ABSTRACT

VIEL, SHIRA COLEMAN POLSTER. Cluster Algebras and Mutation-Linear Algebra: Folding, Dominance, and the Orbifolds Model. (Under the direction of Nathan Reading.)

We study folding and dominance relationships between exchange matrices and the implications thereof on a variety of associated cluster algebra and mutation-linear algebra objects. The orbifolds model serves as a primary tool.

This dissertation begins by examining the folding of an arbitrary exchange matrix B to another exchange matrix C under a symmetry: we show that the mutation fan for B folds to a refinement of that for C and that symmetric bases for the mutation-linear structures defined by B fold to spanning sets for those defined by C. Furthermore, we prove that the folded mutation fan for B coincides with that for C if and only if a symmetric positive basis for $\mathbb{R}^B$ folds to a positive basis for $\mathbb{R}^C$, and establish that these results hold when B is of finite type.

Next, we study the orbifolds model of cluster algebras, a generalization of the marked surfaces model to skew-symmetrizable exchange matrices. We broaden the construction, from the marked surfaces setting, of mutation fans, mutation-linear bases, and universal geometric coefficients for orbifolds that satisfy certain properties. We then establish these properties for particular classes of orbifolds, including those of finite and affine type. Finally, we use the combinatorics of the model, which encompass the notions of folding and unfolding, to prove that the mutation fan and positive bases fold in this context.

We conclude the dissertation by using the orbifolds model to prove examples of several phenomena associated with a dominance relationship between exchange matrices B and B'. In particular, we show that these phenomena arise as a consequence of resection, a simple combinatorial operation on orbifolds which both broadens and extends the notion of surface resection. We prove that when B and B' are realized in this manner, the mutation fan for B refines that for B' and the identity map from $\mathbb{Q}^B$ to $\mathbb{Q}^{B'}$ is mutation-linear. If B is acyclic and arises from a triangulation of a disk with one orbifold point, we further establish an explicit injective, $\mathbf{g}$-vector preserving homomorphism from the principal-coefficients cluster algebra for B' into that for B.

Cluster Algebras and Mutation-Linear Algebra: Folding, Dominance, and the Orbifolds Model

by
Shira Coleman Polster Viel

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Mathematics

Raleigh, North Carolina

2018

APPROVED BY:

Patricia Hersh	Kailash Misra
Justin Post	Seth Sullivant

Nathan Reading
Chair of Advisory Committee

DEDICATION

To Isaac.

BIOGRAPHY

Shira grew up in Shaker Heights, Ohio and studied mathematics and theatre as an undergraduate at Rice University. She worked at a non-profit math educational software company before completing Smith College's post-baccalaureate program in mathematics and then beginning graduate studies at North Carolina State University. At the time of this writing, she looks forward to an appointment as an Assistant Professor of the Practice in the Department of Mathematics at Duke University. Outside of doing and teaching mathematics, Shira enjoys baking, solving puzzles, and being outdoors (preferably with a disc).

ACKNOWLEDGEMENTS

The author thanks her advisor, Nathan Reading, for being an outstanding mentor, teacher, and all-around inspirational figure, and for the numerous beneficial conversations and editorial suggestions without which this dissertation would not exist.

The author thanks Anna Felikson for her always helpful and timely correspondence, and the members of the Reading research group, including Emily Barnard, Chetak Hossain, Shirley Law, Emily Meehan, and Salvatore Stella, for many useful conversations.

The author thanks the Graduate School at North Carolina State University for her Doctoral Dissertation Completion Grant and Mike Carter for his writing advice.

Finally, the author thanks her officemates, friends, family, and Brian for their ongoing encouragement, support, and company.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

Introduction

1.1 Cluster algebras

A ***cluster algebra*** is a commutative ring of rational functions with a certain distinguished set of generators. Rather than being presented at the outset in their entirety, the generators are recursively defined.

Cluster algebras were introduced in [10] by Fomin and Zelevinsky, who noticed that the coordinate rings of many algebraic varieties carry a cluster algebra structure. (Two archetypal examples of cluster algebras which are subrings of fields of rational functions in one variable are $\mathbb{C}[SL_2]$ and $\mathbb{C}[SL_3/N]$, where SL_k denotes the group of $k \times k$ matrices over $\mathbb{C}$ with determinant 1, and N is the maximal unipotent subgroup of all unipotent upper triangular matrices. More general examples include the homogeneous coordinate rings of Grassmanian and Schubert varieties.) Although the conception of cluster algebras in [10–13] was as an algebraic and combinatorial framework for studying total positivity and dual canonical bases in semisimple groups, connections and applications beyond Lie theory and algebraic geometry have arisen in areas of mathematics as diverse as combinatorics, quiver representations, tropical geometry, Teichmüler theory, mirror symmetry, Poisson geometry, and integrable systems.

This dissertation focuses on the rich intrinsic combinatorial structure of cluster algebras, and in particular, of ***cluster algebras of geometric type***. Detailed background is provided in Section 2.2.1, but we give a brief outline of the definition here. A cluster algebra of geometric type of ***rank*** n is a subring of an ***ambient field*** $\mathbb{K}$ of rational functions in n variables, uniquely determined by a choice of an initial ***seed***. A seed consists of two pieces of data: first, a collection of n distinguished, algebraically-independent elements of $\mathbb{K}$, called ***cluster variables***, and second, an ***extended exchange matrix*** $\tilde{B}$, formed from an $n \times n$ skew-symmetrizable integer matrix B, called an ***exchange matrix***, extended by a collection of ***coefficient rows*** which specify ***coefficients***. New seeds are constructed iteratively from the initial seed through a process called ***mutation***: from any seed, one can mutate in any direction $k \in [n]$ (where $[n]$ denotes the set $\{1, \ldots, n\}$). This produces another seed, obtained from the first by replacing the k^{th} cluster variable with a new element of $\mathbb{K}$ and replacing $\tilde{B}$ with a new extended exchange matrix $\mu_k(\tilde{B})$. (The equation relating the new cluster variable to the old is called an ***exchange relation***: it involves only the cluster variables and coefficients in the old seed and is fully determined by the old extended exchange matrix $\tilde{B}$. The operation μ_k on $\tilde{B}$ by which the new extended exchange matrix is obtained is called ***matrix mutation*** in direction k.) The cluster algebra is defined to be the subring of $\mathbb{K}$ generated by the complete set of cluster variables in all seeds. These generators are all rational functions in the initial cluster variables with coefficients given by monomials in the initial coefficients, and are grouped into ***clusters*** of size n (one in each seed): overlapping, algebraically-independent subsets of $\mathbb{K}$. Up to isomorphism, the algebra is completely determined by the extended exchange matrix $\tilde{B}$, so is often denoted by $\mathcal{A}(\tilde{B})$.

A cluster algebra is of ***finite type*** if it has finitely many cluster variables: a complete classification of such cluster algebras is achieved in [11]. Whether or not a given cluster algebra is of finite type depends exclusively on the exchange matrix, so exchange matrices are referred to as being of finite type themselves. As proven in [11], these finite type exchange matrices correspond to to (generalized) Cartan matrices of finite type. (Specifically, if an exchange matrix is of finite type, then one can find a sequence of matrix mutations such that the resulting exchange matrix B is ***acyclic*** (that is, the rows and columns may be reindexed so that $b_{ij} > 0$ whenever $i > j$). One then passes to the corresponding Cartan matrix by placing 2's along the diagonal and making all off-diagonal entries negative.) The classification of cluster algebras of finite type therefore coincides with the Cartan-Killing classification of semisimple Lie algebras and finite root systems by (directed) finite Dynkin diagrams. (For example, the literature often refers to cluster algebras of type A_n, B_n, C_n, or D_n, meaning the (finite type) cluster algebras whose exchange matrices correspond to Cartan matrices of the given type.) More generally, [6, 11] describe how every (skew-symmetrizable) exchange matrix can be encoded by a ***diagram***, a generalization of a quiver or directed graph. An exchange matrix (and any associated cluster algebra) is of ***affine type*** if there exists a sequence of matrix mutations such that the resulting acyclic exchange matrix corresponds to a Cartan matrix of affine type.

1.1.1 Folding

Given the correspondence between Dynkin diagrams and cluster algebras of finite and affine type, it is natural that the tools from Kac-Moody Lie theory of folding a diagram or root system under an automorphism (see, for example, [16, 27]) extend to cluster algebras. The literature on cluster algebra folding includes [1–4, 6, 15, 18].

This dissertation takes its conventions from [4] in defining ***exchange matrix folding***: details are provided in Section 2.3, but the key ideas are as follows. Suppose $B = [b_{ij}]$ is an $n \times n$ exchange matrix equipped with an ***automorphism*** $\sigma \in \mathfrak{S}_n$. That is, $b_{\sigma(i)\sigma(j)} = b_{ij}$ for each $i, j \in [n]$. If the symmetry σ satisfies two ***admissability criteria*** on B, and furthermore, these criteria are preserved under symmetric sequences of mutations, then σ is a ***stable automorphism*** and B admits a folding under σ. The folding of B under σ, denoted by $\pi_\sigma(B)$, is again an exchange matrix, with rows and columns indexed by σ-orbits $\bar{\imath}$ of $[n]$ and entries defined by taking sums down columns of B within a given σ-orbit. (The matrix B is also referred to as the ***unfolding*** of $\pi_\sigma(B)$ under σ: indeed, [1, 6, 15] are formulated in terms of unfolding.)

It is proven in [4] that there is a nice relationship between the ***coefficient-free*** cluster algebras $\mathcal{A}_0(B)$ and $\mathcal{A}_0(\pi_\sigma(B))$, with the latter arising as a cluster-structure preserving quotient of a subalgebra of the former. (The coefficient-free cluster algebra $\mathcal{A}_0(B)$ is defined by taking $\tilde{B} = B$, with no coefficient rows.) This relationship is extended to the ***principal-coefficients*** cluster algebras $\mathcal{A}_\bullet(B)$ and $\mathcal{A}_\bullet(\pi_\sigma(B))$, where the coefficient rows of the initial extended exchange matrices are taken to be the identity matrix, in [1].

Some classical Lie foldings include the finite type foldings $A_{2n-1} \to C_n$ and $D_{n+1} \to B_n$, and the affine type foldings $A^{(1)}_{2n-1} \to C^{(1)}_n$ and $D^{(1)}_{n+1} \to B^{(1)}_n$ (notation from [17, Tables Fin and Aff 1]). It is proven in [2] that every admissable automorphism on an acyclic exchange matrix is stable, so these classical foldings all extend to foldings of exchange matrices.

1.1.2 Questions of interest

Among the many areas of interest in the study of cluster algebras, we highlight three addressed herein.

Models of cluster algebras

The recursive mutation process is complex, so much effort has been dedicated to building models. In [8, 9], Fomin, Shapiro and Thurston developed the idea of cluster algebras arising from ***marked surfaces***, whereby the combinatorics of an algebra are modeled using the geometry/topology of certain types of curves on a surface with marked points. Cluster variables correspond to curves which connect marked points, called ***arcs***, and clusters correspond to maximal pairwise-compatible collections of arcs, called ***triangulations***. Exchange matrices are given by the skew-symmetric ***signed-adjacency matrices*** of triangulations, with mutation corresponding to ***flips*** of arcs in the triangulation. Coefficients are specified by another special class of curves in the surface, called ***allowable curves***.

The marked surfaces model was extended by Felikson, Shapiro, and Tumarkin in [5,7] to ***orbifolds*** as a way to provide geometric realizations of cluster algebras with skew-symmetrizable exchange matrices. Roughly speaking, one might think of an orbifold in this setting as the quotient space, or "folding," of a marked surface under action by a symmetry of the surface. This results in both marked points as well as a new collection of points called ***orbifold points***, the images of fixed points of the symmetry. Indeed, the orbifolds model was introduced in the context of [6], which classified cluster algebras of ***finite mutation type*** using a notion of exchange matrix folding. One can also fold orbifolds to other orbifolds, and it is shown in [5] that this topological folding corresponds to the folding of the associated signed-adjacency matrices. (For example, the finite type folding $A_{2n-1} \to C_n$ corresponds to the folding of a disk with $2n$ points on its boundary (and no marked points or orbifold points in its interior) under a central 2-fold symmetry to a disk with n points on its boundary and one orbifold point in its interior.) Detailed background on the marked surface and orbifold models is provided in Section 3.2; orbifold folding is covered in depth in Section 3.5.

This dissertation advances the development of the orbifolds model by describing the generalization of two properties of surfaces introduced in [22], the ***Curve Separation Property*** and the ***Null Tangle Property***, to the orbifolds setting (see Section 3.4.) We establish both properties for certain classes of orbifolds in Theorems 3.4.5 and 3.4.22, respectively, and confirm that, as for surfaces, the allowable curves on triangulated orbifolds with these properties can be used to construct various objects of interest associated to the corresponding signed-adjacency matrices and cluster algebras (see Sections 1.1.2 and 1.2 below).

Relationships between cluster algebras

A common goal in mathematics is to define and exploit relationships between objects of interest: here, we consider both relationships among cluster algebras with the same exchange matrix as well as between cluster algebras with different exchange matrices. An example of the former type of relationship is discussed in the next section. One example of the latter type of relationship is when one exchange matrix is a folding of the other (see Section 1.1.1 above); another example, which we now discuss, is when one exchange matrix ***dominates*** the other.

Given two $n \times n$ exchange matrices $B = [b_{ij}]$ and $B' = [b'_{ij}]$, the matrix B dominates B' if the entries b_{ij} and b'_{ij} weakly agree in sign and $|b_{ij}| \geq |b'_{ij}|$ for each $i, j \in [n]$. The dominance relationship is introduced in [24], which describes several interesting phenomena that may occur as consequences thereof. One such phenomenon is the explicit realization of an injective, **g*-vector*** preserving ring homomorphism from the principal-coefficients cluster algebra $\mathcal{A}_\bullet(B')$ into $\mathcal{A}_\bullet(B)$. (The $\mathbf{g}$-vector is a $\mathbb{Z}^n$-grading of the cluster algebra.)

This dissertation contributes to the very new study of dominance phenomena by describing a simple combinatorial operation on triangulated orbifolds, called ***resection***, which induces a dominance relationship on the associated signed-adjacency matrices. Resection involves both an extension of a similar operation on marked surfaces from [24] to orbifolds as well as the introduction of a new operation unique to the orbifold setting. Consider an orbifold consisting of a disk with n marked points on its boundary and a single orbifold point in the disk's interior. We show in Section 4.4 that any exchange matrix B' dominated by the acyclic signed-adjacency matrix B of a triangulation of such an orbifold can be obtained by resection. Theorem 4.1.5 then establishes the explicit injective, $\mathbf{g}$-vector preserving ring homomorphism from $\mathcal{A}_\bullet(B')$ into $\mathcal{A}_\bullet(B)$. This completes the proof of [24, Expected Theorem 1.17], which claims that such homomorphisms may be established for any exchange matrix B' dominated by an acyclic exchange matrix B of finite type. We also use the resection operation in Theorems 4.1.1 and 4.1.2 to prove examples of two other dominance phenomena (see Section 1.2 below).

Universal geometric cluster algebras

Continuing the vein of relationships between cluster algebras, there is a desire to find a universal object among all cluster algebras with the same exchange matrix B. That is, does there exist a choice of coefficient rows (with entries in some subring R of $\mathbb{R}$) determining an extended exchange matrix $\tilde{B}_{\text{un}}$ such that any other cluster algebra $\mathcal{A}(\tilde{B})$ with exchange matrix B may be obtained from $\mathcal{A}(\tilde{B}_{\text{un}})$ by ***coefficient specialization***? If so, then $\mathcal{A}_{\text{un}}(B) = \mathcal{A}(\tilde{B}_{\text{un}})$ is called the ***universal geometric cluster algebra for*** B (over R) and the coefficient rows of $\tilde{B}_{\text{un}}$ are called ***universal geometric coefficients***. The definition originates in [13], and is broadened in [21] to encompass the possibilities of both infinitely many coefficient rows as well as coefficient entries in subrings R of $\mathbb{R}$ other than $\mathbb{Z}$.

Theorem 3.4.17 of this dissertation shows that given a triangulation of an orbifold with the Null Tangle Property, universal geometric coefficients for the associated signed-adjacency matrix (over $\mathbb{Z}$ or $\mathbb{Q}$) are given by the ***shear coordinates*** of allowable curves. (Shear coordinates are integer vectors encoding the interaction between a curve and the arcs in a triangulation.) This generalizes the analogous result for marked surfaces in [22]. Thanks to the establishment of the Null Tangle Property for certain classes of orbifolds in Theorem 3.4.22, classes that include the orbifolds of finite and affine type, this provides a means for constructing universal geometric coefficients in these cases. (An orbifold of finite (resp. affine) type is one whose triangulations have finite (resp. affine) type signed-adjacency matrices.) While the finite type cases are known [21, Section 10], the realization in affine type is new, and provides significant evidence towards [21, Conjecture 10.15]. Furthermore, we show in Corollary 3.6.3 that if one orbifold folds to another and both have the Null Tangle Property, then the set of shear coordinates of allowable curves on the first orbifold fold to universal geometric coefficients for the signed-adjacency matrix of the second.

1.2 Mutation-linear algebra

Closely related to cluster algebras of geometric type is the notion of ***mutation-linear algebra***. Although the study of mutation-linear algebra was initiated in [21] and continued in [22], it was not explicitly named until [24].

Let B be an $n \times n$ exchange matrix, and given an arbitrary vector $\mathbf{a} \in \mathbb{R}^n$, define $\tilde{B}$ to be the exchange matrix B extended by a single coefficient row $\mathbf{a}$. Recall that one can perform matrix mutation on $\tilde{B}$ in any sequence $\mathbf{k}$ of directions in $[n]$: reading off the coefficient row of the resulting matrix defines a map $\eta_{\mathbf{k}}^B$ on $\mathbb{R}^n$ which is continuous, piecewise linear, and invertible. The collection of all such homeomorphisms from $\mathbb{R}^n$ to itself (as $\mathbf{k}$ varies) is called the set of ***mutation maps*** associated to B. Mutation-linear algebra is the study of ***B-coherent linear relations***, linear relations among vectors which are preserved under the application of mutation maps. The name comes from the analogy with linear algebra, which can be thought of as the study of linear relations with no restrictions. Given a ring $R \in \{\mathbb{Z}, \mathbb{Q}, \mathbb{R}\}$, we write R^B to indicate the ***mutation-linear structure*** on R^n provided by restricting from all linear relations on R^n to the subset of B-coherent linear relations with coefficients in R.

As in linear algebra, there are mutation-linear algebra notions of ***bases for*** R^B, collections of vectors in R^n which are spanning and independent with respect to B-coherent linear relations, and ***mutation-linear maps*** from R^B to $R^{B'}$, maps between mutation-linear structures given by different exchange matrices which preserve B-coherent linear relations. Another important object in mutation-linear algebra is the ***mutation fan*** $\mathcal{F}_B$, a complete fan in $\mathbb{R}^n$ which encodes the combinatorics and piecewise-linear geometry of the mutation maps associated to B. This dissertation makes use of many of the explicit connections between the mutation-linear algebra of R^B and geometric cluster algebras with exchange matrix B and coefficients in R established in [21, 22, 24]. For example, the fan of $\mathbf{g}$-vectors for B^T is a subfan of the mutation fan $\mathcal{F}_B$, and a collection of vectors in R^n constitute universal geometric coefficients for B over R if and only if they form a basis for R^B. Detailed background on mutation-linear algebra is provided in Section 2.2.2; the connections with cluster algebras are discussed in Section 2.2.3.

We further the study of mutation-linear algebra by exploring the questions of whether bases for R^B and the mutation fan $\mathcal{F}_B$ fold when the exchange matrix B folds, and by generalizing to orbifolds the constructions for marked surfaces of bases for R^B and the mutation fan $\mathcal{F}_B$. In particular, given that B folds to $\pi_\sigma(B) = C$ under a certain symmetry σ, Theorems 2.4.9 and 2.4.15 show that in general, symmetric bases for R^B fold to spanning sets for R^C, and the mutation fan $\mathcal{F}_B$ folds to a refinement of the mutation fan $\mathcal{F}_C$. We then prove in Theorem 2.4.23 that in fact the folding of $\mathcal{F}_B$ coincides with $\mathcal{F}_C$ if and only if a symmetric positive basis for R^B folds to a positive basis for R^C. Theorem 2.4.24 establishes both of these results when B is of finite type, and Section 3.6 establishes both results when B and C are the signed-adjacency matrices of orbifolds with a suitable folding relationship. (The finite type results can be deduced in the setting of Cambrian fans, as in [25], using the folding of root systems, or in the cluster algebra setting using the folding of **g**-vectors: for example, see [1]. However, the proofs provided here are geometric in nature and were obtained independently.) Connections between mutation-linear algebra and the orbifolds model are drawn in Theorems 3.4.3 and 3.4.17. These results respectively show that the shear coordinates of allowable curves may be used to construct the rational part of the mutation fan if the orbifold has the Curve Separation Property, and bases for R^B if the orbifold has the Null Tangle Property (or is a null orbifold).

Finally, as mentioned above in Section 1.1.2, we prove examples of two mutation-linear dominance phenomena. Namely, Theorems 4.1.1 and 4.1.2 state that if B is the signed-adjacency matrix of an orbifold and B' is the signed-adjacency matrix of another orbifold obtained by resection of the first, then $\mathcal{F}_B$ refines $\mathcal{F}_{B'}$. If furthermore both orbifolds have the Null Tangle Property (or are null orbifolds), then the identity map from $\mathbb{Q}^B$ to $\mathbb{Q}^{B'}$ is mutation linear.

Chapter 2

Folding and mutation-linear algebra

2.1 Introduction

This chapter studies the folding of exchange matrices in the context of mutation-linear algebra.

An exchange matrix B is an $n \times n$ skew-symmetrizable integer matrix: the fundamental combinatorial datum specifying a cluster algebra. Recall from Section 1.2 that the term 'mutation' in 'mutation-linear algebra' refers to matrix mutation in the sense of cluster algebras (see Definition 2.2.4). Whereas linear algebra is the study of linear relations among vectors in a module over a ring, *mutation*-linear algebra, a subject which originated in [21] and was named in [24], is the study of B*-coherent* linear relations: essentially, linear relations which are preserved under mutation (see Definition 2.2.19). As in linear algebra, there is a corresponding notion of bases: sets of vectors which are spanning and independent in a mutation-linear sense (see Definition 2.2.22). Such a basis is positive if it is spanning with non-negative coefficients only. Furthermore, finding such a set, which we refer to as a basis for R^B, is equivalent to constructing the universal geometric cluster algebra over B [21, Theorem 4.4]. Closely related to these constructions is a complete fan $\mathcal{F}_B$ in $\mathbb{R}^n$ called the mutation fan for B, which encodes the geometry of mutation (see Definition 2.2.29).

An automorphism of the exchange matrix $B = [b_{ij}]$ is a permutation $\sigma \in \mathfrak{S}_n$ which satisfies $b_{\sigma(i),\sigma(j)} = b_{ij}$ for each $i, j \in [n]$. If the σ-orbits of entries of B satisfy two 'admissability' conditions (see Definition 2.3.2) then σ is admissable on B; if in addition these conditions are preserved under orbit-mutation, then σ is stable (see Definition 2.3.4). (Orbit-mutation is mutation in all indices in a given σ-orbit: see Definition 2.3.3.) There is no guarantee that a given admissable automorphism on B is stable, and stability can be difficult to check as there may be infinitely many matrices obtained by mutation of B. However, if σ *is* a stable automorphism of B, then, as described in [4], one can fold B under σ to obtain a new exchange matrix $\pi_\sigma(B)$ with rows and columns indexed by the σ-orbits of $[n]$. In particular, the entries of $\pi_\sigma(B)$ are given by sums, over σ-orbits, of entries of B. Most importantly, matrix mutation of the folded matrix $\pi_\sigma(B)$ corresponds to orbit-mutation of B (see (2.13)). We address the questions of whether the mutation-linear notions of bases for R^B, the mutation fan $\mathcal{F}_B$, and universal geometric coefficients for B "fold" under σ to bases for $R^{\pi_\sigma(B)}$, the mutation fan $\mathcal{F}_{\pi_\sigma(B)}$, and universal geometric coefficients for $\pi_\sigma(B)$, respectively.

The main results of this chapter are stated below. While there are limitations to our folding arguments in the general case, the equivalence between positive basis folding and mutation fan folding provided by Theorem 2.1.3 is a powerful tool, as for all exchange matrices for which bases have been constructed save one, these bases are positive bases. We show in this chapter that the mutation fan (and therefore positive bases) fold when B is of finite type; in Chapter 3, we show that this is also the case when B and $\pi_\sigma(B)$ are modeled by surfaces and orbifolds in the sense of [5,8]. (As noted in Section 1.2, our proofs of the finite type results are geometric in nature, but these results can alternatively be deduced using root system folding in the setting of Cambrian fans as in [25], or using **g**-vector folding as in [1].)

Given an $n \times n$ exchange matrix B with stable automorphism $\sigma \in \mathfrak{S}_n$ and a choice of ring $R \in \{\mathbb{Z}, \mathbb{Q}, \mathbb{R}\}$, let $\pi_\sigma(B)$ denote the exchange matrix obtained by folding B under σ. Viewing σ as a linear map on $\mathbb{R}^n$ acting by coordinate permutation, let V^σ denote the fixed-point subspace under this action, and let $\phi : V^\sigma \xrightarrow{\cong} \mathbb{R}^{\dim(V^\sigma)}$ denote a naturally defined vector space isomorphism (see (2.23)). For any ***σ-symmetric*** collection of vectors $\mathcal{B} \subseteq R^n$ (that is, a collection $\mathcal{B}$ satisfying $\sigma(\mathcal{B}) = \mathcal{B}$), let $\bar{\mathcal{B}}$ denote the set of σ-orbits $\bar{\mathbf{b}}$ of $\mathcal{B}$, and reuse $\pi_\sigma : \bar{\mathcal{B}} \to V^\sigma$ to denote the map taking sums over σ-orbits: $\pi_\sigma(\bar{\mathbf{b}}) = \sum_{\mathbf{b}_i \in \bar{\mathbf{b}}} \mathbf{b}_i$. Finally, recall that $\mathcal{F}_B^R$ denotes the R-part of the mutation fan $\mathcal{F}_B$ when a positive basis for R^B exists.

Theorem 2.1.1. *If $\mathcal{B} = (\mathbf{b}_i : i \in I) \subseteq R^n$ is a σ-symmetric (positive) basis for R^B, then $\phi(\pi_\sigma(\bar{\mathcal{B}}))$ is a (positive) spanning set for $R^{\pi_\sigma(B)}$.*

Theorem 2.1.2. *The folded mutation fan $\phi(\mathcal{F}_B \cap V^\sigma)$ refines the mutation fan $\mathcal{F}_{\pi_\sigma(B)}$.*

Theorem 2.1.3. *Suppose that $\mathcal{B} = (\mathbf{b}_i : i \in I)$ is a σ-symmetric positive basis for R^B and let $\bar{\mathcal{B}}_\mathcal{F} \subseteq \bar{\mathcal{B}}$ denote the subset of σ-orbits of $\mathcal{B}$ which span cones in $\mathcal{F}_B^R$. Then $\phi(\pi_\sigma(\bar{\mathcal{B}}_\mathcal{F}))$ is a positive basis for $R^{\pi_\sigma(B)}$ if and only if $\phi(\mathcal{F}_B^R \cap V^\sigma) = \mathcal{F}_{\pi_\sigma(B)}^R$.*

Theorem 2.1.4. *Suppose B is of finite type. Then $\phi(\mathcal{F}_B \cap V^\sigma) = \mathcal{F}_{\pi_\sigma(B)}$. If a positive basis exists for R^B, then so too does a σ-symmetric positive basis $\mathcal{B}$, and the following collection is a positive basis for $R^{\pi_\sigma(B)}$ that constitutes universal geometric coefficients for $\pi_\sigma(B)$:*

$$\left\{\phi(\pi_\sigma(\bar{\mathbf{b}})) : \mathbf{b} \in \bar{\mathcal{B}} \text{ spans a cone in } \mathcal{F}_B\right\}$$

The chapter is organized as follows. We present technical background on cluster algebras and mutation-linear algebra in Section 2.2 and then define exchange matrix folding in Section 2.3. We state consequences of applying the latter to the former in Section 2.4.1, whereby we deduce Theorem 2.1.1 (see Theorem 2.4.9 and Corollary 2.4.10). We then consider the same consequences from the geometric standpoint of the mutation fan in Section 2.4.2, yielding Theorem 2.1.2 and Theorem 2.1.3 (see Theorems 2.4.15 and 2.4.23, respectively). We conclude the chapter by using polyhedral geometry to prove the finite-type results of Theorem 2.1.4 (see Theorem 2.4.24, Corollary 2.4.25, and Corollary 2.4.26).

2.2 Background

In this section we define the basic notions of cluster algebras of geometric type and mutation-linear algebra. These definitions will be referenced throughout this chapter and later chapters.

While cluster algebras of geometric type were first introduced in [13], Section 2.2.1 presents the broadened definition given in [21]. This is done with an eye towards the subsequent discussion of mutation-linear algebra in Section 2.2.2, which is also drawn from [21]. However, we intentionally mode our exposition after [21, Section 2], which in turn was intentionally modeled after [13, Section 2], so that readers may easily compare the definitions here directly with the original source material. (Readers may also wish to refer to Remark 2.2.7, which discusses the two adjustments made in detail.)

2.2.1 Cluster algebras of geometric type

Fix a ring R (either the integers $\mathbb{Z}$, the rationals $\mathbb{Q}$, or the reals $\mathbb{R}$), henceforth referred to as the ***underlying ring***.

Remark 2.2.1. For maximal generality, [21, Definition 2.1] permits the underlying ring R to be either $\mathbb{Z}$ or any subfield of $\mathbb{R}$ which contains $\mathbb{Q}$ as a subfield. However, we restrict our attention here to the three primary cases of interest.

Definition 2.2.2 (Tropical semifield over R)**.** Let $(u_i : i \in I)$ be a collection of formal symbols indexed by the set I. We call I the ***indexing set*** and permit it to be infinite, and call the symbols u_i ***tropical variables***. Let $\mathbb{P} = \mathrm{Trop}_R(u_i : i \in I)$ denote the abelian group with elements given by formal products of the form $\prod_{i\in I} u_i^{a_i}$ for $a_i \in R$ and multiplication defined as

$$\prod_{i\in I} u_i^{a_i} \cdot \prod_{i\in I} u_i^{b_i} = \prod_{i\in I} u_i^{a_i+b_i}.$$

Note that there is group isomorphism $\mathbb{P} = \mathrm{Trop}_R(u_i : i \in I) \cong \prod_{i\in I} R$. Endow the underlying ring R with the discrete topology and $\mathbb{P}$ with the product topology as a product of copies of R. (The imposition of these topologies is necessary to define the coefficient specialization which is featured in Definition 2.2.9 of universal geometric coefficients: see [21, Section 3] for details.) Observe that $\mathbb{P} = \mathrm{Trop}_R(u_i : i \in I)$ is uniquely defined by the cardinality of I.

Let the triple $(\mathrm{Trop}_R(u_i : i \in I), \oplus, \cdot)$ denote the ***tropical semifield*** built by defining an ***auxiliary addition*** $\oplus$ in $\mathrm{Trop}_R(u_i : i \in I)$ as

$$\prod_{i\in I} u_i^{a_i} \oplus \prod_{i\in I} u_i^{b_i} = \prod_{i\in I} u_i^{\min(a_i,b_i)}.$$

This auxiliary addition is commutative, associative, and distributive with respect to multiplication.

Let $\mathbb{ZP}$ denote the group ring, over $\mathbb{Z}$, of $\mathbb{P} = \mathrm{Trop}_R(u_i : i \in I)$: that is, the ring of Laurent polynomials in the variables u_i with coefficients in $\mathbb{Z}$. This ring will play the role of the coefficient ring for a cluster algebra $\mathcal{A}$ of geometric type (see Definition 2.2.6). Likewise, let $\mathbb{QP}$ denote the group ring of $\mathbb{P}$ over $\mathbb{Q}$; $\mathbb{QP}$ will be the coefficient ring for the ***ambient field*** $\mathbb{K}$ of $\mathcal{A}$. In particular, we take $\mathbb{K}$ to be isomorphic to the field of rational functions in n independent variables with coefficients in $\mathbb{QP}$, and refer to n as the ***rank*** of $\mathcal{A}$. Note that these definitions consider $\mathbb{P}$ as a multiplicative group: the auxiliary addition $\oplus$ is ignored. Given its role in defining coefficient rings, we call $\mathbb{P} = \mathrm{Trop}_R(u_i : i \in I)$ the ***coefficient semifield***.

Definition 2.2.3 (Labeled geometric seed of rank n)**.** A ***(labeled) geometric seed*** of rank n in $\mathbb{K}$ is a pair $(\mathbf{x}, \tilde{B})$ defined as follows:

- $\mathbf{x} = (x_1, \ldots, x_n)$ is an n-tuple of elements of $\mathbb{K}$ which freely generate $\mathbb{K}$: that is, the entries x_i are algebraically independent and $\mathbb{K} = \mathbb{QP}(x_1, \ldots, x_n)$. We refer to $\mathbf{x}$ as the ***cluster*** of the seed, and its entries x_i as ***cluster variables***.

- $\tilde{B}$ is a function from $([n] \cup I) \times [n]$ to R, called an ***extended exchange matrix***, whose top rows indexed by $[n]$ form a square ***skew-symmetrizable*** integer matrix B. (By skew-symmetrizable, we mean that there exists a ***skew-symmetrizing*** diagonal matrix D with positive integer entries d_i such that DB is ***skew-symmetric***: that is, $d_i b_{ij} = -d_j b_{ji}$ for all $i, j \in [n]$.) We refer to B as the ***exchange matrix*** or ***principal part*** of $\tilde{B}$, and to the rows of $\tilde{B}$ indexed by I as the ***coefficient rows***, or ***complementary part*** of $\tilde{B}$.

  The complementary part of $\tilde{B}$ determines the ***coefficients*** $\mathbf{y} = (y_1, \ldots, y_n)$ of the seed as follows. For each $i \in I$, denote the coefficient row of $\tilde{B}$ indexed by i as $\mathbf{b}_i = (b_{i1}, \ldots, b_{in})$. Then for each $j \in [n]$, define $y_j \in \mathbb{P}$ by $y_j = \prod_{i\in I} u_i^{b_{ij}}$.

Definition 2.2.4 (Mutation of geometric seeds)**.** Each index $k \in [n]$ defines an involution μ_k, called ***seed mutation in direction*** k, on the set of labeled geometric seeds of rank n by setting $\mu_k(\mathbf{x}, \tilde{B}) = (\mathbf{x}', \tilde{B}')$ defined as follows. (For $a \in \mathbb{R}$, the notation $[a]_+$ represents $\max(a, 0)$.)

- The new cluster $\mathbf{x}' = (x'_1, \dots, x'_n)$ is defined by setting $x'_j = x_j$ for $j \neq k$ and defining x'_k by the following ***exchange relation***:

$$\begin{aligned} x_k x'_k &= \frac{y_k \prod_{i=1}^n x_i^{[b_{ik}]_+} + \prod_{i=1}^n x_i^{[-b_{ik}]_+}}{y_k \oplus 1} \\ &= \prod_{i \in I} u_i^{[b_{ik}]_+} \prod_{i=1}^n x_i^{[b_{ik}]_+} + \prod_{i \in I} u_i^{[-b_{ik}]_+} \prod_{i=1}^n x_i^{[-b_{ik}]_+}. \end{aligned} \tag{2.1}$$

- The new extended exchange matrix $\tilde{B}' = (b'_{ij})_{i \in ([n] \cup I),\ j \in [n]}$ has entries

$$\begin{aligned} b'_{ij} &= \begin{cases} -b_{ij} & \text{if } i = k \text{ or } j = k; \\ b_{ij} + \operatorname{sgn}(b_{kj})[b_{ik}b_{kj}]_+ & \text{otherwise} \end{cases} \\ &= \begin{cases} -b_{ij} & \text{if } i = k \text{ or } j = k; \\ b_{ij} + b_{ik}b_{kj} & \text{if } j \neq k, b_{ik} \geq 0, \text{ and } b_{kj} \geq 0; \\ b_{ij} - b_{ik}b_{kj} & \text{if } j \neq k, b_{ik} \leq 0, \text{ and } b_{kj} \leq 0; \\ b_{ij} & \text{otherwise.} \end{cases} \end{aligned} \tag{2.2}$$

Observe that the new extended exchange matrix $\tilde{B}'$ does not depend on the cluster $\mathbf{x}$, and its principal part B' is skew-symmetrizable with the same skew-symmetrizing matrix as the principal part B of $\tilde{B}$. We therefore also use μ_k to denote (extended) ***matrix mutation in direction*** k, and write $\mu_k(\tilde{B}) = \tilde{B}'$ and $\mu_k(B) = B'$ (see Example 2.2.17).

For any finite sequence $\mathbf{k} = k_q, \dots, k_1$ of indices in $[n]$, the notation $\mu_{\mathbf{k}}$ stands for the following composition of mutations:

$$\mu_{\mathbf{k}} = \mu_{k_q} \circ \mu_{k_{q-1}} \circ \cdots \circ \mu_{k_1}. \tag{2.3}$$

Likewise the notation $\mu_{\mathbf{k^{-1}}}$ stands for the composition $\mu_{k_1} \circ \mu_{k_2} \circ \cdots \circ \mu_{k_q}$, so that the compositions $\mu_{\mathbf{k^{-1}}} \circ \mu_{\mathbf{k}} = \mu_{\mathbf{k}} \circ \mu_{\mathbf{k^{-1}}}$ denote the identity map.

Definition 2.2.5 (Cluster and Y-patterns of geometric type)**.** Fix a vertex t_0 in $\mathbb{T}_n$, the n-regular tree whose edges are labeled by the integers in the set $[n]$ such that each vertex is incident to precisely one edge of each label. A ***cluster pattern (of geometric type)*** is a map $t \mapsto (\mathbf{x}_t, \tilde{B}_t)$ on $\mathbb{T}_n$ uniquely specified by an assignment of a given ***initial seed*** $(\mathbf{x}_{t_0}, \tilde{B}_{t_0})$ to t_0 and the requirement that whenever two vertices t and t' are connected by an edge labeled k, the map satisfies $(\mathbf{x}_{t'}, \tilde{B}_{t'}) = \mu_k(\mathbf{x}_t, \tilde{B}_t)$. Recalling that the definition of extended matrix mutation is independent of an associated cluster, we may also define a ***Y-pattern (of geometric type)***: a map $t \mapsto \tilde{B}_t$ uniquely specified by an assignment of a given extended exchange matrix $\tilde{B}$ to

t_0 and the same requirement that edges with label k correspond to mutation in direction k.

The cluster variables, matrix entries, and coefficients of the seed $(\mathbf{x}_t, \tilde{B}_t)$ associated to vertex $t \in \mathbb{T}_n$ are denoted by $\mathbf{x}_t = (x_{1;t}, \ldots, x_{n;t})$, $\tilde{B}_t = (b^t_{ij})$, and $\mathbf{y}_t = (y_{1;t}, \ldots, y_{n;t})$, respectively.

Definition 2.2.6 (Cluster algebra of geometric type)**.** Given a cluster pattern $t \mapsto (\mathbf{x}_t, \tilde{B}_t)$, we define the associated ***cluster algebra (of geometric type)*** $\mathcal{A}$ to be the $\mathbb{ZP}$-subalgebra of $\mathbb{K}$ which is generated by all cluster variables which occur in the pattern: $\mathcal{A} = \mathbb{ZP}[\mathcal{X}]$ where

$$\mathcal{X} = \{x_{i;t} : t \in \mathbb{T}_n, i \in [n]\}.$$

Since the cluster pattern is uniquely determined by the initial seed $(\mathbf{x}_{t_0}, \tilde{B}_{t_0}) = (\mathbf{x}, \tilde{B})$, we write $\mathcal{A} = \mathcal{A}_R(\mathbf{x}, \tilde{B})$ (the subscript R is an indication of our underlying ring). Indeed, up to isomorphism, the cluster algebra is uniquely determined by just the extended exchange matrix $\tilde{B}$, so we also write simply $\mathcal{A} = \mathcal{A}_R(\tilde{B})$. If the subscript R is omitted, it is understood that the underlying ring is $\mathbb{Z}$.

Remark 2.2.7. Definitions 2.2.3, 2.2.4, 2.2.5, and 2.2.6 appear as [21, Definitions 2.3, 2.4, 2.7, and 2.8] respectively, and as noted in [21, Remark 2.9], are respectively comparable to [13, Definitions 2.3, 2.4, 2.9, and 2.11] restricted to the the special case of cluster algebras of geometric type described in [13, Definition 2.12]. The broadening of the definition of cluster algebras of geometric type remarked upon at the start of the section comes into play through Definition 2.2.2, which is precisely [21, Definition 2.2]. As compared to [13, Definition 2.2], this definition allows for an extended exchange matrix $\tilde{B}$ to have infinitely many coefficient rows, and permits these rows to have non-integer entries if the underlying ring R is choen to be $\mathbb{Q}$ or $\mathbb{R}$. These adjustments, along with the imposition of the product topology on the tropical semifield $\mathbb{P}$, lead to the following definitions of coefficent specialization and universal geometric coefficients, which appears as [21, Definitions 3.1 and 3.2] and are comparable to [13, Definitions 12.1 and 12.3].

Definition 2.2.8 (Coefficient specialization)**.** Suppose $(\mathbf{x}, \tilde{B})$ and $(\mathbf{x}', \tilde{B}')$ are seeds of rank n with the same exchange matrix $B = B'$ whose tropical semifields $\mathbb{P}$ and $\mathbb{P}'$ are defined over the same underlying ring R. A ***coefficient specialization*** from $\mathcal{A}_R(\mathbf{x}, \tilde{B})$ to $\mathcal{A}_R(\mathbf{x}', \tilde{B}')$ is a ring homomorphism $\varphi : \mathcal{A}_R(\mathbf{x}, \tilde{B}) \to \mathcal{A}_R(\mathbf{x}', \tilde{B}')$ which satisfies the following conditions:

(i) $\varphi(x_i) = x'_i$ for all $i \in [n]$;

(ii) the restriction of φ to $\mathbb{P}$ is continuous (with respect to the product topology) and linear (with respect to R);

(iii) $\varphi(y_{j;t}) = y'_{j;t}$ and $\varphi(y_{j;t} \oplus 1) = y'_{j;t} \oplus 1$ for all $j \in [n]$ and $t \in \mathbb{T}_n$.

Definition 2.2.9 (Cluster algebra with universal geometric coefficients)**.** Let $\tilde{B}$ be an extended exchange matrix with principal part B. We say that $\tilde{B}$ is ***universal over*** R, that its coefficients rows are ***universal geometric coefficients*** for B over R, and that the associated cluster algebra of geometric type $\mathcal{A} = \mathcal{A}_R(\tilde{B})$ is the ***universal geometric cluster algebra for*** B ***over*** R, denoted $\mathcal{A}_{\text{un}}(B)$, if $\mathcal{A}$ satisfies the following condition: For every other cluster algebra $\mathcal{A}' = \mathcal{A}_R(\tilde{B}')$ of geometric type with underlying ring R and the same initial exchange matrix B, there exists a unique coefficient specialization from $\mathcal{A}$ to $\mathcal{A}'$.

On the opposite end of the spectrum from the universal geometric cluster algebra for B is the coefficient-free cluster algebra.

Definition 2.2.10 (Coefficient-free cluster algebra)**.** We define the ***coefficient-free cluster algebra*** $\mathcal{A}_0(B)$ for exchange matrix B by taking $\tilde{B} = B$ as our initial extended exchange matrix in Definition 2.2.6. In particular, this implies we are choosing as our coefficient semifield (see Definition 2.2.2) the trivial semifield consisting of a single element 1. Thus the ambient field $\mathbb{K}$ is isomorphic to $\mathbb{Q}[x_1^{\pm 1}, \ldots, x_n^{\pm 1}]$, the field of rational functions in n independent variables with coefficients in $\mathbb{Q}$, and $\mathcal{A}_0(B)$ is a $\mathbb{Z}$-subalgebra of $\mathbb{K}$.

Somewhere in between the universal geometric cluster algebra $\mathcal{A}_{\text{un}}(B)$ and the coefficient-free cluster algebra $\mathcal{A}_0(B)$ is $\mathcal{A}_\bullet(B)$, the cluster algebra with principal coefficients. The definition below is drawn from [13, Section 3 and Section 6].

Definition 2.2.11 (Cluster algebra with principal coefficients and **g**-vectors)**.** Given an $n \times n$ exchange matrix B, let $\tilde{B}$ denote the $2n \times n$ extended exchange matrix whose top n rows (the principal part) are the exchange matrix B and whose bottom n rows (the complementary part) are the $n \times n$ identity matrix. A cluster pattern (or Y-pattern) with $\tilde{B}$ in its initial seed t_0 is said to have ***principal coefficients*** at the initial seed, and we denote by $\mathcal{A}_\bullet(B)$ the cluster algebra with principal coefficients whose initial exchange matrix is B. The set of cluster variables in $\mathcal{A}_\bullet(B)$ is denoted by $\text{Var}_\bullet(B)$.

Let $t \mapsto (\mathbf{x}_t^{B;t_0}, \tilde{B}_t)$ denote the cluster pattern associated to $\mathcal{A}_\bullet(B)$, with initial cluster $\mathbf{x}_0^{B;t_0} = (x_1, \ldots, x_n)$, initial extended exchange matrix $\tilde{B}_0 = \tilde{B}$ as described above, and initial coefficients $\mathbf{y}_0^{B;t_0} = (y_1, \ldots, y_n)$, where by construction, $y_j = u_j$ for each $j \in [n]$. Then $\mathcal{A}_\bullet(B)$ is a $\mathbb{Z}^n$-graded subalgebra of $\mathbb{Z}[x_1^{\pm 1}, \ldots, x_n^{\pm 1}; y_1, \ldots, y_n]$ under the $\mathbb{Z}^n$-grading given by

$$\deg(x_i) = \mathbf{e}_i, \qquad \deg(y_j) = -\mathbf{b}_j, \tag{2.4}$$

where $\mathbf{e}_1, \ldots, \mathbf{e}_n$ are the standard basis vectors for $\mathbb{Z}^n$, and $\mathbf{b}_1, \ldots, \mathbf{b}_n$ are the columns of B. Each cluster variable $x_{i;t}^{B;t_0} \in \text{Var}_\bullet(B)$ is homogeneous with respect to this grading, with (multi-)degree given by the **g**-***vector*** $\mathbf{g}_{i;t}^{B;t_0} = \deg(x_{i;t}^{B;t_0})$.

Before moving on to the the basics of mutation-linear algebra, we define several special classes of cluster algebras and exchange matrices. As discussed in Section 2.1, we obtain the strongest form of our mutation-linear folding results in this chapter for exchange matrices/cluster algebras of finite type. Further, the exchange matrices/cluster algebras modeled by surfaces and orbifolds which are the topic of Chapter 3 and also yield strong folding results are mutation-finite, and for some of the results in Chapter 4 we assume acyclicity. We then conclude our presentation of geometric cluster algebra background by discussing a relationship between exchange matrices and the (generalized) Cartan matrices associated to semisimple Lie algebras and root systems (see [16, 17]). Recall from Section 1.1 that Lie theory was the motivation for the definition and study of cluster algebras initiated in [10]: the relationship also provides the definition for another special class of cluster algebras, those of affine type.

Definition 2.2.12 (Finite type)**.** A cluster algebra $\mathcal{A}_R(\tilde{B})$ is of ***finite type*** if and only if its associated cluster pattern has finitely many non-equivalent seeds. Otherwise, it is of ***infinite type***. Being of finite type is a property that depends only on the principal part B of the initial extended exchange matrix $\tilde{B}$ and is independent of the choice of coefficient rows by [11, Theorem 1.4]. Thus we say an exchange matrix B is of finite type if and only if any cluster algebra $\mathcal{A}_R(\tilde{B})$ with initial exchange matrix B is of finite type.

Definition 2.2.13 (Mutation equivalence and mutation finite type)**.** We say that two $n \times n$ exchange matrices are ***mutation equivalent*** if there exists a sequence $\mathbf{k}$ of indices in $[n]$ such that applying $\mu_{\mathbf{k}}$ to one matrix yields the other. The ***mutation class*** of an exchange matrix is the set of all matrices which are mutation equivalent to it. If this set is finite, we say the matrix is ***mutation-finite*** or of ***finite mutation type***. A cluster algebra $\mathcal{A}_R(\tilde{B})$ is mutation-finite if and only if the principal part B of $\tilde{B}$ is a mutation-finite exchange matrix. Exchange matrices and cluster algebras of finite type are by necessity mutation-finite.

Definition 2.2.14 (Acyclicity)**.** An $n \times n$ exchange matrix B is ***acyclic*** if, possibly after applying a permutation $\sigma \in \mathfrak{S}_n$ to reindex the rows and columns, it satisfies the following condition: if $b_{ij} > 0$ then $i < j$.

Definition 2.2.15 (Cartan companion of B)**.** The ***Cartan companion*** of an $n \times n$ exchange matrix B is the $n \times n$ matrix $A = A(B)$ with entries defined as follows.

$$a_{ij} = \begin{cases} 2 & \text{if } i = j; \\ -|b_{ij}| & \text{otherwise.} \end{cases}$$

Observe that DA is symmetric, where D is the skew-symmetrizing matrix for B (see Definition 2.2.3). Thus A is ***symmetrizable***, and is a ***Cartan matrix*** in the usual sense of [16, 17].

A Cartan matrix is of ***finite type*** if it is positive-definite (all eigenvalues are positive), and of ***affine type*** if it is positive-semidefinite (all eigenvalues are non-negative) and every proper principal submatrix is of finite type (see [17, Section 4.3]). The classification of cluster algebras of finite type was achieved in [11, Theorem 1.4] by showing that an exchange matrix B (and any geometric cluster algebra $\mathcal{A}_R(\tilde{B})$ with initial exchange matrix B) is of finite type if and only if it is mutation equivalent to a matrix B' with finite type Cartan companion $A(B')$. Thus in particular, every exchange matrix of finite type is mutation equivalent to an acyclic matrix. We say that an exchange matrix B (and any geometric cluster algebra $\mathcal{A}_R(\tilde{B})$ with initial exchange matrix B) is of ***affine type*** if B is mutation equivalent to an acyclic exchange matrix B' with affine type Cartan companion $A(B')$.

2.2.2 Mutation-linear algebra

Fix an $n \times n$ exchange matrix B (see Definition 2.2.3) and an underlying ring $R \in \{\mathbb{Z}, \mathbb{Q}, \mathbb{R}\}$.

Definition 2.2.16 (Mutation maps $\eta_{\mathbf{k}}^B$)**.** Each index $k \in [n]$ defines a continuous, piecewise linear, invertible map η_k^B on $\mathbb{R}^n$ as follows. Let $\mathbf{a} = (a_1, \ldots, a_n) \in \mathbb{R}^n$ and define $\tilde{B}$ to be the $(n+1) \times n$ extended exchange matrix with principal part B and single coefficient row $\mathbf{a}$. Then $\mathbf{a}' = \eta_k^B(\mathbf{a})$ is the coefficient row of $\tilde{B}' = \mu_k(\tilde{B})$ with entries $a'_j = b'_{(n+1)j}$ given by (2.2), or equivalently, by

$$a'_j = \begin{cases} -a_k & \text{if } j = k; \\ a_j + \operatorname{sgn}(a_k)[a_k b_{kj}]_+ & \text{otherwise} \end{cases} \quad \text{for each } j \in [n].$$

Likewise, for each finite sequence $\mathbf{k} = k_1, \ldots, k_q$ of indices in $[n]$, define $\mathbf{a}' = \eta_{\mathbf{k}}^B(\mathbf{a})$ to be the coefficient row of $\mu_{\mathbf{k}}(\tilde{B})$: that is, setting $B_1 = B$ and $B_{i+1} = \mu_{k_i}(B_i)$ for each $i \in [q]$,

$$\eta_{\mathbf{k}}^B = \eta_{k_q}^{B_q} \circ \eta_{k_{q-1}}^{B_{q-1}} \circ \cdots \circ \eta_{k_1}^{B_1}. \tag{2.5}$$

(See (2.3) and Example 2.2.17.) We refer collectively to all such homeomorphisms $\eta_{\mathbf{k}}^B$ from $\mathbb{R}^n$ to itself as the ***mutation maps*** associated to B.

Example 2.2.17. We demonstrate the results of performing matrix mutation on a given 4×3 extended exchange matrix $\tilde{B}$ with principal part B and coefficient row $\mathbf{a}$ as an illustration of Definitions 2.2.4 and 2.2.16.

$$\tilde{B} = \begin{bmatrix} 0 & 0 & -1 \\ 0 & 0 & -1 \\ 1 & 1 & 0 \\ -2 & 3 & 4 \end{bmatrix} \xrightarrow{\mu_1} \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ -1 & 1 & 0 \\ 2 & 3 & 2 \end{bmatrix} \xrightarrow{\mu_3} \begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \\ 2 & 5 & -2 \end{bmatrix} = \mu_{31}(\tilde{B})$$

$$\mathbf{a} = [-2,\ 3,\ 4] \xrightarrow{\eta_1^B} [2,\ 3,\ 2] \xrightarrow{\eta_3^{\mu_1(B)}} [2,\ 5,\ -2] = \eta_{31}^B(\mathbf{a})$$

Remark 2.2.18. The use of the real numbers $\mathbb{R}$ rather than the underlying ring R in Definition 2.2.16 (and below in Definition 2.2.19) is intentional, as later in this section we will use these mutation maps to build a complete fan in $\mathbb{R}^n$ which encodes their geometry, called the mutation fan (see Definition 2.2.29). Observe that since our underlying ring R is a subring of $\mathbb{R}$, each mutation map $\eta_{\mathbf{k}}^B$ restricts to a piecewise-linear homeomorphism from R^n to itself.

Definition 2.2.19 (B-coherent linear relation)**.** Given collections of vectors $(\mathbf{a}_i : i \in S)$ in $\mathbb{R}^n$ and scalars $(r_i : i \in S)$ in the ring R, both indexed by the same finite set S, the formal sum $\sum_{i\in S} r_i\mathbf{a}_i$ is a ***B-coherent linear relation with coefficients in*** R (briefly, ***B-coherent***) if the following two equalities hold for every finite sequence $\mathbf{k}$ of indices in $[n]$, including the empty sequence, where the notation $\mathbf{min}(\mathbf{a}, \mathbf{0})$ represents the vector $(\min(a_1, 0), \ldots, \min(a_n, 0))$:

$$\sum_{i\in S} r_i \eta_{\mathbf{k}}^B(\mathbf{a_i}) = \mathbf{0}. \tag{2.6}$$

$$\sum_{i\in S} r_i \mathbf{min}(\eta_{\mathbf{k}}^B(\mathbf{a_i}), \mathbf{0}) = \mathbf{0}. \tag{2.7}$$

The collection of all B-coherent linear relations with coefficients in R is denoted by $\mathcal{V}_{\eta^B}$. A B-coherent linear relation $\sum_{i\in S} r_i\mathbf{a}_i$ is ***trivial*** if it is an ***empty relation*** ($S = \emptyset$) or if the scalars $r_i = 0$ for all $i \in S$.

Remark 2.2.20. Observe that a B-coherent linear relation is in particular a linear relation in the usual sense due to the requirement that (2.6) hold when $\mathbf{k}$ is the empty sequence. Likewise, a trivial B-coherent linear relation is in particular a trivial linear relation. We also note that (2.7) essentially requires that the relation respect sign: for example, the relation $\mathbf{v} + (-\mathbf{v})$ is not B-coherent for $\mathbf{v} \neq \mathbf{0}$ since it fails to satisfy (2.7) when $\mathbf{k}$ is the empty sequence.

Proposition 2.2.21. *[21, Proposition 4.12] Let $\sum_{i\in S} r_i\mathbf{a}_i$ be a B-coherent linear relation. Suppose, for some $i \in S$, for some $j \in [n]$, and for some sequence $\mathbf{k}$ of indices in $[n]$, that the j^{th} entry of $\eta_{\mathbf{k}}^B(\mathbf{a}_i)$ is strictly positive (resp. strictly negative) and that the j^{th} entry of every vector $\eta_{\mathbf{k}}^B(\mathbf{a}_{i'})$ with $i' \in S \setminus \{i\}$ is nonpositive (resp. nonnegative). Then $r_i = 0$.*

Partial linear algebra

Let $\operatorname{Rel}(R^n, R)$ denote the set of ***linear relations*** on the module R^n with coefficients in R. That is, the collection of formal expressions $\sum_{i\in S} r_i\mathbf{a}_i$ that evaluate to $\mathbf{0}$, where $r_i \in R$ and $\mathbf{a}_i \in R^n$ for each i in some finite indexing set S. Linear algebra can be viewed as the study of such relations: for example, a set $A \subseteq R^n$ is independent if every linear relation among vectors in A is trivial, and spanning if for every $\mathbf{a} \in M$, there exists a linear relation $\mathbf{a}+\sum_{i\in S\subseteq A}(-r_i)x_i$. In this light, ***mutation-linear algebra*** is the study of the subset $\mathcal{V}_{\eta^B} \subseteq \operatorname{Rel}(R^n, R)$ consisting of the B-coherent linear relations of Definition 2.2.19. More precisely, it is the study of the ***mutation-linear structure*** $R^B = (R^n, R, \mathcal{V}_{\eta^B})$ given by B-coherent linear relations among vectors in R^n with coefficients in R. This structure, introduced in [21], is formalized in [24, Section 2] as (the only known interesting example of) a more general notion called partial linear structure. We describe the general setting here, as we will later add an additional example, that of orbit-mutation-linear structure (see Remark 2.4.3).

Definition 2.2.22 (Partial linear structure)**.** Given a module M over an arbitrary ring R with identity and a subset $\mathcal{V} \subseteq \operatorname{Rel}(M, R)$ of the linear relations on M with coefficients in R, the triple $(M, R, \mathcal{V})$ is a ***partial linear structure*** if and only if it satisfies the following conditions:

(i) **(Empty relation.)** The set $\mathcal{V}$ contains the empty relation.

(ii) (**Irrelevance of zeros**.) Let 0 denote the zero element in R. Then for any $x \in M$, the relation $0x + \sum_{i\in S} r_ix_i$ is in $\mathcal{V}$ if and only if $\sum_{i\in S} r_ix_i$ is in $\mathcal{V}$.

(iii) (**Combining like terms.**) For any element $c = a+b$ of R, the relation $ax+bx+\sum_{i\in S} r_ix_i$ is in $\mathcal{V}$ if and only if $cx + \sum_{i\in S} r_ix_i$ is in $\mathcal{V}$.

(iv) (**Scaling.**) If $c \in R$ and if $\sum_{i\in S} r_ix_i$ is in $\mathcal{V}$, then $\sum_{i\in S} u_ix_i$ is in $\mathcal{V}$, where each $u_i = cr_i$.

(v) (**Formal addition.**) If $\sum_{i\in S} r_ix_i$ and $\sum_{j\in T} u_jy_j$ are in $\mathcal{V}$, then so too is the formal sum $\sum_{i\in S} r_ix_i + \sum_{j\in T} u_jy_j$.

The relations in $\mathcal{V} \subseteq \operatorname{Rel}(M, R)$ are called ***valid*** linear relations: by replacing $\operatorname{Rel}(M, R)$ by $\mathcal{V}$, we obtain partial linear-algebraic versions of the usual linear-algebraic definitions. In particular, a set $A \subseteq M$ is ***independent*** in $(M, R, \mathcal{V})$ if every *valid* linear relation among elements of A is trivial, ***spans*** $(M, R, \mathcal{V})$ if for every $\mathbf{a} \in M$, there exists a *valid* linear relation $\mathbf{a} + \sum_{i\in S\subseteq A}(-r_i)x_i$, and is a ***basis*** for $(M, R, \mathcal{V})$ if it is both independent and spanning. If there is a notion of positivity in R (as there is in our context, where R is a subring of $\mathbb{R}$), we say A is a ***positive spanning set*** for $(M, R, \mathcal{V})$ if for every $\mathbf{a} \in M$, there exists a valid linear relation $\mathbf{a} + \sum_{i\in S\subseteq A}(-r_i)x_i$ with *positive* scalars $r_i > 0$, and a ***positive basis*** if it is also independent.

The standard non-constructive proof, using Zorn's Lemma, of the existence of a basis for an arbitrary vector space can be used to prove the following generalization of [21, Proposition 4.6] to all partial linear structures over fields. The first part of the result appears as [24, Theorem 2.4].

