfrak{X}(G)$ being minimal, or to \mathfrak{X} having no badgons, see Corollary 7.8.5. Thus, our goal is to show \mathfrak{X} has no badgons.

By the induction assumption, the triple diagram $\mathfrak{X}' = \mathfrak{X}(G')$ has no badgons. It follows that any potential badgon in \mathfrak{X} must involve the white endpoint w of edge e ; otherwise this feature would have already been present in \mathfrak{X}' . In particular, this means that w is trivalent, i.e., not an elbow of G .

A monogon in \mathfrak{X} would have to have its vertex at w . Three of the six half-strands at w run straight to or from the boundary, so we need to use two of the remaining three; moreover, those two half-strands have to be oppositely oriented. There are two such cases to consider. In one case, i would be a fixed point of $\tilde{\pi}$, contradicting our choice of (i, j) . In the other case, i would be a fixed point of $\tilde{\pi}' = \tilde{\pi}(G')$, which can also be ruled out since in that case, i would not participate in any bridge in G' , making it impossible to produce the bottom vertex of the vertical edge pointing downwards from w .

Finally, suppose that \mathfrak{X} contains a parallel digon. One of the vertices of the digon has to be w ; let w' denote the other vertex. Since the two sides of the digon are oriented in the same way at w , it follows that these sides lie on the strands S_i and S_j that start near the vertices i and j , respectively.

By our choice of bridge (i, j) , every h with $i < h < j$ is a fixed point of $\tilde{\pi}$, but i and j are not fixed points. It follows that S_i (resp. S_j) does not terminate at i (resp. j), and neither terminates between the boundary vertices i and j . We explained in the monogon case that S_j cannot terminate at i ; one can similarly argue that S_i cannot terminate at j . Also, neither S_i nor S_j intersects itself. Moreover the “tails” of S_i and S_j that start at the second vertex w' of the digon do not intersect each other (since otherwise a parallel digon would have been present in \mathfrak{X}'). It follows that either $\tilde{\pi}(i) < i$ and $\tilde{\pi}(j) > j$, or $\tilde{\pi}(i) > \tilde{\pi}(j) > j$, or $\tilde{\pi}(j) < \tilde{\pi}(i) < i$. In each case, we get $\tilde{\pi}_{\text{aff}}(i) > \tilde{\pi}_{\text{aff}}(j)$, which contradicts the way we chose i and j . \square

Corollary 7.10.4. *Let $\tilde{\pi}$ be a decorated permutation on b letters. Then there exists a reduced plabic graph whose decorated trip permutation is $\tilde{\pi}$.*

Proof. Use either Theorem 7.10.3 or the construction in Definition 7.6.8 (together with Proposition 7.5.12 and Theorem 7.7.3). \square

Corollary 7.10.5. *Let G be a reduced plabic graph with the decorated trip permutation $\tilde{\pi}$. If $\tilde{\pi}$ has b letters and a anti-excedances, then the number of faces in G is $a(b - a) - \ell(\tilde{\pi}_{\text{aff}}) + 1$.*

Proof. The number of faces is invariant under local moves. Therefore, by Theorem 7.4.25, it suffices to establish this formula for a particular reduced plabic graph with decorated trip permutation $\tilde{\pi}$. By Theorem 7.10.3, we can use a bridge decomposition of $\tilde{\pi}_{\text{aff}}$. Since each bridge adds one face to the graph, the claim follows by Proposition 7.10.2. \square

Let $\tilde{\pi}_{a,b}$ denote the decorated permutation on b letters defined by

$$(7.10.1) \quad \tilde{\pi}_{a,b} = (a + 1, a + 2, \dots, b, 1, 2, \dots, a).$$

Exercise 7.10.6. Let $\tilde{\pi}$ be a decorated permutation on b letters that has a anti-excedances. Show that if $\ell(\tilde{\pi}_{\text{aff}}) = 0$, then $\tilde{\pi} = \tilde{\pi}_{a,b}$.

Corollary 7.10.7. *Let G be a reduced plabic graph whose decorated trip permutation $\tilde{\pi}_G$ has b letters and a anti-excedances. Then G has at most $a(b - a) + 1$ faces. Moreover it has $a(b - a) + 1$ faces if and only if $\tilde{\pi}_G = \tilde{\pi}_{a,b}$.*

Proof. Immediate from Corollary 7.10.5 and Exercise 7.10.6. \square

Remark 7.10.8. The *permutohedron* \mathcal{P}_n [23, Exercise 4.64a] is a polytope whose $n!$ vertices are labeled by permutations in the symmetric group \mathcal{S}_n . Shortest paths in the 1-skeleton of \mathcal{P}_n encode reduced expressions in \mathcal{S}_n , and its 2-dimensional faces correspond to their local (braid) transformations, cf. Exercise 7.4.7. Similarly, paths in the 1-skeleton of the *bridge polytope* [26] encode bridge decompositions of the decorated permutation $\tilde{\pi}_{a,b}$; its 2-dimensional faces correspond to local moves in plabic graphs.

