

LEGENDRIAN WEAVES

– N-GRAPH CALCULUS, FLAG MODULI AND APPLICATIONS –

ROGER CASALS AND ERIC ZASLOW

ABSTRACT. We study a class of Legendrian surfaces in contact five-folds by encoding their wavefronts via planar combinatorial structures. We refer to these surfaces as Legendrian weaves, and to the combinatorial objects as N -graphs. First, we develop a diagrammatic calculus which encodes contact geometric operations on Legendrian surfaces as multi-colored planar combinatorics. Second, we present an algebro-geometric characterization for the moduli space of microlocal constructible sheaves associated to these Legendrian surfaces. Then we use these N -graphs and the flag moduli description of these Legendrian invariants for several new applications to contact and symplectic topology.

Applications include showing that any finite group can be realized as a subquotient of a 3-dimensional Lagrangian concordance monoid for a Legendrian surface in $(J^1\mathbb{S}^2, \xi_{st})$, a new construction of infinitely many exact Lagrangian fillings for Legendrian links in (\mathbb{S}^3, ξ_{st}) , and performing \mathbb{F}_q -rational point counts that distinguish Legendrian surfaces in (\mathbb{R}^5, ξ_{st}) . In addition, the manuscript develops the notion of Legendrian mutation, studying microlocal monodromies and their transformations. The appendix illustrates the connection between our N -graph calculus for Lagrangian cobordisms and Elias-Khovanov-Williamson’s Soergel Calculus.

Contents

1.	Introduction	3
1.1.	Summary of Contributions	4
1.2.	Main Results	5
2.	N-graphs and Legendrian Weaves	12
2.1.	N -graphs	13
2.2.	Singularities of wavefronts	13
2.2.1.	The A_1^2 germ	14
2.2.2.	The A_1^3 germ	14
2.2.3.	The D_4^- germ	15
2.3.	Legendrian Weaves	15
2.4.	Smooth Topology of Weaves	18
2.5.	Combinatorial Homology	20
3.	Combinatorial Constructions	23
3.1.	N -Triangulations	23
3.2.	Local Models	24

2010 *Mathematics Subject Classification*. Primary: 53D10. Secondary: 53D15, 57R17.

3.3.	Global Model	25
3.4.	Bicubic graphs	25
4.	Diagrammatic Calculus For Legendrian Weaves	27
4.1.	Surface Reidemeister Moves	27
4.2.	Legendrian Satellite Weaves	34
4.3.	Cusp Edges and Swallowtail Singularities	35
4.4.	Legendrian Front Calculus with Cusp Singularities	37
4.5.	Legendrian Surgeries	41
4.6.	Example of a Closed Legendrian Weave	48
4.7.	N -Graph Stabilization	50
4.8.	Legendrian Mutations	52
4.9.	Diagrammatic Rules for N -graph Mutations	57
4.10.	Sufficiency For Stabilized Legendrians	58
5.	Flag Moduli Spaces	62
5.1.	Preliminaries on the Flag Variety	62
5.2.	Description of the Flag Moduli Space of an N -graph	63
5.3.	Sheaf Description of Flag Moduli and Invariance	64
5.4.	Local Flag Moduli Computations	67
5.5.	Flag Moduli under Legendrian Surgeries	70
5.6.	Non-characterstic Property of Stabilization	71
6.	Applications and Vexillary Computations	72
6.1.	First Pair of Computations	72
6.1.1.	The Legendrian Clifford Torus	72
6.1.2.	The Double of the 4-Triangle	73
6.2.	Symmetry groups for Legendrian weaves	75
6.3.	Flag Moduli and Bipartite Graphs	77
6.4.	Non-isotopic Links of Legendrian Spheres	78
7.	Microlocal Monodromies and Lagrangian Fillings	83
7.1.	Exact Lagrangian Cobordisms	83
7.1.1.	The geometric setup	84
7.1.2.	Free N -Graphs	84
7.1.3.	Explicit Examples of Lagrangian Fillings	86
7.2.	Microlocal monodromies and cluster structures	89
7.2.1.	Microlocal monodromies as cluster coordinates	91
7.2.2.	Legendrian Mutations are cluster transformations	92
7.3.	N -graph Realization of Quiver Mutations	96
8.	Moduli Space for N-triangles and Non-Abelianization	101

8.1. Flag moduli space of the N -triangle	102
8.1.1. Tetrahedral Triangulations at $N = 3$ and $N = 4$	105
8.2. A Computation of the Non-Abelianization Map	106
Example: Tetrahedron with 3-Triangulation.	107
Appendix A. Soergel Calculus and Legendrian Weaves	109
References	112

Legendrian fronts arise naturally in several areas: in topology, as Cerf diagrams of families of smooth functions; in differential equations, as Stokes diagrams of an irregular singularity; and in analysis, as wavefront sets of distributions, generalizing the original context of wavefronts in geometric optics. This article studies Legendrian surfaces through the combinatorics of their wavefronts and develops the theory of N -graphs, planar structures encoding front singularities. The moduli space of simple sheaves microsupported on the Legendrian surface becomes an incidence problem for flags of vector spaces, as dictated by the N -graph. We exploit the connections between the combinatorics of N -graphs, algebraic geometry and cluster algebras to obtain results in contact and symplectic topology.

1. INTRODUCTION

Legendrian knots in contact 3-manifolds [Etn05, Gei08] are central to the study of 3-dimensional contact geometry [Ben83, Eli93, Gom98]. The study of Legendrian knot invariants makes extensive use of their planar front projections, both in the context of Floer theory [EGH00, Che02, Ng03] and microlocal analysis [KS85, GKS12, STZ17]. Higher-dimensional Legendrian submanifolds have proven equally instrumental in the study of higher-dimensional symplectic and contact topology, including the development of Legendrian Kirby Calculus [Eli90, Gom98, CMP19] and Lagrangian skeleta [RSTZ14, Nad17a, Sta18].

In the case of 6-dimensional symplectic manifolds and their 5-dimensional contact boundaries [CE12, CM19], spatial front projections for Legendrian surfaces are available [Ad90, AdG01]. First, this article develops a multi-colored planar diagrammatic calculus for the manipulation of such Legendrian surfaces in 5-dimensional contact manifolds and their Lagrangian projections in 4-dimensional symplectic manifolds. This diagrammatic calculus is first used for the efficient computation of microlocal Legendrian isotopy invariants, as we prove and illustrate throughout the manuscript. Then we provide several new applications, including new results in higher-dimensional contact geometry and low-dimensional symplectic topology. We also expect that this concrete description will prove itself useful for further results, such as the computation of symplectic invariants of Weinstein manifolds [GPS19a, Section 6.4] and homological mirror symmetry [Nad17b, TZ18], see also [CM19, Section 4.4] and Remark 6.1.

Finally, even if Legendrian weaves are a specific class of Legendrian surfaces, we can actually use them to prove new results, such as Theorems 1.6 and 1.8 below, and the diagrammatic calculus presented here is already being used successfully in a variety of recent developments [CG20, Cas20, CGGS, CN20, GSW20a, GSW20b]. In studying this manuscript, we hope that the reader will find these Legendrian surfaces as useful and fascinating as we have.

1.1. Summary of Contributions. Let G be an N -graph¹ drawn on a smooth surface C . The notion of an N -graph, combinatorial in nature, is first defined in Section 2. In a nutshell, our main contributions are as follows:

- A. Diagrammatic Calculus and Legendrian Weaves.** The construction of a Legendrian surface $\Lambda(G)$ in the five-dimensional jet space (J^1C, ξ_{st}) associated to the N -graph G , along with a description of Legendrian Surface Reidemeister moves in terms of combinatorial N -graphs moves. Likewise, we show that Legendrian surgeries and Legendrian mutations, which we introduce, can be reflected by the diagrammatics of N -graphs. This is part of a general calculus of multi-colored planar diagrams that, as we show, captures Legendrian surfaces and 3-dimensional Lagrangian cobordisms between them. The translation from five-dimensional contact topology to such planar diagrammatics allows us to study contact topology through combinatorics and graph theory. In fact, we use this combinatorial perspective to construct Lagrangian and Legendrian surfaces that prove new results in contact topology.
- B. The Microlocal Sheaf Theory of N -Graphs.** A Legendrian surface $\Lambda \subseteq (J^1C, \xi_{\text{st}})$ specifies a category of constructible sheaves on $C \times \mathbb{R}$ with singular support constrained by Λ . When $\Lambda = \Lambda(G)$ for an N -graph G , we show that the moduli stack of objects $\mathcal{M}(G)$ has a combinatorial description in terms of flag varieties, which we introduce in Section 5. This space solves an incidence moduli problem for flags of subspaces in an N -dimensional k -vector space V , with k a field, as dictated by the N -graph G . This stack is typically an algebraic variety and can be studied by algebraic geometric and representation-theoretic techniques. Following [GKS12, STZ17, TZ18], this space is shown to be a Legendrian invariant for surfaces $\Lambda(G)$ and can be used to distinguish Legendrian isotopy types. In addition, we explicitly give formulas for the microlocal monodromies along certain cycles of $H_1(\Lambda(G), \mathbb{Z})$ in terms of generalized cross-ratios of flags, and their transformation under Legendrian mutations.
- C. Applications of N -Graph Calculus.** First, in Section 6 we use the diagrammatics in (A) to study the flag moduli spaces $\mathcal{M}(G)$ in (B), including their rational point counts over finite fields \mathbb{F}_q . This allows us to distinguish many Legendrian surfaces, up to Legendrian isotopy and, independently, show that for any finite group \mathbb{G} , there exists a Legendrian surface in (\mathbb{R}^5, ξ) whose 3-dimensional Lagrangian concordance monoid has \mathbb{G} as a subquotient. Second, Section 7 explains how to apply N -graph calculus to systematically study Lagrangian fillings of Legendrian links in $(\mathbb{S}^3, \xi_{\text{st}})$. In particular, we use Legendrian mutations to give new families of Legendrian links which admit infinitely many Lagrangian fillings.

Finally, given N -triangulations (C, τ) of the smooth surface C , we construct N -graphs $G(\tau)$ such that the Lagrangian projections of the Legendrian surfaces $\Lambda(G(\tau))$ relate to the Goncharov-Kenyon conjugate surfaces [GK13, Section 1.1.1] associated to an N -triangulation.² In Section 3 we provide the construction of $G(\tau)$. In Section 8, we provide an example of how Hitchin’s non-Abelianization map is described from this viewpoint. This provides a context for the symplectic study of the cluster structures associated to moduli spaces of framed local systems of Fock-Goncharov [FG06a, FG06b], and certain classes of Gaiotto-Moore-Neitzke’s spectral networks [GMN10, GMN13, Nei14, GMN14]. In particular, the microlocal sheaf theory of

¹Informally, an N -graph $G \subseteq C$ is a collection of trivalent graphs on C decorated with labels $i \in [1, N]$ such that graphs with successive labels can only intersect at hexavalent vertices, where the six radiating half-edges on the surface must interlace. See Definition 2.2 for details, and note that a 2-graph is simply an embedded trivalent graph.

²See also [Gon17, Section 2.1], and [STWZ19, Section 4.2] describes the conjugate surface as a Lagrangian.

$\Lambda(G(\tau))$ connects, through the moduli space $\mathcal{M}(G(\tau))$ in (B), with their spaces of flag configurations [Gon17, Section 3].

1.2. **Main Results.** We now elaborate upon these topics and state our results.




FIGURE 1. 3-graph (left) and 4-graph (right) on the 2-sphere \mathbb{S}^2 . These correspond to Legendrian surfaces in the contact 5-space $(J^1\mathbb{S}^2, \xi_{\text{st}})$, respectively of genus 3 and 4. The geometric meaning of these two figures is explained in detail in Section 2, in Subsections 2.1 and 2.3.

Diagrammatic Calculus and Legendrian Weaves. Weinstein manifolds [CE12, CE14, CM19], the symplectic counterpart of Stein manifolds, place Legendrian submanifolds at the forefront of higher-dimensional contact and symplectic topology. In this manuscript, we define and study a new class of Legendrian surfaces $\Lambda(G)$ in contact 5-manifolds, associated to an N -graph G , building on our previous works [CM18, TZ18]. Prior work on Legendrian surfaces [ENS18, She19] has focused on the class of Legendrian tori $\Lambda_K \subseteq (T^\infty\mathbb{S}^3, \xi_{\text{std}})$ arising as the conormal torus of a smooth knot $K \subseteq \mathbb{S}^3$. The Legendrian surfaces $\Lambda(G)$ we study provide a second infinite family of Legendrian submanifolds whose contact topology and sheaf invariants can be understood. Their geometry is governed by the combinatorial data of the N -graph G . Figure 1 depicts two examples of N -graphs, representing Legendrian surfaces of genus 3 (left) and 4 (right).

We study three geometric operations for Legendrian surfaces in 5-dimensional contact manifolds. These are Legendrian isotopies [Ad90, CE12, Gei08], exact Lagrangian cobordisms [Ad76, BST15, EHK16], and Legendrian mutations, which we define in Section 4. Lagrangian cobordisms of indices 1 and 2 correspond to Legendrian 0- and 1-surgeries. We establish a correspondence between each of these three types of geometric operations and the combinatorics of N -graphs. In addition, we describe a combinatorial stabilization of an N -graph, which can be understood as a five-dimensional analogue of the Markov stabilization of a Legendrian braid [Rol76, PS97]. Part of these results are summarized in the following two theorems (see Section 4 for details), which are developed in the text:

Theorem 1.1 (Diagrammatics for Legendrian Weave Calculus I). *Let G be a local N -graph. The combinatorial moves in Figures 2 and 3 are Legendrian isotopies for $\Lambda(G)$.* \square

Theorem 1.2 (Diagrammatics for Legendrian Weave Calculus II). *Let G be a local N -graph. The combinatorial moves in Figure 4 are Legendrian surgeries, of indices 0, 1 and 2, Legendrian mutations and connected sums with the standard and Clifford tori.* \square

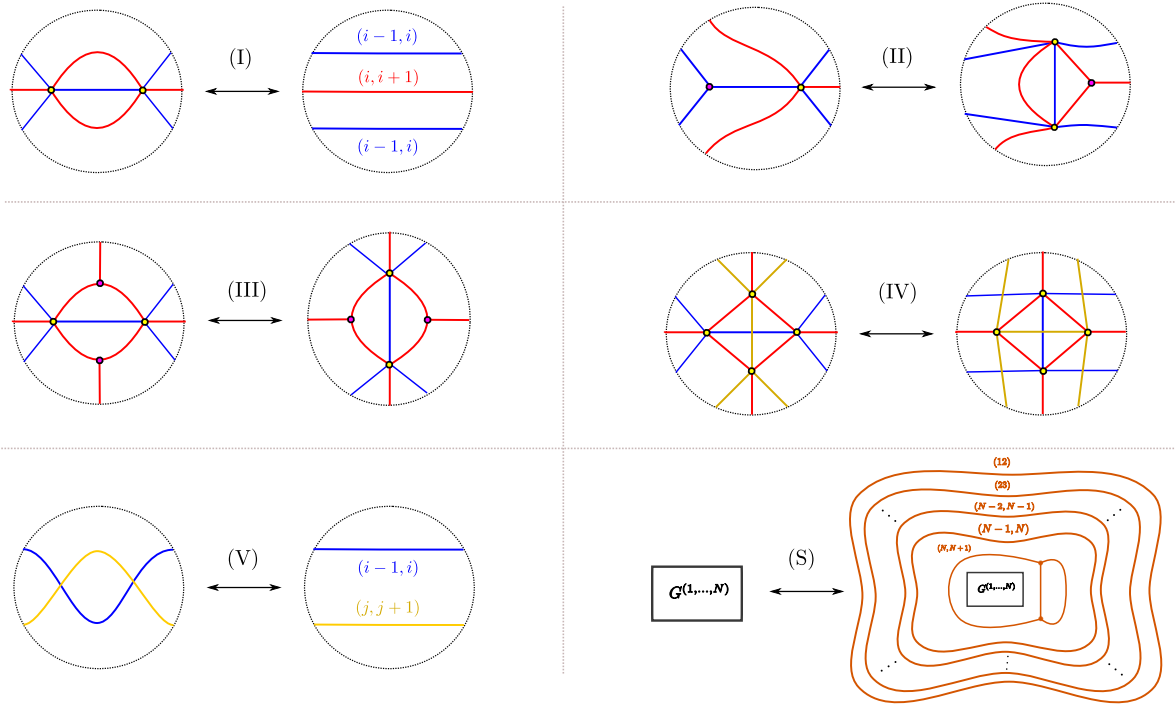


FIGURE 2. Combinatorial Moves for Legendrian Isotopies of Surfaces $\Lambda(G)$. Moves I–V are local Legendrian isotopies in the 1-jet space $(J^1\mathbb{R}^2, \xi_{\text{st}})$. Move S in the lower right is local in $(J^1\mathbb{S}^2, \xi_{\text{st}})$ after satelliting to the Legendrian unknot $\Lambda_0 \subseteq (\mathbb{R}^5, \xi_{\text{st}})$. See Section 2 for precise details on the geometric concepts represented by these pictures.

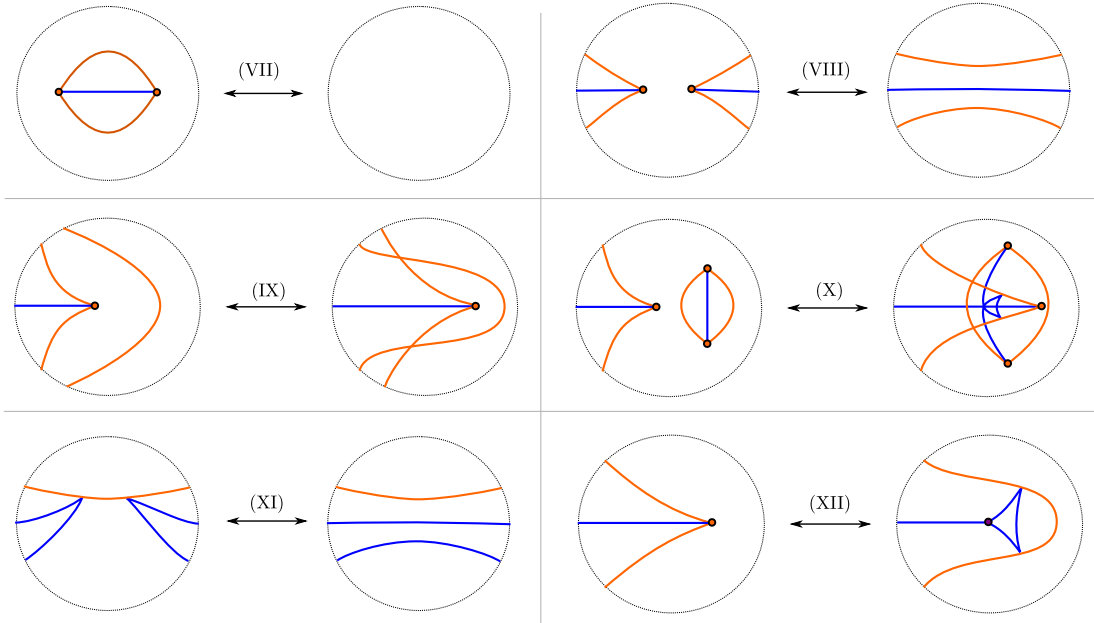


FIGURE 3. Combinatorial Moves for Legendrian Isotopies of Surfaces $\Lambda(G)$. These are homotopies of spatial wavefronts involving A_3 -swallowtail singularities. Section 2 explains the geometric meaning of these pictures.

Theorems 1.1 and 1.2 provide an efficient diagrammatic calculus to manipulate the Legendrian surfaces $\Lambda(G)$ associated to N -graphs G . We refer to the Legendrian surfaces $\Lambda(G)$ as




FIGURE 4. Table of Combinatorial Moves for Surfaces $\Lambda(G)$ corresponding to Legendrian Surgeries, mutations and tori connected sums. See Section 2 for details.

Legendrian weaves, due to the resemble of their Legendrian fronts to a weaving pattern – see Definition 2.7. Theorems 1.1 and 1.2 are geometric in nature and are proven by manipulating Legendrian fronts for Legendrian surfaces in five dimensions. This is the content of Section 4, as part of our study of generic three-dimensional front singularities and their homotopies. In addition, Section 3 provides several combinatorial constructions of Legendrian surfaces $\Lambda \subseteq (\mathbb{S}^5, \xi_{\text{st}})$ which are used in our applications in Sections 6, 7 and 8.

Remark 1.3. The Legendrian weaves $\Lambda(G) \subseteq (J^1C, \xi_{\text{st}})$ associated to an N -graph $G \subseteq C$ admit spatial wavefronts $\pi(\Lambda(G)) \subseteq C \times \mathbb{R}$ with front singularities solely of types³ A_1^2 , A_1^3 and D_4^- , following V.I. Arnol’d’s notation [Ad76, Ad90]. That said, their satellites $\iota(\Lambda(G)) \subseteq (Y, \xi_{\text{st}})$ typically acquire A_2 , A_2A_1 and A_3 singularities. Satellites will be introduced and discussed in Section 4.2; as an example, the satellite of $\Lambda(G)$ along the standard Legendrian unknot $\Lambda_0 \subseteq (\mathbb{R}^5, \xi_{\text{st}})$ necessarily develops A_2 -singularities. In addition, the standard 5-dimensional Legendrian Reidemeister surface moves include the creation of A_3 singularities, and the interaction of A_2 and A_3 singularities yield a D_4^+ singularity. These Legendrian singularities and 3-dimensional Reidemeister moves will also be discussed in Section 4. \square

The Microlocal Sheaf Theory of N -Graphs. The relationship between sheaf theory and contact and symplectic geometry [NZ09, Nad09, GKS12, GS14] provides invariants of Lagrangian and Legendrian submanifolds up to Hamiltonian and contact isotopies [STZ17, STWZ19, CG20]. These invariants are an alternative to the more analytical Floer-theoretic methods [EES05b, EENS13a, EENS13b], and have recently been shown to contain equivalent data [GPS19a, GPS19b, GPS19c].

Let G be an N -graph on C , $\Lambda(G) \subseteq J^1(C)$ its Legendrian surface, and $\mathcal{C}(G)$ the category of simple constructible sheaves on $C \times \mathbb{R}$ microlocally supported along $\Lambda(G)$. In Section

³The A_1^2 -singularity corresponds to a crossing, and the A_1^3 -singularity is given by three planes intersecting transversely at a point. The A_2 -singularity corresponds to a simple cusp, A_3 -singularities are swallowtails, and A_2A_1 -singularities are obtained by intersecting a cusp with a linear space.

5, we describe the moduli space of objects in $\mathcal{C}(G)$ in terms of the combinatorics of G . Specifically, we define the *flag moduli space* $\mathcal{M}(G)$ of an N -graph $G \subseteq C$, an algebraic stack – often a variety – as being described by explicit relations among elements in the flag variety $GL(N, k)/B$, where B is the Borel subgroup of upper triangular matrices. Already when $N = 2$, the number of rational \mathbb{F}_q -points of $\mathcal{M}(G)$, for a finite field \mathbb{F}_q is, up to a factor, the chromatic polynomial of the dual graph evaluated at $q + 1 = |(GL(2, \mathbb{F}_q)/B)(\mathbb{F}_q)|$ [TZ18], and hence the moduli stack $\mathcal{M}(G)$ geometrizes a familiar graph-theoretic construction.

For general N , this algebraic space $\mathcal{M}(G)$ is the moduli space of an incidence problem between flags and their stabilizing monodromies. It has two particular virtues. First, $\mathcal{M}(G)$ changes explicitly under certain combinatorial moves of the N -graph G — thus, each time we can simplify G with our moves from Theorems 1.1 and 1.2, we get closer to solving the moduli problem via purely diagrammatic techniques. Second, $\mathcal{M}(G)$ is an invariant of the Legendrian isotopy class of $\Lambda(G) \subseteq J^1(C)$. In short, $\mathcal{M}(G)$ is defined purely in terms of the combinatorics of the N -graph G , in a manner we understand, and we show it geometrically describes the following invariant:

Theorem 1.4. *Let C be a closed, smooth surface and $G \subseteq C$ an N -graph. The flag moduli space $\mathcal{M}(G)$ is isomorphic to the moduli space of microlocal rank-one sheaves⁴ on $C \times \mathbb{R}$ microlocally supported along $\Lambda(G) \subseteq (J^1C, \xi_{st})$. \square*

After the work of Guillermou-Kashiwara-Schapira [GKS12], which constructs an equivalence of sheaf categories from a Legendrian isotopy, we conclude that the algebraic isomorphism type of the moduli stack $\mathcal{M}(G)$ is a Legendrian isotopy invariant of the Legendrian surface $\Lambda(G) \subseteq (J^1C, \xi_{st})$. In fact, it will remain a Legendrian isotopy invariant for certain satellites along $C \subseteq (\mathbb{R}^5, \xi_{st})$, yielding a Legendrian invariant for $\Lambda(G) \subseteq (\mathbb{R}^5, \xi_{st})$. Theorem 1.4, proven in Section 5, is a generalization to $N \geq 2$ of [TZ18, Section 4] and the 2-dimensional surface analogue of the results in [STZ17, STWZ19], where the computation of the moduli space of microlocal rank-1 sheaves for 1-dimensional Legendrian braid closures in \mathbb{R}^3 is expressed in algebraic combinatorial terms.

Applications of N -Graph Calculus. Sections 6, 7 and 8 exhibit a gallery of computations and uses of the flag moduli space $\mathcal{M}(G)$, including the study of $\mathcal{M}(G)$ as a complex variety and its finite \mathbb{F}_q -counts. For instance, our techniques readily prove the following sample result:

Theorem 1.5 (Flag Moduli for Ladder Graphs). *Let $\mathcal{L}_n \subseteq \mathbb{S}^2$ be the $(2n)$ -runged ladder 3-graph of Figure 5, and let \mathbb{F}_q a finite field. Then the flag moduli space $\mathcal{M}(\mathcal{L}_n)$ has orbifold point count*

$$|\mathcal{M}(\mathcal{L}_n)(\mathbb{F}_q)| = \frac{q^{2n-3} - q^{n-2} + q^{n-1} + q - 1}{(q - 1)^2}$$

In particular, the Legendrian 3-links of 2-spheres $\Lambda(\mathcal{L}_n)$ and $\Lambda(\mathcal{L}_m)$ are Legendrian isotopic if and only if $n = m$. \square

The infinitely many Legendrian surfaces $\Lambda(\mathcal{L}_n)$ in Theorem 1.5, $n \in \mathbb{N}$, are pairwise smoothly isotopic. The distinct finite \mathbb{F}_q -counts of their flag moduli space $\mathcal{M}(\Lambda(\mathcal{L}_n))(\mathbb{F}_q)$ give a direct proof that they are not Legendrian isotopic as Legendrian surfaces in (\mathbb{R}^5, ξ_{st}) . Also, adding the ladder 3-graphs in Theorem 1.5 into a face of an arbitrary N -graph G typically changes the flag moduli space of $\mathcal{M}(G)$ and thus produces another Legendrian surface, smoothly isotopic but *not* Legendrian isotopic to $\Lambda(G)$.

⁴Microlocal rank-one sheaves are also called microlocally simple or just simple [KS85, Chapter 7].

In general, the computation of these Legendrian invariants translates into an incidence moduli problem, which can itself be simplified with our diagrammatic techniques, and then possibly solved with methods from algebraic geometry. In particular, we will understand the effect of combinatorial moves for N -graphs G on the Legendrian invariants $\mathcal{M}(\Lambda(G))$. This will frequently allow for the computation of this moduli stack and distinguish Legendrian weaves up to Legendrian isotopy. This yields a wide range of results in the vein of Theorem 1.5, as we will illustrate. From this perspective, Legendrian weaves, which are in general surfaces of any genus, constitute an attractive complement to the family of knot conormals.




FIGURE 5. The bipartite Ladder 3-Graph \mathcal{L}_n , where the right and left sides are identified after $2n$ rungs.

We now illustrate a second application of our flag moduli stacks, detailed in Section 6. Let $\Lambda \subset (\mathbb{S}^5, \xi_{st})$ be an embedded Legendrian surface and let $\mathcal{L}(\Lambda)$ be the space of embedded Legendrian surfaces which are Legendrian isotopic Λ , with base point Λ . Let $\mathbb{L}(\Lambda)$ be the monoid of 3-dimensional exact Lagrangian concordances in the symplectization $(\mathbb{S}^5 \times \mathbb{R}(t), e^t \lambda_{st})$, up to Hamiltonian isotopy, based at Λ . The flag moduli spaces $\mathcal{M}(G)$ will be used to show the following result:

Theorem 1.6. *Let \mathbb{G} be an arbitrary finite group. Then there exists a Legendrian surface $\Lambda_{\mathbb{G}} \subseteq (\mathbb{S}^5, \xi_{st})$ such that*

- (i) \mathbb{G} is a subquotient of the fundamental group $\pi_1(\mathcal{L}(\Lambda_{\mathbb{G}}))$,
- (ii) \mathbb{G} is a subquotient of the 3-dimensional Lagrangian concordance monoid $\mathbb{L}(\Lambda_{\mathbb{G}})$.

In fact, the latter is the image of the former via the graph map $gr : \pi_1(\mathcal{L}(\Lambda)) \rightarrow \mathbb{L}(\Lambda)$.

Theorem 1.6 essentially states that the study of the 3-dimensional Lagrangian concordance monoid can be as complicated as any finite group. The proof of Theorem 1.6 will exhibit the advantage of using combinatorial constructions on an N -graph G to extract contact and symplectic information in 5- and 6-dimensions. Note that for 1-dimensional max-tb Legendrian torus links, T. Kálmán provided finite cyclic subgroups of the 2-dimensional Lagrangian concordance monoid [K05], and J. Sabloff and M. Sullivan provided⁵ finite cyclic subgroups of the 3-dimensional Lagrangian monoid for certain Legendrian surfaces [SS16]. Sections 5, 6 and 7 contain several computations and applications of the flag moduli spaces $\mathcal{M}(G)$.

Remark 1.7. The Legendrian DGA of a Legendrian knot in (\mathbb{R}^3, ξ_{st}) can be computed algorithmically. The computation of Floer-theoretic invariants of general Legendrian submanifolds in arbitrary higher-dimensions represents a challenge [DR11, EES05a, EES05b] — see [RS19a, RS19b] for progress in this direction. The class of Legendrian 2-tori arising as knot conormals is understood [Ng11, ENS18, She19] and our results, in line with Theorem 1.5 and Theorem 1.6, aim at achieving both a geometric and sheaf-theoretic understanding for the class of Legendrian weaves $\Lambda(G)$. \square

⁵The results of [SS16] are stronger in higher-dimensions, but for Legendrian surfaces the only finite subgroups of the special orthogonal group $SO(2)$ must be cyclic — see [SS16, Remark 4.7].

For a third class of applications, consider an N -graph $G \subseteq \mathbb{D}^2$ with boundary. The Lagrangian projection of the Legendrian weave $\Lambda(G)$ yields⁶ an exact Lagrangian filling of a Legendrian link in (\mathbb{S}^3, ξ_{st}) , associated to ∂G . In Section 7 we will construct different N -graphs G_1, G_2 with $\partial G_1 = \partial G_2$, and explain how microlocal monodromies can be used to show that the Lagrangian projections of the Legendrian weaves $\Lambda(G_1)$ and $\Lambda(G_2)$ are *not* Hamiltonian isotopic relative to their 1-dimensional Legendrian boundaries. In fact, N -graph calculus, in combination with Legendrian mutations, allows us to construct *infinitely* many distinct embedded Lagrangian fillings for certain Legendrian knots. The following family of Legendrian links is studied in detail in Subsection 7.3:

Theorem 1.8. *Let $\Lambda_{s,t} = \Lambda(\beta_{s,t}) \subseteq (\mathbb{S}^3, \xi_{st})$ be the Legendrian link given by the standard satellite of the positive braid*

$$\beta_{s,t} = (\sigma_1^3 \sigma_2)(\sigma_1^3 \sigma_2^2)^s \sigma_1^3 \sigma_2 (\sigma_2^2 \sigma_1^3)^t (\sigma_2 \sigma_1^3)(\sigma_2^{t+1} \sigma_1^2 \sigma_2^{s+2}), \quad s, t \in \mathbb{N}, s, t \geq 1.$$

Then $\Lambda_{s,t} \subseteq (\mathbb{S}^3, \xi_{st})$ admits infinitely many embedded exact Lagrangian fillings in $(\mathbb{D}^4, \lambda_{st})$ realized as 3-graphs $G_{s,t} \subseteq \mathbb{D}^2$ and their Legendrian mutations.

The 3-graphs representing the infinitely many Lagrangian fillings in Theorem 1.8 are diagrammatically interesting, with their complexity increasing as we geometrically realize the iterates in an infinite sequence of quiver mutations. For instance, Figure 6 depicts an example of a Lagrangian filling associated to such a 3-graph, obtained after five mutations. Fortunately, the local mutations rules that we develop in Section 4.9 will allow us to control certain infinite sequences of N -graphs mutations and construct infinite sequences of pairwise distinct Lagrangian fillings.




FIGURE 6. The 3-graph for one of the infinitely many Lagrangian fillings of the Legendrian link $\Lambda_{1,1} \subseteq (\mathbb{S}^3, \xi_{st})$, as featured in Theorem 1.8. Iterative 3-graph mutations will yield new 3-graphs $G \subseteq \mathbb{D}^2$ representing pairwise non-Hamiltonian isotopic Lagrangian fillings of $\Lambda_{1,1}$.

⁶A combinatorial criterion for embeddedness, which will be useful, is described in Lemma 7.4.

Theorem 1.8 is an appropriate complement to the recent results [CG20], as the construction of the infinitely many Lagrangian fillings in Theorem 1.8 is obtained directly by Legendrian mutations.⁷ In more generality, Section 7 develops the relation between the cluster algebra associated to the intersection quiver of a Lagrangian filling and the Legendrian mutations from Section 4.8. In particular, N -graph calculus can serve as an effective tool to show that a given Legendrian link admits infinitely many Lagrangian fillings, in case the quiver is of infinite mutation type⁸ and its vertices are represented by mutable 1-cycles in the N -graph G . In fact, any Legendrian link $\Lambda(\beta) \subseteq (\mathbb{S}^3, \xi_{\text{st}})$ associated to a positive braid $\beta \in \text{Br}_N^+$ admits a Lagrangian filling – oftentimes many – given by an N -graph $G \subseteq \mathbb{D}^2$.

A final application of N -graph calculus for Legendrian weaves develops the connection of symplectic topology to V. Fock and A. Goncharov’s cluster varieties of framed local systems [FG06b] (see also [Gon17, STWZ19]), and should relate to the spectral networks of Gaiotto-Moore-Neitzke [GMN10, GMN13, GMN14]. For that, consider $N \in \mathbb{N}$ and τ an ideal N -triangulation of the smooth punctured surface C . In Section 3, we present a new construction that associates an N -graph $G(\tau)$ to an ideal N -triangulation (C, τ) . In particular, each ideal N -triangulation τ yields a Legendrian surface $\Lambda(G(\tau)) \subseteq (J^1C, \xi_{\text{st}})$. In general, different N -triangulations lead to smoothly isotopic Legendrian surfaces which are *not* Legendrian isotopic, and they are distinguished by their flag moduli space $\mathcal{M}(G(\tau))$. This also relies on the connection between microlocal monodromies and cluster algebras.




FIGURE 7. The Legendrian weave associated to a 4-triangle (left) and to a 5-triangle (right). The open Legendrian surface for the 4-triangle has genus one and two boundary components. The Legendrian surface for the 5-triangle has genus two and three boundary components. \square

The N -graph $G(\tau)$ and the Legendrian weave $\Lambda(G(\tau))$ are both constructed with a local model on an N -triangle. Figure 7 depicts a Legendrian weave associated to the 4- and 5-graphs dual to 4- and 5-triangles. We will prove that their local flag moduli space is a complex torus by using Theorem 1.1 and the flag moduli space results from Section 5. The precise statement, proven in Section 8, reads as follows:

⁷In contrast, the construction for torus links given by the first author in [CG20] uses Lagrangian concordances of infinite order. In that context, see also the upcoming work [GSW20a] which will show that the square of the Donaldson-Thomas transformation [GS18] is a Lagrangian concordance, oftentimes of infinite order.

⁸This is generically the case.

Theorem 1.9. *Let $G(t_N)$ be the N -graph associated to an N -triangle t_N , and let k a field. The flag moduli space of $G(t_N)$ is a $\binom{N-1}{2}$ -dimensional complex torus, i.e.*

$$\mathcal{M}(t_N, G(t_N); k) \cong (k^*)^{\binom{N-1}{2}}.$$

The combinatorial number $\binom{N-1}{2}$ appears geometrically as the rank of the first homology class of the Legendrian weave $\Lambda(G(t_N))$. Now, the class of Legendrian weaves $\Lambda(G(\tau))$ arising from ideal N -triangulations τ of punctured surfaces is of central interest in the study of moduli spaces of framed local systems for the Lie group $\mathrm{GL}(N, \mathbb{C})$ [FG06b]. Indeed, the Legendrian surface $\Lambda(G(\tau))$ is a compactification of the Legendrian lift of the Goncharov-Kenyon Lagrangian conjugate surface $L_\tau \subseteq (T^*C, \lambda_{\mathrm{st}})$, see [Gon17, STWZ19]. Thus, the non-Abelianization technique, expressing higher-rank local systems in S in terms of rank-one local systems on L_τ , can also be recovered by studying these Legendrian weaves $\Lambda(\tau)$ – see Section 8.2 for an explicit computation. In particular, the set of Legendrian surfaces $\{\Lambda(\tau)\}_\tau$ provides a symplectic geometric realization of the set of cluster charts in this moduli spaces of framed local systems. This parallels the work of [STWZ19] on conjugate surfaces. See Section 8 for details.

Basic Notation and Color Code. The germs of singularities of caustics and wavefronts are referred to according to the classical notation from the theory of singularities, following V.I. Arnol’d [Ad90]. Given a subset $X \subseteq Y$ of a smooth manifold Y , we denote by $\mathcal{O}p(X)$ an arbitrarily small but fixed open neighborhood of it, following M. Gromov [Gro86].

Regarding colors, the two colors **blue** and **red** are associated to edges with adjacent transpositions, i.e. edges with consecutive transpositions $(i-1, i)$, $(i, i+1)$, for a choice $2 \leq i \leq N-1$. The same holds for colors **red** and **yellow** used together. The *three* colors **blue**, **red** and **yellow** together denote edges labeled by three consecutive transpositions $(i-1, i)$, $(i, i+1)$ and $(i+1, i+2)$, respectively, for a choice $2 \leq i \leq N-2$. In a diagram with the two colors **blue** and **yellow**, without **red**, these two colors denote any edges with disjoint transpositions. The color **orange** will exclusively be used to denote cusp edges, corresponding to edges of A_2 -singularities. Finally, we use **purple** dots (or black dots) for D_4^- singularities, **yellow** dots for A_1^3 singularities and **orange** dots for A_3 -swallowtail singularities. \square

Acknowledgements. We thank Honghao Gao and Kevin Sackel for their thorough reading of the initial version of this manuscript, and Honghao Gao, Eugene Gorsky and Harold Williams for many valuable comments. We also thank the referees for their suggestions and comments. We are grateful to Dylan Thurston for providing key examples of quiver mutations, and to Ben Elias for discussions on Soergel calculus. We also thank J. Etnyre, O. Lazarev, I. Le, L. Ng, J. Sabloff, L. Traynor and D. Treumann for discussions, questions and interest in this work. R. Casals is supported by the NSF grant DMS-1841913, a BBVA Research Fellowship and the Alfred P. Sloan Foundation. E. Zaslow is supported by the NSF grant DMS-1708503. \square

2. N -GRAPHS AND LEGENDRIAN WEAVES

In this section we introduce the notion of an N -graph G and construct the Legendrian surface $\Lambda(G)$ associated to it. The interaction between the combinatorics of G and the contact geometric invariants of $\Lambda(G)$ is the starting focus of this article. The reader is referred to [BM08, Die17] for introductory material on graph theory and to [Etn05, Gei08] for the basics of contact topology.

2.1. **N -graphs.** Let C be a smooth surface and $N \in \mathbb{N}$ a natural number. An embedded graph $G \subseteq C$ is said to be trivalent if all its vertices have degree three. Such a vertex is depicted on the left in Figure 8.




FIGURE 8. Trivalent vertex (left) and Hexagonal Point (right).

Definition 2.1. Let J and K be two trivalent graphs embedded in C , having an isolated intersection point at a common vertex $v \in J \cap K$. The intersection v is said to be *hexagonal* if the six half-edges in C incident to v interlace, i.e. alternately belong to J and K . \square

The right diagram in Figure 8 depicts a hexagonal vertex, where the graph J is labeled $(i-1, i)$ in blue and K is labeled $(i, i+1)$ in red. These hexagonal intersection points will be referred to as *hexagonal $(i, i+1)$ -points*.

Definition 2.2. An N -graph G on a smooth surface C is a set $G = \{G_i\}_{1 \leq i \leq N-1}$ of $N-1$ embedded trivalent graphs $G_i \subseteq C$, possibly empty or disconnected, such that G_i is allowed to intersect G_{i+1} only at hexagonal points, $1 \leq i \leq N-2$. \square

Two examples of N -graphs on the plane $C = \mathbb{R}^2$ are depicted in Figure 1. The (trivalent) vertices are depicted by *purple* or *black* dots and the hexagonal intersection points by *yellow* dots. Note that $G_i, G_j \subseteq C$ are allowed to intersect (anywhere) if $j \neq i, i \pm 1$, and they may intersect non-transversely.

Remark 2.3. We can think of an N -graph as an immersed graph with colored edges, the color i corresponding to the graph G_i , $1 \leq i \leq N-1$. Edges labeled by numbers differing by two or more may pass through one another (hence the immersed property, which is met generically), but not at a vertex. In particular, a 3-graph is a bicolored graph with monochromatic trivalent vertices and interlacing hexagonal vertices. \square

Consider $\tau(N) := \{(i, i+1) \in S_N : 1 \leq i \leq N-1\} \subseteq S_N$ the subset of simple transpositions and denote $\tau_i := (i, i+1)$. We label the edges of an N -graph $G = \{G_i\}$ which belong to the graph G_i with the transposition τ_i , as we have done in Figure 1. These edges will also be referred to as τ_i -edges, or i -edges. By definition, the trivalent vertices belonging to the graph G_i have three incident τ_i -edges. The hexagonal points in $G_i \cap G_{i+1}$ have six edges incident to it, alternately labeled with the transpositions τ_i and τ_{i+1} in $\tau(N)$. Figure 8 depicts the local model for the trivalent vertices of the cubic graph G_{i-1} and a hexagonal intersection point in $G_i \cap G_{i-1}$. Observe that a 2-graph is, by definition, an embedded trivalent graph.

The study of N -graphs brings the combinatorial ingredients of the article, and we provide in Section 3 several combinatorial constructions of N -graphs. For now, we introduce its geometric counterpart, the Legendrian surface associated to an N -graph.

2.2. **Singularities of wavefronts.** The Legendrian surface $\Lambda(G)$ associated to an N -graph $G \subseteq C$ is an embedded Legendrian in the 1-jet space (J^1C, ξ_{st}) . The Legendrian surface $\Lambda(G)$

is described by using germs of Legendrian wavefronts [Ad90, Section 3.1] in the Darboux chart $(\mathbb{R}^5, \xi_{\text{st}})$, where the contact 4-distribution ξ_{st} is defined as

$$\xi_{\text{st}} = \ker \alpha_{\text{st}}, \text{ where } \alpha_{\text{st}} := dz - y_1 dx_1 - y_2 dx_2,$$

and $(x_1, x_2, y_1, y_2, z) \in \mathbb{R}^5$ are Cartesian coordinates in \mathbb{R}^5 . This is the local model for any contact 4-distribution in the neighborhood of a point [Gei08, Theorem 2.5.1]. Since $\lambda_{\text{st}} = y_1 dx_1 + y_2 dx_2$ is the Liouville form of the cotangent bundle $(T^*\mathbb{R}^2, \omega_{\text{st}})$, this Darboux chart $(\mathbb{R}^5, \xi_{\text{st}})$ is contactomorphic to the 1-jet space $(J^1\mathbb{R}^2, \ker\{dz - \lambda_{\text{st}}\})$.

The Legendrian fibration $\pi : \mathbb{R}^5 \rightarrow \mathbb{R}^3$, $\pi(x_1, x_2, y_1, y_2, z) = (x_1, x_2, z)$ allows us to assign a smoothly embedded Legendrian surface $\Lambda(\Sigma) \subseteq \mathbb{R}^5$ in the domain of π to certain singular surfaces $\Sigma \subseteq \mathbb{R}^3$ in its target. The coordinates (y_1, y_2) of the Legendrian $\Lambda(\Sigma)$ assigned to Σ are

$$y_1 = x_1\text{-slope of the tangent plane } T_{(x_1, x_2, z)}\Sigma,$$

$$y_2 = x_2\text{-slope of the tangent plane } T_{(x_1, x_2, z)}\Sigma.$$

In a local parametrization $\sigma : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ of Σ , $\sigma(u, v) = (u, v, z(u, v))$, this reads

$$y_1 = \partial_u z(u, v), \quad y_2 = \partial_v z(u, v).$$

This assignment is dictated by the vanishing of the contact 1-form α_{st} along $\Lambda = \Lambda(\Sigma)$. The three-dimensional case is explained in detail in [Gei08, Section 3.2], the general case is discussed in [AdG01, Chapter 5], [EES05a, Section 3.2] and [CM19, Section 2]. The germs of singularities of Σ that lift to an embedded Legendrian Λ , and equivalently, the singularities of the map $\pi|_{\Lambda}$, are restricted. These are known as singularities of fronts, or equivalently, Legendrian singularities [AdG01]. By definition, singular surfaces Σ obtained as the image of an embedded Legendrian submanifold via a Legendrian mapping are referred to as (wave)fronts.

Remark 2.4. The classification of generic singularities of spatial fronts $\Sigma \subseteq \mathbb{R}^3$ is stated in [AdG01, Theorem 3.1.1], and that of generic singularities of a 1-parametric family of spatial fronts $\Sigma \subseteq \mathbb{R}^3$ is explained in [AdG01, Theorem 3.4.2]. \square

The main spatial wavefronts Σ that we use in the course of this article use three different germs of singularities of Legendrian fronts: A_1^2, A_1^3 and D_4^- , which we now describe. We emphasize that these are singularities of the wavefront projections only: the corresponding local Legendrian surfaces are all smooth.

2.2.1. *The A_1^2 germ.* This germ is obtained as a product of a 2-dimensional planar front times an interval. It is described by the germ of the singular surface

$$\Sigma(A_1^2) = \{(x_1, x_2, z) \in \mathbb{R}^3 : (x_1^2 - z^2) = 0\}$$

at the origin. This wavefront is informally called an A_1^2 -crossing, or a *crossing*, and the set of points $\{(x_1, x_2, z) \in \Sigma(A_1^2) : x_1 = 0, z = 0\}$ is referred to as an edge, or segment, of A_1^2 -crossings. This spatial front is depicted on the left in Figure 9. Its Legendrian lift $\Lambda(\Sigma(A_1^2)) \subseteq (\mathbb{R}^5, \xi_{\text{st}})$ consists of two disjoint embedded Legendrian 2-disks.

2.2.2. *The A_1^3 germ.* The wavefront A_1^3 is given by the germ at the origin of the singular surface

$$\Sigma(A_1^3) = \{(x_1, x_2, z) \in \mathbb{R}^3 : (x_1 - z)(x_1 + z)(z - x_2) = 0\} \subseteq \mathbb{R}^3,$$

This spatial front is depicted in the center of Figure 9. Considered as a germ, the origin is the A_1^3 -wavefront singularity, and the codimension-1 singular strata consists of six half-lines of A_1^2 singularities. The Legendrian lift $\Lambda(\Sigma(A_1^3))$ of the A_1^3 germ to $(\mathbb{R}^5, \xi_{\text{st}})$ consists of three disjoint embedded Legendrian 2-disks.

2.2.3. *The D_4^- germ.* The third germ $\Sigma(D_4^-) = \text{Im}(\delta_4^-) \subseteq \mathbb{R}^3$ of a Legendrian singularity that we use is given by the germ at the origin for the image of the map

$$\delta_4^- : \mathbb{R}^2 \longrightarrow \mathbb{R}^3, \quad \delta_4^-(x, y) = \left(x^2 - y^2, 2xy, \frac{2}{3}(x^3 - 3xy^2) \right).$$

The D_4^- -singularity of the spatial wavefront $\text{Im}(\delta_4^-)$ is at $(0, 0, 0) \in \mathbb{R}^3$. The front $\text{Im}(\delta_4^-)$ itself also has three half-lines of A_1^2 -crossings, intersecting at the origin. This is depicted in the right of Figure 9. The Legendrian lift $\Lambda(\text{Im}(\delta_4^-)) \subseteq (\mathbb{R}^5, \xi_{\text{st}})$ of the D_4^- spatial front is an embedded Legendrian 2-disk. We refer the reader to [Ad90, TZ18] for more descriptions — see also Remark 2.5 below.




