Microlocal Theory of Legendrian Links and Cluster Algebras

ROGER CASALS AND DAPING WENG

ABSTRACT. We show the existence of quasi-cluster $\mathcal{A}$-structures and cluster Poisson structures on moduli stacks of sheaves with singular support in the alternating strand diagram of grid plabic graphs by studying the microlocal parallel transport of sheaf quantizations of Lagrangian fillings of Legendrian links. The construction is in terms of contact and symplectic topology, showing that there exists an initial seed associated to a canonical relative Lagrangian skeleton. In particular, mutable cluster $\mathcal{A}$-variables are intrinsically characterized via the symplectic topology of Lagrangian fillings in terms of dually $\mathbb{L}$-compressible cycles. New ingredients are introduced throughout this work, including the initial weave associated to a grid plabic graph, cluster mutation along a non-square face of a plabic graph, the concept of the sugar-free hull, and the notion of microlocal merodromy. Finally, a contact geometric realization of the DT-transformation is constructed for shuffle graphs, proving cluster duality for the cluster ensembles.

Més lluny, heu d'anar més lluny
dels arbres caiguts que ara us empresonen,
i quan els haureu guanyat
tingueu ben present no aturar-vos.

L. Llach, *Viatge a Ítaca*

Contents

1. Introduction

The object of this article will be to show the existence of intrinsically symplectic quasi-cluster K_2-structures and quasi-cluster Poisson structures on moduli stacks of sheaves with singular support in the alternating strand diagram of a complete grid plabic graph. The construction of such quasi-cluster structures is achieved via contact and symplectic topology, based on the recently developed machinery of Legendrian weaves, and we show that there exists a canonical initial quasi-cluster seed associated to a relative Lagrangian skeleton. This is the first manuscript proving the existence of such cluster structures for these general moduli stacks, and entirely in symplectic geometric terms, as well as introducing the first symplectic topological definition of cluster $\mathcal{A}$-variables associated to Lagrangian fillings of Legendrian links. In particular, our constructions admit natural contact and symplectic invariance and functoriality properties, and the cluster variables can be named and computed after performing Hamiltonian isotopies.

Several new ingredients are introduced for this purpose, among them are the initial weave of a grid plabic graph, cluster mutations along non-square faces, the concept of sugar-free hulls, and the notion of microlocal merodromy. Microlocal merodromies capture microlocal

parallel transport along a relative cycle and they are crucial in defining a set of initial cluster $\mathcal{A}$-variables. From a contact geometry viewpoint, the construction (and obstruction) of embedded Lagrangian disks whose boundaries lie on an embedded exact Lagrangian filling have a central role. This allows for geometric characterizations of mutable and frozen vertices, which arise from relative homology groups of triples, and naturally explains the appearance of quasi-cluster structures.
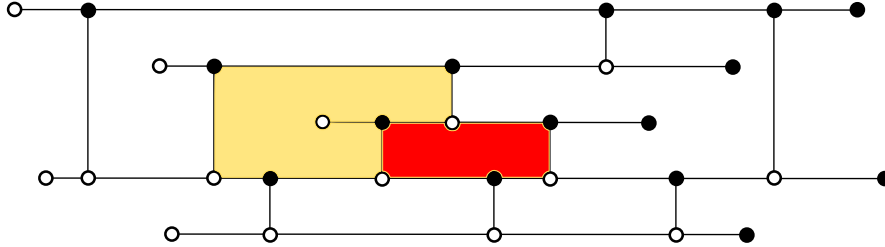

FIGURE 1. The quasi-cluster K_2-structure we construct for this grid plabic graph is on the coordinate ring of the moduli of decorated sheaves on $\mathbb{R}^2$ with singular support in a max-tb Legendrian representative of the $m(9_6)$ knot.

1.1. **Scientific Context.** Cluster algebras, first introduced by S. Fomin and A. Zelevinsky [FZ02, FZ03, BFZ05] in the context of Lie theory, are commutative rings endowed with a set of distinguished generators that have remarkable combinatorial structures. Cluster varieties, a geometric enrichment of cluster algebras introduced by V. Fock and A. Goncharov [FG06b, FG06a], are affine varieties equipped with an atlas of torus charts whose transition maps obey certain combinatorial rules. Cluster varieties come in a dual pairs consisting of a cluster K_2-variety, also known as a cluster $\mathcal{A}$-variety, and a cluster Poisson variety, also known as a cluster $\mathcal{X}$-variety. In particular, the coordinate ring of a cluster $\mathcal{A}$-variety coincides with an upper cluster algebra, see Berenstein-Fomin-Zelevinsky [BFZ05].

Since their introduction, cluster algebras and cluster varieties have appeared in many contexts, such as Teichmüller theory [FG06b, FST08, GSV05], birational geometry [GHK15, GHKK18, HK18], the Riemann-Hilbert correspondence [All21, Nei14, GMN10], exact WKB analysis [IN14, IN16], and the study of positroid and Richardson varieties [GL19, SSBW19], among others. The first appearance of cluster mutations in symplectic geometry occurred in the study of wall-crossing formulas, following the work of D. Auroux, K. Fukaya, M. Kontsevich, P. Seidel, Y. Soibelman and others, e.g. see [Aur07, Aur09, PT20] and references therein. We also thank A. Goncharov for pointing out to us his recent work with M. Kontsevich [GK21] focusing on non-commutative clusters, which also aligns well with the developments we present here.

The first hint that cluster $\mathcal{X}$-structures might naturally exist in the symplectic study of Legendrian knots was provided in [STWZ19], where it was computed how certain absolute monodromies around a square plabic face change under a square move. See also the generalization presented in [STW16]. Although these monodromies are candidates for cluster $\mathcal{X}$-variables, there are nevertheless two obstacles to actually prove the existence of a cluster $\mathcal{A}$-structure. First, many plabic faces are typically not square and may contain lollipops; thus, one needs a new construction that both associates a cluster $\mathcal{X}$-variable to them and allows for a geometric mutation to be performed. Second, more fundamental, the regularity problem: even if all faces are square, these absolute monodromies are *not* global regular functions, and it is not possible to deduce the existence of a cluster structure purely from these microlocal monodromies. These obstacles are unavoidable if one is either restricted to plabic graphs or absolute cycles, both of which are limiting constraints in that approach.

Moreover, in either case, there is still missing an entire cluster K_2-structure, i.e. no upper cluster algebra is constructed, which is a more central object (see Remark 1.3).

Our new approach uses Legendrian weaves, which are more versatile than plabic graphs, and actually builds cluster $\mathcal{A}$-variables from relative cycles, which is stronger than the absolute analogue. In particular, we overcome both obstacles above, resolving the regularity problem, and finally prove the existence of cluster $\mathcal{A}$–structures and, consequently, cluster $\mathcal{X}$-structures in entirely symplectic topological terms. Some of our previous work has been using ideas from the theory of cluster algebras for new applications to contact and symplectic geometry, see e.g. [Cas22, CGGS20, CGGS21, CZ22, GSW20a], including the discovery of infinitely many Lagrangian fillings for many Legendrian links [CG22]. This article builds in the converse direction, using contact and symplectic topology to construct (upper) cluster algebras, and using symplectic topological results to deduce algebraic properties.[1]

Note that what can be deduced from our previous works [CGGS20, CG22, CZ22, GSW20a, GSW20b] is that certain moduli spaces that appear in contact topology are sometimes abstractly isomorphic to certain affine varieties, which themselves can independently be endowed[2] with cluster structures, but currently there does not exist any symplectic construction or characterization of cluster $\mathcal{A}$-variables or general cluster $\mathcal{X}$-variables, nor a symplectic geometric proof of the existence of cluster structures on these moduli spaces, nor even a geometric understanding of frozen variables. In particular, none of these previous constructions is known to have any Hamiltonian or Legendrian invariance properties, which are crucial in contact and symplectic topology. In fact, in all previous constructions even the initial seeds cannot be named after a Hamiltonian isotopy (e.g. even after a Reidemeister I or II move) and no symplectic computation or interpretation of cluster $\mathcal{A}$-variables existed. The present work finally resolves this matter and, as we shall see, interesting symplectic features appear with regards to both mutable and frozen variables.

1.2. **Main Results.** Let $\Lambda \subset (T^*_\infty \mathbb{R}^2, \xi_{st})$ be a Legendrian link in the ideal contact boundary of the cotangent bundle of the plane $\mathbb{R}^2$, and $T \subset \Lambda$ a set of marked points. The precise details and definitions for these contact geometric objects are provided in Section 2. Let $L \subset (T^*\mathbb{R}^2, \lambda_{st})$ be an embedded exact Lagrangian filling of Λ. By definition, an embedded closed curve $\gamma \subset L$ is said to be $\mathbb{L}$-compressible if there exists a properly embedded Lagrangian 2-disk $D \subset (T^*\mathbb{R}^2 \setminus L)$ such that $\partial \overline{D} \cap L = \gamma \subset \mathbb{R}^4$. A collection $\{\gamma_1, \ldots, \gamma_\ell\}$ of such curves, with a choice of $\mathbb{L}$-compressing disk for each curve, is said to be an $\mathbb{L}$-compressing system for L if the curves form a maximal linearly independent subset in $H_1(L)$. In line with this, we will use Lagrangian disk surgeries, as defined in [Pol91, Yau17].

Consider also the moduli stack $\mathfrak{M}(\Lambda, T)$ of decorated microlocal rank-one constructible sheaves on $\mathbb{R}^2$ with singular support contained in Λ, as defined in Subsection 2.8.3, following [KS90, GKS12a], which is invariant under contact isotopies. Let $\mathbb{G} \subset \mathbb{R}^2$ be a complete grid plabic graph and $\Lambda = \Lambda(\mathbb{G}) \subset T^*_\infty \mathbb{R}^2$ its associated Legendrian link, as defined in Section 2. See Subsection 2.3 for the definition of the sugar-free hull $\mathbb{S}_f$ of a face f in $\mathbb{G}$ and Subsection 4.8 for completeness. Note that the concept of sugar-free hulls, and whether a region is sugar-free, only depends on the behaviour at non-convex corners, see Definition 2.6.

The main result of the article, stated in Theorem 1.1, is the existence and explicit symplectic construction of a quasi-cluster $\mathcal{A}$-structure on $\mathfrak{M}(\Lambda, T)$. In particular, the cluster $\mathcal{A}$-variables of the initial seed as well as all the once-mutated seeds are obtained by a new microlocal parallel transport along certain relative cycles on exact Lagrangian fillings of Λ. This microlocal parallel transport is associated to a sheaf quantization of each exact Lagrangian

---

[1] In fact, we can show that $\mathcal{A} = \mathcal{U}$ for these cluster varieties: see upcoming work of the first author and collaborators.

[2] Explicitly, double Bott-Samelson cells for [GSW20a], and positroids for [CG22, STWZ19]. These instances are, in any case, particular cases of the moduli stacks that we associate to grid plabic graphs.

4

filling, following [GKS12a, CZ22], and we refer to it as a *microlocal merodromy*, see Section 4. The central result of the manuscript is stated as follows.

Theorem 1.1 (Main Result). *Let $\mathbb{G} \subset \mathbb{R}^2$ be a complete grid plabic graph, $\Lambda = \Lambda(\mathbb{G}) \subset (\mathbb{R}^3, \xi_{st})$ its associated Legendrian link, $T \subset \Lambda$ a set of marked points, with at least one marked point per component of Λ, and $\mathfrak{M}(\Lambda, T)$ the stack of decorated microlocal rank-one constructible sheaves on $\mathbb{R}^2$ with singular support contained in Λ.*

Then, there exists a canonical embedded exact Lagrangian filling $L = L(\mathbb{G}) \subset (\mathbb{R}^4, \omega_{st})$ of Λ and a canonical $\mathbb{L}$-compressing system $\mathfrak{S} = \{\gamma_1, \ldots, \gamma_\ell\}$ for L, indexed by the sugar-free hulls of $\mathbb{G}$, such that for any completion of $\mathfrak{S}$ into a basis $\mathfrak{B}$ of $H_1(L, T)$ the following hold:

(i) *The microlocal merodromies A_{η_i}, defined on (and by using) the open chart $(\mathbb{C}^\times)^{b_1(L,T)} \subset \mathfrak{M}(\Lambda, T)$ associated to L, extend to global regular functions*

$$A_{\eta_i} : \mathfrak{M}(\Lambda, T) \longrightarrow \mathbb{C}, \quad i.e. \ A_{\eta_i} \in \mathcal{O}(\mathfrak{M}(\Lambda, T)),$$

where $\mathfrak{B}^\vee = \{\eta_1, \ldots, \eta_s\}$ is the dual basis in $H_1(L \setminus T, \Lambda \setminus T)$.

(ii) *The microlocal merodromies $\{A_{\eta_1}, \ldots, A_{\eta_\ell}\}$ associated to the relative cycles that are dual to an $\mathbb{L}$-compressible absolute cycle in $\mathfrak{S}$ are irreducible functions in $\mathcal{O}(\mathfrak{M}(\Lambda, T))$, whereas the merodromies $\{A_{\eta_{\ell+1}}, \ldots, A_{\eta_{b_1(L,T)}}\}$ are non-vanishing functions, i.e. units in $\mathcal{O}(\mathfrak{M}(\Lambda, T))$.*

(iii) *Let $L'_k \subset (\mathbb{R}^4, \omega_{st})$ be the Lagrangian filling obtained via Lagrangian disk surgery on L at the $\mathbb{L}$-compressing disk for $\gamma_k \in \mathfrak{S}$, and $\eta'_k \in H_1(L'_k \setminus T, \Lambda \setminus T)$ the image of η_k under the surgery. Then the merodromy $A_{\eta'_k}$ extends to a global regular function*

$$A_{\eta'_k} : \mathfrak{M}(\Lambda, T) \longrightarrow \mathbb{C}, \quad i.e. \ A_{\eta'_k} \in \mathcal{O}(\mathfrak{M}(\Lambda, T)),$$

and satisfies the cluster $\mathcal{A}$-mutation formula

$$A_{\eta'_k} A_{\eta_k} = \prod_{\eta_i \to \eta_k} A_{\eta_i} + \prod_{\eta_k \to \eta_j} A_{\eta_j}$$

with respect to the intersection quiver $Q(\mathfrak{B})$ of the basis elements $\mathfrak{B} \subset H_1(L, T)$.

Finally, the moduli variety $\mathfrak{M}(\Lambda, T)$ admits a cluster $\mathcal{A}$-structure with quiver $Q(\mathfrak{B})$ in the initial seed associated to the Lagrangian filling L, where the mutable vertices (dually) correspond to the absolute cycles in the $\mathbb{L}$-compressing system $\mathfrak{S}$ for L. Furthermore, different choices of completion of $\mathfrak{S}$ into a basis $\mathfrak{B}$ give rise to quasi-equivalent cluster $\mathcal{A}$-structures.

