Legendrian Loop Actions on the Lagrangian Concordance Monoid

By

James Hughes

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

MATHEMATICS

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Roger Casals

Eugene Gorsky

Andrew Waldron

Committee in Charge

2023

Contents

Abstract iii

Acknowledgments iv

Chapter 1. Introduction 1

1.1. Context and main results 2

1.2. Background 9

Chapter 2. Lagrangian fillings in A-type and their Kálmán loop orbits 44

2.1. Isotopies of exact Lagrangian Cobordisms and Kálmán loop orbits of $\lambda(A_n)$ 44

2.2. Proof of Theorems 1.1.1 and 1.1.3 44

2.3. Algebraic Proof of Theorem 1.1.3 54

2.4. Combinatorial Characterizations 63

Chapter 3. Weave realizability for D-type 75

3.1. Weave Realizability for $\lambda(D_n)$ 76

Chapter 4. Legendrian loops and mapping class groups 98

4.1. Legendrian loops as generators of cluster modular groups 98

4.2. Nielsen-Thurston classification of Legendrian loops 111

Appendix A. Mutation sequence computation via Legendrian weaves 119

Bibliography 128

Legendrian Loop Actions on the Lagrangian Concordance Monoid

Abstract

In this work we investigate the conjectural ADE classification of exact Lagrangian fillings of Legendrian links.We begin by showing that two methods of constructing exact Lagrangian fillings – Legendrian weaves and decomposable exact Lagrangian cobordisms without Reidemeister I or II moves – yield Hamiltonian isotopic exact Lagrangian fillings. Using the method of Legendrian weaves, we construct and distinguish exact Lagrangian fillings in D_n-type. We then investigate Legendrian loops, Legendrian isotopies fixing a Legendrian link pointwise at time one. Legendrian loops act on the set of exact Lagrangian fillings by concatenating the trace of the Legendrian isotopy. We investigate this action first in type A_n, and then more generally, leveraging techniques from the theory of cluster algebras and connections to the theory of mapping class groups. In particular, we give a complete description of the orbital structure of the cluster modular group action on exact Lagrangian fillings of Legendrian $(2, n)$ torus links.

Beyond type A_n, we compile and extend known results interpreting Legendrian loops as generators of cluster modular groups for affine and extended affine type cluster algebras. We show that Legendrian loops virtually generate these cluster modular groups. By extending an analogy between cluster modular groups to mapping class groups, we provide new tests for detecting when a Legendrian loop produces infinitely many distinct exact Lagrangian fillings. Finally, we discuss possible avenues towards producing generating sets for cluster modular groups using Legendrian loops.

Acknowledgments

First of all, a big thank you to Roger Casals for teaching me everything I know about contact topology and for all of his advice and guidance throughout the last four years. Several parts of this dissertation would most likely not have existed without his support and encouragement. Thanks also to Youngjin Bae, Chris Fraser, Lenhard Ng, Linhui Shen, and José Simental for helpful conversations. In particular, thank you to Lenny for the question that originally motivated much of Chapter 2.

I am grateful to the NSF which partially supported this work through grant DMS-1942363 and to UC Davis for a Dean's Summer Graduate Fellowship in the summer of 2022.

I would also like to thank the many people responsible for shaping my experience in grad school: first, Eugene Gorsky, for teaching interesting reading courses nearly every quarter. Almost none of the material from those courses has made its way into this work, but I hope that I am a more well-rounded mathematician because of his efforts. Thanks also to Elysée Wilson-Egolf for all of their efforts towards improving the lives of graduate students here at UC Davis. In addition, thank you to Edgar Jaramillo-Rodriguez, Kayden Mimmack, and Eli Moore for getting me through the first year of grad school and for helping keep me sane throughout. Finally, thanks to Daping Weng and Orsola Capovilla-Searle for being wonderful lab mates and collaborators, and for putting up with all of my questions.

I owe an enormous debt of gratitude to many of the teachers and professors responsible for getting me here. In particular, thank you to Jennifer Taback, for planting the seed of the idea of becoming a mathematician. Her example continues to be an inspiration to me.

Thank you to my family, especially my mom, Patty Hughes, who made sure that I didn't fail my very first algebra class and who was excited that I got to attend a grad school where I could think about doughnuts all day, even if she's still not sure why they're interesting.

Lastly, a special thank you to my wife, Sarah Sears, for her incredible patience and encouragement. I endlessly appreciate her ability to put up with my pacing and my frequent loss of ability to form coherent sentences. Thank you for all of your support, even when geography does its best to intervene.

CHAPTER 1

Introduction

A contact manifold is an odd-dimensional manifold that is equipped with a plane field satisfying certain conditions that generalize physical properties from optics or thermodynamics. In $\mathbb{R}^3$, the **standard contact structure** ξ_{st} is the 2-plane field given as the kernel of the 1-form $\alpha_{st} = dz - ydx$. Symplectic manifolds, the even-dimensional analogues of contact manifolds, arise from generalizations of classical mechanics. A symplectic structure is given by equipping the manifold M with a closed, non-degenerate 2-form. In $\mathbb{R}^4$, the **standard symplectic form** is the 2-form $\omega_{st} = dx_1 \wedge dy_1 + dx_2 \wedge dy_2$.

It is often useful to study manifolds by analyzing their submanifolds. For a contact manifold (M^{2n+1}, ξ), submanifolds that respect the contact structure ξ are either transverse to the contact hyperplanes at every point, or tangent at every point. An n-dimensional submanifold satisfying the latter condition is called **Legendrian**. Similarly, for a submanifold L of a symplectic manifold (M^{2n}, ω), we can have that the restriction $\omega|_L$ either uniformly vanishes or is always positive. We call an n-dimensional submanifold $L^n \subseteq M^{2n}$ **Lagrangian** if the symplectic form ω vanishes when restricted to L.

In this work we consider **Legendrian links** in standard contact $(\mathbb{R}^3, \xi_{st})$, i.e. embeddings of $\mathbb{S}^1$ into $\mathbb{R}^3$ always tangent to ξ_{st}. Any smooth link in $(\mathbb{R}^3, \xi_{st})$, equivalently, $(\mathbb{S}^3, \xi_{st})$, admits infinitely many Legendrian isotopy classes. However, all of the links studied in this work will admit a unique representative that maximizes a classical invariant known as the Thurston-Bennequin number (tb). We restrict ourselves only considering such representatives from this point onward.

One tool in the study of Legendrian links – and symplectic topology more generally – is an understanding of their **exact Lagrangian fillings**. These are orientable Lagrangian surfaces $L \subseteq (\mathbb{D}^4, \omega_{st})$ with boundary equal to a given Legendrian $\lambda \subseteq (\mathbb{S}^3, \xi_{st})$ and satisfying the condition that $\omega|_L = df$ for some function f on L. We consider exact Lagrangian surfaces up to Hamiltonian isotopy, equivalently, exact Lagrangian isotopy. The exactness condition can be thought of

as ensuring that the Lagrangian surface respects the Legendrian condition at the boundary. In the smooth setting, such surfaces can be arbitrarily complicated. However, the exact Lagrangian condition is rigid enough to make a conjectural classification of these objects possible, at least for certain classes of Legendrian links. This conjectural classification and various techniques involved in understanding it comprise the main area of study of this work.

1.1. Context and main results

In this work, all of the specific examples of Legendrian links we consider are Legendrian isotopic to the rainbow closure of a positive braid β. We will specifically note when we have results that yield implications for more general classes of Legendrian links. As pictured in Figure 1.1, the rainbow closure $\lambda(\beta)$ of a positive braid β is Legendrian isotopic to the $(-1)-$closure of $\beta\Delta^2$ where Δ denotes a half twist of the braid. In order to construct exact Lagrangian fillings of these links, we consider two main constructions. The first consists of representing the exact Lagrangian filling as a series of elementary Lagrangian cobordisms involving traces of Legendrian isotopies and a saddle cobordism. The second construction, more recently developed, comes from the method of Legendrian weaves. **Legendrian weaves** are a geometric construction of Casals and Zaslow used to combinatorially encode the information of a Legendrian surface via the singularities of its (front) projection [**CZ21**, Section 2]. Legendrian weaves are crucial tools in proving the existence of additional structure on the set of exact Lagrangian fillings, as we will explain below. As our first result, we show that the two constructions considered in this work agree.

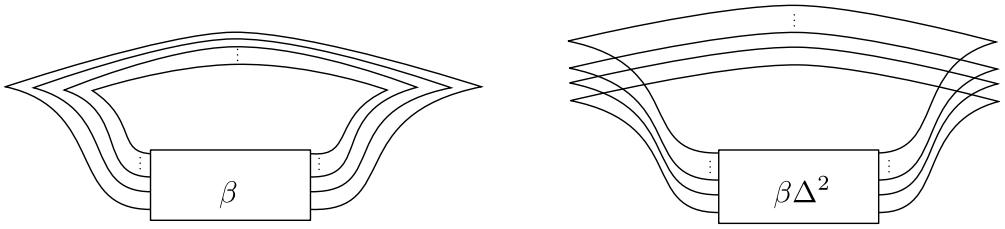

FIGURE 1.1. Front projections of the Legendrian isotopic links given as the rainbow closure (left) and (-1)-framed closure (right) of the positive braids β and $\beta\Delta^2$. Here Δ denotes a half twist of the braid.

THEOREM 1.1.1 (Theorem 1.1, [**Hug21**]). *For any exact Lagrangian filling of $\lambda(\beta)$ constructed via a sequence of pinching cobordisms and traces of Reidemeister III moves, there is a unique Hamiltonian isotopic weave filling up to Hamiltonian isotopy.*

The restriction on the type of elementary cobordisms appearing in Theorem 1.1.1 follows naturally from the Legendrian weave construction and in practice does not exclude any known constructions of exact Lagrangian fillings.

Our next results concern constructions of exact Lagrangian fillings for particular Legendrian links. Let us denote by σ_i the ith Artin generator of the n-stranded braid group Br_n. Among rainbow closures of positive braids, we have several distinguished families, including $\lambda(A_n) = \lambda(\sigma_1^{n+1}), \lambda(D_n) = \sigma_1^{n-2}\sigma_2\sigma_1^2\sigma_2$ and $\lambda(E_6), \lambda(E_7)$, and $\lambda(E_8)$, corresponding to positive braids closures of $\sigma_1\sigma_2^2\sigma_1^2\sigma_2^{i-3}$ for $i = 6, 7$, and 8, respectively. Smoothly, these are the links of the complex algebraic singularities labeled by the same Dynkin type.

In general, one is able to construct and distinguish exact Lagrangian fillings using modern techniques and invariants, but showing that one has obtained all possible fillings remains an elusive problem. In particular, the uniqueness of the exact Lagrangian disk filling of the standard Legendrian unknot is the only complete classification currently known [**EP96**]. Later work in constructing [**EHK16**] and distinguishing [**Pan17**] decomposable exact Lagrangian fillings of $\lambda(A_n)$ gave a Floer theoretic proof of the existence of at least a Catalan number $C_n = \frac{1}{n+1}\binom{2n}{n}$ of Hamiltonian isotopy classes of such fillings. A Catalan number of exact Lagrangian fillings of $\lambda(A_n)$ was also constructed using Legendrian weaves and it was remarked that these two sets of fillings most likely coincided [**TZ18**, Section 2.3]. As a corollary of Theorem 1.1.1, we immediately obtain a proof of this statement. This corollary provides additional evidence for the following conjectural classification of exact Lagrangian fillings by Casals.

CONJECTURE (Conjecture 5.1, [**Cas21**]). *Let λ be a rainbow closure of a positive braid. Up to Hamiltonian isotopy there are precisely*

- *C_n exact Lagrangian fillings of $\lambda(A_n)$;*
- *$(3n-2)C_{n-1}$ exact Lagrangian fillings of $\lambda(D_n)$;*
- *833, 4160, 25080, exact Lagrangian fillings of $\lambda(E_6), \lambda(E_7)$ and $\lambda(E_8)$;*

- *and infinitely many exact Lagrangian fillings of λ otherwise.*

The precise numbers appearing above reflect a connection between exact Lagrangian fillings and cluster theory, as will be explained in much greater detail below. Our next result, concerning exact Lagrangian fillings of $\lambda(D_n)$, provides further support for Conjecture 1.1.

THEOREM 1.1.2 (Theorem [**Hug22**]). *There are at least $(3n-2)C_{n-1}$ exact Lagrangian fillings of $\lambda(D_n)$ up to Hamiltonian isotopy.*

In fact, the proof of Theorem 1.1.2 yields a stronger result, namely that each of the exact Lagrangian fillings is related to one of the others by a specific series of local geometric moves known as weave mutations.

1.1.1. Legendrian loops. In order to discuss the existence of infinitely many exact Lagrangian fillings where predicted by Conjecture 1.1 we turn to the study of **Legendrian loops**. These are isotopies of a Legendrian link that fix it pointwise at time one and act on the set of exact Lagrangian fillings of Legendrian links via concatenating the trace of the Legendrian isotopy to a given filling. They were first introduced by Kálmán in his study of the fundamental group of the space of embeddings of Legendrian torus links [**Kál05**]. More recently, Legendrian loops have appeared as crucial tools in constructing infinitely many exact Lagrangian fillings of large families of Legendrian links, including all rainbow closures of positive braids not of Dynkin type [**CG22**, **CN21**, **CS22**, **GSW20b**].

Nearly all of the works cited above show the existence of faithful $\mathbb{Z}$ actions on the Lagrangian concordance monoid. However, the very first proof of infinitely many exact Lagrangian fillings by Casals and Gao provides hints of a richer structure, showing in particular the existence of faithful $\mathrm{PSL}_2(\mathbb{Z})$ and $\mathrm{MCG}(\mathbb{S}^2, 4)$ actions on the set of fillings of Legendrian torus links $\lambda(3, 6)$ and $\lambda(4, 4)$, respectively. Here $\mathrm{MCG}(\mathbb{S}^2, 4)$ denotes the mapping class group of the 4-punctured sphere. The remainder of this work will focus on understanding additional structures of Legendrian loop actions on the Lagrangian concordance monoid.

As an initial attempt at providing more granular data about the action of particular loops, we start with the simplest case of $\lambda(A_n)$. We define the Kálmán loop to be the Legendrian loop pictured as taking the leftmost crossing of the positive braid $\beta = \sigma_1^n$ and dragging it up past the

left cusps, around top two strands of the front projection $\lambda(\beta)$, and down past the right cusps until it becomes the rightmost crossing. Studying the action of this Legendrian loop yields the following.

THEOREM 1.1.3. *The action of the Kálmán loop on the set $\mathcal{P}_n$, the Catalan number of exact Lagrangian fillings of A-type satisfies:*

The number of Kálmán loop orbits of fillings of $\lambda(A_n)$ is

$$\frac{C_{n+1}}{n+3} + \frac{C_{(n+1)/2}}{2} + \frac{2C_{n/3}}{3}$$

where the $C_{(n+1)/2}$ and $C_{n/3}$ terms appear if and only if the indices are integers.

This theorem gives a complete description of the orbital structure of the action of the Kálmán loop on the conjectural set of exact Lagrangian fillings of $\lambda(A_n)$. Chapter 2 also includes additional combinatorial characterizations for identifying exact Lagrangian fillings of various orbit sizes. Such a detailed understanding of Legendrian loop actions is not currently feasible in general, but with the help of cluster theory we can still pursue a richer understanding.

1.1.2. Clusters and cluster modular groups. In order to investigate Legendrian loop actions and distinguish exact Lagrangian fillings, we require the use of certain invariants of Legendrian links. There are two main packages of modern invariants that we use to distinguish Legendrian links and their exact Lagrangian fillings. The first is a sheaf-theoretic moduli arising from the microlocal theory of sheaves developed by Kashiwara and Schapira [**KS85**]. The second is the Legendrian contact differential graded algebra, often known simply as 'the DGA', which comes from a Floer-theoretic count of pseudoholomorphic disks with boundary. As it is often difficult to extract information from the DGA, we will often consider the augmentation variety, which in our setting can be thought of as the zeroeth homology of the DGA. Both invariants are functorial with respect to exact Lagrangian cobordism [**EHK16**, **Li21**]; given two Legendrians λ_1, λ_2, an exact Lagrangian cobordism from λ_2 to λ_1 induces a morphism $\mathcal{M}_1(\lambda_1) \to \mathcal{M}_1(\lambda_2)$ from the sheaf moduli of λ_1 to the sheaf moduli of λ_2. An analogous statement holds for the augmentation variety. In particular, concatenating the trace of a Legendrian loop φ to λ induces an automorphism $\tilde{\varphi}$ of $\mathcal{M}_1(\lambda)$.

For the remainder of our discussion here, we will focus our attention on the sheaf moduli in order to discuss cluster structures. Analogous statements hold for Floer-theoretic invariants with

an appropriate choice of additional data. Informally, a **cluster variety** is an algebraic variety with an atlas of toric charts $(\mathbb{C}^\times)^n$ such that the gluing maps are a particular form of birational transformation. Cluster theory appears in many disparate contexts, including high energy physics, hyperbolic geometry, and number theory. In our setting, an exact Lagrangian filling L of λ induces a toric chart $(\mathbb{C}^\times)^{b_1(L)}$ in the corresponding $\mathcal{M}_1(\lambda)$. Casals and Weng showed that for certain classes of Legendrians λ (including $\lambda = \lambda(\beta)$), these toric charts glue together in such a way as to yield the structure of a cluster variety for $\mathcal{M}_1(\lambda)$. A local surgery operation (see Section 1.2.1.3) taking one Lagrangian filling to another induces the birational transformation that glues together the corresponding toric charts.

Cluster varieties admit a finite type ADE classification matching the conjectural classification of exact Lagrangian fillings above. In fact, Conjecture 1.1 can be rephrased as stating that exact Lagrangian fillings of λ are in bijection with distinct cluster charts of $\mathcal{M}_1(\lambda)$. Cluster varieties also have several families that admit infinitely many cluster charts but are nevertheless relatively well behaved. These include affine type and extended affine type, obtained from Dynkin type in a straightforward way, as well as surface-type cluster varieties, which correspond to triangulations of surfaces with boundary and marked points.

In our contact-geometric setting, a loop φ of a Legendrian λ necessarily preserves the set of exact Lagrangian fillings of λ, implying that the induced action $\tilde{\varphi}$ permutes the set of toric charts induced by these fillings in a way that respects the cluster structure. From this, it follows that $\tilde{\varphi}$ is a **cluster automorphism** of $\mathcal{M}_1(\lambda)$. The set of all cluster automorphisms forms a group known as the **cluster modular group**. Our goal of studying Legendrian loop actions on the Lagrangian concordance monoid can therefore be rephrased in terms of understanding the cluster modular group $\mathcal{G}(\mathcal{M}_1(\lambda))$.

A systematic study of cluster modular groups was first undertaken by Assem, Schiffler, and Shramchenko in [**ASS12**], where they computed group presentations for several well-behaved families, including Dynkin, affine, and surface type cluster algebras. Later work of Fraser explored cluster modular groups for cluster structures on (top-dimensional positroid cells) of Grassmannians [**Fra18**]. Most recently, Kaufman and Greenberg [**KG21**] have consolidated and extended several known results using a particular combinatorial construction that we explain in Section 4.1.

Few computations of cluster modular groups exist beyond the cases listed above. However, in all of the cases that arise in this work, there is a central element of the cluster modular group known as the Donaldson-Thomas transformation, denoted DT. This automorphism has been explicitly constructed as the induced automorphism of a Legendrian isotopy composed with an involutive contactomorphism [**CW22**, Theorem 5.8] and plays a central role in establishing the existence of a cluster structure for $\mathcal{M}_1(\lambda)$. Somewhat more concretely, for rainbow closures of positive braids, the automorphism DT^2 is induced by the Legendrian loop rotating all of the crossings around the Legendrian [**GSW20b**].

In Section 4.1 we compile results about Legendrian loops and describe them as generators of cluster modular groups for certain families of cluster varieties. In addition to the links we have already encountered, the families of Legendrian links we consider are defined as follows:

- $\lambda(\tilde{D}_n)$ is the (-1) closure of $\beta = (\sigma_2\sigma_1\sigma_3\sigma_2)^2\sigma_1^{n-4}\Delta^2$ for $n \geq 4$.
- $\lambda(\tilde{E}_n)$ is the (-1) closure of $\beta = (\sigma_2\sigma_1\sigma_3\sigma_2)^2\sigma_3\sigma_1^{n-5}\Delta^2$ for $n \in \{6, 7, 8\}$
- $\lambda(E_7^{(1,1)})$ is the (-1) closure of $\beta = (\sigma_2\sigma_1\sigma_3\sigma_2)^2\sigma_3^2\sigma_1^2\Delta^2$
- $\lambda(E_8^{(1,1)})$ is the (-1) closure of $\beta = (\sigma_2\sigma_1\sigma_3\sigma_2)^2\sigma_3\sigma_1^4\Delta^2$

Here $\tilde{D}_n$ is read as "affine" D_n while the $(1, 1)$ appearing in $E_7^{(1,1)}$ indicates extended affine type. The Legendrians above are all Legendrian isotopic to rainbow closures of positive braids, but defining them as (-1) closures affords additional clarity when investigating Legendrian loops. Note that $\lambda(E_7^{(1,1)})$ and $\lambda(E_8^{(1,1)})$ are Legendrian isotopic to the Legendrian torus links $\lambda(4, 4)$ and $\lambda(3, 6)$, respectively.

Let us denote by $\mathcal{H}$ the set of links above along with the finite type links $\lambda(A_n), \lambda(D_n), \lambda(E_6), \lambda(E_7)$, and $\lambda(E_8)$. We also consider a particular subset

$$\mathcal{H}' := \{\lambda(A_n), \lambda(D_n), \lambda(E_6), \lambda(E_7), \lambda(E_8), \lambda(\tilde{E}_7), \lambda(\tilde{E}_8), \lambda(E_8^{(1,1)})\}.$$

By comparing Legendrian loop actions to known presentations of cluster modular groups we show that we can generate a finite index subgroup of the cluster modular group with Legendrian loops. If we add DT to our generating set, we can often obtain a stronger result.

THEOREM 1.1.4. *For $\lambda \in \mathcal{H}$, the cluster modular group $\mathcal{G}(\mathcal{M}_1(\lambda))$ is virtually generated by Legendrian loops. Moreover, if $\lambda \in \mathcal{H}'$, the group $\mathcal{G}(\mathcal{M}_1(\lambda))$ is generated by Legendrian loops and* DT.

The methods we use suggest ways to produce generators of cluster modular groups, even in cases where the full group is not known. We discuss some additional applications of this insight in Sections 4.1 and 4.2.

1.1.3. Connections to Mapping Class Groups. Cluster theory also gives us additional tools to analyze Legendrian loops. Fock and Goncharov originally motivated much of their work on cluster varieties by exploring higher Teichmuller theory for surfaces with boundary and marked points [**FG06b**]. In this context, cluster charts correspond to triangulations of a fixed surface Σ and the cluster modular group is intimately connected to the mapping class group $\mathrm{MCG}(\Sigma)$. Work of Ishibashi and collaborators [**Ish19**, **Ish20**, **IK21**, **AIK22**] has extended this connection by investigating classical results from the study of mapping class groups in a cluster-theoretic context.

In Section 4.2, we intepret Ishibashi's cluster Nielsen-Thurston classification of cluster automorphisms in the context of Legendrian loops. We define (cluster) periodic, reducible, and pseudo-Anosov Legendrian loops and discuss a fixed point property of their action on the positive real part of $\mathcal{M}_1(\lambda)$.

THEOREM 1.1.5. *A Legendrian loop is infinite order if its induced action on the positive real part of $X(\beta)$ or $\mathcal{M}_1(\Lambda(\beta))$ has no fixed points.*

In addition, we show that Legendrian loops coincide with specific generators of the cluster modular group known as cluster Dehn twists. This allows us to recover results of Kaufman and Greenberg that are analogous to a classical theorem from the study of mapping class groups.

THEOREM 1.1.6 ([**KG21**], Corollary 6.5). *For $\lambda \in \mathcal{H}$, the cluster modular group $\mathcal{G}(\mathcal{M}_1(\lambda))$ is virtually generated by cluster Dehn twists. Moreover, if $\lambda \in \mathcal{H}'$, then $\mathcal{G}(\mathcal{M}_1(\lambda))$ is generated by cluster Dehn twists.*

Organization: Section 1.2 below reviews the necessary background material related to constructions of exact Lagrangian fillings of Legendrian links, their invariants, and cluster theory. The

remainder of this dissertation is then presented in alphabetical order, with results pertaining to A_n-type in Chapter 2, results pertaining to D_n-type in Chapter 3, and results pertaining to everything else in Chapter 4. More specifically, Chapter 2 contains proofs of Theorems 1.1.1 and 1.1.3 as well as additional combinatorial characterizations of exact Lagrangian fillings of $\lambda(A_n)$. In Chapter 3, we prove Theorem 1.1.2. Finally, Chapter 4 discusses results relating Legendrian loops to known cluster modular groups, including a proof of Theorem 1.1.4. The chapter closes with proofs of Theorems 1.1.5, and 1.1.6, contact-geometric interpretations of Ishibashi's analogy linking cluster modular groups and mapping class groups.

1.2. Background

We begin with the necessary background material from contact and symplectic topology. The standard contact structure ξ_{st} in $\mathbb{R}^3$ is the 2-plane field given as the kernel of the 1-form $\alpha = dz - ydx$. A link $\lambda \subseteq (\mathbb{R}^3, \xi_{st})$ is Legendrian if λ is always tangent to ξ_{st}. As λ can be assumed to avoid a point, we can equivalently consider Legendrians λ contained in the contact 3-sphere $(\mathbb{S}^3, \xi_{st})$ [**Gei08**, Section 3.2]. We consider Legendrian links up to Legendrian isotopy, i.e. ambient isotopy through a family of Legendrians.

The symplectization of contact $\mathbb{R}^3$ is the symplectic manifold $(\mathbb{R}_t \times \mathbb{R}^3, d(e^t\alpha))$. Given two Legendrian links λ_- and λ_+, an exact Lagrangian cobordism $L \subseteq (\mathbb{R}_t \times \mathbb{R}^3, d(e^t\alpha))$ from λ_- to λ_+ is a cobordism Σ such that there exists some $T > 0$ satisfying the following:

(1) $d(e^t\alpha)|_\Sigma = 0$
(2) $\Sigma \cap ((-\infty, T] \times \mathbb{R}^3) = (-\infty, T] \times \lambda_-$
(3) $\Sigma \cap ([T, \infty) \times \mathbb{R}^3) = [T, \infty) \times \lambda_+$
(4) $e^t\alpha|_\Sigma = df$ for some function f on Σ.

An exact Lagrangian filling of the Legendrian link $\lambda \subseteq (\mathbb{R}^3, \xi_{\text{st}})$ is an exact Lagrangian cobordism L from $\emptyset$ to λ that is embedded in the symplectization $\mathbb{R}_t \times \mathbb{R}^3$. Equivalently, we consider L to be embedded in the symplectic 4-ball with boundary ∂L contained in contact $(\mathbb{S}^3, \xi_{\text{st}})$.

We will depict a Legendrian link $\lambda \subseteq (\mathbb{R}^3, \xi_{st})$ in either of two projections; the front projection $\Pi : (\mathbb{R}^3, \xi_{\text{st}}) \to \mathbb{R}^2$ given by $\Pi(x, y, z) = (x, z)$ or the Lagrangian projection $\pi : (\mathbb{R}^3, \xi_{\text{st}}) \to \mathbb{R}^2$ given by $\pi(x, y, z) = (x, y)$. In the Lagrangian projection, crossings of $\pi(\lambda)$ correspond precisely to

Reeb chords of λ, integral curves of the Reeb vector field ∂_z that start and end on λ. In the front projection, the Legendrian condition $T_x\lambda \subseteq \ker(dz - ydx)$ implies that $y = \frac{dz}{dx}$. Therefore, Reeb chords are given by pairs of points $(x_1, z_1), (x_2, z_2)$ with $x_1 = x_2$ and $\frac{dz_1}{dx_1} = \frac{dz_2}{dx_2}$. The key geometric content in proving Theorem 1.1.1 in Section 2.1 will involve a careful comparison of Reeb chords in the front and Lagrangian projections of slicings of elementary cobordisms.

1.2.1. Legendrian weaves. Let us now describe Legendrian weaves, a geometric construction of Casals and Zaslow that can be used to produce exact Lagrangian fillings of a Legendrian link $\lambda(\beta)$ [**CZ21**]. The key idea of Legendrian weaves is to combinatorially encode a *Legendrian* surface Λ in the 1-jet space $J^1\mathbb{D}^2 = T^*\mathbb{D}^2 \times \mathbb{R}_z$ by the singularities of its front projection in $\mathbb{D}^2 \times \mathbb{R}_z$. The Lagrangian projection of Λ then yields an exact Lagrangian surface in $T^*\mathbb{D}^2$.

The contact geometric setup of the Legendrian weave construction is as follows. Let β be a positive braid and let $\Delta := \sigma_1\sigma_2\sigma_1 \dots \sigma_n\sigma_{n-1} \dots \sigma_1$ denote a positive half twist in the braid group Br_n. We construct a filling of $\lambda(\beta)$ – equivalently, the (-1)-framed closure of $\beta\Delta^2$, pictured in Figure 1.1 (right) – by first describing a local model for a Legendrian surface Λ in $J^1\mathbb{D}^2 = T^*\mathbb{D}^2 \times \mathbb{R}_z$. We equip $T^*\mathbb{D}^2$ with the symplectic form $d(e^r\alpha)$ where $\ker(\alpha) = \ker(dy_1 - y_2 d\theta)$ is the standard contact structure on $J^1(\partial\mathbb{D}^2)$ and r is the radial coordinate. This choice of symplectic form ensures that the flow of $e^r\alpha$ is transverse to $J^1\mathbb{S}^1 \cong \mathbb{R}^2 \times \partial\mathbb{D}^2$ thought of as the cotangent fibers along the boundary of the 0-section. The Lagrangian projection of Λ is then a Lagrangian surface in $(T^*\mathbb{D}^2, d(e^r\alpha))$. Moreover, since $\Lambda \subseteq (J^1\mathbb{D}^2, \ker(dz - e^r\alpha))$ is a Legendrian, we immediately obtain the function $z : \pi(\Lambda) \to \mathbb{R}$ satisfying $dz = e^r\alpha|_{\pi(\Lambda)}$, demonstrating that $\pi(\Lambda)$ is exact.

The boundary of $\pi(\Lambda)$ is taken to be a positive braid β in $J^1\mathbb{S}^1$ so that we regard it as a Legendrian link in a contact neighborhood of $\partial\mathbb{D}^2$. As the 0-section of $J^1\mathbb{S}^1$ is Legendrian isotopic to a max-tb standard Legendrian unknot, we can take $\partial\pi(\Lambda)$ to equivalently be the standard satellite of the standard Legendrian unknot. Diagramatically, this implies that the braid β in $J^1\mathbb{S}^1$ can be given as the (-1)-framed closure of β in contact $\mathbb{S}^3$.

1.2.1.1. *N-Graphs and Singularities of Fronts.* To construct a Legendrian weave surface Λ in $J^1\mathbb{D}^2$, we combinatorially encode the singularities of its front projection in a colored graph. Local models for these singularities of fronts are classified by work of Arnold [**Ad90**, Section 3.2]. The

three singularities that appear in our construction describe elementary Legendrian cobordisms and are pictured in Figure 1.2.

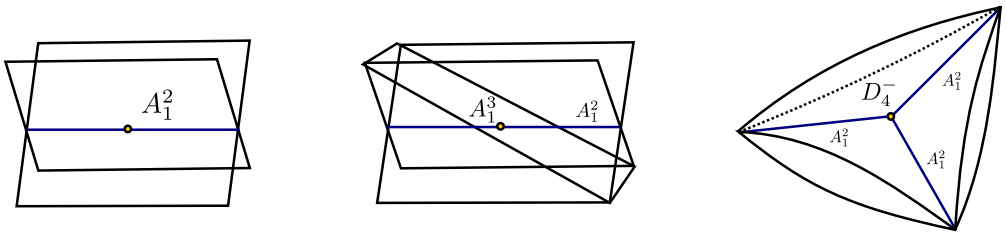

FIGURE 1.2. Singularities of front projections of Legendrian surfaces. Labels correspond to notation used by Arnold in his classification.

Since the boundary of our singular surface $\Pi(\Lambda)$ is the front projection of an N-stranded positive braid, $\Pi(\Lambda)$ can be pictured as a collection of N sheets away from its singularities. We describe the behavior at the singularities as follows:

(1) The A_1^2 singularity occurs when two sheets in the front projection intersect. This singularity can be thought of as the trace of a constant Legendrian isotopy in the neighborhood of a crossing in the front projection of the braid $\beta\Delta^2$.
(2) The A_1^3 singularity occurs when a third sheet passes through an A_1^2 singularity. This singularity can be thought of as the trace of a Reidemeister III move in the front projection.
(3) A D_4^- singularity occurs when three A_1^2 singularities meet at a single point. This singularity can be thought of as the trace of a 1-handle attachment in the front projection.

Having identified the singularities of fronts of a Legendrian weave surface, we encode them by a colored graph $\Gamma \subseteq \mathbb{D}^2$. The edges of the graph are labeled by Artin generators of the braid and we require that any edges labeled σ_i and σ_{i+1} meet at a hexavalent vertex with alternating labels while any edges labeled σ_i meet at a trivalent vertex. To obtain a Legendrian weave $\Lambda(\Gamma) \subseteq (J^1\mathbb{D}^2, \xi_{\text{st}})$ from an N-graph Γ, we glue together the local germs of singularities according to the edges of Γ. First, consider N horizontal sheets $\mathbb{D}^2 \times \{1\} \sqcup \mathbb{D}^2 \times \{2\} \sqcup \cdots \sqcup \mathbb{D}^2 \times \{N\} \subseteq \mathbb{D}^2 \times \mathbb{R}$ and an N-graph $\Gamma \subseteq \mathbb{D}^2 \times \{0\}$. We construct the associated Legendrian weave $\Lambda(\Gamma)$ as follows [**CZ21**, Section 2.3].

- Above each edge labeled σ_i, insert an A_1^2 crossing between the $\mathbb{D}^2 \times \{i\}$ and $\mathbb{D}^2 \times \{i+1\}$ sheets so that the projection of the A_1^2 singular locus under $\pi : \mathbb{D}^2 \times \mathbb{R} \to \mathbb{D}^2 \times \{0\}$ agrees with the edge labeled σ_i.
- At each trivalent vertex v involving three edges labeled by σ_i, insert a D_4^- singularity between the sheets $\mathbb{D}^2 \times \{i\}$ and $\mathbb{D}^2 \times \{i+1\}$ in such a way that the projection of the D_4^- singular locus agrees with v and the projection of the A_2^1 crossings agree with the edges incident to v.
- At each hexavalent vertex v involving edges labeled by σ_i and σ_{i+1}, insert an A_1^3 singularity along the three sheets in such a way that the origin of the A_1^3 singular locus agrees with v and the A_1^2 crossings agree with the edges incident to v.

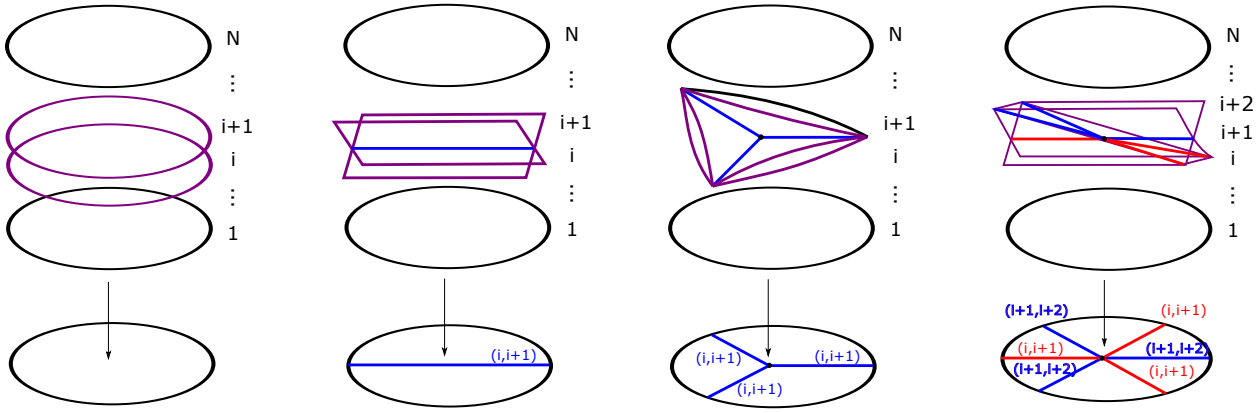

FIGURE 1.3. The weaving of singularities of fronts along the edges of the N-graph (courtesy of Roger Casals and Eric Zaslow, used with permission). Gluing these local models according to the N-graph Γ yields the weave $\Lambda(\Gamma)$.

If we take an open cover $\{U_i\}_{i=1}^m$ of $\mathbb{D}^2 \times \{0\}$ by open disks, refined so that any disk contains at most one of these three features, we can glue together the resulting fronts according to the intersection of edges along the boundary of our disks. Specifically, if $U_i \cap U_j$ is nonempty, then we define $\Pi(\Lambda(U_1 \cup U_2))$ to be the front resulting from considering the union of fronts $\Pi(\Lambda(U_1)) \cup \Pi(\Lambda(U_j))$ in $(U_1 \cup U_2) \times \mathbb{R}$.

DEFINITION 1.2.1. *The Legendrian weave* $\Lambda(\Gamma) \subseteq (J^1\mathbb{D}^2, \xi_{st})$ *is the Legendrian lift of the front* $\Pi(\Lambda(\cup_{i=1}^m U_i))$ *given by gluing the local fronts of singularities together according to the* N*-graph* Γ.

The immersion points of a Lagrangian projection of a weave surface Λ correspond precisely to the Reeb chords of Λ. In particular, if Λ has no Reeb chords, then $\pi(\Lambda)$ is an embedded exact Lagrangian filling of $\partial(\Lambda)$. In the Legendrian weave construction, Reeb chords correspond to critical points of functions giving the difference of heights between sheets. Every weave surface in this work admits an embedding where the distance between the sheets in the front projection grows monotonically in the direction of the boundary, ensuring that there are no Reeb chords.

1.2.1.2. *Homology of Weaves.* In this subsection, we describe the homology of a Legendrian weave $\Lambda(\Gamma)$. The smooth topology of $\Lambda(\Gamma)$ is that of an N-fold branched cover over $\mathbb{D}^2$ with simple branched points corresponding to each of the trivalent vertices of the N-graph Γ. Assuming that $\Lambda(\Gamma)$ is connected, the genus of $\Lambda(\Gamma)$ is then computed using the Riemann-Hurwitz formula:

$$g(\Lambda(\Gamma)) = \frac{1}{2}(v(\Gamma) + 2 - N\chi(\mathbb{D}^2) - |\partial\Lambda(\Gamma)|)$$

where $v(\Gamma)$ is the number of trivalent vertices of Γ and $|\partial\Lambda(\Gamma)|$ denotes the number of boundary components of Γ.

EXAMPLE 1.2.1. *If we apply this formula to the 3-graph $\Gamma_0(D_4)$, pictured in Figure 1.13 (left), we have 6 trivalent vertices and 3 link components, so the genus is computed as $g(\Lambda(\Gamma_0(D_4))) = \frac{1}{2}(6 + 2 - 3 - 3) = 1$.*

We now describe a recipe for finding elements of $H_1(\Lambda(\Gamma)); \mathbb{Z})$ combinatorially in terms of the N-graph Γ. We first consider an edge connecting two trivalent vertices. Closely examining the sheets of our surface, we can see that each such edge corresponds to a 1-cycle, as pictured in Figure 1.4 (left). We refer to such a 1-cycle as a short I-cycle. Similarly, any three edges of the same color that connect a single hexavalent vertex to three trivalent vertices correspond to a 1-cycle, as pictured in 1.5 (left). We refer to such a 1-cycle as a short Y-cycle. See figures 1.4 (right) and 1.5 (right) for a diagram of these 1-cycles in the front $\Pi(\Lambda(\Gamma))$. We can also consider a sequence of edges starting and ending at trivalent vertices and passing directly through any number of hexavalent vertices, as pictured in Figure 1.6. Such a cycle is referred to as a long I-cycle. Finally, we can combine any number of I-cycles and short Y-cycles to describe an arbitrary 1-cycle as a tree with leaves on trivalent vertices and edges passing directly through or branching at hexavalent vertices.

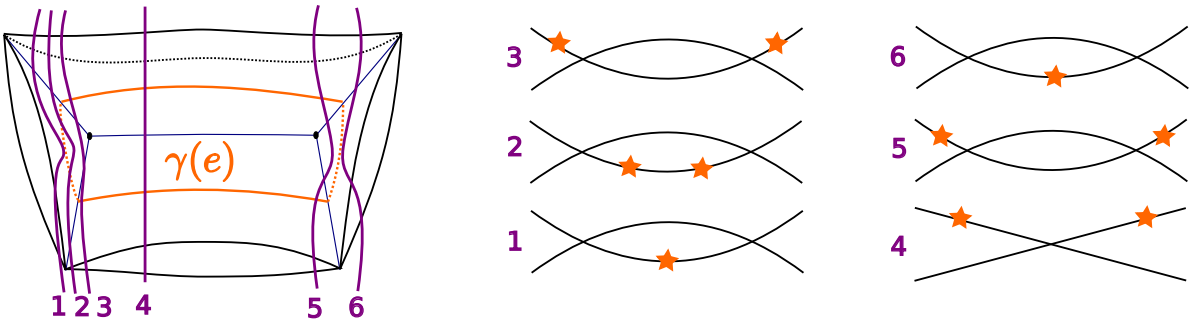

FIGURE 1.4. A short I-cycle $\gamma(e)$ for the edge $e \in G$ pictured in the front $\Pi(\Lambda(\Gamma))$ (left) and a vertical slicing of $\Pi(\Lambda(\Gamma))$ (right).

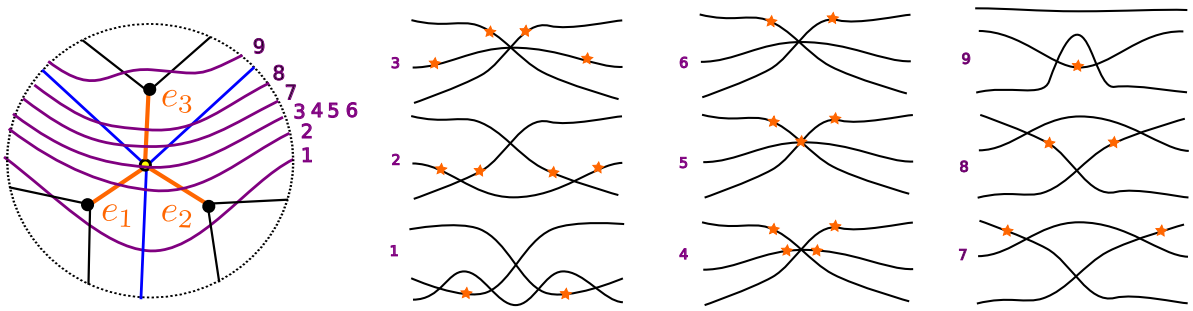

FIGURE 1.5. A short Y-cycle $\gamma(e)$ defined by the edges $e_1, e_2, e_3 \in G$ pictured in the front $\Pi(\Lambda(\Gamma))$ (left) and a vertical slicing of $\Pi(\Lambda(\Gamma))$ (right).

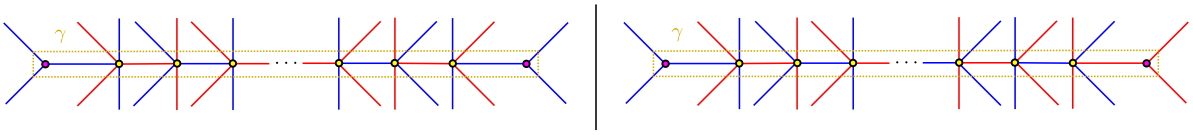

FIGURE 1.6. A pair of long I-cycles, both denoted by γ. The cycle on the left passes through an even number of hexavalent vertices, while the cycle on the right passes through an odd number.

The intersection form $\langle \cdot, \cdot \rangle$ on $H_1(\Lambda(\Gamma))$ plays a key role in distinguishing our Legendrian weaves. If we consider a pair of 1-cycles $\gamma_1, \gamma_2 \in H_1(\Lambda(\Gamma))$ with nonempty geometric intersection in Γ, as pictured in Figure 1.7, we can see that the intersection of their projection onto the graph Γ differs from their intersection in $\Lambda(\Gamma)$. Specifically, we can carefully examine the sheets that the 1-cycles

cross in order to see that γ_1 and γ_2 intersect only in a single point of $\Lambda(\Gamma)$. If we fix an orientation on γ_1 and γ_2, then we can assign a sign to this intersection based on the convention given in Figure 1.7. We refer to the signed count of the intersection of γ_1 and γ_2 as their algebraic intersection and denote it by $\langle \gamma_1, \gamma_2 \rangle$. For the remainder of this manuscript, we will fix a counterclockwise orientation for all of our cycles and adopt the convention that any two cycles γ_1 and γ_2, intersecting at a trivalent vertex as in Figure 1.7 have algebraic intersection $\langle \gamma_1, \gamma_2 \rangle = -1$.

Notation: For the sake of visual clarity, we will often represent an element of $H_1(\Lambda(\Gamma); \mathbb{Z})$ by a colored edge. This also ensures that the geometric intersection more accurately reflects the algebraic intersection. The original coloring of the edges can be readily obtained by examining Γ and its trivalent vertices. □

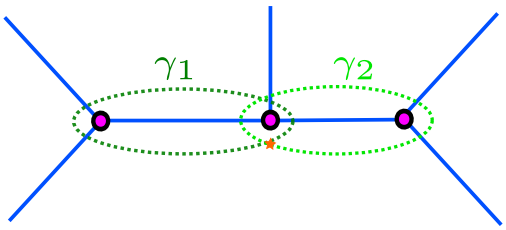

FIGURE 1.7. Intersection of two cycles, γ_1 and γ_2. The intersection point is indicated by an orange star. If we orient both cycles counterclockwise, then we will set $\langle \gamma_1, \gamma_2 \rangle = -1$ as our convention.

In our correspondence between N-graphs and weaves, we must consider how a Legendrian isotopy of the weave $\Lambda(\Gamma)$ affects the N-graph Γ and its combinatorially encoded homology basis. We can restrict our attention to certain isotopies, referred to as Legendrian Surface Reidemeister moves. These moves create specific changes in the Legendrian front $\Pi(\Lambda(\Gamma))$ studied by Arnol'd [**Ad90**]. From [**CZ21**], we have the following theorem relating surface Reidemeister moves to the corresponding N-graphs.

THEOREM 1.2.1 ([**CZ21**], Theorem 4.2). *Let Γ and Γ' be two N-graphs related by one of the moves shown in Figure 1.8. Then the associated weaves $\Lambda(\Gamma)$ and $\Lambda(\Gamma')$ are Legendrian isotopic relative to their boundaries.* □

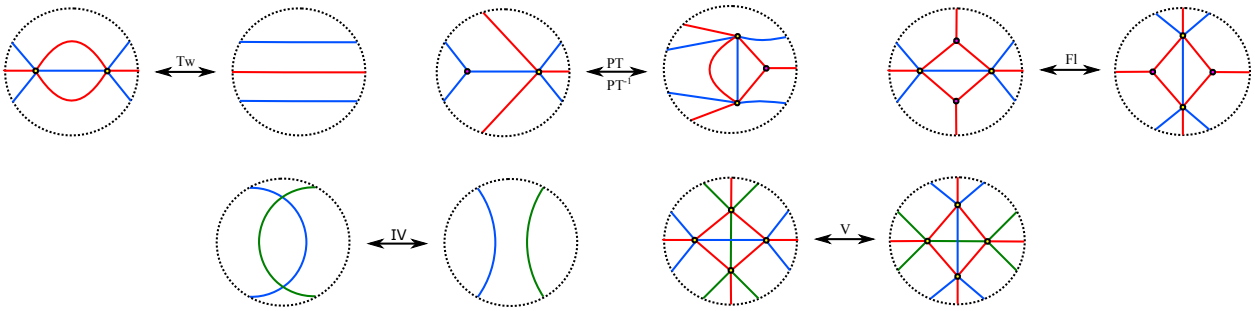

FIGURE 1.8. Legendrian Surface Reidemeister moves for N-graphs. Clockwise from top left, a candy twist, a push-through, a flop, and two additional moves, denoted by I, II, III, IV, and V respectively.

See Figure 1.9 for a description of the behavior of elements of $H_1(\Lambda(\Gamma);\mathbb{Z})$ under Moves I, II and III. The behavior under Move V can be deduced from Move III. In the pair of graphs in Figure 1.9 (center), we have denoted a push-through by II or II^{-1} depending on whether we go from left to right or right to left. We will sometimes refer to Move II^{-1} as a reverse push-through.

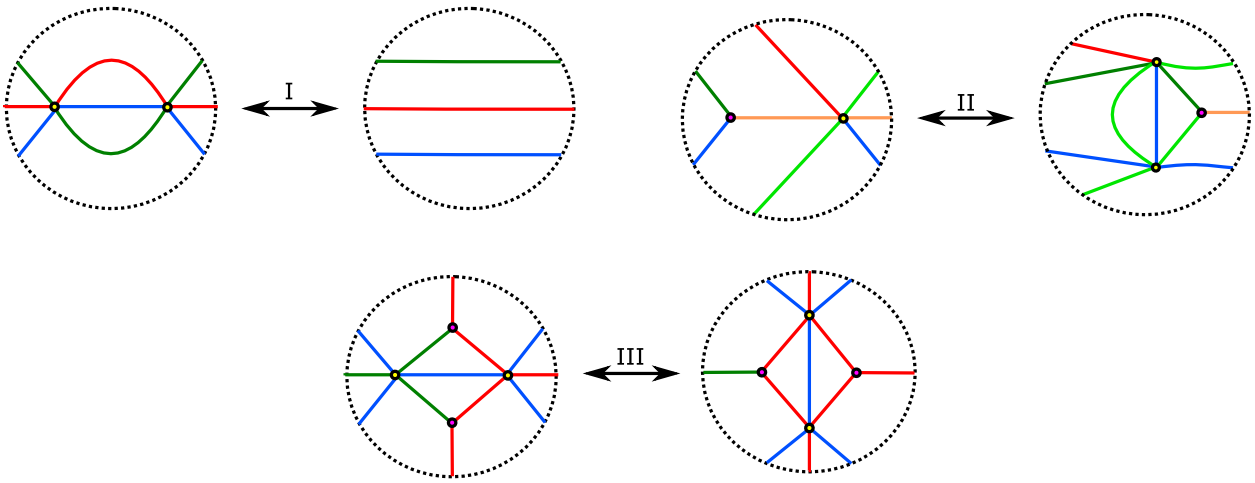

FIGURE 1.9. Behavior of certain homology cycles under Legendrian Surface Reidemeister moves.

REMARK 1.2.1. *It is also possible to verify the computations in Figure 1.9 by examining the relative homology class of a cycle; given a basis of the relative homology $H_1(\Lambda(\Gamma), \partial\Lambda(\Gamma);\mathbb{Z})$, the intersection form on that basis allows us to determine a given cycle by Poincaré-Lefschetz duality.*

1.2.1.3. *Mutations of weaves.* In this section we describe Legendrian mutation, a geometric operation that we can use to generate additional Legendrian weaves. Given a Legendrian weave $\Lambda(\Gamma)$ and a 1-cycle $\gamma \in H_1(\Lambda(\Gamma); \mathbb{Z})$, the Legendrian mutation $\mu_\gamma(\Lambda(\Gamma))$ outputs a Legendrian weave smoothly isotopic to $\Lambda(\Gamma)$ but that is generally not Legendrian isotopic to $\Lambda(\Gamma)$.