FIGURE 9. The A_1^2 spatial front (left), the germ of the A_1^3 Legendrian singularity (center) and the D_4^- Legendrian wavefront (right).

The connection of the above three Legendrian singularities with the Weyl groups, justifying their nomenclature, can be found in [AdG01, Section 3.3]. It might be relevant to notice that D_4^- is not the germ of a singularity for a *generic* Legendrian wavefront, but still a valid singularity for a given spatial wavefront. In addition, it is known that the singularity D_4^- is generic in 1-parameter families of Legendrian fronts [Ad90, Section 3.3]. As a result, most of the Legendrians we construct are non-generic, in their isotopy class, with respect to the fixed Legendrian projection. This rigidification simplifies the analysis and combinatorics.

Remark 2.5. The D_4^- Legendrian singularity has the property that its singular strata, excluding A_1^2 singularities, is a point, which lies in real codimension 2. This is not the case for the majority of Legendrian surface singularities, such as the Legendrian A_3 -swallowtail, cusp-edges A_2A_1 and the purse wavefront D_4^+ , the former two even being generic. (These singularities feature in Section 4.) The geometric reason for this codimension-2 phenomenon is the existence of the *holomorphic* Legendrian surface singularity

$$t : \mathbb{C} \longrightarrow (J^1(\mathbb{C}, \mathbb{C}), \ker\{dw_1 - w_2dw_3\}), \quad (w_1, w_2, w_3) = t(w) = \left(w^2, w, \frac{2}{3}w^3 \right),$$

whose real part is the real Legendrian singularity D_4^- . This holomorphic map is the complexification of the real simple cusp singularities appearing in generic front projections of embedded Legendrian knots in a Darboux chart $(\mathbb{R}^3, \xi_{\text{st}})$. \square

We also use the A_2, A_2A_1 front singularities, geometrically represented by a simple cusp in \mathbb{R}^2 times an interval, and its intersection with a 2-plane. These A_2 -singularities do not directly arise from an N -graph $G \subseteq C$, but rather from satelliting the smooth surface C to a Legendrian surface in a contact 5-manifold (Y, ξ) , typically $(\mathbb{S}^5, \xi_{\text{st}})$.

2.3. Legendrian Weaves. Let $G \subseteq C$ be an N -graph, as introduced in Subsection 2.1 above. The principle that associates a Legendrian $\Lambda(G)$ to the N -graph G is that G dictates the configuration of A_1^2 singularities (crossings) of its Legendrian wavefront. This is possible

because the singularities introduced in Subsection 2.2 are uniquely determined by their A_1^2 front singularities. Let us explain the construction in detail.

First, we choose the ambient contact manifold, where the embedded Legendrian surface $\Lambda(G)$ belongs, to be the 1-jet space of the smooth surface C . That is,

$$\Lambda(G) \subseteq (J^1C, \xi_{\text{st}}) = (\{(x, z) \in T^*C \times \mathbb{R}\}, \ker\{dz - \lambda_{\text{st}}\}),$$

where $\lambda_{\text{st}} \in \Omega^1(T^*C)$ is the Liouville form [Gei08, Section 1.4], and see [Ad90, Example 2] and [Gei08, Example 2.5.11] for details on the 1-jet space. The local germs described in Subsection 2.2 above and the Legendrian front projection $\pi : (J^1C, \xi_{\text{st}}) \rightarrow C \times \mathbb{R}$ allow us to assign a Legendrian $\Lambda(\Sigma) \subseteq (J^1C, \xi_{\text{st}})$ to a spatial wavefront $\Sigma \subseteq C \times \mathbb{R}$ in the target, as follows.

The construction of the front $\Lambda(G) \subseteq (J^1C, \xi_{\text{st}})$ is obtained by gluing local wavefront models in $U_i \times \mathbb{R}$, $i \in I$, $U_i \cong \mathbb{D}^2$, which are the targets of front projections in the Darboux charts $(J^1U_i, \xi_{\text{st}}) \cong (J^1\mathbb{D}^2, \xi_{\text{st}})$, for $i \in I$. This is formalized in the following definition:

Definition 2.6. Let $\mathbb{D}_N = \mathbb{D}^2 \times \{1\} \cup \dots \cup \mathbb{D}^2 \times \{N\} \subseteq \mathbb{D}^2 \times \mathbb{R}$. We consider \mathbb{D}_N as a disconnected, horizontal wavefront. Let $P \subseteq \mathbb{D}^2 \times \{0\}$ be one of the following four local models of an N -graph $G \subseteq \mathbb{D}^2$:

1. A unique i -edge in \mathbb{D}^2 , as drawn at the bottom of the second column in Figure 10.
2. A unique trivalent i -vertex, as shown at the bottom of the third column in Figure 10.
3. A unique hexagonal $(i, i + 1)$ -point, depicted in the fourth column in Figure 10.
4. The empty set.

Here, recall that an i -edge is an edge belonging to the graph $G_i \subseteq G$ of the N -graph $G \subseteq \mathbb{D}^2$, for $1 \leq i \leq (N - 1)$. By definition, the Legendrian wavefront $\mathbb{D}_N(P) \subseteq \mathbb{D}^2 \times \mathbb{R}$ associated to P is obtained as follows:

- If P is a i -edge, insert an A_1^2 -intersection along the two sheets $\mathbb{D}^2 \times \{i\}$ and $\mathbb{D}^2 \times \{i + 1\}$ of the wavefront \mathbb{D}_N . This A_1^2 intersection must be inserted such that the image of the A_1^2 singular locus coincides with P under the projection $\mathbb{D}^2 \times \mathbb{R} \rightarrow \mathbb{D}^2$ onto the first factor.
- If P is a trivalent i -vertex, introduce a D_4^- -singularity between the two sheets $\mathbb{D}^2 \times \{i\}$ and $\mathbb{D}^2 \times \{i + 1\}$ in the wavefront \mathbb{D}_N^2 . This D_4^- singularity must be introduced such that, under the projection $\mathbb{D}^2 \times \mathbb{R} \rightarrow \mathbb{D}^2$ onto the first factor, the image of the A_1^2 -crossings coincides with the three edges of P and the D_4^- singular point is mapped to the unique trivalent vertex of P .
- If P is a hexagonal $(i, i + 1)$ -point, insert an A_1^3 -intersection along the three disjoint sheets $\mathbb{D}^2 \times \{i\}$, $\mathbb{D}^2 \times \{i + 1\}$ and $\mathbb{D}^2 \times \{i + 2\}$ of the wavefront \mathbb{D}_N^2 . The pattern for the A_1^3 -wavefront must be inserted such that, under the projection $\mathbb{D}^2 \times \mathbb{R} \rightarrow \mathbb{D}^2$ onto the first factor, the origin in the A_1^3 -singularity maps to the unique vertex of P , and the six half-lines of A_1^2 -crossings map to the six edges emanating from the vertex.

These wavefronts are depicted in Figure 10. For P empty we use the front $\mathbb{D}^2 \times \{1\} \cup \dots \cup \mathbb{D}^2 \times \{N\} \subseteq \mathbb{D}^2 \times \mathbb{R}$. We refer to the wavefronts $\mathbb{D}_N(P)$ as being obtained from the wavefront \mathbb{D}_N by weaving according to the pattern P . \square

Definition 2.6 describes how to weave the wavefront $\mathbb{D}_N \subseteq \mathbb{D}^2 \times \mathbb{R}$, which we have fixed, according to a pattern $P \subseteq \mathbb{D}^2 \times \{0\}$. To glue models, let $\{U_i\}_{i \in I}$ be a finite cover of C by open 2-disks $U_i \cong \mathbb{D}^2$, refined as necessary so that each U_i contains no more than one non-empty feature P of the N -graph G . Now, let us consider two 2-disks $U_1, U_2 \subseteq C$ and two corresponding patterns P_1, P_2 therein.




FIGURE 10. The leftmost wavefront is \mathbb{D}_N , then from left to right we find $\mathbb{D}_N(P)$ where P is an edge, a trivalent vertex and a hexagonal vertex.

Suppose that the patterns P_1 and P_2 coincide along the intersection $U_1 \cap U_2$. Then we say that $P_1 \cup P_2$ defines a pattern in $U_1 \cup U_2$. By definition, the wavefront $\Sigma(P_1 \cup P_2)$ associated to $P_1 \cup P_2$ is obtained by considering the set-theoretical union of $\mathbb{D}_N(P_1)$ and $\mathbb{D}_N(P_2)$ in $(U_1 \cup U_2) \times \mathbb{R}$. For brevity of notation, we will say that $\Sigma(P_1 \cup P_2)$ is obtained by weaving $\mathbb{D}_N^2 \cup \mathbb{D}_N^2 \subseteq (U_1 \times \mathbb{R}) \cup (U_2 \times \mathbb{R})$ according to the pattern $P_1 \cup P_2$. Finally, the Legendrian surface associated to an N -graph is defined as follows:

Definition 2.7. Let C be a smooth surface and $G \subseteq C$ an N -graph, the Legendrian weave

$$\Lambda(G) \subseteq (J^1C, \xi_{\text{st}})$$

is the embedded Legendrian surface whose wavefront $\Sigma(G) \subseteq C \times \mathbb{R}$ is obtained by weaving the wavefront $C \times \{1\} \cup \dots \cup C \times \{N\} \subseteq C \times \mathbb{R}$ according to the pattern $G \subseteq C$. \square

Let $\{\varphi_t\}_{t \in [0,1]} \subseteq \text{Diff}^c(C)$, $\varphi_0 = \text{Id}$, be a compactly supported isotopy of the smooth surface C . Then the Legendrian surfaces $\Lambda(\varphi_t(G)) \subseteq (J^1C, \xi_{\text{st}})$, as described in Definition 2.7, are Legendrian isotopic, relative to their boundaries. Hence, for the purposes of this article, our N -graphs $G \subseteq C$ are considered up to such planar isotopies. Similarly, Legendrian fronts in \mathbb{R}^3 are to be considered up to homotopy of fronts.

Thanks to Definition 2.7, the wealth of contact topology invariants [EGH00, EES05a, GKS12, STZ17, CM19] can be used to define algebraic structures associated to N -graphs $G \subseteq C$. For instance, the articles [CM18, TZ18] show that the chromatic polynomial of (the dual of) a trivalent graph G – which is a 2-graph – is contained in the Floer-theoretical invariants of the Legendrian weave $\Lambda(G)$. Conversely, from a contact topology perspective, the connection to combinatorics and algebraic geometry provides a new tool for computing contact invariants of higher-dimensional Legendrian submanifolds. This will be the focus of subsequent sections.

Remark 2.8. The one-dimensional analogue of a Legendrian weave is a Legendrian braid, i.e. a positive braid. Indeed, an N -graph in a one-manifold I is defined to be a set of points, each point labeled with a permutation in $\tau(N) \subseteq S_N$. The only planar front singularity that we can use is A_1^2 , corresponding to a crossing, necessarily positive. Thus, 1-dimensional weaving consists of introducing positive crossings to the N strands

$$\mathbb{D}_N^1 = I \times \{1\} \cup \dots \cup I \times \{N\} \subseteq I \times \mathbb{R}$$

and concatenating them side by side. This is precisely the front for an N -strand positive braid [PS97], which lifts to a Legendrian link in $(J^1S^1, \xi_{\text{st}})$ [Gei08, Section 3.3.1]. The Legendrian weaves introduced in Definition 2.7 are thus the Legendrian surface generalization of Legendrian braids. \square

2.4. Smooth Topology of Weaves. Let G be an N -graph in a surface C , in this subsection we address the smooth topology of the Legendrian surface $\Lambda(G) \subseteq (J^1C, \xi_{st})$.⁹ The smooth invariants of $\Lambda(G)$ are the first homology $H_1(\Lambda(G), \mathbb{Z})$, in particular its genus $g(\Lambda(G)) \in \mathbb{N}$, and the number of boundary components $|\partial\Lambda(G)|$. For simplicity, we assume that C is a closed surface, and thus $\partial\Lambda(G) = \emptyset$. We also assume that G is a connected N -graph, i.e. the union of the graphs $G_i, i \in I$, is a connected topological subspace of C .

The surface $\Lambda(G)$ is a branched N -fold cover over C simply branched over the trivalent vertices of G . Indeed, the image of $\Lambda(G)$ by the projection $J^1C \rightarrow T^*C$ along the Reeb \mathbb{R} -direction yields an immersed surface $L(G) \subseteq T^*C$, and the canonical projection $T^*C \rightarrow C$ restricts to $L(G)$ as an N -fold branched cover. The branch set is the image of the set of D_4^- singularities. As a result, the genus of $\Lambda(G)$ is provided by the Riemann-Hurwitz formula

$$\chi(\Lambda(G)) = N\chi(C) - v(G), \text{ i.e. } g(\Lambda(G)) = \frac{1}{2}(v(G) + 2 - N\chi(C))$$

where $v(G)$ is the number of (trivalent) vertices of G .

Remark 2.9. If the surface C has boundary, each boundary component of ∂C contributes to a piece of the boundary $\partial\Lambda(G)$ of the Legendrian surface $\Lambda(G) \subseteq (J^1C, \xi_{st})$. Let $\kappa \in \mathbb{N}$ be the number of cycles in the (minimal length) factorization of the monodromy of the branched cover along a given boundary component of ∂C . Then, that one boundary component of C contributes to κ distinct boundary components for the Legendrian surface $\Lambda(G)$. \square

Example 2.10. The Legendrian weaves $\Lambda(G_1), \Lambda(G_2) \subseteq (J^1(\mathbb{S}^2), \xi_{st})$ associated to the 3- and 4-graphs in Figure 1 are closed Legendrian surfaces of genus 3 and 4, respectively. Should the graphs $G_1, G_2 \subseteq \mathbb{R}^2$ be considered in the 2-plane \mathbb{R}^2 , instead of the 2-sphere \mathbb{S}^2 , the Legendrian surfaces $\Lambda(G_1), \Lambda(G_2) \subseteq (J^1(\mathbb{R}^2), \xi_{st})$ have genus 3 and 4, with 3 and 4 boundary components, respectively. \square

Now, the \mathbb{Z}_2 -monodromy of $\Lambda(G)$ along a non-trivial 1-cycle of the base C is trivial, and thus the contributions of the graph G to $H_1(\Lambda(G), \mathbb{Z})$, as expressed by the above formula, can be considered by studying planar pieces. Let us then assume that $g(C) = 0$ and construct 1-cycles in $H_1(\Lambda(G), \mathbb{Z})$ in terms of the edges of the N -graph.




FIGURE 11. Two combinatorial descriptions of 1-cycles in $H_1(\Lambda(G), \mathbb{Z})$.

There are two direct descriptions of 1-cycles $\gamma \in H_1(\Lambda(G), \mathbb{Z})$:

1. Each edge e of the graph G connecting two trivalent vertices defines a 1-cycle $\gamma(e) \in H_1(\Lambda(G), \mathbb{Z})$. The projection of this 1-cycle onto the pattern P with two trivalent vertices is depicted in orange on the left of Figure 11. In order to construct $\gamma(e)$ from the orange curve, lift a point in the orange curve to the annulus $\Lambda(P)$, to either

⁹This is necessary for our applications, especially in the study of microlocal monodromies and Lagrangian fillings in Section 7 and the non-Abelianization map in Section 8.

one of the two sheets, and uniquely follow the lift along the orange curve. Since the lift is isotopic to one of the boundary components of the annulus, it generates $H_1(\Lambda(P), \mathbb{Z}) \cong \mathbb{Z}$. The 1-cycle $\gamma(e)$ is drawn directly in the wavefront projection in Figures 12. We refer to this type of 1-cycles as *monochromatic edges* or (short) 1-cycles.




FIGURE 12. The first type of 1-cycle $\gamma(e)$ drawn in the wavefront (left) and in a vertical slicing (right). Each slice on the left is labeled by a number. The 1-cycle $\gamma(e)$ appears as five-pointed stars in each slice as shown on the right.

There is a simple extension of this construction, depicted in Figure 13. Consider a trivalent vertex $v \in G_i$ and a linear chain of edges e_1, e_2, \dots, e_k in G such that e_1 connects v to a hexagonal vertex, e_i connect two hexagonal vertices for $2 \leq i \leq k-1$ and e_k connects the free hexagonal vertex in e_{k-1} to a trivalent vertex. Suppose further that e_j and e_{j+1} meet at opposite rays of the hexagonal vertex between them, $1 \leq j \leq k-1$. Then the orange curves in the patterns all lift to 1-cycles which are essential¹⁰ in the surfaces $\Lambda(P)$ for the corresponding patterns P . These 1-cycles are referred to as *long edges* or *long 1-cycles*.




FIGURE 13. Descriptions of 1-cycles in $H_1(\Lambda(G), \mathbb{Z})$ of the first type, generalizing $\gamma(e)$ on the left of Figure 11. The lift of the orange curves generate the first homology $H_1(\Lambda(P), \mathbb{Z}) \cong \mathbb{Z}$ for the corresponding patterns P .

2. Three edges e_1, e_2, e_3 of a graph G_i connecting a hexagonal vertex with three trivalent vertices in G_i defines a cycle $\gamma(e_1, e_2, e_3) \in H_1(\Lambda(G), \mathbb{Z})$. This is depicted on the right in Figure 11. The 1-cycle $\gamma(e_1, e_2, e_3)$ is drawn in the wavefront projection in Figure 14. We refer to this type of 1-cycles $\gamma(e_1, e_2, e_3)$ as a *Y-cycle*.

We can also combine the above two constructions to associate a 1-cycle to any tree with leaves on trivalent vertices that passes directly through any hexagonal vertices, i.e. entering

¹⁰The topology of $\Lambda(P)$ is that of an annulus union disjoint 2-disks.




FIGURE 14. The second type of 1-cycle $\gamma(e_1, e_2, e_3)$ drawn in a slicing of the wavefront associated to the pattern on the left.

and exiting along opposing edges, see Figure 100 for an example. For such a tree, we refer to the pieces corresponding to edges as l-pieces, or edges, and the pieces that go through a hexagonal vertex as Y-pieces. In addition, we can decorate such 1-cycles with a number, indicating higher multiplicity¹¹. If we require the curves in the Legendrian surface to be connected, then higher multiplicity in general requires these curves to be immersed.

Remark 2.11. Let $\Lambda(G) \subseteq (J^1C, \xi_{st})$ be a connected surface, and $G \subseteq C$ a connected N -graph. The trivalent vertices of the N -graph $G \subseteq C$ can be assumed to belong to G_1 . This follows once we impose certain equivalence relations on the set of N -graphs, which is done in Section 4. \square

2.5. Combinatorial Homology. Let $G \subseteq C$ be an N -graph. We present a combinatorial model for the (chain-level) simplicial homology of $\Lambda(G)$. This can be achieved in general, but for this subsection we assume that G is a planar 3-graph, i.e. $C = \mathbb{S}^2$ and $N = 3$. We will think of G as bicolored — see Remark 2.3. This will ease notation, while containing the essential idea for higher $N \in \mathbb{N}$ and higher-genus C . Note that the results in this subsection will not be used in the rest of the manuscript, we have included them for completeness.

The edges, faces and vertices of G lift to edges, faces and vertices of the Legendrian surface $\Lambda := \Lambda(G)$. Let us suppose that G and Λ are connected, and that the faces of G define a polyhedral decomposition (F, E, V) of the sphere. This decomposition lifts to a polyhedral decomposition of Λ , as follows. Each face, edge and hexagonal vertex of G has three lifts to Λ ; each trivalent vertex has two lifts. This yields

$$\chi(\Lambda) = 3 \cdot 2 - v,$$

where $v = |V(G)|$ is the number of trivalent vertices. For a point $P \in \mathbb{S}^2$, we write P_1, P_2, P_3 for the (up to) three pre-images in non-decreasing order of the z -coordinate. If P is on G , we must choose a nearby point to define the ordering of z coordinates of sheets. If P is a trivalent vertex with label (12), in blue, then $P_1 = P_2$ while $P_2 = P_3$ for a label (23), in red. Lifts of edges and faces are labeled analogously. The chain complex C_\bullet associated to this polyhedral decomposition of Λ computes the homology $H_*(\Lambda; \mathbb{Z})$. There is a simplified chain complex that computes $H_1(\Lambda; \mathbb{Z})$, which we now explain.

Lift each edge $e = (P, Q)$ labeled $(i, i + 1)$ to a one-chain as follows (here $i = 1$ or 2). In the (any) orientation of the plane, if A is the sheet with lower z value in the region to the left of \overline{PQ} and B is the sheet with lower z value to the right of \overline{PQ} then lift e to the chain $\overline{P_B Q_B} - \overline{P_A Q_A}$; this only depends on e and not the ordering of P and Q . Write \hat{e} for this lift of e . Extending by linearity, we get a map $\mathbb{Z}^E \rightarrow C_1$.

The embedded bicolored graph G is the union $G = G_B \cup G_R$ of embedded blue and a red graphs intersecting at hexagonal vertices, where $G_B = (F_B, E_B, V_B)$, F_B denotes the number of faces of the graph G_B , and E_B the number of edges and V_B the number of vertices. (Similarly for $G_R = (F_R, E_R, V_R)$.) We define a complex A_\bullet as follows. $A_2 := \mathbb{Z}^{F_B} \oplus \mathbb{Z}^{F_R}$,

¹¹Higher multiplicities will rarely feature in this manuscript, only in relation to Theorem 7.14.

$A_1 := \mathbb{Z}^E = \mathbb{Z}^{E_B} \oplus \mathbb{Z}^{E_R}$, and A_0 is the image $\partial \hat{A}_1$, where \hat{A}_1 is the image of A_1 in C_1 under $e \mapsto \hat{e}$. A monochromatic face $f \in \mathbb{Z}^{F_B} \subset A_2$ has a lift to C_2 as $f_1 - f_2$, whereas $f \in \mathbb{Z}^{F_R} \subset A_2$ lifts to $f_2 - f_3$. Summarizing, we have

$$\begin{array}{ccccc} C_2 & \longrightarrow & C_1 & \longrightarrow & C_0, \\ \uparrow & & \uparrow & & \uparrow \\ A_2 & & A_1 & \longrightarrow & A_0 \end{array}$$

where the map $A_1 \rightarrow A_0$ sends e to $\partial \hat{e}$. The missing differential $A_2 \rightarrow A_1$ is defined as follows. For a monochromatic face $f \in \mathbb{Z}^{F_B}$ or $\mathbb{Z}^{F_R} \subset A_2$,

$$\partial f = \sum_{\text{boundary edges}} e - \sum_{\text{interior edges}} e,$$

which we extend by linearity.

Proposition 2.12. A_\bullet is a chain complex and $A_\bullet \rightarrow C_\bullet$ is a chain map.

Proof. Let $f \in A_2$ be a blue face. A similar argument will work for red faces. We need to check that A_\bullet is a chain complex, i.e. the differential squares to zero. This reads

$$\sum_{\text{boundary edges}} \partial \hat{e} - \sum_{\text{interior edges}} \partial \hat{e} = 0,$$

for a face f . The left hand side of this equality is a (formal) sum of some of the vertices of the graph, some trivalent, some hexagonal. Thus, this imposes a condition at all the interior and exterior vertices of f . In fact, the condition is null at an interior vertex, since it must be monochromatic and hence trivalent, and $\partial \hat{e}$ is zero over any trivalent vertex. Likewise for an exterior trivalent vertex, there is nothing to check and it remains to discuss exterior hexagonal vertices. For an exterior hexagonal vertex, a local study is needed, as follows.

Let h be a hexagonal vertex and let us study the differential restricted to it. Let e_1, e_2, e_3 be three attached blue half-edges, with $e_4 = e'_1, e_5 = e'_2, e_6 = e'_3$ the opposite red half-edges, respectively. Let h_1, h_2, h_3 be the three preimages of h . We can restrict the differential $A_1|_h \rightarrow A_0|_h$ to edges intersecting h and points over h , and in the chosen basis it takes the form

$$(2.1) \quad \partial|_h = \begin{pmatrix} -1 & 0 & 1 \\ 1 & -1 & 0 \\ 0 & 1 & -1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \\ 0 & -1 & 1 \end{pmatrix}.$$

The kernel is generated by $e_1 + e'_1, e_2 + e'_2, e_3 + e'_3, e_1 + e_2 + e_3$. Note that the last generator could also have been taken to be $e_1 + e_2 - e'_3$. The first three represent long two-colored edges passing straight through the hexagonal vertex, while the last is a monochromatic Y shape. In more detail, the element $e_1 + e'_1$ of the kernel is, diagrammatically, given by a (bi-colored) edge passing through the hexagonal vertex. Similarly for $e_2 + e'_2$ and $e_3 + e'_3$, they represent straight edges passing through the hexagonal vertex, starting blue and then turning red, or viceversa. The element $e_1 + e_2 + e_3$ of the kernel is given by the Y-shaped union of the three blue edges coming out of a hexagonal vertex. The element $e_1 + e_2 - e'_3$ also belongs to the kernel, and it represents two blue edges *and* the red edge in between (with a minus sign). The terms arising in $\partial^2 f$ which involve a boundary hexagonal vertex are given by the image of such configurations in the kernel, and thus they must (and do) vanish. This concludes the calculation that (A_\bullet, ∂) is a chain complex.

To check that $A_\bullet \rightarrow C_\bullet$ is a chain map, we must show that for $f \in \mathbb{Z}^{FB} \subset A_2$, we have

$$\partial f_1 - \partial f_2 = \sum_{e \text{ exterior}} \hat{e} - \sum_{e \text{ interior}} \hat{e}.$$

This is shown by direct calculation. □

Let us now prove the following lemma before showing that A_\bullet is quasi-isomorphic to C_\bullet in degree one, and thus computes the first homology $H_1(\Lambda; \mathbb{Z})$.

Lemma 2.13. *In the notation above, $H_2(A_2) = 0$.*

Proof. This says that $\partial : A_2 \rightarrow A_1$ is injective. Suppose $\partial f = 0$. Let h be a hexagonal vertex, which must exist since both Γ and Λ are assumed connected. Label the edges adjacent to h by $e_1, e_2, e_3, e_4 = e'_1, e_5 = e'_2, e_6 = e'_3$ as in the proof of Proposition 2.12. For $1 \leq c \leq 6$, let f_c be the unique (opposite color) monochromatic face containing e_c in its interior, and again we notate $f_4 = f'_1$, etc. Now for $i = 1, 2, 3$, write i, j, k for cyclically ordered elements of $\{1, 2, 3\}$, i.e. $j = i + 1 \pmod{3}$, etc. Then e_i is an exterior edge of f'_j and f'_k and by definition an interior edge of f_i . If we write $f = \sum_{i=1}^3 a_i f_i + a'_i f'_i + \dots$, then we must have $a_i + a_j = a'_k$ for all i , and therefore $\sum a'_i = 2 \sum a_i$. By the same token, $\sum a_i = 2 \sum a'_i$, and therefore all a_i and a'_i are zero.


The faces f_c with coefficients $a_c \neq 0$ must therefore have no hexagonal vertices on their boundary or interior. That said, the union U of such faces must have a boundary, and therefore the coefficient of any face on the boundary of U must be zero. By iterating this argument, all coefficients are zero. □

The 3-graphs associated to a 3-triangulation, and the 3-graph moves named *candy twists* and *push-through*, will be defined in Section 4. We will use them now just in this particular proposition as part of this isolated subsection.¹² Now we establish the point of this subsection:

Proposition 2.14. *Let Γ be a 3-graph for a 3-triangulation, or any graph related by candy twists or push-through moves. Then*

$$H_1(A_\bullet) \cong H_1(C_\bullet) \cong H_1(\Lambda; \mathbb{Z}).$$

Before the proof, a *warning*: $H_1(A_\bullet) \not\cong H_1(C_\bullet)$ in general. Here is an example of a weave with topology of the twice-punctured plane.



Despite $b_1 = 2$, there is only one 1-cycle in $H_1(A_\bullet)$, represented by the tree with four leaves – the sum of edges \hat{e} darkened in the picture. A choice for another generating 1-cycle is clear: it is a branch cut connecting the two trivalent vertices in the top (or bottom) – pictured as a dotted black curve. This class can be represented in C_\bullet , but the chain connecting the two hexagonal vertices is not in A_\bullet . One could accommodate such chains with further notational complexity, but we will not require them for our applications.

¹²This subsection on combinatorial homology is included for completeness, but it will not be used in the rest of the article.

Proof. We need to prove the first equality only. Since $A_\bullet \rightarrow C_\bullet$ is a chain map, we need only compare the dimensions of their first homology groups. We prove this first for the 3-graph Γ_T of a 3-triangulation $T = (F_T, E_T, V_T)$, then show that the result is invariant under candy twist and push-through moves.

By definition, $\partial : A_1 \rightarrow A_0$ is surjective, so since by the lemma, $\partial : A_2 \rightarrow A_1$ is injective, we know $\dim H_1(A_\bullet) = -\chi(A_\bullet)$. On the other hand, we know $\chi_\Lambda = 6 - v = 6 - 3|F_T| = 2 - h_1(\Lambda)$, or $h_1(\Lambda) = 3|F_T| - 4$. We recall that each face of T has three blue vertices. It also has one hexagonal vertex which is a vertex of the blue and red graphs comprising Γ_T . It is similarly easy to compute that $|F_B| = |V_T| + |E_T|$, $|F_R| = |V_T|$, $|E_B| = 2|E_T| + 3|F_T|$, $|E_R| = |E_T|$, $|V_B| = 3|F_T|$, $|V_R| = 0$. Now $|A_2| = |F_B| + |F_R| = 2|V_T| + |E_T|$, $|A_1| = |E_B| + |E_R| = 3|E_T| + 3|F_T|$, and $|A_0| = 2|F_T|$ is computed by noting that each hexagonal vertex contributes two possible dimensions to $|A_0|$ via the rank-two matrix in Equation (2.1), and these dimensions are realized as boundaries, while each trivalent vertex contributes nothing. We get $h_1(A_\bullet) = -\chi(A_\bullet) = -2\chi_T + 3|F_T| = 3|F_T| - 4 = h_1(\Lambda)$, as claimed.

It remains to compute what happens after push-through or a candy-twist move. In fact, since the result only depends on the Euler characteristic of A_\bullet , we only need to show that this is invariant under candy twist and push-through. But these change the dimensions of (A_2, A_1, A_0) by $(2, 6, 4)$ and $(1, 3, 2)$, respectively: a local argument shows again that the two possible dimensions that a hexagonal vertex contributes to A_0 are in fact realized, and the result follows. \square

3. COMBINATORIAL CONSTRUCTIONS

In this section we introduce two combinatorial constructions for N -graphs, focusing primarily on how to associate an N -graph to a given N -triangulation. The notion of an N -triangulation was introduced in [FG06b, Section 1.15], and has since had a central role in higher Teichmüller theory [Gon17, GS18]. Legendrian weaves associated to an N -triangulation, via our construction, place contact topology in the context of the recent developments in exact WKB analysis [GMN13, GMN14, Kuw20] and quiver Fukaya categories [BS15, Smi15].

3.1. N -Triangulations. Let $N \in \mathbb{N}$ be a natural number, and consider the triangle

$$t_N := \{(x, y, z) \in \mathbb{R}^3 : x + y + z = N \quad x, y, z \geq 0\}.$$

Subtriangulate this triangle t_N with the planes

$$(\{x = s\} \cup \{y = s\} \cup \{z = s\}) \cap t_N, \quad 0 \leq s \leq N,$$

which we refer to as an N -subdivision of the triangle t_1 , following [FG06b, GMN14]. This subtriangulation has N^2 triangles.




FIGURE 15. The Legendrian weaves $\Lambda(G(t_3))$, $\Lambda(G(t_4))$ and $\Lambda(G(t_5))$ associated to a 3-, 4- and 5-triangles.

Now, let (C, T) be a triangulation T of a smooth closed surface C and subdivide each triangle $t_1 \in T$ according to the N -subdivision above. This yields a triangulation T_N of the surface C . By definition, an N -triangulation on C is any triangulation isotopic to T_N for some triangulation (C, T) .

3.2. Local Models. The N -graph associated to an N -triangulation is obtained by gluing local models for the N -graph $G(t_N)$ associated to each triangle t_N . We provide a definition of this local N -graph, in terms of the following construction. The reader content with using Figure 15 as a definition is invited to defer reading these technical descriptions.

Construction. Consider the triangles in t_N which point up, i.e. have a unique vertex with highest z -value. For each of these $\binom{N}{2}$ triangles, we insert a τ_1 -trivalent vertex dual to it — that is, a trivalent vertex associated with the permutation (12) and such that the edges of this piece of 2-graph intersects orthogonally with the edges of each triangle. By definition, the rest of the N -graph $G(t_N)$ is then *uniquely* determined by extending the edges from these τ_1 vertices such that wherever three τ_i -edges collide, we insert a hexagonal vertex with three edges in τ_i and three edges in τ_{i+1} . That is, the two rules to generate the N -graph for an N -triangulation are:

- (i) Insert exactly one (12)-trivalent vertex at the center of each upward pointing triangle,
- (ii) In the collision of three τ_i -edges, a (τ_i, τ_{i+1}) hexagonal vertex is inserted.

We stress that the original triangles are not part of the N -graph.

This construction of $G(t_N)$ can be considered as a *dynamical* description, in contrast with the static definition given by the second construction below Remark 3.1. Indeed, in this first construction one starts by placing the τ_1 -vertices and lets the edges grow symmetrically from these trivalent vertices, such that each edge intersects the interior edges of the N -triangulation at the middle point. These edges must collide in the interior of the triangle, and these collisions are resolved via the insertion of hexagonal vertices, creating τ_{i+1} -edges. This insertion of hexagonal vertices is iteratively performed when the τ_{i+1} -edges collide, creating τ_{i+2} -edges, and the process terminates when exactly three τ_{N-1} -edges are created at a unique hexagonal vertex.

Thus, given the triangle t_N , we obtain a local model for an N -graph. The boundary conditions for this local model are such that the N -graphs associated to two N -triangles t_N and t'_N , which share an edge of the underlying t_1 and t'_1 , match together.

Remark 3.1. This description, according to these two rules above, captures the properties of the spectral network associated to the WKB singular foliation for an $SU(2)$ quadratic differential lifted via the unique N -dimensional irreducible representation of $SU(2)$ — see Sections 2 and 4 in [GMN14]. The dynamical component, induced by the growing of the edges from vertices, corresponds with the time evolution of the differential equation defining the WKB system. \square

We can also give a second succinct description of $G(t_N)$ as follows. Following Definition 2.2, it suffices to describe the image of the graphs G_i , $1 \leq i \leq N - 1$. The trivalent graph G_i will be given by the vertices and edges of an hexagonal regular lattice: a finite number of vertices lying inside the triangle $t_N = \{x + y + z = N - 1\}$ and with external edges extending to the boundary of t_N . Let ϑ_i be the set of points of t_N all of whose coordinates lie in $\mathbb{Z}_{\geq 0} + i/3$. Then the vertices of the trivalent graph G_i are precisely the points in $\vartheta_i \cup \vartheta_{i+1}$. Note that the intersection between G_i and G_{i+1} is precisely given by the points in ϑ_{i+1} , and the trivalent vertices are uniquely specified by the hexagonal lattice condition.

Remark 3.2. Both these constructions provide a Legendrian front for the Legendrian lift of certain exact Lagrangian spectral curve for a local spectral network. In particular, this shows that the BPS graphs studied in [GLPY17], introduced as an interpolation between

spectral networks and BPS quivers, are in fact the set of A_1^2 singularities of the Legendrian front $\Sigma(G(t_N))$ between the first two sheets. \square

Note that the boundary of this local N -graph $G(t_N)$ can be compactly described as follows. Consider the permutation

$$\Delta_N := \prod_{i=1}^{N-1} \left(\prod_{j=i}^1 \tau_j \right) = \tau_1 \cdot (\tau_2 \tau_1) \cdot (\tau_3 \tau_2 \tau_1) \cdot \dots \cdot (\tau_{N-1} \tau_{N-2} \cdot \dots \cdot \tau_2 \tau_1) \in S_N,$$

which is the projection to the Coxeter group S_N of the Garside element of the braid group B_N in N -strands, i.e. a braid half-twist in N -strands. Then the edges of the N -graph associated to t_N along each of the three edges of t_1 are precisely given by the ordered terms in Δ_N . That is, there exists an isotopy of the N -graph such that as one travels along an edge of t_1 , the edges of the N -graph that we encounter are first τ_1 , then τ_2 and τ_1 , then τ_3 , τ_2 , and τ_1 and iteratively until reaching τ_1 for the $(N-1)$ th time. This is equivalent to the association $\tau(r_i)$, $i = 1, \dots, 2N-3$, in the construction of $G(t_N)$ above.

For context, these permutations along the boundary are particularly relevant for the study of Legendrian surface weaves with boundary, whose Lagrangian projections yield interesting Lagrangian fillings of their Legendrian boundary links. The braid description of these Legendrian links is determined precisely by these permutations – see Section 7. We see again, confer Remark 2.8, that it is useful to think of Legendrian weaves as two-dimensional Legendrian braids: their one-dimensional boundaries are positive braids.

3.3. Global Model. Given that the boundary conditions for the N -graphs in the local models for t_N allow for gluing, we define the N -graph $G \subseteq C$ associated to a global N -triangulation of C to be the N -graph obtained by concatenating the local models $G(t_N)$ along each triangle t_N in the N -triangulation. We study the flag moduli space invariants for these N -graphs and their associated Legendrian weaves in Sections 6 and 8. Note that the genus of these Legendrian weaves increases as $N \in \mathbb{N}$, or the number of triangles, increases.

Remark 3.3. Trivalent vertices are dual to triangulations of surfaces. In particular, triangulations of surfaces with a large group of symmetries yield particularly interesting 2-graphs. From this perspective, Riemann surfaces with a conformal automorphism group of large order give rise to highly symmetric 2-graphs. For instance, Riemann surfaces associated to tilings of the hyperbolic plane \mathbb{H}^2 with Schläfli symbol $\{n, 3\}$ are highly symmetry, with $\{7, 3\}$ being the Klein quartic, $\{8, 3\}$ giving Bolza’s surface and $\{12, 3\}$ the $M(3)$ surface. We expect the flag moduli space associated to the Legendrian surfaces of these 2-graphs, as defined in Section 5, to be algebraic spaces with correspondingly large symmetry. We begin an exploration of this kind with our Theorem 6.3 in Section 6. \square

3.4. Bicubic graphs. Here is a second construction of 3-graphs in a smooth surface C , strictly disjoint from the class of 3-graphs arising from 3-triangulations.

By definition, a graph is *bicubic* if it is both trivalent (cubic) and bipartite. Now consider an embedded bicubic graph $G \subseteq C$, and replace each vertex of G with a hexagonal vertex, doubling the edges as in Figure 16.

The bipartite condition on the graph guarantees that these local models can be glued together, uniquely up to isotopy, yielding a 3-graph in C . Note that this 3-graph is entirely built from hexagonal vertices, and no trivalent vertex is used. As a result, the topology of the Legendrian weave associated to such a 3-graphs is always that of a 3-component link of Legendrian 2-spheres. We will study a family of such 3-graphs in Section 6.




FIGURE 16. 3-graph associated to an edge in a bicubic graph.

Example 3.4. The bicubic graph $G \subseteq \mathbb{S}^2$ associated to the 1-skeleton of a 3-dimensional cube, depicted in Figure 17, yields a 3-component Legendrian link $\Lambda(G) \subseteq (J^1(\mathbb{S}^2), \xi_{st})$. The flag moduli space $\mathcal{M}(G)$ will show that these three Legendrian spheres, even after satellited to a Darboux ball (\mathbb{R}^5, ξ_{st}) are Legendrian knotted (and smoothly unknotted). \square




FIGURE 17. 3-graph associated to the 1-skeleton of the cube and the 3-component Legendrian link of 2-spheres.

Remark 3.5. Not every 3-graph which is exclusively formed by hexagonal vertices arises from a bicubic graph, even up to candy-twist equivalence (see Section 4). In particular, two vertices may have just a single edge connecting them, with no vertices connected by three edges. Figure 18 shows such an example.




FIGURE 18. 3-graph with only hexagonal vertices which does not arise from a bicubic graph, even up to Move I.

This example can be generalized in several ways: by adding more interior squares with one edge connecting adjacent vertices and/or replacing the squares with $2n$ -gons. \square

Example 3.6. (An Explosion of Examples.) Bicubic graphs can be readily generated as follows. Let P' be a polytope, not necessarily regular, and G' its edge graph, i.e. G' is the one skeleton of P' . Suppose that P' has v' vertices, e' edges and f' faces. By definition, the

explosion of the polytope P' is the polytope P formed by first truncating at the vertices and then truncating the resulting polytope along the original edges of P' . Then the 1-skeleton of P is cubic and has a unique bipartite coloring, up to an overall black-white swap, so therefore is bicubic. Note that P has $v = 4e'$ vertices, $e = 6e'$ edges, and $f = v' + e' + f'$ faces.

Even degenerate polytopes P' give interesting examples. For instance, if P' is the degenerate polytope with two n -gon faces ($v' = n, e' = n, f' = 2$), then P is a $2n$ -gon prism ($v = 4n, e = 6n, f = 2n + 2$). The cube edge graph described in Example 3.4 is the bicubic graph which arises when P' has just two bigon faces. \square

4. DIAGRAMMATIC CALCULUS FOR LEGENDRIAN WEAVES

Let $G \subseteq C$ be an N -graph. The geometric objects that we are interested in are the Legendrian weaves $\Lambda(G) \subseteq (J^1C, \xi_{\text{st}})$ and their invariants up to Legendrian isotopy. In this section we introduce a series of combinatorial operations that can be performed to an N -graph G , and we show how they affect the Legendrian isotopy type of $\Lambda(G)$. The geometric understanding of the Legendrian isotopy type through this diagrammatic calculus allows us to significantly simplify computations of algebraic invariants associated to $\Lambda(G)$ in Section 5. Algebraic computations, using the results in this section, are detailed in Sections 6 and 7. Let us begin with the combinatorial moves in G that preserve the Legendrian isotopy type of $\Lambda(G)$.

4.1. Surface Reidemeister Moves. Let $\Lambda \subseteq (J^1(C), \xi_{\text{st}})$ be a Legendrian surface, a Legendrian isotopy $\{\Lambda_t\}_{t \in [0,1]}$ will generically induce singularities of the Legendrian fibration $J^1C \rightarrow C \times \mathbb{R}$. As a result, the front sets $\Sigma(\Lambda_t)$ and their singularities will restructure as the parameter $t \in [0, 1]$ ranges along a 1-parameter family. These modifications of the Legendrian fronts are referred to as *perestroikas*, or Reidemeister moves [Ad90, Chapter 3].

Remark 4.1. The three classical 1-dimensional Reidemeister moves have been the main method of study for smooth knots in geometric topology, since first introduced [Rei27, AB27]. The corresponding seven moves for smooth surfaces are known as Roseman moves, after [Ros98, Theorem 1]. The corresponding Legendrian Reidemeister, and Legendrian Roseman moves, for Legendrian knots, and Legendrian surfaces, follow from the classification of (stable) wavefront singularities in dimensions $\dim(\Lambda) \leq 3$ [Ad75, Theorem 13]. We will refer to Legendrian Roseman moves as surface Legendrian Reidemeister moves. \square

The combinatorial operations inducing surface Legendrian Reidemeister moves are the content of the following theorem. In the moves, the local pieces of the N -graphs are actually 3- or 4-graphs. The color code follows our standard notation: blue and red are adjacent colors (corresponding to adjacent transpositions), red and yellow are adjacent colors, and blue and yellow are disjoint colors.

Theorem 4.2. *Let G_1, G_2 be one of the pairs of N -graphs depicted in Figures 19, 20, 21, 22, 23, 24 and 25. Then the associated Legendrian surface $\Lambda(G_1)$ is Legendrian isotopic to $\Lambda(G_2)$ relative to their boundaries. That is, Moves I, II, III, IV, V, VI and VI' are local surface Legendrian Reidemeister moves.*




FIGURE 19. (Move I) The first pair of local N -graphs G_1 , on the left, and G_2 on the right. We refer to this move as *candy twist*.




FIGURE 20. (Move II) The second pair of local N -graphs G_1 , on the left, and G_2 on the right. We refer to this move as the *push-through*, since the trivalent vertex gets pushed through the hexagonal vertex.




FIGURE 21. (Move III) The third pair of local N -graphs G_1 , on the left, and G_2 on the right. We refer to this move as the *flop*.




FIGURE 22. (Move IV) The fourth pair of local N -graphs G_1 , on the left, and G_2 on the right. Note we must have $N \geq 4$. This move implies the A_3 generalized Zamolodzhikov relation depicted in Figure 106.

Proof. Let us start with Move I, the candy twist, as depicted in Figure 19. It illustrates the method of proof for these surface Legendrian Reidemeister moves. There are essentially three equivalent viewpoints: exhibiting the Legendrian isotopy as N -graphs, visualizing the surface wavefronts explicitly in $(J^1\mathbb{R}^2, \xi_{st})$, or studying these surface wavefronts as families of (possibly singular) Legendrian links. In the first perspective, we need to justify that all the N -graphs lift to embedded Legendrian surfaces. In the second, the challenge is visualizing the actual front and ensuring that all the singularities lift to Legendrian embeddings. In




FIGURE 23. (Move V) The fifth pair of local N -graphs G_1 , on the left, and G_2 on the right, with $N \geq 4$. The blue and yellow colors are associated to disjoint transpositions.




FIGURE 24. (Move VI) The sixth pair of local N -graphs G_1 , on the left, and G_2 on the right, with $N \geq 4$.




FIGURE 25. (Move VI') Variation on the sixth pair of local N -graphs G_1 , on the left, and G_2 on the right, with $N \geq 4$.

the third perspective, visualization is simplified, with the trade-off of having to draw several movies of links. The second perspective is the strongest, as it readily implies the other two. The first perspective is drawn in Figure 26. The left and rightmost 3-graphs lift to Legendrian weaves, yielding embedded Legendrian surfaces (with boundary). The diagram in the center of Figure 26 does *not* immediately lift to an embedded Legendrian surface, as the six-valent vertex is *not* a hexagonal vertex – the colors of the edges around it are not alternating, which is the condition for the hexagonal vertices introduced in Section 2.




FIGURE 26. The 3-graph movie showing that the candy move - Move I - is a Legendrian isotopy. The geometric meaning of the central picture (not a 3-graph) is explained in the text.

Nevertheless, the center diagram in Figure 26 does in fact come from a Legendrian wavefront whose Legendrian lift is an embedded surface. Indeed, we have depicted such a front in the second front of Figure 27.




FIGURE 27. The homotopy of Legendrian wavefronts associated to Move I.

The movie of wavefronts in Figure 27 geometrically constructs the homotopy of Legendrian fronts which lifts to the Legendrian isotopy corresponding to Move I. The three fronts in Figure 27 lift to *embedded* Legendrian surfaces, as the singularities are all Legendrian and there are no vertical tangent planes. The singularities at the beginning of Figure 27 are segments of A_1^2 -crossings, and two isolated A_1^3 points. The singularities at the end of Figure 27 are just segments of A_1^2 -crossings. The singularity in the middle of the movie, not corresponding to an A_1^2 segment, is *not* a stable front singularity, but it does lift to an embedded Legendrian surface, and thus the homotopy of fronts actually represents a Legendrian isotopy. Indeed, the tangent spaces at that singularity intersect transversely, and hence their lifts are disjoint. This concludes that Move I combinatorially represents a surface Legendrian Reidemeister move.

Remark 4.3. For completeness, in Figure 28 we have drawn the homotopy of surface fronts from Figure 27 as a movie (of movies). It is thus a 2-homotopy of Legendrian links. These three movies of links, one per each column, are obtained by slicing each of the respective fronts in Figure 27 from left to right. This is the third viewpoint we mentioned above. \square

Let us now justify Move II, where a D_4^- -singularity pushes-through an A_1^3 -singularity. The resulting front has a D_4^- -singularity and *two* A_1^3 -singularities. The clearest proof that this is a Legendrian isotopy comes from carefully drawing and examining the right homotopy of fronts. In this case, the required movie of fronts is depicted in Figure 29. These Legendrian fronts start with the front whose A_1^2 -singularities yield the 3-graph G_1 on the left of Move II, and end with the front whose A_1^2 -singularities yield the 3-graph G_2 on the right of Move II.

These fronts describe a neighborhood \mathbb{R}^3 of a D_4^- -singularity with a 2-plane $\Pi \subseteq \mathbb{R}^3$ which starts away from the D_4^- -singularity. This 2-plane Π is drawn with a tilt in its slope. The homotopy of fronts consists of this 2-plane Π moving towards the D_4^- -singularity and *crossing* through it. There exists a unique moment in this isotopy in which the D_4^- -singularity is




FIGURE 28. The proof that Move I is a Legendrian isotopy by (transversely) slicing each of the Legendrian wavefronts in Figure 27.