7.11. Edge labels of reduced plabic graphs

Definition 7.11.1. Let G be a plabic graph. Let us label the edges of G by subsets of integers that indicate which one-way trips traverse a given edge; more precisely, for each boundary vertex i , we include i in the label of every edge contained in the trip that starts at i . By Remark 7.4.15, each edge will be labeled by at most two integers. See Figure 7.53.

We say that G has the *resonance property* if after labeling the edges of G as in Definition 7.11.1, the following condition is satisfied at each internal vertex v that is not a lollipop:

- there exist numbers $i_1 < \dots < i_m$ such that the edges incident to v are labeled by the two-element sets $\{i_1, i_2\}, \{i_2, i_3\}, \dots, \{i_{m-1}, i_m\}, \{i_1, i_m\}$, appearing in clockwise order.

In particular, each edge of G that is not incident to a lollipop is labeled by a two-element subset. See Figure 7.53.




Figure 7.53. A reduced plabic graph from Figures 7.1(b) and 7.21. Its edge labeling exhibits the resonance property, see Definition 7.11.1. For example, the edge labels around the vertex v (resp., v'), listed in clockwise order, are $\{1, 2\}, \{2, 4\}, \{1, 4\}$ (resp., $\{1, 2\}, \{2, 5\}, \{1, 5\}$).

Remark 7.11.2. If a plabic graph has an internal leaf that is not a lollipop, then the resonance property fails. At a bivalent vertex v , the resonance condition is satisfied if and only if the two trips passing through v are distinct and none of them is a roundtrip.

Remark 7.11.3. If a plabic graph G is trivalent (apart from lollipops), then the resonance property is equivalent to the following requirement at each interior vertex v (other than a lollipop):

- the three edges incident to v have labels $\{a, b\}, \{a, c\}$, and $\{b, c\}$, for some $a < b < c$, and moreover this (lexicographic) ordering of labels corresponds to the counterclockwise direction around v .

For example, in Figure 7.53, the edge labels around the vertex v (resp., v') are, in lexicographic order, $\{1, 2\}, \{1, 4\}, \{2, 4\}$ (resp., $\{1, 2\}, \{1, 5\}, \{2, 5\}$). The three edges carrying these labels appear in the counterclockwise order around v (resp., v').

Exercise 7.11.4. Verify that none of the plabic graphs shown in Figure 7.48 (draw a disk around each of the fragments) satisfy the resonance property.

Theorem 7.11.5 ([16, Theorem 10.5]). *Let G be a plabic graph without internal leaves other than lollipops. Then G is reduced if and only if it has the resonance property.*

Theorem 7.11.5 is proved below in this section, following a few remarks and auxiliary lemmas.

Remark 7.11.6. Theorem 7.11.5 can be used to test whether a given plabic graph (potentially having internal leaves) is reduced or not. To this end, use the moves (M2) and (M3) to get rid of collapsible trees. If an internal leaf (not a lollipop) remains, then the graph is not reduced. Otherwise, one can apply the criterion in Theorem 7.11.5.

Remark 7.11.7. We find the resonance criterion of Theorem 7.11.5 easier to check than the “bad features” criterion of Theorem 7.8.6.

Remark 7.11.8. Certain reduced plabic graphs were realized as tropical curves in [16], where it was shown that the resonance property corresponds to the *balancing condition* for tropical curves.

Lemma 7.11.9. *The resonance property is preserved under the local moves (M1)–(M3), except when the decontraction move (M3) creates a new leaf.*

Proof. The square move (M1) only changes the labels of the sides of the square, see Figure 7.54. Moreover, the labels around each vertex match the labels around the opposite vertex after the square move, with the same cyclic order. Hence this move preserves the resonance property.

The case of the local move (M2) is easy, cf. Remark 7.11.2.

For the case of the local move (M3), see Figure 7.55. □




Figure 7.54. Transformation of edge labels under a square move (M1).




Figure 7.55. Transformation of edge labels under a local move (M3) at a 4-valent black vertex. (Alternatively, make all the vertices white.)