DEFINITION 1.2.2. *Two Legendrian surfaces $\Lambda_0, \Lambda_1 \subseteq (\mathbb{R}^5, \xi_{st})$ with equal boundary $\partial\Lambda_0 = \partial\Lambda_1$, are mutation-equivalent if and only if there exists a compactly supported Legendrian isotopy $\{\tilde{\Lambda}_t\}$ relative to the boundary, with $\tilde{\Lambda}_0 = \Lambda_0$ and a Darboux ball (B, ξ_{st}) such that*

(i) *Outside the Darboux ball, we have $\tilde{\Lambda}_1|_{\mathbb{R}^5 \setminus B} = \Lambda_1|_{\mathbb{R}^5 \setminus B}$*

(ii) *There exists a global front projection $\pi : \mathbb{R}^5 \to \mathbb{R}^3$ such that the pair of fronts $\Pi|_{B \cap \tilde{\Lambda}_1}$ and $\Pi|_{B \cap \Lambda_1}$ coincides with the pair of fronts in Figure 1.10 below.*

□

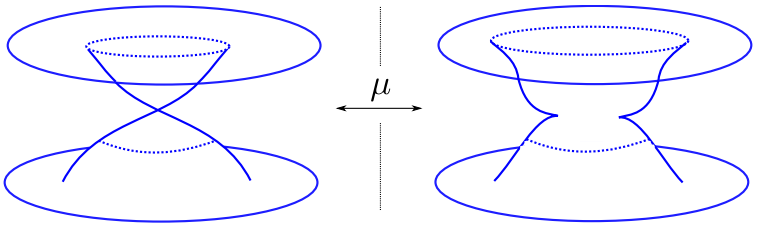

FIGURE 1.10. Local fronts for two Legendrian cylinders non-Legendrian isotopic relative to their boundaries.

We briefly note that these two fronts lift to non-Legendrian isotopic Legendrian cylinders in $(\mathbb{R}^5, \xi_{\text{st}})$, relative to the boundary, and that the 1-cycle we input for our operation is precisely the 1-cycle defined by the cylinder corresponding to Λ_0.

Combinatorially, we can describe mutation in terms of the N-graph associated to a weave. Figure 1.11 (left) depicts mutation at a short I-cycle, while Figure 1.11 (right) depicts mutation at a short Y-cycle. In the $N = 2$ setting, we can identify 2-graphs with triangulations of an $n-$gon, in which case mutation at a short I-cycle corresponds to a Whitehead move. In order to describe mutation at a short Y-cycle, we can first reduce the short Y-cycle case to a short I-cycle, as shown in Figure 1.12, before applying our mutation. See [**CZ21**, Section 4.9] for a more general description of mutation at long I- and Y-cycles in N-graphs.

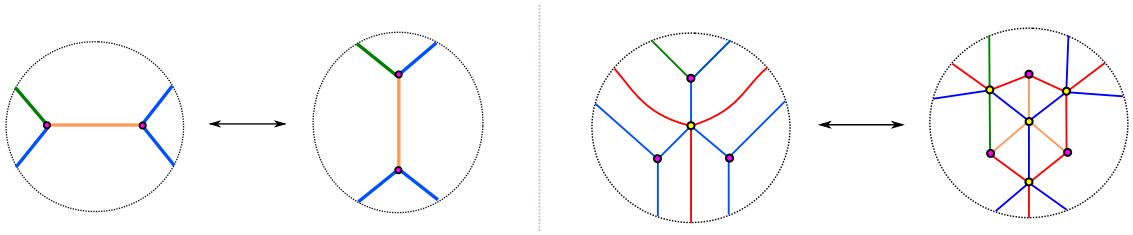

FIGURE 1.11. Mutations at a short I-cycle (left) and short Y-cycle (right). In both cases, the dark green edge depicts the effect of mutation on any cycle intersecting the orange cycle.

The geometric operation above coincides with the combinatorial manipulation of the N-graphs. Specifically, we have the following theorem.

THEOREM 1.2.2 ([**CZ21**], Theorem 4.2.1). *Given two N-graphs, Γ and Γ' related by either of the combinatorial moves described in Figure 1.11, the corresponding Legendrian weaves $\Lambda(\Gamma)$ and $\Lambda(\Gamma')$ are mutation-equivalent relative to their boundary.* □

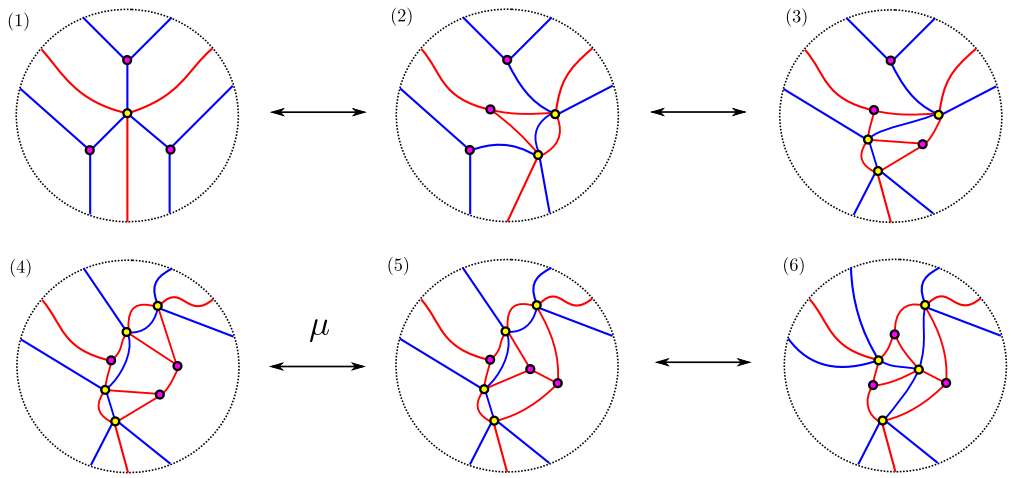

FIGURE 1.12. Mutation at a short Y-cycle given as a sequence of Legendrian Surface Reidemeister moves and mutation at a short I-cycle. The Y-cycle in the initial graph is given by the three blue edges that each intersect the yellow vertex in the center.

1.2.1.4. *Quivers from Weaves.* We complete our discussion of Legendrian weaves by describing quivers and how they arise via the intersection form on $H_1(\Lambda(\Gamma);\mathbb{Z})$. A quiver is a directed graph

without loops or oriented 2-cycles. In the Legendrian weave setting, the data of a quiver can be extracted from a given weave and a basis of its first homology via the intersection form. The intersection quiver is defined as follows: for every basis element $\gamma_i \in H_1(\Lambda(\Gamma);\mathbb{Z})$ we have a vertex v_i in the quiver; there are k arrows pointing from v_j to v_i if $\langle \gamma_i, \gamma_j \rangle = k$ for $k > 0$. See Figure 1.13 (left) for an example of the quiver $Q(\Lambda(\Gamma_0(D_4)), \{\gamma_i^{(0)}\})$ defined by $\Lambda(\Gamma_0(D_4))$ and the indicated basis for $H_1(\Lambda(\Gamma_0(D_4);\mathbb{Z})$.

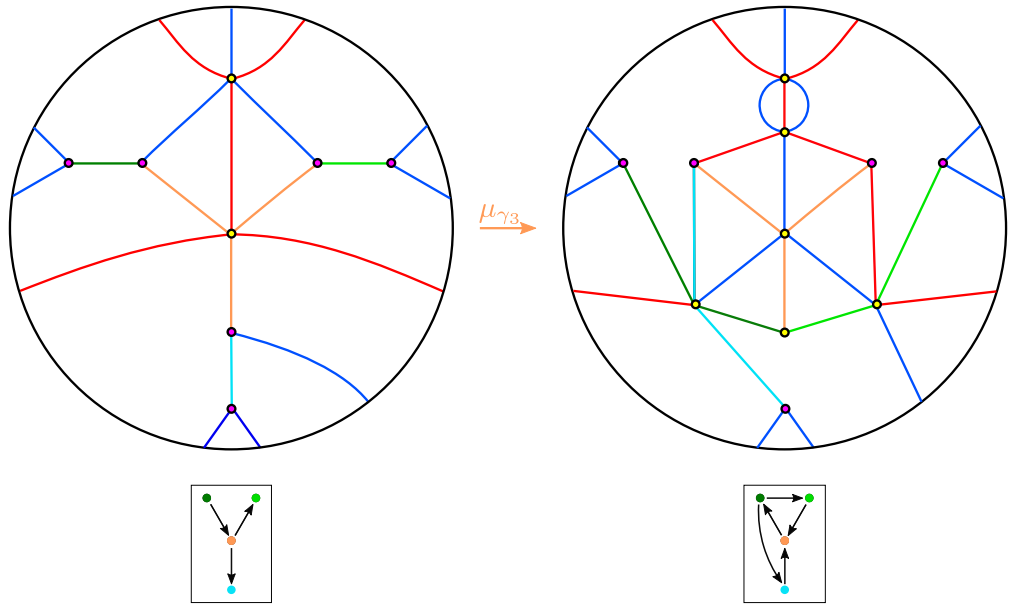

FIGURE 1.13. 3-graphs $\Gamma_0(D_4)$ (left) and $\Gamma_1(D_4)$ (right) obtained by mutation at the homology cycle γ_3, pictured with their associated intersection quivers $Q(\Gamma_0(D_4), \{\gamma_i^{(0)}\})$ and $Q(\Gamma_1(D_4), \{\gamma_i^{(0)}\})$. The Legendrian surface $\Lambda(\Gamma_0(D_4)$ is smoothly a 3-punctured torus, so that $H_1(\Lambda(\Gamma_0(D_4));\mathbb{Z})$ has rank 4. The basis $\{\gamma_i^{(0)}\}$ for $H_1(\Lambda(\Gamma_0(D_4));\mathbb{Z})$ is depicted by the colored cycles drawn in the graph.

The combinatorial operation of quiver mutation at a vertex v is defined as follows, see e.g. [**FWZ20a**, Definition 2.1.2]. First, for every pair of incoming edges and outgoing edges, we add an edge starting at the tail of the incoming edge and ending at the head of the outgoing edge. Next, we reverse the direction of all edges adjacent to v. Finally, we cancel any directed 2-cycles. If we started with the quiver Q, then we denote the quiver resulting from mutation at v by $\mu_v(Q)$. See Figure 1.13 (bottom) for an example. Under this operation, we can naturally identify the vertices

of Q with $\mu_v(Q)$, just as we can identify the homology bases of a weave before and after Legendrian mutation.

REMARK 1.2.2. *The crucial difference between algebraic and geometric intersections is captured in the step canceling directed 2-cycles. This cancellation is implemented by default in a quiver mutation, as the arrows of the quiver only capture algebraic intersections. In contrast, the geometric intersection of homology cycles after a Legendrian mutation will, in general, not coincide with the algebraic intersection. This dissonance will be explored in detail in the proof of Theorem 1.1.2 in Chapter 3.* □

The following theorem relates the two operations of quiver mutation and Legendrian mutation:

THEOREM 1.2.3 ([**CZ21**], Section 7.3). *Given an N-graph Γ, Legendrian mutation at an embedded cycle γ induces a quiver mutation for the associated intersection quivers, taking $Q(\Gamma, \{\gamma_i\})$ to $\mu_\gamma(Q(\Gamma, \{\gamma_i\}))$.* □

See Figure 1.13 for an example showing the quiver mutation of $Q(\Gamma_0(D_4), \{\gamma_i^{(0)}\})$, $i \in [1, 4]$, corresponding to Legendrian mutation applied to $\Lambda(\Gamma_0(D_4))$.

1.2.2. The microlocal theory of sheaves. In this section we define the moduli of microlocal rank-one sheaves with prescribed singular support, one of two families of Legendrian invariants that we consider in this manuscript. This sheaf-theoretic invariant of a Legendrian is defined with the auxiliary data of a front projection. The singular support of a sheaf encodes the codirection along which a sheaf fails to propagate. Informally, the moduli we consider is the space of all constructible sheaves (with some additional conditions) for which the singular support is contained within a given Legendrian front. In the setting we work in, this reduces to computations involving spaces of flags of vector spaces. Standard references are [**KS90**, **STZ17**]. See also the appendices of [**CL22**] for a careful explanation of the relationships between various categories involved in defining this moduli in greater generality.

1.2.2.1. *Singular support.* We first consider the dg-category of complexes of sheaves of $\mathbb{C}-$modules on M with constructible cohomology sheaves, denoted $Sh(M)$. For our purposes, M can either be taken to be S^1 or D^2. The following definition introduces the notion of the singular support of $\mathcal{F}$.

Intuitively, this notion captures the (co)directions along which the map induced by restriction fails to be a quasi-isomorphism. Let $R\Gamma$ denote the derived sections functor.

DEFINITION 1.2.3. *The singular support $SS(\mathcal{F}) \subseteq T^*M$ of a complex of sheaves $\mathcal{F} \subseteq Sh(M)$ is the closure of the set of covectors $(x, \xi) \in T^*M$ such that there exists an open neighborhood U_x about x, a smooth function $\varphi \in C^\infty(M)$ satisfying $d\varphi = \xi$, $\varphi(x) = 0$, and*

$$\mathrm{Cone}(R\Gamma((\varphi < -\delta) \cap U_x, \mathcal{F}) \to R\Gamma((\varphi < \delta) \cap U_x, \mathcal{F})) \ncong 0.$$

We denote the unit cotangent bundle of M by $T^\infty M$ and note that it carries a canonical contact structure inherited from the symplectic structure on T^*M. As defined, the singular support $SS(\mathcal{F})$ is a conical Lagrangian submanifold of the cotangent bundle T^*M. If we quotient by the $\mathbb{R}_+$ scaling action, then $SS(\mathcal{F})$ becomes a Legendrian submanifold of $T^\infty(M)$. We identify the unit cotangent bundle $T^{\infty,-}(M \times \mathbb{R})$ with the first jet space $J^1(M)$.

Given a Legendrian link λ in $J^1\mathbb{S}^1$ the front projection $\Pi(\lambda)$ stratifies $\mathbb{S}^1 \times \mathbb{R}$. We can then ask the following question: for what sheaves $\mathcal{F} \in Sh(\mathbb{S}^1 \times \mathbb{R})$ do we have $SS(F) \subseteq \Pi(\lambda)$? For a front without cusps and with binary Maslov potential in $J^1(S^1)$, we can restrict our attention to sheaves concentrated in degree 0 [**STZ17**, Proposition 5.17]. In this work, we will further simplify to only consider locally constant sheaves of vector spaces that are 0 in a neighborhood of $\mathbb{S}^1 \times \{-\infty\}$. For a Legendrian λ in J^1M, we denote by $\mathcal{M}(\lambda)$ the space of sheaves $\mathcal{F} \in Sh(M \times \mathbb{R})$ satisfying these conditions.

THEOREM 1.2.4 ([**GKS12**]). *Given a Legendrian isotopy $\{\lambda_t\}$, $t \in [0, 1]$, there is a isomorphism $\mathcal{M}(\lambda_0) \cong \mathcal{M}(\lambda_1)$.*

Note that the result in loc. cit. provides a means of explicitly computing the induced isomorphism between sheaf moduli. We will not require any such computations using Theorem 1.2.4 in this work. To describe $\mathcal{M}(\lambda)$ in the case of Legendrian links in $(J^1\mathbb{S}^1, \xi_{st})$, one (in practice) examines the stratification of $S^1 \times \mathbb{R}$ induced by the front projection $\Pi(\lambda)$ and locally computes the conditions on maps $f : V^\bullet \to W^\bullet$ between different regions of $\mathbb{S}^1 \times \mathbb{R} \backslash \Pi(\lambda)$ in neighborhoods of arcs, crossings, and cusps [**STZ17**]. The following definition allows us to simplify our space of sheaves yet further.

DEFINITION 1.2.4. *The microlocal rank of a sheaf $\mathcal{F}$ at a point $x \in \mathcal{M}(\lambda)$ is the rank of $Cone(V^\bullet \to W^\bullet)$.*

We denote by $\mathcal{M}_1(\lambda)$ the moduli of microlocal-rank one sheaves with singular support on the Legendrian λ. As noted above, $\mathcal{M}_1(\lambda)$ is functorial with respect to exact Lagrangian cobordism [**Li21**]. For a Legendrian weave $\Lambda(\Gamma)$, away from a D_4^- singularity, the front is locally of the form arc $\times [a,b]$ or crossing $\times [a,b]$. A computation of the singular support around a D_4^- singularity shows that no new singular support condition is imposed there [**TZ18**, Section 4.1.2], see also [**CZ21**, Theorem 5.3]. As a result, an embedded, orientable Legendrian weave $\Lambda(\Gamma)$ induces a map $Loc(\Lambda(\Gamma)) \hookrightarrow \mathcal{M}_1(\partial\Lambda(\Gamma))$. In the following subsection, we explicitly describe how to compute $\mathcal{M}_1(\partial\Lambda(\Gamma))$ and a toric chart induced by $\Lambda(\Gamma)$.

1.2.2.2. *Flag moduli.* For a front without cusps and with binary Maslov potential in $J^1(S^1)$, we can describe the sheaf moduli as a space of flags with certain transversality conditions given by the front. If β is a positive braid, then the (-1) closure of β is Legendrian isotopic to a front in $J^1\mathbb{S}^1$ without cusps and with binary Maslov potential.

Let $\lambda \subseteq J^1\mathbb{S}^1$ be the (-1) closure of an N-stranded positive braid $\beta \in Br_N^+$. The crossing singularities of the front $\Pi(\beta) \subseteq S^1 \times \mathbb{R}$ divide $S^1 \times \mathbb{R}$ into regions $[a_i, a_{i+1}] \times \mathbb{R}$. Given our description of $\mathcal{M}_1(\lambda)$ above, each sheaf $\mathcal{F} \in \mathcal{M}_1(\lambda)$ has as a stalk a locally constant vector space in each region of $[a_i, a_{i+1}] \times \mathbb{R}$, divided by the strands of β. The microlocal rank-one condition implies that the rank of these vector spaces increases by one as we pass from a region to the one immediately above it. This sequence of vector spaces $0 \subseteq V_i^1 \subseteq \cdots \subseteq V_i^n = \mathbb{C}^n$ can be thought of as a flag $V_i^\bullet$ in $\mathbb{C}^N$. The singular support condition at the crossing corresponding to the Artin generator σ_j implies that the two flags $V_i^\bullet$ and $V_{i+1}^\bullet$ differ at the jth position; in other words, V_i^j is transverse to V_{i+1}^j. The space $\mathcal{M}_1(\lambda)$ can then be understood as the space of flags satisfying these transversality conditions, modulo a choice of basis.

For a Legendrian weave $\Lambda(\Gamma)$ defined by the N-graph Γ, we have a similar story. As described in [**CZ21**, Section 5.3] the data of $\mathcal{M}_1(\Lambda(\Gamma))$ is equivalent to providing:

(i) An assignment to each face F (connected component of $\mathbb{D}^2 \backslash G$) of a flag $V^\bullet(F)$ in the vector space $\mathbb{C}^N$.

(ii) For each pair F_1, F_2 of adjacent faces sharing an edge labeled by σ_i, we require that the corresponding flags satisfy

$$V^j(F_1) = V^j(F_2), \qquad 0 \leq j \leq N, j \neq i, \qquad \text{and} \qquad V^i(F_1) \neq V^i(F_2).$$

Finally, we consider the moduli space of flags satisfying (i) and (ii) modulo the diagonal action of $GL_N(\mathbb{C})$ on $V^\bullet$. The precise statement [**CZ21**, Theorem 5.3] we require is that the flag moduli space, denoted $\mathcal{C}(\Gamma)$ is isomorphic to the space of microlocal rank-one sheaves $\mathcal{M}_1(\Lambda(\Gamma))$. Since $\mathcal{M}_1(\Lambda(\Gamma))$ is an invariant of $\Lambda(\Gamma)$ up to Hamiltonian isotopy, it follows that $\mathcal{C}(\Gamma)$ is an invariant as well.

To better understand local systems on $\Lambda(\Gamma)$, we give examples of the flag moduli space in a neighborhood of homology cycles of $\Lambda(\Gamma)$. In the short I-cycle case, when the edges are labeled by σ_1, the moduli space is determined by four lines $a \neq b \neq c \neq d \neq a$, as pictured in Figure 1.14 (left). If the edges are labeled by σ_2, then the data is given by four planes $A \neq B \neq C \neq D \neq A$. Around a short Y-cycle, the data of the flag moduli space is given by three distinct planes $A \neq B \neq C \neq A$ contained in $\mathbb{C}^3$ and three distinct lines $a \subsetneq A, b \subsetneq B, c \subsetneq C$ with $a \neq b \neq c \neq a$, as pictured in Figure 1.14 (right).

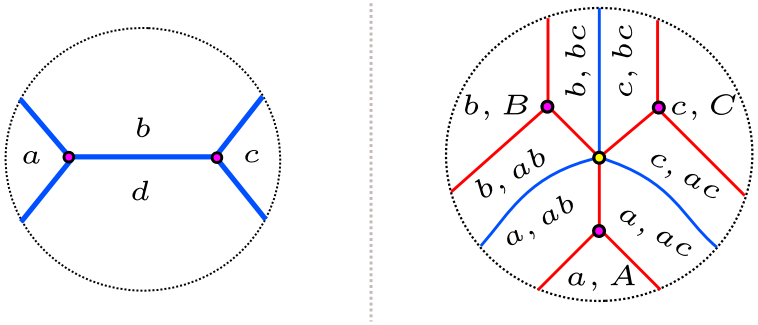

FIGURE 1.14. The data of the flag moduli space given in the neighborhood of a short I-cycle (left) and a short Y-cycle (right). Lines are represented by lowercase letters, while planes are written in uppercase. The intersection of the two lines a and b is written as ab.

To describe a cluster algebra structure on $\mathcal{C}(\Gamma)$, we need to specify the toric chart $(\mathbb{C}^\times)^{b_1(\Lambda(\Gamma))}$ associated to the quiver $Q(\Lambda(\Gamma), \{\gamma_i\})$ via the microlocal monodromy functor μ_{mon}. This is a

functor from the category $\mathcal{C}(\Gamma)$ to the category of rank one local systems on $\Lambda(\Gamma)$. As described in [**STZ17**, **STWZ19**], the functor μ_{mon} takes a 1-cycle as input and outputs the isomorphism of sheaves given by the monodromy about the cycle. Since it is locally defined, we can compute the microlocal monodromy about an I-cycle or Y-cycle using the data of the flag moduli space in a neighborhood of the cycle. If we have a short I-cycle γ with flag moduli space described by the four lines a, b, c, d, as in Figure 1.14 (left), then the microlocal monodromy about γ is given by the cross ratio

$$\frac{a \wedge b}{b \wedge c} \frac{c \wedge d}{d \wedge a}$$

Similarly, for a short Y-cycle with flag moduli space given as in Figure 1.14 (right), the microlocal monodromy is given by the triple ratio

$$\frac{B(a)C(b)A(c)}{B(c)C(a)A(b)}$$

where we interpret the plane B as a covector in $\mathbb{C}^3$ to define the pairing $B(a)$. As described in [**CZ21**, Section 7.2], the microlocal monodromy about a 1-cycle gives rise to an X-cluster variable at the corresponding vertex in the quiver. Under mutation of the 3-graph, the cross ratio and triple ratio transform as cluster X-coordinates. Specifically, if we start with a 3-graph with cluster variables x_j, then the cluster variables x'_j of the 3-graph after mutating at γ_i are given by the equation

$$x'_j = \begin{cases} x_j^{-1} & i = j \\ x_j(1 + x_i^{-1})^{-\langle \gamma_i, \gamma_j \rangle} & \langle \gamma_i, \gamma_j \rangle > 0 \\ x_j(1 + x_i)^{-\langle \gamma_i, \gamma_j \rangle} & \langle \gamma_i, \gamma_j \rangle < 0 \end{cases}$$

See Figure 1.15 for an example.

1.2.3. Preliminaries on cluster theory. In this subsection, we define cluster structures and discuss how to use the ingredients described above to define them. We start by introducing cluster algebras and cluster varieties and then discuss some specific cluster algebras related to tagged triangulations of surfaces. The subsection ends with a description of cluster modular groups. See [**FWZ20a**, **FWZ20b**] for an introductory reference on cluster theory.

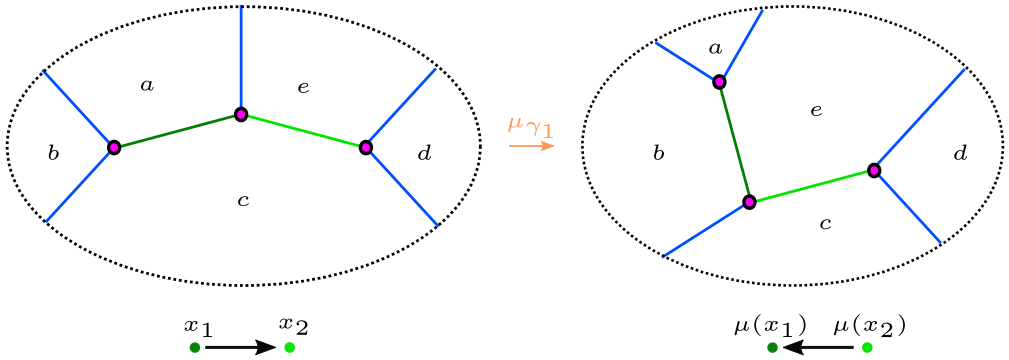

FIGURE 1.15. Prior to mutating at γ_1, we have $\langle \gamma_1, \gamma_2 \rangle = -1$. Computing the cross ratios for γ_1 and $\mu_1(\gamma_1)$ we can see that the cross ratio transforms as $\mu_1(\gamma_1) = \frac{b\wedge c}{c\wedge e}\frac{e\wedge a}{a\wedge b} = x_1^{-1}$ under mutation. Similarly, computing the cross ratios for γ_1 and $\mu_1(\gamma_2)$ and applying the relation $e \wedge b \cdot a \wedge c = b \wedge c \cdot e \wedge a + a \wedge b \cdot c \wedge e$, we have $\mu_1(x_2) = \frac{e\wedge a}{a\wedge c}\frac{c\wedge d}{d\wedge e}\left(1 + \frac{a\wedge b}{b\wedge c}\frac{c\wedge e}{e\wedge a}\right)$.

1.2.3.1. *Definition of a cluster algebra.* We follow [**FWZ20a**] for describing a cluster algebra associated to an initial quiver Q_0. To an initial quiver Q_0 with n vertices, we associate an initial set of variables $a_1, \dots a_n$, one for each vertex. Together, the $n-$tuple $\mathbf{a} = (\mathbf{a_1}, \dots, \mathbf{a_n})$ and the quiver Q_0 form a **cluster seed** $(\mathbf{a}, \mathbf{Q_0})$. We designate a subset of vertices Q_0^{mut} to be the mutable part of the quiver. The vertices in $Q_0 \backslash Q_0^{mut}$ are designated as frozen and we require that there are no arrows between them. There are two types of cluster algebras, type A or type X, depending on the precise form of the birational map relating different cluster variables.

DEFINITION 1.2.5. *Let $(\mathbf{a}, \mathbf{Q})$ be a cluster seed and $k \in Q^{mut}$ be a mutable vertex. The cluster A **seed mutation** μ_k is an operation taking as input the seed $(\mathbf{a}, \mathbf{Q})$ and outputs the new seed $(\mathbf{a}', \mathbf{Q}')$ where Q' is related to Q by quiver mutation at vertex k and $\mathbf{a}'$ is related to $\mathbf{a}$ by $a_i' = a_i$ for all $i \neq k$ and*

$$a_k a_k' = \prod_{i \to k} a_i^{\text{number of edges } i \to k} + \prod_{i \leftarrow k} a_i^{\text{number of edges } i \leftarrow k}$$

where $i \to k$ denotes an edge from i to k in Q.

Note that seed mutation is an involution, so that $\mu_k^2(\mathbf{a}, \mathbf{Q}) = (\mathbf{a}, \mathbf{Q})$.

Denote by $\mathcal{F}$ the field of rational functions $\mathbb{C}(a_1, \dots a_n)$ and consider an initial seed $(\mathbf{a}, \mathbf{Q_0}) \subseteq \mathcal{F}$.

DEFINITION 1.2.6. *The type A cluster algebra generated by $(\mathbf{a}, \mathbf{Q_0})$ is the $\mathbb{C}$-algebra generated by all cluster variables arising in arbitrary mutations of the initial seed.*

The type X cluster algebra is generated from an initial seed $(\mathbf{x}, \mathbf{Q_0})$ by the mutation formula

$$x_j' = \begin{cases} x_j^{-1} & i = j \\ x_j(y_k+1)^{\text{number of edges } j \leftarrow k} & j \neq k, j \leftarrow k \\ x_j(x_k^{-1}+1)^{-\text{number of edges } j \rightarrow k} & j \neq k, j \rightarrow k \end{cases}$$

A cluster algebra is of finite type if it has only finitely many distinct cluster seeds. Otherwise, it is of infinite type. Cluster algebras admit an ADE classification.

THEOREM 1.2.5 (Theorem 1.4, [**FZ03**]). *Cluster algebras are of finite type if and only if their quiver is mutation equivalent to a Dynkin diagram of finite type with any orientation given to its edges*

For the simply-laced cases, this restricts our attention to ADE-type. We discuss some of the combinatorial ingredients for understanding cluster algebras of types A_n and D_n in Subsection 1.2.3.2. Beyond finite-type cluster algebras, the next simplest families are clusters arising from quivers of finite mutation type. These are cluster algebras with an underlying quiver that is mutation equivalent to only finitely many quivers. These are classified in [**FST12**]. Among the finite mutation type cluster algebras, we have types $\tilde{A}_n, \tilde{D}_n$, and others corresponding to triangulations of surfaces (See e.g. [**FST08**]), as well as types $\tilde{E}_6, \tilde{E}_7, \tilde{E}_8, E_6^{(1,1)}, E_7^{(1,1)}, E_8^{(1,1)}$ and two additional exceptional quivers.

For the purposes of understanding invariants of Legendrians via cluster theory, we briefly discuss the notion of a cluster variety introduced by Fock and Goncharov [**FG06a**,**FG06b**]. In short, cluster varieties are varieties constructed by gluing together algebraic tori (corresponding to cluster seeds) using maps corresponding to cluster mutation so that the coordinate ring of regular functions on an algebraic variety with a cluster structure is an (upper) cluster algebra. An abbreviated statement of the main theorem of Casals and Weng tells us that cluster varieties arise naturally from symplectic geometry. The class of Legendrians they consider arises from a combinatorial construction known as a (complete) grid plabic graph. This class includes all Legendrian links considered in this work.

THEOREM 1.2.6 (Theorem 1.1 [**CW22**]). *For λ a Legendrian arising from a complete grid plabic graph, the coordinate ring $\mathbb{C}[\mathcal{M}_1(\lambda)]$ is a cluster algebra. Moreover, there is an explicitly*

constructed Legendrian weave filling L of λ with intersection quiver and sheaf moduli $\mathcal{M}_1(L)$ giving the data of the initial seed.

Note that Casals and Weng obtain cluster X-variables using the microlocal monodromy functor described in Section 1.2.2. They also obtain cluster A-variables by studying relative homology cycles of the Legendrian weave, and computing microlocal merodromies along these relative cycles [**CW22**, Section 4]. Frozen variables correspond to marked points or homology cycles in a Legendrian weave filling L that do not bound embedded Lagrangian disks in the complement $\mathbb{D}^4 \backslash L$. In the setting of rainbow closures of positive braids, one can obtain a basis of embedded mutable homology cycles from a Legendrian weave and may generally avoid considering frozen variables, as we do in this work.

1.2.3.2. *Combinatorics of (tagged) triangulations.* The combinatorics of (tagged) triangulations of punctured surfaces play a key role in defining and understanding many of the simpler classes of cluster algebras. For A_n, D_n and $\tilde{D}_n$ type, many of the computations of cluster modular groups are most accessible through the combinatorics of tagged triangulations. More generally, we can define surface-type cluster algebras as cluster algebras whose underlying quiver comes from a (tagged) triangulation of a punctured surface. In this subsection, we discuss the necessary combinatorial ingredients for understanding tagged triangulations of surfaces in the context of cluster theory.

Let $\mathbb{S}_{n_1,\ldots,n_k}$ be a surface with n_i marked points on the ith boundary component. We allow for $n_i = 0$ and interpret this as a puncture. Following [**FST08**], we define (tagged) arcs, and (tagged) triangulations.

Definition 1.2.7. *An arc $\gamma \in \mathbb{S}_{n_1,\ldots,n_k}$ is a curve in $\mathbb{S}$ such that:*

- *the endpoints of γ lie on marked points;*
- *the interior of γ does not intersect itself;*
- *the interior of γ is disjoint from $\partial\mathbb{S}$ and marked points;*
- *γ does not cut out an unpunctured monogon or an unpunctured bigon.*

The last condition ensures that no arc is contractible to a point or into the boundary of $\mathbb{S}$. We consider arcs equivalent up to isotopy. Two (isotopy classes of) arcs are said to be compatible if there are two arcs in their respective isotopy classes that do not intersect in the interior of $\mathbb{S}_{n_1,\ldots,n_k}$.

A triangulation of $\mathbb{S}_{n_1,\ldots,n_k}$ is then a maximal pairwise compatible collection of (isotopy classes of) arcs.

From a triangulation $\mathcal{T}$, we can produce a cluster algebra as follows. For every edge γ_i in $\mathcal{T}$, we assign a vertex v_i in our quiver $Q_{\mathcal{T}}$. The vertices corresponding to boundary edges of $\mathcal{T}$ are declared to be frozen. We add an edge between v_i and v_j if there is a face where γ_j is counterclockwise from γ_i and at least one of v_i or v_j is mutable. Note that we must cancel any oriented two-cycles once we have accounted for all of the edges in this fashion. To each vertex in the quiver, we assign a cluster variable which can be thought of as measuring the length of the arc in an appropriate hyperbolic geometric context [**FT18**]. Mutation is given by exchanging one diagonal of a quadrilateral for the other.

As defined, triangulations of $\mathbb{S}_{n_1,\ldots,n_k}$ do not realize every possible cluster seed in the corresponding cluster algebra if the number of punctures is at least 1. This is because of the appearance of self-folded triangles, which produce arcs that cannot be mutated at. In order to represent all possible cluster seeds as a triangulation, we require additional decorations.

We arbitrarily divide an arc into two ends and allow each end to be either tagged or untagged. To produce a tagged triangulation, we introduce additional compatibility relations.

Definition 1.2.8 (Definition 7.4 [**FST08**]). *Two tagged arcs γ_1 and γ_2 in $\mathbb{S}_{n_1,\ldots,n_k}$ are compatible if:*

- *the untagged arcs corresponding to γ_1 and γ_2 are compatible;*
- *if the untagged arcs corresponding to γ_1 and γ_2 represent distinct isotopy classes and they share an endpoint a, then the tagging at the ends of γ_1 and γ_2 incident to a coincide;*
- *if the untagged arcs corresponding to γ_1 and γ_2 lie in the same isotopy class, then at least one end of γ_1 must be tagged in the same way as the same end of γ_2.*

In general, (tagged) triangulations of a disk with either 0, 1, or 2 punctures correspond to cluster algebras of type A_n, D_n, and $\tilde{D}_n$, respectively. To obtain a quiver from a tagged triangulation, we treat tagged arcs as normal arcs. For an arc γ sharing the same endpoints as γ', we use the face obtained by deleting γ' to compute the direction of the arrows in the quiver. See Figure 4.6 (left) for an example.

1.2.3.3. *Cluster modular groups.* Given a cluster variety, one can consider maps that act on the variety by permuting the cluster tori, thereby preserving the cluster structure. More explicitly, a cluster automorphism of a cluster algebra $\mathcal{A}$ is a permutation of the cluster variables of $\mathcal{A}$ that preserves cluster seeds and commutes with mutation. A result of Assem-Schiffler-Shramchenko implies that any cluster automorphism can be defined by the image of a single seed and necessarily preserves the underlying quiver up to a simultaneous change of orientation on all of the arrows [**ASS12**, Proposition 2.4]. Therefore, we have the following definition.

DEFINITION 1.2.9. *The (orientation preserving) cluster modular group $\mathcal{G}(\mathcal{A})$ of a cluster algebra $\mathcal{A}$ is the group of maps π permuting cluster variables and commuting with mutation such that the induced map on quivers $Q(\mathbf{x}) \to Q(\pi(\mathbf{x}))$ is an (orientation preserving) quiver automorphism.*

EXAMPLE 1.2.2. *For an A_2 cluster algebra, the cluster modular group $\mathcal{G}(A_2)$ is isomorphic to $\mathbb{Z}_5$ and is generated by a $2\pi/5$ rotation of the triangulation corresponding to the initial quiver. Viewing each diagonal $D_{i,j}$ of the triangulation as the image of a Plücker coordinate $\Delta_{i,j}$ in the top-dimensional positroid strata of the Grassmannian $Gr(2,5)$, we can see that this cluster automorphism is given by the map $\Delta_{i,j} \mapsto \Delta_{i-1,j-1}$ for all $1 \leq i < j \leq 5$.*

For all classes of cluster algebras discussed in this work, any cluster automorphism φ can be given as a finite sequence of mutations $\mu_{v_1}\mu_{v_2}\dots\mu_{v_m}$ and a permutation $\pi \in S_n$ of the quiver vertex labels. In order to avoid confusion, we fix the notation $\mu_{v_1}\mu_{v_2}\dots\mu_{v_m}$ as denoting a sequence of mutations starting with μ_{v_1} and ending with μ_{v_m}. When we need to specify the particular data of a cluster automorphism φ, we denote it by the tuple $\varphi = (\mu_{v_1}\mu_{v_2}\dots\mu_{v_m}, \pi)$ with the permutation π expressed in cycle notation. Following the conventions of [**ASS12**, **KG21**], we allow for cluster automorphisms φ defined solely as a permutation of the quiver vertex labels without any mutations.

EXAMPLE 1.2.3. *Consider an A_2 quiver with vertices labeled 1 and 2 and an edge from 1 to 2. The cluster automorphism $\varphi = (\mu_1, (12))$ also generates $\mathcal{G}(A_2)$ and corresponds to a rotation of an initial triangulation of the pentagon by $6\pi/5$.*

Cluster modular groups have been computed for finite, affine, and surface-type cluster algebras [**ASS12**]. The case of extended affine types was also computed in [**KG21**], building on work of

Chris Fraser investigating cluster modular groups of Grassmannians [**Fra18**]. Results relevant to this manuscript are summarized in Tables 4.1 and 4.2.

We highlight here the results related to the theory of mapping class groups and tagged triangulations, as they will reappear in crucial arguments below. We define the mapping class group $\mathrm{MCG}(\mathbb{S}_{n_1,\dots,n_k})$ of a surface $\mathbb{S}_{n_1,\dots,n_k}$ with n_i marked points on the ith boundary component as the group of orientation-preserving homeomorphisms of Σ fixing the set of punctures up to homeomorphisms isotopic to the identity. The tagged mapping class group $\mathrm{MCG}_{tag}(\mathbb{S}_{n_1,\dots,n_k})$ is defined to be the semidirect product of $\mathrm{MCG}(\mathbb{S}_{n_1,\dots,n_k})$ with $\mathbb{Z}_2^p$ where p is the number of punctures of Σ. The product structure is given by the action of simultaneously swapping tags at all arcs incident to a particular punctures.

Denote by $\mathcal{A}(\mathbb{S}_{n_1,\dots,n_k})$ the surface-type cluster algebra associated to $\mathbb{S}_{n_1,\dots,n_k}$. The following theorem relates the tagged mapping class group of a surface and the cluster modular group of the associated cluster algebra.

THEOREM 1.2.7 (Proposition 8.5, [**BS15**]). *Assume that $\mathbb{S}_{n_1,\dots,n_k}$ is not a once or twice-punctured disk with four or fewer marked points on the boundary. Then $\mathrm{MCG}_{tag}(\mathbb{S}_{n_1,\dots n_k}) \cong \mathcal{G}(\mathcal{A}(\mathbb{S}_{n_1,\dots,n_k}))$.*

We introduce the following notion in order to compare cluster automorphisms defined on different initial seeds.

DEFINITION 1.2.10. *Two cluster automorphisms φ_1 and φ_2 are conjugate if they act identically on the set of cluster charts.*

In Subsection 4.1.2.3 we will use the combinatorics of tagged triangulations to show that cluster automorphisms induced by Legendrian loops are conjugate to cluster automorphisms coming directly from quiver combinatorics.

1.2.4. The Legendrian contact DGA and the augmentation variety. In this subsection, we describe the Legendrian contact DGA, a Floer-theoretic invariant of Legendrian knots and their exact Lagrangian fillings. We first give the Ekholm-Honda-Kálmán construction for exact Lagrangian cobordisms and then describe the necessary Floer-theoretic background in the context of understanding exact Lagrangian fillings of $\lambda(A_n)$.

1.2.4.1. *The pinching cobordism and pinching sequence fillings.* The following definition gives a condition for being able to perform a local move resolving a crossing in the Lagrangian projection. We refer to such a cobordism as a pinching cobordism See also [**EHK16**, Definition 6.2] and [**CN21**, Section 2.1].

DEFINITION 1.2.11. *A crossing in the Lagrangian projection* $\Pi(\lambda)$ *of a Legendrian* λ *is contractible if there is a Legendrian isotopy of* λ *inducing a planar isotopy of* $\Pi(\lambda)$ *making the length of the corresponding Reeb chord arbitrarily small.*

We now describe the precise topological construction of the elementary cobordisms defining pinching sequence fillings, i.e. exact Lagrangian fillings of built out of pinching cobordisms and Legendrian isotopies. Consider a neighborhood of a contractible crossing depicted in the Lagrangian projection. Attaching a 1-handle at the crossing yields an exact Lagrangian cobordism in the symplectization $(\mathbb{R}_t \times \mathbb{R}^3, d(e^t(dz - ydx)))$ [**EHK16**, Section 6.5]. In the Lagrangian projection, this 1-handle attachment is diagrammatically given as a 0-resolution of the crossing, as depicted in Figure 1.16. If λ is the rainbow closure of a positive braid, as is the case for $\lambda(A_{n-1})$, then every crossing of the braid is contractible [**CN21**, Proposition 2.8].

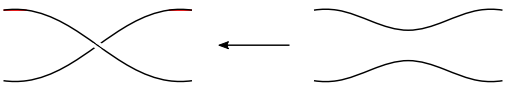

FIGURE 1.16. A local model of a pinching cobordism as a 0-resolution of a contractible crossing in the Lagrangian projection. The direction of the arrow indicates a cobordism from the concave end to the convex end.

Let us consider $\lambda \subseteq (\mathbb{R}^3, \xi_{\text{st}})$ and its front projection $\Pi(\lambda)$. In order to describe a pinching cobordism in terms of a projection of λ, we introduce the Ng resolution. This is a Legendrian isotopy λ_t such that $\lambda_0 = \lambda$ and the Lagrangian projection $\pi(\lambda_1)$ can be obtained from the front projection $\Pi(\lambda_1)$ by smoothing all left cusps and replacing all right cusps with small loops [**Ng03**]. See Figure 1.17 for an example. A pinching cobordism in the front projection of the link $\lambda(\beta)$ is then given by first taking the Ng resolution of $\lambda(\beta)$, performing a 0-resolution at a crossing in the Lagrangian projection as specified above, and then undoing the Ng resolution.

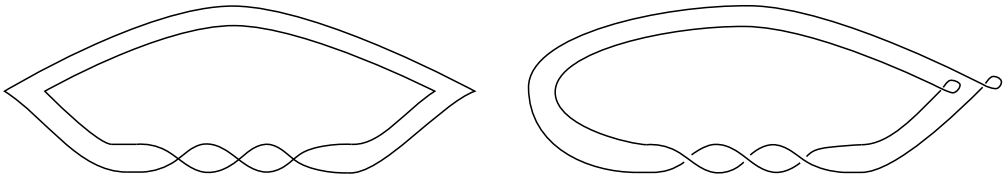

FIGURE 1.17. A front projection of $\lambda(A_2)$ (left) and its Ng resolution (right). The three leftmost crossings are contractible.

Given that $\lambda(A_{n-1})$ has n crossings, a pinching sequence filling can be characterized by a permutation σ in S_n. Such a permutation specifies an order in which to apply these elementary cobordisms to the n contractible crossings in the Ng resolution of $\lambda(A_{n-1})$. Given a permutation σ, we will denote it in one line notation $\sigma(1)\dots\sigma(n)$. If σ is of the form $\sigma(1)\dots i\, j\dots k\dots\sigma(n)$ for $i > k > j$, the permutation $\sigma' = \sigma(1)\dots j\, i\dots k\dots\sigma(n)$ obtained by interchanging i and j gives an order of resolving crossings that yields the same Floer-theoretic invariant[1] [**Pan17**]. This leads us to consider only a subset of permutations in S_n.

DEFINITION 1.2.12. *A 312-avoiding permutation is a permutation $\sigma \in S_n$ such that any triple of letters i, j, k appearing in order in σ does not satisfy the inequality $i > k > j$.*

Distinct 312-avoiding permutations yield distinct Hamiltonian isotopy invariants of exact Lagrangian fillings, i.e. restricting the indexing set from S_n to 312-avoiding permutations yields the existence of at least a Catalan number of fillings of $\lambda(A_{n-1})$ up to Hamiltonian isotopy [**Pan17**, Theorem 1.1].

1.2.4.2. *The Legendrian contact DGA.* For a Legendrian link λ, the Legendrian contact differential graded algebra (DGA) is a powerful Floer-theoretic invariant of λ [**Che02**]. We denote the DGA of λ by $\mathcal{A}(\lambda; R)$, or by $\mathcal{A}(\lambda)$ when we wish to suppress the dependence on the coefficient ring R. In our description below we will generally take R to be $\mathbb{Z}[H_1(\lambda)]$ or $\mathbb{Z}[H_1(L)]$ for an exact

[1]The two fillings corresponding to σ and σ' yield identical augmentations ϵ_σ and $\epsilon_{\sigma'}$ of the DGA $\mathcal{A}(\lambda(A_{n-1}))$. This is because the presence of the crossing labeled j prevents the existence of any holomorphic strip with positive punctures occurring at both crossings i and k. Therefore, resolving crossing k (resp. i) has no effect on the generator z_i (resp. z_k) in the DGA $\mathcal{A}(\lambda(A_{n-1}))$.

Lagrangian cobordism L. We refer the interested reader to [**EN19**] for a more general introduction to the Legendrian contact DGA and [**CN21**, Section 3] for a discussion of different choices of coefficient rings.

To describe the algebra $\mathcal{A}(\lambda; R)$, we introduce the auxiliary data of a set of marked points on λ and a corresponding set of capping paths. We label a marked point on the component λ_a of λ by $t_a^{\pm 1}$. In order for the grading to be well defined, we require one marked point for every component of λ and that the oriented tangent vectors to $\Pi(\lambda) \subseteq \mathbb{R}^2_{xy}$ at the marked points $t_1, \dots, t_m$ are all parallel. Note that the collection $t_1, \dots t_m$ of marked points can be thought of as encoding $H_1(\lambda)$. Given a Reeb chord z with ends z_a and z_b lying on components λ_a and λ_b, a capping path γ_z is the concatenation of paths following the orientation of λ from z_a to t_a and t_b to z_b. Here we require that z_a corresponds to the undercrossing of z in the Lagrangian projection. The data of the DGA is then given as follows.

Generators: For a knot λ, the Legendrian contact DGA is freely generated over the Laurent polynomial ring $R = \mathbb{Z}[t_1^{\pm 1}, \dots t_m^{\pm 1}]$ by the Reeb chords of $\lambda \subseteq (\mathbb{R}^3, \xi_{\text{st}})$. In the Lagrangian projection, we can equivalently think of these generators as the crossings of λ.

Grading: We restrict our attention to the case of Legendrian links where each component has rotation number zero, as is the case for every link considered in this manuscript. In this setting, each t_i and t_i^{-1} is assigned grading 0. We define the grading for a Reeb chord z as follows. As we traverse the capping path γ_z, the unit tangent vector to $\Pi(\lambda)$ makes a number of counterclockwise revolutions. We can perturb λ in such a way that the tangent vectors at a crossing of $\Pi(\lambda)$ are always orthogonal and the number $r(\gamma_z)$ of such revolutions is always an odd multiple of $\frac{1}{4}$. The grading of z is then defined to be $|z| := 2r(\gamma_z) - \frac{1}{2}$. Grading is extended to products of generators $|yz|$ additively by $|yz| = |y| + |z|$.

In the case of rainbow closures of a positive braid β, every Reeb chord that corresponds to a crossing of β in the Ng resolution has degree zero while the remaining Reeb chords at the right of the diagram have degree one.

Differential: The differential is given by counts of certain holomorphic disks in the following way. We first decorate each quadrant of a crossing of $\Pi(\lambda)$ with two signs, a Reeb sign and an orientation sign. The Reeb sign is specified as pictured in Figure 1.18 (left), where opposite

quadrants have the same sign and adjacent quadrants have different signs. The orientation sign is given as in Figure 1.18 (right), where the shaded regions are decorated with orientation sign $-$ and unshaded regions are decorated with orientation sign $+$.

The differential considers immersions u from a punctured disk into $\mathbb{R}^2$ with boundary punctures on $\Pi(\lambda)$ up to reparametrization. We refer to any puncture appearing at a quadrant with a positive (resp. negative) Reeb sign as a positive (resp. negative) puncture. We restrict to immersions that have a single positive puncture and arbitrarily many negative punctures. For any such immersion u, denote by $w(u)$ the product of generators given by the negative boundary punctures. If the boundary of u passes through any marked point t_i, then we obtain $w'(u)$ as the product of $w(u)$ by $t_i^{\pm 1}$. The power is assigned according to whether the orientation of u at the relevant marked point agrees (t_i) with the orientation of λ or does not agree (t_i^{-1}) with the orientation of λ. To each disk u, we also assign the quantity $\mathrm{sgn}(u) \in \{\pm 1\}$ given by the product of the orientation signs appearing at boundary punctures of u. The differential at z is then given by

$$\partial(z) = \sum \mathrm{sgn}(u) w'(u)$$

where the sum is taken over all immersed disks u with a single positive puncture at z. We extend the differential to products $z_1 z_2$ by the Leibniz rule $\partial(z_1 z_2) = \partial(z_1) z_2 + (-1)^{|z_1|} z_1 \partial(z_2)$

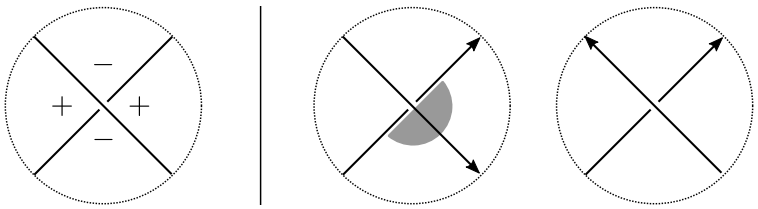

FIGURE 1.18. Reeb signs (left) and orientation signs (right) at a crossing of $\Pi(\lambda)$. The quadrants shaded dark gray carry negative orientation signs, while the unshaded quadrants are positive.