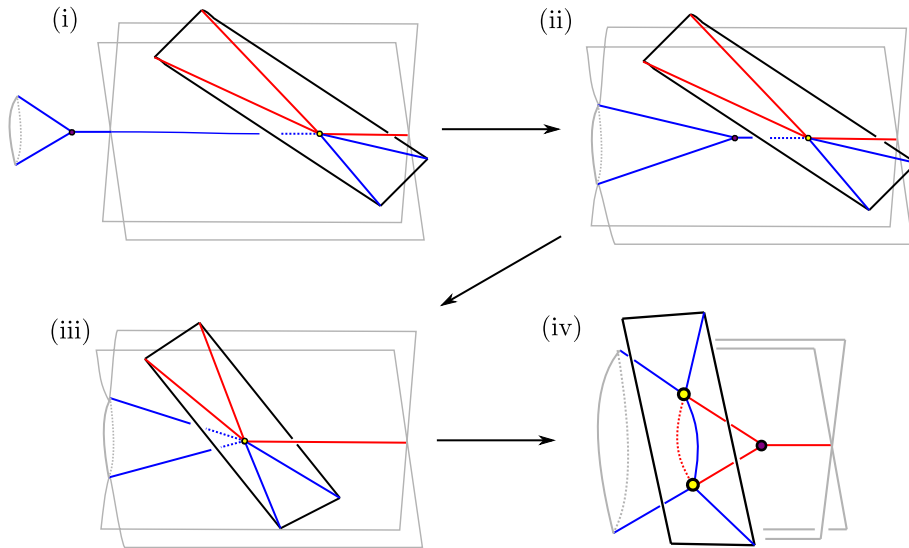


FIGURE 29. The homotopy of Legendrian fronts inducing Move II.

contained in the 2-plane Π . The A_1^2 -singularities right before that moment give rise to G_1 for Move II, and right after this moment the A_1^2 -singularities give rise to G_2 for Move II. Since the 2-plane Π is not vertical, and the tangent 2-planes of the different branches at the D_4^- -singularity in all moments are distinct, each of the fronts in this homotopy lift to embedded Legendrian surfaces. Thus, the movie of fronts in Figure 29 shows that there exists a Legendrian isotopy with A_1^2 -singularities as dictated by Move II, and $\Lambda(G_1)$ and $\Lambda(G_2)$ are Legendrian isotopic relative to their boundaries. This concludes Move II.

For Move III, we can proceed analogously by drawing a homotopy of fronts which lifts to a Legendrian isotopy. Nevertheless, Move III can actually be deduced as a combination of Moves I and II. We leave it as an exercise for the reader to visualize the spatial Legendrian fronts, and instead explain how to deduce Move III from the previous two moves, as follows.

Starting with one side of Move III, push both trivalent vertices through in the clockwise direction using Move II. This is depicted in the first two steps of Figure 30.




FIGURE 30. Deducing Move III from Moves I and II. The first step pushes a red trivalent vertex through a hexagonal vertex. The second step does the same for the other red trivalent vertex. The third step is a simplification undoing a candy twist. The dashed green lines indicate push-through moves that are about to occur.

This creates additional hexagonal vertices and the two trivalent vertices do change color. Perform Move II twice more, pushing-through these trivalent vertices again, and then cancel two pairs of hexagonal vertices with a candy twist (Move I) to obtain the right hand side of Move III. Alternatively, first undo the candy twist as in the third step of Figure 30: this yields a 3-graph which is identical to a partial rotation of the initial 3-weave with red and blue switched. Iterating this again, i.e. pushing the two blue trivalent vertices through, as indicated by the dashed green lines in Figure 30, and undoing a candy twist yields the right hand side of Move III.

Let us now show that Move IV is a Legendrian isotopy. The corresponding spatial wavefronts consist of configurations of four 2-planes. The graph G_1 on the left of Move IV is obtained as the A_1^2 -singularities, i.e. intersections, of the union of the four 2-planes

$$\begin{aligned} \pi_x &= \{(x, y, z) \in \mathbb{R}^3 : x + 0.0001z = 0\}, & \pi_y &= \{(x, y, z) \in \mathbb{R}^3 : y + 0.0001z = 0\}, \\ \pi_z &= \{(x, y, z) \in \mathbb{R}^3 : z = 0\}, & \pi_1 &= \{(x, y, z) \in \mathbb{R}^3 : x + y + z = 1\}. \end{aligned}$$

These intersections and 2-planes are depicted, with the corresponding colors, in Figure 31. Now consider the 2-planes $\pi_t = \{(x, y, z) \in \mathbb{R}^3 : x + y + z = t\}$, $t \in [-1, 1]$. The homotopy of spatial wavefronts is locally given by the union $\pi_x \cup \pi_y \cup \pi_z \cup \pi_t$, $t \in [-1, 1]$.

This homotopy is not relative to the boundary, as the 2-planes π_t , $t \in [-1, 1]$, change the boundary conditions — but this is easily corrected by only pushing a compact piece of π_t , $t \in [-1, 1]$ through the triple intersection point $\pi_x \cap \pi_y \cap \pi_z$. The A_1^2 -singularity pattern of the resulting spatial wavefront is precisely as in the right graph G_2 in Move IV, as required.




FIGURE 31. Front for the start of Move IV. The lines depict the intersections of the union of the four 2-planes $\pi_x \cup \pi_y \cup \pi_z \cup \pi_1$.

Let us now address Move V, which depicts the local transition between two 4-graphs G_1, G_2 in Figure 23. The corresponding spatial fronts consist of four 2-planes $\pi_1, \pi_2, \pi_3, \pi_4 \subseteq \mathbb{R}^3$, where the only non-empty intersections are $\pi_1 \cap \pi_2$, corresponding to the blue segment in G_1 (and G_2), and $\pi_3 \cap \pi_4$, corresponding to the yellow segment in G_1 , and G_2 .

The fact that the fronts giving G_1 and G_2 are homotopic as Legendrian fronts is proven in Figure 32. Each of the columns in the figure represents a spatial surface front, with the links in the columns corresponding to slices. The corresponding intersections, dictating the A_1^2 -singularities, are marked with the same color as in Figure 23. The union of these slices in Figure 32 yield spatial fronts which lift to embedded Legendrian surfaces, and thus the movie of columns in Figure 32 exhibits a Legendrian isotopy from $\Lambda(G_1)$ to $\Lambda(G_2)$. Therefore, Move V is a surface Legendrian Reidemeister move. Move VI in Figure 24 follows with the same argument as for Move V, with a segment of A_1^2 -singularities passing above, and disjointly, a D_4^- -singularity — and likewise for Move VI'. This concludes the proof of Theorem 4.2. \square

Remark 4.4. The Legendrian Reidemeister moves in Theorem 4.2 provide a symplectic geometric realization of A-type Soergel calculus. Moves I and V should be compared to [EW16, Figure 4.4]. Move II and Move VI are known as two-color associativity of type $A_1 \times A_1$, with Coxeter exponent $m_{st} = 2$, and of type A_2 , with Coxeter exponent $m_{st} = 3$, and Move IV corresponds to the A_3 relation [EW16, Figure 4.7]. It should be emphasized that the notation in Soergel calculus follows the notation for (rank three) parabolic subgroup of finite Coxeter groups, whereas we use the notation for Lie algebras whose irregular Weyl orbits yield spatial wavefronts. See Appendix A for further details. \square

Theorem 4.2 contains the Reidemeister moves that we use in the course of the article. They are all the possible (generic) Legendrian surface moves with only D_4^- and A_1^2 Legendrian singularities in the endpoints of the Legendrian isotopy. The complete set of surface Reidemeister moves [Ad90, Section 3.3] also includes the moves associated to the A_4 and D_4^+ -singularities, which will require the interaction of A_2 -cusp edges A_2 and A_3 -swallowtails.

Theorem 4.2 allows one to make local modifications to an N -graph G_1 and obtain an N -graph G_2 such that the Legendrian surfaces $\Lambda(G_1) \cong \Lambda(G_2) \subseteq (J^1C, \xi_{st})$ are Legendrian isotopic. For the case $C = \mathbb{S}^2$, we define in Subsection 4.7 an additional combinatorial move, which we

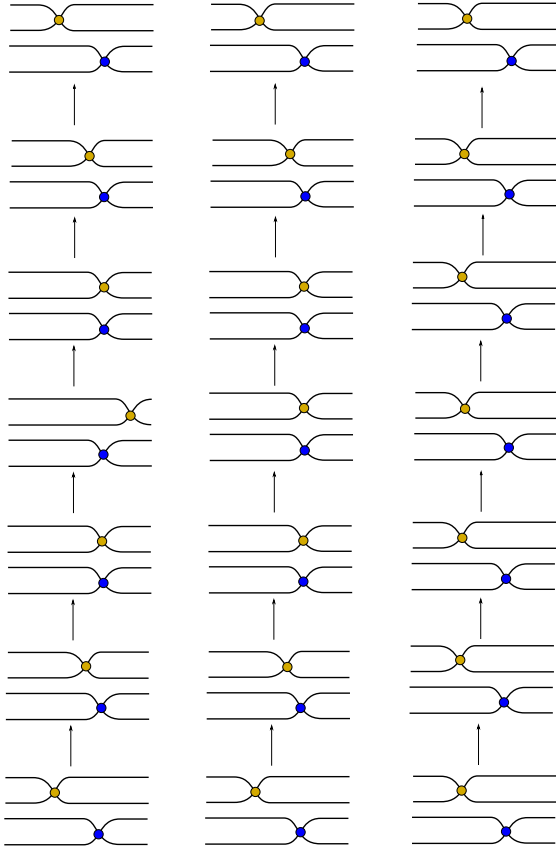


FIGURE 32. Each column represent slices of a spatial front. The A_1^2 -singularities of the left column gives rise to G_1 in Figure 23, and the A_1^2 -singularities of the right column gives rise to G_2 .

refer to as a *stabilization*, going from an N -graph G_1 to a $(N + 1)$ -graph G_2 . This requires a discussion on satellite constructions for Legendrian weaves, which is useful on its own, and also needed for Subsection 4.5.

4.2. Legendrian Satellite Weaves. Let $G \subseteq C$ be an N -graph. The Legendrian surface $\Lambda(G)$ defined by the weaving construction lies in the contact 5-manifold (J^1C, ξ_{st}) . Now, consider a contact 5-manifold (Y, ξ) and a Legendrian embedding $\iota : C \rightarrow (Y, \xi)$. The Weinstein Neighborhood Theorem [Wei71, Section 7] for Legendrian submanifolds gives a contactomorphism

$$\tilde{\iota} : (J^1C, \xi_{\text{st}}) \rightarrow (\mathcal{O}p(\iota(C)), \xi|_{\mathcal{O}p(\iota(C))}),$$

where $\mathcal{O}p(A)$ is a sufficiently small neighborhood of $A \subseteq Y$, and such that the restriction to the zero section $C \subseteq J^1C$ is the initial Legendrian embedding ι . In particular, any Legendrian $\Lambda \subseteq (J^1C, \xi_{\text{st}})$ yields a Legendrian $\tilde{\iota}(\Lambda) \subseteq (Y, \xi)$. Thus, the contact 1-jet spaces serve as *local* contact manifolds, and a Legendrian embedding of C in an arbitrary ambient contact 5-manifold allows one to embed a Legendrian weave there as well. In this context, the Legendrian surface $\tilde{\iota}(\Lambda) \subseteq (Y, \xi)$ is called the *ι -satellite* of $\Lambda \subseteq (J^1C, \xi_{\text{st}})$ and the Legendrian surface $\iota(C) \subseteq (Y, \xi)$ is called the *companion*. This terminology parallels the theory of satellite knots, as introduced in [Sch53], and see also [NR13, EV18]. Notice that the smooth topology of Λ and its satellite $\tilde{\iota}(\Lambda)$ is identical, only the ambient contact manifold (and thus the Legendrian embedding type) are affected by this Legendrian satellite construction.

Example 4.5. Let $(Y, \xi) = (\mathbb{S}^5, \xi_{st})$, $C = \mathbb{S}^2$, and let $\iota = \iota_0$ be the Legendrian embedding of the standard Legendrian unknot $\iota_0 : \mathbb{S}^2 \rightarrow \mathbb{S}^5$. Given any Legendrian $\Lambda \subseteq (J^1\mathbb{S}^2, \xi_{st})$, we will refer to $\tilde{\iota}_0(\Lambda) \subseteq (\mathbb{S}^5, \xi_{st})$ as the standard satellite of Λ . Since $(\mathbb{S}^5 \setminus \{\text{pt}\}, \xi_{st}) \cong (\mathbb{R}^5, \xi_{st})$, and the image $\tilde{\iota}_0(\Lambda)$ will avoid some point, this surface can be equivalently considered in $\tilde{\iota}_0(\Lambda) \subseteq (\mathbb{R}^5, \xi_{st}) \cong (J^1(\mathbb{R}^2), \xi_{st})$. It can thereupon be described by its front projection to $\mathbb{R}^3 = \mathbb{R}^2 \times \mathbb{R}$. This is depicted in Figure 33. \square




FIGURE 33. A Legendrian surface $\Lambda \subseteq (J^1\mathbb{S}^2, \xi_{st})$ drawn in the front projection $\mathbb{S}^2 \times \mathbb{R}$ (Left). The satellite $\tilde{\iota}_0(\Lambda)$ of Λ along the standard Legendrian unknot $\iota_0 : \mathbb{S}^2 \rightarrow (\mathbb{R}^5, \xi_{st})$, drawn in the front projection $\mathbb{R}^2 \times \mathbb{R}$ (Right). These pictures are schematic and ought to be rotated symmetrically along their central vertical axis so that the wavefronts for Λ and $\tilde{\iota}_0(\Lambda)$ are indeed surfaces in a 3-dimensional ambient space.

In case no Legendrian embedding ι is specified and $C = \mathbb{S}^2$, the notation $\iota(\Lambda)$ will implicitly refer to the *standard* satellite $\tilde{\iota}_0(\Lambda) \subseteq (\mathbb{R}^5, \xi_{st})$ as in Example 4.5 and Figure 33. It is often the case that the Legendrians $\Lambda(G) \subseteq (J^1C, \xi_{st})$ that we introduce in this work do not have an *a priori* name nor they have been previously studied. Interestingly, for a certain variety of graphs $G \subseteq \mathbb{S}^2$ we will see how their standard Legendrian satellites are actually related to well-known Lagrangian surfaces, e.g. see Subsection 6.1.

In addition, and in line with Markov's Theorem for smooth 1-dimensional braids [Bir74, PS97], the satellite operation is also required for a meaningful *stabilization* operation. Finally, note also that even if $\Lambda(G) \subseteq (J^1C, \xi_{st})$ has no A_2 -cusp edges, the spatial wavefronts for its standard satellite $\tilde{\iota}_0(\Lambda(G))$ will *always* have A_2 -cusp edges, as any front for the standard Legendrian unknot $\Lambda_0 \subseteq (\mathbb{R}^5, \xi_{st})$ must have A_2 -cusp edges. We now discuss A_2 -cusp edges and A_3 -swallowtail singularities, which are required for such a *stabilization* operation and Theorems 4.10 and 4.21 below, regarding Legendrian surgeries and Legendrian mutations.

4.3. Cusp Edges and Swallowtail Singularities. Let $G \subseteq C$ be an N -graph, the Legendrian weave $\Lambda(G) \subseteq (J^1C, \xi_{st})$ associated to G is determined by its front $\pi(\Lambda(G)) \subseteq C \times \mathbb{R}$. By definition, these fronts only have D_4^- , A_1^2 and A_1^3 singularities. The latter two are stable, i.e. a generic Legendrian isotopy $\Lambda_t \subseteq (J^1C, \xi_{st})$, $t \in [0, 1]$, such that $\Lambda(G) = \Lambda_0$, will have each of the A_1^2 and A_1^3 singularities of the front $\pi(\Lambda_0)$ persist for $\pi(\Lambda_t)$, $t \in (0, 1]$. In contrast, D_4^- is not: the fronts $\pi(\Lambda_t)$, $t \in (0, \varepsilon]$, will *not* have any D_4^- -singularity for $\varepsilon \in \mathbb{R}^+$ small enough.

The generic (stable) singularities of fronts in 3-dimensional space are A_1^2 , A_1^3 , A_2 , A_2A_1 and A_3 , as shown in [Ad90, Section 3.2]. These singularities are depicted in Figure 35. The appearance of A_2 , A_2A_1 and A_3 singularities in a generic front forces us to extend our combinatorial diagrammatics, as our Legendrian isotopies will (typically) be generic. In the figures for this subsection, and only this subsection, we will draw edges around a hexagonal vertex with




FIGURE 34. The Legendrian front of an A_3 -swallowtail singularity (left). The planar diagrammatic depiction in our calculus (right).

the same color – this will simplify our diagrams, which are no longer N -graphs due to the presence of A_2 -cusp edges.




FIGURE 35. The generic Legendrian singularities of wavefronts in 3-space. The depicted A_3 -singularity is known as the A_3 -swallowtail, and the center A_2 -singularity in the first row is referred to as the A_3 -cusp edge. Note that the two D_4^\pm -singularities are *not* generic.

We extend the diagrammatics with the following rule: *orange* segments will denote A_2 -cusp edges of singularities, and *orange* dots will stand for A_3 -swallowtail singularities. Figure 34 depicts on its left a genuine spatial front for the A_3 -swallowtail singularity. The singularities of this front consist of a segment of A_1^2 -crossings, shown in blue, two A_2 -cusp edges, in orange, and a unique A_3 -swallowtail point. The planar diagram through which we represent this front is shown on the right of Figure 34. It is simply a vertical view of the front (from above or below) with the A_1^2 , A_2 and A_3 -singularities marked.

Remark 4.6. For the same reasons that we label A_1^2 singularities with transpositions, in order to indicate which two sheets are crossing, we should label A_2 -cusp edges with the corresponding information. This is necessary information in order to recover the actual (homotopy type of the) Legendrian front, and thus the Legendrian itself. That said, in this article, it should be clear from context where such A_2 -cusp edges lie, so these labels will be omitted. □

The D_4^- -singularities are the central pieces in the construction of our Legendrian weaves $\Lambda(G) \subseteq (J^1C, \xi_{\text{st}})$. It is important to emphasize that D_4^- is *not* a generic singularity of a real spatial front, despite the fact that its complexification is a stable holomorphic Legendrian singularity. In particular, in our upcoming study of Legendrian surgeries, we will need generic Legendrian isotopies starting at $\Lambda(G)$, whose fronts will break the non-generic D_4^- into *generic* singularities of real spatial wavefronts.

The generic deformation of the D_4^- -singularity is depicted in Figure 36 (left). It contains three A_3 -swallowtails arranged in a triangle and connected by A_2 -cusp edges. Following our convention above, the associated planar diagram is shown in Figure 36 (right).




FIGURE 36. The spatial wavefront for a generic perturbation of the D_4^- -singularity (left). The associated planar diagram for this stable spatial wavefront (right). Note that the A_1^2 -edges around the hexagonal vertex all drawn with the same color (blue), following the convention in this subsection.

4.4. Legendrian Front Calculus with Cusp Singularities. Let us continue our development of a diagrammatic front calculus for Legendrian surfaces, this time including A_2 -cusp edges and A_3 -swallowtails. Proposition 4.7 below is used to prove Proposition 4.9 and also Theorem 4.10, in the upcoming Subsection 4.5.

Proposition 4.7. *Let $G \subseteq C$ be an N -graph, $N \in \mathbb{N}$. The four moves in Figure 37 are achieved by compactly supported Legendrian isotopies, relative to the boundary.*

Proof. Moves VII and VIII, on the creation and fusion of two A_3 -swallowtails singularities are immediate from the 3-dimensional First Reidemeister Move R1. Indeed, the left-to-right 1-dimensional Legendrian slices in Move VII correspond to a concatenation of R1 and its inverse, i.e. an R1 is performed, corresponding to the appearance of the leftmost A_3 -swallowtail, and then the same R1 is undone, corresponding to the appearance of the rightmost A_3 -swallowtail. This movie of 1-dimensional Legendrian slices can be isotoped to a movie with no R1 fronts, whose (big) front corresponds to the right of Move VII, with no swallowtails. For Move VIII, the R1 moves are performed in reverse order. That is, the left-to-right 1-dimensional Legendrian slices correspond to the inverse of an R1 move (a pair of cusps being undone) and then the exact same R1 move. This homotopy of 1-dimensional Legendrian fronts can be itself homotoped to a constant homotopy, which the local N -graph depicted in the right of Move VIII.

For Move IX, we proceed with our slicing techniques. The 1-dimensional vertical left-to-right slices of the two fronts for Move IX are depicted in the left and right columns of Figure 38. In the left column, the Reidemeister R1 move is performed for the upper piece of the 1-dimensional Legendrian knot. In the right column, the R1 move is performed for the lower

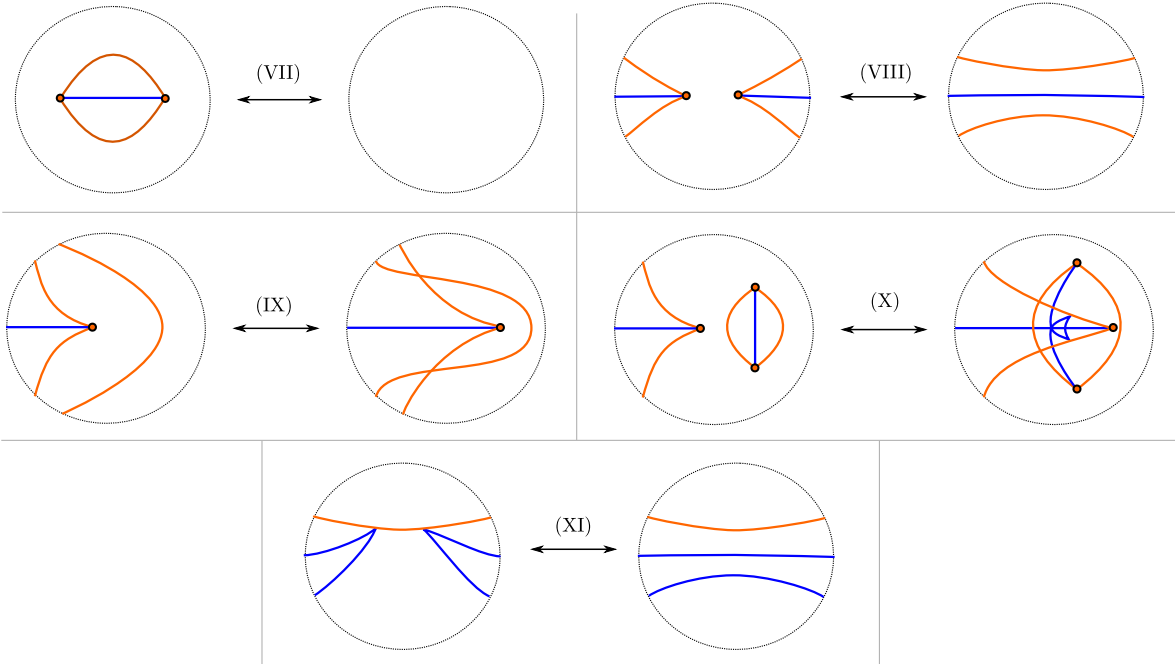


FIGURE 37. The five Legendrian front moves in Proposition 4.7. The moves are referred to as Move VII (upper left), Move VIII (upper right), Move IX (center left), Move X (center right) and Move XI (lower center).

piece of the 1-dimensional Legendrian knot. The homotopy of Legendrian surface fronts is achieved by the center column in Figure 38, where both R1 are performed simultaneously.




FIGURE 38. The homotopy of surface fronts showing that Move IX is a Legendrian Reidemeister move. The left-to-right slices for the left diagram in Move IX are depicted in the left column, whereas the slices for the right diagram in Move IX are depicted in the right column.

Since the homotopy of fronts preserves the boundary conditions, this lifts to a Legendrian isotopy of embedded Legendrian surfaces, thus proving that Move IX is a Legendrian Reidemeister move. The fact that Move IX is a Legendrian Reidemeister move also follows carefully from visualizing the critical fronts associated to the generating family

$$D_4^+ : F(x, y, \xi_1, \xi_2, \xi_3) = x^2y + y^3 + \xi_1y^2 + \xi_2y + \xi_3x,$$

which leads to the above families in Figure 38.

Move X consists of a sliding for a A_3 -swallowtail along an A_1^2 -crossing line, as depicted in the top row of Figure 39, in Figures 39.(a) and 39.(b). The realistic surface fronts are depicted in the bottom row of Figure 39, in Figures 39.(A) and 39.(B), where the A_3 -swallowtail singularity has been moved past the A_1^2 -segment of singularities.

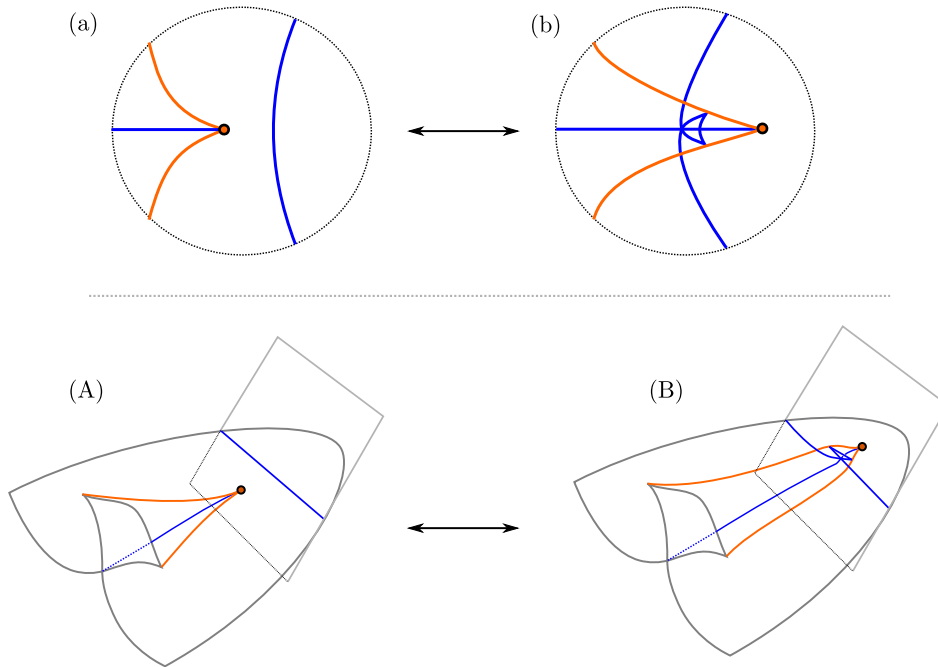


FIGURE 39. The front depiction of the non-trivial part in Move X. The A_3 -swallowtail singularity slides across a orthogonal A_1^2 -line, changing sheets as it slides through.

The sliding lifts to a Legendrian isotopy, as the interaction between the A_3 -swallowtail and the A_1^2 -line only sees a critical moment, where a $A_3A_1^2$ singularity appears. At this critical stage, the slopes are all distinct and non-vertical, thus the A_3 -swallowtail is allowed to move past with a homotopy of fronts. This concludes that Move X is a Legendrian Reidemeister move.

Finally, Move XI is proven in Figure 40. The middle singularity corresponds to the generic spatial front A_2A_1 -singularity. In short, Move XI is obtained by performing a homotopy which interpolates between a constant movie of Legendrian links, and a movie consisting of doing a Reidemeister R2 move and then undoing it, as in the left column of Figure 40. \square




FIGURE 40. The homotopy of fronts for Move XI. The left front diagram of Move XI is obtained as the union of the slices in the left column, whereas the right front in Move XI is the union of the slices in the right column.

Remark 4.8. It would appear that Reidemeister moves for Legendrian knots have been mastered by the vast majority of contact topologists. This does not seem to be the case in higher dimensions, including the Legendrian singularities appearing in surface fronts. Should the reader be interested in that, [Ben86, Ad90] provides a starting presentation of the generic singularities of surface fronts. Our present manuscript develops the diagrammatic calculus adding to that classification, which allows us to manipulate fronts in a versatile manner. The combination of the results of this article, along with [Ad90], should permit the reader to be fluent in the manipulation of wavefronts for Legendrian surfaces in contact 5-manifolds. \square

Let us now address the move shown in Figure 41, which we prove in the following:

Proposition 4.9. *The combinatorial move depicted in Figure 41 is realized by a compactly supported Legendrian isotopy of surfaces in a 5-dimensional Darboux ball $(J^1\mathbb{R}^2, \xi_{st})$, relative to the boundary.*




FIGURE 41. (Move XII) This move allows us to exchange A_3 -swallowtail singularities with D_4^- -singularities in the presence of a A_2 -cusp edge.

Proof. Let us start with the left front in Figure 41. Apply Move VII to create a canceling pair of A_3 -swallowtails, as shown in the beginning of Figure 42. Now slide the A_3 -swallowtail by

performing a Move X, and use the D_4^+ -singularity, i.e. Move IX to exchange the A_2 -cusp edge where the A_3 -swallowtail connects. This is depicted in the first and second steps of Figure 42. The next two steps in Figure 42 consists of Legendrian isotopies where no singularities interact with each other, it is a plain homotopy of fronts with the *same* singularities. Finally, the last step consists in joining the three existing A_3 -swallowtails into a *single* D_4^- -singularity, as depicted at the end of Figure 42.




FIGURE 42. The homotopy of fronts for Move XII. The initial A_3 -swallowtail requires two additional swallowtails to become a D_4^- -singularity, and certain intermediate moves. The homotopy realizing this can be read in this picture.

□

4.5. Legendrian Surgeries. The theory of Legendrian surgeries was initiated in [Ad76, Ad79] in the study of critical points of the time function with respect to a Legendrian wavefront. Its modern description in terms of Lagrangian handle attachments is described in [BST15, Theorem 4.2] and [DR16, Section 4]. A Legendrian surgery on $\Lambda \subseteq (Y, \xi)$ is an operation which inputs an isotropic sphere within Λ , bounding ambiently, and outputs a Legendrian $\tilde{\Lambda} \subseteq (Y, \xi)$. The Legendrians Λ and $\tilde{\Lambda}$ are not even homotopy equivalent, and thus Legendrian surgery is a useful method to create *new* Legendrians by modifying the topology of a given Legendrian Λ .

In the context of Legendrian surfaces, there are different types of Legendrian surgeries [Ad90, Figure 48]. The following result characterizes the combinatorial operations that correspond to Legendrian 0-surgeries, 1-surgeries and Legendrian connected sums.

Theorem 4.10 (Legendrian Surgeries). *Let $G \subseteq C$, $G_1 \subseteq C_1$ be N -graphs and $G_2 \subseteq C_2$ an M -graph, for $N, M \in \mathbb{N}$. The following statements hold:*

1. (*0-Surgery*) *The combinatorial move of adding an i -edge and two vertices along an existing i -edge corresponds to a Legendrian 0-surgery. This move is shown in the*

upper right diagram in Figure 43.

2. (1-Surgery) The combinatorial move of removing an i -edge between two trivalent vertices corresponds to a Legendrian 1-surgery. This move is shown in the lower left of Figure 43.

3. (Connect Sum) The kissing of two trivalent vertices $v_1 \in G_1$ and $v_2 \in G_2$, where $G_1 \subseteq C_1, G_2 \subseteq C_2$ are two disjoint graphs, corresponds to a connect sum

$$\iota(\Lambda(G_1)) \# \iota(\Lambda(G_2)) \subseteq (\mathbb{R}^5, \xi_{st}),$$

for any satellite $\iota : \Lambda \rightarrow (\mathbb{R}^5, \xi_{st})$. This is shown in the upper left of Figure 43.

4. (Clifford Sum) The combinatorial move of substituting a trivalent vertex by a triangle corresponds to a connected sum of $\iota(\Lambda(G))$ with a Clifford 2-torus $\mathbb{T}_c^2 \subseteq (\mathbb{R}^5, \xi_{st})$. This move is shown in the lower right of Figure 43.

The 0-surgeries, 1-surgeries are local in any $\Lambda(G) \subseteq (J^1C, \xi_{st})$. In contrast, the connected sum in the third item requires to geometrically satellite the Legendrian weaves $\Lambda(G_1) \subseteq (J^1C_1, \xi_{st})$ and $\Lambda(G_2) \subseteq (J^1C_2, \xi_{st})$ via any Legendrian embedding

$$\iota : C_1 \cup C_2 \rightarrow (\mathbb{R}^5, \xi_{st}).$$




FIGURE 43. The Legendrian Surgery Moves in Theorem 4.10

Theorem 4.10 will be proven below. The Legendrian weaves in the statements involve only D_4^- and A_1^2 (and A_2 -cusp edges for the connected sum, due to the satellite operation). Nevertheless the manipulation of their fronts in the proof of Theorem 4.10 requires the use of further Legendrian front moves, involving A_3 -swallowtails and A_2A_1 -singularities and their interaction with the A_2, A_1^2 and D_4^- -germs, as developed in Subsection 4.4 above.

Remark 4.11. (i) Should the reader be solely interested in the satellited Legendrian surface $\iota(\Lambda(G)) \subseteq (\mathbb{R}^5, \xi_{st})$, the connected sum operation in Theorem 4.10.(3) is the strongest of the four statements (and the hardest to prove). Indeed, the satellite analogue of Items 1,2 and 4 follow from Item 3. That said, Items 1,2 do *not* follow from Item 3 locally.

(ii) Note also that the ι -satellite of the Legendrian 0-surgery depicted in Move (1) of Figure 43, and Theorem 4.10.(1), corresponds to a Legendrian connected sum with the standard

Legendrian 2-torus in (\mathbb{R}^5, ξ_{st}) . This is the 2-torus whose front is obtained by \mathbb{S}^1 -front spinning of the saucer front for the standard Legendrian unknot in (\mathbb{R}^3, ξ_{st}) . See [DR11, Section 4.1], and Figure 6 therein, and also [BST15, DR16]. \square

We recall that, by definition, the index of an elementary exact Lagrangian cobordism is the Morse index of its unique critical point, see [BST15, Section 4.1] and [DR16, Section 4]. Note that elementary index- k exact Lagrangian cobordisms are also referred to as Lagrangian k -handle attachments. In particular, the Legendrian convex end of an elementary index- k exact Lagrangian cobordism is a Legendrian $(k - 1)$ -surgery on the Legendrian concave end. In combination with Theorem 4.2, Theorem 4.10 yields the following two moves:

Corollary 4.12. *The two N -graph moves in Figure 44 corresponds to a Legendrian 1-surgery, i.e. upon performing $(2')$, or $(2'')$, there exists an elementary index-2 exact Lagrangian cobordism from the Legendrian weave on the left to the Legendrian weave on the right.*

In fact, in Move $(2')$ the Lagrangian 2-disk is attached along the 1-cycle represented by the (bi)chromatic horizontal edge between the two trivalent vertices. In Move $(2'')$ the Lagrangian 2-disk is attached along the 1-cycle represented by the (blue) tripod at the hexagonal vertex uniting the three trivalent vertices.




FIGURE 44. The two Legendrian Surgery Moves in Corollary 4.12, both representing Lagrangian 2-handle attachments.

Proof of Theorem 4.10. We start by proving that adding an i -edge with two trivalent vertices to an existing i -edge effects a Legendrian 0-surgery, i.e. a Lagrangian 1-handle attachment. The homotopy of spatial fronts is depicted in Figure 45, according to the conventions in Subsection 4.3. The detailed description reads as follows. We first generically perturb the two D_4^- -singularities in the first spatial front, which yields the second front. Performing Move VIII and then Move I yields the third and fifth fronts, respectively, in Figure 45. Note that the homotopy from the third to the fourth front does not involve any change in the singularities of fronts, as the blue segment of A_1^2 -singularities intersecting the orange A_2 -cusp segment lies strictly below it in 3-space. The homotopy from the fifth to the sixth front emphasizes the yellow band where the (reverse) 1-surgery is to be performed. The step from the sixth to the seventh fronts is precisely the reverse surgery: the A_2 -cusp edges in the seventh front are surgered along the yellow band [Ad76, BST15], in the sixth front, to obtain the fifth front. The seventh front is homotopic to the eighth front by Move VII.

Let us now show that removing an i -edge corresponds to a Lagrangian 2-handle attachment, i.e. a Legendrian 1-surgery. The homotopy of fronts is depicted in Figure 46. Starting with the first front, generically perturbing yields the second front and two applications of Move

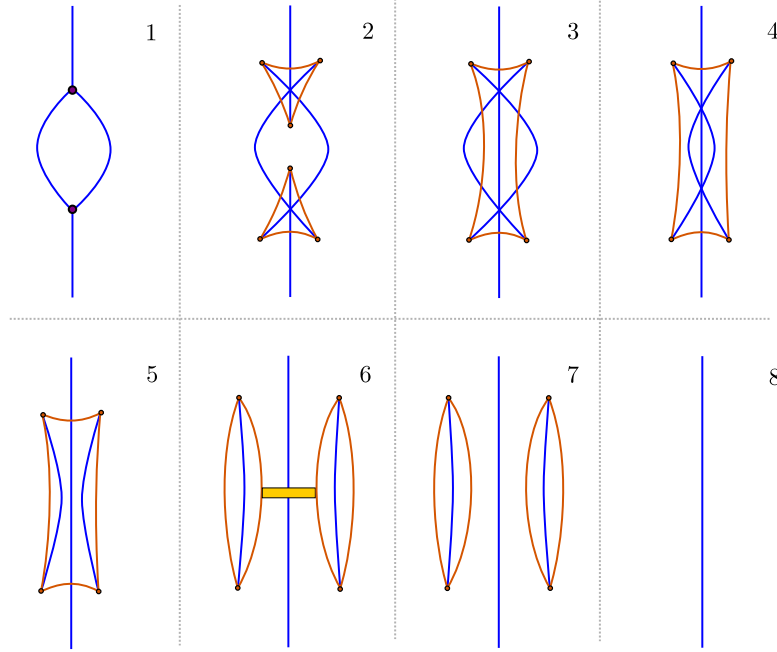


FIGURE 45. The diagrammatic homotopy of spatial fronts associated to the Legendrian 1-surgery move. It shows that the first front is a Legendrian 0-surgery on the eight front, i.e. the result of a Lagrangian 1-handle attachment.

VIII give the third front. In the fourth front we have shown the Legendrian 2-disk (in yellow) along which we perform the 1-surgery [Ad90, BST15], the result of which is the fifth front. Indeed, the 1-surgery opens up the inner circle of A_2 -cusp edges and adds two horizontal (Legendrian) 2-disks. As a result, the effect on its diagrammatic representation is removing the inner circle of A_2 -cusps, as shown in the fifth front. The application of Move I gives the sixth front, which is readily homotopic to the seventh front. The eighth front is then obtained by performing a Move VII.

Now, we prove that joining two trivalent vertices in distinct graphs $G_1 \subseteq C_1, G_2 \subseteq C_2$ is realized by a Legendrian surface connected sum, which is a Lagrangian 1-handle attachment (a Legendrian 0-surgery) whose attaching 0-sphere has its two points belonging to different boundary components. The required homotopy of fronts is shown in Figure 47. In this case, we must satellite the Legendrian weaves $\Lambda(G_1), \Lambda(G_2)$ to a Darboux ball (\mathbb{R}^5, ξ_{st}) . From the perspective of spatial fronts, we must locally add a A_1^2 -curve and two A_2 -cusp edges as depicted in the first front of Figure 47. The Legendrian 0-surgery is performed from the first front to the second, along the Legendrian band given by the red dotted line. The homotopy from the second front to the third consists of four applications of Move XI. Then, we use Move XII to obtain the fourth front. The fifth front is achieved by applying Move VII, and the sixth front consists of two applications of Move XI.

Finally, substituting a trivalent vertex by a triangle corresponds to a connected sum with the four vertex graph $G_c \subseteq \mathbb{S}^2$ in the left of Figure 48. One then shows that the spatial front of the Legendrian weave $\iota(\Lambda(G_c)) \subseteq (\mathbb{R}^5, \xi_{st})$ is front equivalent to the front on the right of Figure 48, which is known to be the Legendrian lift of the Clifford torus [DR11, CM19]. In brief, this can be shown by first identifying the Legendrian 2-torus associated to the Clifford graph with

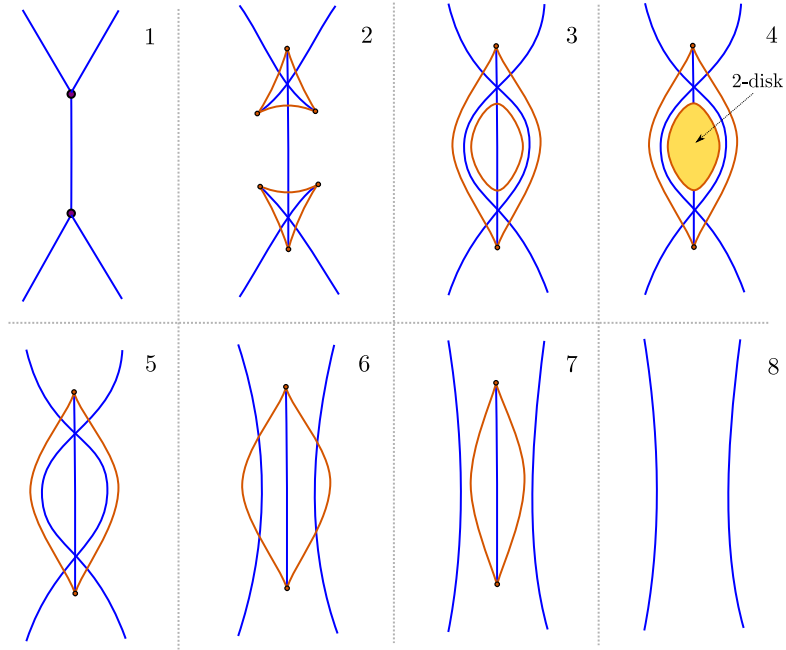


FIGURE 46. The diagrammatic homotopy of spatial fronts associated to the Legendrian 1-surgery move. It shows that the eighth front is a Legendrian 1-surgery on the first front, i.e. the result of a Lagrangian 2-handle attachment. Note that diagrams 3 and 4 in the first row are the same. The difference is that in diagram 4 we have depicted (in yellow) the 2-disk along which the Legendrian 1-surgery is performed, which overlaps with part of diagram 3 (and thus this part is not depicted in diagram 4).




FIGURE 47. The diagrammatic homotopy of spatial fronts associated to the Legendrian connected sum.

the vanishing cycle associated to the superpotential $W : \mathbb{C}^3 \rightarrow \mathbb{C}$, $W(z_1, z_2, z_3) = z_1 z_2 z_3$.

This superpotential has a singular Lagrangian thimble

$$L = \{(z_1, z_2, z_3) \in \mathbb{C}^3 : W(z_1, z_2, z_3) \in \mathbb{R}_{\geq 0}, |z_1| = |z_2| = |z_3|\},$$

whose intersection with the contact unit 5-sphere $(\mathbb{S}^5, \xi_{\text{st}}) \subseteq \mathbb{C}^3$ is a Legendrian 2-torus L_W . It is shown in [Nad17b, Section 3.3] that the Clifford graph is a front for this 2-torus L_W . In order to obtain the cone front from Figure 48 (on the right), one stereographically projects from $(\mathbb{S}^5, \xi_{\text{st}}) \subseteq \mathbb{C}^3$ to $(\mathbb{R}^5, \xi_{\text{st}})$ with the contactomorphism provided in [Gei08, Proposition 2.1.8] and draws the (image of L_W in the) front projection. The resulting front for L_W is precisely the one drawn on the right of Figure 48.




FIGURE 48. The Clifford graph $G_c \subseteq \mathbb{S}^2$ and a simplified spatial front for the satellited Legendrian $\iota(\Lambda(G_c)) \subseteq (\mathbb{R}^5, \xi_{\text{st}})$.

□

Proof of Corollary 4.12. In Figure 44, Move (2') follows by applying a sequence of Moves II to the leftmost trivalent vertex, pushing that vertex through *all* the hexagonal vertices – until it is connected to the rightmost trivalent vertex with a monochromatic edge – and then using Move (3) in Theorem 4.10. Move (2'') is more interesting, and its proof is shown in Figure 49.

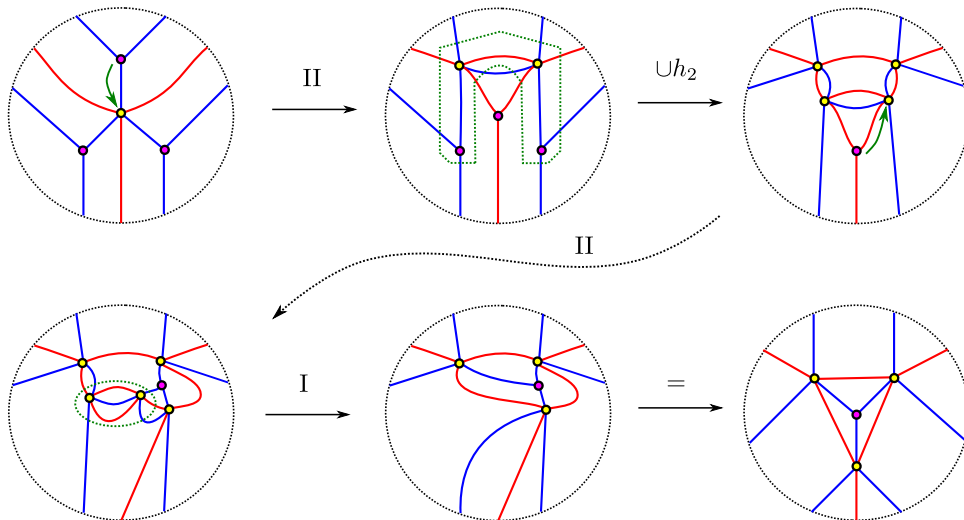


FIGURE 49. The Lagrangian 2-handle attachment in Move (2'') decomposed as a sequence of surface Reidemeister moves, from Theorem 4.2, and Move (2) in Theorem 4.10, in the guise of Corollary 4.12.

□

Theorem 4.10 provides a useful and efficient way to describe Legendrian surfaces in terms of N -graph combinatorics. Its statement is as strong as possible, in that the conclusion is on the Legendrian isotopy type of the associated Legendrian weaves. The computation of algebraic invariants then follows as a consequence of our geometric understanding.

In particular, we have following.

Corollary 4.13. *Let $G \subseteq C$ be an N -graph and $v \in G$ a trivalent vertex. The blow-up combinatorial move on G , given by an insertion of a triangle at the vertex v , is a twisted 0-surgery on $\iota(\Lambda(G))$. \square*

The blow-up procedure was first studied in [TZ18, Section 5]. It is depicted in Figure 43 (lower right). By definition, a twisted 0-surgery is a connected sum with a non-standard Legendrian torus in (\mathbb{S}^5, ξ_{st}) . For now, we refer to [DR11, Section 4] for more details.

A consequence of Corollary 4.13 is that the Legendrian isotopy type of $\iota(\Lambda)$ is independent of the choice of vertex $v \in G$, because a twisted 0-surgery is independent of the choice of 0-sphere at which it is performed (since all pair of points are isotopic in a connected surface). This question was initially asked in [TZ18] in the study of the dependence of the sheaf invariants in terms of v . Since the Legendrian isotopy type of $\iota(\Lambda)$ is independent of v , the algebraic invariants are also independent of v .

Finally, note that the Legendrian 0-surgery in Theorem 4.10.(1) can be understood as a Legendrian connected sum with the 2-graph $G \subseteq \mathbb{S}^2$ shown in Figure 50 (Left). In fact, the standard Legendrian satellite $\iota(\Lambda(G))$ for this 4-vertex 2-graph is the standard Legendrian 2-torus, a Legendrian front of which is shown in Figure 50 (Right). Indeed, they are both obtained from the standard Legendrian unknot by a 0-surgery (which yields a unique Legendrian isotopy class of Legendrian 2-tori) and thus they must be Legendrian isotopic.




FIGURE 50. A 2-graph G in the 2-sphere \mathbb{S}^2 (Left) and a Legendrian front for its Legendrian weave $\iota(\Lambda(G))$ (Right). This is the standard Legendrian 2-torus $\mathbb{T}_{st}^2 \subseteq (\mathbb{R}^5, \xi_{st})$, given by Legendrian front spinning of the 1-dimensional standard Legendrian unknot $\Lambda_0 \subseteq (\mathbb{S}^3, \xi_{st})$.

Remark 4.14. The Legendrian 0- and 1-surgeries in Theorem 4.10 physically correspond to partial puncture degenerations in the context of spectral networks [GMN13, GMN14]. Indeed, the Legendrian weaves obtained as the Legendrian lift of the Lagrangian hyperkähler rotation of the spectral curve of a diagonalizable Higgs field are related by the Legendrian surgeries in Theorem 4.10. For instance, the process of a full puncture [1,1,1] degenerating to a simple [2,1] puncture in a punctured 3-sphere is precisely a Legendrian 0-surgery [GLPY17, Section 6]. \square

The Reidemeister moves in Subsection 4.1 and the stabilization operation in Subsection 4.7 preserve the Legendrian isotopy type of the (satellite) Legendrian weaves. The Legendrian surgeries discussed in Theorem 4.10 generically change the topology of $\Lambda(G)$. The natural next step is to modify the Legendrian isotopy type of $\Lambda(G)$ without changing its topology, which we will discuss in Subsection 4.8. For now, we study an explicit example and present the stabilization operation.