Theorem 2.2.23. *Suppose R is a field. There exists a basis for the partial linear structure $(M, R, \mathcal{V})$, and given a spanning set A for $(M, R, \mathcal{V})$, there exists a basis for $(M, R, \mathcal{V})$ contained in A.*

Remark 2.2.24. Specializing Definition 2.2.22 to the mutation-linear structure $R^B = (R^n, R, \mathcal{V}_{\eta^B})$ where the set of valid relations $\mathcal{V}_{\eta^B}$ consists of B-coherent linear relations with coefficients in R, we obtain the notions of an ***independent set for*** R^B, a ***(positive) spanning set for*** R^B, and a ***(positive) basis for*** R^B. When these objects were introduced in [21], they were referred to as an "R-independent set for B", a "(positive) R-spanning set for B", and "(positive) R-basis for B", respectively. (Similar language was used when referring to cone bases: see Definition 2.2.34.) We use the newer terminology, from [24], as it is more suggestive of the partial linear structure introduced there.

Observe that any module homomorphism $\lambda : M \to M'$ between two modules M and M' over the same ring R induces a map $\lambda : \text{Rel}(M, R) \to \text{Rel}(M', R)$ on linear relations by defining

$$\lambda\left(\sum_{i \in S} r_i x_i\right) = \sum_{i \in S} r_i \lambda(x_i). \tag{2.8}$$

Definition 2.2.25 (Morphism of partial linear structures)**.** Given partial linear structures $(M, R, \mathcal{V})$ and $(M', R', \mathcal{V}')$, a module homomorphism $\lambda : M \to M'$ is ***linear with respect to the partial linear structures*** if the induced map on linear relations defined by (2.8) preserves valid relations: that is, if $\lambda\left(\sum_{i \in S} r_i x_i\right) \in \mathcal{V}'$ for every $\sum_{i \in S} r_i x_i \in \mathcal{V}$.

If λ is bijective and $\lambda^{-1} : M' \to M$ is also linear, then we say that λ is an ***isomorphism of partial linear structures***. Specializing to mutation-linear structures, given two exchange matrices B and B' we call a linear map $\lambda : R^B \to R^{B'}$ ***mutation-linear*** and a bijective linear map with linear inverse a ***mutation-linear isomorphism***.

The mutation fan

As mentioned in Remark 2.2.18, the combinatorics of the piecewise-linear mutation maps $\eta^B_{\mathbf{k}}$ associated to an $n \times n$ exchange matrix B (defined in Definition 2.2.16) are encoded by the mutation fan $\mathcal{F}_B$, a geometric object in $\mathbb{R}^n$ which we now define. We also quote an assortment of associated results from [21, Section 5]. Before defining the mutation fan itself, we recall some definitions from convex geometry.

Definition 2.2.26 (Convex set, cone, face, fan)**.** A subset of $\mathbb{R}^n$ is ***convex*** if every line segment connecting two points in the set is fully contained in the set. The ***convex hull*** of a collection of points is the smallest convex set which contains it. A ***convex cone*** is a subset of $\mathbb{R}^n$ that is closed under addition and positive scaling: in particular, a convex cone is a convex set. Every closed cone can be defined as the nonnegative $\mathbb{R}$-linear span of a collection of vectors; such a cone is ***polyhedral*** if this collection of vectors is finite. A polyhedral cone is ***rational*** if these vectors are integer vectors, and ***simplicial*** if the vectors are linearly independent. A subset F of a convex set is called a ***face*** if it is convex and if any line segment in the set whose relative interior intersects F is fully contained in F. Each face of a cone is itself a cone, and the intersection of an arbitrary set of faces is a face. A one-dimensional face of a cone is called a ***ray*** of the cone: in particular, a cone is the convex hull of its rays, and if a cone is polyhedral, it has finitely many rays. A ***fan*** is a collection $\mathcal{F}$ of closed convex cones such that if $F \in \mathcal{F}$ and G is a face of F, then $G \in \mathcal{F}$, and such that the intersection of any two cones in $\mathcal{F}$ is a face of each of the two. The one-dimensional cones in $\mathcal{F}$ are called rays of the fan.

Remark 2.2.27. While in some contexts a fan is required to have finitely many cones, here we follow [21] and allow infinitely many. Also, note that the requirement that a pairwise intersection between cones in a fan must be a face of each implies that the same is true for arbitrary intersections. That is, if $\mathcal{F}$ is a fan, A is some indexing set, and for each $\alpha \in A$, F_α is a cone in $\mathcal{F}$, then $\bigcap_{\alpha \in A} F_\alpha$ is a face of F_α for each $\alpha \in A$. This is because for any $\beta \in A$ we may rewrite $\bigcap_{\alpha \in A} F_\alpha = \bigcap_{\alpha \in A}(F_\alpha \cap F_\beta)$. Each pairwise intersection $F_\alpha \cap F_\beta$ must be a face of F_β since $\mathcal{F}$ is a fan, and the intersection of an arbitrary set of faces of F_β is again a face of F_β.

Definition 2.2.28 (Fan characteristics)**.** A fan is ***complete*** if the union of its cones is the entire ambient space. A fan is ***polyhedral*** if all its cones are polyhedral. Similarly, a polyhedral fan is ***simplicial*** if all its cones are simplicial, and ***rational*** if all its cones are rational. A ***subfan*** of a fan $\mathcal{F}$ is a subset of $\mathcal{F}$ that is itself a fan. If $\mathcal{F}$ and $\mathcal{F}'$ are complete fans such that every cone in $\mathcal{F}'$ is a union of cones in $\mathcal{F}$, or equivalently, if $\mathcal{F}$ and $\mathcal{F}'$ are complete and every cone of $\mathcal{F}$ is contained in a cone of $\mathcal{F}'$, then we say that $\mathcal{F}$ ***refines*** $\mathcal{F}'$ and $\mathcal{F}'$ ***coarsens*** $\mathcal{F}$. A cone is ***maximal*** in a fan if it is not strictly contained in any other cones in the fan, and ***full-dimensional*** if it is of the same dimension as the ambient space. Given a fan $\mathcal{F}$ in $\mathbb{R}^n$ and a subspace V of $\mathbb{R}^n$, the ***restriction of*** $\mathcal{F}$ ***to*** V defined by $\mathcal{F} \cap V = \{F \cap V : F \in \mathcal{F}\}$ is a fan in V.

Definition 2.2.29 (The mutation fan $\mathcal{F}_B$)**.** Define an equivalence relation $\equiv^B$ on $\mathbb{R}^n$ by setting $\mathbf{a} \equiv^B \mathbf{b}$ if and only if $\mathbf{sgn}(\eta^B_{\mathbf{k}}(\mathbf{a})) = \mathbf{sgn}(\eta^B_{\mathbf{k}}(\mathbf{b}))$ for every finite sequence $\mathbf{k}$ of indices in $[n]$. (Given a vector $\mathbf{a} = (a_1, \ldots, a_n) \in \mathbb{R}^n$, the notation $\mathbf{sgn}(\mathbf{a})$ represents the vector of signs $(\operatorname{sgn}(a_1), \ldots, \operatorname{sgn}(a_n))$, where $\operatorname{sgn}(a_i) = a_i/|a_i|$ for $a_i \neq 0$ and $\operatorname{sgn}(0) = 0$.) The equivalence classes of $\mathbb{R}^n$ under $\equiv^B$ are called B***-classes***, and the closures of these classes, which are closed convex cones, are called B***-cones***. By [21, Theorem 5.13], the collection of B-cones together with their faces form a complete fan in $\mathbb{R}^B$. This fan is called the ***mutation fan*** for B and denoted by $\mathcal{F}_B$. Note that all known examples of mutation fans are polyhedral; however the possibility of a non-polyhedral B-cone has not been ruled out.

Proposition 2.2.30. *[21, Proposition 5.3] Every mutation map $\eta_{\mathbf{k}}^B$ is linear on every B-cone.*

Proposition 2.2.31. *[21, Proposition 5.30] A set $C \subseteq \mathbb{R}^n$ is contained in some B-cone if and only if the set $\eta_{\mathbf{k}}^B(C)$ is **sign-coherent** for every sequence $\mathbf{k}$ of indices in $[n]$. (That is, if their k^{th} coordinates weakly agree in sign for each $k \in [n]$.)*

Proposition 2.2.32. *[21, Proposition 7.2] Given any $\sigma \in \mathfrak{S}_n$, let $\sigma(B)$ denote the exchange matrix whose $(ij)^{th}$ entry is $b_{\sigma(i)\sigma(j)}$ and for $\mathbf{a} = (a_i)_{i \in [n]} \in \mathbb{R}^n$ let $\sigma(\mathbf{a}) = (a_{\sigma^{-1}(i)})_{i \in [n]}$. Then,*

$$\mathcal{F}_{\sigma(B)} = \sigma(\mathcal{F}_B)$$

where $\sigma(\mathcal{F}_B)$ denotes the collection of cones $\sigma(F) = \{\sigma(\mathbf{a}) : \mathbf{a} \in F\}$ where F is a cone in $\mathcal{F}_B$.

Proposition 2.2.33. *[21, Proposition 7.3] $\mathcal{F}_{\mu_{\mathbf{k}}(B)} = \eta_{\mathbf{k}}^B(\mathcal{F}_B)$, where $\eta_{\mathbf{k}}^B(\mathcal{F}_B)$ denotes the collection of cones $\eta_{\mathbf{k}}^B(F) = \{\eta_{\mathbf{k}}^B(\mathbf{a}) : \mathbf{a} \in F\}$ such that F is a cone in $\mathcal{F}_B$.*

Definition 2.2.34 (Cone basis for R^B)**.** A collection of vectors $(\mathbf{b}_i : i \in I)$ in R^n indexed by some set I is a ***cone basis for*** R^B if and only if the following two conditions hold:

(i) If F is a B-cone, then $R^n \cap F$ is contained in the R-linear span of $\{\mathbf{b}_i : i \in I\} \cap F$.

(ii) The collection $(\mathbf{b}_i : i \in I)$ is an independent set for R^B.

The following is a combination of [21, Propositions 5.9 and 6.4] and [24, Proposition 2.18]:

Proposition 2.2.35. *If $(\mathbf{b}_i : i \in I)$ is a cone basis for R^B, then it is a basis for R^B. Furthermore, for each $\mathbf{a} \in R^B$, there exists some finite subset $S \subseteq I$ such that $\mathbf{a} + \sum_{i \in S}(-r_i)\mathbf{b}_i$ is a B-coherent linear relation with coefficients in R where the basis vectors $\{\mathbf{b}_i : i \in S\}$ are all contained in some common B-cone.*

Proposition 2.2.36. *[21, Proposition 6.7] Given a collection $(\mathbf{b}_i : i \in I)$ of vectors in R^n, the following conditions are equivalent:*

(i) $(\mathbf{b}_i : i \in I)$ is a positive basis for R^B.

(ii) $(\mathbf{b}_i : i \in I)$ is a positive cone basis for R^B.

(iii) $(\mathbf{b}_i : i \in I)$ is an independent set for R^B with the following property: If F is a B-cone, then $R^n \cap F$ is contained in the nonnegative R-linear span of $\{\mathbf{b}_i : i \in I\}$.

Definition 2.2.37 (R-part of a fan)**.** Suppose $\mathcal{F}$ is a fan in $\mathbb{R}^n$. We call a fan $\mathcal{F}'$ the ***R*-part** of $\mathcal{F}$ if it satisfies the following conditions:

(i) Each cone in $\mathcal{F}'$ is the non-negative linear span of finitely many vectors in R^n. (Thus in particular, $\mathcal{F}$ is polyhedral. If the underlying ring R is $\mathbb{Z}$ or $\mathbb{Q}$, then $\mathcal{F}$ is rational.)

(ii) Every cone in $\mathcal{F}'$ is contained in a cone of $\mathcal{F}$.

(iii) For each cone F of $\mathcal{F}$, there is a unique largest cone (under containment) among cones og $\mathcal{F}'$ which are contained in F, and this largest cone contains $R^n \cap F$.

While the R-part $\mathcal{F}'$ of a fan $\mathcal{F}$ in $\mathbb{R}^n$ need not exist, if it does it is unique. Furthermore, every cone of $\mathcal{F}$ spanned by vectors in R^n is a cone in $\mathcal{F}'$, and the full-dimensional cones in $\mathcal{F}$ and $\mathcal{F}'$ coincide by [21, Proposition 6.10]. We now prove a straightforward lemma about the R-part of the restriction of a fan to a subspace (see Definition 2.2.28).

Lemma 2.2.38. *Let $\mathcal{F}$ be a fan in $\mathbb{R}^n$ with R-part $\mathcal{F}'$ and let V be a subspace of $\mathbb{R}^n$ defined by a (finite) system of homogenous linear equations with coefficients in R. Then the R-part of the fan $\mathcal{F} \cap V$ is given by $\mathcal{F}' \cap V$.*

Proof. We show that since $\mathcal{F}'$ satisfies each of the three conditions of Definition 2.2.37 with respect to $\mathcal{F}$, then $\mathcal{F}' \cap V$ satisfies them with respect to $\mathcal{F} \cap V$. Let $F' \cap V$ be any cone in $\mathcal{F}' \cap V$. By Condition (i), F' is the non-negative linear span of finitely many vectors in R^n, and by hypothesis V is defined by (finitely many) linear equations with coefficients in R, so $G = F' \cap V$ is the non-negative linear span of finitely many vectors in $R^n \cap V$. By Condition (ii), F' is contained in a cone F in $\mathcal{F}$; thus $G = F' \cap V$ is contained in the cone $F \cap V$ of $\mathcal{F} \cap V$. Finally, let $F \cap V$ be any cone in $\mathcal{F} \cap V$. By Condition (iii), there is a unique largest cone F' among cones in $\mathcal{F}'$ contained in F, and $F' \supseteq R^n \cap F$. It follows that $F' \cap V$ is a cone in $\mathcal{F}' \cap V$ contained in $F \cap V$, and furthermore that $F \cap V \supseteq R^n \cap F \cap V$. Suppose $H' \cap V$ is another cone of $\mathcal{F}' \cap V$ which is contained in $F \cap V$. Then as established earlier, $H' \cap V$ is the non-negative linear span of finitely many vectors in $R^n \cap V$, and in particular of finitely many vectors in $R^n \cap F \cap V$. Since all such vectors are contained in $F' \cap V$, it follows that $H' \cap V$ is contained in $F' \cap V$ as well. □

Definition 2.2.39 ($\mathcal{F}_B^R$)**.** Suppose $(\mathbf{b}_i : i \in I)$ is a positive basis for R^B. Define $\mathcal{F}_B^R$ to be the collection of all cones spanned by sets of the form $\{\mathbf{b}_i : i \in I\} \cap F$ for some B-cone F, together with all the faces of such cones. Then $\mathcal{F}_B^R$ is a simplicial fan and is the R-part of $\mathcal{F}_B$ by [21, Proposition 6.11]. The construction does not depend on the choice of $(\mathbf{b}_i : i \in I)$, as by [21, Proposition 6.2], positive bases are unique up to scaling by positive units.

Proposition 2.2.40. *[21, Corollary 6.12] If a positive basis exists for $\mathbb{Z}^B$, then it is unique and consists of the smallest nonzero integer vector in each ray of $\mathcal{F}_B^{\mathbb{Z}}$. If R is a field and a positive basis exists for R^B, then a collection of vectors is a positive basis for R^B if and only if it consists of exactly one nonzero vector in each ray of $\mathcal{F}_B^R$, or equivalently, one nonzero vector in each ray of $\mathcal{F}_B \cap R^n$.*

In addition to the various properties of the mutation fan from [21] cited above, we will eventually need one additional result, an easy corollary to Proposition 2.2.32 which follows from the application of Definition 2.2.37. We omit the details of the proof: it suffices to verify the three conditions of the definition.

Corollary 2.2.41. *For $\sigma \in \mathfrak{S}_n$, let $\sigma(B)$ and $\sigma(\mathbf{a})$ be as in Proposition 2.2.32. Let $\mathcal{F}'$ denote the R-part of $\mathcal{F}_B$ and let $\sigma(\mathcal{F}')$ denotes the collection of cones $\sigma(F) = \{\sigma(\mathbf{a}) : \mathbf{a} \in F\}$ such that F is a cone in $\mathcal{F}'$. Then $\sigma(\mathcal{F}')$ is the R-part of $\sigma(\mathcal{F}_B)$, and if a positive basis exists for R^B,*

$$\mathcal{F}^R_{\sigma(B)} = \sigma(\mathcal{F}^R_B)$$

2.2.3 Connections

We conclude this discussion of background material by providing explicit connections from [21] between the studies of geometric cluster algebras and mutation-linear algebra, both of which rely fundamentally on the mutation of exchange matrices.

Remark 2.2.42. At the time of the publication of [21], several of the results which are referenced here relied on the assumption of the "Standard Hypotheses" on B: namely, that for every exchange matrix B' in the mutation class of either B or $-B$, the non-negative orthant $O^n = (\mathbb{R}_{\geq 0})^n$ is a B'-cone. However, [21, Proposition 8.9] shows that these hypotheses hold for B if and only if B satisfies the standard conjecture known as the "sign-coherence of $\mathbf{c}$-vectors" [19, (1.8)], which is proven in [14, Corollary 5.5] to hold for *all* exchange matrices. Thus we omit any references to the Standard Hypotheses when quoting from [21].

Recall that in this section we have taken as given both an $n \times n$ exchange matrix B and an underlying ring $R \in \{\mathbb{Z}, \mathbb{Q}, \mathbb{R}\}$. The first result is that universal geometric coefficients for B over R (Definition 2.2.9) and bases for R^B (Remark 2.2.24) are equivalent.

Theorem 2.2.43. *[21, Theorem 4.4] Let $\tilde{B}$ be an extended exchange matrix with principal part B whose coefficient rows have entries in R. Then these coefficient rows are universal geometric coefficients for B over R if and only if they form a basis for R^B.*

The next few results relate $\mathbf{g}$-vectors for B^T, the matrix transpose of B (Definition 2.2.11) to the mutation fan $\mathcal{F}_B$ for B (Definition 2.2.29).

Definition 2.2.44 ($\mathbf{g}$-vector cones)**.** For each cluster $\mathbf{x}_t^{B;t_0}$ in $\mathcal{A}_\bullet(B)$, define the associated ***$\mathbf{g}$-vector cone*** $\mathrm{Cone}_t^{B;t_0}$ to be the nonnegative linear span of the $\mathbf{g}$-vectors $\left\{\mathbf{g}_{i;t'}^{B;t_0} : i \in [n]\right\} \in \mathbb{Z}^n$ of the cluster variables in $\mathbf{x}_t^{B;t_0}$.

Definition 2.2.45 ($\mathbf{g}$-vector fan)**.** The set of $\mathbf{g}$-vector cones $\left\{\mathrm{Cone}_t^{B;t_0} : t \in \mathbb{T}_n\right\}$ together with their faces form a simplicial fan in $\mathbb{R}^n$, called the ***$\mathbf{g}$-vector fan*** for B [21, Definition 8.6 and Theorem 8.7]. In particular, the $\mathbf{g}$-vector fan for B^T is the subfan $\mathcal{F}_B^0$ of the mutation fan $\mathcal{F}_B$ formed of all full-dimensional cones F in $\mathcal{F}_B$ that are ***transitively adjacent*** to the nonnegative orthant O_n. (We say F and O_n are transitively adjacent if there is a sequence $O_n = F_0, F_1, \ldots, F_q = F$ of full-dimensional cones in $\mathcal{F}_B$ (possibly with $q = 0$) such that C_{i-1} and C_i are ***adjacent*** for each $i \in [q]$: i.e., have disjoint interiors but share a face of codimension 1). If $t_0, t_1, \ldots, t_q = t$ is a path in $\mathbb{T}_n$ where for each $i \in [q]$ the edge from t_{i-1} to t_i is labeled $k_i \in [n]$, then $\mathrm{Cone}_t^{B^T;t_0} = \eta_{k_1,\ldots,k_q}^{\mu_{k_q,\ldots,k_1}(B)}(O_n)$ [21, Proposition 8.13].

When B is of finite type, the $\mathbf{g}$-vector fan is complete (see, e.g., [21, Theorem 10.6]). Thus the following result is an immediate consequence of Definition 2.2.45 and the results cited therein:

Proposition 2.2.46. *[21, Section 10] When B is of finite type, the mutation fan $\mathcal{F}_B$ coincides with the $\mathbf{g}$-vector fan for B^T, and is therefore rational and simplicial. Furthermore, all maximal cones in $\mathcal{F}_B$ are full-dimensional and correspond bijectively with clusters in $\mathcal{A}_\bullet(B^T)$, where the spanning rays of a maximal cone are precisely the rays in the $\mathbf{g}$-vectors of the cluster variables in the corresponding cluster.*

Remark 2.2.47. Recall from Definition 2.2.37 that the full-dimensional cones in $\mathcal{F}_B$ are precisely the full dimensional cones in $\mathcal{F}_B^R$: thus for B of finite type, Proposition 2.2.46 implies that $\mathcal{F}_B^R = \mathcal{F}_B$ for any choice of underlying ring R.

Theorem 2.2.48. *[21, Theorem 10.12] When B is of finite type, the $\mathbf{g}$-vectors associated to B^T constitute a positive basis for R^B.*

2.3 Folding

This section begins with a precise definition of the folding of an exchange matrix under an automorphism, following [2, 4, 18]. We then provide several known results about the folding of the associated coefficient-free and principal coefficients cluster algebras from [1, 4].

Matrix folding

Definition 2.3.1 (Orbits)**.** Let $\sigma \in \mathfrak{S}_n$. The ***$\sigma$-orbits*** of $[n]$ are the equivalence classes defined by the relation $i \sim j$ if and only if $i = \sigma^k(j)$ for some $k \in \mathbb{Z}$. The notation $\bar{\imath} = \{j \in [n] : i \sim j\}$ represents the σ-orbit containing $i \in [n]$.

Definition 2.3.2 (Admissable automorphism)**.** Let B be an $n \times n$ exchange matrix and let $\sigma \in \mathfrak{S}_n$. We call σ an ***automorphism*** of B (or equivalently, say that B is ***symmetric with respect to*** σ) if $b_{\sigma(i)\sigma(j)} = b_{ij}$ for all $i, j \in [n]$. If furthermore,

(i) the $\bar{\imath} \times \bar{\imath}$ block of B is a zero matrix for each $i \in [n]$, and

(ii) all entries in the $\bar{\imath} \times \bar{\jmath}$ block of B agree weakly in sign for all $i, j \in [n]$,

then we say that σ is a ***admissable automorphism*** of B (see Example 2.3.5).

As a consequence of Condition (i) in Definition 2.3.2 and the definition of matrix mutation in (2.2), if σ is an admissable automorphism of B then the result of applying a composition of matrix mutations μ_{k_i} over a given σ-orbit $\bar{k} = \{k_1, \ldots, k_\ell\}$ of $[n]$ is independent of order:

$$\mu_{k_\ell} \circ \cdots \circ \mu_{k_1}(B) = \mu_{k_{\tau(\ell)}} \circ \cdots \circ \mu_{k_{\tau(1)}}(B) \text{ for any } \tau \in \mathfrak{S}_\ell. \tag{2.9}$$

Indeed, the same reasoning implies that matrix mutations within a given σ-orbit $\bar{k} = \{k_1, \ldots, k_\ell\}$ commute for any extended exchange matrix $\tilde{B}$ with principal part B:

$$\mu_{k_\ell} \circ \cdots \circ \mu_{k_1}(\tilde{B}) = \mu_{k_{\tau(\ell)}} \circ \cdots \circ \mu_{k_{\tau(1)}}(\tilde{B}) \text{ for any } \tau \in \mathfrak{S}_\ell. \tag{2.10}$$

Thus the product $\Pi_{k\in\bar{k}}\mu_k$ is well-defined. (See [4, Lemma 2.12] for a proof of (2.9) that is easily extended to prove (2.10).)

Definition 2.3.3 (Orbit-mutation)**.** Given an admissable automorphism σ of B and σ-orbit $\bar{k}$ of $[n]$, the σ***-orbit-mutation*** $\mu^\sigma_{\bar{k}}$ in direction $\bar{k}$ is the composition

$$\mu^\sigma_{\bar{k}} = \prod_{k\in\bar{k}} \mu_k. \tag{2.11}$$

As in (2.3), given a finite sequence $\bar{\mathbf{k}} = \bar{k}_1, \ldots, \bar{k}_q$ of σ-orbits, we use the notation $\mu^\sigma_{\bar{\mathbf{k}}}$ to denote the composition $\mu^\sigma_{\bar{\mathbf{k}}} = \mu^\sigma_{\bar{k}_q} \circ \cdots \circ \mu^\sigma_{\bar{k}_1}$.

Definition 2.3.4 (Stable automorphism)**.** An automorphism σ of B is ***stable*** if it is admissable and remains so under orbit-mutation. That is, if and only if σ is admissable on $\mu^\sigma_{\bar{\mathbf{k}}}(B)$ for any finite sequence of σ-orbits $\bar{\mathbf{k}}$. It is easily checked that if σ is a stable automorphism of B, so too is σ^q for any $q \in \mathbb{Z}$.

Example 2.3.5. Consider the 5×5 exchange matrix B below left with admissable automorphism $\sigma = (12)(3)(45)$. For convenience in checking Conditions (i) and (ii) of Definition 2.3.2, we have divided B into blocks indexed by the three σ-orbits $\bar{1} = \{1, 2\}, \bar{3} = \{3\}$ and $\bar{5} = \{4, 5\}$.

$$\begin{array}{c} \bar{1} \\ \\ B = \ \bar{3} \\ \\ \bar{5} \end{array} \left[\begin{array}{cc|c|cc} 0 & 0 & -1 & 2 & 0 \\ 0 & 0 & -1 & 0 & 2 \\ \hline 1 & 1 & 0 & -1 & -1 \\ \hline -2 & 0 & 1 & 0 & 0 \\ 0 & -2 & 1 & 0 & 0 \end{array}\right] \xrightarrow{\mu_{\bar{3}}=\mu_3} \begin{array}{c} \bar{1} \\ \\ \bar{3} \\ \\ \bar{5} \end{array} \left[\begin{array}{cc|c|cc} 0 & 0 & 1 & 1 & -1 \\ 0 & 0 & 1 & -1 & 1 \\ \hline 1 & 1 & 0 & -1 & -1 \\ \hline -1 & 1 & -1 & 0 & 0 \\ 1 & -1 & -1 & 0 & 0 \end{array}\right] = B'$$

When we perform orbit-mutation on B in direction $\bar{3}$, the resulting exchange matrix B' no longer satisfies Condition (ii) with respect to σ: in particular, the entries in the $\bar{1} \times \bar{5}$ block of B' do not agree weakly agree in sign. (Skew-symmetrically, neither do the entries in the $\bar{5} \times \bar{1}$ block.) Thus σ is *not* a stable automorphism of B.

Remark 2.3.6. As mentioned in Section 2.1 and illustrated in Example 2.3.5, not all admissable automorphisms are stable. Furthermore, if an exchange matrix B is not mutation-finite (see Definition 2.2.13), then it is impossible to prove the stability of a given admissable automorphism computationally. However, every admissable automorphism on an *acyclic* exchange matrix (see Definition 2.2.14) is stable [2, Theorem 3.1.11], which implies, in particular, that the classical foldings from Lie theory of finite and affine-type root systems all extend to foldings of the corresponding acyclic exchange matrices.

If a given admissable automorphism $\sigma \in \mathfrak{S}_n$ of exchange matrix B is stable, then it defines a new, "folded" exchange matrix $\pi_\sigma(B)$ with rows and columns indexed by the σ-orbits of $[n]$ in a natural way:

Definition 2.3.7 (Matrix folding)**.** Given stable automorphism $\sigma \in \mathfrak{S}_n$ of exchange matrix B, define the ***folding*** π_σ of B under σ as follows:

$$\begin{aligned} \pi_\sigma(B) =& (c_{\bar{\imath}\bar{\jmath}})_{\substack{\sigma\text{-orbits} \\ \bar{\imath},\bar{\jmath} \text{ of } [n]}} \text{ where} \\ c_{\bar{\imath}\bar{\jmath}} =& \sum_{i \in \bar{\imath}} b_{ij} \text{ for any } j \in \bar{\jmath}. \end{aligned} \tag{2.12}$$

The following simple computation shows that the matrix entries $c_{\bar{\imath}\bar{\jmath}} = \sum_{i\in\bar{\imath}} b_{ij}$ of $\pi_\sigma(B)$ in (2.12) are well-defined because σ is an automorphism of B: Let $j_1, j_2 \in \bar{\jmath}$. Then $j_2 = \sigma^q(j_1)$ for some $q \in \mathbb{N}$, and we have the following chain of equalities, where the last equality is obtained by reindexing the sum over the orbit $\bar{\imath}$ by setting $i' = \sigma^q(i)$:

$$\sum_{i\in\bar{\imath}} b_{ij_1} = \sum_{i\in\bar{\imath}} b_{\sigma^q(i)\sigma^q(j_1)} = \sum_{i\in\bar{\imath}} b_{\sigma^q(i)j_2} = \sum_{i'\in\bar{\imath}} b_{i'j_2}.$$

Clearly the folded matrix $\pi_\sigma(B)$ is an integer matrix; it is also skew-symmetrizable and therefore an exchange matrix. (See [4, Lemma 2.5] for the proof of skew-symmetrizability.) The matrix B is also referred to as the ***unfolding*** of $\pi_\sigma(B)$ under σ. We write m for the number of σ-orbits, so that $\pi_\sigma(B)$ is $m \times m$.

The key feature of folding an exchange matrix B under a stable automorphism σ is the relationship between orbit-mutation of B and mutation of $\pi_\sigma(B)$. Roughly speaking, the stability of σ ensures that matrix mutation commutes with folding. More precisely, for any sequence $\overline{\mathbf{k}}$ of σ-orbits, σ is a stable automorphism of $\mu^\sigma_{\overline{\mathbf{k}}}(B)$, and

$$\pi_\sigma(\mu^\sigma_{\overline{\mathbf{k}}}(B)) = \mu_{\overline{\mathbf{k}}}(\pi_\sigma(B)). \tag{2.13}$$

This equality follows from the fact that $\mu^\sigma_{\overline{\mathbf{k}}}(B)$ satisfies Condition (ii) of Definition 2.3.2 (see [4, Theorem 2.24] for the complete proof). We show in the next section that an analogous folding relationship holds for the mutation maps associated to B and $\pi_\sigma(B)$ (see (2.24)).

Cluster algebra folding

Before proceeding to consider the consequences of folding on mutation-linear algebra, we relate some known facts about cluster algebra folding. The description of the folding of coefficent-free cluster algebras comes from [4], and the extension to cluster algebras with principal coefficients is from [1]. Suppose that B is an $n \times n$ exchange matrix which folds to the $m \times m$ exchange matrix $\pi_\sigma(B)$ under $\sigma \in \mathfrak{S}_n$.

Let $\mathcal{A}_0(B)$ denote the coefficient-free cluster algebra for B with initial seed $\big((x_i)_{i\in[n]}, B\big)$, and let $\mathcal{A}_0(\pi_\sigma(B))$ denote the coefficient-free cluster algebra with initial seed $\Bigg((x_{\bar{\imath}})_{\substack{\sigma\text{-orbits} \\ \bar{\imath} \text{ of } [n]}}, \pi_\sigma(B)\Bigg)$. Define an action σ on the ambient field $\mathbb{Q}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ of $\mathcal{A}_0(B)$ by setting

$$\sigma(x_i) = x_{\sigma(i)}. \tag{2.14}$$

Now reuse the symbol π_σ for the following map, where $\bar{\imath}$ denotes, as usual, the σ-orbit of i:

$$\begin{aligned} \pi_\sigma : \mathbb{Q}[x_i^{\pm 1}]_{i\in[n]} &\to \mathbb{Q}[x_{\bar{\imath}}^{\pm 1}]_{\substack{\sigma\text{-orbits} \\ \bar{\imath} \text{ of } [n]}} \\ x_i &\mapsto x_{\bar{\imath}}. \end{aligned} \tag{2.15}$$

Let $\mathcal{A}_0^\sigma(B)$ denote the subalgebra of $\mathcal{A}_0(B)$ generated by all clusters in seeds $\mu^\sigma_{\overline{\mathbf{k}}}\big((x_i)_{i\in[n]}, B\big)$ reachable from the initial seed by orbit-mutation. The following is a rephrasing of [4, Theorem 2.24 and Corollary 2.25]:

Theorem 2.3.8. *The action of σ on $\mathbb{Q}[x_i^{\pm 1}]_{i\in[n]}$ defined by* (2.14) *induces an action on the* *cluster algebra $\mathcal{A}_0(B)$, where for each seed $((x_{i;t})_{i\in[n]}, B_t)$ in the cluster pattern for $\mathcal{A}_0(B)$,*

$$\sigma\left((x_{i;t})_{i\in[n]}, B_t\right) = \left((x_{\sigma(i);t})_{i\in[n]}, \sigma(B_t)\right).$$

Then for every sequence $\bar{\mathbf{k}}$ of σ-orbits $\bar{\imath}$ of $[n]$, the seed $\mu^{\sigma}_{\bar{\mathbf{k}}}\left((x_i)_{i\in[n]}, B\right) = \left((x_{i;t})_{i\in[n]}, B_t\right)$ is fixed under action by σ. If for each such seed we define

$$\pi_\sigma\left(\{x_{i;t}\}_{i\in[n]}, B_t\right) = \left((\pi_\sigma(x_{i;t}))_{i\in[n]}, \pi_\sigma(B_t)\right),$$

where the action of π_σ on cluster variables is given by (2.15), then

$$\pi_\sigma\left(\mu^{\sigma}_{\bar{\mathbf{k}}}\left((x_i)_{i\in[n]}, B\right)\right) = \mu_{\bar{\mathbf{k}}}\left((x_{\bar{\imath}})_{\substack{\sigma\text{-orbits}\\ \bar{\imath}\text{ of }[n]}}, \pi_\sigma(B)\right). \tag{2.16}$$

Thus every seed in $\mathcal{A}_0(\pi_\sigma(B))$ is the folding of a σ-invariant seed in $\mathcal{A}_0(B)$, and hence $\mathcal{A}_0(\pi_\sigma(B))$ can be identified with a $\mathbb{Z}$-subalgebra of $\pi_\sigma(\mathcal{A}_0(B))$. Specifically, $\pi_\sigma(\mathcal{A}^\sigma_0(B)) = \mathcal{A}_0(\pi_\sigma(B))$.

In general, $\pi_\sigma(\mathcal{A}_0(B))$ is larger than $\pi_\sigma(\mathcal{A}^\sigma_0(B))$: there are some seeds in $\mathcal{A}_0(B)$ which do not fold to seeds in $\mathcal{A}_0(\pi_\sigma(B))$. However when B is of finite type, $\pi_\sigma(\mathcal{A}_0(B))$ and $\pi_\sigma(\mathcal{A}^\sigma_0(B))$ coincide. The following is [4, Theorem 3.13]:

Theorem 2.3.9. *If B is of finite type, then*

$$\pi_\sigma(\mathcal{A}_0(B)) = \pi_\sigma(\mathcal{A}^\sigma_0(B)) = \mathcal{A}_0(\pi_\sigma(B)).$$

In particular, $\pi_\sigma(B)$ is also of finite type.

Let $\mathcal{A}_\bullet(B)$ denote the principal-coefficients cluster algebra for B with initial cluster $(x_i)_{i\in[n]}$, and let $\mathcal{A}_\bullet(\pi_\sigma(B))$ denote the principal-coefficients cluster algebra for $\pi_\sigma(B)$ with initial cluster $(x_{\bar{\imath}})_{\substack{\sigma\text{-orbits}\\ \bar{\imath}\text{ of }[n]}}$. Extend (2.14) to define the action of σ on the ambient field $\mathbb{Q}[x_i^{\pm1}; y_j]_{i,j\in[n]}$ of $\mathcal{A}_\bullet(B)$ by setting $\sigma(y_j) = y_{\sigma(j)}$, and extend (2.15) to define $\pi_\sigma : \mathbb{Q}[x_i^{\pm1}; y_j]_{i,j\in[n]} \to \mathbb{Q}[x_{\bar{\imath}}^{\pm1}; y_{\bar{\jmath}}]_{\substack{\sigma\text{-orbits}\\ \bar{\imath},\bar{\jmath}\text{ of }[n]}}$ by setting $\pi_\sigma(y_j) = y_{\bar{\jmath}}$. Continuing the extension of the coefficient-free case, let $\mathcal{A}^\sigma_\bullet(B)$ denote the subalgebra of $\mathcal{A}_\bullet(B)$ generated by all clusters in seeds $\mu^\sigma_{\bar{\mathbf{k}}}((x_i)_{i\in[n]}, \tilde{B})$ reachable from the initial seed by orbit-mutation. The following is a rephrasing of [1, Theorem 2.9]:

Theorem 2.3.10. *The action of σ on $\mathbb{Q}[x_i^{\pm1}; y_i]_{i\in[n]}$ induces an action on $\mathcal{A}_\bullet(B)$, and every seed $\mu^\sigma_{\bar{\mathbf{k}}}((x_i)_{i\in[n]}, \tilde{B}) = ((x_{i;t})_{i\in[n]}, \tilde{B}_t)$ is σ-invariant, where, if $\tilde{B}_t$ has principal part B_t and complementary part D_t, we define*

$$\sigma(\tilde{B}_t) = \sigma\begin{bmatrix} B_t \\ D_t \end{bmatrix} = \begin{bmatrix} \sigma(B_t) \\ \sigma(D_t) \end{bmatrix}.$$

Defining $\pi_\sigma(\tilde{B}_t) = \begin{bmatrix}\pi_\sigma(B_t)\\ \pi_\sigma(D_t)\end{bmatrix}$ *analogously, the restriction of* π_σ *to* $\mathcal{A}^\sigma_\bullet(B)$ *is a* $\mathbb{Z}^n$*-graded surjective homomorphism which preserves cluster structure:*

$$\pi_\sigma : \mathcal{A}^\sigma_\bullet(B) \to \mathcal{A}_\bullet(\pi_\sigma(B))$$
$$((x_{i;t})_{i\in[n]}, \tilde{B}_t) \mapsto ((x_{\bar{\imath};t})_{\substack{\sigma\text{-orbits}\\ \bar{\imath}\text{ of }[n]}}, \pi_\sigma(\tilde{B}_t))$$

2.4 Mutation-linear algebra folding

Throughout this section let $R \in \{\mathbb{Z}, \mathbb{Q}, \mathbb{R}\}$ be the underlying ring, let B be an $n\times n$ exchange matrix with stable automorphism $\sigma \in \mathfrak{S}_n$ and let $C = \pi_\sigma(B)$ be the $m \times m$ folding of B under σ, with rows and columns indexed by the m σ-orbits $\bar{\imath}$ of $[n]$.

2.4.1 Basis folding

We begin by defining orbit-mutation maps $\eta^{\sigma,B}_{\mathbf{k}}$ analogous to the matrix orbit-mutations $\mu^\sigma_{\mathbf{k}}$ of Definition 2.3.3, and an m-dimensional subspace V^σ of $\mathbb{R}^n$ on which these orbit-mutation maps act essentially the same as the mutation maps $\eta^C_{\mathbf{k}}$ on $\mathbb{R}^m$. Our main result, Theorem 2.4.9, is that any basis for R^B which is preserved under the action of σ folds to a spanning set for R^C by taking sums over σ-orbits. We conclude the section by discuss our conjecture that to obtain a basis for R^C, one must restrict consideration to a certain subset of σ-orbits: namely, those which are contained in B-cones. Further justification for this conjecture is then provided in Section 2.4.2.

Recall from Definition 2.2.16 that $\eta^B_{\mathbf{k}}(\mathbf{a})$ is the coefficient row of $\mu_{\mathbf{k}}(\tilde{B})$, where $\tilde{B}$ is the extended exchange matrix with principal part B and single coefficient row $\mathbf{a}$. It therefore follows immediately from (2.10) that for any σ-orbit $\bar{k}$, the composition of all mutation maps η^B_k for $k \in \bar{k}$ is independent of order, so we may write this composition as a product.

Definition 2.4.1 (Orbit-mutation map)**.** The σ***-orbit-mutation map*** $\eta^{\sigma,B}_{\bar{k}}$ in direction $\bar{k}$ is the composition

$$\eta^{\sigma,B}_{\bar{k}} = \prod_{k\in\bar{k}} \eta^B_k. \tag{2.17}$$

Analogously to Equation (2.3), given a finite sequence $\bar{\mathbf{k}} = \bar{k}_1, \dots, \bar{k}_q$ of σ-orbits, we use the notation $\eta^{\sigma,B}_{\bar{\mathbf{k}}}$ to denote the composition $\eta^{\sigma,B}_{\bar{\mathbf{k}}} = \eta^{\sigma,B_q}_{\bar{k}_q} \circ \eta^{\sigma,B_{q-1}}_{\bar{k}_{q-1}} \circ \cdots \circ \eta^{\sigma,B_1}_{\bar{k}_1}$ where $B_1 = B$ and $B_{i+1} = \mu^\sigma_{\bar{k}_i}(B_i)$ for each $i \in [q]$. We refer collectively to all such maps as the σ***-orbit-mutation maps*** associated to the pair (B, σ).

Definition 2.4.2 (Orbit-B-coherent linear relation)**.** Let $(\mathbf{v}_i : i \in S)$ be a set of vectors in $\mathbb{R}^n$ and let $(r_i : i \in S)$ be a set of scalars in R, both indexed by the same finite set S. Analogously to Definition 2.2.19, the formal sum $\sum_{i \in S} r_i \mathbf{v}_i$ is an ***orbit-B-coherent linear relation with coefficients in*** R (briefly, ***orbit-B-coherent***) if the following equalities hold for every finite sequence $\bar{\mathbf{k}}$ of σ-orbits, including the empty sequence:

$$\sum_{i \in S} r_i \eta_{\bar{\mathbf{k}}}^{\sigma,B}(\mathbf{v}_i) = \mathbf{0}, \text{ and} \tag{2.18}$$

$$\sum_{i \in S} r_i \mathbf{min}(\eta_{\bar{\mathbf{k}}}^{\sigma,B}(\mathbf{v}_i), \mathbf{0}) = \mathbf{0}. \tag{2.19}$$

Denote the collection of all orbit-B-coherent linear relations with coefficients in R by $\mathcal{V}_{\eta^{\sigma,B}}$.

Remark 2.4.3. Define $R^{\sigma,B} = (R^n, R, \mathcal{V}_{\eta^{\sigma,B}})$. It is easily checked, in the same manner as for $R^B = (R^n, R, \mathcal{V}_{\eta^B})$, that $R^{\sigma,B}$ is a partial linear structure in the sense of Definition 2.2.22: we refer to is as an ***orbit-mutation-linear structure***. Futhermore, since an orbit-mutation map $\eta_{\bar{k}}^{\sigma,B}$ is in particular a mutation map in the usual sense (namely, $\eta_{k_1,\ldots,k_\ell}^B$ where $\bar{k} = \{k_1, \ldots, k_\ell\}$), the identity map from R^B to $R^{\sigma,B}$ is linear with respect to partial linear structures in the sense of Definition 2.2.25.

While thus far we have viewed $\sigma \in \mathfrak{S}_n$ as acting on the entries of B, as described in the hypotheses of Proposition 2.2.32 and Corollary 2.2.41, we may also view σ as a linear map on $\mathbb{R}^n$ acting via coordinate permutation as follows:

$$\begin{aligned} \sigma : \mathbb{R}^n &\xrightarrow{\cong} \mathbb{R}^n \\ (a_1, \ldots, a_n) &\mapsto (a_{\sigma^{-1}(1)}, \ldots, a_{\sigma^{-1}(n)}) \end{aligned} \tag{2.20}$$

Employing the terminology introduced in Section 2.1, we refer to a set of vectors $A \subseteq \mathbb{R}^n$ as ***symmetric with respect to*** σ, or ***σ-symmetric***, if $\sigma(A) = A$. While (2.20) may seem more like a definition of the map $\sigma^{-1} : \mathbb{R}^n \to \mathbb{R}^n$, this formulation matches [21,24] and the adjustment is harmless; if the set $A \subseteq \mathbb{R}^n$ is σ-symmetric, then it is σ^{-1}-symmetric as well. Indeed, the set A is σ^q-symmetric for any $q \in \mathbb{Z}$, so we may define the following generalized version of the map in (2.20):

$$\begin{aligned} \sigma^q : \mathbb{R}^n &\xrightarrow{\cong} \mathbb{R}^n \text{ for } q \in \mathbb{Z} \\ (a_1, \ldots, a_n) &\mapsto (a_{\sigma^{-q}(1)}, \ldots, a_{\sigma^{-q}(n)}) \end{aligned} \tag{2.21}$$

The following is a weaker version of [24, Proposition 2.26]:

Proposition 2.4.4. *For each $q \in \mathbb{Z}$, the map $\sigma^q|_{R^B} : R^B \to R^B$ defined by restricting the map σ^q of* (2.21) *to R^n is a mutation-linear isomorphism.*

Definition 2.4.5 (Fixed space and module)**.** Define the ***σ-fixed space*** V^σ to be the 1-

eigenspace of σ under the coordinate permutation action defined in (2.20). That is,

$$V^\sigma = \{\mathbf{a} \in \mathbb{R}^n : \sigma(\mathbf{a}) = \mathbf{a}\} \tag{2.22}$$

Suppose $\mathbf{v} = (v_i)_{i\in[n]}$ is a vector in V^σ. Then for each of the m σ-orbits $\bar{\imath}$ of $[n]$, the coordinates v_i for $i \in \bar{\imath}$ are all equal to one another: call this common scalar $v_{\bar{\imath}}$. It follows that V^σ has dimension m, and we use ϕ to denote the following natural vector space isomorphism:

$$\begin{aligned} \phi : V^\sigma &\xrightarrow{\cong} \mathbb{R}^m \\ \mathbf{v} &\mapsto (v_{\bar{\imath}})_{\sigma\text{-orbits } \bar{\imath}} \end{aligned} \tag{2.23}$$

Denote the inverse linear transformation by $\phi^{-1} : \mathbb{R}^m \xrightarrow{\cong} V^\sigma$, where for each $i \in [n]$ the i^{th} coordinate of $\mathbf{v} = \phi^{-1}(\mathbf{w})$ is defined by $v_i = w_{\bar{\imath}}$ for each $i \in \bar{\imath}$.

Our next goal is to to give a mutation map analog of (2.13), which established a folding relationship between matrix orbit-mutation of B and matrix mutation of $C = \pi_\sigma(B)$. We first verify that the subspace V^σ of $\mathbb{R}^n$ is closed under orbit-mutation maps, and then show that the orbit-mutation maps $\eta^{\sigma B}_{\bar{\mathbf{k}}}$ "upstairs" on V^σ correspond to mutation maps $\eta^{C}_{\bar{\mathbf{k}}}$ "downstairs" on $\mathbb{R}^m$ under the vector space isomorphism ϕ defined in (2.23).

Lemma 2.4.6. *Let $\bar{\mathbf{k}}$ be a finite sequence of σ-orbits. The orbit-mutation map $\eta^{\sigma,B}_{\bar{\mathbf{k}}}$ fixes V^σ as a set, and*

$$\phi(\eta^{\sigma,B}_{\bar{\mathbf{k}}}(\mathbf{v})) = \eta^{C}_{\bar{\mathbf{k}}}(\phi(\mathbf{v})) \text{ for all } \mathbf{v} \in V^\sigma \tag{2.24}$$

Proof. We argue by induction on the length q of the sequence $\bar{\mathbf{k}} = \bar{k}_1, \ldots, \bar{k}_q$, our primary method of proof in this chapter. The result holds trivially when $\bar{\mathbf{k}}$ is the empty sequence. If $q > 0$, then for each vector $\mathbf{v} \in V^\sigma$ and index $i \in [n]$, consider the i^{th} entry $[\eta^{\sigma,B}_{\bar{k}_1}(\mathbf{v})]_i$ of $\eta^{\sigma,B}_{\bar{k}_1}(\mathbf{v})$. By Definitions 2.3.2 and 2.4.5, for each $k \in \bar{k}_1$ all matrix entries b_{ki} are of weakly the same sign and all vector entries v_k are equal. Thus by Definition 2.2.16 and (2.2),

$$[\eta^{\sigma,B}_{\bar{k}_1}(\mathbf{v})]_j = \begin{cases} -v_j & \text{if } j \in \bar{k} \\ v_j + \sum_{k\in\bar{k}} v_k b_{kj} & \text{if } j \notin \bar{k}_1, v_k \geq 0, \text{ and } b_{kj} \geq 0 \text{ for each } k \in \bar{k}_1 \\ v_j - \sum_{k\in\bar{k}} v_k b_{kj} & \text{if } j \notin \bar{k}, v_k \leq 0, \text{ and } b_{kj} \leq 0 \text{ for each } k \in \bar{k}_1 \\ v_j & \text{otherwise} \end{cases}$$

Employing the notation $v_{\bar{k}_1} = v_k$ for each $k \in \bar{k}_1$ from Definition 2.4.5 and applying Definition 2.3.7 of matrix folding, we obtain the following:

$$[\eta^{\sigma,B}_{\bar{k}_1}(\mathbf{v})]_j = \begin{cases} -v_{\bar{\jmath}} & \text{if } \bar{\jmath} = \bar{k}_1 \\ v_{\bar{\jmath}} + v_{\bar{k}_1} c_{\bar{k}_1 \bar{\jmath}} & \text{if } \bar{\jmath} \neq \bar{k}_1, v_{\bar{k}_1} \geq 0, \text{ and } c_{\bar{k}_1 \bar{\jmath}} \geq 0 \\ v_{\bar{\jmath}} - v_{\bar{k}_1} c_{\bar{k}_1 \bar{\jmath}} & \text{if } \bar{\jmath} \neq \bar{k}_1, v_{\bar{k}_1} \leq 0, \text{ and } c_{\bar{k}_1 \bar{\jmath}} \leq 0 \\ \bar{v}_{\bar{\jmath}} & \text{otherwise} \end{cases}$$

Clearly $\eta^{\sigma,B}_{\bar{k}_1}(\mathbf{v}) \in V^\sigma$ and $\phi(\eta^{\sigma,B}_{\bar{k}_1}(\mathbf{v})) = \eta^C_{\bar{k}_1}(\phi(\mathbf{v}))$. The result then follows by induction: the sequence $k_2, \dots, k_q$ has length less than q, and by (2.13), $\mu^\sigma_{\bar{k}_1}(B)$ folds to $\mu_{\bar{k}_1}(C)$ under stable automorphism σ. Thus

$$\phi(\eta^{\sigma,B}_{\bar{\mathbf{k}}}(\mathbf{v})) = \phi\left(\eta^{\sigma,\mu^\sigma_{\bar{k}_1}(B)}_{\bar{k}_2,\dots,\bar{k}_q}\left(\eta^B_{\bar{k}_1}(\mathbf{v})\right)\right) = \eta^{\mu_{\bar{k}_1}(C)}_{\bar{k}_2,\dots,\bar{k}_q}\left(\phi\left(\eta^B_{\bar{k}_1}(\mathbf{v})\right)\right) = \eta^{\mu_{\bar{k}_1}(C)}_{\bar{k}_2,\dots,\bar{k}_q}\left(\eta^C_{\bar{k}_1}(\phi(\mathbf{v}))\right) = \eta^C_{\bar{\mathbf{k}}}(\phi(\mathbf{v}))$$

□

Since V^σ is closed under the application of orbit-mutation maps, by restricting $\mathcal{V}_{\eta^{\sigma,B}}$ to B-coherent linear relations (with coefficients in R) among vectors in V^σ one obtains a partial linear structure $(V^\sigma, R, \mathcal{V}_{\eta^{\sigma,B}})$. Clearly, $V^\sigma \cap R^n = \{\mathbf{a} \in R^n : \sigma(\mathbf{a}) = \mathbf{a}\}$ is also closed under the application of orbit-mutation maps; thus restricting $\mathcal{V}_{\eta^{\sigma,B}}$ to B-coherent linear relations (with coefficients in R) among vectors in $V^\sigma \cap R^n$ yields the partial linear structure $(V^\sigma \cap R^n, R, \mathcal{V}_{\eta^{\sigma,B}})$. We show that restricting the vector space isomorphism $\phi : V^\sigma \to \mathbb{R}^n$ of (2.23) to $V^\sigma \cap R^n$ yields an isomorphism of partial linear structures in the sense of Definition 2.2.25 between $(V^\sigma \cap R^n, R, \mathcal{V}_{\eta^{\sigma,B}})$ and $R^C = (R^m, R, \mathcal{V}_{\eta^C})$.

Proposition 2.4.7. *The map $\phi|_{V^\sigma \cap R^n} : (V^\sigma \cap R^n, R, \mathcal{V}_{\eta^{\sigma,B}}) \to R^C$ is an isomorphism of partial linear structures.*

Proof. Clearly, $\phi|_{V^\sigma \cap R^n}$ is a bijection between $V^\sigma \cap R^n$ and R^m. Thus it suffices to show that both ϕ and ϕ^{-1} (and therefore $\phi|_{V^\sigma \cap R^n}$ and $\phi^{-1}|_{R^m}$) are linear with respect to partial linear structure. Let S be a finite indexing set and consider the linear relation $\sum_{i \in S} r_i \mathbf{v}_i$ among vectors $(\mathbf{v}_i : i \in S) \subseteq V^\sigma$ with coefficients $(r_i : i \in S) \subseteq R$. We show that $\sum_{i \in S} r_i \mathbf{v}_i \in \mathcal{V}_{\eta^{\sigma,B}}$ if and only if $\phi\left(\sum_{i \in S} r_i \mathbf{v}_i\right) \in \mathcal{V}_{\eta^C}$.

Let $\bar{\mathbf{k}}$ be any finite sequence of σ-orbits. Applying Lemma 2.4.6,

$$\phi\left(\sum_{i \in S} r_i \eta^{\sigma,B}_{\bar{\mathbf{k}}}(\mathbf{v}_i)\right) = \sum_{i \in S} r_i \phi\left(\eta^{\sigma,B}_{\bar{\mathbf{k}}}(\mathbf{v}_i)\right) = \sum_{i \in S} r_i \eta^C_{\bar{\mathbf{k}}}\left(\phi(\mathbf{v}_i)\right).$$

Since ϕ is injective, its kernel is trivial. Thus $\sum_{i \in S} r_i \mathbf{v}_i$ satisfies (2.18), the first condition of orbit-B-coherence, if and only if $\sum_{i \in S} r_i \phi(\mathbf{v}_i)$ satisfies (2.6), the first condition of C-coherence.

We verify the equivalence of satisfying (2.19) and (2.7), the second conditions of orbit-B- and C-coherence, respectively, in a similar manner. Note that for any $\mathbf{v} \in V^\sigma$, $\mathbf{min}(\mathbf{v}, \mathbf{0}) \in V^\sigma$, and it is easy to check that the map $\mathbf{v} \mapsto \mathbf{min}(\mathbf{v}, \mathbf{0})$ commutes with ϕ. Applying Lemma 2.4.6,

$$\phi\left(\sum_{i \in S} r_i \mathbf{min}(\eta_{\mathbf{k}}^{\sigma,B}(\mathbf{v}_i), \mathbf{0})\right) = \sum_{i \in S} r_i \mathbf{min}(\phi(\eta_{\mathbf{k}}^{\sigma,B}(\mathbf{v})), \mathbf{0}) = \sum_{i \in S} r_i \mathbf{min}(\eta_{\overline{\mathbf{k}}}^{C}(\phi(\mathbf{v}_i)), \mathbf{0}).$$

Again, the fact that ϕ has a trivial kernel yields the desired result. $\square$

We are now ready to provide a constructive proof that a ***σ-symmetric basis for*** R^B (that is, a basis for R^B which is fixed as a set under the action of σ) folds to a spanning set for R^C. The construction is natural: just as the folded exchange matrix $C = \pi_\sigma(B)$ was built by taking sums of entries of B over σ-orbits, we build our spanning set for R^C by taking sums over σ-orbits of basis vectors for R^B. In particular, recall from Section 2.1, that for each $\mathbf{a} \in \mathbb{R}^n$ the notation $\bar{\mathbf{a}}$ represents the σ-orbit of $\mathbf{a}$, and given a σ-symmetric set $A = (\mathbf{a}_i : i \in I)$ of vectors in $\mathbb{R}^n$, the notation $\bar{A}$ represents the collection of σ-orbits of A. We then reuse π_σ to denote the following map from $\bar{A}$ to the σ-fixed subspace V^σ:

$$\begin{aligned} \pi_\sigma : \bar{A} &\to V^\sigma \\ \bar{\mathbf{a}} &\mapsto \sum_{\mathbf{a}_i \in \bar{\mathbf{a}}} \mathbf{a}_i. \end{aligned} \tag{2.25}$$

Remark 2.4.8. As an alternative to (2.25), the map π_σ could have been defined on the σ-symmetric set A rather than on the collection $\bar{A}$ of its σ-orbits. An advantage of this setup would be the view of the map $\pi_\sigma : \mathbb{R}^n \to V^\sigma$ as orthogonal projection without normalization for vector length. However the choice to define π_σ on the collection of σ-orbits was made with an eye towards the primary goal of this section, to fold a σ-symmetric basis for R^B to a basis for R^C. Had π_σ been defined on the full σ-symmetric basis, there would be no hope of independence since any two vectors in the same σ-orbit would have the same image.

Theorem 2.4.9. *Let $\mathcal{B} = (\mathbf{b}_i : i \in I) \subseteq R^n$ be a σ-symmetric basis for R^B with set of σ-orbits $\bar{\mathcal{B}}$. Then $\phi(\pi_\sigma(\bar{\mathcal{B}}))$ is a spanning set for R^C.*

Proof. Let $\mathbf{a} \in R^m$. Since $\mathcal{B}$ spans R^B, by Definition 2.2.22 there exist a finite subset of indices $S \subseteq I$ and a collection of scalars $(r_i : i \in S) \subseteq R$ such that the following is a B-coherent linear relation (with coefficients in R) among vectors in $\mathcal{B}$:

$$\phi^{-1}(\mathbf{a}) + \sum_{i \in S} (-r_i)\mathbf{b}_i$$

Define $\bar{\mathcal{B}}_S = \{\bar{\mathbf{b}}_i : i \in S\} \subseteq \bar{\mathcal{B}}$, the set of orbits of vectors in $\mathcal{B}$ indexed by S, and for each orbit $\bar{\mathbf{b}} \in \bar{\mathcal{B}}_S$ define $S_{\bar{\mathbf{b}}} = \{i \in I : \bar{\mathbf{b}}_i = \bar{\mathbf{b}}\}$, the subset of I indexing vectors in $\mathcal{B}$ in the σ-orbit $\bar{\mathbf{b}}$. Finally, define $S' = \bigcup_{\bar{\mathbf{b}} \in \bar{\mathcal{B}}_S} S_{\bar{\mathbf{b}}}$. For each index $i \in S'$, define the scalar $t_i \in R$ by $t_i = r_i$ if

$i \in S$, and $t_i = 0$ otherwise. Then $\sum_{i\in S}(-r_i)\mathbf{b}_i = \sum_{i\in S'}(-t_i)\mathbf{b}_i$, and furthermore, this equality is preserved under the application of mutation maps $\eta^B_{\mathbf{k}}$. Thus the following is a B-coherent linear relation (with coefficients in R) as well:

$$\phi^{-1}(\mathbf{a}) + \sum_{i\in S'}(-t_i)\mathbf{b}_i \tag{2.26}$$

Now, we argue that because $\mathcal{B}$ is an independent set for R^B, for each orbit $\bar{\mathbf{b}} \in \bar{\mathcal{B}}_S$, each vector $\mathbf{b}_i \in \bar{\mathbf{b}}$ appears with the same coefficient in (2.26). Consider two distinct indices $i_1, i_2 \in S_{\bar{\mathbf{b}}}$. Since $\mathbf{b}_{i_1}$ and $\mathbf{b}_{i_2}$ are elements of the common orbit $\bar{\mathbf{b}}$, there exists some $q \in \mathbb{N}$ such that $\sigma^q(\mathbf{b}_{i_1}) = \mathbf{b}_{i_2}$. For each $i \in S'$, let $i_{\sigma^q} \in S'$ denote the index of $\sigma^{-q}(\mathbf{b}_i)$: for example, $(i_2)_{\sigma^q} = i_1$. By (2.26), $\sum_{i\in S'}(-t_i)\mathbf{b}_i = -\phi^{-1}(\mathbf{a}) = \phi^{-1}(-\mathbf{a})$. Since $\phi^{-1}(-\mathbf{a}) \in V^\sigma$ is fixed under action by σ^q, so too is $\sum_{i\in S'}(-t_i)\mathbf{b}_i$. Thus,

$$\sum_{i\in S'}(-t_i)\mathbf{b}_i = \sigma^q\left(\sum_{i\in S'}(-t_i)\mathbf{b}_i\right) = \sum_{i\in S'}(-t_i)\sigma^q(\mathbf{b}_i) = \sum_{i\in S'}(-t_{i_{\sigma^q}})\mathbf{b}_i$$

Grouping like terms and isolating $\mathbf{b}_{i_2}$, we obtain the following linear relation:

$$(-t_{i_2} + t_{i_1})\mathbf{b}_{i_2} + \sum_{\substack{i\in S'\\ i\neq i_2}}(-t_i + t_{i_{\sigma^q}})\mathbf{b}_i \tag{2.27}$$

In fact, the relation in (2.27) is B-coherent as it is obtained as the difference between two other B-coherent linear relations: namely, $\phi^{-1}(\mathbf{a}) - \sum_{i\in S'}(-t_i)\mathbf{b}_i$ from (2.26), and its image under action by σ^q, $\phi^{-1}(\mathbf{a}) - \sum_{i\in S'}(-t_{i_{\sigma^q}})\mathbf{b}_i$. (The latter relation is B-coherent by Proposition 2.4.4.) But the only vectors $(\mathbf{b}_i : i \in S')$ appearing with non-zero coefficients in (2.27) are elements of $\mathcal{B}$, a basis for R^B: hence since $\mathcal{B}$ is independent, this B-coherent relation must be trivial. In particular, $t_{i_1} = t_{i_2}$.