The grid plabic graph $\mathbb{G}$ actually provides several natural completions of the $\mathbb{L}$-compressing system $\mathfrak{S}$ to a basis $\mathfrak{B}$, as explained in Section 3. The canonical exact Lagrangian filling $L = L(\mathbb{G})$ associated with $\mathbb{G}$ is obtained as the Lagrangian projection of the Legendrian surface whose front is given by the weave $\mathfrak{w}(\mathbb{G})$ associated with $\mathbb{G}$, which is constructed in Section 3. The weave $\mathfrak{w}(\mathbb{G})$ is used crucially in the argument so as to obtain a sheaf quantization of $L(\mathbb{G})$ and prove Items (i) through (iii) as required. In addition to the existence of the cluster $\mathcal{A}$-structures on $\mathfrak{M}(\Lambda, T)$, another upshot of Theorem 1.1 is that the initial and the once-mutated cluster $\mathcal{A}$-variables can be named entirely in terms of symplectic topology, in an intrinsic and geometric manner. The resulting quasi-cluster $\mathcal{A}$-structure and these $\mathcal{A}$-variables can be equally considered and computed after a Hamiltonian isotopy.

In terms of the dichotomy between geometry and algebra, Theorem 1.1 shows that the ring $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ behaves *as if* it were always possible to perform an arbitrary sequence of Lagrangian disk surgeries starting at $L(\mathbb{G})$ with the curve configuration from the $\mathbb{L}$-compressing system $\mathfrak{S}$. It is known that geometric obstructions to further surger the Lagrangian skeleton can arise as one performs a series of Lagrangian surgeries (geometric mutations), e.g. through

the appearance of immersed curves, or algebraic intersection numbers differing from geometric ones, and yet the existence of the cluster $\mathcal{A}$-structure built in Theorem 1.1 shows that it is not possible to detect such obstructions by studying $\mathcal{O}(\mathfrak{M}(\Lambda, T))$. The following table schematically relates different ingredients involved in the proof of Theorem 1.1:

Ingredients in the symplectic construction of upper cluster algebra for $\mathcal{O}(\mathfrak{M}(\Lambda, T))$		
Grid Plabic Graph $\mathbb{G}$	**Symplectic Topology in** $T^*\mathbb{R}^2$	**Cluster Theory**
Alternating strand diagram (with marked points T)	Legendrian Link $\Lambda \subseteq T^*_\infty \mathbb{R}^2$ (with marked points T)	D^--stack $\mathfrak{M}(\Lambda, T)$ from dg-category $\mathrm{Sh}_\Lambda(\mathbb{R}^2)$
Goncharov-Kenyon conjugate surface associated to $\mathbb{G}$	Weave for Lagrangian filling L ($\Longrightarrow$ Sheaf quantization $\mathcal{F}(L)$)	Open toric chart $T_L = (\mathbb{C}^\times)^{b_1(L)} \subseteq \mathfrak{M}(\Lambda, T)$
Sugar-free hull of $\mathbb{G}$	$\mathbb{L}$**-compressible curve** $\gamma \subseteq L$ with dual relative cycle $[\eta] \in N = H_1(L \setminus T, \Lambda \setminus T)$	T_L-coordinate that extends to a **global regular** function $A_\eta : \mathfrak{M}(\Lambda, T) \longrightarrow \mathbb{C}$
Set S of sugar-free hulls	Mutable sublattice $\mathbb{Z}^{|S|} \subseteq N$	Mutable variables $\{A_\eta\}$ in T_L
Non sugar-free region of $\mathbb{G}$ (e.g. a non sugar-free face)	Immersed curve $\vartheta \subseteq L$ with dual relative cycle ϕ in N (ϑ represented by immersed Y-tree in weave)	T_L-coordinate extending to *non-vanishing* global regular function $A_\phi : \mathfrak{M}(\Lambda, T) \longrightarrow \mathbb{C}$
Subset of non s.-f. regions chosen via Hasse diagram (different choices allowed)	Sublattice $\mathbb{Z}^{b_1(L)-|S|} \subseteq N$ complement to sublattice $\mathbb{Z}^{|S|}$ (different complements)	Frozen variables $\{A_\phi\}$ in T_L (quasi-cluster equivalent)
Intersection form on absolute H_1 of conjugate surface	Intersection form on $M = H_1(L, T)$ (and thus on dual $N = M^*$)	Quiver $Q(\{A_\eta\}, \{A_\phi\})$ for T_L (different from naive $Q(\mathbb{G})$)
"Mutation" at sugar-free hull (not necessarily a square face, result often not plabic graph but represented by weave)	**Lagrangian surgery** $L' = \mu_\gamma(L)$ and relative cycle $\eta' = \mu_\gamma(\eta)$ in L'. Sheaf quantization $\mathcal{F}(L')$ via **weave mutation** at Y-tree for γ.	$T_{L'}$-coordinate extending to **global regular** function $A_{\eta'} : \mathfrak{M}(\Lambda, T) \longrightarrow \mathbb{C}$ given by **cluster $\mathcal{A}$-mutation** at η

There are several items from Theorem 1.1 that can be helpful to unpack. First, by a modification of the Guillermou-Jin-Treumann map [JT17], the Lagrangian filling L yields an open toric chart $(\mathbb{C}^\times)^{b_1} \subset \mathfrak{M}(\Lambda, T)$, where $b_1 = \mathrm{rk}(H_1(L \setminus T, \Lambda \setminus T)) = \mathrm{rk}(H_1(L, T))$. The group $H^1(L; \mathbb{C}^\times) = Hom(H_1(L; \mathbb{Z}), \mathrm{GL}_1(\mathbb{C}))$ accounts for the $\mathbb{C}^\times$-local systems on $L(\mathbb{G})$, and the modification accounts for the relative piece given by the marked points T; see Section 2.8 for details. By construction, microlocal merodromies are *a priori* functions on this particular chart $(\mathbb{C}^\times)^{b_1}$, and they visibly depend on L. In fact, in many cases they are (restrictions of) rational functions with non-trivial denominators and do not extend to global regular functions. Nevertheless, Theorem 1.1 shows that, remarkably, there is a particular set of such functions, indexed by a basis completion of the $\mathbb{L}$-compressing system $\mathfrak{S}$, whose elements extend to regular functions from $(\mathbb{C}^\times)^{b_1}$ to the entire moduli $\mathfrak{M}(\Lambda, T)$.

Second, the frozen cluster $\mathcal{A}$-variables in Theorem 1.1 have two geometric, markedly distinct, origins: absolute cycles in $H_1(L)$, and relative cycles with endpoints in T, which are themselves not dual to any absolute cycle. The appearance of the former type of frozen variables, associated to absolute cycles, is an entirely new phenomenon, starting the study of $\mathbb{L}$-(in)compressible curves in Lagrangian fillings. (E.g. we show that a Chekanov $m(5_2)$ already displays such features.) At least to date, all known instances of frozen variables of geometric origin were related to marked points, in line with the latter type of frozens. The existence of a cluster structure on $\mathfrak{M}(\Lambda, T)$ with a particular quiver Q has neat applications to symplectic geometry, e.g. studying the possible relative Lagrangian skeleta containing L for the Weinstein relative pair $(\mathbb{C}^2, \Lambda)$; see below for more.

Third, Item (iii) in Theorem 1.1 is geometrically keeping track of certain relative cycles before and after a Lagrangian surgery: the data being analyzed is the change of a specific local system along that relative cycle (which itself changes topologically). This local system is obtained by applying the microlocal functor, with the target being the Kashiwara-Schapira stack $\mu\mathrm{Sh}_\Lambda$, to a sheaf quantization of L. In our proof of Theorem 1.1, the sheaf quantization is obtained thanks to the construction of the weave $\mathfrak{w}(\mathbb{G})$, which represents a (front of the) Legendrian lift of L. In fact, Section 3 will provide a diagrammatic method to draw those relative cycles before and after a weave mutation, and Section 4 provides a Lie-theoretic procedure to compute with such (microlocal) local systems. Note also that the geometric mutations are associated with sugar-free hulls, which are not necessarily square faces and might include lollipops: the fact that the calculus of weaves allows for these general mutations is crucial so as to conclude that the coordinate ring of $\mathfrak{M}(\Lambda, T)$ is an upper cluster algebra.

Finally, the symplectic geometry perspective naturally leads to a quasi-cluster $\mathcal{A}$-structure, rather than a cluster $\mathcal{A}$-structure. Indeed, the weave $\mathfrak{w}(\mathbb{G})$ canonically gives the $\mathbb{L}$-compressing system $\mathfrak{S}$, which yields a linearly independent subset of $H_1(L \setminus T, \Lambda \setminus T)$. Nonetheless, there are cases in which this subset does not span and a choice of basis completion is precisely what introduces the quasi-cluster ambiguity. In particular cases, such as $\mathbb{G}$ being a plabic fence, the $\mathbb{L}$-compressing system already gives basis and hence $\mathfrak{M}(\Lambda(\mathbb{G}), T)$ carries a natural cluster $\mathcal{A}$-structure, but for a generic grid plabic graph $\mathbb{G}$ there is no a priori reason for that to be the case; the natural algebraic structure arising from symplectic geometry is only unique up to quasi-cluster equivalence.

Theorem 1.1 also implies a series of new computations and results in 3-dimensional contact topology. Indeed, in many interesting cases, such as those where the cluster algebra equals the upper cluster algebra [Mul13, Mul14], the existence of a cluster $\mathcal{A}$-structure on the moduli space[3] $\mathfrak{M} = \mathfrak{M}(\Lambda, T)$, as proven in Theorem 1.1, leads to:

(1) The computation of its deRham cohomology ring $H^*(\mathfrak{M}, \mathbb{C})$, including the refinement of its mixed Hodge structure. These computations are done in [LS22] for the locally acyclic cases, and these cohomology rings (and their mixed Hodge structure) are beautifully computed for all max-tb torus links $T(k, n)$ in [GL21]. As a simple instance, following [LS22], this readily implies that the cohomology ring $H^*(\mathfrak{M}_n, \mathbb{C})$ is of Hodge-Tate type and isomorphic to $\mathbb{C}[x]/(x^{n/2})$, $|x| = 2$, where $\mathfrak{M}_n$ is the moduli space for the max-tb $T(2, 2n + 1)$ torus knot. For the max-tb torus links $T(2, 2n)$, the corresponding non-zero Betti numbers are $b_i = 1$, $i \in [0, 2n]$. Similarly, the Poincaré polynomial of $\mathfrak{M}(T(3, 4))$, resp. $\mathfrak{M}(T(3, 5))$, is $1 + x^2 + 2x^4 + x^6$, resp. $1 + x^2 + 2x^4 + 2x^6 + x^8$, and the corresponding Deligne splitting is also understood. The cohomology rings of the moduli $\mathfrak{M}(\Lambda, T)$, equivalently of augmentation varieties, were not much understood and Theorem 1.1 can be a versatile tool for that.

(2) The existence of a holomorphic (pre)symplectic structure for the moduli space $\mathfrak{M}$. This allows for many classical techniques, such as quantization, to be applied to the coordinate ring $\mathcal{O}(\mathfrak{M})$, see [GSV10]. We emphasize that the cluster $\mathcal{A}$-variables associated to a seed are exponential Darboux coordinates for the symplectic 2-form. Note also that a holomorphic symplectic structure on the augmentation variety was recently constructed in [CGGS20] by different means (using the Cartan 3-form and Bott-Shulman forms). Upcoming work with our collaborators will show that these two holomorphic symplectic structures coincide whenever they can be compared.

---

[3]If not made explicitly, the set of marked points T is taken to have one marked point per component of Λ.

(3) In the Louise case [LS22], it is possible to compute the eigenvalues of the Frobenius automorphism on ℓ-adic cohomology and perform finite point counts $\#\mathfrak{M}(\mathbb{F}_q)$ performed over finite fields $\mathbb{F}_q$, $q = p^k$ and p large enough. These ought to be compared with the contact and symplectic results in [HR15, NRSS17].

Another byproduct of our result, thinking in terms of cluster ensembles [FG06a], is that there also exists a cluster $\mathcal{X}$-variety. Let $\mathcal{M}_1(\Lambda, T)$ be the undecorated stack associated to $\mathfrak{M}(\Lambda, T)$, with framing data at T. Theorem 1.1 implies the following result:

Corollary 1.2. *Let* $\mathbb{G} \subset \mathbb{R}^2$ *be a complete grid plabic graph,* $\Lambda = \Lambda(\mathbb{G}) \subset (\mathbb{R}^3, \xi_{st})$ *its associated Legendrian link and* $T \subset \Lambda$ *marked points. Then there exists a quasi-cluster* $\mathcal{X}$*-structure on* $\mathcal{M}_1(\Lambda, T)$.

In fact, each completion of the $\mathbb{L}$*-compressing system* $\mathfrak{S}$ *to a basis* $\mathfrak{B}$ *of* $H_1(L, T)$*, gives a cluster* $\mathcal{X}$*-structure on* $\mathcal{M}_1(\Lambda, T)$*. The initial quiver* Q *is defined by the intersections in* $\mathfrak{B}$ *and the initial cluster* $\mathcal{X}$*-variables are microlocal monodromies associated with elements of* $\mathfrak{B}$*. In addition, the mutable cluster* $\mathcal{X}$*-variables are those associated with curves in the* $\mathbb{L}$*-compressing system* $\mathfrak{S}$ *and different choices of completion of* $\mathfrak{S}$ *to a basis* $\mathfrak{B}$ *give quasi-equivalent cluster* $\mathcal{X}$*-structures on* $\mathcal{M}_1(\Lambda, T)$*.*

Corollary 1.2 is a new result, and completes the proof of existence of a cluster $\mathcal{X}$-structure, first hinted at in [STWZ19]. Indeed, the particular computation of square faces done in that work is independently recovered, and generalized, by the above corollary.[4] It is crucial to understand that there is currently no proof of Corollary 1.2 on its own. Namely, we are only able to deduce the existence of a cluster $\mathcal{X}$-structure once we have proven the existence of a cluster $\mathcal{A}$-structure in Theorem 1.1: the mathematical reason is that the results used from Berenstein-Fomin-Zelevinsky [BFZ05] are only applicable to cluster $\mathcal{A}$-structures.

Remark 1.3. Note that, as said above, the cluster $\mathcal{A}$-structures constructed in Theorem 1.1 can be used to compute cohomology rings, the Deligne splitting, build holomorphic symplectic structures, perform finite point counts and more. In contrast, none of this is possible in general if we just have a cluster $\mathcal{X}$-structure: this further illustrates the significant differences between cluster $\mathcal{A}$-structures and cluster $\mathcal{X}$-structures, and the versatility of the former in comparison to the latter.