4.6. Example of a Closed Legendrian Weave. Let us illustrate our spatial front calculus in an example. Consider the triangulation of $C = \mathbb{S}^2$ given by a tetrahedron, and the 3-graph G associated to this triangulation according to Section 3. This 3-graph is shown in Figure 51 (upper left). The 3-graph G is depicted in the plane as an *unfolded* triangulation, thus the triangles should be identified according to the faces of the tetrahedron: the outer three vertices of the dashed triangle are identified, and the dashed lines are glued accordingly. In particular, the 3-graph G has twelve trivalent vertices and four hexagonal vertices. The question is to describe the Legendrian isotopy type of this Legendrian surface $\iota(\Lambda(G)) \subseteq (\mathbb{R}^5, \xi_{\text{st}})$. In addition, we would like to compute Legendrian invariants, such as the augmentation variety of 3-dimensional Lagrangian fillings in $(\mathbb{D}^6, \omega_{\text{st}})$. In this context, understanding the Legendrian isotopy type readily implies the computation of this Legendrian invariant.

We will exploit Theorem 4.2 and Theorem 4.10 to understand this Legendrian weave, and note that the closed surface $\iota(\Lambda(G)) := \tilde{\iota}_0(\Lambda(G))$ has genus 4. First, we describe the sequence of Legendrian moves and surgeries in Figure 51. In Diagram (1) on the upper left, first note that there are three blue triangles each having one vertex in the central triangle, one each in two outer triangles, and passing through one glued edge. There is another blue triangle with one vertex on each of the outer triangles. By Theorem 4.10, we conclude that Diagram (1) corresponds geometrically to a connected sum of the weave from Diagram (2) with four copies of the Clifford 2-torus \mathbb{T}_c^2 . The 3-graph of Diagram (2) is still complicated, so we use Theorem 4.2 to simplify. First apply Move III, flopping the four vertices in the upper right of the 3-graph. This brings us to Diagram (3). Now do a Move I to undo the newly appearing candy twist.




FIGURE 51. Simplification of a 3-graph with Theorem 4.2 and Theorem 4.10.

This brings us to Diagram (4). So we have proven that the standard satellite $\iota(\Lambda(G))$ is Legendrian isotopic to $\iota(\Lambda(G')) \#_{i=1}^4 \mathbb{T}_c^2$, where G' is the 3-graph in Diagram (4) of Figure 51. It now suffices to understand the Legendrian $\iota(\Lambda(G')) \subseteq (\mathbb{R}^5, \xi_{\text{st}})$.

Assertion: Let $G' \subseteq \mathbb{S}^2$ be the 3-graph in Figure 52 (upper left). The Legendrian 2-sphere $\iota(\Lambda(G'))$ is Legendrian isotopic to the standard Legendrian unknot $\Lambda_0 \subseteq (\mathbb{R}^5, \xi_{\text{st}})$.

Proof of the assertion: By Theorem 4.10, we can undo the two bigons in Diagram (5) of Figure 52, and understand them as two connect sums with the standard Legendrian 2-torus \mathbb{T}_{st}^2 , defined as any Lagrangian 1-handle attachment to the standard Legendrian unknot $\Lambda_0 \subseteq (\mathbb{R}^5, \xi_{\text{st}})$.




FIGURE 52. Diagrammatic proof that the standard satellite of the Legendrian 2-sphere associated to the 3-graph in Diagram (5) is the standard Legendrian unknot two-sphere $\Lambda_0 \subseteq (\mathbb{R}^5, \xi_{\text{st}})$.

By applying Move I in Theorem 4.2 to the 3-graph in Diagram (6), we arrive at the 3-graph G'' in Diagram (7) of Figure 52, which simplifies to the three concentric circles of alternating colors in Diagram (8). The Legendrian weave $\iota(\Lambda(G'')) \subseteq (\mathbb{R}^5, \xi_{\text{st}})$ is readily seen to be the standard 3-component unlink $\Lambda_0 \cup \Lambda_0 \cup \Lambda_0 \subseteq (\mathbb{R}^5, \xi_{\text{st}})$. Indeed, in 3-dimensional contact topology, the standard satellite of an N -stranded braid along the unknot – with its standard saucer front – creates a w_0^{-1} of crossings at each side of the braid, where $w_0 \in S_N$ is the longest element. That is, a S^0 -worth of w_0^{-1} crossings. In particular, a positive braid given by w_0^2 , which consist of a S^0 -worth of w_0 -crossings, will get satellited to the standard Legendrian N -component unlink. See e.g. [CN20, Section 2.2]. By S^1 -symmetrically rotating this picture, we conclude that an N -weave in S^2 given by $\binom{N}{2}$ concentric circles whose colors exactly give w_0 will be satellited along the standard 2-dimensional unknot to a standard Legendrian N -component unlink. Here the case at hand is $N = 3$ and it suffices to note that red-blue-red represents $w_0 \in S_3$.

In conclusion, $\iota(\Lambda(G'))$ is obtained by performing Lagrangian 1-handle attachments to $\Lambda_0 \cup \Lambda_0 \cup \Lambda_0 \subseteq (\mathbb{R}^5, \xi_{\text{st}})$, and thus $\iota(\Lambda(G'))$ must be the standard Legendrian unknot. \square

The conclusion of the above discussion is that the Legendrian isotopy type of the Legendrian surface $\iota(\Lambda(G)) \subseteq (\mathbb{R}^5, \xi_{\text{st}})$ associated to 3-triangulation of the tetrahedron, i.e. Diagram (1) of Figure 51, is that of the connected sum of four copies of the Clifford 2-torus \mathbb{T}^2 . Hence, we

now have a complete geometric understanding of $\iota(\Lambda(G))$. In particular, this readily implies [Siv11, DR11] that the \mathbb{C} -moduli of objects of the category of microlocal rank-one sheaves in \mathbb{R}^3 supported in $\iota(\Lambda(G))$ is isomorphic to $(\mathbb{C} \setminus \{0, 1\})^4$. \square

4.7. N -Graph Stabilization. The Reidemeister moves introduced in Theorem 4.2 constitute combinatorial operations on a given N -graph G which yield the same Legendrian isotopy type for the associated Legendrian weave $\Lambda(G)$, as a Legendrian in (J^1C, ξ_{st}) . In particular, the resulting graph is still an N -graph.

In this section we discuss a different type of combinatorial move, where the number of sheets $N \in \mathbb{N}$ is increased. This operation, which we call *stabilization*, inputs an N -graph $G \subseteq C$ and outputs an $(N + 1)$ -graph $s(G) \subseteq C$. The main property of stabilization, proven in Theorem 4.17 below, is that it preserves the Legendrian isotopy type of the standard Legendrian satellite $\iota(\Lambda(G)) \subseteq (\mathbb{R}^5, \xi_{\text{st}})$, and as a result it is a non-characteristic operation.

Remark 4.15. The relative homology class of the surface $\Lambda(G) \subseteq J^1C$ has order N , and thus no combinatorial operation that modifies the number $N \in \mathbb{N}$ of sheets for a Legendrian weave will ever yield a Legendrian isotopic surface in the 1-jet space J^1C . Therefore, preserving the Legendrian isotopy type for the (standard) *satellite* is the optimal statement for a stabilization operation. \square

Let us describe the Legendrian weave stabilization. Given an N -graph G , the first step is to introduce a *ladybug* trivalent graph B in $(N, N + 1)$ as depicted in blue in the left of Figure 53 in such a way that G is completely contained in one face¹³ of B , i.e. G is inside one of the wings of the ladybug B . The second step is the introduction of *descending halos* centered at an $(N + 1)$ -graph G , which consists of a nested set of $N - 1$ circles of A_1^2 -crossings indexed by the permutations $(N - 1, N), (N - 2, N - 1), \dots, (23), (12)$ reading outward. This is depicted in the right of Figure 53.

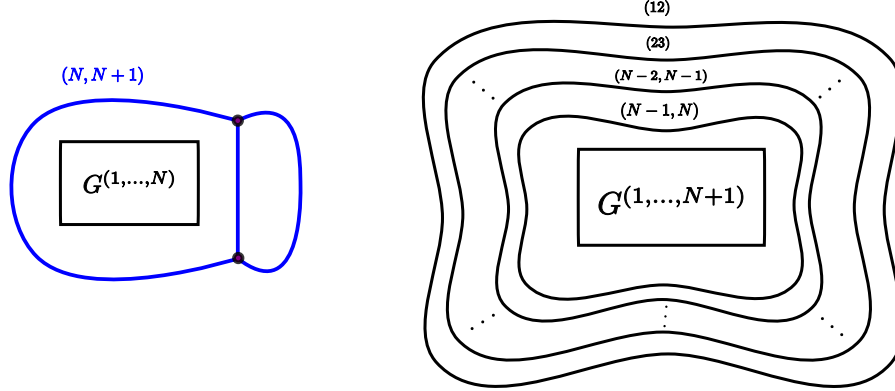


FIGURE 53. Ladybug graph B around G (left) and halos centered at G (right).

The concatenation of these two operations leads to the following:

Definition 4.16. Let $G \subseteq C$ be an N -graph. The stabilization of G is the $(N + 1)$ -graph $s(G) \subseteq C$ obtained from G by placing a ladybug B around G , labeled with the transposition $(N, N + 1)$, and a sequence of descending halos centered at the $(N + 1)$ -graph $G \cup B$. \square

Figure 54 depicts the stabilization for the cases $N = 2, 3$. The ladybug graph B is shown in blue.

¹³The construction is independent of the choice of such face.




FIGURE 54. Stabilization of a 2-graph (left) and of a 3-graph (right).

The stabilization in Definition 4.16 is the Legendrian surface generalization of the Type II Markov move for smooth N -strand braids [Mar35, Bir74]. The main property of graph stabilization is the following geometric result:

Theorem 4.17. *Let $G \subseteq \mathbb{S}^2$ be an N -graph. Then the standard satellites $\iota(\Lambda(G))$ and $\iota(\Lambda(s(G)))$ are Legendrian isotopic in (\mathbb{S}^5, ξ_{st}) .*

Proof. Let us provide a detailed proof for the case $N = 2$, where the stabilization is a 3-graph. The argument for higher $N \geq 3$ is identical. Consider the standard satellite closure $\iota(\Lambda(s(G)))$, which yields the diagram on the left of Figure 55 – we refer the reader to Figure 33 for the front of the standard satellite closure. The standard satellite closure of a 3-graph introduces three circles of A_1^2 -crossings, drawn in dark grey, and three circular cusp edges, drawn in orange.¹⁴ Perform a Legendrian isotopy which exchanges the (12)-circle of A_1^2 -crossings with the adjacent (34)-circle of A_1^2 -crossings; this gives the diagram in the right of Figure 55. This move is possible thanks to the cusp sliding shown in the first two columns of Figure 56.




FIGURE 55. Exchange of (12) and (34) circles of A_1^2 -crossings.

Then use the innermost cusp circle and perform a Move XI, also denoted R_1^2 as it consists of two Reidemeister I moves, to remove two of the A_1^2 -crossings as in the left of Figure 57, this corresponds in the slice to the third column of Figure 56. Iterate with an R_1^2 in the same

¹⁴For a general N -graph, a front for the standard satellite closure of the Legendrian weave contains N additional sheets, $(N + 1), (N + 2), \dots, 2N$. The bottom N sheets $1, \dots, N$ are woven according to G , and the top horizontal N sheets are parallel. The bottom and top sheets are then connected by circles worth of A_1^2 -crossings, according to the half-twist $\Delta \in \text{Br}_N^+$, and N circles worth of A_2 -cusp edges – see Figure 33.




FIGURE 56. The left three diagrams depict slices in the dotted segments for Figure 55. The rightmost diagram depicts a slice for the dotted segment in the right of Figure 57.

cusped edge with the (34)-circle of crossings and the ladybug piece B , arriving at rightmost diagram in Figure 57.

Finally, eliminate the two half-moons in the cusped edge and isotope the cusped edge above the graph $G^{(1,2)}$, which is possible thanks to the configuration shown at the rightmost column of Figure 56. The resulting diagram is that on the left of Figure 58, which is Legendrian isotopic to the diagram on its right. by applying two Moves XII, from Figure 41, and an inverse Move VII from Figure 37.




FIGURE 57. Performing an R_1^2 -move with (34) and the ladybug.

□

In this manuscript, Reidemeister moves in Subsection 4.1 and the Stabilization in Theorem 4.17 form the set of combinatorial moves that is available to us when manipulating an N -graph, if the Legendrian isotopy type of the associated (satellite) Legendrian weave is to be preserved.

4.8. Legendrian Mutations. We now discuss the N -graph combinatorics of Legendrian mutations, a new geometric operation that we define in this manuscript. This operation inputs a Legendrian surface $\Lambda \subseteq (\mathbb{R}^5, \xi_{st})$ and an isotropic 1-cycle $\gamma \subseteq \Lambda$, and outputs a Legendrian surface $\mu_\gamma(\Lambda) \subseteq (\mathbb{R}^5, \xi_{st})$. The Legendrian surface $\mu_\gamma(\Lambda) \subseteq (\mathbb{R}^5, \xi_{st})$ will be




FIGURE 58. From $N = 2$ to $N = 3$ (left) and $N = 3$ to $N = 4$ (right).

ambiently (relatively) smoothly isotopic to Λ , and oftentimes not Legendrian isotopic to Λ . The choice of notation aims at emphasizing its relation to the wall-crossing phenomenon [GMN10, KS10, KS14], Lagrangian mutation [Pol91, Aur07, Aur09] and [FOOO09, Chapter 10], and quiver mutations [FZ02, Via14].

Definition 4.18. Let $G \subseteq C$ be an N -graph and $e \in G$ and i -edge between two trivalent vertices. The mutation of G along e is the N -graph $\mu_e(G)$ obtained by performing the exchange depicted in Figure 60 (left), also shown in Figure 4 (3). \square

By Theorem 4.21 below, the Legendrian weaves $\Lambda(G)$ and $\Lambda(\mu_e(G))$ will be mutation-equivalent, according to the upcoming 4.19 – this motivates Definition 4.18 from the perspective of contact topology. Note that the operation in Definition 4.18 is the simplest possible mutation, corresponding to the combinatorics associated to a Whitehead move, i.e. an edge flip in the context of triangulations dual to 2-graphs. Indeed, consider the two unique non-degenerate triangulations T_1, T_2 of the square, the dual 2-graphs G_1, G_2 differ precisely by a mutation along their unique internal edge.

Correspondingly, the standard satellites of their associated Legendrian weaves are two Legendrian cylinders with coinciding Legendrian boundary, smoothly isotopic relative to their boundary but which are *not* Legendrian isotopic relative to their boundary.

In general, given a 1-cycle $\gamma \in \Lambda(G)$ which is expressed combinatorially in G , it is possible to describe the mutation of G along such 1-cycle γ . The mutated graph $\mu_\gamma(G)$ can either be defined in an *ad hoc* way, or rather be understood as a graph which is equivalent via Reidemeister moves, as in Subsection 4.1, to the mutated graph $\mu_e(\gamma)(G')$. Here G' is Reidemeister equivalent to G and $e(\gamma)$ is an i -edge between trivalent vertices such that $[e(\gamma)] = [\gamma] \in H_1(\Lambda(G), \mathbb{Z})$ under the canonical identification $H_1(\Lambda(G), \mathbb{Z}) \cong H_1(\Lambda(G'), \mathbb{Z})$ given by a Legendrian isotopy. Here is the definition:

Definition 4.19 (Legendrian Mutation). Two Legendrian surfaces $\Lambda_0, \Lambda_1 \subseteq (\mathbb{R}^5, \xi_{st})$ are mutation-equivalent if and only if there exists a compactly supported Legendrian isotopy $\{\tilde{\Lambda}_t\}_{t \in [0,1]}$ relative to the boundary $\partial\Lambda_0$, with $\tilde{\Lambda}_0 = \Lambda_0$, and a Darboux ball (B, ξ_{st}) such that

- (i) The two restrictions $\tilde{\Lambda}_1|_{(\mathbb{R}^5 \setminus B)} = \Lambda_1|_{(\mathbb{R}^5 \setminus B)}$ coincide away from this Darboux ball,
- (ii) There exists a global front projection $\pi : \mathbb{R}^5 \rightarrow \mathbb{R}^3$ such that each of the spatial fronts $\pi|_B(\tilde{\Lambda}_1)$ and $\pi|_B(\Lambda_1)$ respectively coincide with each of the two fronts in Figure 59.

\square




FIGURE 59. Legendrian mutation in a local spatial wavefront.

The two fronts depicted in Figure 59 coincide at their boundaries and lift to Legendrian cylinders. These Legendrian cylinders are *not* Legendrian isotopic relative to their boundary. Indeed, compactifying the upper sheet of the fronts with an A_2 -cusp edge and a flat 2-disk, and the lower sheet with a different A_2 -cusp edge and a flat 2-disk, yields the standard Legendrian unknot $\Lambda_0 \subseteq (\mathbb{R}^5, \xi_{\text{st}})$ for the left front in Figure 59, and a loose Legendrian 2-sphere $\mathfrak{s}(\Lambda_0)$ for the right front in Figure 59. The Legendrians $\Lambda_0, \mathfrak{s}(\Lambda_0) \subseteq (\mathbb{R}^5, \xi_{\text{st}})$ are not Legendrian isotopic [EES05a, EES05b].

A strong motivation for the study of the above mutations is the production of Legendrian surfaces which are *not* Legendrian isotopic, even though they belong to the same formal Legendrian isotopy class [Gro86, EM02]. In order to distinguish Legendrian isotopy classes we will be using *flag moduli spaces*, which synthesize Legendrian invariants coming from the study of microlocal sheaves in terms of algebraic geometry.

Remark 4.20. The conic Legendrian singularity for the front in Figure 59 (left) is *not* a generic singularity. It is explained in detail in [DR11, CM19], and its generic perturbation contains four A_3 -swallowtail singularities. \square

Theorem 4.21 (Legendrian Mutations). *Let G_1, G_2 be one of the pairs of N -graphs depicted in Figure 60. Then the associated Legendrian surface $\Lambda(G_1)$ is a Legendrian mutation of $\Lambda(G_2)$ relative to their boundaries.*




FIGURE 60. The Legendrian Mutation Moves in Theorem 4.21

Proof. Let us start by showing that the exchange move in Figure 60 (left) corresponds to a Legendrian mutation, as in Definition 4.19. By [CMP19, Theorem 6.3], the Lagrangian projections $\Pi(\Lambda_0), \Pi(\Lambda_1) \subseteq \mathbb{R}^4$ of the Legendrian lifts of the fronts $\pi(\Lambda_0), \pi(\Lambda_1) \subseteq \mathbb{R}^3$ in Figure 59 correspond to the two Polterovich surgeries associated to the normal crossing of two Lagrangian planes $\mathbb{R}^2 \times \{0\}, \{0\} \times \mathbb{R}^2 \subseteq (\mathbb{R}^4, \omega_{\text{st}})$. The Lagrangian projection of the Legendrian lifts for each two g -graphs in the exchange move in Figure 60 (left) are exact Lagrangian fillings L_1, L_2 of the Hopf link $\Lambda_{\text{Hopf}} \subseteq (\mathbb{S}^3, \xi_{\text{st}}) \cong \partial(\mathbb{R}^4, \omega_{\text{st}})$. Indeed, the 2-stranded braid word at the boundary of the 2-weave is σ_1^4 , as there are four blue edges arriving at the boundary, and then note that the (-1) -framed closure of σ_1^4 in $(\mathbb{R}^3, \xi_{\text{st}})$ is the Hopf link. See Section 7 for more details on Lagrangian fillings. Thus, it suffices to

show that $L_1, L_2 \subseteq (\mathbb{R}^4, \xi_{\text{st}})$ are the positive and negative Polterovich surgeries of the two Lagrangian planes $\mathbb{R}^2 \times \{0\}, \{0\} \times \mathbb{R}^2 \subseteq (\mathbb{R}^4, \omega_{\text{st}})$ at their intersection points. Indeed, Figure 61 (center) depicts the 2-graph for the *singular* Legendrian whose Lagrangian projections is the Lagrangian union $(\mathbb{R}^2 \times \{0\}) \cup (\{0\} \times \mathbb{R}^2)$.




FIGURE 61. The 2-graphs associated to a Legendrian Mutation. The middle 2-graph yields a spatial front which lifts to a *singular* Legendrian surface, consisting of the union of two 2-planes intersecting at a point.

The 2-graph in Figure 61 (center) describes a topological surface which is the union of 2-planes intersecting at a point, both for the Lagrangian surfaces in $(\mathbb{R}^4, \omega_{\text{st}})$ and the Legendrian surfaces in $(\mathbb{R}^5, \xi_{\text{st}})$. Topologically, the front in Figure 61 (center) is the cone over the annular projection of the $(2, 4)$ -braid, with singular crossings¹⁵.

Finally, the Lagrangian projections of the Legendrian lifts of Figure 61 (left) and Figure 61 (right) are realized as Polterovich surgeries of the corresponding Lagrangian projection in Figure 61 (center). Since the Legendrian lifts of Polterovich surgeries are Legendrian mutations [CMP19, Theorem 6.3], this concludes the first part of Theorem 4.21.

Let us now show that the exchange move in Figure 60 (right) also corresponds to a unique Legendrian mutation. This is proven directly through the homotopy of fronts in Figure 62.

Indeed, the first step in Figure 62, starting from the upper left, consists of applying Move II, pushing a trivalent vertex through a hexagonal vertex. The second and third steps are also a direct application of a Move II, pushing the remaining two trivalent vertices through the newly created *two* hexagonal vertices. The fourth move, starting at the left of the second row, is a mutation of 2-graphs. This yields the 3-graph at the center of the second row, the arrow being labeled by the letter μ . Finally, we apply a Move III, flopping the four vertices nearest to the center, in order to achieve the 3-graph at the right of Figure 60 (right). This shows that the exchange move in Figure 60 (right) is a Legendrian mutation.

□

For our applications to Lagrangian fillings, it is important to understand how 1-cycle representatives of classes in $H_1(\Lambda(G), \mathbb{Z})$ change under the mutations depicted in Figure 60. Following Subsection 2.4, we focus on 1-cycles represented by monochromatic edges – or more generally *long* edges – and by Y-cycles. Figure 63 explicitly shows how to transport certain 1-cycles along the mutation. (See Section 2, specifically Subsection 2.4, for the definition of 1-cycles.) In addition, mutation along a long edge is dictated by the following:

¹⁵This is consistent with the fact that the Hopf link is the boundary of two transversely intersecting planes in the 4-ball \mathbb{D}^4 . For the max-tb Legendrian Hopf link, these two planes should be taken to be Lagrangian.

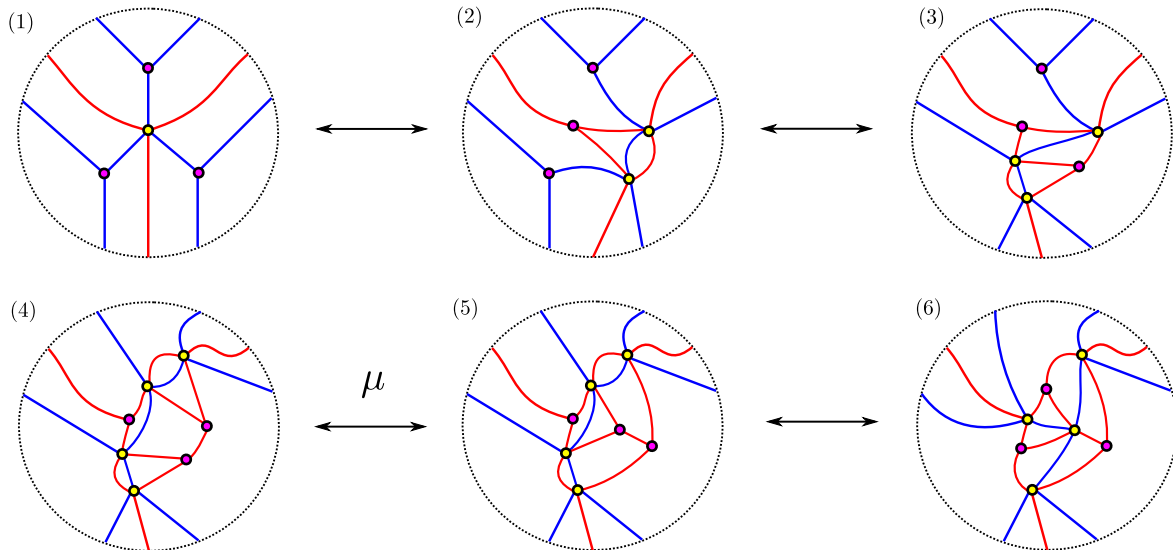


FIGURE 62. The Legendrian mutation for 3-graphs as a sequence of Legendrian isotopies and 2-graph mutation.




FIGURE 63. The 2-graph mutation with the additional information of the 1-cycles, before and after the 2-graph mutation (Left). The Y-cycle and an incident 1-cycle transforming before and after a mutation along the Y-cycle (Right).

Corollary 4.22. *Let $[\gamma] \in H_1(\Lambda(G), \mathbb{Z})$ be represented by a long edge in an N -graph G , as shown in the first row of Figure 64. Then the Legendrian mutation $\mu_\gamma(\Lambda(G))$ is the Legendrian weave associated to the graph $\mu_\gamma(G)$ as depicted in the second row of Figure 64.*




FIGURE 64. The two cases, left and right, of a Legendrian mutation along a 1-cycle γ represented by a long-edge.

Theorem 4.21 and Corollary 4.22 describe mutations along Y-cycles and l-cycles, either monochromatic or long edges. In general, we might be interested in mutating along a cycle γ which is a tree, both with Y-pieces and l-piece, as introduced in Section 2.4. Thus, we now

develop *local* rules for Legendrian mutations that will allow us to mutate along any such cycle γ . These rules also imply Corollary 4.22. All these rules are obtained and proven in the same manner: one simplifies the weave with equivalence moves – using Section 4 – until the cycle to be mutated becomes a *short* l-cycle. Then we apply the short l-cycle mutation in Figure 60 (left) and rearrange the weave with moves to the required configuration. For instance, for Corollary 4.22 (and so Figure 64), we push-through the left-most trivalent vertex through all the hexagonal vertices until the long l-cycle becomes a short l-cycle. Then we mutate at the short cycle, and push-through one of the trivalent vertices back to the left.

4.9. Diagrammatic Rules for N -graph Mutations. Let γ be a 1-cycle in an N -graph, given by a tree with Y-pieces and l-pieces. In this subsection we gather the necessary rules for performing a general mutation along γ and also diagrammatically carrying a 1-cycle after the mutation at γ . The rules are *local*, either near a hexagonal vertex or a trivalent vertex, and there are three cases that we need to draw: Legendrian mutation being performed at a Y-piece, at a l-piece, and mutation near a trivalent vertex.

First, we draw the rules for the effect of mutating at a cycle which contains Y-pieces:

- (i) Figure 65 shows how the Y-cycle at which we mutate transforms, this cycle is depicted in green. Note that the resulting cycle locally contains only one Y-piece.
- (ii) Figure 66 explains how to transform the *other* Y-cycle, in ochre (a darker yellow), under mutation at the green Y-cycle in Figure 65.
- (iii) Figure 67 then depicts the transformation of edge l-cycles through a hexagonal vertex under mutation at the green Y-cycle in Figure 65.
- (iv) Finally, Figure 68 provides the last information needed for carrying any cycle upon mutating at the green Y-cycle in Figure 65. These are the three ways in which a 1-cycle must be continued if the 1-cycle is coming from the extremes of one of the sides.

Second, the rules for mutating at a long edge of an l-piece of a 1-cycle:

- (v) Figure 69 shows how to transform an l-piece upon mutation at the green l-piece.
- (vi) Figure 70 then depicts the transformation of a Y-piece of a cycle, in ochre, upon mutation at the green l-piece in Figure 69.

Finally, the local rules for mutating near a trivalent vertex are shown in Figure 71. These rules are derived by performing Legendrian Reidemeister moves, especially Move II, until the given cycle at which we want to mutate becomes a monochromatic (short) edge. Then a monochromatic edge mutation is performed, as in Theorem 4.21, and Legendrian Reidemeister moves are performed back to the starting configuration. The two non-canceling applications of a push-through move, before and after a monochromatic edge mutation, are responsible for the tripling behavior seen in the diagrams.




FIGURE 65. Case Mutation at Y-cycle: Internal Mutation along Y-piece in green.




FIGURE 66. Case Mutation at Y-cycle in Figure 65: Effect for ochre Y-cycle of Internal Mutation along Y-cycle in green in Figure 65.

4.10. Sufficiency For Stabilized Legendrians. Finally, we conclude this section by introducing the following combinatorial idea, motivated by the topology of Legendrian surfaces in 5-dimensional contact manifolds.

Definition 4.23. An N -graph $G \subseteq C$ is said to have a *bridge* if there exists two disjoint 2-disks $D_1, D_2 \subseteq C$ such that the complement $G \setminus (G \cap D_1 \cup G \cap D_2)$ consists of $(N - 1)$ disjoint strands with labels $\tau_1, \tau_2, \dots, \tau_{N-1}$ consecutive with respect to a transverse oriented curve in $C \setminus (D_1 \cup D_2)$. \square

For the $N = 2$ case, where G is a trivalent graph, a bridge for G according to Definition 4.23 coincides with the standard graph-theoretic notion of a bridge [BM08, Die17]. A general N -graph $G \subseteq C$ with a bridge is depicted in Figure 72 (left), and an example of a 4-graph with a bridge is shown in Figure 72 (right).

The geometric motivation for this definition is based on the theory of loose Legendrian surfaces, also known as stabilized Legendrians [Mur12]. This class of loose Legendrians are known to satisfy an h -principle and has proven to be very useful in the study of Weinstein manifolds [CE12, CM19]. The reader is referred to [CE12, Mur12] for further details. For the present manuscript, we will assume known its definition and state the following property:

Proposition 4.24. *Let $G \subseteq C$ be an N -graph with a bridge. Then $\iota(\Lambda(G))$ is a loose Legendrian surface.*

Proof. The proof is a simple argument in the theory of spatial fronts. Indeed, consider the 1-dimensional front slice along the dashed orange line in Figure 72. The braid shown along

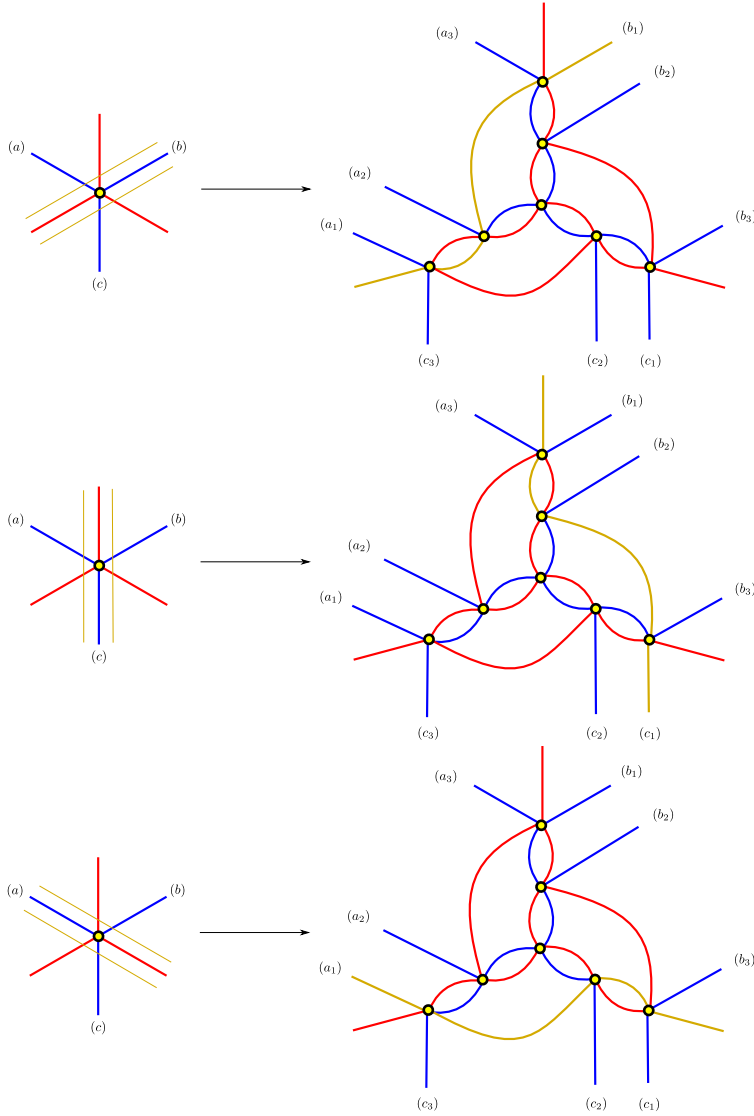


FIGURE 67. Case Mutation at Y-cycle in Figure 65: Effect for ochre l-cycle of Internal Mutation along Y-cycle in green.

this slice is depicted in Figure 73 (left). Its closure as a satellite of the standard Legendrian unknot is shown in Figure 73 (center). This Legendrian link is isotopic, via a sequence of Reidemeister II moves, to the Legendrian link given by the front in Figure 73 (right). The loose chart is exhibited in yellow in this figure. Note that this chart has arbitrarily large thickness due to the dilation freedom in (\mathbb{R}^5, ξ_{st}) and the fact that our front is global. This proves that $\iota(\Lambda(G))$ is a loose Legendrian if G has a bridge. \square

Proposition 4.24 immediately has the following consequence.

Corollary 4.25. *Let $G \subseteq C$ be an N -graph with a bridge. Then $\iota(\Lambda(G)) \subseteq (\mathbb{S}^5, \xi_{st})$ admits no exact Lagrangian filling $L \subseteq (\mathbb{D}^6, \omega_{st})$.* \square

Corollary 4.25 should be contrasted with the fact that many of the Legendrian surfaces $\iota(\Lambda(G)) \subseteq (\mathbb{S}^5, \xi_{st})$ admit exact Lagrangian fillings. For instance, it follows from Theorem 4.10 that any 2-graph G obtained from the unique two-vertex 2-graph by adding *bigons*, i.e. a 1-surgery, yields a Legendrian surface $\iota(\Lambda(G))$ which admits exact Lagrangian fillings. On




FIGURE 68. Case Mutation at Y-cycle in Figure 65: Effect for side l-cycles of Internal Mutation along Y-cycle in green.

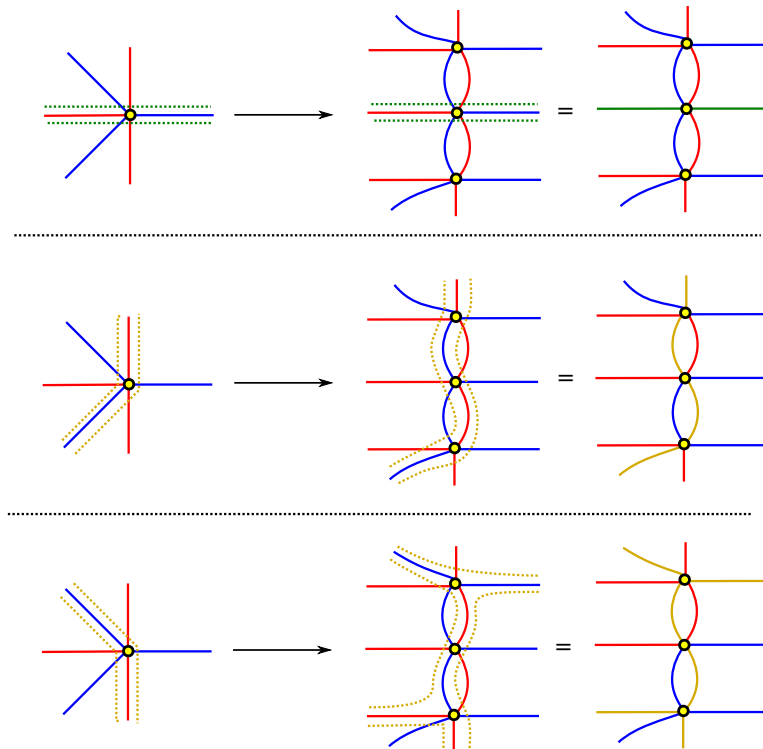


FIGURE 69. Case Mutation at l-cycle in green (upper Left). In second and third row: effect of this mutation for ochre l-cycle of Internal Mutation along l-cycle in green.




FIGURE 70. Case Mutation at horizontal l-cycle as in Figure 69: Effect for ochre Y-cycle of Internal Mutation along l-cycle in green in Figure 69 (left).




FIGURE 71. Case Mutation near trivalent vertex for green cycle (first row). Second and third rows: Effect for ochre Y-cycle of Internal Mutation at green cycle near trivalent vertex.




FIGURE 72. Structure of an N -graph with a bridge (left) and instance of a 4-graph with a bridge (right).

the other hand, *simple* 2-graphs do not – see [TZ18, Theorem 1.3].

Example 4.26 (Exact Lagrangian Cobordisms To a Loose Legendrian). Consider the Legendrian Clifford 2-torus $\mathbb{T}_c^2 \subseteq (\mathbb{S}^5, \xi_{st})$ associated, via the standard satellite, to the 2-graph in Figure 74 (Left). By applying our combinatorial Legendrian surgery from Theorem 4.10, Figure 43.(2), we obtain an (index-2) exact Lagrangian cobordism from \mathbb{T}_c^2 to the Legendrian 2-sphere Λ_l associated Figure 74 (Right). By Proposition 4.24, the Legendrian Λ_l is a loose Legendrian surface. This proves that the Legendrian Clifford 2-torus $\mathbb{T}_c^2 \subseteq (\mathbb{S}^5, \xi_{st})$ is a




FIGURE 73. The front for the Legendrian link obtained in a 3-dimensional slice of a bridge (left). The front for the corresponding satellite closure (center) and a homotopy front exhibiting a loose chart (right).

subloose Legendrian surface, and we will show in Section 6 that \mathbb{T}_c^2 is *not* a loose Legendrian. In particular, this also proves that $\mathbb{T}_c^2 \subseteq (\mathbb{S}^5, \xi_{\text{st}})$ admits no 3-dimensional exact Lagrangian fillings $L \subseteq (\mathbb{D}^6, \lambda_{\text{st}})$ in the standard symplectic 6-disk. The points in the non-empty flag moduli associated to \mathbb{T}_c^2 will in fact be geometrically represented by *non-exact* Lagrangian fillings. \square




FIGURE 74. An exact Lagrangian cobordism from a non-loose Legendrian 2-torus to a loose Legendrian 2-sphere.

5. FLAG MODULI SPACES

In this section we introduce one of the central algebraic invariants in this article, the flag¹⁶ moduli space $\mathcal{M}(G)$ of an N -graph G and its associated Legendrian weave. We will prove that these flag moduli spaces are moduli spaces of constructible sheaves associated to a Legendrian weave, but we first present their explicit and self-contained definition.

5.1. Preliminaries on the Flag Variety. Let $N \in \mathbb{N}$ be a natural number and R a commutative ground ring, which will oftentimes be a field. We denote by GL_N the general linear group, a scheme whose value over R is $GL(N, R)$, and likewise for PGL_N , the projective general linear group. By definition, a (full or complete) flag is an element

$$\mathcal{F}^\bullet \in \{\mathcal{F}^0 \subset \mathcal{F}^1 \subset \mathcal{F}^2 \subset \dots \subset \mathcal{F}^{N-1} \subset \mathcal{F}^N : \dim \mathcal{F}^i = i, 0 \leq i \leq N\},$$

i.e. a sequence of nested linear subspaces $\mathcal{F}^i \subseteq R^N = R \oplus \binom{N}{i} \oplus R$, $0 \leq i \leq N$. Let $B \subseteq GL_N$ be the Borel subgroup¹⁷ of upper triangular matrices preserving the standard coordinate flag. Since GL_N acts transitively on the set of bases, the space that parametrizes such full flags is the homogeneous space $\mathcal{B} = GL_N/B$. This is an algebraic variety, known as the flag variety.

¹⁶“Vexillary” is the appropriate adjectival form of “flag”. Hence, it should technically be named the *vexillary* moduli space. The word is possibly too obscure, and we thus favor *flag* moduli space, as in *flag* variety.

¹⁷This is a maximal Zariski closed and connected solvable algebraic subgroup. Since B is a minimal parabolic subgroup of GL_N it preserves the *most* geometric linear structure in R^N , which is precisely a flag \mathcal{F}^\bullet .

The relative position of two flags $(\mathcal{F}^\bullet, \mathcal{G}^\bullet) \in \mathcal{B} \times \mathcal{B}$ is encoded algebraically by the Bruhat decomposition

$$\mathrm{GL}_N = \bigsqcup_{w \in S_N} BwB,$$

where the symmetric group $S_N = W(\mathrm{GL}_N)$ is identified with the Weyl group. That is, the orbits of the diagonal action of GL_N on a pair of flags are indexed by the symmetric group S_N . The dimension of the Bruhat cell BwB is the length $\ell(w)$ of the permutation $w \in S_N$. By definition, \mathcal{F}^\bullet and \mathcal{G}^\bullet are in *transverse* position (or *totally transverse* or *completely transverse*) if their relative position is $w_0 \in S_N$, where w_0 denotes the longest element in the Coxeter group S_N . Note that $\ell(w_0) = \binom{N}{2}$, $w_0 \in S_N$, and that *totally transverse* is the generic relative position between two points in the flag variety \mathcal{B} . In particular, an elementary transposition $\tau_i \in S_N$ determines a relative position between two flags \mathcal{F}^\bullet and \mathcal{G}^\bullet in which only their i th vector spaces differ, and no others.

We will require a slight generalization of the above when the surface C is not simply connected: compatible local systems of flags, rather than flags of subspaces of a fixed vector space. This will not be required for our applications in Sections 6, 7 and 8, so the reader is welcome to skip this paragraph. Let $E \rightarrow X$ be a local system on a topological space X . By a local system of flags, we mean a complete filtration (flag) \mathcal{E}^\bullet of E by local systems \mathcal{E}^k such that the monodromy preserves the filtration. In this sense, the flag itself makes global sense. Let $U \subset X$ be a subspace and let \mathcal{F}^\bullet be a flag of sub-local systems on U , so that $\mathcal{F}^k \subseteq E$ for all $0 \leq k \leq N$. We say that \mathcal{F}^\bullet is *compatible* with \mathcal{E}^\bullet if the monodromies are: specifically, for $\gamma \in \pi_1(U, u)$ and $v \in \mathcal{F}^k$, $i_k(\gamma \cdot v) = i_{k,*}(\gamma) \cdot i_k(v)$, where the symbol \cdot denotes (ambiguously) the action of any group on a vector space. Note that by monodromy invariance, we may speak of the relative position of two compatible sub-local systems of flags \mathcal{F}^\bullet and \mathcal{G}^\bullet on subspaces U and U' of X .

With these algebraic preliminaries, we turn to describing the flag moduli space associated to an N -graph.

5.2. Description of the Flag Moduli Space of an N -graph. Let G be an N -graph on a connected surface C , thought of as the union of the embedded graphs G_i . By a *face* of G we mean the closure of a connected component of the complement $C \setminus G$.

We first give a general description of the flag moduli space for C not necessarily simply connected. We will not use this in our applications, so the reader is welcome to skip to the simpler Definition 5.2, which is equivalent when C is simply connected.

Let $\Sigma(G) \subset C \times \mathbb{R}$ be the wavefront of the Legendrian weave, woven according to $G \subseteq C$. Call a *region* a connected component of the complement $(C \times \mathbb{R}) \setminus \Sigma(G)$.

Definition 5.1. Let C be a connected surface and let $G \subseteq C$ be an N -graph. The *framed flag moduli space* $\widetilde{\mathcal{M}}(C, G)$ associated to G is comprised of the following data.

- i) A rank- N local system $E \rightarrow C$, equivalently a vector space V and a representation of the based fundamental group $\pi_1(C)$ on V .
- ii) For each face F of the N -graph G , a compatible local system of flags $\mathcal{F}^\bullet(F)$.
- iii) For each pair of adjacent faces F_1, F_2 , sharing an i -edge e , their two associated compatible local systems of flags $\mathcal{F}^\bullet(F_1), \mathcal{F}^\bullet(F_2)$ are in relative position $\tau_i \in S_N$, and along the common edge e we have chosen isomorphisms

$$\mathcal{F}^j(F_1) \cong \mathcal{F}^j(F_2), \quad 0 \leq j \leq N, j \neq i,$$

and no other information, as $\mathcal{F}^i(F_1) \not\cong \mathcal{F}^i(F_2)$.

- iv) By gluing, these isomorphisms define local systems in each region, since the j th step of a flag of local systems \mathcal{F}^j compatible with E defines a local system on the region between the j th and $(j + 1)$ st sheets — and these are not separated by a τ_i crossing of sheets when $j \neq i$. We require that such local systems in regions, each of which are sub-local systems of E via upward generalization morphisms, are compatible with E .¹⁸

The group PGL_N acts on the space $\widetilde{\mathcal{M}}(C, G)$ diagonally, i.e. as isomorphisms of E and on all flags of local systems at once. By definition, the *flag moduli space* of the N -graph G is the quotient stack

$$\mathcal{M}(C, G) := \widetilde{\mathcal{M}}(C, G) / \mathrm{PGL}_N.$$

We simply write $\mathcal{M}(G)$ when C is understood. □

Definition 5.2. Let C be a connected, simply connected surface and let $G \subseteq C$ be an N -graph. The *framed flag moduli space* $\widetilde{\mathcal{M}}(C, G)$ associated to G is comprised of tuples of flags, specifically:

- i) There is a flag $\mathcal{F}^\bullet(F)$ assigned to each face F of the N -graph G .
- ii) For each pair of adjacent faces $F_1, F_2 \subseteq C \setminus G$, sharing an i -edge, their two associated flags $\mathcal{F}^\bullet(F_1), \mathcal{F}^\bullet(F_2)$ are in relative position $\tau_i \in S_N$, i.e. they must satisfy

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

The group GL_N acts on the space $\widetilde{\mathcal{M}}(C, G)$ diagonally, i.e. on all flags at once. By definition, the *flag moduli space* of the N -graph G is the quotient stack

$$\mathcal{M}(C, G) := \widetilde{\mathcal{M}}(C, G) / \mathrm{PGL}_N.$$

We simply write $\mathcal{M}(G)$ when C is understood. □

We will equivalently exchange between the linear and projective perspective for a full flag. In the projective setting, flags \mathcal{F}^\bullet (or local systems of flags) are understood as a sequence of nested projective planes $\mathbb{P}(\mathcal{F})^\bullet$, given by the projectivization of the linear spaces of the linear flag \mathcal{F}^\bullet . For a ground field R , the moduli space $\mathcal{M}(C, G; R)$ is representable by an Artin stack of finite type [LO08, LO09], and is typically an algebraic variety (unless G is so symmetric that an admissible configuration of flags might be fixed by PGL_N).

In Subsection 5.3 we explain why the moduli space $\mathcal{M}(C, G; R)$ is an invariant of the Legendrian isotopy type of the associated Legendrian weave $\Lambda(G) \subseteq (J^1C, \xi_{\mathrm{st}})$. The algebraic questions we are interested in this article are about the different properties and computations of the moduli $\mathcal{M}(C, G; R)$ — for instance the cardinality of $|\mathcal{M}(C, G; \mathbb{F}_q)|$ over a finite field or how $\mathcal{M}(C, G; R)$ changes upon performing the combinatorial moves in Section 4, including Legendrian mutations and surgeries. To ease notation, we will denote flags \mathcal{F}^\bullet by \mathcal{F} .

5.3. Sheaf Description of Flag Moduli and Invariance. Let C be a smooth surface, R a commutative ring, and $\mathrm{Sh}(C \times \mathbb{R})$ the category of *constructible sheaves*, i.e. the R -linear dg-derived category of complexes of sheaves of R -modules on $C \times \mathbb{R}$ with constructible cohomology sheaves. For algebraic preliminaries on (derived) dg-categories we refer the reader to [Kel94, Tab05, Toe07, LO10], and for simplicity we will choose R a field. In this section, we use the identification $J^1(C) \cong T^{\infty, -}(C \times \mathbb{R})$ of the first jet bundle of C with downward covectors of $C \times \mathbb{R}$ — see [NRS⁺15, Section 2.1]. Now given an N -graph $G \subseteq C$, the Legendrian $\Lambda(G) \subset J^1(C) \cong T^{\infty, -}(C \times \mathbb{R}) \subset T^\infty(C \times \mathbb{R})$ can be used to define the subcategory $\mathrm{Sh}_{\Lambda(G)}(C \times \mathbb{R}) \subset \mathrm{Sh}(C \times \mathbb{R})$ whose objects are constructible sheaves whose

¹⁸This condition is not local in the N -graph, G .

singular support at contact infinity is contained in $\Lambda(G) \subset T^\infty(C \times \mathbb{R})$ — see [TZ18, Section 4].

We write $\mathcal{C}(C, G) := \text{Sh}_{\Lambda(G)}^1(C \times \mathbb{R})_0 \subset \text{Sh}_{\Lambda(G)}(C \times \mathbb{R})$ for the subcategory of microlocal rank-one sheaves which are zero in a neighborhood of $C \times \{-\infty\}$, or $\mathcal{C}(G)$ for short. This has a simple description, which we now explain. The dg-category $\text{Sh}_{\Lambda(G)}(C \times \mathbb{R})$ is itself a subcategory of sheaves constructible with respect to the stratification defined by the front projection $\Sigma(G)$, and thus has a combinatorial description. By [KS90, Theorem 8.1.11], it is equivalent to the dg-category of functors from the poset of strata to $k\text{-mod}$ (chain complexes) — see also [Nad09, Section 2.3] and [STZ17, Section 3.3]. The subcategory cut out by $\mathcal{C}(G)$ is the one whose objects are isomorphic to ones with the following properties: the chain complex assigned to a neighborhood of $C \times \{-\infty\}$ is zero; the complexes in each region of $(C \times \mathbb{R}) \setminus \Sigma(G)$ are rank-one local systems (or just vector spaces if C is simply connected); the morphisms assigned to all downward restriction maps are isomorphisms; and the upward morphisms from small open sets intersecting $\Sigma(G)$ to the regions above them which do not are codimension-one inclusions.