EXAMPLE 1.2.4. *The DGA of $\lambda(A_2)$ is freely commutatively generated over $\mathbb{Z}[t, t^{-1}]$ by generators a_1, a_2, z_1, z_2, z_3, labeled in Figure 1.19. The gradings are given by $|a_1| = |a_2| = 1$ and $|z_i| = |t_1| = 0$. The differential on generators a_i is given by*

$$\partial(a_i) = \begin{cases} z_1 + z_3 + z_1 z_2 z_3 + t_1^{-1} & i = 1 \\ 1 + t_1 + z_2 + t_1 z_3 z_2 + t_1 z_1 z_2 + t_1 z_1 z_2 z_3 z_2 & i = 2 \end{cases}$$

The differential on the remaining generators vanishes for degree reasons. Note that setting $t_1 = -1$ *implies* $\partial(a_2) = -z_2 \partial(a_1)$ □

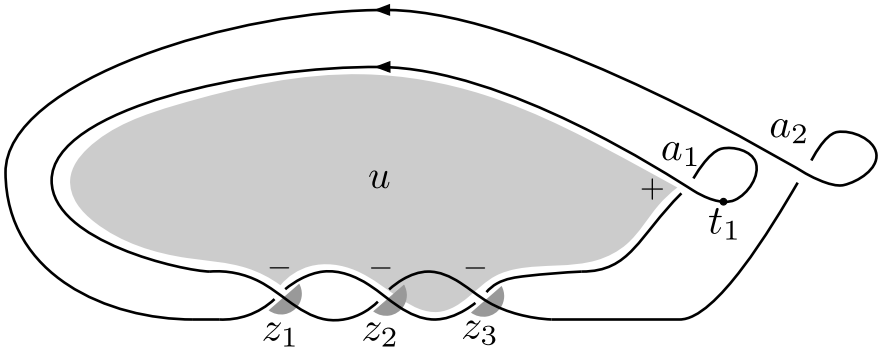

FIGURE 1.19. The Lagrangian projection of the Legendrian trefoil decorated with Reeb signs and orientation signs. The light gray disk labeled u has a positive puncture at a_1 and a single negative puncture at z_1. Thus, it corresponds to the term z_1 appearing in $\partial(a_1)$.

1.2.4.3. *Augmentations.* The Legendrian contact DGA can be difficult to extract information from, so it is often useful to consider augmentations of the DGA. Augmentations are DGA maps from $\mathcal{A}(\lambda)$ to some ground ring. Here we consider the ground ring of Laurent polynomials in $n-1$ variables with coefficients in $\mathbb{Z}$, understood as a DGA with trivial differential and concentrated in degree 0. In this subsection, we define augmentations and the related Legendrian isotopy invariant, the augmentation variety.

Augmentations of $\mathcal{A}(\lambda)$ are intimately tied to exact Lagrangian fillings of λ. This relationship can be understood through the functoriality of the DGA with respect to exact Lagrangian cobordisms. More precisely, Ekholm-Honda-Kálmán show that an exact Lagrangian cobordism L from λ_- to λ_+ induces a DGA map Φ_L from $\mathcal{A}(\lambda_+; \mathbb{Z}_2)$ to $\mathcal{A}(\lambda_-; \mathbb{Z}_2)$ [**EHK16**, Theorem 1.2]. Their result was upgraded to make use of $\mathbb{Z}_2[H_1(L)]$ coefficients with an appropriate choice of marked points encoding $H_1(L)$ by Pan [**Pan17**, Proposition 2.6]. Pan's use of $H_1(L)$ coefficients is crucial

for her ability to distinguish C_n Lagrangian fillings of $\lambda(A_{n-1})$, as Ekholm-Honda-Kálmán are only able to identify $(2^{n+1} - (-1)^{n+1})/3$ distinct Lagrangian fillings working over $\mathbb{Z}_2$ [**EHK16**, Theorem 1.6]. The following result of Karlsson further improves Pan's coefficient ring to consider the augmentations over $\mathbb{Z}[H_1(L)]$.

PROPOSITION 1.2.1. *[**Kar20**, Theorem 2.5] An exact Lagrangian cobordism L from λ_- to λ_+ induces a DGA map $\Phi_L : \mathcal{A}(\lambda_+; \mathbb{Z}[H_1(\lambda_+)]) \to \mathcal{A}(\lambda_-; \mathbb{Z}[H_1(L)])$.*

See also [**CN21**, Section 3.3] for a discussion on Karlsson's choice of signs, as well as a geometric understanding of the induced map. As a result of Proposition 1.2.1, we can think of an augmentation of λ as a map induced from $\mathcal{A}(\lambda; \mathbb{Z}[H_1(\lambda)])$ to the DGA of the empty set induced by an exact Lagrangian filling of λ.

DEFINITION 1.2.13. *An augmentation ϵ_L induced by a Lagrangian filling L of λ is a DGA map*

$$\epsilon_L : \mathcal{A}(\lambda, \mathbb{Z}[H_1(\lambda)]) \to \mathbb{Z}[H_1(L)]$$

where we think of $\mathbb{Z}[H_1(L)]$ as a DGA concentrated in degree zero with trivial differential.

The functoriality of the DGA motivates the study of augmentations of $\mathcal{A}(\lambda)$ in order to better understand Lagrangian fillings of λ. The space of all augmentations of $\mathcal{A}(\lambda)$, denoted by $\mathrm{Aug}(\lambda)$, is an invariant of λ. In the case where λ is the rainbow closure of a positive braid, $\mathrm{Aug}(\lambda)$ has the structure of an affine algebraic variety and is known as the augmentation variety. We may tensor our coefficients ring with $\mathbb{C}$ in order to consider augmentations over a field.[2] When the grading of $\mathcal{A}(\lambda)$ is concentrated in non-negative degrees, as is the case for rainbow closures of positive braids, then $\mathrm{Aug}(\lambda) \cong \mathrm{Spec}\, H_0(\mathcal{A}(\lambda))$, see e.g. [**GSW20a**, Corollary 2.9]. Since Spec is contravariant, ϵ_L induces a map $\mathrm{Spec}(\mathbb{C}[s_1^{\pm 1}, \ldots, s_{b_1(L)}^{\pm 1}]) \to \mathrm{Spec}\, H_0(\mathcal{A}(\lambda; \mathbb{C}[H_1(\lambda)]))$, where we have identified the ground ring of Laurent polynomials with complex coefficients $\mathbb{C}[H_1(L)]$ with the group ring $\mathbb{C}[s_1^{\pm 1}, \ldots, s_{b_1(L)}^{\pm 1}]$. We interpret this map as the inclusion of a toric chart $(\mathbb{C}^\times)^{b_1(L)}$ into the augmentation variety

$$\mathrm{Spec}(\mathbb{C}[s_1^{\pm 1}, \ldots, s_{b_1(L)}^{\pm 1}]) \cong (\mathbb{C}^\times)^{b_1(L)} \hookrightarrow \mathrm{Aug}(\lambda).$$

[2]To clarify, complexifying is solely for the purpose of simplifying the algebro-geometric discussion in this paragraph. For all other purposes relating to computations with the DGA, we will continue to use integer coefficients.

The image of degree-zero generators under an augmentation give local coordinate functions on the corresponding toric chart. In order to describe these local coordinate functions, we discuss Pan's explicit computation of induced DGA maps in the context of Lagrangian fillings of $\lambda(A_{n-1})$ with a lift to $\mathbb{Z}[H_1(L)]$ following [**CN21**, Section 4.2]. For a pinching cobordism, the induced map is given by a certain count of holomorphic disks, similar to the differential. The homology coefficients are determined by the intersection of these disks with relative homology classes in $H_1(L, \lambda_- \sqcup \lambda_+)$, which is identified with $H_1(L)$ via Poincaré duality.

Pan describes a set of generators for $H_1(L, \lambda_- \sqcup \lambda_+)$ for a sequence of pinching cobordisms. In this setting, a relative homology cycle $\gamma_{\sigma(i)}$ starts from the saddle point originally labeled $z_{\sigma(i)}$ and extends downwards to λ_- where it meets the boundary in $s_{\sigma(i)}$ and $s_{\sigma(i)}^{-1}$. In order to consider signs, we orient this relative cycle so that the two halves of the cycle are labeled by $s_{\sigma(i)}$ and $-s_{\sigma(i)}^{-1}$, as in Figure 1.20. In a slicing of the symplectization, $\gamma_{\sigma(i)}$ meets the Lagrangian projection of $\lambda(A_k)$ in two points labeled $s_{\sigma(i)}$ and $-s_{\sigma(i)}^{-1}$, so that in practice, these generators reduce the computation of the coefficients to a combinatorial count of marked points.

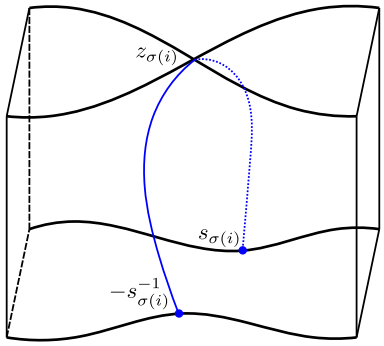

FIGURE 1.20. A local model of a relative cycle encoding the homology of a pinching cobordism L. At the top of the figure, the length of the Reeb chord $z_{\sigma(i)}$ is 0, i.e. the two strands intersect at the point $z_{\sigma(i)}$. The bottom of the figure depicts a 0-resolution of the crossing where marked points labeled $-s_{\sigma(i)}^{-1}$ and $s_{\sigma(i)}$ encode $H_1(L)$.

Given a pinching cobordism $L_{\sigma(i)}$ at the Reeb chord $z_{\sigma(i)}$ as part of a Lagrangian filling L_σ, the induced map Φ_i on the generator z_j is computed as a sum over all immersed disks with positive punctures at both $z_{\sigma(i)}$ and z_j. As before, we denote by $w'(u)$ the product of negative punctures of

the immersed disk u and intersections of u with marked points counted with orientation. Likewise, $\mathrm{sgn}(u)$ is given by the product of the orientation signs appearing at boundary punctures of u.

DEFINITION 1.2.14. *The DGA map Φ_i induced by a pinching cobordism at the Reeb chord $z_{\sigma(i)}$ is given by*

$$\Phi_i(z_j) = z_j + \sum \mathrm{sgn}(u) w'(u).$$

The Reeb chord $z_{\sigma(i)}$ is sent to $s_{\sigma(i)}$ by Φ_i.

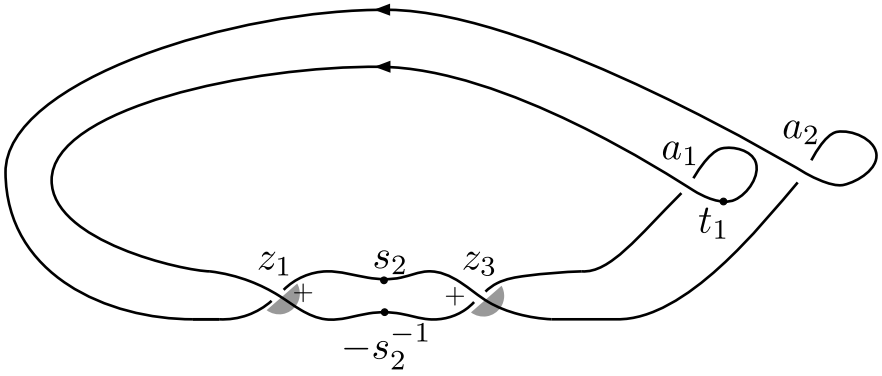

FIGURE 1.21. A Legendrian Hopf link obtained from the knot pictured in Figure 1.19 by pinching at the Reeb chord labeled by z_2.

EXAMPLE 1.2.5. *Pinching at the Reeb chord labeled by z_2 of the Legendrian trefoil pictured in Figure 1.19 yields the Legendrian Hopf link, pictured in Figure 1.21 with the addition of marked points s_2 and $-s_2^{-1}$. The induced map Φ_1 on the DGA is given by*

$$z_1 \mapsto z_1 - s_2^{-1}, \qquad z_2 \mapsto s_2, \qquad z_3 \mapsto z_3 - s_2^{-1}.$$

Performing another pinch at z_1 induces the map Φ_2 given by

$$z_1 \mapsto s_1, \qquad s_2 \mapsto s_2, \qquad z_3 \mapsto z_3 + s_1^{-1} s_2^{-2}.$$

The map on z_3 is determined by the disk with positive punctures at z_1 and z_3 passing through s_2 and $-s_2^{-1}$. □

Pan gives a purely combinatorial description of the map Φ_i induced by opening the crossing labeled $z_{\sigma(i)}$ [**Pan17**, Definition 3.2]. First, define the set

$$T^i_\sigma = \{j \in \{1,\dots,n\} | \sigma^{-1}(j) > \sigma^{-1}(i) \text{ and if } i<k<j \text{ or } j<k<i, \text{ then } \sigma^{-1}(k) < \sigma^{-1}(i)\}.$$

For $j \in T^{\sigma(i)}_\sigma$ and $1 \leq i \leq n$, the DGA map is given by

$$\Phi_i(z_j) = z_j + s^{-1}_{\sigma(i)} \prod_{\substack{j<k<\sigma(i) \text{ or} \\ \sigma(i)<k<j}} s_k^{-2}$$

and for $j = \sigma(i)$, $\Phi_i(z_j) = s_j$. Otherwise, we take Φ_i to be the identity.

LEMMA 1.2.1. *The lift of Pan's combinatorial formula for Φ_i to $\mathbb{Z}[H_1(L)]$ is given by*

$$\Phi_i(z_j) = z_j + (-1)^{|j-\sigma(i)|+1} s^{-1}_{\sigma(i)} \prod_{\substack{j<k<\sigma(i) \ or \\ \sigma(i)<k<j}} s_k^{-2}.$$

PROOF. To upgrade Pan's formula to $\mathbb{Z}[H_1(L)]$ coefficients, we note that the number of pairs of marked points that appear between z_j and $z_{\sigma(i)}$ is precisely $|j-\sigma(i)|$. Since each pair of marked points contributes a -1 factor to $\mathrm{sgn}(u)$ and the disk u with positive punctures at z_j and $z_{\sigma(i)}$ picks up an additional -1 factor from the orientation sign of the leftmost positive puncture, we arrive at the formula given above. □

To compute an augmentation of $\mathcal{A}(\lambda; \mathbb{Z}[T])$, we also need to describe the map induced by the minimum cobordism. This minimum cobordism is given by filling a standard Legendrian unknot λ_U with an exact Lagrangian disk. It induces a map Φ_{min} sending the Reeb-chord generator of $\mathcal{A}(\lambda_U; \mathbb{Z}[T])$ to 0. The map on the marked point generators can be deduced from the fact that Φ_{min} is a DGA map and therefore we must have $\Phi_{min} \circ \partial = 0$. In the context of a filling L_σ of $\lambda(A_{n-1})$, this tells us that $s_1 \dots s_n + t_1 = 0$. For knots, we also obtain $-(s_1 \dots s_n)^{-1} + 1 = 0$, implying that t_1 is mapped to -1. For links, we have $(s_1 \dots s_n)^{-1} + t_2 = 0$, implying only that $t_1 t_2 = 1$. Pan avoids this ambiguity by computing augmentations of $\mathcal{A}(\lambda(A_{2n-1}; H_1(\lambda(A_{2n-1}))$ induced by fillings L_σ of $\lambda(A_n)$ where $\sigma(1) = n+1$. This is equivalent to setting $t_1 = -1$, from which we obtain $t_2 = -1$.

Therefore, the DGA map induced by the minimum cobordism is given by the following formula.

$$\Phi_{min}(s_n) = s_1^{-1} \dots s_{n-1}^{-1}$$

For $s_i \neq s_n$, Φ_{min} is the identity map. For marked points t_i, we have $\Phi_{min}(t_i) = -1$.

Together, the maps Φ_i and Φ_{min} give us the ingredients to define the augmentation induced by a pinching sequence filling of $\lambda(A_{n-1})$.

DEFINITION 1.2.15. *The augmentation ϵ_{L_σ} induced by the Lagrangian filling L_σ of $\lambda(A_{n-1})$ is given by the DGA map*

$$\epsilon_{L_\sigma} := \Phi_{min} \circ \Phi_n \circ \dots \circ \Phi_1$$

To simplify our computations involving the DGA and the Kálmán loop, we will always set $t_1 = -1$ and $t_2 = -1$ for the remainder of this manuscript. By our definition of ϵ_{L_σ}, this does not affect the augmentation induced by L_σ.

1.2.4.4. *Braid matrices.* For $\lambda(A_{n-1})$, the polynomials defining the augmentation variety have a combinatorial description as a specific entry in a product of matrices. These matrices originally appeared in [**Kál06**] as a means for encoding the immersed disks contributing to the differential. More recently, they were used in [**CGGS20**] to give a holomorphic symplectic structure on the augmentation variety. We adopt the conventions of [**CGGS20**] in defining the braid matrix.

DEFINITION 1.2.16. *The braid matrix $B(z_i)$ is given by* $B(z_i) := \begin{pmatrix} 0 & 1 \\ 1 & z_i \end{pmatrix}$.

Intuitively, one can think of the matrix $B(z_i)$ as encoding whether or not an immersed disk has a negative puncture at the crossing labeled by z_i. A product of braid matrices can be used to compute the differential of $\mathcal{A}(\lambda(A_{n-1}; \mathbb{Z})$ as follows. First, label the crossings of $\lambda(A_{n-1})$ by $a_1, a_2, z_1, \dots z_n$, as in Figure 1.19. From [**Kál06**, Section 3.1], we have that $\partial(a_1) = -1 + [\prod_{i=1}^{n} B(z_i)]_{(2,2)}$ where the subscript denotes the $(2,2)$ entry of the product and the -1 appears due to our choice of conventions. See also [**CN21**, Proposition 5.2] for a similar computation. An analogous computation to the case of $\lambda(A_2)$ implies that the differential of a_2 is given by $\partial(a_2) = -\Delta_{2,n+1}(\partial(a_1))$.

The computation of the differential via braid matrices also allows us to express the augmentation variety $\mathrm{Aug}(\lambda(A_{n-1}))$ in a similar manner. As augmentations are DGA maps, they satisfy $\epsilon \circ \partial =$

$\partial \circ \epsilon$. Since ϵ respects the grading, it must vanish on generators of nonzero degree, implying that for such generators a, we have $\partial \circ \epsilon(a) = 0$. Therefore, any augmentation ϵ of $\mathcal{A}(\lambda(A_{n-1}))$ satisfies $\epsilon \circ \partial(a_1) = \partial(a_2) = 0$. Since $\partial(a_2)$ is a multiple of $\partial(a_1)$, the vanishing of $\partial(a_1)$ is both necessary and sufficient to satisfy the vanishing condition for $\epsilon \circ \partial$, so that the augmentation variety is cut out by the vanishing of the equation $\partial(a_1) = 0$.

LEMMA 1.2.2. *The augmentation variety* $\mathrm{Aug}(\lambda(A_{n-1}))$ *is the zero set of the polynomial*

$$X_n := -1 + \left[\prod_{i=1}^{n} B(z_i)\right]_{(2,2)}$$

In addition to computing the augmentation variety, braid matrices also define regular functions $\Delta_{i,j}$ on $\mathrm{Aug}(\lambda(A_{n-1}))$ that will play an important role in Section 2.3.

DEFINITION 1.2.17. *The regular function* $\Delta_{i,j} \in \mathbb{Z}[X_n]$ *is given by* $\Delta_{i,j} := \left[\prod_{k=i}^{j-2} B(z_k)\right]_{(2,2)}$.

We specify the value of $\Delta_{i,i+1}$ to be 1. We collect some useful identities relating the $\Delta_{i,j}$ functions to the theory of continuants below.

EXAMPLE 1.2.6. *Consider the Legendrian trefoil,* $\lambda(A_2)$. *The augmentation variety* $\mathrm{Aug}(\lambda(A_2))$ *is the zero set of the polynomial* $X_3 = -1 + z_1 + z_3 + z_1 z_2 z_3$. *The regular functions* $\Delta_{i,j}$ *are of the form* $\Delta_{i,i+2} = z_i$ *or* $\Delta_{i,i+3} = 1 + z_i z_{i+1}$, *for* $1 \leq i \leq 3$. □

1.2.4.5. *Continuants.* Continuants are a family of polynomials $K_n(x_1, \ldots x_n)$ studied by Euler in his work on continued fractions [**Eul64**]. Continuants are defined by the following recursive formula:

$$K_n(x_1, \ldots x_n) = x_1 K_{n-1}(x_2, \ldots x_n) + K_{n-2}(x_3, \ldots x_n)$$

$K_0() = 1, K_1(x_1) = x_1$.

As mentioned above, the regular functions $\Delta_{i,j}$ are related to continuants.

LEMMA 1.2.3. *Let* $n = j - 2 - i$ *and* $x_k = z_{i+k-1}$. *Then*

$$K_n(x_1, \ldots x_n) = \Delta_{i,j}.$$

PROOF. The following is a classical property of continuants (see e.g. [**Fra49**, Section 1]) that allows us to understand the defining recursion relation in terms of braid matrices.

$$\begin{pmatrix} K_{n-2}(x_2,\dots,x_{n-1}) & K_{n-1}(x_2,\dots,x_n) \\ K_{n-1}(x_1,\dots x_{n-1}) & K_n(x_1,\dots,x_n) \end{pmatrix} = B(x_1)\dots B(x_n).$$

This follows inductively from applying the recursion relation in computing the matrix product

$$\begin{pmatrix} 0 & 1 \\ 1 & x_1 \end{pmatrix} \begin{pmatrix} K_{n-3}(x_3,\dots,x_{n-1}) & K_{n-2}(x_3,\dots,x_n) \\ K_{n-2}(x_2,\dots x_{n-1}) & K_{n-1}(x_2,\dots,x_n) \end{pmatrix} = \begin{pmatrix} K_{n-2}(x_2,\dots,x_{n-1}) & K_{n-1}(x_2,\dots,x_n) \\ K_{n-1}(x_1,\dots x_{n-1}) & K_n(x_1,\dots,x_n) \end{pmatrix}.$$

Therefore,

$$K_n(x_1,\dots,x_n) = \left[\prod_{k=1}^{n} B(x_k)\right]_{(2,2)}$$

Replacing x_k with z_{i+k-1} yields the desired identification. □

As a consequence, we obtain the continuant recursion relation in the context of the $\Delta_{i,j}$ functions.

$$\Delta_{i,j} = z_i\Delta_{i+1,j} + \Delta_{i+2,j} \tag{1.1}$$

Continuants satisfy several identities, the most general of which is Euler's identity for continuants. We present this identity in the context of the $\Delta_{i,j}$ functions:

$$\Delta_{1,\mu+\nu+2}\Delta_{\mu+1,\mu+\kappa+2} - \Delta_{1,\mu+\kappa+2}\Delta_{\mu+1,\mu+\nu+2} = (-1)^{\kappa+1}\Delta_{1,\mu+1}\Delta_{\mu+\kappa+2,\mu+\nu+2}$$

for $\mu \geq 1, \kappa \geq 0, \nu \geq \kappa+1$ [**Ust06**]. We require a special case of this identity for our algebraic proof of Theorem 1.1.3. Namely, when $\mu = 1, \kappa = k-3 \geq 0, \nu = n-1 \geq k-2$, we obtain

$$\Delta_{1,n+2}\Delta_{2,k} - \Delta_{1,k}\Delta_{2,n+2} = (-1)^k\Delta_{1,2}\Delta_{k,n+2}. \tag{1.2}$$

1.2.4.6. *The Kálmán loop.* In [**Kál05**], Kálmán defined a geometric operation on Legendrian torus links that induces an action on their exact Lagrangian fillings. In the case of $\lambda(A_{n-1})$, this operation consists of a Legendrian isotopy that is visualized by dragging the leftmost crossing clockwise around the link until it becomes the rightmost crossing. The graph of this isotopy is an

exact Lagrangian cylinder in the symplectization of $(\mathbb{R}^3, \xi_{\text{st}})$. Concatenating this cylinder with a Lagrangian filling L of $\lambda(A_{n-1})$ yields another filling, generally not Hamiltonian isotopic to L. As computed in [**Kál05**, Proposition 9.1], this induces a map on the DGA $\mathcal{A}(\lambda(A_{n-1}); \mathbb{Z}_2)$, which in turn induces an automorphism ϑ on the augmentation variety $\text{Aug}(\lambda(A_{n-1}))$. Following [**CN21**, Section 3], we can compute this induced action with integer coefficients. The additional information of this integral lift consists solely of a choice of signs for terms in the image of ϑ, as can be seen in Kálmán's example computation over $\mathbb{Z}[t, t^{-1}]$ in the case of the $\lambda(A_2)$ [**Kál05**, Section 5].[3] The map ϑ on generators z_i is then given by

$$\vartheta(z_i) = \begin{cases} -\Delta_{2,n+2} & i = 1 \\ z_{i-1} & 2 \leq i \leq n \end{cases}$$

In the $\Delta_{i,j}$ functions, this is expressed as $\vartheta(\Delta_{i,j}) = \Delta_{i-1,j-1}$ for $i > 1$, and

$$\vartheta(\Delta_{1,j}) = -\left[B(\Delta_{2,n+2}) \prod_{i=1}^{j-3} B(z_i) \right]_{(2,2)}.$$

[3]Note that Kálmán uses a different choice of sign conventions than Casals and Ng. By [**CN21**, Proposition 3.14], these different sign conventions yield equivalent induced augmentations.

CHAPTER 2

Lagrangian fillings in A-type and their Kálmán loop orbits

2.1. Isotopies of exact Lagrangian Cobordisms and Kálmán loop orbits of $\lambda(A_n)$

In this chapter we investigate exact Lagrangian fillings of Legendrian torus links $\lambda(2,n) \cong \lambda(A_{n-1})$ and their behavior under the action of the Kálmán loop. We start with a proof of Theorem 1.1.1 and use that to relating Legendrian weave fillings to exact Lagrangian fillings obtained via a sequence of elementary cobordisms. In Section 2.3, we give an alternative proof of the orbital structure of the induced action of the Kálmán loop on the Legendrian contact DGA. The classically known Euler's identity for continuants makes a key appearance in establishing the required algebraic description of the induced action. Finally, Section 2.4 contains combinatorial characterizations of exact Lagrangian fillings with a given orbit size and explains how to realize the Kálmán loop action in terms of both Legendrian weaves and 312-avoiding permutations.

2.2. Proof of Theorems 1.1.1 and 1.1.3

In this section we prove that a pinching sequence filling is Hamiltonian isotopic to a corresponding Legendrian weave filling. We first relate the elementary cobordisms used to construct these fillings.

PROPOSITION 2.2.1. *The pinching cobordism and D_4^- cobordism described below are Hamiltonian isotopic relative to their boundaries.*

We prove this by giving a local model for the D_4^- cobordism as a sequence of diagrams in both the front and Lagrangian projections and then describing an exact Lagrangian isotopy between the two cobordisms that fixes the boundary. Since compactly supported Lagrangian isotopy is equivalent to Hamiltonian isotopy [**FOOO09**, Theorem 3.6.7], this implies the proposition. We then use Proposition 2.2.1 to prove Theorem 1.1.1 in the general case of Lagrangian fillings of $\lambda(\beta)$.

To relate pinching sequence fillings and weave fillings of $\lambda(A_{n-1})$, we will describe a combinatorial bijection between 312-avoiding permutations σ and triangulations $\mathcal{T}_\sigma$ dual to Legendrian weaves. Denote by L_σ or $L_{\mathcal{T}_\sigma}$ the exact Lagrangian filling of $\lambda(A_{n-1})$ defined by the given combinatorial input. In the specific case of $\lambda(A_{n-1})$, Theorem 1.1.1 then implies the following corollary.

COROLLARY 2.2.1. *The pinching sequence filling L_σ is Hamiltonian isotopic to the weave filling $L_{\mathcal{T}_\sigma}$.*

The vertical weave construction we use in the proof of Corollary 2.2.1 also allows us to argue that a 312-avoiding permutation yields a unique pinching sequence filling up to Hamiltonian isotopy, as we explain below. Finally, we conclude the section with a proof of Theorem 1.1.3 as a further corollary of Theorem 1.1.1.

2.2.0.1. *The D_4^- cobordism.* As preparation for relating Legendrian weaves to elementary exact Lagrangian cobordisms, we now give a precise definition of a D_4^- cobordism. As the D_4^- singularity is not a generic Legendrian front singularity, we consider a generic perturbation of the D_4^- singularity, as described in [**CZ21**, Remark 4.6] and pictured in Figure 2.1 (top). A slicing of the Legendrian front, depicted in Figure 2.1 (bottom), gives a movie of fronts describing the cobordism as follows. Near a Reeb chord trapped between two crossings, we apply a Reidemeister I move and Legendrian isotopy to shrink the Reeb chord. We then add a 1-handle to remove this Reeb chord and apply another pair of Reidemeister I moves to simplify to a diagram with one fewer crossing than we started with. The trace of this movie of fronts forms a surface in $J^1[a,b]$ and yields an exact Lagrangian cobordism in symplectic $\mathbb{R}^4$ by taking the Lagrangian projection of its embedding in contact $\mathbb{R}^5$. By convention, we will identify the remaining crossing with the leftmost crossing of the original pair.

REMARK 2.2.1. *Note that in our definition of the D_4^- cobordism above, the direction of the arrow in Figure 2.1 indicates a Lagrangian cobordism from the concave end to the convex end. In contrast, the Legendrian surface given as the Legendrian lift of the (perturbed) D_4^- singularity has no inherent directionality. Nevertheless, in our embedding of the Legendrian weave into $J^1\mathbb{D}^2$, thought of as the contactization of the symplectization of $J^1\mathbb{S}^1$, we choose to orient it so that two*

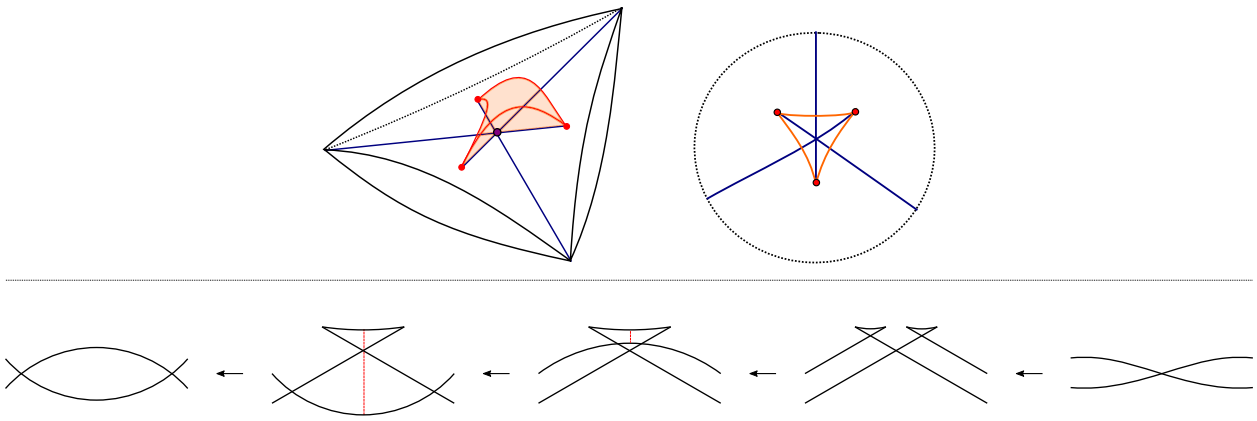

FIGURE 2.1. A local model of a generic perturbation of the D_4^- singularity as the front projection of a Legendrian surface (top, courtesy of Roger Casals and Eric Zaslow, used with permission) and as a movie of 1-dimensional fronts (bottom). The Reeb chord is depicted as a dashed red line. We first apply a Reidemeister I move before adding a 1-handle and applying two more Reidemeister I moves to arrive at a diagram with a single crossing.

of the edges appearing in the D_4^- singularity point towards the boundary in order to project to an embedded exact Lagrangian surface. Indeed, one can see that in cases where we do not require this, our Legendrian weave $\Lambda(\Gamma)$ develops additional Reeb chords that then project to immersed points of $\pi(\Lambda(\Gamma))$. This subtlety will also be clarified by the fact that the vertical weave construction in the following subsection allows us to unambiguously associate a decomposable exact Lagrangian to a Legendrian weave. □

2.2.1. Local models of exact Lagrangian saddle cobordisms. We now have the necessary ingredients for a proof of Proposition 2.2.1.

PROOF OF PROPOSITION 2.2.1. We give two local models of a D_4^- cobordism, depicted in Figures 2.2 and 2.3 as movies in the front (top) and Lagrangian (bottom) projections. The first local model depicts the removal of a Reeb chord trapped between a pair of crossings and a 0-resolution of the rightmost crossing. The second local model depicts the removal of a Reeb chord originally appearing to the left of the leftmost crossing and a 0-resolution of this crossing. This is accomplished by first applying a Legendrian isotopy to create a pair of crossings with this Reeb chord trapped between them and proceeding as in the first local model.

The main difficulty in our comparison of these local models to the pinching cobordism is to unambiguously relate the Reeb chord removed in the D_4^- cobordism to the Reeb chord removed in the pinching cobordism. This means that we must carefully manipulate the slope of the Legendrian in the front projection to ensure that no new Reeb chords are introduced throughout the process. The local models allow us to verify by inspection that no new Reeb chords appear at any point in this cobordism, as the slopes of the front projection are specified so that no new intersections appear in the Lagrangian projection.

Armed with a local model for the slicing of the D_4^- cobordism, we now describe an exact Lagrangian isotopy between this local model and the pinching cobordism. Starting in the front projection of $\lambda(\beta)$, a slicing of the pinching cobordism as defined in Subsection 1.2.4.1 consists of applying the Ng resolution, resolving a crossing, and then undoing the Ng resolution. Restricting to a neighborhood of a crossing allows us to describe the desired isotopy.

First, consider a contractible Reeb chord with a neighborhood resembling one of the two models shown in Figures 2.2 and 2.3. In such a neighborhood, the exact Lagrangian isotopy between the two cobordisms is visible when examining the Lagrangian projection of the local models depicted in Figures 2.2 and 2.3 (bottom). Indeed, after applying the Ng resolution, the only difference between these local models and the pinching cobordism in the Ng resolution is the rotating of the strand before resolving. Therefore, the movie of movies realizing the exact Lagrangian isotopy from the D_4^- cobordism to the pinching cobordism consists of incrementally applying the Legendrian isotopy of the Ng resolution, rotating the crossing before pinching, and then undoing the Ng resolution.

Now consider a Reeb chord that does not admit a neighborhood resembling one of our two local models. In this case, we can rotate all of the crossings that appear to the left of the Reeb chord past the cusps by applying the local model appearing in Figure 2.4 so that we obtain a neighborhood resembling the initial figure in 2.3. We then apply the local model given in Figure 2.3 to resolve this crossing. Finally, we rotate the remaining crossings back, and by analogous reasoning to above, the resulting cobordism is Hamiltonian isotopic to the pinching cobordism at z. □

Now that we have established the equivalence between the pinching cobordism and the D_4^- cobordism, Theorem 1.1.1 follows as a corollary.

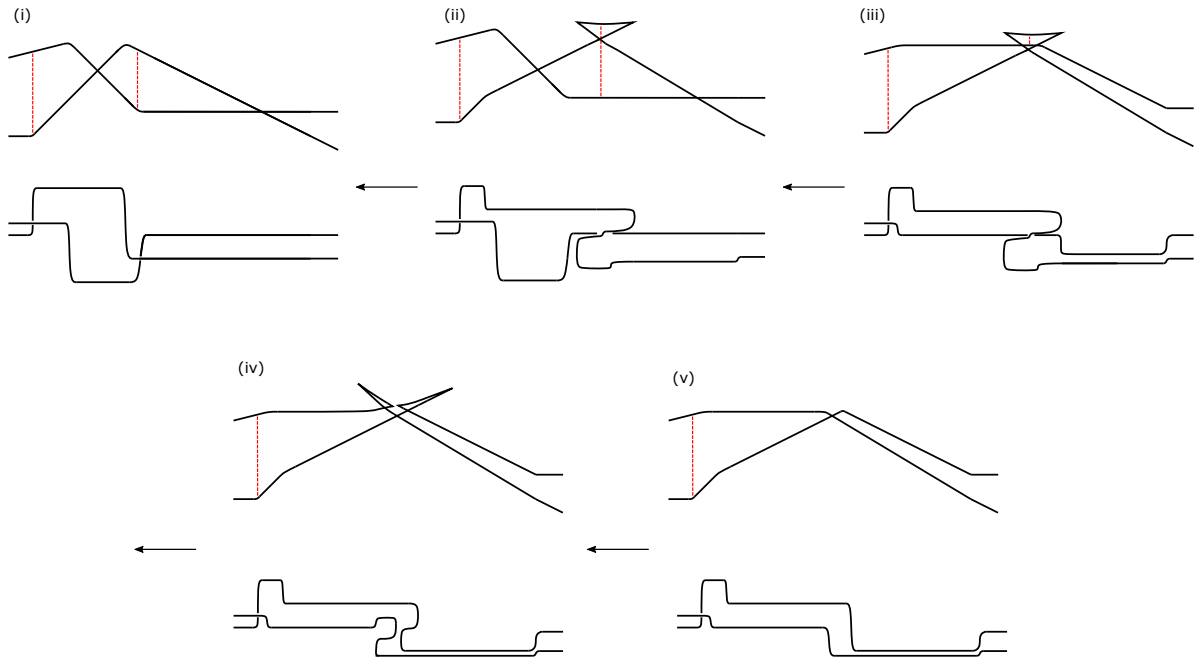

FIGURE 2.2. Local model of a D_4^- cobordism applied to a pair of crossings in the front (top) and Lagrangian (bottom) projections. Reeb chords are depicted by red dashed lines. The direction of the arrows indicate a cobordism from the concave end to the convex end.

PROOF OF THEOREM 1.1.1. By construction, any weave filling is a decomposable Lagrangian filling made up of elementary cobordisms corresponding to Reidemeister III moves and D_4^- cobordisms. By Proposition 2.2.1, the D_4^- cobordism is Hamiltonian isotopic to a pinching cobordism. Therefore, a Legendrian weave filling is Hamiltonian isotopic to a decomposable Lagrangian filling made up of Reidemeister III moves and pinching cobordisms. □

2.2.2. Exact Lagrangian fillings of $\lambda(A_n)$. To complete the proof of Corollary 2.2.1 we describe a bijection between 312-avoiding permutations of S_n and triangulations of the $n+2$-gon and show that it gives a one-to-one correspondence between fillings that resolves crossings in the same order.

For a 312-avoiding permutation σ, we denote the corresponding triangulation by $\mathcal{T}_\sigma$ and a diagonal between vertex i and vertex j of $\mathcal{T}_\sigma$ by $D_{i,j}$. Adopting the terminology of [**Reg13**], we refer to a triangle in $\mathcal{T}_\sigma$ with sides $D_{i,i+2}, D_{i,i+1}, D_{i+1,i+2}$, two of which lie on the $(n+2)$-gon, as

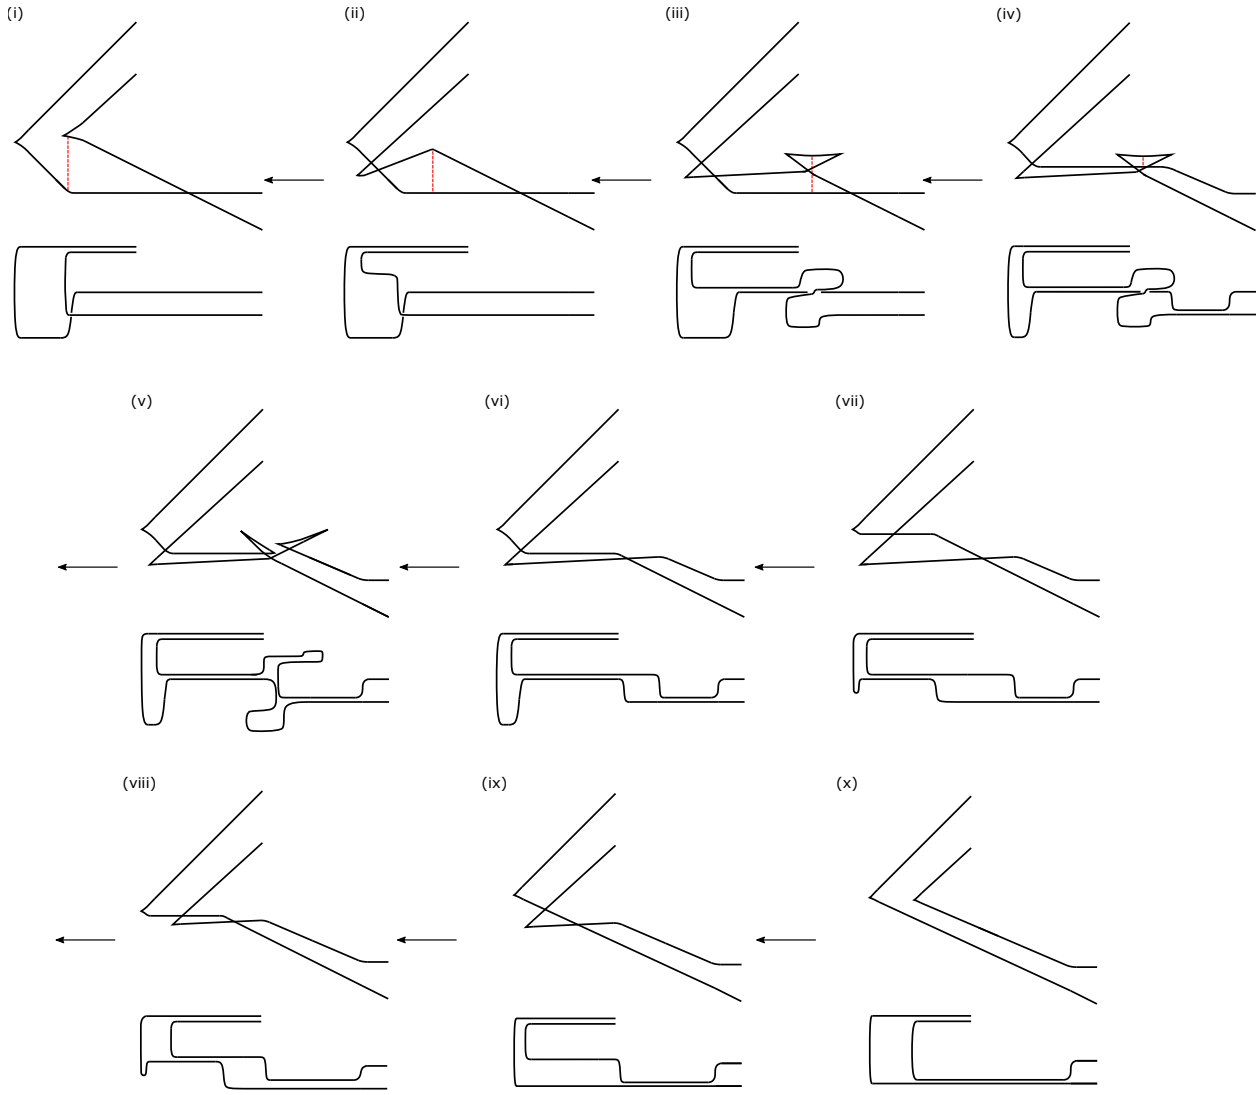

FIGURE 2.3. Local model of the leftmost crossing in the front (top) and Lagrangian (bottom) projections with a single Reeb chord depicted by a red dashed line. We first apply a Reidemeister II move in order to artificially introduce an additional crossing so that there is a single Reeb chord trapped between the new crossing and the original crossing. The D_4^- cobordism is performed in (iii)-(vi) and the remaining part of the cobordism undoes the Reidemeister II move without creating any new Reeb chords.

an ear of the triangulation. Note that any triangulation must have at least two ears and that the middle vertex of an ear necessarily has no diagonal incident to it.

Given a triangulation of the $(n+2)$-gon, the clip sequence bijection is defined as follows. First, label the vertices in clockwise order from 1 to $n+2$. Remove the middle vertex of the ear with

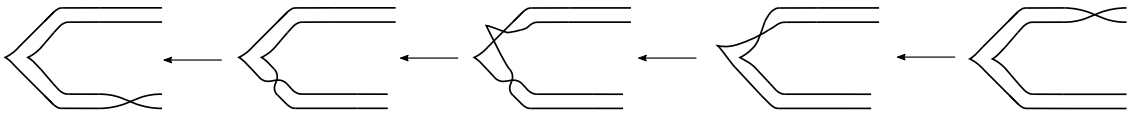

FIGURE 2.4. Local model for rotating crossings past cusps in order to isolate a desired crossing. The direction of the arrows again indicates a cobordism from the concave end to the convex end when concatenated with a cobordism described by the local model in Figure 2.3.

the smallest label, record the label and delete all edges of the $(n+2)$-gon incident to the vertex. Repeat this process with the ear whose middle vertex is now the smallest of the remaining vertices in the resulting triangulation of the $(n+1)$-gon. Continue this process until no triangles remain. The main result of [**Reg13**] is that this map is defines a bijection between the set of 312-avoiding permutations in S_n and triangulations of the $(n+2)$-gon.

The clip sequence bijection allows us to explicitly define a weave filling with the input of a 312-avoiding permutation σ.

DEFINITION 2.2.1. *The Lagrangian filling $L_{\mathcal{T}_\sigma}$ is the weave filling defined by the 2-graph dual to the triangulation $\mathcal{T}_\sigma$.*

See Figure 2.5 for a computation of the 312-avoiding permutation corresponding to the triangulation dual to the 2-graph example given above.

2.2.2.1. *Vertical weaves.* In order to relate weave fillings to decomposable fillings described by a sequence of elementary cobordisms we will make use of an equivalent way of describing weaves, combinatorially presented in [**CGGS20**] and slightly modified here. This construction arranges the N-graph vertically, with $\partial\mathbb{D}^2$ at the top and rest of the N-graph appearing below. This construction has the advantage of allowing us to unambiguously associate elementary Lagrangian cobordisms.

Let $\Gamma \subseteq \mathbb{D}^2$ be an N-graph and $\Lambda(\Gamma) \subseteq J^1(\mathbb{D}^2)$ be the associated weave. In order to produce the associated vertical weave $\Lambda_V(\Gamma)$, we foliate the disk by copies of $\mathbb{S}^1$, as shown in Figure 2.6 (left). We then consider a diffeomorphism φ taking $\mathbb{D}^2\backslash\{pt\}$ to $\mathbb{S}^1 \times (-\infty, 0]$. We define φ in such a way that the image of the foliation of $\mathbb{D}^2\backslash\{pt\}$ is a foliation of $\mathbb{R} \times (-\infty, 0]$ by horizontal lines that are identified at $\pm\infty$ to form a foliation of $\mathbb{S}^1 \times (-\infty, 0]$. The diffeomorphism φ induces a contactomorphism $\tilde{\varphi} : J^1(\mathbb{D}^2) \to J^1(\mathbb{D}^2)$.

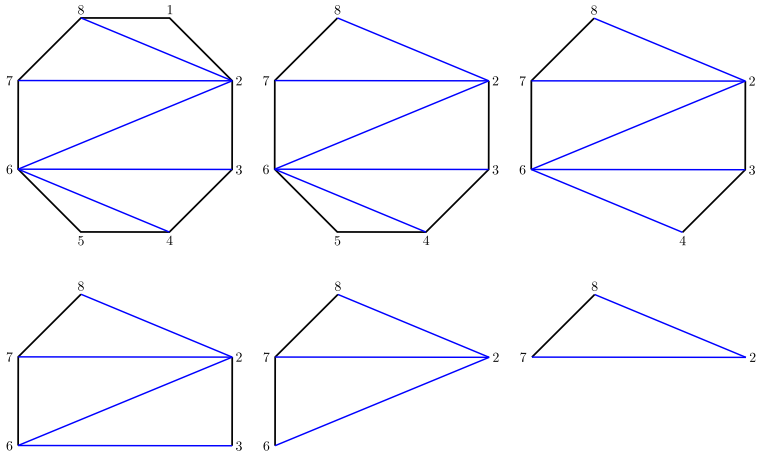

FIGURE 2.5. An example computation of the clip sequence bijection. Starting with our initial triangulation, we remove and record the smallest numbered vertex with no incident diagonals. From the sequence pictured, we get the 312-avoiding permutation $\sigma = 1\,5\,4\,3\,6\,2$. The diagonal $D_{2,8}$ yields the function $\Delta_{1,3}$ after adding 1 to both indices and reducing mod 8.

DEFINITION 2.2.2. *The vertical weave $\Lambda_V(\Gamma)$ is the Legendrian weave encoded (in the sense of Definition 1.2.1) by the $N-$graph $\varphi(\Gamma)$.*

After a planar isotopy of Γ, corresponding to a Legendrian isotopy of $\Lambda(\Gamma)$, we can assume that there are no pairs of hexavalent or trivalent vertices appearing in the same horizontal strip $\mathbb{R} \times \{t\}$. The purpose of this modification is to unambiguously decompose a weave filling into elementary Lagrangian cobordisms in order to relate it to a decomposable Lagrangian filling in the symplectization of contact $\mathbb{R}^3$. Other than our manipulation of the ambient contact manifold, the vertical weave construction is identical to the Legendrian weaves described in Section 1.2.1. See Figure 2.6 for an example comparing Legendrian weave fillings of $\lambda(A_5)$.

We list our choice of conventions for vertical weave fillings of $\lambda(A_{n-1})$ below for ease of reference.

- In a vertical weave, we encode $\lambda(\beta)$ with the braid word $\Delta\beta\Delta$ appearing at $\mathbb{R} \times \{0\}$.
- In a vertical weave, the edge exiting below a trivalent vertex with incoming edges i and $i+1$ inherits the label i.
- In a 2-graph Γ dual to a triangulation $\mathcal{T}$, the edge of Γ most immediately clockwise from vertex i of $\mathcal{T}$ is labeled by i.

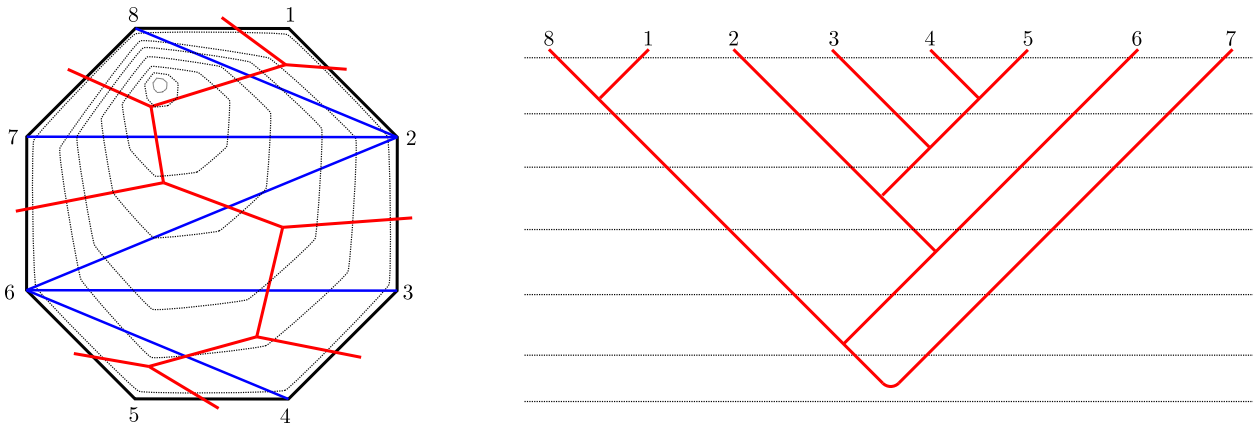

FIGURE 2.6. A pair of 2-graphs representing the same weave filling of $\lambda(A_5)$. On the left, the 2-graph Γ is inscribed in its dual triangulation of the octagon. On the right, the corresponding vertical weave is the image of the diffeomorphism φ. The edges of the vertical 2-graph are labeled by the nearest counterclockwise label of the dual triangulation. The dotted lines on the left give a foliation of D^2, corresponding to the foliation of $\mathbb{R} \times (-\infty, 0]$ depicted on the right.

Note that our choice of labeling edges differs slightly from the conventions of [**CGGS20**]. The choice of labeling given there corresponds to resolving the leftmost crossing of the pair in the D_4^- cobordism. With our choice of conventions, we can see that the clip sequence bijection yields Hamiltonian isotopic fillings.