The two moduli $\mathfrak{M}(\Lambda(\mathbb{G}), T)$ and $\mathcal{M}_1(\Lambda(\mathbb{G}), T)$ in Theorem 1.1 and Corollary 1.2 form a cluster ensemble. In Section 5, we focus on shuffle grid plabic graphs and prove that these cluster varietes always admit a Donaldson-Thomas (DT) transformation. See [KS10, GS18] for the necessary preliminaries on DT-transformations. In fact, we realize this cluster automorphism geometrically, as a composition of a Legendrian isotopy of $\Lambda(\mathbb{G})$ and the strict contactomorphism $t : (x, y, z) \longmapsto (-x, y, -z)$ of $(\mathbb{R}^3, \ker\{dz - y dx\})$. In particular, we conclude the following result:

Corollary 1.4. *Let* $\mathbb{G}$ *be a shuffle grid plabic graph. Consider the contactomorphism* t *and the half Kálmán loop Legendrian isotopy* $K^{1/2}$*. Then the composition* $t \circ K^{1/2}$ *induces the (unique) cluster Donaldson-Thomas transformation of* $\mathcal{M}_1(\Lambda(\mathbb{G}))$*.*

In particular, the cluster duality conjecture holds for the cluster ensemble $(\mathfrak{M}(\Lambda(\mathbb{G}), T), \mathcal{M}_1(\Lambda(\mathbb{G}), T))$*.*

The explicit sequence of mutations realizing the DT-transformation is presented in Section 5. We show it is a reddening sequence. Examples prove that it is not necessarily a maximal green sequence.

Finally, the contact and symplectic geometric results and techniques we use and develop to prove Theorem 1.1 are invariant under Hamiltonian isotopies, not necessarily compactly

---

[4]E.g. for all positroids in $\mathrm{Gr}(k, n)$, $k \geq 3, n \geq 6$, Corollary 1.2 constructs all infinitely many $\mathcal{X}$-clusters via microlocal techniques, whereas [STWZ19] built only finitely many candidates.

supported. Given that the cluster coordinates in Theorem 1.1 and Corollary 1.2 are all intrinsically named through symplectic geometric means, they can be named, and computed, after a compactly supported Hamiltonian isotopy is applied to $L(\mathbb{G})$ or a contact isotopy is applied to $\Lambda(\mathbb{G})$. This is a distinctive crucial feature which had been missing in our previous works [CGGS20, CGGS21, CZ22, GSW20a], where even the initial seed could not typically be defined (nor computed) after a Legendrian isotopy.[5]

Notation. We denote by $[a, b]$ the discrete interval $[a, b] := \{k \in \mathbb{N} : a \leq k \leq b\}$, if $a \leq b$, $a, b \in \mathbb{N}$. In this article, S_n denotes the group of permutations of n elements, $n \in \mathbb{N}$, and s_i its i-th simple transposition, $i \in [1, n-1]$. We abbreviate $s_{[b,a]} := s_b s_{b-1} \ldots s_{a+1} s_a$ and $s_{[b,a]}^{-1} := s_a s_{a+1} \ldots s_{b-1} s_b$, for $a < b$, $a, b \in \mathbb{N}$, and $s_{[b,a]}$ and $s_{[b,a]}^{-1}$ are empty if $b < a$. Let $w_{0,n} \in S_n$ be the longest word in the symmetric group S_n, we will sometimes write $w_0 \in S_n$ if n is clear by context. The standard word $\mathtt{w}_{0,n}$ for $w_{0,n}$ is defined to be the reduced expression $\mathtt{w}_{0,n} := s_{[1,1]} s_{[2,1]} s_{[3,1]} \cdots s_{[n-1,1]}$.

2. Grid Plabic Graphs and Legendrian Links

In this section we introduce the starting characters in the manuscript. On the combinatorial side, we introduce the notion of a grid plabic graph $\mathbb{G}$, in Subsection 2.1, and that of sugar-free hulls in Subsection 2.3. On the geometric side, we introduce a front for the Legendrian link $\Lambda(\mathbb{G})$ associated to the alternating strand diagram of a grid plabic graph $\mathbb{G}$, in Subsection 2.4, and set up the necessary moduli spaces from the microlocal theory of sheaves in Subsection 2.8. Several explicit examples are provided in Subection 2.5.

2.1. **Grid Plabic Graphs.** The input object in our results is the following type of graphs.

Definition 2.1. An embedded planar bicolored graph $\mathbb{G} \subset \mathbb{R}^2$ is said to be a *grid plabic graph* (or *GP-graph* for short) if it satisfies the following conditions.

(i) The vertices of $\mathbb{G} \subset \mathbb{R}^2$ belong to the standard integral lattice $\mathbb{Z}^2 \subset \mathbb{R}^2$, and they are colored in either black or white.

(ii) The edges of $\mathbb{G} \subset \mathbb{R}^2$ belong to the standard integral grid $(\mathbb{Z} \times \mathbb{R}) \cup (\mathbb{R} \times \mathbb{Z}) \subset \mathbb{R}^2$. Edges that are contained in $\mathbb{Z} \times \mathbb{R}$ are said to be *vertical*, and edges that are contained in $\mathbb{R} \times \mathbb{Z}$ are said to be *horizontal*.

---

[5]The pull-back structures from [STWZ19, Section 3] had the same issue.

(iii) A maximal connected union of horizontal edges is called a *horizontal line*. Each horizontal line must end at a univalent white vertex on the left and a univalent black vertex on the right. These univalent vertices are called *lollipops*.

(iv) Each vertical edge must end at trivalent vertices of opposite colors, and the end points of a vertical edge must be contained in the interior of a horizontal line.

Figure 2 depicts an example of a GP-graph $\mathbb{G}$, and Subsection 2.5 presents more instances.
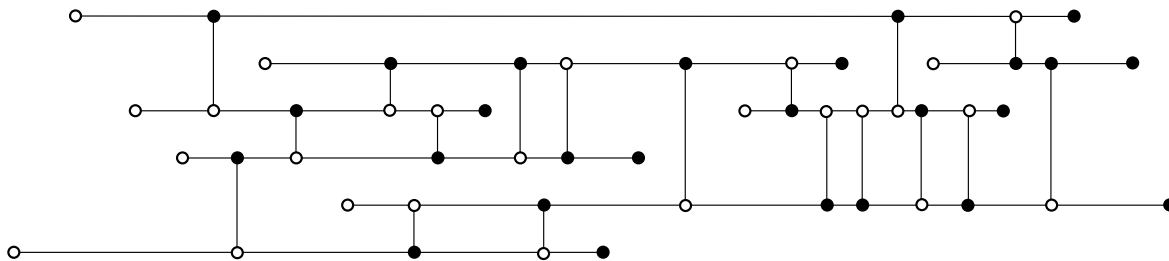

FIGURE 2. A GP-graph in accordance with Definition 2.1.

Remark 2.2. In Definition 2.1, it would have been fine to allow for bivalent vertices in a straight line. The Legendrian isotopy type of the zig-zag diagram, as introduced in Subsection 2.4, does not change when inserting such vertices (it performs a Reidemeister II move in the front), nor does the Hamiltonian isotopy type of the Lagrangian filling associated to the conjugate surface. Thus, all the structure discussed in this article remains the same and we might as well discard bivalent vertices in a straight line.

Remark 2.3. The condition in Definition 2.1.(iii) is merely technical and we strongly believe the results of this article can be generalize without this condition, i.e. also allowing black lollipops on the left and white lollipops on the right. The mathematical difference is that the Legendrian links then live in the ideal contact boundary $T^*_\infty \mathbb{R}^2 \cong (J^1 S^1, \xi_{st})$, but might represent a non-trivial homology class. The conceptual ideas we introduce in this article also work in that case, once the zero section $S^1 \subset J^1 S^1$ is satellited to the standard unknot in $(\mathbb{R}^3, \xi_{st})$, but the number of distinct cases and additional technicalities increases. Adding the condition Definition 2.1.(iii) already allows us to prove a new result in a vast amount of cases and keep the technicality and length of the article at a good balance.

2.2. **Column types and associated transpositions.** The intersection of a GP-graph $\mathbb{G} \subset \mathbb{R}^2$ with a subset of the form $\{(x,y) \in \mathbb{R}^2 : l < x < r\} \subset \mathbb{R}^2$, for some $l, r \in \mathbb{R}$, $l < r$, is said to be a column of $\mathbb{G}$. Any GP-graph $\mathbb{G}$ is composed by the horizontal concatenation of three types of non-empty columns called *elementary columns*. These three types of elementary columns are depicted in Figure 3 and can be described as follows.

- Type 1: a column is said to be Type 1 if it solely consists of parallel horizontal lines, i.e. it contains no vertices.

- Type 2: a column is said to be Type 2, or a crossing, if it contains exactly two oppositely colored vertices of $\mathbb{G}$ and a (unique) vertical edge between them.

- Type 3: a column is said to be Type 3, or a lollipop, if it contains exactly one lollipop. Note that the lollipop can be either white or black.

Remark 2.4. We label the horizontal $\mathbb{G}$-edges in Type 1 and 2 columns by consecutively increasing natural numbers from bottom to top. The horizontal lines of a Type 3 column

Type 1 | Type 2 (Crossings) | Type 3 (Lollipops)

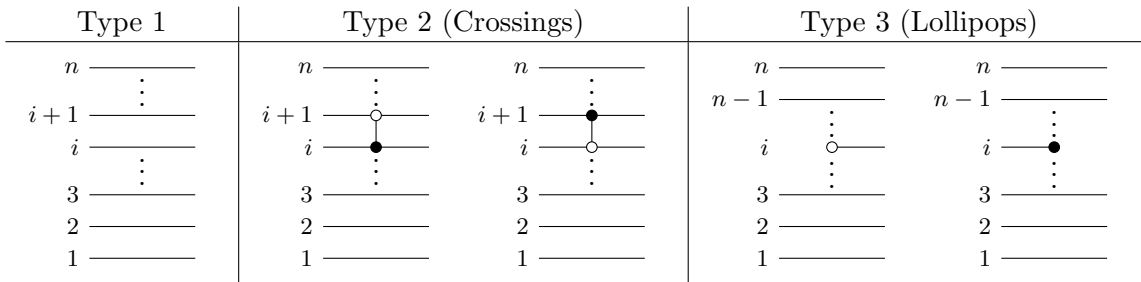n ⋯

$i+1$

i

3

2

1

FIGURE 3. The three types of elementary columns in a GP-graph.

are labeled in a similar way, but using the right side of the column in the case of a white lollipop and using the left side of the column in the case of a black lollipop. Without loss of generality, we always assume that there is a Type 1 column on each side of a column of Type 2 or 3.

Let $S_\mathbb{N}$ be the (infinite) group of permutations on the set $\mathbb{N}$. It is generated by simple transpositions $s_i = (i, i+1)$, $i \in \mathbb{N}$. Within $S_\mathbb{N}$, we define $S_{[a,b]} \cong S_{b-a+1}$ to be subgroup consisting of bijections that map i back to itself for all $i \notin [a, b]$. As we scan from left to right across the elementary columns of $\mathbb{G}$, we associate a copy of $S_{[a,b]}$ for some $[a, b]$ with each column of Type 1 or 2 by following the rules:

- We start with the empty set before the leftmost white lollipop, and we associate $S_{[1,1]}$ with the Type 1 column right after the leftmost white lollipop.

- The symmetric group $S_{[a,b]}$ does not change as we scan through a Type 1 or 2 column.

- If the symmetric group is $S_{[a,b]}$ before a Type 3 column with a white lollipop, then the symmetric group after this Type 3 column is $S_{[a,b+1]}$.

- If the symmetric group is $S_{[a,b]}$ before a Type 3 column with a black lollipop, then the symmetric group after this Type 3 column is $S_{[a+1,b]}$.

In summary, when passing through a white lollipop we move from a copy of S_k to a copy of S_{k+1} by adding a simple transposition at the end (with a larger subindex), and when passing through a black lollipop we move from a copy of S_{k+1} to a copy of S_k by dropping the first transposition (with smaller subindex).

Example 2.5. *Figure 4 depicts the GP-graph $\mathbb{G}$ in Figure 2 divided into columns, such that each columns has the same copy of the symmetric group; the symmetric group is written in red for each such column. The simple transpositions generating each of these copies are listed in orange, at the bottom of the figure, right below each dashed vertical yellow line. This example illustrates the description above, whereby a white lollipop adds a generating transposition and a black lollipops deletes one.*

2.3. **Sugar-free Hulls.** By definition, a *face* of a GP-graph $\mathbb{G}$ is any bounded connected component of $\mathbb{R}^2 \setminus \mathbb{G}$. A face is said to contain a lollipop if its closure in $\mathbb{R}^2$ contains a univalent vertex of $\mathbb{G}$. A *region* of a GP-graph $\mathbb{G}$ is a union of faces whose closure in $\mathbb{R}^2$ is connected; in particular, a face is a region and the union of any pair of adjacent faces is a region. For instance, the yellow and red area depicted in Figure 1 are both faces and the yellow face contains a lollipop; their union is a region (which will be the sugar-free hull of the yellow face).

11

S_1 S_2 S_3 S_4 S_5 S_6 S_5 S_4 S_3 S_4 S_3 S_4 S_3 S_2 S_1

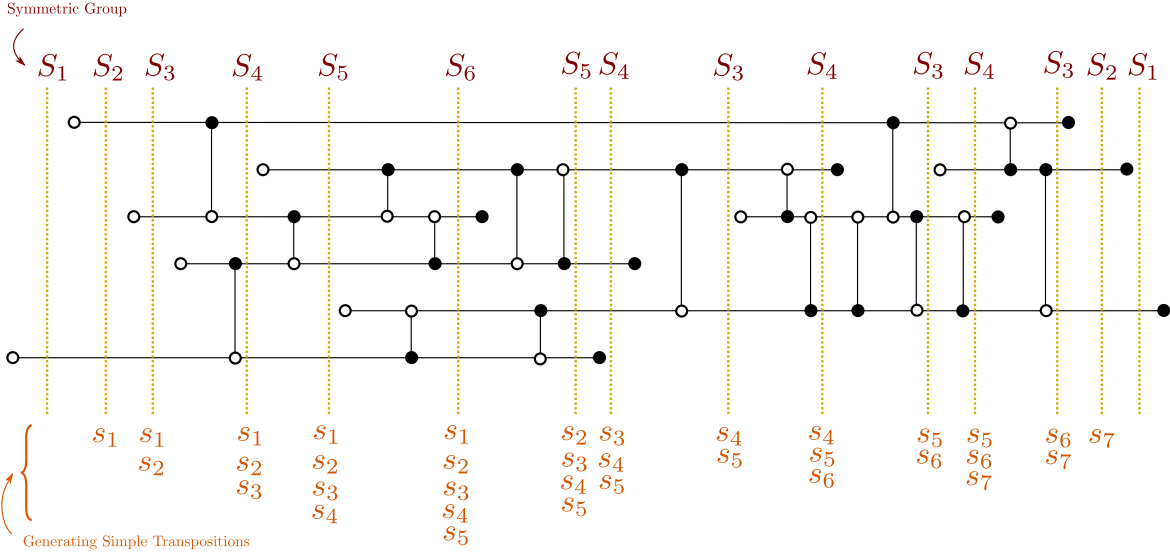|									
s_1	s_1	s_1	s_1	s_1	s_2 s_3	s_4	s_4	s_5 s_5	s_6 s_7
	s_2	s_2	s_2	s_2	s_3 s_4	s_5	s_5	s_6 s_6	s_7
		s_3	s_3	s_3	s_4 s_5		s_6	s_7	
			s_4	s_4	s_5				
				s_5					

Generating Simple Transpositions

FIGURE 4. The GP-graph in Figure 2 divided into vertical regions, each of which contains exactly one lollipop. The symmetric group associated to each yellow dashed line is labeled on top and its generating (simple) transpositions are listed at the bottom.