Thus for each orbit $\bar{\mathbf{b}} \in \bar{\mathcal{B}}_S$, we may define, without ambiguity, the coefficient $t_{\bar{\mathbf{b}}} = t_i$ for each $i \in S_{\bar{\mathbf{b}}}$, and rewrite the B-coherent linear relation given in (2.26) as follows:

$$\phi_\sigma^{-1}(\mathbf{a}) + \sum_{\bar{\mathbf{b}}\in\bar{\mathcal{B}}_S}\sum_{i\in S_{\bar{\mathbf{b}}}}(-t_{\bar{\mathbf{b}}})\mathbf{b}_i.$$

Rewriting once more using the definition of π_σ from (2.25) yields

$$\phi_\sigma^{-1}(\mathbf{a}) + \sum_{\bar{\mathbf{b}}\in\bar{\mathcal{B}}_S}(-t_{\bar{\mathbf{b}}})\pi_\sigma(\bar{\mathbf{b}}).$$

Finally, applying ϕ gives the relation below among vectors in $\phi(\pi_\sigma(\bar{\mathcal{B}}))$, and this relation is

C-coherent (with coefficients in R) by Proposition 2.4.7:

$$\mathbf{a} + \sum_{\bar{\mathbf{b}} \in \bar{\mathcal{B}}_S} (-t_{\bar{\mathbf{b}}}) \phi(\pi_\sigma(\bar{\mathbf{b}})).$$

We thereby conclude by Definition 2.2.22 that $\phi(\pi_\sigma(\bar{\mathcal{B}}))$ spans R^C. □

The following corollary, on the folding of a positive basis for R^B (see Definition 2.2.22) is immediate from the proof of Theorem 2.4.9:

Corollary 2.4.10. *If $\mathcal{B} = (\mathbf{b}_i : i \in I) \subseteq R^n$ is a σ-symmetric positive basis for R^B, then $\phi(\pi_\sigma(\bar{\mathcal{B}}))$ is a positive spanning set for R^C.*

Recall from Theorem 2.2.23 that for any exchange matrix B and underlying ring $R \in \{\mathbb{Q}, \mathbb{R}\}$, there exists a basis for R^B. It is natural, in light of Theorem 2.4.9, to ask if a *σ-symmetric* basis for R^B also exists. While we cannot prove existence in general using the non-constructive argument from the proof of Theorem 2.2.23, we *can* show existence (for any choice of R, including $\mathbb{Z}$) so long as we have a positive basis for R^B.

Proposition 2.4.11. *If $\mathcal{B} = (\mathbf{b}_i : i \in I)$ is a positive basis for R^B, then there exists a σ-symmetric positive basis $\mathcal{B}^\sigma$ for R^B:*

Proof. Proposition 2.2.40 implies that $\mathcal{B}$ must consist of a nonzero vector in each ray in $\mathcal{F}_B \cap R^n$. By Proposition 2.2.32, $\sigma(\mathcal{F}_B) = \mathcal{F}_B$, and therefore $\sigma(\mathcal{F}_B \cap R^n) = \mathcal{F}_B \cap R^n$. In particular, the set of rays in $\mathcal{F}_B \cap R^n$ is fixed under action by σ. If $R = \mathbb{Z}$, then $\mathcal{B}$ is unique and consists of the smallest nonzero integer vector in each ray: thus $\mathcal{B}$ must be σ-symmetric itself, and we take $\mathcal{B}^\sigma = \mathcal{B}$. Otherwise, there exist nonzero scalars $t_i \in R$ such that the collection $\mathcal{B}^\sigma = (t_i \mathbf{b}_i : i \in I)$ is σ-symmetric, and since $\mathcal{B}^\sigma$ consists of a single nonzero vector in each ray, by Proposition 2.2.40 it forms a positive basis for R^B. □

Another natural question arising from Theorem 2.4.9 is whether, given a σ-symmetric basis $\mathcal{B}$ for R^B, the set $\phi(\pi_\sigma(\bar{\mathcal{B}}))$ is independent in R^C and therefore a basis. We suspect that the answer in general is no, and that a better candidate for a basis for R^C is the following subset of $\phi(\pi_\sigma(\bar{\mathcal{B}}))$:

$$\left\{ \phi(\pi_\sigma(\bar{\mathbf{b}})) : \bar{\mathbf{b}} \in \bar{\mathcal{B}} \text{ is contained in a } B\text{-cone} \right\}. \tag{2.28}$$

The first justification for this conjecture is provided below in Corollary 2.4.12, which shows that this set is a spanning set for R^C if $\mathcal{B}$ is a σ-symmetric *cone* basis for R^B. As mentioned in [21, Remark 6.5], for every exchange matrix B for which a basis for R^B has been constructed, the basis is a cone basis. An even stronger justification is provided in the next section, where we prove in Theorem 2.4.23 that this set is a positive basis for R^C if and only if the R-part of

$\mathcal{F}_B$ folds to the R-part of $\mathcal{F}_C$. So while we cannot show, in general, that a σ-symmetric basis folds to a basis, we *can* prove positive basis folding as a consequence of mutation fan folding. This will provide valuable consequences when B is of finite type and when B arises as the signed-adjacency matrix of a triangulated orbifold, as we show that the mutation fan folds in both cases.

Corollary 2.4.12. *Let $\mathcal{B} = (\mathbf{b}_i : i \in I) \subseteq R^n$ be a σ-symmetric cone basis for R^B with set of σ-orbits $\bar{\mathcal{B}}$. Then the subset of $\phi(\pi_\sigma(\bar{\mathcal{B}}))$ defined in* (2.28) *is a spanning set for R^C.*

Proof. Let $\mathbf{a} \in R^m$. By the proof of Theorem 2.4.9, there exists a finite subset $S' = \cup_{\bar{\mathbf{b}} \in \bar{\mathcal{B}}_S} S_{\bar{\mathbf{b}}}$ of I such that the following relation among vectors $(\mathbf{b}_i : i \in S')$ is B-coherent:

$$\phi_\sigma^{-1}(\mathbf{a}) + \sum_{\bar{\mathbf{b}} \in \bar{\mathcal{B}}_S} (-t_{\bar{\mathbf{b}}}) \pi_\sigma(\bar{\mathbf{b}}).$$

Furthermore, since $\mathcal{B}$ is a basis for R^B and therefore independent, this relation is unique: that is, it is the *only* B-coherent relation (with coefficients in R) of the form $\phi^{-1}(\mathbf{a}) + \sum_{i \in S}(-r_i)\mathbf{b}_i$ for some finite subset $S \subseteq I$. It therefore follows from Proposition 2.2.35 that the basis vectors indexed by S' must all be contained in some common B-cone. In particular, each σ-orbit $\bar{\mathbf{b}} \in \mathcal{B}_S$ is contained in a B-cone, so $\phi(\mathcal{B}_S)$ is contained in the subset of $\phi(\pi_\sigma(\bar{\mathcal{B}}))$ defined in (2.28). □

2.4.2 Mutation fan folding

The first result of this section establishes a geometric analog of the mutation-linear basis folding of Theorem 2.4.9 by showing that (the R-part of) the mutation fan $\mathcal{F}_B$ folds to a refinement of (the R-part of) the mutation fan $\mathcal{F}_C$. The second result shows that the folded (R-part of) $\mathcal{F}_B$ coincides with (the R-part of) $\mathcal{F}_C$ if and only if a σ-symmetric positive basis for $\mathbb{R}^B$ (resp. R^B) folds, in the sense of (2.28), to a positive basis for $\mathbb{R}^C$ (resp. R^C). Finally, we prove that the two fans coincide when B is of finite type, and thereby deduce various associated results about the folding of mutation-linear bases and universal geometric coefficients.

The general case

We start by defining an orbit-mutation version $\equiv^{\sigma,B}$ of the equivalence relation $\equiv^B$ on $\mathbb{R}^B$ from Definition 2.2.29. Namely, set $\mathbf{a}_1 \equiv^{\sigma,B} \mathbf{a}_2$ if and only if $\mathbf{sgn}(\eta_{\bar{\mathbf{k}}}^{\sigma,B}(\mathbf{a}_1)) = \mathbf{sgn}(\eta_{\bar{\mathbf{k}}}^{\sigma,B}(\mathbf{a}_2))$ for every finite sequence $\mathbf{k}$ of indices in $[n]$. The following observation is immediate from the fact that every orbit-mutation map $\eta_{\bar{\mathbf{k}}}^{\sigma,B}$ is in particular a mutation map $\eta_{\mathbf{k}}^B$ (see Remark 2.4.3):

Lemma 2.4.13. *Let $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}^n$. If $\mathbf{a_1} \equiv^B \mathbf{a_2}$, then $\mathbf{a}_1 \equiv^{\sigma,B} \mathbf{a}_2$.*

It is likewise immediate from (2.23) that for any $\mathbf{v}_1, \mathbf{v}_2 \in V^\sigma$, $\mathbf{sgn}(\mathbf{v}_1) = \mathbf{sgn}(\mathbf{v}_2)$ if and only if $\mathbf{sgn}(\phi(\mathbf{v}_1)) = \mathbf{sgn}(\phi(\mathbf{v}_2))$. Thus, we may also make the following observation:

Lemma 2.4.14. *Let* $\mathbf{v}_1, \mathbf{v}_2 \in V^\sigma$*. Then* $\mathbf{v}_1 \equiv^{\sigma,B} \mathbf{v}_2$ *if and only if* $\phi(\mathbf{v}_1) \equiv^C \phi(\mathbf{v}_2)$*.*

Theorem 2.4.15. *The fan* $\phi(\mathcal{F}_B \cap V^\sigma)$ *refines* $\mathcal{F}_C$*. Further, the R-part of* $\phi(\mathcal{F}_B \cap V^\sigma)$ *refines the R-part of* $\mathcal{F}_C$*.*

Proof. It suffices to prove the first part of the theorem: the second part then follows easily by application of Definition 2.2.37.

Observe that since $\mathcal{F}_B$ is a complete fan in $\mathbb{R}^n$ and $\phi : V^\sigma \to \mathbb{R}^m$ is a vector space isomorphism (see (2.23)), $\phi(\mathcal{F}_B \cap V^\sigma)$ is complete in $\mathbb{R}^m$, and we know that the mutation fan $\mathcal{F}_C$ is complete. Thus by Definition 2.2.28, we must show that each cone in $\phi(\mathcal{F}_B \cap V^\sigma)$ is contained in a cone in $\mathcal{F}_C$. It suffices to show such containment for each maximal cone F' of $\phi(\mathcal{F}_B \cap V^\sigma)$, which, by construction of $\mathcal{F}_B$ in Definition 2.2.29, is of the form $F' = \phi(F \cap V^\sigma)$ for some B-cone F. By definition, F is the closure of some B-class $F_{\mathbf{a}} = \{\mathbf{v} \in \mathbb{R}^n : \mathbf{v} \equiv^B \mathbf{a}\}$. Then,

$$F_{\mathbf{a}} \cap V^\sigma = \{\mathbf{v} \in V^\sigma : \mathbf{v} \equiv^B \mathbf{a}\} \subseteq \{\mathbf{v} \in V^\sigma : \mathbf{v} \equiv^{\sigma,B} \mathbf{a}\} \tag{2.29}$$

where the set inclusion is a consequence of Lemma 2.4.13. Applying $\phi : V^\sigma \to \mathbb{R}^m$ to (2.29) and appealing to Lemma 2.4.14,

$$\begin{aligned}
\phi(F_{\mathbf{a}} \cap V^\sigma) &\subseteq \phi\left(\{\mathbf{v} \in V^\sigma : \mathbf{v} \equiv^{\sigma,B} \mathbf{a}\}\right) \\
&= \{\phi(\mathbf{v}) : \mathbf{v} \in V^\sigma \text{ with } \mathbf{v} \equiv^{\sigma,B} \mathbf{a}\} \\
&= \{\phi(\mathbf{v}) : \mathbf{v} \in V^\sigma \text{ with } \phi(\mathbf{v}) \equiv^C \phi(\mathbf{a})\} \\
&= \{\phi(\mathbf{v}) \in \mathbb{R}^m : \phi(\mathbf{v}) \equiv^C \phi(\mathbf{a})\}
\end{aligned}$$

Taking closures yields the desired result: letting G denote the C-cone obtained by taking the closure of the C-class containing $\phi(\mathbf{a})$, a C-cone, we have $F' = \phi(F \cap V^\sigma) \subseteq G \in \mathcal{F}_C$. ☐

For the remainder of this section, we will work with simplicial fans, so qw take a moment to present a few additional definitions and results from convex geometry to complement those introduced in Definitions 2.2.26 and 2.2.28. Proofs of the first two facts below are found, for example, in [26, Corollary 18.1.3] and [26, Theorem 18.2].

Definition 2.4.16 (Relative interior)**.** The ***affine hull*** of a convex set F is the union of all lines defined by two distinct points of F. The ***relative interior*** of F, written $\text{relint}(F)$, is the interior of F as a subset of its affine hull, and the ***relative boundary*** of F, written $\text{relbd}(F)$, is the set difference between the closure and relative interior of F.

Lemma 2.4.17. *Any proper face of a convex set in is contained in its relative boundary.*

Lemma 2.4.18. *Let* F *be a non-empty convex set. The collection of all relative interiors of non-empty faces of* F *form a partition of* F*.*

Lemma 2.4.19. *The relative boundary of a closed convex set is the union of its proper faces.*

Proof. Let U denote the union of the proper faces of a convex set F. By Lemma 2.4.17, U is contained in the relative boundary of F. For the reverse containment, let $\mathbf{x} \in \operatorname{relbd}(F)$. Then since F is closed, $\mathbf{x} \in F$, so by Lemma 2.4.18, $\mathbf{x} \in \operatorname{relint}(T)$ for some face T of F. Since $\mathbf{x} \notin \operatorname{relint}(F)$, T must be a proper face, and hence $\mathbf{x} \in U$. □

Recall from Definition 2.2.28 that a complete fan $\mathcal{F}$ refines another complete fan $\mathcal{G}$ if and only if every cone of $\mathcal{F}$ is contained in a cone of $\mathcal{G}$. In particular, this implies that every ray of $\mathcal{F}$ is a ray of $\mathcal{G}$. Recall as well from Definition 2.2.26 that if $\mathcal{F}$ is a simplicial fan, every cone of $\mathcal{F}$ is the convex hull of finitely many rays spanned by linearly independent vectors.

Proposition 2.4.20. *Let $\mathcal{F}$ and $\mathcal{G}$ be complete fans in some module V with $\mathcal{F}$ refining $\mathcal{G}$. If $\mathcal{G}$ is simplicial and every ray in $\mathcal{F}$ is a ray in $\mathcal{G}$, then $\mathcal{F} = \mathcal{G}$.*

Proof. It suffices to show that $\mathcal{G}$ is a subfan of $\mathcal{F}$ by showing every cone in $\mathcal{G}$ is a cone in $\mathcal{F}$. Let $G \in \mathcal{G}$ and $\mathbf{x} \in \operatorname{relint}(G)$. Because $\mathcal{F}$ refines $\mathcal{G}$, there exists some cone $F \in \mathcal{F}$ such that $\mathbf{x} \in F \subseteq G$. The cone F is the convex hull of some (a priori, possibly infinite) subset S of rays in V. Thus each ray $\mathbf{r} \in S$ is a face of F and therefore a ray in $\mathcal{F}$, implying, since $\mathcal{F}$ refines $\mathcal{G}$, that $\mathbf{r} \in \mathcal{G}$. But $F \subseteq G$ implies $\mathbf{r} \subseteq G$, so $\mathbf{r} = \mathbf{r} \cap G$ is in fact a 1-dimensional face, or ray, in G. Thus F is the convex hull of some set of rays in G, implying that F is a face of G since G is simplicial. But $\mathbf{x} \in F \cap \operatorname{relint}(G)$, so by Lemma 2.4.19, F cannot be a proper face. Thus $G = F$ is a cone in $\mathcal{F}$. □

Thus equipped with tools from convex geometry, we prepare to characterize the rays in the R-part of the folded mutation fan $\phi(F_B \cap V^\sigma)$, assuming the R-part of $\mathcal{F}_B$ is simplicial. This implies in turn, that the R-part of $\mathcal{F}_B \cap V^\sigma$, which coincides with the intersection of the R-part of $\mathcal{F}_B$ with V^σ, is simplicial, as is the R-part of $\phi(F_B \cap V^\sigma)$. Recall that $\sigma(B) = B$, so by Proposition 2.2.32 (resp. Corollary 2.2.41), the set of rays $\mathcal{R}$ in the R-part of $\mathcal{F}_B$ is fixed under the action of σ on $\mathbb{R}^n$. Therefore for each ray $\mathbf{r} \in \mathcal{R}$, we may choose a non-zero vector $\mathbf{a_r}$ in $\mathbf{r}$ such that the set $A_\mathcal{R} = \{\mathbf{a_r} : \mathbf{r} \in \mathcal{R}\}$ is a σ-symmetric subset of $\mathbb{R}^n$ (resp. R^n). Denote by $\bar{\mathbf{a}}_\mathbf{r}$ and $\bar{A}_\mathcal{R}$ the σ-orbit of each vector $\mathbf{a_r} \in A_\mathcal{R}$ and the set of all such orbits, respectively.

Lemma 2.4.21. *Suppose that the R-part of $\mathcal{F}_B$ is simplicial and let $A_\mathcal{R}$ be a σ-symmetric collection of non-zero vectors in the rays $\mathcal{R}$ of the R-part of $\mathcal{F}_B$. For each σ-orbit $\bar{\mathbf{a}} \in \bar{A}_\mathcal{R}$, define $\mathbf{s}_{\bar{\mathbf{a}}}$ to be the non-negative R-linear span of $\pi_\sigma(\bar{\mathbf{a}}) = \sum_{\mathbf{a}_i \in \bar{\mathbf{a}}} \mathbf{a}_i$. Then the rays $\mathcal{S}$ of the R-part of the fan $\mathcal{F}_B \cap V^\sigma$ are given by*

$$\mathcal{S} = \{\mathbf{s}_{\bar{\mathbf{a}}} : \bar{\mathbf{a}} \in \bar{A}_\mathcal{R} \text{ spans a cone in the } R\text{-part of } \mathcal{F}_B\}. \tag{2.30}$$

Thus the rays of the R-part of $\phi(\mathcal{F}_B \cap V^\sigma)$ are given by $\phi(\mathcal{S})$.

Proof. Let $\mathbf{s}$ be a ray of the R-part of $\mathcal{F}_B \cap V^\sigma$. Since by Lemma 2.2.38 the R-part of $\mathcal{F}_B \cap V^\sigma$ is given by the intersection of the R-part of $\mathcal{F}_B$ with V^σ, $\mathbf{s}$ is in particular the one-dimensional intersection of a cone in the R-part of $\mathcal{F}_B$ with V^σ. Choose F to be the minimal such cone (under containment): that is, such that there does not exist a proper face F' of F with $\mathbf{s} = F' \cap V^\sigma$. By the minimality of F, $\mathbf{s} \subseteq \operatorname{relint}(F)$. Since $\mathbf{s} \in V^\sigma$, $\mathbf{s} = \sigma(\mathbf{s}) \subseteq \operatorname{relint}(\sigma(F))$, so Lemma 2.4.18 implies that $F = \sigma(F)$. In particular, $\mathcal{R}_F = \sigma(\mathcal{R}_F)$, where $\mathcal{R}_F \subseteq \mathcal{R}$ denotes the rays of F. Let $A_{\mathcal{R}_F} \subseteq A_{\mathcal{R}}$ denote the associated σ-symmetric collection of spanning vectors. We claim by simpliciality that $A_{\mathcal{R}_F}$ consists of a single σ-orbit $\bar{\mathbf{a}}$, as otherwise either the dimension of $\mathbf{s} = F \cap V^\sigma$ must be greater than one or F must not be minimal, both of which are contradictions. Thus since $\mathbf{s} \in \operatorname{relint}(F)$ and $\bar{\mathbf{a}} = A_{\mathcal{R}_F}$ spans F, we must have $\mathbf{s} = \mathbf{s}_{\bar{\mathbf{a}}}$, the non-negative linear span of $\pi_\sigma(\bar{\mathbf{a}}) = \sum_{\mathbf{a}_i \in \bar{\mathbf{a}}} \mathbf{a}_i$.

For the reverse inclusion, suppose $\bar{\mathbf{a}} \in \bar{A}_{\mathcal{R}}$ is a σ-orbit of vectors in $A_{\mathcal{R}}$ which spans a cone F in the R-part of $\mathcal{F}_B$. Since $\mathbf{s}_{\bar{\mathbf{a}}}$ denotes the non-negative R-linear span of $\pi_\sigma(\bar{\mathbf{a}}) = \sum_{\mathbf{a}_i \in \bar{\mathbf{a}}} \mathbf{a}_i$, clearly $\mathbf{s}_{\bar{\mathbf{a}}} \subseteq F \cap V^\sigma$. We show that $F \cap V^\sigma \subseteq \mathbf{s}_{\bar{\mathbf{a}}}$, and thereby conclude that $\mathbf{s}_{\bar{\mathbf{a}}} = F \cap V^\sigma$ is a ray in the R-part of $\mathcal{F}_B \cap V^\sigma$. Suppose $\mathbf{v} \in F \cap V_\sigma$. Then $\mathbf{v} = \sum_{\mathbf{a}_i \in \bar{\mathbf{a}}} v_i \mathbf{a}_i$ for some scalars $v_i \in R$. Since $\mathbf{v} \in V^\sigma$ and $\bar{\mathbf{a}}$ is a σ-orbit, all the v_i must be equal: thus $\mathbf{v} \in \mathbf{s}_{\bar{\mathbf{a}}}$. □

Now, suppose a positive basis for R^B exists. Then by Definition 2.2.39, the R-part of $\mathcal{F}_B$ is given by $\mathcal{F}_B^R$, a simplicial fan. Further, by Proposition 2.4.11, there exists a σ-symmetric positive basis for R^B, which we denote by $\mathcal{B}$. By Proposition 2.2.36, $\mathcal{B}$ is in particular a positive cone basis. Thus the subset of $\phi(\pi_\sigma(\bar{\mathcal{B}}))$ given in (2.28) consisting only of the foldings of σ-orbits $\bar{\mathbf{b}} \in \bar{\mathcal{B}}$ which are contained in B-cones is a spanning set for R^C by Corollary 2.4.12. As the last necessary preliminary result before we show that positive bases fold if and only if the mutation fan folds, we give an equivalent formulation of this set.

Lemma 2.4.22. *Suppose that $\mathcal{B} = (\mathbf{b}_i : i \in I)$ is a σ-symmetric positive basis for R^B with set of σ-orbits $\bar{\mathcal{B}}$. Then the vectors in a given σ-orbit $\bar{\mathbf{b}} \in \bar{\mathcal{B}}$ are contained in a common B-cone if and only if they span a cone in $\mathcal{F}_R^B$.*

Proof. Clearly, if the vectors in a σ-orbit $\bar{\mathbf{b}} \in \bar{\mathcal{B}}$ span a cone in the R-part of $\mathcal{F}_B$, then they are contained in a cone in $\mathcal{F}_B$, and therefore a B-cone.

Conversely, suppose all the vectors in a given σ-orbit $\bar{\mathbf{b}} \in \bar{\mathcal{B}}$ are contained in some common B-cone F. Recall from Proposition 2.2.40 that each vector in $\bar{\mathbf{b}}$ spans a unique ray in $\mathcal{F}_B^R$. By Definition 2.2.37, there exists a unique largest cone F' in $\mathcal{F}_B^R$ which is contained in F and contains $F \cap R^n$. Thus the vectors in $\bar{\mathbf{b}}$ are contained in F', and therefore span rays of F' by Definition 2.2.26. But $\mathcal{F}_B^R$ is simplicial, so $\bar{\mathbf{b}}$ spans a face of F', and therefore a cone in $\mathcal{F}_B^R$. □

Theorem 2.4.23. *Suppose that $\mathcal{B} = (\mathbf{b}_i : i \in I)$ is a σ-symmetric positive basis for R^B and let $\bar{\mathcal{B}}_{\mathcal{F}} \subseteq \bar{\mathcal{B}}$ denote the subset of σ-orbits of $\mathcal{B}$ which span cones in the R-part of $\mathcal{F}_B$. Then*

$\phi(\pi_\sigma(\bar{\mathcal{B}}_\mathcal{F}))$ is a positive basis for R^C if and only if $\phi(\mathcal{F}_B^R \cap V^\sigma) = \mathcal{F}_C^R$.

Proof. Suppose first that $\phi(\mathcal{F}_B^R \cap V^\sigma) = \mathcal{F}_C^R$. By Proposition 2.2.40, $\mathcal{B}$ consists of a non-zero vector in each ray of $\mathcal{F}_B^R$. It is easily checked by applying Definition 2.2.37 that since $\mathcal{F}_B^R$ is the R-part of $\mathcal{F}_B$, the fan $\phi(\mathcal{F}_B^R \cap V^\sigma)$ is the R-part of $\phi(\mathcal{F}_B \cap V^\sigma)$. Both are simplicial since a positive basis for R^B exists. Thus by Lemma 2.4.21, the set $\phi(\pi_\sigma(\bar{\mathcal{B}}_\mathcal{F}))$ consists of one non-zero vector in each ray of $\phi(\mathcal{F}_B^R \cap V^\sigma)$. Once more invoking Proposition 2.2.40, it follows that $\phi(\pi_\sigma(\bar{\mathcal{B}}_\mathcal{F}))$ is a positive basis for R^C.

Conversely, suppose that $\phi(\pi_\sigma(\bar{\mathcal{B}}_\mathcal{F}))$ is a positive basis for R^C. By Definition 2.2.39, $\mathcal{F}_B^R$ is simplicial, implying $\phi(\mathcal{F}_B^R \cap V^\sigma)$ is simplicial, and since a positive basis exists for R^C, $\mathcal{F}_C^R$ is simplicial as well. Both $\phi(\mathcal{F}_B^R \cap V^\sigma)$ and $\mathcal{F}_C^R$ are complete fans in with $\phi(\mathcal{F}_B^R \cap V^\sigma)$ refining $\mathcal{F}_C^R$ as a consequence of Theorem 2.4.15. Thus by Proposition 2.4.20 to show that every ray in $\phi(\mathcal{F}_B^R \cap V^\sigma)$ is a ray in $\mathcal{F}_C^R$. By Lemma 2.4.21, the rays in $\phi(\mathcal{F}_B^R \cap V^\sigma)$ are spanned by the vectors in $\phi(\pi_\sigma(\bar{\mathcal{B}}_\mathcal{F}))$. But these vectors precisely span the rays in $\mathcal{F}_C^R$ by Proposition 2.2.40 since $\phi(\pi_\sigma(\bar{\mathcal{B}}_\mathcal{F}))$ is a positive basis for $\mathbb{R}^C$. □

Finite type

We now restrict our attention to the case when B is of finite type. By Proposition 2.3.9, this implies that $C = \pi_\sigma(B)$ is of finite type as well. Proposition 2.2.46 gives us a powerful tool: in this case, the mutation fans $\mathcal{F}_B$ and $\mathcal{F}_C$ coincide with the $\mathbf{g}$-vector fans for B^T and C^T, respectively, rational simplicial fans all of whose maximal cones are full-dimensional and obtained via mutation maps from the respective non-negative orthants $O_n = (\mathbb{R}_{\geq 0})^n$ and $O_m = (\mathbb{R}_{\geq 0})^m$, themselves maximal cones in $\mathcal{F}_B$ and $\mathcal{F}_C$ (see Definition 2.2.45). There is no need to consider R-parts: for any choice of $R \in \{\mathbb{Z}, \mathbb{Q}, \mathbb{R}\}$, the "$R$-part" of the mutation fan for a finite type exchange matrix is simply the mutation fan itself.

Theorem 2.4.24. *If B is of finite type, then $\phi(\mathcal{F}_B \cap V^\sigma) = \mathcal{F}_C$.*

Proof. By Theorem 2.4.15, we know that $\phi(\mathcal{F}_B \cap V^\sigma)$ refines $\mathcal{F}_C$. Furthermore, since B is of finite type, Proposition 2.2.46 implies that $\mathcal{F}_B$ is simplicial, so by Proposition 2.4.20, it suffices to show that every ray in $\phi(\mathcal{F}_B \cap V^\sigma)$ is a ray in $\mathcal{F}_C$. In particular, we show that every ray in $\mathcal{F}_B \cap V^\sigma$ is a ray in $\phi^{-1}(\mathcal{F}_C)$: since ϕ is a vector space isomorphism, this is equivalent.

Suppose by way of contradiction that $\mathbf{r}$ is a ray in $\mathcal{F}_B \cap V^\sigma$ but is not a ray in $\phi^{-1}(\mathcal{F}_C)$. Then $\mathbf{r}$ is contained in, but not a face of, some cone in $\phi^{-1}(\mathcal{F}_C)$, implying that $\phi(\mathbf{r})$ is contained in, but not a face of, some cone in $\mathcal{F}_C$. By Proposition 2.3.9, C is of finite type, and hence by Proposition!2.2.46, $\mathcal{F}_C$ coincides with the complete, simplicial, $\mathbf{g}$-vector fan for C^T. So $\phi(\mathbf{r})$ is in fact contained in, but not a face of, some *full-dimensional* cone G in $\mathcal{F}_C$, and by Definition 2.2.45, there exists a sequence $\bar{\mathbf{k}}$ of σ-orbits such that $\eta_{\bar{\mathbf{k}}}^C(G) = O_m$, the non-negative orthant in $\mathbb{R}^m$. Thus $\eta_{\bar{\mathbf{k}}}^C(\phi(\mathbf{r}))$ is contained in, but not a ray of, O_m.

By (2.24), $\phi^{-1}(\eta_{\bar{\mathbf{k}}}^{C}(\phi(\mathbf{r}))) = \eta_{\bar{\mathbf{k}}}^{\sigma,B}(\mathbf{r})$, and clearly, $\phi^{-1}(O_m) = O_n \cap V^{\sigma}$. Hence $\eta_{\bar{\mathbf{k}}}^{\sigma,B}(\mathbf{r})$ is contained in, but not a ray of, $O_n \cap V^{\sigma}$. On the other hand, since $\mathbf{r}$ is a ray in $\mathcal{F}_B \cap V_{\sigma}$, by Proposition 2.2.33 $\eta_{\bar{\mathbf{k}}}^{\sigma,B}(\mathbf{r})$ must be a ray in $\mathcal{F}_{\mu_{\mathbf{k}}(B)} \cap V^{\sigma}$. But O_n is a cone in $\mathcal{F}_{\mu_{\mathbf{k}}(B)}$ by Remark 2.2.42, implying that $O_n \cap V^{\sigma}$ is a cone in $\mathcal{F}_{\mu_{\mathbf{k}}(B)} \cap V^{\sigma}$. This is a contradiction: by Definition 2.2.26, since both $\eta_{\bar{\mathbf{k}}}^{\sigma,B}(\mathbf{r})$ and $O_n \cap V^{\sigma}$ are cones in $\mathcal{F}_{\mu_{\mathbf{k}}(B)} \cap V^{\sigma}$, their intersection $\eta_{\bar{\mathbf{k}}}^{\sigma,B}(\mathbf{r}) \cap (O_n \cap V^{\sigma}) = \eta_{\bar{\mathbf{k}}}^{\sigma,B}(\mathbf{r})$ must be a face of both. □

The following results are immediate from Theorems 2.4.23 and 2.2.43, respectively.

Corollary 2.4.25. *If B is of finite type and $\mathcal{B} = (\mathbf{b}_i : i \in I)$ is a positive basis for R^B, then there exists a σ-symmetric positive basis $\mathcal{B}^{\sigma}$ for R^B, and the following collection is a positive basis for R^C, where $\bar{\mathcal{B}}^{\sigma}$ denotes the set of σ-orbits $\bar{\mathbf{a}}$ of $\mathcal{B}^{\sigma}$:*

$$\big(\phi\left(\pi_{\sigma}(\bar{\mathbf{a}})\right) : \bar{\mathbf{a}} \in \bar{\mathcal{B}}^{\sigma} \text{ spans a cone in } \mathcal{F}_B\big) \tag{2.31}$$

Corollary 2.4.26. *If B is of finite type and $\mathcal{B} = (\mathbf{b}_i : i \in I)$ is a positive basis for R^B, then the set in* (2.31) *constitutes universal geometric coefficients for C over R.*

Chapter 3

Marked surfaces, orbifolds, and folding

3.1 Introduction

An orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ consists of an oriented surface $\mathbf{S}$ and two disjoint, finite sets of points in $\mathbf{S}$: marked points $\mathbf{M}$ and orbifold points $\mathbf{Q}$. If $\mathbf{Q} = \emptyset$, then $\mathcal{O}$ is a marked surface. (There are some additional technical restrictions on $\mathcal{O}$: see Definition 3.2.1.) We consider two special classes of curves in $\mathcal{O}$, called (tagged) arcs (see Definition 3.2.6), and (allowable) curves (see Definition 3.2.10). A maximal compatible collection of (tagged) arcs, always of the same cardinality n, is called a (tagged) triangulation, and a compatible collection of curves with positive weights is called a (quasi-)lamination. The interaction between a given lamination L and the arcs in a given triangulation T is encoded by a shear coordinate vector $\mathbf{b}(T, L) \in \mathbb{Q}^n$. Together, these components define a cluster algebra of geometric type (see Section 2.2.1), where T determines an $n \times n$ exchange matrix $B(T)$, arcs in T correspond to cluster variables, arc "flips" defining new triangulations of $\mathcal{O}$ (see Definition 3.2.5) correspond to matrix mutation, and coefficients rows are given by shear coordinates of laminations.

The goals of this chapter are two-fold. First, we consider mutation fans, mutation-linear bases, and universal geometric coefficients for cluster algebras arising from orbifolds, generalizing many constructions and results for marked surfaces from [22]. In particular, in Sections 3.2 and 3.3, we use [5, 7–9] to describe cluster algebras arising from a triangulated orbifold $\mathcal{O}, T$ in detail, and then extend the notions of rational quasi-laminations on a surface and the rational quasi-lamination fan $\mathcal{F}_{\mathbb{Q}}(T)$ from [22] to the orbifold setting (see Definitions 3.2.10 and 3.3.19). In Section 3.4, we generalize two properties of surfaces, the Curve Separation Property (see Definition 3.4.1) and Null Tangle Property (see Definition 3.4.9), to orbifolds, and prove the following results, which appear in the text as Theorems 3.4.5 and 3.4.22, respectively.

Theorem 3.1.1. *If an orbifold $\mathcal{O}$ has either no punctures and no orbifold points, a unique puncture and no orbifold points, or no punctures and a unique orbifold point, then $\mathcal{O}$ has the Curve Separation Property.*

Theorem 3.1.2. *If $\mathcal{O}$ is a sphere with b boundary components, p punctures, and q orbifold points with $b + p + q \leq 3$, then $\mathcal{O}$ has the Null Tangle Property.*

The next three results are generalizations of [21, Theorems 4.10, 4.12, and 7.3], respectively, from marked surfaces to orbifolds, and their proofs are nearly identical. These appear in the text as Theorems 3.4.3, 3.4.4, and 3.4.17, respectively. One consequence of these results is a geometric realization of the universal geometric coefficients of classical affine type conjectured in [21, Conjecture 10.15].

Theorem 3.1.3. *Suppose T is a tagged triangulation of $\mathcal{O}$, with all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture. The rational quasi-lamination fan $\mathcal{F}_{\mathbb{Q}}(T)$ is a rational, simplicial fan. $\mathcal{F}_{\mathbb{Q}}(T)$ is the rational part of the mutation fan $\mathcal{F}_{B(T)}$ if and only if $\mathcal{O}$ has the Curve Separation Property.*

Theorem 3.1.4. *Suppose T is a tagged triangulation of $\mathcal{O}$, with all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture. Let R be either $\mathbb{Z}$ or $\mathbb{Q}$. If $\mathcal{O}$ has the Curve Separation Property,*

1. *The shear coordinates of allowable curves in $\mathcal{O}$ form a positive basis for $R^{B(T)}$ so long as they are an independent set for $R^{B(T)}$.*
2. *The shear coordinates of allowable curves in $\mathcal{O}$ form a positive basis for $R^{B(T)}$ so long as a positive basis for $R^{B(T)}$ exists.*

Theorem 3.1.5. *Suppose T is a tagged triangulation of $\mathcal{O}$, with all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture. Let R be either $\mathbb{Z}$ or $\mathbb{Q}$. If $\mathcal{O}$ is not a null orbifold, then the following are equivalent:*

1. *$\mathcal{O}$ has the Null Tangle Property.*
2. *The shear coordinates of allowable curves form a basis for $R^{B(T)}$.*
3. *The shear coordinates of allowable curves form a positive basis for $R^{B(T)}$.*
4. *The shear coordinates of allowable curves form universal geometric coefficients for $B(T)$ over R.*

If $\mathcal{O}$ is a null orbifold, then it fails Conclusion 1 and satisfies Conclusions 2, 3, and 4.

The second part of this chapter is dedicated to using the combinatorics of the orbifolds model to prove that the mutation-linear notions established above fold in the sense of Chapter 2. Section 3.5 follows [5] in describing the folding of orbifolds and its correspondence with exchange matrix folding (see Definitions 3.5.1, 3.5.2, and 3.5.3). The desired results, stated below, are then proven in Section 3.6. Our primary result appears in the text as Theorem 3.6.1.

Theorem 3.1.6. *Let $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ be an orbifold with tagged triangulation T which folds, under the symmetry σ, to orbifold $\pi_\sigma(\mathcal{O})$ with tagged triangulation $\pi_\sigma(T)$. Then $\mathcal{F}_{\mathbb{Q}}(T)$ folds to $\mathcal{F}_{\mathbb{Q}}(\pi_\sigma(T))$ under σ.*

As consequences of Theorems 3.1.4 and 3.1.5, respectively, Theorem 3.1.6 then gives us the following two results, which appear in the text as Corollaries 3.6.2 and 3.6.3.

Corollary 3.1.7. *Suppose both $\mathcal{O}$ and $\pi_\sigma(\mathcal{O})$ have the Curve Separation Property and that T has all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture. Then the rational part of $\mathcal{F}_{B(T)}$ folds to the rational part of $\mathcal{F}_{B(\pi_\sigma(T))}$ under σ.*

Corollary 3.1.8. *5 Suppose both $\mathcal{O}$ and $\pi_\sigma(\mathcal{O})$ have the Null Tangle Property, that T has all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture, and R is $\mathbb{Z}$ or $\mathbb{Q}$. Then the set of shear coordinates of allowable curves in $\mathcal{O}$ folds to a positive (cone) basis for $R^{B(\pi_\sigma(T))}$, and therefore to universal geometric coefficients for $B(\pi_\sigma(T))$ over R.*

3.2 Marked surfaces and orbifolds

We begin by describing cluster algebras arising from triangulated surfaces and orbifolds, following [5,7–9]. An orbifold is a topological space which is specified by local conditions: broadly, where a manifold looks locally like Euclidean space, an orbifold looks locally like the quotient space of Euclidean space under the linear action of a finite group [28, Chapter 13]. In the cluster algebra context discussed here we take a very specific view of orbifolds as generalizations of marked surfaces. Indeed, the orbifold definitions below are adapted from [5], where they were introduced to generalize the marked surfaces model of [8,9] to accommodate skew-symmetrizable exchange matrices.

Throughout, we note the specialization of each orbifold definition to the more familiar realm of marked surfaces. Assertions left unproven here are proven in [5,8,9]. Differences between our approach and [5,8,9], most of which correspond to the differences between [22] and [8,9], are described in Remarks 3.2.8, 3.2.11, and 3.3.5.

Definition 3.2.1 (Orbifolds and marked surfaces)**.** An ***orbifold*** is a triple $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ where $\mathbf{S}$ is an oriented surface with (possibly empty) boundary $\partial\mathbf{S}$ and $\mathbf{M}$ and $\mathbf{Q}$ are disjoint, finite sets of points in $\mathbf{S}$. In particular, we view $\mathbf{S}$ as arising from a compact, oriented surface *without* boundary by removing a finite collection of open disks with disjoint closures, so that the boundary $\partial\mathbf{S}$, if it exists, is a finite collection of circles called ***boundary components***, one for each deleted disk. By the ***boundary*** $\partial\mathcal{O}$ of $\mathcal{O}$, we mean $\partial\mathbf{S}$.

We require that $\mathbf{M}$, the set of ***marked points***, be nonempty, with at least one marked point in each boundary component of $\mathcal{O}$. Marked points in the interior of $\mathcal{O}$ are called ***punctures***. On the other hand, the set $\mathbf{Q}$ of ***orbifold points*** may be empty, and must be fully contained in the interior of $\mathcal{O}$. Each orbifold point $q \in \mathbf{Q}$ has a weight $w(q)$ of either 2 or 1/2. We denote the sets of points of each type as $\mathbf{Q_2}$ and $\mathbf{Q_{1/2}}$, respectively, and use the respective symbols $\odot$ and $\times$ to represent them. The symbol $\otimes$ denotes an orbifold point of arbitrary weight.

We forbid the connected components of $\mathcal{O}$ without orbifold points from being any of the following marked surfaces: unpunctured monogons, digons, or triangles, or once-punctured monogons. Furthermore, we require all spheres (with no boundary components) to have at least 4 points in $\mathbf{M} \sqcup \mathbf{Q}$. If such a sphere has exactly four points in $\mathbf{M} \sqcup \mathbf{Q}$, then at least two of them must be punctures in $\mathbf{M}$. (Note that this explicitly excludes the four exceptional cases in [5, Table 3.5].) We also require that every connected component of $\mathcal{O}$ contain at least one marked point, and forbid monogons which contains a single orbifold point and no punctures. A ***null orbifold*** is an empty quadrilateral or a digon containing a unique point in $\mathbf{M} \sqcup \mathbf{Q}$ in its interior (a generalization of the null surfaces defined in [24]).

If $\mathbf{Q} = \mathbf{Q_2} \sqcup \mathbf{Q_{1/2}} = \emptyset$, then the pair $(\mathbf{S}, \mathbf{M})$ is a bordered surface with marked points, as in [8, Definition 2.1], or equivalently, a ***marked surface***.

Remark 3.2.2. Generally, $\mathbf{S}$ is considered to have a Riemannian structure. While the surfaces described above admit a Riemann metric, we do not need it here. We also do away with the usual convention that $\mathbf{S}$ be connected. This is a criterion for "irreducibility" and again, is not needed for our applications. Further, Chapter 4 will introduce an operation on $\mathcal{O}$ which may disconnect it, so we must explicitly permit disconnected surfaces.

3.2.1 Arcs and triangulations

Definition 3.2.3 (Arcs and triangulations)**.** An ***arc*** γ in an orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ is a curve in $\mathbf{S}$, considered up to isotopy relative to $\mathbf{M} \sqcup \mathbf{Q}$, whose endpoints are in $\mathbf{M} \sqcup \mathbf{Q}$. In particular, either both endpoints of γ are points in $\mathbf{M}$, in which case γ is called an ***ordinary arc***, or one endpoint belongs to $\mathbf{M}$ while the other is an orbifold point $q \in \mathbf{Q} = \mathbf{Q_2} \sqcup \mathbf{Q_{1/2}}$, in which case γ is called a ***non-ordinary arc***. Note that arcs cannot have both endpoints in $\mathbf{Q}$. A non-ordinary arc is referred to as a ***double arc*** if $q \in \mathbf{Q_2}$, and a ***pending arc*** if $q \in \mathbf{Q_{1/2}}$. The ***weight*** of an arc γ is 1 if γ is ordinary and the weight of its unique orbifold endpoint if it is not. That is, the weight of a double arc is 2 and the weight of a pending arc is 1/2. (To emphasize the

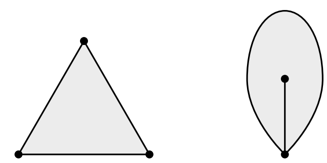

Figure 3.1: Types of triangles admissable in a triangulation of a marked surface (two or three of the vertices of the lefthand triangle may coincide; if two edges coincide, one obtains the self-folded triangle on the right)

difference between ordinary and non-ordinary arcs, we render non-ordinary arcs in bold. See, e.g., Figure 3.2.)

We forbid γ from intersecting either itself or $\partial\mathcal{O} \cup \mathbf{M} \cup \mathbf{Q}$ except possibly at its endpoints, and further forbid γ from being contractible into a single boundary segment. (That is, if γ is contractible into $\partial\mathcal{O}$, then γ has two distinct endpoints and the portion of $\partial\mathcal{O}$ between these endpoints must contain at least 1 marked point.) If γ bounds a monogon, then this monogon must contain a puncture or at least two orbifold points in its interior.

Two arcs are ***compatible*** if (there is an isotopy representative of each such that) the two do not intersect, except possibly at shared endpoints in $\mathbf{M}$. Thus in particular, two compatible non-ordinary arcs cannot share an orbifold point. A ***triangulation*** T is a maximal collection of pairwise compatible arcs. The arcs which define T divide the orbifold into ***triangles*** with three vertices and three edges, some of which may coincide. If two (by necessity, ordinary) edges of a triangle coincide, we refer to it as a ***self-folded triangle*** (see the second image in Figure 3.1). Two distinct arcs which form the edges of (at least one) common triangle are called ***adjacent***. The types of triangles containing non-ordinary arcs are depicted in Figure 3.2. We at times refer to these as ***non-ordinary triangles*** in contrast to those depicted in Figure 3.1. Examples of orbifold triangulations are provided in Figures 3.3, 3.4, and 3.5. Every orbifold point $q \in \mathbf{Q}$ is incident to exactly one non-ordinary arc in T, and each non-ordinary arc is an edge of exactly one triangle. On the other hand, each ordinary arc in T is contained in exactly two triangles unless it constitutes the ***fold edge*** of a self-folded triangle.

If $\mathcal{O}$ is a marked surface, these definitions specialize to [8, Definitions 2.2, 2.4, 2.6].

Definition 3.2.4 (Signed adjacency matrix)**.** Suppose T is a triangulation of an orbifold $\mathcal{O}$ consisting of n arcs $\gamma_1, \dots, \gamma_n$. The ***signed adjacency matrix*** $B(T) = (b_{ij})$ is an $n \times n$ integer matrix with rows and columns indexed by the arcs in T, where each entry b_{ij} for $i, j \in [n]$ encodes the adjacency of arcs γ_i and γ_j and is defined as follows.

If γ_i constitutes two sides of a self-folded triangle (the arc with two distinct endpoints in a triangle of the type depicted on the right in Figure 3.1), set $\pi_T(\gamma_i)$ to be this triangle's

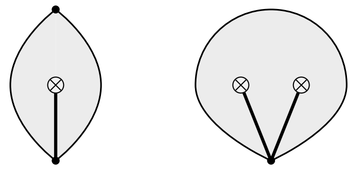

Figure 3.2: Types of non-ordinary triangles admissable in a triangulation of an orbifold (orbifold points of arbitrary weight are denoted by $\otimes$; non-ordinary arcs are rendered in bold)

remaining side (the enclosing loop). Otherwise, set $\pi_T(\gamma_i) = \gamma_i$. Let w_i denote the weight of arc γ_i: since the two arcs in a self-folded triangle are ordinary, w_i is also the weight of $\pi_T(\gamma_i)$. Likewise for γ_j.

Now, for each non-self-folded triangle Δ in T, define ϵ_{ij}^{Δ} to be 0 unless Δ has sides $\pi_T(\gamma_i)$ and $\pi_T(\gamma_j)$, in which case $\epsilon_{ij}^{\Delta} = 1$ if $\pi_T(\gamma_j)$ follows $\pi_T(\gamma_i)$ in the clockwise order, and $\epsilon_{ij}^{\Delta} = -1$ if $\pi_T(\gamma_j)$ follows $\pi_T(\gamma_i)$ in the counter-clockwise order. Further, define

$$b_{ij}^{\Delta} = \begin{cases} 1 & \text{if } w_i < w_j \\ 1 & \text{if } w_i = w_j = 1 \\ 2 & \text{if } w_i = w_j \neq 1 \\ w_i/w_j & \text{if } w_i > w_j \end{cases}$$

Finally, define the matrix entry b_{ij} to be the sum, over all non-self-folded triangles Δ in T, of the signed contributions $\epsilon_{ij}^{\Delta} b_{ij}^{\Delta}$:

$$b_{ij} = \sum_{\substack{\text{non-self-folded} \\ \Delta \text{ in } T}} \epsilon_{ij}^{\Delta} b_{ij}^{\Delta}$$

By construction, the integer matrix $B = B(T)$ is ***skew-symmetrizable***: that is, there exists a (unique) diagonal matrix D with positive integer entries $(d_1, \ldots, d_n)$ (each a power of 2) such that BD is ***skew-symmetric*** (i.e., the transpose of BD equals its negative) and such that the greatest common divisor of the entries of D is 1. Thus $B(T)$ is an exchange matrix. Furthermore, all entries $b_{ij} \in \{0, \pm 1, \pm 2, \pm 4\}$, and if $\mathcal{O}$ is a marked surface so that all of the arcs in T are ordinary, then the definition specializes to [8, Definition 4.1] and all entries $b_{ij} \in \{0, \pm 1, \pm 2\}$. If the orbifold $\mathcal{O}$ has multiple connected components, then $B(T)$ is a block-diagonal matrix with a diagonal block for each such component. The matrix $B(T)$ is the zero matrix if and only if $B(T')$ is the zero matrix for *every* triangulation T' of $\mathcal{O}$, which is true if and only if every connected component of $\mathcal{O}$ is a null orbifold. Examples of orbifold triangulations and the associated signed adjacency matrices are provided in a later section: see

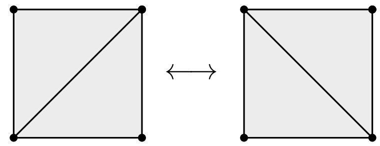

Figure 3.3: An ordinary arc flip

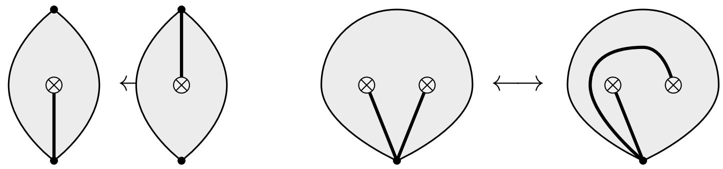

Figure 3.4: Flips of non-ordinary arcs

Figures 3.21 and 3.22.

Definition 3.2.5 (Arc flips)**.** Given a triangulation T of an orbifold $\mathcal{O}$, for each arc $\gamma \in T$ that does not constitute two sides of a self-folded triangle, there exists a unique arc $\gamma' \neq \gamma$ such that $\gamma' \cup (T \setminus \gamma)$ forms a new triangulation T' of $\mathcal{O}$. The operation of replacing γ with γ' is called an arc ***flip***. A flip of an ordinary arc contained in two non-self-folded triangles is depicted in Figure 3.3. Non-ordinary arc flips are depicted in Figure 3.4, and flips of ordinary arcs adjacent to non-ordinary arcs are depicted in Figure 3.5.

Ordinary arcs which constitute two sides of a self-folded triangle cannot be flipped, but passing to tagged arcs and triangulations as in [8, Definitions 7.1,7.2] resolves this issue:

Definition 3.2.6 (Tagged arcs and tagged triangulations)**.** A ***tagged arc*** in $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ is an arc that does not cut out a once-punctured monogon, with each ***end*** of the arc (a location near an endpoint) tagged either ***plain*** or ***notched*** ($\bowtie$). We require that an arc be tagged plain at any marked point endpoints in the boundary $\partial\mathcal{O}$ as well as at any orbifold point endpoints

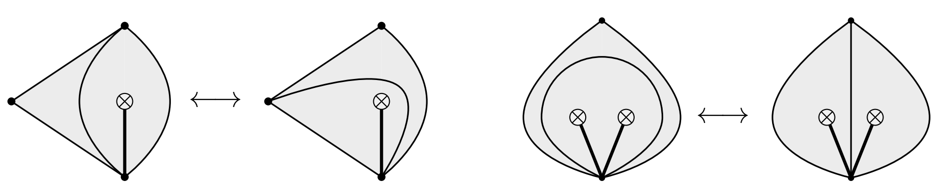

Figure 3.5: Flips of ordinary arcs adjacent to non-ordinary arcs

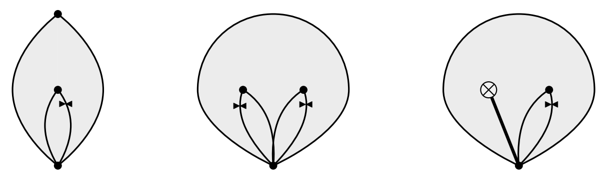

Figure 3.6: Additional types of triangles admissable in a tagged triangulation of an orbifold (the left and righthand images depict two triangles, one of which is self-folded; the central image depicts three triangles, two of which are self-folded)

in $\mathbf{Q} = \mathbf{Q_2} \sqcup \mathbf{Q_{1/2}}$, and that both ends of an arc with coinciding endpoints be tagged the same way. Two tagged arcs are ***compatible*** if either the underlying untagged arcs are distinct and compatible and any shared endpoints have the same tagging, or if the underlying untagged arcs are ordinary and identical, with their tagging agreeing at exactly one endpoint. Arcs satisfying the second condition are referred to here as a ***conjugate pair***. (We emphasize again that compatible arcs cannot share an orbifold point, so in particular, only ordinary tagged arcs can be part of a conjugate pair.)

A ***tagged triangulation*** is a maximal collection of distinct pairwise compatible tagged arcs. In addition to the non-self-folded ordinary triangle depicted on the left in Figure 3.1 and the non-ordinary triangles depicted in Figure 3.2, a tagged triangulation of an orbifold may also contain triangles of the type depicted in Figure 3.6. Arc taggings can be assigned arbitrarily at any marked point vertex which is neither in a boundary component nor incident to the oppositetly-tagged ends of a conjugate pair.

There is a canonical map τ from (untagged) arcs to tagged arcs in surfaces described in [8, Definition 7.2] which extends easily to orbifolds. Given an (untagged) arc γ in $\mathcal{O}$, $\tau(\gamma)$ is the same curve tagged plain at both ends unless γ bounds a once-punctured monogon. In the latter case, γ must be ordinary, and if the coinciding endpoints of γ are at $a \in \mathbf{M}$ and the bounded puncture is $b \in \mathbf{M}$, then $\tau(\gamma)$ is the arc connecting a to b inside the mongon, tagged plain at a and notched at b. Applying τ to each arc in an (untagged) triangulation yields a tagged triangulation, where (untagged) self-folded triangles of the type depicted on the right in Figure 3.1 are mapped, under τ, to self-folded triangles formed by a conjugate pair as depicted in Figure 3.6.

Define the ***signed-adjacency matrix*** $B(T)$ of a tagged triangulation T as follows. First, everywhere there is a puncture in $\mathbf{M}$ incident only to notched ends of arcs in T, change all of these notched ends to plain ends. The resulting triangulation T' has no notched endpoints on arcs except on those which are members of conjugate pairs, and therefore $T' = \tau(T^0)$ for some (untagged) triangulation T^0. Then. set $B(T) = B(T^0)$ as defined in Definition 3.2.4.

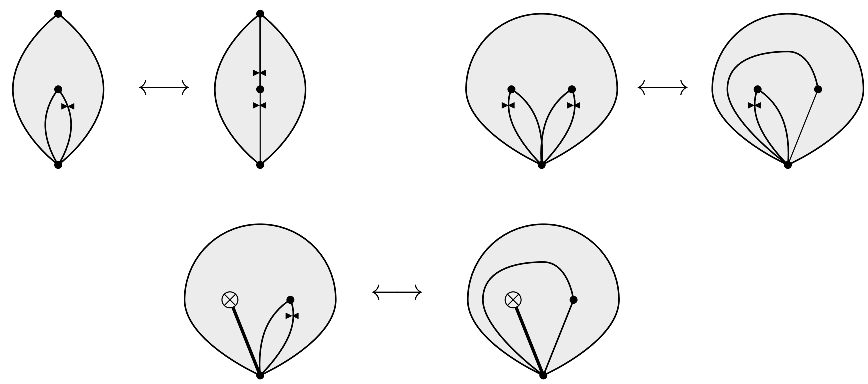

Figure 3.7: Flips of arcs in a conjugate pair

All tagged arcs in a tagged triangulation can be flipped: in particular, flips of arcs in a conjugate pair are depicted in Figure 3.7.

Remark 3.2.7. Let $\mathcal{O}$ be an orbifold. By [5, Theorem 4.2], any two (untagged) triangulations of $\mathcal{O}$ are related by a sequence of (untagged) arc flips. This is a generalization of [8, Proposition 3.8], which establishes the same result for marked surfaces. In addition, [8, Proposition 7.10] states that any two *tagged* triangulations of a marked surface are related by a sequence of tagged arc flips, unless the surface has no boundary components and exactly one puncture. In this case, the collection of tagged triangulations is split into two disjoint sets: any two triangulations in which all arcs are tagged plain are related by a sequence of arc flips, and likewise for any two triangulations in which all arcs are tagged notched. However, there is no sequence of arc flips that will relate a plain triangulation to a notched one.

While tagged arc flips in orbifolds are not explicitly addressed in [5], we quickly verify that the same issue arises in this more general setting. That is, if $\mathcal{O}$ has no boundary components and exactly one puncture, then *regardless* of its number of orbifold points there is no sequence of arc flips that will relate a triangulation of $\mathcal{O}$ with all arcs tagged plain at the puncture to one with all arcs tagged notched. (Recall that arcs must be tagged plain at orbifold point endpoints.) The tagging at the puncture can be switched if and only if there exists a tagged triangulation of $\mathcal{O}$ which contains a conjugate pair of arcs with opposite taggings at that puncture. Since arcs in a conjugate pair are by necessity ordinary and incident to two *distinct* marked points, no such tagged triangulation may be constructed. Thus, as with surfaces with no boundary and exactly one puncture, the collection of tagged triangulations of $\mathcal{O}$ is split into two disjoint sets, where there is no way to obtain a triangulation in one from a triangulation in the other through a sequence of tagged arc flips. However, it is easy to use [5, Theorem 4.2] to verify that the triangulations within each set are all related to one another by a sequence of tagged arc flips, and furthermore, that if one instead assumes that $\mathcal{O}$ has at least one boundary component or more than one puncture (or both), *any* two tagged triangulations of $\mathcal{O}$ are related in this way. Hence whenever we wish to ensure that two tagged triangulations T and T' on a given orbifold $\mathcal{O}$ are related to one another by a sequence of tagged arc flips, we specify that both T and T' have all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture.