PROOF OF COROLLARY 2.2.1. Let σ be a 312-avoiding permutation indexing a pinching sequence filling L_σ of $\lambda(A_{n-1})$ and consider the vertical weave corresponding to the triangulation $\mathcal{T}_\sigma$. By construction, a 0-resolution at the crossing i in $\lambda(A_{n-1})$ corresponds to a trivalent vertex where the incident rightmost edge is labeled by i. By Proposition 2.2.1, these denote Hamiltonian isotopic exact Lagrangian cobordisms applied to corresponding Reeb chords. Thus, the filling L_σ is Hamiltonian isotopic to the weave filling dual to the triangulation $\mathcal{T}_\sigma$. □

It is claimed without proof in [**EHK16**, Section 8.1] that, in addition to yielding the same Floer-theoretic invariant, there is a Hamiltonian isotopy between pinching sequence fillings represented by permutations $\sigma = \dots i\, k \dots j \dots$ and $\sigma' = \dots k\, i \dots j \dots$ in S_n. This claim implies that a 312-avoiding permutation represents a unique equivalence class of Lagrangian filling up to Hamiltonian isotopy. The claim follows from Corollary 2.2.1 and the lemma below.

LEMMA 2.2.1. *Let $(x_i, z_i), (x_j, z_j)$, and (x_k, z_k) denote the coordinates of three trivalent vertices in the 2-graph Γ satisfying $x_i < x_j < x_k$ and $z_j < z_k < z_i$. The planar isotopy between Γ and the 2-graph Γ' with trivalent vertices at $(x_i, z_k), (x_j, z_j)$, and (x_k, z_i) lifts to a compactly supported Hamiltonian isotopy of the fillings L_Γ and L'_Γ fixing the boundary.*

PROOF. By construction, the planar isotopy between Γ and Γ' lifts to a Legendrian isotopy between the weaves $\Lambda(\Gamma)$ and $\Lambda(\Gamma')$ in $J^1(\mathbb{D}^2)$. Note that this planar isotopy can be taken to be the identity at the boundary $\partial\Lambda(\Gamma)$. Considering the Lagrangian projection of this sequence of weaves yields a compactly supported exact Lagrangian isotopy between the Lagrangian fillings L_Γ and L'_Γ. By [**FOOO09**, Theorem 3.6.7], this implies the existence of a compactly supported Hamiltonian isotopy between the two fillings. □

By Corollary 2.2.1, the exact Lagrangian isotopy of the weave filling extends to pinching sequence fillings. Thus, our result together with [**Pan17**, Theorem Theorem 1.1] implies that there are exactly a Catalan number C_n of pinching sequence fillings[1] of $\lambda(A_{n-1})$ up to Hamiltonian isotopy.

We conclude this section with a proof of the orbital structure described in Theorem 1.1.3 as a corollary of Theorem 1.1.1. Namely, the orbital structure of the Kálmán loop action on pinching sequence fillings of $\lambda(A_{n-1})$ can be obtained from the Hamiltonian isotopy between the pinching sequence filling L_σ and weave filling $L_{\mathcal{T}_\sigma}$.

PROOF OF THEOREM 1.1.3. Let L_σ be a filling of $\lambda(A_{n-1})$ and consider the Hamiltonian isotopic weave filling $L_{\mathcal{T}_\sigma}$ with corresponding 2-graph Γ dual to the triangulation $\mathcal{T}_\sigma$. The Kálmán loop action on weave fillings is geometrically described as a cylinder rotating the entire 2-graph Γ by $\frac{2\pi}{n+2}$ radians counterclockwise. This can be readily observed from the fact that crossings of $\lambda(A_{n-1})$ are represented by edges of the dual graph intersecting the boundary of the $(n+2)$-gon. Therefore, the correspondence between triangulations $\mathcal{T}_\sigma$ and weave fillings $L_{\mathcal{T}_\sigma}$ implies that the orbital structures of triangulations under the action of rotation and weave fillings under the action of the Kálmán loop coincide.

[1]Note that a precise classification of fillings currently only exists for the Legendrian unknot. In general, it is not known whether every filling is constructible, i.e. can be given as a series of elementary cobordism.

The number of orbits of the set of triangulations of the $(n+2)$-gon under the action of counterclockwise rotation is given by the formula

$$\frac{C_n}{n+2} + \frac{C_{n/2}}{2} + \frac{2C_{(n-1)/3}}{3}$$

where, as previously, the terms with $C_{n/2}$ and $C_{n/3}$ only appear if the indices are integers. These terms correspond, respectively, to triangulations with no rotational symmetry, rotational symmetry by π, and rotational symmetry by $\frac{2\pi}{3}$. No other rotational symmetry of a triangulation is possible. The orbit sizes are $n+2$, $\frac{n+2}{2}$ and $\frac{n+2}{3}$, where again the corresponding orbit size only occurs if the relevant fraction is an integer. □

Note here the appearance of $\lambda(A_{n-1})$ as the (-1)-framed closure of the braid σ^{n+2} in the description of the weave filling. This geometrically describes why the Kálmán loop action on the rainbow closure of σ^n has order $n+2$ as an action on the $n+2$ crossings of the (-1)-framed closure.

2.3. Algebraic Proof of Theorem 1.1.3

In this section we give an algebraic proof of Theorem 1.1.3 by examining the Kálmán loop action on the augmentation variety $\mathrm{Aug}(\lambda(A_{n-1}))$ of the Legendrian link $\lambda(A_{n-1})$. As discussed in Subsection 1.2.4.2, an embedded exact Lagrangian filling yields the inclusion of an algebraic torus into the augmentation variety $\mathrm{Aug}(\lambda(A_{n-1}))$. From [**Pan17**], we have an explicit computation of a set of coordinate functions $\{s_1, \dots s_{n-1}\}$ on an induced toric chart coming from a pinching sequence filling L; namely, this set of coordinates is in bijection with the relative cycles associated to the unstable manifolds of the saddle critical points for L. Naively, we might hope to distinguish the Hamiltonian isotopy classes of the Lagrangian fillings under the Kálmán loop action by studying the associated toric charts and their s_i coordinate functions. In practice, these *local* coordinate functions are somewhat difficult to compare under this particular action. Instead, we consider the action of the Kálmán loop on the set of *global* regular functions $\{\Delta_{i,j}\}$ with $\Delta_{i,j} \in \mathbb{Z}[\mathrm{Aug}(\lambda(A_{n-1}))]$, defined in Subsection 1.2.4.4. In fact, $\Delta_{i,j} \in \mathbb{Z}[z_1, \dots, z_n]$ are globally defined polynomials, which restrict to global regular functions on the augmentation variety $\mathrm{Aug}(\lambda(A_{n-1})) \subseteq \mathbb{Z}^n$.

When considering the restriction of the $\Delta_{i,j}$ functions to the toric chart induced by the augmentation ϵ_σ, Theorem 2.3.1 below establishes that the correspondence between diagonals $D_{i-1,j-1}$

of the triangulation $\mathcal{T}_\sigma$ and the functions $\Delta_{i,j}$ is a $\mathbb{Z}_{n+2}$-equivariant map. We then show in Subsection 2.3.2 that the $\Delta_{i,j}$ functions corresponding to diagonals of a triangulation $\mathcal{T}_\sigma$ restrict to a coordinate basis of the toric chart defined by L_σ. In addition, we give an explicit formula for these coordinate functions as monomials in the s_i local coordinates. It follows that the induced action on the set of augmentations ϵ_σ in the augmentation variety $\mathrm{Aug}(\lambda(A_{n-1}))$ is equivalent to the action of rotation on triangulations of the $(n+2)$-gon, from which we can conclude the orbital structure as given in Theorem 1.1.3. See Subsection 2.3.3 for a cluster-algebraic motivation for the $\Delta_{i,j}$ functions and triangulations of the $(n+2)$-gon.

2.3.1. The Kálmán loop action on $\{\Delta_{i,j}\}$. Let us start by describing the action of the Kálmán loop on the global regular functions $\Delta_{i,j}$ using Euler's identity for continuants. All indices in this section are modulo $n+2$. Recall that we denote by $\vartheta \in \mathrm{Aut}(\mathbb{Z}[\mathrm{Aug}(\lambda(A_{n-1}))])$ the automorphism induced by the Kálmán loop acting on the augmentation variety $\mathrm{Aug}(\lambda(A_{n-1})) = \{(z_1,\ldots,z_n)|X_n = 0\} \subseteq \mathbb{Z}^n$, where $X_n \in \mathbb{Z}[z_1,\ldots,z_n]$ is the polynomial defined by $X_n = -1 + \Delta_{1,n+2}$. The action of the Kálmán loop on the set of global regular functions $\{\Delta_{i,j}\}$ is described in the following algebraic restatement of Theorem 1.1.3.

THEOREM 2.3.1. *The global regular functions $\Delta_{i,j}$ in $\mathbb{Z}[\mathrm{Aug}(\lambda(A_{n-1}))]$ satisfy the equation*

$$\vartheta(\Delta_{1,k+1}) + (-1)^k \Delta_{k,n+2} = -\Delta_{2,k} X_n \tag{2.1}$$

as global polynomials in ambient $\mathbb{Z}^n$ for $2 < k < n+2$.

As a corollary, we see that the action of ϑ on the augmentation variety $\mathrm{Aug}(\lambda(A_{n-1}))$ coincides with the action of rotation on triangulations of the $(n+2)$-gon.

COROLLARY 2.3.1. *As regular functions on $\mathrm{Aug}(\lambda(A_{n-1}))$ the $\Delta_{i,j}$ satisfy*

$$\vartheta(\Delta_{i,j}) = \begin{cases} (-1)^{j-1}\Delta_{j-1,n+2} & i = 1 \\ \Delta_{i-1,j-1} & i \neq 1 \end{cases}$$

and the map $\Delta_{i,j} \to D_{i-1,j-1}$ is a $\mathbb{Z}_{n+2}$-equivariant map.

Proof of Corollary 2.3.1. The case of $i > 1$ is discussed in Subsection 1.2.4.6 and follows from the formula for $\vartheta(\Delta_{i,j})$ given there. In the case of $i = 1$, restricting to $\mathrm{Aug}(\lambda(A_{n-1})) = \{X_n = 0\}$ causes the right hand side of Equation 2.1 to vanish. Therefore, by Theorem 2.3.1, the Kálmán loop action on the restriction of $\Delta_{1,j}$ to $\mathrm{Aug}(\lambda(A_{n-1}))$ is $\vartheta(\Delta_{1,j}) = (-1)^j \Delta_{j-1,n+2}$. Under rotation, the diagonal $D_{i-1,j-1}$ maps to $D_{i-2,j-2}$. It follows that the correspondence between $\Delta_{i,j}$ restricted to the toric chart induced by ϵ_σ and a diagonal $D_{i-1,j-1}$ of the triangulation $\mathcal{T}_\sigma$ is a $\mathbb{Z}_{n+2}$-equivariant map. □

We now give a proof of the behavior of the $\Delta_{i,j}$ as ambient polynomials in $\mathbb{Z}^n$. Note here the appearance of Euler's identity for continuants in the form of Equation 1.2.

Proof of Theorem 2.3.1. We first rewrite the left hand side of the desired equation using the continuant recursion relation (1.1) and the action of ϑ.

$$\begin{aligned}\vartheta(\Delta_{1,k+1}) + (-1)^{k-1}\Delta_{k,n+2} &= \vartheta(z_1\Delta_{2,k+1} + \Delta_{3,k+1}) + (-1)^{k-1}\Delta_{k,n+2} \\ &= -\Delta_{2,n+2}\Delta_{1,k} + \Delta_{2,k} + (-1)^{k-1}\Delta_{k,n+2}.\end{aligned}$$

We substitute this expression into the left hand side of the desired equation from Theorem 2.3.1 to obtain

$$-\Delta_{2,n+2}\Delta_{1,k} + \Delta_{2,k} + (-1)^{k-1}\Delta_{k,n+2} = -\Delta_{2,k}(\Delta_{1,n+2} - 1).$$

In order to verify that this equation holds, we will apply the special case of Euler's identity for continuants given in Equation 1.2. To do so, we distribute the right hand side and subtract $\Delta_{2,k}$ from both sides to get

$$-\Delta_{2,n+2}\Delta_{1,k} + (-1)^{k-1}\Delta_{k,n+2} = -\Delta_{2,k}\Delta_{1,n+2}.$$

This expression is equivalent to

$$\Delta_{1,n+2}\Delta_{2,k} - \Delta_{1,k}\Delta_{2,n+2} = (-1)^k\Delta_{k,n+2},$$

which is the identity given in Equation (1.2). Thus, we have established Theorem 2.3.1. □

2.3.2. The Kálmán loop action on the augmentation variety. We now prove that the $\Delta_{i,j}$ functions corresponding to the diagonals of the triangulation $\mathcal{T}_\sigma$ define a coordinate basis on the toric chart induced by the filling L_σ. To do so, we first show that the $\Delta_{i,j}$ functions can be written as monomials in the local s_i coordinate functions defined by the augmentation ϵ_σ. We then define a bijection between the $\Delta_{i,j}$ corresponding to the triangulation $\mathcal{T}_\sigma$ and the s_i variables on the toric chart induced by L_σ. Throughout the remainder of this section, let σ denote a 312-avoiding permutation corresponding to a pinching sequence filling and $D_{i,j}$ be a diagonal of the triangulation $\mathcal{T}_\sigma$. The goal of this subsection will be to prove the following proposition.

PROPOSITION 2.3.1. *The Laurent polynomial ring $\mathbb{Z}[\Delta_{i,j}^{\pm 1}]$ corresponding to the diagonals of the triangulation $\mathcal{T}_\sigma$ is isomorphic to the ring of regular functions on the toric chart induced by the augmentation ϵ_σ.*

The technical lemma introduced below will be used to prove the first part of Proposition 2.3.1.

LEMMA 2.3.1. *For any diagonal $D_{i-1,j-1}$ in the triangulation $\mathcal{T}_\sigma$, the image of the regular function $\Delta_{i,j}$ in the toric chart induced by the augmentation ϵ_σ is given by $\epsilon_\sigma(\Delta_{i,j}) = s_i \dots s_{j-2}$.*

Assuming the lemma, we first prove Proposition 2.3.1.

PROOF OF PROPOSITION 2.3.1. We first define a bijection φ between the set of triangles in the triangulation $\mathcal{T}_\sigma$ and the local toric coordinates $s_1, \dots, s_{n-1}$ induced by the augmentation ϵ_σ. Let T be a triangle in $\mathcal{T}_\sigma$ with sides $D_{i-1,j-1}, D_{j-1,k-1}$ and $D_{i-1,k-1}$. We define the map φ by

$$\varphi(T) := (\Delta_{i,j})^{-1}(\Delta_{j,k})^{-1}\Delta_{i,k}.$$

where we recall that $\Delta_{i,i+1} = 1$ by definition. By Lemma 2.3.1, we have

$$(\Delta_{i,j})^{-1}(\Delta_{j,k})^{-1}\Delta_{i,k} = (s_i \dots s_{j-2})^{-1}(s_j \dots s_{k-2})^{-1} s_i \dots s_{k-2} = s_{j-1}.$$

To see that φ is injective, consider two triangles T and T' belonging to the triangulation $\mathcal{T}_\sigma$ with sides $\{D_{i-1,j-1}, D_{j-1,k-1}, D_{i-1,k-1}\}$ and $\{D_{i'-1,j'-1}, D_{j'-1,k'-1}, D_{i'-1,k'-1}\}$, respectively. Assume that $\varphi(T) = \varphi(T')$. Then $s_{j-1} = s_{j'-1}$, and therefore $j = j'$. Since T and T' share a middle vertex, and belong to the same triangulation, they must be the same triangle. We can conclude

immediately that φ is bijective because it is an injective map between two sets of $n-1$ elements. Thus, the set of $\Delta_{i,j}$ functions corresponding to diagonals $\mathcal{T}_\sigma$ form a coordinate basis for the toric chart induced by the augmentation ϵ_σ.

□

We now give a proof of Lemma 2.3.1 by carefully examining the effect of the DGA map Φ on the braid matrices defining $\Delta_{i,j}$.

Proof of Lemma 2.3.1. Consider $\Delta_{i,j}$ corresponding to some diagonal $D_{i-1,j-1}$ of a triangulation $\mathcal{T}_\sigma$. By definition, we have

$$\Phi(\Delta_{i,j}) = \left[\prod_{k=i}^{j-2} B(\Phi(z_k))\right]_{(2,2)}.$$

Therefore, Lemma 2.3.1 is equivalent to the claim that the $(2,2)$ entry of $\prod_{k=i}^{j-2} B(\epsilon_\sigma(z_k))$ is precisely $\prod_{k=i}^{j-2} s_k$. To verify this statement, we show inductively that applying $\Phi_l \circ \cdots \circ \Phi_1$ yields a product of $B(z_k)$ for $k \in \{i, \ldots, j-2\} \backslash \{\sigma(1) \ldots, \sigma(l)\}$ with a particular collection of diagonal matrices, upper triangular and lower triangular matrices.

Define the matrices

$$C(s) := \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix} \qquad U(s) := \begin{pmatrix} -s^{-1} & 1 \\ 0 & s \end{pmatrix} \qquad L(s) := \begin{pmatrix} -s^{-1} & 0 \\ 1 & s \end{pmatrix} \qquad D(s) := \begin{pmatrix} -s^{-1} & 0 \\ 0 & s \end{pmatrix}$$

Denote by A^T the transpose of the matrix A. The following identities are immediate.

$$B(z \pm s) = B(z)C(\pm s) \tag{2.2}$$

$$B(z \pm s) = C^\mathsf{T}(\pm s)B(z) \tag{2.3}$$

$$C(s \pm t) = C(s)C(\pm t) \tag{2.4}$$

$$C(-s^{-1})B(s)C^\mathsf{T}(-s^{-1}) = D(s) \tag{2.5}$$

$$C(-s^{-1})B(s) = L(s) \tag{2.6}$$

$$B(s)C^\mathsf{T}(-s^{-1}) = U(s) \tag{2.7}$$

Equipped with this set of identities, we proceed with the proof of Lemma 2.3.1. First, we may assume that $\sigma(1), \ldots, \sigma(k)$ all lie in the set $\{i, \ldots j-2\}$. Indeed, for $\sigma(l)$ not in $\{i, \ldots j-2\}$, the map Φ_l is the identity on the polynomial $\Phi_{l-1} \circ \cdots \circ \Phi_1(\Delta_{i,j})$. This follows from the observation that if $D_{i-1,j-1}$ is in the triangulation $\mathcal{T}_\sigma$, then i appears before $i-1$ and $j-2$ appears before $j-1$ in σ under the clip sequence bijection. Therefore, $\sigma^{-1}(i) < \sigma^{-1}(i-1)$ and $\sigma^{-1}(j-2) < \sigma^{-1}(j-1)$, which implies that no elements of the set T^l_σ appear in terms of $\Phi_{l-1} \circ \cdots \circ \Phi_1(\Delta_{i,j})$.

Denote by M_l^+ and M_l^- the maximum and minimum of the set $\{i, \ldots, j-2\} \backslash \{\sigma(1), \ldots, \sigma(l-1)\}$. We claim that the result of applying Φ_l to $\Phi_{l-1} \circ \cdots \circ \Phi_1(\Delta_{i,j})$ results in the replacement of $B(z_{\sigma(l)})$ in the product $\prod_{k=i}^{j-2} B(\Phi_{l-1} \circ \cdots \circ \Phi_1(z_k))$ with one of three possibilities depending on l:

(1) For $\sigma(l) = M_l^-$, the map Φ_l replaces $B(z_{\sigma(l)})$ by the upper triangular matrix $U(s_{\sigma(l)})$.
(2) For $\sigma(l) = M_l^+$, the map Φ_l replaces $B(z_{\sigma(l)})$ by the lower triangular matrix $L(s_{\sigma(l)})$.
(3) For $M_l^- < \sigma(l) < M_l^+$, the map Φ_l replaces $B(z_{\sigma(l)})$ by the diagonal matrix matrix $D(s_{\sigma(l)})$.

We prove this claim by induction. For the base case, we consider the three possibilities listed above.

(1) If $\sigma(1) = M_1^- = i$, then we have

$$\begin{aligned}\Phi_1(\Delta_{i,j}) &= B(s_i)B(z_{i+1} - s_i^{-1})B(z_{i+2}) \ldots B(z_{j-2}) \\ &= B(s_i)C^{\mathsf{T}}(-s_i^{-1})B(z_{i+1}) \ldots B(z_{j-2}) \\ &= U(s_i)B(z_{i+1}) \ldots B(z_{j-2})\end{aligned}$$

where the second equality follows from Equation (2.3) and the final one from Equation (2.7).

(2) If $\sigma(1) = M_1^+ = j-2$, then Equations (2.2) and (2.6) imply that

$$\Phi_1(\Delta_{i,j}) = B(z_i) \ldots B(z_{j-3})C(-s_{j-2}^{-1})B(s_{j-2}) = B(z_i) \ldots B(z_{j-3})L(s_{j-2}).$$

(3) If $i < \sigma(1) < j-2$, then we apply Equations (2.2), (2.3), and (2.5) to $\Phi_1(\Delta_{i,j})$ to obtain

$$\begin{aligned}\Phi_1(\Delta_{i,j}) &= B(z_i)\dots B(z_{\sigma(1)-1})C(-s_{\sigma(1)}^{-1})B(s_{\sigma(1)})C^{\mathsf{T}}(-s_{\sigma(1)}^{-1})B(z_{\sigma(1)+1})\dots B(z_{j-2})\\ &= B(z_i)\dots B(z_{\sigma(1)-1})D(s_{\sigma(1)})B(z_{\sigma(1)+1})\dots B(z_{j-2}).\end{aligned}$$

Assume inductively that applying the composition $\Phi_{l-1}\circ\cdots\circ\Phi_1$ replaces each $B(z_{\sigma(k)})$ for $1 \le k \le l-1$ with either $U(s_{\sigma(k)})$, $L(s_{\sigma(k)})$, or $D(s_{\sigma(k)})$ depending on whether $\sigma(k) = M_k^-$, $\sigma(k) = M_k^+$, or $M_k^- < \sigma(k) < M_k^+$ respectively. We consider the same three cases for Φ_l:

(1) If $\sigma(l) = M_l^-$, then T_σ^l has a single element l' and by the combinatorial formula for Φ_l, we have

$$B(\Phi(z_{\sigma(l')})) = C^{\mathsf{T}}(\pm s_{\sigma(l)}^{-1}s_{\sigma(l)+1}^{-2}\dots s_{\sigma(l')-1}^{-2})B(z_{\sigma(l')})$$

where the sign is given by $(-1)^{|\sigma(l)-(\sigma(l')-1)|}$. By the inductive hypothesis, we have that the matrices appearing between $B(\Phi(z_{\sigma(l)})) = B(s_{\sigma(l)})$ and $C^{\mathsf{T}}(s_{\sigma(l)}^{-1}s_{\sigma(l)+1}^{-2}\dots s_{\sigma(l')-1}^{-2})$ are of the form $D(s_{\sigma(l)+1})\dots D(s_{\sigma(l')-1})$. We then compute

$$\begin{aligned}B(s_{\sigma(l)})\left(\prod_{m=\sigma(l)+1}^{\sigma(l')-1} D(s_m)\right)C^{\mathsf{T}}(\pm s_{\sigma(l)}^{-1}s_{\sigma(l)+1}^{-2}\dots s_{\sigma(l')-1}^{-2}) &= \begin{pmatrix}\pm s_{\sigma(l)}^{-1}\dots s_{\sigma(l')-1}^{-1} & s_{\sigma(l)+1}\dots s_{\sigma(l')-1}\\ 0 & s_{\sigma(l)}\dots s_{\sigma(l')-1}\end{pmatrix}\\ &= U(s_{\sigma(l)})\left(\prod_{m=\sigma(l)+1}^{\sigma(l')-1} D(s_m)\right),\end{aligned}$$

showing that applying Φ_l replaces $B(z_l)$ by $U(s_l)$.

(2) If $\sigma(l) = M_l^+$, then T_σ^l again has a single element l'' and we have

$$C(\pm s_{\sigma(l'')+1}^{-2}\dots s_{\sigma(l)-1}^{-2}s_{\sigma(l)}^{-1})\left(\prod_{m=\sigma(l'')+1}^{\sigma(l)-1} D(s_m)\right)B(s_{\sigma(l)}) = \left(\prod_{m=\sigma(l'')+1}^{\sigma(l)-1} D(s_m)\right)L(s_{\sigma(l)}).$$

(3) Finally, if $M_l^- < \sigma(l) < M_l^+$, then T_σ^l has two elements, denote them by l' and l'' with $l' > l''$. Then we must consider the product of

$$C(\pm s_{\sigma(l'')+1}^{-2}\dots s_{\sigma(l)-1}^{-2}s_{\sigma(l)}^{-1})\left(\prod_{m=\sigma(l'')+1}^{\sigma(l)-1} D(s_m)\right)$$

with

$$B(s_{\sigma(l)})\left(\prod_{m=\sigma(l)+1}^{\sigma(l')-1} D(s_m)\right)C^{\mathsf{T}}(\pm s_{\sigma(l)}^{-1}s_{\sigma(l)+1}^{-2}\dots s_{\sigma(l')-1}^{-2})$$

where the two signs of entries of C and C^{T} need not agree. We apply our computation from the previous case and simplify

$$\left(\prod_{m=\sigma(l'')+1}^{\sigma(l)-1} D(s_m)\right)L(s_{\sigma(l)})\left(\prod_{m=\sigma(l)+1}^{\sigma(l')-1} D(s_m)\right)C^{\mathsf{T}}(\pm s_{\sigma(l)}^{-1}s_{\sigma(l)+1}^{-2}\dots s_{\sigma(l')-1}^{-2})$$

$$= \left(\prod_{m=\sigma(l'')+1}^{\sigma(l)-1} D(s_m)\right)L(s_{\sigma(l)})\begin{pmatrix} \mp(s_{\sigma(l)+1}\dots s_{\sigma(l')-1})^{-1} & 0 \\ \pm(s_{\sigma(l)}s_{\sigma(l)+1}\dots s_{\sigma(l')-1})^{-1} & s_{\sigma(l)+1}\dots s_{\sigma(l')-1}\end{pmatrix}$$

This yields the product $\prod_{m=\sigma(l'')+1}^{\sigma(l')-1} D(s_m)$, as desired.

Thus, by induction, Φ_l replaces $B(z_l)$ with an upper triangular, lower triangular, or diagonal matrix for $1 < l < j-2-i$. Therefore, when we arrive at the final Φ_{j-2-i} map, we have $B_{s_{\sigma(j-2-i)}}$ multiplied on the left by the product of some number of upper triangular and diagonal matrices with the corresponding s variable appearing in the $(2,2)$ entry and multiplied on the right by some number of diagonal and lower triangular matrices with the same condition. Since multiplication by a diagonal matrix preserves the property of being upper or lower triangular, the result is a product of the form $UB(s_{\sigma(j-2-i)})L$ where U and L are upper and lower triangular matrices. Therefore the $(2,2)$ entry of this product is the product of the $(2,2)$ entries of each of the factors. It follows that the $(2,2)$ entry of $\Phi(\Delta_{i,j})$ is $s_i \dots s_{j-2}$, as desired. □

Together, Corollary 2.3.1 and Proposition 2.3.1 allow us to give an algebraic proof of Theorem 1.1.3.

Proof of Theorem 1.1.3. By Proposition 2.3.1, the set of $\Delta_{i,j}$ corresponding to a triangulation $\mathcal{T}_\sigma$ gives a basis for the toric chart induced by the augmentation ϵ_σ. Therefore, the image of ϵ_σ under ϑ corresponds to the image of the set of $\Delta_{i,j}$ corresponding to $\mathcal{T}_\sigma$. By Corollary 2.3.1, we know that the induced action of the Kálmán loop on the augmentation variety of $\lambda(A_{n-1})$ is equivalent to the action of rotation on the $(n+2)$-gon. □

2.3.3. Relation to cluster theory. The appearance of the $\Delta_{i,j}$ functions and the combinatorics of the $(n+2)$-gon is explained by a cluster structure on the augmentation variety, the existence of which was recently proven by Gao-Shen-Weng in [**GSW20a**]. In brief, a cluster variety is an algebraic variety containing a set of toric charts (cluster charts) with coordinate functions (cluster variables) that transform according to a specific operation (cluster mutation) under the chart maps. See [**FWZ20a**, **FWZ20b**] for more on cluster algebras.

For a Legendrian λ given as the rainbow closure of a positive braid, [**GSW20a**] describes a cluster structure on $\mathrm{Aug}(\lambda)$ by proving a natural isomorphism to double Bott-Samelson cells. In particular, the cluster structure on $\mathrm{Aug}(\lambda(A_{n-1}))$ is a cluster algebra of A_{n-1}-type. A_{n-1}-type cluster algebras were originally defined and studied by Fomin and Zelevinksy in the context of regular functions on the affine cone of the Grassmanian $\mathcal{G}r^{\times}(2, n+2)$ [**FZ03**]. If we consider the Plücker coordinate $P_{i,j}$ of the (ordinary) Grassmanian $Gr(2, n+2)$, then its image in the affine cone is precisely the function $\Delta_{i,j}$. The combinatorics of the relationship between cluster charts is captured by the flip graph, where a single cluster seed is given by all $\Delta_{i,j}$ corresponding to diagonals $D_{i,j}$ of a triangulation. In the context of this manuscript, [**GSW20a**] implies the existence of cluster coordinates on $\mathrm{Aug}(\lambda(A_{n-1}))$ while Proposition 2.3.1 gives a precise formula. The special case of Euler's identity for continuants that appears in the algebraic proof of Theorem 1.1.3 can be understood as a three-term Plücker relation describing mutation at the diagonal $\Delta_{1,k}$ or $\Delta_{2,n+2}$.

Also of interest in the cluster setting is the fact that the Kálmán loop induces a cluster automorphism of the augmentation variety $\mathrm{Aug}(\lambda(A_{n-1}))$. Subsection 2.4.2 explicitly realizes this automorphism as a sequence of mutations. For an A_{n-1}-type cluster algebra, Assem, Schiffler, and Shramchenko showed that the cluster automorphism group is $\mathbb{Z}_{n+2}$ [**ASS12**]. Theorem 2.3.1 implies that the order of the Kálmán loop action on $\mathrm{Aug}(\lambda(A_{n-1}))$ is precisely $n+2$, so we immediately deduce the following corollary.

Corollary 2.3.2. *The induced action of the Kálmán loop on* $\mathrm{Aug}(\lambda(A_{n-1}))$ *is a generator of the* A_{n-1}*-type cluster modular group.*

2.4. Combinatorial Characterizations

In this section, we describe the combinatorial properties of the Kálmán loop action on a pinching sequence filling L_σ of $\lambda(A_{n-1})$ purely in terms of the corresponding 312-avoiding permutation σ. We first present an explicit algorithm for determining the orbit size of L_σ from σ in Subsection 2.4.1. The end of the subsection includes a table where orbit sizes are computed for the case $n = 4$, corresponding to triangulations of the hexagon. We then give a recipe for constructing a geodesic path in the flip graph that describes a counterclockwise rotation of the triangulation $\mathcal{T}_\sigma$. Since the weave filling $L_{\mathcal{T}_\sigma}$ is Hamiltonian isotopic to the pinching sequence filling L_σ by Theorem 1.1.1, this geodesic path describes the Kálmán loop action on L_σ as a sequence of edge flips. Finally, we discuss the behavior of 312-avoiding permutations under a single edge flip in the flip graph. Together, these last two results give a combinatorial characterization of the Kálmán loop action on fillings purely in terms of 312-avoiding permutations. As in previous sections, all indices are computed modulo $n + 2$.

2.4.1. Orbit size. To produce an algorithm for determining orbit size of a pinching sequence filling, we give explicit criteria in Lemmas 2.4.3 and 2.4.4 for when a filling of $\lambda(A_{n-1})$ has orbit size $\frac{n+2}{2}$ or $\frac{n+2}{3}$ under the action of the Kálmán loop. If it does not satisfy either of these criteria, then it necessarily has orbit size $n+2$. We start by describing the permutations that arise from an orbit of size $\frac{n+2}{2}$.

Consider some 312-avoiding permutation $\sigma \in S_n$. In order for the filling L_σ to have orbit size $\frac{n+2}{2}$, the triangulation $\mathcal{T}_\sigma$ must have rotational symmetry through an angle of π. Therefore, $\mathcal{T}_\sigma$ has a diameter $D_{i,i+\frac{n+2}{2}}$ and the triangulated polygons on either side of this diameter must be mirror images. We consider the diameter as an external edge of two $(\frac{n+2}{2}+1)$-gons, one containing both vertices labeled $n+1$ and $n+2$, and the other containing at most one of them. For any 312-avoiding permutation σ such that $\mathcal{T}_\sigma$ has a diameter $D_{i,i+\frac{n+2}{2}}$, we define the 312-avoiding permutation τ in $S_{\frac{n+2}{2}}$ corresponding to half of the triangulation $\mathcal{T}_\sigma$ as follows.

Definition 2.4.1. *The permutation τ in the letters $i+1, i+1, \ldots, i+\frac{n+2}{2}-1$ is the 312-avoiding permutation obtained from applying the clip sequence bijection to the triangulation $\mathcal{T}_\tau$ of the $(\frac{n+2}{2}+1)$-gon containing at most one of the vertices labeled by $n+1$ and $n+2$.*

We can always unambiguously identify τ from the permutation σ.

LEMMA 2.4.1. *Let $\sigma \in S_n$ be any 312-avoiding permutation such that $\mathcal{T}_\sigma$ has rotational symmetry through an angle of π. The permutation τ is the first subword of σ of length $\frac{n}{2}$ letters forming a subinterval of the integers $\{1, \ldots, n\}$.*

PROOF. Let $\sigma \in S_n$ correspond to a triangulation $\mathcal{T}_\sigma$ with rotational symmetry through an angle of π. Under the clip sequence bijection, there may be letters of σ that appear before τ. Therefore, to identify τ as a subword of σ we search for the first 312-avoiding permutation of length $\frac{n}{2}$ that appears in σ. A diameter $D_{i,i+\frac{n+2}{2}}$ forces the condition that any letters appearing before τ will be less than i, so that even if i appears directly after τ, there is no ambiguity in identifying τ.

Explicitly, we identify τ by first checking if the set $\{\sigma(1), \ldots \sigma(\frac{n}{2})\}$ of the first $\frac{n}{2}$ letters of σ is equal to a subinterval of the integers $\{1, \ldots, n\}$ of length $\frac{n}{2}$. If not, we check the $\{\sigma(2), \ldots, \sigma(\frac{n}{2}+1)\}$. We continue in this way until we have either identified the subword τ or exhausted all possibilities. If no such subword exists, then σ does not have the assumed rotational symmetry. □

We now state a preparatory lemma regarding details of the clip sequence bijection that may give some insight into the structure of the orbit size algorithm below. We consider the most general case where $\mathcal{T}_\tau$ is a subtriangulation of $\mathcal{T}_\sigma$ with vertices $i, \ldots, i+k$ for $i+k \leq n+1$.

LEMMA 2.4.2. *Let $k \in \mathbb{N}$ satisfy $j < i+k \leq n+1$. The 312-avoiding permutation τ ends in the letter j if and only if the subtriangulation $\mathcal{T}_\tau$ contains the triangle labeled by vertices i, j, and $i+k$. In this case, all letters taking values strictly between i and j appear before any other letters in τ.*

PROOF. The first claim follows from the definition of the clip sequence bijection because the diagonal $D_{i,i+k}$ must appear in the final triangle remaining after removing the previous $n-1$ vertices. Therefore, j is the final letter of τ, if and only if it is also the third vertex of this triangle.

The second claim follows by similar reasoning to the case of the diameter, as the existence of the diagonal $D_{i,j}$ implies that there must be some ear $D_{l,l+2}$ with $i < l < j-2$. Therefore, $l+1$ appears before j and we can repeat this argument for the subtriangulation of $\mathcal{T}_\tau$ obtained by removing the vertex $l+1$. □

We now give explicit criteria for determining whether the filling L_σ has orbit size $\frac{n+2}{2}$ solely in terms of σ.

LEMMA 2.4.3. *The following algorithm detects whether a 312-avoiding permutation σ in S_n yields a filling L_σ of orbit size $\frac{n+2}{2}$ under the action of the Kálmán loop.*

(1) Identify τ from σ as in the proof of lemma 2.4.1.

(2) Define σ' to be an empty string and set $\tau' = \tau$. Find the smallest j for which $j \geq \frac{n+2}{2}$ and some $k > j$ appears before j in τ. For the first such k appearing in τ', append $k - \frac{n+2}{2}$ to σ', remove k from τ' and repeat until no such letters remain in τ'. Append τ to σ'.

(3) While τ' ends in the largest (resp. smallest) number remaining in τ' not equal to $\frac{n+2}{2} - 1$ (resp. $\frac{n+2}{2}$), then append the next largest (resp. next smallest) number in $\{1, \ldots, n\} \backslash \sigma'$ less than the smallest number (resp. greater than the largest number) of τ' to σ' and delete the final number of τ'.

(4) If τ' does not end in its largest or smallest remaining number, then add $\frac{n+2}{2}$ to all numbers less than the final number and append to σ' in the order they appear. Delete the corresponding numbers from τ'.

(5) Now τ' ends in its smallest remaining number, so return to Step (3) and repeat until only one number remains in τ'. The final number of σ' is then determined by the unique number remaining in $\{1, \ldots, n\} \backslash \sigma'$.

(6) σ has orbit size $\frac{n+2}{2}$ if it is equal to σ'.

EXAMPLE 2.4.1. *Consider the 312-avoiding permutation $\sigma = 1\,5\,4\,3\,6\,2$. We can identify $\tau = 5\,4\,3$ as the first length 3 subword appearing in σ and the diameter of the triangulation $\mathcal{T}_\sigma$ is therefore $D_{2,6}$. Applying the above algorithm to τ, we see that Step (2) yields $\sigma' = 1\,5\,4\,3$ because 5 precedes 4. Then 3 is the smallest number appearing in τ, so we append 6 to σ'. Finally, we append 2, to get $\sigma = \sigma'$, indicating that the filling labeled by σ has orbit size 4 under the Kálmán loop.* □

PROOF. Let $\sigma \in S_n$ be a 312-avoiding permutation with orbit size $\frac{n+2}{2}$. Denote the diameter of $\mathcal{T}_\sigma$ as $D_{i,i+\frac{n+2}{2}}$ for some $1 \leq i \leq \frac{n+2}{2} - 1$ and the permutation corresponding to the triangulation of the $(\frac{n+2}{2} + 1)$-gon given by the vertices $i, \ldots, i + \frac{n+2}{2}$ by τ. We will show that the algorithm detects when the triangulation $\mathcal{T}_\sigma$ is obtained from the triangulation $\mathcal{T}_\tau$ by gluing $\mathcal{T}_\tau$ to a rotation of $\mathcal{T}_\tau$

by π along the diagonal $D_{i,i+\frac{n+2}{2}}$. The lemma then follows from the observation that τ is uniquely determined from σ.

Under the clip sequence bijection, we delete the smallest vertex with no incident diagonal at each step and append the label to the permutation. Therefore, any letter k of σ appearing before τ is less than i. Moreover, any diagonal $D_{j,k}$ or $D_{k,j}$ (should it exist) incident to the vertex k has endpoint j in the set $\{n+2, 1, \ldots, i\}$. Therefore any triangle with vertices j, k,and l with j, k, l given in clockwise order must also have $j, l \in \{n+2, 1, \ldots, i\}$. The rotational symmetry of $\mathcal{T}_\sigma$ implies that the triangle with vertices $j+\frac{n+2}{2}, k+\frac{n+2}{2}, l+\frac{n+2}{2}$ appears in $\mathcal{T}_\tau$. It follows from Lemma 2.4.2 that in τ the letter $k+\frac{n+2}{2}$ precedes j for some $j \geq \frac{n+2}{2}$ and that all such k appear before τ in σ. Therefore, Step (2) produces all letters of σ that appear before τ.

To determine the letters following τ in σ, we first consider the case where one of the diameter vertices, i or $i+\frac{n+2}{2}$, has no incident diagonals with endpoint taking values in the set of vertices labeled by letters appearing after τ in σ. If this is the case, then the appropriate diameter vertex label immediately follows τ in σ under the clip sequence bijection. We also observe that when i (respectively, $i+\frac{n+2}{2}$) is such a vertex, then there is a triangle in $\mathcal{T}_\sigma$ with vertices $i, i-1$, and $i+\frac{n+2}{2}$ (resp. $i, i+\frac{n+2}{2}, i+\frac{n+2}{2}+1$). Therefore, the rotational symmetry of $\mathcal{T}_\sigma$ implies that we have a triangle with vertices $i, i+\frac{n+2}{2}-1$ and $i+\frac{n+2}{2}$ (resp. $i, i+1$, and $i+\frac{n+2}{2}$) in $\mathcal{T}_\tau$. By Lemma 2.4.2, the vertex $i+\frac{n+2}{2}-1$ (resp. $i+1$) appears as the final letter in τ. The vertex $i-1$ (resp. $i+\frac{n+2}{2}+1$) then appears immediately following τ. The same reasoning applies if we replace the diameter $D_{i,i+\frac{n+2}{2}}$ with the diagonal $D_{i-1,i+\frac{n+2}{2}}$, $D_{i,i+\frac{n+2}{2}+1}$, or any such longest remaining diagonal arising under the clip sequence bijection in this way, so long as $n+1$ or $n+2$ do not appear as endpoints of this diagonal.

If both diameter vertices have diagonals incident to them with endpoints in the remaining vertices, then the letter following τ under the clip sequence bijection labels the smallest vertex greater than $i+\frac{n+2}{2}$ with no incident diagonals. By previous reasoning, we know that the diameter is one side of a triangle with vertices $i, k, i+\frac{n+2}{2}$ in $\mathcal{T}_\sigma$. The rotational symmetry of $\mathcal{T}_\sigma$ implies that the triangle labeled by $i, k-\frac{n+2}{2}, i+\frac{n+2}{2}$ appears in $\mathcal{T}_\tau$. It follows from Lemma 2.4.2 that k appears as the final letter of τ and any letter j with $j<k$ appearing before k in τ

This process continues until we have eliminated all numbers from τ except for either $n+1$ or $n+2$. This unambiguously determines the final number of our permutation. By construction, we have shown that the above algorithm yields the 312-avoiding permutation σ with $\mathcal{T}_\sigma$ constructed by gluing a rotated copy of $\mathcal{T}_\tau$ to $\mathcal{T}_\tau$. $\square$

We now consider the case of a 312-avoiding permutation σ with orbit size $\frac{n+2}{3}$. In order to exhibit the appropriate rotational symmetry, the triangulation $\mathcal{T}_\sigma$ must have a central triangle labeled by vertices $i, i+\frac{n+2}{3}, i+\frac{2(n+2)}{3}$, dividing the triangulation up into three identical triangulations of $(\frac{n+2}{3}+1)-$gons. Two of these polygons do not contain the pair of vertices $n+1$ and $n+2$, so a permutation σ with $\mathcal{T}_\sigma$ having rotational symmetry through an angle of $\frac{2\pi}{3}$ must have two subwords τ_1 and τ_2 of length $\frac{n+2}{3}-1$ that differ by $\frac{n+2}{3}$ and are immediately followed by $i+\frac{n+2}{3}$. We determine the third subword from τ_1 using the same reasoning as in the $\frac{n+2}{2}$ orbit size case.

LEMMA 2.4.4. *The following algorithm detects whether a 312-avoiding permutation σ in S_n yields a filling L_σ of orbit size $\frac{n+2}{3}$ under the action of the Kálmán loop.*

(1) Determine τ_1 by finding the first subword of length $\frac{n+2}{3}-1$ in σ with letters $i, \ldots, i+\frac{n+2}{3}-1$ for some i. If no such τ_1 exists, then σ does not have orbit size $\frac{n+2}{3}$.

(2) Set σ' to be the empty word. For any numbers greater than $\frac{n+2}{3}$ that appear after $\frac{n+2}{3}$ or some other number greater than $\frac{n+2}{3}$, add $\frac{2(n+2)}{3} \pmod{n+2}$ to them and append the result to σ'. Append τ_1 to σ'. Add $\frac{n+2}{3}$ to each entry of τ_1 to get τ_2 and append to σ'. Append $i+\frac{n+2}{3}$ to σ'. Delete the corresponding number from τ_1.

(3) So long as τ_1 ends in the largest (resp. smallest) number remaining in τ_1 not equal to $\frac{n+2}{3}-1$ (resp. $\frac{n+2}{3}$), then append the next largest (resp. next smallest) number of $\{1, \ldots, n\} \backslash \sigma'$ less than the smallest number (resp. greater than the largest number) of τ_1 to σ' and delete the final number in τ_1.

(4) If τ_1 does not yet end in the largest or smallest remaining number, add $\frac{n+2}{3}$ to all numbers less than the final number and append. Delete the corresponding numbers from τ_1.

(5) Now τ_1 ends in the smallest remaining number, so return to Step (3) and continue until one number remains in τ_1. The final number of σ' is then determined by the unique number appearing in $\{1, \ldots, n\} \backslash \sigma'$.

(6) σ has orbit size $\frac{n+2}{3}$ if it is equal to σ'.

EXAMPLE 2.4.2. *We can identify* $\sigma = 2\,1\,5\,4\,3\,6\,7$ *as a permutation with orbit size* $\frac{7+2}{3} = 3$ *using the above algorithm. First identify* τ_1 *as the first length 2 subword with two consecutive letters, i.e.,* $\tau_1 = 2\,1 \in S_2$. *Then* $\tau_2 = 5\,4$ *and the string* $2\,1\,5\,4\,3$ *must appear in* σ *in order for it to have orbit size 3. We can also determine that no letters appear before* τ_1 *because 1 and 2 already appear in our word. Since* τ_1 *ends with the smallest letter of the triangulation* $\mathcal{T}_{\tau_1}$, *we append 6. The final remaining number is 7, so we see that* $\sigma' = \sigma$ *and therefore* σ *has orbit size 3.* □

We conclude this subsection with a table of orbit sizes of pinching sequence fillings of $\lambda(A_3)$, i.e. the case $n = 4$.

Permutation	Orbit Size
1 2 3 4	6
1 2 4 3	3
1 3 2 4	2
1 3 4 2	6
1 4 3 2	3
2 1 3 4	3
2 1 4 3	6
2 3 1 4	3
2 3 4 1	6
2 4 3 1	2
3 2 1 4	6
3 2 4 1	3
3 4 2 1	3
4 3 2 1	6

2.4.2. Rotations of Triangulations. In this subsection, we describe a counterclockwise rotation of the $(n+2)$-gon through an angle of $\frac{2\pi}{n+2}$ as a sequence of edge flips in a given triangulation. As remarked in [**TZ18**, Remark 1.6], Legendrian mutation of a Legendrian weave surface dual to a triangulation corresponds to exchanging diagonals of a quadrilateral in the original triangulation

to form a new triangulation. Such an exchange of diagonals is depicted in Figure 2.9, and we refer to it as an edge flip. See Subsection 2.3.3 for more on the cluster-algebraic interpretation of this operation in terms of cluster mutation. The flip graph or associahedron is then defined to have vertices given by triangulations and an edge between two vertices if the triangulations are related by a single edge flip. The diameter of the flip graph was first investigated via geometric methods by Thurston, Sleator and Tarjan in [**STT88**] and later combinatorially by Pournin in [**Pou14**]. In general, there is no known algorithm for determining geodesics of the flip graph. Below, we present a description of the Kálmán loop as a sequence of edge flips in the flip graph and describe the result of a single edge flip on a 312-avoiding permutation, thus providing a characterization of the Kálmán loop action as a geodesic path in the flip graph.

We refer to any triangle with edges made up solely of diagonals $D_{i,i+j}$ for $j \geq 2$ as an internal triangle, and we denote the number of internal triangles in a triangulation $\mathcal{T}_\sigma$ by t_σ.

We will say that a diagonal $D_{i,j}$ is (counter)clockwise to another diagonal $D_{i,j'}$ if the vertex j is (counter)clockwise to j'. Similarly, $D_{i,j}$ is (counter)clockwise to $D_{i',j}$ if i is (counter)clockwise to i'. Given a triangulation $\mathcal{T}_\sigma$, the following algorithm describes a sequence of $n-1+t_\sigma$ edge flips that produce a rotation of $\mathcal{T}_\sigma$ by $\frac{2\pi}{n+2}$ radians in the counterclockwise direction.

(1) For any diagonals $D_{i,j}$ with no incident diagonal counterclockwise to it, perform an edge flip at $D_{i,j}$ to get $D_{i-1,j-1}$. Continue to flip any such diagonals not previously flipped until no such diagonals remain.

(2) Choose an internal triangle T with a diagonal $D_{i,j}$ not previously flipped and admitting no incident diagonal $D_{i',j}$ counterclockwise to it.

Perform an edge flip at $D_{i,j}$ and then flip any diagonals not previously flipped that have no incident counterclockwise diagonals.

(3) If a diagonal $D_{i',j'}$ of T does have incident counterclockwise diagonals, then perform an edge flip at the counterclockwise-most of these incident diagonals. Flip any diagonals not previously flipped that now admit no incident counterclockwise diagonals.

(4) Repeat Step (3) until no diagonals counterclockwise to $D_{i',j'}$ remain. Perform an edge flip at $D_{i',j'}$. Once the second and third diagonals of T have been flipped, perform an edge flip at the initial diagonal previously belonging to T.

(5) Repeat Steps (3) and (4) starting with the remaining diagonals in the triangle corresponding to the counterclockwise diagonal flipped in Step (3). Continue until all possible diagonals have been flipped at.