Lemma 7.11.10. *Any plabic graph obtained via the bridge decomposition construction (see Definition 7.10.1) has the resonance property.*

Proof. We will show that, more concretely, the edge labels around trivalent vertices in such a plabic graph G follow one of the patterns described in Figure 7.56. We will establish this result by induction on β , the number of bridges, following the strategy used in the proof of Theorem 7.10.3.




Figure 7.56. Edge labels near trivalent vertices in a bridge decomposition. At a white vertex, shown on the left, either $r < i < j$ or $i < j < r$. At a black vertex, shown on the right, either $s < i < j$ or $i < j < s$.

Let G be a bridge decomposition of $\tilde{\pi}_{\text{aff}}$, associated to the sequence of transpositions $\sigma_1, \dots, \sigma_\beta$. Thus $\tilde{\pi} = \tilde{\pi}_G = \sigma_\beta \cdots \sigma_1$. Here $\sigma_1 = (ij)$, where i and j satisfy (7.9.4)–(7.9.7). Let G' be the bridge decomposition associated to $\sigma_2, \dots, \sigma_\beta$, so that G is obtained from G' by adding a single bridge in position (i, j) at the top of G' .

Suppose the result is true for G' . We need to verify it for G . Let $r = \tilde{\pi}^{-1}(i)$ and $s = \tilde{\pi}^{-1}(j)$. Adding the bridge in position (i, j) at the top of G' adds at most two trivalent vertices: it adds a white (respectively, black) trivalent vertex provided that $r \neq j$ (respectively, $s \neq i$). The cases when one of the vertices on the bridge is bivalent are easy to verify, so we are going to assume that $r \neq j$ and $s \neq i$. In this case, the local configuration around positions i and j in G' and G is as shown in Figure 7.57.




Figure 7.57. The local configuration around positions i and j in G' and G .

Recall that $i < j$ and moreover any h such that $i < h < j$ is a fixed point of $\tilde{\pi}$. For the reasons indicated in the proof of Theorem 7.10.3, adding the bridge (i, j) at the top of G' has the effect of replacing the label i (resp., j) by j (resp., i) in every edge label outside of the bridge.

If G' has an edge with the label ij , then G has a bad double crossing involving the trips originating at i and j . This however is impossible since G is reduced, by Theorem 7.10.3. Therefore G' has no edge with label ij . Furthermore, G' has no edge with a label h for $i < h < j$. Since all trivalent vertices of G' satisfy the resonance condition of Figure 7.56, the same remains true after switching the labels i and j . Thus, all trivalent vertices of G that were present in G' satisfy this resonance condition.

Finally, the two new trivalent vertices in G satisfy this condition because $i < j$ and we can exclude $i < r < j$ and $i < s < j$ because of (7.9.6). \square

Proof of Theorem 7.11.5. We first establish the “if” direction. Let G be a plabic graph without internal leaves that are not lollipops. Suppose that G has the resonance property. We want to show that G is reduced.

Assume the contrary. By Proposition 7.4.9, G can be transformed by local moves that don’t create internal leaves into a plabic graph G' containing a hollow digon. Since G has the resonance property, so does G' , by Lemma 7.11.9. This yields a contradiction because the labels around a hollow digon do not satisfy the resonance property, see Figure 7.58.




Figure 7.58. A plabic graph containing a hollow digon fails to satisfy the resonance property.

We next establish the “only if” direction. Suppose that G is reduced, with $\tilde{\pi}_G = \tilde{\pi}$. We know from Theorem 7.10.3 that there is a bridge decomposition G' —a reduced plabic graph—with trip permutation $\tilde{\pi}$. By Lemma 7.11.10, G' has the resonance property. By Theorem 7.4.25, $G \sim G'$.

Since G has no internal leaves, there exists a trivalent plabic graph $G_3 \sim G$, see Lemma 7.3.2. Moreover it can be seen from the proof of this lemma that G and G_3 are related via moves that do not create internal leaves, so by Lemma 7.11.9 G has the resonance property if and only if G_3 does. Similarly, by removing bivalent vertices from G' , we obtain a trivalent graph $G'_3 \sim G'$ that has the resonance property. Since G_3 and G'_3 are trivalent and move-equivalent to each other, they are related via a sequence of (M1) and (M4) moves, see Theorem 7.3.5. By Lemma 7.11.9, these local moves preserve the resonance property, so G_3 has the resonance property and therefore G has it as well. \square

7.12. Face labels of reduced plabic graphs

In this section, we use the notion of a trip introduced in Definition 7.4.12 to label each face of a reduced plabic graph by a collection of positive integers. These face labels generalize the labeling of diagonals in a polygon by Plücker coordinates (cf. Section 1.2) as well as the labeling of faces in (double or ordinary) wiring diagrams by chamber minors (cf. Sections 1.3–1.4). In the following chapter, we will relate the face labels of reduced plabic graphs to Plücker coordinates that form an extended cluster for the standard cluster structure on a Grassmannian or, more generally, on a Schubert or positroid subvariety within it.