The combinatorial model for this description leads to the flag moduli space $\mathcal{M}(G)$ of isomorphism classes of objects in $\mathcal{C}(G)$. Indeed, the flag moduli space $\mathcal{M}(G)$ associated to an N -graph $G \subseteq C$, as introduced in Definition 5.2, relates to the category $\mathcal{C}(G) := \text{Sh}_{\Lambda(G)}^1(C \times \mathbb{R})_0$ according to the following result, which itself generalizes [TZ18, Section 4.3] to N -graphs with $N \geq 3$:

Theorem 5.3. *The flag moduli space $\mathcal{M}(C, G; \mathbb{R})$ is isomorphic to the moduli space of objects in $\mathcal{C}(G) := \text{Sh}_{\Lambda(G)}^1(C \times \mathbb{R})_0$, the subcategory of microlocal rank-1 objects in $\text{Sh}_{\Lambda(G)}(C \times \mathbb{R})$ supported away from $C \times \{-\infty\}$.*

Proof. We first assume that C is simply connected. The argument parallels that of [STZ17, Sections 6.2 and 6.3], with the additions required by the strictly two-dimensional behavior. The moduli space of objects is defined locally, meaning that it is the fiber product over its restriction-to-boundary maps of the moduli spaces $\widetilde{\mathcal{M}}$ in neighborhoods of C . We can assume that these neighborhoods of C are chosen small enough so that they are contractible and contain no more than one “feature” of the given N -graph G . That is, for some such neighborhood $U \subseteq C$, either $U \cap G$ is empty or contains part of an edge, a single trivalent vertex, or a single hexagonal vertex. We then have a local study for each of these cases.

In the case where U is empty or contains part of an edge, the front of the Legendrian weave over U is either N parallel sheets or N sheets with a single crossing labeled τ_i , and can be identified with $\sigma \times \mathbb{R}$, where σ is a front of a one-dimensional Legendrian knot being either N parallel lines or N lines with a single crossing. Then, since the \mathbb{R} factor is contractible, we can identify the moduli space $\widetilde{\mathcal{M}}$ over U using the one-dimensional study in [STZ17, Sections 6.2 and 6.3], concluding that it is either the flag variety or pairs of τ_i -transverse flags, respectively.

The moduli $\text{Sh}_{\Lambda(G)}^1(C \times \mathbb{R})$ is local with respect to $G \subseteq C$ and the topology of the surface C , i.e. it is globally described as fibered products for the local pieces of $G \subseteq C$. It therefore remains to show that $\text{Sh}_{\Lambda(G)}^1(C \times \mathbb{R})$ coincides with $\mathcal{M}(C, G)$ for the local graphs $G_{tri} \subseteq \mathbb{D}^2$ and $G_{hex} \subseteq \mathbb{D}^2$, respectively given by a trivalent vertex and a hexagonal vertex, as introduced in Section 2. We do these in turn.

The trivalent vertex case was studied in [TZ18, Section 4] for 2-graphs, and we will make the needed adjustments to N -graphs. The computation for the local N -graph G_{tri} consists of an analysis of the moduli of constructible sheaves supported at the D_4^- -wavefront singularity, as directly carried out in [TZ18]. The boundary conditions for an object in $\text{Sh}_{\Lambda(G_{tri})}^1(\mathbb{D}^2 \times \mathbb{R})$ consist of a triple of flags $(\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3)$ such that $\mathcal{F}_i \in S_{\tau_k}(\mathcal{F}_j)$ for $i \neq j$, $1 \leq i, j \leq 3$, if the edges of G_{tri} are labeled by τ_k . This can be seen by combining the result for a neighborhood

of a single crossing edge above, taking the fiber product over the spaces of flags in the empty neighborhoods in-between. Then [TZ18, Section 4.1] implies that these are all the required conditions (and strata) and thus $\mathrm{Sh}_{\Lambda(G_{tri})}^1(\mathbb{D}^2 \times \mathbb{R})$ coincides with $\mathcal{M}(\mathbb{D}^2, G_{tri})$. Note that the analysis in [TZ18, Subsection 4.1.2] restricts to the case where the local model is a 2-graph $G_{tri} \subseteq \mathbb{D}^2$, it is readily seen that this model suffices for the analysis of the local model N -graph $G_{tri} \subseteq \mathbb{D}^2$.




FIGURE 75. Constructible sheaf in \mathbb{R}^3 microlocally supported on the A_3 -swallowtail singularity (left). The sheaf convolution given by the Guillermou-Kashiwara-Schapira [GKS12] quantization upon performing a Reidemeister R1 move (right).

Alternatively, it is possible to directly conclude the analysis of the D_4^- -singularity by performing a generic perturbation of the D_4^- -wavefront, as depicted in Figure 36, and studying the category of constructible sheaves supported at a A_3 -swallowtail singularity. Indeed, Figure 75 (left) shows the conditions for a constructible sheaf microlocally supported along the front of an A_3 -swallowtail singularity, which consists of a choice of injective map $f : C_1 \rightarrow C_2$, where $C_1 \cong R^k$ and $C_2 \cong R^{k+1}$, for some $k \in \mathbb{N}$. The crucial fact is that the (stalk of the) sheaf in the remaining 3-dimensional open strata I is uniquely determined to be the cone of the map $(f, -f) : C_1 \rightarrow C_2 \oplus C_2$. This is a consequence of the Guillermou-Kashiwara-Schapira quantization [GKS12, Theorem 3.7] of Legendrian isotopies: since the A_3 -swallowtail is the big wavefront [Ad90] of the first Reidemeister move for 1-dimensional Legendrian fronts, it follows that the sheaves in the strata I are uniquely determined by $f : C_1 \rightarrow C_2$ by the sheaf kernel associated to the first Reidemeister move. It is readily seen [STZ17] that the result of the convolution with such a kernel yields the sheaf transformation in Figure 75 (right). By the non-characteristic property of the category of microlocal sheaves [GKS12], the sheaves microlocally supported on the wavefront of the D_4^- -singularity is equivalent to that for a generic perturbation of such D_4^- -singularity. The generic perturbation consists of three A_3 -swallowtails and the conditions for the constructible sheaves on these stratification follow from the above analysis. In conclusion, we obtain an isomorphism $\mathrm{Sh}_{\Lambda(G_{tri})}^1(\mathbb{D}^2 \times \mathbb{R}) \cong \mathcal{M}(\mathbb{D}^2, G_{tri})$.

Let us now address the hexagonal vertex G_{hex} . Since the Legendrian weave $\Sigma(G_{hex})$ is the big wavefront of the third Reidemeister move for 1-dimensional Legendrian fronts, it suffices to understand the kernel of its quantization. Figure 76 shows the local transformation for constructible sheaves near the third Reidemeister move [STZ17, Section 4.4.3].

In Figure 76, the C_i , $1 \leq i \leq 5$ and E_1, E_2 are complexes of vector spaces, which we can actually assume to be vector spaces [STZ17, Section 3.3]. If $C_1 \cong R^k$, for some $k \in \mathbb{N}$,




FIGURE 76. The explicit flag exchange given by the Guillermou-Kashiwara-Schapira [GKS12] quantization upon performing a Reidemeister R3 move.

the microlocal rank 1 condition implies that $E_1, E_2, C_2, C_3 \cong R^{k+1}$, $C_4, C_5 \cong R^{k+2}$ and $C_6 \cong R^{k+3}$. The four flags at one of the sides of the hexagonal vertex are

$$\begin{aligned} \mathcal{F}_1^{(1)} &= C_1 \longrightarrow C_3 \longrightarrow C_5 \longrightarrow C_6, & \mathcal{F}_2^{(1)} &= C_1 \longrightarrow E_1 \longrightarrow C_5 \longrightarrow C_6, \\ \mathcal{F}_3^{(1)} &= C_1 \longrightarrow E_1 \longrightarrow C_4 \longrightarrow C_6, & \mathcal{F}_4^{(1)} &= C_1 \longrightarrow C_2 \longrightarrow C_4 \longrightarrow C_6, \end{aligned}$$

and the four flags at the other side of the hexagonal vertex are

$$\begin{aligned} \mathcal{F}_1^{(2)} &= C_1 \longrightarrow C_3 \longrightarrow C_5 \longrightarrow C_6, & \mathcal{F}_2^{(2)} &= C_1 \longrightarrow C_3 \longrightarrow E_2 \longrightarrow C_6, \\ \mathcal{F}_3^{(2)} &= C_1 \longrightarrow C_2 \longrightarrow E_2 \longrightarrow C_6, & \mathcal{F}_4^{(2)} &= C_1 \longrightarrow C_2 \longrightarrow C_4 \longrightarrow C_6. \end{aligned}$$

The three crossings in Figure 76 (left) imply, from left to right, that

$$\mathcal{F}_1^{(1)} \in S_{\tau_{k+1}}(\mathcal{F}_2^{(1)}), \mathcal{F}_2^{(1)} \in S_{\tau_{k+2}}(\mathcal{F}_3^{(1)}), \mathcal{F}_3^{(1)} \in S_{\tau_{k+1}}(\mathcal{F}_4^{(1)}).$$

Similarly, the three crossings in Figure 76 (right) imply, from left to right, that

$$\mathcal{F}_1^{(2)} \in S_{\tau_{k+2}}(\mathcal{F}_2^{(2)}), \mathcal{F}_2^{(2)} \in S_{\tau_{k+1}}(\mathcal{F}_3^{(2)}), \mathcal{F}_3^{(2)} \in S_{\tau_{k+2}}(\mathcal{F}_4^{(2)}).$$

These are precisely the conditions for the flag moduli space $\mathcal{M}(\mathbb{D}^2, G_{hex})$ in Definition 5.2, and hence $\text{Sh}_{\Lambda(G_{hex})}^1(\mathbb{D}^2 \times \mathbb{R}) \cong \mathcal{M}(\mathbb{D}^2, G_{hex})$.

This concludes the argument for the case where C is simply connected. We now turn to the case where C is not simply connected. There are no further local conditions. The only additional concerns regard compatibilities of local systems.

Let $\Sigma(G) \subset C \times \mathbb{R}$ be the wavefront of the Legendrian weave, and recall that we call a *region* a connected component of the complement $(C \times \mathbb{R}) \setminus \Sigma(G)$. A constructible sheaf in $\mathcal{C}(G)$ restricts to a local system on each region, since there is no singular support away from the wavefront. There are two distinguished regions R_{top} and R_{bot} containing neighborhoods of $C \times \{\infty\}$ and $C \times \{-\infty\}$, respectively. A constructible sheaf in $\mathcal{C}(G)$ restricts to 0 in R_{bot} (by definition) and to a local system on $R_{\text{top}} \sim C$ that we assign to be the data E from Definition 5.1(i). Now, as explained in Definition 5.1(iv), the data of a point in $\widetilde{\mathcal{M}}(G)$ defines a local system in each region. Commutativity of sheaf restriction maps requires that a section which is parallel transported around a region and then included into E arrives at the same place as a section which is included first and then parallel transported around R_{top} , and this is the requirement of Definition 5.1(iv). \square

5.4. Local Flag Moduli Computations. Let us prove the following useful lemmas on the flag moduli, which can be implicitly used when performing computations on $\mathcal{M}(C, G; R)$. In this section, and subsequent computations, we will consider a ground field $R = k$, with $k = \mathbb{C}$ and finite fields $k = \mathbb{F}_q$ as the main fields of interest.

We start with the study of the flag moduli space at a trivalent vertex, as depicted in the left of Figure 8, and characterize that local flag moduli space.

Lemma 5.4. *Consider the neighborhood $\mathcal{O}p_3(N)$ of a τ_i -trivalent vertex in an N -graph. Then the local moduli of flags $\mathcal{M}(\mathcal{O}p_3, G; k)$ is set-theoretically a point, and the PGL_N -action on $\widetilde{\mathcal{M}}(\mathcal{O}p_3, G; k)$ has stabilizer $(k^*)^{N-2} \times k^{\binom{N}{2}-1}$.*

Proof. Let $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3$ be the three flags in $\mathcal{O}p_3(N)$. The GL_N -action is transitive on the space of flags, and thus \mathcal{F}_1 can be mapped to the standard flag \mathcal{S}_1 , defined by

$$\mathcal{S}_1^j = \{x_{j+1} = x_{j+2} = \dots = x_{N-1} = x_N = 0\}, \text{ where } k = \mathrm{Spec} k[x_1, \dots, x_N].$$

The GL_N -action allows us to also map the two flags \mathcal{F}_2 and \mathcal{F}_3 , respectively, to \mathcal{S}_2 and \mathcal{S}_3 , defined by

$$\begin{aligned} \mathcal{S}_2^j &= \mathcal{S}_3^j = \mathcal{S}_1^j, \quad 0 \leq j \leq N, \quad j \neq i, \\ \mathcal{S}_2^j &= \{x_i = x_{i+2} = \dots = x_{N-1} = x_N = 0\}, \\ \mathcal{S}_3^j &= \{x_i - x_{i+1} = x_{i+2} = \dots = x_{N-1} = x_N = 0\}. \end{aligned}$$

This implies that the quotient of the moduli $\widetilde{\mathcal{M}}(\mathcal{O}p_3, G; R)$ by the gauge group PGL_N is set-theoretically a point. In order to recover its structure as a quotient stack, it suffices to identify the stabilizer of the triple of flags $\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3$. For that, notice that the stabilizer of \mathcal{S}_1 is the projectivization of the Borel subgroup of upper triangular matrices, isomorphic to $(k^*)^{N-1} \times k^{\binom{N}{2}}$. The condition of fixing the flag \mathcal{S}_2 transversely cuts out a k -coordinate in the interior of the upper triangle, since it sets the $(i, i+1)$ entry equals to zero. This cuts the stabilizer down to $(k^*)^{N-1} \times k^{\binom{N}{2}-1}$, and finally stabilizing \mathcal{S}_3 imposes the equality of the two diagonal entries (i, i) and $(i+1, i+1)$, thus transversely cutting down a k^* . The resulting stabilizer is $(k^*)^{N-2} \times k^{\binom{N}{2}-1}$, as claimed. \square

In its simplest instance of $N = 2$, this is the statement that three distinct points in the projective line $\mathbb{P}^1(k)$ can be sent to $\{0, 1, \infty\}$ with trivial stabilizer. A lesson from Lemma 5.4 is that for any N , near at least one trivalent vertex of an N -graph, we are allowed to use the gauge group PGL_N and fix the flags around that vertex. The (proof of the) lemma also provides the (geometric) degrees of freedom left after this choice.

Example 5.5. *Consider the 2-graph G associated to the triangulation of $C = S^2$ with two triangles. This 2-graph G , dual to the triangulation, has two vertices, three edges and three faces. Then the flag moduli space $\mathcal{M}(S^2, G; \mathbb{C})$ consists of a point $\{*\}$. In fact, this point $\{*\}$ of the flag moduli space geometrically corresponds to the conjecturally unique Lagrangian 3-disk filling of the standard Legendrian unknot $\Lambda_0 \subseteq (S^5, \xi_{st})$. \square*

Lemma 5.4 is a statement about a particular triple of flags. It ought to be noted that a *generic* triple of flags is part of a moduli space of dimension $\binom{N-1}{2}$, with birational coordinates given by generalized triple ratios – see Section 7 and [FG06b, Section 9]. The flags appearing in the context of our N -graphs are in general a combination of non-generic flags, arising from the local vertices, with a flag being modified at exactly one degree when crossing an edge.

Let us now address our second local model at a vertex, that of a hexagonal vertex, as depicted in the right of Figure 8.

Lemma 5.6. *Consider the neighborhood $\mathcal{O}p_6(N)$ of a hexagonal vertex, with edges τ_i, τ_{i+1} and consecutively ordered flags $\mathcal{F}_j, j \in \mathbb{Z}/6\mathbb{Z}$. Then any pair of opposite flags $\mathcal{F}_k, \mathcal{F}_{k+3}$ determines the others.*

Proof. By symmetry, it suffices to show that the flags $\mathcal{F}_1, \mathcal{F}_4$ determine \mathcal{F}_5 and \mathcal{F}_6 . We assume that \mathcal{F}_4 and \mathcal{F}_5 are separated by a τ_{i+1} edge — a similar argument will work if it is of type τ_i . By the prescribed transversality, we have $\mathcal{F}_6^j = \mathcal{F}_1^j$ and $\mathcal{F}_5^j = \mathcal{F}_4^j$ for $j \neq i+1$. Now since $\mathcal{F}_5^i \neq \mathcal{F}_6^i$, and $\mathcal{F}_1^{i+2} = \mathcal{F}_4^{i+2}$, there exists a unique linear subspace $V \subseteq \mathcal{F}_1^{i+2}$

which contains $\mathcal{F}_5^i, \mathcal{F}_6^i$. So we must have $\mathcal{F}_5^{i+1} = \mathcal{F}_6^{i+1} = V$, uniquely determining the flags \mathcal{F}_5 and \mathcal{F}_6 . \square

A direct application of Lemma 5.6 is the invariance of the moduli of flags under the N -graph Reidemeister Move I from our Theorem 4.2 above:

Corollary 5.7. *The flag moduli space $\mathcal{M}(C, G; R)$ is invariant under the candy twist.* \square

The candy twist – Move I – is the move depicted in Figure 19 above, and the proof of Corollary 5.7 follows immediately from Lemma 5.6, since the interior faces of the local model are uniquely determined by two opposing boundary flags, and they in turn determine the remaining ones. Corollary 5.7 also follows from Theorem 1.1 and the Legendrian invariance proven in [GKS12, Theorem 3.7]. The invariance of the moduli of flags under the other moves in Theorem 4.2 can be proven similarly by direct means.

Lemma 5.6 discusses the flags in a neighborhood of a hexagonal vertex and allows for a computation of the local flag moduli space $\mathcal{M}(\mathcal{O}p_6(N); R)$ at a hexagonal vertex, since it reduces it to the study of a quadruple of flags.

Example 5.8. *Let us illustrate this point by computing $\mathcal{M}(\mathcal{O}p_6(3); \mathbb{C})$, which we claim is isomorphic to a point stabilized by the subgroup $(\mathbb{C}^*)^2 \subseteq \mathrm{PGL}(3, \mathbb{C})$. Indeed, the incidence problem at a hexagonal vertex is given by six flags*

$$\begin{aligned} \mathcal{F}_1 &= (p_1, l_1), & \mathcal{F}_2 &= (p_2, l_1), & \mathcal{F}_3 &= (p_2, l_2), \\ \mathcal{F}_4 &= (p_3, l_2), & \mathcal{F}_5 &= (p_3, l_3), & \mathcal{F}_6 &= (p_1, l_3) \end{aligned}$$

where p_i and l_i , for $1 \leq i \leq 4$, are points and lines in $\mathbb{P}^2(\mathbb{C})$ and the notation (p_i, l_i) stands for the projectivized flag $p_i \in l_i$. Since the three points p_1, p_2, p_3 are pairwise distinct, $\mathrm{PGL}(3, \mathbb{C})$ acts on them transitively, and their stabilizer is the (projectivization) of a maximal torus in $\mathrm{GL}(3, \mathbb{C})$, which is isomorphic to $(\mathbb{C}^*)^2$. Lemma 5.6 provides a more direct route: it suffices to observe that the $\mathrm{PGL}(3, \mathbb{C})$ -stabilizer of the two completely transverse flags $\mathcal{F}_1, \mathcal{F}_4$ is the set of diagonal matrices in $\mathrm{PGL}(3, \mathbb{C})$, i.e. $(\mathbb{C}^*)^2$. \square

It is an exercise to extend the argument for Lemma 5.4 above in this context and show that:

Lemma 5.9. *Consider the neighborhood $\mathcal{O}p_6(N)$ of a (τ_i, τ_{i+1}) -hexagonal vertex in an N -graph. Then the local moduli of flags $\mathcal{M}(\mathcal{O}p_6, G; k)$ is set-theoretically a point, and the PGL_N -action on $\widetilde{\mathcal{M}}(\mathcal{O}p_6, G; k)$ has stabilizer $(k^*)^2 \times \left((k^*)^{N-3} \times k^{\binom{N}{2}-3} \right)$. \square*

Having computed the local models at trivalent and hexagonal vertices, in Lemmas 5.4 and 5.9, we now address the local flag moduli space around a τ_i -edge connecting two trivalent vertices for $1 \leq i \leq N - 1$. Thanks to our discussion in Subsection 2.4 on the homology of the associated Legendrian weaves, we know that this is the flag moduli space associated to a Legendrian cylinder. In contrast to Lemmas 5.4 and 5.9 above, we will now discover that the local flag moduli space around a monochromatic edge is (set-theoretically) non-trivial.

Lemma 5.10 (Flag Cross-ratio). *Let G be an N -graph, and $e \in G$ a monochromatic edge between two trivalent vertices. The local flag moduli space $\mathcal{M}(\mathcal{O}p(e), G; k)$ in a neighborhood $\mathcal{O}p(e)$ is isomorphic to k^* with stabilizer $(k^*)^{N-2} \times k^{\binom{N}{2}-1}$, under the PGL_N -action.*

Lemma 5.10 appears in the study of cluster coordinates for 2-graphs in the works [FG06b, TZ18], yet a treatment of it here, in the context of N -graphs, seems in order. The interesting part in Lemma 5.10 is the existence of a non-trivial flag moduli space around the edge $e \in G$. The stabilizer only appears due to the dependence on N . Note also that, by using Lemma 5.6, the statement in Lemma 5.10 can readily be generalized for a long edge e , i.e. an l -cycle between two trivalent vertices, as described in Section 2.

Proof. For an edge $e \in G$ between two trivalent vertices, it suffices to discuss the case of a monochromatic edge, since the push-through move preserves the flag moduli. In this case, let v_1, v_2 be the two endpoints of e . By Lemma 5.4, the local flag moduli space around v_1 can be fixed to be a point with stabilizer $(k^*)^{N-2} \times k^{\binom{N}{2}-1}$. In this normalization, the flag moduli space around v_2 is determined in two of the sectors, and thus it is uniquely described by the remaining choice of flag. This is tantamount to the choice of a fourth point in $\mathbb{P}^1(k) \setminus \{0, 1, \infty\}$, which yields a modulus of k^* . \square

In general, the existence of a non-trivial 1-cycle $\gamma \in H_1(\Lambda(G), \mathbb{Z})$ provides the flag moduli space with a k^* factor, which can be geometrically interpreted as being a contribution of the microlocal monodromy of the associated local system induced in the Legendrian surface $\Lambda(G)$, as we explain in Section 7. The following example illustrates this point in the case of a Y-cycle in G .

Example 5.11. *Let us compute the local flag moduli space in an $N = 3$ neighborhood of a Y-cycle, as depicted in Figure 11 (Right). The configurations of points for this incidence problem are given by the following conditions:*

- (a) *Three distinct points p_0, p_1, p_2 , and three points $q_i \in l_i = \langle p_i, p_{i+1} \rangle$, where the index $0 \leq i \leq 2$ is understood modulo 3,*
- (b) *The triples $\{p_i, p_{i+1}, q_i\}$, $0 \leq i \leq 2$, are triples of distinct points.*

The action of PGL_3 allows us to set $p_0 = [1 : 0 : 0]$, $p_1 = [0 : 1 : 0]$ and $p_2 = [0 : 0 : 1]$ with a $(k^)^2$ Cartan stabilizer, and this stabilizer can then be used to fix $q_0 = [1 : 1 : 0] \in l_0$ and $q_1 = [0 : 1 : 1]$. The remaining choice of q_2 yields the k^* contribution to the flag moduli space since it is a choice of a point $q_2 \in l_2$ distinct from p_1, p_2 . \square*

This concludes our local computations of flag moduli spaces $\mathcal{M}(C, G)$. We now study the behavior of the invariant $\mathcal{M}(C, G)$ under Legendrian surgery, and Sections 6 and 8 will develop global computation of flag moduli spaces. Given an N -graph $G \subseteq C$, we ease notation by writing $\mathcal{M}(G)$ for $\mathcal{M}(C, G)$.

5.5. Flag Moduli under Legendrian Surgeries. Let G, G' be N -graphs such that G' is obtained by Legendrian surgery on G , as described in Theorem 4.10. The following result relates the flag moduli spaces $\mathcal{M}(G)$ and $\mathcal{M}(G')$ before and after Legendrian surgery.

Theorem 5.12. *Let k be a field and G an N -graph. For any τ_i -edges of G , the flag moduli space $\mathcal{M}(G)$ satisfies the following local relations:*

$$\begin{aligned}
 (i) \quad & \mathcal{M}\left(\text{Diagram 1}\right) = k^* \times \mathcal{M}\left(\text{Diagram 2}\right) \\
 (ii) \quad & \mathcal{M}\left(\text{Diagram 3}\right) = (k \setminus \{0, 1\}) \times \mathcal{M}\left(\text{Diagram 4}\right) \\
 (iii) \quad & \mathcal{M}\left(\text{Diagram 5}\right) = \emptyset
 \end{aligned}$$

FIGURE 77. The change of the flag moduli spaces $\mathcal{M}(G)$ under combinatorial changes in a piece of an N -graph G .

Proof. The relations can be verified with our description of the flag moduli space in Subsection 5.2. We can also argue directly thanks to the geometry developed in Section 4. Indeed, the moduli of objects of the category of constructible sheaves microlocally supported at a Legendrian connected sum $\Lambda_1 \# \Lambda_2$ is a direct product of the moduli of objects microlocally supported at Λ_1 and those microlocally supported at Λ_2 . By Theorem 4.10, the right and left graphs G_r, G_l for the Relations (i) and (ii) geometrically correspond to Legendrian connected sums with the standard Legendrian 2-torus \mathbb{T}_{st}^2 , and the Legendrian Clifford 2-torus \mathbb{T}_c^2 , respectively. The flag moduli for the former is k^* , and for the latter it is $k \setminus \{0, 1\}$, which concludes (i), (ii). Finally, the relation (iii) follows from Proposition 4.24, as there do not exist constructible sheaves microlocally supported at a loose Legendrian. \square

Note that, by construction, there exists a 3-dimensional exact Lagrangian cobordism $L(G, G')$ from $\Lambda(G)$ to $\Lambda(G')$, in the symplectization of (J^1C, ξ_{st}) . Thus, from the standard results in Floer theory [EES05b, EGH00], we expect¹⁹ a map from $\mathcal{M}(G) \times H_1(L(G, G'), k)$ to $\mathcal{M}(G')$. Theorem 5.12 gives a strong indication of what these maps should be, i.e. for (i), (ii), $\mathcal{M}(G')$ is a k^* - or a $(k \setminus \{0, 1\})$ -bundle over $\mathcal{M}(G)$, with the map being a section for this bundle projection.

5.6. Non-characteristic Property of Stabilization. We conclude Section 5 with an interesting and direct computation of flag moduli spaces. First, note that the proof of Theorem 4.17, showing that the standard satellites of $\Lambda(G)$ and $\Lambda(s(G))$ are Legendrian isotopic, and Theorem 5.3 imply the isomorphism

$$\mathcal{M}(\Lambda(G)) \cong \mathcal{M}(\Lambda(s(G))),$$

where $s(G)$ is the stabilization we introduced in Subsection 4.7. We will nevertheless provide a self-contained sheaf-theoretical proof of that equivalence, which we now illustrate in the case $N = 2$.

Proof of flag moduli space equivalence $N = 2$. In that case, the moduli of objects in the category $\mathcal{M}(\Lambda(s(G)))$ parametrizes flags in \mathbb{P}^2 up to $\text{PGL}(3, \mathbb{C})$ equivalence abiding the constraints imposed by the 3-graph on the left of Figure 78. We assume that the 2-graph $G = G^{1,2}$, before stabilizing, contains at least a vertex.




FIGURE 78. The flag configuration for $N = 2$ stabilization.

The graph $G^{(1,2)}$ imposes constraints on the points lying in a line $l_1 \subseteq \mathbb{P}^2$, the ladybug changes this line to distinct lines l_2, l_3 , also different from l_1 , and the descending (12)-halo provides the freedom of a point $p \in l_3$. The fact that $G^{(1,2)}$ is contained in a wing of the

¹⁹To our knowledge, these maps have yet to be studied in the context of microlocal sheaf theory. The expectation that they exist comes from the fact that the flag moduli space $\mathcal{M}(G)$ should correspond to an augmentation variety for $\Lambda(G)$, and these maps are known to exist between augmentation varieties.

ladybug implies that $l_1 \cap l_2 \cap l_3$ is a point, which for now we denote ∞ . Let us show that this moduli space coincides with the moduli space of points in l_1 imposed by $G^{(1,2)}$. For that, note that the stabilizer of three non-collinear points $p_1, p_2, p_3 \in \mathbb{P}^2$ is isomorphic to $\mathbb{C}^* \times \mathbb{C}^*$; indeed, it is isomorphic to the space of invertible diagonal matrices in $\mathrm{PGL}(3, \mathbb{C})$. Geometrically, each of the \mathbb{C}^* allows us to move any point in one of the three possible lines spanned by two of the three points $\{p_1, p_2, p_3\}$ around that line, on the complement of these two spanning points.

Hence we can start by using the $\mathrm{PGL}(3, \mathbb{C})$ and fix the points $1, p, \infty \in \mathbb{P}^2$ in the configuration shown in the right of Figure 78, which determine the lines l_1, l_3 . From the $\mathbb{C}^* \times \mathbb{C}^*$ we can use the first \mathbb{C}^* in order to send the third point in l_1 imposed by $G^{(12)}$ to $0 \in l_1$, and the second \mathbb{C}^* to choose a point in the line $l = \langle 1, p \rangle$, which in turn determines a line $l_2 \subseteq \mathbb{P}^2$ by taking its span with $\infty \in l_1 \cap l_3$. This fixes the configuration of lines l_1, l_2, l_3 and the points $0, 1, \infty, p \subseteq \mathbb{P}^2$ with $\{0, 1, \infty\} \subseteq l_1$, and that is precisely the three points being fixed by the $\mathrm{PGL}(2, \mathbb{C})$ symmetry acting in $G^{(12)}$. \square

This argument is self-contained, yet hopefully illustrates how in general the *geometric* conclusion from Theorem 4.17, and the invariance of the flag moduli space \mathcal{M} under Legendrian isotopy, are stronger and neater tools than the strict algebraic invariance of the flag moduli space. Let us now move forward with the following Sections 6, 7 and 8, which display several applications of the techniques developed in Sections 2, 3, 4 and 5, and in particular prove Theorems 1.5, 1.6, 1.8 and 1.9 stated in the introduction.

6. APPLICATIONS AND VEXILLARY COMPUTATIONS

In this section we study applications of our diagrammatic calculus for Legendrian weaves $\Lambda(G)$ associated to an N -graph G , and their flag moduli spaces $\mathcal{M}(G)$. In particular, we will prove Theorem 1.5 and Theorem 1.6.

6.1. First Pair of Computations. Let us start with two simple examples of Legendrian weaves and their flag moduli: the Legendrian Clifford torus and the double $t_4 \cup t_4 \subseteq \mathbb{S}^2$ of the 4-triangle t_4 in the 2-sphere \mathbb{S}^2 .

6.1.1. The Legendrian Clifford Torus. Let us consider the 2-graph $G = (\partial\Delta^3)^{(1)} \subseteq \mathbb{S}^2$ in Figure 79, which has already featured in the proof of Theorem 4.10. The flag moduli space $\mathcal{M}(G)$ is readily seen to be the pair of pants $\mathbb{P}^1 \setminus \{0, 1, \infty\}$. Indeed, there are four contractible connected components in $\mathbb{S}^2 \setminus G$, which implies that

$$\widetilde{\mathcal{M}}(G) = \{(p_1, p_2, p_3, p_4) \in (\mathbb{P}^1)^4 : p_i \neq p_j, \quad i \neq j\}$$

where $\mathbb{P}^1 \cong \mathrm{GL}(2, \mathbb{C})/B$ is the flag variety of lines in \mathbb{C}^2 . Since $\mathrm{PGL}(2, \mathbb{C})$ acts 3-transitively on \mathbb{P}^1 , we can assume that $(p_2, p_3, p_4) = (0, 1, \infty)$, and the quotient $\widetilde{\mathcal{M}}(G)/\mathrm{PGL}(3, \mathbb{C})$ is given by

$$\mathcal{M}(G) = \{\lambda \in \mathbb{P}^1 : \lambda \neq 0, 1, \infty\}.$$

This flag moduli space is shown in Figure 79 (left), which is uniquely determined by the choice of $\lambda \in \mathbb{P}^1 \setminus \{0, 1, \infty\}$.

Let us illustrate the Legendrian geometry in this case. The Euler characteristic of the Legendrian weave $\Lambda(G)$ is $\chi(\Lambda(G)) = 2 \cdot \chi(\mathbb{S}^2) - 4 = 0$, and thus $\Lambda(G)$ is a closed 2-torus. A different front for $\Lambda(G)$ is depicted in Figure 79 (right), where the cone singularity [CM19, Section 2] is used, in line with the description in [DR11, Section 3]. The flag moduli space $\mathcal{M}(G)$ for the 2-graph G is read in this front as the moduli space of constructible sheaves in \mathbb{R}^3 microlocally supported with rank-1 in the front Figure 79 (right). This latter moduli is




FIGURE 79. The tetrahedral 2-graph G as a planar projection of the 1-skeleton $(\Delta^3)^{(1)}$ of the tetrahedron $\partial\Delta$ (left). A front projection for the Legendrian 2-torus $\iota(\Lambda(G))$ (right).

given with the data of a 1-dimensional vector space \mathbb{C} in the bounded region in the interior of the front and a linear monodromy map $\lambda : \mathbb{C} \rightarrow \mathbb{C}$. The monodromy must be an isomorphism, and thus $\lambda \in \text{GL}(1, \mathbb{C}) \cong \mathbb{C}^*$, and also satisfy the additional constraint imposed by the cone singularity. By generically perturbing this singularity, it is readily seen that the condition is that the monodromy λ does not have 1 as an eigenvalue, which in this case reduces to $\lambda \in \mathbb{C} \setminus \{0, 1\} \cong \mathbb{P}^1 \setminus \{0, 1, \infty\}$. This is precisely the flag moduli space $\mathcal{M}(G)$. \square

Remark 6.1 ([Nad17b, TZ18]). This particular wavefront allows for a direct Legendrian analysis of the Landau-Ginzburg model $(\mathbb{C}^3, z_1 z_2 z_3)$, as follows. The regular fiber $F \subseteq \mathbb{C}^3$ of the superpotential is isomorphic to $F \cong (\mathbb{C}^*)^2$, and its Lagrangian skeleton is thus an exact 2-torus $\mathbb{T}^2 \subseteq F$, i.e. the vanishing cycle for the (non-isolated) singularity W . Its Legendrian lift

$\Lambda := \{(z_1, z_2, z_3) \in \mathbb{C}^3 : |z_1| = |z_2| = |z_3| = 1/3, \arg(z_1) + \arg(z_2) + \arg(z_3) = 0\} \subseteq (\mathbb{S}^5, \xi_{\text{st}})$, has vanishing (singular) thimble the conic Lagrangian

$$L = \{(z_1, z_2, z_3) \in \mathbb{C}^3 : W(z_1, z_2, z_3) \in \mathbb{R}^+, |z_1| = |z_2| = |z_3|\}.$$

By performing a real blow-up at the origin, we introduce a real 2-sphere \mathbb{S}^2 at the origin and a projection map $\pi : \Lambda \rightarrow \mathbb{S}^2$ from our Legendrian 2-torus onto this exceptional 2-sphere \mathbb{S}^2 . In coordinates, the map $\pi(z_1, z_2, z_3) = (\Re(z_1), \Re(z_2), \Re(z_3))$ is just given by taking the real parts of the complex coordinates and realizes the Legendrian surface $\Lambda \subseteq (\mathbb{S}^5, \xi_{\text{st}})$ as the Legendrian weave $\iota(\Lambda(G))$ associated to the four-vertex 2-graph $G \subseteq \mathbb{S}^2$, given by the 1-skeleton of the tetrahedron. Thus, the mirror of the Landau-Ginzburg model $(\mathbb{C}^3, z_1 z_2 z_3)$ is the Legendrian 2-torus in $(J^1\mathbb{S}^2, \xi_{\text{st}})$ which satellites to the Clifford 2-torus $\mathbb{T}_c^2 \subseteq (\mathbb{S}^5, \xi_{\text{st}})$. This leads to the description of the A-model Landau-Ginzburg model $(\mathbb{C}^3, z_1, z_2, z_3)$, given by the category $\mu\text{Sh}_L(\mathbb{C}^3)$ of wrapped sheaves, as the bounded dg-category of finitely-generated torsion complexes on the flag moduli space $\mathcal{M}(\mathbb{T}_c^2) \cong \mathbb{P}^1 \setminus \{0, 1, \infty\}$. \square

6.1.2. *The Double of the 4-Triangle.* Let us consider the 4-graph $G(t_4)$ associated to a 4-triangle t_4 , as depicted in Figure 80 (left), and described in Section 3. Let $G = G(t_4) \cup_{\partial} G(t_4) \subseteq \mathbb{S}^2$ be the 4-graph obtained by gluing two copies of this 4-graph along their boundaries, i.e. G is the 4-graph associated to the 4-triangulation of \mathbb{S}^2 with two underlying t_1 -triangles. The 4-graph G is depicted in Figure 80 (right), where the circle at the boundary is identified to a unique point, which is a hexagonal vertex.

For the computation of the flag moduli space $\mathcal{M}(G)$, we employ our geometric techniques in Section 4. Theorem 4.10 allows us to remove the initial three (blue) τ_1 -bigons, by considering a direct sum with three copies of the standard Legendrian 2-torus \mathbb{T}_{st}^2 , see Section 4. By applying Move I in Theorem 4.2 three times, we obtain the 3-graph in Figure 81 (left). Further removing three of the bigons, we reach the 3-graph G_0 in Figure 81 (right). The framed flag moduli space $\widetilde{\mathcal{M}}(G_0)$ for the 3-graph G_0 is given by the choice of two flags




FIGURE 80. The local 4-graph $G(t_4)$ associated to a 4-triangle t_4 (left). The global 4-graph $G(\tau_4)$ given by the 4-triangulation τ_4 of the 2-sphere \mathbb{S}^2 with two triangles.

$\mathcal{F}_1 = (p_1, l_1, \pi_1), \mathcal{F}_2 = (p_2, l_2, \pi_2) \in \text{GL}_4/B$ in projective 3-space, and a choice of three points $p_3, p_4, p_5 \in \mathbb{P}_k^3$ such that

- (l_1, π_1) and (l_2, π_2) are completely transverse, i.e. $l_1 \not\subset \pi_2$ and $l_2 \not\subset \pi_1$, and $p_1 \neq p_2$,
- $p_3 \in l_1, p_3 \neq p_1$,
- $p_4 \in l_2, p_4 \neq p_2$,
- $p_5 \in \pi_1 \cap \pi_2, p_3 \neq p_1$.

In particular, $\mathcal{M}(G_0) \cong \widetilde{\mathcal{M}}(G_0)/\text{PGL}_4$, and the flag moduli space $\mathcal{M}(G_0)$ is described by the data above. By Theorem 4.10, and the fact that each bigon contributes to k^* once the Legendrian weave is connected, we deduce that our original flag moduli space must be isomorphic to $\mathcal{M}(G) \cong \mathcal{M}(G_0) \times (k^*)^4$.




FIGURE 81. The 4-graph $G(\tau_4)$ in Figure 80 after three index 1 anti-surgeries - accounted by the connected sums with \mathbb{T}_{st}^2 - and simplified with Move I (right). The 4-graph obtained by three additional index anti-surgeries (right).

This simplification, from the original 4-graph G to G_0 , allows for a direct description above of the flag moduli space $\mathcal{M}(G)$, from which further information can be readily extracted. For instance, the \mathbb{F}_q -rational count for $\mathcal{M}(G)(\mathbb{F}_q)$ is immediately:

$$|\mathcal{M}(G)(\mathbb{F}_q)| = \frac{(q-1)^5}{(q^4-1)(q^4-q)(q^4-q^2)(q^4-q^3)} \cdot q^3 \cdot \frac{(q^4-1)(q^3-1)(q^2-1)}{(q-1)^3} (q+1)q,$$

as $|\mathrm{PGL}(4, \mathbb{F}_q)| = (q^4-1)(q^4-q)(q^4-q^2)(q^4-q^3)(q-1)^{-1}$, the rightmost multiplicative factor is the count for the two flags $\mathcal{F}_1, \mathcal{F}_2$, and the q^3 factors stands for the final choice of (p_3, p_4, p_5) . \square

We conclude this initial gallery of computations with the following:

Example 6.2 (Concentric Circles). Let $\tau = (\tau_{i_1}, \tau_{i_2}, \dots, \tau_{i_n})$ be an ordered collection of n simple transpositions $\tau_{i_j} \in S_{N-1}$, $1 \leq j \leq n$, $n \in \mathbb{N}$. Consider the N -graph $G(\tau) \subseteq \mathbb{S}^2$ described by n concentric circles $C_i \subseteq \mathbb{S}^2$, $1 \leq i \leq n$, with center on the North Pole, and strictly increasing radius. This N -graph is depicted in Figure 82 (left).




FIGURE 82. The N -graph $G(\tau)$ associated to the sequence of transpositions $\tau = (\tau_{i_1}, \tau_{i_2}, \dots, \tau_{i_n})$ (left). The 4-graph $G(\tau)$ associated to $\tau = (\tau_1, \tau_2, \tau_3, \tau_1, \tau_2, \tau_1)$ (right).

The Legendrian weave $\Lambda(G(\tau)) \subseteq (J^1\mathbb{S}^2, \xi_{\mathrm{st}})$ is a radial version of the N -stranded positive braid closure of $\beta = \sigma_{i_1}\sigma_{i_2} \dots \sigma_{i_n}$. Smoothly, it is a link of N two-spheres \mathbb{S}^2 . The moduli space of rank-one sheaves in \mathbb{R}^2 supported along the positive braid β is the open Bott-Samelson variety $O(\beta)$ [STZ17, Tri19, CG20]. By Section 5.2, since $C = \mathbb{S}^2$ is simply connected, there is no further monodromy information and $\mathcal{M}(G(\tau)) = O(\beta)$. In particular, the links with different n have a different number of points over \mathbb{F}_q and cannot be Legendrian isotopic. We note further that [STZ17, Theorem 6.34] relates this number to the HOMFLY-PT polynomial of the (topological) knot in \mathbb{R}^3 defined by the braid β . \square

6.2. Symmetry groups for Legendrian weaves. Let \mathbb{G} be an arbitrary finite group and $\Lambda \subseteq (\mathbb{S}^5, \xi_{\mathrm{st}})$ a Legendrian surface, with underlying smooth surface $S(\Lambda)$. Let $\mathcal{L}(\Lambda)$ be the space of embedded Legendrian surfaces in $(\mathbb{S}^5, \xi_{\mathrm{st}})$ Legendrian isotopic to the Legendrian surface Λ , with base point Λ . In addition, let $\mathbb{L}(\Lambda)$ be the monoid of 3-dimensional exact Lagrangian concordances in the symplectization $(\mathbb{S}^5 \times \mathbb{R}(t), e^t \lambda_{\mathrm{st}})$, up to Hamiltonian isotopy, based on the Legendrian surface $\Lambda \subseteq (\mathbb{S}^5, \xi_{\mathrm{st}})$. Let $\varphi_t : S(\Lambda) \rightarrow (\mathbb{S}^5, \xi_{\mathrm{st}})$ be a \mathbb{S}^1 -family of Legendrian embeddings, $t \in \mathbb{S}^1$. Then the graph map

$$\mathrm{gr} : \pi_1(\mathcal{L}(\Lambda)) \rightarrow \mathbb{L}(\Lambda), \quad [\varphi_t] \mapsto (\varphi_t(S(\Lambda)), t),$$

allows us to relate loops of Legendrian surfaces with Lagrangian concordances.

These spaces $\mathcal{L}(\Lambda), \mathbb{L}(\Lambda)$ are challenging to study. Already in the 1-dimensional case of Legendrian links $\Lambda \subseteq (\mathbb{S}^3, \xi_{\mathrm{st}})$, it was only established recently that there exist Legendrian links such that the fundamental groups $\pi_1(\mathcal{L}(\Lambda))$ can admit (infinite order) non-Abelian

subgroups [CG20, Corollary 1.6], and $\mathbb{L}(\Lambda)$ actually contains elements of infinite order [CG20, Corollary 1.7]. To our knowledge, the only previous result about the fundamental group $\pi_1(\mathcal{L}(\Lambda))$ or the monoid $\mathbb{L}(\Lambda)$ for $\Lambda \subseteq (\mathbb{R}^5, \xi_{\text{st}})$ a Legendrian surface was proven in [SS16], where Legendrian surfaces $\Lambda_{\mathbb{Z}_n}$, $n \in \mathbb{N}$, were built such that $\pi_1(\mathcal{L}(\Lambda_{\mathbb{Z}_n}))$ admits the finite cyclic group \mathbb{Z}_n as a subgroup. Legendrian weaves and their flag moduli space are well-suited to address these questions. We present the following result for Legendrian surfaces in $(\mathbb{S}^5, \xi_{\text{st}})$:

Theorem 6.3. *Let \mathbb{G} be an arbitrary finite group. Then there exists a Legendrian surface $\Lambda_{\mathbb{G}} \subseteq (\mathbb{S}^5, \xi_{\text{st}})$ such that*

- (i) \mathbb{G} is a subquotient of the fundamental group $\pi_1(\mathcal{L}(\Lambda_{\mathbb{G}}))$,
- (ii) \mathbb{G} is a subquotient of the 3-dimensional Lagrangian concordance monoid $\mathbb{L}(\Lambda_{\mathbb{G}})$.

In fact, the latter is the image of the former via the graph map $gr : \pi_1(\mathcal{L}(\Lambda)) \rightarrow \mathbb{L}(\Lambda)$.

Proof. The argument is structured in two parts. First, we describe a construction of a 2-graph G' given a triangulation of a surface. Second, we use this construction to prove the statement. The second part has itself two steps: in the first step, the statement is proven only for those finite groups \mathbb{G} which are Hurwitz groups \mathbb{H} .²⁰ In the second step, the case of Hurwitz groups is used to conclude the statement for an arbitrary finite group.

First, we begin by describing a construction of 2-graphs. Let (C, T) be a closed smooth surface, T a triangulation with $e(T)$ edges, and $G(T)$ the trivalent 2-graph dual to the triangulation T . Consider the 2-graph G' obtained by adding a bigon at each edge of $G(T)$, using Move 4 in Figure 4. By Theorem 4.10, specifically Remark 4.11.(ii), the Legendrian $\iota(\Lambda(G'))$ is obtained by performing a connected sum of $\iota(\Lambda(G(T)))$ with $e(T)$ copies of the standard Legendrian torus $\mathbb{T}_{\text{st}}^2 \subseteq (\mathbb{S}^5, \xi_{\text{st}})$. Then [DR11, Proposition 4.6], or Theorem 5.12, implies that the complex flag moduli space $\mathcal{M}(G')$ is isomorphic to the product $\mathcal{M}(G(T)) \times (\mathbb{C}^*)^{e(T)}$, and thus $H^*(\mathcal{M}(G'), \mathbb{Q}) \cong H^*(\mathcal{M}(G(T), \mathbb{Q}) \otimes H^*((\mathbb{C}^*)^{e(T)}, \mathbb{Q}))$ by the Künneth formula.

Second, we will now prove the statement in the case that \mathbb{G} is assumed to be an arbitrary but fixed Hurwitz group \mathbb{H} . By virtue of Hurwitz' theorem [Hur92, LT99], there exists a compact Riemann surface $C = C(\mathbb{H})$ whose automorphism group is (isomorphic to) \mathbb{H} ; this surface C is called a Hurwitz surface in the literature. The topological surface underlying the Riemann surface C admits a triangulation $T(\mathbb{H})$ with symmetry group \mathbb{H} . In particular, the dual graph $G = G(T(\mathbb{H}))$ also has symmetry group \mathbb{H} . Let us now consider the 2-graph G' , associated to G as in the paragraph above, where the edge bigons are added such that \mathbb{H} is still a subgroup of the symmetry group of G' . Note that, by construction, \mathbb{H} acts faithfully on the set of edges of the triangulation T , and thus \mathbb{H} also acts faithfully on the $1 \otimes H^*((\mathbb{C}^*)^{e(T)}, \mathbb{Q}) \subseteq H^*(\mathcal{M}(G'), \mathbb{Q})$ piece of the cohomology of the flag moduli space $\mathcal{M}(G')$.