Sugar-free   Sugar content
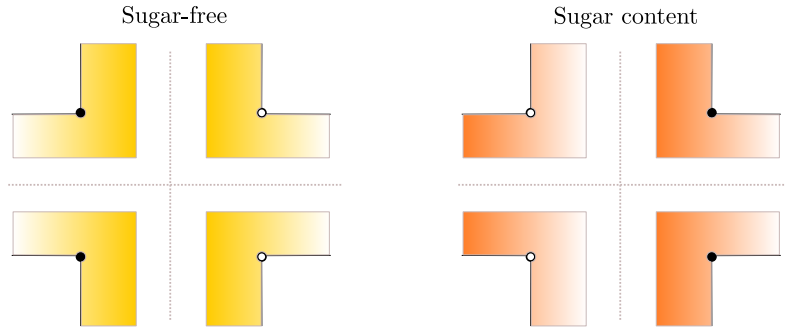
FIGURE 5. The four corners depicted on the left, in yellow, are allowed in a sugar-free region. The four corners depicted on the right, in orange, are not allowed in a sugar-free region, they have sugar content. Similarly, 270° corners with a lollipop always have sugar content.

The *boundary* ∂R of a region R is the topological (PL-smooth) boundary of its closure $\overline{R} \subset \mathbb{R}^2$. The boundary ∂R of a region necessarily consists of straight line segments meeting at corners that have either 90° or 270° angles. By definition, a 270° corner is said to be *left-pointing* if it is of the form $\ulcorner$ or $\llcorner$, and a 270° corner is said to be *right-pointing* if it is of the form $\urcorner$ or $\lrcorner$. Equipped with this terminology, we introduce the following notion:

Definition 2.6. Given a grid plabic graph $\mathbb{G}$, a region R is said to be *sugar-free* if all left-pointing 270° corners along ∂R are white and all right-pointing 270° corners along ∂R are black. See Figure 5 for a picture with the allowed (and disallowed) corners. The *sugar-free hull* $\mathbb{S}(f)$ of a face f of a plabic graph $\mathbb{G}$ is defined to be the intersection of all sugar-free regions R containing f. In particular, sugar-free hulls are sugar-free regions.

Remark 2.7. The terminology "sugar-free" is in part due to the fact that if a lollipop is attached to a 270° corner of same color, we can delete this lollipop without changing the Legendrian isotopy type of the Legendrian link $\Lambda(\mathbb{G})$ (see Subsection 2.4). This is not the case if the lollipop is attached to a 270° corner of the opposite color.

Example 2.8. *Figure 6 depicts three GP-graphs, with a few sugar-free faces, in yellow, and also faces which are not sugar-free. Note that the sugar-free hull $\mathbb{S}_f$ of a face f might be empty, e.g. the sugar-free hull of the green face in the graph on the second row is empty. Figure 7 depicts the sugar-free hulls of f_1, f_2, f_3 and f_4.*
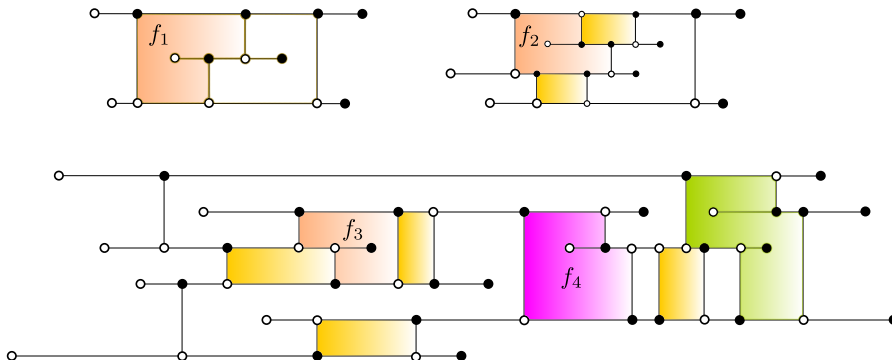
FIGURE 6. Three GP-graphs. Faces f_1, f_2, f_3, f_4 are not sugar-free, and nor is the green face in the graph on the second row. The yellow faces are sugar-free.
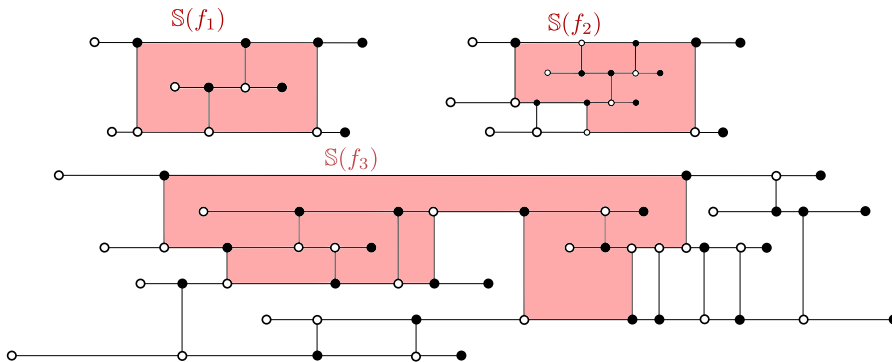
FIGURE 7. The three GP-graphs from Figure 6 with the sugar-free hulls of the faces f_1, f_2, f_3 (as labeled in Figure 6). The sugar-free hull of f_4 coincides with that of f_3. The sugar-free hulls of the yellow faces in Figure 6 are the yellow faces themselves and the sugar-free hull of the green face is empty.

Remark 2.9. Square faces are a pillar of the relation between Postnikov's plabic graphs and cluster algebras. A square face of a plabic graph can be mutated on, whereas mutating at a non-square face is often impossible with a plabic graph. In this article, we show how to mutate at a sugar-free full of a GP-graph $\mathbb{G}$, regardless of whether it is a square face, a non-square face, or a union of faces. This will use the technology of weaves, developed in Section 3, and see also Section 4. Thus, sugar-free hulls are set at the combinatorial center of our construction.

The boundary of a sugar-free region can be described explicitly, as follows:

Lemma 2.10. *Let R be a sugar-free region in a GP-graph $\mathbb{G}$. Then ∂R must be decomposed as a concatenation of staircases of the following four types:*

13
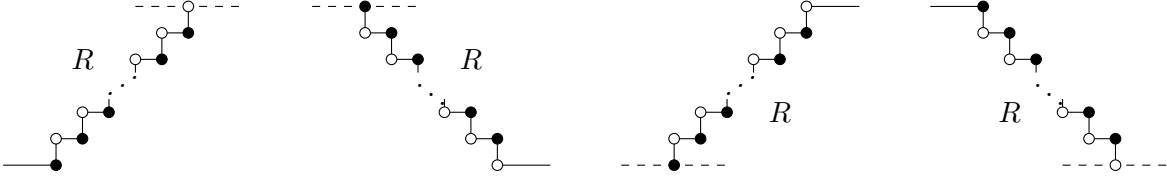
FIGURE 8. Four types of staircase building blocks for the boundary ∂R of a sugar-free region $R \subset \mathbb{G}$. In each instance, the letter R marks the location of the region in the plane. The dashed lines indicate that ∂R can continue in either of the two branches

Proof. This characterization follows immediately from the fact that all vertical bars must be of different colors at the two ends. □

Lemma 2.10 has the following application.

Lemma 2.11. *Let $\mathbb{G}$ be a GP-graph, $R \subset \mathbb{G}$ be a sugar-free region and C a column in $\mathbb{G}$, of any type. Then the intersection $R \cap C$ has at most one connected component.*

*Proof.* By definition, the region R is connected. Thus, in order for the intersection $R \cap C$ to have more than one connected component, R needs to make a (horizontal) U-turn at some point and ∂R must contain a part that is of the shape "R (" or ") R", where the parentheses indicate the U-turn and the letter R indicates the side of the region. However, such a shape cannot be built using the four types of staircases in Lemma 2.10 and therefore $R \cap C$ can have at most one connected component. □

Note that if a region R were not simply-connected, then there must exist a column C such that $R \cap C$ has more than one connected component. Thus, Lemma 2.11 has the following consequence, despite the fact that there may exist non-simply connected faces in the GP-graph.

Corollary 2.12. *Sugar-free regions are simply-connected.*

2.4. **Legendrian Links.** In this subsection we introduce the Legendrian link $\Lambda(\mathbb{G}) \subset (\mathbb{R}^3, \xi_{\mathrm{st}})$ associated to a GP-graph $\mathbb{G} \subset \mathbb{R}^2$ and explain how to algorithmitcally draw a specific front by scanning $\mathbb{G}$ left to right. Let us begin with the concise definition of $\Lambda(\mathbb{G})$, which reads as follows:

Definition 2.13. Let $\mathbb{G} \subset \mathbb{R}^2$ be a GP-graph. The Legendrian link $\Lambda(\mathbb{G}) \subset (\mathbb{R}^3, \xi_{\mathrm{st}})$ is the Legendrian lift of the alternating strand diagram of $\mathbb{G}$, understood as a co-oriented front in $\mathbb{R}^2$, considered inside a Darboux ball in $(T^*_\infty \mathbb{R}^2, \xi_{\mathrm{st}})$.

There are at least two points to discuss with regards to Definition 2.13. First, the co-orientation of the alternating strand diagram. Second, the contact geometric fact that there exists a Darboux ball in $(T^*_\infty \mathbb{R}^2, \xi_{\mathrm{st}})$ that actually contains $\Lambda(\mathbb{G})$, which is connected to the condition on $\mathbb{G}$ imposed in Definition 2.1.(iii).

Alternating strand diagrams were introduced in [Pos06, Definition 14.1] for a reduced plabic graph. In general, we associate such diagrams to a GP-graph $\mathbb{G} \subset \mathbb{R}^2$ according to the two following local models:
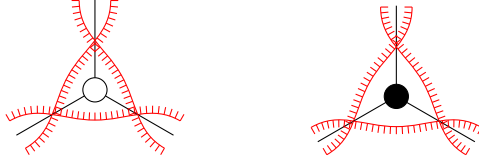
14

FIGURE 9. The local models for an alternating strand diagram associated to a GP-graph $\mathbb{G}$. The small hairs indicate the co-orienting direction, which is needed to specify a Legendrian lift.

The alternating strand diagram near a lollipop (or a bivalent vertex) is the same as in [Pos06], and the co-orientation in these pieces is implied by the co-orientations above.

By definition, the Legendrian lift of a co-oriented immersed curve on the plane $\mathbb{R}^2$ is a Legendrian link inside the ideal contact boundary $(T^*_\infty \mathbb{R}^2, \xi_{\mathrm{st}})$, or equivalently the unit cotangent bundle of the plane. The contact structure is the kernel of the restriction of the Liouville 1-form on $T^*\mathbb{R}^2$ to this hypersurface. In general, such Legendrian links cannot be contained in a Darboux ball, e.g. if they represent a non-zero homology class. For our GP-graph $\mathbb{G}$, the Legendrian lift $\Lambda(\mathbb{G})$ is naturally contained in a Darboux ball, as we now explain. Let us choose Cartesian coordinates $(u, v) \in \mathbb{R}^2$. Then, the contact structure on $T^*_\infty \mathbb{R}^2_{u,v}$ can be identified as the kernel of the contact 1-form

$$\alpha_{\mathrm{st}} := \cos\theta\, du + \sin\theta\, dv,$$

where $\theta \in [0, 2\pi)$ is the angle between a given covector $a\, du + b\, dv$ and du, $a, b \in \mathbb{R}$ and $a^2 + b^2 \neq 0$. Note that $T^*_\infty \mathbb{R}^2$ is diffeomorphic to $\mathbb{R}^2 \times S^1$, as $\mathbb{R}^2$ is parallelizable, and $\theta \in S^1$ is just recording that circle coordinate. In fact, we can consider the 1-jet space $(J^1 S^1, \xi_{\mathrm{st}})$ with its standard contact structure $\ker\{\beta_{\mathrm{st}}\}$, $\beta_{\mathrm{st}} := dz - y\, d\theta$, where $y \in \mathbb{R}$ is the coordinate along the cotangent fiber and $z \in \mathbb{R}$ the Reeb coordinate, as $J^1 S^1 := T^*S^1 \times \mathbb{R}$. Then, there exists a strict contactomorphism $\varphi : (T^\infty \mathbb{R}^2, \alpha_{\mathrm{st}}) \to (J^1 S^1, \beta_{\mathrm{st}})$ whose pull-back is given by

$$\varphi^*(\theta) = \theta, \qquad \varphi^*(y) = -u\sin\theta + v\cos\theta, \qquad \varphi^*(z) = u\cos\theta + v\sin\theta.$$

The key point at this stage is that for any open interval $I \subset S^1$, $(J^1 I, \xi_{\mathrm{st}})$ is contactomorphic to a standard Darboux ball $(\mathbb{R}^3, \xi_{\mathrm{st}})$. In consequence, if a co-oriented immersed curve $\mathfrak{f} \subset \mathbb{R}^2$ in $\mathbb{R}^2$ has a Gauss map that misses one given angle θ_0, the Legendrian lift of $\mathfrak{f} \subset \mathbb{R}^2$ is contained in $(J^1(S^1 \setminus \theta_0), \xi_{\mathrm{st}})$, which is contactomorphic to a Darboux ball. This is precisely what happens with the alternating strand diagram of a GP-graph $\mathbb{G} \subset \mathbb{R}^2$. Indeed, the alternating strand diagram of $\mathbb{G}$ can be isotoped so that the angle $\theta = 0$ is missed, say, and thus $\Lambda(\mathbb{G})$ naturally lives inside a Darboux ball.

Now that we have discussed Definition 2.13 and the two points above, we gear towards constructing a particular type of (wave)front for the Legendrian link $\Lambda(\mathbb{G})$ which will be useful to describe our moduli spaces in terms of Lie-theoretic objects, using tuples of flags. Note that the alternating strand diagram *is* a front, but it is not suited for our computations.