Remark 7.12.1. Let G be a reduced plabic graph. Let T_i be the one-way trip in G that begins at a boundary vertex i and ends at a boundary vertex j .

If $i \neq j$, then we claim that there are two kinds of faces in G : those on the left side of the trip T_i and those on the right side of it. If G is normal, then this claim follows from the fact (see Theorem 7.8.6) that G does not contain essential self-intersections. For a general reduced plabic graph, the claim can be deduced from the case of normal graphs using the procedure described in Definition 7.8.2.

If $i = j$, then by Proposition 7.4.22, the boundary vertex i is the root of a tree that collapses to a lollipop. If this lollipop is white (resp., black), then we declare that all faces of G lie on the left (resp., right) side of the trip T_i .

Definition 7.12.2. Let G be a reduced plabic graph with boundary vertices $1, \dots, b$. We define two natural face labelings of G , cf. Figure 7.59:

- in the *source labeling* $\mathcal{F}_{\text{source}}(G)$, each face f of G is labeled by the set

$$I_{\text{source}}(f) = \{i \mid f \text{ lies to the left of the trip starting at vertex } i\};$$

- in the *target labeling* $\mathcal{F}_{\text{target}}(G)$, each face f of G is labeled by the set

$$I_{\text{target}}(f) = \{i \mid f \text{ is to the left of the trip ending at vertex } i\}.$$

Remark 7.12.3. The edge labeling and the face labeling of a reduced plabic graph G are related as follows: if two faces f and f' of G are separated by a single edge whose edge label is $\{i, j\}$, then the face label of f' is obtained from that of f by either removing i and adding j , or removing j and adding i .

Theorem 7.12.4. Let G be a reduced plabic graph with b boundary vertices. Let a denote the number of anti-excedances in the trip permutation π_G . Let us label the faces of G using either the source or the target labeling. Then every face of G will be labeled by an a -element subset of $\{1, \dots, b\}$.




Figure 7.59. (a) The source labeling $\mathcal{F}_{\text{source}}(G)$ of a reduced plabic graph G . (b) The target labeling $\mathcal{F}_{\text{target}}(G)$. Here $\tilde{\pi}_G = (5, \underline{2}, \bar{3}, 6, 4, 1)$. Every face is labeled by a subset of cardinality 3, in agreement with Theorem 7.12.4, cf. Example 7.9.2.

Proof. By Theorem 7.11.5, every reduced plabic graph G has the resonance property, which in particular means that every edge label of G consists of two distinct numbers. It then follows from Remark 7.12.3 that every face label of G has the same cardinality. It remains to show that this cardinality is a , the number of anti-excedances of $\tilde{\pi} = \tilde{\pi}_G$.

Furthermore, it is sufficient to establish the latter claim for one particular reduced plabic graph with the trip permutation $\tilde{\pi}_G$, e.g., for a bridge decomposition of $\tilde{\pi}_{\text{aff}}$. Indeed, any two reduced plabic graphs with the same trip permutation are related by local moves, and any such move preserves all labels except at most one, see Exercise 7.12.5.

To prove the theorem for bridge decompositions, we use induction on the number of bridges β . In the base case $\beta = 0$, the bridge decomposition G consists of a white lollipops, $b - a$ black lollipops, and no bridges. Thus G has a single face, labeled by the a -element subset indicating the positions of the white lollipops.

Consider a bridge decomposition G built from a bridge decomposition G' by adding a bridge in position (i, j) at the top of G' , as in Figure 7.57. Both G and G' have trip permutations with a anti-excedances (cf. Remark 7.9.13), so by the induction assumption, the faces in G' have cardinality a . Since G inherits most of its faces from G' , and all face labels of G have the same cardinality, this cardinality is equal to a . \square

Exercise 7.12.5. Verify that applying a move (M2) or (M3) does not affect the face labels of a plabic graph, whereas applying the square move (M1) changes the face labels as shown in Figure 7.60.




Figure 7.60. The effect of the square move (M1) on the face labeling. Here i, j, k, l are the (source or target) labels of the trips that traverse the outer edges towards the central square; S is an arbitrary set of labels disjoint from $\{i, j, k, l\}$; and abS is a shorthand for the set $\{a, b\} \cup S$.