Now, the generators x, y of the triangle group $T(2, 3, 7)$ are geometrically given by rotations of the Poincaré hyperbolic disk, namely x is a rotation of angle π about the vertices of the $(2, 3, 7)$ -Schwarz triangle and y corresponds to a rotation of angle $2\pi/3$. Since a rotation ρ is smoothly isotopic to the identity, as a diffeomorphism, there exists a contact isotopy $\phi_t(\rho)$, $t \in [0, 1]$, of $(J^1 C, \xi_{\text{st}})$ such that $\phi_0(\rho) = id$ and $\phi_1(\rho)$ set-wise fixes the weave front associated to G' , and thus the Legendrian surface $\Lambda(G')$ associated to it. This contact isotopy $\phi_t(\rho)$ defines an element of $\pi_1(\mathcal{L}(\Lambda(G')))$, and its graph an element of $\mathbb{L}(\Lambda(G'))$. The flag moduli space $\mathcal{M}(G')$ is a Legendrian isotopy invariant of the Legendrian surface $\Lambda(G') \subseteq (\mathbb{S}^5, \xi_{\text{st}})$, and this contact isotopy induces an automorphism of $\mathcal{M}(G')$. In particular, there are Legendrian isotopies $\phi_t(x)$ and $\phi_t(y)$ associated to the generators x, y of any Hurwitz group \mathbb{H} , x rotating π and y rotating $2\pi/3$. Thus the subgroup $\langle \phi_t(x), \phi_t(y) \rangle \subseteq \pi_1(\mathcal{L}(\Lambda(G')))$

²⁰A Hurwitz group \mathbb{H} is any finite group which can be generated by an element x of order 2 and an element y of order 3 whose product xy has order 7. Equivalently, a Hurwitz group is any finite nontrivial quotient of the $(2, 3, 7)$ -triangle group $T(2, 3, 7) := \langle x, y \mid x^2 = y^3 = (xy)^7 = 1 \rangle$.

acts by automorphisms in $\mathcal{M}(G')$. Since \mathbb{H} acts faithfully in the cohomology $H^*(\mathcal{M}(G'))$, as pointed out above, \mathbb{H} is a subquotient of $\pi_1(\mathcal{L}(\Lambda(G')))$, namely, it is a quotient of the subgroup $\langle \phi_i(x), \phi_i(y) \rangle$. The argument for $\mathbb{L}(\Lambda(G'))$ is identical, and this concludes the required statement for Hurwitz groups \mathbb{H} .

Finally, to conclude the general statement, let \mathbb{G} be an arbitrary finite group and assume the result holds for Hurwitz groups, which is proven above. Then \mathbb{G} is a subgroup of the alternating group A_n for large enough $n \in \mathbb{N}$. By [Con84, Section 3], see also [LT99], A_n is a Hurwitz group $\mathbb{G}(C)$ for $n \geq 168$, and thus \mathbb{G} injects into such a Hurwitz group $\mathbb{G}(C)$.²¹ The argument above thus implies that \mathbb{G} is a subquotient for $\pi_1(\mathcal{L}(\Lambda(G')))$ and $\mathbb{L}(\Lambda(G'))$. Hence, the choice of weave $\Lambda(G')$ completes the proof of Theorem 6.3. \square

We do not know whether or not a result analogous to Theorem 6.3 holds for 1-dimensional Legendrian knots $\Lambda \subseteq (\mathbb{S}^3, \xi_{\text{st}})$. That could be a good question in low-dimensional contact topology. Any answer – positive or negative – would be of interest.

There is a complement to Theorem 6.3 for certain groups \mathbb{G} of infinite order, including non-Abelian groups such as $\text{PSL}(2, \mathbb{Z})$, by using results of the first author. Indeed, the Legendrian weave associated to the 4-graph $G(\tau)$ with the eighteen concentric circles


$$\tau = (\tau_1, \tau_2, \tau_1, \tau_2, \tau_1, \tau_2, \tau_1, \tau_2, \tau_1, \tau_2, \tau_1, \tau_2, \tau_1, \tau_2, \tau_1, \tau_2, \tau_1, \tau_2)$$

represents a 3-component Legendrian link $\Lambda(G(\tau))$ of 2-spheres. The geometric Br_3 -braid action constructed in [CG20], modulo its center $Z(\text{Br}_3)$, acts faithfully on the flag moduli space $\mathcal{M}(\Lambda(G(\tau)))$. This flag moduli space is described in Example 6.2. Then [CG20, Theorem 1.1] shows that the modular group $\text{PSL}(2, \mathbb{Z})$ acts faithfully on the cluster charts for the space obtained by forgetting the monodromies in the Grothendieck resolution $\mathcal{M}(\Lambda(G(\tau)))$. Hence, $\text{PSL}(2, \mathbb{Z})$ is a subquotient of $\pi_1(\mathbb{L}(\Lambda(G(\tau))))$ and $\mathcal{L}(\Lambda(G(\tau)))$ for these Legendrian weaves $\Lambda(G(\tau))$.

6.3. Flag Moduli and Bipartite Graphs. In Section 3, we introduced the construction of a 3-graph $G \subseteq C$ associated to an embedded eponymous bipartite graph G . This subsection explains how to compute flag moduli spaces for such 3-graphs.

We will employ a useful notation, local to this subsection. If $a, b \in V^3$ are distinct vectors in a 3-dimensional vector space V , we denote by ab the unique 2-plane spanned by a, b . Similarly, given two 2-planes $A, B, \subseteq V^3$, the intersection $A \cap B$ will be denoted by AB .

At a hexagonal vertex, traveling between opposite faces requires crossing three edges of alternating colors, and thus opposite faces are assigned completely transverse flags $\mathcal{A} = (a, A) = aA$ and $\mathcal{B} = (b, B) = bB$. Note that a single such pair \mathcal{A}, \mathcal{B} determines the remaining four regions, by Lemma 5.6: if crossing *red, blue, red* from \mathcal{A} to \mathcal{B} , the flags in succession are $(a, A), (AB, A), (AB, B), (b, B)$. If crossing *blue, red, blue*, the flags are $(a, A), (a, ab), (b, ab), (b, B)$. This is depicted as follows:



Now consider an edge of the bicubic graph G . In the associated 3-graph, this edge generates two hexagonal vertices which are connected by two adjacent edges of different colors. This local configuration is said to be a hexagonal edge. Let us denote the two flags on opposite

²¹Note that A_n , for $n \leq 168$, is a subgroup of A_m , for a greater $m \geq n$, and thus all cases A_n are covered.

regions along the axis connecting the hexagonal vertices by $\mathcal{A} = aA$ and $\mathcal{C} = cC$. Let \mathcal{B} be the flag in the interior region of the hexagonal edge, transverse to both \mathcal{A} and \mathcal{C} . There are two further conditions on the flag \mathcal{B} :

$$AB \subset C, \quad c \subset ab.$$

The Weyl group $W(A_2) \cong S_3$ is the symmetric group on three elements, and thus there are six possible relative positions for the two flags $\mathcal{A}, \mathcal{C} \in \text{GL}_3/B$. Here we consider the case of a finite field $k = \mathbb{F}_q$. In a hexagonal edge, the relative position of the two outer flags \mathcal{A}, \mathcal{C} is restricted:

Lemma 6.4. *The two outer flags \mathcal{A}, \mathcal{C} in a hexagonal edge must coincide or be completely transverse. In addition, with \mathcal{A}, \mathcal{C} fixed, number the of choices of flag \mathcal{B} in the interior of the hexagonal edge is q^3 , in the case $\mathcal{A} = \mathcal{C}$, and $q - 1$, in the case $\mathcal{A} \neq \mathcal{C}$.*

Proof. Let us analyze their possible relative positions, labeled according to the elements $W(A_2) = \{0, 1, 2, 12, 21, 121\}$:

- Type 0: $\mathcal{C} = \mathcal{A}$. Then the conditions are automatic, and \mathcal{B} is simply transverse to $\mathcal{A} = \mathcal{C}$. There are q^3 such choices.
- Type 1: $c = a, C \neq A$. The second condition is then automatic, but $C \supset a = c$ and $C \supset AB$ means $C = A$. This is a contradiction.
- Type 2: $c \neq a, C = A$. The first condition is then automatic, but $c \subset C = A$ and $c \subset ab$ means $c = a$. This is a contradiction.
- Type 12: $a \neq c, C \neq A$ but $a \subset C$. Then $a \subset C$ and $AB \subset C$ means $C = A$ This is a contradiction.
- Type 21: $a \neq c, C \neq A$ but $c \subset A$. Then $c \subset A$ and $c \subset ab$ means $c = a$. This is a contradiction.
- Type 121: In this case, the flag \mathcal{B} is determined by *either* equivalent choice: a line b in ac not equal to a or c (then B is the plane bAC) or a plane B containing AC not equal to A or C (then b is acB). The number of such choices is $q - 1$.

Therefore, this flag \mathcal{B} has either q^3 or $q-1$ internal degrees of freedom, respectively, after fixing the outer flags \mathcal{A} and \mathcal{C} to be either equal or completely transverse. The other configurations have no solutions. \square

We now apply Lemma 6.4 and the discussion above to prove Theorem 1.5 in the introduction.

6.4. Non-isotopic Links of Legendrian Spheres. Let $n \in \mathbb{N}$ and consider the bipartite Ladder Graph $L_n \subseteq \mathbb{S}^2$ depicted in Figure 83 (bottom). The number $n \in \mathbb{N}$ denotes half the number of square faces, and the right and left sides of the bipartite graph are identified in \mathbb{S}^2 . In particular, $\mathbb{S}^2 \setminus L_n$ has $2n + 2$ connected components, $2n$ squares and two 2-disks, at the north and south poles of \mathbb{S}^2 . We consider its associated 3-graph $\mathcal{L}_n \subseteq \mathbb{S}^2$, as described in Section 3, which is shown in Figure 83 (bottom). The Legendrian weave $\Lambda(\mathcal{L}_n) \subseteq (J^1(\mathbb{S}^2), \xi_{\text{st}})$ consists of a 3-component link of Legendrian 2-spheres, independent of $n \in \mathbb{N}$.

Note that the Legendrian link $\Lambda(\mathcal{L}_n) \subseteq (J^1(\mathbb{S}^2), \xi_{\text{st}})$ is smoothly isotopic to the surface unlink, as the codimension of this smooth embedding is three. We now show that the Legendrian isotopy type of the Legendrian link $\Lambda(\mathcal{L}_n) \subseteq (J^1(\mathbb{S}^2), \xi_{\text{st}})$ is different for each $n \in \mathbb{N}$. This




FIGURE 83. The bipartite Ladder Graph L_n , where the right and left sides are identified after n rungs (bottom). The 3-graph \mathcal{L}_n associated to L_n (top).

will be achieved by counting the number of points of their flag moduli spaces $\mathcal{M}(\mathcal{L}_n)$ over a finite field. The precise statement reads:

Theorem 6.5 (Theorem 1.5). *Let $\mathcal{L}_n \subseteq \mathbb{S}^2$ be the $(2n)$ -runged ladder graph and \mathbb{F}_q a finite field, q a prime power. Then the flag moduli space $\mathcal{M}(\mathcal{L}_n)$ has orbifold point count*

$$|\mathcal{M}(\mathcal{L}_n)(\mathbb{F}_q)| = \frac{q^{2n-3} - q^{n-2} + q^{n-1} + q - 1}{(q - 1)^2}.$$

Hence, the Legendrian surface links $\Lambda(\mathcal{L}_n)$ and $\Lambda(\mathcal{L}_m)$ are Legendrian isotopic iff $n = m$.

Proof. Let us consider the two flags $\mathcal{A}, \mathcal{C} \in \text{GL}(3, \mathbb{C})/B$ located in the strata corresponding to the neighborhoods of the north and south poles. We have shown these flags in Figure 84. The flags in the vertical regions will be denoted \mathcal{B}_i , $0 \leq i \leq 2n - 1$, with the cyclic condition $\mathcal{B}_0 = \mathcal{B}_{2n}$.




FIGURE 84. The flag configuration at a point of the flag moduli space $\mathcal{M}(\mathcal{L}_n)$ where $\mathcal{A} = (a, A)$, $\mathcal{C} = (c, C)$ are the inner and outer flags. Observe that the choice of \mathcal{A}, \mathcal{C} partially fills the flags in the horizontal eye-shaped regions.

By Lemma 6.4, the existence of the flags \mathcal{B}_i in the vertical hexagonal edges, $0 \leq i \leq 2n - 1$, as in Figure 84, implies that the relative position of \mathcal{A}, \mathcal{C} must either be trivial, i.e. $\mathcal{A} = \mathcal{C}$, or completely transverse, i.e. the projective lines $A \neq C$ are distinct, and $a \notin C$ and $c \notin A$. The \mathbb{F}_q -count is divided into these two cases.

First, let us consider the case where \mathcal{A} and \mathcal{C} are completely transverse, i.e. they belong to the Bruhat $\text{GL}(3, \mathbb{C})$ -orbit labeled by $w = (12)(23)(12) \in W(A_2)$. We claim that after

choosing the flag $\mathcal{B}_0 = (b, B)$, the remaining flags \mathcal{B}_i , $1 \leq i \leq 2n - 1$ are uniquely determined. The resulting flag configuration is shown in Figure 85.




FIGURE 85. Flag configuration at a point of the flag moduli space $\mathcal{M}(\mathcal{L}_n)$ in the case $\mathcal{A} = (a, A)$ is completely transverse to $\mathcal{C} = (c, C)$. For these configurations, the choice of flag (b, B) uniquely determines the point in the flag moduli.

Let us prove this. Since \mathcal{A} and \mathcal{C} are completely transverse, they determine the flags $(AC, A), (a, ac)$ in the horizontal eye-shaped spaces in the upper row, and the flags $(AC, C), (c, ac)$ in the corresponding horizontal spaces along the bottom. The additional choice of $\mathcal{B}_0 = (b, B)$ determines the flags $(AC, B), (b, ac)$ in the left and right regions adjacent to that of \mathcal{B}_0 . Note that $B \neq ac$ and $b \in B \cap ac$. Similarly, $b \neq AC$ and the two points $AC, b \in \mathbb{P}^2$ span the line B . The flag \mathcal{B}_1 must have $b \in \mathbb{P}^2$ as its point, and its line must contain $AC, b \in \mathbb{P}^2$. Hence the flag $\mathcal{B}_1 = (b, B)$ is uniquely determined, and coincides with \mathcal{B}_0 . By an analogous reasoning, \mathcal{B}_1 determines the flag (AC, B) on the adjacent region at its right, and hence the line in \mathcal{B}_2 must be B . Since the point in \mathcal{B}_2 must be the intersection $B \cap ac$, we conclude $\mathcal{B}_2 = (b, B)$ and thus $\mathcal{B}_2 = \mathcal{B}_1 = \mathcal{B}_0$. Iteratively applying these two steps, we show that $\mathcal{B}_i = \mathcal{B}_0$ for all $1 \leq i \leq 2n - 1$. The cyclic condition $\mathcal{B}_0 = \mathcal{B}_{2n}$ is automatically verified in this case. In conclusion, in this completely transverse case, the choices are the three flags $\mathcal{A}, \mathcal{B}_1, \mathcal{C}$, being pairwise completely transverse. This configuration is depicted in Figure 86 (left).




FIGURE 86. The projective flags \mathcal{A}, \mathcal{C} and $\mathcal{B} = (b, B)$ in the case \mathcal{A}, \mathcal{C} are completely transverse (left). The configuration of projective flags in the case $\mathcal{A} = \mathcal{C}$, where admissible flags $(b_1, B_1), (b_2, B_1), (b_2, B_2), (b_3, B_2), (b_3, B_3)$ are depicted (right).

The counts over a finite field are

$$|\mathrm{PGL}(3, \mathbb{F}_q)| = \frac{(q^3 - 1)(q^3 - q)(q^3 - q^2)}{q - 1}, \quad |\mathbb{P}^2(\mathbb{F}_q)| = |\mathbb{P}^2(\mathbb{F}_q)^*| = \frac{(q^3 - 1)}{(q - 1)} = q^2 + q + 1,$$

and a projective line $\mathbb{P}^1(\mathbb{F}_q)$ has $q + 1$ points. Also, note that there are $|\mathbb{P}^1(\mathbb{F}_q)| = q + 1$ choices of lines through a point. Now, the choice of the flag $\mathcal{A} = (a, A)$ gives a count of $|\mathbb{P}^2(\mathbb{F}_q)| \cdot |\mathbb{P}^1(\mathbb{F}_q)|$. The choice of the completely transverse flag $\mathcal{C} = (c, C)$ gives q^3 , as we must have $a \notin C$, and $c \in C$ but $c \neq A \cap C$. The line B in the third transverse flag $\mathcal{B}_0 = (b, B)$ must contain the point $A \cap C$, and its point $b = B \cap ac$ is uniquely determined by the choice of such B . Since B must be distinct from A and C , we get $q - 1$ choices for the line B . This yields a total count of

$$\frac{((1 + q + q^2)(1 + q)) \cdot (q^3) \cdot (q - 1)}{(1 + q + q^2)(q^3 - q)(q^3 - q^2)} = \frac{1}{q - 1},$$

for the case where the flags \mathcal{A}, \mathcal{C} are completely transverse. Thus, $\mathcal{A}, \mathcal{B}_0, \mathcal{C}$ can be fixed, mutually completely transverse, and a factor of $(q - 1)^{-1}$ remains.

Second, let us consider the case where $\mathcal{A} = \mathcal{C}$. In this case, the flags \mathcal{B}_i , $1 \leq i \leq n$, will not all be equal. We proceed with the same systematic analysis as before. The initial choice is $\mathcal{B}_1 = (b_1, B_1)$, and this determines the flags $(b_1, ab_1), (AB_1, B_1)$ in the left and right adjacent regions of \mathcal{B}_1 . In turn, this determines the line in \mathcal{B}_2 to be $B_1 \subseteq \mathbb{P}^2$. The point in \mathcal{B}_2 remains undetermined at this stage, and this is a choice of $b_2 \in \mathcal{B}_2$, with a count of q , since $b_2 \in B_1$ and $b_2 \neq A \cap B_1$. This is depicted in Figure 85. The choice of the point $b_2 \in \mathcal{B}_2$ readily determines the point in the flag \mathcal{B}_3 , whose line is undetermined. There are exactly q choices for a line $B_2 \subseteq \mathbb{P}^2$ in \mathcal{B}_3 , as it must contain b_2 and be different from B_1 . This is an iterative process, where the count of choices that determine the flag \mathcal{B}_i , $2 \leq i \leq n$ is exactly q , either because of the choice of a point or a line. The flag configuration is depicted in Figure 87.




FIGURE 87. Flag configuration at a point of the flag moduli space $\mathcal{M}(\mathcal{L}_n)$ in the case $\mathcal{A} = \mathcal{C}$. For these configurations, the sequence of flags (b_i, B_i) are part of the choice that determine the points in the flag moduli.

At this stage of the case $\mathcal{A} = \mathcal{C}$, we need to impose the cyclic condition $\mathcal{B}_0 = \mathcal{B}_{2n}$ given by the ladder graph. This is not automatic, and it will actually reduce the naive count of q^{2n} for the choices of \mathcal{B}_i , $0 \leq i \leq 2n - 1$. Let us use the $\mathrm{PGL}(3, \mathbb{F}_q)$ symmetry to fix the flags $\mathcal{A} = \mathcal{C}$ and \mathcal{B}_0 . We will now use affine coordinates, so the flag \mathcal{A} will be understood as a line $a \subseteq \mathbb{F}_q^3$

and a plane $A \subseteq \mathbb{F}_q^3$. Thus, we assume that the line $a \subseteq \mathbb{F}_q^3$ in \mathcal{A} is spanned by $\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$ and the plane $A \subseteq \mathbb{C}^3$ is the kernel of the covector $(0, 0, 1)$, and the flag \mathcal{B}_0 is given by the pair

$\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$, $(1, 0, 0)$. Note that this flag configuration has a residual isotropy group isomorphic to $(\mathbb{F}_q^\times)^2$, and we will divide our count for fixed $\mathcal{A}, \mathcal{B}_0$ by the isotropy factor of $(q-1)^2$.

Let us parametrize the remaining degrees of freedom for flags \mathcal{B}_i , $1 \leq i \leq \mathcal{B}_{2n-1}$ by the choice of coordinates $x_i \in \mathbb{F}_q$ and $a_i \in \mathbb{F}_q$, respectively used for each line b_i and plane B_i , $1 \leq i \leq n$. By labeling lines and planes by their normalized vectors and covectors, we obtain the description:

$$\begin{array}{ll}
\mathcal{B}_0 : b_0 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} & B_0 = (1, 0, 0) \\
\mathcal{B}_1 : b_1 = \begin{pmatrix} 0 \\ x_1 \\ 1 \end{pmatrix} & B_0 = (1, 0, 0) \\
\mathcal{B}_2 : b_1 = \begin{pmatrix} 0 \\ x_1 \\ 1 \end{pmatrix} & B_1 = (1, a_1, -a_1x_1) \\
\mathcal{B}_3 : b_2 = \begin{pmatrix} -a_1x_2 \\ x_1 + x_2 \\ 1 \end{pmatrix} & B_1 = (1, a_1, -a_1x_1) \\
\mathcal{B}_4 : b_2 = \begin{pmatrix} -a_1x_2 \\ x_1 + x_2 \\ 1 \end{pmatrix} & B_2 = (1, a_1 + a_2, -a_1x_1 - a_2(x_1 + x_2)) \\
\mathcal{B}_5 : b_3 = \begin{pmatrix} -a_1x_2 - (a_1 + a_2)x_3 \\ x_1 + x_2 + x_3 \\ 1 \end{pmatrix} & B_2 = (1, a_1 + a_2, -a_1x_1 - a_2(x_1 + x_2)) \\
\mathcal{B}_6 : b_3 = \begin{pmatrix} -a_1x_2 - (a_1 + a_2)x_3 \\ x_1 + x_2 + x_3 \\ 1 \end{pmatrix} & B_3 = (1, a_1 + a_2 + a_3, -a_1x_1 \\
& \qquad \qquad \qquad -a_2(x_1 + x_2) - a_3(x_1 + x_2 + x_3)) \\
\vdots & \vdots \\
\mathcal{B}_{2k} : b_k = \begin{pmatrix} -\sum_{i=2}^k \left(\sum_{j=1}^{i-1} a_j \right) x_i \\ \sum_{j=1}^k x_j \\ 1 \end{pmatrix}, & B_k = \left(1, \sum_{j=1}^k a_j, -\sum_{i=1}^k a_i \left(\sum_{j=1}^i x_j \right) \right)
\end{array}$$

Since the dot product $B_k \cdot b_k = 0$ for all $1 \leq k \leq n$, the 2-planes B_k contain the points b_k , as required. Define the new variables

$$\alpha_i = \sum_{j=1}^i a_j, \quad y_i = x_{i+1}, \quad X = \sum_{j=1}^n x_j,$$

and the vectors $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_{n-1})$, $y = (y_1, \dots, y_{n-1})$. This is an allowed change of variables, as it is a triangular and invertible transformation. The equation $\mathcal{B}_0 = \mathcal{B}_{2n}$ gives four equalities. Two of the equalities are $\alpha_n = 0$, $X = 0$. The third equation reads

$$\alpha \cdot y = 0, \quad \text{i.e.} \quad \sum_{i=1}^{n-1} \alpha_i y_i = 0.$$

The fourth equation, imposed by the vanishing of the third coordinate of B_{2n} is dependent on the first three equations, as $b_{2n} \in B_{2n}$. We are now in position to count solutions of this system over \mathbb{F}_q :

- (i) Suppose that the vector $\alpha \in (\mathbb{F}_q)^{n-1}$ is non-vanishing. There are $(q^{n-1} - 1)$ such possibilities for α . Then the equation $\alpha \cdot y = 0$ imposes exactly one linear relation among the y_i variables, $1 \leq i \leq n - 1$. This yields a choice of q^{n-2} possibilities for the vector y . The contribution in this case is thus $(q^{n-1} - 1)q^{n-2}$.
- (ii) Suppose that instead $\alpha = 0$ is the zero vector. Then the equation $\alpha \cdot y = 0$ is vacuous. The choice of an arbitrary vector $y \in \mathbb{F}_q^{n-1}$ completes the count with a factor of q^{n-1} .

In conclusion, the case $\mathcal{A} = \mathcal{C}$ yields a total count of

$$\frac{(q^{n-1} - 1)q^{n-2} + q^{n-1}}{(q - 1)^2}.$$

Finally, adding together the two cases for the relative position of the two flags \mathcal{A}, \mathcal{C} , we obtain a finite field count of

$$|\mathcal{M}(\mathcal{L}_n)(\mathbb{F}_q)| = \frac{1}{(q - 1)} + \frac{(q^{n-1} - 1)q^{n-2} + q^{n-1}}{(q - 1)^2} = \frac{q^{2n-3} - q^{n-2} + q^{n-1} + q - 1}{(q - 1)^2}.$$

□

Note also that the proof of Theorem 6.5 shows that the moduli space of n -gons \mathcal{M}_n [MGOT12, OST13] admits an embedding into our flag moduli space $\mathcal{M}(\mathcal{L}_n)(\mathbb{C})$. In the next section, we will consider N -graphs $G \subseteq \mathbb{D}^2$ with *non-empty* boundary $\partial G \neq \emptyset$, which feature prominently in our study of Lagrangian fillings through N -graphs G .

7. MICROLOCAL MONODROMIES AND LAGRANGIAN FILLINGS

This section explains how to use N -graphs G in order to study 2-dimensional exact Lagrangian cobordisms between 1-dimensional Legendrian links in $(\mathbb{S}^3, \xi_{\text{st}})$ – in particular, the study of their exact Lagrangian fillings. Briefly, the Legendrian mutations we developed in Section 4 will be used to construct Lagrangian fillings, and we use microlocal monodromies – and the connection to cluster algebras – to distinguish them. The proof of Theorem 1.8, using these two steps to build infinitely many distinct Lagrangian fillings for a class of Legendrian knots, is also given here.

7.1. Exact Lagrangian Cobordisms. This manuscript has heretofore focused on the study of Legendrian surfaces in an ambient 5-dimensional contact manifold. In fact, the theory of N -graphs and Legendrian weaves that we have developed is also useful for studying exact Lagrangian fillings of 1-dimensional Legendrian links $\Lambda \subseteq (\mathbb{S}^3, \xi_{\text{st}})$ and, more generally, exact Lagrangian cobordisms between such Legendrian links. This is also the context in which applications to both Spectral Networks and Soergel Calculus should arise.

There are two advantages to studying exact Lagrangian fillings L of $\partial L \subseteq (\mathbb{S}^3, \xi_{\text{st}})$ from the perspective of N -graphs. First, the manipulation of their Hamiltonian isotopy class $L \subseteq (\mathbb{D}^4, \omega_{\text{st}})$ becomes combinatorial, as do operations such as Polterovich surgery (see Theorem 4.10). Second, the computation of cluster coordinates for the augmentation variety $\text{Aug}(\Lambda)$ associated to the Legendrian link $\partial L = \Lambda \subseteq (\mathbb{S}^3, \xi_{\text{st}})$ is accessible.

Remark 7.1. The cluster structures in the coordinate rings of $\text{Aug}(\Lambda)$ have proven to be an effective method for proving new results for Legendrian knots in the 3-sphere [STWZ19, CG20]. We do not know how to prove these cited results using Floer-theoretic methods (such

as the Legendrian DGA [Che02, Etn05]), nor is there currently a Floer-theoretic description²² for the cluster coordinates induced by an exact Lagrangian filling $L \subseteq (\mathbb{D}^4, \omega_{\text{st}})$. \square

In this section we present the context in which Legendrian weaves $\Lambda(G)$ provide exact Lagrangian cobordisms. This is a viewpoint that we will use extensively in the remainder of the article, including Section 8 and Appendix A.

7.1.1. *The geometric setup.* Let $(\mathbb{R}^5, \xi_{\text{st}})$ have coordinates $(x, y, z, s, t) \in \mathbb{R}^5$, contact 1-form $\alpha_{\text{st}} = e^s(dz - y_1 dx_1) - dt$, and let $\pi : (\mathbb{R}^5, \xi_{\text{st}}) \rightarrow (\mathbb{R}^4, \lambda_{\text{st}})$ be the projection $\pi(x, y, z, s, t) = (x, y, z, s)$. Consider the contact 3-planes $(\mathbb{R}_l^3, \xi_{\text{st}}) := \{t = l, s = l\} \subseteq \mathbb{R}^5$ and choose two Legendrians $\Lambda_1 \subseteq (\mathbb{R}_1^3, \xi_{\text{st}})$ and $\Lambda_2 \subseteq (\mathbb{R}_2^3, \xi_{\text{st}})$. Suppose that $\Lambda \subseteq (\mathbb{R}^5, \xi_{\text{st}})$ is a Legendrian surface with isotropic boundaries $\partial\Lambda = \Lambda_1 \sqcup \Lambda_2$, and $\Lambda_1 = \Lambda \cap (\mathbb{R}_1^3, \xi_{\text{st}})$, $\Lambda_2 = \Lambda \cap (\mathbb{R}_2^3, \xi_{\text{st}})$.

The crucial geometric fact is that the projection $\pi(\Lambda) \subseteq (\mathbb{R}^4, \lambda_{\text{st}})$ is an immersed exact Lagrangian, whose immersion points are in bijection with the Reeb chords of $\Lambda \subseteq (\mathbb{R}^5, \alpha_{\text{st}})$. In particular, if the Legendrian surface $\Lambda \subseteq (\mathbb{R}^5, \xi_{\text{st}})$ has *no* Reeb chords, then the Lagrangian image $\pi(\Lambda) \subseteq (\mathbb{R}^4, \lambda_{\text{st}})$ is an embedded exact Lagrangian with boundary $\Lambda_1 \sqcup \Lambda_2$. It is readily verified that $\pi(\Lambda)$ is an exact Lagrangian cobordism from Λ_1 to Λ_2 (and *not* viceversa). The particular case of $\Lambda_1 = \emptyset$ yields exact Lagrangian fillings of Λ_2 .

In line with the constructions in this article, the Legendrians $\Lambda_1, \Lambda_2 \subseteq (\mathbb{R}^3, \xi_{\text{st}})$ that we study arise from positive braids – see [CG20, Section 2] – and thus can be described as satellites of the standard Legendrian unknot $\Lambda_{\text{st}} \subseteq (\mathbb{R}^3, \xi_{\text{st}})$. The description in the paragraph above is then modified as follows. Consider $(J^1(\mathbb{S}^1 \times [1, 2]), \xi_{\text{st}})$, two Legendrian links

$$\Lambda_1 \subseteq (J^1(\mathbb{S}^1 \times \{1\})), \quad \Lambda_2 \subseteq (J^1(\mathbb{S}^1 \times \{2\})),$$

and a Legendrian surface $\Lambda \subseteq (J^1(\mathbb{S}^1 \times [1, 2]), \xi_{\text{st}})$ such that

$$\Lambda \cap (J^1(\mathbb{S}^1 \times \{1\})) = \Lambda_1, \quad \Lambda \cap (J^1(\mathbb{S}^1 \times \{2\})) = \Lambda_2.$$

Now, suppose that the surface Λ has *no* Reeb chords, then the Lagrangian projection $\pi(\Lambda) \subseteq (J^1\mathbb{S}^1 \times \mathbb{R}, \lambda_{\text{st}})$ in the symplectization of $(J^1\mathbb{S}^1, \xi_{\text{st}})$ provides an exact Lagrangian cobordism from Λ_1 to Λ_2 . The case in which $\Lambda_1 = \emptyset$ can be compactified to $(J^1\mathbb{D}^2, \xi_{\text{st}})$ in the $(J^1(\mathbb{S}^1 \times \{1\}), \xi_{\text{st}})$ end, which symplectically corresponds to adding a standard symplectic 4-disk $(\mathbb{D}^4, \xi_{\text{st}})$ in the concave end of the symplectization, i.e. as an exact symplectic filling of $(\mathbb{S}^3, \xi_{\text{st}})$. Diagrammatically, this implies that we can describe exact Lagrangian fillings of a positive Legendrian braid $\Lambda_2 = \Lambda(\beta) \subseteq (\mathbb{S}^3, \xi_{\text{st}})$ in $(\mathbb{D}^4, \omega_{\text{st}})$ by drawing N -graphs in \mathbb{D}^2 whose free edges meet the boundary according to a positive braid word β . Here $\Lambda(\beta)$ denotes the standard satellite of the Legendrian in $(J^1\mathbb{S}^1, \xi_{\text{st}})$ whose front in $\mathbb{S}^1 \times \mathbb{R}$ is given by the positive braid (word) β .

In short, exact Lagrangian fillings between Legendrian links can be studied via the spatial wavefronts of their Legendrian lifts to the contactization, and the techniques we have developed for Legendrian surfaces can be applied. In particular, we can use our diagrammatic N -graph calculus to study and distinguish exact Lagrangian cobordisms.

7.1.2. *Free N -Graphs.* Let \mathcal{G}_β be the set of N -graphs on a 2-disk \mathbb{D}^2 with boundary braid word β . As stated above, in order to construct *embedded* exact Lagrangian fillings $L \subseteq (\mathbb{D}^4, \omega_{\text{st}})$ for $\Lambda(\beta) \subseteq (\mathbb{S}^3, \xi_{\text{st}})$ as N -graphs $G \subseteq \mathbb{D}^2$ in \mathcal{G}_β , we must have that the Legendrian weave $\Lambda(G) \subseteq (\mathbb{R}^5, \xi_{\text{st}})$ has no Reeb chords. Let us introduce the following:

Definition 7.2. An N -graph $G \subseteq \mathbb{D}^2$ is said to be free if its associated Legendrian front $\Sigma(G)$ can be woven with no Reeb chords. \square

In this section many of the N -graphs $G \subseteq \mathbb{D}^2$ can be checked to be free by direct inspection.

²²As far as we know, this remains an open question even if the exact Lagrangian filling is given by a pinching sequence [EHK16, Pan17b, Pan17a].

Example 7.3. Let $G \subseteq \mathbb{D}^2$ be a 2-graph such that $(\mathbb{D}^2 \setminus G)/(\partial\mathbb{D}^2 \cap (\mathbb{D}^2 \setminus G))$ is simply-connected. Then G is free if and only if G has no faces contained in the interior of \mathbb{D}^2 . Figure 88 shows four examples of 2-graphs.




FIGURE 88. Two free 2-graphs (i) and (ii), shown on the Left. Two 2-graphs, (iii) and (iv), whose woven front must have a Reeb chord (Right). Each of the fronts associated to the non-free two 2-graph can be woven with exactly one Reeb chord, as indicated. In both cases, the green lines depict the two sheets of a woven front and the orange segments indicate the distance between these sheets. On the left, these length of the distance grows as we approach the boundary, whereas for the 2-graph (iii) there must be a maximum for this distance, forcing a Reeb chord.

The two 2-graphs (i), (ii) on Figure 88 (Left) are free. For that, consider a smooth 1-dimensional foliation of $\mathbb{D}^2 \setminus G$ whose leaves are open intervals and such that the closure of each leaf intersects $\partial\mathbb{D}^2$. The radial-like yellow foliations depicted in Figure 88 (Left) suffice. Then choose a woven front for such 2-graphs such that the differences between the heights of the two sheets of the front strictly increase along each of the leaves of this foliation, being 0 at G and having positive value at $\partial\mathbb{D}^2$. These woven fronts do not have Reeb chords, as the functions giving the differences of heights between the sheets do not have critical points. In contrast, such foliations do not exist for the two 2-graphs (iii), (iv) on Figure 88 (Right), as $\mathbb{D}^2 \setminus G$ contains a region whose closure is contained in the interior of \mathbb{D}^2 . It can be shown that any front woven with respect to (iii) or (iv) must have a Reeb chord and there exists a woven front with a minimal number of Reeb chords, one per each interior face of G . \square

From the perspective of Lagrangian fillings, the 2-graph (i) in Figure 88 is an embedded (exact) Lagrangian filling for the 2-component standard unlink, which is the union of two disjoint Lagrangian disks $\mathbb{D}^2 \cup \mathbb{D}^2$. The 2-graph (ii) yields the embedded Lagrangian filling for the standard unknot, which is the standard flat Lagrangian disk $\mathbb{D}^2 \subseteq \mathbb{D}^4$. This stands in contrast with the immersed Lagrangian fillings represented by (iii) and (iv). The 2-graph (iii) is an immersed exact Lagrangian annulus with boundary the 2-component standard unlink, and (iv) is an immersed exact Lagrangian once-punctured 2-torus filling the standard Legendrian unknot. In general, the following criterion is useful:

Lemma 7.4. *Let $G \subseteq \mathbb{D}^2$ be a free N -graph. Then the N -graph $\mu(G) \subseteq \mathbb{D}^2$, obtained from G by performing a Legendrian mutation at any 1-cycle or Y -tree of G , is also free.*

Proof. Consider the 2-graph mutation at a monochromatic i -edge of an N -graph G . Let $\mathcal{O}_p(e)$ be a neighborhood of a monochromatic edge e in a free N -graph. The 2-graph mutation along the 1-cycle γ_e can then be performed by the exchange in Figure 89, which builds on Figure 60 (Left). Since both 2-graphs G and $\mu_e(G)$ in the exchange coincide in a neighborhood of the boundary, we can force that the front woven with respect to $\mu_e(G)$ coincides identically – not just up to homotopy of Legendrian fronts – with the given front $\Sigma(G)$ woven with respect to G . Let us choose a 1-dimensional foliation in \mathbb{D}^2 with respect to G , as in Example 7.3, such that the difference between the heights of any pair of sheets in the woven front strictly

increase (or decreases) as we move along the sheets of the foliations away from G . (This foliation exists because G is free.) We have depicted such a foliation for G in Figure 89.

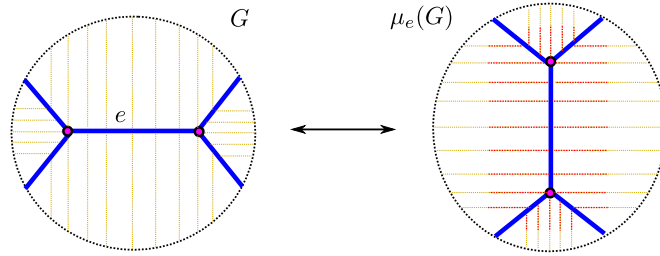


FIGURE 89. Mutation for an N -graph G along a monochromatic i -edge e . The mutated graph $\mu_e(G)$ admits a woven front $\Sigma(\mu_e(G))$ which coincides with any front $\Sigma(G)$ woven with respect to G near the boundary of the neighborhood $\mathcal{O}p(e)$. The yellow foliation near the boundary fixes the difference between the i th and $(i + 1)$ th sheets in both fronts $\Sigma(G)$ and $\Sigma(\mu_e(G))$. This foliation is extended to the interior in two different ways, yellow or red, depending on the graph being G or $\mu_e(G)$.

In order to guarantee that $\mu_e(G)$ is free, we construct a front $\Sigma(\mu_e(G))$ woven with respect to $\mu_e(G)$ as follows: this new front is identical to that of G near the boundary of the neighborhood of the monochromatic edge, and the j -th sheets for $\Sigma(\mu_e(G))$ coincide with those of $\Sigma(G)$ except for the sheets corresponding to $j = i, i + 1$. The i th and $(i + 1)$ th sheets of $\Sigma(\mu_e(G))$ are woven according to $\mu_e(G)$ such that the difference in heights between the i th and the $(i + 1)$ th sheets increases (or decreases) strictly along the 1-dimensional red foliation as we move away from $\mu_e(G)$ as shown in Figure 89 (Right). Since the red foliation is drawn to coincide with the yellow foliation at the boundary of the neighborhood $\mathcal{O}p(e)$, this is consistent with the sheets coinciding in that neighborhood. Given that the leaves of the 1-dimensional red foliation are intervals with a free end, it is possible to build such a front, meeting the condition that the difference of heights between i th and $(i + 1)$ th strictly increases (or decreases). In addition, we can draw the front $\Sigma(\mu_e(G))$ such that the slopes of each sheet are arbitrarily close to the slopes of $\Sigma(G)$. This guarantees that $\mu_e(G)$ is free as required.

For a general N -graph mutation along a l- or Y-cycle, it suffices to observe that Subsections 4.8 and 4.9 show that such mutations are given by a composition of Legendrian Reidemeister moves, as presented in Subsection 4.1, and mutations along monochromatic edges. Legendrian Reidemeister moves are local, relative to the boundary, and can be performed without ever introducing Reeb chords. Thus an N -graph mutation $\mu(G)$ of a free G is free if the statement holds for 2-graph mutations, which we have already proven above. \square

Lemma 7.4 allows us to perform Legendrian mutations to the N -graph and obtain potentially new embedded exact Lagrangian fillings. Examples of this are now illustrated. We will implicitly apply Lemma 7.4 in Subsection 7.3, in order to realize cluster mutations as N -graph mutations of embedded exact Lagrangian fillings.

7.1.3. Explicit Examples of Lagrangian Fillings. For the case of free 2-graphs on a disk \mathbb{D}^2 , this immediately yields that the max-tb Legendrian $(2, n)$ -torus positive link $\Lambda(2, n)$ has at least a Catalan C_n number worth of exact Lagrangian fillings [EHK16, Pan17b, STZ17, TZ18]. This is because C_n counts binary trees, which are equivalent to free 2-graphs. These exact Lagrangian fillings are distinguished, up to Hamiltonian isotopy, through the use of cluster

coordinates – see Subsection 7.2.1. Now, the ability to increase $N \in \mathbb{N}$ greatly expands²³ the class of Legendrian links for which their Lagrangian fillings can be studied with N -graph calculus, including all Legendrian positive braids $\Lambda(\beta)$, $\beta \in \text{Br}_N^+$ for any $N \in \mathbb{N}$.

Example 1: Recently, the first examples of Legendrian links with infinitely many exact Lagrangian fillings were described in the article [CG20]. We exhibit them here in terms of 3-graphs. For any $(p, q) \in \mathbb{N}^2$, the max-tb Legendrian (p, q) -torus positive link $\Lambda(p, q) \subseteq (\mathbb{S}^3, \xi_{\text{st}})$ is the satellite of the braid $\Delta(\sigma_1\sigma_2 \cdots \sigma_{p-1})^q\Delta$ along the standard Legendrian unknot, where $\Delta = \Delta_p \in S_p$ is the p -stranded half-twist. Let us now illustrate how to diagrammatically visualize these infinitely many Lagrangian for the Legendrian link $\Lambda(3, 6)$.

Remark 7.5. Similar p -graphs can be drawn for $\Lambda(p, q)$ for all $(p, q) \in \mathbb{N}^2$ and they produce infinitely many Lagrangian fillings if $p \geq 3, q \geq 6$ or $(p, q) = (4, 4), (4, 5)$. Alternatively, infinitely many exact Lagrangian fillings for $\Lambda(p, q) \subseteq (\mathbb{S}^3, \xi_{\text{st}})$, $p \geq 4, q \geq 7$ can also be readily constructed from those of $\Lambda(3, 6)$ [CG20, Corollary 1.5]. \square

Consider the braid word $\beta = (\sigma_1\sigma_2)^9 = \Delta(\sigma_1\sigma_2)^6\Delta$ in the 1-jet space $(J^1\mathbb{S}^1, \xi_{\text{st}})$, where $\Delta = \Delta_3 \in S_3$ is the 3-stranded half-twist $\Delta = \sigma_1\sigma_2\sigma_1$. This braid β can be depicted as a set of points in the circle \mathbb{S}^1 labeled with two colors, corresponding to σ_1, σ_2 . Figure 90 shows this braid β in two circles, the inner circle $\mathbb{S}^1 \times \{1\}$ and outer circle $\mathbb{S}^1 \times \{2\}$ in the annulus $\mathbb{S}^1 \times [1, 2]$. These two marked circles are labeled by $\Lambda(3, 6)$, as the Legendrian link associated to the marking β is the $(3, 6)$ -Legendrian link $\Lambda(3, 6)$.




FIGURE 90. Legendrian weave whose Lagrangian projection defines an infinite order element in the fundamental group of the space of Legendrian links isotopic to $\Lambda(3, 6)$. In particular, this Lagrangian concordance has infinite order in the Lagrangian concordance monoid. Infinitely many Lagrangian fillings for $\Lambda(3, 6)$, and all torus links $\Lambda(n, m)$, $n \geq 3, m \geq 6$, are obtained by concatenating this 3-graph.

The 3-graph $G \subseteq \mathbb{S}^1 \times [1, 2]$ depicted in Figure 90 describes a Legendrian surface $\Lambda(G) \subseteq (J^1\mathbb{S}^1 \times [1, 2], \xi_{\text{st}})$ with boundary $\Lambda(3, 6) \sqcup \Lambda(3, 6)$. By increasing the slope in the radial

²³This is particularly relevant for the study of exact Lagrangian fillings, as it is expected that any $\Lambda(\beta)$ with $\beta \in \text{Br}_N^+$ has only *finitely* many exact Lagrangian fillings, and we will show in Theorem 7.14 that this is *not* the case already for $N = 3$.

direction, the Legendrian surface $\Lambda(G)$ can be assumed to have no Reeb chords, and thus $\pi(\Lambda(G))$ is an exact Lagrangian cobordism from $\Lambda(3,6)$ to itself. Since the graph G has no trivalent vertices, $\Lambda(G)$ has the topology of $\Lambda(3,6) \times [1,2]$ and it is in fact an exact Lagrangian concordance. The remarkable property of the 3-graph G , and its Lagrangian projection $\pi(\Lambda(G))$, is stated in the following:

Theorem 7.6 ([CG20]). *The 3-graph exact Lagrangian concordance in Figure 90 has infinite order. In particular, for any fixed exact Lagrangian filling of $\Lambda(3,6)$, iterated concatenation of this 3-graph yields infinitely many Lagrangian fillings of the Legendrian link $\Lambda(3,6) \subseteq (\mathbb{S}^3, \xi_{st})$. \square*

In fact, it is possible to describe the entire faithful modular $\mathrm{PSL}(2, \mathbb{Z})$ -representation in [CG20] with the diagrammatics of 3-graphs. Similarly, the diagrammatics of 4-graphs give explicit spatial wavefronts for the $M_{0,4}$ -worth of the (Legendrian lift of the) Lagrangian fillings for the Legendrian link $\Lambda(4,4) \subseteq (\mathbb{S}^3, \xi_{st})$. The non-triviality, and infinite order, of this Lagrangian concordance is detected by studying its action on the cluster structure of the coordinate ring of the moduli space of isomorphism classes of simple objects in $\mathrm{Sh}_{\Lambda(3,6)}(\mathbb{R}^2)$.

Example 2: Let us address the following question. Given a positive braid β , and the Legendrian link $\Lambda = \Lambda(\Delta\beta\Delta)$, how do we diagrammatically produce an N -graph which represents an embedded exact Lagrangian filling for $\Lambda \subseteq (\mathbb{S}^3, \xi_{st})$?

Let us begin with a simple example, with $\beta = \Delta^2 = (\sigma_1\sigma_2)^3$ the full-twist, which is smoothly the $(3,3)$ -torus link. The game is to draw τ_i -edges along the boundary $\partial\mathbb{D}^2$ of a (planar) 2-disk \mathbb{D}^2 according to the braid word β and complete these edges to an N -graph G inside \mathbb{D}^2 . The only rule is that the Legendrian weave $\Lambda(G)$ should not have Reeb chords, or else it would yield an *immersed* Lagrangian filling, and thus we require G to be free.

Consider the free 3-graph G_1 in Figure 91 (upper Left). This represents an embedded exact Lagrangian filling L_1 of the max-tb Legendrian $(3,3)$ -torus link $\Lambda(3,3) = \Lambda(\Delta\beta\Delta) = \Lambda(\Delta^4)$. We can now apply the Legendrian mutation moves in Theorem 4.21 in order to produce another Lagrangian filling L_2 which is *not* Hamiltonian isotopic to the exact Lagrangian filling L_1 . (Note that L_1 and L_2 are smoothly isotopic relative to their boundaries, and L_2 will be also embedded thanks to Lemma 7.4.) In Figure 91 we perform a Lagrangian disk surgery on L_1 along a Lagrangian 2-disk which bounds the 1-cycle in $H_1(L_1, \mathbb{Z})$ graphically given by the Y -cycle in surrounded by the dashed green curve.

At this stage we can manipulate L_2 with Theorem 4.2, in this case Figure 91 (upper right) to 91 (bottom left) shows how to apply Move II to push-through a hexagonal vertex through a trivalent vertex (as indicated by the green arrow). This is an interesting move because it makes a new 1-cycle for L_2 readily visible, as represented by the blue monochromatic edge in 91 (bottom left) surrounded by a dashed green curve. We can perform Lagrangian surgery at this monochromatic edge, as in Theorem 4.21, to obtain another exact Lagrangian filling L_3 , also embedded by Lemma 7.4. It is immediate that L_1 and L_3 are not *not* Hamiltonian isotopic to L_2 , as the cluster coordinates associated to these 3-graphs, as explained in Subsection 7.2.1, show that L_1 and L_3 are not Hamiltonian isotopic. In conclusion, the 3-graphs in Figure 91 represent three distinct embedded exact Lagrangian fillings for $\Lambda(3,3)$.