Remark 3.2.8. Before proceeding, we stop to reconcile some differences between our approach and [5], wherein the orbifold model is introduced. The object $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ which we call an orbifold in Definition 3.2.1 is a "weighted orbifold" $\mathcal{O}^w$ in the language of [5, Section 4], which initially defines an "orbifold" as the same object but with no weights on the orbifold points in $\mathbf{Q}$. Given a weighted orbifold $\mathcal{O}^w$ with tagged triangulation T, [5, Section 5] then constructs an additional object, $\hat{\mathcal{O}}$, called the "associated orbifold," with triangulation $\tilde{T}$, by replacing each orbifold point $q \in \mathbf{Q}$ of weight 2 by a "special marked point," and replacing the unique (non-ordinary) arc in T which was incident to q by a conjugate pair of (ordinary) arcs tagged differently at this point which must be flipped simultaneously. The final step in the construction is to replace this conjugate pair of arcs by a single "double arc" tagged plain at the special marked point. For simplicity, we forego all of these distinctions, and introduce the term "double arcs" from the outset to denote non-ordinary arcs incident to orbifold points of weight 2. Further, we use the same symbol, ⦿, for orbifold points of weight 2 that [5] uses for special marked points. We reserve the symbol $\times$ for orbifold points of weight $1/2$, and the term "pending arcs" for non-ordinary arcs incident to such points. We further introduce the symbol $\otimes$ for an orbifold point of arbitrary weight. (In [5], $\times$ is used to denote any orbifold point, and all non-ordinary arcs are called pending.) None of these changes are substantive, and essentially amount to changes in notation and terminology: all of the underlying definitions in Section 3.2.1 coincide with those here.

Translating terminology as explained in Remark 3.2.8, the following result is equivalent to [5, Theorem 4.19]:

Theorem 3.2.9. *Let T be a tagged triangulation of an orbifold $\mathcal{O}$. Then the triangulation T' of $\mathcal{O}$ is obtained from T by a flip at an arc $\gamma \in T$ if and only if $B(T') = \mu_\gamma(B(T))$, where the indexing of the exchange matrices by arcs in the triangulations is maintained. Furthermore, an exchange matrix B' is mutation-equivalent to $B(T)$ if and only if $B' = B(T')$ for some tagged triangulation T' of $\mathcal{O}$ obtained from T by a sequence of arc flips.*

3.2.2 Allowable curves and quasi-laminations

In addition to tagged arcs and triangulations, we will need an additional class of curves in orbifolds, called allowable curves, and a corresponding notion of pairwise compatibility. If $\mathcal{O}$ is a marked surface the definitions below specialize to [22, Definition 4.1]. Just as the collections defined there are almost, but not quite, the unbounded measured laminations on surfaces of [9, Definition 12.1], the collections defined here can be thought of as a similar adjustment to the laminations on orbifolds described in [5, Definition 6.1]. We discuss this adjustment in Remark 3.2.11 after presenting the definitions.

Definition 3.2.10 (Allowable curves and quasi-laminations)**.** An ***allowable curve*** λ in an orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ is a non-self-intersecting curve in $\mathbf{S}$, considered up to isotopy relative to $\mathbf{M} \sqcup \mathbf{Q}$, which is either

- a ***closed curve***, or
- a non-closed curve which at each end does one of the following:
  - spirals into a puncture in $\mathbf{M}$ (in either a clockwise or counterclockwise direction: we call this a ***spiral point***),
  - terminates in an unmarked point in the boundary $\partial\mathcal{O}$, or
  - terminates in an orbifold point in $\mathbf{Q}$.

We forbid λ from having any self-intersections, from being contractible in $\mathbf{S} \setminus (\mathbf{M} \sqcup \mathbf{Q})$, and from being contractible to either a puncture, an orbifold point, two orbifold points of weight 1/2 (see the leftmost image in Figure 3.8), or a portion of the boundary $\partial\mathcal{O}$ containing zero or one marked points. Mirroring our terminology for arcs, we refer in general to a curve λ with an end terminating in an orbifold point as ***non-ordinary*** (and render it in bold), and all other curves as ***ordinary***. If λ terminates in an orbifold point at both ends, we require these orbifold points to be distinct and both of weight 1/2, and refer to λ as a ***semi-closed curve***. If λ terminates in an orbifold point q at exactly one end, we refer to it as a ***pending curve*** if $q \in \mathbf{Q_{1/2}}$ or as a ***double curve*** if $q \in \mathbf{Q_2}$.

Finally, we also forbid λ from cutting out, either by itself or together with a portion of the boundary not containing a marked point, a disc containing a unique point in $\mathbf{M} \sqcup \mathbf{Q}$. (See the second through fifth images Figure 3.8 for pictures of these forbidden non-closed curves.) Note that while only clockwise spiraling is depicted in the fourth and fifth images, counterclockwise spiraling is forbidden as well. (Closed curves which cut out a disk containing a unique point in $\mathbf{M} \sqcup \mathbf{Q}$ are already forbidden, since allowable curves cannot be contractible to a puncture or an orbifold point.)

Two allowable curves are ***compatible*** if and only if they satisfy one of following conditions:

(i) there is an isotopy representative of each such that the two representatives do not intersect (where coinciding spiral points in the same direction are considered non-intersecting), or

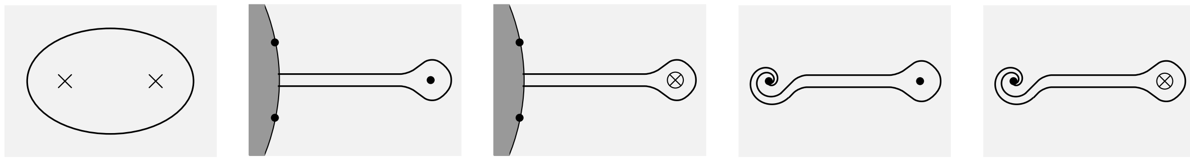

Figure 3.8: Excluded curves for quasi-laminations

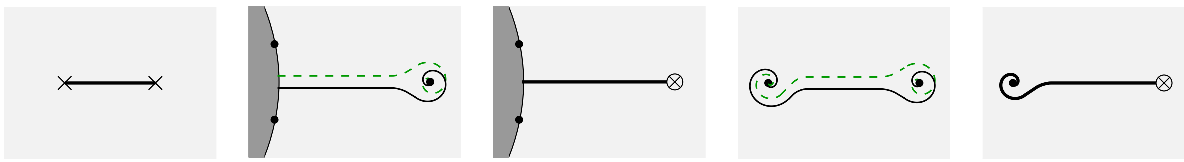

Figure 3.9: Allowable curves and curve intersections for quasi-laminations

(ii) the two curves are ordinary and the same (up to isotopy) except that their spiral direction disagrees at exactly one end (for example, see the first and third images in Figure 3.9). Again mirroring our terminology for arcs, we refer to such curves as a ***conjugate pair***.

Note that in particular, two compatible curves cannot share an orbifold endpoint.

An ***integral quasi-lamination*** L is a collection of pairwise compatible allowable curves λ, distinct up to isotopy, each with a positive integer weight w_λ. A ***rational quasi-lamination*** is defined in the same way, but the positive weight w_λ need only be rational. The set of curves appearing in a given quasi-lamination L is called the ***support*** of L, and usually denoted Λ_L. (In general, Λ is used to denote a collection of (unweighted) pairwise compatible allowable curves.) We speak of both $w_\lambda\lambda \in L$ and $\lambda \in L$. For $k \in \mathbb{Z}_+$, kL denotes the quasi-lamination $\{kw_\lambda\lambda : w_\lambda\lambda \in L\}$.

Remark 3.2.11. Returning to the discussion of Remark 3.2.8, we pause to reconcile the differences between the notions of allowable curves and quasi-laminations in $\mathcal{O}$ described in Definition 3.2.10 and the curves and laminations given in [5, Definition 6.1]. In contrast to the discrepancies in terminology of the previous section, there are substantive changes here: however, as mentioned in this section's introduction, they are still fairly minor and are analogous to the adjustments which [22, Definition 4.1] makes to [9, Definition 12.1]. Given an (integral unbounded measured) lamination $\mathcal{L}$, one may obtain an integral quasi-lamination in the following manner. First, replace the curves depicted in Figure 3.8, which are not allowable but satisfy the conditions placed on curves in [5], by the corresponding (collections of) allowable curves depicted in Figure 3.9. If the resulting collection contains k copies of the same curve, replace them with a single copy with weight k. It is clear from Figures 3.8 and 3.9 that these distinct allowable curves with integer weights are pairwise compatible, and therefore form an integral quasi-lamination.

3.3 Shear coordinates on orbifolds

The final ingredient of the orbifolds and marked surfaces model for cluster algebras is a means of encoding the interaction between a tagged triangulation and a quasi-lamination on $\mathcal{O}$: this is a map from quasi-laminations to rational vectors is called shear coordinates.

3.3.1 Computing shear coordinates

The definition below is [9, Definition 13.1], except we apply it to quasi-laminations on orbifolds rather than laminations on marked surfaces.

Definition 3.3.1 (Shear coordinates of quasi-laminations). Let T be a tagged triangulation of an orbifold $\mathcal{O}$ and let L be a quasi-lamination. The ***shear coordinate vector*** (or simply ***shear coordinates***) of L with respect to T is a vector $\mathbf{b}(T, L) = (b_\gamma(T, L) : \gamma \in T)$ indexed by arcs $\gamma \in T$, where each entry $b_\gamma(T, L)$ is defined as a weighted sum of quantities $b_\gamma(T, \lambda)$ for allowable curves $\lambda \in L$. Specifically,

$$b_\gamma(T, L) = \sum_{\lambda \in L} w_\lambda b_\gamma(T, \lambda) \tag{3.1}$$

where w_λ is the weight of the curve λ in the quasi-lamination L and $b_\gamma(T, \lambda)$ is a quantity that we now describe.

Select an isotopy-class representative of λ which minimizes the number of points at which λ intersects γ. Define $b_\gamma(T, \lambda)$ as the sum, over each intersection of γ with λ, of a number in the set $\{0, \pm 1, \pm 2, \pm 4\}$:

- If γ is ordinary, an intersection of λ with γ contributes 0 unless λ intersects the two triangles sharing γ as shown in Table 3.1. If γ is part of a conjugate pair, we indicate the contribution to $b_\gamma(T, \lambda)$ for both tagging possibilities on γ.

- If γ is non-ordinary, an intersection of λ with γ contributes 0 unless λ intersects the single triangle containing γ as shown in Table 3.2.

Intersections which contribute non-zero amounts are called ***non-trivial intersections***. Similarly, those which contribute zero are called ***trivial***.

Tables 3.1 and 3.2 introduce the color convention of [22, Figure 2], whereby we color λ black, color γ purple, and color the other relevant arcs in T red and blue so that, when traveling clockwise along a triangle, we meet the colors in the order blue, purple, red. We further introduce the following additional conventions:

- Non-ordinary arcs and curves are drawn in bold. Unless specified, all other arcs and curves may be ordinary or non-ordinary.
- The images in the top two rows of Table 3.1 each contain two marked points: these need not be distinct.
- The latter rows of Table 3.1 handle the cases when γ is part of a conjugate pair: while the contributions for both tagging possibilities are given, the coloring shown assumes γ is tagged notched. (In particular, to emphasize the connection to [22, Figure 2], we color the arc with a notched tagging purple and color the arc with plain tagging both red and blue, since it is playing the role of two sides of the self-folded triangle.)
- An intersection between dashed portions of an arc γ (rendered in green) and a curve λ (rendered in black) may be one of 6 possible types, as illustrated in Figure 3.10.

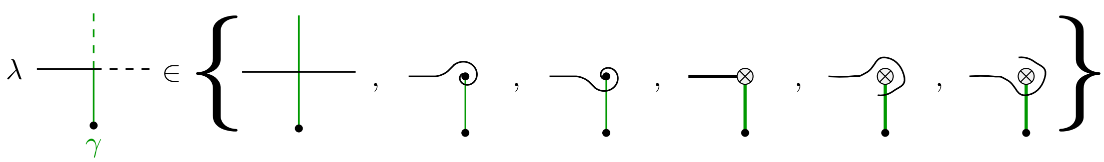

Figure 3.10: Convention for curve and arc intersections used in Tables 3.1 and 3.2

Remark 3.3.2. Observe from Tables 3.1 and 3.2 that whether or not an arc or curve is ordinary has no impact on the *sign* of the shear coordinate contribution arising from their intersection, just the magnitude. That is, an intersection between an allowable curve λ and a tagged arc γ in a tagged triangulation T contributes positively, negatively, or zero regardless of whether λ and γ are ordinary or non-ordinary.

3.3.2 Properties of shear coordinates

The following observations about shear coordinates in orbifolds follow easily from examination of Tables 3.1 and 3.2. We use them both to reconcile Definition 3.3.1 with [5, Definition 6.2] (see Remark 3.3.5), as well as to prove the primary result of this subsection, that there exists a bijection between rational (resp., integral) quasi-laminations and $\mathbb{Q}^n$ (resp., $\mathbb{Z}^n$) (see Theorem 3.3.6).

Arc γ	**Curve λ**	**Positive intersection**	**Negative intersection**
ordinary	ordinary, pending, or semi-closed	$+1$ 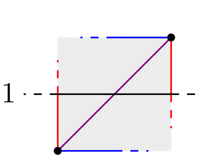	-1 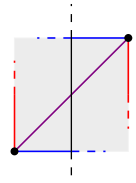
ordinary	double	$+2$ 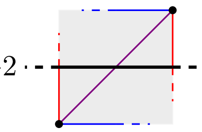	-2 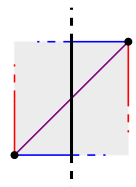
ordinary conjugate pair	ordinary, pending, or semi-closed	notched: $+1$ plain: $+1$ 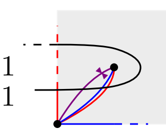	notched: -1 plain: -1 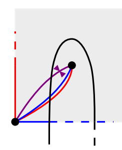
ordinary conjugate pair	double	notched: $+2$ plain: $+2$ 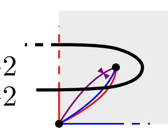	notched: -2 plain: -2 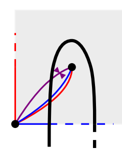
ordinary conjugate pair	ordinary or pending	notched: $+1$ plain: 0 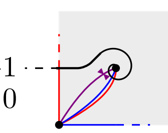 notched: 0 plain: $+1$ 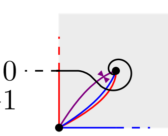	notched: -1 plain: 0 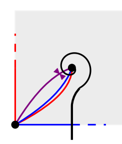 notched: 0 plain: -1 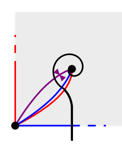
ordinary conjugate pair	double	notched: $+2$ plain: 0 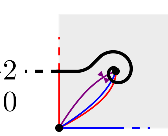 notched: 0 plain: $+2$ 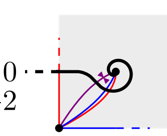	notched: -2 plain: 0 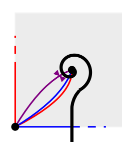 notched: 0 plain: -2 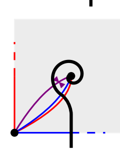

Table 3.1: Shear coordinate computation, ordinary arcs: ends of arcs at punctures not incident to a conjugate pair of arcs with opposite taggings can be tagged arbitrarily; in the latter rows shear coordinates are given with respect to both members of a conjugate pair, but coloring shown is with respect to the arc tagged notched

Arc γ	**Curve λ**	**Positive intersection**	**Negative intersection**
pending or double	pending, double, or semi-closed	$+1$ 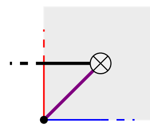	-1 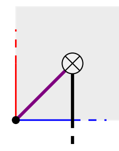
pending	ordinary, pending, or semi-closed	$+2$ 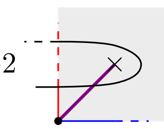	-2 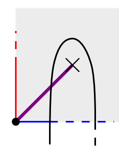
pending	double	$+4$ 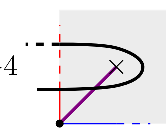	-4 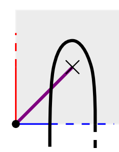
double	ordinary, pending, or semi-closed	$+1$ 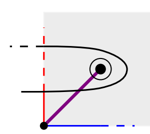	-1 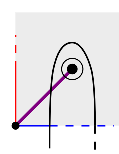
double	double	$+2$ 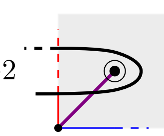	-2 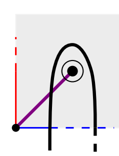

Table 3.2: Shear coordinate computation, non-ordinary arcs: ends of arcs at punctures not incident to a conjugate pair of arcs with opposite taggings can be tagged arbitrarily

Lemma 3.3.3. *Let T be a tagged triangulation of orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ with conjugate pair of arcs $\gamma', \gamma'' \in T$ tagged differently at puncture $v \in \mathbf{M}$.*

1. *Let L be a quasi-lamination on $\mathcal{O}$.*
   (a) *$b_{\gamma'}(T, L) = b_{\gamma''}(T, L)$ if and only if the set of curves in L which spiral into v is either empty or consists of a conjugate pair λ', λ'' with the same weight $w_{\lambda'} = w_{\lambda''} = w_\lambda$.*
   (b) *$|b_{\gamma'}(T, L) - b_{\gamma''}(T, L)| = 1$ if and only if the set of curves in L which spiral into v consists of either a single curve with weight 1 or a conjugate pair λ', λ'' with weights satisfying $|w_{\lambda'} - w_{\lambda''}| = 1$.*
2. *Let λ', λ'' be a conjugate pair of curves in $\mathcal{O}$ with opposite spiral directions at v.*
   (a) *$b_\gamma(T, \lambda') = b_\gamma(T, \lambda'')$* for all *$\gamma \in T \setminus \{\gamma', \gamma''\}$.*
   (b) *$b_{\gamma'}(T, \lambda'), b_{\gamma''}(T, \lambda'), b_{\gamma'}(T, \lambda'')$, and $b_{\gamma''}(T, \lambda'')$ all weakly agree in sign, and each pair-wise difference among them is of absolute value 1.*
   (c) *The following four sums are equal:*

$$b_{\gamma'}(T, \lambda') + b_{\gamma''}(T, \lambda') = b_{\gamma'}(T, \lambda'') + b_{\gamma''}(T, \lambda'')$$
$$b_{\gamma'}(T, \lambda') + b_{\gamma'}(T, \lambda'') = b_{\gamma''}(T, \lambda') + b_{\gamma''}(T, \lambda'').$$

Lemma 3.3.4. *Let T be a tagged triangulation of orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ and let $q \in \mathbf{Q}$. Then q is incident to exactly one non-ordinary arc: choose an indexing $\gamma_1, \gamma_2, \ldots, \gamma_n$ of the arcs in T such that this non-ordinary arc is γ_n. Form a new orbifold $\mathcal{O}'_q = (\mathbf{S}, \mathbf{M} \cup \{q\}, \mathbf{Q} \setminus \{q\})$ with tagged triangulation T'_q from $\mathcal{O}, T$ by replacing the orbifold point at q with a puncture, and substituting the non-ordinary arc γ_n incident to q with a conjugate pair of arcs γ'_n, γ''_n that coincide with γ_n except γ'_n is tagged notched at q' and γ''_n is tagged plain. Let λ be a curve in $\mathcal{O}$.*

1. *If λ does not terminate in q, then λ is an allowable curve in $\mathcal{O}$ if and only if it is an allowable curve in $\mathcal{O}'_q$. Further,*
   (a) *If $q \in \mathbf{Q_2}$,*

$$b_{\gamma_i}(T, \lambda) = \begin{cases} b_{\gamma_i}(T'_q, \lambda) & \text{if } i \neq n \\ b_{\gamma'_n}(T'_q, \lambda) = b_{\gamma''_n}(T'_q, \lambda) = \frac{1}{2}\left(b_{\gamma'_n}(T'_q, \lambda) + b_{\gamma''_n}(T'_q, \lambda)\right) & \text{if } i = n. \end{cases}$$

   (b) *If $q \in \mathbf{Q_{1/2}}$,*

$$b_{\gamma_i}(T, \lambda) = \begin{cases} b_{\gamma_i}(T'_q, \lambda) & \text{if } i \neq n \\ b_{\gamma'_n}(T'_q, \lambda) + b_{\gamma''_n}(T'_q, \lambda) = 2b_{\gamma'_n}(T, \lambda) = 2b_{\gamma''_n}(T, \lambda) & \text{if } i = n. \end{cases}$$

2. *If λ terminates in q, then let λ', λ'' be the (unique) pair of curves in $\mathcal{O}'_q$ that coincide with λ except that they spiral into q in opposite directions. Then λ is an allowable curve in $\mathcal{O}$ if and only if λ' and λ'' both allowable curves in $\mathcal{O}'_q$. Further,*

   (a) *If $q \in \mathbf{Q_2}$, then λ is a double curve, λ, λ'' are a (compatible, ordinary) conjugate pair of curves, and*

$$b_{\gamma_i}(T,\lambda) = \begin{cases} b_{\gamma_i}(T'_q,\lambda') + b_{\gamma_i}(T'_q,\lambda'') = 2b_{\gamma_i}(T'_q,\lambda') = 2b_{\gamma_i}(T'_q,\lambda'') & \text{if } i \neq n \\ b_{\gamma'_n}(T'_q,\lambda') + b_{\gamma''_n}(T'_q,\lambda') = b_{\gamma'_n}(T'_q,\lambda'') + b_{\gamma''_n}(T'_q,\lambda'') & \text{if } i = n. \end{cases}$$

   (b) *If $q \in \mathbf{Q_{1/2}}$ and λ is a pending curve, then λ', λ'' are a conjugate pair of curves. If λ is semi-closed, then λ, λ'' are non-compatible pending curves. Regardless,*

$$b_{\gamma_i}(T,\lambda) = \begin{cases} b_{\gamma_i}(T'_q,\lambda') = b_{\gamma_i}(T'_q,\lambda'') & \text{if } i \neq n \\ b_{\gamma'_n}(T'_q,\lambda') + b_{\gamma''_n}(T'_q,\lambda') = b_{\gamma'_n}(T'_q,\lambda'') + b_{\gamma''_n}(T'_q,\lambda'') & \text{if } i = n.. \end{cases}$$

3. *If λ terminates in q and is semi-closed, with its other end terminating in distinct orbifold point $r \neq q$ (recall by necessity that both $q, r \in \mathbf{Q_{1/2}}$), then let λ''' be the (unique) closed curve in $\mathcal{O}'_q$ separating q and r from the other points in $\mathbf{M} \sqcup \mathbf{Q}$. Then λ is an allowable curve in $\mathcal{O}$ if and only if λ''' is an allowable curve in $\mathcal{O}'_q$. Further,*

$$b_{\gamma_i}(T,\lambda) = \begin{cases} \frac{1}{2} b_{\gamma_i}(T',\lambda''') & \text{if } i \neq n \\ b_{\gamma'_n}(T,\lambda''') = b_{\gamma''_n}(T,\lambda''') & \text{if } i = n. \end{cases}$$

Remark 3.3.5. A more concisely stated (albeit less direct) method for computing shear coordinates on orbifolds may be adapted from [5, Definition 6.2], with a few small adjustments to account for our use of quasi-laminations rather than laminations. We describe this method in the following paragraph; application of Lemma 3.3.4 shows that it yields precisely the same vectors as Definition 3.3.1.

Let $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ be an orbifold with tagged triangulation T and quasi-lamination L. Construct a surface $\mathcal{O}'$ with tagged triangulation T' by substituting each digon or monogon in T containing an orbifold point by a digon or monogon containing only marked points as shown in Figure 3.11. Note that this amounts to constructing $\mathcal{O}'_q, T'_q$ as in Lemma 3.3.4 successively for each $q \in \mathbf{Q}$: each ordinary arc $\gamma \in T$ is also an arc in T', while each non-ordinary arc $\gamma \in T$ corresponds to a conjugate pair of arcs $\gamma', \gamma'' \in T'$. Construct a lamination L' on $\mathcal{O}'$ by substituting each pending or semi-closed curve in $\mathrm{supp}(L)$ by a curve with the same weight that coincides with it except for spiraling counterclockwise into the marked point(s) corresponding to its orbifold endpoint(s), and substituting each double curve in $\mathrm{supp}(L)$ by a

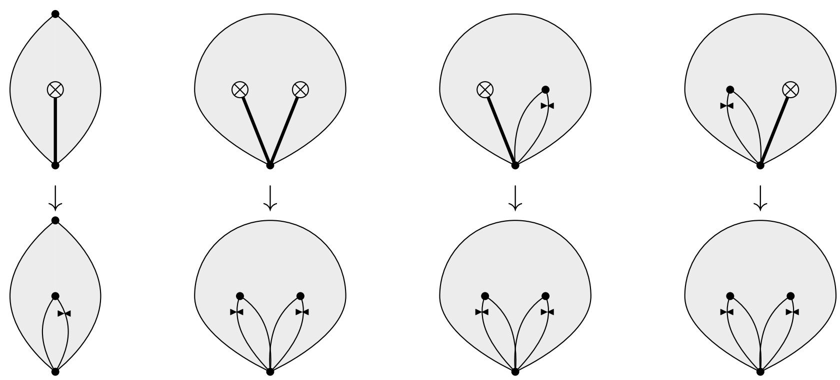

Figure 3.11: Constructing the surface $\mathcal{O}'$ and triangulation T' to compute shear coordinates on $\mathcal{O}$ with respect to T as in Remark 3.3.5.

pair of conjugate pair of curves, both with the same weight, that coincide with it except for spiraling in opposite directions into the marked point corresponding to its orbifold endpoint. Then the shear coordinates $b_\gamma(T, L)$ may be defined in terms of the shear coordinates of the quasi-lamination L' on the surface $\mathcal{O}'$ with tagged triangulation T' as follows, where the latter terms are computed using [9, Definitions 12.2 and 13.1].

- If γ is ordinary, $b_\gamma(T, L) = b_\gamma(T', L')$.
- If γ is pending, $b_\gamma(T, L) = b_{\gamma'}(T', L') + b_{\gamma''}(T', L')$.
- If γ is double, $b_\gamma(T, L) = \frac{1}{2}\left[b_{\gamma'}(T', L') + b_{\gamma''}(T', L')\right]$.

If the quasi-lamination L contains neither conjugate pairs of curves nor double curves, which are not permissable in the laminations of [5, Definition 6.1], then this definition coincides with [5, Definition 6.2] (after making the adjustments described in Remarks 3.2.8 and 3.2.11.) Furthermore, it is easy to check that if all conjugate pairs of curves and double curves in the quasi-lamination L (as depicted in Figure 3.9) are replaced with the corresponding curves which are permissable in [5] (as depicted in Figure 3.8) to produce a lamination $\tilde{L}$, then $\mathbf{b}(T, L) = \mathbf{b}(T, \tilde{L})$.

We are now ready to prove the main result of this section, a generalization of [22, Theorem 4.4] to orbifolds. Just as [22, Theorem 4.4] was a reframing of [9, Theorem 13.6] in terms of quasi-laminations, our result may be thought of as a similar reframing of [5, Theorem 6.7]. Due to the differences between our definitions and those in [5], we prove the result from scratch, although the method is identical.

Theorem 3.3.6. *Fix a tagged triangulation T of orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$, where T has n arcs. Then the map $L \mapsto \mathbf{b}(T, L)$ is a bijection between integral (respectively, rational) quasi-laminations and $\mathbb{Z}^n$ (respectively, $\mathbb{Q}^n$).*

Proof. We proceed by induction on $|\mathbf{Q}|$, the number of orbifold points on $\mathcal{O}$. If $\mathbf{Q} = \emptyset$, then $\mathcal{O} = (\mathbf{S}, \mathbf{M})$ is a marked surface, and the result follows from [22, Theorem 4.4].

Assume $|\mathbf{Q}| > 0$ and the result holds for all orbifolds with fewer than $|\mathbf{Q}|$ orbifold points. Let $q \in \mathbf{Q}$ and choose an indexing $\gamma_1, \gamma_2, \ldots, \gamma_n$ of the arcs in T such that the unique non-ordinary arc incident to q is γ_n. We make use of the triangulated orbifold $\mathcal{O}'_q = (\mathbf{S}, \mathbf{M} \cup \{q\}, \mathbf{Q} \setminus \{q\})$, $T'_q = T \setminus \{\gamma_n\} \cup \{\gamma'_n, \gamma''_n\}$ constructed from $\mathcal{O}, T$ in Lemma 3.3.4 and deal with the cases $q \in \mathbf{Q_2}, q \in \mathbf{Q_{1/2}}$ separately, first showing that the map $L \mapsto \mathbf{b}(T, L)$ is surjective and then showing it is injective.

Suppose that $q \in \mathbf{Q_2}$ and let $(k_1, \ldots, k_n) \in \mathbb{Z}_n$ (respectively, $\mathbb{Q}^n$). By inductive assumption, there exists a unique quasi-lamination L'_q on $\mathcal{O}'_q$ such that

$$\begin{aligned} &b_{\gamma_i}(T'_q, L'_q) = k_i \text{ for } i \in [n-1],\\ &b_{\gamma'_n}(T'_q, L'_q) = k_n, \text{ and}\\ &b_{\gamma''_n}(T'_q, L'_q) = k_n \end{aligned}$$

By Lemma 3.3.3, the set of curves in L' which spiral into q consists of a conjugate pair λ', λ'' with the same weight w (where we use $w = 0$ to mean that no curves spiral into q'). Let λ be the (unique) non-ordinary allowable curve in $\mathcal{O}$ that terminates in q but otherwise coincides with λ' and λ'', and set $L = \tilde{L} \cup \{w\lambda\}$, where $\tilde{L} = L' \setminus \{w\lambda', w\lambda''\}$. (If $w = 0$, then $L = \tilde{L} = L'$.) Applying Lemma 3.3.4, we conclude that L is a quasi-lamination on $\mathcal{O}$ with shear coordinate vector $\mathbf{b}(T, L) = (k_1, \ldots, k_n)$:

$$\begin{aligned} b_{\gamma_i}(T, L) &= b_{\gamma_i}(T, \tilde{L}) + w b_{\gamma_i}(T, \lambda) \text{ for all } i \in [n-1]\\ &= b_{\gamma_i}(T'_q, \tilde{L}) + w\left[b_{\gamma_i}(T'_q, \lambda') + b_{\gamma_i}(T'_q, \lambda'')\right]\\ &= b_{\gamma_i}(T'_q, L'_q)\\ &= k_i.\\ b_{\gamma_n}(T, L) &= b_{\gamma_n}(T, \tilde{L}) + w b_{\gamma_n}(T, \lambda)\\ &= \frac{1}{2}\left(b_{\gamma'_n}(T, \tilde{L}) + b_{\gamma''_n}(T, \tilde{L})\right)\\ &\quad + w\left[\frac{1}{2}\left(b_{\gamma'_n}(T'_q, \lambda') + b_{\gamma''_n}(T'_q, \lambda') + b_{\gamma'_n}(T'_q, \lambda'') + b_{\gamma''_n}(T'_q, \lambda'')\right)\right]\\ &= \frac{1}{2}\left(b_{\gamma'_n}(T, L'_q) + b_{\gamma''_n}(T, L'_q)\right)\\ &= k_n. \end{aligned}$$

We now show that L is unique, thereby proving injectivity. Suppose L_1, L_2 are quasi-laminations on $\mathcal{O}$ with $\mathbf{b}(T, L_1) = \mathbf{b}(T, L_2)$. We build (unique) quasi-laminations L'_1, L'_2 on $\mathcal{O}'_q$ with $\mathbf{b}(T'_q, L'_1) = \mathbf{b}(T'_q, L'_2)$ as follows:

- If L_i does not contain a non-ordinary curve terminating in q, then take $L'_i = L_i$. Lemma 3.3.4 imply that L'_i is indeed a quasi-lamination on $\mathcal{O}'_q$, and that

$$\begin{aligned} &b_{\gamma'_i}(T'_q, L'_i) = b_{\gamma_i}(T, L_i) \text{ for all } i \in [n-1], \\ &b_{\gamma'_n}(T'_q, L'_i) = b_{\gamma_n}(T, L_i), \text{ and} \\ &b_{\gamma''_n}(T'_q, L'_i) = b_{\gamma_n}(T, L_i). \end{aligned}$$

- Otherwise, $L_i = \tilde{L}_i \cup \{w\lambda\}$ for double curve λ, and we take $L'_i = \tilde{L}_i \cup \{w_\lambda \lambda', w\lambda'\}$ where λ', λ'' are the conjugate pair of curves with opposite spiral directions at q'. Lemma 3.3.4 implies that L'_i is once more a quasi-lamination on $\mathcal{O}'_q$, and that it has the same shear coordinates as the construction above:

$$\begin{aligned} b_{\gamma'_i}(T'_q, L'_i) &= b_{\gamma_i}(T'_q, \tilde{L}_i) + w b_{\gamma_i}(T'_q, \lambda') + w b_{\gamma_i}(T'_q, \lambda'') \text{ for all } i \in [n-1] \\ &= b_{\gamma_i}(T, \tilde{L}_i) + w \cdot \frac{1}{2} b_{\gamma_i}(T, \lambda) + w \cdot \frac{1}{2} b_{\gamma_i}(T, \lambda) \\ &= b_{\gamma_i}(T, \tilde{L}_i) + w b_{\gamma_i}(T, \lambda) \\ &= b_{\gamma_i}(T, L_i). \\ b_{\gamma'_n}(T'_q, L'_i) &= b_{\gamma'_n}(T'_q, \tilde{L}_i) + w b_{\gamma'_n}(T'_q, \lambda') + w b_{\gamma'_n}(T'_q, \lambda'') \\ &= b_{\gamma_n}(T, \tilde{L}_i) + w b_{\gamma_n}(T, \lambda) \\ &= b_{\gamma_n}(T, L_i). \\ b_{\gamma''_n}(T'_q, L'_i) &= b_{\gamma''_n}(T'_q, \tilde{L}_i) + w b_{\gamma''_n}(T'_q, \lambda') + w b_{\gamma''_n}(T'_q, \lambda'') \\ &= b_{\gamma_n}(T, \tilde{L}_i) + w b_{\gamma_n}(T, \lambda) \\ &= b_{\gamma_n}(T, L_i). \end{aligned}$$

By this construction, L_1 and L_2 are distinct quasi-laminations on $\mathcal{O}$ if and only if L'_1 and L'_2 are distinct quasi-laminations on $\mathcal{O}'_q$. However since $\mathbf{b}(T'_q, L'_1) = \mathbf{b}(T'_q, L'_2)$, our inductive assumption implies that $L'_1 = L'_2$. Thus $L_1 = L_2$, and we conclude that the desired result holds when $q \in \mathbf{Q_2}$.

Alternatively, suppose that $q \in \mathbf{Q_{1/2}}$, and again let $(k_1, \ldots, k_n) \in \mathbb{Z}_n$ (respectively, $\mathbb{Q}^n$). By inductive assumption, there exists a unique quasi-lamination L'_q on $\mathcal{O}'_q$ such that

$$b_{\gamma_i}(T'_q, L'_q) = k_i \text{ for } i \in [n-1],$$

$$b_{\gamma_n'}(T_q', L_q') = \left\lfloor \frac{k_n}{2} \right\rfloor, \text{ and}$$
$$b_{\gamma_n''}(T_q', L_q') = \left\lceil \frac{k_n}{2} \right\rceil.$$

If k_n is even so that $b_{\gamma_n'}(T_q', L_q') = b_{\gamma_n''}(T_q', L_q')$, then Lemma 3.3.3 implies as above that the set of curves in L' which spiral into q consists of a conjugate pair λ', λ'' with the same weight w. Again, let λ be the (unique) non-ordinary allowable curve in $\mathcal{O}$ that terminates in q but otherwise coincides with λ' and λ'', and set $L = \tilde{L} \cup \{2w\lambda\}$. Application of Lemma 3.3.4 shows that L is a quasi-lamination on $\mathcal{O}$ and $\mathbf{b}(T, L) = (k_1, \ldots, k_n)$:

$$\begin{aligned}
b_{\gamma_i}(T, L) &= b_{\gamma_i}(T, \tilde{L}) + 2wb_{\gamma_i}(T, \lambda) \text{ for all } i \in [n-1] \\
&= b_{\gamma_i}(T_q', \tilde{L}) + 2wb_{\gamma_i}(T_q', \lambda') \\
&= b_{\gamma_i}(T_q', \tilde{L}) + wb_{\gamma_i}(T_q', \lambda') + wb_{\gamma_i}(T_q', \lambda'') \\
&= b_{\gamma_i}(T_q', L_q') \\
&= k_i. \\
b_{\gamma_n}(T, L) &= b_{\gamma_n}(T, \tilde{L}) + 2wb_{\gamma_n}(T, \lambda) \\
&= b_{\gamma_n'}(T, \tilde{L}) + b_{\gamma_n''}(T, \tilde{L}) + 2w\left(b_{\gamma_n'}(T_q', \lambda') + b_{\gamma_n''}(T_q', \lambda')\right) \\
&= b_{\gamma_n'}(T, \tilde{L}) + b_{\gamma_n''}(T, \tilde{L}) \\
&\quad + w\left(b_{\gamma_n'}(T_q', \lambda') + b_{\gamma_n''}(T_q', \lambda')\right) + w\left(b_{\gamma_n'}(T_q', \lambda'') + b_{\gamma_n''}(T_q', \lambda'')\right) \\
&= b_{\gamma_n'}(T, L_q') + b_{\gamma_n''}(T, L_q') \\
&= k_n.
\end{aligned}$$

If k_n is odd, then $|b_{\gamma_n'}(T_q', L_q') - b_{\gamma_n''}(T_q', L_q')| = 1$, and Lemma 3.3.3 implies that the set of curves in L' which spiral into q' consists of a conjugate pair λ', λ'' with $w_{\lambda''} = w_{\lambda'} + 1$. Once more, let λ be the (unique) non-ordinary allowable curve in $\mathcal{O}$ that terminates in q but otherwise coincides with λ' and λ'', and set $L = \tilde{L} \cup \{(w_{\lambda'} + w_{\lambda''})\lambda\}$. Application of Lemma 3.3.4 shows that L is again a quasi-lamination on $\mathcal{O}$ and $\mathbf{b}(T, L) = (k_1, \ldots, k_n)$:

$$\begin{aligned}
b_{\gamma_i}(T, L) &= b_{\gamma_i}(T, \tilde{L}) + (w_{\lambda'} + w_{\lambda''})b_{\gamma_i}(T, \lambda) \text{ for all } i \in [n-1] \\
&= b_{\gamma_i}(T, \tilde{L}) + w_{\lambda'}b_{\gamma_i}(T, \lambda) + w_{\lambda''}b_{\gamma_i}(T, \lambda) \\
&= b_{\gamma_i}(T_q', \tilde{L}) + w_{\lambda'}b_{\gamma_i}(T_q', \lambda') + w_{\lambda''}b_{\gamma_i}(T_q', \lambda'') \\
&= b_{\gamma_i}(T_q', L_q') \\
&= k_i. \\
b_{\gamma_n}(T, L) &= b_{\gamma_n}(T, \tilde{L}) + (w_{\lambda'} + w_{\lambda''})b_{\gamma_n}(T, \lambda)
\end{aligned}$$

$$
\begin{aligned}
&= b_{\gamma_n'}(T,\tilde{L}) + b_{\gamma_n''}(T,\tilde{L}) + (w_{\lambda'} + w_{\lambda''})\left(b_{\gamma_n'}(T_q',\lambda') + b_{\gamma_n''}(T_q',\lambda')\right) \\
&= b_{\gamma_n'}(T,\tilde{L}) + b_{\gamma_n''}(T,\tilde{L}) + w_{\lambda'}\left(b_{\gamma_n'}(T_q',\lambda') + b_{\gamma_n''}(T_q',\lambda')\right) \\
&\quad + w_{\lambda''}\left(b_{\gamma_n'}(T_q',\lambda'') + b_{\gamma_n''}(T_q',\lambda'')\right) \\
&= b_{\gamma_n'}(T,L_q') + b_{\gamma_n''}(T,L_q') \\
&= k_n.
\end{aligned}
$$

Finally, we show that L is unique, thereby proving injectivity in this case. Suppose L_1, L_2 are quasi-laminations on $\mathcal{O}$ with $\mathbf{b}(T,L_1) = \mathbf{b}(T,L_2)$. Then $\mathbf{b}(T,2L_1) = \mathbf{b}(T,2L_2)$, where for each $i = 1,2$, we let $2L_i$ denote the quasi-lamination $\{2w_\lambda\lambda : w_\lambda\lambda \in L_i\}$, so that $\mathbf{b}(T,2L_i) = 2\mathbf{b}(T,L_i)$. We build (unique) quasi-laminations L_1', L_2' on $\mathcal{O}_q'$ with $\mathbf{b}(T_q',L_1') = \mathbf{b}(T_q',L_2')$ as follows:

- If L_i does not contain a non-ordinary curve terminating in q, then take $L_i' = 2L_i$. Lemma 3.3.4 imply that L_i' is indeed a quasi-lamination on $\mathcal{O}_q'$, and that

$$
\begin{aligned}
b_{\gamma_i'}(T_q',L_i') &= b_{\gamma_i}(T,2L_i) \text{ for all } i \in [n-1]. \\
b_{\gamma_n'}(T_q',L_i') &= \frac{1}{2} b_{\gamma_n}(T,2L_i) = b_{\gamma_n}(T,L_i). \\
b_{\gamma_n''}(T_q',L_i') &= \frac{1}{2} b_{\gamma_n}(T,L_i) = b_{\gamma_n}(T,L_i).
\end{aligned}
$$

- If $L_i = \tilde{L}_i \cup \{w_\lambda\lambda\}$ for pending curve λ, take $L_i' = 2\tilde{L}_i \cup \{w_\lambda\lambda', w_\lambda\lambda'\}$. Lemma 3.3.4 implies that L_i' is once more a quasi-lamination on $\mathcal{O}_q'$, and that it has the same shear coordinates as the construction above:

$$
\begin{aligned}
b_{\gamma_i'}(T_q',L_i') &= b_{\gamma_i}(T_q',2\tilde{L}_i) + w_\lambda b_{\gamma_i}(T_q',\lambda') + w_\lambda b_{\gamma_i}(T_q',\lambda'') \text{ for all } i \in [n-1] \\
&= b_{\gamma_i}(T,2\tilde{L}_i) + w_\lambda b_{\gamma_i}(T,\lambda) + w_\lambda b_{\gamma_i}(T,\lambda) \\
&= b_{\gamma_i}(T,\tilde{L}_i) + 2w_\lambda b_{\gamma_i}(T,\lambda) \\
&= b_{\gamma_i}(T,2L_i) \\
b_{\gamma_n'}(T_q',L_i'). &= b_{\gamma_n'}(T_q',2\tilde{L}_i) + w_\lambda b_{\gamma_n'}(T_q',\lambda') + w_\lambda b_{\gamma_n'}(T_q',\lambda'') \\
&= \frac{1}{2} \cdot b_{\gamma_n}(T,2\tilde{L}_i) + w_\lambda b_{\gamma_n}(T,\lambda) \\
&= b_{\gamma_n}(T,L_i). \\
b_{\gamma_n''}(T_q',L_i') &= b_{\gamma_n'}(T_q',2\tilde{L}_i) + w_\lambda b_{\gamma_n'}(T_q',\lambda') + w_\lambda b_{\gamma_n'}(T_q',\lambda'') \\
&= \frac{1}{2} \cdot b_{\gamma_n}(T,2\tilde{L}_i) + w_\lambda b_{\gamma_n}(T,\lambda) \\
&= b_{\gamma_n}(T,L_i).
\end{aligned}
$$

- If $L_i = \tilde{L}_i \cup \{w_\lambda \lambda\}$ for semi-closed curve λ whose other end terminates at $r \in \mathbf{Q_{1/2}}$, take $L_i' = 2\tilde{L}_i \cup \{w_\lambda \lambda'''\}$. where λ''' is the closed curve in $\mathcal{O}_q'$ separating q' and r from all other points in $\mathbf{M} \cup \mathbf{Q}$. Lemma 3.3.4 implies that L_i' is once more a quasi-lamination on $\mathcal{O}_q'$, and that it has the same shear coordinates as the two cases above:

$$\begin{aligned}
b_{\gamma_i'}(T_q', L_i') &= b_{\gamma_i}(T_q', 2\tilde{L}_i) + w_\lambda b_{\gamma_i}(T_q', \lambda''') \text{ for all } i \in [n-1] \\
&= b_{\gamma_i}(T, 2\tilde{L}_i) + 2w_\lambda b_{\gamma_i}(T, \lambda) \\
&= b_{\gamma_i}(T, 2L_i). \\
b_{\gamma_n'}(T_q', L_i') &= b_{\gamma_n'}(T_q', 2\tilde{L}_i) + w_\lambda b_{\gamma_n'}(T_q', \lambda''') \\
&= \frac{1}{2} \cdot b_{\gamma_n}(T, 2\tilde{L}_i) + w_\lambda b_{\gamma_n}(T, \lambda) \\
&= b_{\gamma_n}(T, L_i). \\
b_{\gamma_n''}(T_q', L_i') &= b_{\gamma_n'}(T_q', 2\tilde{L}_i) + w_\lambda b_{\gamma_n'}(T_q', \lambda''') \\
&= \frac{1}{2} \cdot b_{\gamma_n}(T, 2\tilde{L}_i) + w_\lambda b_{\gamma_n}(T, \lambda) \\
&= b_{\gamma_n}(T, L_i).
\end{aligned}$$

By this construction, L_1 and L_2 are distinct quasi-laminations on $\mathcal{O}$ if and only if L_1' and L_2' are distinct quasi-laminations on $\mathcal{O}_q'$. However since $\mathbf{b}(T_q', L_1') = \mathbf{b}(T_q', L_2')$, our inductive assumption implies that $L_1' = L_2'$. Thus $L_1 = L_2$, and we conclude our desired result when $q \in \mathbf{Q_{1/2}}$, and therefore in general. □

3.3.3 The orbifold model of geometric cluster algebras

We conclude this section by summarizing the details of the orbifold model of geometric cluster algebras (see Section 2.2.1) and previewing the connection to mutation-linear algebra (see Section 2.2.2).

Let $\mathcal{O}$ be an orbifold. Define a ***multi-quasi-lamination*** $\mathbf{L}$ on $\mathcal{O}$ to be a collection of quasi-laminations $\mathbf{L} = (L_i)_{i \in I}$, indexed by some (possibly infinite) set I. For each tagged triangulation T of $\mathcal{O}$ (consisting of n arcs), define $\tilde{B}(T, \mathbf{L})$ to be the $([n] \cup I) \times [n])$ extended exchange matrix with principal (top) part given by the signed-adjacency matrix $B(T)$ of T, and coefficient rows (bottom part) given by $\mathbf{b}_i = \mathbf{b}(T, L_i)$ for each $i \in L$.

This definition is [5, Definition 6.4], a generalization of [9, Definition 12.5] to orbifolds, modified to use quasi-laminations and permit I to be infinite. (See Remarks 3.2.8, 3.2.11, and 3.3.5 for further details on the relationship between our approach and that in [5].) We thereby obtain the following two results: the first is [5, Lemma 6.5], a generalization of [9, Theorem 13.5] to orbifolds; the second is [5, Theorem 9.1 and Corollary 9.2], generalizations of [9, Theorem 15.6] to orbifolds.

Theorem 3.3.7. *Fix a quasi-lamination L and multi-quasi-lamination $\mathbf{L}$ on $\mathcal{O}$. If T and T' are tagged triangulations such that T' is obtained from T by a flip at an arc γ, then*

$$\mathbf{b}(T', L) = \eta_\gamma^{B(T)}(\mathbf{b}(T, L)) \text{ and} \tag{3.2}$$

$$\tilde{B}(T', \mathbf{L}) = \mu_\gamma(\tilde{B}(T, \mathbf{L})), \tag{3.3}$$

where $\eta_\gamma^{B(T)}$ and μ_γ are the mutation map and the matrix mutation map, respectively, at the index for the tagged arc γ in T.

Theorem 3.3.8. *Let $\mathcal{O}$ be an orbifold. Fix a tagged triangulation T_0 of $\mathcal{O}$ with n arcs and a multi-quasi-lamination $\mathbf{L} = (L_i)_{i \in I}$ on $\mathcal{O}$. (If $\mathcal{O}$ has no boundary and exactly one puncture, assume T_0 has all tags plain: see Remark 3.2.7.) There exists a unique cluster algebra of geometric type $\mathcal{A}(\tilde{B}(T_0, \mathbf{L}))$ with the following properties:*

- *the coefficient semifield is the tropical semifield $\mathbb{P}_{\mathbf{L}} = Trop(u_i : i \in I)$;*
- *the cluster variables $x_{\mathbf{L}}(\alpha)$ are labeled by the tagged arcs γ on $\mathcal{O}$ (resp. the tagged arcs with all tags plain if $\mathcal{O}$ has no boundary and exactly one puncture);*
- *the seeds $(\mathbf{x}(T), \tilde{B}(T, L))$ are labeled by the tagged triangulations T of $\mathcal{O}$ (resp. the triangulations with all arcs tagged plain if $\mathcal{O}$ has no boundary and exactly one puncture);*
- *each cluster $\mathbf{x}(T)$ consists of cluster variables $x_{\mathbf{L}}(\alpha)$ labeled by the tagged arcs $\alpha \in T$;*
- *the coefficient rows $\mathbf{b}_i(T)$ of $\tilde{B}(T, L)$ are the shear coordinate vectors $\mathbf{b}_i(T) = \mathbf{b}(T, L_i)$;*
- *if two vertices $t, t' \in \mathbb{T}_n$ are connected by an edge, then the tagged triangulations T and T' which label the associated seeds are related by an arc flip.*

Recall from Definition 2.2.11 that the cluster algebra $\mathcal{A}_\bullet(B(T_0))$ with principal coefficients is constructed by taking initial extended exchange matrix $\tilde{B}$ whose principal (top) part consists of the signed-adjacency matrix $B(T_0)$ and whose coefficient rows (the bottom part) are the $n \times n$ identity matrix. In light of Theorem 3.3.6, we may realize $\mathcal{A}_\bullet(B(T_0))$ as $\mathcal{A}(\tilde{B}(T_0, \mathbf{L}^{T_0}))$ for some particular choice of multi-quasi-lamination $\mathbf{L}^{T_0} = (L_\gamma)_{\gamma \in T_0}$ by choosing n quasi-laminations L_γ on $\mathcal{O}$ indexed by the tagged arcs $\gamma \in T_0$ such that $\mathbf{b}(T_0, L_\gamma)$ satisfies

$$b_\gamma(T_0, L_\gamma) = 1 \text{ and}$$
$$b_\alpha(T_0, L_\gamma) = 0 \text{ for all other arcs } \alpha \in T.$$

The quasi-laminations L_γ which comprise $\mathbf{L}^{T_0}$ do not in fact depend on the tagged triangulation T_0, just on the arcs γ, and are called ***elementary quasi-laminations***. They are constructed for marked surfaces in [9, Definition 17.2], and for orbifolds with no punctures and all orbifold points of weight 1/2 in [7, Section 3.3]. in which case they are in particular laminations in the sense of [9, Definition 12.1] and [5, Definition 4.1], respectively, as they consist of single (allowable) curves (of weight 1). The construction is easily extended to arbitrary orbifolds: however, we must permit double curves (see Remark 3.2.11).

Definition 3.3.9 (Elementary quasi-laminations)**.** Let α be a tagged arc in $\mathcal{O}$. The ***elementary quasi-lamination*** L_α associated to $\mathcal{O}$ consists of a single allowable curve (which, for simplicity, we also denote by L_α) that coincides with the arc γ except within (a) small neighborhood(s) around the endpoint(s) of α. The behavior of L_γ within each endpoint neighborhood is specified as follows:

- If the endpoint is a marked point p on a boundary component, then L_α terminates at an unmarked point x on the same boundary component, such that the path along the boundary from p to x, keeping $\mathcal{O}$ on the **right**, does not leave the small ball.
- If the endpoint is a puncture p, then L_α spirals into p, **counterclockwise** if α is tagged plain at p and **clockwise** if α is tagged notched at p.
- If the endpoint is an orbifold point q, then L_α also terminates at q.

Taking the opposite orientation throughout Definition 3.3.9 (in particular, substituting the three bold words by **left, clockwise**, and **counterclockwise**, respectively), we recover the map κ from [22, Section 5]. One can think of the allowable curves L_α and $\kappa(\alpha)$ as being related by a "reflection" about the arc α. Indeed, the shear coordinate vectors $\mathbf{b}(T_0, \kappa(\gamma))$ satisfy

$$\begin{aligned} &b_\gamma(T_0, \kappa(\gamma)) = -1 \text{ and} \\ &b_\alpha(T_0, \kappa(\gamma)) = 0 \text{ for all other arcs } \alpha \in T_0. \end{aligned} \tag{3.4}$$

The following lemma, an orbifold analog of [22, Lemma 5.1], is immediate.

Lemma 3.3.10. *The maps $\alpha \mapsto L_\alpha$ and $\alpha \mapsto \kappa(\alpha)$ from tagged arcs to allowable curves are both one-to-one and surjective onto the set of allowable curves that are neither closed nor semi-closed. Furthermore, the following statements are equivalent:*

- *Tagged arcs α and γ are compatible.*
- *Allowable curves L_α and L_γ are compatible.*
- *Allowable curves $\kappa(\alpha)$ and $\kappa(\gamma)$ are compatible.*

Remark 3.3.11. It is worth noting that Lemma 3.3.10 does *not* hold if one considers, rather than the allowable curves of Definition 3.2.10, the set of curves that can participate in laminations in the sense of [9, Definition 12.1] and [5, Definition 4.1]. In particular, the non-closed ordinary curves depicted in Figure 3.8 cannot be realized as $\kappa(\gamma)$ (or L_γ) for any tagged arcs γ in $\mathcal{O}$. This was a primary motivation for the introduction of allowable curves in [22], and the reason we continue to use them here.

Corollary 3.3.12. *If $B(T)$ is of finite type, then there are no allowable curves on $\mathcal{O}$ which are closed or semi-closed. Thus the maps $\alpha \mapsto L_\alpha$ and $\alpha \mapsto \kappa(\alpha)$ are bijections between the set of tagged arcs in $\mathcal{O}$ and the set of all allowable curves in $\mathcal{O}$.*

By Theorem 3.3.8, there is a bijection between the tagged arcs in $\mathcal{O}$ (resp. the tagged arcs in $\mathcal{O}$ with all tags plain if $\mathcal{O}$ has no boundary and exactly one puncture) and the cluster variables in $\mathcal{A}_\bullet(B(T_0)) = \mathcal{A}(\tilde{B}(T_0, \mathbf{L}^{T_0}))$. We write $\alpha \mapsto x_{\mathbf{L}^{T_0}}(\alpha)$ for this bijection. Recall from Definition 2.2.11 that each cluster variable $x_{\mathbf{L}^{T_0}}(\alpha)$ has an associated vector in $\mathbb{Z}^n$ called a $\mathbf{g}$-vector, which we may view as being indexed by the tagged arcs in T_0 via the bijection between tagged arcs in T_0 and cluster variables in the initial cluster $\mathbf{x}(T_0)$. We denote the $\mathbf{g}$-vector of $x_{\mathbf{L}^{T_0}}(\alpha)$ by $\mathbf{g}_{\mathbf{L}^{T_0}}(\alpha)$. The following result is [22, Proposition 5.2]. We show that an analogous result holds for orbifolds.

Proposition 3.3.13. *Fix a tagged triangulation T_0 of a marked surface $(\mathbf{S}, \mathbf{M})$ and let α be a tagged arc, not necessarily in T_0. (If $(\mathbf{S}, \mathbf{M})$ has exactly one puncture and no boundary components, then take all tags on T_0 and on α to be plain.) Then*

$$\mathbf{g}_{\mathbf{L}^{T_0}}(\alpha) = -\mathbf{b}(T_0, \kappa(\alpha)). \tag{3.5}$$

The proof of Proposition 3.3.13 relies on the following fact, which we state using the original notation of $\mathcal{A}_\bullet(B(T_0))$ from Definition 2.2.11 rather than the orbifold notation of its geometric realization via $\mathcal{A}(\tilde{B}(T_0, \mathbf{L}^{T_0}))$. (This result is a weak version of [13, Conjecture 7.12], which is true for all skew-symmetrizable exchange matrices as a consequence of [14].)

Proposition 3.3.14. *[22, Conjecture 5.3] Let $t_1, t_2 \in \mathbb{T}_n$ be two vertices connected by an edge labeled k, and let B_1 and B_2 be exchange matrices such that $B_2 = \mu_k(B_1)$. For any $t \in \mathbb{T}_n$ and $i \in [n]$, the $\mathbf{g}$-vectors $\mathbf{g}_{i;t}^{B_1;t_1}$ and $\mathbf{g}_{i;t}^{B_2;t_2}$ are related by*

$$\mathbf{g}_{i;t}^{B_2;t_2} = \eta_k^{B_1^T}(\mathbf{g}_{i;t}^{B_1;t_1}), \tag{3.6}$$

where B_1^T is the transpose of B_1.

The following is [7, Definition 8.3], rephrased with the adjustments of Remark 3.2.8.

Definition 3.3.15 (Reversed orbifold)**.** Let $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ be an orbifold with tagged triangulation T_0. The ***reversed orbifold*** $\mathcal{O}^* = (\mathbf{S}, \mathbf{M}, \mathbf{Q}^*)$ is obtained from $\mathcal{O}$ by swapping the weights on all orbifold points $q \in \mathbf{Q}$ (between 1/2 and 2). The corresponding triangulation T_0^* on $\mathcal{O}^*$ is obtained from T_0 by swapping the types of all non-ordinary arcs $\gamma \in T_0$ (between pending and double).

By [7, Remark 8.4], the signed-adjacency matrix of the triangulation T_0^* on the reversed orbifold $\mathcal{O}^*$ is the negative of the transpose of the signed adjacency matrix of T_0 on $\mathcal{O}$. That is,

$$B(T_0^*) = -B(T_0)^T. \tag{3.7}$$

Furthermore, it is easily verified using Tables 3.1 and 3.2 that for any tagged arcs $\alpha, \gamma \in T_0$ with corresponding arcs $\alpha^*, \gamma^* \in T_0^*$,

$$b_{\alpha^*}(T_0^*, \kappa(\gamma^*)) = b_\alpha(T_0, \kappa(\gamma)) = -\delta_{\alpha,\gamma}. \tag{3.8}$$

Remark 3.3.16. Since all signed adjacency matrices $B(T_0)$ of triangulations T_0 of marked surfaces are skew-symmetric, they satisfy the identity $B(T_0) = -B(T_0)^T$. The same is not true if T_0 is a triangulation of an orbifold $\mathcal{O}$. While the signed-adjacency matrix $B(T_0)$ is still skew-symmetrizable, it generally has a non-identity skew-symmetrizing matrix D (see Definition 2.2.3). However, $-B(T_0)^T$ is a ***rescaling*** of $B(T_0)$: that is, setting $B' = -B^T$, the respective matrix entries satisfy the following two conditions:

$$\operatorname{sgn}(b_{ij}) = \operatorname{sgn}(b'_{ij}) \text{ and } b_{ij}b_{ji} = b'_{ij}b'_{ji} \text{ for all } i, j \in [n]. \tag{3.9}$$

Applying (3.8), we conclude that $B(T_0^*)$ is a rescaling of $B(T_0)$. (See [21, Definition 7.4 and Proposition 7.5] for the definition of rescaling.)

Let B be an exchange matrix. The following is [21, Equation (7.1)]:

$$\eta_{\mathbf{k}}^B(\mathbf{a}) = -\eta_{\mathbf{k}}^{-B}(-\mathbf{a}). \tag{3.10}$$

The proof of Proposition 3.3.13 uses (3.10) to rewrite (3.6) as follows, in the case when $B_2 = B(T_2)$ and $B_1 = B(T_1)$ are the (skew-symmetric) signed adjacency matrices of triangulations T_2 and T_1 of a marked surface $(\mathbf{S}, \mathbf{M})$ and T_2 is obtained from T_1 by a flip of the arc γ:

$$-\mathbf{g}_{i;t}^{B_2;t_2} = \eta_k^{B_1}(-\mathbf{g}_{i;t}^{B_1;t_1}).$$

Or equivalently, for any tagged arc α in $(\mathbf{S}, \mathbf{M})$,

$$-\mathbf{g}_{\mathbf{L}^{T_2}}(\alpha) = \eta_\gamma^{B(T_1)}(-\mathbf{g}_{\mathbf{L}^{T_1}}(\alpha)).$$

It then applies (3.2) to show that the vectors $-\mathbf{g}_{L^{T_0}}(\alpha)$ and $\mathbf{b}(T_0, \kappa(\alpha))$ satisfy the same recurrence. Since (3.4) implies they are equal for tagged arcs $\alpha \in T_0$ corresponding to initial cluster variables, the result follows that they are equal for all tagged arcs.