EXAMPLE 2.4.3. *If the triangulation $\mathcal{T}_\sigma$ only contains diagonals of the form $D_{i,j_1}, \ldots, D_{i,j_{n-1}}$, then the instructions above reduce to simply performing edge flips in reverse order of indexing, starting with $D_{i,j_{n-1}}$ and ending with D_{i,j_1}. See Figure 2.7 for a more involved example with three internal triangles.*

REMARK 2.4.1. *Theorem 2.4.1 appears previously in work of Cormier, Dillery, Resh, Serhiyenko, and Whelan in the context of automorphisms of cluster algebras and quiver combinatorics. The argument below was written without knowledge of their work but follows a similar line of reasoning. The interested reader is referred to their paper [**CDR$^+$16**] for a more detailed discussion of the topic phrased in terms of maximal green sequences of cluster algebras.*

THEOREM 2.4.1 (Theorem 1.1, [**CDR$^+$16**]). *The number of edge flips required to realize a counterclockwise rotation of a triangulation $\mathcal{T}_\sigma$ of the $(n+2)$-gon by $\frac{2\pi}{n+2}$ is $n-1+t_\sigma$. The above instructions describe a sequence of $n-1+t_\sigma$ edge flips realizing such a rotation.*

PROOF. We first argue that the number of flips needed to rotate a triangulation is at least $n-1+t_\sigma$. Since no diagonal of our original triangulation is a diagonal of our rotated triangulation, a rotation of the triangulation $\mathcal{T}_\sigma$ requires at least $n-1$ edge flips, i.e. as many edge flips as diagonals of $\mathcal{T}_\sigma$. However, in an internal triangle, it is not possible to apply a single edge flip to any of the three sides (or any other diagonal) so that the result is a side of the rotated triangle, or indeed any diagonal of the rotated triangulation. This is because each of the three sides prevents the side immediately counterclockwise to it from rotating in a counterclockwise direction. If none of the internal triangles share a side, then the claim follows. Otherwise, we argue that any two triangles sharing an edge still require at least two extra edge flips to rotate. The only possible way we could have fewer is if we could perform an edge flip at the shared side and then rotate the two triangles with a single edge flip of each of the remaining sides. However, if we apply an edge flip

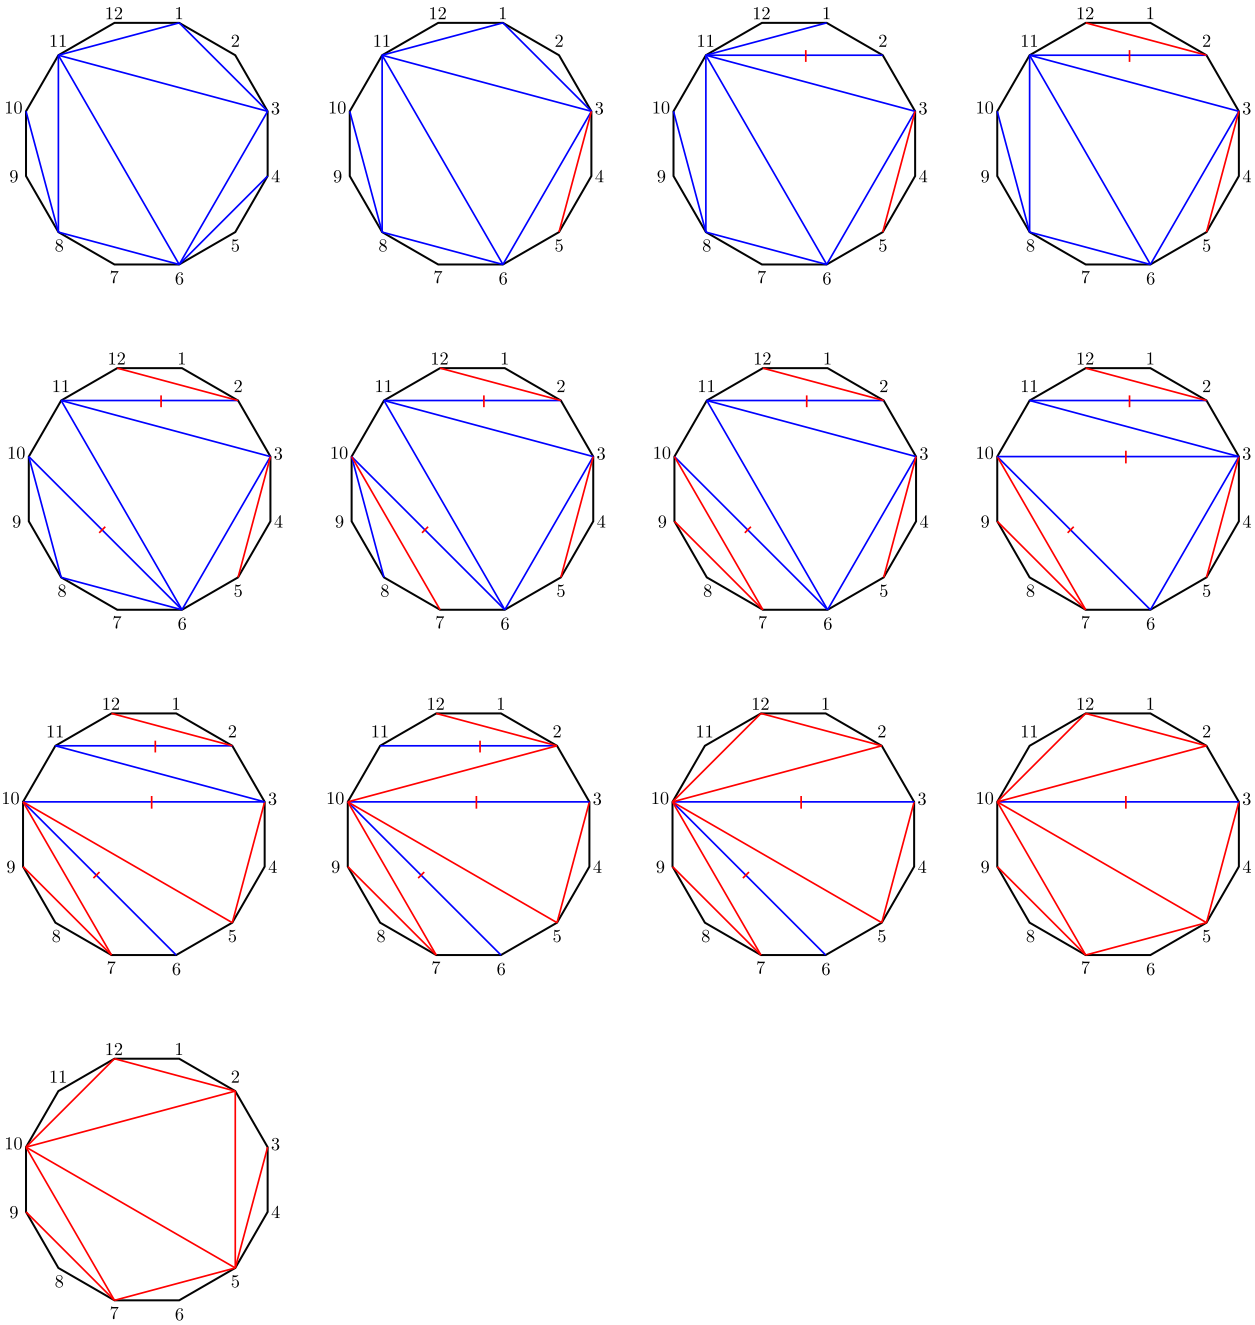

FIGURE 2.7. Counterclockwise rotation of a triangulation of the dodecagon by 10-1+3=12 edge flips. The red diagonals are diagonals of the rotated triangulation, while the blue diagonals with a red mark are diagonals that are the result of a previous edge flip but are not diagonals of the rotated triangulation.

at the shared side, then the remaining sides of the two triangles prevent the opposite pair from

achieving the desired rotation. Therefore, we must have at least $n-1+t_\sigma$ edge flips for a rotation of $\frac{2\pi}{n+2}$.

The algorithm given above describes a path in the flip graph of length $n-1+t_\sigma$ since we have two edge flips for a single diagonal in each internal triangle and one for every other diagonal. It remains to show that the result is a rotation of the initial triangulation $\mathcal{T}_\sigma$. In Step (1), an edge flip at a diagonal $D_{i,j}$ results in the diagonal $D_{i-1,j-1}$ precisely because there are no diagonals counterclockwise to it and therefore $D_{i,j}$ is a diagonal of the quadrilateral with sides $D_{i,j-1}, D_{j-1,j}, D_{i-1,j}, D_{i-1,i}$. It follows that each edge flip in Step (1) results in a diagonal of the rotated triangulation. If the triangulation $\mathcal{T}_\sigma$ has no internal triangles, then applying Step (1) to each of the $n-1$ diagonals results in the desired rotation.

Suppose that $\mathcal{T}_\sigma$ has at least one internal triangle. In Step (2), an edge flip at the diagonal $D_{i,j}$ in an internal triangle $\{D_{i,j}, D_{j,k}, D_{i,k}\}$ with no diagonal counterclockwise to it, results in the diagonal $D_{j-1,k}$. Once the remaining diagonals of the triangle have no incident counterclockwise diagonals, Step (4) applies an edge flip to them so that $D_{i,k}$ becomes $D_{i-1,j-1}$ and $D_{j,k}$ becomes $D_{j-1,k-1}$. Step (4) then flips $D_{j-1,k}$ to $D_{i-1,k-1}$. Crucially, the order of edge flips ensures that during Steps (2)-(4), we strictly decrease the number of counterclockwise incident diagonals to $D_{j,k}$ and $D_{i,k}$ at each step. After rotating our initial triangle, we can continue this process with the next internal triangle.

It remains to show that in Step (2), a diagonal $D_{i,j}$ of an internal triangle with no incident counterclockwise diagonals $D_{i',j}$ always exists If $D_{i,j}$ has an incident counterclockwise diagonal not belonging to an internal triangle, then Step (1) will apply an edge flip at such a diagonal so that it is no longer counterclockwise to $D_{i,j}$. If $D_{i,j}$ has an incident counterclockwise diagonal that belongs to an internal triangle, then there is some counterclockwise-most diagonal $D_{i,j'}$ also belonging to an internal triangle. Note that an edge flip at $D_{i,j'}$ removes one of the diagonals counterclockwise to $D_{i,j}$, so we can repeat this argument until we have performed an edge flip at all such diagonals. □

REMARK 2.4.2. *For triangulations that allow for a choice of ordering edge flips, it follows from a theorem of Pournin's [**Pou14**, Theorem 2] that naively proceeding with any of the equivalent options will still yield a geodesic. We can reinterpret this in the cluster algebraic setting as the fact*

that distant mutations commute. In this context, any geodesic path gives the mutations describing the cluster automorphism induced by the Kálmán loop. □

2.4.3. Edge flips in terms of permutations. In this subsection, we describe an edge flip at a diagonal $D_{j,l}$ of the triangulation $\mathcal{T}_\sigma$ in terms of the 312-avoiding permutation σ.

Let $\sigma \in S_n$ be a 312-avoiding permutation with corresponding triangulation given by the clip sequence bijection. Consider a quadrilateral with sides $D_{i,j}, D_{j,k}, D_{k,l}$ and $D_{i,l}$ appearing in the triangulation $\mathcal{T}_\sigma$. Figure 2.8 depicts this quadrilateral with two possible diagonals, $D_{i,k}$ and $D_{j,l}$ separating it into two triangles. An edge flip at one of these diagonals yields the other.

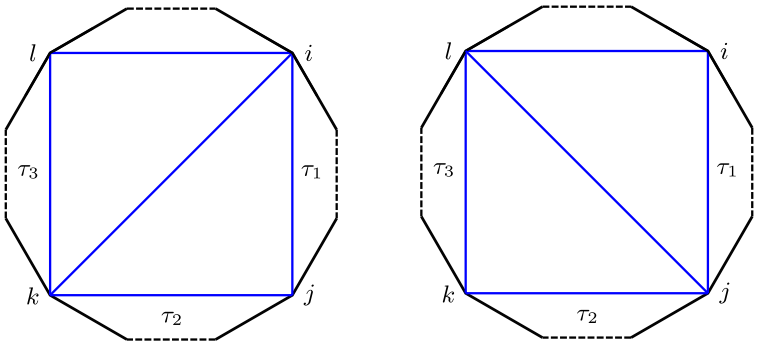

FIGURE 2.8. Schematic of an edge flip depicting the triangulation $\mathcal{T}_\sigma$ (left) and the result of applying an edge flip to $\mathcal{T}_\sigma$ at $D_{i,k}$ (right). The dotted lines represent arbitrarily many edges of the $(n+2)$-gon and the indices are chosen so that either $1 \leq i < j < k < l \leq n+1$ or $j < k < l = n+1, i = n+2$. The labels τ_1, τ_2, and τ_3 represent subwords of σ corresponding to different section of $\mathcal{T}_\sigma$. If any of the edges of the quadrilateral lie on the $(n+2)$-gon, then we consider the corresponding τ_i to be the empty word.

As in the orbit size algorithm, we can determine the structure of σ based on the existence of the edges of the quadrilateral. Specifically, σ admits subwords τ_1, τ_2, and τ_3, where the subword τ_1 contains the letters $i+1, \ldots, j-1$, the subword τ_2 contains letters $j+1, k-1$, and the subword τ_3 contains letters $k+1, \ldots l-1$. From this construction, we can deduce the effect on σ of a single edge flip at $D_{i,j}$.

THEOREM 2.4.2. *Given a triangulation $\mathcal{T}_\sigma$ containing a quadrilateral $ijkl$ with diagonal $D_{i,k}$, the 312-avoiding permutation σ is of the form $\ldots \tau_1 \tau_2 j \tau_3 k \ldots$. An edge flip at the diagonal $D_{i,k}$ yields a permutation of the form $\ldots \tau_1 \tau_2 \tau_3 k j \ldots$.*

PROOF. The theorem follows from the observation that each τ_i must contain at least one ear – a triangle of with edges $D_{i,i+1}, D_{i,i+2}, D_{i+1,i+2}$ – of the triangulation $\mathcal{T}_\sigma$. Therefore, under the clip sequence bijection, the word τ_i appears before τ_j if $i < j$. Moreover, the vertex labels j, k appear only after the two quadrants immediately adjacent to the vertex have been deleted under the clip sequence process. Thus, the two 312-avoiding permutations corresponding to the triangulation $\mathcal{T}_\sigma$ and the triangulation resulting from applying an edge flip are precisely of the form described. □

EXAMPLE 2.4.4. *Consider the permutation* $\sigma = 1\,5\,4\,3\,6\,2$. *If we wish to apply an edge flip to the diagonal* $D_{2,6}$, *then we can identify the vertex labels of the relevant quadrilateral as* $i = 2, j = 3, k = 6$, *and* $l = 7$. *This immediately tells us that* τ_1 *and* τ_3 *are both empty and* τ_2 *is the subword* $5\,4$. *Therefore, Theorem 2.4.2 above implies that we simply interchange* j *and* k *to get the resulting permutation* $\mu(\sigma) = 1\,5\,4\,6\,3\,2$. *See Figure 2.9 for the triangulations* $\mathcal{T}_\sigma$ *and the triangulation resulting from the edge flip.* □

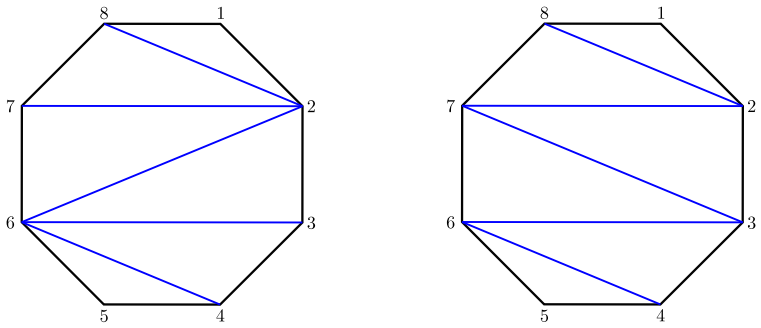

FIGURE 2.9. An edge flip at the diagonal $D_{2,6}$ in the triangulation $\mathcal{T}_{1\,5\,4\,3\,6\,2}$ yields the permutation $1\,5\,4\,6\,3\,2$.

Together with Theorem 2.4.1, the above computation gives an explicit combinatorial construction of Kálmán loop in terms of geodesics paths of the flip graph and the corresponding behavior of 312-avoiding permutations.

CHAPTER 3

Weave realizability for D-type

In this chapter, we prove Theorem 1.1.2 by showing that each quiver mutation is weave realizable. By definition, a sequence of quiver mutations for $Q(\Gamma_0(D_n), \{\gamma_i^{(0)}\})$ is weave realizable if each quiver mutation in the sequence can be realized as a Legendrian weave mutation at an embedded homology cycle of the 3-graph. For an alternative construction of the exact Lagrangian fillings of $\lambda(D_n)$, as well as $\lambda(E_n), \lambda(\tilde{D}_n)$, and $\lambda(\tilde{E}_n)$, see [**ABL22**].

The following definitions relate the algebraic intersections of cycles to geometric intersections in the context of 3-graphs.

DEFINITION 3.0.1. *A 3-graph Γ with associated homology basis $\{\gamma_i\}$, $i \in [1, b_1(\Lambda(\Gamma)]$ of $H_1(\Lambda(\Gamma); \mathbb{Z})$ is* sharp at a cycle γ_j *if, for any other cycle $\gamma_k \in \{\gamma_i\}$, the geometric intersection number of γ_j with γ_k is equal to the algebraic intersection $\langle \gamma_j, \gamma_k \rangle$.*

Γ is locally sharp *if, for any cycle $\gamma \in \{\gamma_i\}$, there exists a sequence of Legendrian Surface Reidemeister moves taking Γ to some other 3-graph Γ' such that Γ' is sharp at the corresponding cycle $\gamma' \in H_1(\Lambda(\Gamma'); \mathbb{Z})$.*

A 3-graph Γ with a set of cycles Γ is sharp *if Γ is sharp at all $\gamma_i \in \{\gamma_i\}$.* □

For 3-graphs that are not sharp, it is possible that a sequence of mutations will cause a cycle to become immersed. This is the only obstruction to weave realizability. Therefore, sharpness is a desirable property for our 3-graphs, as it simplifies our computations and helps us avoid creating immersed cycles. We will not be able to ensure sharpness for all $\Gamma(D_n)$ that arise as part of our computations, (e.g., see the type III.i normal form in Figure 3.2) but we will be able to ensure that each of our 3-graphs is locally sharp.

3.1. Weave Realizability for $\lambda(D_n)$

The following result is slightly stronger than the statement of Theorem 1.1.2, as we are able to show that each 3-graph in our sequence of mutations is locally sharp.

THEOREM 3.1.1. *Let $\mu_{v_1}, \ldots, \mu_{v_k}$ be a sequence of quiver mutations, with initial quiver $Q(\Gamma_0(D_n), \{\gamma_i^{(0)}\})$. Then, there exists a sequence $\Gamma_0(D_n), \ldots, \Gamma_k(D_n)$ of 3-graphs such that*

i. *$\Gamma_{j-1}(D_n)$ is related to $\Gamma_j(D_n)$ by mutation at a cycle γ_j and by Legendrian Surface Reidemeister moves I, II, and III. The cycle γ_j represents the vertex v_j in the intersection quiver and it is given by one of the cycles in the initial basis $\{\gamma_i^{(0)}\}$ after mutation and Reidemeister moves.*

ii. *$\Gamma_j(D_n)$ is sharp at γ_j.*

iii. *$\Gamma_j(D_n)$ is locally sharp.*

iv. *The basis of cycles for $\Gamma_j(D_n)$, obtained from the initial basis $\{\gamma_i^{(0)}\}$ by mutation and Reidemeister moves, consists entirely of short Y-cycles and short I-cycles.*

The conditions ii-iv allow us to continue to iterate mutations after applying a small number of simplifications at each step. Theorem 1.1.2 thus follows from Theorem 3.1.1.

PROOF. We proceed by organizing the 3-graphs arising from any sequence of mutations of $\Gamma_0(D_n)$ into four types, in line with the organization scheme introduced by Vatne for quivers of D_n-type [**Vat10**]. Vatne's classification of quivers in the mutation class of D_n-type uses the configuration of a certain subquiver to define the different types. Outside of that subquiver, there are a number of disjoint subquivers of A_n-type that are referred to as A_n tail subquivers. We will refer to the corresponding cycles in the 3-graph as A_n tail subgraphs, or simply A_n tails when it is clear from context whether we are referring to the quiver or the 3-graph. For each type, Vatne describes the results of quiver mutation at different vertices, which can depend on the existence of A_n tail subquivers. See Figures 3.3, 3.9, 3.13, and 3.17 for the four types and their mutations.

Notation. As mentioned in the previous section, cycles are pictured as colored edges for the sake of visual clarity. Throughout this section, we denote all of the dark green cycles by γ_1, light green cycles by γ_2, orange cycles by γ_3, light blue cycles by γ_4, pink cycles by γ_5, purple cycles by γ_6, and olive cycles by γ_7. With this notation, γ_i will correspond to the vertex labeled by v_i in the quivers given below.

***A_n* Tails.** We briefly describe the behavior of the A_n tail subquivers, as given in [**Vat10**], in terms of weaves. Any of the n vertices in an A_n tail subquiver can have valence between 0 and 4. Cycles in the quiver are oriented with length 3. If a vertex v has valence 3, then two of the edges form part of a 3-cycle, while the third edge is not part of any 3-cycle. If v has valence 4, then two of the edges belong to one 3-cycle and the remaining two edges belong to a separate 3-cycle.

Any A_n tail of the quiver can be represented by a sharp configuration of n l-cycles in the 3-graph. See Figure 3.1 for an identification of l-cycles with quiver vertices of a given valence. Mutation at any vertex v_i in the quiver corresponds to mutation at the l-cycle γ_i in the 3-graph, so it is readily verified that mutation preserves the number of l-cycles and requires no application of Legendrian Surface Reidemeister moves to simplify. The sequences of mutations given in the remainder of the proof As a consequence, any sequence of A_n tail mutations is weave realizable, and a sharp 3-graph remains sharp after mutation at A_n tail l-cycles that only intersect other A_n tail l-cycles.

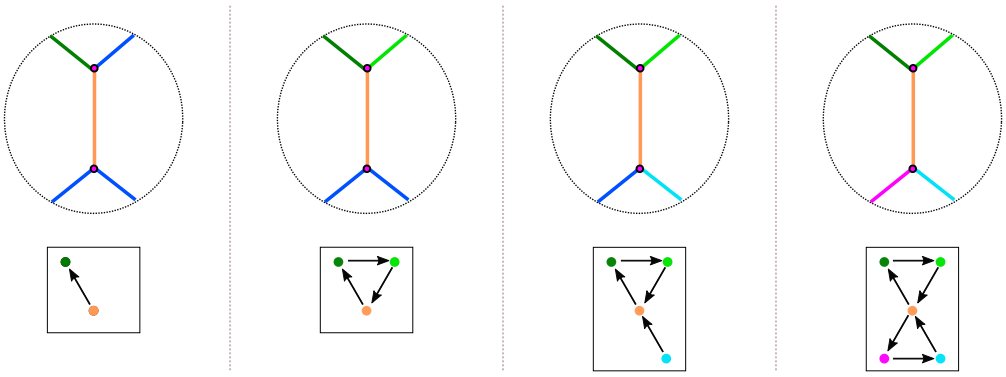

FIGURE 3.1. All possible arrangements of l-cycles in an A_n tail of the 3-graph corresponding to a given vertex in the A_n tail subquiver of valence between 0 and 4.

Normal Forms. For each of the four types of D_n quivers described in [**Vat10**], we give a set of specific subgraphs of $\Gamma(D_n)$, which we refer to as normal forms. These normal forms are pictured in Figure 3.2. We indicate the possible existence of A_n tail subgraphs by an unfilled circle. In our discussion below, we will say that an edge of the 3-graph carries a cycle if it is part of a homology cycle. We will generally use this terminology to specify which edges cannot carry a cycle.

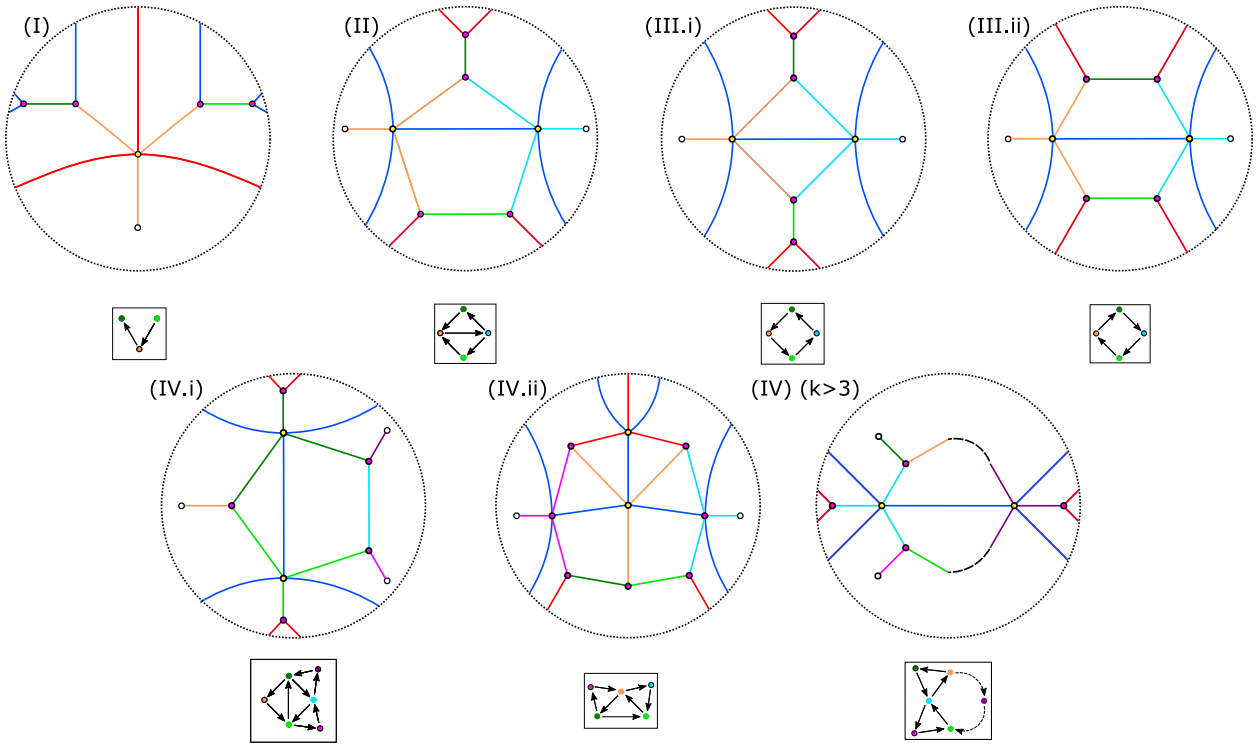

FIGURE 3.2. Normal forms labeled by their type. The possible addition of l-cycles corresponding to A_n tails of the quiver are represented by unfilled circles appended to the end of edges that do not intersect the boundary.

For each possible quiver mutation, we describe the possible mutations of the 3-graph and show that the result matches the quiver type and retains the properties listed in Theorem 3.1.1 above. In addition, the Legendrian Surface Reidemeister moves we describe ensure that the A_n tail subgraphs continue to consist solely of short l-cycles. If the mutation results in a long l-cycle or pair of long l-cycles connecting our A_n tail to the rest of the 3-graph, we can simplify by applying a sequence of n push-throughs to ensure that these are all short l-cycles. It is readily verified that we can always

do this and that no other simplifications of the A_n tails are required following any other mutations. We include A_n tail cycles only where relevant to the specific mutation. In our computations below, we generally omit the final steps of applying a series of push-throughs to make any long I or Y-cycles into short I or Y-cycles. Figure 3.8 provides an example where these push-throughs are shown for both an I-cycle and a Y-cycle.

In order to simplify the overall presentation of the normal forms and the computations below, we allow for the following variations in the Type I and Type IV cases. In the Type I case, mutating at either of the short I-cycles γ_1 or γ_2 in the Type I normal form produces one of four possible configurations of the cycles γ_1, γ_2, and γ_3 in a 3-graph corresponding to a Type I quiver. Since these mutations are readily computed, we simplify our presentation by giving a single normal form rather than four, and describing the relevant mutations of two of the four possible 3-graphs in figures 3.4, 3.5, 3.6, and 3.7. The remaining cases can be seen by swapping the cycle(s) to the left of the short Y-cycle with the cycle(s) to the right of it. This symmetry corresponds to reversing all of the arrows in the quiver. In general, we will implicitly appeal to similar symmetries of the normal form 3-graphs to reduce the number of cases we must consider. In the Type IV case, the edge(s) corresponding to γ_3, γ_5 or γ_6 need not carry a cycle. See the discussion of Type IV quiver mutations below for a more detailed description.

Type I. We start with 3-graphs, always endowed with a homology basis, whose associated intersection quivers are a Type I quiver. See Figure 3.3 for the relevant quiver mutations.

i. (Type I to Type I) There are two possible Type I to Type I mutations of 3-graphs depicted in Figure 3.4 (left) and (right). As shown in Figure 3.4 (left), mutation at γ_1 only affects the sign of the intersection of γ_1 with the γ_3. This reflects the fact that the corresponding quiver mutation has only reversed the orientation of the edge between v_1 and v_3. Mutating at any other I-cycle is equally straightforward and yields a Type I to Type I mutation as well.

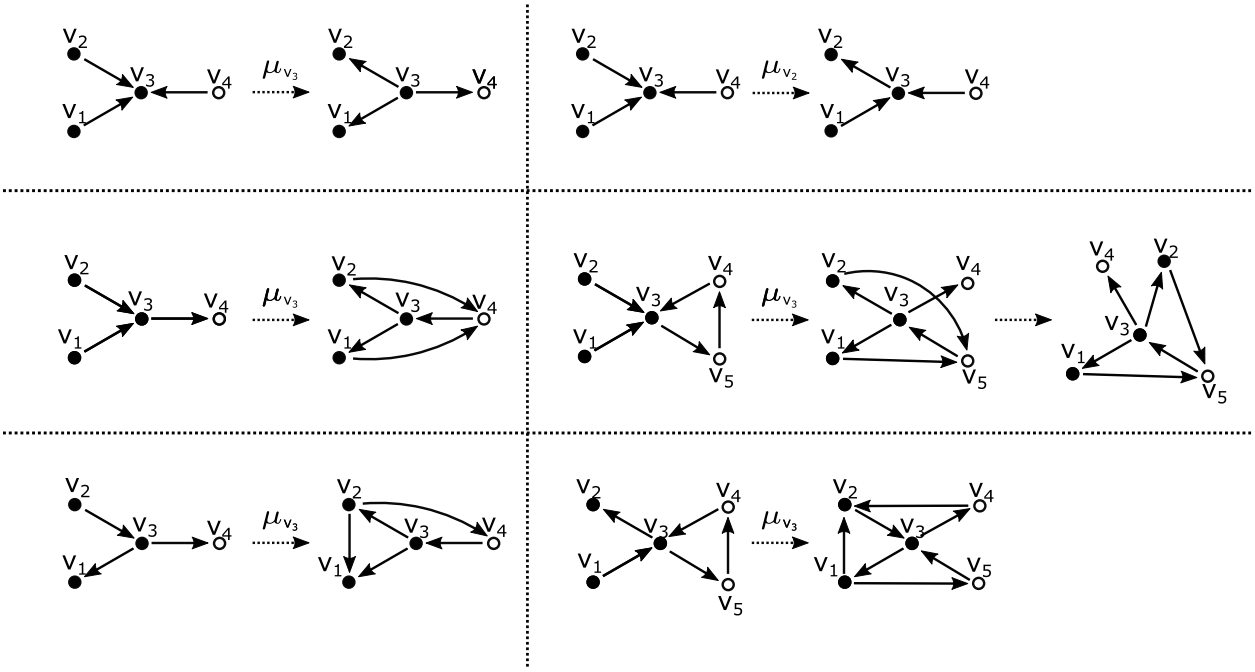

FIGURE 3.3. From top to bottom, two Type I to Type I quiver mutations, Type I to Type II quiver mutations, and Type I to Type IV quiver mutations. The arrow labeled by μ_{v_i} indicates mutation at the vertex v_i. Unfilled circles represent potential A_n tails. In each line, the first quiver mutation shows the case where v_3 is only adjacent to one A_n tail vertex, while the second quiver mutation shows the case where v_3 is adjacent to two A_n tail vertices. Note that reversing the direction of all of the arrows simultaneously before mutating gives additional possible quiver mutations of the same type.

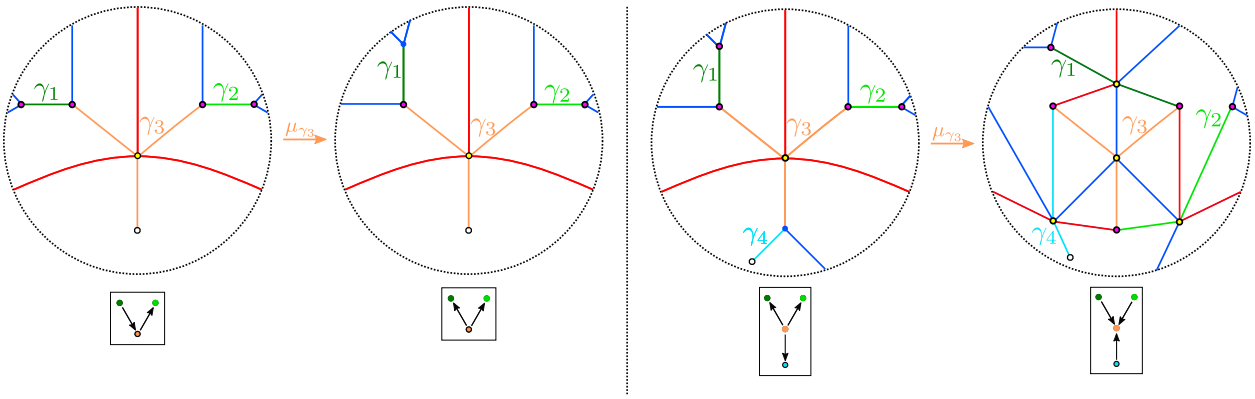

FIGURE 3.4. Type I to Type I mutation. Arrows labeled by μ indicate mutation at a cycle of the same color.

ii. (Type I to Type I) For the second possible Type I to Type I mutation, we proceed as pictured in Figure 3.4 (right). Mutation at γ_3 does not create any new additional geometric

or algebraic intersections. Instead, it takes positive intersections to negative intersections and vice versa. This is reflected in the quivers pictured underneath the 3-graphs, as the orientation of edges has reversed under the mutation. As explained above, we could simplify the resulting 3-graph by applying a push-through move to each of the long I-cycles to get a sharp 3-graph where the homology cycles are made up of a single short Y-cycle and some number of short I-cycles.

iii. (Type I to Type II) In Figure 3.5 we consider the cases where the Y-cycle γ_3 intersects one I-cycle (top) or two I-cycles (bottom) in the A_n tail subgraph. Mutation at γ_3 introduces an intersection between γ_2 and γ_4 that causes the second 3-graph in of each mutation sequences to no longer be sharp. Applying a push-through to γ_2 resolves this intersection so that the geometric intersection between γ_2 and γ_4 matches their algebraic intersection. This simplification ensures that the result of μ_{γ_3} is a sharp 3-graph that matches the Type II normal form. If we compare the mutations in the top and bottom sequences, we can see that the presence of the A_n tail cycle γ_5 does not affect the computation.

iv. (Type I to Type IV.i) We now consider the first of two Type I to Type IV mutations, shown in Figure 3.6. Starting with the configuration of cycles at the left of each sequence and mutating at γ_3 causes γ_1 and γ_2 to cross. Applying a push-through to γ_1 or to γ_2 (not pictured) simplifies the resulting intersection and yields a Type IV.i normal form made up of the cycles $\gamma_1, \gamma_2, \gamma_3$, and γ_4. The sequences on the top and bottom of Figure 3.6 differ only by the presence of the A_n tail cycle γ_5.

v. (Type I to Type IV.ii) In Figure 3.7, we consider the cases where γ_1 intersects one I-cycle (top) or two I-cycles (bottom) in the A_n tail subgraph, as we did in the Type I to Type II case. As in the Type I to Type II case, we must apply a push-through to resolve the new intersections between that cause the second 3-graph in each sequence to fail to be

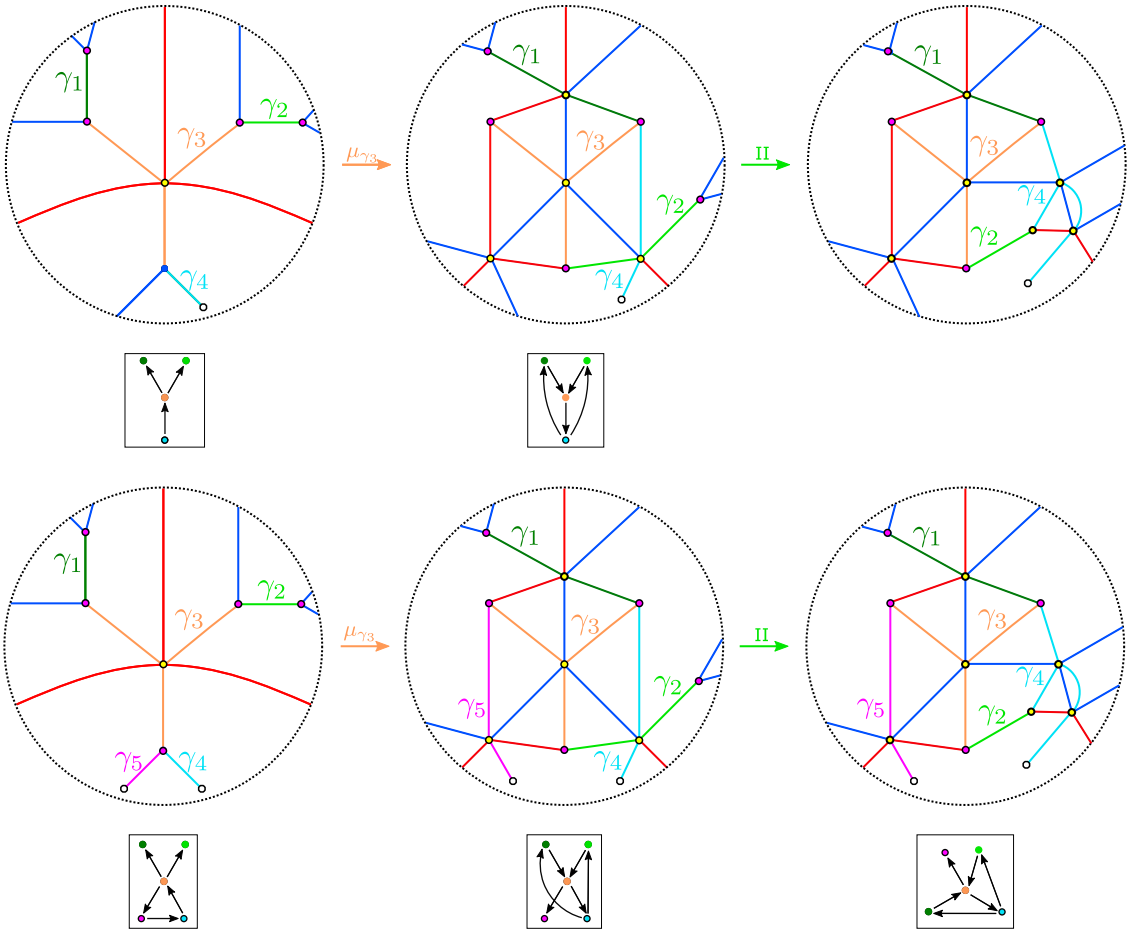

FIGURE 3.5. Type I to Type II mutations. Legendrian Surface Reidemeister are moves labeled as in Theorem 2, Figure 1.8.

sharp. When we include both γ_4 and γ_5 in the sequence on the right, we get two new intersections after mutating, and therefore require two push-throughs. Note that in the IV.ii case, we must first apply the push-through to γ_1 and γ_2 in order to ensure that we can apply a push-through to any additional cycles in the A_n tail subgraph. This causes the Y-cycles of the graph to correspond to different vertices in the quiver than in the Type IV.i normal form, which is the main reason we distinguish between the normal forms for Type IV.i and Type IV.ii.

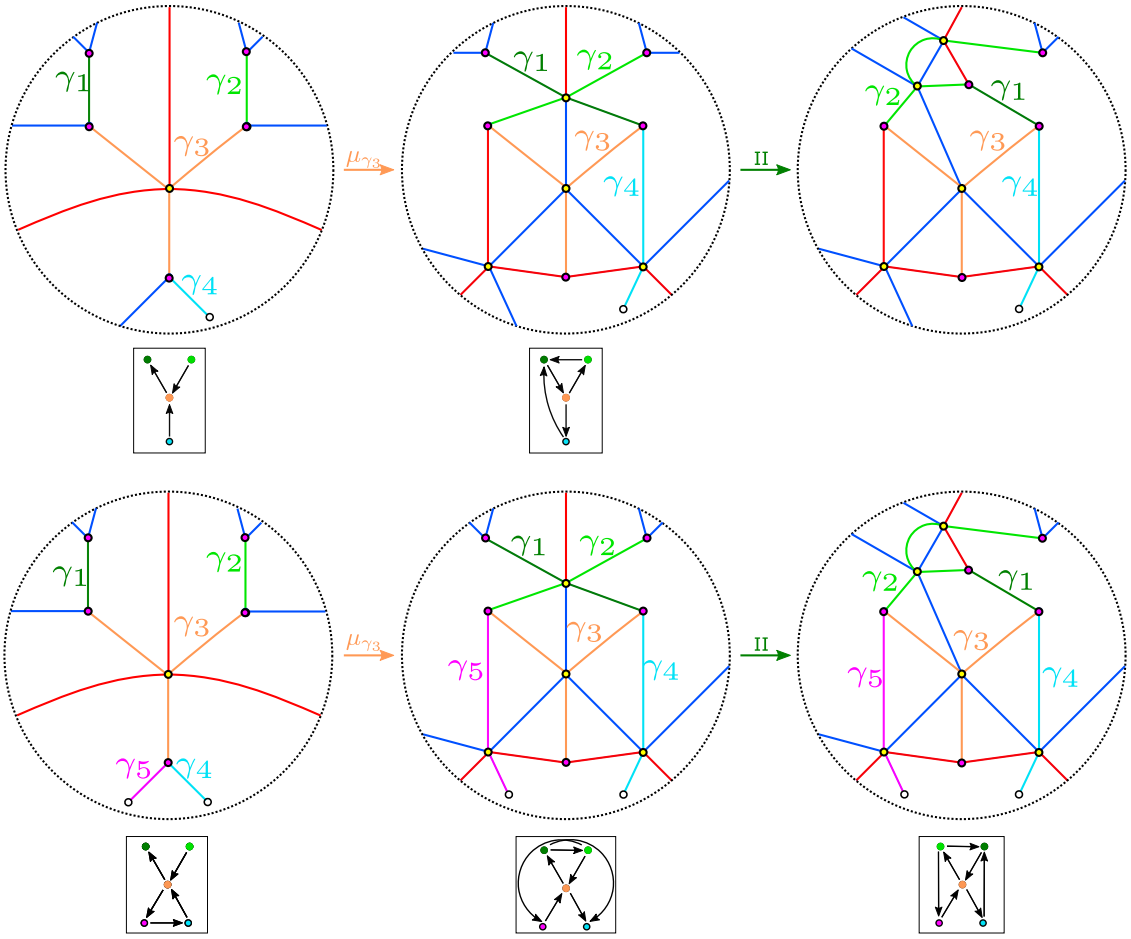

FIGURE 3.6. Type I to Type IV.i mutations.

In Figure 3.8 we show how to apply push-throughs to completely simplify the long I- and Y-cycles pictured in the Type I to Type IV.ii graph. As mentioned above, these push-throughs are identical to any other computation required to simplify our resulting 3-graphs to a set of short I- and Y-cycles.

The above cases describe all possible mutations of the Type 1 normal form. Each of these mutations yields a sharp 3-graph with short I-cycles and Y-cycles, as desired.

Type II. We now consider mutations of our Type II normal form. See Figure 3.9 for the relevant quivers. As shown in the figure, performing a quiver mutation at the 2-valent vertices labeled by v_1 or v_2 yields a Type III quiver, while a quiver mutation at the vertices labeled v_3 or

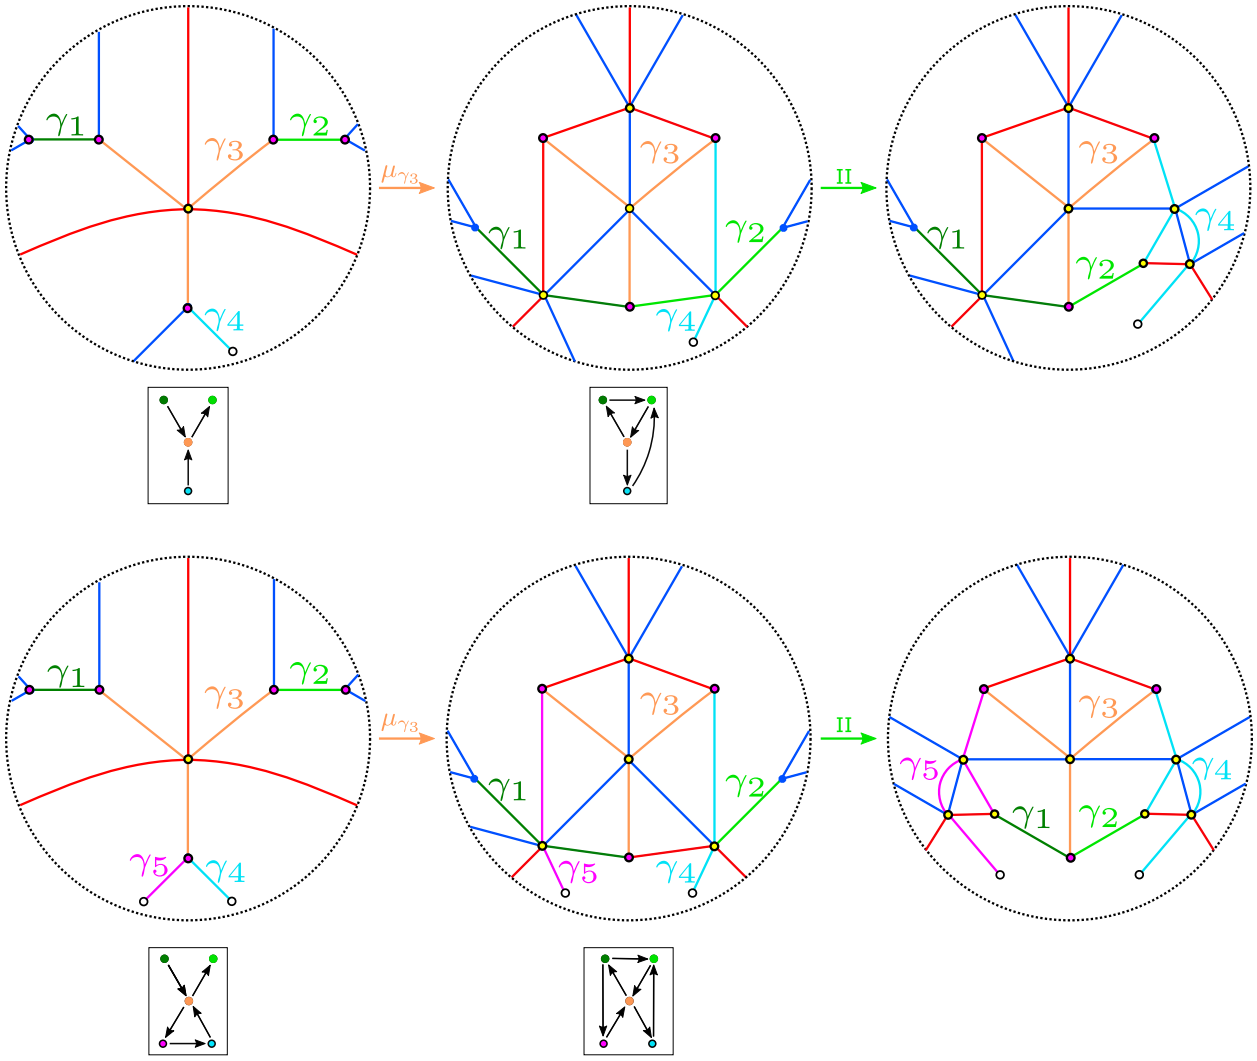

FIGURE 3.7. Type I to Type IV.ii mutations.

v_4 yields either another Type II quiver or a Type I quiver, depending on the intersection of v_3 or v_4 with any A_n tail subquivers.

i. (Type II to Type I) We first consider the sequence of 3-graphs pictured in Figure 3.10. Mutation at γ_4 results in a new geometric intersection between γ_2 and γ_3 even though their algebraic intersection is zero. We can resolve this by applying a reverse push-through at

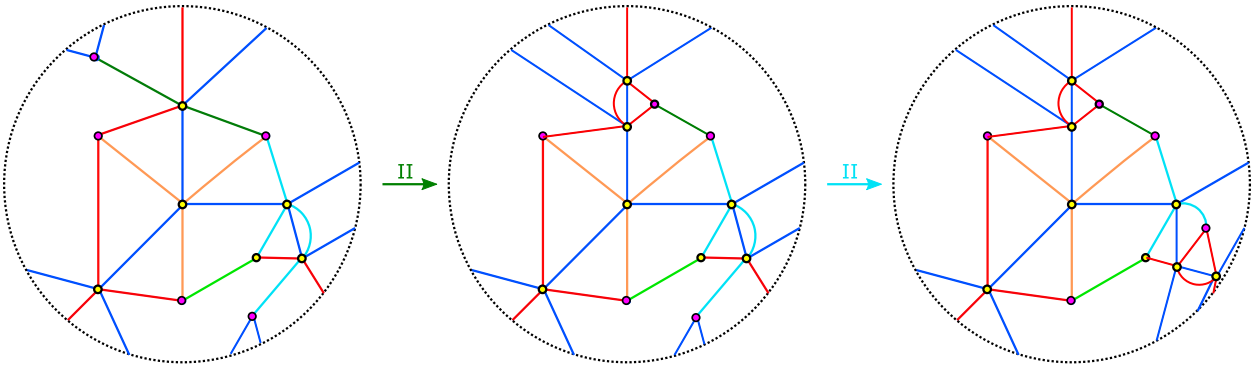

FIGURE 3.8. Push-through examples. The first push-through move simplifies the long I-cycle labeled γ_1, while the second simplifies the long Y-cycle labeled γ_4.

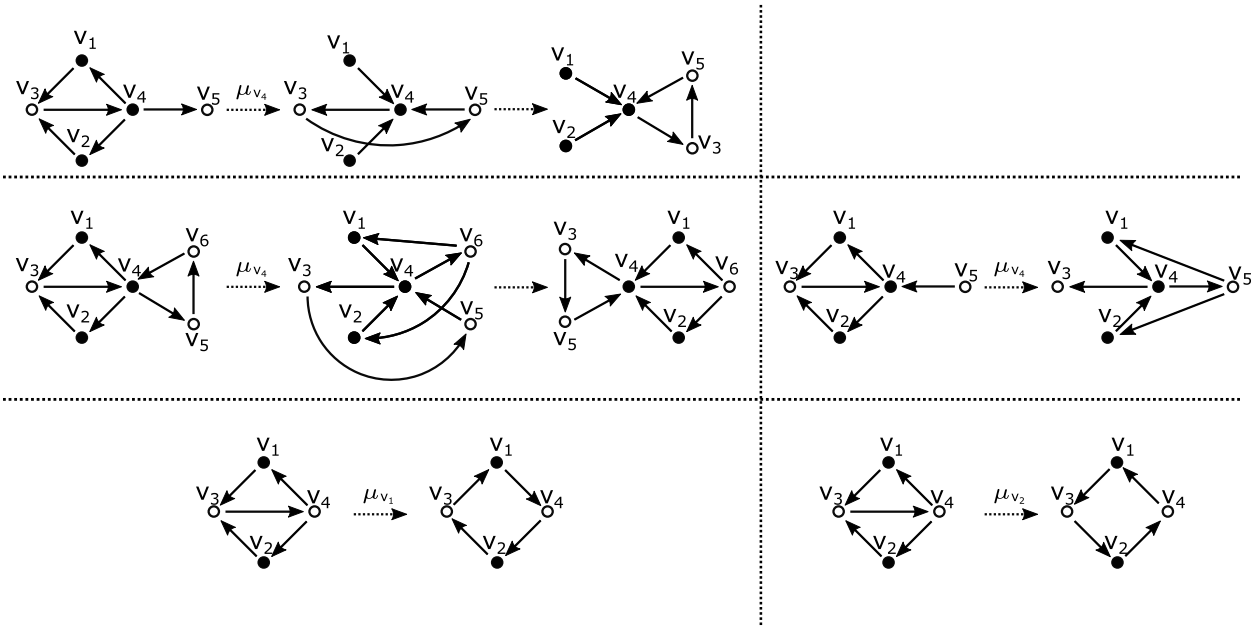

FIGURE 3.9. From top to bottom, Type II to Type I, Type II to Type II, and Type II to Type III quiver mutations.

the trivalent vertex where γ_2 and γ_3 meet. The resulting 3-graph is sharp, as γ_2 and γ_3 no longer have any geometric intersection. This computation is identical if γ_3 were to intersect a single A_n tail cycle and we mutated at γ_3 instead. Note that here we require the red edge adjacent labeled e to not carry a cycle, as specified by our normal form.

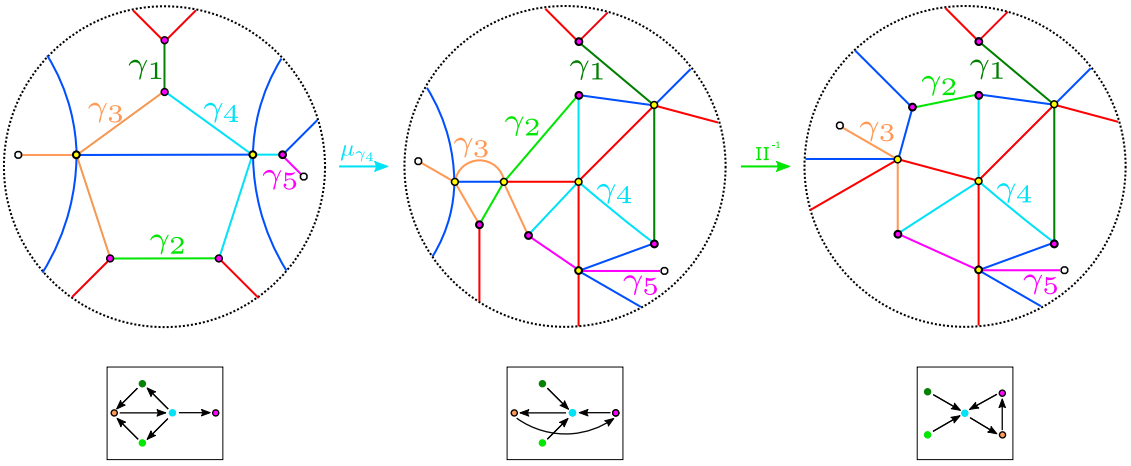

FIGURE 3.10. Type II to Type I mutations. The red e labels an edge in the 3-graph that does not carry a cycle.