Example 3: Let us illustrate what a *generic* 3-graph diagram like for a positive braid $\beta \in \mathrm{Br}_3^+$. The pictures in the case of $\beta \in \mathrm{Br}_N^+$, $N \geq 3$ are alike, with as many as $(N-1)$ -colors instead. Let us consider a random braid

$$\beta = (\sigma_1\sigma_2\sigma_1)\sigma_2^2\sigma_1^2\sigma_2^3\sigma_1^3\sigma_2^3(\sigma_1\sigma_2\sigma_1),$$




FIGURE 91. Four 3-graphs representing exact embedded Lagrangian fillings for the maximal-tb $(3, 3)$ -torus link $\Lambda(3, 3)$.

which has no particular significance to us. To obtain exact Lagrangian fillings, we draw blue and red edges around a circle $\mathbb{S}^1 \subseteq \mathbb{R}^2$, according to σ_1 or σ_2 , and construct 3-graphs with no Reeb chords and these boundary constraints. Figure 92 shows four free 3-graphs G_i , $i \in [1, 4]$, such that the Lagrangian projections $\pi(\iota(\Lambda(G_i))) \subseteq (\mathbb{D}^4, \omega_{\text{st}})$ are embedded exact Lagrangian fillings which are distinct up to Hamiltonian isotopy for $i \neq j$, $i, j \in [1, 4]$.

Remark 7.7. From our experience drawing 3-graphs, the pictures in Figure 92 accurately represent the generic appearance of exact Lagrangian fillings described by free 3-graphs. We presently do not know any example of a Lagrangian filling for a positive braid which does *not* arise as an N -graph, for some $N \in \mathbb{N}$. \square

Remark 7.8. There exists a technique for producing many such free N -graphs G , filling β -boundary conditions at a circle and thus representing embedded exact Lagrangian fillings. This is ongoing work by the first author, which in particular proves that *any* Legendrian link $\Lambda(\beta)$ arising from a positive braid $\beta \in \text{Br}_N^+$ admits an embedded Lagrangian filling whose Legendrian lift is a Legendrian weave. In precise terms, it can be proven that for each triangulation of a $|\beta|$ -gon, one can assign a free N -graph which represented an embedded Lagrangian filling of β , where $|\beta|$ is the length of the positive braid β . \square

7.2. Microlocal monodromies and cluster structures. In this section, we demonstrate how notions of cluster theory are borne out with N -graphs. This is an important ingredient in showing that microlocal monodromies can be used to distinguish exact Lagrangian fillings, as we do in Section 7.3 and as has been mentioned previously.

To orient the discussion, we recall that the cluster structures on the Fock-Goncharov moduli spaces of framed local systems described in [FG06b] were given a sheaf-theoretic description in [STW16, STWZ19]. In these works, the spectral surface associated to a bipartite graph, as defined in [Gon17, Section 2.2], is described symplectically as an exact Lagrangian filling of the zigzag Legendrian curve. In the case of bipartite graphs associated to an N -triangulation, as in [Gon17, Section 1], the zigzag curves isotope to concentric circles around the vertices of the triangulation, and the singular support of such a configuration translates to the data of a local system with a monodromy-invariant flag at each vertex. Sheaf quantization [GKS12] then implies that local systems on the exact filling embed as a cluster chart of objects, the




FIGURE 92. Four exact embedded Lagrangian fillings for the braid β in Example 3. Their satellites in $(\mathbb{R}^4, \omega_{\text{st}})$ are smoothly isotopic relative to their boundaries, but *not* Hamiltonian isotopic.)

chart being provided by the bipartite graph (and its dual quiver), and the cluster coordinates given by microlocal monodromies. The intersection form in $H_1(L, \mathbb{Z})$, or its negative, corresponds to the skew-symmetric bilinear form in cluster theory. For us, the crucial point is that we can represent all these Lagrangian fillings by N -graphs, as in the diagrammatics of Subsection 7.1, and the cluster coordinates can be read directly from the N -graph, as we will now explain.

Remark 7.9. In [TZ18], the case of Legendrian surfaces defined by trivalent 2-graphs was studied, giving a sheaf-theoretic description of the constructions in [DGG16]. In this setting, the microlocal monodromy functor μmon induces, at the level of moduli of objects, a morphism from the sheaf moduli space to the cluster chart defined by the triangulation dual to the 2-graph. The image is a (holomorphic) Lagrangian in a (holomorphic) symplectic leaf, as in [DGG16], in a manner compatible with quantization of algebra of functions.²⁴ Furthermore, in that work, the potential describing the local exact structure of the Lagrangian was interpreted as a generator of BPS states or disk invariants, following the analysis of

²⁴In work in progress with Linhui Shen, the second author will develop the relation to cluster theory more systematically, and prove Lagrangianity of the moduli space.

Aganagic-Vafa [AV00, AV12]. Here we generalize some of the constructions to N -graphs, $N \geq 2$. \square

In this article, the Legendrian surfaces are described by N -graphs, a more complex construction, but we will now explain how the basic features should persist. That is, the microlocal monodromy functor allows us to read cluster coordinates for the moduli spaces of isomorphism classes of simple objects in $Sh_\Lambda(\mathbb{R}^2)$, equivalently augmentation varieties, directly from N -graphs with boundary Λ . Examples of these constructions are provided below.

7.2.1. *Microlocal monodromies as cluster coordinates.* By definition, microlocal monodromy is a functor

$$\mu mon : Sh_\Lambda \rightarrow Loc(\Lambda)$$

from the category Sh_Λ of sheaves microsupported on the Legendrian surface Λ , as defined in Subsection 5.3, to the category of local systems on Λ [STZ17]. This functor carries microlocal rank-one sheaves $F \in Sh_\Lambda^1$, i.e. simple sheaves, to rank one local systems on the surface Λ . Since it is locally defined, the monodromy of the local system $\mu mon(F)$ around a loop $\gamma \in H_1(\Lambda)$ can be evaluated by restricting the constructible sheaf F to an annular tubular neighborhood of γ . Below, these annuli are depicted as thin purple loops. In short, the calculation for Legendrian weaves can be done using the microlocal monodromy functor μmon as it is used for knots, as described in [STZ17].

The main point in these computations is that the stalk $\mu mon(F)|_\lambda$ at a point $\lambda \in \Lambda$ is the cone of the restriction map corresponding to λ , and for flags this is the inclusion of subspaces, whence cones become cokernels. The transversality of adjacent flags ensures that these cokernels propagate as a local system. Let us now perform these calculations for 1-cycle $\gamma \in H_1(\Lambda, \mathbb{Z})$, starting at the 1-cycle represented by a monochromatic edge.

Let us consider a monochromatic edge with label τ_i , as depicted in Figure 93.




FIGURE 93. Neighborhood of a monochromatic edge e with the data determining a constructible sheaf F . As we show, the microlocal monodromy $\mu mon(F)$ along the 1-cycle $\gamma(e)$ is given by the cross-ratio $\langle a, b, c, d \rangle$.

Near such a monochromatic edge, a sheaf object in a simply connected face is specified by the data of a quadruple of flags. Each of these flags has the same subspaces \mathcal{F}^j in each region for $j \neq i$, and for $j = i$ we additionally require the data in each region of a line l in the two-dimensional space $V := \mathcal{F}^{i+1}/\mathcal{F}^{i-1}$. This is the data of four lines $a, b, c, d \subseteq V$, as specified in Figure 93. Restricted to the purple oval shown, we have a cylindrical braid of type $\beta = \sigma_i^4$, where σ_i is the lift of the transposition τ_i from the Coxeter group S_N to the braid group Br_N . Given the prescribed transversality imposed by the flag moduli of an N -graph, we further know that the cyclic chain of inequalities $a \neq b \neq c \neq d \neq a$ holds. We thus have the chain of isomorphisms of cokernels

$$a \cong V/b \cong c \cong V/d \cong a,$$

which computes the microlocal monodromy. In this case, the isomorphism that we obtain is the cross ratio

$$\langle a, b, c, d \rangle = \frac{a \wedge b}{b \wedge c} \cdot \frac{c \wedge d}{d \wedge a}$$

of the four lines a, b, c, d , and it is equal to the cluster coordinate associated to 1-cycle γ as prescribed in [FG06b, Section 9].

Let us now consider the cluster coordinate associated to a Y-cycle, which is a new type of 1-cycle, as it only appears for $N \geq 3$. Figure 94 depicts a Y-cycle, drawn as a purple circle, along with the data determining a constructible sheaf in a neighborhood of this 1-cycle.




FIGURE 94. Neighborhood of a Y-cycle with the data determining a constructible sheaf F . As we compute, the microlocal monodromy $\mu\text{mon}(F)$ along the associated 1-cycle γ is given by the triple ratio of the three transverse flags.

Following the notation in Section 6, we denote by ab the unique plane containing the two lines a and b , while AB denotes the intersection of the planes A and B . The braid associated to the Y-cycle γ , as drawn by the purple circle in Figure 94, is given by $\beta = (\sigma_i \sigma_{i+1} \sigma_i)^3$, where σ_i corresponds to the crossing coming from a τ_i -edge. By considering the three-dimensional vector space $V := \mathcal{F}^{i+2} / \mathcal{F}^{i-1}$, a given flag is specified by a line and a plane in V . Since the word $\sigma_i \sigma_{i+1} \sigma_i$ represents the half-twist Δ for flags on V , and $\tau_i \tau_{i+1} \tau_i$ is the Coxeter element in S_3 , the complete data specifying a constructible sheaf near the Y-cycle is given by three transverse flags $(a, A), (b, B), (c, C)$ in V . In this notation, the line is written in lower case and the covector defining the the plane in upper case, thus (a, A) determines a flag. Now, the microlocal monodromy functor μmon along γ is computed as the composition of the isomorphisms

$$a \cong V/B \cong c \cong V/A \cong b \cong V/C \cong a.$$

Let $v_a \in a, v_b \in b, v_c \in c, v_d \in d$ be non-zero vectors defining the corresponding one-dimensional lines. Then the parallel transport from a to c in this basis is given by the quotient $B(a)/B(c)$, where $B(a)$ is the pairing between the vector v_a and the covector B . Iterating these isomorphisms, we conclude that the microlocal monodromy along the Y-cycle is given by

$$\langle (a, A), (b, B), (c, C) \rangle := \frac{B(a)C(b)A(c)}{B(c)C(a)A(b)}.$$

This expression is precisely the triple product of transverse flags as defined in [FG06b], and thus we have shown that the microlocal monodromy along a Y-cycle determines a cluster coordinate.

7.2.2. Legendrian Mutations are cluster transformations. The coordinate transformations upon Legendrian mutations can also be computed, as we will demonstrate in an example. The conclusion is that Legendrian mutations induce cluster transformations. The case of a monochromatic edge follows from the analysis in [TZ18, STWZ19], and we now study the mutation at a Y-cycle. To do so, consider the local geometry shown in Figure 95. We want to compute how the cluster coordinate associated to the unique monochromatic (blue) edge

– as in Subsection 7.2.1 – changes as we perform a Legendrian mutation along the Y-cycle specified by the unique hexagonal vertex.




FIGURE 95. The geometric setup before performing a Legendrian mutation at the Y-cycle, where the cluster coordinate associated to the monochromatic edge is given by the cross-ratio $\langle B, bc, ab, B' \rangle$.

The monochromatic blue edge has monodromy equal to the cross ratio $z := \langle B, bc, ab, B' \rangle$ of the four planes in the projective line of planes containing b . (This can be computed directly or by intersecting the four lines with any transverse line – see Subsection 7.2.1.) Now, after Legendrian mutation at the Y-cycle, the resulting 3-graph is shown Figure 96.




FIGURE 96. The result of applying a Legendrian mutation to Figure 95 along the Y-cycle, along with the data of a constructible sheaf.

The 1-cycle determined by the blue monochromatic edge in Figure 95 becomes a (bichromatic edge) 1-cycle contained in the 3-graph shown in Figure 97, which is itself a piece of Figure 96, in its upper-right corner²⁵:

²⁵The trivalent blue vertex in Figure 97 is the unique trivalent blue vertex in Figure 96. The trivalent red vertex in Figure 97 is the rightmost trivalent red vertex in Figure 96.




FIGURE 97. Local geometry near the 1-cycle after mutation.

By applying Move II, we can push the red trivalent vertex in Figure 97 through the hexavalent vertex. This allows us to represent the 1-cycle as a monochromatic edge again, as shown in Figure 98:




FIGURE 98. The constructible sheaf near the 1-cycle after Legendrian mutation and Move II. The new coordinate is thus the cross-ratio $\langle B, bAC, ab, B' \rangle$.

The required conclusion, stating that the new cross-ratio $z' = \langle B, bAC, ab, B' \rangle$ is obtained by a cluster transformation, follows from this:

Lemma 7.10. *Let $x = \langle (a, A), (b, B), (c, C) \rangle$ be the triple ratio of flags and $z = \langle B, bc, ab, B' \rangle$ the cross-ratio of lines. Denote by $z' = \langle B, bAC, ab, B' \rangle$ the new microlocal monodromy. Then*

$$z' = z(1 + x).$$

Proof. By PGL_3 invariance, we may assume that

$$a = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, A = (0, 0, 1), b = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, B = (1, 0, 0), c = \begin{pmatrix} 1 \\ -1 \\ 1 \end{pmatrix}, C = (1, 1 + x, x).$$

Since the cross-ratio z is prescribed, we find that $B' = (z, 1, 0)$, and similarly

$$AC = \begin{pmatrix} 1 + x \\ -1 \\ 0 \end{pmatrix}.$$

This implies that $bAC = (1, 1 + x, 0)$, and thus $z' = z(1 + x)$. \square

Note that $z' = z(1 + x)$ in Lemma 7.10 is the transformation expected for a cluster-X transformation.²⁶ This concludes that a Legendrian mutation at the Y-cycle induces a *cluster* transformation for the microlocal monodromy coordinate at the monochromatic blue edge in Figure 95. The computation is analogous if we choose a different blue monochromatic edge to be added near the Y-cycle. In particular, if we had chosen instead the blue edge attaching

²⁶The rule for a cluster-X transformation upon mutating at loop k is that the monodromy z_i transforms to $1/z_k$ if $i = k$ and otherwise $z'_i = z_i(1 + z_k^{-\mathrm{sgn}\epsilon_{ik}})^{-\epsilon_{ik}}$, where $\epsilon_{i,k}$ is the skew-symmetric cluster form. We get agreement on the nose if we make this form the *negative* of the intersection pairing.

at the lower-right of the Y-cycle and pointing upward, and again called its monodromy z , then we would have $A' = (0, z, 1)$ and would obtain

$$z' = \langle ab, aBC, A, A' \rangle = z \left(1 + \frac{1}{x} \right)^{-1},$$

in agreement with the cluster transformation.²⁷

Example 7.11. Flip of a N -triangulation. Let (C, τ) be a punctured surface C , τ an ideal triangulation and τ' an ideal triangulation obtained from τ by a flip. Denote by t_N , resp. t'_N , the N -triangulation refinement of τ , resp. τ' . It is an exercise [Gon17, Proposition 1.1] to show that the Legendrian weave $\Lambda(G(t'_N))$ differs from $\Lambda(G(t_N))$ by a sequence of $\binom{N+1}{3}$ 2-graph mutations, i.e. $\Lambda(G(t'_N))$ can be obtained from $\Lambda(G(t_N))$ by performing $\binom{N+1}{3}$ Legendrian mutations along 1-cycles represented by monochromatic edges.

For instance, [Gon17, Figure 9] translates into four monochromatic edge mutations for a flip in a $N = 3$ triangulation, as we have depicted in Figure 99. We can see how to perform the corresponding moves for 3-triangulations with 3-graphs. Indeed, referring to the notation in Figure 105, perform a monochromatic edge mutation at z and w , then perform Move III, a flop of the two trivalent and two hexagonal vertices in the center, and proceed with a mutation at the remaining two monochromatic edges. In conclusion, the constructions of this paper can therefore be used to give a geometric understanding of the intermediate quivers arising when flipping N -triangulations.




FIGURE 99. Flip in a 3-triangulation realized as four monochromatic edge mutations. In general, the Legendrian weaves associated to two N -triangulations which differ by a flip of the underlying (1-)triangulation differ by a sequence of $\binom{N+1}{3}$ such 2-graph mutations. Note the Move III flop isotopy in-between the two mutation pairs.

²⁷We remark that the case $x = -1$ is not a generic configuration of flags, since in this case $c \in ab$, and thus not in the domain of the birational cluster map.

7.3. N-graph Realization of Quiver Mutations. In this subsection we explain how to use N -graphs in order to construct infinitely many Lagrangian fillings for certain Legendrian links in the standard contact 3-sphere $(\mathbb{S}^3, \xi_{\text{st}})$. These Lagrangian fillings are distinguished by the microlocal monodromies/cluster coordinates in Subsection 7.2.

Let $\beta \in \text{Br}_N^+$ be a positive braid, with a fixed braid word $w(\beta)$. Consider an N -graph $G \subseteq \mathbb{D}^2$ on the 2-disk such that the labels of the edges of G near $\partial\mathbb{D}^2$, read cyclically, form the word $w(\beta)$. Following Subsection 7.1, the Lagrangian projection $L(G) = \pi(\iota(\Lambda(G))) \subseteq (\mathbb{R}^4, \omega_{\text{st}})$ of $\iota(\Lambda(G)) \subseteq (\mathbb{R}^5, \xi_{\text{st}})$ is an exact Lagrangian filling of the Legendrian link $\Lambda(\beta) \subseteq (\mathbb{S}^3, \xi_{\text{st}})$ associated to the positive braid β . All the N -graphs $G \subseteq \mathbb{D}^2$ which feature in this subsection will be free, and thus the Lagrangian projections are embedded, equivalently $\Lambda(G)$ has no Reeb chords.

Now consider a free N -graph $G \subseteq \mathbb{D}^2$, $b_1(G) := \text{rk}(H_1(\Lambda(G), \mathbb{Z}))$ and a basis

$$B = \{[\gamma_1], \dots, [\gamma_{b_1(G)}]\}$$

for $H_1(\Lambda(G), \mathbb{Z}) \cong \mathbb{Z}^{b_1(G)}$, equivalently a basis for the first homology group of its Lagrangian projection $L(G)$. For a choice of basis B , we denote by $Q(B)$ the intersection quiver of the 1-cycles γ_i , $i \in [1, b_1(G)]$. The vertices v_i of the quiver $Q(B)$ are in bijection with elements of the homology basis B , and the number of arrows between two distinct vertices v_i, v_j is given by the geometric intersection number $|\gamma_i \cap \gamma_j|$. The direction of each arrow is given by the sign of each geometric intersection, and there are no loops, i.e. no edges from a vertex v_i to itself. The quiver obtained by mutation of a quiver Q at the vertex v_i will be denoted $\mu_i(Q)$.

We will study the realization of quiver mutations, algebraic in nature, as Legendrian mutations of free N -graphs, which are geometric. Suppose that there exists a subset B_μ of classes $[\gamma_i]$, $i \in [1, k]$, for some $k \leq b_1(G)$, such that $\gamma_i \in B_\mu$ is represented by a 3-graph cycle with *no* multiplicity. That is, each 1-cycle γ_i is represented by either a Y-cycle, a tree, a monochromatic edge 1-cycle or a long edge. Let $\{x_1, \dots, x_k\}$ be the cluster coordinates associated to $\{\gamma_1, \dots, \gamma_k\}$ via microlocal monodromies, as in Subsection 7.2.1.

Remark 7.12. In general, this set of cluster coordinates $\{x_1, \dots, x_k\}$ is only a partial subset of the entire cluster seed $\{x_1, \dots, x_{b_1(G)}\}$ for $H_1(L(G), \mathbb{Z})$. The ability to work with a subset is an advantage that allows for our methods to be applied in more generality. From the viewpoint of cluster algebras, the vertices of $Q(B)$ which are *not* in $Q(B_\mu)$ are to be considered as frozen vertices, and the variables $\{x_{k+1}, \dots, x_{b_1(G)}\}$ as frozen coordinates. \square

By Subsection 4.8, and Lemma 7.4, we can perform a Legendrian mutation at $\gamma_i \in B_\mu$ and obtain a free N -graph $\mu_i(G)$. The intersection quiver $Q(\mu_i(B))$ associated to the mutated basis $\mu_i(B)$ is the mutated quiver $\mu_i(Q(B))$. The 1-cycle in the mutated graph $\mu_i(G)$ corresponding to $\gamma_j \in B$, under mutation at γ_i , is denoted by $\mu_i(\gamma_j)$. By Subsection 7.2.2, the cluster coordinate associated to $\mu_i(\gamma_j)$ is given by the j -th coordinate in the cluster transformation of $\{x_1, \dots, x_{b_1(G)}\}$ at x_i . Therefore, the exact Lagrangian filling represented by the free N -graph $\mu_i(G)$ has intersection quiver $\mu_i(Q(B))$ and cluster coordinates obtained by mutation of the cluster seed $\{x_1, \dots, x_{b_1(G)}\}$ for $L(G)$ at x_i . In conclusion, if the 1-cycles are represented by trees, performing *one* quiver (or cluster seed) mutation as a Legendrian mutation is possible, and the microlocal monodromies after the Legendrian mutation accurately reflect cluster mutation.

Remark 7.13. The challenging aspect of the *geometric* side is that iterating this procedure is not necessarily possible, or at least readily accessible. This aspect is *not* reflected in

the algebra of quiver mutations (or cluster coordinate mutations) since, by definition, two opposite edges between vertices are canceled²⁸. \square

The technology of 3-graphs and their mutations, as developed in Subsection 4.8, allows us to iterate Legendrian mutations in an abundance of cases, including arbitrarily high genus. We will illustrate explicit cases in which an infinite sequence of quiver mutations can be realized as an infinite sequence of N -graph mutations. These cases can be inserted in (infinitely many) other examples, and the first consequence is the production of new families of Legendrian links with *infinitely many* exact Lagrangian fillings:

Theorem 7.14. *Let $\Lambda_{s,t} = \Lambda(\beta_{s,t}) \subseteq (\mathbb{S}^3, \xi_{st})$ be the Legendrian link given by the standard satellite of the positive braid*

$$\beta_{s,t} = (\sigma_1^3 \sigma_2)(\sigma_1^3 \sigma_2^2)^s \sigma_1^3 \sigma_2 (\sigma_2^2 \sigma_1^3)^t (\sigma_2 \sigma_1^3)(\sigma_2^{t+1} \sigma_1^2 \sigma_2^{s+2}), \quad s, t \in \mathbb{N}, s, t \geq 1.$$

Then $\Lambda_{s,t} \subseteq (\mathbb{S}^3, \xi_{st})$ admits infinitely many embedded exact Lagrangian fillings in $(\mathbb{D}^4, \lambda_{st})$ realized as 3-graphs $G_{s,t} \subseteq \mathbb{D}^2$ and their Legendrian mutations.

Proof. The argument is uniform for all $s, t \in \mathbb{N}$ and all the difficulties, and their solutions, are already present for the simplest case.²⁹ Let us thus assume $s = t = 1$ for now. First, we need to construct a free 3-graph $G = G_{1,1}$ which represents a Lagrangian filling for the Legendrian link $\Lambda(\beta)$ associated to $\beta = \beta_{1,1}$. This 3-graph is shown in Figure 100:




FIGURE 100. The 3-graph G and the initial Quiver Q .

The exact Lagrangian $L(G)$ associated to G is a genus-4 surface with two boundary components, and thus $b_1(G) = 9$. Let us consider the subset $B_\mu = \{\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5\}$ given by the following 1-cycles: γ_1 is represented by the yellow 1-cycle in Figure 100, which is a tree of Y-pieces, and $\gamma_2, \gamma_3, \gamma_4, \gamma_5$ are represented by monochromatic edges, in purple in Figure 100. In addition, we consider the 1-cycle γ_6 represented by the monochromatic edge, in green.

²⁸In previous attempts to geometrically iterate Lagrangian mutations, such as [STW16, Section 2], this obstruction manifests itself as embedded curves becoming immersed upon performing Dehn twists, a problem which presently has no known solution.

²⁹We thank Dylan Thurston for useful discussions on quivers and their mutations. In particular, for providing the infinite sequence of mutations that we use in this proof.

The intersection quiver $Q = Q(B_\mu \cup \{\gamma_6\})$ is given by the quiver drawn in Figure 100. The quiver Q is of infinite mutation type, as it is associated³⁰ to hyperbolic Coxeter diagram [Law17, Table 1]. In fact, we claim that the sequence of quiver mutations μ_{s_n} , where

$$s_n = \begin{cases} 1 & \text{if } n \equiv 1 \\ 2 & \text{if } n \equiv 2 \\ 3 & \text{if } n \equiv 3 \\ 4 & \text{if } n \equiv 4 \\ 5 & \text{if } n \equiv 5 \end{cases} \pmod{5},$$

is an infinite sequence of quiver mutations. Indeed, each time we apply the sequence of mutations $\mu_5\mu_4\mu_3\mu_2\mu_1$, the number of arrows from the vertex v_1 to v_6 increases by two, and the number of arrows from v_i to v_6 , for $i \in [2, 5]$ increases by one. In particular, at the k th iteration there are $2k + 1$ arrows from v_1 to v_6 and k arrows from v_i to v_6 , $i \in [2, 5]$. Now we reach the core of the issue, which is realizing this infinite sequence of quiver mutations as Legendrian mutations of 3-graphs. For that, we observe the following two properties:

Let us perform a Legendrian mutation along the Y-tree which represents the 1-cycle γ_1 . The resulting 3-graph, which is free by Lemma 7.4, is depicted in Figure 101:




FIGURE 101. Mutated 3-graph $\mu_1(G)$ at the 1-cycle γ_1 (yellow) corresponding to vertex 1 in Q , and its associated intersection quiver $\mu_1(Q)$.

Now, upon this Legendrian mutation at γ_1 the 1-cycles $\gamma_2, \gamma_3, \gamma_4, \gamma_5$ are *still* represented by monochromatic edges. These new 1-cycles $\mu_1(\gamma_2), \mu_1(\gamma_3), \mu_1(\gamma_4), \mu_1(\gamma_5)$ are circled in purple in Figure 101. The figure also displays the mutated quiver $\mu_i(Q(B_\mu \cup \{\gamma_6\}))$ and the cycle

³⁰Precisely, the quiver Q corresponds to the rank 6 paracompact hyperbolic Coxeter group $\bar{L}_5 = [3^{1,1,1,1,1}]$.

γ_6 in green. Similarly, upon this Legendrian mutation, the 1-cycle $\mu_1(\gamma_1)$ is *still* represented by an embedded Y-tree, as depicted in yellow in Figure 101.

These properties hold true as we now perform Legendrian mutations at the monochromatic edges $\gamma_2, \gamma_3, \gamma_4, \gamma_5$. The free 3-graph resulting from these four mutations is drawn in Figure 102:




FIGURE 102. Mutated 3-graph $\mu_5\mu_4\mu_3\mu_2\mu_1(G)$ and its associated intersection quiver $\mu_5\mu_4\mu_3\mu_2\mu_1(Q)$.

The claim is that we can iterate the sequence of mutations $\mu_5\mu_4\mu_3\mu_2\mu_1$ geometrically as Legendrian mutation of the free 3-graph, and these two properties hold. That is, at any stage in the sequence of mutations μ_{s_n} we have that

- (i) the 1-cycles $\gamma_2, \gamma_3, \gamma_4, \gamma_5$ are represented by monochromatic edges,
- (ii) the 1-cycle γ_1 is represented by an embedded Y-tree, with no multiplicities.

In fact, the Y-tree representing γ_1 always has exactly four Y-pieces. These four pieces have been surrounded by a dashed pink circle in Figures 100 through 104. The two items above can be readily verified, as follows. The behavior of the mutated 3-graph near each of the the monochromatic edges is as depicted in Figure 103:

It thus follows that $\gamma_i, i \in [2, 5]$, remains a monochromatic edge upon *any* iteration of the 3-graph mutation $\mu_5\mu_4\mu_3\mu_2\mu_1$. Similarly, according to the Legendrian mutation rules of Subsection 4.8, each Y-piece of the Y-tree representing γ_1 itself mutates to a Y-piece, and mutating at $\gamma_2, \gamma_3, \gamma_4, \gamma_5$ preserves this property. Thus the pattern persists upon any




FIGURE 103. The effect of the sequence of mutations $\mu_5\mu_4\mu_3\mu_2\mu_1$ near the monochromatic edges γ_i , $i \in [2, 5]$. The only mutations from these five that affect γ_i are μ_1 and μ_i .

iteration. The two properties (i) and (ii) now allow us to perform the sequence of mutations μ_{s_n} up to *any* point in the sequence. For instance, the sequence of mutations $\mu_1\mu_5\mu_4\mu_3\mu_2\mu_1$ applied to G lead to the 3-graph in Figure 104:




FIGURE 104. Mutated 3-graph $\mu_1\mu_5\mu_4\mu_3\mu_2\mu_1(G)$ and its associated intersection quiver $\mu_1\mu_5\mu_4\mu_3\mu_2\mu_1(Q)$.

In order to pairwise distinguish the exact Lagrangian fillings associated to the sequence of 3-graphs $(\mu_{s_n}\mu_{s_{n-1}}\cdots\mu_1)(G)$, up to Hamiltonian isotopy, we use the microlocal monodromies $\{x_1, x_2, x_3, x_4, x_5, x_6\}$ along the 1-cycles γ_i , $i \in [1, 6]$, and their mutations. By Subsection 7.2.2, the cluster seed $\{x_1, x_2, x_3, x_4, x_5, x_6\}$ associated to the quiver Q mutates to the cluster seed associated to $(\mu_{s_n}\mu_{s_{n-1}}\cdots\mu_1)(Q)$ upon performing the Legendrian 3-graph mutations $(\mu_{s_n}\mu_{s_{n-1}}\cdots\mu_1)(G)$. Since the quivers $(\mu_{s_n}\mu_{s_{n-1}}\cdots\mu_1)(Q)$ are distinct, and so are the associated cluster seeds, it follows that the associated Lagrangian fillings are distinct. This concludes the proof for $s = t = 1$.

The general case $s, t \in \mathbb{N}$ is proven with the same argument. Indeed, the free 3-graph $G_{1,1}$ in Figure 100 generalizes to a 3-graph whose boundary is $\beta_{s,t}$, just by adding s copies of the leftmost pattern in $G_{1,1}$, to the left, and t copies of the rightmost pattern in $G_{1,1}$, to

the right. In this general case, it is still true that γ_1 is represented³¹ by a Y-tree and the remaining $\{\gamma_2, \gamma_3, \dots, \gamma_s, \gamma_{s+1}, \dots, \gamma_{s+t+3}\}$ cycles are represented by monochromatic edges. The argument is then identical, with the infinite sequence of mutations given by

$$s_n = i, \quad \text{if } n \equiv i \pmod{s+t+3}, \quad 1 \leq i \leq s+t+3.$$

The reader can directly verify that this is an infinite sequence of mutations, as the multiplicity of the arrows to the cycle γ_{s+t+4} – generalizing the green cycle γ_6 in Figure 100 – increases as we apply the mutations $\mu_{s+t+3}\mu_{s+t+2}\cdots\mu_2\mu_1$. \square

Remark 7.15. For $s = t = 1$, note that the sequence μ_{s_n} never mutates at the 1-cycle γ_6 , i.e. at the sixth vertex v_6 in Q . It is nevertheless crucial to include γ_6 in the quiver as well as the cluster variable x_6 , with its subsequent mutations. Note that the 1-cycle γ_6 is initially represented by an embedded curve in the 3-graph, but this curve develops immersed points as we iterate the sequence of mutations μ_{s_n} according to Subsection 4.8. This still allows us to define the cluster coordinate associated to it *but* we would not be able to mutate along such a 1-cycle just with the rules developed in Subsection 4.8. (This is just a side remark, since the argument for Theorem 7.14 does *not* require mutating at v_6 .) \square

The Legendrian links in Theorem 7.14 are relatively simple. For instance, the Legendrian knot associated to $\beta_{1,1}$ is genus-4 two-component link. One of the components is an unknot and the other is the $(2, 7)$ -torus knot 7_1 . Note that $\Lambda(\beta_{1,1})$ is (smoothly) distinct from the $(3, 6)$ -torus link that the first author studied in [CG20], which also has genus-4. Thus, not only does Theorem 7.14 bring a *new* method to construct infinitely many Lagrangian fillings, but it in fact provides new Legendrian links with infinitely many Lagrangian fillings.

Remark 7.16. Note that the \overline{L}_5 quiver that we used in Theorem 7.14 appears as a subquiver of the intersection quiver for several other positive braids. Following L. Lewark’s positive braid table³² each of the following positive genus-6 braids, $14n_{5644}$, $15n_{118169}$, $16n_{144958}$, $16n_{149517}$, $16n_{173894}$, $16n_{175324}$ and $16n_{339638}$, to name a few, contain \overline{L}_5 in their intersection quiver. We believe that an argument similar to Theorem 7.14 should prove that the maximal-tb representative of each of these links has infinitely many exact Lagrangian fillings. \square

Finally, the contrast between Theorem 7.14 and [CG20, Corollary 1.5] is interesting. The former constructs an infinite family of Lagrangian fillings for a Legendrian link by directly using Legendrian mutations, which are themselves distinguished by their effect – as cluster mutations – on the microlocal monodromies. The latter result [CG20] is entirely about constructing infinite order Lagrangian concordances, coming from Legendrian loops of positive braids, and the infinite family of Lagrangian fillings is a byproduct of such construction. In particular, N -graph calculus should apply to much more general Legendrian links, and does not require knowing about the existence of an infinite order element in their Lagrangian concordance monoid.

8. MODULI SPACE FOR N -TRIANGLES AND NON-ABELIANIZATION

In this final section, we focus on N -graphs associated to N -triangulations, as introduced in Section 3. This class of N -graphs G yields Legendrian weaves $\Lambda(G)$ whose Lagrangian projections are related to the Goncharov-Kenyon conjugate Lagrangian surfaces [GK13, STWZ19]. These Lagrangian surfaces have also appeared in the context of Gaiotto-Moore-Neitzke’s spectral networks [GMN13, Nei14]. In particular, we prove Theorem 1.9, which computes the flag moduli space $\mathcal{M}(G)$ for G any N -triangle t_N , matching the algebraic results in [GMN14, Section 8] and [FG06b, Section 9].

³¹In this case the Y-tree has $(s+t+2)$ Y-pieces, $s+1$ to the left and $t+1$ to the right of the base root.

³²Lukas Lewark’s Positive Knots Table: Braids and Trees at “<http://lewark.de/lukas/braids.html>”.

8.1. Flag moduli space of the N -triangle. Let us compute the flag moduli space associated to the N -graph $G(t_N)$ of an N -triangle t_N , as we defined in Section 3.2 (see Figure 15). The result reads as follows:

Theorem 8.1. *Let $G(t_N)$ be the N -graph associated to an N -triangle t_N . The flag moduli space of $G(t_N)$ is a $\binom{N-1}{2}$ -dimensional complex torus, i.e.*

$$\mathcal{M}(t_N, G(t_N); k) \cong (k^*)^{\binom{N-1}{2}}.$$

This rest of this subsection is devoted to the proof of Theorem 8.1. The statement of Theorem 8.1 is an instance of how incidence geometry problems connect to the contact topology of Legendrian surfaces. Indeed, although our proof is entirely within projective geometry, the conclusion from Theorem 8.1 ought to be read as the fact that the moduli space $\mathcal{M}(t_N, G(t_N); k)$ is parametrized by the toric coordinates provided by the holonomies $\text{Hom}(H_1(\Lambda(G(t_N)), \mathbb{Z}), k^*)$. For $k = \mathbb{C}$, this complex torus should be related to the complex torus appearing in Fock-Goncharov [FG06b, FG06a] in their study of cluster varieties, see [Kuw20, Theorem 8.3].

Theorem 8.1 can also be interpreted as follows. The triangle t_N is topologically a disk \mathbb{D}^2 with boundary a circle $\partial\mathbb{D}^2 = \mathbb{S}^1$. The Legendrian weave

$$\Lambda(G(t_N)) \subset (J^1(\mathbb{D}^2), \xi_{\text{st}})$$

has a Lagrangian projection $L := \pi(\Lambda(G(t_N)))$, which is an exact Lagrangian submanifold, where $\pi : J^1(\mathbb{D}^2) \rightarrow T^*\mathbb{D}^2$ is the projection along the standard (vertical) Reeb flow. The Lagrangian L has boundary in $T^*\mathbb{D}^2|_{\mathbb{S}^1} \cong J^1(\mathbb{S}^1)$, and it is checked that ∂L is the cylindrical Legendrian braid Δ^3 , where Δ is the half-twist positive braid corresponding to a longest word in the Weyl group, i.e. the Garside element. Since $G(t_N)$ is free, L is an embedded exact Lagrangian filling of ∂L . Now, by looking at the boundary circle \mathbb{S}^1 and considering the moduli space à la [STZ17], we conclude that the moduli space of Lagrangian fillings should carry a cluster structure: the flag moduli space $\mathcal{M}(t_N, G(t_N); k)$ is *one* such chart.

In fact, by an argument akin to Lemma 5.6, the flags at two vertices of the triangle t_N determine the flags along the edge they bound, and therefore the flags along the boundary circle $\partial\mathbb{D}^2 = \mathbb{S}^1$ must be determined by the flags at the vertices, themselves three mutually completely transverse flags in the flag variety \mathcal{B} . This space of triples of mutually transverse flags is one of the Richardson varieties³³ \mathcal{R} . Now, by the PGL_N action, two totally transverse flags can be put in standard position B, B^- , with residual symmetry the Cartan H of diagonal matrices up to scale. Then the moduli space \mathcal{R}/H is a cluster variety and the exact Lagrangian filling L provides a cluster chart via its moduli of local systems

$$\text{Loc}(L) \cong \text{Hom}(H_1(\Lambda(G(t_N)), \mathbb{Z}), k^*) \cong (k^*)^{\binom{N-1}{2}}.$$

This torus can be checked to agree with that of 8.1. Note also that, following Section 7 many other cluster charts and exact Lagrangian fillings can be found by performing N -graph mutations. Let us now prove our result:

Proof of Theorem 8.1. Let us argue by induction on N , where the base case $N = 2$ follows from the fact that PGL_2 acts transitively on triples of distinct points. Let us assume that

$$\mathcal{M}(t_N, G(t_N); k) \cong (k^*)^{\binom{N-1}{2}}$$

for the N -graph of an N -triangle. Consider an $(N + 1)$ -triangle with one side being an arbitrary fixed preferred base, and thus the row associated to this base contains $2N - 3$ triangles. It is combinatorially apparent that the complement of this row in t_{N+1} is in fact an N -triangle t_N , and thus we can construct t_{N+1} by adding such base row to t_N . This combinatorial splitting is translated into a containment of an N -graph $G(t_N)$ within $G(t_{N+1})$.

³³We thank Ian Le for many discussions on the Richardson variety.

Let us describe such splitting in the $(N + 1)$ -graph by providing its construction starting from the N -graph $G(t_N)$.

Start with the N -graph $G(t_N)$ – see Section 3.2 – and consider the $\binom{N}{2}$ edges intersecting the base side of t_N . The edges are depicted vertically and the base side horizontally – see Figure 15. These are τ_i -edges, $i = 1, \dots, N - 1$, with exactly $(N - 1 - k)$ τ_k -edges. The $(N + 1)$ -graph $G(t_{N+1})$ can be described in the following N stages:

1. First, insert an $(N - 1, N)$ -hexagonal point in the unique τ_{N-1} edge in the base side of $G(t_N)$. The τ_N -edge aligned with the previously existing τ_{N-1} -edge is continued down vertically. The remaining two τ_N -edges are extended horizontally, respectively to the left and to the right, and the remaining two τ_{N-1} edges are continued down diagonally, in south-east and south-west direction respectively.
2. Second, continue down the τ_i -edges, $i = 1, \dots, N$, until the two τ_{N-1} -edges intersect with the two originally existing τ_{N-2} -edges. In the moment of collision, insert two $(N - 2, N - 1)$ -hexagonal vertices at the intersection point matching the two incoming τ_{N-1} and τ_{N-2} trajectories. We extend the two τ_{N-1} -edges adjacent to the incoming τ_{N-2} -edges horizontally to the left and to the right.
The remaining two pairs of three edges, each with two τ_{N-2} -edges and a τ_{N-1} -edge, are continued down, with the τ_{N-1} -edges continued vertically and the τ_{N-2} -edges continued diagonally in the south-east or south-west directions, accordingly.
3. Iteratively, we proceed as follows in the l th stage, $2 \leq l \leq N - 1$. We continue down the τ_i -edges, $i = 1, \dots, N$, and at this stage the only edges being continued diagonally down are τ_{N-l+1} -edges. There are $2l - 2$ of such edges, which can be gathered in two groups, internal and external.

By definition, there are two external edges, which are the leftmost and rightmost τ_{N-l+1} -edges, respectively continuing south-west and south-east. For these two external edges, we insert two $(N - l, N - l + 1)$ -hexagonal and describe the N -graph as described in Stage 2. The $2l - 4$ internal τ_{N-l+1} -edges, which continue down diagonally, ought to intersect with τ_{N-l} -edges, which continue down vertically.

For the $2(l - 2)$ internal edges, there are $(l - 2)$ such intersections, since an intersection occurs for each pair. For each such an intersection, insert a $(N - l, N - l + 1)$ -hexagonal vertex, and continue the outgoing three edges down as described by the local model for the hexagonal vertex. Hence, for each of these hexagonal vertices, the outgoing τ_{N-l+1} -edge continues vertically down whereas the two τ_{N-l} -edges continue down diagonally. Thus, at the l th stage we have inserted exactly l $(N - l, N - l + 1)$ -hexagonal points.

4. In the N th stage, all τ_i -edges, $i \leq 2 \leq N$ continue down vertically and we are left with $2(N - 1)$ τ_1 -edges continuing diagonally. In line with the previous stages, there are two external τ_1 -edges and $2(N - 2)$ internal edges. Insert two τ_1 -trivalent vertices at the end of the two external τ_1 -edges. The internal edges will meet in consecutive pairs at $N - 2$ intersection points. In this final stage we insert a τ_1 -vertex in each of these intersection points, and continue the remaining τ_1 -edge vertically down.

Let us now compute the flag moduli space $\mathcal{M}(t_N, G(t_{N+1}); k)$ using this inductive construction of $G(t_{N+1})$. A crucial fact to be used is Lemma 5.6, i.e. at a hexagonal vertex, four consecutive flags uniquely determine the remaining two flags. Let us assume that we have chosen a point in $\mathcal{M}(t_N, G(t_N); k)$ and we thus have the data of a flag \mathcal{F} in \mathbb{P}^N for each open

region in $\mathbb{D}^2 \setminus G(t_N)$. This data needs to be considered in the moduli space of flags, given that $G(t_{N+1})$ is an $(N + 1)$ -graph, and thus we fix an embedding $i_0 : \mathbb{P}^N \rightarrow \mathbb{P}^{N+1}$ and the corresponding inclusion $\mathrm{PGL}_N \subseteq \mathrm{PGL}_{N+1}$. Let us then start the construction $G(t_{N+1})$ from $G(t_N)$ by stages, as described above, and prove the statement in Theorem 8.1.

In the first stage, the flag data at the inserted $(N - 1, N)$ -hexagonal point in the τ_{N-1} -edge is uniquely determined by a choice of a codimension-2 projective subspace H^2 in \mathbb{P}^N , transverse to $i_0(\mathbb{P}^N)$. Note that the intersection of H^2 and $i_0(\mathbb{P}^N)$ is uniquely determined by the flag data coming from $\mathcal{M}(t_N, G(t_N); k)$. We claim that this choice in the first stage can be absorbed by the symmetry group PGL_{N+1} .

In order to understand the symmetry group, it is convenient to represent an element in PGL_{N+1} via the projective matrix

$$\left[\begin{array}{cccc|c} a_{1,1} & a_{1,2} & \cdots & a_{1,N} & a_{1,N+1} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,N} & a_{2,N+1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ a_{N,1} & a_{N,2} & \cdots & a_{N,N} & a_{N,N+1} \\ \hline a_{N+1,1} & a_{N+1,2} & \cdots & a_{N+1,N} & a_{N+1,N+1} \end{array} \right],$$

where the subgroup $\mathrm{PGL}_N \subseteq \mathrm{PGL}_{N+1}$ is defined by

$$\mathrm{PGL}_N = \{A \in \mathrm{PGL}_{N+1} : a_{N+1,N+1} = 1, a_{i,N+1} = a_{N+1,i} = 0, 1 \leq i \leq N\}.$$

In these coordinates, we can assume that the subgroup $K \subseteq \mathrm{PGL}_{N+1}$ fixing our fixed hyperplane $i_0(\mathbb{P}^n) \subseteq \mathbb{P}^{N-1}$ is cut out by the equations

$$K := \{A \in \mathrm{PGL}_{N+1} : a_{N+1,i} = 0, 1 \leq i \leq N\}.$$

As a result, the remaining PGL_{N+1} -symmetries (once the flag moduli space $\mathcal{M}(t_N, G(t_N); k)$ is fixed, and thus the symmetries of PGL_N have been used) consist of projective transformations of the form

$$\left[\begin{array}{cccc|c} a_{1,1} & a_{1,2} & \cdots & a_{1,N} & c_1 \\ a_{2,1} & a_{2,2} & \cdots & a_{2,N} & c_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ a_{N,1} & a_{N,2} & \cdots & a_{N,N} & c_N \\ \hline 0 & 0 & \cdots & 0 & c_{N+1} \end{array} \right],$$

where $a_{i,j}$ are fixed, $1 \leq i, j \leq N$, $c_i \in k$, $1 \leq i \leq N$, and $c_{N+1} \in k^*$ are free. Then, in this coordinate system, we can assume that the choice of the codimension-2 projective subspace H^2 uses the gauge provided by $c_1, c_2 \in k$.

In the second stage, two $(N - 2, N - 1)$ -hexagonal vertices are inserted. For each of them, the flag data is fixed by induction in three out of the six regions near the hexagonal vertex. Hence, there is a choice of a codimension-3 projective subspace H^3 in each of these two vertices. Let us fix one of these choices by using the free coordinate $c_3 \in k$ and notice that

the other choice has an a priori moduli of k . Nevertheless, the τ_{N-3} -edge that interacts with the τ_{N-2} edges in the third stage forces that moduli to be k^* , since the two newly chosen flags must be τ_{N-3} -transverse. Thus in the second stage we have used the symmetry provided by $c_3 \in k$ and we are left with a k^* contribution to the flag moduli.

In the l th stage, $3 \leq l \leq N - 1$, we proceed inductively as follows. We partition the $(N-l, N-l+1)$ -hexagonal vertices inserted in this stage into two groups: external, containing two of them, and internal, containing $(l-2)$ of them. By definition, the two external $(N-l, N-l+1)$ -hexagonal vertices are the leftmost and rightmost vertices. Each of these two external vertices have flags fixed in three out of the six regions, by the process in the $(l-1)$ st stage. Thus, as in the second stage, there is exactly one choice of flag at each of these $(N-l, N-l+1)$ -hexagonal vertices which determines each of their respective neighborhoods. This corresponds to a choice of codimension- l projective subspace H^{l+1} in accordance with the incidence conditions imposed by the given flags. Proceeding as in the second stage, we fix one of these choices with the free variable c_{l+1} and the remaining choice contributes k^* to the flag moduli.

The $(l-2)$ internal $(N-l, N-l+1)$ -hexagonal vertices have flags fixed in four out of the six regions, given the process in the $(l-1)$ st stage. By Lemma 5.6, these hexagonal vertices are uniquely determined in their neighborhoods. Thus, although the k^* contribution of one of the external hexagonal vertices interacts with an internal vertex, no contributions to the flag moduli space come directly from the internal vertices.

The argument then develops iteratively in the above manner until the $(N-1)$ st stage is completed. The last N th stage consists of the insertion of N τ_1 -trivalent vertices. Following the same pattern as before, only the two external trivalent vertices contribute to the flag moduli, since each of the internal trivalent vertices have their three surrounding flags determined at the $(N-1)$ st stage. In this last stage, the variables c_i , $1 \leq i \leq N$ have been fixed and the only remaining degree of free symmetry is $c_{N+1} \in k^*$. Let us use such symmetry to fix the choice in one of the two external trivalent vertices, and thus the contributions of this last stage to the flag moduli space is the k^* choice of the remaining point coming from the remaining external trivalent vertex.

The conclusion in the statement Theorem 8.1 now follows by gathering the contributions of the flag moduli space at each stage. Indeed, the first stage has no contribution, whereas each of the $(N-1)$ stages, from the second to the last N th stage, has a k^* flag moduli space contribution. By the inductive hypothesis, the desired flag moduli space is

$$\mathcal{M}(t_{N+1}, G(t_{N+1}); k) \cong \mathcal{M}(t_N, G(t_N); k) \times (k^*)^{N-1} \cong (k^*)^{\binom{N-1}{2}} \times (k^*)^{N-1} \cong (k^*)^{\binom{N}{2}},$$

which corresponds to the statement, as required. \square

Remark 8.2. Note that the inductive combinatorial description of $G(t_{N+1})$ in terms of $G(t_N)$ used in the proof of Theorem 8.1 can be used to provide a third alternative definition of the local N -graph $G(t_N)$, in addition to the descriptions introduced in Subsection 3.2. \square

8.1.1. *Tetrahedral Triangulations at $N = 3$ and $N = 4$.* Let us study the Legendrian weaves $\Lambda(G(\tau))$ and flag moduli space $\mathcal{M}(\mathbb{S}^2, G(\tau))$ associated to 3- and 4-graphs $G(\tau)$ for the tetrahedral 3- and 4- triangulations τ of the 2-sphere \mathbb{S}^2 . The case $N = 2$ has been discussed in Subsection 6.1 above, where $\Lambda(G) \cong \mathbb{T}_c^2$ is the Legendrian Clifford Torus and $\mathcal{M}(G) \cong \mathbb{P}^1 \setminus \{0, 1, \infty\}$. Let us denote the pair of pants $\mathbb{P}^1 \setminus \{0, 1, \infty\}$ by \mathcal{H} .