Similarly we may apply (3.10) to rewrite (3.6) as follows, in the case when $B_2 = B(T_2)$ and $B_1 = B(T_1)$ are the (skew-symmetric) signed adjacency matrices of triangulations T_2 and T_1 of an *orbifold* $\mathcal{O}$, T_2 is obtained from T_1 by a flip of the arc γ, and $B(T_1^*) = \Sigma^{-1}B(T_1)\Sigma$:

$$-\mathbf{g}_{\mathbf{L}^{T_2}}(\alpha) = \eta_\gamma^{B(T_1^*)}(-\mathbf{g}_{\mathbf{L}^{T_1}}(\alpha)) \text{ for any tagged arc } \alpha \text{ in } \mathcal{O}.$$

By (3.2), the vectors $-\mathbf{g}_{L^{T_0}}(\alpha)$ and $\mathbf{b}(T_0^*, \kappa(\alpha^*))$ satisfy the same recurrence, and by (3.8) they are equal for tagged arcs $\alpha \in T_0$. We thereby obtain our desired orbifold analog of Proposition 3.3.13:

Proposition 3.3.17. *Fix a tagged triangulation T_0 of an orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ and let α be a tagged arc, not necessarily in T_0. (If $\mathcal{O}$ has exactly one puncture and no boundary components, then take all tags on T_0 and on α to be plain.) Then,*

$$\mathbf{g}_{\mathbf{L}^{T_0}}(\alpha) = -\mathbf{b}(T_0^*, \kappa(\alpha^*)). \tag{3.11}$$

Remark 3.3.18. Proposition 3.3.17 corresponds to [7, Lemma 8.6], where the curve $\kappa(\alpha^*)$ is denoted by L_γ^*, the "reversed elementary lamination" on $\mathcal{O}^*$. However, given the restricted outlook of most of [7] to orbifolds with no punctures and all orbifold points of weight 1/2, we felt it worthwhile to justify the conclusion directly. The result is asserted in [7] as a consequence of the duality between **c**-vectors and **g**-vectors of cluster algebras, which corresponds to (3.6).

In light of the results above, which relate **g**-vectors to shear coordinates by way of the action of mutation maps, it seems plausible that a relation exists between quasi-laminations on an orbifold $\mathcal{O}$ with respect to a fixed triangulation T and the mutation-linear algebra notions of Definition 2.2.2 associated to the signed-adjacency matrix $B(T)$. (Recall as well from Proposition 2.2.46 that if $B(T)$ is of finite type, the **g**-vector fan for the transpose of $B(T)$ coincides with the mutation fan for $B(T)$.) This connection is the subject of Section 3.4, and relies on the following object, which we show is closely related to the mutation fan $\mathcal{F}_{B(T)}$.

Definition 3.3.19 (Rational quasi-lamination fan)**.** Fix a tagged triangulation T of an orbifold $\mathcal{O}$. For each set Λ of pairwise compatible allowable curves in $\mathcal{O}$, define C_Λ be the nonnegative $\mathbb{R}$-linear span of the integer shear coordinate vectors $\{\mathbf{b}(T, \lambda) : \lambda \in \Lambda\}$. The ***rational quasi-lamination fan*** for T, denoted by $\mathcal{F}_\mathbb{Q}(T)$, is the collection of all such cones C_Λ. We call $\mathcal{F}_\mathbb{Q}(T)$ the ***rational quasi-lamination fan*** for T.

3.4 Orbifolds and mutation-linear algebra

The goal of this section is to relate quasi-laminations in an orbifold to the mutation-linear algebraic notions associated to the exchange matrix $B(T)$ arising from a tagged triangulation. Indeed, one might think of this section as an orbifold generalization of the primary results in [22], which considered mutation fans and universal geometric coefficients for cluster algebras arising from marked surfaces. To construct these respective objects using quasi-laminations, a marked surface had to have the Curve Separation Property and Null Tangle Property, respectively. We show that the same is true of orbifolds.

3.4.1 The Curve Separation Property

Let $\mathcal{O}$ be an orbifold with tagged triangulation T. Our main results are that if $\mathcal{O}$ has the Curve Separation Property, then as in the surfaces case, the rational quasi-lamination fan $\mathcal{F}_{\mathbb{Q}}(T)$ is the rational part of the mutation fan $\mathcal{F}_{B(T)}$ (in the sense of Definition 2.2.37 with $R = \mathbb{Q}$), and shear coordinates of allowable curves have the potential to form positive bases for $\mathbb{Z}^{B(T)}$ and $\mathbb{Q}^{B(T)}$. The implications of this conclusion are established at the end of the section, where we prove that the Curve Separation Property holds for a particular class of orbifolds which includes those of finite type. In fact, there are no orbifolds known not to have this property, and we plan to establish it for additional classes in future work.

What follows are the respective generalizations of [22, Definition 4.8, 4.5, 4.9, and 4.6], [22, Proposition 4.6], [22, Theorem 4.10], and [22, Theorem 4.12] to orbifolds. Note that the statements of the results are identical to those in [22], except that we have substituted the more general $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ for $(\mathbf{S}, \mathbf{M})$. We even place the same hypothesis on the triangulation T in the case when $\mathbf{S}$ has no boundary components and $\mathbf{M}$ consists of exactly one puncture. The necessity for this hypothesis is addressed in Remark 3.2.7.

Definition 3.4.1 (Curve Separation Property)**.** An orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ has the ***Curve Separation Property*** if for any incompatible allowable curves λ and ν in $\mathcal{O}$, there exists a tagged triangulation T (with all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture: see Remark 3.2.7) and an arc γ in T such that the shear coordinates $b_\gamma(T, \lambda)$ and $b_\gamma(T, \nu)$ of λ and ν at γ have strictly opposite signs.

Proposition 3.4.2. *Suppose $\mathcal{O}$ satisfies the Curve Separation Property, let T be any tagged triangulation of $\mathcal{O}$ (with all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture) and let Λ be a set of allowable curves. Then the curves in Λ are pairwise compatible if and only if the vectors $\{\mathbf{b}(T, \lambda) : \lambda \in \Lambda\}$ are contained in some $B(T)$-cone.*

Proof. We argue in an analogous manner as in the proof of [22, Proposition 4.6]. Suppose that two allowable curves $\lambda, \nu \in \Lambda$ are not compatible. Then the Curve Separation Property implies that there exists a tagged triangulation T' of $\mathcal{O}$ (with all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture) and an arc $\gamma \in T'$ such that $b_\gamma(T', \lambda)$ and $b_\gamma(T', \nu)$ have strictly opposite signs. By Remark 3.2.7, T' may be obtained from T by a sequence of tagged arc flips, so by Theorem 3.3.7, there exists some mutation map $\eta_{\mathbf{k}}^{B(T)}$ such that $\eta_{\mathbf{k}}^{B(T)}(\mathbf{b}(T, \lambda)) = \mathbf{b}(T', \lambda)$ and $\eta_{\mathbf{k}}^{B(T)}(\mathbf{b}(T, \nu)) = \mathbf{b}(T', \nu)$. It then follows from Proposition 2.2.31 that because of the strict sign difference between the coordinates $b_\gamma(T', \lambda)$ and $b_\gamma(T', \nu)$, the vectors $\mathbf{b}(T, \lambda)$ and $\mathbf{b}(T, \nu)$ are not contained in any common $B(T)$-cone.

Conversely, suppose that all allowable curves in Λ are pairwise compatible. Let $\lambda, \nu \in \Lambda$, let T' be any tagged triangulation of $\mathcal{O}$, and let γ be any tagged arc in T'. We show that $\mathbf{b}_\gamma(T', \lambda)$ and $\mathbf{b}_\gamma(T', \lambda)$ weakly agree in sign, and then invoke Theorem 3.3.7 and Proposition 2.2.31 to conclude that $\mathbf{b}_\gamma(T, \lambda)$ and $\mathbf{b}_\gamma(T, \nu)$ must be contained in some common $B(T)$-cone. Referring to Definition 3.3.1 and comparing the 'Positive Intersection' and 'Negative Intersection' columns of Tables 3.1 and 3.2, we see that the only way for two curves to have shear coordinates of strictly opposite signs at a given arc is if they intersect, either by crossing or sharing an orbifold point. Since λ and ν are distinct and compatible, by Definition 3.2.10 they can only intersect if they form a conjugate pair of curves. But in this case Lemma 3.3.3 says $\mathbf{b}_\gamma(T', \lambda)$ and $\mathbf{b}_\gamma(T', \lambda)$ must still weakly agree in sign, as desired. □

Theorem 3.4.3. *Let T be any tagged triangulation of orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ (with all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture). The collection $\mathcal{F}_{\mathbb{Q}}(T)$ is a rational, simplicial fan. Furthermore, $\mathcal{F}_{\mathbb{Q}}(T) = \mathcal{F}_{B(T)}^{\mathbb{Q}}$, the rational part of $\mathcal{F}_{B(T)}$, if and only if $\mathcal{O}$ has the Curve Separation Property.*

Proof. Again, our proof follows that of the surfaces result we generalize, [22, Theorem 4.10], first showing that the rational cones in $\mathcal{F}_{\mathbb{Q}}(T)$ are simplicial, then showing they form a fan, and finally establishing the equivalence between the Curve Separation Property and $\mathcal{F}_{\mathbb{Q}}(T)$ being the rational part of $\mathcal{F}_{B(T)}$. Assume T consists of n arcs.

To begin, suppose some cone $C_\Lambda \in \mathcal{F}_{\mathbb{Q}}(T)$ is not simplicial, where we recall from Definition 3.3.19 that Λ is a collection of pairwise compatible allowable curves in $\mathcal{O}$. Then the shear coordinate vectors $\{\mathbf{b}(T, \lambda) : \lambda \in \Lambda\}$ which span C_Λ over $\mathbb{R}_+$ are not linearly independent. In particular, because they are integer vectors, the set is not linearly independent over $\mathbb{Q}$, so there exists some rational point in C_Λ which can be written as a $\mathbb{Q}_+$-linear combination of vectors in $\{\mathbf{b}(T, \lambda) : \lambda \in \Lambda\}$ in two distinct ways. This point is therefore obtained as the shear coordinate vector of two distinct rational quasi-laminations supported on a subset of Λ, contradicting Theorem 3.3.6.

Next, consider any two cones $C_\Lambda, C_{\Lambda'} \in \mathcal{F}_{\mathbb{Q}}(T)$. We show $C_\Lambda \cap C_{\Lambda'}$ is a cone in $\mathcal{F}_{\mathbb{Q}}(T)$ as well, namely, $C_{\Lambda \cap \Lambda'}$, to conclude that $\mathcal{F}_{\mathbb{Q}}(T)$ is a simplicial fan; it is rational by construction. By definition, $C_\Lambda \cap C_{\Lambda'} \supseteq C_{\Lambda \cap \Lambda'}$. For the reverse containment, let $\mathbf{a} \in C_\Lambda \cap C_{\Lambda'} \cap \mathbb{Q}^n$. By

Theorem 3.3.6, $\mathbf{a}$ is the shear coordinate vector of some unique rational quasi-lamination which is therefore supported on a subset $\tilde{\Lambda} \subseteq \Lambda \cap \Lambda'$ of pairwise compatible allowable curves. It follows that $C_\Lambda \cap C_{\Lambda'} \cap \mathbb{Q}^n \subseteq C_{\Lambda\cap\Lambda'}$. Since $C_\Lambda \cap C_{\Lambda'}$ is a rational cone, we thereby conclude that $C_\Lambda \cap C_{\Lambda'} = C_{\Lambda\cap\Lambda'}$.

Now, suppose $\mathcal{O}$ does not have the Curve Separation Property. Then there exist incompatible allowable curves λ, ν in $\mathcal{O}$ such that the shear coordinates $\mathbf{b}(T', \lambda)$ and $\mathbf{b}(T', \nu)$ have weakly the same sign for all triangulations T'. This implies by Theorem 3.3.7 that $\eta_{\mathbf{k}}^{B(T)}(\mathbf{b}(T, \lambda))$ and $\eta_{\mathbf{k}}^{B(T)}(\mathbf{b}(T, \lambda))$ are of weakly the same sign for every sequence $\mathbf{k}$ of indices in $[n]$, which in turn implies by Proposition 2.2.31 that they are contained some common $B(T)$-cone. But since λ and ν are incompatible, $\mathbf{b}(T, \lambda)$ and $\mathbf{b}(T, \nu)$ span rays of $\mathcal{F}_\mathbb{Q}(T)$ which are not contained in a common cone of $\mathcal{F}_\mathbb{Q}(T)$. This contradicts Condition (iii) of Definition 2.2.37, so $\mathcal{F}_\mathbb{Q}(T)$ is not the rational part of a fan.

Conversely, suppose $\mathcal{O}$ does have the Curve Separation Property. Since $\mathcal{F}_\mathbb{Q}(T)$ is a rational simplicial fan, it satisfies Condition (i) of Definition 2.2.37, and Condition (ii) holds by Proposition 3.4.2. To establish Condition (iii), let C be an arbitrary $B(T)$-cone and let Λ denote the set of allowable curves λ whose shear coordinates $\mathbf{b}(T, \lambda)$ are contained in C. By definition, C_Λ is contained in C since C is a convex cone, and by Proposition 3.4.2, the curves in Λ are pairwise compatible, so C_Λ is a cone in $\mathcal{F}_\mathbb{Q}(T)$. Furthermore, if $C_{\Lambda'}$ is any other cone contained in C, then $\Lambda' \subseteq \Lambda$ by construction, so $C_{\Lambda'} \subseteq C_\Lambda$ and C_Λ is the unique largest cone among cones of $\mathcal{F}_\mathbb{Q}(T)$ contained in C.

We conclude by showing that C_Λ contains every rational point in C using induction on the dimension of C. The assertion holds trivially when C is the zero cone: assume it holds when C is of dimension no greater than k for some $k \in \mathbb{Z}_{\geq 0}$, and suppose C is of dimension $k+1$. Let $\mathbf{a} \in \mathbb{Q}^n \cap C$. By Theorem 3.3.6, $\mathbf{a} = \mathbf{b}(T, L)$ for some rational quasi-lamination L in $\mathcal{O}$. Let Λ' denote the support of L: by Condition (ii) of Definition 2.2.37, which we have already established, $C_{\Lambda'}$ is contained in some $B(T)$-cone C', and $\mathbf{a} \in C \cap C'$. If $C' = C$, then by the first part of Condition (iii), established above, $a \in C_{\Lambda'} \subseteq C_\Lambda$. Otherwise, $C \cap C'$ is a $B(T)$-cone since $\mathcal{F}_B$ is a fan, and has dimension no greater than k. Again invoking the first part of Condition (iii), let $C_{\tilde{\lambda}}$ be the largest cone of $\mathcal{F}_\mathbb{Q}(T)$ contained in $C \cap C'$. By induction, $\mathbf{a} \in C_{\tilde{\lambda}}$, and since $C_{\tilde{\lambda}} \subseteq C$, we again conclude that $\mathbf{a} \in C_\Lambda$, as desired. $\mathbf{a} \in C_{\tilde{\lambda}} \subseteq C_\Lambda$. □

For the next result, which establishes the implications of the Curve Separation Property on the connection between shear coordinate vectors and positive bases, we do not provide our own rephrasing of the proof, rather referring the reader to the original proof in the surfaces case in [22, Theorem 4.12]. However we note the necessity of invoking the following orbifold results instead of their respective surface analogs:

- Theorem 3.4.3 rather than [22, Theorem 4.10]

- Theorem 3.3.6 rather than [22, Theorem 4.4]

Note that [22, Proposition 2.10] and [22, Proposition 4.13] are general mutation-linear algebra results from [21]; they are quoted here as Proposition 2.2.36 and Proposition 2.2.40, respectively.

Theorem 3.4.4. *Let T be a tagged triangulation of $\mathcal{O}$ (with all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture) and let R be $\mathbb{Z}$ or $\mathbb{Q}$. If $\mathcal{O}$ has the Curve Separation Property, then*

1. *If the collection of shear coordinate vectors of allowable curves in $\mathcal{O}$ are an independent set for $R^{B(T)}$, then they form a positive basis for $R^{B(T)}$.*

2. *If a positive basis for $R^{B(T)}$ exists, then the shear coordinate vectors of allowable curves in $\mathcal{O}$ are a positive basis.*

We finish this section by proving the following theorem, which is part of current work to show that an orbifold $\mathcal{O}$ has the Curve Separation Property if it has one or more boundary components or two or more punctures. The expected result generalizes [22, Theorem 6.1] to orbifolds with the same set of hypotheses. In particular, the hypotheses rule out the case where $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ has no boundary components and exactly one puncture, regardless of the number of orbifold points. One additional case of the Curve Separation Property is established in [20].

Theorem 3.4.5. *If an orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ has either no punctures and no orbifold points, a unique puncture and no orbifold points, or no punctures and a unique orbifold point, then $\mathcal{O}$ has the Curve Separation Property.*

Our proof of Theorem 3.4.5 follows [22, Section 6] closely, and we will appeal often to statements, constructions, and figures therein. The reader may find it helpful to have a copy open for easy reference. Since we will be dealing with shear coordinates, the reader may also appreciate having Tables 3.1 and 3.2 close at hand.

What follows is a summary of the methods and observations set forth in [22, Section 6] which we employ here. To prove that $\mathcal{O}$ has the Curve Separation Property, we take as given two incompatible allowable curves λ and ν in $\mathcal{O}$, and construct a tagged triangulation T containing tagged arc γ such that the shear coordinates $b_\gamma(T, \lambda)$ and $b_\gamma(T, \nu)$ of λ and ν, respectively, with respect to γ have strictly opposite signs. Since allowable curves may not have any self-intersections, all contributions to these shear coordinates from intersections with γ have the same sign. Thus it suffices to find one intersection of λ with γ and one intersection of ν with γ that contribute opposite signs. We consider many possible cases on λ and ν: in most cases, we construct, rather than a tagged triangulation T, an (untagged) triangulation T^0, in which case we tacitly take $T = \tau(T^0)$ (see Definition 3.2.6). We also in general construct T or T^0 such that $b_\gamma(T, \lambda) > 0$ while $b_\gamma(T, \lambda) < 0$; however, as in [22], at times we have to switch this convention.

In many cases we include a figure to illustrate the salient part of the triangulation T or T^0: these figures follow the color and shading conventions established in [22, Section 6]. In particular, we continue to use the red and blue arcs and light-gray shading from Tables 3.1 and 3.2, and color the arc γ purple and the curve λ solid and black. When it appears, the curve ν is dashed and green. Areas beyond the boundary are shaded dark gray, and we indicate that a loop is not contractible by shading inside with a light gray striped pattern. Often, arcs in T or T^0 are constructed by ***following closely*** along existing curves: a curve follows another closely if there are no punctures, orbifold points, or boundary components between them. For legibility in the figures, we often straighten such curves somewhat rather than drawing them tightly next to one other. Finally, note that certain points and arcs which appear distinct in the figures might possibly coincide: these possibilities are explicitly addressed in the proof, and often appeal to the assumption, in Definition 3.2.1, that the surface $\mathbf{S}$ in $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ is oriented.

The following three cases are considered: these are directly comparable to the first three cases considered in [22]. We have generalized them to our setting by replacing the original "spiral point" with the more general "endpoint not on a boundary segment", which is equivalent, for orbifolds, to "spiral point or orbifold point." Due to this generalization, we must consider additional subcases. However, we maintain the numbering used in [22].

Case 1: Both endpoints of λ are on boundary segments.

Case 2: λ has exactly one endpoint on a boundary segment.

Case 3: λ is a closed curve.

We make a few general statements which will be used throughout.

Remark 3.4.6. As in [22], we must show that the given allowable curves λ and ν must interact with our constructed triangulation T^0 in a specific manner. Namely, we often argue that a given curve and arc must intersect. We explicitly use 'intersect' rather than 'cross' in light of Figure 3.10, in order to encompass the case when a curve ν and an arc α share an common orbifold point. For this same reason, when citing similar arguments from [22] generalized to our orbifold setting, we implicitly read 'crossing' as 'intersecting.'

Remark 3.4.7. Recall from Definition 3.2.3 that arcs in an orbifold are forbidden from defining a digon containing no points in $\mathbf{M} \cup \mathbf{Q}$, from bounding a monogon containing no points in $\mathbf{M} \cup \mathbf{Q}$, and from bounding a monogon containing a unique orbifold point and no punctures. (Note that (untagged) arcs *are* permitted to bound once-punctured monogons, as (untagged) triangulations may contain self-folded triangles).

The first two configurations are also forbidden for arcs in surfaces; thus the arguments we cite from [22] explicitly deal with these possibilities when constructing T^0. We handle the third

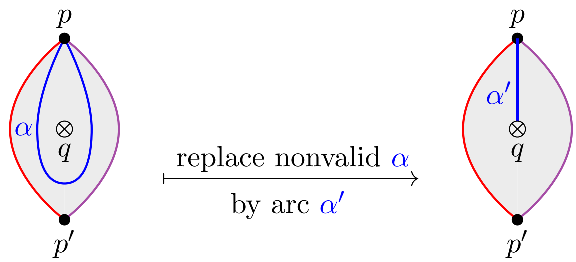

Figure 3.12: An illustration for Remark 3.4.7: handling an arc α which, as originally constructed, bounds a monogon containing a unique orbifold point q and no punctures.

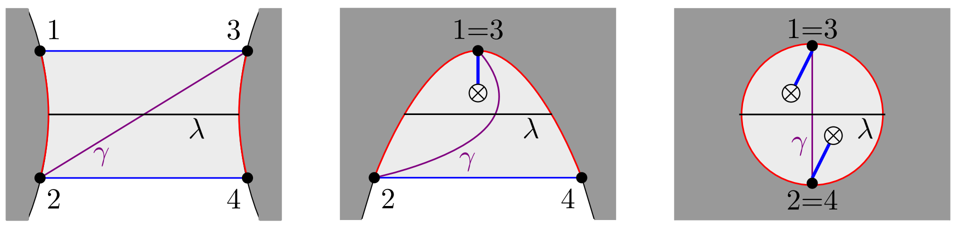

Figure 3.13: Illustrations for **Case 1**a

forbidden configuration, where a constructed arc α bounds a ***"once-orbifolded" monogon***, as follows: Since the endpoints coincide, the common endpoint must be a marked point $p \in \mathbf{M}$. Replace α by the non-ordinary arc α', in the monogon, which connects p and q. This operation is illustrated in Figure 3.12: the result is a valid orbifold triangle (see Figure 3.6).

We now argue that if a curve crossed α, then it must intersect the newly defined arc α'. This allows us to cite arguments from [22] about the necessity of a given curve interacting with the constructed triangulation T^0 in a specific way even if we have replaced α by α'. This necessity is clear when we recall from Definition 3.2.3 that arcs are considered up to isotopy relative to $\mathbf{M} \cup \mathbf{Q}$. Referring, for example, to Figure 3.12, it the blue curve α in the lefthand triangle is crossed by some other curve and this crossing cannot be removed by an isotopic deformation relative to p, q, and p', then the other curve must either cross between p and q or else terminate at q. Regardless, the same curve intersects the blue arc α' in the righthand triangle.

Case 1. Both endpoints of λ are on boundary segments.

Case 1a. *λ has endpoints on two distinct boundary segments.* Proceed as in [22, Case 1a]. We need only rule out the possibilities that either or both of the blue arcs bound once-orbifolded monogons, which we do by appealing to Remark 3.4.7. See Figure 3.13.

Case 1b. *λ has both endpoints on the same boundary segment.* Proceed as in [22, Case 1b], invoking Remark 3.4.7 if the first blue arc constructed bounds a once-orbifolded monogon.

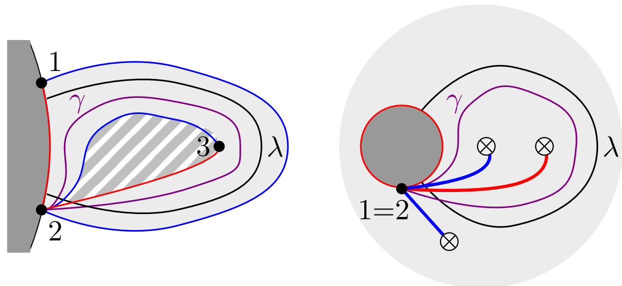

Figure 3.14: Illustrations for **Case 1**b

Definition 3.2.10 implies that λ is not contractible to the boundary, and furthermore because λ cannot be one of the forbidden curves of type illustrated in Figure 3.8, the purple arc γ cuts out neither a once-punctured nor a once-orbifolded monogon. Thus we may construct a triangulation T^0 in which there is a second triangle with the purple arc γ as an edge. In general this triangle is ordinary as in the lefthand image of Figure 3.14; however, it may be non-ordinary. For example, see the righthand image in Figure 3.14, which depicts the following case: $\mathcal{O}$ is a sphere with one boundary component, a single marked point on the boundary, and 3 orbifold points, and λ, together with a portion of the boundary, cuts out a digon with two orbifold points. This case requires the application of Remark 3.4.7 as well as the necessity of the purple arc γ being contained in both an ordinary and non-ordinary triangle. However, the arguments of [22] regarding ν still apply, as all non-ordinary triangles are composed of three distinct arcs.

Case 2. λ has exactly one endpoint on a boundary segment.

Case 2a. *The other endpoint of λ is a spiral point.* Proceed as in [22, Case 2], invoking Remark 3.4.7 if the first blue arc constructed bounds a once-orbifolded monogon. The purple arc γ does not bound a once-orbifolded monogon. Notice, as in [22], that the only curve which intersects λ but is also compatible with it is the curve with whom λ forms a conjugate pair, and its shear coordinate with respect to γ is zero. See Figure 3.15.

Case 2b. *The other endpoint of λ is an orbifold point.* Begin as in [22, Case 2] by coloring the boundary segment red and drawing the first blue arc (invoking Remark 3.4.7 if it bounds a once-orbifolded monogon). Label the orbifold point as 3. Draw a purple non-ordinary arc γ from point 2 to 3, following the red segment closely from point 2 to λ and then following λ closely to 3. Complete these arcs to a triangulation T^0. As in [22], there are two possibilities for ν. The first possibility is that some portion of ν intersects the blue arc (or originates on the blue boundary segment) and then crosses the purple arc γ before intersecting the outer blue

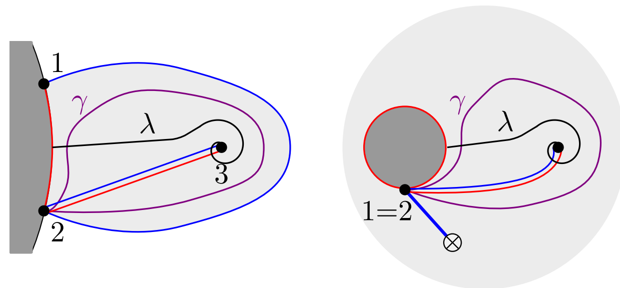

Figure 3.15: Illustrations for **Case 2**a

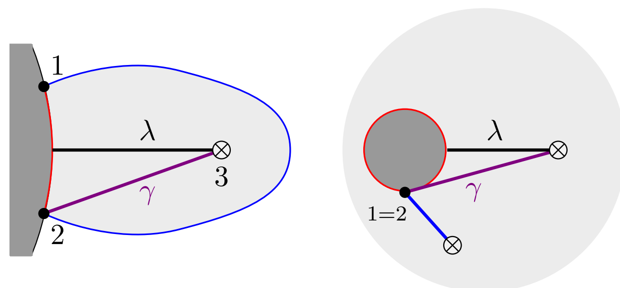

Figure 3.16: Illustrations for **Case 2**b

arc again (or terminating on the blue boundary segment). The second possibility is that after ν intersects the blue arc (or originates on the blue boundary segment) it terminates at point 3, the shared orbifold endpoint of λ and the purple arc γ. Regardless, ν picks up a negative shear coordinate contribution at γ while the shear coordinate of λ is positive. See Figure 3.16.

Case 3. λ is a closed curve.

Case 3a. *There exists a marked point left of λ and a marked point right of λ.* Proceed as in [22, Case 3a], invoking Remark 3.4.7 if the inner and/or outer blue arc bounds a once-orbifolded monogon. (Possibly points 1 and 2 coincide, in which case the blue arcs may coincide as well. However, this can only happen if the blue arc(s) are ordinary: otherwise, $\mathcal{O}$ is a sphere with one puncture and one or two orbifold points, neither of which is permitted under Definition 3.2.1.) See Figure 3.17, which also illustrates some possibilities for ν.

Case 3b. *λ is a closed curve not falling into Case 3a and ν is ordinary.* By Definition 3.2.1, $\mathcal{O}$ contains at least one marked point. Thus we may proceed as in [22, Case 3b], supposing, without loss of generality, that all marked points are left of λ, so that $\mathcal{O}$ decomposes at λ as a connected sum. In Figure 3.18, we depict the "left" summand, containing marked points, outside of λ and the "right" summand, containing no marked points, inside λ. Since λ is not contractible, the

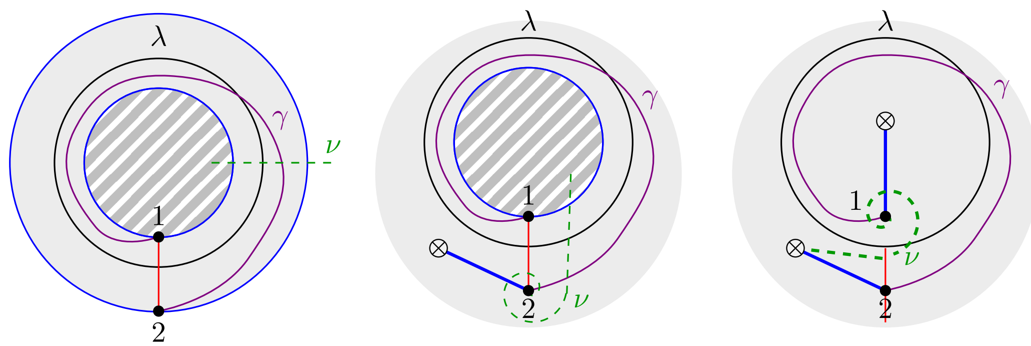

Figure 3.17: Illustrations for **Case 3**a

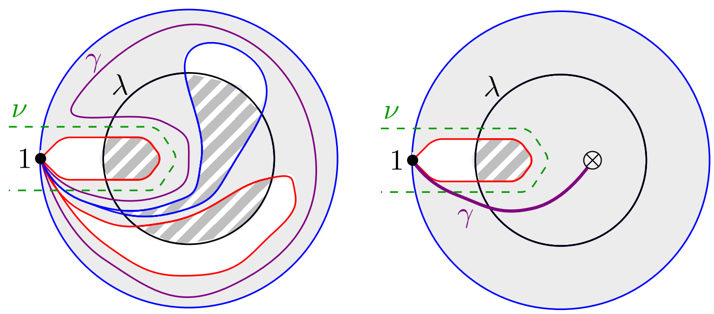

Figure 3.18: Illustrations for **Case 3**b

right summand is either topologically nontrivial or is a disc with at least two orbifold points (at least three if all are of weight 1/2: see Definition 3.2.10). If any of the blue, red, and purple arcs bound once-orbifolded monogons, apply Remark 3.4.7. These arcs can be completed to a triangulation T^0. In particular, if the purple arc did not bound a once-orbifolded monogon, then there exists a triangulation T^0 in which there is a second, non-self-folded triangle having the purple arc γ as an edge. The argument is then completed in [22] (see the lefthand image in Figure 3.18). If the originally-constructed purple arc *did* bound a once-orbifolded monogon, then the purple arc γ in T^0 is non-ordinary and therefore contained in a unique triangle (see the righthand image in Figure 3.18). It is significantly easier to compute shear coordinates in this case, and again ν picks up a negative shear coordinate contribution with respect to γ while the shear coordinate of λ is positive. (The case when the blue and/or red arcs bound once-orbifolded monogons are not pictured, but are easy to imagine.)

We remark that while surface examples $\mathcal{O} = (\mathbf{S}, \mathbf{M})$ of Case 3b require that $\mathbf{S}$ has some nontrivial topology, this is not the case for orbifolds. In particular, consider the digon with two orbifold points of weight 2 depicted in Figure 3.19.

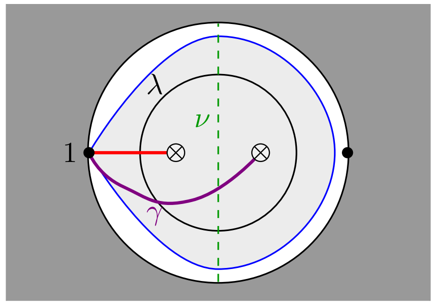

Figure 3.19: An orbifold example of **Case 3**b

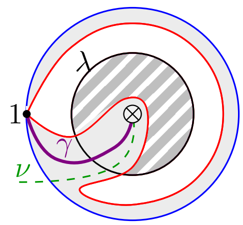

Figure 3.20: Illustration for **Case 3**c

Case 3c. λ *is a closed curve not falling into Case 3a and* ν *is non-ordinary.* As in Case 3b, we may again suppose, without without loss of generality, that all marked points are left of λ, so that $\mathcal{O}$ decomposes at λ as a connected sum, depicted in Figure 3.20 in a similar manner to Figure 3.18 with the same assumptions on the "left" and "right" summands. Choose some marked point on the left and label it as 1. If both endpoints of ν lie in the left summand, then we construct the triangulation described in Case 3b and use the same argument to complete the proof. Otherwise, since ν and λ are incompatible, somewhere along its extent, ν crosses λ to enter the right summand, and then terminates at an orbifold point in the right summand. Label this orbifold point as 2. As in the previous cases, begin by drawing a blue arc from 1 to itself that closely follows λ, replacing it with a blue boundary segment if it defines an unpunctured digon, and invoking Remark 3.4.7 if it bounds a once-orbifolded monogon. Next, draw a purple arc γ from point 1 to point 2 that closely follows the blue arc to without crossing it to its intersection with ν, then closely follows ν without crossing it to point 2. Complete the triangle by drawing a red arc from point 1 to itself as in Figure 3.20 that closely follows γ without crossing it to point 2, then turns to curve around point 2 and closely follow γ without crossing it to λ, then turns to closely follow λ without crossing it until reaching the earlier part of the red arc, then following the red arc closely back to point 1. The red arc cannot bound an unpunctured monogon, as then λ is not allowable by Definition 3.2.10. If the red arc bounds a once-orbifolded monogon, then we invoke Remark 3.4.7. Extend to a triangulation T^0.

3.4.2 The Null Tangle Property

We define the Null Tangle Property in the orbifolds setting, and show, as for surfaces, that it implies the Curve Separation Property. Furthermore, the strongest consequence of the Null Tangle Property for surfaces, [22, Theorem 7.3], is easily shown to hold for orbifolds as well. That is, if an orbifold has the Null Tangle Property, then the shear coordinates of allowable curves with respect to any tagged triangulation T form a positive basis for $R^{B(T)}$, and therefore universal geometric coefficients for $B(T)$ over R. Finally, in the previous section we established the Curve Separation Property for all but one class of orbifolds in Theorem 3.4.5. As discussed in [22], while the complexity of the proof and the fact that the Null Tangle Property implies the Curve Separation property suggest that it may be quite difficult to establish the former in general, we extend [22, Theorem 7.4] by proving the Null Tangle Property holds for a family of orbifolds which includes those of classical finite and affine type.

The following two definitions are comparable to [22, Definition 7.1 and Definition 7.2]. The propositions that follow, which connect null tangles to $B(T)$-coherent relations (see Definition 2.2.19), the Curve Separation Property (see Definition 3.4.1), and finally to bases for $R^{B(T)}$ (see Remark 2.2.24), are comparable, respectively, to [22, Propositions 7.9-7.13], [22, Corollary 7.14], and [22, Theorem 7.3], with the correction for null orbifolds given in [24, Theorem 3.14] and discussed in [24, Remark 3.15]. The proofs thereof are nearly identical, and we note any necessary adjustments.

Definition 3.4.8 (Weighted tangle of curves)**.** A ***(weighted) tangle (of curves)*** Ξ in an orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ is a finite collection of allowable curves λ (see Definition 3.2.10) which are distinct up to isotopy, each with an integer ***weight*** $w_\lambda \in \mathbb{Z}$. (In contrast with the quasi-laminations of Definition 3.2.10, we do not require the curves in a tangle to be compatible, nor do we require that they have positive weights.) The set of curves in Ξ with nonzero weights is called the ***support*** of the tangle. A ***straight tangle*** is a tangle in which all curves are pairwise compatible. (Again, this is a different object than a quasi-lamination as weights of curves in a straight tangle are still permitted to be negative). A tangle is ***trivial*** if all curves have weight zero. The ***weighted union*** of tangles Ξ_1 and Ξ_2 is built from the multiset union of Ξ_1 and Ξ_2 by doing the following: for each curve λ which appears in both Ξ_1 (with weight w_λ^1) and Ξ_2 (with weight w_λ^2), replace the two copies of λ in the union with a single copy of weight $w_\lambda^1 + w_\lambda^2$. The ***disorder*** of a tangle Ξ is the smallest number k such that, after deleting all curves of weight zero, Ξ is a weighted union of k straight tangles.

Definition 3.4.9 (Shear coordinates and the Null Tangle Property)**.** The shear coordinates $\mathbf{b}(T, \Xi) = (b_\gamma(T, \Xi))_{\gamma \in T}$ of a tangle Ξ with respect to a tagged triangulation T of $\mathcal{O}$ are computed in an analogous manner to those of a quasi-lamination in Definition 3.3.1, by setting

$$b_\gamma(T,\Xi) = \sum_{\lambda\in\Xi} w_\lambda b_\gamma(T,\lambda) \text{ for each arc } \gamma \in T.$$

If $\mathcal{O}$ has one or more boundary components or two or more punctures (or both), then Ξ is a ***null tangle*** if $\mathbf{b}(T,\Xi) = \mathbf{0}$ for tagged triangulations T. Otherwise, $\mathcal{O}$ has no boundary components and exactly one puncture, and we require only that $\mathbf{b}(T,\Xi) = \mathbf{0}$ for all tagged triangulations T in which all arcs are tagged plain. The orbifold has the ***Null Tangle Property*** if the only null tangles in $\mathcal{O}$ are trivial tangles.

Proposition 3.4.10. *Suppose that no component of $\mathcal{O}$ is a null orbifold. A tangle Ξ in $\mathcal{O}$ is a null tangle if and only if for any tagged triangulation T of $\mathcal{O}$ (with all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture), the linear relation $\sum_{\lambda\in\Xi} w_\lambda \mathbf{b}(T,\lambda)$ is $B(T)$-coherent.*

Proof. By Theorem 3.3.7, a tangle Ξ is null if and only if $\sum_{\lambda\in\Xi} w_\lambda \mathbf{b}(T,\lambda)$ satisfies the first of the two requirements, (2.6), of $B(T)$-coherent linear relations (see Definition 2.2.19). The remainder of the proof is identical to that of [22, Proposition 7.9]: since the signed adjacency matrix of a triangulation of an orbifold which has no null component has no row consisting entirely of zeros, [22, Proposition 2.3] implies that (2.6) is sufficient for $B(T)$-coherence. Thus Ξ is null if and only if $\sum_{\lambda\in\Xi} w_\lambda \mathbf{b}(T,\lambda)$ is $B(T)$-coherent. ☐

Proposition 3.4.10 implies the following rephrasing of the Null Tangle Property in terms of independence in $R^{B(T)}$:

Proposition 3.4.11. *Let T be a tagged triangulation of $\mathcal{O}$ (with all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture) and let R be $\mathbb{Z}$ or $\mathbb{Q}$. Then the shear coordinates of allowable curves in $\mathcal{O}$ constitute an independent set for $R^{B(T)}$ if and only if every component of $\mathcal{O}$ either has the Null Tangle Property or is a null orbifold.*

Remark 3.4.12. While Proposition 3.4.10 is true for any choice underlying ring $R \in \{\mathbb{Z},\mathbb{Q},\mathbb{R}\}$, Proposition 3.4.11 does not hold for $R = \mathbb{R}$. Recall from Definition 2.2.22 that the set A of shear coordinates of allowable curves is independent in $R^{B(T)}$ if and only if every $B(T)$-coherent linear relation among them with coefficients in R is trivial. Clearly when $R = \mathbb{Z}$, Definition 3.4.8 implies that every linear relation among elements of A is of the form $\sum_{\lambda\in\Xi} w_\lambda \mathbf{b}(T,\lambda)$ for some tangle Ξ, and Proposition 3.4.10 then says that such a relation is $B(T)$-coherent if and only if Ξ is null. Furthermore, when $R = \mathbb{Q}$, if $\sum_{\lambda\in A'} r_\lambda \lambda$ is a $B(T)$-coherent linear relation among elements of some (by necessity finite) subset $A' \subseteq A$, then there exists some scalar $r \in \mathbb{Z}$ such that $rr_\lambda \in \mathbb{Z}$ for all $\lambda \in A'$ (in particular, set r to be the least common multiple of the denominators of

$r_\lambda \in \mathbb{Q}$). By the 'Scaling' property of partial linear structures (see Definition 2.2.22), $\sum_{\lambda\in A'} rr_\lambda$ is a $B(T)$-coherent linear relation among vectors in A with coefficients in $\mathbb{Z}$, and can be realized as $\sum_{\lambda\in A'} w_\lambda \lambda$ for some null tangle Ξ with support A' and weights $w_\lambda := rr_\lambda$. On the other hand, there is no guarantee that such a scalar r exists if the coefficients $r_\lambda \in \mathbb{R}$.

Using Proposition 3.4.10, we give a specialization of Proposition 2.2.21 in terms of tangles rather than $B(T)$-coherent relations. As mentioned in the introduction to the section, this is a generalization of [22, Proposition 7.11], with a correction of a small typo: Ξ must be taken to be a *null* tangle.

Proposition 3.4.13. *Let Ξ be a null tangle in $\mathcal{O}$. Suppose there exists some tagged triangulation T of $\mathcal{O}$ (with all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture), some tagged arc $\gamma \in T$, and some curve $\lambda \in \Xi$ such that $b_\gamma(T, \lambda)$ is strictly positive (respectively, strictly negative) while $b_\gamma(T, \nu)$ is nonpositive (respectively, nonnegative) for every other curve $\nu \in \Xi$. Then $w_\lambda = 0$.*

As a consequence of the 'Irrelevance of zeros' property of partial linear structures (see Definition 2.2.22), it is sufficient in the hypotheses of Proposition 3.4.13 to consider only curves $\nu \in \Xi$ with nonzero weights, yielding the following corollary.

Recall from Definition 3.4.8 that the disorder of a tangle Ξ is the smallest number k such that, after deleting all curves of weight zero, Ξ is a weighted union of k straight tangles. Thus the Null Tangle Property is equivalent to the statement that the disorder of every null tangle is 0: that is, if Ξ is a null tangle, after deleting all curves of weight zero, no curves are left. Hence one method for establishing the property for a given orbifold $\mathcal{O}$ is to rule out the existence of any null tangles of disorder k for all $k \in \mathbb{Z}_+$. Indeed, this is the method used in [22] which we extend here in Theorem 3.4.22, because, as noted in [22], the Curve Separation Property is closely related to ruling out null tangles with certain small disorders.

The proofs of the next four results are identical to those for the comparable assertions in the surfaces case, which appear in [22]. However, we must invoke the following orbifold results instead of their respective surface analogs:

- Theorem 3.3.7 rather than [22, Theorem 4.3]
- Proposition 3.4.2 rather than [22, Proposition 4.6]
- Theorem 3.4.3 rather than [22, Theorem 4.10]
- Theorem 3.4.4 rather than [22, Theorem 4.12]
- Proposition 3.4.11 rather than [22, Propostition 7.10]
- Corollary 3.4.16 rather than [22, Corollary 7.14]

(Note that [22, Propositions 2.4 and 2.7] are results about arbitrary mutation fans from [21]; they are quoted here as Proposition 2.2.30 and Proposition 2.2.31, respectively.)

Proposition 3.4.14. *Null tangles of disorder 1 do not exist in any orbifold. If an orbifold $\mathcal{O}$ has the Curve Separation Property, then it does not admit any null tangles of disorder 2.*

Proposition 3.4.15. *If no null tangles of disorder 2 or 3 exist in $\mathcal{O}$, then the orbifold has the Curve Separation Property.*

Corollary 3.4.16. *If $\mathcal{O}$ has the Null Tangle Property it also has the Curve Separation Property.*

Theorem 3.4.17. *Suppose R is $\mathbb{Z}$ or $\mathbb{Q}$ and let T be a tagged triangulation of $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ (with all arcs tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture). If $\mathcal{O}$ is not a null orbifold, then the following are equivalent:*

1. *$\mathcal{O}$ has the Null Tangle Property.*
2. *The shear coordinates of allowable curves form a basis for $R^{B(T)}$.*
3. *The shear coordinates of allowable curves form a positive basis for $R^{B(T)}$.*
4. *The shear coordinates of allowable curves form universal geometric coefficients for $B(T)$ over R.*

If $\mathcal{O}$ is a null orbifold, then it fails Conclusion 1 and satisfies Conclusions 2, 3, and 4.

Remark 3.4.18. As mentioned at the start of the section, Theorem 3.4.17 is actually a direct analog to [24, Theorem 3.14], a correction to [22, Theorem 7.3] which takes null surfaces into account. Conclusion 4 in Theorem 3.4.17 is not included in [22, Theorem 7.3] and [**?**], however it is equivalent to Conclusion 2 (and therefore Conclusions 1 and 3) by Theorem 2.2.43.

As noted in [22], while establishing the Null Tangle Property in general may be difficult, some of the easier cases of the proof of Theorem 3.4.5 extend to statements about null tangles, which we will use to establish the Null Tangle Property for a small class of orbifolds in Theorem 3.4.22. The following results are comparable to [22, Propositions 7.15-7.17], and the proofs are also similar.

Proposition 3.4.19. *No curve in the support of a null tangle in an orbifold connects two distinct boundary segments.*

Proof. Let λ be an allowable curve which connects two distinct boundary segments. Construct a tagged triangulation T with distinguished (purple) arc γ as described in **Case 1**a in the proof of Theorem 3.4.5 (see Figure 3.13). Observe that λ is the unique allowable curve whose shear coordinate at γ with respect to T is strictly positive. Thus by Proposition 3.4.13, λ cannot appear in the support of a null tangle. □

Proposition 3.4.20. *No curve in the support of a null tangle in an orbifold has exactly one endpoint on a boundary segment.*

Proof. Let λ be an allowable curve with exactly one endpoint on a boundary segment. If the other endpoint of λ is a spiral point, construct a tagged triangulation T with distinguished (purple) arc γ as described in **Case 2**a in the proof of Theorem 3.4.5 (see Figure 3.15). If the other endpoint of λ is an orbifold point, construct a tagged triangulation T with distinguished (purple) arc γ as described in **Case 2**b (see Figure 3.16). In both cases, we again observe that λ is the unique allowable curve whose shear coordinate at γ with respect to T is strictly positive and apply Proposition 3.4.13. □

Proposition 3.4.21. *If $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ is a sphere with b boundary components, p punctures, and q orbifold points with $b + p + q \leq 3$, then no curve in the support of a null tangle in $\mathcal{O}$ has both endpoints on the same boundary segment.*

Proof. The hypothesis is that $\mathcal{O}$ is one of the following 10 orbifolds, 6 of which are in particular marked surfaces since $q = 0$:

- A disk with either

  $\mathcal{O}_1$: 0 punctures and 0 orbifold points, a marked surface (Type A_n),

  $\mathcal{O}_2$: 1 puncture and 0 orbifold points, a marked surface (Type D_n),

  $\mathcal{O}_3$: 0 punctures and 1 orbifold point (Types B_n and C_n),

  $\mathcal{O}_4$: 1 puncture and 1 orbifold point,

  $\mathcal{O}_5$: 2 punctures and 0 orbifold points, a marked surface (Type $\tilde{D}_n$), or

  $\mathcal{O}_6$: 0 punctures and 2 orbifold points (Types $\tilde{B}_n$ and $\tilde{C}_n$).

- An annulus with either

  $\mathcal{O}_7$: 0 punctures and 0 orbifold point, a marked surface (Type $\tilde{A}_n$),

  $\mathcal{O}_8$: 1 puncture and 0 orbifold points, a marked surface, or

  $\mathcal{O}_9$: 0 punctures and 1 orbifold point.

- A sphere with

  $\mathcal{O}_{10}$: 3 boundary components (a marked surface).

  (Recall that spheres with no boundary components are required by Definition 3.2.1 to have at least four points in $\mathbf{M} \cup \mathbf{Q}$: such orbifolds do not satisfy the hypothesis.)

Let Ξ be a null tangle in $\mathcal{O}$ and let λ be a curve in Ξ with both endpoints on the same boundary segment. Note that since λ is allowable and therefore not one of the forbidden curves in Figure 3.8, $\mathcal{O} \notin \{\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3\}$. The proof of [22, Proposition 7.17] handles the other marked surface cases when $\mathcal{O} \in \{\mathcal{O}_5, \mathcal{O}_7, \mathcal{O}_8, \mathcal{O}_{10}\}$. We handle cases $\mathcal{O}_4, \mathcal{O}_6$ and $\mathcal{O}_9$ below using the method of [22].

Construct a tagged triangulation T of $\mathcal{O}$ with (purple) arc γ as described in **Case 1**b in the proof of Theorem 3.4.5 and shown in Figure 3.14. Suppose by way of contradiction that there exists a curve $\lambda \in \Xi$, distinct from ν, for which $b_\gamma(T, \nu) > 0$ and $w_\nu \neq 0$. Then we observe from the figure that ν must have an endpoint on the same boundary segment as λ, implying by Propositions 3.4.19 and 3.4.20 that in fact *both* endpoints of ν must be on the same boundary segment as λ.

If $\mathcal{O} = \mathcal{O}_4$ is a disk with 1 puncture and 1 orbifold point or $\mathcal{O} = \mathcal{O}_6$ is a disk with no punctures and 2 orbifold points, then each boundary segment admits at most one allowable curve with both endpoints on that boundary segment: namely, the one which, together with the unmarked portion of the boundary between its endpoints, encloses the two interior points in $\mathbf{M} \cup \mathbf{Q}$. (If $\mathcal{O}$ has only one marked point on its boundary, this curve is not allowable.) Thus either λ (and ν) don't exist or else λ is unique, both contradictions. Applying Proposition 3.4.13, we conclude that λ is not in the support of Ξ.

If $\mathcal{O} = \mathcal{O}_9$ is an annulus with 0 punctures and 1 orbifold point, then each boundary segment admits at most two allowable curves with both endpoints on that boundary segment: namely, those which, together with the unmarked portion of the boundary between their endpoints, enclose either the other boundary component or both the other boundary component and the orbifold point. (Similar to the case above, if one or both of the two boundary segments of $\mathcal{O}$ has only one marked point, then one or both of these curves is not allowable.) Regardless of which curve (if either exists) is λ, the construction of T ensures that the other curve has a shear coordinate of zero at γ. Again, we have our desired contradiction and Corollary 3.4.13 implies that λ is not in the support of Ξ. □

Theorem 3.4.22. *If $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ is a sphere with b boundary components, p punctures, and q orbifold points with $b + p + q \leq 3$ and is not a null orbifold, then $\mathcal{O}$ has the Null Tangle Property.*

Proof. The hypothesis on $\mathcal{O}$ is a specialization of that in Proposition 3.4.21, and hence $\mathcal{O}$ is one of the 10 types of orbifolds listed in its proof. All 10 have the Curve Separation Property by Theorem 3.4.5. We further require that $\mathcal{O}$ not be a null orbifold, which, from Definition 3.2.1 means it is neither an empty quadrilateral nor a digon containing a unique point in $\mathbf{M} \sqcup Q$ in its interior.

The result is established in the marked surfaces cases when $\mathcal{O} \in \{\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_5, \mathcal{O}_7, \mathcal{O}_8, \mathcal{O}_{10}\}$

by [22, Theorem 7.4]. If $\mathcal{O} = \mathcal{O}_3$ is a disk with 0 punctures and 1 orbifold point, then Propositions 3.4.19 and 3.4.20 imply the Null Tangle Property. In the remaining three cases, Propositions 3.4.19, 3.4.20, and 3.4.21 show that the support of a null tangle cannot contain any curves that involve boundary segments. Proposition 3.4.14 implies that null tangles of disorder 1 and 2 do not exist. We use the strategy laid out in the proof of [22, Theorem 7.4] and sort the remaining allowable curves in $\mathcal{O}$ (see Definition 3.2.10 into 2 or fewer straight tangles. Since Proposition 3.4.14 implies that null tangles of disorder 1 and 2 do not exist, null tangles must therefore have disorder zero, and we conclude that the Null Tangle Property holds.

If $\mathcal{O} = \mathcal{O}_4$ is a disk with 1 puncture and 1 orbifold point, then there are three remaining allowable curves: one closed curve and two non-ordinary curves which spiral into the puncture in opposite directions. The two non-ordinary curves form a conjugate pair and are therefore pairwise compatible, so there are only two straight tangles.

If $\mathcal{O} = \mathcal{O}_6$ is a disk with 0 punctures and 2 orbifold points, there is one remaining allowable curve, and therefore only one straight tangle. Namely, if both orbifold points of $\mathcal{O}$ are of weight 1/2, there exists a unique semi-closed curve and no closed curves; otherwise, there exists a unique closed curve.

If $\mathcal{O} = \mathcal{O}_9$ is an annulus with 0 punctures and 1 orbifold point, there are two remaining allowable curves: the closed curve which encloses just the inner boundary segment, and the closed curve which encloses both the inner boundary segment and the orbifold point. These two curves are compatible, so there is just one straight tangle. □

Remark 3.4.23. In light of Theorem 3.4.17, it is highly desirable to establish the Null Tangle Property for a given orbifold $\mathcal{O}$. We may then conclude that the shear coordinates of allowable curves in $\mathcal{O}$ with respect to a given tagged triangulation T (where all arcs in T are tagged plain if $\mathcal{O}$ has no boundary components and exactly one puncture) form a (positive) basis for $R^{B(T)}$, and therefore constitute universal geometric coefficients for $B(T)$ over R. Recall from Remark 3.3.16 that $B(T^*)$ is a rescaling of $B(T)$, where $B(T^*)$ is the signed adjacency matrix of the reversed triangulation T^* on the reversed orbifold $\mathcal{O}^*$ obtained from $\mathcal{O}, T$ by swapping the weights on all orbifold points in $\mathcal{O}$. By [21, Proposition 7.8(6)], there is an explicit bijection between (positive) bases for $R^{B(T^*)}$ and (positive) bases for $R^{B(T)}$. Thus it is sufficient to establish the Null Tangle Property and compute the shear coordinates of allowable curves in $\mathcal{O}$ to construct (positive) bases for $R^{B(T^*)}$ (and universal geometric coefficients for $B(T^*)$).

3.5 Orbifold folding and unfolding

We now prepare for our use, in the next section, of the surface and orbifold models to attack the folding questions of Chapter 2. Recall from Section 3.2 that in general, orbifolds are constructed as the quotient spaces of finite group actions on manifolds [28, Chapter 13]. In the very specific setting of the cluster algebra model of [5] used in this paper, orbifolds can be thought of as the quotients, or foldings, of marked surfaces under symmetries which are composed of two types of involutions. These two involutions, called local symmetry and prime symmetry, correspond to the two types of orbifold points, of weight 2 and weight 1/2, respectively. Note that only prime symmetries are true geometric symmetries; local symmetries are combinatorial. There is no reason to restrict folding to marked surfaces: any orbifold can be folded under a composition of local and prime symmetries.

We consider, in particular, foldings of *triangulated* orbifolds $\mathcal{O}, T$ under symmetries σ (composed of local and prime symmetries) of both the orbifold and the tagged triangulation, and let $\pi_\sigma(\mathcal{O}), \pi_\sigma(T)$ denote the folding of $\mathcal{O}, T$ under σ. We then reuse σ to denote the element of the symmetric group $\mathfrak{S}_\mathfrak{n}$ defined by taking $\sigma(i)$ such that $\sigma(\gamma_i) = \gamma_{\sigma(i)}$ for each of the n arcs $\gamma_i \in T$. As described in [5, Section 12], σ is a stable automorphism of $B(T)$, and

$$\pi_\sigma(B(T)) = B(\pi_\sigma(T)). \tag{3.12}$$

Definition 3.5.1 (Local folding)**.** Let $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ be an orbifold with puncture $\hat{q} \in \mathbf{M}$ and let σ be the (combinatorial) symmetry on $\mathcal{O}$ which switches the tagging of any tagged arcs incident to $\hat{q}$ and the spiral directions of any allowable curves spiraling into $\hat{q}$. Let T be a tagged triangulation of $\mathcal{O}$ which is preserved under the action of σ: by necessity, T contains a conjugate pair of arcs γ, γ' which are tagged differently at $\hat{q}$. We call T a ***σ-symmetric triangulation***, and call σ a ***local symmetry*** (about $\hat{q}$) of the triangulated orbifold $\mathcal{O}, T$. ***Folding $\mathcal{O}, T$ under σ*** results in the orbifold $\pi_\sigma(\mathcal{O}) = (\mathbf{S}, \mathbf{M} \setminus \{\hat{q}\}, \mathbf{Q} \cup \{\pi_\sigma(\hat{q})\})$ with tagged triangulation $\pi_\sigma(T)$, where $\pi_\sigma(\hat{q})$ is an orbifold point of weight 2 in the location of $\hat{q}$, and $\pi_\sigma(T)$ is obtained from T by replacing the conjugate pair of arcs γ, γ' in T which form a single σ-orbit $\bar{\gamma}$ by a double arc $\pi_\sigma(\bar{\gamma})$ that coincides with γ and γ' (including tagging) except that it terminates at $\pi_\sigma(\hat{q})$. Figure 3.21 illustrates the portions of $\mathcal{O}, T$ and $\pi_\sigma(\mathcal{O}), \pi_\sigma(T)$ near $\hat{q}$ and $\pi_\sigma(\hat{q})$, respectively, in a typical case, along with the associated portions of the signed-adjacency matrices.

Definition 3.5.2 (Prime folding)**.** Let $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ be an orbifold and let σ be a nontrivial, involutive homeomorphism of $\mathbf{S}$ which fixes only a finite set $\tilde{\mathbf{Q}}$ of points in $\mathbf{S}$ which satisfy $\tilde{\mathbf{Q}} \cap (\mathbf{M} \cup \mathbf{Q} \cup \partial\mathcal{O}) = \emptyset$, and which permutes $\mathbf{M}$ and permutes $\mathbf{Q}$, preserving weights on orbifold points. Then σ acts as a two-fold rotation near each point in $\tilde{\mathbf{Q}}$. Let T be a tagged triangulation of $\mathcal{O}$ which is preserved under the action of σ. We call T a ***σ-symmetric triangulation***, and call σ a ***prime symmetry*** (about $\tilde{\mathbf{Q}}$) of the triangulated orbifold $\mathcal{O}, T$. The quotient space identifying each point x in $\mathbf{S}$ with $\sigma(x)$ is a manifold, and we let π_σ denote the natural map to this manifold. Assuming $\mathbf{S}$ is connected (and there is no harm in doing so), π_σ is 2−to−1

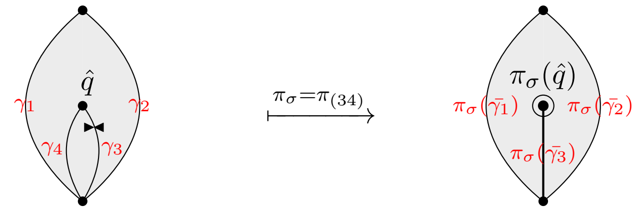

$$\begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \end{array} \left[\begin{array}{c|c|cc} 0 & 1 & -1 & -1 \\ \hline -1 & 0 & 1 & 1 \\ \hline 1 & -1 & 0 & 0 \\ 1 & -1 & 0 & 0 \end{array}\right] \qquad \begin{array}{c} \bar{1} \\ \bar{2} \\ \bar{3} \end{array} \begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 2 & -2 & 0 \end{bmatrix}$$

Figure 3.21: A portion of $\mathcal{O}, T$ around $\hat{q} \in \mathbf{M}$ (and the associated portion of the signed-adjacency matrix $B(T)$) under a local symmetry σ about $\hat{q}$, and the image under π_σ

everywhere except on $\tilde{\mathbf{Q}}$. The application of π_σ is called ***folding*** $\mathcal{O}, T$ ***under*** σ, resulting in the orbifold $\pi_\sigma(\mathcal{O}) = (\pi_\sigma(\mathbf{S}), \pi_\sigma(\mathbf{M}), \pi_\sigma(\mathbf{Q}) \cup \pi_\sigma(\tilde{\mathbf{Q}}))$ with tagged triangulation $\pi_\sigma(T)$, where each point in $\pi_\sigma(\tilde{\mathbf{Q}})$ is an orbifold point of weight $1/2$, all orbifold points in $\pi_\sigma(\mathbf{Q})$ inherit their weights in $\mathbf{Q}$, each pair of arcs $\gamma, \gamma' \in T$ which form a single σ-orbit $\bar{\gamma}$ folds to a single arc $\pi_\sigma(\bar{\gamma}) \in \pi_\sigma(T)$ with corresponding taggings, and each symmetric arc in T which is fixed by σ folds to a pending arc with the inherited tagging at its non-orbifold endpoint. There are four possible options for portions of the triangulated orbifold $\mathcal{O}, T$ which contain neighborhoods of points $\tilde{q} \in \tilde{\mathbf{Q}}$. Figure 3.22 depicts representatives of the three where the two triangles containing $\tilde{q}$ are both ordinary: the symbol $\square$ denotes the point $\tilde{q} \in \tilde{\mathbf{Q}}$. The fourth option, when $\tilde{q}$ is contained in two once-orbifolded digons, is easily visualized by substituting each conjugate pair of arcs in the second row of Figure 3.22 by a single non-ordinary arc. Note that it is not possible for $\tilde{q}$ to be contained in two monogons, each containing two points in $\mathbf{M} \cup \mathbf{Q}$ in their interior, as then $\tilde{q}$ would need to be the single marked point shared by the arcs bounding each monogon.