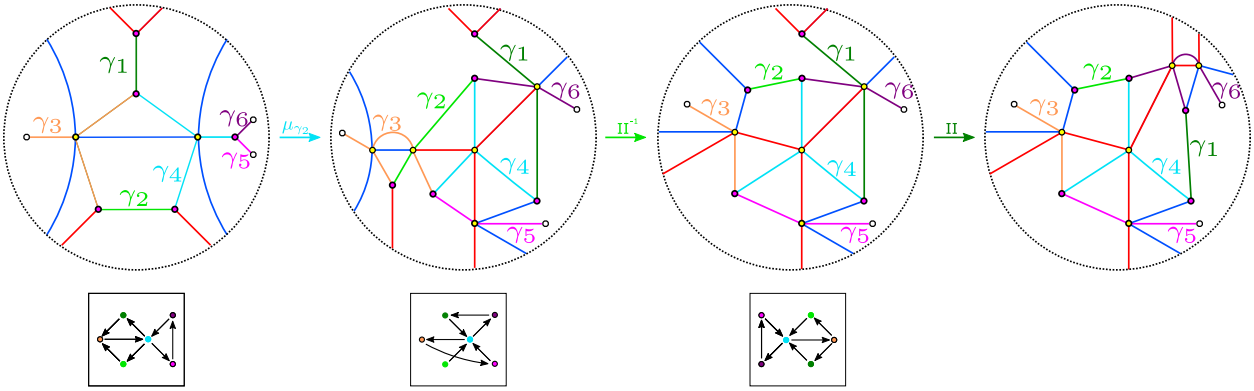

FIGURE 3.11. Type II to Type II mutations.

ii. (Type II to Type II) We now consider the sequence shown in Figure 3.11. After mutating at γ_4, we have the same intersection between γ_2 and γ_3 as in the previous case. We again resolve this intersection by applying a reverse push-through at the same trivalent vertex. In this case, we also have an intersection between γ_1 and γ_6, which we resolve via push-through of γ_1. As a result, γ_6 becomes a Y-cycle, and the Type II normal form is now made up of the cycles γ_1, γ_2, γ_4, and γ_6, while γ_3 becomes an A_n tail cycle.

iii. (Type II to Type III.i) Mutation at γ_1 or γ_2 in the Type II normal form yields either of the Type III normal forms. In the sequence on the left of Figure 3.12, mutation at γ_2 leads

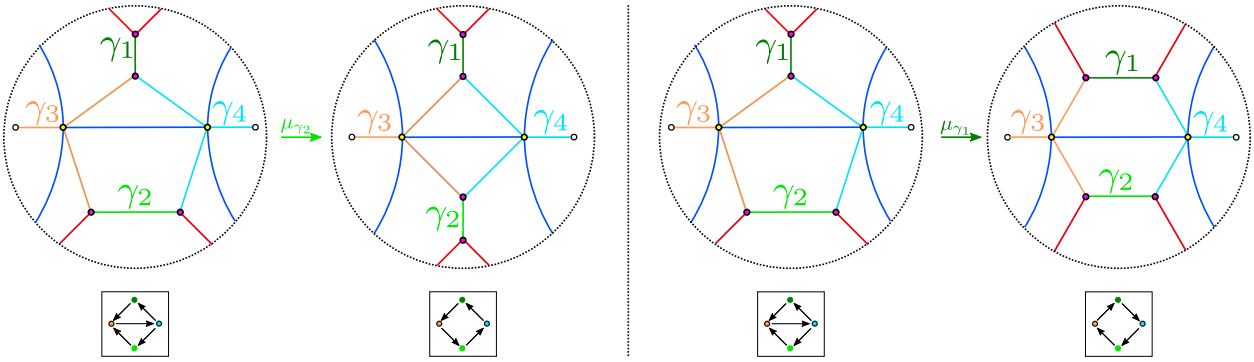

FIGURE 3.12. Type II to Type III mutations.

to a geometric intersection between γ_3 and γ_4 at two trivalent vertices. Since the signs of these two intersections differ, the algebraic intersection $\langle \gamma_3, \gamma_4 \rangle$ is zero, so the resulting 3-graph is not sharp. However, it is sharp at γ_1 and γ_2, and applying a flop to the 3-graph removes the geometric intersection between γ_3 and γ_4 at the cost of introducing the same intersection between γ_1 and γ_2. Therefore, applying the flop does not make the 3-graph sharp, but it does show that the 3-graph resulting from our mutation is locally sharp at every cycle.

iv. (Type II to Type III.ii) In the sequence on the right of Figure 3.12, mutation at γ_1 yields a sharp 3-graph that matches the Type III.ii normal form.

Type III: Figure 3.13 illustrates the Type III quiver mutations. Figures 3.14, 3.15, and 3.16 depict the corresponding Legendrian mutations of the Type III normal forms.

i. (Type III.i to Type II) We first consider the sequence of 3-graphs in Figure 3.14 (left). Mutating at γ_1 or γ_2 immediately yields a Type II normal form. Mutating at γ_1 and γ_2 in succession yields a Type III.ii normal form. Note that if the 3-graph were not sharp at γ_1 or γ_2 we would first need to apply a flop. We can always apply this move because the 3-graph is locally sharp at each of its cycles. See the Type III.i to Type IV.i subcase

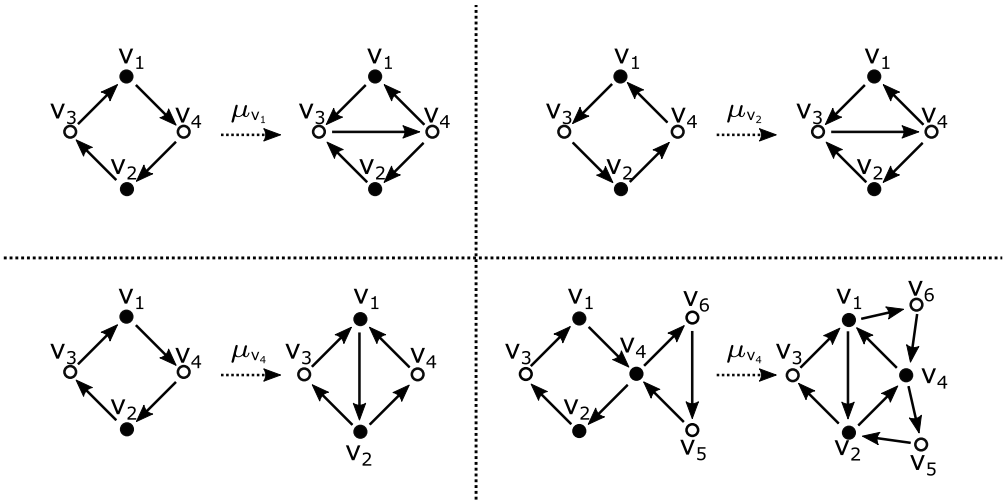

FIGURE 3.13. Type III to Type II quiver mutations (top) and Type III to Type IV quiver mutations (bottom).

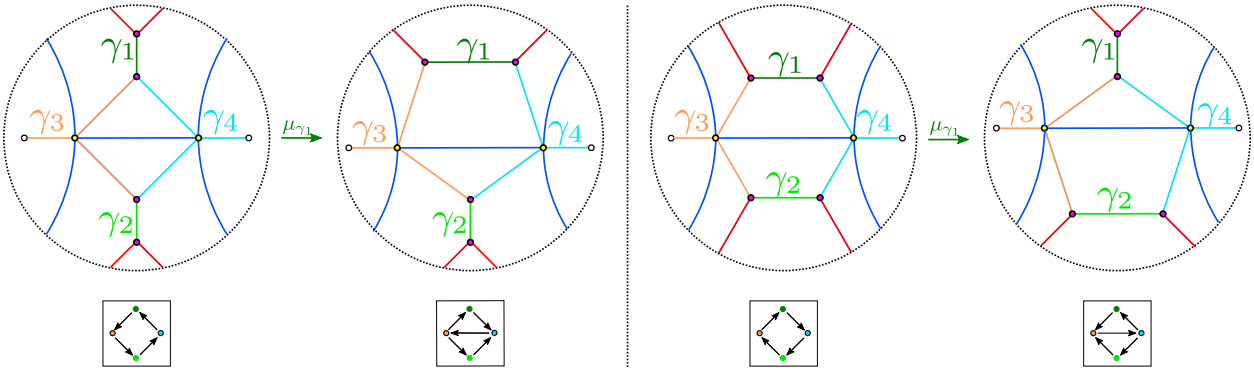

FIGURE 3.14. Type III.i to Type II mutations (left) and Type III.ii to Type II mutations (right).

below for an example where we demonstrate this move.

ii. (Type III.ii to Type II) In the sequence on the right of Figure 3.14, mutation at either γ_1 or γ_2 yields a Type II normal form. Mutation at γ_1 and γ_2 in succession yields a Type III.i normal form. Therefore, applying these two moves in succession can take us between both of our Type III normal forms.

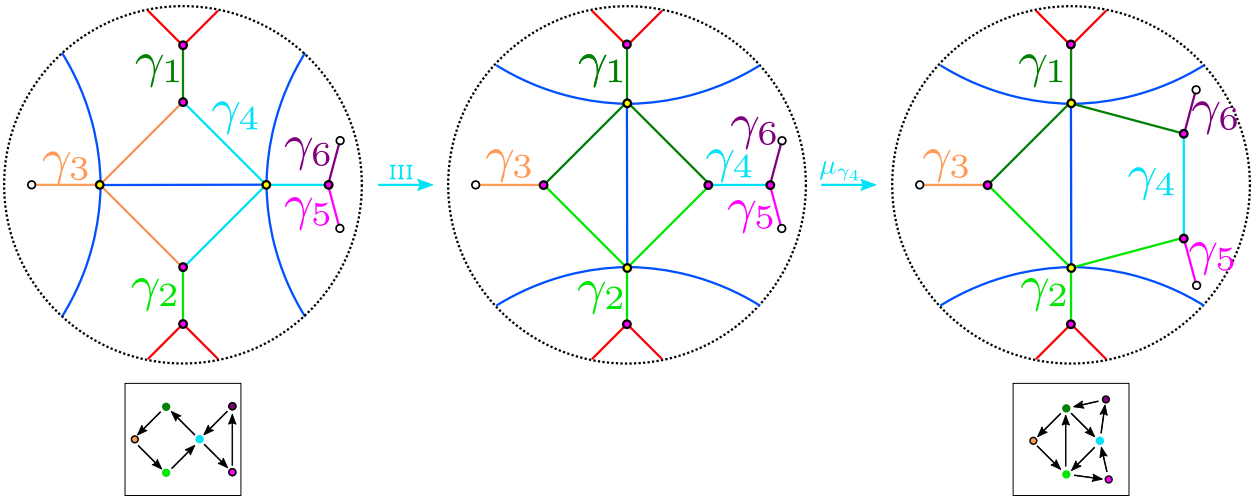

FIGURE 3.15. Type III.i to Type IV mutations.

iii. (Type III.i to Type IV) We now consider the sequence of 3-graphs in Figure 3.15. Since the initial 3-graph is not sharp at γ_4, we must first apply a flop before mutating. After applying this flop, γ_4 is a short I-cycle and the 3-graph is sharp at γ_4. Mutating at γ_4 then yields a Type IV.i normal form. The short I-cycles γ_5 and γ_6 are included to indicate where any A_n tail cycles would be sent under this mutation.

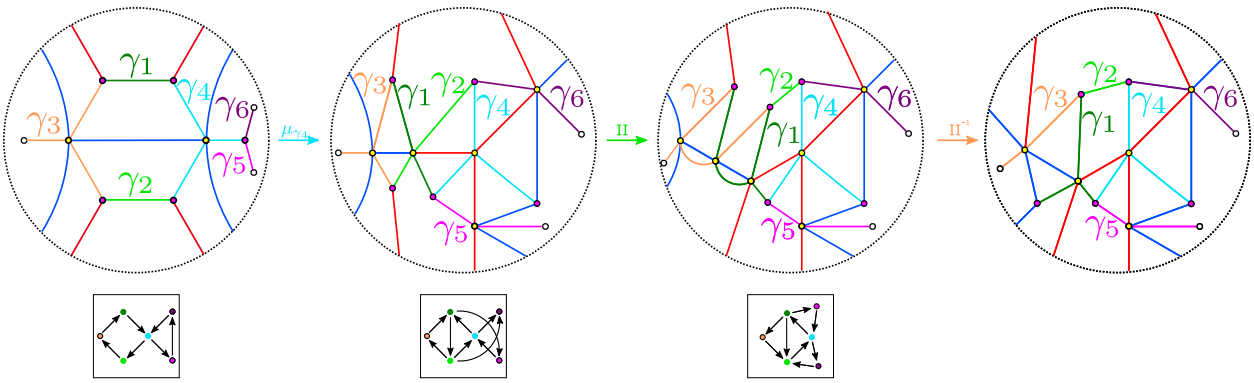

FIGURE 3.16. Type III.ii to Type IV mutations.

iv. (Type III.ii to Type IV) In Figure 3.16, mutation at γ_4 causes γ_1 and γ_2 to cross while still intersecting γ_3 and γ_4 at either end. We resolve this by first applying a push-through

to γ_2 and then applying a reverse push-through to the trivalent vertex where γ_1 and γ_3 intersect a red edge. This results in a sharp 3-graph with γ_1, γ_2, γ_3, and γ_4 making up the Type IV normal form. We again include γ_5 and γ_6 as cycles belonging to a potential A_n tail subgraph in order to show where the A_n tail cycles are sent under this mutation.

Type IV: Figure 3.17 illustrates all of the relevant Type IV quivers and their mutations. In general, the edges of a Type IV quiver have the form of a single $k-$cycle with the possible existence of 3-cycles or outward-pointing "spikes" at any of the edges along the $k-$cycle. At the tip of each of these spikes is a possible A_n tail subquiver. We will refer to a vertex at the tip of any of the spikes (e.g., the vertex v_3 in Figure 3.17) as a spike vertex and any vertex along the $k-$cycle will be referred to as a $k-$cycle vertex. A homology cycle corresponding to a spike vertex will be referred to as a spike cycle. Mutating at a spike vertex increases the length of the internal $k-$cycle by one, while mutating at a $k-$cycle vertex decreases the length by 1, so long as $k > 3$. Figures 3.18, 3.19, 3.20, and 3.21 illustrate the corresponding mutations of 3-graphs for Type IV to Type I and Type IV to Type III when $k = 3$.

i. (Type IV.i to Type I) We first consider the sequence of 3-graphs in Figure 3.18. Mutation at γ_1 causes γ_2 and γ_4 to cross. Application of a reverse push-through at the trivalent vertex where γ_2 and γ_4 intersect a red edge removes this crossing and yields a Type I normal form where γ_1 is the sole Y-cycle.

ii. (Type IV.ii to Type I) Mutation at γ_3 in Figure 3.19 yields a 3-graph with geometric intersections between γ_1 and γ_5 and between γ_2 and γ_4. The application of reverse push-throughs at the trivalent vertex intersections of γ_1 with γ_5 and γ_2 with γ_4 removes these geometric intersections, resulting in a Type I normal form where γ_1 is the sole Y-cycle. We also apply a candy twist (Legendrian Surface Reidemeister move I) to simplify the

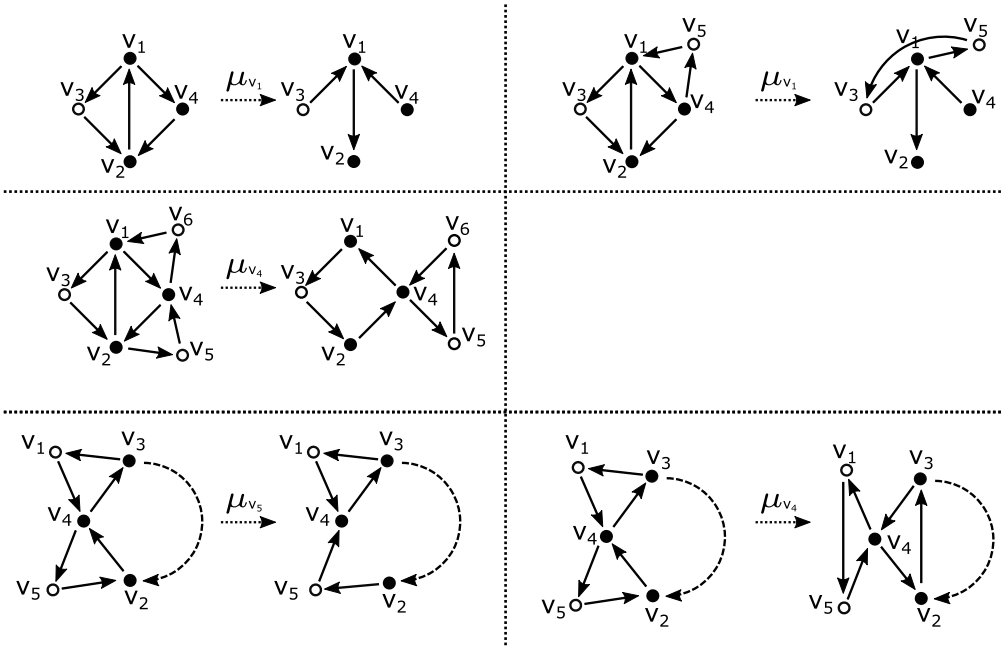

FIGURE 3.17. From top to bottom, Type IV to Type I, Type IV to Type III, Type IV spike vertex (left) and cycle vertex (right) quiver mutations. The presence or absence of the A_n tail vertices v_6 and v_7 in the quiver mutation depicted in the first column, third row correspond to the presence or absence of spikes appearing in the resulting quiver.

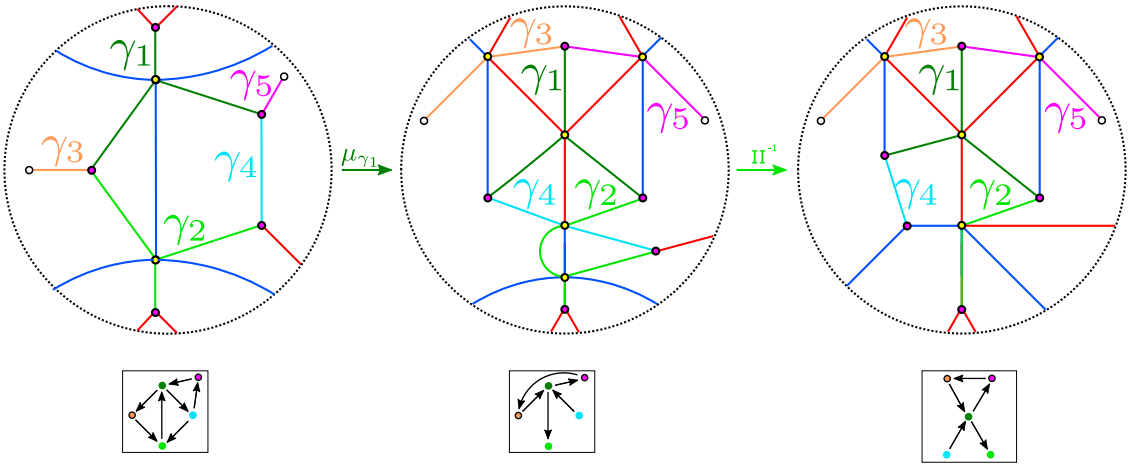

FIGURE 3.18. Type IV.i to Type I mutations.

intersection at the top of the resulting 3-graph.

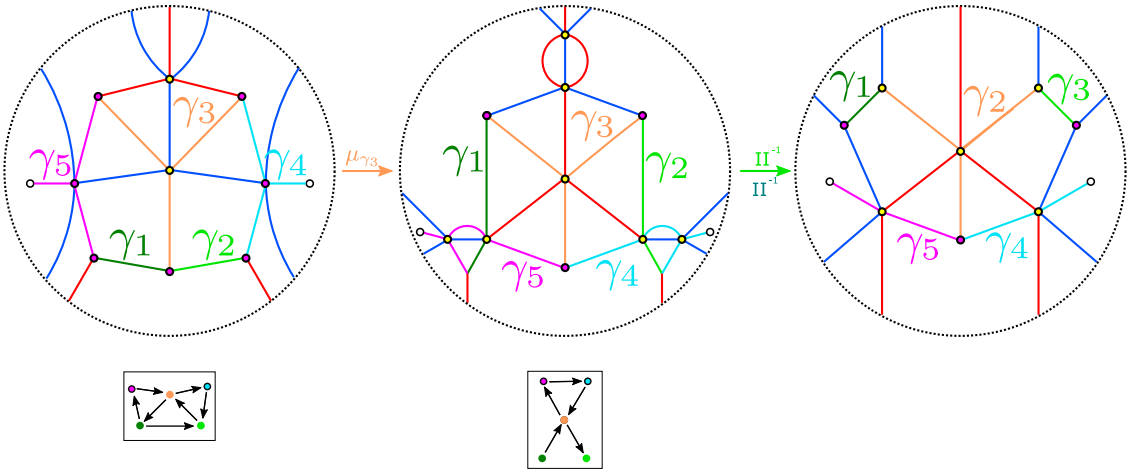

FIGURE 3.19. Type IV.ii to Type I mutations.

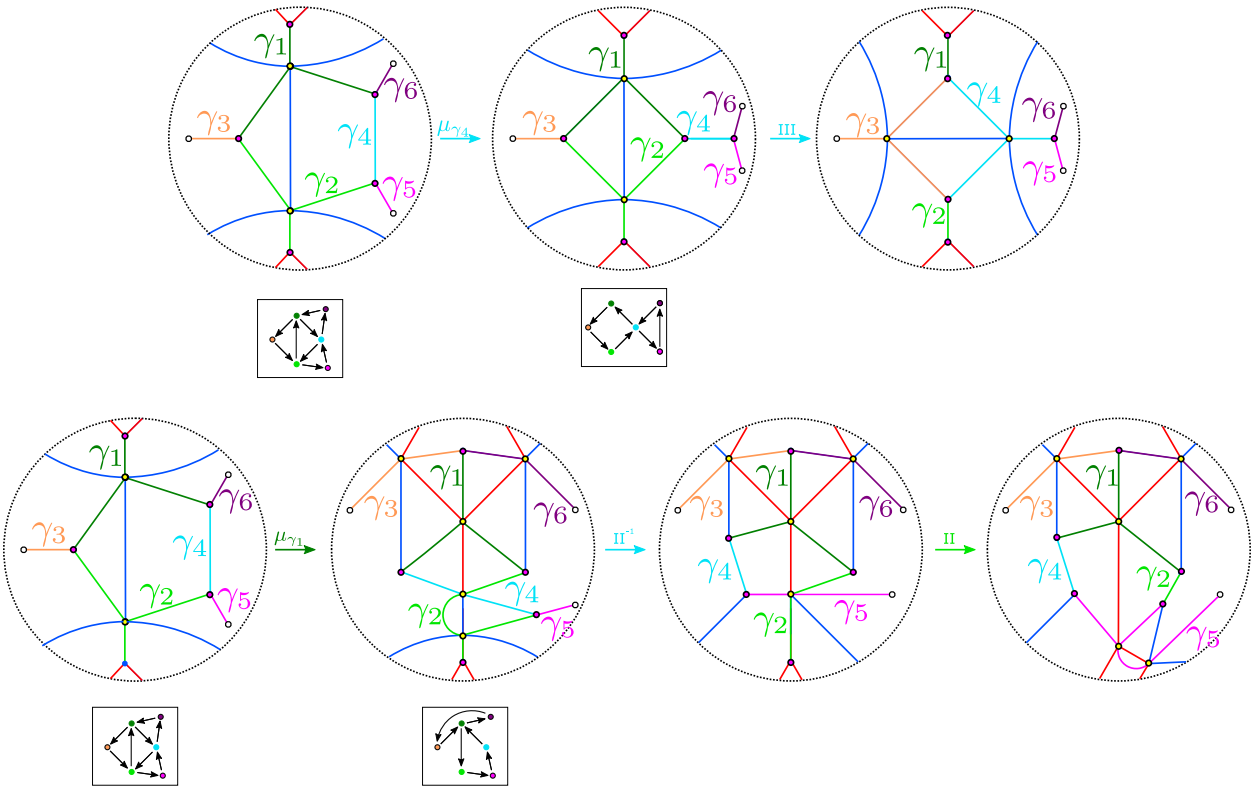

FIGURE 3.20. Type IV.i to Type III mutations.

iii. (Type IV.i to Type III) We now consider the two sequences of 3-graphs in Figure 3.20. Mutation at any of γ_1, γ_2, γ_3, or γ_4 in the Type IV.i normal form yields a Type III normal

form. Specifically, mutation at γ_4 yields a Type III.i normal form that requires no simplification, while mutation at γ_3 (not pictured) yields a Type III.ii normal form that also requires no simplification. The computation for mutation at γ_1 is pictured in the sequence on the right and is identical to the computation for mutation at γ_2. The first step of the simplification is the same as the Type IV.i to Type I subcase described above. However, we require the application of an additional push-through to remove the geometric intersection between γ_2 and γ_5. This makes γ_5 into a Y-cycle and results in a Type III normal form.

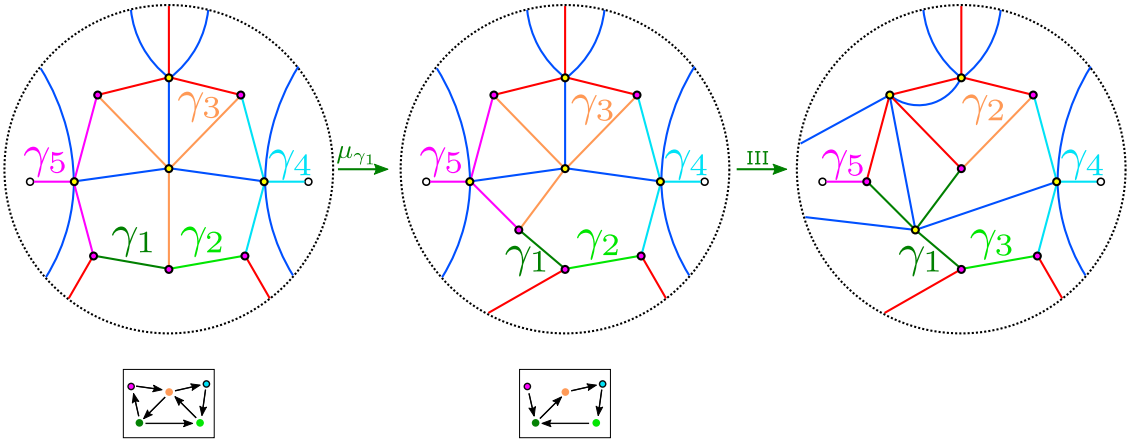

FIGURE 3.21. Type IV.ii to Type III mutations.

iv. (Type IV.ii to Type III) Mutation at γ_1 in our Type IV.ii normal form, depicted in Figure 3.21, results in a pair of geometric intersections between γ_3 and γ_5. Application of a flop removes these geometric intersections and results in a sharp 3-graph with Y-cycles γ_1 and γ_4, which matches our Type III.ii normal form. Note that the computations for mutations involving a Type IV.ii 3-graph with a single spike cycle are identical.

The remaining three subcases are all Type IV to Type IV mutations.

v. (Type IV.ii to Type IV) Figure 3.22 depicts mutation of a Type IV.ii normal form at a spike cycle. Mutating at γ_5 results in an additional geometric intersection between γ_1 and

γ_3. We first apply a reverse push-through at the trivalent vertex where γ_1, γ_2 and γ_3 meet. This introduces an additional geometric intersection between γ_2 and γ_3, that we resolve by applying a push-through to γ_3. Application of a reverse push-through to the trivalent vertex where γ_1 and γ_5 intersect a red edge resolves the final geometric intersection between γ_1 and γ_5. The Y-cycles of the resulting 3-graph correspond to $k-$cycle vertices of the quiver. As shown below, none of the other Type IV to Type IV mutations result in Y-cycles corresponding to spike vertices. Therefore, assuming we have simplified after each of our mutations in the manner described above, the only possible way a Type IV.ii 3-graph arises is by mutating from the initial Type I graphs in Figure 3.7. Hence, all other Type IV 3-graphs only have Y-cycles corresponding to $k-$cycle vertices in the quiver. The computations involving a Type IV.ii 3-graph with a single spike cycle are again identical.

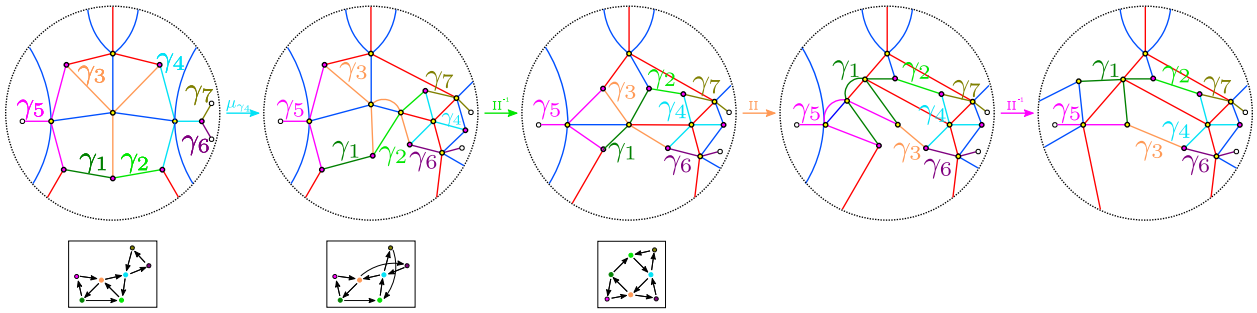

FIGURE 3.22. Type IV.ii graph mutation at a spike cycle.

vi. (Type IV to Type IV) Figure 3.23 depicts Type IV to Type IV mutations when the length of the quiver $k-$cycle is greater than 3. When mutating at a homology cycle corresponding to a $k-$cycle vertex of the quiver, we have two possibilities. Figure 3.23 (top) shows the case where γ_4 intersects another Y-cycle γ_2, which corresponds to a $k-$cycle vertex in the quiver. Figure 3.23 (bottom) considers the case where γ_4 only intersects I-cycles. In both of these cases we must apply a reverse push-through to the trivalent vertex where γ_3 and γ_4 intersect a red edge in order to simplify the 3-graph. This particular simplification requires that neither of the two edges adjacent to the leftmost edge of γ_4 carry a cycle

before we mutate. A similar computation (not pictured) involving the Y-cycle γ_2 would also require that neither of the two edges adjacent to the bottommost edge of γ_2 carry a cycle. Crucially, our computations show that Type IV to Type IV mutation preserve this property, i.e., that both of the Y-cycles have an edge that is adjacent to a pair of edges which do not carry a cycle. When $k = 4$, the resulting 3-graph resulting from the computations in the top line will have a short I-cycle adjacent to γ_2 and γ_3, while the 3-graph resulting from the computations in the bottom line will have a short Y-cycle adjacent to γ_2 and γ_3.

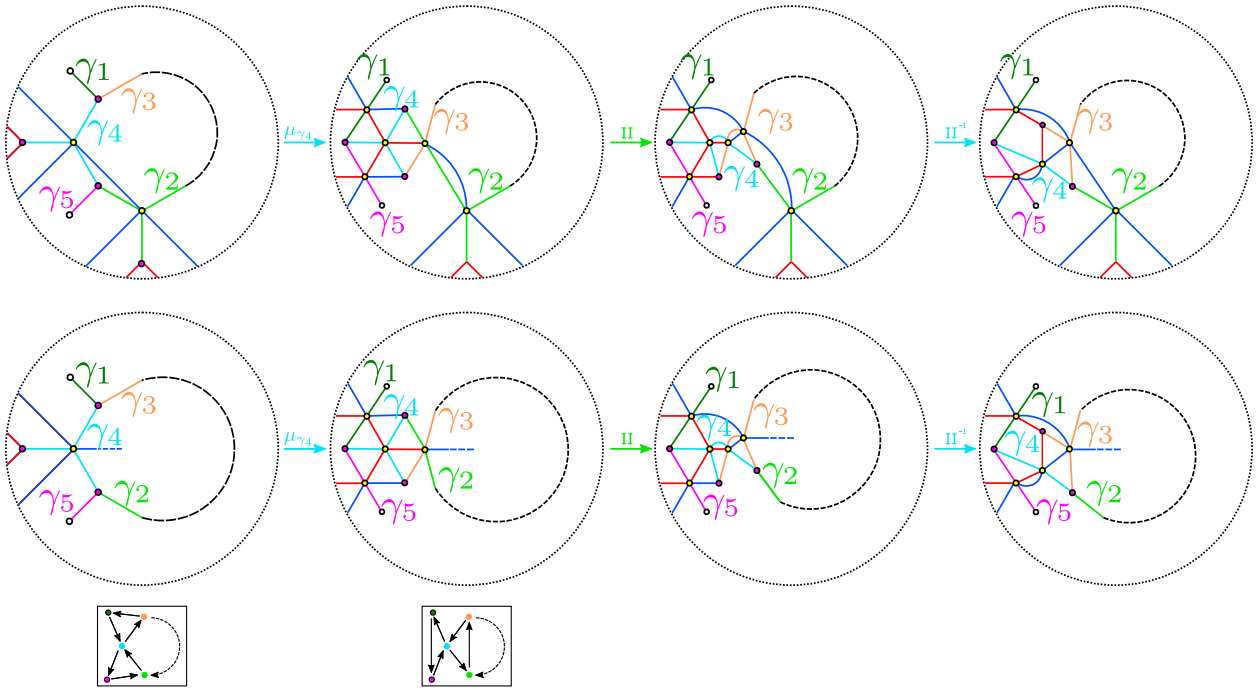

FIGURE 3.23. Type IV to Type IV mutations at homology cycles corresponding to $k-$cycle vertices in the quiver. Mutating at γ_2, γ_3, or γ_4 (corresponding to $k-$cycle vertices in the quiver) in the 3-graphs on the left decreases the length of the $k-$cycle in the quiver by 1.

vii. (Type IV to Type IV) Figure 3.24 depicts mutation at a spike cycle. Since we have already discussed the Type IV.ii spike cycle subcase above, we need only consider the case where

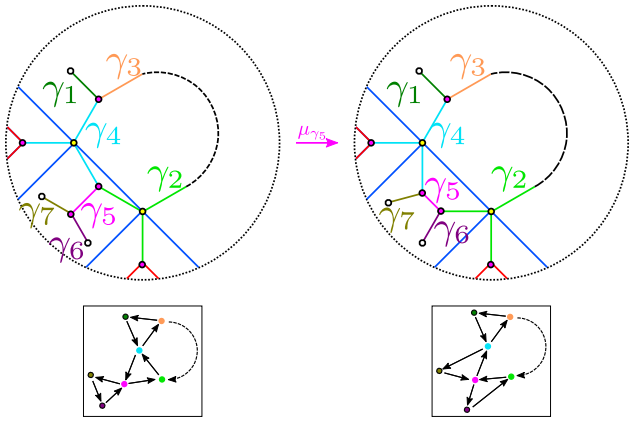

FIGURE 3.24. Type IV to Type IV mutations at spike cycles. Mutating at the spike cycles γ_1 or γ_5 in the 3-graphs on the left increases the length of the $k-$cycle in the intersection quiver by 1.

each of the spike cycles is a short I-cycle. γ_7 and γ_6 are included to help indicate where A_n tail cycles are sent under this mutation. The computation for mutating at a spike edge for Type IV.i (i.e. the $k = 3$ case) is identical to the $k > 3$ case. We have omitted the case where each of the cycles involved in our mutation is an I-cycle, but the computation is again a straightforward mutation of a single I-cycle that requires no simplification.

In each of the Type IV to Type IV subcases above, mutating at a Y-cycle or an I-cycle and applying the simplifications as shown preserves the number of Y-cycles in our graph. Therefore, our computations match the normal form we gave in Figure 3.2 with $k-2$ short I-cycles in the normal form 3-graph not belonging to any A_n tail subgraphs.

This completes our classification of the mutations of normal forms. In each case, we have produced a 3-graph of the correct normal form that is locally sharp and made up of short Y-cycles and I-cycles. Thus, any sequence of quiver mutations for the intersection quiver $Q(\Gamma_0(D_n), \{\gamma_i^{(0)}\})$ of our initial $\Gamma_0(D_n)$ is weave realizable. Hence, given any sequence of quiver mutations, we can apply a sequence of Legendrian mutations to our original 3-graph to arrive at a 3-graph with intersection quiver given by applying that sequence of quiver mutations to $Q(\Gamma_0(D_n), \{\gamma_i^{(0)}\})$, as desired.

□

Having proven weave realizability for $\Gamma_0(D_n)$, we conclude with a proof of the following corollary.

COROLLARY 3.1.1. *Every cluster chart of the moduli of microlocal rank-1 sheaves $\mathcal{C}(\Gamma_0(D_n))$ is induced by at least one embedded exact Lagrangian filling of $\lambda(D_n) \subset (\mathbb{S}^3, \xi_{st})$. In particular, there exist at least $(3n-2)C_{n-1}$ exact Lagrangian fillings of the link $\lambda(D_n)$ up to Hamiltonian isotopy, where C_n denotes the nth Catalan number.*

3.1.1. Proof of Corollary 3.1.1. We take our initial cluster seed in $\mathcal{C}(\Gamma)$ to be the cluster seed associated to $\Gamma_0(D_n)$. The cluster variables in this initial seed exactly correspond to the microlocal monodromies along each of the homology cycles of the initial basis $\{\gamma_i^{(0)}\}$. The intersection quiver $Q(\Gamma_0(D_n), \{\gamma_i^0\})$ is the D_n Dynkin diagram and thus the cluster seed is D_n-type. By definition, any other cluster seed in the D_n-type cluster algebra is obtained by a sequence of quiver mutations starting with the quiver $Q(\Gamma_0(D_n), \{\gamma_i^0\})$ and its associated cluster variables. Theorem 1.1.2 implies that any quiver mutation of $Q(\Gamma_0(D_n), \{\gamma_i^0\})$ can be realized by a Legendrian mutation in $\Lambda(\Gamma_0(D_n))$, so we have proven the first part of the corollary. The remaining part of the corollary follows from the fact that the D_n-type cluster algebra is known to be of finite mutation type with $(3n-2)C_{n-1}$ distinct cluster seeds. □

CHAPTER 4

Legendrian loops and mapping class groups

In this chapter, we investigate Legendrian loops as elements of cluster modular groups. While the granular data of the orbital structure of exact Lagrangian fillings under the action of a Legendrian loop is not as accessible beyond the case of $\lambda(A_n)$, we are able to give a contact-geometric interpretation of known results from the cluster literature to build on the initial work of Casals and Gao. In addition to realizing specific group presentations, we also gain insights into various Legendrian loop actions via Ishibashi's work comparing cluster modular groups and mapping class groups. In particular, we apply a cluster Nielsen-Thurston classification for cluster modular group elements to our study of Legendrian loops. This approach yields new techniques for showing that a Legendrian loops produces infinitely many fillings by studying fixed points of its induced action on the positive real part of the cluster variety and by understanding a process known as cluster reduction. The richness of the theory of mapping class groups promises further applications for this approach.

4.1. Legendrian loops as generators of cluster modular groups

In this section, we give presentations for cluster modular groups and finite-index subgroups of Legendrian links corresponding to a subset of simply-laced finite, affine, and extended affine cluster types. We start by describing the quiver combinatorics used in [**KG21**] to give presentations for cluster modular groups of affine and extended affine types. We then define plabic fences, which we will use to compute sequences of mutations associated to nearly all of the Legendrian loops we consider. With this combinatorial approach, we are able to show that the induced action of our Legendrian loops matches with the presentations of cluster modular groups given.

4.1.1. Presentations of cluster modular groups via $T_{\mathbf{n}}$ quivers. In [**KG21**], the authors give a particular presentation for quivers that are associated to affine and extended affine cluster algebras. They use this presentation to explicitly describe the generators of known cluster modular

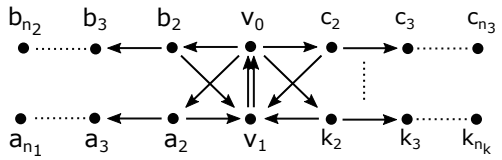

FIGURE 4.1. A $T_{n_1,\dots n_k}$ quiver. Deleting the vertex v_0 results in a quiver with central vertex v_1 and tails of length $n_i - 1$ (not including the vertex v_1) where the vertices alternate as either sources or sinks for the two incident edges.

groups and in some cases derive previously unknown cluster modular groups. We make use of their work by showing that the initial quivers from plabic fences for affine and extended affine type yield nearly identical combinatorial presentations. We begin by defining Kaufman and Greenberg's quivers for simply-laced type.[1]

DEFINITION 4.1.1. *Given a vector* $\mathbf{n} = (n_1, \dots, n_k)$ *of natural numbers* $n_i \geq 2$*, a* $T_{\mathbf{n}}$ *quiver is a quiver with a pair of special vertices* v_1 *and* v_0 *and a collection of* k *'tails' of vertices of length* $n_1, \dots, n_k$*, as pictured in Figure 4.1.*[2]

In this work, we will always consider the case of $k = 3$. The particular cluster types we consider are listed in Tables 4.1 and 4.2 along with the corresponding $T_{\mathbf{n}}$ quiver. $T_{\mathbf{n}}$ quivers admit a particular class of quiver automorphisms, which we will denote $\tau_1, \dots, \tau_k$. Recall that we denote a cluster automorphism by a tuple with the first entry a sequence of quiver mutations and the second entry a permutation describing the relabeling of the quiver vertices. For a tail of length n_i in $T_{\mathbf{n}}$, we denote $i_{\text{odd}} = \{i_j | 3 \leq j \leq n_i, j \text{ odd}\}$ and $i_{\text{even}}\{i_j | 3 \leq j \leq n_i, j \text{ even}\}$. The automorphism τ_i is then given by

$$\tau_i = (\mu_{i_{\text{odd}}} \mu_{i_{\text{even}}} \mu_{i_2} \mu_{v_0} \mu_{v_1}, (i_2\, i_0\, i_1))$$

where $\mu_{i_{\text{odd}}}$ denotes a sequence of mutations starting at i_j for $j < n_i$ the largest number in i_{odd} and ending at the vertex i_3. The mutation sequence $\mu_{i_{\text{even}}}$ is defined analogously. Denote by Γ_τ the subgroup of $\mathcal{G}$ generated by $\tau_1, \dots, \tau_k$. In addition to the τ_i, we also have graph automorphisms $\sigma \in \text{Aut}(T_{\mathbf{n}})$ acting on Γ_τ by swapping tails i and j of length $n_i = n_j$. The utility of the $T_{\mathbf{n}}$ quiver construction is demonstrated by the following theorem.

[1]Kaufman and Greenberg also use folding to obtain quivers corresponding to non-simply-laced type, but contact-geometric interpretation of cluster automorphisms in this setting is not as clear.

[2]The quiver pictured in Figure 4.1 is actually opposite to the one considered in [**KG21**]. The pictured orientation is chosen to match our previous conventions.

THEOREM 4.1.1 ([**KG21**], Theorems 5.2 and 6.1). *For a cluster algebra $\mathcal{A}$ of affine type, $\mathcal{G}(\mathcal{A})) \cong \Gamma_\tau \rtimes \mathrm{Aut}(T_{\mathbf{n}})$. For a cluster algebra $\mathcal{A}$ of extended affine type, $\Gamma_\tau \rtimes \mathrm{Aut}(T_{\mathbf{n}})$ is isomorphic to a finite index subgroup of $\mathcal{G}(\mathcal{A})$.*

Kaufman and Greenberg also conjecture that $\mathcal{G}(\mathcal{A})) \cong \Gamma_\tau \rtimes \mathrm{Aut}(T_{\mathbf{n}})$ for cluster algebras $\mathcal{A}$ that admit a $T_{\mathbf{n}}$ quiver and are not of extended affine type [**KG21**, Conjecture 4.7].

We now give a more explicit description of the group structure of Γ_τ. The following algebraic statement will have a geometric interpretation in terms of Legendrian loops.

THEOREM 4.1.2 (Theorem 4.11, [**KG21**]). *Γ_τ is an abelian group with relations $\tau_i^{n_i} = \tau_j^{n_j}$.*

Denote by $\gamma := \tau_i^{n_i}$. By the above theorem, any choice of i results in a well-defined element of Γ_τ. The Donaldson-Thomas transformation can then be defined by $\mathrm{DT} = \gamma^2 \prod_{i=1}^{k} \tau_i \gamma^{-1}$ [**KG21**, Theorem 4.14]. For the cluster algebras we investigate here, this simplifies to $\mathrm{DT} = \tau_1\tau_2\tau_3\gamma^{-1}$.

4.1.2. Legendrian loops in affine and extended affine type. In this subsection, we explain how to geometrically realize Kaufman and Greenberg's presentation of cluster modular groups for affine type. We introduce plabic fences as a combinatorial means of obtaining a sequence of mutations induced by a Legendrian loop. We then describe fronts with initial quivers that are mutation equivalent to $T_{\mathbf{n}}$ quivers and Legendrian loops that induce automorphisms conjugate to the τ_i generators of Γ_τ.

4.1.2.1. *Plabic fences.* Due to the variety of weave equivalence moves, it is often difficult to determine a sequence of mutations that induces the same cluster automorphism as a Legendrian loop. Even when we can compute the sequence of mutations via other combinatorial means, it can be challenging to show that the mutation sequence in the weave agrees with the Legendrian loops. For example, see Appendix A for a somewhat involved computation of a short mutation sequence induced by a Legendrian loop. Therefore, in order to determine a sequence of mutations corresponding to the induced action of most of the Legendrian loops we consider, we require a combinatorial way to relate quivers to weaves. In A_n-type, we were able to understand the Kálmán loop via triangulations dual to 2-weaves, but arbitrary cluster types do not admit a similar construction. Instead, we use the combinatorics of plabic graphs.

Plabic graphs are a general combinatorial object related to cluster theory and were first studied by Postnikov [**Pos06**]. The term 'plabic fence' that we use for the particular graphs we work with was coined in [**FPST22**] and refers to a particular form of plabic graph, as we describe below. We will follow the recipe of [**CW22**] for producing a Legendrian weave from a plabic fence and use the graph to encode the combinatorics of Legendrian loop mutations.

DEFINITION 4.1.2. *A plabic fence is a planar graph with univalent or trivalent vertices colored either black or white constructed as follows:*

(1) *Stack n horizontal lines of the same length on top of each other, each starting with a white vertex on the left and ending with a black vertex on the right.*
(2) *Add vertical edges between adjacent pairs of horizontal lines with trivalent vertices where they meet colored so that each endpoint of the vertical edge is a different color.*

From the plabic fence $\mathbb{G}$, we can extract a quiver $Q_{\mathbb{G}}$ as follows:

(1) Assign a mutable vertex to each face of $\mathbb{G}$.
(2) For every edge $e \in \mathbb{G}$ connecting two faces, add an arrow between the corresponding vertices oriented so that the white endpoint of e is to the right of the edge when traveling in the direction of orientation.

Given a plabic fence $\mathbb{G}$, we can associate a Legendrian link $\lambda_{\mathbb{G}}$ and an initial Legendrian weave surface with boundary $\lambda_{\mathbb{G}}$, following the recipe from [**CW22**, Sections 2 and 3]. To describe the process for obtaining a positive braid β from $\mathbb{G}$ such that $\lambda_{\mathbb{G}} \cong \lambda(\beta)$, first label the horizontal lines from bottom to top with the numbers 1 through n and label the vertical edges between horizontal lines i and $i+1$ by σ_i. We refer to a vertical edge by the color of its top vertex, so that a vertical edge between lines i and $i+1$ is white if the vertex on line i is colored white. In the front projection, one should think of the black vertical edges as denoting crossings appearing above cusps, while white vertical edges correspond to crossings below cusps. More precisely, from $\mathbb{G}$, we obtain braid words β_1 from white vertical edges and β_2 from black vertical edges by scanning $\mathbb{G}$ from left to right. The Legendrian link is Legendrian isotopic to the rainbow closure of the braid $\beta_1\beta_2^{\circ}$ or equivalently $\beta_2^{\circ}\beta_1$ where β_2° is obtained from β_2 by reading the braid word from right to left and replacing σ_i with σ_{n-i} for all i.

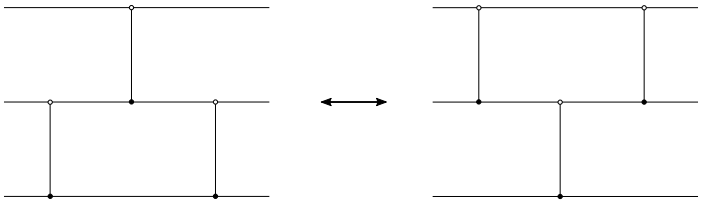

FIGURE 4.2. Local move of a plabic fence corresponding to the braid move exchanging $\sigma_i\sigma_{i+1}\sigma_i$ and $\sigma_{i+1}\sigma_i\sigma_{i+1}$.

From $\lambda_{\mathbb{G}}$, Casals and Weng construct an initial weave $\Lambda(\mathbb{G})$ [**CW22**, Definition 3.24]. By [**CL22**, Theorem 1.2], the combinatorial data of the seed corresponding to the plabic graph $\mathbb{G}$ agrees with this choice of initial weave $\Lambda(\mathbb{G})$. More precisely, the Legendrian $\lambda_{\mathbb{G}}$ given here is Legendrian isotopic to the Legendrian obtained as the conormal lift of zig-zag strands and the conjugate Lagrangian surface associated to an initial seed is Hamiltonian isotopic to the Lagrangian projection of the initial weave. As a result, the intersection quiver of $\Lambda(\mathbb{G})$ agrees with the initial quiver $Q_{\mathbb{G}}$ coming from the plabic fence.

4.1.2.2. *Mutation sequences induced by Legendrian loops.* In order to compute the sequence of mutations induced by Legendrian loops we decompose our Legendrian loops into a series of simple Legendrian isotopies and describe how to realize them in the plabic fence. The first Legendrian isotopy we consider is a Reidemeister III move. In terms of our braid, this swaps $\sigma_i\sigma_{i+1}\sigma_i$ and $\sigma_{i+1}\sigma_i\sigma_{i+1}$, so the plabic fences $\mathbb{G}_\beta$ and $\mathbb{G}_{\beta'}$ corresponding to the two braid words β and β' differ by the local move pictured in Figure 4.2. Combinatorially, we can see that $Q_{\mathbb{G}_\beta}$ differs from $Q_{\mathbb{G}_{\beta'}}$ by a mutation at the vertex corresponding to the face. In the sheaf moduli, the Legendrian isotopy induces a map between $\mathcal{M}_1(\Lambda(\mathbb{G}_\beta))$ and $\mathcal{M}_1(\Lambda(\mathbb{G}_{\beta'})$ that is an isomorphism, but generally not an automorphism. The following lemma states that this isomorphism agrees with cluster mutation.