For that, we consider the front $\mathfrak{f}(\mathbb{G}) \subset \mathbb{R}^2$ obtained by dividing the GP-graph $\mathbb{G}$ into elementary columns and then use the assignments as depicted in Figures 10 and 11. Namely, to an elementary column of Type 1 with n strands, we assign a front consisting of $2n$ parallel horizontal strands. For an elementary column of Type 2 with n strands and a vertical bar at the ith position, we assign a front consisting of $2n$ parallel horizontal strands with a crossing at the ith position either at the top n strands or the bottom n strands, depending on whether the vertical bar had a white vertex at the top or at the bottom. Figure 10 depicts these three cases, with Types 1 and 2. The case of an elementary column of Type 3 involves inserting a right (resp. left) cusp at the ith position (plus some additional crossings) if there is a black (resp. white) lollipop inserted at the ith position. Figure 11 depicts the two possible cases for a Type 3 column.
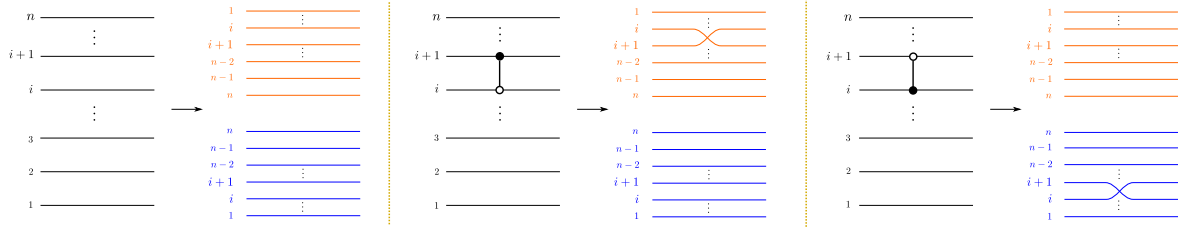
15

FIGURE 10. The rules to construct the front $\mathfrak{f}(\mathbb{G})$ from the elementary columns of a GP-graph $\mathbb{G}$. In this case, Type 1 and Type 2 columns are depicted, with the GP-graph $\mathbb{G}$ on the left and the front $\mathfrak{f}(\mathbb{G})$ on the right. We have colored the top n-strands of the front in orange and the bottom n-strands of the front in blue for clarification purposes.
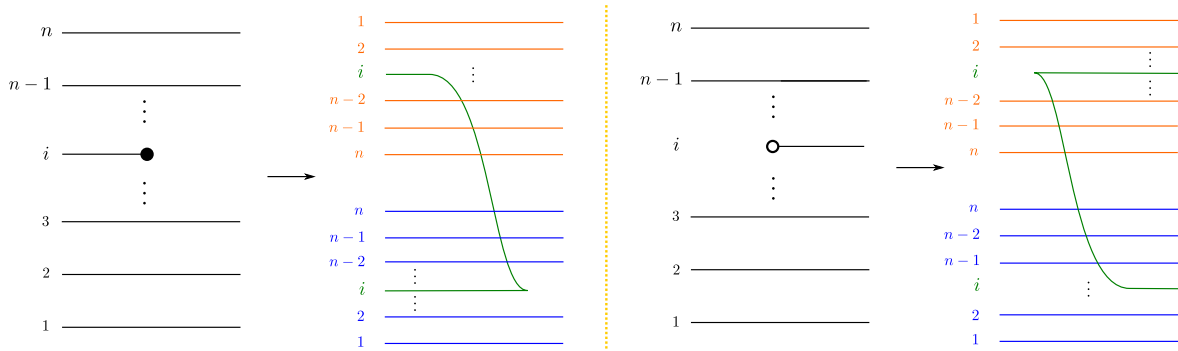
FIGURE 11. The rules to construct the front $\mathfrak{f}(\mathbb{G})$ from the elementary columns of a GP-graph $\mathbb{G}$. In this case, the two kinds of Type 3 columns are depicted, with the GP-graph $\mathbb{G}$ on the left and the front $\mathfrak{f}(\mathbb{G})$ on the right. We have colored the top n-strands of the front in orange, the bottom n-strands of the front in blue, and the newly inserted strand with a cusp in green, to help visualize the front.

Note that the $2n$ strands in the front are labeled in a specific manner in Figures 10 and 11, starting the count from the outer strand and increasing towards the middle. This choice of labeling is the appropriate one: in this way, when only left cusps have appeared, which is always the case at the beginning if we read $\mathbb{G}$ left to right, the ith top strand (in orange) and the ith bottom strand (in blue) coincide. Now, the front $\mathfrak{f}(\mathbb{G}) \subset \mathbb{R}^2$ lifts to a Legendrian link $\Lambda(\mathfrak{f}(\mathbb{G})) \subset (\mathbb{R}^3, \xi_{\mathrm{st}})$. We observe that in this case, the lift can be considered directly into $\mathbb{R}^3$, as the front is co-oriented upwards and there are no vertical tangencies. Let us show that this Legendrian link is actually Legendrian isotopic to $\Lambda(\mathbb{G})$.

Proposition 2.14. *Let $\mathbb{G} \subset \mathbb{R}^2$ be a GP-graph. Then the two Legendrian links $\Lambda(\mathbb{G}) \subset (\mathbb{R}^3, \xi_{st})$ and $\Lambda(\mathfrak{f}(\mathbb{G})) \subset (\mathbb{R}^3, \xi_{st})$ are Legendrian isotopic.*

Proof. Let us start with the GP-graph $\mathbb{G} \subset \mathbb{R}^2$ and consider the polar coordinates $(\rho, \theta) \in \mathbb{R}^2$. First, we perform a planar isotopy such that $\mathbb{G}$ is sent into the third quadrant of $\mathbb{R}^2$ and satisfies the following two properties:

- All the horizontal edges of $\mathbb{G}$ become polar circular arcs, i.e. cut out by equations

$$\{(\rho, \theta) \in \mathbb{R}^2 : \rho = \rho_0\} \cap \{(\rho, \theta) \in \mathbb{R}^2 : \theta_0 < \theta < \theta_1\},$$

16

where $\rho_0 \in \mathbb{R}^+$ is a positive constant and $\theta_0, \theta_1 \in S^1$ are two angles.

- All its vertical edges of $\mathbb{G}$ become polar radial line segments, i.e. cut out by equations

$$\{(\rho, \theta) \in \mathbb{R}^2 : \rho_0 < \rho < \rho_1\} \cap \{(\rho, \theta) \in \mathbb{R}^2 : \theta = \theta_0\},$$

where $\rho_0, \rho_1 \in \mathbb{R}^+$ are positive constants and $\theta_0 \in S^1$ is a fixed angle.

By compressing $\mathbb{G}$ radially, with a scaling diffeomorphism $\rho \longmapsto c \cdot \rho$ for some $c \in \mathbb{R}^+$, if necessary, we can draw the alternating strand diagram associated to $\mathbb{G}$ by only using the following three building blocks:

1. A polar circular arc along the horizontal edges of $\mathbb{G}$. Both pieces in Figure 12 have building blocks of this type.

2. A straight line segment connecting two polar circular arcs and tangent to the inner one; such line segments are drawn near vertical edges. The left diagram of Figure 12 depicts two such straight line segments, crossing each other exactly at the vertical edge.

3. A local rotation, i.e. an arc of a small circle that connects the two nearby alternating strands. This typically occurs at lollipops and vertical edges. Figure 12 depicts examples where such a local rotation occurs.
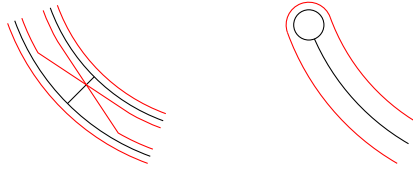
FIGURE 12. Examples of local rotations occurring at a vertical edge (the non-smooth corners on the left) and at a lollipop.

Now we consider the Legendrian lift of this particular type of alternating strand diagrams, apply the contactomorphism $(T^*_\infty \mathbb{R}^2, \xi_{\mathrm{st}}) \cong (J^1 S^1, \xi_{\mathrm{st}})$ explained above (restricted to the appropriate Darboux ball in the domain) and analyze how the image of this Legendrian lift projects under the standard front projection $(J^1 I, \xi_{\mathrm{st}}) \longrightarrow I \times \mathbb{R}$ of the (image of the) Darboux chart. The fronts for each of the three building blocks above transforms as follows:

(1) A polar circular arc at radius ρ becomes a horizontal line segment on $\mathbb{R}^2_{(\theta,z)}$ at height $z = \pm \rho$. This is because of the equality:

$$z = u \cos \theta + v \sin \theta = \pm(\rho \cos^2 \theta + \rho \sin^2 \theta) = \pm \rho.$$

Note that the sign depends on the conormal direction: it is a positive sign $(+)$ if it points radially outward and a negative sign $(-)$ if it points radially inward.

(2) A straight line segment in $\mathbb{R}^2_{(u,v)}$ can be parametrized as $(u_0 + at, v_0 + bt)$, where $t \in \mathbb{R}$ runs in an interval, and we choose (u_0, v_0) to be the point at which the line segment is tangent to the inner circle. Let ψ be the polar angle of (u_0, v_0). Then, the conormal direction of this line segment is constant at $\pm \psi$, and the tangent vector (a, b) must satisfy $a \cos \psi + b \sin \psi = 0$. Consequently, the z-coordinate is

$$z = (u_0 + at) \cos \psi + (v_0 + bt) \sin \psi = u_0 \cos \psi + v_0 \sin \psi,$$

which is also a constant. Therefore, the image of this straight line segment in $\mathbb{R}^2_{(\theta,z)}$ is a fixed point.

17

(3) Finally, for a local rotation, we can contract the circle to be small enough such that the rotation takes place approximately at the center, whose coordinates are $(u, v) = (\rho \cos \psi, \rho \sin \psi)$. Then, the z-coordinate becomes

$$z = \rho \cos \psi \cos \theta + \rho \sin \psi \sin \theta = \rho \cos(\theta - \psi).$$

That is, the image of a local rotation is qualitatively equal a part of the graph of a cosine function. This leads to the cusp.

The combination of these three computations above assemble into the rules established in Figure 10 and 11. Given that the planar isotopy and scaling of the GP-graph $\mathbb{G}$ induces a front homotopy for $\Lambda(\mathbb{G})$, the procedure above – setting $\mathbb{G}$ in the appropriate position – is equivalent to applying a Legendrian isotopy to $\Lambda(\mathbb{G})$. Since the image of the above three building blocks under the contactomorphism yields the front $\mathfrak{f}(\mathbb{G})$, we conclude that the two Legendrian links $\Lambda(\mathbb{G})$ and $\Lambda(\mathfrak{f}(\mathbb{G}))$ are Legendrian isotopic. $\qquad\square$

Example 2.15. *Let us consider the GP-graph $\mathbb{G}$ depicted on the left of Figure 13. Its associated Legendrian link $\Lambda(\mathbb{G})$ is the unique maximal-tb unknot in the standard contact $(\mathbb{R}^3, \xi_{st})$. The precise front that we obtain for this Legendrian unknot according to the procedure in the proof of Proposition 2.14 is depicted on the right of Figure 13.* $\qquad\square$
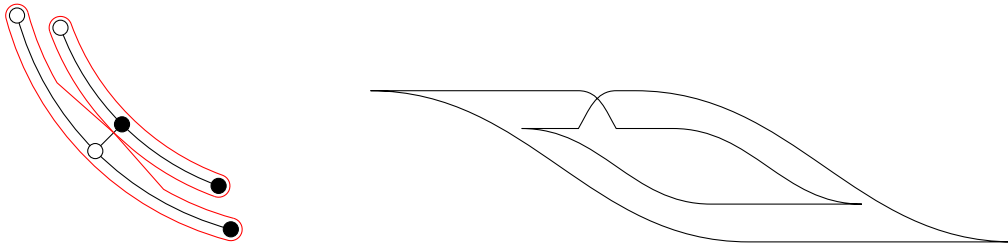
FIGURE 13. (Left) A GP-graph $\mathbb{G}$ isotoped and scaled according to the proof of Proposition 2.14. (Right) The associated front according to that same proof, or equivalently using the rules from Figures 10 and 11.

2.5. **Instances of GP-graphs $\mathbb{G}$ and their Legendrian Links $\Lambda(\mathbb{G})$.** In this subsection we discuss a few examples of GP-graphs $\mathbb{G}$ that lead to particularly interesting and well-studied Legendrian links.
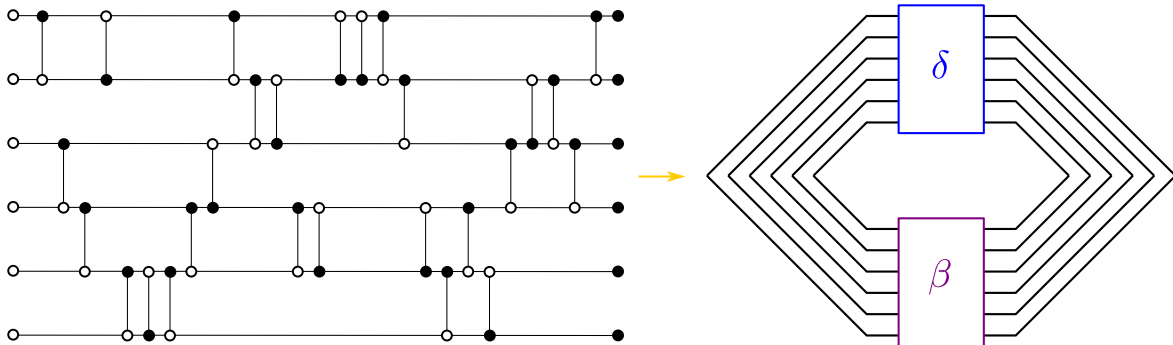
FIGURE 14. A front for the Legendrian link associated to the GP-graph on the left is drawn on the right, where $\beta, \delta \in \mathrm{Br}_6^+$ are the positive braid words $\beta = \sigma_5 \sigma_1 \sigma_3 \sigma_4 \sigma_3 \sigma_5^2 \sigma_2 \sigma_1 \sigma_4$ and $\delta = \sigma_1 \sigma_3 \sigma_4 \sigma_5^2 \sigma_4 \sigma_1 \sigma_2 \sigma_4 \sigma_1 \sigma_2 \sigma_5 \sigma_4 \sigma_3 \sigma_2 \sigma_3 \sigma_1$.
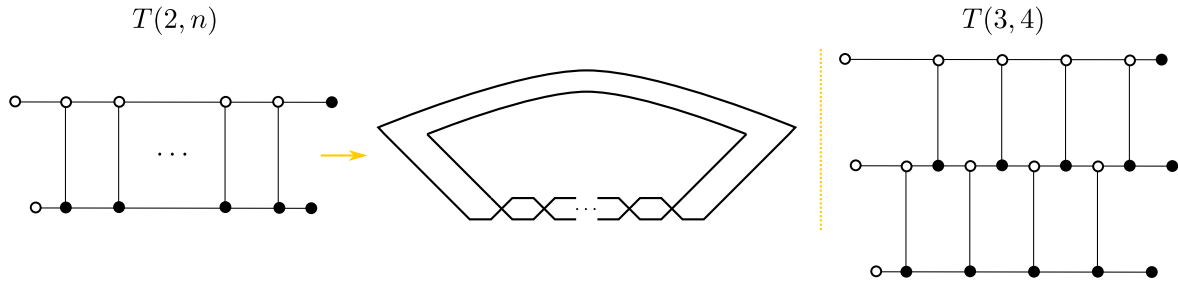
FIGURE 15. The Legendrian link associated to the GP-graph on the left is the max-tb representative of the $(2, n)$-torus link $T(2, n)$. A Legendrian front is depicted at its right. The Legendrian link associated to the GP-graph on the right is the max-tb representative of the $(3, 4)$-torus link $T(3, 4)$.