Definition 3.5.3 (Orbifold folding and unfolding)**.** Let $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ be an orbifold with tagged triangulation T and let σ be a composition of local and prime symmetries of $\mathcal{O}, T$. Then we call σ a ***folding symmetry*** of $\mathcal{O}, T$, and we say that $\mathcal{O}, T$ ***folds to*** $\pi_\sigma(\mathcal{O}), \pi_\sigma(T)$ ***under*** σ, and likewise that $\pi_\sigma(\mathcal{O}), \pi_\sigma(T)$ ***unfolds to*** $\mathcal{O}, T$ ***under*** σ***.***

Remark 3.5.4. Note that [5] defines *unfoldings* of orbifolds rather than foldings, as the authors were interested in proving that unfoldings to marked surfaces exist. Indeed, they construct such (full) unfoldings of all orbifolds other than spheres with a unique orbifold point of weight $1/2$ in [5, Theorem 12.9]. For our purposes of folding various mutation-linear algebra notions, we begin with symmetric triangulated orbifolds $\mathcal{O}, T$, and then define the associated quotient map as folding. Our subsequent definition of unfolding agrees with that in [5], including the terms 'prime' and 'local.' Likewise, [5,6] define exchange matrix *unfolding* rather than the folding. The

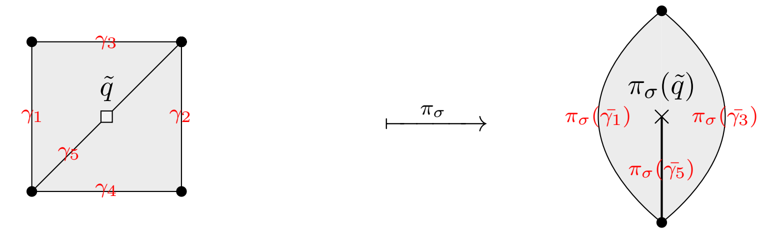

$$\begin{array}{c} 1\\2\\3\\4\\5 \end{array}\left[\begin{array}{cc|cc|c} 0 & 0 & 1 & 0 & -1\\ 0 & 0 & 0 & 1 & -1\\ \hline -1 & 0 & 0 & 0 & 1\\ 0 & -1 & 0 & 0 & 1\\ \hline 1 & 1 & 0 & -1 & -1 \end{array}\right] \qquad \begin{array}{c} \bar{1}\\\bar{3}\\\bar{5} \end{array}\left[\begin{array}{ccc} 0 & 1 & -2\\ -1 & 0 & 2\\ 1 & -1 & 0 \end{array}\right]$$

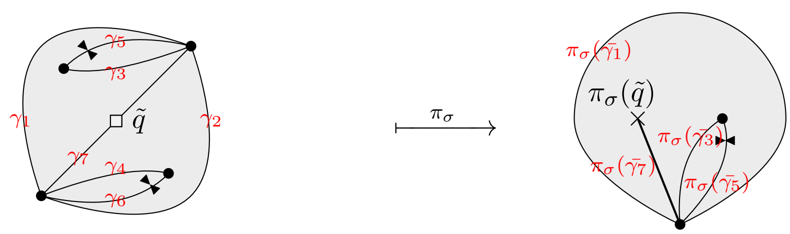

$$\begin{array}{c} 1\\2\\3\\4\\5\\6\\7 \end{array}\left[\begin{array}{cc|cc|cc|c} 0 & 0 & 1 & 0 & 1 & 0 & -1\\ 0 & 0 & 0 & 1 & 0 & 1 & -1\\ \hline -1 & 0 & 0 & 0 & 0 & 0 & 1\\ 0 & -1 & 0 & 0 & 0 & 0 & 1\\ \hline -1 & 0 & 0 & 0 & 0 & 0 & 1\\ 0 & -1 & 0 & 0 & 0 & 0 & 1\\ \hline 1 & 1 & 0 & -1 & -1 & -1 & -1 \end{array}\right] \qquad \begin{array}{c} \bar{1}\\\bar{3}\\\bar{5}\\\bar{7} \end{array}\left[\begin{array}{cccc} 0 & 1 & 1 & -2\\ -1 & 0 & 0 & 2\\ -1 & 0 & 0 & 2\\ 1 & -1 & -1 & 0 \end{array}\right]$$

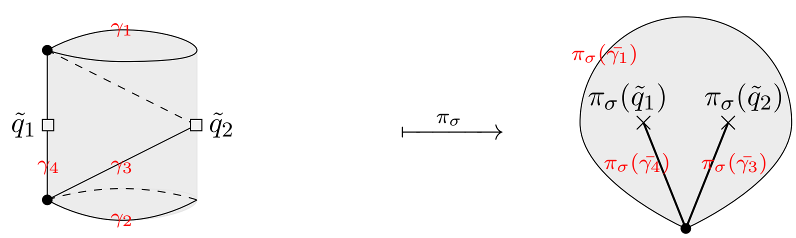

$$\begin{array}{c} 1\\2\\3\\4 \end{array}\left[\begin{array}{cc|c|c} 0 & 0 & 1 & -1\\ 0 & 0 & 1 & -1\\ \hline -1 & -1 & 0 & 2\\ \hline 1 & 1 & -2 & 0 \end{array}\right] \qquad \begin{array}{c} \bar{1}\\\bar{3}\\\bar{4} \end{array}\left[\begin{array}{ccc} 0 & 2 & -2\\ -1 & 0 & 2\\ 1 & -2 & 0 \end{array}\right]$$

Figure 3.22: Reading down the left, possible ordinary portions of $\mathcal{O}, T$ surrounding neighborhoods of points $\tilde{q} \in \tilde{\mathbf{Q}}$ (and the associated portions of the signed-adjacency matrices $B(T)$) where σ is a prime symmetry about $\tilde{\mathbf{Q}}$, with their images under π_σ on the right

matrix unfolding of Definition 2.3.7, drawn from [4], can be thought of as a specialization of the definitions of [5, 6] (which were communicated to the authors by Zelevinsky) as we explicitly require a symmetry condition.

3.6 Orbifolds and mutation-linear algebra folding

Given an orbifold $\mathcal{O}$ with tagged triangulation T, recall from Definition 3.3.19 that $\mathcal{F}_{\mathbb{Q}}(T)$ denotes the rational quasi-lamination fan for T, a rational fan in $\mathbb{R}^{|T|}$ whose cones C_Λ are the non-negative spans of the shear coordinate vectors of collections Λ of pairwise compatible allowable curves in $\mathcal{O}$. The primary aim of this section is to prove the following theorem, that $\mathcal{F}_{\mathbb{Q}}(T)$ folds under a folding symmetry σ of $\mathcal{O}, T$.

Theorem 3.6.1. *Suppose $\mathcal{O}, T$ folds to $\pi_\sigma(\mathcal{O}), \pi_\sigma(T)$ under folding symmetry σ. Recall from Definition 2.4.5 that V^σ denotes the fixed space of $\mathbb{R}^{|T|}$ under coordinate permutation action by σ, and ϕ denotes the vector space isomorphism between V^σ and $\mathbb{R}^{|\pi_\sigma(T)|}$. Then,*

$$\phi(\mathcal{F}_{\mathbb{Q}}(T) \cap V^\sigma) = \mathcal{F}_{\mathbb{Q}}(\pi_\sigma(T)).$$

The following two corollaries are then immediate as a consequence of Theorems 3.4.3 and 3.4.17, respectively, and are the respective orbifold analogs of the finite type results Theorem 2.4.24 and Corollary 2.4.25.

Corollary 3.6.2. *Suppose both $\mathcal{O}, T$ and $\pi_\sigma(\mathcal{O}), \pi_\sigma(T)$ have the Curve Separation Property, and let $\mathcal{F}'_{B(T)}$ and $\mathcal{F}'_{B(\pi_\sigma(T))}$ denote the rational parts of the mutation fans $\mathcal{F}_{B(T)}$ and $\mathcal{F}'_{B(\pi_\sigma(T))}$, respectively. Then,*

$$\phi(\mathcal{F}'_{B(T)} \cap V^\sigma) = \mathcal{F}'_{B(\pi_\sigma(T))}.$$

Corollary 3.6.3. *Suppose both $\mathcal{O}, T$ and $\pi_\sigma(\mathcal{O}), \pi_\sigma(T)$ have the Null Tangle Property and R is $\mathbb{Z}$ or $\mathbb{Q}$. Let $\mathcal{B}$ denote the set of shear coordinates of allowable curves in $\mathcal{O}$ (which form a positive (cone) basis for $R^{B(T)}$ and universal geometric coefficients for $B(T)$ over R). Then, the following collection is a positive (cone) basis for $R^{B(\pi_\sigma(T))}$ (and universal geometric coefficients for $\pi_\sigma(B(T))$ over R), where $\bar{\mathcal{B}}$ denotes the set of σ-orbits $\bar{\mathbf{a}}$ of $\mathcal{B}$:*

$$\left(\phi\left(\pi_\sigma(\bar{\mathbf{a}})\right) : \bar{\mathbf{a}} \in \bar{\mathcal{B}} \text{ spans a cone in } \mathcal{F}_{\mathbb{Q}}(T)\right).$$

We now prepare for the proof of Theorem 3.6.1. Suppose $\mathcal{O}, T$ folds to $\pi_\sigma(\mathcal{O}), \pi_\sigma(T)$ under folding symmetry σ. Let Λ be a collection of pairwise compatible allowable curves in $\mathcal{O}$ and let L be a quasi-lamination supported on Λ. Define the action of σ on Λ and L in the natural

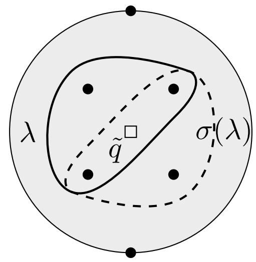

Figure 3.23: Allowable curve λ is not compatible with allowable curve $\sigma(\lambda)$, where σ is the prime symmetry about $\tilde{q} \in \tilde{Q}$, denoted by the symbol $\square$ (note that for local symmetries, λ and $\sigma(\lambda)$ are always compatible)

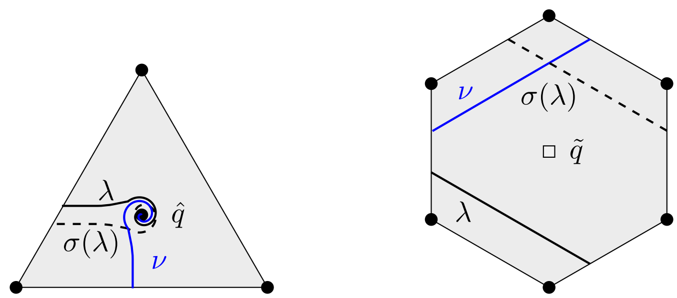

Figure 3.24: Allowable curves λ, ν in $\mathcal{O}$ are compatible, but $\sigma(\lambda)$ and ν are not for, from left to right, local and prime symmetries σ about $\hat{q}$ and $\tilde{q}$

manner:

$$\begin{aligned} &\sigma(\Lambda) = \{\sigma(\lambda) : \lambda \in \Lambda\}, \text{ and} \\ &\sigma(L) = \{w_\lambda \cdot \sigma(\lambda) : w_\lambda \lambda \in L\}. \end{aligned} \tag{3.13}$$

It is clear from Definitions 3.5.1, 3.5.2, and 3.5.3 that λ is an allowable curve in $\mathcal{O}$ if and only if $\sigma(\lambda)$ is allowable, and that two allowable curves λ and ν in $\mathcal{O}$ are compatible if and only if $\sigma(\lambda)$ and $\sigma(\nu)$ are compatible. Thus $\sigma(\Lambda)$ (resp. $\sigma(L)$) consists of pairwise compatible curves. However in general, the curves in the full σ-orbit $\bar{\Lambda}$ of Λ (resp. $\bar{L}$ of L) need not be pairwise compatible. (See Figure 3.23 for an example of a curve λ which is incompatible with $\sigma(\lambda)$, and Figure 3.24 for an example of compatible curves λ and ν with $\sigma(\lambda)$ incompatible with ν.)

Define Λ^σ (resp. L^σ), the ***σ-symmetric part*** of Λ (resp. L), to be the unique maximal subset of Λ (resp. L) which is fixed under action by σ, or equivalently, the union of all σ-orbits of allowable curves which are fully contained in Λ (resp. L). It is possible that $\Lambda^\sigma = \emptyset$ (resp., $L^\sigma = \emptyset$). Furthermore, if $\Lambda^\sigma = \emptyset$ then by necessity $L^\sigma = \emptyset$, as L^σ is supported on Λ^σ. If $\Lambda^\sigma = \Lambda$ (resp. $L^\sigma = L$), then we say that Λ (resp. L) is ***σ-symmetric***.

For each curve λ in Λ^σ, the σ-orbit $\bar{\lambda}$ of λ must itself be a σ-symmetric collection of pairwise compatible allowable curves, which we refer to as a ***compatible σ-orbit***. Thus we may define $\pi_\sigma(\bar{\lambda})$ to be the quotient of $\bar{\lambda}$ under the folding symmetry σ, a single allowable curve in $\pi_\sigma(\mathcal{O})$. Observe that curves in the same σ-orbit $\bar{\lambda}$ must appear with the same weight $w_{\bar{\lambda}}$ in L^σ. Extend π_σ to Λ^σ and L^σ in the natural way, by defining

$$\begin{aligned} \pi_\sigma(\Lambda^\sigma) &= \{\pi_\sigma(\bar{\lambda}) : \lambda \in \Lambda^\sigma\}, \text{ and} \\ \pi_\sigma(L^\sigma) &= \{w_{\bar{\lambda}} \cdot \pi_\sigma(\bar{\lambda}) : w_{\bar{\lambda}}\lambda \in L^\sigma\}. \end{aligned} \tag{3.14}$$

The key to the proof of Theorem 3.6.1 is establishing a correspondence between the shear coordinates of σ-symmetric rational quasi-laminations on $\mathcal{O}, T$ and the shear coordinates of rational quasi-laminations on $\pi_\sigma(\mathcal{O}), \pi_\sigma(T)$. Towards that goal, we first establish a correspondence between the quasi-laminations themselves.

Proposition 3.6.4. *The quotient map π_σ is a bijection between σ-symmetric collections of pairwise compatible allowable curves (resp. quasi-laminations) in $\mathcal{O}$ and collections of pairwise compatible allowable curves (resp. quasi-laminations) in $\pi_\sigma(\mathcal{O})$.*

Proof. Recall from Definition 3.5.3 that $\sigma = \sigma_1 \circ \sigma_2 \circ \cdots \circ \sigma_k$, where each σ_i is a single local or prime symmetry. The claim holds trivially when σ is the identity (that is, when the sequence is empty), so by induction on k, it is sufficient to consider the two cases when σ consists of either a single local or prime symmetry.

We begin by establishing a bijection between compatible σ-orbits of allowable curves in $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ and allowable curves in $\pi_\sigma(\mathcal{O}) = (\pi_\sigma(\mathbf{S}), \pi_\sigma(\mathbf{M}), \pi_\sigma(\mathbf{Q}))$.

If σ is a local symmetry about some puncture $\hat{q} \in \mathbf{M}$, then all σ-orbits of allowable curves in $\mathcal{O}$ are in particular compatible σ-orbits. These orbits come in two types:

1. A single curve in $\mathcal{O}$ which does not spiral into $\hat{q}$.

2. A conjugate pair of curves in $\mathcal{O}$ which spiral into $\hat{q}$ in opposite directions.

These (compatible) σ-orbits correspond, under π_σ, to the two types of allowable curves in $\pi_\sigma(\mathcal{O})$, listed below in the matching order:

1. A curve in $\pi_\sigma(\mathcal{O})$ which does not terminate at the orbifold point $\pi_\sigma(\hat{q}) \in \pi_\sigma(\mathbf{Q})$.

2. A double curve in $\pi_\sigma(\mathcal{O})$ with one endpoint at $\pi_\sigma(\hat{q})$.

If σ is a prime symmetry with fixed-point set $\tilde{\mathbf{Q}}$, then it is possible that not all σ-orbits of allowable curves in $\mathcal{O}$ are compatible (see, e.g., Figure 3.23). However, every *compatible* σ-orbit is of one of six types. First, we list the three types of compatible orbits of curves which are not closed, as these are easy to visualize:

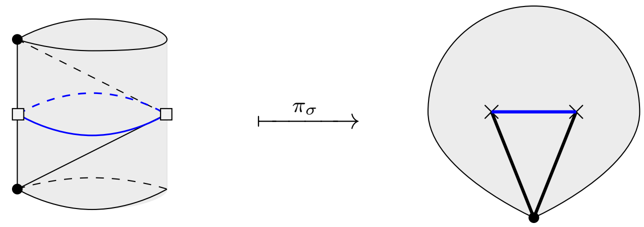

Figure 3.25: A compatible prime σ-orbit of Type 5 (rendered in blue) and its image under π_σ

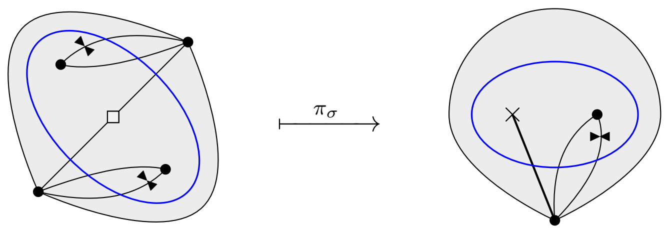

Figure 3.26: A compatible prime σ-orbit of Type 6 (rendered in blue) and its image under π_σ

1. Two disjoint non- or semi-closed curves in $\mathcal{O}$ neither of which intersects any points in $\tilde{\mathbf{Q}}$.

2. A single non-closed ordinary curve in $\mathcal{O}$ whose interior intersects a unique point in $\tilde{\mathbf{Q}}$.

3. A single semi-closed curve in $\mathcal{O}$ whose interior intersects a unique point in $\tilde{\mathbf{Q}}$.

The latter three types of compatible orbits are obtained by examining the possible actions of σ on closed curves in $\mathcal{O}$. Such curves are topological circles, and hence σ, which by Definition 3.5.2 is an involutive homeomorphism which fixes only the points in $\tilde{\mathbf{Q}}$, can act in one of three possible ways: "swapping," reflection, and rotation. (The key observation for the latter two actions, which yield orbits of size one, is that a homeomorphism of a topological circle can have either zero or two fixed points.) These actions yield the following respective types of orbits:

4. Two disjoint closed curves in $\mathcal{O}$ neither of which intersects any points in $\tilde{\mathbf{Q}}$.

5. A single closed curve in $\mathcal{O}$ that intersects two distinct points in $\tilde{\mathbf{Q}}$. (See Figure 3.25.)

6. A single closed curve in $\mathcal{O}$ that does not intersect any points in $\tilde{\mathbf{Q}}$ and cuts out a disk containing a point in $\tilde{\mathbf{Q}}$. (See Figure 3.26.)

These six types of compatible σ-orbits correspond, under π_σ, to each of the possible types of allowable curves in $\pi_\sigma(\mathcal{O})$, listed below in the matching order:

1. A non- or semi-closed curve in $\pi_\sigma(\mathcal{O})$ which does not terminate at any of the orbifold points in $\pi_\sigma(\tilde{\mathbf{Q}}) \subseteq \pi_\sigma(\mathbf{Q})$.

2. A pending curve in $\pi_\sigma(\mathcal{O})$ with one endpoint at a point in $\pi_\sigma(\tilde{\mathbf{Q}})$.

3. A semi-closed curve in $\pi_\sigma(\mathcal{O})$ with exactly one endpoint at a point in $\pi_\sigma(\tilde{\mathbf{Q}})$.

4. A closed curve in $\pi_\sigma(\mathcal{O})$ that does not cut out a disk containing a point in $\pi_\sigma(\tilde{\mathbf{Q}})$.

5. A semi-closed curve in $\pi_\sigma(\mathcal{O})$ with both endpoints in $\pi_\sigma(\tilde{\mathbf{Q}})$. (See Figure 3.25. This correspondence is the rationale behind the term 'semi-closed': see [7, Remark 4.3].)

6. A closed curve in $\pi_\sigma(\mathcal{O})$ that cuts out a disk containing a point in $\pi_\sigma(\tilde{\mathbf{Q}})$. (See Figure 3.26.)

By definition, a σ-symmetric collection of pairwise compatible allowable curves in $\mathcal{O}$ must be comprised of compatible σ-orbits, each of which corresponds to an allowable curve in $\pi_\sigma(\mathcal{O})$ by the work above. We now check that two allowable curves in $\pi_\sigma(\mathcal{O})$ are pairwise compatible if and only if the union of their corresponding compatible σ-orbits in $\mathcal{O}$ is a pairwise compatible collection of curves.

First, suppose σ is a local symmetry. Clearly, the pairwise-compatibility of the union of a collection of orbits consisting of single curves is equivalent to pairwise-compatibility between the curves themselves. More generally, any σ-symmetric pairwise compatible collection of curves in $\mathcal{O}$ either contains no curves spiraling into $\hat{q}$, or else contains a unique conjugate pair of curves with opposite spirals at $\hat{q}$ (a single compatible orbit). Likewise, a pairwise compatible collection of curves in $\pi_\sigma(\mathcal{O})$ either contains no curves terminating at $\pi_\sigma(\hat{q})$ or else contains a unique such curve. Curves in $\pi_\sigma(\mathcal{O})$ which are compatible with such a double curve correspond to (compatible orbits of size one of) curves in $\mathcal{O}$ which are compatible with both members of the conjugate pair to which it lifts.

When σ is a prime symmetry, we appeal to the fact that in this case, $\mathcal{O}$ is a ramified Galois covering of $\pi_\sigma(\mathcal{O})$ with branching points in $\tilde{\mathbf{Q}}$ of order two [5, Section 12.1]. Furthermore, two compatible allowable curves cannot share an orbifold point. Thus, intersections among a σ-symmetric collection of allowable curves in $\mathcal{O}$ correspond to intersections between their images in $\pi_\sigma(\mathcal{O})$, which in particular implies that such a collection in $\mathcal{O}$ is *not* pairwise compatible if and only if its image also not compatible.

Having established the result for σ-symmetric collections of pairwise compatible allowable curves, we extend it to quasi-laminations by invoking (3.14), which ensures that all curves in a compatible σ-orbit in a quasi-lamination have the same weight. □

Turning our attention to shear coordinates, the following observation is immediate since σ is a symmetry of $\mathcal{O}, T$.

Lemma 3.6.5. *Let L a quasi-lamination on $\mathcal{O}$. Then,*

$$b_\gamma(T, L) = b_{\sigma(\gamma)}(\sigma T, \sigma(L)) = b_{\sigma(\gamma)}(T, \sigma(L)). \tag{3.15}$$

Lemma 3.6.6. *A quasi-lamination L on $\mathcal{O}$ is σ-symmetric if and only if $\mathbf{b}(T,L) \in V^\sigma$.*

Proof. As discussed in the proof of Proposition 3.6.4, it suffices to consider the case when σ consists of a single local or prime symmetry. In this case, σ is an involution on $\mathcal{O}, T$, and therefore also involutive when viewed as an element of $\mathfrak{S}_\mathfrak{n}$ acting on $\mathbb{R}^n$ by coordinate permutation.

Suppose L is σ-symmetric, and for each tagged arc $\gamma \in T$, let $[\sigma(\mathbf{b}(T,L))]_\gamma$ denote the coordinate of $\sigma(\mathbf{b}(T,L))$ indexed by γ. We show $[\sigma(\mathbf{b}(T,L))]_\gamma = b_\gamma(T,L)$ to conclude that $\mathbf{b}(T,L) \in V^\sigma$:

$$\begin{aligned}
[\sigma(\mathbf{b}(T,L))]_\gamma &= [\sigma(\mathbf{b}(T,\sigma(L)))]_\gamma \text{ since } L \text{ is } \sigma\text{-symmetric} \\
&= [\mathbf{b}(T,\sigma(L))]_{\sigma^{-1}(\gamma)} \text{ by (2.20)} \\
&= b_{\sigma^{-1}(\gamma)}(T,\sigma(L)) \\
&= b_{\sigma(\gamma)}(T,\sigma(L)) \text{ since } \sigma \text{ is an involution} \\
&= b_\gamma(T,L) \text{ by (3.15).}
\end{aligned}$$

Conversely, suppose $\mathbf{b}(T,L) \in V^\sigma$ and assume that $L \neq \sigma(L)$. By Theorem 3.3.6, it follows that $\mathbf{b}(T,L) \neq \mathbf{b}(T,\sigma L)$. In particular, there exists an arc $\gamma \in T$ such that $b_\gamma(T,L) \neq b_\gamma(T,\sigma L)$. However by Equation (3.15) and the fact that σ is an involution,

$$b_\gamma(T,\sigma(L)) = b_{\sigma(\gamma)}(T,\sigma^2(L)) = b_{\sigma(\gamma)}(T,L).$$

Thus $b_\gamma(T,L) \neq b_{\sigma(\gamma)}(T,L)$, contradicting the hypothesis that $\sigma(\mathbf{b}(T,L)) = \mathbf{b}(T,L)$. □

Recall from Proposition 3.6.4 that π_σ induces a bijection between σ-symmetric quasi-laminations on $\mathcal{O}$ and quasi-laminations on $\pi_\sigma(\mathcal{O})$. By Lemma 3.6.6, the shear cooordinates of σ-symmetric quasi-laminations on $\mathcal{O}, T$ are themselves σ-symmetric: we show that they map, under the isomorphism $\phi : V^\sigma \xrightarrow{\cong} \mathbb{R}^m$ of (2.23), to the shear coordinates of the corresponding quasi-laminations on $\pi_\sigma(\mathcal{O}), \pi_\sigma(T)$.

Proposition 3.6.7. *Let L be a σ-symmetric quasi-lamination on $\mathcal{O}$. Then,*

$$\phi(\mathbf{b}(T,L)) = \mathbf{b}(\pi_\sigma(T), \pi_\sigma(L)). \tag{3.16}$$

Proof. As discussed in the proof of Proposition 3.6.4, it suffices to consider the case when σ consists of a single local or prime symmetry. Furthermore, since L consists of a collection of compatible σ-orbits where all the curves in a given orbit $\bar{\lambda}$ appear with the same weight $w_{\bar{\lambda}}$ in L, it suffices to prove the following:

$$b_\gamma(T,\bar{\lambda}) = b_{\pi_\sigma(\bar{\gamma})}(\pi_\sigma(T), \pi_\sigma(\bar{\lambda})) \text{ for each tagged arc } \gamma \in T. \tag{3.17}$$

First, suppose that σ is a local symmetry. We consider four cases, and appeal throughout to the classification of compatible orbits and curves into types from the proof of Proposition 3.6.4.

(a) Suppose $\bar{\gamma} = \{\gamma\}$ and $\bar{\lambda} = \{\lambda\}$ is of Type 1. Then the contributions from intersections between λ and the triangle(s) in T containing γ are of exactly the same interaction type, sign, and magnitude as the intersections between $\pi_\sigma(\bar{\lambda})$ and the triangles(s) in $\pi_\sigma(T)$ containing $\pi_\sigma(\bar{\gamma})$, since the open sets of $\mathcal{O}$ in which each intersection between λ and γ occur are identical to their images in $\pi_\sigma(\mathcal{O})$, and both λ and γ are of the same type (e.g., ordinary, pending, etc.) as their images.

(b) Suppose $\bar{\gamma} = \{\gamma\}$ and $\bar{\lambda} = \{\lambda, \lambda'\}$ is of Type 2. Since λ and λ', a conjugate pair, the two curves coincide outside of a small neighborhood around $\hat{q}$. Further, γ is disjoint from $\hat{q}$, so the intersections between λ and the triangle(s) in T containing γ are exactly the same as those between between λ' and the triangle(s) in T containing γ, and are of exactly the same interaction type and sign as those between $\pi_\sigma(\bar{\lambda})$ and the triangles in $\pi_\sigma(T)$ containing $\pi_\sigma(\gamma)$. However, since $\pi_\sigma(\bar{\lambda})$ is a double arc while λ and λ' are ordinary, and $\pi_\gamma(\bar{\lambda})$ and λ are of the same type, examination of Tables 3.1 and 3.2 shows that each intersection in $\pi_\sigma(\mathcal{O})$ is of twice the magnitude as the corresponding intersection in $\mathcal{O}$. Thus,

$$b_{\pi_\sigma(\bar{\gamma})}(\pi_\sigma(T), \pi_\sigma(\bar{\lambda})) = 2 \cdot b_\gamma(T, \lambda) = b_\gamma(T, \lambda) + b_\gamma(T, \lambda') = b_\gamma(T, \bar{\lambda}).$$

(c) Suppose $\bar{\gamma} = \{\gamma, \gamma'\}$ and $\bar{\lambda} = \{\lambda\}$ is of Type 1. Use similar reasoning as above to to establish an interaction- and sign-preserving correspondence between intersections of λ with γ and of $\pi_\sigma(\bar{\lambda})$ with $\pi_\sigma(\bar{\gamma})$. Comparing the third and fourth rows of Table 3.1 with the fourth and fifth rows, respectively, of Table 3.2, we see that this correspondence is also magnitude-preserving.

(d) Finally, suppose $\bar{\gamma} = \{\gamma, \gamma'\}$ and $\bar{\lambda} = \{\lambda, \lambda'\}$ is of Type 2 so that both $\pi_\sigma(\bar{\gamma})$ and $\pi_\sigma(\bar{\lambda})$ are double. Then, comparing the fifth row of Table 3.1 with the first row of Table 3.2, the contribution of the intersection at $\hat{q}$ between $\bar{\lambda}$ and γ is the same as the contribution of the intersection at $\pi_\sigma(\hat{q})$ between $\pi_\sigma(\bar{\lambda})$ and $\pi_\sigma(\bar{\gamma})$, either ± 1. There is an interaction- and sign-preserving correspondence between the remainder of the intersections of λ with γ, λ' with γ, and $\pi_\sigma(\bar{\gamma})$ with $\pi_\sigma(\bar{\lambda})$, where, comparing the third row of Table 3.1 with the fifth row of Table 3.2, we see that each intersection in $\pi_\sigma(\mathcal{O})$ is of twice the magnitude as the corresponding intersections in $\mathcal{O}$, and therefore equal to their sum.

Now, suppose that σ is a prime symmetry. We have six cases to consider, again based on the classification of compatible orbits and curves into types from the proof of Proposition 3.6.4.

(a) Suppose $\bar{\gamma} = \{\gamma, \gamma'\}$ and $\bar{\lambda} = \{\lambda, \lambda'\}$ is of Type 1 or 4. Because $\mathcal{O}, T$ is a ramified covering of $\pi_\sigma(\mathcal{O})$ [5, Section 12.1] and the intersections between the curves λ and λ' with the arc γ in $\mathcal{O}$ take place outside of neighborhoods of points in $\tilde{\mathbf{Q}}$ (see Figure 3.22, there is an interactions- and sign-preserving correspondence between intersections of $\pi_\sigma(\bar{\lambda})$ with $\pi_\sigma(\bar{\gamma})$ and the union of intersections of λ with γ and intersections of λ' with γ. Since $\pi_\sigma(\bar{\lambda})$ is of the same type as both λ and λ', and likewise $\pi_\sigma(\bar{\gamma})$ is of the same type as γ, the correspondence is also magnitude-preserving.

(b) Suppose $\bar{\gamma} = \{\gamma, \gamma'\}$ and $\bar{\lambda} = \{\lambda\}$ is of Type 3 or 6. Then $\pi_\sigma(\bar{\gamma})$ is of the same type as γ, and $\pi_\sigma(\bar{\lambda})$ and λ are of the same type. As in case (a), the intersections of λ with γ in $\mathcal{O}$ take place outside of neighborhoods of points in $\tilde{\mathbf{Q}}$ (see Figure 3.22), so we have an interaction-, sign- and magnitude-preserving correspondence between these intersections and intersections of $\pi_\sigma(\lambda)$ with $\pi_\sigma(\gamma)$.

(c) Suppose $\bar{\gamma} = \{\gamma, \gamma'\}$ and $\bar{\lambda} = \{\lambda\}$ is of Type 2 or 5. As in cases (a) and (b), the intersections of λ with γ in $\mathcal{O}$ take place outside of neighborhoods of points in $\tilde{\mathbf{Q}}$ (see Figure 3.22), so we have an interaction- and sign-preserving correspondence between these intersections and intersections of $\pi_\sigma(\lambda)$ with $\pi_\sigma(\gamma)$. While λ and $\pi_\sigma(\bar{\lambda})$ are not of the same type, they are either ordinary and pending or ordinary (closed) and semi-closed, respectively, examining Tables 3.1 and 3.2 confirms that none of these types affects the magnitude of an intersection's contribution.

(d) Suppose $\bar{\gamma} = \{\gamma\}$ and $\bar{\lambda} = \{\lambda, \lambda'\}$ is of Type 1 or 4. Then $\pi_\sigma(\bar{\lambda})$ is of the same type as both λ and λ', and $\pi_\sigma(\bar{\gamma})$ is a pending arc while γ is an ordinary arc. By (3.15), $b_\gamma(T, \lambda) = b_\gamma(T, \lambda')$, and since $\mathcal{O}$ is a ramified covering of $\pi_\sigma(\mathcal{O})$, trivial intersections in $\pi_\sigma(\mathcal{O})$ between $\pi_\sigma(\bar{\lambda})$ and $\pi_\sigma(\bar{\gamma})$ correspond to pairs of trivial intersections in $\mathcal{O}$, between λ and γ and λ' and γ. Likewise for non-trivial intersections.

   The non-zero contributions to $b_{\pi_\sigma(\bar{\gamma})}(\pi_\sigma(T), \pi_\sigma(\bar{\lambda}))$ come from intersections between the curve and the arc in $\pi_\sigma(\mathcal{O})$ of the type shown in the second and third rows of Table 3.2, in the neighborhood of the orbifold point $\pi_\sigma(\tilde{q})$ at which $\pi_\sigma(\bar{\gamma})$ has an endpoint (see Figure 3.22). Each such intersection in $\pi_\sigma(\mathcal{O})$ unfolds to two non-trivial intersections in $\mathcal{O}$ of the type shown in the first two rows of Table 3.1 of the same sign and half the magnitude: namely, a non-trivial intersection of λ with γ and a non-trivial intersection of λ' with γ. Thus $b_{\pi_\sigma(\bar{\gamma})}(\pi_\sigma(T), \pi_\sigma(\bar{\lambda})) = b_\gamma(T, \bar{\lambda})$.

(e) Suppose $\bar{\gamma} = \{\gamma\}$ and $\bar{\lambda} = \{\lambda\}$ is of Type 6. As in case (d), $\pi_\sigma(\bar{\gamma})$ is a pending arc while γ is an ordinary arc, and there is the same sign-preserving correspondence between intersections in $\pi_\sigma(\mathcal{O})$ and pairs of intersections in $\mathcal{O}$, with all non-trivial intersections in $\pi_\sigma(\mathcal{O})$ of the type shown in the second row of Table 3.2, corresponding to two non-trivial intersections in $\mathcal{O}$ of the type shown in the first row of Table 3.1. In this case the intersections in $\mathcal{O}$ are all of λ with γ, since $\bar{\lambda} = \{\lambda\}$, but the correspondence is still magnitude-preserving as well since $\pi_\sigma(\bar{\lambda})$ and λ are both ordinary closed curves.

(f) Suppose $\bar{\gamma} = \{\gamma\}$ and $\bar{\lambda} = \{\lambda\}$ is of Type 2, 3, or 5 . As in case (e), $\pi_\sigma(\bar{\gamma})$ is a pending arc while γ is an ordinary arc. While $\pi_\sigma(\bar{\gamma})$ and λ may not be of the same type, neither is a double curve, so as noted in case (c), the type does not affect the magnitude of an intersection's shear coordinate contribution.

There are two types of intersections between $\pi_\sigma(\bar{\lambda})$ and $\pi_\sigma(\bar{\gamma})$: those in their interiors, which are handled as in case (e), and those which occur at shared orbifold points $\pi_\sigma(\tilde{q}) \in \pi_\sigma(\tilde{Q})$. These latter intersections are all of the type shown in the first row of Table 3.2, and are the foldings of intersections between λ and γ of the type shown in the first row of Table 3.2, which have the same sign and magnitude.

We conclude that (3.17), and hence (3.16), hold. □

Proof of Theorem 3.6.1. We begin by characterizing the cones in $\mathcal{F}_{\mathbb{Q}}(T) \cap V^\sigma$. Each is of the form $C_\Lambda \cap V^\sigma$ for some cone C_Λ in $\mathcal{F}_{\mathbb{Q}}(T)$, where Λ is a collection of pairwise compatible allowable curves in $\mathcal{O}$. Recall from Definition 3.3.19 that $C_\Lambda \cap \mathbb{Q}^{|T|}$ is the set of shear coordinates of all rational quasi-laminations supported on subsets of Λ. Then by Lemma 3.6.6, $C_\Lambda \cap V^\sigma \cap \mathbb{Q}^{|T|}$ is the set of shear coordinates of all σ*-symmetric* rational quasi-laminations supported on subsets of Λ. But σ-symmetric quasi-laminations must have σ-symmetric support, so in particular $C_\Lambda \cap V^\sigma \cap \mathbb{Q}^{|T|}$ is the set of shear coordinates of all σ-symmetric rational quasi-laminations supported on subsets of Λ^σ, the σ-symmetric part of Λ. Thus,

$$C_\Lambda \cap V^\sigma \cap \mathbb{Q}^{|T|} = C_{\Lambda^\sigma} \cap V^\sigma \cap \mathbb{Q}^{|T|}.$$

Since both $C_\Lambda \cap V^\sigma$ and $C_{\Lambda_\sigma} \cap V^\sigma$ are rational cones, $C_\Lambda \cap V^\sigma = C_{\Lambda^\sigma} \cap V^\sigma$.

Since ϕ is a vector space isomorphism, we have established a bijection between cones $\phi(C_\Lambda \cap V^\sigma)$ in $\phi(\mathcal{F}_{\mathbb{Q}}(T) \cap V^\sigma)$ and σ-symmetric collections Λ^σ of pairwise compatible curves in $\mathcal{O}$. By Proposition 3.6.4, the latter set is in bijection with collections $\pi_\sigma(\Lambda^\sigma)$ of pairwise compatible curves in $\pi_\sigma(\mathcal{O})$, which by definition of the rational quasi-lamination fan are in bijection with cones $C_{\pi_\sigma(\Lambda^\sigma)}$ in $\mathcal{F}_{\mathbb{Q}}(\pi_\sigma(T))$. In particular, applying Proposition 3.6.7,

$$\begin{aligned}
\phi(C_\Lambda \cap V^\sigma) \cap \mathbb{Q}^{|\pi_\sigma(T)|} &= \phi(C_\Lambda \cap V^\sigma \cap \mathbb{Q}^{|T|}) \\
&= \phi(C_{\Lambda^\sigma} \cap V^\sigma \cap \mathbb{Q}^{|T|}) \\
&= \{\phi(\mathbf{b}(T, L)) : L \text{ is a } \sigma\text{-symmetric quasi-lamination on } \mathcal{O}
\end{aligned}$$

$$
\begin{aligned}
&\quad \text{supported on } \Lambda_L \subseteq \Lambda^\sigma\} \\
&= \{\mathbf{b}(\pi_\sigma(T), \pi_\sigma(L)) : \pi_\sigma(L) \text{ is a quasi-lamination on } \pi_\sigma(\mathcal{O}) \\
&\quad \text{supported on } \pi_\sigma(\Lambda_L) \subseteq \pi_\sigma(\Lambda^\sigma)\} \\
&= C_{\pi_\sigma(\Lambda^\sigma)} \cap \mathbb{Q}^{|\pi_\sigma(T)|}.
\end{aligned}
$$

Since both $\phi(C_\Lambda \cap V^\sigma)$ and $C_{\pi_\sigma(\Lambda^\sigma)}$ are rational cones, we conclude that they are equal. Hence the cones in $\phi(\mathcal{F}_\mathbb{Q}(T) \cap V^\sigma)$ are precisely the cones in $\mathcal{F}_\mathbb{Q}(\pi_\sigma(T))$, and the two fans coincide. $\square$

Chapter 4

Dominance and resection of orbifolds

4.1 Introduction

Given two $n \times n$ exchange matrices $B = [b_{ij}]$ and $B' = [b'_{ij}]$, we say that B ***dominates*** B' if the entries b_{ij} and b'_{ij} weakly agree in sign and $|b_{ij}| \geq |b'_{ij}|$ for each $i, j \in [n]$. The dominance relationship is introduced in [24], which describes four interesting mutation-linear algebra and principal-coefficients cluster algebra phenomena that often occur when B dominates B'.

Phenomenon I: *In many cases when B dominates B', the identity map from the mutation-linear structure $\mathbb{R}^B$ to the mutation-linear structure $\mathbb{R}^{B'}$ is mutation-linear.*

See Section 2.2.2 for background on mutation-linear structures and mutation-linear maps.

Phenomenon II: *In many cases when B dominates B', the mutation fan $\mathcal{F}_B$ refines the mutation fan $\mathcal{F}_{B'}$.*

See Definition 2.2.29 for the definition of the mutation fan for an exchange matrix.

Phenomenon III: *In many cases when B dominates B', the cluster scattering fan* $\mathrm{ScatFan}(B)$ *refines the cluster scattering fan* $\mathrm{ScatFan}(B')$.

The ***cluster scattering fan*** is defined in [23] by the (principal coefficients) ***cluster scattering diagram*** of [14]. The construction is reviewed in [24, Section 4].

Phenomenon IV: *In many cases when B dominates B', the map $\nu_{\mathbf{z}}$ defined in Section 4.4 is an injective,* $\mathbf{g}$*-vector-preserving ring homomorphism from the principal-coefficients cluster algebra $\mathcal{A}_\bullet(B')$ into $\mathcal{A}_\bullet(B)$. In a smaller set of cases, this map further sends each ray theta function for B' to a ray theta function for B'.*

See Section 2.2.1 for background on cluster algebras of geometric type, and Definition 2.2.11 for descriptions of cluster algebras with principal coefficients and $\mathbf{g}$-vectors, respectively.

Ray theta functions are defined in [24, Section 1.4] as a generalization of the set of cluster variables in the principal-coefficients cluster algebra. If B is of finite type, the second part of the phenomenon can be restated to assert that the map sends cluster variables to cluster variables.

The goal of this chapter is to prove theorems which provide examples of these phenomena using the orbifolds model of Chapter 3. In particular, [24, Definition 3.17] introduces a simple operation, called resection, at arcs in triangulations of marked surfaces. This operation induces a dominance relation on signed adjacency matrices, and provides examples of each of the Phenomena in certain cases when B and B' are related in this manner. In Section 4.2, we show that this notion of resection can be broadened to encompass resection at ordinary arcs in orbifolds in a natural way (see Definition 4.2.1), and introduce a new operation of resection at non-ordinary arcs (see Definition 4.2.2). We then show in Proposition 4.2.8 that any composition of these operations induces a dominance relation on signed-adjacency matrices.

Section 4.3 is devoted to proving the following two results, which generalize [24, Theorems 1.2 and 1.7] by providing examples of "rational versions" of Phenomena I and II, respectively. (The bulk of the work is in proving Theorem 4.1.2; analogously to [24], Theorem 4.1.1 then follows from the generalizations in Section 3.4 of results in [22] to orbifolds.)

Theorem 4.1.1. *Let $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ be an orbifold with tagged triangulation T consisting of n arcs, and let $\mathcal{O}' = (\mathbf{S}', \mathbf{M}', \mathbf{Q}')$ be obtained by a resection of $\mathcal{O}$ that is compatible with T. Let T' denote the triangulation induced on $\mathcal{O}'$ by T. If every component of $\mathcal{O}$ and $\mathcal{O}'$ either has the Null Tangle Property or is a null surface, then the identity map from $\mathbb{Q}^{B(T)}$ to $\mathbb{Q}^{B(T')}$ is mutation-linear.*

Theorem 4.1.2. *Let $\mathcal{O}$ be an orbifold with tagged triangulation T consisting of n arcs, and let $\mathcal{O}'$ be obtained by a resection of $\mathcal{O}$ that is compatible with T. Let T' denote the triangulation induced on $\mathcal{O}'$ by T. The rational quasi-lamination fan $\mathcal{F}_{\mathbb{Q}}(T)$ refines the rational quasi-lamination fan $\mathcal{F}_{\mathbb{Q}}(T')$. If every component of $\mathcal{O}$ and $\mathcal{O}'$ has the Curve Separation Property, then $\mathcal{F}^{\mathbb{Q}}_{B(T)} \cap \mathbb{Q}^n$ refines $\mathcal{F}^{\mathbb{Q}}_{B(T')} \cap \mathbb{Q}^n$, where $\mathcal{F}^{\mathbb{Q}}_{B(T)}$ and $\mathcal{F}^{\mathbb{Q}}_{B(T')}$ denote the rational parts of the mutation fans $\mathcal{F}_{B(T)}$ and $\mathcal{F}_{B(T')}$, respectively.*

When an exchange matrix is of finite type, it is known that both its cluster scattering fan and its mutation fan coincide with the $\mathbf{g}$-vector fan for its transpose. (The equality of the cluster scattering fan with the $\mathbf{g}$-vector fan for the transpose can be deduced from [14] as discussed in [24, Section 1.3]. The equality of the mutation fan with the $\mathbf{g}$-vector fan for the transpose, proven in [21], appears here as Proposition 2.2.46.) Since the $\mathbf{g}$-vector fan is rational and all orbifolds of finite type have the Curve Separation Property by Theorem 3.4.5 and [22, Theorem 6.1], we obtain the following corollary to Theorem 4.1.2, a generalization of [24, Corollary 1.15] which provides an example of Phenomenon III:

Corollary 4.1.3. *Let $\mathcal{O}$ be an orbifold with tagged triangulation T consisting of n arcs, and let $\mathcal{O}'$ be obtained by a resection of $\mathcal{O}$ that is compatible with T. Let T' denote the triangulation induced on $\mathcal{O}'$ by T. If $\mathcal{O}$, and therefore $\mathcal{O}'$, are of finite type, then the cluster scattering fan* $\mathrm{ScatFan}(B(T))$ *refines* $\mathrm{ScatFan}(B(T'))$.

Remark 4.1.4. We hope in future work to understand $\mathrm{ScatFan}(B(T))$ when $B(T)$ is the signed-adjacency matrix of a triangulation T of an orbifold; current work is underway with Muller and Reading to characterize the scattering fans of marked surfaces. By [24, Theorem 1.10], for *any* exchange matrix B, $\mathrm{ScatFan}(B)$ refines $\mathcal{F}_B$.

We prove our third main result, an example of Phenomenon IV, in Section 4.4.

Theorem 4.1.5. *Suppose $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ is a disk with one orbifold point, no punctures, and at least two marked points, and that T is a triangulation of $\mathcal{O}$ such that the signed-adjacency matrix $B = B(T)$ is acyclic. If $B(T)$ dominates B', then both parts of Phenomenon IV occur.*

Theorem 4.1.5 completes the proof of the following result, which appears as [24, Expected Theorem 1.17] by handling the cases when the exchange matrix B is of type B_n or C_n:

Theorem 4.1.6. *Suppose B is an acyclic matrix of finite type, and suppose B' is another exchange matrix dominated by B. Then the map $\nu_{\mathbf{z}}$ is an injective,* $\mathbf{g}$*-vector preserving homomorphism from $\mathcal{A}_\bullet(B')$ to $\mathcal{A}_\bullet(B)$ and sends each cluster variable in $\mathcal{A}_\bullet(B')$ to a cluster variable in $\mathcal{A}_\bullet(B)$.*

4.2 Resection of orbifolds

In this section we define our primary tool, resection of orbifolds. Definition 4.2.1 is a generalization of the resection of marked surfaces in [24, Definition 3.17] to ordinary arcs in orbifolds. We introduce resection at non-ordinary arcs in Definition 4.2.2; the remaining definitions then generalize [24, Definitions 3.18, 3.19, and 3.20]. If the orbifold $\mathcal{O}$ under consideration is in particular a marked surface, then the definitions coincide.

Definition 4.2.1 (Resection at an ordinary arc)**.** Suppose α is an ordinary (untagged) arc in orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ with endpoints at $p_1, p_2 \in \mathbf{M}$ as in the lefthand image of Figure 4.1 (endpoints p_1 and p_2 need not be distinct). Add a new marked point p_α in the interior of $\mathcal{O}$ close to α, and then draw a curve in $\mathcal{O}$ which connects p_1 to p_2 and forms a digon whose interior contains p_α but no points in $\mathbf{M} \cup \mathbf{Q}$. Draw two additional curves inside the digon, connecting p_1 and p_2 to p_α. This cuts the digon into two triangles, neither of which has any marked or orbifold points in its interior. Finally, remove the interior of the triangle that does not have α as an edge: the result is illustrated in the righthand image of Figure 4.1. If neither endpoint p_1 nor

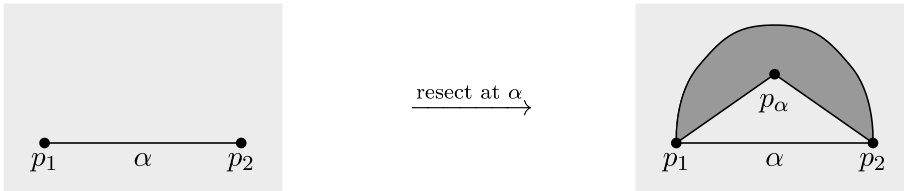

Figure 4.1: Resection of an orbifold at an ordinary arc

p_2 is on a boundary component, we are done. Otherwise, complete the procedure by cutting the resulting surface at each endpoint p_i which *is* on a boundary component in order to satisfy the requirement in Definition 3.2.1 that the boundary components of an orbifold be circles. In this case each such p_i becomes two distinct marked points on the boundary. The resulting marked surface is called a ***resection of*** $\mathcal{O}$ ***at*** α, and we follow [24] in using the verb "resect" and the noun "resection" to describe passing from $\mathcal{O}$ to the resected orbifold, which we denote by $\mathcal{O}'$. Note that $\mathcal{O}'$ may be disconnected even if $\mathcal{O}$ is connected; we explicitly permitted disconnected orbifolds in Definition 3.2.1 for precisely this reason. In general, for each ordinary arc α in $\mathcal{O}$ there are two possible resections at α, one for each side of the arc. However, we disallow any resections on $\mathcal{O}$ that result in a connected component of $\mathcal{O}'$ which is an empty triangle, as these are not permitted for orbifolds in Definition 3.2.1. (The term "empty" means the triangle contains no punctures orbifold points. None of the other forbidden connected components can arise as a result of resection at an ordinary arc.)

Definition 4.2.2 (Resection at a non-ordinary arc)**.** Suppose α is a non-ordinary (untagged) arc in orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ connecting marked point $p \in \mathbf{M}$ to orbifold point $q \in \mathbf{Q}$ as in the lefthand image of Figure 4.2 . Cut $\mathcal{O}$ "inside" α, thereby eliminating q and turning α' into a regular arc with both endpoints at p. Glue a once-punctured monogon along the exposed side of α'. Add a new marked point p_α inside the monogon near the former location of q. Draw two non-coinciding (although isotopic) curves inside the mongon, both connecting p to p_α. This cuts the monogon into a digon and a triangle, neither of which has any marked or obifold points in its interior. Finally remove the interior of the digon: the result is illustrated in the righthand image of Figure 4.2. If p is not on a boundary component, we are done. Otherwise, complete the procedure by cutting the resulting orbifold at p to satisfy the requirement that boundary components be circles. In this case p becomes two distinct marked points on the boundary. The resulting orbifold $\mathcal{O}'$ is called the ***resection of*** $\mathcal{O}$ ***at*** α. Observe that for each non-ordinary arc α in $\mathcal{O}$ there is a unique resection at α.

Remark 4.2.3. Note that Definition 4.2.2 of resection at a non-ordinary arc does not depend on whether the given arc α is double or pending (that is, if the orbifold endpoint q is of weight

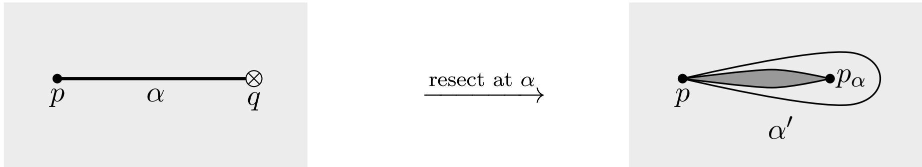

Figure 4.2: Resection of an orbifold at a non-ordinary arc

2 or 1/2, respectively). However, the idea of "cutting inside" α to turn it into an ordinary arc with two coinciding endpoints at marked point p makes some intuitive sense in both cases, which we now attempt to illuminate. Let T be a tagged triangulation of $\mathcal{O}$ containing the arc α and suppose that $\mathcal{O}, T$ arises as a folding $\pi_\sigma(\tilde{\mathcal{O}}), \pi_\sigma(\hat{T})$ of some other orbifold $\hat{\mathcal{O}}$ with tagged triangulation $\hat{T}$ under a folding symmetry σ (see Definition 3.5.3), where $\pi_\sigma^{-1}(q)$ is not an orbifold point of $\hat{\mathcal{O}}$. (Recall from Remark 3.5.4 that *full* unfoldings to marked surfaces are constructed in [5] for all orbifolds other than spheres with a unique orbifold point of weight 1/2, so this assumption is not unreasonable.) For simplicity, suppose σ consists of a single local or prime symmetry about $\pi_\sigma^{-1}(q)$.

- If q is of weight 2 and σ is a local symmetry (see Definition 3.5.1), then $\pi^{-1}(q)$ is a single puncture in $\hat{\mathcal{O}}$, $\pi^{-1}(p)$ is a single marked point in $\hat{OO}$, and $\pi^{-1}(\alpha)$ consists of a conjugate pair of arcs in $\hat{T}$ with opposite taggings at $\pi^{-1}(q)$. We then view α as being obtained by removing those taggings and gluing the two arcs in $\pi^{-1}(\alpha)$ together. By performing resection on $\mathcal{O}$ at α, we are cutting those two pieces apart in their interior, turning the double arc into an ordinary arc with coinciding endpoints.

- On the other hand, if q is of weight 1/2 and σ is a prime symmetry (see Definition 3.5.2), then $\pi^{-1}(q)$ is a single unmarked interior point of $\hat{\mathcal{O}}$, $\pi^{-1}(p)$ consists of two distinct marked points in $\hat{\mathcal{O}}$, and $\pi^{-1}(\alpha)$ is an ordinary arc in $\hat{T}$ connecting one element of $\pi^{-1}(p)$ to the other, with its "midpoint" at $\pi^{-1}(q)$. We then view α as being obtained by folding $\pi^{-1}(\alpha)$ in half and then gluing the two halves together. By performing resection on $\mathcal{O}$ at α, we are cutting those two halves apart in their interior, restoring the pending arc to a 'full' ordinary arc with coinciding endpoints.

Definition 4.2.4 (Resection at a collection of arcs)**.** In general, a ***resection*** of an orbifold $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ is an orbifold $\mathcal{O}' = (\mathbf{S}', \mathbf{M}', \mathbf{Q}')$ obtained by performing any number of resections at ordinary arcs and oribfold-resections at non-ordinary arcs, possibly resecting the same ordinary arc on both sides. We also use the term resection to refer to the operation itself. (Note that $\mathbf{M}' \supset \mathbf{M}$ while $\mathbf{Q}' \subseteq \mathbf{Q}$, with strict containment $\mathbf{Q}' \subset \mathbf{Q}$ if any resections at non-ordinary arcs were performed.) This operation is well-defined up to isotopy (that is, independent of the order of composition of resections at arcs) so long as the arcs which are resected are pairwise-compatible and none of the resections at ordinary arcs are of the type forbidden in Definition 4.2.1 (which result in $\mathcal{O}'$ having an empty triangle as a connected component).

Definition 4.2.5 (Resection compatible with a triangulation)**.** Let $\mathcal{O}$ be an orbifold with (untagged) triangulation T. A resection of $\mathcal{O}$ is ***compatible with*** T if it satisfies the following conditions:

1. The arcs that are resected are all arcs in T.

2. For each ordinary arc that is resected, the new marked point p_α is placed in the interior of a triangle of T which is bounded by α, and the new curves which define the resection are also drawn in the interior of this triangle (with the possible exception of their endpoints).

3. If an ordinary arc α that is resected has an endpoint at a puncture p, then either both endpoints of α are at p or else there exists at least one additional arc β with an endpoint at p that is also resected. In the latter case, we further require that the new marked points p_α and p_β be placed in the interiors of different triangles of T.

Note that Definition 4.2.5 implicitly requires that a resection not cut off an empty triangle (see Definition 4.2.1). Furthermore, these conditions, which all come directly from [24, Definition 3.19], imply that a resected ordinary arc α may not be the fold edge of a self-folded triangle of T, and must therefore be contained in two distinct triangles of T. (Recall that non-ordinary arcs are always contained in unique triangles with three distinct edges.) No additional requirements are needed to accommodate orbifold-resection.

Definition 4.2.6 (Triangulation induced on the resected orbifold)**.** Let $\mathcal{O}$ be an orbifold with (untagged) triangulation T. If $\mathcal{O}'$ is a resection of $\mathcal{O}$ that is compatible with T, then each ordinary arc in T is an ordinary arc in $\mathcal{O}'$, and each non-ordinary arc in T that was not resected is a non-ordinary arc in $\mathcal{O}$. Each non-ordinary arc in T that *was* resected corresponds to a unique ordinary arc in $\mathcal{O}'$. The collection of these arcs cut $\mathcal{O}$ into triangles, called the ***triangulation induced on*** $\mathcal{O}'$ ***by*** T.

Remark 4.2.7. While resection of $\mathcal{O}, T$ is defined with respect to (untagged) arcs and triangulations (see Definition 3.2.3), the results of Section 3.4 relating quasi-laminations on $\mathcal{O}, T$ to the mutation-linear algebra notions associated to the signed-adjacency matrix $B(T)$ are phrased in terms of tagged triangulations. We therefore pause to recall the discussion in Definition 3.2.6 of the methods for moving from one to another.

There exists a bijective map τ between (untagged) triangulations and the set of tagged triangulations where the only notched taggings occur at punctures incident to a conjugate pair of arcs with opposite taggings. (Recall that such pairs are the images, under τ, of self-folded triangles in the (untagged) triangulation: see the central two images in Figure **??**.) To transform an arbitrary tagged triangulation into an untagged triangulation, simply reverse all taggings at punctures incident only to arcs tagged notched at that puncture. The resulting tagged triangulation satisfies the condition described above: the only notched taggings occur at conjugate pairs of arcs with opposite taggings at that puncture. Replacing each conjugate pair of arcs with the corresponding self-folded triangle under τ^{-1} results in the desired (untagged) triangulation. Passing between a tagged triangulation and an (untagged) triangulation in this way does not change the associated signed-adjacency matrix: indeed, the signed-adjacency matrix of a tagged triangulation is *defined* as the signed-adjacency matrix of the (untagged) triangulation which is obtained in the manner described above.

Define ***resection of*** $\mathcal{O}$ ***compatible with* tagged *triangulation*** T as any resection of $\mathcal{O}$ compatible with the (untagged) triangulation obtained from T by the method described above. (For convenience, we refer to the resulting (untagged) triangulation by T as well.) The **tagged *triangulation induced on*** $\mathcal{O}'$ is then obtained by first applying τ to the triangulation T', and then, at any punctures in $\mathcal{O}'$ which correspond to punctures in $\mathcal{O}$ at which all arcs were tagged the same, applying those taggings to the arcs' images in T'. Again, we abuse terminology and call the resulting tagged triangulation T'.