LEMMA 4.1.1 (Proposition 6.10, [**CLSBW23**]). *The initial seeds of $\mathcal{M}_1(\Lambda(\mathbb{G}_\beta)$ and $\mathcal{M}_1(\Lambda(\mathbb{G}_{\beta'})$ obtained by a left-to-right pinching sequence are related by mutation at the cluster variable corresponding to the face in the plabic fence.*

In addition to the Reidemeister III move, we also perform a sequence of isotopies corresponding to rotating the leftmost crossing of our braid clockwise around our Legendrian link until it becomes the rightmost crossing. This modifies the braid word by conjugation. In a plabic fence with all

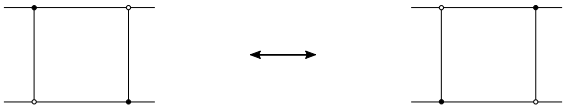

FIGURE 4.3. A square move in a plabic fence, corresponding to mutation at the vertex represented by the face.

white vertical edges, we can interpret this isotopy as flipping the leftmost white edge to black, moving it past all of the white vertical edges in the row to its right, and then flipping it back to white. Each time we move a black edge past a white one, we perform the local move pictured in Figure 4.3, known as a square move. As with the Reidemeister III move, the local move induces an isomorphism between sheaf moduli that corresponds to a cluster mutation between the initial seeds.

LEMMA 4.1.2. *[**CW22**, Section 5.3] The initial seeds of $\mathcal{M}_1(\Lambda(\mathbb{G}))$ and $\mathcal{M}_1(\Lambda(\mathbb{G}'))$ corresponding to plabic fences $\mathbb{G}$ and $\mathbb{G}'$ related by the local move pictured in Figure 4.3 are related by mutation at the cluster variable corresponding to the face in the plabic fence.*

Together, Lemmas 4.1.1 and 4.1.2 give a sequence of mutations corresponding to a Legendrian loop. We will use this to connect Legendrian loops to known descriptions of cluster automorphisms.

One automorphism of particular importance, the Donaldson-Thomas transformation (DT) is not known to be induced by a Legendrian loop. Instead, Casals and Weng describe a Legendrian isotopy and a strict contactomorphism that induce DT [**CW22**, Section 5]. This procedure is roughly described as rotating all of the crossings from the strands below the cusps to the strands above the cusps and then performing the strict contactomorphism $x \mapsto -x, z \mapsto -z$. Starting with a plabic fence of all white edges, the first component of DT can be combinatorially realized by repeatedly flipping the leftmost white edge to black and moving it past the remaining white edges in its row by square moves until it becomes the leftmost black edge. The contactomorphism then takes all black edges to white edges, so that we return to the initial plabic fence. Recording the square moves involved in this process gives an explicit mutation sequence.

REMARK 4.1.1. *In order to match the conventions of Legendrian loops rotating crossings in a clockwise orientation, our description above is reversed from that of [**CW22**]. The automorphism we describe above actually corresponds to* DT^{-1}.

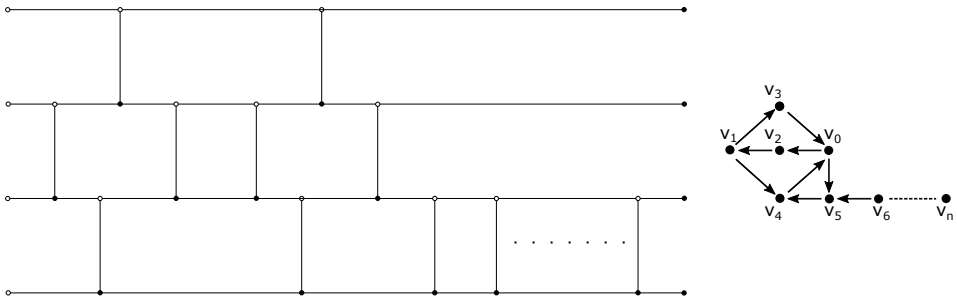

FIGURE 4.4. Plabic fence $\mathbb{G}$ corresponding to $\lambda(\tilde{D}_n)$. The dots correspond to $n-6$ additional vertical edges

4.1.2.3. *Legendrian loops conjugate to τ_i.* Using Lemmas 4.1.1 and 4.1.2, we can compute an explicit mutation sequence for the families of Legendrian loops we are interested in. This computation, together with work of Kaufman and Greenberg on combinatorially understanding cluster modular groups of affine and extended affine types will allow us to describe cluster automorphisms as Legendrian loops.

Let us define a family of positive braids β_{k,n_1,n_2} by

$$\beta_{k,n_1,n_2} := (\sigma_2\sigma_1\sigma_3\sigma_2)^{2k}\sigma_1^{n_1-2}\sigma_3^{n_2-2}$$

for $k, n_1, n_2 \in \mathbb{N}$ and $n_1, n_2 \geq 2$. For Legendrian links corresponding to affine and extended affine type cluster algebras, we will require $k = 1$ and n_1 and n_2 agreeing with the corresponding values for $T_{\mathbf{n}}$ quivers. See Tables 4.1 and 4.2 below. Denote by $\Lambda(k, n_1, n_2)$ the initial weave filling of $\lambda(\beta_{k,n_1,n_2})$.

LEMMA 4.1.3. *The intersection quiver $Q_{\Lambda(1,n_1,n_2)}$ is mutation equivalent to $T_{n_1,n_2,2}$.*

PROOF. A plabic fence corresponding to $\lambda(\beta_{1,n_1,n_2})$ can be obtained from the plabic fence in Figure 4.4 by the addition of $n_2 - 2$ vertical edges between the top two horizontal lines. We then get an initial quiver by adding $n_2 - 2$ vertices to the quiver pictured in Figure 4.4 in a manner identical to the tail labeled by vertices $v_6, \ldots, v_n$. Mutating at vertices v_3 and v_4 of $Q_{\Lambda(1,n_1,n_2)}$ yields a $T_{n_1,n_2,2}$ quiver up to the alternating behavior of the arrows in the tails. This alternation can be obtained by observing that the tails are $A-$type subquivers. □

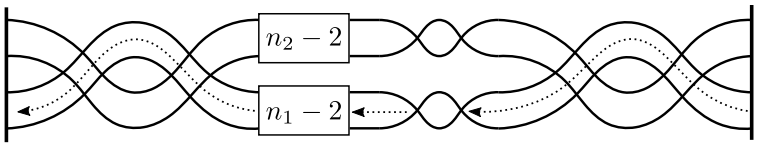

FIGURE 4.5. The Legendrian loop ϑ_1 for the braid β_{1,n_1,n_2} pictured in $J^1\mathbb{S}^1$. The loop takes one of the $n_1 - 2$ crossings appearing in the box and drags it around the front, following the path of the dotted arrows.

To present Legendrian loops of $\lambda(\beta_{k,n_1,n_2})$, we will consider the Legendrian isotopic link given as the (-1) closure of the braid $(\sigma_2\sigma_1\sigma_3\sigma_2)^{2k}\sigma_1^{n_1-2}\sigma_3^{n_2-2}\sigma_1^2\sigma_3^2(\sigma_2\sigma_1\sigma_3\sigma_2)^2$. Note that in either front, the Legendrian link $\lambda(\beta_{k,n_1,n_2})$ can be described as a pair of Legendrian $(2,n)$ torus links λ_1 and λ_2 that are linked together in a nontrivial way. With this description, we immediately obtain two ϑ loops, the first, defined by commuting a single σ_1 crossing around λ_1 and the second defined by commuting a single σ_3 crossing around λ_2. See Figure 4.5 for an example. Note that when $n_i = 2$, ϑ_i can still be identified because of our choice of braid word for Δ^2.

The following result geometrically realizes the algebraic statements of Theorem 4.1.2. In the statement below, we use $\cong$ to denote Hamiltonian isotopy of traces of Legendrian loops fixing their boundaries.

LEMMA 4.1.4. *The Legendrian loops ϑ_1 and ϑ_2 satisfy $\vartheta_1 \circ \vartheta_2 \cong \vartheta_2 \circ \vartheta_1$ and $\vartheta_1^{n_1} \cong \vartheta_2^{n_2}$.*

PROOF. Consider the trace of the isotopy $\vartheta_2 \circ \vartheta_1$. This exact Lagrangian cobordism can be described by first performing ϑ_1 during time t_1 and then performing ϑ_2 during time t_2 for $0 \leq t_1 < \frac{1}{2}$ and $\frac{1}{2} \leq t_2 \leq 1$. The isotopy required for commutativity is simply the isotopy defined by gradually increasing t_1 and gradually decreasing t_2.

The second part of the lemma follows from our description of the Legendrian loops ϑ_1 and ϑ_2 as being performed on crossings of separate copies of Legendrian (2, n) torus links. The loop $\vartheta_1^{n_1}$ can then be understood as passing all of the crossings of λ_1 around the link λ_2, which is equivalent to passing all of the crossings of λ_2 around the link λ_1. □

We now show that our description of Legendrian loops coincides with Γ_τ.

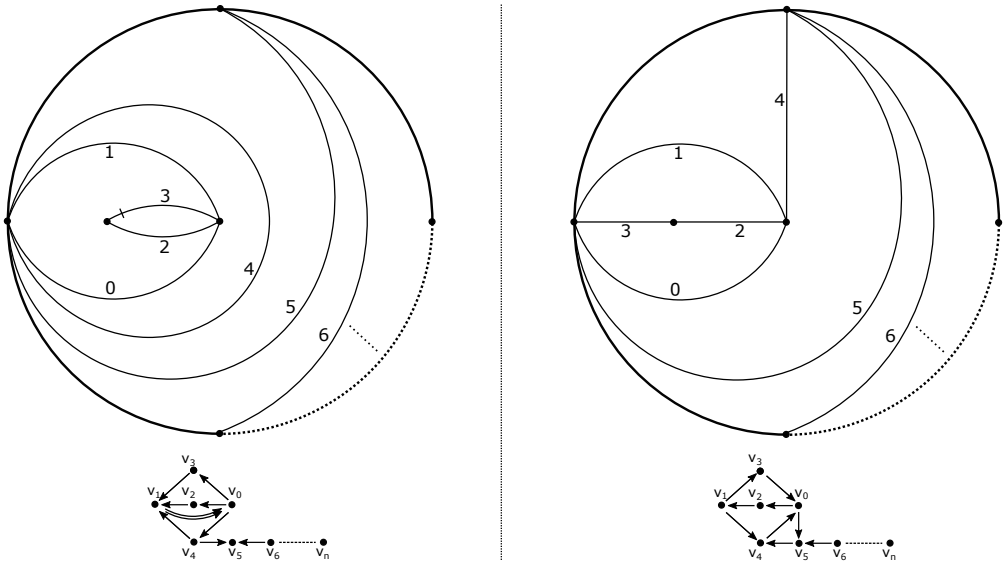

FIGURE 4.6. A pair of tagged triangulations of $\mathbb{D}^2_{n-2,0,0}$. The triangulation on the left corresponds to the quiver $T_{n-2,2,2}$, while the triangulation on the right corresponds to the quiver $\mathbb{G}(\lambda(1, n_1, n_2))$ and is obtained from the triangulation on the left by mutation at edges labeled 3 and 4. The dotted lines on the left represent a zig-zag pattern of $n-6$ edges, while the dotted lines on the right represent $n-6$ edges all sharing the top vertex.

LEMMA 4.1.5. *For any Legendrian $\lambda(\beta_{1,n_1,n_2})$, the cluster automorphisms $\tilde{\vartheta}_1$ and $\tilde{\vartheta}_2$ induced by the corresponding Legendrian loops are conjugate to τ_1 and τ_2.*

PROOF. We first establish the lemma in the case of $\lambda(\tilde{D}_n) \cong \lambda(\beta_{1,n-2,2})$ using the combinatorics of tagged triangulations of a twice-punctured disk. We then leverage these combinatorics for the case of $n_2 > 2$.

Consider the plabic fence $\mathbb{G}(\tilde{D}_n)$ depicted in Figure 4.4, corresponding to the Legendrian link $\lambda(\tilde{D}_n) \cong \lambda_{1,n-2,2}$. The sequence of mutations corresponding to ϑ_1 can be determined from $\mathbb{G}(\tilde{D}_n)$ using Lemmas 4.1.1 and 4.1.2. In particular, ϑ_1 induces the mutation sequence $\tilde{\vartheta}_1 = (5\,0\,3\,1\,1\,4\,6 \ldots n, (1\,4\,5\,0\,3))$ and ϑ_2 induces the mutation sequence $\tilde{\vartheta}_2 = (1\,4, (1\,4\,0\,3))$. This is shown by explicit computation in the initial weave filling $\Lambda(\mathbb{G}(\tilde{D}_n))$. See Appendix A for the computation. Note that the use of Legendrian weaves for this computation appears to be necessary, as the crossings involved in the Legendrian loop ϑ_2 do not appear in the rainbow closure $\lambda((\sigma_2\sigma_1\sigma_3\sigma_2)^2\sigma_1^{n-3})$ and therefore are not captured by the combinatorics of the plabic fence.

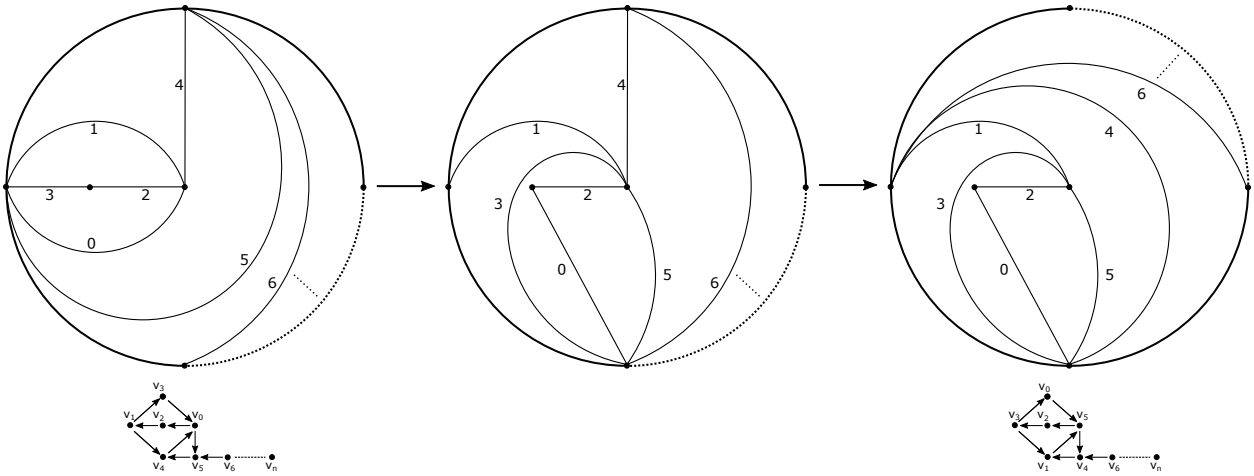

FIGURE 4.7. Sequence of mutations induced by the Legendrian loop ϑ_1 in the tagged triangulation corresponding to $Q_{\mathbb{G}(\tilde{D}_n)}$. The second triangulation is obtained from the first by mutating at edges labeled $5, 0$, and 3 in order. The third triangulation is obtained from the second performing the remaining mutations of ϑ_1.

Figure 4.6 depicts two tagged triangulations of a twice-punctured disk. The triangulation $\mathcal{T}_1$ (left) corresponds to a quiver identical to $Q_{\mathbb{G}(\tilde{D}_n)}$, while the triangulation $\mathcal{T}_2$ (right) corresponds to a quiver identical to $T_{n-2,2,2}$. The generator τ_1 corresponds to a rotation of the boundary of the disk by $2\pi/n$ [**KG21**, Lemma B.3]. Therefore, to show that τ_1 and $\tilde{\vartheta}_1$ are conjugate, we need to show that $\tilde{\vartheta}_1$ also corresponds to a rotation of the boundary of the disk by $2\pi/n$. This is done explicitly in Figure 4.7.

Similar to the case of τ_1, the generator τ_2 corresponds to a half twist about the two punctures in the triangulation $\mathcal{T}_2$. Figure 4.8 shows the computations for $\tilde{\vartheta}_2$.

For $n_2 > 2$, we observe that the sequence of mutations induced by $\tilde{\vartheta}_1$ or $\tilde{\vartheta}_2$ fixes the quiver vertices corresponding to the other tail. Freezing or deleting these vertices therefore yields a $\tilde{D}_n$ quiver and we can simply apply our above reasoning to show that cluster automorphisms are conjugate. □

4.1.2.4. *Proof of Theorem 1.1.4.* We are now able to give a proof of Theorem 1.1.4, restated below for additional clarity.

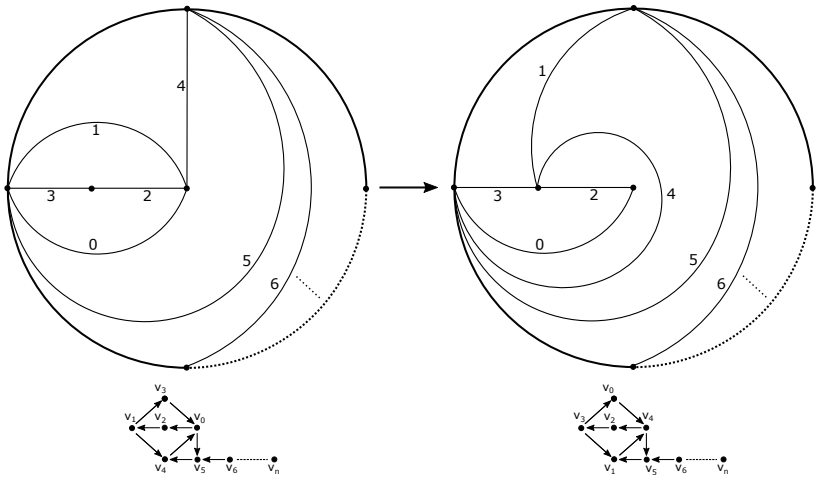

FIGURE 4.8. Sequence of mutations induced by the Legendrian loop ϑ_2 in the tagged triangulation corresponding to $Q_{\mathbb{G}(\tilde{D}_n)}$.

THEOREM 4.1.3. *For $\lambda \in \mathcal{H}$, the cluster modular group of $\mathcal{M}_1(\lambda)$ is virtually generated by Legendrian loops. Moreover, if $\lambda \in \mathcal{H}'$, the cluster modular group of $\mathcal{M}_1(\lambda)$ is generated by Legendrian loops and* DT.

We break the proof into three cases: (1) affine type (2) extended affine type, and (3) finite type.

PROOF. (1) Recall that the cluster modular group $\mathcal{G}$ of affine type cluster algebras is isomorphic to $\Gamma_\tau \rtimes \mathrm{Aut}(T_{\mathbf{n}})$ and that Γ_τ is generated by τ_1, τ_2, and τ_3. By Lemma 4.1.5, we can realize τ_1 and τ_2 as Legendrian loops, and the relation $\tau_1^{n_1} = \tau_2^{n_2} = \tau_3^{n_3}$ implies that the subgroup generated by τ_1 and τ_2 is a finite index subgroup of $\mathcal{G}$.

In the case of $\tilde{E}_7$ or $\tilde{E}_8$, an inspection of the quiver $T_{\mathbf{n}}$ verifies that that $\mathrm{Aut}(T_{\mathbf{n}})$ is trivial. Therefore, we need only show that we can generate Γ_τ by Legendrian loops and DT. By [**KG21**, Theorem 4.14], we have that $\mathrm{DT} = \tau_1\tau_2\tau_3\gamma^{-1}$. Solving for $\tau_3 = \mathrm{DT}\,\tau_1^{-1}\tau_2^{-1}\gamma$ therefore gives the remaining generator of Γ_τ.

(2) By [**CG22**, Theorems 1.1 and 1.3], we have Legendrian loops generating $\mathrm{PSL}_2(\mathbb{Z})$ and $\mathrm{MCG}(\mathbb{S}^2, 4)$ subgroups of $\mathcal{G}(\mathcal{M}_1(\lambda(3,6)))$ and $\mathcal{G}(\mathcal{M}_1(\lambda(4,4))$, respectively. Fraser shows that these subgroups are isomorphic to the quotient $\mathcal{G}/\mathcal{Z}(\mathcal{G})$. The center of each of these cluster modular groups is generated by DT. However, Kaufman and Greenberg identify a quiver automorphism $\sigma \in \mathrm{Aut}(T_{4,4,2})$ corresponding to the case of $\lambda(4,4)$. Therefore, the

Legendrian loops of Casals and Gao virtually generate the cluster modular group in this case.

(3) For $\lambda(E_6), \lambda(E_7), \lambda(E_8)$, the cluster modular group is generated solely by DT, so the statement follows immediately from [**CW22**, Theorem 5.8].

For $\lambda(A_n)$, the Kálmán loop has order $n+3$ and generates $\mathcal{G}(\mathcal{M}_1(\lambda(A_n)))$, as stated in Corollary 2.3.2.

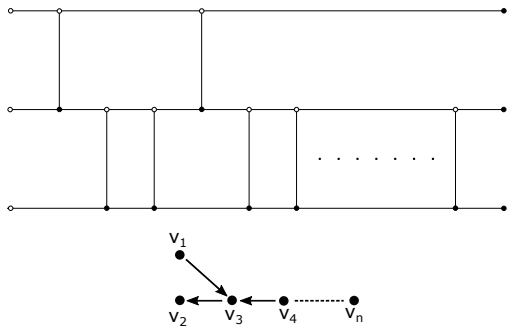

FIGURE 4.9. Plabic fence $\mathbb{G}(D_n)$ with corresponding quiver. The dots represent $n-5$ additional vertical edges.

For $\lambda(D_n)$, we use the combinatorics of tagged triangulations of once-punctured n-gons. By [**ASS12**, Theorem 1.2], the generators of the cluster modular group correspond to rotation of the n-gon by $2\pi/n$ and simultaneous changing of the tags at the puncture when $n \geq 5$. Let $\mathbb{G}(D_n)$ be the plabic fence pictured in Figure 4.9. The corresponding Legendrian $\lambda(\mathbb{G}(D_n))$ is Legendrian isotopic to the link $\lambda(D_n)$ defined in Section 1.1. Depicting $\lambda(\mathbb{G}_n)$ as the (-1) closure of $\sigma_2\sigma_1^2\sigma_2\sigma_1^n\sigma_2\sigma_1^2\sigma_2$ allows us to define a Legendrian loop ϑ by commuting a σ_1 past each of the $\sigma_2\sigma_1^2\sigma_2$ subwords. As in the $\tilde{D}_n$ case, we can compute the induced sequence of mutations. In Figure 4.10 we give a tagged triangulation with a quiver equivalent to $Q_{\mathbb{G}(D_n)}$ and show by explicit computation that the sequence of mutations induced by ϑ corresponds to rotation of the punctured n-gon by $2\pi/n$. In addition, DT induces a sequence of mutations that corresponds to a rotation of the once-punctured n-gon by $2\pi/n$ and a simultaneous changing of all the taggings at the puncture. Therefore, the pair ϑ and DT generate the cluster modular group $\mathcal{G}(\mathcal{M}_1(\lambda(D_n)))$.

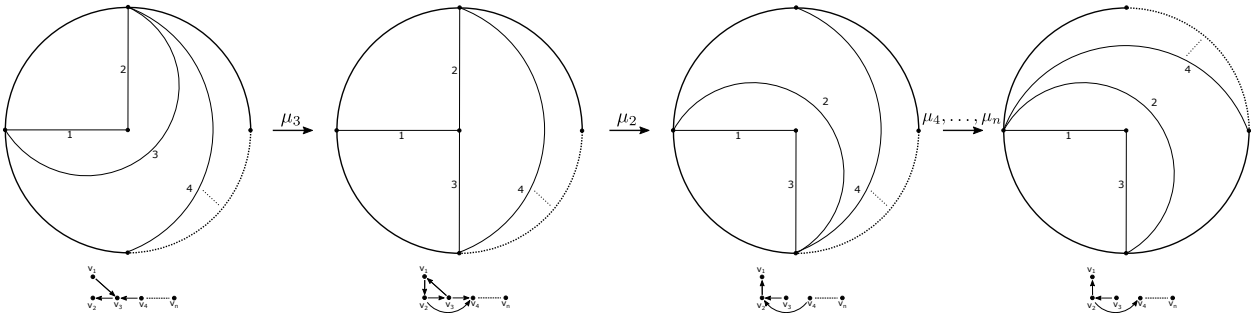

FIGURE 4.10. Once-punctured disks depicting the mutation sequence $(3\,2\,4, \ldots n, (1\,2\,3))$ induced by the Legendrian loop ϑ of $\lambda(D_n)$.

Cluster Type	Cluster modular group	Generated by	$T_{\mathbf{n}}$ quiver type
A_n	$\mathbb{Z}_{n+3}$	ρ	–
D_n $n > 4$	$\mathbb{Z}_n \times \mathbb{Z}_2$	ϑ_1, DT	–
E_6, E_7, E_8	$\mathbb{Z}_{14}, \mathbb{Z}_{10}, \mathbb{Z}_{16}$	DT	–
$\tilde{E}_7$	$(\mathbb{Z}_2 \times \mathbb{Z})$	DT, ϑ_1	(4, 3, 2)
$\tilde{E}_8$	$\mathbb{Z}_8$	DT	(5, 3, 2)
$E_8^{(1,1)}$	G_1		(6, 3, 2)

TABLE 4.1. Cluster modular groups generated by Legendrian loops. The final column contains the information of the vector $\mathbf{n}$ in cases where there is a $T_{\mathbf{n}}$ quiver in the mutation class. The letter ρ denotes the Kálmán loop. The group G_1 is isomorphic to $\mathrm{PSL}_2(\mathbb{Z}) \rtimes \mathbb{Z}_6$

.

Cluster Type	Cluster modular group	$T_{\mathbf{n}}$ quiver type
D_4	$\mathbb{Z} \times S_3$	–
$\tilde{D}_n$ n odd	$(\mathbb{Z}_2 \times \mathbb{Z}_2) \rtimes (\mathbb{Z}_2 \times \mathbb{Z})$	(n-2, 2, 2)
$\tilde{D}_n$ n even	$(\mathbb{Z}_2 \times \mathbb{Z}_2) \rtimes \mathbb{Z}$	(n-2, 2, 2)
$\tilde{E}_6$	$(S_3 \times \mathbb{Z})$	(3, 3, 2)
$E_7^{(1,1)}$	G_2	(4, 4, 2)

TABLE 4.2. Cluster modular groups virtually generated by Legendrian loops. The group G_2 is isomorphic to $\mathrm{MCG}(\mathbb{S}^2, 4) \rtimes (\mathbb{Z}_4 \times \mathbb{Z}_2)$.

□

REMARK 4.1.2. *By Lemma 4.1.5 and the fact that τ_3 can be obtained from τ_1, τ_2 and* DT*, we are able to generate Γ_τ for all $T_{n_1,n_2,2}$. Therefore, for $n_1 \neq n_2$, we obtain the conjectured cluster modular group, while for $n_1 = n_2$, we obtain an index 2 subgroup (see [**KG21**, Conjecture 4.7]). Moreover,*

we claim without proof that in the case of $n_1 = n_2$, the order two element generating $\mathrm{Aut}(T_{\mathbf{n}})$ *is induced by a Legendrian isotopy swapping the two sublinks of $\lambda(\beta_{1,n_1,n_2})$ given as satellites of different strands.*

4.2. Nielsen-Thurston classification of Legendrian loops

In this section, we consider qualitative properties of Legendrian loops by investigating connections between cluster modular groups and mapping class groups. We start by describing a Nielsen-Thurston classification of cluster automorphisms due to Ishibashi and then use this framework to give some general statements about fixed point properties of Legendrian loop actions. We conclude with a pair of examples.

We take the following characterization of cluster automorphisms to be our Nielsen-Thurston-like classification. Our definition differs from Ishibashi's original definition in order to give a more natural description in the context of Legendrian loops. See [**Ish19**] for alternate characterizations in terms of fixed points of automorphisms acting on the cluster complex.

DEFINITION 4.2.1. *A cluster automorphism $\varphi \in \mathcal{G}(\mathcal{M}_1(\lambda))$ is*

(1) *periodic if it is of finite order.*

(2) *cluster reducible if it fixes a set of cluster variables.*

(3) *cluster pseudo-Anosov if no power of φ is cluster reducible.*

Note that the cyclic subgroup generated by any cluster automorphism will correspond to at least one of these classes. We say that a Legendrian loop is (cluster) periodic, reducible, or pseudo-Anosov if its induced cluster automorphism is of the corresponding type. Below, we provide examples of periodic and reducible Legendrian loops.

EXAMPLE 4.2.1. *Consider the sheaf moduli $\mathcal{M}_1(\lambda(k,n))$ of the Legendrian torus link $\lambda(k,n)$. This cluster variety is known to have the same mutable part as the top-dimensional positroid cell of the Grassmannian $Gr(k,n+k)$, which itself admits a cluster structure. The Kálmán loop induces the cyclic shift ρ, which acts on column vectors in the matrix representation of the top-dimensional positroid cell by $v_i \mapsto v_{i-1}$. By construction, ρ has order $k+n$, which implies that the Kálmán loop is always periodic.*

Note that for torus links $\lambda(k,n)$, we have $\mathrm{DT}^2 = \rho^{k+n}$, so that DT is also periodic in this case.

EXAMPLE 4.2.2. *Let $\lambda = \lambda_{k,n_1,n_2}$ with $k \in 2\mathbb{N}$ and $n_1, n_2 \geq 2$. The cluster automorphisms $\tilde{\vartheta}_1$ or $\tilde{\vartheta}_2$ induced by the Legendrian loops defined above fix at least $\frac{k}{2}$ cluster variables and are therefore cluster reducible.*

In the case of cluster pseudo-Anosovs, the condition that no power of a given cluster automorphism fixes no set of cluster variables in any seed is more challenging to verify. We conjecture that the image of the braid group element $\sigma_1\sigma_2^{-1} \in \mathrm{MCG}(\mathbb{S}^2, 4) \subseteq \mathcal{G}(\mathcal{M}_1(\lambda(4,4)))$ gives a cluster pseudo-Anosov automorphism, as $\sigma_1\sigma_2^{-1}$ is a pseudo-Anosov mapping class in $\mathrm{MCG}(\mathbb{S}^2, 4)$.

4.2.1. Cluster Reduction. In this section, we discuss the process of cluster reduction, analogous to the concept of a reduction system in the theory of mapping class groups. From a reducible mapping class φ, one can obtain a mapping class on a simpler surface by cutting along curves fixed by φ. The analogous process in cluster theory allows us to gain additional information about certain Legendrian loops.

We define a cluster reducible automorphism to be proper reducible if in addition to fixing some collection of cluster variables setwise, it also fixes at least one cluster variable. Note that any cluster reducible automorphism yields a proper reducible automorphism after raising it to a high enough power. Analogous to the notion of a reduction system in the theory of mapping class groups, Ishibashi defines the process of cluster reduction. Given a cluster reducible cluster automorphism φ, one freezes quiver vertices corresponding to cluster variables fixed by some power of φ. This induces a cluster automorphism of a cluster algebra with a smaller mutable part.

EXAMPLE 4.2.3. *For $\lambda(\tilde{D}_n)$, the Legendrian loop ϑ_1 fixes a single cluster variable. The cluster reduction of this quiver produces a new quiver with the fixed cluster variable either frozen or deleted. Upon inspection, the mutable part of the quiver obtained by cluster reduction corresponds to a (tagged) triangulation of an annulus with n marked points on the outer boundary and two marked points on the inner boundary. Note that this triangulation, pictured in Figure 4.11, can be readily obtained from the triangulation of twice punctured disk we started with. The induced cluster automorphism is a $2\pi/n$ rotation of the outer boundary component.*

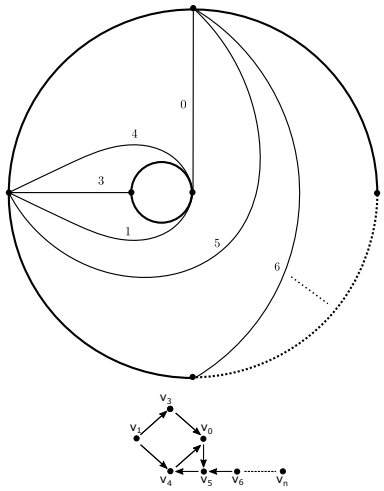

FIGURE 4.11. Triangulation of $\mathbb{D}^2_{n_1-2,2}$ corresponding to the cluster reduction of ϑ_1. Note that replacing the inner boundary component by a single edge recovers the triangulation of the twice-punctured disk in Figure 4.6 (right).

REMARK 4.2.1. *Note that the process of cluster reduction presented here is entirely algebraic. From a contact geometric perspective, it does not appear to be possible to isolate an arbitrary homology cycle γ in a Legendrian weave in order to remove it and still obtain an embedded exact Lagrangian surface. One can instead consider modifying the ambient symplectic manifold of the Lagrangian projection of the weave so that the cycle γ no longer bounds an embedded Lagrangian disk with which to perform a mutation. This modification might reasonably be obtained by removing the Lagrangian 2-disk bounding γ from the Lagrangian skeleta of $(\mathbb{D}^4, \omega_{st})$ constructed in [**Cas21**, Section 1.1]. The full details of this construction lie beyond the scope of this work.* □

The previous example motivates the notion of a cluster Dehn twist, which we introduce below.. Denote by Q_i the quiver with two mutable vertices and i edges from vertex v_1 to vertex v_2. The quiver Q_i admits a cluster automorphism $tw_i = (\mu_1, (12))$. We have the following definition due to Ishibashi [**Ish20**].

DEFINITION 4.2.2. *A cluster Dehn twist is a cluster automorphism φ such that after a finite number of cluster-reductions, the induced automorphism $\tilde{\varphi}$ satisfies $\tilde{\varphi}^n = tw_k^m$ for some nonzero integers m, n and $i \geq 2$.*

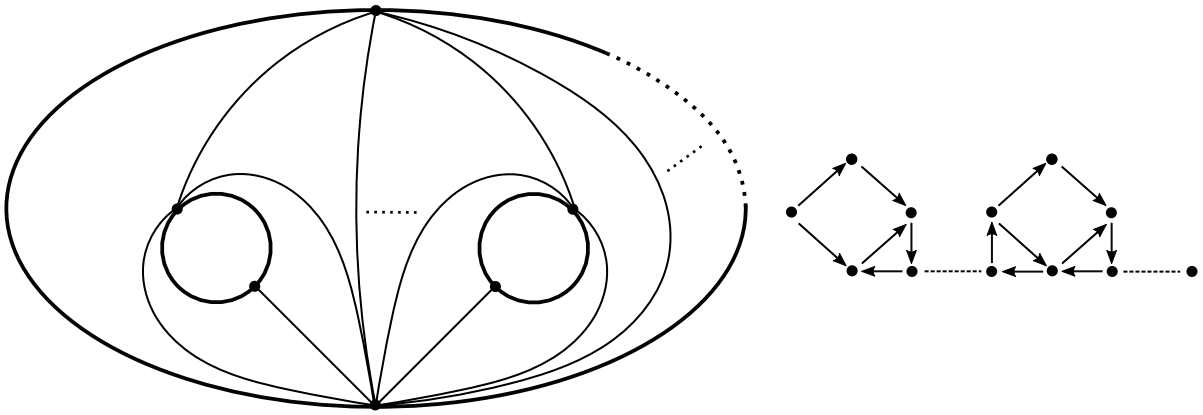

FIGURE 4.12. The mutable subquiver (right) resulting from cluster reduction of ϑ_1 acting on $\mathcal{M}_1(\lambda(\beta_k, n_1, n_2))$ and its corresponding triangulation.

Similar to above, we say that a Legendrian loop is a cluster Dehn twist if its induced action is. Note that a cluster Dehn twist is necessarily of infinite order because the induced cluster automorphism is of infinite order. Any Dehn twist (or half twist) of a tagged triangulation is a cluster Dehn twist [**Ish20**]. Indeed, the quiver Q_2 corresponds to an annulus with one marked point on each boundary component and tw_2 corresponds to a Dehn twist in this annulus. Example 4.2.3 give us an example of a Legendrian loop that is a cluster Dehn twist. More generally, we have the following.

THEOREM 4.2.1. *The Legendrian loops ϑ_1 and ϑ_2 acting on $\mathcal{M}_1(\lambda(\beta_{k,n_1,n_2}))$ are cluster Dehn twists.*

PROOF. Let $Q(k, n_1, n_2)$ be the quiver corresponding to the plabic fence $\mathbb{G}(\beta_{k,n_1,n_2})$. Freezing the quiver vertices fixed by ϑ_1 results in a quiver corresponding to the triangulation of the surface $\mathbb{S}_{n_1+k-1,2,\ldots,2}$, i.e. a disk with n_1+k-1 marked points on the boundary and k additional boundary components, each with 2 marked points, as pictured in Figure 4.12. The induced action of ϑ_1 is a $\frac{2\pi}{n_1+\frac{k}{2}-1}$ rotation about the outer boundary of the surface. The argument for ϑ_2 is exactly analogous when $n_2 > 2$. □

REMARK 4.2.2. *The technique for this proof follows an argument of Fraser related to cluster modular groups of Grassmannians.*[3] *Based on his argument, we expect that this technique can be*

[3]Chris Fraser, Personal Communication, 2/22/22.

applied more generally to braids of the form $\Delta^{2k}\sigma_i^{n_i}\beta'$ where β' contains no σ_{i-1}, σ_i, or σ_{i+1}. More precisely, we conjecture that the cluster reduction of the induced action of an analogous Legendrian loop is a partial Dehn twist about the boundary of an annulus with some number of boundary components and marked points.

As a corollary of Theorem 4.2.1, we show that the Legendrian loops ϑ_1 and ϑ_2 considered above produce infinitely many exact Lagrangian fillings.

COROLLARY 4.2.1. *The Legendrian loops ϑ_1 and ϑ_2 have infinite order on $\mathcal{M}_1(\lambda(\beta_{k,n_1,n_2}))$.*

Note that in the case of $k = 2$, we also have that DT is a cluster Dehn twist, as $\mathrm{DT}^2 = \vartheta_1^{n_1}$. As a result, we obtain an additional corollary, extending the analogy with mapping class groups. See also [**KG21**, Corollary 6.5].

COROLLARY 4.2.2. *For any $\lambda \in \mathcal{H}$, the group $\mathcal{G}(\mathcal{M}_1(\lambda))$ is virtually generated by cluster Dehn twists. Moreover, if $\lambda \in \mathcal{H}'$, then $\mathcal{G}(\mathcal{M}_1(\lambda))$ is generated by cluster Dehn twists.*

4.2.2. Fixed points. Continuing the analogy between mapping class groups and cluster modular groups, we study the fixed points of Legendrian loop actions on $\mathcal{M}_1(\lambda)$. Define the positive real part $\mathcal{M}_1(\lambda)_{>0}$ of the cluster variety $\mathcal{M}_1(\lambda)$ to be the space given by requiring all of the cluster variables of the initial seed to be strictly positive real numbers. This definition gives a well-defined notion of positivity, because the cluster variables in a cluster seed are efficient positivity tests, i.e. positivity is preserved under mutation. Note that $\mathcal{M}_1(\lambda)_{>0}$ is homeomorphic to an open ball of dimension equal to the first Betti number of a filling of λ.The cluster modular group acts on this space by permuting cluster charts. Interpreting [**Ish19**, Theorem 2.2] in the context of contact geometry yields the following statement.

LEMMA 4.2.1. *Any finite order Legendrian loop φ of a Legendrian λ has a fixed point in $\mathcal{M}_1(\lambda)_{>0}$.*

We define a group action of a group G on a topological space X to be properly discontinuous if for every compact subset $K \subseteq X$, the set $\{g \in G | gK \cap K \neq \emptyset\}$ is finite. Ishibashi shows that if the cluster modular group action on the cluster variety is properly discontinuous, then the implication

of Lemma 4.2.1 can be upgraded to an equivalence. We accordingly upgrade our lemma to the following theorem.

THEOREM 4.2.2. *Assume that the cluster modular group acts properly discontinuously on $\mathcal{M}_1(\lambda)_{>0}$. A Legendrian loop φ has a fixed point in $\mathcal{M}_1(\lambda)_{>0}$, if and only if the induced automorphism $\tilde{\varphi}$ has finite order.*

The cluster modular group action is properly discontinuous in the case of surface-type cluster algebras [**Ish19**, Theorem 3.8]. As a result, we can apply our reasoning from Theorem 4.2.1 to obtain the following statement about Legendrian loops of $\lambda(\beta_{k,n_1,n_2})$.

THEOREM 4.2.3. *The induced actions of the Legendrian loops φ_1 and φ_2 of $\lambda(\beta_{k,n_1,n_2})$ have no fixed points in $\mathcal{M}_1(\lambda(\beta_{k,n_1,n_2}))_{>0}$.*

PROOF. In the proof of Theorem 4.2.1, we showed that the cluster reduction of ϑ_i induces a fractional Dehn twist on a surface-type cluster algebra $\mathcal{A}'$. Since this induced automorphism is of infinite order, [**Ish19**, Theorem 3.8] implies that it has no fixed points in $\mathcal{A}'_{>0}$ obtained by freezing the cluster variables fixed by ϑ_i. Since $\mathcal{A}'$ was obtained from $\mathcal{M}_1(\lambda(\beta)$ by freezing the fixed cluster variables, it follows that no point of $\mathcal{M}_1(\lambda(\beta)_{>0}$ is fixed by ϑ_i. □

We can also use Theorem 4.2.1 to detect Legendrian loops of infinite order. In fact, we can make a slightly stronger statement. The following description mimics a definition of Casals and Ng given for Legendrian loops in the augmentation variety.

DEFINITION 4.2.3. *The induced action $\tilde{\varphi}$ of a Legendrian loop is entire on a toric chart $\mathcal{C}_L = (\mathbb{C}^\times)^n$ induced by an exact Lagrangian filling L of λ if for $k \neq l \in \mathbb{Z}$ we have*

$$\varphi^k(\mathcal{C}_L) \neq \varphi^l(\mathcal{C}_L).$$

As a direct consequence of Lemma 4.2.1, we obtain the following.

THEOREM 4.2.4. *The induced action of a Legendrian loop on $\mathcal{M}_1(\lambda)$ is entire if it has no fixed points in $\mathcal{M}_1(\lambda)_{>0}$.*

PROOF. Let φ be a Legendrian loop of λ with no fixed points in $\mathcal{M}_1(\lambda)_{>0}$. Since φ has no fixed points, this in particular implies that no power of it can fix an entire cluster chart. Therefore, $\varphi^k(\mathcal{C}_L) \neq \varphi^l(\mathcal{C}_L)$ for any $k \neq l \in \mathbb{Z}$. □

Note that the fixed point property of a Legendrian loop is independent of any choice of initial chart, in contrast with the methods used by Casals and Ng to obtain entirety of Legendrian loops in the augmentation variety.

4.2.3. Examples. We present examples of the fixed point behavior of Legendrian loops of $\lambda(A_n)$ and $\lambda(E_8^{(1,1)})$.

EXAMPLE 4.2.4. *Consider the initial seed $x_1 \leftarrow x_2$ in $\mathcal{M}_1(\lambda(A_2))$. The cluster automorphism $\tilde{\rho}$ induced by the Kálmán loop has a single fixed point $x_1 = x_2 = \frac{1+\sqrt{5}}{2}$. In the case of A_n-type, we know that since the quiver corresponds to triangulations of a disk with $n+3$ marked points on the boundary, the cluster modular group action is properly discontinuous on $\mathcal{M}_1(\lambda(A_n))_{>0}$. Therefore, this fixed point recovers the fact that ρ has finite order.*

EXAMPLE 4.2.5. *Now consider the initial seed of $\lambda(3,6)$ corresponding to the front $\lambda(\beta)$ for $\beta = (\sigma_1\sigma_2)^6$. The loop σ_1 defined in [**CG22**] has no fixed points in $\mathcal{M}_1(\lambda(3,6))_{>0}$, recovering the fact that σ_1 (conjugate to ϑ_1 by [**KG21**, Theorem 6.1]) has infinite order and implying that it is entire on any seed. This is verified by showing that there are no positive real solutions to the system of equations obtained from performing the Legendrian loop and setting the corresponding cluster variables equal to each other.*

$$a_1 = \frac{a_2 + a_3 + a_1 a_4}{a_2}$$

$$a_2 = a_4$$

$$a_3 = \frac{a_2 + a_3}{a_1}$$

$$a_4 = \frac{a_3 a_6 + a_4 a_7}{a_5}$$

$$a_5 = \frac{a_1 a_3 a_6 + (a_1 a_4 + (a_2 + a_3) a_5) a_7}{a_1 a_3 a_5}$$

$$a_7 = \frac{a_1 a_3 a_5 a_6 a_8 + (a_1 a_3 a_6^2 + ((a_1 a_4 + (a_2 + a_3) a_5) a_6) a_7) a_9 + (a_1 a_3 a_6 a_7 + (a_1 a_4 + (a_2 + a_3) a_5) a_7^2) a_{10}}{a_1 a_3 a_5 a_7 a_8}$$

$$a_8 = a_{10}$$

$$a_9 = \frac{a_1 a_3 a_5 a_8 + (a_1 a_3 a_6 + (a_1 a_4 + (a_2 + a_3) a_5) a_7) a_9}{a_1 a_3 a_5 a_7}$$

$$a_{10} = a_9$$

Note that the cluster variable a_6 is fixed by the induced (cluster reducible) automorphism. Inputting the above system of equations into a computer algebra system such as Sage verifies that no positive real solutions exists.

APPENDIX A

Mutation sequence computation via Legendrian weaves

In this appendix, we show that the Legendrian loop ϑ_2 defined for $\lambda(\tilde{D}_n)$ induces the cluster automorphism $\tilde{\vartheta}_2 = (1\,4, (1\,4\,0\,3))$, computed as a sequence of mutations in the initial quiver coming from the plabic fence $\mathbb{G}(\tilde{D}_n)$ pictured in Figure 4.9. We do this by mutating at cycles γ_1 and γ_4 in the initial weave $\Lambda(\mathbb{G}(\tilde{D}_6))$ – constructed following [**CW22**, Section 3.3] – and then showing that the resulting weave simplifies to the concatenation of the trace of ϑ_2 with the initial weave, up to relabeling the homology cycles. The example readily generalizes to $\lambda(\tilde{D}_n)$ by replacing the dashed blue short ⅼ-cycle by $n-5$ blue short ⅼ-cycles.

In the figures below, homology cycles are color coded as follows: γ_0 is light blue, γ_1 is orange, γ_2 is light green, γ_3 is purple, γ_4 is pink, γ_5 is yellow, and γ_6 is the dashed blue short ⅼ-cycle. When an edge of the 4-graph carries two cycles, as in the second and third weaves of Figure A.1, we choose one of the colors for the edge itself and then highlight the edge in the color corresponding to the additional homology cycle. The numerals correspond to Legendrian surface Reidemeister moves; we freely apply Move IV by passing edges labeled by σ_3 (colored dark green) and edges labeled by σ_1 (colored blue) over each other. In cases where multiple push-through moves are used, we omit some intermediate steps when the computation is otherwise straightforward.

Note that the choice of braid word for Δ^2 appearing in the weave $\Lambda(\mathbb{G}(\tilde{D}_n))$ differs from the choice of braid word for Δ^2 appearing in the Legendrian front pictured in Figure 4.5. The front for $\partial\Lambda(\mathbb{G}(\tilde{D}_n))$ differs from the front appearing in Figure 4.5 by a sequence of Reidemeister III moves collecting a pair of σ_1 and σ_3 crossings at the left of the braid word for Δ^2. The front appearing in Figure 4.5 displays the Legendrian loops in a more recognizable way, but because the Reidemeister moves only involve crossings in the Δ^2 portion of the braid, the sheaf moduli obtained from the two fronts are equal.

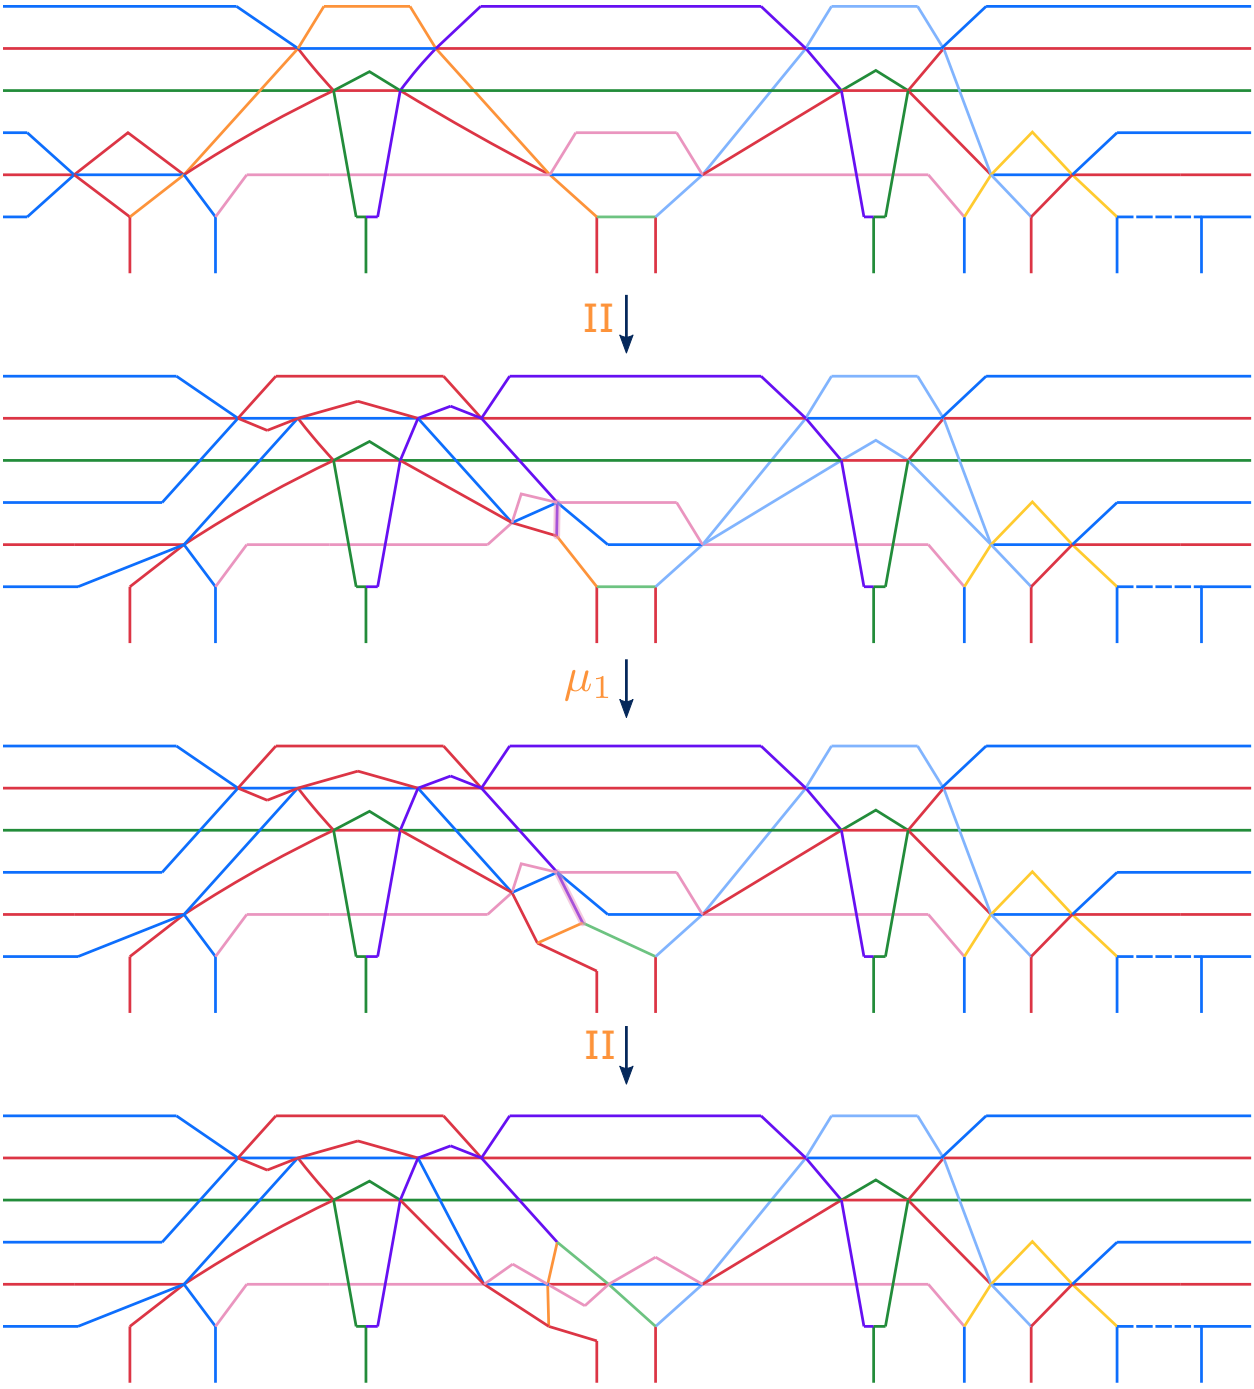

FIGURE A.1. From top to bottom: a sequence of three push-throughs to isolate γ_1 as a short l-cycle; mutation at γ_1; a push through to remove the edge carrying cycles γ_1 and γ_4.