Plabic fences. Consider a GP-graph $\mathbb{G} \subset \mathbb{R}^2$ whose white lollipops all belong to the line $\{-1\} \times \mathbb{R}$, and all black lollipops belong to the line $\{1\} \times \mathbb{R}$. Figures 14 and 15 depict instances of such GP-graphs. These GP-graphs are called plabic fences in [FPST17, Section 12], following L. Rudolph's fence terminology. It follows from Proposition 2.14, and the rules from Figures 10 and 11, that the Legendrian link $\Lambda(\mathbb{G})$ associated to a plabic fence $\mathbb{G} \subset \mathbb{R}^2$ is Legendrian isotopic to the (Legendrian lift of the) rainbow closure of a positive braid.

In fact, given such a plabic fence $\mathbb{G} \subset \mathbb{R}^2$ with n horizontal lines, consider the positive braid word $\beta \in \mathrm{Br}_n^+$ whose kth crossing is σ_j if and only if the kth vertical edge *with black on bottom* of $\mathbb{G}$ (starting from the left) is between the jth and $(j + 1)$st horizontal strands. Similarly, consider the positive braid word δ whose mth crossing is σ_{n-j} if and only if the mth vertical edge *with white on bottom* of $\mathbb{G}$ is between the jth and $(j + 1)$st horizontal strands. Then, Figure 14 (right) depicts a front for the Legendrian link $\Lambda(\mathbb{G})$, which is readily homotopic to the rainbow closure of the positive braid word $\beta\delta^\circ$ (or equivalently $\delta^\circ\beta$), where δ° denotes the reverse positive braid of δ.

Remark 2.16. In the case that $\mathbb{G} \subset \mathbb{R}^2$ is the plabic fence associated to a positive braid word for the (unique) max-tb (k, n) torus link, $k, n \in \mathbb{N}$, the cluster $\mathcal{A}$-variety in Theorem 1.1 is the top open positroid cell of the Grassmannian $\mathrm{Gr}(k, n + k)$ (with a suitable choice of marked points). In this particular case, the unfrozen part of the initial quiver is dual to the plabic fence, as all absolute cycles in the associated weave $\mathfrak{w}(\mathbb{G})$ will be representable by long I-cycles. The frozen quiver vertices would be related to the choice of marked points. More generally, a GP-graph $\mathbb{G}$ whose associated Legendrian $\Lambda(\mathbb{G})$ is a positroid link, as studied in [CGGS21], yields the corresponding open positroid variety as the cluster $\mathcal{A}$-variety in Theorem 1.1.

The family Legendrian links $\Lambda(\mathbb{G})$ obtained from GP-graphs extends beyond rainbow closures of positive braids. The following family of examples readily illustrates this point.

Legendrian twist knots. Let us consider the family of GP-graphs $\mathbb{G}_n$, indexed by $n \in \mathbb{N}$, that we have depicted in Figure 16. Each GP-graph $\mathbb{G}_n$ has two long horizontal bars and it is obtained by inserting a staircase with n steps between two vertical bars, themselves located at the leftmost and rightmost position. Figure 16 draws $\mathbb{G}_1$ and $\mathbb{G}_2$ in the first row, and $\mathbb{G}_3$ on the left of the second row. By using Figures 10 and 11, fronts for the associated Legendrian knots $\Lambda(\mathbb{G}_n)$ are readily drawn: Figure 17 depicts fronts for $\Lambda(\mathbb{G}_1)$ and $\Lambda(\mathbb{G}_2)$. In general, we conclude that $\Lambda(\mathbb{G}_n)$ is a max-tb Legendrian representatives of a twist knot, with zero rotation number. Note that Legendrian twist knots are classified in [ENV13]. In particular, the Legendrian knot $\Lambda(\mathbb{G}_1)$ associated to the GP-graph depicted in the upper left
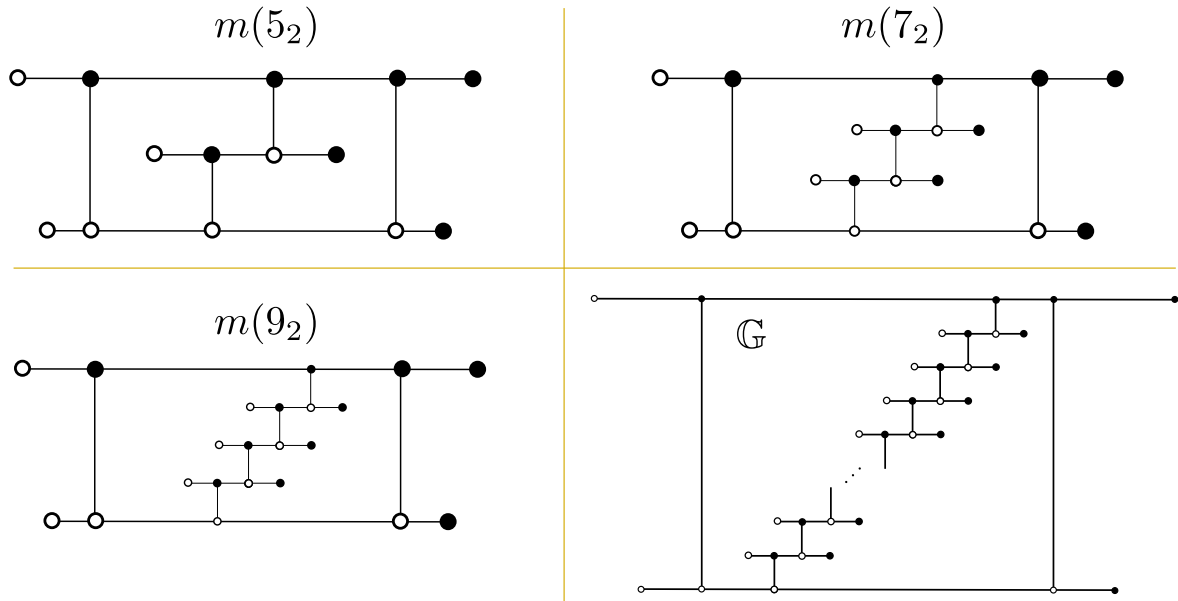
19

FIGURE 16. GP-graphs whose alternating strand diagrams have the smooth type of (mirrors of) twist knots. The GP-graph on the upper left (resp. upper right) yields a max-tb Legendrian representative of $m(5_2)$ (resp. $m(7_2)$). The bottom row illustrates the case for $m(9_2)$ and the general case, where a GP-graph is built by iteratively inserting – in a staircase manner – the local piece inside the GP-graph in the upper left.

of Figure 16 is the unique max-tb Legendrian representative of $m(5_2)$ with a binary Maslov index. This is one half of the well-known Chekanov pair.[6]

The moduli varieties associated to these Legendrian knots $\Lambda(\mathbb{G}_n)$ are interesting. Indeed, on the one hand, they present the simplest instance – and as far as we know, the first instance – of a cluster structure on the moduli variety such that the quiver has a frozen vertex associated to an *absolute* cycle in the Lagrangian filling. (The only frozen variables that we were aware of in the symplectic context were related to marked points, not absolute cycles.) This has a neat symplectic geometric consequence, as it allows one to obstruct the existence of embedded Lagrangian disks on the complement of a Lagragian filling by studying the cluster structure in the sheaf moduli (or an augmentation variety). On the other hand, the GP-graphs $\mathbb{G}_n$ have two faces, none of which is square, not even a regular polygon, because of the lollipops. Nevertheless, we will explain how there exists a *mutable* vertex in the quiver, which is associated to the sugar-free hull of one (either) of the faces. In addition, the quivers for the graphs $\mathbb{G}_n$ all coincide, independent of n: it is a quiver with one mutable vertex, one frozen vertex, and a single arrow connecting them. For each $n \in \mathbb{N}$, these two vertices are in bijection with the generators of the first homology group of a once punctured Lagrangian torus that bounds $\Lambda(\mathbb{G}_n)$. Thus, the family of Legendrian links $\Lambda(\mathbb{G}_n)$ becomes the first known family of Legendrian links whose sheaf moduli (or augmentation varieties) have isomorphic cluster $\mathcal{A}$-structures, despite the links not even being smoothly isotopic. □

Remark 2.17. (Connected Sums) The Legendrian knot associated to the GP-graph $\mathbb{G}$ given by taking the GP-graph in the lower right of Figure 16 and removing the rightmost and leftmost vertical bars, i.e. $\mathbb{G}$ is just has the staircase configuration connecting the two long

---

[6]The other max-tb representative of $m(5_2)$ is not isotopic to $\Lambda(\mathbb{G})$ for any GP-plabic graph $\mathbb{G}$.
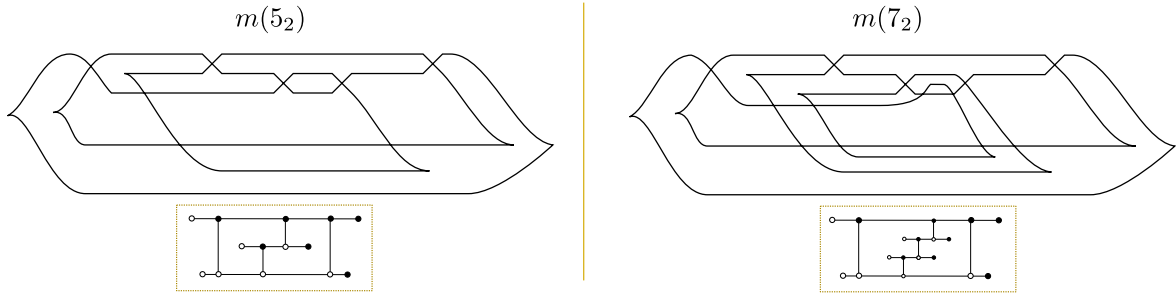
FIGURE 17. The two Legendrian fronts obtained from the GP-graphs in the top row of Figure 16, according to our recipe translating from GP-graphs to Legendrian front diagrams. The corresponding GP-graphs are drawn in a small yellow box below each front. The front diagram for the general case (whose GP-graph is depicted in the lower right of Figure 16) is readily inferred from these two pictures: a knotted spiralling pattern is iteratively added to the center region of the front.

horizontal bars, is the unique Legendrian unknot with $tb = -1$. In fact, given two disjoint GP-graphs $\mathbb{G}_1, \mathbb{G}_2$, we can form a new GP-graph $\mathbb{G}_1 \# \mathbb{G}_2$ by concatenating one of the rightmost black lollipops of $\mathbb{G}_1$ with one of the leftmost white lollipops of $\mathbb{G}_2$, as depicted in Figure 18 (left). From the perspective of the associated Legendrian links, this performs a Legendrian connected sum, i.e. $\Lambda(\mathbb{G}_1 \# \mathbb{G}_2) \cong \Lambda(\mathbb{G}_1) \# \Lambda(\mathbb{G}_2)$. More generally, this operation of merging two aligned lollipops is also available if they are located in the interior of a GP-graph; the effect of merging these lollipops into a straight horizontal segment on the associated Legendrian link is that of a Legendrian 0-surgery. □
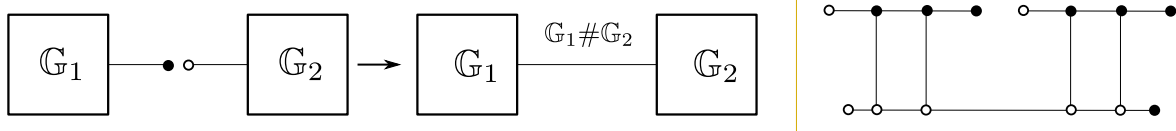
FIGURE 18. (Left) A Legendrian connect sum operation in terms of GP-graphs. (Right) A GP-graph $\mathbb{G}$ with $\Lambda(\mathbb{G})$ a max-tb Legendrian representative of the granny knot, obtained as the connected sum of two right-handed trefoils.

Shuffle graphs. Consider a permutation $\sigma \in S_n$, $i \in [1, n]$, a positive integer constant $M \in \mathbb{N}$, and the following set $\mathcal{H}_M(\sigma) := \{l_1^\sigma, l_2^\sigma, \ldots, l_n^\sigma\}$ of horizontal lines

$$l_i^\sigma := \{(x, y) \in \mathbb{R}^2 : y = i, -\sigma(i) \cdot M \leq x \leq \sigma(i) \cdot M\}, \quad i \in [1, n].$$

Definition 2.18. A GP-graph $\mathbb{G} \subset \mathbb{R}^2$ is said to be a *shuffle graph* if there exists $n, M \in \mathbb{N}$ and $\sigma \in S_n$ such that $\mathbb{G}$ can be obtained by adding vertical edges to $\mathcal{H}_M(\sigma)$ of the same pattern (i.e., all of which have either a black vertex on top, or a white vertex on top).

The two families above, of plabic fences and the GP-graphs in Figure 16 for Legendrian twist knots, are instances of shuffle graphs. Shuffle graphs $\mathbb{G} \subset \mathbb{R}^2$ have the property that the Legendrian $\Lambda(\mathbb{G}) \subset (\mathbb{R}^3, \xi_{st})$ is Legendrian isotopic to the Legendrian lift of the (-1)-closure of a positive braid of the form $\beta\Delta$, where $\Delta \in Br_n^+$ is the half-twist and $\beta \in Br_n^+$ has Demazure product $\text{Dem}(\beta) = w_0 = w_{0,n} \in S_n$. In fact, it is a simple exercise to verify that any (-1)-closure of a 3-stranded $\beta\Delta$, where $\beta, \Delta \in Br_3^+$ and $\text{Dem}(\beta) = w_0 \in S_3$, arises as $\Lambda(\mathbb{G})$ for some shuffle graph $\mathbb{G}$.

Remark 2.19. Note that the (-1)-closure of a positive braid of the form $\beta\Delta \in Br_n^+$ admits a Lagrangian filling if and only if $\text{Dem}(\beta) = w_0 \in S_n$, as follows from combining [CGGS20] (for the *only if* implication) and [CN21] (for the *if* implication). Section 3 implies that $\Lambda(\mathbb{G})$ admits a Lagrangian filling for any GP-graph $\mathbb{G}$ (this was already stated in Theorem 1.1), and thus the property $\text{Dem}(\beta) = w_0 \in S_n$ is necessary.

In view of this, and [CGGS20, CGGS21, CN21], it is convenient to introduce the following:

Definition 2.20. A positive braid $\beta \in Br_n^+$ is said to be Δ-*complete* if it is cyclically equivalent to a positive braid of the form $\Delta\gamma$ where Δ is the half twist and $\text{Dem}(\gamma) = w_{0,n}$.

Then the above paragraph states that the braid $\beta = \beta(\mathbb{G})$ associated to a shuffle graph $\mathbb{G}$ is Δ-complete. Note also tha the Legendrian links $\Lambda(\mathbb{G})$ associated to shuffle graphs also have natural Legendrian isotopies associated to them, in line with [CG22, CN21, GSW20b], which in all likelihood yield interesting cluster transformations for the associated cluster $\mathcal{A}$-variety (e.g. Donaldson-Thomas transformations), generalizing the case of plabic fences proven in [GSW20b].