The section concludes with a proof that resection on orbifolds induces a dominance relationship on signed adjacency matrices, generalizing [24, Proposition 3.23]. The argument here is structured similarly: we can almost, but not quite, resect one arc at a time, but at times must appeal to Condition 3 of Definition 4.2.5.

Proposition 4.2.8. *Given an orbifold* $\mathcal{O}$ *with tagged triangulation* T*, perform a resection of* $\mathcal{O}$ *compatible with* T *and let* T' *be the triangulation induced by* T *on the resected orbifold* $\mathcal{O}'$*. Then* $B(T)$ *dominates* $B(T')$*.*

Proof. We consider several cases for resection at an arc α in T which is not the fold edge of a self-folded triangle. In each, we let T' be the triangulation obtained by resecting at α (as well as, if required by condition 3 of Definition 4.2.5, resecting at an additional arc $\beta \in T$), and show that $B(T)$ dominates $B(T')$. The conclusion, that $B(T)$ dominates $B(T')$ for any resection of $\mathcal{O}$ consisting of a collection of such resections at arcs, is actually stronger than the claimed result, as it encompasses additional resections besides those which are compatible with T. In cases (i)-(v), we assume α is ordinary, and in cases (vi) and (vii), we assume it is non-ordinary.

(i) Suppose α is ordinary and incident to neither the puncture in a once-punctured digon nor the puncture in a monogon containing one puncture and one orbifold point in its interior. Furthermore suppose that the new marked point p_α is placed in the interior of a

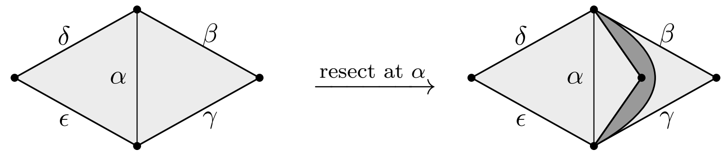

Figure 4.3: Illustration for case (i) in the proof of Proposition 4.2.8

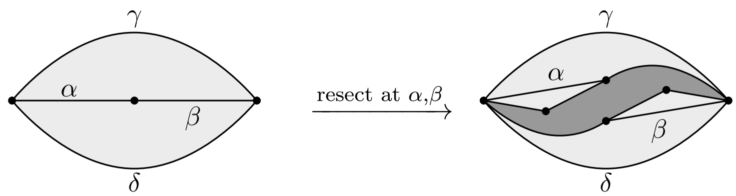

Figure 4.4: Illustration for case (ii) in the proof of Proposition 4.2.8

triangle of T with distinct ordinary edges α, β, and γ. The parts of T and T' right around α are illustrated in Figure 4.3 (cf. [24, Figure 6]), and the case is covered by the proof of [24, Proposition 3.23]. While not explicitly mentioned there, it is possible that β and or γ is the non-fold edge of a self-folded triangle: in this case, in addition to the entries of $B(T)$ indexed by β (and or γ) decreasing in absolute value by 1 to obtain $B(T')$, so too do the entries indexed by the fold edge. Note as well that either or both of the arcs δ and ϵ may be non-ordinary: this does not affect the proof.

(ii) Suppose α is ordinary and is incident to the puncture in a once-punctured digon. Then by necessity, p_α must be placed in the interior of a triangle with distinct ordinary edges α, β, and γ, where β is also incident to the puncture. We again appeal to the proof of [24, Proposition 3.23], which in this case invokes condition 3 of Definition 4.2.5 and requires resection at β. The situation is illustrated in Figure 4.4 (cf. [24, Figure 7]). Similarly to case (i), it is possible that γ and or δ is the non-fold edge of a self-folded triangle: again, this would imply that the entries of $B(T)$ indexed by the fold edge(s) will also decrease in absolute value by 1 to obtain $B(T')$.

(iii) Suppose α is ordinary and is incident to the puncture in a monogon containing one puncture and one orbifold point. T must contain two additional arcs inside the monogon bounded by γ: an ordinary arc β and a non-ordinary arc δ. Since α is not the fold edge of a self-folded triangle, there exists a unique ordinary arc $\beta \in T$ which is also incident to the puncture, and by condition 3 of Definition 4.2.5, both α and β must be resected according to condition 2 with p_α and p_β placed in the interiors of different triangles of T. By symmetry, we may assume that p_α is placed in the interior of the triangle with

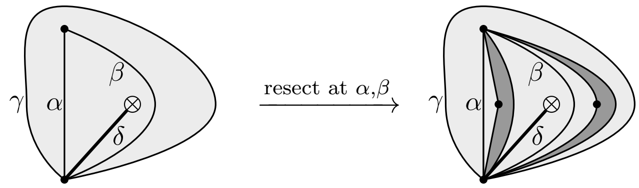

Figure 4.5: Illustration for case (iii) in the proof of Proposition 4.2.8

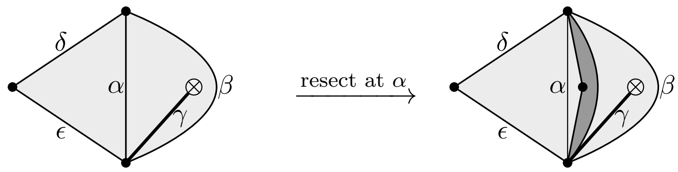

Figure 4.6: Illustration for case (iv) in the proof of Proposition 4.2.8

distinct ordinary edges α, β, γ and p_β is placed in the interior of the triangle with non-ordinary edge δ, and let T' denote the resulting resection. The situation is illustrated in Figure 4.5. Then $B(T)$ dominates $B(T')$, where $B(T')$ is obtained from $B(T)$ as follows: set the $\alpha\delta$- and $\alpha\beta$-entries equal to zero. If γ is an arc (rather than a boundary segment), then decrease the absolute values of the $\beta\gamma$- and $\gamma\beta$-entries by 1. The $\alpha\beta$- and $\beta\alpha$-entries in both signed adjacency matrices are zero.

(iv) Suppose α is ordinary and not incident to the puncture in a monogon containing one puncture and one orbifold point. Furthermore suppose that the new marked point p_α is placed in the interior of a triangle of T with ordinary edges α and β and non-ordinary edge γ. (By necessity, α cannot be incident to the puncture in a once-punctured digon in this case.) Then $B(T)$ dominates $B(T')$. The situation is illustrated in Figure 4.6: let ϵ and δ denote the edges of the other triangle containing α. If δ and β were to coincide, then ϵ would bound a mongon containing one puncture and one orbifold point with α incident to the puncture, a contradiction of our hypothesis. Likewise, ϵ and β can't coincide. Thus resection sets the $\alpha\gamma$- and $\gamma\alpha$- entries of $B(T)$ equal to zero, and, if β is an arc, decreases the absolute value of the $\alpha\beta$- and $\beta\alpha$- entries by 1.

(v) For our last ordinary case, suppose α is ordinary and p_α is placed in the interior of a triangle with two non-ordinary edges. (By necessity, α cannot be incident to either the puncture in a once-punctured digon nor the puncture in a monogon with one puncture and one orbifold point in this case.) Then $B(T)$ dominates $B(T')$, where resection sets the $\alpha\gamma$-, $\gamma\alpha$-, $\alpha\beta$-, and $\beta\alpha$-entries equal to zero. The situation is illustrated in Figure 4.7.

(vi) Now, suppose α is non-ordinary and is the unique non-ordinary edge of a triangle in T.

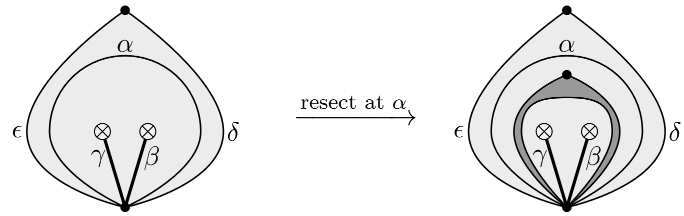

Figure 4.7: Illustration for case (v) in the proof of Proposition 4.2.8

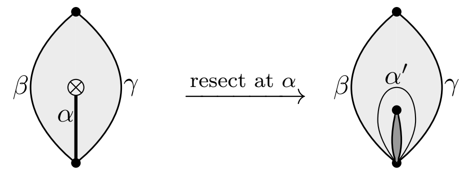

Figure 4.8: Illustration for case (vi) in the proof of Proposition 4.2.8

Then $B(T)$ dominates $B(T')$, where the row and column in $B(T)$ indexed by α corresponds to the row and column in $B(T')$ indexed by α', the image of α in $\mathcal{O}'$ under resection at α. The situation is illustrated in Figure 4.8.

If α is a double arc, then the $\alpha\beta$-entry of $B(T)$ is 2 and the $\alpha\gamma$-entry is -2, while the $\beta\alpha$-entry is -1 and the $\gamma\alpha$ entry is 1. Resection affects only the $\alpha\beta$- and $\alpha\gamma$- entries, setting the corresponding $\alpha'\beta$ and $\alpha'\gamma$-entries of $B(T')$ to 1 and -1, respectively.

If instead α is a pending arc, then the $\alpha\beta$-entry of $B(T)$ is 1 and the $\alpha\gamma$-entry is -1, while the $\beta\alpha$-entry is -2 and the $\gamma\alpha$ entry is 2. Resection now affects only the $\beta\alpha$- and $\gamma\alpha$-entries, setting the corresponding $\beta\alpha'$- and $\gamma\alpha'$-entries of $B(T)$ to -1 and 1, respectively.

(vii) Finally, suppose α is non-ordinary and is one of two non-ordinary edges of a triangle in T. Then $B(T)$ dominates $B(T')$, where again the row and column in $B(T)$ indexed by α corresponds to the row and column in $B(T')$ indexed by α', the image of α in $\mathcal{O}'$ under resection at α. The situation is illustrated in Figure 4.9. (Note that while the figure illustrates the case when, traveling around the triangle in a clockwise direction, one encounters the arcs in the order α, β, γ, the written description below also applies to the case when the other non-ordinary arc β follows α in the counterclockwise direction.)

If α is a double arc, then resection changes only the $\alpha\gamma$- and $\alpha\beta$-entries of $B(T)$. In particular, one obtains the $\alpha'\gamma$- entry of $B(T')$ by decreasing the absolute value of the $\alpha\gamma$-entry of $B(T)$ by 1. The $\alpha'\beta$-entry is obtained by decreasing the absolute value of the $\alpha\beta$-entry by 1 if β is a double arc or by 2 if β is pending.

If instead α is a pending arc, then resection changes only the $\gamma\alpha$- and $\beta\alpha$-entries of $B(T)$.

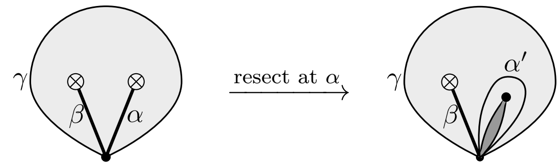

Figure 4.9: Illustration for case (vii) in the proof of Proposition 4.2.8

In particular, one obtains the $\gamma\alpha'$- entry of $B(T')$ by decreasing the absolute value of the $\alpha\gamma$-entry of $B(T)$ by 1. The $\beta\alpha'$-entry is obtained by decreasing the absolute value of the $\beta\alpha$-entry by 1 if β is a double arc or by 2 if β is pending. □

Remark 4.2.9. Given a triangulation T of an orbifold $\mathcal{O}$, it is generally not possible to obtain every exchange matrix which is dominated by $B(T)$ by resecting $\mathcal{O}$. One potential issue arises in (ordinary) resection of surfaces [24, Remark 3.24], and carries over to resection at ordinary arcs in orbifolds. In particular, in case (i) of the proof of Proposition 4.2.8, illustrated in Figure 4.3, if both β and γ are (ordinary) arcs, rather than boundary segments, then the entries $b_{\alpha\beta}, b_{\alpha\gamma}, b_{\beta\alpha}$, and $b_{\gamma\alpha}$ of $B(T)$ are all changed by resection at α. It is impossible to change only the entries $b_{\alpha\beta}$ and $b_{\beta\alpha}$ through resection, even though the result would be an exchange matrix which is dominated by $B(T)$.

A second issue arises which is particular to resection at non-ordinary arcs. As mentioned in Definition 3.2.4, the signed adjacency matrix $B(T)$ of an orbifold triangulation always has entries in the set $\{0, \pm 1, \pm 2, \pm 4\}$, where an entry $b_{\alpha\beta}$ of absolute value 4 corresponds to an adjacency between a double arc $\alpha \in T$ and a pending arc $\beta \in T$. (The entry $b_{\beta\alpha}$ has absolute value 1.) In such a situation, the only way to change the entry $b_{\alpha\beta}$ is by resection at α: however, as discussed in case (vii) of the proof of Proposition 4.2.8 and illustrated in 4.9, such a resection decreases the absolute value of $b_{\alpha\beta}$ by 2. It is impossible to decrease the absolute value of $b_{\alpha\beta}$ by 1 to obtain an entry $b_{\alpha'\beta}$ of $B(T')$ of absolute value 3, even though the result would be an exchange matrix dominated by $B(T)$.

4.3 Mutation fan refinement and mutation-linearity

In this section, we prove Theorem 4.1.2, and then use the result to prove Theorem 4.1.1.

Proof of Theorem 4.1.2. As a consequence of Theorem 3.3.6, there exists a bijective map from rational quasi-laminations L on $\mathcal{O}$ to rational quasi-laminations L' on $\mathcal{O}'$ that preserves shear coordinates. We determine this bijection, and then use it to construct, for each cone C_Λ in $\mathcal{F}_{\mathbb{Q}}(T)$, a corresponding cone $C'_{\Lambda'}$ in $\mathcal{F}_{\mathbb{Q}}(T')$ such that $C_\Lambda \subseteq C'_{\Lambda'}$. In particular, recall from Definition 3.3.19 that $C_\Lambda \cap \mathbb{Q}^n$ is the set of shear coordinates of all non-negative rational weightings

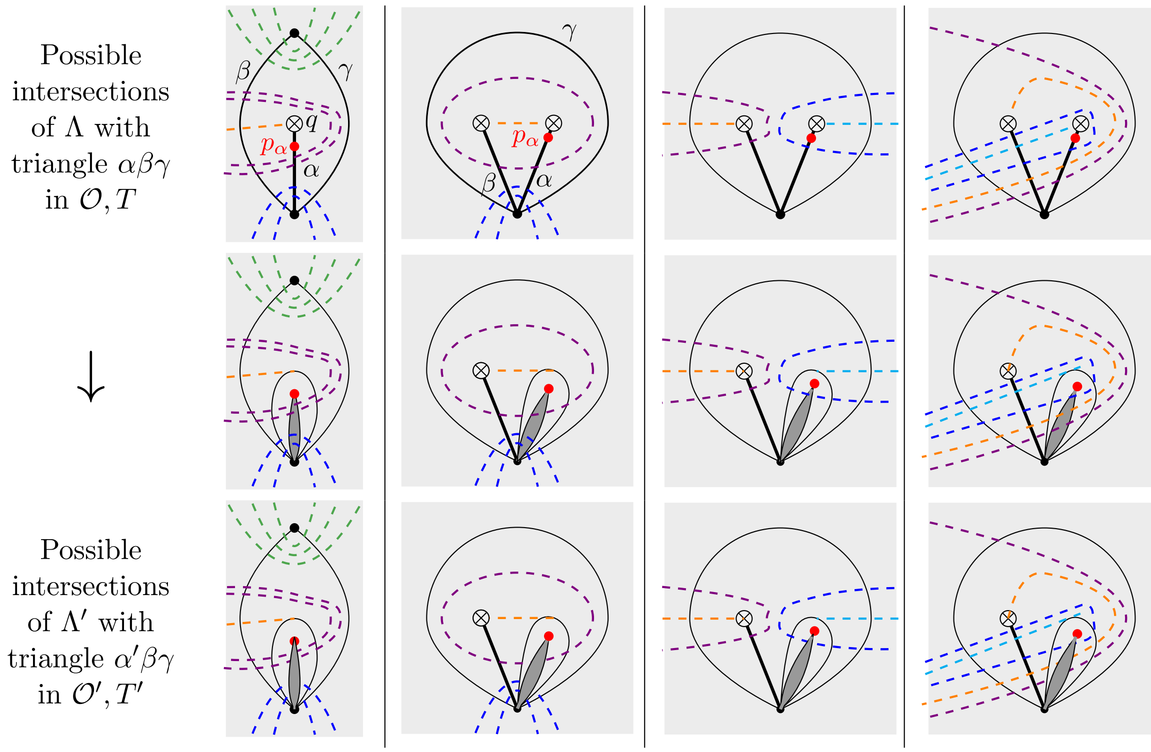

Figure 4.10: Proceeding down each column, first steps in the identity map on quasi-laminations for resection at a non-ordinary arc

on the collection Λ of pairwise-compatible allowable curves in $\mathcal{O}$. We construct a collection Λ' of pairwise-compatible allowable curves in $\mathcal{O}'$ such that for any nonnegative rational weighting on Λ (yielding rational quasi-lamination L), there exists a nonnegative rational weighting on Λ' (yielding rational quasi-lamination L') such that $\mathbf{b}(T', L') = \mathbf{b}(T, L)$. This then implies that $C_\Lambda \cap \mathbb{Q}^n \subseteq C'_{\Lambda'} \cap \mathbb{Q}^n$, and since C_Λ and $C'_{\Lambda'}$ are both rational cones, that $C_\Lambda \subseteq C'_{\Lambda'}$, as desired.

Recall that the resection consists of a collection of resections at arcs of T satisfying conditions 2 and 3 of Definition 4.2.5, and these resections may be performed in any order. We may almost, but not quite, construct Λ' on $\mathcal{O}', T'$ from Λ on $\mathcal{O}, T$ by resecting at one arc $\alpha \in T$ at a time; we will at some point need to appeal to condition 3 of Definition 4.2.5. We consider resection at a non-ordinary arc α first, and then at an ordinary arc.

Assume α is non-ordinary. Let q denote the orbifold point endpoint of α, and denote the other edges of the unique triangle in T of which α is an edge by β and γ, respectively, as one travels around the triangle in the clockwise direction. All of these edges are distinct, and we refer to the triangle by its edges as triangle $\alpha\beta\gamma$. Assume, as in Definition 3.3.1, that pairwise compatible isotopy representatives of the allowable curves in Λ have been chosen to minimize the number of intersections between Λ and the triangle $\alpha\beta\gamma$, and consider all such intersections.

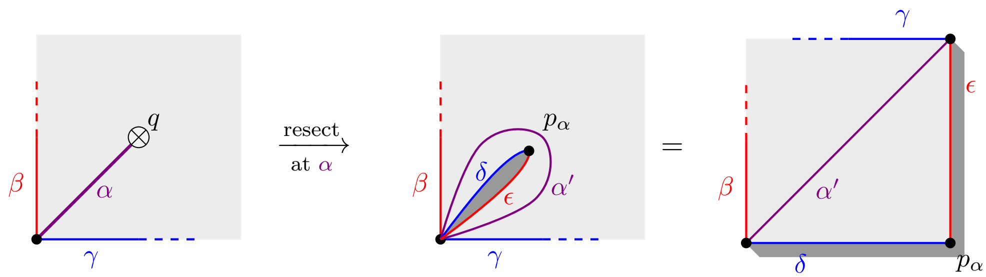

Figure 4.11: Shear coordinates and resection: cf. Tables 3.1 and 3.2

Note that a single curve in Λ may intersect the triangle more than once, and possible infinitely many times if it spirals into a vertex of the triangle.

In the first steps of constructing Λ' from Λ, we only alter the portions of the allowable curves in the interior of triangle $\alpha\beta\gamma$. Up to isotopy, we may assume that the new marked point p_α placed when performing resection at α is situated as illustrated in the pictures in the top row of Figure 4.10. (Note that there are other possibilities for ways in which a given compatible collection of curves can intersect α non-trivially: for example, in the leftmost image, we could have alternatively depicted the purple and orange curves as exiting the triangle through γ. However, the four images are general enough to make clear where the point p_α must be placed.) We may further assume that, after performing resection at α and thereby obtaining arc $\alpha' \in T'$, that the curves intersect the new boundary component as illustrated in the corresponding pictures in the middle row of Figure 4.10. Now, remove from each curve its intersection with the new boundary component, as illustrated in the bottom row of Figure 4.10, thereby potentially cutting it into smaller curves.

This procedure results in a collection Λ' of curves in $\mathcal{O}'$. If each curve in Λ' inherits the weight of the curve in L from whence it came, we obtain a collection of curves L' with non-negative rational weightings. These curves may not be allowable, nor are they necessarily pairwise-compatible; however, if we permit ourselves to compute the shear coordinates of L' with respect to T' regardless, we would like to obtain the shear coordinates of the original collection L with respect to T. For this to be the case, we claim that we must alter certain curves inside the triangle to account for the fact that $\alpha \in T$ was non-ordinary while $\alpha' \in T$ is ordinary. In particular, if α was a double arc ($q \in \mathbf{Q}_2$), then we show that we need only alter one type of curve in Λ'. If α was a pending arc ($q \in \mathbf{Q}_{1/2}$), then we show that we must alter a second type of curve as well. Label the two boundary segments which make up the new boundary component of $\mathcal{O}'$: name as ϵ the segment encountered when traveling along the boundary from p_α keeping $\mathcal{O}'$ on the left, and name as δ the other segment, encountered when traveling along the boundary from p_α keeping $\mathcal{O}'$ on the right. (See the central image in Figure 4.11.)

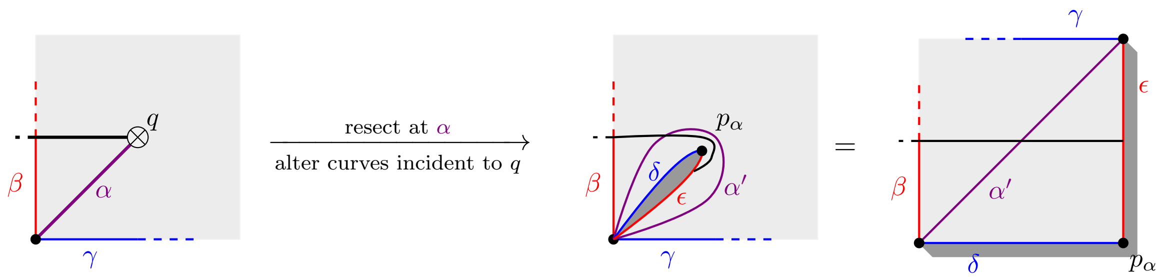

Figure 4.12: Adjustment of non-ordinary curves incident to q in the identity map on quasi-laminations

- Let $\lambda' \in \Lambda'$ be a curve arising from a non-ordinary curve $\lambda \in \Lambda$ with endpoint at q. That is, λ' connects β (resp. γ) to a "floating end" on the arc α'. (For example, see the orange and cyan curves in Figure 4.10.) Extend λ' from this floating end by a short curve which terminates at an unmarked point x on boundary segment ϵ (resp. δ) without intersecting any arcs in T or curves in Λ' as illustrated in Figure 4.12. Comparing the lefthand image of the figure with the top row of Table 3.2 and the righthand image with the top row of Table 3.1, we see that $\mathbf{b}_{\alpha'}(T', \lambda') = \mathbf{b}_{\alpha}(T, \lambda)$, and the shear coordinates of L' with respect to every other arc in the triangulation are unaffected by the adjustment.

- Suppose α is a pending arc, and let $\lambda' \in \Lambda'$ be a curve which connects β to an unmarked point y on boundary segment δ (resp. γ to an unmarked point on ϵ) after crossing α', and arose under resection from a curve $\lambda \in \Lambda$ which intersected β, then α, then β (resp. γ, then α, then γ). In particular, resection at α cut λ into two pieces: in addition to λ', the other curve created, λ'', connects β to an unmarked point on ϵ. Replace the portion of λ' between α' and y by a short curve which terminates at an unmarked point x on boundary segment ϵ (resp. δ) without intersecting any arcs in T or curves in Λ' as illustrated in Figure 4.13. Comparing the lefthand image of the figure with the bottom four rows of Table 3.2 and the righthand image with the top two rows of Table 3.1, we see that $\mathbf{b}_{\alpha'}(T', \lambda') + \mathbf{b}_{\alpha'}(T', \lambda'') = \mathbf{b}_{\alpha}(T, \lambda)$, and the shear coordinates of L' with respect to every other arc in the triangulation are unaffected by the adjustment. (We do *not* perform this adjustment if α was a double arc.)

Reusing Λ' for the collection of curves resulting from these alterations, and L' for the weighted collection, we now have $\mathbf{b}(T', L') = \mathbf{b}(T, L)$, as desired.

The next step is to alter any curves in Λ' which are not allowable. Since the original collection Λ consisted of pairwise compatible curves with isotopy representatives chosen to minimize intersections with triangle $\alpha\beta\gamma$, each such non-allowable curve must be one of the following:

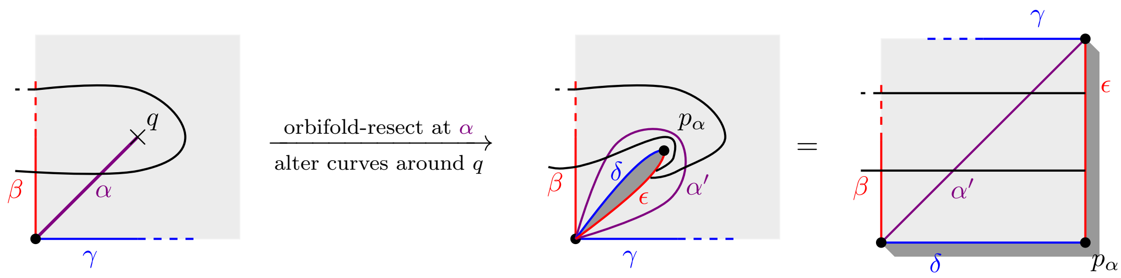

Figure 4.13: Further adjustment of curves if $q \in \mathbf{Q}_{1/2}$ in the identity map on quasi-laminations

- A curve with two endpoints on a boundary component that is contractible to a portion of the boundary component containing one marked point. (For example, see the blue curves in the second column of Figure 4.10 if either β or γ (or both) are boundary segments.)

- A curve with both endpoints on the same boundary segment that cuts out, together with the portion of the boundary between its endpoints, a disc containing a unique point in $\mathbf{M} \cup \mathbf{Q}$. (For example, this will occur if the arc resected in the second column of Figure 4.10 is pending, after the purple closed curve is altered as described in the previous step.)

The first type of non-allowable curve has shear coordinates of zero with respect to every arc in T', so we can delete all such curves from the collection Λ' without changing the shear coordinate vector $\mathbf{b}(T', L')$ of the corresponding quasi-lamination L'. The second type of non-allowable curve cannot be dealt with in this manner, as it may make a nonzero contribution to $\mathbf{b}(T', L')$. Instead, as in [24], we replace each such curve which cuts out a once-punctured disc, depicted in the second image in Figure 3.8, by the corresponding conjugate pair of curves depicted in Figure 3.9, each endowed with the same weight as the curve that was replaced. Similarly, we replace each curve which cuts out a disc containing one orbifold point q' and no punctures, depicted in the third image of Figure 3.8), by the corresponding non-ordinary curve depicted in Figure 3.9 connecting the boundary to q', endowed with the same weight as the replaced curve if $q' \in \mathbf{Q}_2$ and double the weight of the replaced curve if $q' \in \mathbf{Q}_{1/2}$. In both cases, the shear coordinates of the replacement curve(s) equal(s) the shear coordinates of the non-allowable curve which was replaced. Furthermore, the replacement curve(s) are pairwise-compatible with the non-altered allowable curves in the collection (and the conjugate pair of curves are compatible with one another).

It is possible that distinct curves in the new collection may coincide up to isotopy, in which case we delete any such repetitions and adjust weights accordingly. We again reuse Λ' and L' for the collection and weighted collection, respectively. All of these modifications preserve the shear coordinates of L', and now Λ' consists solely of allowable curves in $\mathcal{O}$. It remains only to show that the curves in Λ' are pairwise compatible.

If all of the curves in the original collection Λ were pairwise non-intersecting, then the images of those curves which comprise Λ' will be as well, and are therefore pairwise compatible. Potentially, however, Λ may contain a conjugate pair of curves. If the puncture at which they have opposite spiral directions is not the marked point endpoint of α, then the image of the pair in Λ' is still a pairwise-compatible collection (which contains a conjugate pair of curves). Even if this puncture is the marked point endpoint of α, the image of the pair in Λ' is still pairwise compatible, as in particular, all non-discarded curves are non-intersecting. We therefore conclude, as desired, that Λ' is a collection of pairwise-compatible allowable curves in $\mathcal{O}'$, with a non-negative rational weighting on these curves producing a rational quasi-lamination L' such that $\mathbf{b}(T', L') = \mathbf{b}(T, L)$.

Now, assume α is ordinary. If α falls into either case (i) or (ii) of Proposition 4.2.8 (as illustrated in Figures 4.3 and 4.4, respectively) then the construction of Λ' is almost entirely handled by the proof of [24, Theorem 1.7], the analogous result for surfaces. The only adjustment required for orbifolds is to also replace any non-allowable curve in Λ' of which has both endpoints on the same boundary segment and cuts out, together with the portion of the boundary between its endpoints, a disc containing a single orbifold point q' and no punctures. This replacement is handled precisely as described above. Indeed, the entire construction in the proof of [24, Theorem 1.7] is very similar to that described above, except that it does not require the step where curves are altered to preserve shear coordinates. The proof also separately handles the case when ordinary arcs β and γ coincide. Finally, if α falls into any of the remaining cases of resection at ordinary arcs in Proposition 4.2.8, as depicted in Figures 4.5, 4.6, and 4.7, the construction of Λ' proceeds precisely as described above, except that again we may omit the step where curves are altered to preserve shear coordinates. □

Recall from Section 2.2.2 that given an $n \times n$ exchange matrix B and an underlying ring R (which we take to be $\mathbb{Z}, \mathbb{Q}$ or $\mathbb{R}$), that the notation R^B denotes the partial linear structure $R^B = (R^n, R, \mathcal{V}_{\eta^B})$ (see Definition 2.2.22) where the set of valid relations $\mathcal{V}_{\eta^B}$ consists of B-coherent linear relations with coefficients in R (see Definition 2.2.19). Given a second $n \times n$ exchange matrix B', recall that a linear map $\lambda : R^B \to R^{B'}$ is mutation-linear if and only if it maps B-coherent linear relations to B'-coherent linear relations.

We quote a result from [24] and then use it in conjunction with results from Chapters 2 and 3 to prove that a rational version of Phenomenon I holds in the case of resection of orbifolds as an an immediate consequence of Theorem 4.1.2. The proof is identical to that of [24, Theorem 1.2] with the relevant orbifold results substituted for the surface ones. However, we provide it here as it is a nice illustration of the connection between Phenomena I and II.

Proposition 4.3.1. *[24, Proposition 3.3] Suppose that B and B' are $n \times n$ exchange matrices,* *and let $\mathcal{F}_B^{\mathbb{Q}}$ and $F_{B'}^{\mathbb{Q}}$ denote the rational parts of mutation fans $\mathcal{F}_B$ and $\mathcal{F}_{B'}$, respectively. If the identity map $\mathbb{Q}^B \to \mathbb{Q}^{B'}$ is mutation-linear, then $\mathcal{F}_B^{\mathbb{Q}} \cap \mathbb{Q}^n$ refines $\mathcal{F}_{B'}^{\mathbb{Q}} \cap \mathbb{Q}^n$. Assuming also that $\mathbb{Q}^B$ admits a cone basis, the identity map is mutation-linear if and only if $\mathcal{F}_B^{\mathbb{Q}} \cap \mathbb{Q}^n$ refines $\mathcal{F}_{B'}^{\mathbb{Q}} \cap \mathbb{Q}^n$.*

Proof of Theorem 4.1.1. By Theorem 3.4.17 and Proposition 2.2.36, $\mathbb{Q}^{B(T)}$ admits a positive basis and therefore a cone basis. Since the Null Tangle Property implies the Curve Separation Property by Corollary 3.4.16, it follows from Theorem 3.4.3 that $\mathcal{F}_{\mathbb{Q}}(T)$ is the rational part of $\mathcal{F}_{B(T)}$ and $\mathcal{F}_{\mathbb{Q}}(T')$ is the rational part of $\mathcal{F}_{B(T)}$. Then by Proposition 4.3.1, the identity map $\mathbb{Q}^{B(T)} \to \mathbb{Q}^{B(T')}$ is mutation-linear. □

4.4 Ring homomorphisms between cluster algebras

In this section, we prove Theorem 4.1.5: If $\mathcal{O}$ is a disk with a unique orbifold point and no punctures, and T is a tagged triangulation of $\mathcal{O}$ such that $B(T)$ is acyclic, then if $B = B(T)$ dominates B', the map $\nu_{\mathbf{z}}$ defined below in (4.1) is an injective, **g**-vector preserving ring homomorphism from $\mathcal{A}_\bullet(B')$ into $\mathcal{A}_\bullet(B)$ sending cluster variables to cluster variables.

We begin by reviewing the setup of [24, Theorem 1.20], quoted below as Theorem 4.4.1, the analogous result for once-punctured and empty disks upon which our work is based. All of the exchange matrices which arise in this way are of finite type (and thus any matrices they dominate are as well). In particular, let T denote, respectively, a tagged triangulation of an empty disk, a disk with one puncture, a disk with one orbifold point of weight 1/2, or a disk with one orbifold point of weight 2. Then $B(T)$ is of respective finite type A_n, D_n, C_n, or B_n.

Recall from Section 2.2.1 that $\mathcal{A}_\bullet(B)$ denotes the principal-coefficients cluster algebra associated to B. Assuming B is an $n \times n$ exchange matrix, let $\mathbf{x} = (x_1, \ldots, x_n)$ denote the initial cluster, and $\mathbf{y} = (y_1, \ldots, y_n)$ the initial choice of (principal) coefficients. The cluster algebra $\mathcal{A}_\bullet(B)$ is therefore a subring of the ring of Laurent polynomials in $x_1, \ldots, x_n$ with coefficients given by integer polynomials in $y_1, \ldots, y_n$. Likewise, let $\mathbf{x}' = (x'_1, \ldots, x'_n)$ and $\mathbf{y}' = (y'_1, \ldots, y'_n)$ denote the initial cluster and choice of (principal) coefficients for $\mathcal{A}_\bullet(B')$. Then a ring homomorphism from $\mathcal{A}_\bullet(B')$ to $\mathcal{A}_\bullet(B)$ is entirely determined by its values on the set $\{x'_1, \ldots, x'_n, y'_1, \ldots, y'_n\}$.

Recall from Definition 2.2.11 that **g**-vectors are a $\mathbb{Z}^n$ grading on the principal-coefficients cluster algebra. In the finite type setting we consider here, for each integer vector $\mathbf{a} \in \mathbb{Z}^n$, there exists a unique cluster monomial whose **g**-vector is **a**. (A ***cluster monomial*** is a monomial in cluster variables which all belong to a common cluster.) For each $k \in [n]$, set z_k to be the cluster monomial whose **g**-vector is the k^{th} column of B minus the k^{th} column of B'. Define a ring homomorphism from $\mathcal{A}_\bullet(B')$ to the ring of Laurent polynomials containing $\mathcal{A}_\bullet(B)$ by extending the following set map $\nu_{\mathbf{z}}$ on $\{x'_1, \ldots, x'_n, y'_1, \ldots, y'_n\}$:

$$\begin{aligned} \nu_{\mathbf{z}}(x'_k) &= x_k \text{ for each } k \in [n] \\ \nu_{\mathbf{z}}(y'_k) &= y_k z_k \text{ for each } k \in [n]. \end{aligned} \tag{4.1}$$

If we knew that $\nu_{\mathbf{z}}$ sent every cluster variable in $\mathcal{A}_\bullet(B')$ to a cluster variable ib, or any element of $\mathcal{A}_\bullet(B)$, it would be enough to conclude that that the image of $\nu_{\mathbf{z}}$ is contained in $\mathcal{A}_\bullet(B)$. It is clear by construction that $\nu_{\mathbf{z}}$ sends *initial* cluster variables of $\mathcal{A}_\bullet(B')$ to (initial) cluster variables of $\mathcal{A}_\bullet(B)$, and furthermore that the map is **g**-vector preserving on this set. Yet beyond the initial cluster, there is no immediate reason why the image of $\nu_{\mathbf{z}}$ should be contained in $\mathcal{A}_\bullet(B)$, let alone send cluster variables to cluster variables with the same **g**-vectors. However, [24] shows this is true when B is an acyclic signed adjacency matrix of type A_n or D_n:

Theorem 4.4.1. *[24, Theorem 1.20] Suppose* $(\mathbf{S}, \mathbf{M})$ *is a once-punctured or empty disk and suppose* T *is a triangulation of* $\mathcal{O} = (\mathbf{S}, \mathbf{M})$ *such that* $B(T)$ *is acyclic. If* $B(T)$ *dominates* B'*, then both parts of Phenomenon IV occur.*

Since this result is not at all obvious a priori, we summarize the motivating ideas and argument from [24] in our more general orbifold setting. Recall from Section 3.3.3 that cluster variables in $\mathcal{A}_\bullet(B(T))$ are in bijection with tagged arcs in $\mathcal{O}$, where the initial cluster variables $\{\mathbf{x}\} = \{x_\gamma : \gamma \in T\}$ correspond to the tagged arcs in T. Furthermore, the initial choices of coefficients correspond to the elementary quasi-laminations associated to T. That is, we may write $\{\mathbf{y}\} = \{L_\gamma : \gamma \in T\}$ where L_γ is the elementary quasi-lamination associated to γ consisting of a single allowable curve of weight 1. Thus elements of $\mathcal{A}_\bullet(B(T))$ can be represented as sums of terms, each term of which is a monomial in the variables x_α for tagged arcs α in $\mathcal{O}$ (*all* tagged arcs, not just those in T, and tagged plain when $\mathcal{O}$ has one puncture and no boundary components) with "coefficient" given by a monomial in the elementary quasi-laminations L_γ. Recall further that there is also a bijection, κ, between the set of tagged arcs in T and allowable curves on $\mathcal{O}$ which are not closed or semi-closed. For the surfaces and orbifolds under consideration in this section, namely, empty disks, disks with one puncture, and disks with one orbifold point, there are no closed or semi-closed allowable curves, so κ is a bijection between the set of tagged arcs in T and the complete set of allowable curves in T. Finally, for each tagged arc α in $\mathcal{O}$, by (3.11) the **g**-vector for the corresponding cluster variable x_α is given by $-\mathbf{b}(T^*, \kappa(\alpha^*))$, where T^* denotes the reversed triangulation and α^* denotes the reversed arc. (Since "reversing" swaps the weights on all orbifold points, so if $\mathcal{O}$ is a surface as in Theorem 4.4.1, then $T^* = T$ and $\alpha^* = \alpha$.)

If $B' = B(T')$ for some triangulation T' induced by a resection $\mathcal{O}'$ of $\mathcal{O}$ compatible with T, then all of the above holds for $\mathcal{A}_\bullet(B(T'))$ as well. Furthermore, by Theorem 4.1.2, the rational quasi-lamination fan $\mathcal{F}_{\mathbb{Q}}(T)$ refines $\mathcal{F}_{\mathbb{Q}}(T')$, so in particular, every ray of $\mathcal{F}_{\mathbb{Q}}(T')$ is also a ray of $\mathcal{F}_{\mathbb{Q}}(T)$. But rays in the rational quasi-lamination fan are spanned by the shear coordinates of allowable curves, so this correspondence is a shear-coordinate-preserving injective map from allowable curves on $\mathcal{O}'$ to allowable curves on $\mathcal{O}$. Thus, since allowable curves are in bijection with tagged arcs, which in turn are in bijection with cluster variables, and $\mathbf{g}$-vectors are given by shear coordinates, we have a natural candidate for desired map from $\mathcal{A}_\bullet(B')$ into $\mathcal{A}_\bullet(B)$. Let χ denote this map: in particular, for each tagged arc α' on $\mathcal{O}'$, define $\chi(\alpha') = \alpha$ to be the tagged arc on $\mathcal{O}$ such that $\mathbf{b}((T')^*, \kappa((\alpha')^*)) = \mathbf{b}(T^*, \kappa(\alpha^*))$. Extend χ to a map on $\mathrm{Var}_\bullet(B(T')) \cup \{\mathbf{y}'\}$ which agrees with $\nu_{\mathbf{z}}$ as defined in (4.1) by setting $\chi(L'_{\gamma'}) = L_\gamma z_\gamma$. Then χ maps every exchange relation in $\mathcal{A}_\bullet(B(T'))$ to an equation relating elements of $\mathcal{A}_\bullet(B)$: if each such equation is a *valid* relation in $\mathcal{A}_\bullet(B)$ (that is, if the equation is true), then χ is a ring homomorphism [24, Proposition 5.2]. Furthermore, since χ and $\nu_{\mathbf{z}}$ agree on the set $\mathrm{Var}_\bullet(B(T')) \cup \{\mathbf{y}'\}$, they must coincide. Thus $\chi = \nu_{\mathbf{z}}$ is our desired $\mathbf{g}$-vector-preserving ring homomorphism mapping cluster variables to cluster variables. Injectivity can be checked using another result from [24]:

Proposition 4.4.2. *[24, Proposition 5.3] Suppose B dominates B'. If there is at most one index k such that there exists an index i with $b_{ik} < b'_{ik} \leq 0$, then the homomorphism $\nu_{\mathbf{z}}$ is injective.*

For any exchange matrix B' dominated by B, there exists a sequence of exchange matrices, starting at B and ending at B', such that each matrix is obtained from the latter by either

(i) setting two non-zero symmetric entries in B to both equal zero in B', or

(ii) decreasing the absolute value of a unique entry in B to obtain a non-zero entry in B'.

Both of these operations satisfy the hypotheses of Proposition 4.4.3, and we show below that they correspond to resection operations. The following result from [24] then lets us factor $\nu_{\mathbf{z}}$ through these individual steps (the notation $\nu_{\mathbf{z}}^{B',B}$ makes explicit the factor in question), so that it is sufficient to check each case on its own.

Proposition 4.4.3. *[24, Proposition 5.28] Suppose B is acyclic and dominates B', which dominates B''. If $\nu_{\mathbf{z}}^{B'',B'}$ sends cluster variables of $\mathcal{A}_\bullet(B'')$ to cluster variables of $\mathcal{A}_\bullet(B')$ and $\nu_{\mathbf{z}}^{B',B}$ sends cluster variables of $\mathcal{A}_\bullet(B')$ to cluster variables of $\mathcal{A}_\bullet(B)$, then the composition $\mathcal{A}_\bullet(B'') \xrightarrow{\nu_{\mathbf{z}}} \mathcal{A}_\bullet(B') \xrightarrow{\nu_{\mathbf{z}}} \mathcal{A}_\bullet(B)$ sends cluster variables of $\mathcal{A}_\bullet(B'')$ to cluster variables of $\mathcal{A}_\bullet(B)$. This composition equals $\nu_{\mathbf{z}}^{B'',B} : \mathcal{A}_\bullet(B'') \to \mathcal{A}_\bullet(B)$.*

We next examine the hypotheses of Theorem 4.1.5, and let $\mathcal{O}$ be a disk with a unique orbifold point q and no punctures (with $n+1 \geq 2$ marked points on its boundary), and T a tagged

$\mathcal{O}, T =$ 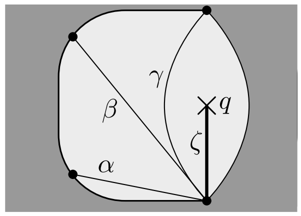

$\mathcal{O}, T =$ 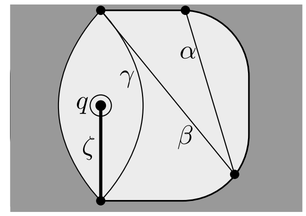

$$B(T) = \begin{array}{c} \alpha \\ \beta \\ \gamma \\ \zeta \end{array} \begin{bmatrix} 0 & -1 & 0 & 0 \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -2 \\ 0 & 0 & 1 & 0 \end{bmatrix} \qquad B(T) = \begin{array}{c} \alpha \\ \beta \\ \gamma \\ \zeta \end{array} \begin{bmatrix} 0 & 1 & 0 & 0 \\ -1 & 0 & -1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & -2 & 0 \end{bmatrix}$$

Figure 4.14: Examples of triangulations of a disk with a unique orbifold point q and no punctures which have acyclic signed-adjacency matrices

triangulation of $\mathcal{O}$ such that the $n \times n$ signed-adjacency matrix $B = B(T)$ is acyclic, as depicted in Figure 4.14. (Since $\mathcal{O}$ has no punctures, T has all arcs tagged plain and therefore applying τ^{-1} yields an (untagged) triangulation with no self-folded triangles.) Since $B(T)$ is acyclic, every triangle of T has at least one edge on the boundary, and the unique non-ordinary triangle in T consists of a boundary segment, an ordinary arc, and a non-ordinary arc. For the remainder of this paragraph, denote the ordinary arc by γ and the non-ordinary arc by ζ. If q is of weight $1/2$ then $|b_{\gamma\zeta}| = 2$ and $|b_{\zeta\gamma}| = 1$, and if q is of weight 2 then $|b_{\gamma\zeta}| = 1$ and $|b_{\zeta\gamma}| = 2$. For all other ordinary arcs in T besides γ, $b_{\alpha\zeta} = b_{\zeta\alpha} = 0$, and for any $\alpha, \beta \in T \setminus \{\zeta\}$, either $b_{\alpha\beta} = b_{\beta\alpha} = 0$ or else $|b_{\alpha\beta}| = |b_{\beta\alpha}| = 1$. If resection is performed at a single ordinary arc $\alpha \in T \setminus \{\zeta\}$, then $\mathcal{O}'$ consists of two connected components: a disk containing q with $m \geq 2$ marked points on its boundary and an empty disk with $n + 1 - m$ marked points. Under resection at ζ, the new orbifold $\mathcal{O}'$ is a single empty disk with $n + 3$ marked points on its boundary. Regardless, T' is acyclic since resection does not introduce any new arc adjacencies, and furthermore each component of $\mathcal{O}'$ satisfies the hypotheses of both Theorem 4.1.5 or Theorem 4.4.1. We now show that there exist resections of $\mathcal{O}$ compatible with T that result in the operations on exchange matrices enumerated above as (i) and (ii).

(i) Suppose B' is obtained by setting any pair $b_{\alpha\beta} = -b_{\beta\alpha}$ of nonzero symmetric entries in $B(T)$ equal to zero in B'. Then $B' = B(T')$ may be realized as the signed adjacency matrix of the induced triangulation T' on $\mathcal{O}'$ obtained by ordinary resection at a single arc in T. If both $\alpha, \beta \in T$ are ordinary, then the result is accomplished by resecting at either α or β, placing the new marked point (p_α or p_β) placed in the interior of the triangle with edges α, β, and a boundary segment: choose the arc such that, after resection, it is not in the same component of $\mathcal{O}$ as the orbifold point. (No other entries of B are affected by the operation since T is acyclic.) If, without loss of generality, α is ordinary and β

is non-ordinary, then perform resection at α, placing the point p_α in the non-ordinary triangle shared by α and β. Again, α is not in the same component of $\mathcal{O}$ as the orbifold point.

(ii) Suppose B' is obtained by decreasing the absolute value of a unique entry in $B(T)$ to obtain a non-zero entry in B'. By our description of $B(T)$, there is only one way to do this: either decreasing $|b_{\gamma\zeta}|$ by 1 if $q \in \mathbf{Q}_{1/2}$, or decreasing $|b_{\zeta\gamma}|$ by 1 if $q \in \mathbf{Q}_2$. Regardless, $B = B(T')$ is realized as the signed adjacency matrix of the induced triangulation T' on $\mathcal{O}'$ obtained by resection at ζ.

We handle the first case in Proposition 4.4.4 using folding ideas from Section 3.5 (orbifold folding) and Section 2.3 (matrix and cluster algebra folding) and applying Theorem 4.4.1. We handle the second case in Proposition 4.4.5, following the outline used in similar proofs in [24, Section 5.4]: first explicitly describing the map χ on tagged arcs and elementary quasi-laminations, and then showing that χ maps exchange relations to exchange relations.

Proposition 4.4.4. *Suppose $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ is a disk with one orbifold point q and no punctures (and at least 2 marked points on its boundary). Let T be a (untagged) triangulation of $\mathcal{O}$. Suppose $\mathcal{O}', T'$ are obtained from $\mathcal{O}, T$ by a resection compatible with T at a single ordinary arc $\alpha \in T$ with distinct endpoints, both on the boundary. Suppose also that the point p_α used to construct the resection is in a triangle of T having exactly one edge on the boundary and that the orbfiold point q is in the component of $\mathcal{O}'$ not containing α. Then both parts of Phenomenon IV occur.*

Proof. Regardless of the weight of the unique orbifold point q in $\mathcal{O}$, there exists a marked surface $\hat{\mathcal{O}} = (\hat{\mathbf{S}}, \hat{\mathbf{M}})$ with tagged triangulation $\hat{T}$ and folding symmetry σ such that $\hat{O}, \hat{T}$ folds to $\mathcal{O} = \pi_\sigma(\hat{\mathcal{O}}), T = \pi_\sigma(\hat{T})$ under σ (see Definition 3.5.3), and thus $\pi_\sigma(B(\hat{T})) = B(T)$ by (3.12). These surfaces satisfy the hypotheses of 4.4.1: namely, if q is of weight $1/2$, then $\hat{\mathcal{O}}$ is an empty disk and σ is a single prime folding about the (unmarked) center point, as depicted in the first column of Figure 4.15. If q is of weight 2, then $\hat{\mathcal{O}}$ is a once-punctured disk and σ is a single local folding about the puncture, as depicted in the first column of Figure 4.16.

These figures also illustrate that unfolding and folding commute with the resection at an ordinary arc α on $\mathcal{O}$ and the corresponding resection on $\mathcal{O}'$ at all arcs in the preimage $\pi_\sigma^{-1}(\alpha)$. That is, letting $\hat{T}'$ denote the induced triangulation on the orbifold $\hat{\mathcal{O}}'$ obtained by resecting $\hat{\mathcal{O}}$ under the lift of the resection on $\mathcal{O}$, σ is a folding symmetry of $\hat{\mathcal{O}}', \hat{T}'$, and $\pi_\sigma(\hat{\mathcal{O}}') = \mathcal{O}'$ and likewise $\pi_\sigma(\hat{T}') = T'$. By Proposition 4.2.8, $B(\hat{T})$ dominates $B(\hat{T}')$, so by Theorem 4.4.1 we conclude that $\hat{\nu}_{\mathbf{z}} : \mathcal{A}_\bullet(B(\hat{T}')) \to \mathcal{A}_\bullet(B(\hat{T}))$ is an injective, $\mathbf{g}$-vector preserving ring homomorphism sending cluster variables to cluster variables.

Let $\mathcal{A}_\bullet^\sigma(B(\hat{T}))$ denote the subalgebra of $\mathcal{A}_\bullet(B(\hat{T}))$ generated by all clusters in seeds reachable from the initial seed by orbit-mutation (see Definition 2.3.3), and likewise $\mathcal{A}_\bullet^\sigma(B(\hat{T}'))$ for

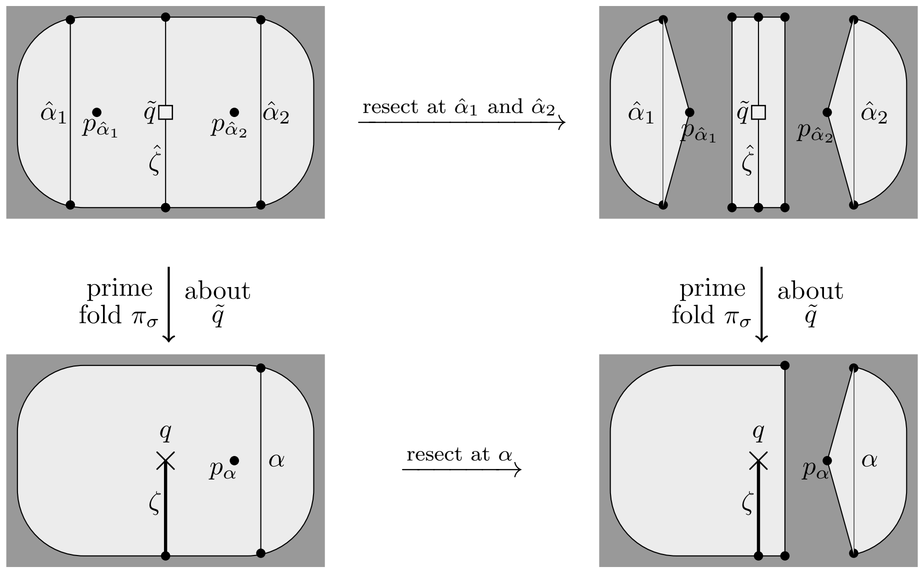

Figure 4.15: Commutativity of prime folding and σ-symmetric resection at ordinary arcs

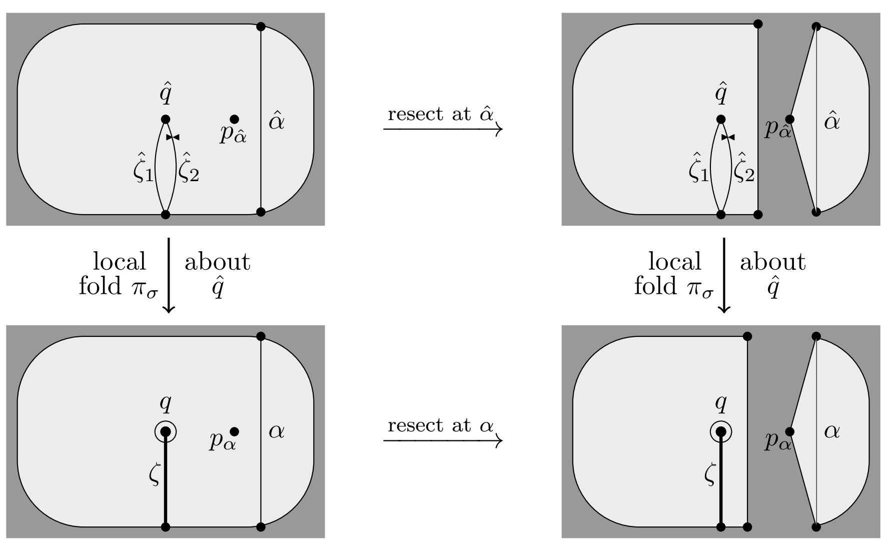

Figure 4.16: Commutativity of local folding and σ-symmetric resection at ordinary arcs

$$
\begin{array}{ccc}
\mathcal{A}_\bullet(B(\hat{T})) & \xrightarrow{\hat{\nu}_{\mathbf{z}}} & \mathcal{A}_\bullet(B(\hat{T}')) \\
\cup\!\text{I} & & \cup\!\text{I} \\
\mathcal{A}^\sigma_\bullet(B(\hat{T})) & \xrightarrow{\hat{\nu}_{\mathbf{z}}|_{\mathcal{A}^\sigma_\bullet(B(\hat{T}))}} & \mathcal{A}^\sigma_\bullet(B(\hat{T}')) \\
\pi_\sigma^{-1}\Big\uparrow \cong & \circlearrowright & \cong \Big\downarrow \pi_\sigma \\
\mathcal{A}_\bullet(B(T)) & \overset{\nu_{\mathbf{z}}}{\dashrightarrow} & \mathcal{A}_\bullet(B(T'))
\end{array}
$$

Figure 4.17: Defining $\nu_{\mathbf{z}}$ as a folding of $\hat{\nu}_{\mathbf{z}}$ under σ

$\mathcal{A}_\bullet(B(\hat{T}'))$. By Theorem 2.3.10, $\pi_\sigma(\mathcal{A}^\sigma_\bullet(B(\hat{T}))) \cong \mathcal{A}_\bullet(B(T))$ and $\pi_\sigma(\mathcal{A}^\sigma_\bullet(B(\hat{T}'))) \cong \mathcal{A}_\bullet(B(T'))$, where π_σ denotes the quotient map under the symmetry σ. Thus we may equivalently define $\nu_{\mathbf{z}}$ on $\mathrm{Var}_\bullet(B(T')) \cup \{\mathbf{y}'\}$ by the following composition:

$$\nu_{\mathbf{z}} = \pi_\sigma \circ \hat{\nu}_{\mathbf{z}}|_{\mathcal{A}^\sigma_\bullet(B(\hat{T}))} \circ \pi_\sigma^{-1}.$$

We conclude from the corresponding properties of $\hat{\nu}_{\mathbf{z}}$ that $\nu_{\mathbf{z}}$ is a $\mathbf{g}$-vector preserving injective ring homomorphism from $\mathcal{A}_\bullet(B(T'))$ into $\mathcal{A}_\bullet(B(T))$ taking cluster variables to cluster variables. The construction is summarized as a commutative diagram in Figure 4.17. □

Proposition 4.4.5. *Suppose $\mathcal{O} = (\mathbf{S}, \mathbf{M}, \mathbf{Q})$ is a disk with one orbifold point q, no punctures, and at least 3 marked points on its boundary. Let T be a (untagged) triangulation of $\mathcal{O}$. Suppose $\mathcal{O}', T'$ are obtained from $\mathcal{O}, T$ by a resection compatible with T consisting of resection along the unique non-ordinary arc $\alpha \in T$, where the unique triangle in T with edge α has exactly one edge on the boundary. Then both parts of Phenomenon IV occur.*

Proof. Throughout, we use the labels given in Figure 4.18, which depicts and labels some of the arcs in T as well as the corresponding arcs in T'. For generality, neither of the edges β and γ of the unique triangle in T containing ζ is depicted as a boundary segment, although exactly one is by hypothesis. The marked point labeled r is the point closest to p_2, moving along the boundary from p_2 while keeping the orbifold on the right. (Note that if β is a boundary segment, then $r = p_1$.)

Let ζ' be any tagged arc in $\mathcal{O}'$. Recall from (3.5) that the $\mathbf{g}$-vector of $x_{\zeta'}$ is given by the negative of the shear coordinates of the curve $\kappa((\zeta')^*)$ with respect to $(T')^*$, where, by Definition 3.3.15, $(\mathcal{O}')^*, (T')^*$ denotes the triangulated reversed orbifold obtained from $\mathcal{O}', T'$ by switching the weights of all orbifold points in $\mathcal{O}'$ and switching the types of all non-ordinary arcs in T', and $(\zeta')^*$ denotes the arc corresponding to (ζ') in $(\mathcal{O}')^*$.) We explicitly identify the tagged arc ζ in $\mathcal{O}$ such that $\kappa(\zeta^*)$ has those same shear coordinates with respect to T^*, and

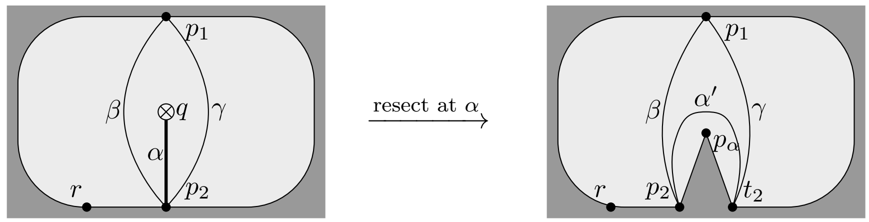

Figure 4.18: Resection at non-ordinary arc α in a disc with one orbifold point and no punctures

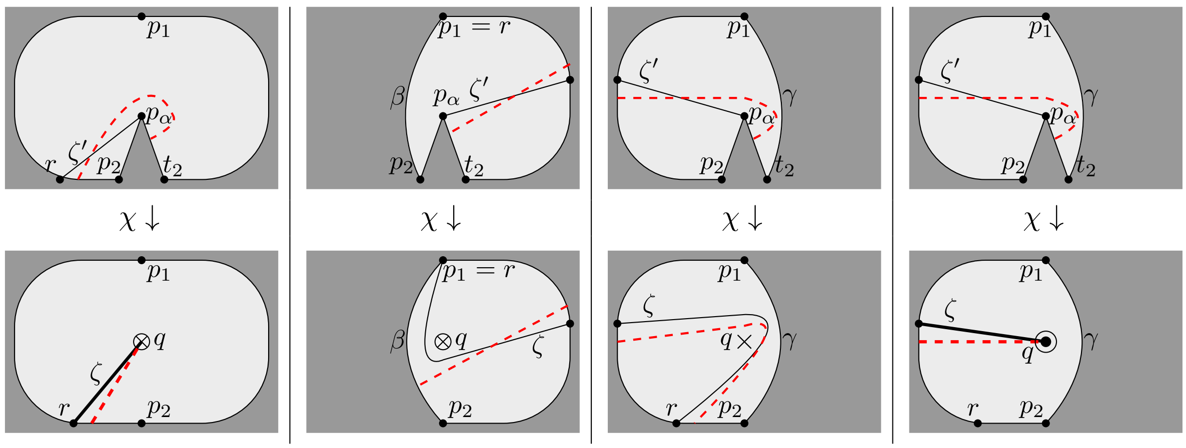

Figure 4.19: Illustrations of the map χ on tagged arcs ζ' in $\mathcal{O}'$ with an endpoint at p_α

define $\zeta = \chi(\zeta')$. (Note that $\alpha = \chi(\alpha')$, where α' is the new ordinary arc in $\mathcal{O}'$ constructed in the resection as depicted in Figure 4.18.)