The face labelings of plabic graphs can be used to recover the labelings of diagonals in a polygon by Plücker coordinates as well as the labelings of chambers in (ordinary or double) wiring diagrams by minors:

Exercise 7.12.6. Let T be a triangulation of a convex m -gon \mathbf{P}_m , and let $G(T)$ be the plabic graph defined in Example 7.2.1. Explain how to label the boundary vertices of $G(T)$ in such a way that the face labeling of $G(T)$ recovers the labeling of diagonals of \mathbf{P}_m by pairs of integers. Cf. Figure 7.61.




Figure 7.61. A triangulation T of an octagon and the corresponding plabic graph $G(T)$, cf. Figure 2.2.

Exercise 7.12.7. Let D be a wiring diagram with m wires. Let $G(D)$ be the plabic graph defined in Example 7.2.4, see also Figure 7.62. Label the boundary vertices of $G(D)$ by the numbers $1, \dots, 2m$ in the clockwise order, starting with a 1 at the lower left boundary vertex of $G(D)$. Label the faces of $G(D)$ using the source labeling $\mathcal{F}_{\text{source}}(G)$. Show that intersecting each face label with the set $\{1, 2, \dots, m\}$ recovers the labeling of D by chamber minors. See Figure 7.62.




Figure 7.62. A wiring diagram D and the plabic graph $G(D)$ with the source labeling of its faces.

Exercise 7.12.8. Let D be a double wiring diagram with m pairs of wires. Let $G(D)$ be the plabic graph defined in Example 7.2.8. Label the boundary vertices of $G(D)$ by the numbers

$$1, 2, \dots, m - 1, m, m', \dots, 2', 1'$$

in clockwise order, starting with the label 1 at the lower left boundary vertex of $G(D)$. Label the faces of $G = G(D)$ using the source labeling $\mathcal{F}_{\text{source}}(G)$, so that each face gets labeled by $I' \cup J$, where $I' \subset \{1', \dots, m'\}$ and $J \subset \{1, \dots, m\}$. Let I denote the set obtained from I' by replacing each i' by i . Show that mapping each face label $I' \cup J$ to the pair $([1, m] \setminus I, \mathbf{J})$ recovers the labeling of D by chamber minors. See Figure 7.63.




Figure 7.63. Double wiring diagram labeling from plabic graphs. The labeling of a double wiring diagram D is obtained from the source labeling of the associated plabic graph $G(D)$ using the recipe described in Exercise 7.12.8.

7.13. Grassmann necklaces and weakly separated collections

Fix two nonnegative integers b and $a \leq b$. We denote by $\binom{[b]}{a}$ the set of all a -element subsets of $\{1, \dots, b\}$.

In this section, we provide an intrinsic combinatorial characterization of the subsets of $\binom{[b]}{a}$ that arise as sets of face labels of reduced plabic graphs. The proofs are omitted.

Definition 7.13.1 ([18]). We say that two a -element subsets $I, J \in \binom{[b]}{a}$ are *weakly separated* if there do not exist $i, j, i', j' \in \{1, \dots, b\}$ such that

- $i < j < i' < j'$ or $j < i' < j' < i$ or $i' < j' < i < j$ or $j' < i < j < i'$;
- $i, i' \in I \setminus J$ and $j, j' \in J \setminus I$.

Put differently, I and J are weakly separated if and only if, after drawing the numbers $1, 2, \dots, b$ clockwise around a circle, there exists a chord separating the sets $I \setminus J$ and $J \setminus I$ from each other.

Theorem 7.13.2 ([5, 19]). *Let I and J be target face labels of two faces in a reduced plabic graph. Then I and J are weakly separated.*

Definition 7.13.3. A collection $\mathcal{C} \subset \binom{[b]}{a}$ of a -element subsets of $[b]$ is called *weakly separated* if any $I, J \in \mathcal{C}$ are weakly separated. Thus, Theorem 7.13.2 asserts that the collection of target face labels of a reduced plabic graph is weakly separated. A weakly separated collection \mathcal{C} is called *maximal* if it is not contained in any other weakly separated collection. See Figure 7.64.




Figure 7.64. The target face labeling of the reduced plabic graph G from Figure 7.1(a). The set of labels $\{12, 23, 34, 45, 15, 24, 14\}$ is a maximal weakly separated collection in $\binom{[5]}{2}$. Here $\tilde{\pi}_G = \tilde{\pi}_{2,5} = (3, 4, 5, 1, 2)$.