Remark 2.21. Empirically, GP-graphs which are not plabic fences tend to yield Legendrians whose smooth types do not admit a positive braid representative. Figure 19 depicts two more simple GP-graphs whose smooth knot types are described in the caption.
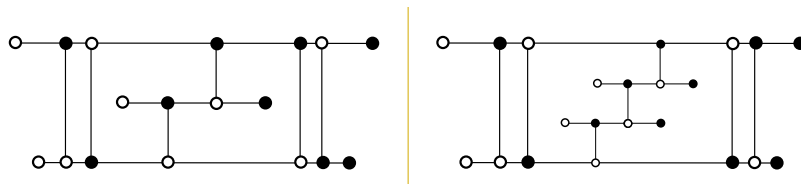
FIGURE 19. These two GP-graphs yield maximal-tb Legendrian representatives in the smooth isotopy classes of 7_3 (left) and 9_4 (right).

Let us conclude this section by pointing out a few properties of the Legendrian links $\Lambda(\mathbb{G})$ which can be of use. First, as we will explain, the Legendrian links $\Lambda(\mathbb{G})$ always bound an orientable exact embedded Lagrangian filling in the symplectization of $(\mathbb{R}^3, \xi_{\text{st}})$, and thus in the standard symplectic Darboux 4-ball. In particular, their Thurston-Bennequin invariant is always maximal and their rotation number vanishes. Second, it follows from the discussion in Subsection 2.4, especially Figures 10 and 11, that $\Lambda(\mathbb{G})$ admits a binary Maslov index and that the smooth type of $\Lambda(\mathbb{G})$ is that of the (-1)-closure of a positive braid. The former is particularly useful for us, as this implies that complexes of sheaves with singular support in $\Lambda(\mathbb{G})$ are quasi-isomorphic to sheaves (concentrated in degree 0) and it is the possible to parametrize the moduli of objects of the appropriate dg-category by an affine variety (or algebraic quotient thereof). Subsection 2.8 sets up the necessary ingredients and framework with regards to the microlocal theory of sheaves as it relates to these Legendrian links $\Lambda(\mathbb{G})$.

2.6. **Lollipop Chain Reaction.** In this subsection we introduce an algorithmic procedure, called a *lollipop chain reaction*, which aims to select faces for a sugar-free hull. The lollipop chain reaction initiates at a face f, and produces a collection of faces that are guaranteed to be inside the sugar-free hull $\mathbb{S}_f$. In many interesting cases of $\mathbb{G}$, such as shuffle graphs, this procedure yields the entire sugar-free hull $\mathbb{S}_f$. These combinatorial tools are used in Subsection 4.7, in the proof of Proposition 4.34. Let us start with the definition of a single *lollipop reaction*:

Definition 2.22. Let w be a white lollipop in a GP-graph $\mathbb{G}$, and let h_1 and h_2 be the two adjacent horizontal $\mathbb{G}$-edges exactly to the left of w (in between which the lollipop appears). A vertical line segment between h_1 and h_2 is said to be *wall*. By definition, the *lollipop reaction* initiated from the lollipop w pushes this wall to the right, and the wall shrinks or expands according to the following rules:

| wall shrinks | wall shrinks | wall expands | wall expands |
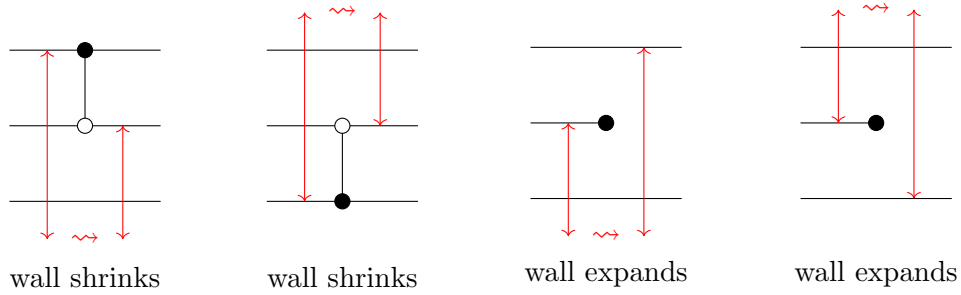
A single lollipop reaction initiated from a black lollipop b is defined in a symmetric fashion: start with a wall going between the two adjacent horizontal lines to the right of b, consider a wall between them and scan to the left. For a black lollipop, the wall shrinks or expands as it moves left according to the following rules:

| wall shrinks | wall shrinks | wall expands | wall expands |
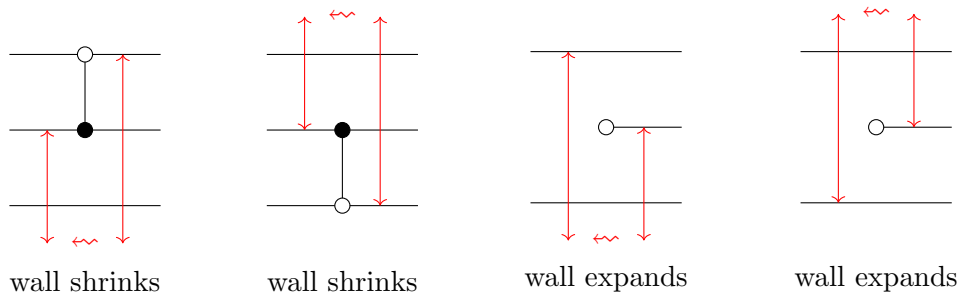
As the wall moves to the right (for a white lollipop) or to the left (for a black lollipop), we select all the faces that this wall scans through. By definition, a lollipop reaction *completes* when the length of the wall becomes zero. The output of a lollipop reaction is selection of faces of the GP-graph which it has scanned through. If the length of the wall becomes infinite (i.e., going to the unbounded region), then the lollipop reaction is said to be *incomplete*, and it outputs nothing.

In order to be effective, these lollipop reactions in general need to be iterated as follows:

Definition 2.23. Let $f \subset \mathbb{G}$ be a face of a GP-graph $\mathbb{G}$. A *lollipop chain reaction* initiated at f is the recursive face selection procedure obtained as follows:

  - Select the face f;
  - For each of the newly selected faces and each inward-pointing lollipop of this face, run a single lollipop reaction and select new faces (if any).

Since the number of faces in $\mathbb{G}$ is finite, this process terminates either when no new faces are selected, for which we say that the lollipop chain reaction is *complete*, or when one of the single chain reactions is incomplete, for which we say that the whole lollipop chain reaction is *incomplete*.

Example 2.24. *Each of the following four pictures is the result of a single lollipop reaction:*
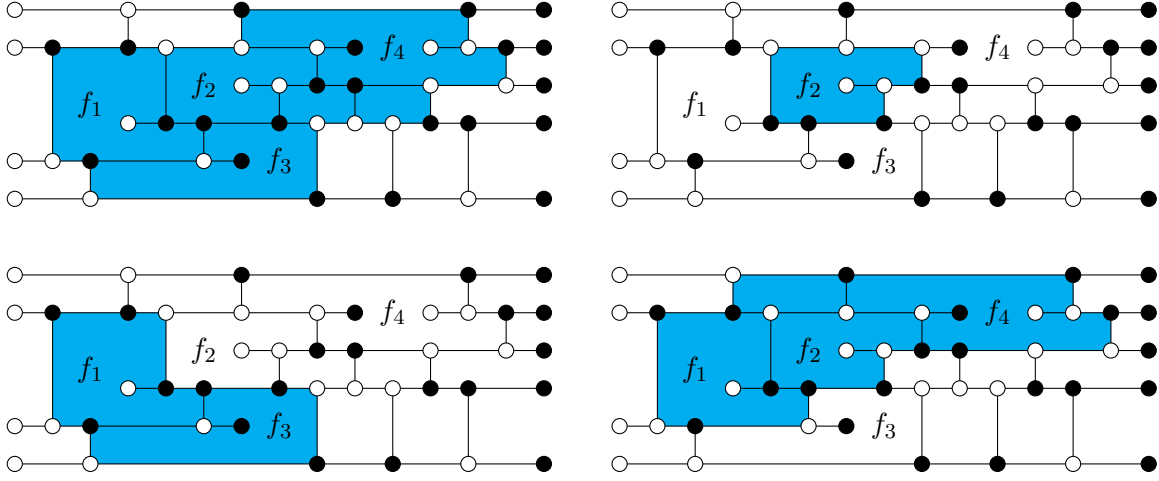
In detail, we have that:

- *For the GP-graph at the top left, the blue faces are the ones selected via a lollipop reaction initiated from the white lollipop of the face f_1,*
- *For the top right GP-graph, the reaction is initiated from the white lollipop of f_2;*
- *For the bottom left GP-graph, it is initiated from the black lollipop of f_3;*
- *For the bottom right, it is initiated from the black lollipop of f_4.*

Note that among these four single lollipop reactions, only the top right one produces a result that is sugar-free. However, if we take the union of all four regions, we get the sugar-free hull of f_1, which is the same as the sugar-free hull of f_3 or that of f_4.

Example 2.25. *In some cases, the iterated procedure does not always produce the full sugar-free hull $\mathbb{S}_f$, as illustrated by the following graph:*
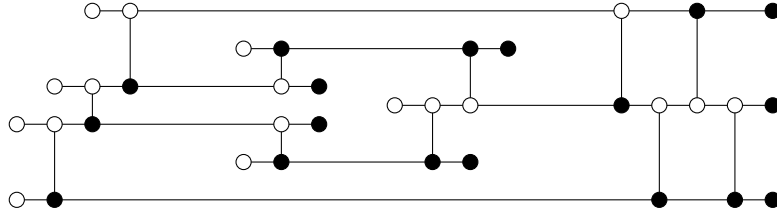
Note that a single lollipop reaction selects faces that are *minimally needed* to avoid sugar-content corners on the immediate right of a white lollipop or the immediate left of a black lollipop. Therefore the outcome of the lollipop chain reaction initiated from a face f must be contained in the sugar-free hull $\mathbb{S}_f$. In other words, if the lollipop chain reaction initiated from f is incomplete, then $\mathbb{S}_f$ does not exist. On the other hand, when sugar-free hulls $\mathbb{S}_f$ exist, lollipop chain reactions do produce sugar-free hulls for a large family of GP-graphs:

Proposition 2.26. *Let $\mathbb{G}$ be a shuffle graph and f a face of $\mathbb{G}$ for which $\mathbb{S}_f$ is non-empty. Then the lollipop chain reaction initiated from f is complete and $\mathbb{S}_f$ coincides with the outcome of this lollipop chain reaction.*

Proof. We observe that in a shuffle graph there cannot be any black lollipop on the left side of a white lollipop, nor can there be any white lollipop on the right side of a black lollipop. Therefore, if there is a white lollipop inside a face f, then the part of the boundary ∂f straightly left of the white lollipop can only consists of a single vertical bar. Similarly, if there is a black lollipop inside a face f, then the part of the boundary ∂f straightly right of the black lollipop can only consists of a single vertical bar as well. Now, if the face f does not have any lollipops, then $\mathbb{S}_f = f$, which is also equal to the outcome of the lollipop chain

24

reaction, as required. It remains to consider faces that do have lollipops inside. Without loss of generality let us suppose that only vertical $\mathbb{G}$-edge of the type ⬤—○ appear in $\mathbb{G}$ and suppose that f contains a white lollipop. Consider the left most white lollipop w of f. Then, at the starting point, the wall for this lollipop goes between two adjacent horizontal lines h_1 and h_2, and ∂f only has a single vertical bar to the left of this wall. Then the lollipop reaction starts moving the wall to the right, and one of the following two situation must occur:

- *The wall never expands.* If this is the case, the wall must be shrinking towards the top as shown below. The result of the lollipop reaction is obviously sugar-free.
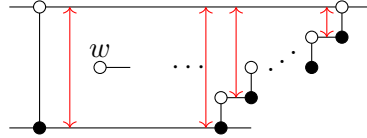
- *The wall expands at some point.* Note that the wall only expands when it passes through a black lollipop b. Let g be the face containing b. Then the part of ∂g straightly right of b only consists of a single vertical edge e, and hence the rightward scanning must end at e. Note that in this case there can be a concavity below the selected faces right before the expansion of the wall.
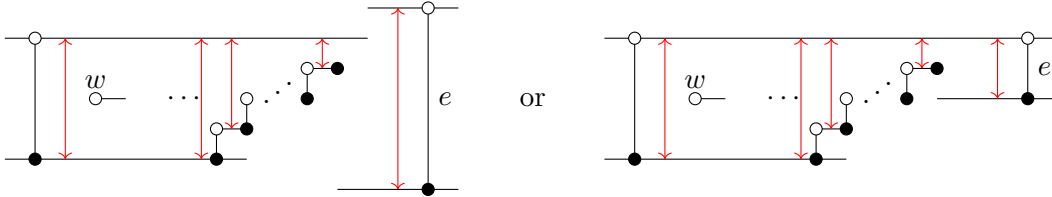
Now we start scanning leftward from the right most black lollipop in the face g. Note that for a leftward scanning, the wall should be shrinking towards the bottom. Also, due to the "last in first out" order on the horizontal lines, the top vertex of the vertical edge e cannot be below the top horizontal line of the previous scanning. If the top vertex of e is above the previous horizontal line, then the bottom vertex of e must be below the previous bottom horizontal line, and the left-ward scanning will not stop until it goes back all the way to the beginning point of the previous scanning. On the other hand, if the top vertex of e is on the previous horizontal line, then the bottom vertex of e can be above the previous bottom horizontal line. But then the horizontal line at the bottom of the vertical edge e must extend to the left and meets the staircase of the previous rightward scanning, and that is where the left-ward scanning stops. In consequence, the lollipop reaction from the rightmost black lollipop of the face g must fill in the lower concavity of the previous rightward scanning.

Note that in the second case above, the leftward scanning can also end in two ways, but we can conclude by induction that in the end all concavities will be filled and hence the resulting union must be sugar-free, as required. $\square$

There are many non-shuffle GP-graphs for which lollipop chain reactions also yield sugar-free hulls, and thus the hypothesis in Proposition 2.26 is sufficient but not necessary.

2.7. $\mathbb{L}$-compressing systems in Lagrangian fillings.

This brief subsection introduces the geometric notion of $\mathbb{L}$-compressing systems, which helps consolidate part of the ideas involved in the statement and proof of Theorem 1.1. Let Λ be a Legendrian link, in $(T^*_\infty \mathbb{R}^2, \xi_{st})$ or $(\mathbb{R}^3, \xi_{st})$, and $L \subset (T^*\mathbb{R}^2, \lambda_{st})$ an embedded exact Lagrangian filling. By definition, an embedded closed curve $\gamma \subset L$ is said to be $\mathbb{L}$-compressible if there exists a properly embedded Lagrangian 2-disk $D \subset (T^*\mathbb{R}^2 \setminus L)$ such that $\partial \overline{D} \cap L = \gamma \subset \mathbb{R}^4$.

Definition 2.27. A collection $\{\gamma_1, \ldots, \gamma_\ell\}$ of properly embedded closed curves $\gamma_i \subset L$, $i \in [1, \ell]$, is said to be an $\mathbb{L}$-compressing system if γ_i is $\mathbb{L}$-compressible for all $i \in [1, \ell]$ and the homology classes $[\gamma_1], \ldots, [\gamma_\ell]$ are a linearly independent in $H_1(L)$.