- Suppose ζ' does not intersect the interior of the quadrilateral $p_1p_2p_\alpha t_2$. (That is, neither endpoint of ζ' is p_α, if β is a boundary segment, neither endpoint is p_2, and if γ is a boundary segment, neither endpoint is t_2.) Then ζ is obtained from ζ' by the natural inclusion (preserving taggings) taking the closure of $\mathcal{O}'$ minus the quadrilateral $p_1p_2p_\alpha t_2$ into $\mathcal{O}$. (Note that this case implies that $\gamma = \chi(\gamma)$ for all tagged arcs $\gamma \in T \setminus \{\alpha\}$.)

- If an endpoint of ζ' is p_α, delete the part of ζ' contained in the quadrilateral $p_1p_2p_\alpha t_2$, include what remains of ζ' into $\mathcal{O}$, and extend as described below and illustrated in Figure 4.19:

  - Suppose the other endpoint of ζ' is r: regardless of whether β or γ is a boundary segment, there is a unique such tagged arc since $\mathcal{O}'$ is an unpunctured disk. (Recall that $r = p_1$ when β is a boundary segment.) Then ζ is obtained by attaching the included arc to the orbifold point q as illustrated in the first column of Figure 4.19. The dashed red curves are $\kappa(\zeta')$ and $\kappa(\zeta)$.

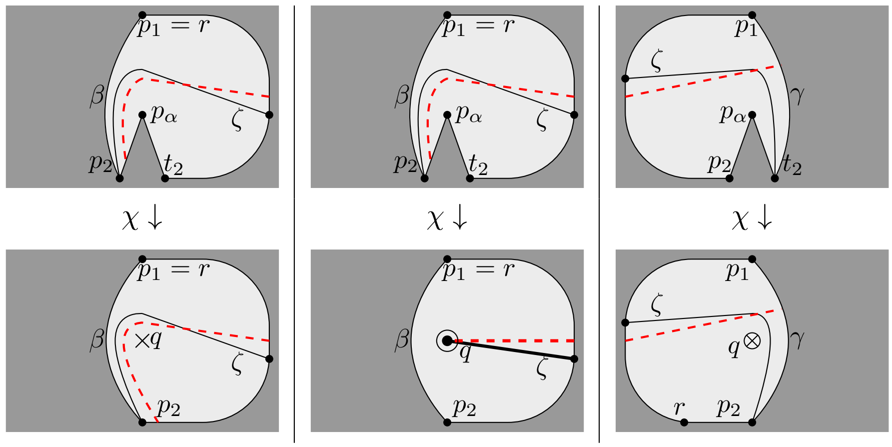

Figure 4.20: Illustrations of the map χ on tagged arcs ζ' in $\mathcal{O}'$ with an endpoint at p_2 or t_2

- Suppose the other endpoint of ζ' is not r and that β is a boundary segment. Then ζ is obtained by attaching the included arc to r such that, together with the portion of the boundary between its endpoints which contains p_2, it cuts out an empty disc. This case is illustrated in the second column of Figure 4.19 (recall $r = p_1$).
- Suppose the other endpoint of ζ' is not r and that γ is a boundary segment. If q is of weight $1/2$, then ζ is obtained by attaching the included arc to r as described above. If q is of weight 2, then ζ is obtained by attaching the included arc to q. These cases are illustrated in the third and fourth columns of Figure 4.19, respectively.

- Suppose β is a boundary segment and one endpoint of ζ' is at p_2, with the other at a point other than t_2, so that in particular $\zeta' \neq \alpha'$. (Recall $\chi(\alpha') = \alpha$.) Then, delete the part of ζ' contained in the quadrilateral $p_1 p_2 p_\alpha t_2$ and include what remains of ζ' into $\mathcal{O}$. If q is of weight $1/2$, then ζ is obtained by attaching the included arc to p_2 such that, together with the portion of the boundary between its endpoints which contains p_1, it cuts out an empty disc. If q is of weight 2, then ζ is obtained by attaching the included arc to q. These cases are illustrated in the first and second columns of Figure 4.20, respectively.

- Suppose γ is a boundary segment and one endpoint of ζ' is at t_2, with the other at a point other than p_2 (i.e., $\zeta' \neq \alpha'$). Then delete the part of ζ' contained in the quadrilateral $p_1 p_2 p_\alpha t_2$ and include what remains of ζ' into $\mathcal{O}$. We then obtain ζ by attaching the included arc to p_2 such that, together with the portion of the boundary between its endpoints which contains p_1, it cuts out an empty disc. This case is illustrated in the third column of Figure 4.20.

Since $\mathcal{O}'$ is a marked surface, $(\mathcal{O}')^*, (T')^* = \mathcal{O}', T'$, so to confirm that the **g**-vectors for the

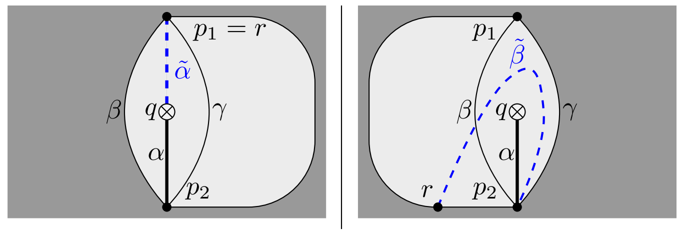

Figure 4.21: Illustrations of the arcs corresponding to elements z_ζ in the map $\chi(L'_{\zeta'}) = L_\zeta z_\zeta$ on elementary quasi-laminations

cluster variables $x_{\zeta'} \in \mathrm{Var}_\bullet(B(T'))$ and $x_\zeta \in \mathrm{Var}_\bullet(B(T))$ agree one need only check that

$$\mathbf{b}(T', \kappa(\zeta')) = \mathbf{b}(T^*, \kappa(\zeta^*)) \text{ for each tagged arc } \zeta' \text{ in } \mathcal{O}'.$$

In particular, one need only explicitly check that the shear coordinate vectors agree at the coordinates corresponding to α and whichever of β or γ is not a boundary component: this is easily done using Figures 4.18-4.20.

Next, we describe the map χ (i.e., $\nu_{\mathbf{z}}$), on $\{\mathbf{y}'\} = \{L'_{\zeta'} : \zeta' \in T'\}$ by describing the elements z_ζ, referring to Figure 4.21. We again have several cases: in each case, there is a unique entry of $B = B(T)$ which differs from the corresponding entry in $B' = B(T')$, so $z_\zeta = 1$ for all but one unique arc $\zeta \in T'$. First, suppose β is a boundary segment. If q is of weight 2, then $b_{\alpha\gamma} = -2$ while $b'_{\alpha'\gamma} = -1$. Thus $z_\gamma = x_{\tilde{\alpha}}$. If q is of weight $1/2$, then $b_{\gamma\alpha} = 2$ while $b'_{\gamma\alpha'} = 1$, so $z_\alpha = x_\gamma$. Now, suppose γ is a boundary segment. If q is of weight 2, then $b_{\alpha\beta} = 2$ while $b'_{\alpha'\beta} = 1$, so $z_\beta = x_\alpha$. If q is of weight $1/2$, then $b_{\beta\alpha} = -2$ while $b'_{\beta\alpha'} = -1$. Thus $z_\alpha = x_{\tilde{\beta}}$.

The final component of the proof, necessary to invoke [24, Propositions 5.2] and thereby conclude that $\nu_{\mathbf{z}}$ is a ring homomorphism, is to show that χ takes every exchange relation of $\mathcal{A}_\bullet(B(T'))$ to a valid relation in $\mathcal{A}_\bullet(B(T))$. Note that in the surfaces case of [24, Proposition 5.29], χ takes every exchange relation to an *exchange* relation: we will see that this does not hold for the orbifolds considered here. To begin, we describe the exchange relations (2.1) in $\mathcal{A}_\bullet(B(T'))$ pictorially by identifying cluster variables with tagged arcs and identifying coefficients with elementary quasi-laminations. For example, if two ordinary tagged arcs ϵ' and δ' on $\mathcal{O}'$ are being exchanged, the situation is as illustrated in Figure 4.22, and the exchange relation is as follows, where $\bar{T}'$ is any triangulation of $\mathcal{O}'$ containing the arcs $\delta', \mu'\eta', \nu'$, and θ':

$$x'_{\epsilon'} x'_{\delta'} = x'_{\mu'} x'_{\eta'} \prod_{\zeta' \in T'} (L'_{\zeta'})^{[b_{\epsilon'}(\bar{T}', L'_{\zeta'})]_+} + x_{\nu'} x_{\theta'} \prod_{\zeta' \in T'} (L_{\zeta'})^{[-b_{\epsilon'}(\bar{T}', L'_{\zeta'})]_+} \tag{4.2}$$

If neither of the tagged arcs in $\mathcal{O}'$ being exchanged intersects the interior of the quadrilateral $p_1 p_2 p_\alpha t_2$, then none of the arcs involved in the exchange relation intersect it either. Thus the

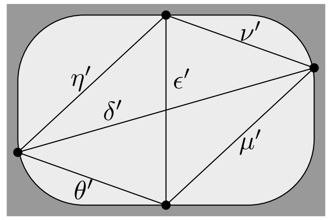

Figure 4.22: Arcs in exchange relations of ordinary arcs in ordinary triangles

corresponding arcs on $\mathcal{O}$ are all obtained, under χ, by the natural inclusion taking the closure of $\mathcal{O}'$ minus the quadrilateral $p_1p_2p_\alpha t_2$ into $\mathcal{O}$. Furthermore, neither $L'_{\alpha'}$ nor whichever of L'_β or L'_γ is defined is involved in the exchange relation, and for all tagged arcs $\zeta' \in T' \setminus \{\alpha', \beta, \gamma\}$, we have $\chi(L_{\zeta'}) = L_\zeta$ and $b_{\epsilon'}(\bar{T}', L'_{\zeta'}) = b_\epsilon(\bar{T}, L_\zeta)$ (where $\bar{T} = \chi(\bar{T}')$). Thus, mapping the exchange relation (4.2) to $\mathcal{O}$ under χ, we obtain precisely the exchange relation that exchanges the corresponding arcs in $\mathcal{O}$. This treatment is recast pictorially in Figure 4.23: the first row shows the exchange relation in $\mathcal{O}'$, and the second row shows the image of the relation in $\mathcal{O}$ under χ, itself an exchange relation since exchange relations among ordinary arcs in ordinary triangles on orbifolds are the same as for surfaces [5, Section 5]. Possibly one of the endpoints of ϵ' coincides with p_1 (or p_2, if γ is a boundary segment) and likewise possibly one of the endpoints of δ' coincides with t_2 (resp. p_1 if γ is a boundary segment). The involved elementary quasi-laminations $L'_{\zeta'}$ and L_ζ are omitted from the illustration (recall from above that $\zeta \notin \{\alpha, \beta, \gamma\}$): since $\zeta = \zeta'$, the quasi-laminations $L'_{\zeta'}$ and L_ζ look nearly identical, with a slight adjustment if an endpoint of ζ' is on the boundary segment t_2p_α if β is a boundary segment, or p_2p_α if γ is a boundary segment. Figure 4.23 establishes a convention used for the remainder of the proof: we denote the tagged arcs in $\mathcal{O}'$ being exchanged by ϵ' and δ', and give the 'positive' summand with respect to ϵ' first, and the 'negative' summand second to match (2.1) and (4.2).

Before considering the remaining exchange relations on $\mathcal{O}'$, we take a moment to illustrate various relations among tagged arcs on $\mathcal{O}$ involving non-ordinary arcs. Figure 4.24 illustrates the coefficient-free exchange relations among tagged arcs on $\mathcal{O}$ which involve non-ordinary arcs. One then adds in whatever coefficients (indexed by elementary quasi-laminations) are necessary to make the relations homogenous in the sense of **g**-vectors. For proofs of these relations, see [5, Section 5]. Observe that two copies of a given tagged arc may arise in such an exchange relation: for example, the relation illustrated in the first row of Figure 4.24 is $x_\delta x_\gamma = x_\mu^2 + x_\nu^2$. As mentioned earlier, not all exchange relations on $\mathcal{O}'$ map to exchange relations on $\mathcal{O}$ under χ. However, we will show that each exchange relation maps to a valid relation on $\mathcal{O}$, and therefore to a valid relation in $\mathcal{A}_\bullet(B(T))$. Figure 4.25 illustrates the coefficient-free valid, non-exchange relations among tagged arcs on $\mathcal{O}$ which we will encounter. (Again, one then adds in whatever coefficients are necessary to make the relations homogenous in the sense of **g**-vectors.) For proofs

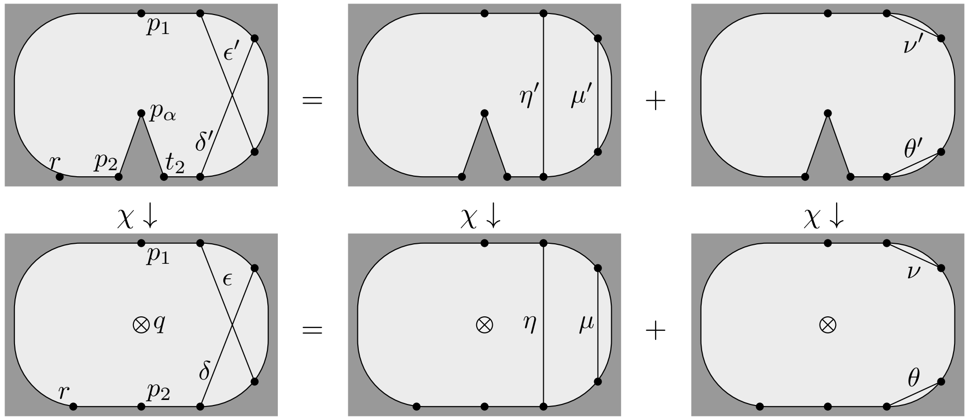

Figure 4.23: The action of χ on exchange relations away from quadrilateral $p_1p_2p_\alpha t_2$: reading from left to right, the first row is an exchange relation in $\mathcal{O}'$, and the second row its image in $\mathcal{O}$ under χ, a convention continued throughout the section

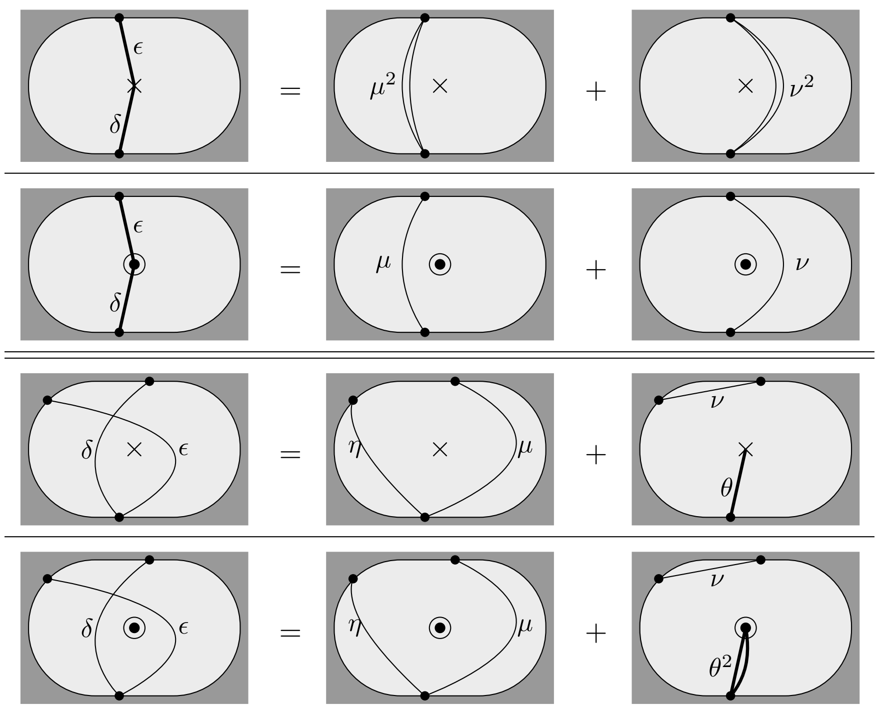

Figure 4.24: Exchange relations on $\mathcal{O}$ involving non-ordinary arcs

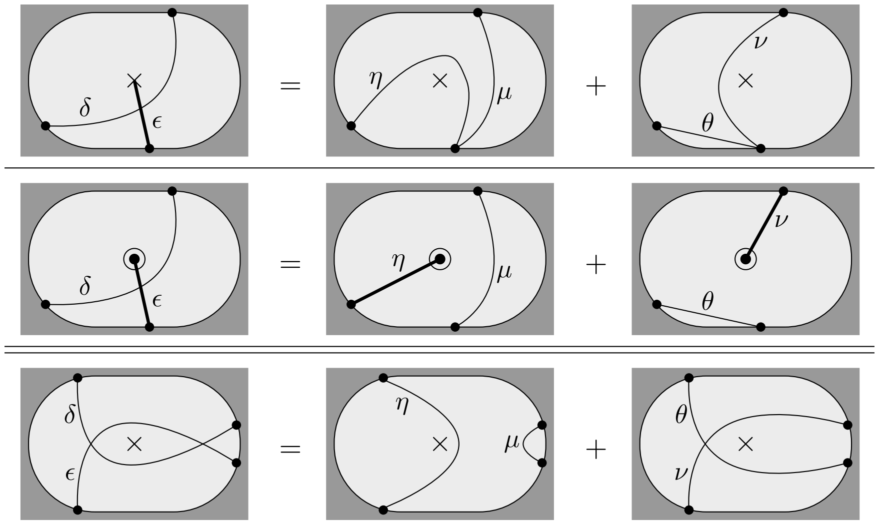

Figure 4.25: Valid non-exchange relations on $\mathcal{O}$ involving non-ordinary arcs

of the relations when q is of weight $1/2$, see [7, Section 5.2]. The relation when q is of weight 2 illustrated in the second row of the figure is easily verified as the image, under a single local folding, of an exchange relation on a once-punctured disc.

We now turn to showing that the images of exchange relations are valid relations by continuing the pictorial treatment of Figure 4.23, and consider several cases. We have already handled the exchange of any two tagged arcs in $\mathcal{O}'$ which do not intersect the interior of the quadrilateral $p_1p_2p_\alpha t_2$; it remains to consider arcs which do. In particular, we consider in turn exchange relations involving α', exchange relations involving arcs other than α' which have an endpoint at p_2 if β is a boundary segment or t_2 is γ is a boundary segment, and finally exchange relations involving arcs with an endpoint at p_α. In our illustrations, we omit all elementary quasi-laminations other than those corresponding to the arcs α, α', and whichever of β or γ is not a boundary segment. As discussed above, for all other $\zeta' \in T'$, we have $\chi(L_{\zeta'}) = L_\zeta$, and it is easily checked that $\mathbf{b}(\bar{T}', L'_{\zeta'}) = \mathbf{b}(\bar{T}, L_\zeta)$ for any tagged triangulation $\bar{T}'$ of $\mathcal{O}$ and its image $\bar{T}$ under χ. We also generally handle separately each of the four combinations of weight 2 or $1/2$ on orbifold point q and of β or γ being a boundary segment, although at times we are able to group some together, or at least depict them in the same image.

Suppose one of the tagged arcs in $\mathcal{O}'$ being exchanged is α'. Then the other arc, which we will call δ', must have one endpoint at p_α and the other at some distinct marked point other than p_2 or t_2. First, suppose the other endpoint of δ' is r: the exchange relation on $\mathcal{O}'$ (regardless of the weight of q or which of β or γ is a boundary segment) is depicted in the first row of Figure 4.26. (Recall that $r = p_1$ when β is a boundary segment.) The second row of the figure

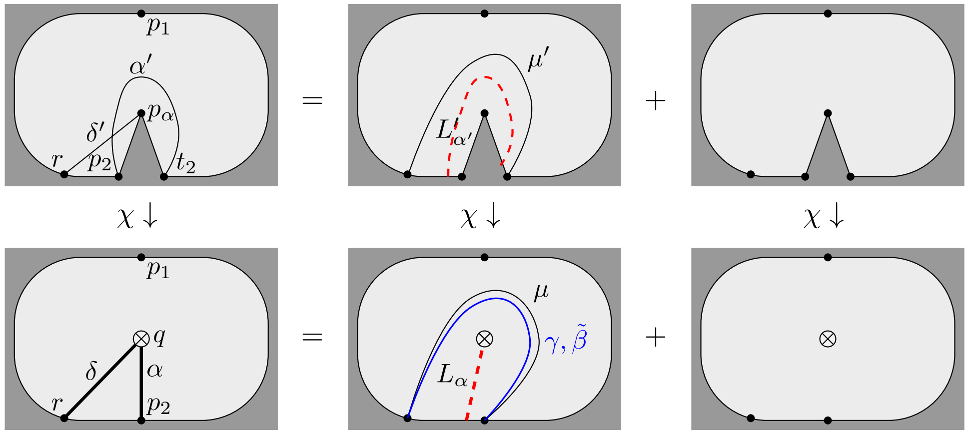

Figure 4.26: The action of χ on exchange relations involving α'
.

depicts the image of the relation under χ: the blue arc, labeled both γ and $\tilde{\beta}$, only appears if orbifold point q is of weight $1/2$, as in this case $z_\alpha = x_\gamma$ if β is a boundary segment and $z_\alpha = x_{\tilde{\beta}}$ if γ is a boundary segment. (Both γ and $\tilde{\beta}$ coincide with μ.) Thus we have an exchange relation on $\mathcal{O}$, of the type displayed in the first row of Figure 4.24 if $q \in \mathbf{Q_{1/2}}$, and in the second row if $q \in \mathbf{Q_2}$. Now, suppose that the other endpoint of δ' is *not* r. The exchange relation on $\mathcal{O}'$ when β is a boundary segment is depicted in the first row of Figure 4.27. The second and third rows of the figure depict the image of the relation under χ when $q \in \mathbf{Q_{1/2}}$ and $q \in \mathbf{Q_2}$, respectively, and can be verified as valid non-exchange relations on $\mathcal{O}$ of the type illustrated in the first and second rows, respectively, of Figure 4.25. The exchange relation on $\mathcal{O}'$ when γ is a boundary segment is depicted in the first row of Figure 4.28. The second and third rows of the figure depict the image of the relation under χ when $q \in \mathbf{Q_{1/2}}$ and $q \in \mathbf{Q_2}$, respectively. The former is a valid non-exchange relation on $\mathcal{O}$ of the type illustrated in the first row of Figure 4.25, and the latter is an exchange relation on $\mathcal{O}$ of the type illustrated in the second row of Figure 4.24.

Next, suppose one of the tagged arcs in $\mathcal{O}'$ being exchanged, $\epsilon' \neq \alpha'$, has an endpoint at p_2 if β is a boundary segment or at t_2 if γ is a boundary segment. Further suppose that the other arc, which we will again call δ', has one endpoint at p_α. The exchange relation on $\mathcal{O}'$ when β is a boundary segment and the other endpoint of δ' is r is depicted in the first row of Figure 4.29. The second and third rows of the figure depict the image of the relation under χ when $q \in \mathbf{Q_{1/2}}$ and $q \in \mathbf{Q_2}$, respectively. The former is a valid non-exchange relation on $\mathcal{O}$ of the type illustrated in the first row of Figure 4.25, and the latter is an exchange relation of the type illustrated in the second row of Figure 4.24. The exchange relation on $\mathcal{O}'$ when β is a boundary segment and the other endpoint of δ' is *not* r is depicted in the fourth row of Figure 4.29. The fifth and sixth rows of the figure depict the image of the relation under χ

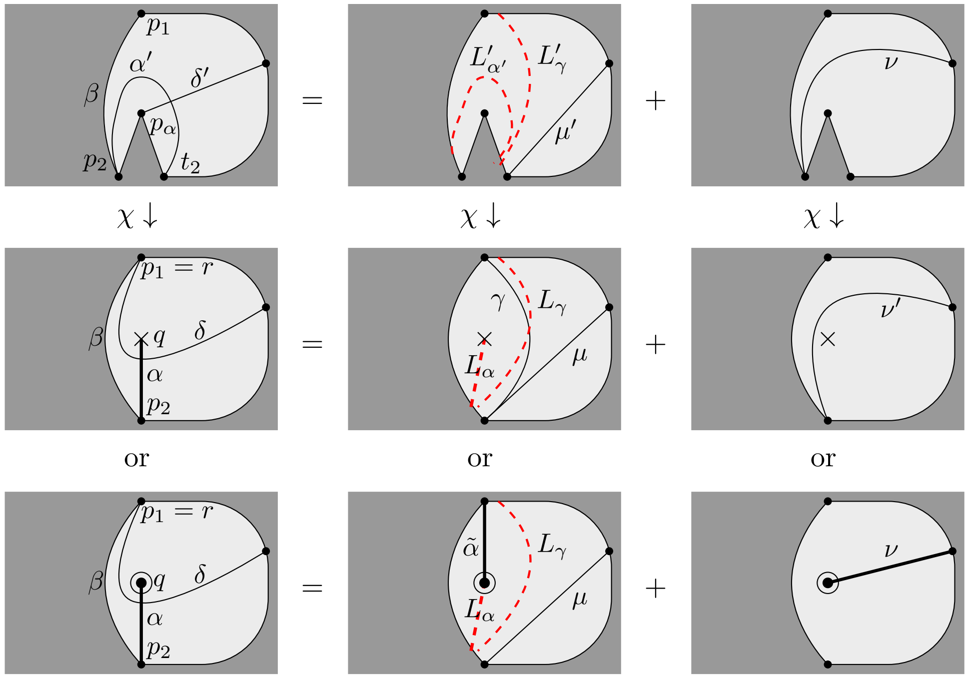

Figure 4.27: The action of χ on exchange relations involving α', contd., when β is a boundary segment

.

when $q \in \mathbf{Q_{1/2}}$ and $q \in \mathbf{Q_2}$, respectively. They are valid non-exchange relations on $\mathcal{O}$ of the type illustrated in the first and second rows of Figure 4.25, respectively. The exchange relation on $\mathcal{O}'$ when γ is a boundary segment is depicted in the first row of Figure 4.30. (In this case, the other endpoint of δ' cannot be at r, as otherwise it would be compatible with $\epsilon' \neq \alpha'$.) The second and third rows of the figure depict the image of the relation under χ when $q \in \mathbf{Q_{1/2}}$ and $q \in \mathbf{Q_2}$, respectively. The former is an ordinary exchange relation on $\mathcal{O}$; the latter is a valid non-exchange relation of the type shown in the second row of Figure 4.25.

We continue with the case when one of the tagged arcs in $\mathcal{O}'$ being exchanged, $\epsilon' \neq \alpha'$, has an endpoint at p_2 if β is a boundary segment or at t_2 if γ is a boundary segment. Now suppose that the other arc, δ', does not have an endpoint at p_α. (Possibly δ' has an endpoint at r; this does not affect the argument below since neither arc has an endpoint at p_α.) The exchange relation on $\mathcal{O}'$ when β is a boundary segment and δ' has an endpoint at t_2 is depicted in the first row of Figure 4.31. The second and third rows of the figure depict the image of the relation under χ when $q \in \mathbf{Q_{1/2}}$ and $q \in \mathbf{Q_2}$, respectively. The former is an exchange relation on $\mathcal{O}$ of the type illustrated in the third row of Figure 4.24; the latter is a valid non-exchange relation of the type illustrated in the second row of Figure 4.25. The exchange relation on $\mathcal{O}'$ when β is a boundary segment and δ' does *not* have an endpoint at t_2 is depicted in the fourth row

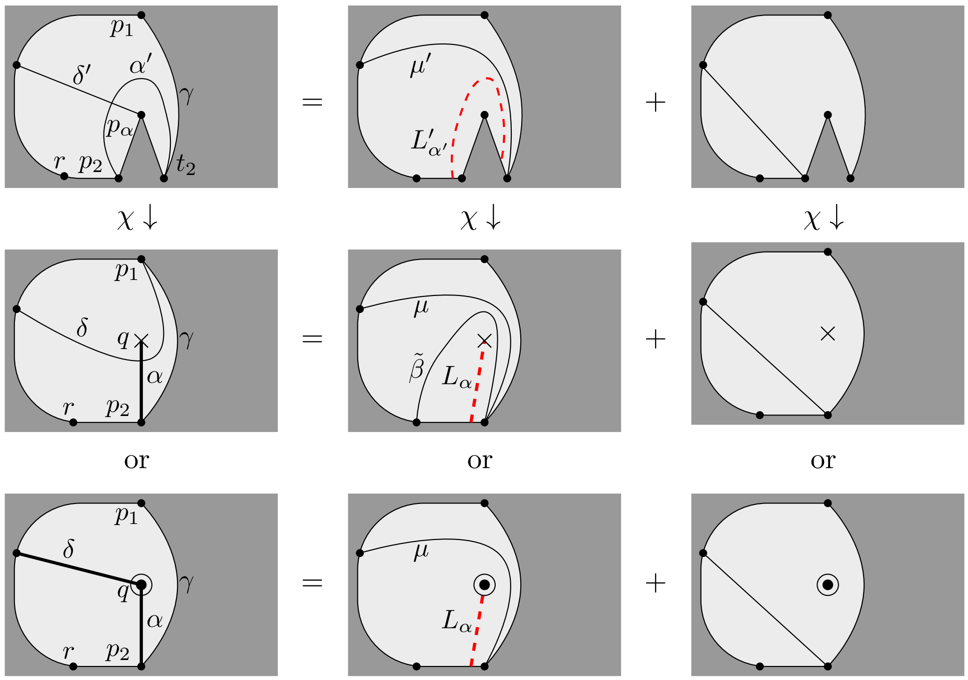

Figure 4.28: The action of χ on exchange relations involving α', contd., when γ is a boundary segment

.

of Figure 4.31. The fifth and sixth rows of the figure depict the image of the relation under χ when $q \in \mathbf{Q_{1/2}}$ and $q \in \mathbf{Q_2}$, respectively. The former is an ordinary exchange relation on $\mathcal{O}$ as depicted in Figure 4.22; the latter is a valid non-exchange relation of the type illustrated in the second row of Figure 4.25. The exchange relation on $\mathcal{O}'$ when γ is a boundary segment and δ' has an endpoint at p_2 is depicted in the first row of Figure 4.32. The second and third rows of the figure depicts the image of the relation under χ when $q \in \mathbf{Q_{1/2}}$ and $q \in \mathbf{Q_2}$, respectively, and are exchange relations on $\mathcal{O}$ of the respective type illustrated in the third and fourth rows of Figure 4.24. The exchange relation on $\mathcal{O}'$ when γ is a boundary segment and δ' does *not* have an endpoint at p_2 is depicted in the fourth row of Figure 4.32. The fifth row of the figure depict the image of the relation under χ regardless of the weight of q, an ordinary exchange relation on $\mathcal{O}$.

We consider our final set of cases, when one of the tagged arcs in $\mathcal{O}'$ being exchanged, ϵ', has an endpoint at p_α. Once more, denote the other arc being exchanged by δ': we may assume δ' does not have an endpoint at p_2 if β is a boundary segment or at t_2 if γ is a boundary segment, as these possibilities have already been handled. Thus we may also assume that ϵ' does not have an endpoint at r. Suppose first that β is a boundary segment. Then we may further assume that the other endpoint of ϵ' is not at r, as there are no remaining exchangeable arcs δ' to consider.

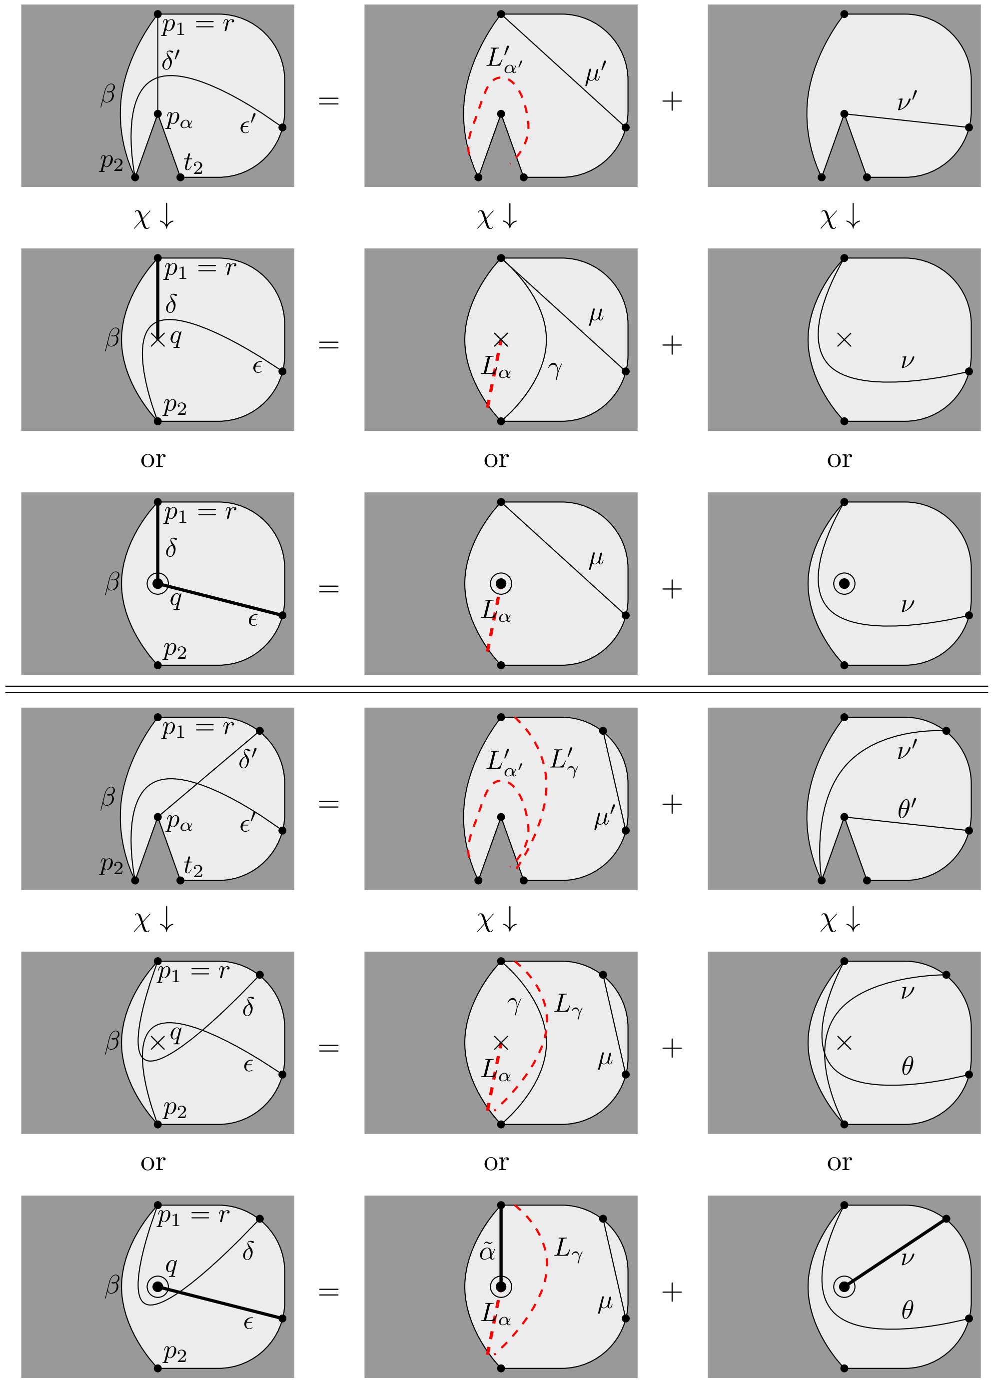

Figure 4.29: The action of χ on exchange relations involving one arc with an endpoint at p_2 and another with an endpoint at p_α when β is a boundary segment

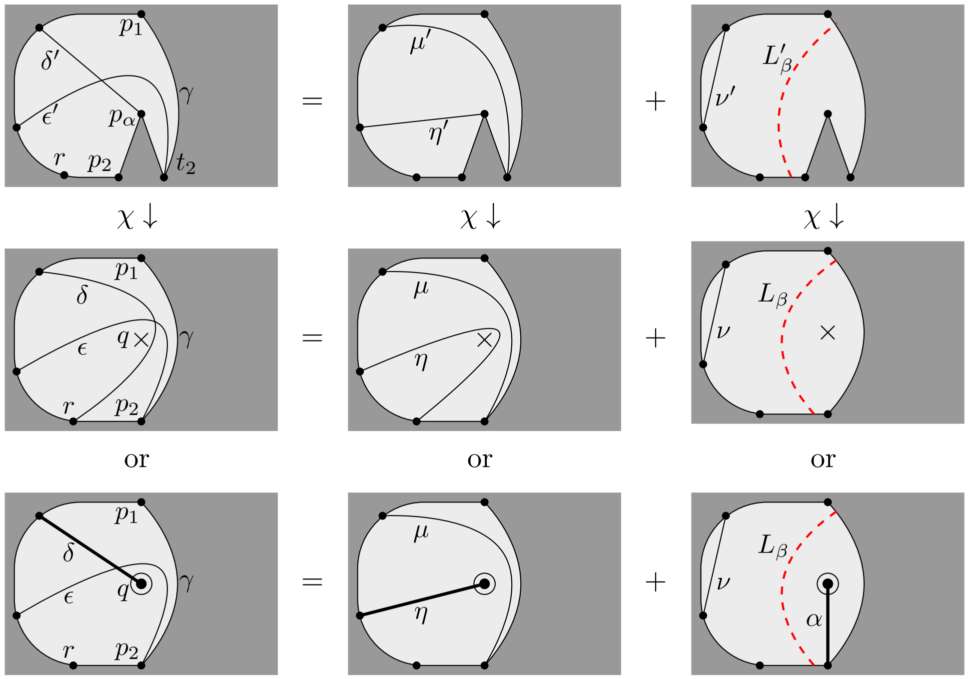

Figure 4.30: The action of χ on exchange relations involving one arc with an endpoint at t_2 and another with an endpoint at p_α when γ is a boundary segment

The exchange relation on $\mathcal{O}'$ when β is a boundary segment and δ' has an endpoint at $p_1 = r$ is depicted in the first row of Figure 4.33. The second row of the figure depicts the image of the relation under χ: the blue arc $\tilde{\alpha}$ only appears if q is of weight 2. This is an exchange relation on $\mathcal{O}$ of the type depicted in the third row of Figure 4.24 if $q \in \mathbf{Q_{1/2}}$ and the fourth row if $q \in \mathbf{Q_2}$. The exchange relation on $\mathcal{O}'$ when β is a boundary segment and δ' does *not* have an endpoint at $p_1 = r$ is depicted in the third row of Figure 4.33. The fourth row of the figure depicts the image of the relation under χ regardless of the weight of q, an ordinary exchange relation in $\mathcal{O}$. Finally, suppose that γ is a boundary segment. The exchange relation on $\mathcal{O}'$ when the other endpoint of ϵ' is at r is depicted in the first row of Figure 4.34. The second and third rows of the figure depict the image of the relation under χ when $q \in \mathbf{Q_{1/2}}$ and $q \in \mathbf{Q_2}$, respectively. The former is an exchange relation of the type illustrated in the third row of Figure 4.24; the latter is a valid non-exchange relations of the type illustrated in the second row of Figure 4.25. The exchange relation on $\mathcal{O}'$ when the other endpoint of ϵ' is *not* at r is depicted in the fourth row of Figure 4.34. The fifth and sixth rows of the figure depict the image of the relation under χ when $q \in \mathbf{Q_{1/2}}$ and $q \in \mathbf{Q_2}$, respectively. The former is an ordinary exchange relation on $\mathcal{O}$, and the latter is a valid non-exchange relation of the type illustrated in the second row of Figure 4.25. □

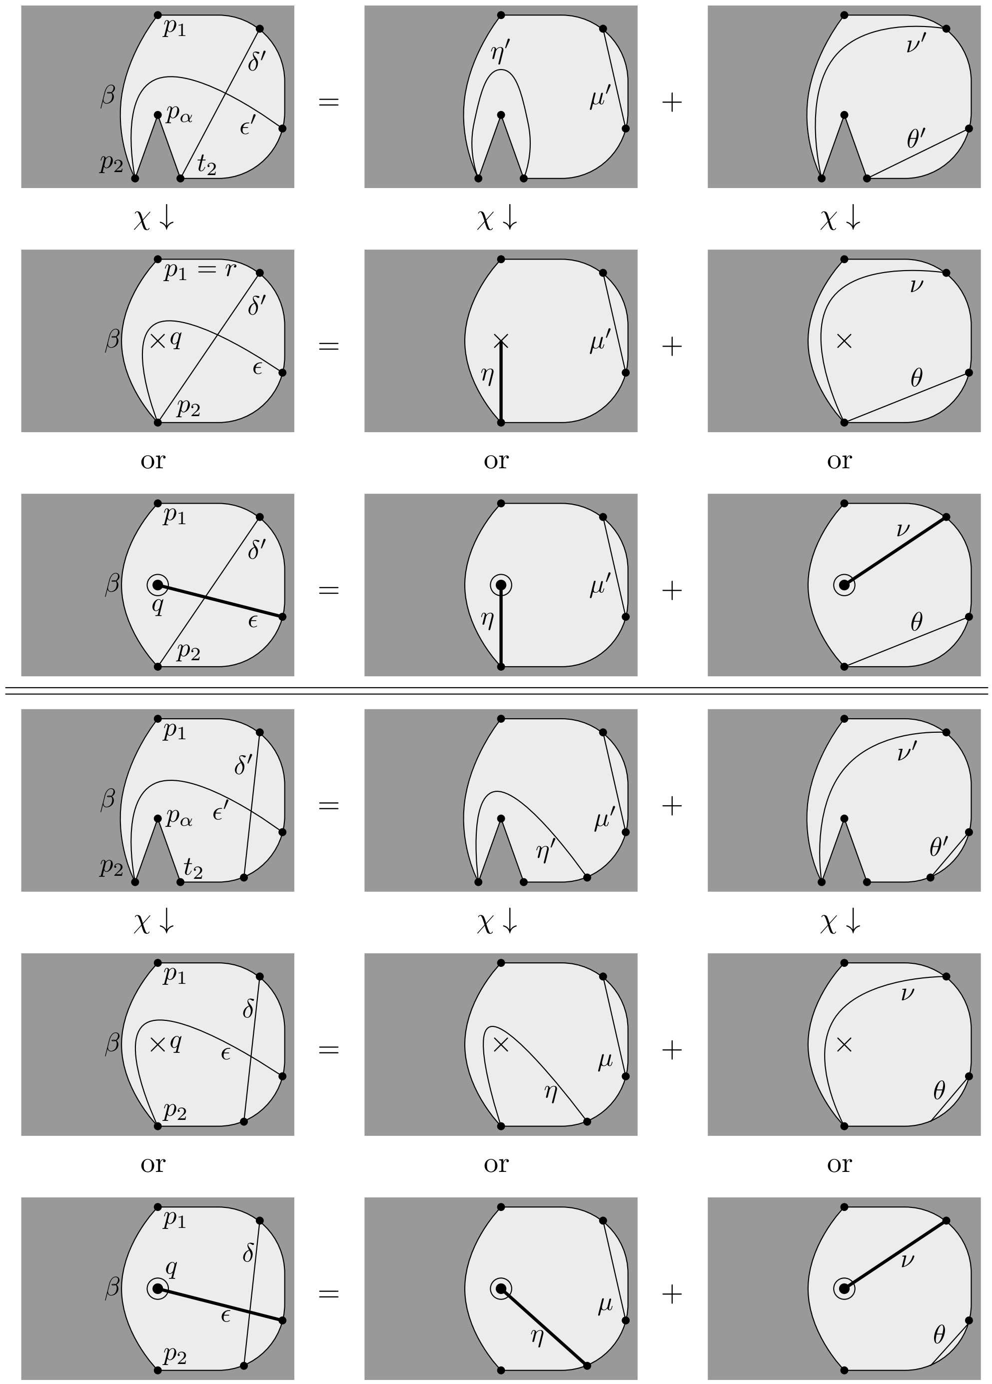

Figure 4.31: The action of χ on exchange relations involving one arc with an endpoint at p_2, contd., when β is a boundary segment

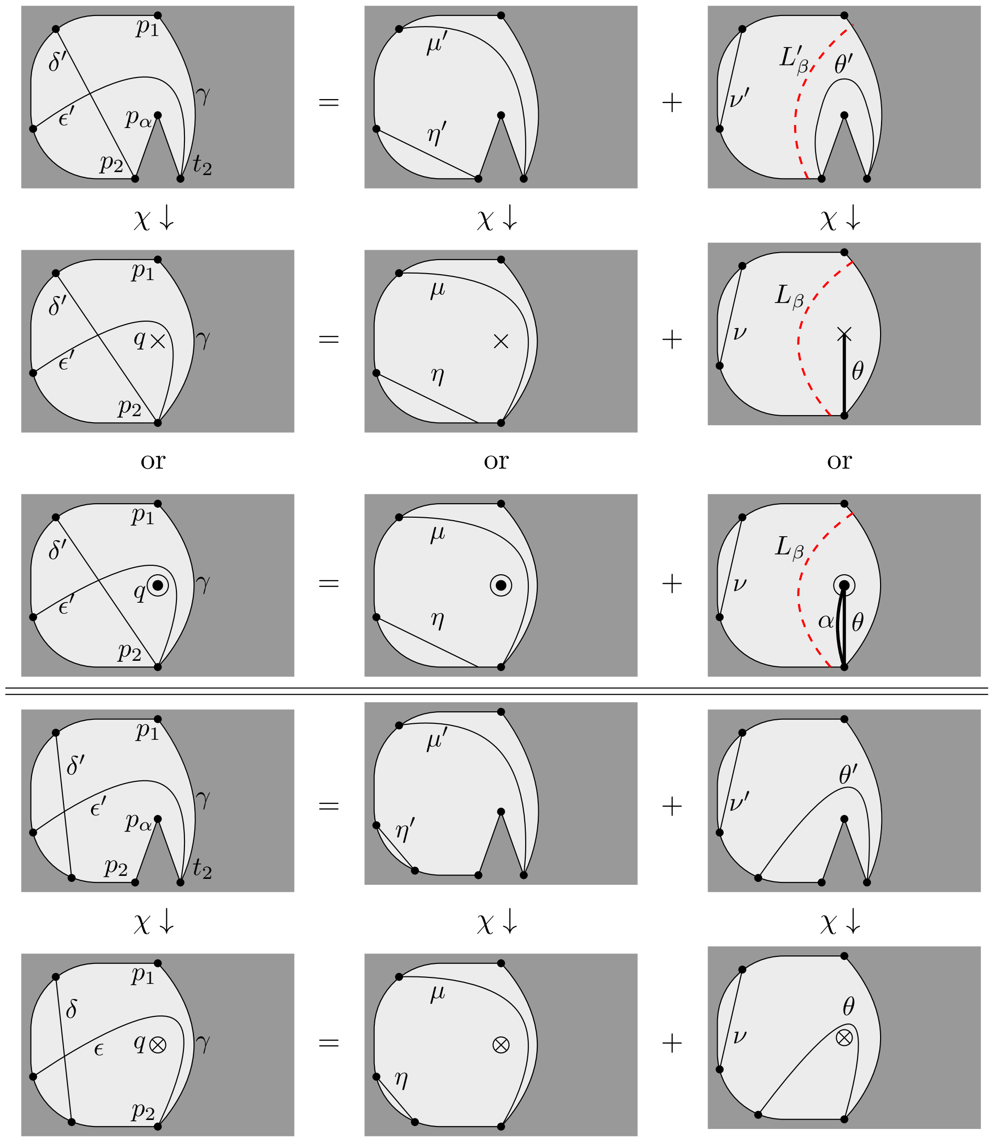

Figure 4.32: The action of χ on exchange relations involving one arc with an endpoint at t_2, contd., when γ is a boundary segment

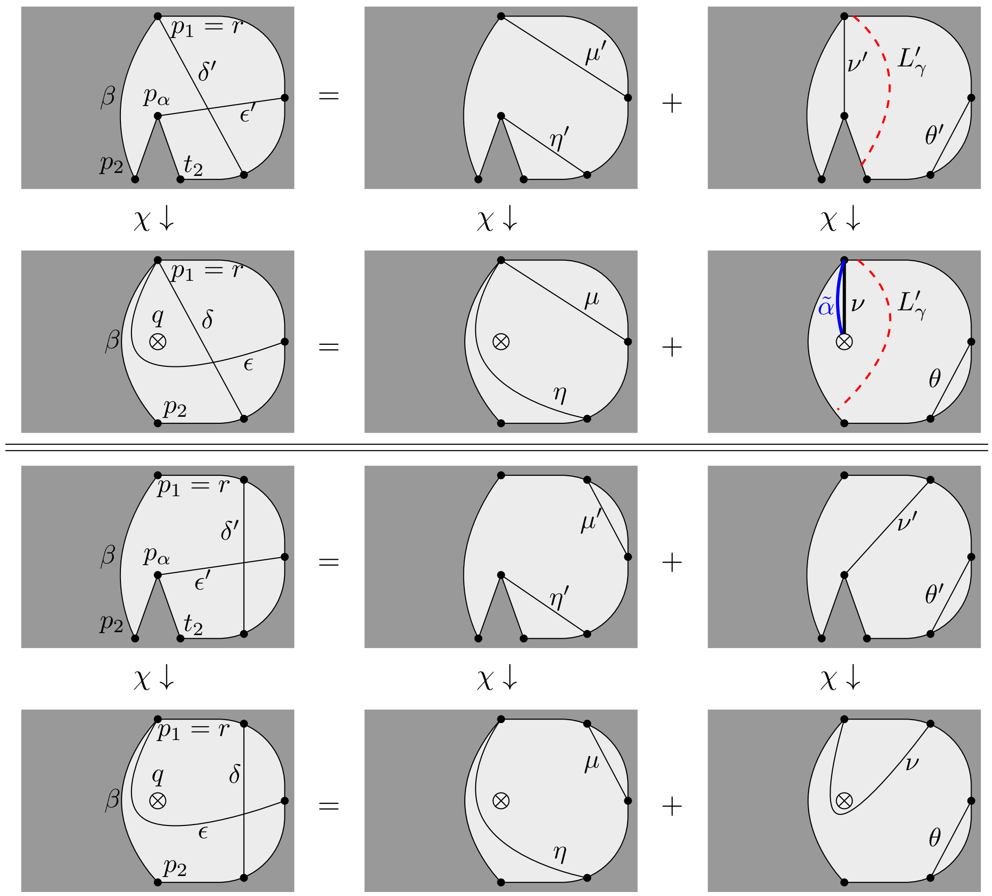

Figure 4.33: The action of χ on exchange relations involving one arc with an endpoint at p_α when β is a boundary segment

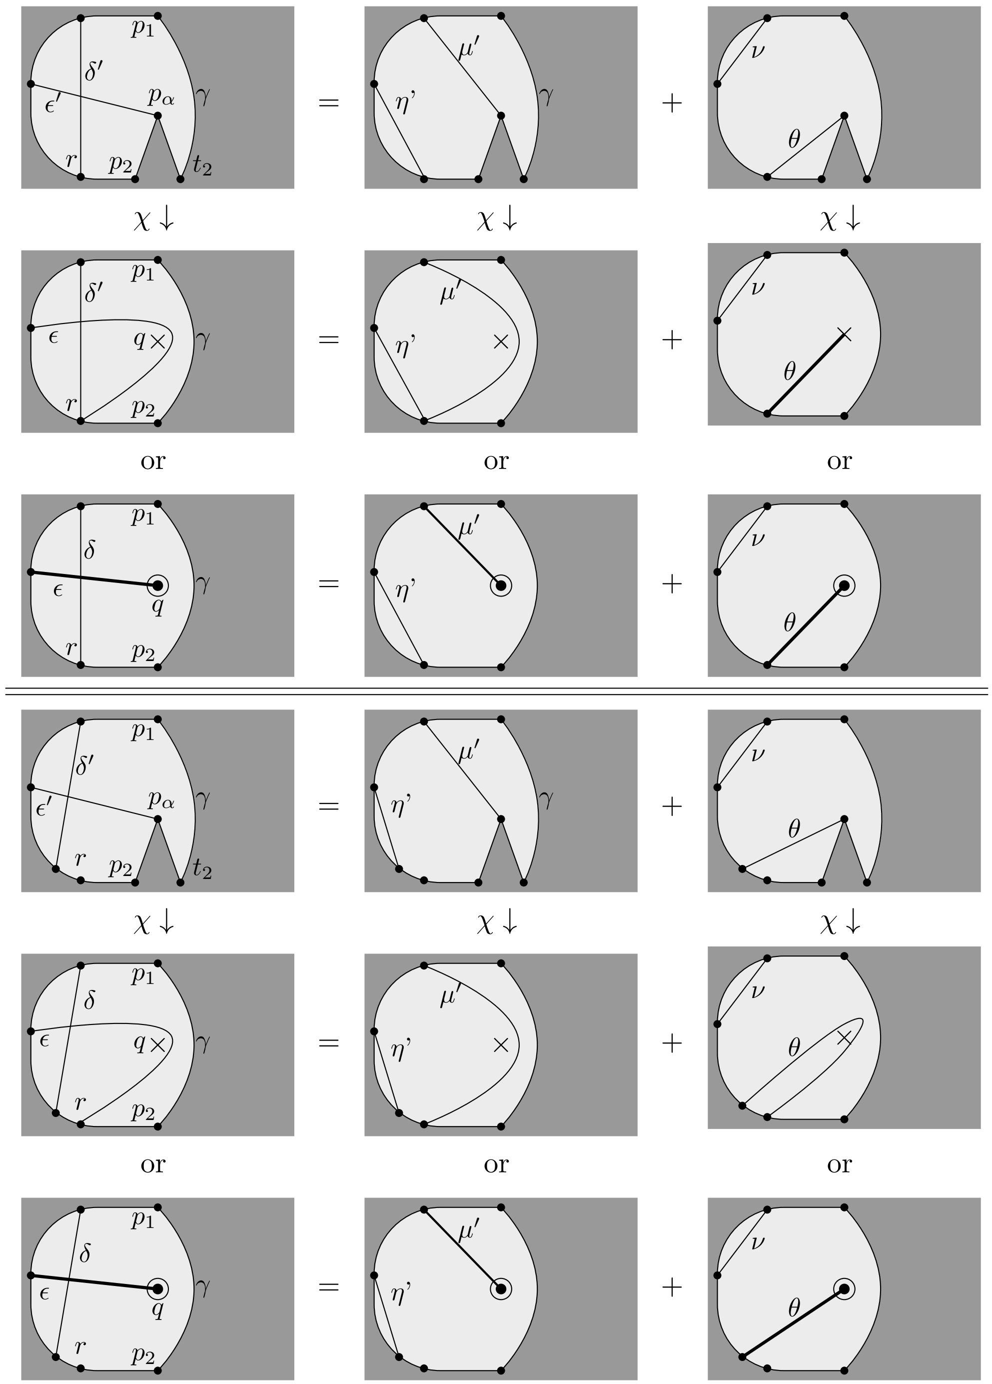

Figure 4.34: The action of χ on exchange relations involving one arc with an endpoint at p_α when γ is a boundary segment

REFERENCES

[1] Peigen Cao, Min Huang, and Fang Li. Categorification of sign-skew-symmetric cluster algebras and some conjectures on g-vectors, v2. *ArXiv e-prints*, April 2017.

[2] Laurent Demonet. Catégorificiation d'algèbres amassées antisymétrisables. *PhD thesis, Université de Caen.*

[3] Laurent Demonet. Algèbres amassées et algèbres préprojectives: le cas non simplement lacé. *C. R. Math. Acad. Sci. Paris*, 346(7-8):379–384, 2008.

[4] G. Dupont. An approach to non-simply laced cluster algebras. *J. Algebra*, 320(4):1626–1661, 2008.

[5] Anna Felikson, Michael Shapiro, and Pavel Tumarkin. Cluster algebras and triangulated orbifolds. *Adv. Math.*, 231(5):2953–3002, 2012.

[6] Anna Felikson, Michael Shapiro, and Pavel Tumarkin. Cluster algebras of finite mutation type via unfoldings. *Int. Math. Res. Not. IMRN*, (8):1768–1804, 2012.

[7] Anna Felikson and Pavel Tumarkin. Bases for cluster algebras from orbifolds. *Adv. Math.*, 318:191–232, 2017.

[8] Sergey Fomin, Michael Shapiro, and Dylan Thurston. Cluster algebras and triangulated surfaces. I. Cluster complexes. *Acta Math.*, 201(1):83–146, 2008.

[9] Sergey Fomin and Dylan Thurston. Cluster algebras and triangulated surfaces. Part II: Lambda lengths. *ArXiv e-prints*, October 2012.

[10] Sergey Fomin and Andrei Zelevinsky. Cluster algebras. I. Foundations. *J. Amer. Math. Soc.*, 15(2):497–529, 2002.

[11] Sergey Fomin and Andrei Zelevinsky. Cluster algebras. II. Finite type classification. *Invent. Math.*, 154(1):63–121, 2003.

[12] Sergey Fomin and Andrei Zelevinsky. Y-systems and generalized associahedra. *Ann. of Math. (2)*, 158(3):977–1018, 2003.

[13] Sergey Fomin and Andrei Zelevinsky. Cluster algebras. IV. Coefficients. *Compos. Math.*, 143(1):112–164, 2007.

[14] M. Gross, P. Hacking, S. Keel, and M. Kontsevich. Canonical bases for cluster algebras, v2. *ArXiv e-prints*, October 2016.

[15] Min Huang and Fang Li. Unfolding of acyclic sign-skew-symmetric cluster algebras and applications to positivity and F-polynomials, v2. *ArXiv e-prints*, November 2016.

[16] James E. Humphreys. *Introduction to Lie algebras and representation theory.* Springer-Verlag, New York-Berlin, 1972. Graduate Texts in Mathematics, Vol. 9.

[17] Victor G. Kac. *Infinite-dimensional Lie algebras*, volume 44 of *Progress in Mathematics.*

[18] Tomoki Nakanishi and Salvatore Stella. Diagrammatic description of c-vectors and d-vectors of cluster algebras of finite type. *Electron. J. Combin.*, 21(1):Paper 1.3, 107, 2014.

[19] Tomoki Nakanishi and Andrei Zelevinsky. On tropical dualities in cluster algebras. In *Algebraic groups and quantum groups*, volume 565 of *Contemp. Math.*, pages 217–226. Amer. Math. Soc., Providence, RI, 2012.

[20] Nathan Reading. Universal geometric coefficients for the once-punctured torus. *Sém. Lothar. Combin.*, 71:Art. B71e, 29, 2013/14.

[21] Nathan Reading. Universal geometric cluster algebras. *Math. Z.*, 277(1-2):499–547, 2014.

[22] Nathan Reading. Universal geometric cluster algebras from surfaces. *Trans. Amer. Math. Soc.*, 366(12):6647–6685, 2014.

[23] Nathan Reading. Scattering diagrams and scattering fans, v1. *ArXiv e-prints*, December 2017.

[24] Nathan Reading. Dominance phenomena: mutation, scattering and cluster algebras, v1. *ArXiv e-prints*, February 2018.

[25] Nathan Reading and David E. Speyer. Cambrian fans. *J. Eur. Math. Soc. (JEMS)*, 11(2):407–447, 2009.

[26] R. Tyrrell Rockafellar. *Convex analysis*. Princeton Mathematical Series, No. 28. Princeton University Press, Princeton, N.J., 1970.

[27] John Stembridge. Folding by automorphisms. `http://www.math.lsa.umich.edu/~jrs/papers/folding.pdf`, August 2008.

[28] Dylan Thurston. The geometry and topology of three-manifolds. chapter 13. `http://library.msri.org/books/gt3m/`, March 2002.