Theorem 7.13.4 ([5, 19]). *For $\mathcal{C} \subset \binom{[b]}{a}$, the following are equivalent:*

- \mathcal{C} is a maximal weakly separated collection;
- \mathcal{C} is the set of target face labels of a reduced plabic graph G with $\tilde{\pi}_G = \tilde{\pi}_{a,b}$ (see (7.10.1)).

In that case, the cardinality of \mathcal{C} is equal to $|\mathcal{C}| = a(b - a) + 1$.

Remark 7.13.5. A general formula for the number of maximal weakly separated collections in $\binom{[b]}{a}$ is unknown. For $a = 2$, the maximal weakly separated collections in $\binom{[b]}{2}$ are in bijection with triangulations of a convex b -gon, so they are counted by the Catalan numbers C_{b-2} , where $C_n = \frac{1}{n+1} \binom{2n}{n}$. For $a = 3$ and $b = 6, \dots, 12$, the number of maximal weakly separated collections in $\binom{[b]}{3}$ is equal to 34, 259, 2136, 18600, 168565, 1574298, 15051702. See [6] for more data.

Theorem 7.13.4 can be generalized to arbitrary reduced plabic graphs. To state this result, we will need the following notion.

Definition 7.13.6 ([21, Definition 16.1]). A *Grassmann necklace* of type (a, b) is a sequence $\mathcal{I} = (I_1, \dots, I_b)$ of subsets $I_i \in \binom{[b]}{a}$ such that, for $i = 1, \dots, b$, we have $I_{i+1} \supset I_i \setminus \{i\}$. (Here the indices are taken modulo b , so that $I_1 \supset I_b \setminus \{b\}$.) Thus, if $i \notin I_i$, then $I_{i+1} = I_i$.

In other words, either $I_{i+1} = I_i$ or I_{i+1} is obtained from I_i by deleting i and adding another element. Note that if $I_{i+1} = I_i$, then either i belongs to all elements I_j of the necklace, or i belongs to none of them.

Example 7.13.7. The sequence $\mathcal{I} = (126, 236, 346, 456, 156, 126)$ is a Grassmann necklace of type $(3, 6)$.

Definition 7.13.8. For $\ell \in \{1, \dots, b\}$, we define the linear order $<_\ell$ on $\{1, \dots, b\}$ as follows:

$$\ell <_\ell \ell + 1 <_\ell \ell + 2 <_\ell \dots <_\ell b <_\ell 1 <_\ell \dots <_\ell \ell - 1.$$

For a decorated permutation $\tilde{\pi}$ on b letters, we say that $i \in \{1, \dots, b\}$ is an ℓ -anti-excedance of $\tilde{\pi}$ if either $\tilde{\pi}^{-1}(i) >_\ell i$ or if $\tilde{\pi}(i) = \bar{i}$. Thus, a 1-anti-excedance is the same as an (ordinary) anti-excedance, as in Definition 7.9.1.

It is not hard to see that the number of ℓ -anti-excedances does not depend on the choice of $\ell \in \{1, \dots, b\}$, so we simply refer to this quantity as the number of anti-excedances.

Lemma 7.13.9. *Decorated permutations on b letters with a anti-excedances are in bijection with Grassmann necklaces \mathcal{I} of type (a, b) .*

Proof. To go from \mathcal{I} to the corresponding decorated permutation $\tilde{\pi} = \tilde{\pi}(\mathcal{I})$, we set $\tilde{\pi}(i) = j$ whenever $I_{i+1} = (I_i \setminus \{i\}) \cup \{j\}$ for $i \neq j$. If $i \notin I_i = I_{i+1}$ then $\tilde{\pi}(i) = \bar{i}$, and if $i \in I_i = I_{i+1}$ then $\tilde{\pi}(i) = \bar{i}$.

Going in the other direction, let $\tilde{\pi}$ be a decorated permutation. For $\ell \in \{1, \dots, b\}$, we denote by I_ℓ the set of ℓ -anti-excedances of $\tilde{\pi}$. Then $\mathcal{I} = \mathcal{I}(\tilde{\pi}) = (I_1, \dots, I_b)$ is the corresponding Grassmann necklace. \square

Example 7.13.10. Let $\mathcal{I} = (126, 236, 346, 456, 156, 126)$, cf. Example 7.13.7. Then $\tilde{\pi}(\mathcal{I}) = (3, 4, 5, 1, 2, \bar{6})$.