Let us consider the 3-graph $G^{(3)} = G(\tau^{(3)}) \subseteq \mathbb{S}^2$ associated to the tetrahedral 3-triangulation $\tau^{(3)}$ of the 2-sphere \mathbb{S}^2 , according to the construction in Section 3. We want to compute its flag moduli space $\mathcal{M}(G)$. This will be done directly by using the N -graph calculus computations in Section 4. Indeed, it is proven in Subsection 4.6 that in this case the (satellite of the)

Legendrian weave $\Lambda(G)$ is Legendrian isotopic to the four-fold connected sum of the Clifford torus \mathbb{T}_C^2 . Hence, we obtain that $\mathcal{M}(\mathbb{S}^2, G^{(3)}) \cong \mathcal{H}^4$. From the description in Theorem 8.1, we are also giving a contact geometric proof of the following

Corollary 8.3 ([FG06b]). *The moduli of four generic flags in \mathbb{C}^3 is isomorphic to \mathcal{H}^4 . \square*

The same argument, using N -graph calculus also allows us to study the flag moduli space $\mathcal{M}(\mathbb{S}^2, G^{(4)})$, where $G^{(4)} = G(\tau^{(4)}) \subseteq \mathbb{S}^2$ is the 4-graph associated to the tetrahedral 4-triangulation $\tau^{(4)}$ of the 2-sphere \mathbb{S}^2 . It is left as an exercise for the reader to use Theorem 8.1 and conclude that $\mathcal{M}(\mathbb{S}^2, G^{(4)})$ is isomorphic, as an algebraic variety, to

$$\mathcal{M}(\mathbb{S}^2, G^{(4)}) = \{(z_1, w_1, \dots, z_5, w_5) \in (\mathbb{C}^*)^9 : (1 - \kappa)w_i z_i - z_i + 1 = 0, 1 \leq i \leq 5\} \times (\mathcal{H})^4,$$

where $\kappa = 1 - z_1 z_2 z_3 z_4 z_5 \in \mathbb{C}^*$. The exercise is solved in [DGG16, Section 6.3.2] in the language of the 3d $N = 2$ superconformal field theory $T_4[\Delta, \Pi]$.

8.2. A Computation of the Non-Abelianization Map. We conclude the main body of the manuscript by exploring the relationship between Legendrian weaves and the works [FG06b, AV00, AV12, Pal15] in some explicit examples. In particular, we present a case in which the non-Abelianization map featured in [GMN13, GMN14] can be realized by the microlocal monodromies associated to constructible sheaves microlocally supported along Legendrian weaves.

The context is described as follows. Let (C, τ_N) be a polygon endowed with an ideal N -triangulation τ_N , and choose a wavefront for $\Lambda(G(\tau_N))$ with no Reeb chords, such that the Lagrangian projection is a smooth exact Lagrangian L embedded in the cotangent bundle $(T^*C, \lambda_{\text{st}})$. This Lagrangian projection L has a sheaf quantization [NZ09] to a rank- N sheaf on C with no singular support, i.e. a local system in C . Now, of course, all local systems on polygons are trivial, *but* the crucial point is that the Lagrangian covering gives a preferred³⁴ basis for the fibers of the local system, which can undergo changes à la handle-slides in the Morse context – see [GKS12]. Here, the Lagrangian covering is given by the restriction $\pi|_L : L \rightarrow C$ of the projection $\pi : T^*C \rightarrow C$ onto the zero section. Now, the N -graphs and the microlocal monodromies, as discussed in Section 7.2, precisely encode these changes. In our context, the non-Abelianization map is the construction that recovers the constructible sheaf from its microlocal monodromy.

We illustrate this in the following example. Figure 105 shows the 3-graph G associated to two adjacent 3-triangles. Suppose that we are given a local system on $\Lambda(G)$. Denote by x, y the two monodromies of the corresponding Legendrian weave around the two Y-cycles, and z, w the two microlocal monodromies along the two l-cycles represented by the two (red) monochromatic edges.




FIGURE 105. The 3-graph associated to an adjacent pair of 3-triangles. The monodromies along the four 1-cycles are labeled x, y, z, w .

³⁴In the Floer-theoretic language of the Fukaya category, the basis elements are the intersections of the exact Lagrangian with the cotangent fibers.

The Legendrian weave $\Lambda(G)$ is a thrice-punctured genus-one surface and these four 1-cycles are a basis for $H_1(\Lambda(G), \mathbb{Z})$. We would like to reconstruct the flag data, specifying a constructible sheaf, from the monodromies x, y, z, w of the local system. Indeed, this will realize the Non-Abelianization map [GMN13] from rank-one local systems on the (spectral, or conjugate) Lagrangian – parametrized by monodromies – to decorated rank-two local systems on the base surface. Since the base surface here is contractible, the only degrees of freedom are the choices of flags at vertices. The map is computed as follows.

Let $(a, A), (b, B), (c, C)$ be the flags at the vertices of the left triangle, and let (d, D) be the remaining flag. We would like a birational map from the monodromies (x, y, z, w) to the choice of flags. By using the PGL_3 -action, we may assume $(a, A), (b, B), (c, C)$ are as in Subsection 7.2.2 above, with triple product x . Then the flag (c, D) is determined by the cross ratios z and w , and the triple product y . For instance, z is the cross ratio $\langle b, BC, AB, BD \rangle$ while we find $w = \langle a, AD, AB, AC \rangle$. These determine D , whence the triple product y fixes $d \in D$. Direct computation shows

$$d = \begin{pmatrix} -\frac{x}{q}(1+x) \\ x(1+y) \\ -py(1+x) \end{pmatrix}, \quad D = (pq, p(1+x), x).$$

This thus recovers [FG06b, Pal15] from the perspective of N -graphs.

Example: Tetrahedron with 3-Triangulation. Let us conclude this subsection by analyzing the genus-4 Legendrian weave Λ in Example 4.6 from the microlocal perspective. We also compute, following [TZ18], the primitive which characterizes (a discrete cover of) $\mathcal{M}(G)$ as an exact Lagrangian subvariety. Following [AV00, AV12], this primitive – the superpotential of an effective 4d theory – is interpreted as a generating function of BPS numbers, and should have integrality properties. We check this for this example.

Consider the tetrahedron with its unique 3-triangulation, as in Example 4.6, which gives rise to a 3-graph G . An object in the category of simple constructible sheaves $Sh_{\Lambda(G)}^1(\mathbb{S}^2 \times \mathbb{R}, \Lambda)_0$ microlocally supported along $\Lambda(G)$ is defined by a four-tuple of transverse flags in $V \cong \mathbb{C}^3$, placed at the vertices of the tetrahedron, as in Figure 94.

Note that there are $4 \cdot 3 = 12$ total nodes, and the Legendrian surface indeed has genus $g = 4$. We therefore have $2g = 8$ cluster variables, specified by the monodromies around each of the eight loops, which themselves are a basis for $H_1(\Lambda(G), \mathbb{Z}) \cong \mathbb{Z}^8$. Four of the monodromies are the triple ratios along the faces. Let us label the faces by the three unordered vertices it contains, e.g. we write x_{123} for the monodromy of the loop determined by the minimal triangle at the center of the face (123): it is the triple ratio of the three flags at vertices 1, 2 and 3. There are $4 \times 3 = 12$ more edge monodromies, but we will find $4 \cdot 2 = 8$ relations among all these 16 total, giving 8 independent monodromies as expected for a genus-4 surface. Let us compute the edge monodromies.

First, following [TZ18], for each edge e we define a corresponding coordinate x_e to be the *negative* of the cross ratio.³⁵ Now there are two relations for each vertex: first, the product of the edge coordinates around the encircling triangular face is unity; second, the product of edge and Y -monodromies encircling the vertex at a greater distance is unity. There are thus 8 independent coordinates, and we can take two from each of the triangles surrounding the four vertices. Let us then write

$$x_{12} = -\frac{w_1 \wedge v_4 \ v_2 \wedge v_3}{v_4 \wedge v_2 \ v_3 \wedge w_1}$$

³⁵We believe the sign appears due to the fact that we should be considering *twisted* local systems, i.e. lifts to the circle bundle of the surface that have monodromy -1 over the circle fibers, as in [FG06b] and [GMN13, Section 10].

for the coordinate associated to the edge of the triangle encircling vertex 1 and traversing the one-simplex of the triangulation between vertices 1 and 2, where v_i , $i \in [1, 4]$, are generators of lines and w_i are generators for planes (thought of as anti-symmetric two-vectors) — and likewise for the other edges. Then the relation for the triangle encircling vertex 1 is $x_{12}x_{13}x_{14} = 1$, and likewise for the other vertices. Recall that we have similarly denoted by x_{123} the inverse of the coordinate corresponding to the Y-cycle in the face containing vertices 1, 2 and 3 — and likewise for the coordinates x_{ijk} , $i, j, k \in [1, 4]$. For the first vertex we have the unital relation

$$x_{123}x_{21}x_{142}x_{41}x_{134}x_{31} = 1,$$

and likewise for the other three vertices. This expresses the flag moduli in terms of generators, given by x_{ij} , x_{ijk} , and relations, as above.

Let us verify that these coordinates define a (holomorphic) Lagrangian embedding of flag moduli space $\mathcal{M}(G)$ associated to the genus-4 Legendrian $\Lambda(G)$ into the moduli space of framed local systems for $C = \mathbb{S}^2$. The symplectic 2-form is computed from the intersection form to be

$$\omega = -d \log x_{12} \wedge d \log x_{13} + d \log x_{23} \wedge d \log x_{24} - d \log x_{34} \wedge d \log x_{31} + d \log x_{41} \wedge d \log x_{42}.$$

We can directly compute the following four relations

$$x_{12} = \frac{-1}{1 + x_{13}}, \quad x_{24} = \frac{-1}{1 + x_{23}}, \quad x_{34} = \frac{-1}{1 + x_{31}}, \quad x_{42} = \frac{-1}{1 + x_{41}},$$

which readily imply that the embedding of the flag moduli space $\mathcal{M}(G)$ in each of the cluster charts for the moduli space of framed local systems is Lagrangian. This holomorphic Lagrangian $\mathcal{M}(G)$ is in fact *exact* and we can compute a primitive function W for the restriction of the Liouville 1-form λ_{st} . This would allow us to write $\mathcal{M}(G)$ as the graph Γ_{dW} of the 1-form dW in this chart. This primitive encodes the BPS states associated to some Lagrangian filling, given by the Lagrangian projection of $\Lambda(G)$, determined by a phase and a framing (implicit here) as in [TZ18, Section 4.8] – see also [AV00, AV12]. For that, we define the variables

$$\begin{aligned} U_1 &= -x_{13}, & V_1 &= -x_{12}, & U_2 &= -x_{23}, & V_2 &= -x_{24}, \\ U_3 &= -x_{31}, & V_3 &= -x_{34}, & U_4 &= -x_{41}, & V_4 &= -x_{42}. \end{aligned}$$

Also, recall that if we have $U + V^{-1} = 1$ with $U = e^u$ and $V = e^v$, then we can write

$$v = -\log(1 - U) = \partial_u \text{Li}_2(U).$$

Hence, since we have $U_i + V_i^{-1} = 1$ for all i , with symplectic 2-form $\omega = \sum_i du_i \wedge dv_i$, we conclude that $\mathcal{M}(G) = \Gamma_{dW}$ where

$$W(U_1, U_2, U_3, U_4) = \sum_{i=1}^4 \text{Li}_2(U_i).$$

This computation for the BPS potential is in line with the results in [TZ18, Section 5].

Finally, let us review how *geometric* methods, as developed in Section 4, would lead to this result. Instead of the algebraic computation above, we could have directly used the diagrammatic calculus, as in Example 4.6, and deduced that our Legendrian weave $\Lambda(G) \cong \#_{i=1}^4 \mathbb{T}_c^2$ is the Legendrian connected sum of four Clifford 2-tori \mathbb{T}_c^2 . Since the generating function of BPS numbers for \mathbb{T}_c^2 is given by one dilogarithm $\text{Li}_2(U)$, by direct computation, and the potential W is additive under connected sum, we could have directly deduced that $W(U_1, U_2, U_3, U_4) = \sum_{i=1}^4 \text{Li}_2(U_i)$. This concludes that our algebraic computation above is consistent with the contact topology of the underlying Legendrian weave.

In this appendix, we provide a construction and a concise speculation regarding the symplectic geometrization of Soergel Calculus via Legendrian weaves. The following discussion owes a good deal to B. Elias and E. Gorsky, as explained in the introduction, to whom we are very grateful. Soergel calculus, as developed by B. Elias, M. Khovanov and G. Williamson [EK10, EW16], provides a *diagrammatic* presentation of the category of Soergel bimodules, which itself categorifies the Hecke algebra. The similarities between Elias’ diagrammatic calculus and our Legendrian weaves are apparent. Legendrian weaves can be understood as a *geometric* approach to the study of the algebra of certain complexes of Soergel bimodules. We explain this below.

Remark A.1. Soergel bimodules are essential to categorifications of knot invariants [Rou06, Soe07, Kho07, EK10]. The link between these and moduli spaces of sheaves for Legendrian braid closures was described in [STZ17, Section 6]. From this perspective, it is not unnatural to seek a connection between planar Soergel structures and planar structure defined by Legendrian weaves, the two-dimensional version of braids. \square

The category of Soergel bimodules is the Karoubi completion of the subcategory of Bott-Samelson bimodules, arising as the equivariant cohomology of a closed Bott-Samelson variety, and thus it suffices to understand the relation to this latter class of bimodules. The key connection between the present work and Soergel bimodules is that a subclass of Legendrian weaves yields exact Lagrangian cobordisms between Legendrian links, which are themselves represented as positive braids. The moduli space of microlocal constructible sheaves supported on a *singular* compactification of a positive braid is a closed Bott-Samelson variety, and our Legendrian weaves, understood as Lagrangian cobordisms – and singularly compactified – induce morphisms between these closed Bott-Samelson varieties.

Thus, we are able to geometrize the diagrammatics of Soergel calculus by considering the D_4^- -singularity for the trivalent vertices in [EK10, EW16], the A_3 -swallowtail singularity for the univalent vertex and the A_1^3 -singularity for their hexagonal vertices. (The Soergel calculus we geometrize corresponds to the $m = 2$ Coxeter exponent.)

Remark A.2. Exact Lagrangian cobordisms are *directed*, due to the convexity directionality in symplectic topology. The dissonance arises from the fact that, as of today, Soergel calculus only considers closed Bott-Samelson varieties, whereas the moduli space of microlocal sheaves supported on a positive braid is an *open* Bott-Samelson variety. Thus, the Soergel calculus is geometrized by *singular* compactifications of our Legendrian weaves, and our Legendrian weave calculus, without compactification, should naturally induce a Soergel calculus for *open* Bott-Samelson varieties. \square

For instance, the A_3 -Zamolodchikov relation from Soergel calculus corresponds to the A_1^4 -Reidemeister move in Legendrian weave calculus, as depicted in Figure 106.

Now, let us consider two positive braids $\beta_1, \beta_2 \in \text{Br}_n^+$, $n \in \mathbb{N}$, and their associated Legendrian (long) links $\Lambda(\beta_1), \Lambda(\beta_2) \subseteq (J^1[0, 1], \xi_{\text{st}})$ [CG20, Section 2]. A Legendrian weave $\Lambda \subseteq (J^1([0, 1] \times [1, 2]), \xi_{\text{st}})$ with no Reeb chords and boundaries $\Lambda(\beta_1)$ at $[0, 1] \times \{1\}$, and $\Lambda(\beta_2)$ at $[0, 1] \times \{2\}$, yields an embedded and exact Lagrangian cobordism $L(\Sigma)$ from $\Lambda(\beta_1)$ to $\Lambda(\beta_2)$ in the symplectization of $(J^1[0, 1], \xi_{\text{st}})$, as in Section 7. In particular, each trivalent vertex $\Sigma(G_{tri})$ and hexagonal vertex $\Sigma(G_{hex})$ yield the following exact Lagrangian cobordism:

- (i) The Lagrangian projection $L(G_{tri})$ of the Legendrian weave $\Lambda(G_{tri})$ is a Lagrangian cobordism from the Legendrian tangle $\Lambda(\beta_1)$ given by one crossing in two strands




FIGURE 106. Contact Isotopy among Legendrian weaves, relative to the boundaries. The lack of Reeb chords allows us to interpret these as exact Lagrangian cobordisms between the positive braids $\sigma_{i+1}\sigma_i\sigma_{i-1}\sigma_{i+1}\sigma_i\sigma_{i+1}$ and $\sigma_{i-1}\sigma_i\sigma_{i-1}\sigma_{i+1}\sigma_i\sigma_{i-1}$. The fact that these Lagrangian cobordisms are Hamiltonian isotopic implies that the morphism induced between the associated Bott-Samelson bimodules must coincide. \square

$\beta_1 = \sigma_i$, to the Legendrian tangle $\Lambda(\beta_2)$ given by two crossing in two strands $\beta_1 = \sigma_i^2$, where $i \in \mathbb{N}$ is labeling the transposition τ_i of the edges of G_{tri} . Smoothly, this is a saddle cobordism obtained by an index-1 handle attachment to the Lagrangian cone $\Lambda(\beta_1) \times [0, \varepsilon]$ in the symplectization, for $\varepsilon \in \mathbb{R}^+$ small.

- (ii) The Lagrangian projection $L(G_{hex})$ of the Legendrian weave $\Lambda(G_{hex})$ is a Lagrangian concordance from the Legendrian tangle $\Lambda(\beta_1)$ given by three crossings in three strands $\beta_1 = \sigma_i\sigma_{i+1}\sigma_i$, to the Legendrian tangle $\Lambda(\beta_2)$ given by $\beta_2 = \sigma_{i+1}\sigma_i\sigma_{i+1}$, where $i \in \mathbb{N}$ is labeling the transpositions τ_i, τ_{i+1} in the edges of G_{hex} . Smoothly, this is a Lagrangian surface obtained by graphing a Reidemeister three move.

For simplicity, let us suppose that the relative homology $H_1(L, \partial_-L; \mathbb{Z})$, which we denote by $H_1(L)$, is a free \mathbb{Z} -module and the surface L is spin, as is verified for the two local cobordisms above. An exact Lagrangian cobordism $L \subseteq (J^1[0, 1], \xi_{st}) \times [1, 2]$ from $\Lambda(\beta_1)$ to $\Lambda(\beta_2)$ yields an algebraic map

$$\Phi_L : \widehat{\mathcal{M}}(\Lambda(\beta_1)) \longrightarrow \mathcal{M}(\Lambda(\beta_2)),$$

where $\widehat{\mathcal{M}}(\Lambda(\beta_1))$ is an algebraic $(\mathbb{C}^*)^{b_1(L)}$ -bundle over $\mathcal{M}(\Lambda(\beta_1))$, and $\mathcal{M}(\Lambda(\beta))$ denotes the moduli space of microlocal rank-1 objects in the dg-category of microlocal sheaves in $\mathbb{S}^1 \times$ microlocally supported on $\Lambda(\beta)$, as described in [CG20, Section 3], [STWZ19, STZ17].

Remark A.3. In the Floer-theoretic context, the map Φ_L is obtained by applying the contravariant functor $Hom(\cdot, k)$ in the category of dg-algebras to the morphism

$$\Phi_L^{Fl} : \mathcal{A}(\Lambda(\beta_2)) \longrightarrow \mathcal{A}(\Lambda(\beta_1)) \otimes_{\mathbb{Z}} \mathbb{Z}[H_1(L)]$$

of the Legendrian Contact dg-algebras $\mathcal{A}(\Lambda(\beta))$ associated to Legendrian links $\Lambda(\beta)$. The Floer theoretic map Φ_L^{Fl} is described in [EHK16, Pan17b], and it is a count of holomorphic strips whose boundary homology classes are encoded in $\mathbb{Z}[H_1(L)]$. To ease the geometry, we have tensored by the flag moduli space map Φ_L above $\mathbb{C}[H_1(L)] \cong \mathbb{Z}[H_1(L)] \otimes_{\mathbb{Z}} \mathbb{C}$ to base change the $\text{Spec}(\mathbb{Z}[H_1(L)])$ -bundle to a complex variety $\widehat{\mathcal{M}}(\Lambda(\beta_1))$. \square

The relation to Soergel calculus now arises because the moduli space of simple microlocal sheaves $\mathcal{M}(\Lambda(\beta_1))$ is (explicitly) isomorphic to the open Bott-Samelson variety associated to β , also known as the Broué-Michel variety of β [STZ17, Tri19, CG20]. Let $R = H^*(\mathcal{B})$

denote the cohomology of the complete flag variety for $\mathrm{GL}(N, \mathbb{C})$, $N \in \mathbb{N}$, and B_{s_i} the Bott-Samelson Soergel $(R \otimes R)$ -bimodule associated to a permutation $s_i \in S_N$ in the Weyl group S_N of $\mathrm{GL}(N, \mathbb{C})$. The Rouquier complex $T_i := [B_{s_i} \rightarrow R]$ will be denoted by T_i , for all $i \in \mathbb{N}$. Consider a braid

$$\beta = \prod_{j=1}^l \sigma_{i_j}, \quad 1 \leq i_j \leq k-1,$$

where σ_{i_1} is the leftmost crossing in the front diagram of the Legendrian braid, and the crossings are read from left to right. Then the (singular) compactly supported cohomology of algebraic variety $\mathcal{M}(\Lambda(\beta))$ is described by the tensor product

$$T_\beta = T_{i_1} \otimes_R T_{i_2} \otimes_R \cdots \otimes_R T_{i_l},$$

of Rouquier complexes.

Remark A.4. Should the reader be interested in the closure of the Legendrian $\Lambda(\beta) \subseteq (J^1 S^1, \xi_{\mathrm{st}})$, instead of the long link $(J^1[0, 1], \xi_{\mathrm{st}})$, the cohomology of the corresponding moduli space $\mathcal{M}(\Lambda(\beta))$ is obtained by applying Hochschild homology to the above complex T_β . In particular, $H^*(\mathcal{M}(\Lambda(\beta)))$ coincides with the triply-graded homology of the knot associated to β , equivalently, Khovanov-Rozansky link homology – see [STZ17, Theorem 6.14]. \square

In conclusion, the geometric map $\Phi_L : \widehat{\mathcal{M}}(\Lambda(\beta_1)) \rightarrow \mathcal{M}(\Lambda(\beta_2))$ functorially induces

$$H_c^*(\Phi_L) : H_c^*(\widehat{\mathcal{M}}(\Lambda(\beta_1))) \rightarrow H_c^*(\mathcal{M}(\Lambda(\beta_2))),$$

which is a map of (products of) Rouquier complexes $\widehat{T}_{\beta_1} \rightarrow T_{\beta_2}$, where \widehat{T}_{β_1} is the compactly supported cohomology of $\widehat{\mathcal{M}}(\Lambda(\beta_1))$, which contains the information of the compactly supported cohomology T_{β_1} of the open Bott-Samelson variety for β_1 .

Now, applying this to the two Lagrangian cobordisms associated to the trivalent vertices G_{tri} and the hexagonal vertices G_{hex} , we obtain the following two maps:

- (i) The map $\Phi_{L(G_{tri})} : T_{s_i} \otimes H^*(\mathbb{S}^1) \rightarrow T_{s_i} \otimes_R T_{s_i}$, where i labels the τ_i -edges of G_{tri} , and we have identified the fiber bundle $\widehat{\mathcal{M}}(\Lambda(\beta_1)) \cong (\Lambda(\beta_1)) \times \mathbb{C}^*$ with the Cartesian product, as in this case the bundle is topologically trivial. The fact that there is *one* copy of $\mathbb{C}^* = \mathbb{S}^1 \times \mathbb{R}$ corresponds to the fact that the Lagrangian cobordism $L(G_{tri})$ has a unique index 1 critical point and its cocore carries the data \mathbb{C}^* .
- (ii) The map $\Phi_{L(G_{hex})} : T_{s_i} \otimes T_{s_{i+1}} \otimes T_{s_i} \rightarrow T_{s_{i+1}} \otimes T_{s_i} \otimes T_{s_{i+1}}$, where in this case $\widehat{\mathcal{M}}(\Lambda(\beta_1)) \cong \mathcal{M}(\Lambda(\beta_1))$ as the Lagrangian $L(G_{hex})$ is a cylinder and $H_1(L) \cong \{0\}$ is trivial.

In conclusion, the above discussion can be summarized according to the following tenet:

Principle A.5. *Let T_{β_1}, T_{β_2} be the Rouquier complexes associated to positive braids β_1, β_2 and $\Psi : T_{\beta_1} \rightarrow T_{\beta_2}$ the morphism given by a graph G_Ψ with only (upwards) trivalent and hexagonal morphisms in (open) Soergel calculus. Then the Lagrangian projection of the Legendrian weave $\Lambda(G_\Psi)$ yields an embedded exact Lagrangian cobordism L from $\Lambda(\beta_1)$ to $\Lambda(\beta_2)$ and a geometric map*

$$\Phi_L : \widehat{\mathcal{M}}(\Lambda(\beta_1)) \rightarrow \Lambda(\beta_2)$$

such that $H_c^*(\Phi_L) = \Psi$. \square

The difference between the principle above and a theorem lies on the correct definition of *open* Soergel calculus, of which we are not aware at this stage. That said, since the trivalent and the hexagonal vertices are two of the main building blocks for closed Soergel calculus, the above construction provides a potential symplectic geometrization of *open* Soergel calculus,

associated to Rouquier complexes, instead of Soergel bimodules. In particular, in the context of *open* Bott-Samelson varieties, the Lagrangian cobordisms above indicate the need for additional data from $H_1(L) = \mathbb{Z}^{|V|}$ in specifying a morphism, where $|V|$ is the number of trivalent vertices. The development of open Soergel calculus, the computations establishing that our geometric maps induce the expected algebraic maps, as well as the Lagrangian description of the univalent vertex, will be the subject of upcoming and more algebraic work.

REFERENCES

- [AB27] J. W. Alexander and G. B. Briggs. On types of knotted curves. *Ann. of Math. (2)*, 28(1-4):562–586, 1926/27.
- [Ad75] Vladimir Igorevich Arnol’ d. Critical points of smooth functions. *Proceedings of the International Congress of Mathematicians (Vancouver, B. C., 1974)*, Vol. 1, pages 19–39, 1975.
- [Ad76] Vladimir Igorevich Arnol’ d. Wave front evolution and equivariant Morse lemma. *Comm. Pure Appl. Math.*, 29(6):557–582, 1976.
- [Ad79] Vladimir Igorevich Arnol’ d. Indexes of singular points of 1-forms on manifolds with boundary, convolutions of invariants of groups generated by reflections, and singular projections of smooth surfaces. *Uspekhi Mat. Nauk*, 34(2(206)):3–38, 1979.
- [Ad90] Vladimir Igorevich Arnol’ d. *Singularities of caustics and wave fronts*, volume 62 of *Mathematics and its Applications (Soviet Series)*. Kluwer Academic Publishers Group, Dordrecht, 1990.
- [AdG01] V. I. Arnol’ d and A. B. Givental’. Symplectic geometry [MR0842908 (88b:58044)]. In *Dynamical systems, IV*, volume 4 of *Encyclopaedia Math. Sci.*, pages 1–138. Springer, Berlin, 2001.
- [Aur07] Denis Auroux. Mirror symmetry and T -duality in the complement of an anticanonical divisor. *J. Gökova Geom. Topol. GGT*, 1:51–91, 2007.
- [Aur09] Denis Auroux. Special Lagrangian fibrations, wall-crossing, and mirror symmetry. In *Surveys in differential geometry. Vol. XIII. Geometry, analysis, and algebraic geometry: forty years of the Journal of Differential Geometry*, volume 13 of *Surv. Differ. Geom.*, pages 1–47. Int. Press, Somerville, MA, 2009.
- [AV00] Mina Aganagic and Cumrun Vafa. Mirror Symmetry, D-Branes and Counting Holomorphic Discs. 2000.
- [AV12] Mina Aganagic and Cumrun Vafa. Large N Duality, Mirror Symmetry, and a Q-deformed A-polynomial for Knots. 2012.
- [Ben83] Daniel Bennequin. Entrelacements et équations de Pfaff. In *Third Schnepfenried geometry conference, Vol. 1 (Schnepfenried, 1982)*, volume 107 of *Astérisque*, pages 87–161. Soc. Math. France, Paris, 1983.
- [Ben86] Daniel Bennequin. Caustique mystique (d’après Arnol’ d et al.). Number 133-134, pages 19–56. 1986. Seminar Bourbaki, Vol. 1984/85.
- [Bir74] Joan S. Birman. *Braids, links, and mapping class groups*. Princeton University Press, Princeton, N.J.; University of Tokyo Press, Tokyo, 1974. Annals of Mathematics Studies, No. 82.
- [BM08] J. A. Bondy and U. S. R. Murty. *Graph theory*, volume 244 of *Graduate Texts in Mathematics*. Springer, New York, 2008.
- [BS15] Tom Bridgeland and Ivan Smith. Quadratic differentials as stability conditions. *Publ. Math. Inst. Hautes Études Sci.*, 121:155–278, 2015.
- [BST15] Frédéric Bourgeois, Joshua M. Sabloff, and Lisa Traynor. Lagrangian cobordisms via generating families: construction and geography. *Algebr. Geom. Topol.*, 15(4):2439–2477, 2015.
- [Cas20] Roger Casals. Lagrangian Skeleta and Plane Curve Singularities. *J. Fixed Point Theory and Applications*, (Viterbo 60), 2020.
- [CE12] Kai Cieliebak and Yakov Eliashberg. *From Stein to Weinstein and back*, volume 59 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2012. Symplectic geometry of affine complex manifolds.
- [CE14] Kai Cieliebak and Yakov Eliashberg. Stein structures: existence and flexibility. In *Contact and symplectic topology*, volume 26 of *Bolyai Soc. Math. Stud.*, pages 357–388. János Bolyai Math. Soc., Budapest, 2014.
- [CG20] Roger Casals and Honghao Gao. Infinitely many Lagrangian fillings. *ArXiv e-prints*, 2001.01334, 2020.
- [CGGS] Roger Casals, Eugene Gorsky, Mikhail Gorsky, and José Simental. *arXiv*, 2012.06931.
- [Che02] Yuri Chekanov. Differential algebra of Legendrian links. *Invent. Math.*, 150(3):441–483, 2002.
- [CM18] Roger Casals and Emmy Murphy. Differential algebra of cubic planar graphs. *Adv. Math.*, 338:401–446, 2018.
- [CM19] Roger Casals and Emmy Murphy. Legendrian fronts for affine varieties. *Duke Math. J.*, 168(2):225–323, 2019.

- [CMP19] Roger Casals, Emmy Murphy, and Francisco Presas. Geometric criteria for overtwistedness. *J. Amer. Math. Soc.*, 32(2):563–604, 2019.
- [CN20] Roger Casals and Lenhard Ng. Braid Loops with infinite monodromy on the Legendrian contact DGA. *Arxiv e-prints*, 2020.
- [Con84] Marston D. E. Conder. Some results on quotients of triangle groups. *Bull. Austral. Math. Soc.*, 30(1):73–90, 1984.
- [DGG16] Tudor Dimofte, Maxime Gabella, and Alexander B. Goncharov. K-decompositions and 3d gauge theories. *J. High Energy Phys.*, (11):151, front matter+144, 2016.
- [Die17] Reinhard Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Berlin, fifth edition, 2017.
- [DR11] Georgios Dimitroglou Rizell. Knotted Legendrian surfaces with few Reeb chords. *Algebr. Geom. Topol.*, 11(5):2903–2936, 2011.
- [DR16] Georgios Dimitroglou Rizell. Legendrian ambient surgery and Legendrian contact homology. *J. Symplectic Geom.*, 14(3):811–901, 2016.
- [EENS13a] Tobias Ekholm, John Etnyre, Lenhard Ng, and Michael Sullivan. Filtrations on the knot contact homology of transverse knots. *Math. Ann.*, 355(4):1561–1591, 2013.
- [EENS13b] Tobias Ekholm, John B. Etnyre, Lenhard Ng, and Michael G. Sullivan. Knot contact homology. *Geom. Topol.*, 17(2):975–1112, 2013.
- [EES05a] Tobias Ekholm, John Etnyre, and Michael Sullivan. The contact homology of Legendrian submanifolds in \mathbb{R}^{2n+1} . *J. Differential Geom.*, 71(2):177–305, 2005.
- [EES05b] Tobias Ekholm, John Etnyre, and Michael Sullivan. Non-isotopic Legendrian submanifolds in \mathbb{R}^{2n+1} . *J. Differential Geom.*, 71(1):85–128, 2005.
- [EGH00] Y. Eliashberg, A. Givental, and H. Hofer. Introduction to symplectic field theory. *Geom. Funct. Anal.*, (Special Volume, Part II):560–673, 2000. GAFA 2000 (Tel Aviv, 1999).
- [EHK16] Tobias Ekholm, Ko Honda, and Tamás Kálmán. Legendrian knots and exact Lagrangian cobordisms. *J. Eur. Math. Soc. (JEMS)*, 18(11):2627–2689, 2016.
- [EK10] Ben Elias and Mikhail Khovanov. Diagrammatics for Soergel categories. *Int. J. Math. Math. Sci.*, pages Art. ID 978635, 58, 2010.
- [Eli90] Yakov Eliashberg. Topological characterization of Stein manifolds of dimension > 2 . *Internat. J. Math.*, 1(1):29–46, 1990.
- [Eli93] Yakov Eliashberg. Legendrian and transversal knots in tight contact 3-manifolds. In *Topological methods in modern mathematics (Stony Brook, NY, 1991)*, pages 171–193. Publish or Perish, Houston, TX, 1993.
- [EM02] Y. Eliashberg and N. Mishachev. *Introduction to the h-principle*, volume 48 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2002.
- [ENS18] Tobias Ekholm, Lenhard Ng, and Vivek Shende. A complete knot invariant from contact homology. *Invent. Math.*, 211(3):1149–1200, 2018.
- [Etn05] John B. Etnyre. Legendrian and transversal knots. In *Handbook of knot theory*, pages 105–185. Elsevier B. V., Amsterdam, 2005.
- [EV18] John Etnyre and Vera Vértesi. Legendrian satellites. *Int. Math. Res. Not. IMRN*, (23):7241–7304, 2018.
- [EW16] Ben Elias and Geordie Williamson. Soergel calculus. *Represent. Theory*, 20:295–374, 2016.
- [FG06a] V. V. Fock and A. B. Goncharov. Cluster x-varieties, amalgamation, and Poisson-Lie groups. In *Algebraic geometry and number theory*, volume 253 of *Progr. Math.*, pages 27–68. Birkhäuser Boston, Boston, MA, 2006.
- [FG06b] Vladimir Fock and Alexander Goncharov. Moduli spaces of local systems and higher Teichmüller theory. *Publ. Math. Inst. Hautes Études Sci.*, (103):1–211, 2006.
- [FOOO09] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono. *Lagrangian intersection Floer theory: anomaly and obstruction. Part II*, volume 46 of *AMS/IP Studies in Advanced Mathematics*. American Mathematical Society, Providence, RI; International Press, Somerville, MA, 2009.
- [FZ02] Sergey Fomin and Andrei Zelevinsky. Cluster algebras. I. Foundations. *J. Amer. Math. Soc.*, 15(2):497–529, 2002.
- [Gei08] Hansjörg Geiges. *An introduction to contact topology*, volume 109 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2008.
- [GK13] Alexander B. Goncharov and Richard Kenyon. Dimers and cluster integrable systems. *Ann. Sci. Éc. Norm. Supér. (4)*, 46(5):747–813, 2013.
- [GKS12] Stéphane Guillermou, Masaki Kashiwara, and Pierre Schapira. Sheaf quantization of Hamiltonian isotopies and applications to nondisplaceability problems. *Duke Math. J.*, 161(2):201–245, 2012.
- [GLPY17] Maxime Gabella, Pietro Longhi, Chan Y. Park, and Masahito Yamazaki. BPS graphs: from spectral networks to BPS quivers. *J. High Energy Phys.*, (7):032, front matter+47, 2017.
- [GMN10] Davide Gaiotto, Gregory W. Moore, and Andrew Neitzke. Four-dimensional wall-crossing via three-dimensional field theory. *Comm. Math. Phys.*, 299(1):163–224, 2010.

- [GMN13] Davide Gaiotto, Gregory W. Moore, and Andrew Neitzke. Spectral networks. *Ann. Henri Poincaré*, 14(7):1643–1731, 2013.
- [GMN14] Davide Gaiotto, Gregory W. Moore, and Andrew Neitzke. Spectral networks and snakes. *Ann. Henri Poincaré*, 15(1):61–141, 2014.
- [Gom98] Robert E. Gompf. Handlebody construction of Stein surfaces. *Ann. of Math. (2)*, 148(2):619–693, 1998.
- [Gon17] A. B. Goncharov. Ideal webs, moduli spaces of local systems, and 3d Calabi-Yau categories. In *Algebra, geometry, and physics in the 21st century*, volume 324 of *Progr. Math.*, pages 31–97. Birkhäuser/Springer, Cham, 2017.
- [GPS19a] Sheel Ganatra, John Pardon, and Vivek Shende. Microlocal Morse theory of wrapped Fukaya categories. 2019.
- [GPS19b] Sheel Ganatra, John Pardon, and Vivek Shende. Covariantly functorial wrapped Floer theory on Liouville sectors. *Publ. Math. Inst. Hautes Études Sci. (to appear)*, 2019.
- [GPS19c] Sheel Ganatra, John Pardon, and Vivek Shende. Sectorial descent for wrapped Fukaya categories. 2019.
- [Gro86] Mikhael Gromov. *Partial differential relations*, volume 9 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3)*. Springer-Verlag, Berlin, 1986.
- [GS14] Stéphane Guillermou and Pierre Schapira. Microlocal theory of sheaves and Tamarkin’s non displacement theorem. In *Homological mirror symmetry and tropical geometry*, volume 15 of *Lect. Notes Unione Mat. Ital.*, pages 43–85. Springer, Cham, 2014.
- [GS18] Alexander Goncharov and Linhui Shen. Donaldson-Thomas transformations of moduli spaces of G-local systems. *Adv. Math.*, 327:225–348, 2018.
- [GSW20a] Honghao Gao, Linhui Shen, and Daping Weng. Augmentations, Fillings, and Clusters. *ArXiv e-prints*, 2008.10793, 2020.
- [GSW20b] Honghao Gao, Linhui Shen, and Daping Weng. Positive Braid Links with Infinitely Many Fillings. *ArXiv e-prints*, 2009.00499, 2020.
- [Hur92] A. Hurwitz. Ueber algebraische Gebilde mit eindeutigen Transformationen in sich. *Math. Ann.*, 41(3):403–442, 1892.
- [K05] Tamás Kálmán. Contact homology and one parameter families of Legendrian knots. *Geom. Topol.*, 9:2013–2078, 2005.
- [Kel94] Bernhard Keller. Deriving DG categories. *Ann. Sci. École Norm. Sup. (4)*, 27(1):63–102, 1994.
- [Kho07] Mikhail Khovanov. Triply-graded link homology and Hochschild homology of Soergel bimodules. *Internat. J. Math.*, 18(8):869–885, 2007.
- [KS85] Masaki Kashiwara and Pierre Schapira. Microlocal study of sheaves. *Astérisque*, (128):235, 1985. Corrections to this article can be found in *Astérisque* No. 130, p. 209.
- [KS90] Masaki Kashiwara and Pierre Schapira. *Sheaves on manifolds*, volume 292 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1990. With a chapter in French by Christian Houzel.
- [KS10] Maxim Kontsevich and Yan Soibelman. Motivic Donaldson-Thomas invariants: summary of results. In *Mirror symmetry and tropical geometry*, volume 527 of *Contemp. Math.*, pages 55–89. Amer. Math. Soc., Providence, RI, 2010.
- [KS14] Maxim Kontsevich and Yan Soibelman. Wall-crossing structures in Donaldson-Thomas invariants, integrable systems and mirror symmetry. In *Homological mirror symmetry and tropical geometry*, volume 15 of *Lect. Notes Unione Mat. Ital.*, pages 197–308. Springer, Cham, 2014.
- [Kuw20] Tatsuki Kuwagaki. Sheaf quantization from exact WKB analysis. 2020.
- [Law17] John W. Lawson. Minimal mutation-infinite quivers. *Exp. Math.*, 26(3):308–323, 2017.
- [LO08] Yves Laszlo and Martin Olsson. The six operations for sheaves on Artin stacks. I. Finite coefficients. *Publ. Math. Inst. Hautes Études Sci.*, (107):109–168, 2008.
- [LO09] Yves Laszlo and Martin Olsson. Perverse t -structure on Artin stacks. *Math. Z.*, 261(4):737–748, 2009.
- [LO10] Valery A. Lunts and Dmitri O. Orlov. Uniqueness of enhancement for triangulated categories. *J. Amer. Math. Soc.*, 23(3):853–908, 2010.
- [LT99] A. Lucchini and M. C. Tamburini. Classical groups of large rank as Hurwitz groups. *J. Algebra*, 219(2):531–546, 1999.
- [Mar35] A.A. Markov. über die freie aquivalenz der geschlossen zöpfe. *Recueil Math. Moscou*, 1:73–78, 1935.
- [MGOT12] Sophie Morier-Genoud, Valentin Ovsienko, and Serge Tabachnikov. 2-frieze patterns and the cluster structure of the space of polygons. *Annales de l’Institut Fourier*, 62(3):937–987, 2012.
- [Mur12] Emmy Murphy. Loose Legendrian Embeddings in High Dimensional Contact Manifolds. 2012.
- [Nad09] David Nadler. Microlocal branes are constructible sheaves. *Selecta Math. (N.S.)*, 15(4):563–619, 2009.

- [Nad17a] David Nadler. Arboreal singularities. *Geom. Topol.*, 21(2):1231–1274, 2017.
- [Nad17b] David Nadler. A combinatorial calculation of the Landau-Ginzburg model $M = \mathbb{C}^3$, $W = z_1 z_2 z_3$. *Selecta Math. (N.S.)*, 23(1):519–532, 2017.
- [Nei14] Andrew Neitzke. Cluster-like coordinates in supersymmetric quantum field theory. *Proc. Natl. Acad. Sci. USA*, 111(27):9717–9724, 2014.
- [Ng03] Lenhard L. Ng. Computable Legendrian invariants. *Topology*, 42(1):55–82, 2003.
- [Ng11] Lenhard Ng. Combinatorial knot contact homology and transverse knots. *Adv. Math.*, 227(6):2189–2219, 2011.
- [NR13] Lenhard Ng and Daniel Rutherford. Satellites of Legendrian knots and representations of the Chekanov-Eliashberg algebra. *Algebr. Geom. Topol.*, 13(5):3047–3097, 2013.
- [NRS⁺15] Lenhard Ng, Dan Rutherford, Vivek Shende, Steven Sivek, and Eric Zaslow. Augmentations are Sheaves. *To appear in Geom. Top.*, 2015.
- [NZ09] David Nadler and Eric Zaslow. Constructible sheaves and the Fukaya category. *J. Amer. Math. Soc.*, 22(1):233–286, 2009.
- [OST13] Valentin Ovsienko, Richard Evan Schwartz, and Serge Tabachnikov. Liouville–arnold integrability of the pentagram map on closed polygons. *Duke Math. J.*, 162(12):2149–2196, 2013.
- [Pal15] Frederic Palesi. Introduction to positive representations and Fock-Goncharov Coordinates. 2015.
- [Pan17a] Yu Pan. *Augmentations and Exact Lagrangian Cobordisms*. ProQuest LLC, Ann Arbor, MI, 2017. Thesis (Ph.D.)—Duke University.
- [Pan17b] Yu Pan. Exact Lagrangian fillings of Legendrian $(2, n)$ torus links. *Pacific J. Math.*, 289(2):417–441, 2017.
- [Pol91] L. Polterovich. The surgery of Lagrange submanifolds. *Geom. Funct. Anal.*, 1(2):198–210, 1991.
- [PS97] V. V. Prasolov and A. B. Sossinsky. *Knots, links, braids and 3-manifolds*, volume 154 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 1997.
- [Rei27] Kurt Reidemeister. Elementare Begründung der Knotentheorie. *Abh. Math. Sem. Univ. Hamburg*, 5(1):24–32, 1927.
- [Rol76] Dale Rolfsen. *Knots and links*. Publish or Perish Inc., Berkeley, Calif., 1976. Mathematics Lecture Series, No. 7.
- [Ros98] Dennis Roseman. Reidemeister-type moves for surfaces in four-dimensional space. In *Knot theory (Warsaw, 1995)*, volume 42 of *Banach Center Publ.*, pages 347–380. Polish Acad. Sci. Inst. Math., Warsaw, 1998.
- [Rou06] Raphaël Rouquier. Categorification of sl_2 and braid groups. In *Trends in representation theory of algebras and related topics*, volume 406 of *Contemp. Math.*, pages 137–167. Amer. Math. Soc., Providence, RI, 2006.
- [RS19a] Daniel Rutherford and Michael Sullivan. Cellular Legendrian contact homology for surfaces, part II. *Internat. J. Math.*, 30(7):1950036, 135, 2019.
- [RS19b] Daniel Rutherford and Michael Sullivan. Cellular Legendrian contact homology for surfaces, part III. *Internat. J. Math.*, 30(7):1950037, 111, 2019.
- [RSTZ14] Helge Ruddat, Nicolò Sibilla, David Treumann, and Eric Zaslow. Skeleta of affine hypersurfaces. *Geom. Topol.*, 18(3):1343–1395, 2014.
- [Sch53] Horst Schubert. Knoten und Vollringe. *Acta Math.*, 90:131–286, 1953.
- [She19] Vivek Shende. The conormal torus is a complete knot invariant. *Forum Math. Pi*, 7:e6, 16, 2019.
- [Siv11] Steven Sivek. A bordered Chekanov-Eliashberg algebra. *J. Topol.*, 4(1):73–104, 2011.
- [Smi15] Ivan Smith. Quiver algebras as Fukaya categories. *Geom. Topol.*, 19(5):2557–2617, 2015.
- [Soe07] Wolfgang Soergel. Kazhdan-Lusztig-Polynome und unzerlegbare Bimoduln über Polynomringen. *J. Inst. Math. Jussieu*, 6(3):501–525, 2007.
- [SS16] Joshua M. Sabloff and Michael G. Sullivan. Families of Legendrian submanifolds via generating families. *Quantum Topol.*, 7(4):639–668, 2016.
- [Sta18] Laura Starkston. Arboreal singularities in Weinstein skeleta. *Selecta Math. (N.S.)*, 24(5):4105–4140, 2018.
- [STW16] Vivek Shende, David Treumann, and Harold Williams. On the combinatorics of exact Lagrangian surfaces. 2016.
- [STWZ19] Vivek Shende, David Treumann, Harold Williams, and Eric Zaslow. Cluster varieties from Legendrian knots. *Duke Math. J.*, 168(15):2801–2871, 2019.
- [STZ17] Vivek Shende, David Treumann, and Eric Zaslow. Legendrian knots and constructible sheaves. *Invent. Math.*, 207(3):1031–1133, 2017.
- [Tab05] Goncalo Tabuada. Une structure de catégorie de modèles de Quillen sur la catégorie des dg-catégories. *C. R. Math. Acad. Sci. Paris*, 340(1):15–19, 2005.
- [Toe07] Bertrand Toen. The homotopy theory of dg-categories and derived morita theory. *Invent. Math.*, 167(3):615–667, 2007.
- [Tri19] Minh-Tam Q. Trinh. Annular Homology of Artin Braids I. 2019.

- [TZ18] David Treumann and Eric Zaslow. Cubic planar graphs and Legendrian surface theory. *Adv. Theor. Math. Phys.*, 22(5):1289–1345, 2018.
- [Via14] Renato Vianna. On exotic Lagrangian tori in $\mathbb{C}\mathbb{P}^2$. *Geom. Topol.*, 18(4):2419–2476, 2014.
- [Wei71] Alan Weinstein. Symplectic manifolds and their Lagrangian submanifolds. *Advances in Math.*, 6:329–346 (1971), 1971.

UNIVERSITY OF CALIFORNIA DAVIS, DEPT. OF MATHEMATICS, SHIELDS AVENUE, DAVIS, CA 95616, USA
Email address: `casals@math.ucdavis.edu`

NORTHWESTERN UNIVERSITY, DEPARTMENT OF MATHEMATICS, EVANSTON, IL 60208-2730, USA
Email address: `zaslow@math.northwestern.edu`