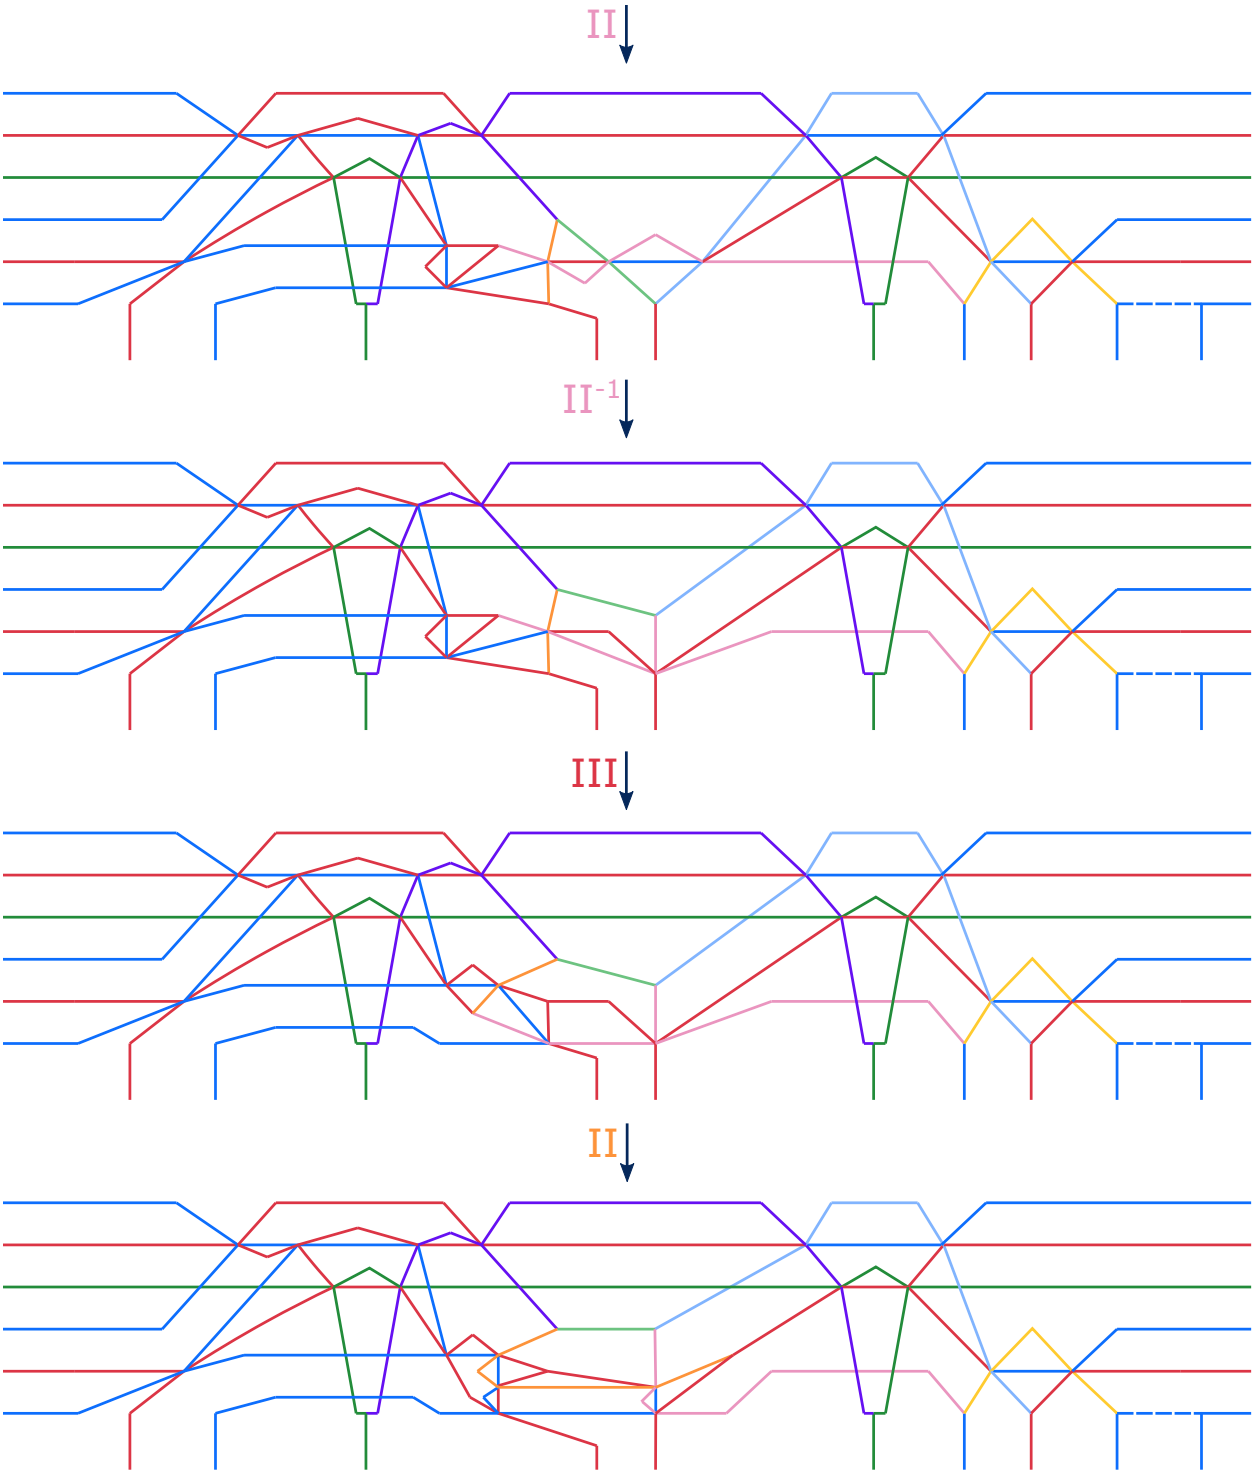

FIGURE A.2. A series of weave equivalence moves designed to produce γ_4 as a long l-cycle.

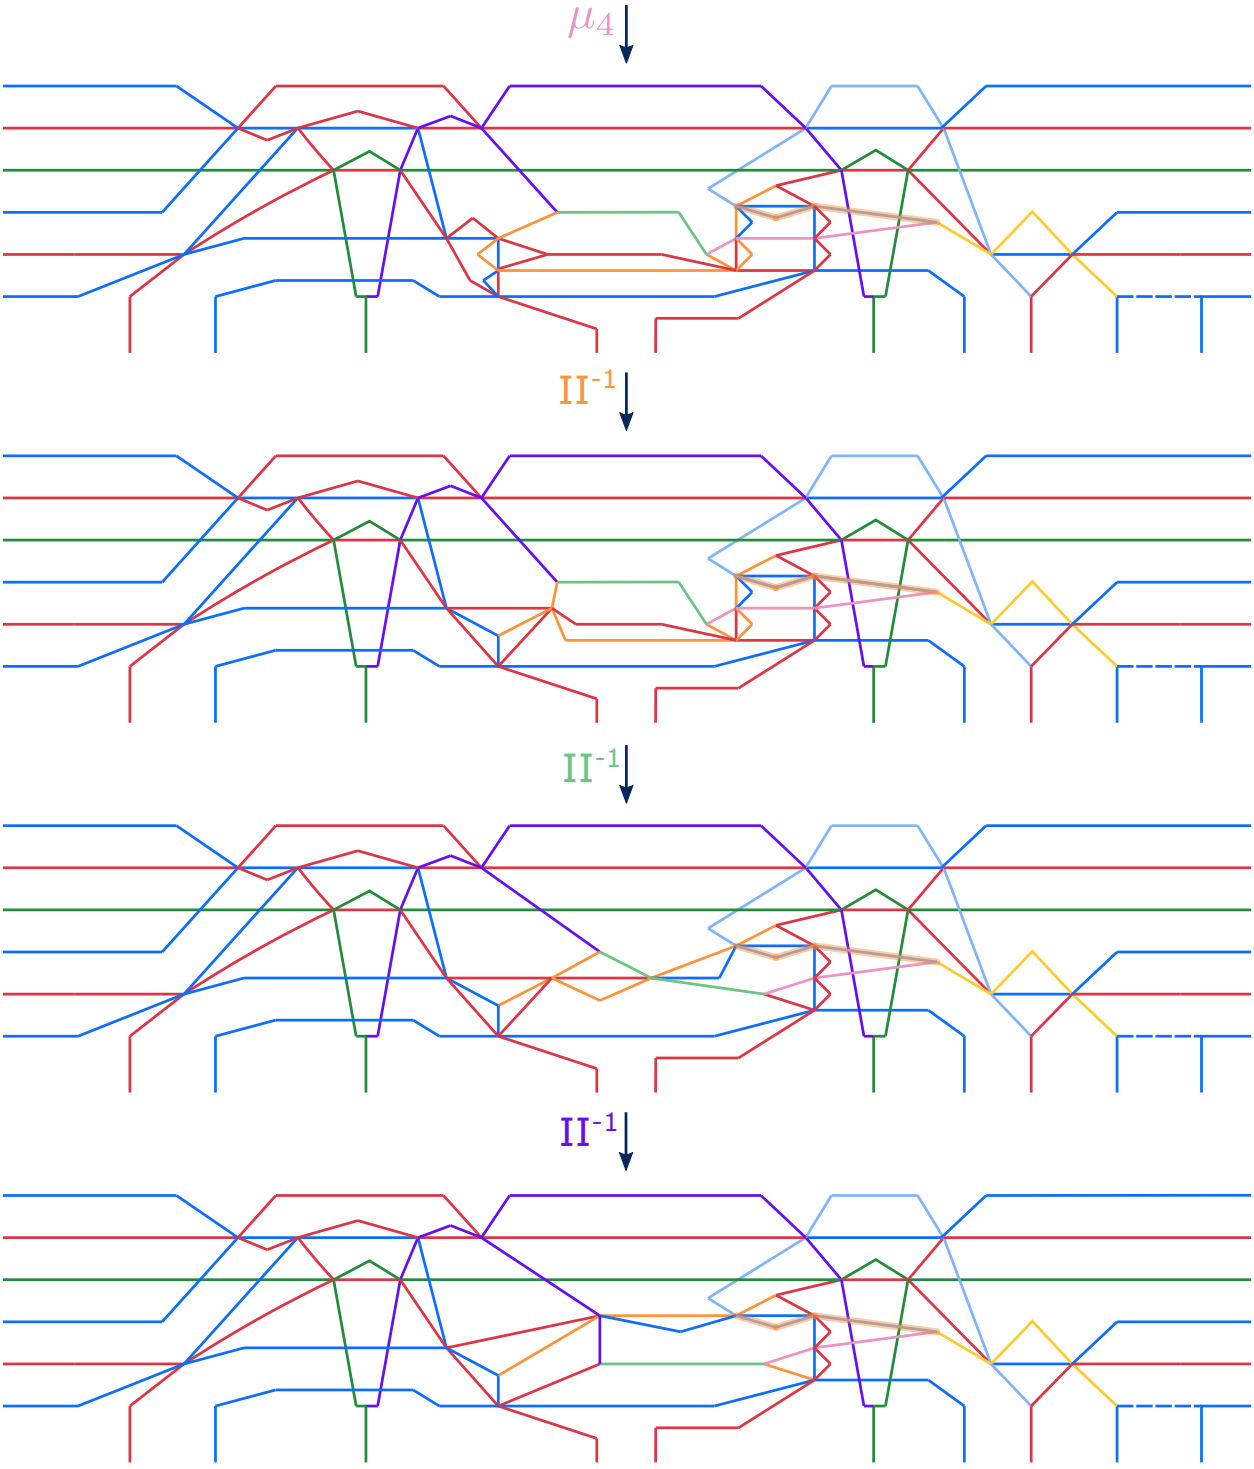

FIGURE A.3. Mutation at γ_4 and a sequence of reverse push-throughs simplifying γ_1.

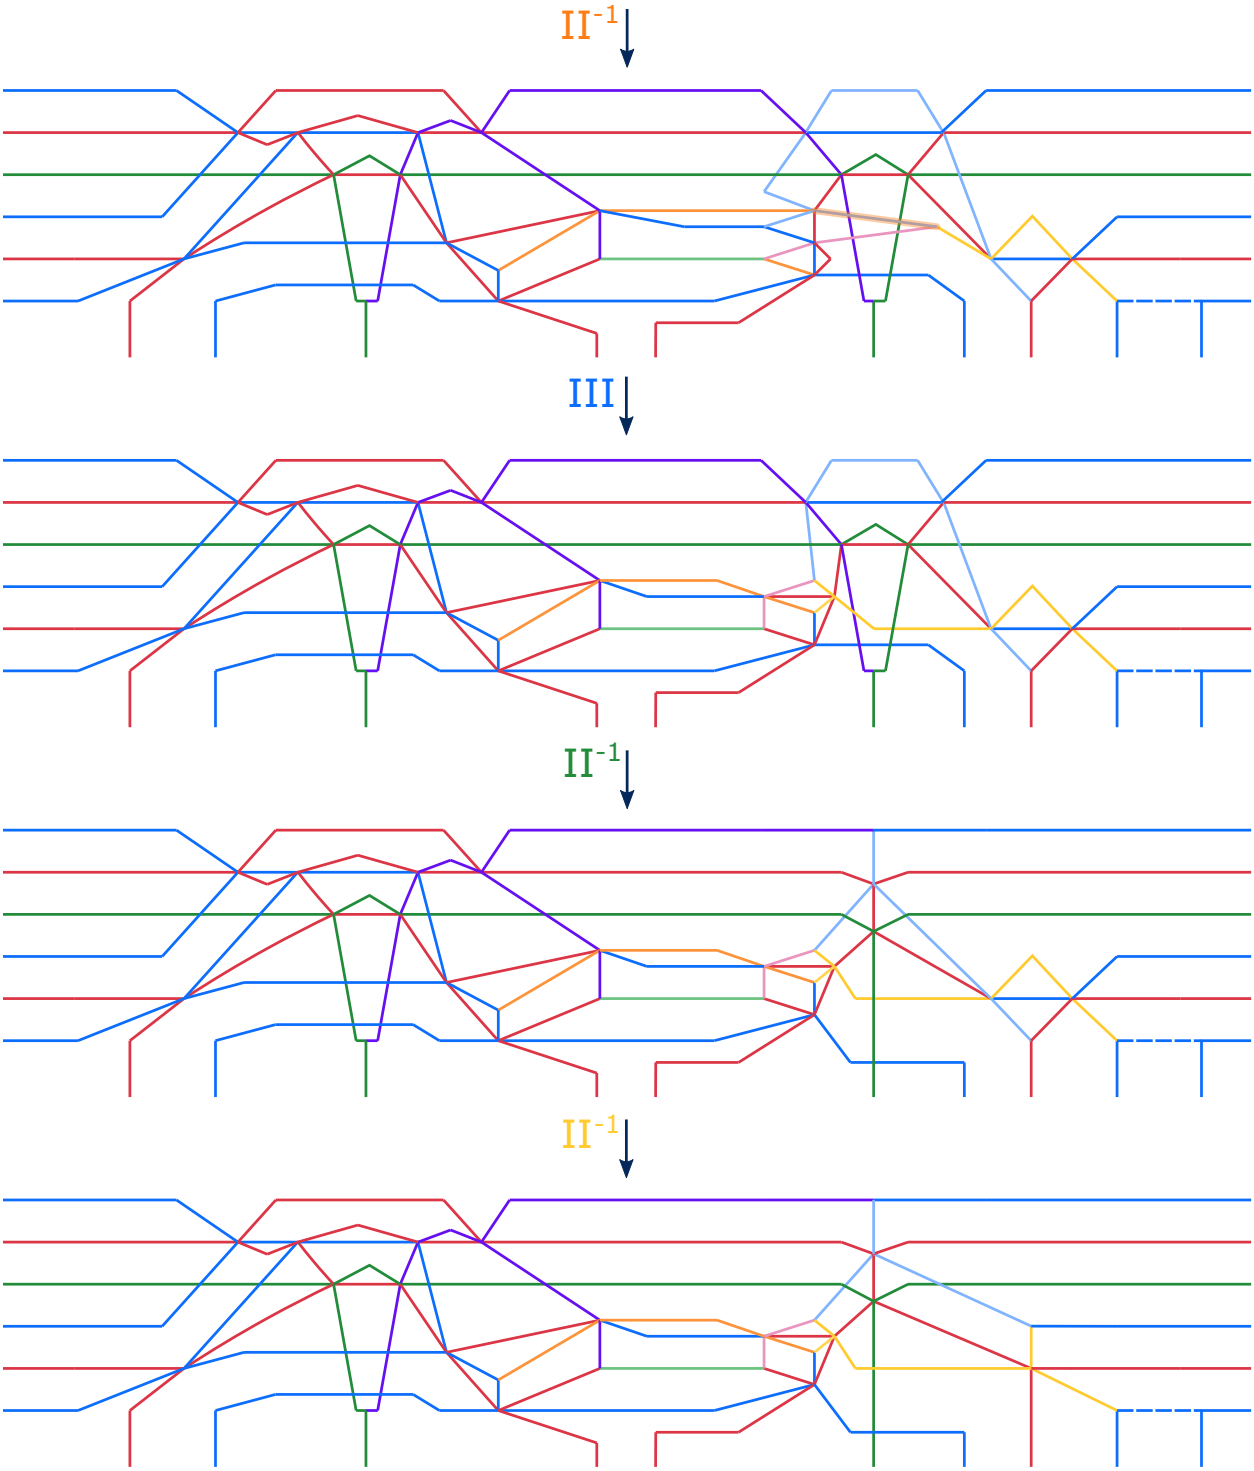

FIGURE A.4. A sequence of weave equivalence moves removing the geometric intersections between γ_1 and γ_0.

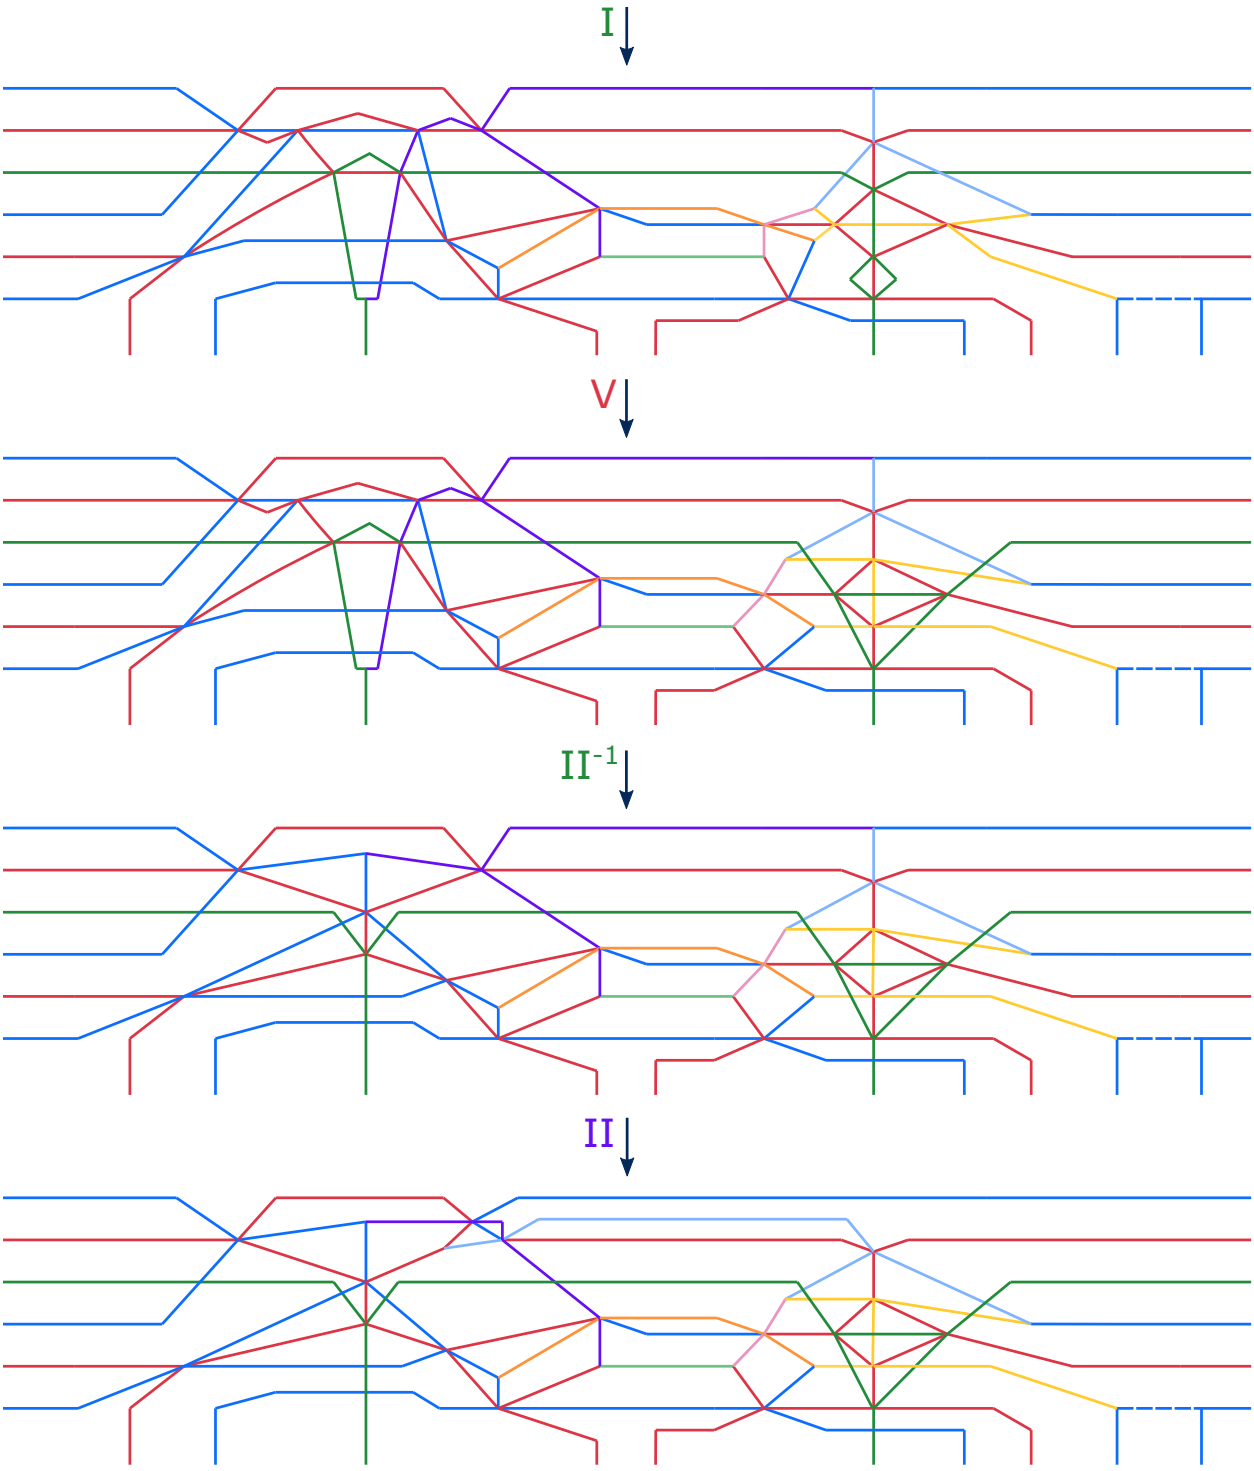

FIGURE A.5. From top to bottom: a candy twist to introduce the necessary $\sigma_2 - \sigma_3$ hexavalent vertices appearing in ϑ_2; Move V; simplification of the left half of the weave starting with a pair of reverse push-throughs involving the leftmost cycle in the weave; a push-through involving γ_0 and γ_3.

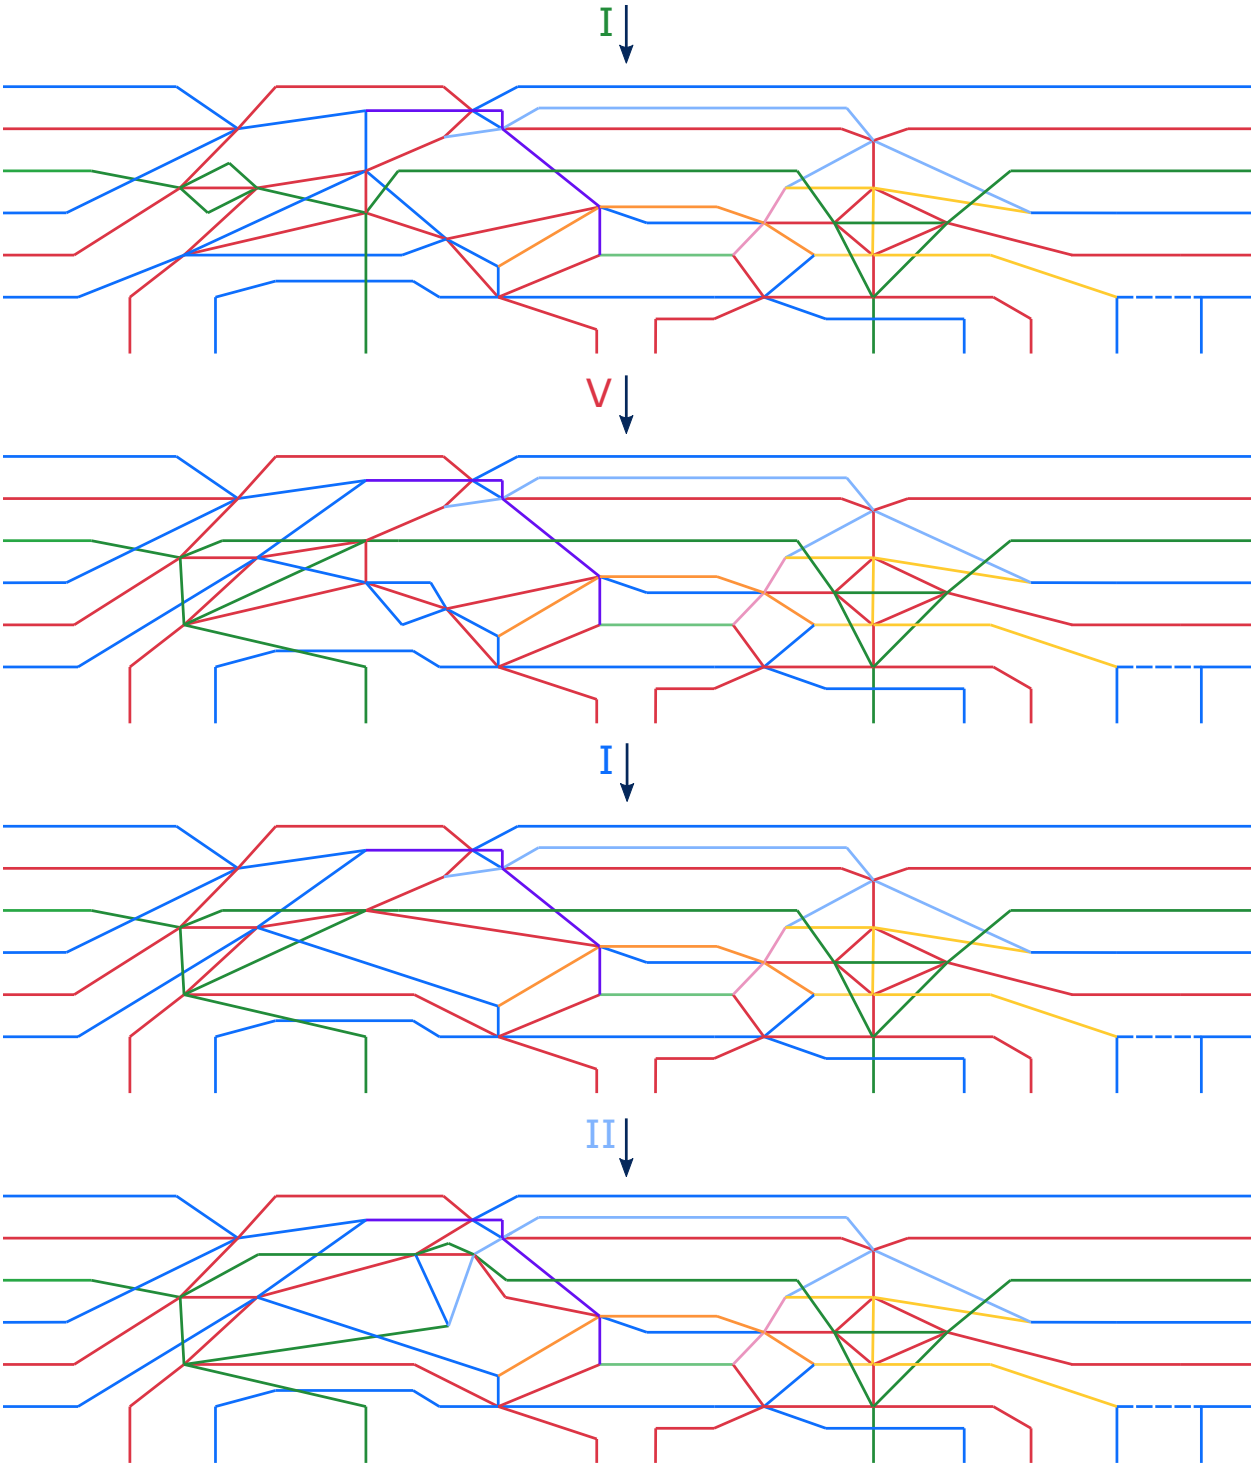

FIGURE A.6. A series of weave equivalence moves continuing to simplify the left half of the weave

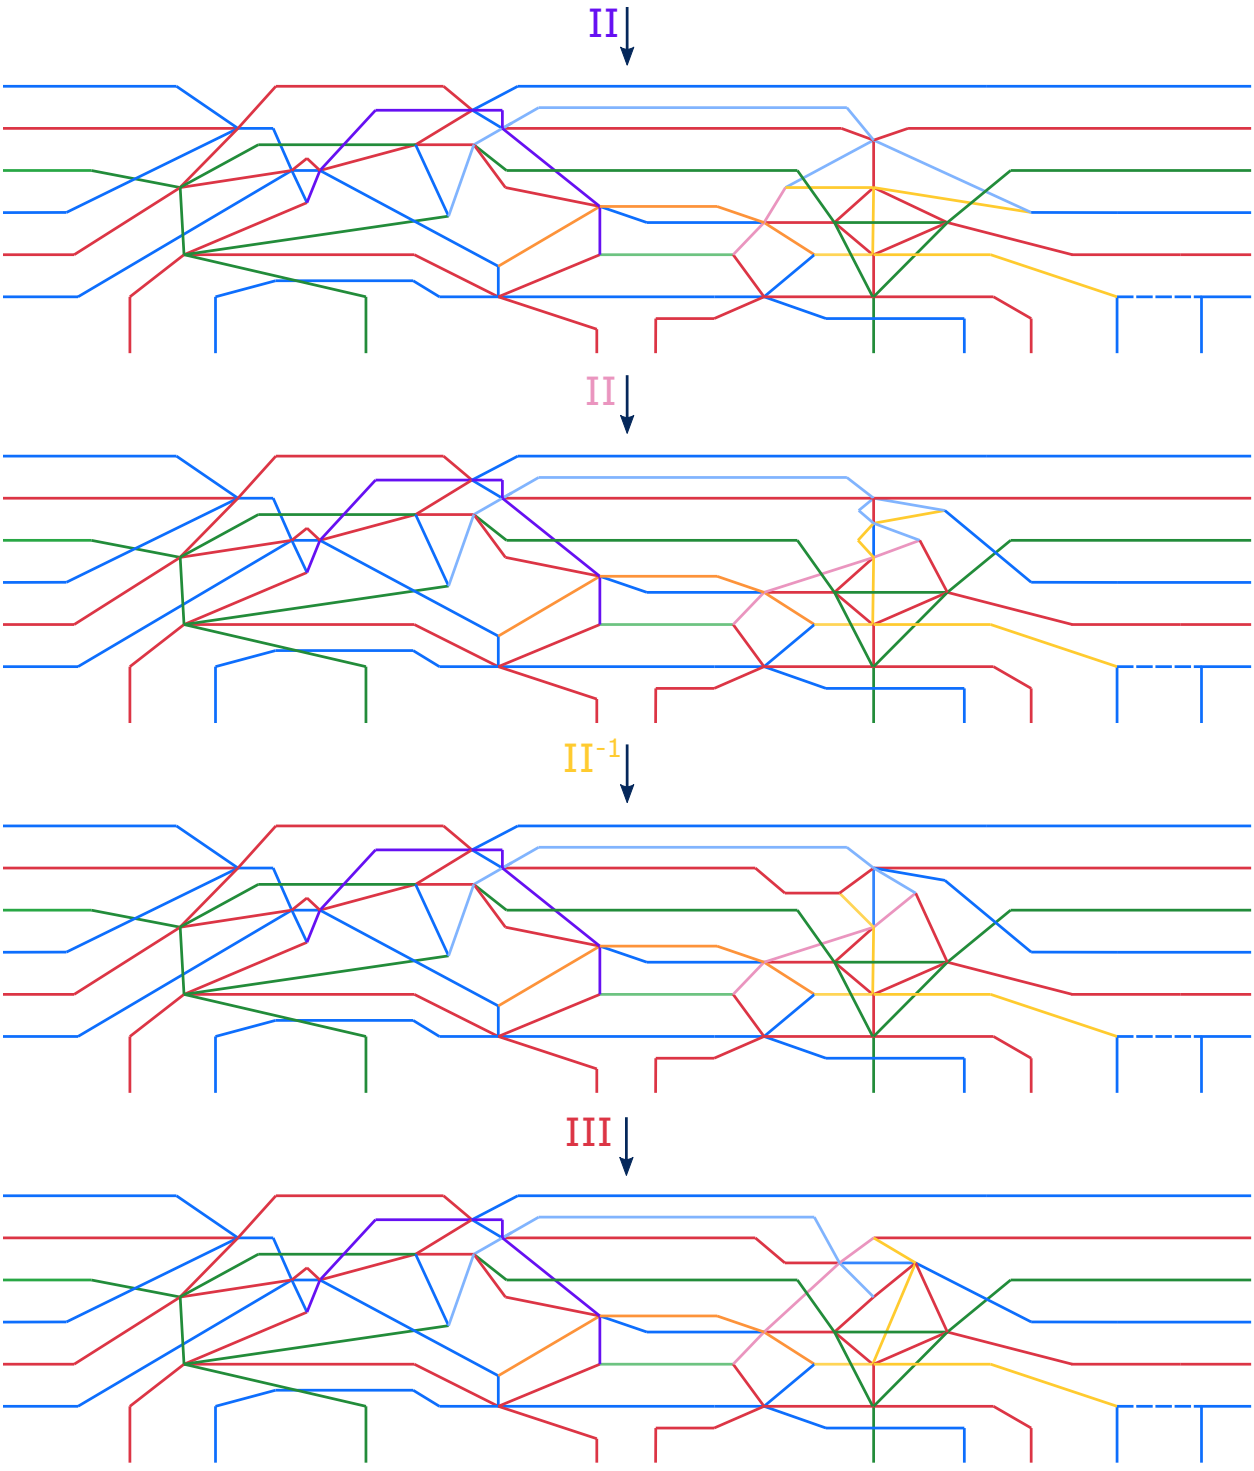

FIGURE A.7. A series of weave equivalence moves continuing to simplify the right half of the weave

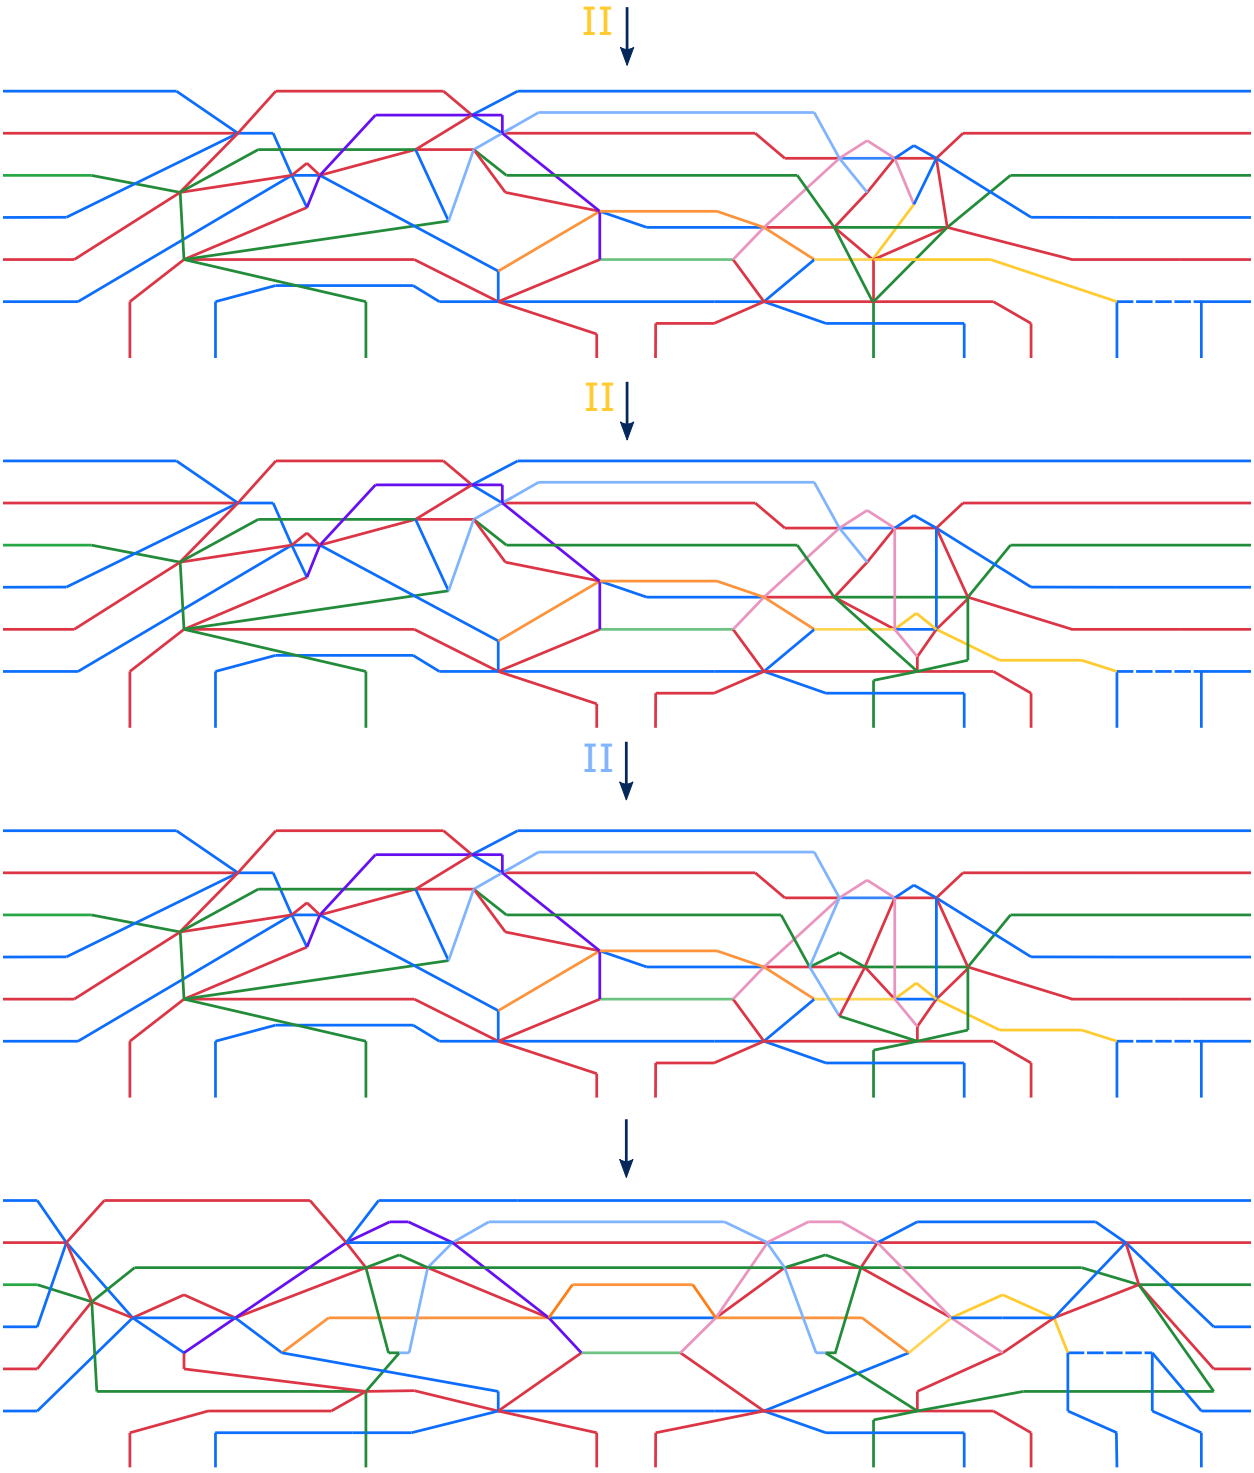

FIGURE A.8. The final weave appearing here differs from the one above it by planar isotopy and freely applying Move IV. It is readily identified with the concatenation of the trace of ϑ_2 to the initial weave appearing in Figure A.1 up to relabeling the cycles by the permutation $(1\,4\,0\,3)$.

Bibliography

[ABL22] B. H. An, Y. Bae, and E. Lee, *Lagrangian fillings for Legendrian links of finite or affine Dynkin type*, 2022, arxiv:2201.00208.

[Ad90] V. I. Arnol′ d, *Singularities of caustics and wave fronts*, Mathematics and its Applications (Soviet Series), vol. 62, Kluwer Academic Publishers Group, Dordrecht, 1990.

[AIK22] T. Asaka, T. Ishibashi, and S. Kano, *Earthquake theorem for cluster algebras of finite type*, arxiv:2206.15226, 2022.

[ASS12] I. Assem, R. Schiffler, and V. Shramchenko, *Cluster automorphisms*, Proc. Lond. Math. Soc. (3) **104** (2012), no. 6, 1271–1302.

[BS15] T. Bridgeland and I. Smith, *Quadratic differentials as stability conditions*, Publ. Math. Inst. Hautes Études Sci. **121** (2015), 155–278.

[Cas21] R. Casals, *Lagrangian skeleta and plane curve singularities*, JFPTA **Viterbo 60** (2021).

[CDR+16] E. Cormier, P. Dillery, J. Resh, K. Serhiyenko, and J. Whelan, *Minimal length maximal green sequences and triangulations of polygons*, J. Algebraic Combin. **44** (2016), no. 4, 905–930.

[CG22] R. Casals and H. Gao, *Infinitely many Lagrangian fillings*, Ann. of Math. (2) **195** (2022), no. 1, 207–249.

[CGGS20] R. Casals, E. Gorsky, M. Gorsky, and J. Simental, *Algebraic weaves and braid varieties*, arXiv:2012.06931, 2020.

[Che02] Y. Chekanov, *Differential algebra of Legendrian links*, Invent. Math. **150** (2002), no. 3, 441–483.

[CL22] R. Casals and W. Li, *Conjugate fillings and legendrian weaves*, 2022.

[CLSBW23] R. Casals, I. Le, M. Sherman-Bennett, and D. Weng, *Legendrian weaves for positroids*, In preparation, 2023.

[CN21] R. Casals and L. Ng, *Braid loops with infinite monodromy on the Legendrian contact DGA*, J. Topology (to appear) (2021).

[CS22] O. Capovilla-Searle, *Infinitely many planar fillings and symplectic milnor fibers*, arxiv:2201.03081, 2022.

[CW22] R. Casals and D. Weng, *Microlocal theory of Legendrian links and cluster algebras*, arxiv:2204.13244, 2022.

[CZ21] R. Casals and E. Zaslow, *Legendrian weaves*, Geom. Topol. (2021).

[EHK16] T. Ekholm, K. Honda, and T. Kálmán, *Legendrian knots and exact Lagrangian cobordisms*, J. Eur. Math. Soc. (JEMS) **18** (2016), no. 11, 2627–2689.

[EN19] J. Etnyre and L. Ng, *Legendrian contact homology in* $\mathbb{R}^3$, https://arxiv.org/pdf/1811.10966.pdf, 2019.

[EP96] Y. Eliashberg and L. Polterovich, *Local Lagrangian 2-knots are trivial*, Ann. of Math. (2) **144** (1996), no. 1, 61–76.

[Eul64] L. Euler, *Specimen algorithmi singularis*, Novi Commentarii academiae scientiraum Petropolitanae **9** (1764), 53–69.

[FG06a] V. V. Fock and A. B. Goncharov, *Cluster x-varieties, amalgamation, and Poisson-Lie groups*, Algebraic geometry and number theory, Progr. Math., vol. 253, Birkhäuser Boston, Boston, MA, 2006, pp. 27–68.

[FG06b] V. Fock and A. Goncharov, *Moduli spaces of local systems and higher Teichmüller theory*, Publ. Math. Inst. Hautes Études Sci. (2006), no. 103, 1–211.

[FOOO09] K. Fukaya, Y.-G. Oh, H. Ohta, and K. Ono, *Lagrangian intersection Floer theory: anomaly and obstruction. Part II*, AMS/IP Studies in Advanced Mathematics, vol. 46, American Mathematical Society, Providence, RI; International Press, Somerville, MA, 2009.

[FPST22] S. Fomin, P. Pylyavskyy, E. Shustin, and D. Thurston, *Morsifications and mutations*, J. Lond. Math. Soc. (2) **105** (2022), no. 4, 2478–2554.

[Fra49] J. S. Frame, *Classroom Notes: Continued Fractions and Matrices*, Amer. Math. Monthly **56** (1949), no. 2, 98–103.

[Fra18] C. Fraser, *Braid group symmetries of Grassmannian cluster algebras*, 2018, 1702.00385.

[FST08] S. Fomin, M. Shapiro, and D. Thurston, *Cluster algebras and triangulated surfaces. I. Cluster complexes*, Acta Math. **201** (2008), no. 1, 83–146.

[FST12] A. Felikson, M. Shapiro, and P. Tumarkin, *Skew-symmetric cluster algebras of finite mutation type*, J. Eur. Math. Soc. (JEMS) **14** (2012), no. 4, 1135–1180.

[FT18] S. Fomin and D. Thurston, *Cluster algebras and triangulated surfaces Part II: Lambda lengths*, Mem. Amer. Math. Soc. **255** (2018), no. 1223, v+97.

[FWZ20a] S. Fomin, L. Williams, and A. Zelevinsky, *Introduction to cluster algebras: Chapters 1-3*, arXiv:1608.05735, 2020.

[FWZ20b] ———, *Introduction to cluster algebras: Chapters 4-5*, arXiv:1707.07190, 2020.

[FZ03] S. Fomin and A. Zelevinsky, *Cluster algebras. II. Finite type classification*, Invent. Math. **154** (2003), no. 1, 63–121.

[Gei08] H. Geiges, *An introduction to contact topology*, Cambridge Studies in Advanced Mathematics, vol. 109, Cambridge University Press, Cambridge, 2008.

[GKS12] S. Guillermou, M. Kashiwara, and P. Schapira, *Sheaf quantization of Hamiltonian isotopies and applications to nondisplaceability problems*, Duke Math. J. **161** (2012), no. 2, 201–245.

[GSW20a] H. Gao, L. Shen, and D. Weng, *Augmentations, fillings, and clusters*, arXiv:2008.10793, 2020.

[GSW20b] ———, *Positive braid links with infinitely many fillings*, arXiv:2009.00499, 2020.

[Hug21] J. Hughes, *Lagrangian fillings in A-type and their Kálmán loop orbits*, arxiv:2109.09662, 2021, 2109.09662.

[Hug22] ———, *Weave realizability for D-type*, Algebr. Geom. Topol. (to appear) (2022).

[IK21] T. Ishibashi and S. Kano, *Algebraic entropy of sign-stable mutation loops*, Geom. Dedicata **214** (2021), 79–118.

[Ish19] T. Ishibashi, *On a Nielsen-Thurston classification theory for cluster modular groups*, Ann. Inst. Fourier (Grenoble) **69** (2019), no. 2, 515–560.

[Ish20] ———, *Presentations of cluster modular groups and generation by cluster Dehn twists*, SIGMA Symmetry Integrability Geom. Methods Appl. **16** (2020), Paper No. 025, 22.

[Kál05] T. Kálmán, *Contact homology and one parameter families of Legendrian knots*, Geom. Topol. **9** (2005), 2013–2078.

[Kál06] ———, *Braid-positive Legendrian links*, Int. Math. Res. Not. (2006), Art ID 14874, 29.

[Kar20] C. Karlsson, *A note on coherent orientations for exact Lagrangian cobordisms*, Quantum Topol. **11** (2020), no. 1, 1–54.

[KG21] D. Kaufman and Z. Greenberg, *Cluster modular groups of affine and doubly extended cluster algebras*, 2021, arxiv:2107.10334.

[KS85] M. Kashiwara and P. Schapira, *Microlocal study of sheaves*, Astérisque (1985), no. 128, 235, Corrections to this article can be found in Astérisque No. 130, p. 209.

[KS90] ———, *Sheaves on manifolds*, Grundlehren der Mathematischen Wissenschaften, vol. 292, Springer-Verlag, Berlin, 1990, With a chapter in French by Christian Houzel.

[Li21] W. Li, *Lagrangian cobordism functor in microlocal sheaf theory*, arxiv:2108.10914, 2021.

[Ng03] L. L. Ng, *Computable Legendrian invariants*, Topology **42** (2003), no. 1, 55–82.

[Pan17] Y. Pan, *Exact Lagrangian fillings of Legendrian* $(2, n)$ *torus links*, Pacific J. Math. **289** (2017), no. 2, 417–441.

[Pos06] A. Postnikov, *Total positivity, Grassmannians, and networks*, 2006, math/0609764.

[Pou14] L. Pournin, *The diameter of associahedra*, Adv. Math. **259** (2014), 13–42.

[Reg13] A. Regev, *A bijection between triangulations and* 312*-avoiding permutations*, (2013).

[STT88] D. D. Sleator, R. E. Tarjan, and W. P. Thurston, *Rotation distance, triangulations, and hyperbolic geometry*, J. Amer. Math. Soc. **1** (1988), no. 3, 647–681.

[STWZ19] V. Shende, D. Treumann, H. Williams, and E. Zaslow, *Cluster varieties from Legendrian knots*, Duke Math. J. **168** (2019), no. 15, 2801–2871.

[STZ17] V. Shende, D. Treumann, and E. Zaslow, *Legendrian knots and constructible sheaves*, Invent. Math. **207** (2017), no. 3, 1031–1133.

[TZ18] D. Treumann and E. Zaslow, *Cubic planar graphs and Legendrian surface theory*, Adv. Theor. Math. Phys. **22** (2018), no. 5, 1289–1345.

[Ust06] A. V. Ustinov, *A short proof of Euler's identity for continuants*, Mat. Zametki **79** (2006), no. 1, 155–156.

[Vat10] D. F. Vatne, *The mutation class of D_n quivers*, Comm. Algebra **38** (2010), no. 3, 1137–1146.