Example 7.13.11. Let $\tilde{\pi}_G = \tilde{\pi}_{a,b}$, cf. (7.10.1). The corresponding Grassmann necklace (cf. Lemma 7.13.9) is given by

$$(7.13.1) \quad \mathcal{I}(\tilde{\pi}_{a,b}) = (\{1, 2, \dots, a\}, \{2, 3, \dots, a, a+1\}, \dots, \{b, 1, 2, \dots, a-1\}).$$

Definition 7.13.12. We extend the linear order $<_\ell$ on $\{1, \dots, b\}$ to a partial order on $\binom{[b]}{a}$, as follows. Let

$$\begin{aligned} I &= \{i_1, \dots, i_a\}, & i_1 <_\ell i_2 <_\ell \dots <_\ell i_a; \\ J &= \{j_1, \dots, j_a\}, & j_1 <_\ell j_2 <_\ell \dots <_\ell j_a. \end{aligned}$$

Then, by definition, $I \leq_\ell J$ if and only if $i_1 \leq_\ell j_1, \dots, i_a \leq_\ell j_a$.

Definition 7.13.13. For a Grassmann necklace $\mathcal{I} = (I_1, \dots, I_b)$ of type (a, b) , we define the associated *positroid* $\mathcal{M}_{\mathcal{I}}$ by

$$\mathcal{M}_{\mathcal{I}} = \{J \in \binom{[b]}{a} \mid I_\ell \leq_\ell J \text{ for all } \ell \in \{1, \dots, b\}\}.$$

As we will see in Chapter 8, positroids are the (realizable) *matroids* that arise from full rank $a \times b$ matrices with all Plücker coordinate nonnegative. Abstractly, one may also define a *positively oriented matroid* to be an oriented matroid on $\{1, 2, \dots, b\}$ whose chirotope takes nonnegative values on any ordered subset $\{i_1 < \dots < i_a\}$. By [1], these two notions are the same, in other words, every positively oriented matroid is realizable.

Example 7.13.14. Let $\mathcal{I} = \mathcal{I}(\tilde{\pi}_{a,b})$, see (7.13.1). Then $\mathcal{M}_{\mathcal{I}} = \binom{[b]}{a}$, i.e., the positroid associated with \mathcal{I} contains all a -element subsets of $\{1, \dots, b\}$.

Definition 7.13.15. Two reduced plabic graphs are called *strongly equivalent* if they have the same sets of face labels.

We note that two plabic graphs which are connected via moves (M2) and (M3) are strongly equivalent.

Recall that $\mathcal{F}_{\text{target}}(G)$ denotes the collection of target-labels of faces of a reduced plabic graph G .

Theorem 7.13.16 ([19, Theorem 1.5]). *Fix a decorated permutation $\tilde{\pi}$ on b letters with a anti-excedances. Let \mathcal{I} be the corresponding Grassmann necklace of type (a, b) , cf. Lemma 7.13.9. Let $\mathcal{M}_{\mathcal{I}}$ be the associated positroid, cf. Definition 7.13.13. Then the map $G \mapsto \mathcal{F}_{\text{target}}(G)$ gives a bijection between*

- the strong equivalence classes of reduced plabic graphs G with decorated trip permutation $\tilde{\pi}_G = \tilde{\pi}$ and
- the collections $\mathcal{C} \subset \binom{[b]}{a}$ that are maximal (with respect to inclusion) among the weakly separated collections satisfying $\mathcal{I} \subseteq \mathcal{C} \subseteq \mathcal{M}_{\mathcal{I}}$.

Remark 7.13.17. Let $\mathcal{I} = \mathcal{I}(\tilde{\pi}_{a,b})$. Then $\mathcal{I} = \mathcal{I}(\tilde{\pi}_{a,b})$ is given by (7.13.1). Each of the b cyclically consecutive subsets in (7.13.1) is weakly separated from every other a -element subset of $\{1, \dots, b\}$, so every maximal weakly separated collection $\mathcal{C} \subset \binom{[b]}{a}$ must contain \mathcal{I} . Furthermore, $\mathcal{M}_{\mathcal{I}} = \binom{[b]}{a}$ (see Example 7.13.14), so any such \mathcal{C} automatically satisfies the inclusions $\mathcal{I} \subseteq \mathcal{C} \subseteq \mathcal{M}_{\mathcal{I}}$. We thus recover Theorem 7.13.4 as a special case of Theorem 7.13.16.

Bibliography

- [1] ARDILA, F., RINCÓN, F., AND WILLIAMS, L. Positively oriented matroids are realizable. *J. Eur. Math. Soc. (JEMS)* 19, 3 (2017), 815–833.
- [2] ARKANI-HAMED, N., BOURJAILY, J., CACHAZO, F., GONCHAROV, A., POSTNIKOV, A., AND TRNKA, J. *Grassmannian geometry of scattering amplitudes*. Cambridge University Press, Cambridge, 2016.
- [3] BOCKLANDT, R. Calabi-Yau algebras and weighted quiver polyhedra. *Math. Z.* 273, 1-2 (2013), 311–329.
- [4] BROOMHEAD, N. Dimer models and Calabi-Yau algebras. *Mem. Amer. Math. Soc.* 215, 1011 (2012), viii+86.
- [5] DANILOV, V. I., KARZANOV, A. V., AND KOSHEVOY, G. A. On maximal weakly separated set-systems. *J. Algebraic Combin.* 32, 4 (2010), 497–531.
- [6] EARLY, N. From weakly separated collections to matroid subdivisions. arXiv:1910.11522.
- [7] FOCK, V., AND GONCHAROV, A. Moduli spaces of local systems and higher Teichmüller theory. *Publ. Math. Inst. Hautes Études Sci.*, 103 (2006), 1–211.
- [8] FOMIN, S., IGUSA, K., AND LEE, K. Universal quivers. arXiv:2003.01244.
- [9] FOMIN, S., PYLYAVSKYY, P., SHUSTIN, E., AND THURSTON, D. Morsifications and mutations. arXiv:1711.10598.
- [10] GALASHIN, P. Plabic graphs and zonotopal tilings. *Proc. Lond. Math. Soc. (3)* 117, 4 (2018), 661–681.
- [11] HARER, J. L. The virtual cohomological dimension of the mapping class group of an orientable surface. *Invent. Math.* 84, 1 (1986), 157–176.
- [12] HATCHER, A. On triangulations of surfaces. *Topology Appl.* 40, 2 (1991), 189–194.
- [13] KAWAMURA, T. Links associated with generic immersions of graphs. *Algebr. Geom. Topol.* 4 (2004), 571–594.
- [14] KENYON, R. An introduction to the dimer model. In *School and Conference on Probability Theory*, ICTP Lect. Notes, XVII. Abdus Salam Int. Cent. Theoret. Phys., Trieste, 2004, pp. 267–304.
- [15] KNUTSON, A., LAM, T., AND SPEYER, D. E. Positroid varieties: juggling and geometry. *Compos. Math.* 149, 10 (2013), 1710–1752.

- [16] KODAMA, Y., AND WILLIAMS, L. KP solitons and total positivity for the Grassmannian. *Invent. Math.* 198, 3 (2014), 637–699.
- [17] KODAMA, Y., AND WILLIAMS, L. K. KP solitons, total positivity, and cluster algebras. *Proc. Natl. Acad. Sci. USA* 108, 22 (2011), 8984–8989.
- [18] LECLERC, B., AND ZELEVINSKY, A. Quasicommuting families of quantum Plücker coordinates. In *Kirillov’s seminar on representation theory*, vol. 181 of *Amer. Math. Soc. Transl. Ser. 2*. Amer. Math. Soc., Providence, RI, 1998, pp. 85–108.
- [19] OH, S., POSTNIKOV, A., AND SPEYER, D. E. Weak separation and plabic graphs. *Proc. Lond. Math. Soc. (3)* 110, 3 (2015), 721–754.
- [20] OH, S. H., AND SPEYER, D. E. Links in the complex of weakly separated collections. *J. Comb.* 8, 4 (2017), 581–592.
- [21] POSTNIKOV, A. Total positivity, Grassmannians, and networks, [arXiv:math/0609764](https://arxiv.org/abs/math/0609764).
- [22] POSTNIKOV, A. Positive Grassmannian and polyhedral subdivisions. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. IV. Invited lectures* (2018), World Sci. Publ., Hackensack, NJ, pp. 3181–3211.
- [23] STANLEY, R. P. *Enumerative combinatorics. Vol. 1*, second ed., vol. 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2012.
- [24] THURSTON, D. P. From dominoes to hexagons. In *Proceedings of the 2014 Maui and 2015 Qinhuangdao conferences in honour of Vaughan F. R. Jones’ 60th birthday* (2017), vol. 46 of *Proc. Centre Math. Appl. Austral. Nat. Univ.*, Austral. Nat. Univ., Canberra, pp. 399–414. Preprint versions: [arXiv:math/0405482](https://arxiv.org/abs/math/0405482) (2004 and 2016).
- [25] WILLIAMS, L. K. Enumeration of totally positive Grassmann cells. *Adv. Math.* 190, 2 (2005), 319–342.
- [26] WILLIAMS, L. K. A positive Grassmannian analogue of the permutohedron. *Proc. Amer. Math. Soc.* 144, 6 (2016), 2419–2436.