An $\mathbb{L}$-compressing system is said to be complete if the homology classes $[\gamma_1], \ldots, [\gamma_\ell]$ span $H_1(L)$. Equivalently, they form a basis of $H_1(L)$, i.e. $\ell = b_1(L)$. Finally, a collection of Lagrangian disks $\{D_1, \ldots, D_\ell\}$ is said to realize an $\mathbb{L}$-compressing system $\{\gamma_1, \ldots, \gamma_\ell\}$ if, for all $i \in [1, \ell]$, $D_i \subset (T^*\mathbb{R}^2 \setminus L)$ is properly embedded and $\partial \overline{D}_i \cap L = \gamma_i$. □

An outcome of the construction in Section 3 will be that a Legendrian link $\Lambda(\mathbb{G})$ coming from a GP-graph $\mathbb{G}$ is canonically endowed with an embedded exact Lagrangian filling $L(\mathbb{G})$ *and* an $\mathbb{L}$-compressing system $\{\gamma_1, \ldots, \gamma_\ell\}$, where ℓ is the number of sugar-free hulls in $\mathbb{G}$. In fact, if $\mathbb{G}$ is a plabic fence as in Subsection 2.5, then this $\mathbb{L}$-compressible system is actually complete.

Example 2.28. *The max-tb Legendrian representative $\Lambda \subset (\mathbb{R}^3, \xi_{st})$ of the smooth $(2, n)$-torus link is known to admit at least a Catalan number C_n of distinct Hamiltonian isotopy classes of Lagrangian fillings, see* [Pan17]. *By using 2-weaves* [CZ22], *it is readily proven that each of these Lagrangian fillings admits a complete $\mathbb{L}$-compressing system.*

From the perspective of smooth topology, we note that these embedded Lagrangian disks in $\mathbb{R}^4$ with boundary on a Lagrangian filling $L \subset \mathbb{R}^4$ have non-trivial normal framing, in the sense that it is *not* possible to perform a 0-surgery along them so as to decrease the genus of L. We thank M. Hedden and T. Mark for underscoring this point. A simple example is given by considering the following two Lagrangian planes in the standard Darboux chart $(\mathbb{R}^4, dq_1 dp_1 + dq_2 dp_2)$

$$L_1 = \{(q_1, q_2, p_1, p_2) \in \mathbb{R}^4 : q_1 = 0, q_2 = 0\}, \quad L_2 = \{(q_1, q_2, p_1, p_2) \in \mathbb{R}^4 : p_1 = 0, p_2 = 0\},$$

which bound a (component-wise) max-tb Hopf link $\Lambda_H \subset (S^3, \xi_{st})$ in the ideal contact boundary. Each of the Polterovich surgeries $L_1 + L_2$ and $L_2 + L_1$ are embedded exact Lagrangian annuli bounding Λ_H and, by construction, the cores of each of these annuli are $\mathbb{L}$-compressible. Nevertheless, since each of the components of Λ_H have Thurston-Bennequin invariant -1, it is not possible to perform a smooth 0-surgery along these bounding Lagrangian disks. (In particular, such an operation would result in two disjoint and smoothly embedded disks in the 4-ball bounding a Hopf link in the 3-sphere boundary, which do not exist.)

Remark 2.29. The classification of $\mathbb{L}$-compressing systems for a given Lagrangian filling, and that of their associated realizing Lagrangian 2-disks, seems like an interesting line of research in symplectic topology, to which we hope to return to in the future. For instance, even the existence of a complete $\mathbb{L}$-compressing systems might be interesting. To wit, the Lagrangian filling of the Legendrian $m(5_2)$ associated to the GP-graph in Figure 16 (upper left) does not seem to admit a complete $\mathbb{L}$-compressing system.

2.8. **Legendrian invariants from the microlocal theory of sheaves.** In this subsection we lay out the necessary ingredients of the microlocal theory of constructible sheaves that we shall use in our contact geometric framework. We describe the general setup in Subsection 2.8.1, based on works of S. Guillermou, M. Kashiwara and P. Schapira [KS85, KS90, GKS12a, GS14] and V. Shende, D. Treumann and E. Zaslow [STZ17].[7] We discuss in Subsection 2.8.2 the specific simplifications that occur for the Legendrian links $\Lambda(\mathbb{G})$. Finally, Subsection 2.8.3 introduces the necessary decorated version of the moduli stacks being discussed.

Remark 2.30. This section is necessary for a full conceptual understanding of the symplectic geometry underlying our constructions and results. Nevertheless, the summarized relevant

---

[7]See also S. Guillermou's notes for his lecture series at the conference "Symplectic topology, sheaves and mirror symmetry" at the IMJ-PRG in Paris (2016) and [STWZ19].

outcome of this section is that we can consider two algebraic varieties invariantly associated to the Legendrian links $\Lambda(\mathbb{G})$, and these varieties are the ones featuring in Theorem 1.1 and Corollary 1.2. In fact, Section 4 gives a description of these varieties which is independent of this section and can be understood by a reader not familiar with the microlocal theory of sheaves. □

Let $\Bbbk$ be a coefficient ring, which in this manuscript will be taken to be either $\Bbbk = \mathbb{Z}$ or a field, typically $\Bbbk = \mathbb{C}$. Consider a smooth manifold M, $\pi_M : T^*M \longrightarrow M$ its cotangent bundle and $T^*_\infty M \longrightarrow M$ its ideal contact boundary, which is contactomorphic to the unit cotangent bundle of M for any choice of Riemannian metric g. In this article we will only need $M = \mathbb{R}^2$ and $\mathbb{R}^3$.

2.8.1. *The general setup.* Let us succinctly review the general results on the microlocal theory of sheaves that we use. These results were pioneered by M. Kashiwara and P. Schapira in [KS90] and, more recently, in collaboration with S. Guillermou in [GKS12a]. The first category that we need is defined as follows:

Definition 2.31. The category $\underline{\mathbb{I}}(\Bbbk_M)$ is the full dg-subcategory of the dg-category of locally bounded complexes of sheaves of $\Bbbk$-modules on M which consist of h-injective complexes of injective sheaves. The homotopy category of $\underline{\mathbb{I}}(\Bbbk_M)$ is denoted by $[\underline{\mathbb{I}}(\Bbbk_M)]$.

The dg-category $\underline{\mathbb{I}}(\Bbbk_M)$ is a strongly pretriangulated dg-category and the six functor formalism lifts to this dg-enhancement $\underline{\mathbb{I}}(\Bbbk_M)$, see [Sch18]. The homotopy category $[\underline{\mathbb{I}}(\Bbbk_M)]$ is triangulated equivalent to the locally bounded derived category of sheaves on M, often denoted by $D^{lb}(\Bbbk_M)$. For an object $F \in \underline{\mathbb{I}}(\Bbbk_M)$, we denote by $\mu\mathrm{supp}(F) \subset T^*M$ its singular support understood as an object in $[\underline{\mathbb{I}}(\Bbbk_M)] \simeq D^{lb}(\Bbbk_M)$. The notion of singular support allows us to define the following dg-categories:

Definition 2.32. Let $S \subset T^*M$ be a subset. The category $\underline{\mathbb{I}}_S(\Bbbk_M)$ is the subcategory of $\underline{\mathbb{I}}(\Bbbk_M)$ consisting of objects $F \in \underline{\mathbb{I}}(\Bbbk_M)$ such that $\mu\mathrm{supp}(F) \subset S$. The category $\underline{\mathbb{I}}_{(S)}(\Bbbk_M)$ is the subcategory $\underline{\mathbb{I}}_S(\Bbbk_M)$ consisting of objects $F \in \underline{\mathbb{I}}_S(\Bbbk_M)$ for which there exists an open neighborhood Ω such that $\mu\mathrm{supp}(F) \cap \Omega \subset S$.

Definition 2.33. Let $\Lambda \subset T^*_\infty M$ be a Legendrian submanifold. We denote by $\underline{\mathbb{I}}_\Lambda(\Bbbk_M)$ and $\underline{\mathbb{I}}_{(\Lambda)}(\Bbbk_M)$ the categories as above with the choice of subset S being the Lagrangian cone of Λ union the zero section $M \subset T^*M$.

Note that the assignment $U \longmapsto \underline{\mathbb{I}}(\Bbbk_U)$ to each open subset $U \subset M$ is a stack of dg-categories, whereas the assignment $U \longmapsto D^{lb}(\Bbbk_U)$ only yields a pre-stack. Similarly, the pre-stack $\underline{\mathbb{I}}_{(\Lambda)}$ defined by

$$\underline{\mathbb{I}}_{(\Lambda)}(U) := \underline{\mathbb{I}}_{(T^*U \cap \Lambda)}(\Bbbk_U), \quad U \subset M \text{ open subset,}$$

where U is an open subset, is a stack. This is an advantage of using the injective dg-enhancements instead of derived categories.

Remark 2.34. Since we work in characteristic zero, the dg-enhancements $\underline{\mathbb{I}}(\Bbbk_M)$, $\underline{\mathbb{I}}_\Lambda(\Bbbk_M)$ and $\underline{\mathbb{I}}_{(\Lambda)}(\Bbbk_M)$ of $D^{lb}(\Bbbk_M)$ are equivalent to stable $(\infty, 1)$-categories, see [Fao17] and work of L. Cohn, and we thus consider all stacks in this manuscript to be derived stacks, always using the formalism of D^--stacks from [TV08].

A central result in symplectic topology [GKS12a] is that the stack $\underline{\mathbb{I}}_{(\Lambda)}$ on M is a Legendrian isotopy invariant of the Legendrian $\Lambda \subset T^*_\infty M$. There are two constructions associated to the stack $\underline{\mathbb{I}}_{(\Lambda)}$, as follows.

(i) **The microlocal functor** $\mathfrak{m}_\Lambda$. The Kashiwara-Schapira stack $\mu\mathrm{Sh}(\Bbbk_\Lambda)$ is the stack on Λ associated to the pre-stack

$$V \longmapsto \mathbb{I}_{(V)}(\Bbbk_M; V), \quad V \subset \Lambda \text{ open subset,}$$

where $\mathbb{I}_{(V)}(\Bbbk_M; V)$ is the Drinfeld dg-quotient of $\mathbb{I}_{(V)}(\Bbbk_M)$ by $\mathbb{I}_{T^*M \setminus (V \cup M)}(\Bbbk_M)$. See [KS90, GKS12a]. The transition from this pre-stack to its associated stack can significantly change the sections of the pre-stack. In particular, this pre-stack can be sensitive to the Legendrian isotopy type of Λ, whereas its associated stack $\mu\mathrm{Sh}(\Bbbk_\Lambda)$ is typically not. For instance, the pre-stack has no global sections when Λ is a stabilized Legendrian knot, whereas $\mu\mathrm{Sh}(\Bbbk_\Lambda)$ does have global sections in that case.[8] The quotient functor gives a functor of stacks

$$\mathfrak{m}_\Lambda : \mathbb{I}_{(\Lambda)} \longrightarrow (\pi_M|_\Lambda)_* (\mu\mathrm{Sh}(\Bbbk_\Lambda)).$$

The use of this functor in our article is twofold: it is used in the case that $\Lambda \subset T^*_\infty \mathbb{R}^2$ is a Legendrian link, and in the case where $\Lambda \subset T^*_\infty \mathbb{R}^3$ is a Legendrian surface obtained as the lift of an exact Lagrangian filling of a Legendrian link.

(ii) **The moduli stack** $\mathcal{M}_{\mathbb{I}_{(\Lambda)}(M)}$. By [Nad16, Theorem 3.21] the global sections $\mathbb{I}_{(\Lambda)}(M)$ is a dg-category equivalent to the category of (pseudo)perfect modules of a finite type category, namely of the category of wrapped constructible sheaves $\mathrm{Sh}^w_\Lambda(M)$ defined in [Nad16, Definition 3.17]. Then, the main result of [TV07] implies that there exists a locally geometric D^--stack $\mathcal{M}_{\mathbb{I}_{(\Lambda)}(M)}$, locally of finite presentation, which acts as the moduli stack of objects in the dg-category $\mathbb{I}_{(\Lambda)}(M)$.

Remark 2.35. For the purpose of the existence of cluster structures, it will suffice to truncate $\mathcal{M}_{\mathbb{I}_{(\Lambda)}(M)}$ to an Artin stack. In fact, the most interesting cases in this manuscript will have this truncation be a classical smooth affine scheme (or a stack quotient thereof by an algebraic group). Note though that, in general, the derived structure (in positive degree) is useful in extracting infinitesimal structures, such as the symplectic structure constructed in [CGGS20, Theorem 1.1]. Similarly, in general, there are Legendrian knots for which the ∞-stack structure (in negative degree) must be non-trivial. To wit, the (unique) max-tb Legendrian representative of $m(5_2)$ without a binary Maslov index has a stack with objects that contain non-zero higher automorphisms; in fact, already $Ext^{-1}(F, F) \neq 0$ for a certain sheaf F, as explained in [STZ17, Prop.7.6], which is a non-trivial 2-automorphism group.

Finally, as explained in [JT17, Section 1.7] given an embedded exact Lagrangian filling $L \subset T^*M$ of a Legendrian submanifold $\Lambda \subset T^*_\infty M$, the microlocal functor $\mathfrak{m}_{\overline{L}}$ applied to the Legendrian lift $\overline{L} \subset J^1(M)$ yields an equivalence of categories between (a subcategory of) $\mathbb{I}_{(\overline{L})}(M)$ and the dg-derived category of local systems on L. This induces an open inclusion

$$\iota_L : \mathbb{R}\mathrm{Loc}(L) \longrightarrow \mathcal{M}_{\mathbb{I}_{(\Lambda)}(M)}$$

where $\mathbb{R}\mathrm{Loc}(L)$ denotes the derived moduli space of local systems on L.

2.8.2. *The concrete models.* In the context of this article, strictly focusing on the construction of cluster structures, the general theory presented above can be simplified as follows.

First, for a Legendrian link $\Lambda \subset T^*_\infty \mathbb{R}^2$ with vanishing rotation number, the global sections of the Kashiwara-Schapira stack $\mu\mathrm{Sh}(k_\Lambda)$ admits a simple object Ξ by [Gui19, Part 10]. In addition, the functor $\mu\mathrm{hom}(\Xi, \cdot)$ is an explicit equivalence between $\mu\mathrm{Sh}(k_\Lambda)$ and the (twisted)

---

[8]This is akin to the distinction between pre-sheaves and sheaves, except that in our case the stacks are $(\infty, 1)$-category valued instead of merely set-valued. In general, sections must be added in order to sheafify.

stack Loc_Λ of (twisted) local systems on Λ.[9] In consequence, the microlocal functor $\mathfrak{m}_\Lambda$ described above yields a functor

$$\mathfrak{m}_{\Lambda,\Xi} : \mathbb{I}_{(\Lambda)}(M) \longrightarrow Loc_\Lambda(\Lambda),$$

where we have considered global sections and identified the codomain of $\mathfrak{m}_\Lambda$ with Loc_Λ via $\mu\mathrm{hom}(\Xi,\cdot)$ and a choice of spin structure.

Second, for a Legendrian link $\Lambda \subset T^*_\infty\mathbb{R}^2$, we only need to consider the moduli substack $\mathcal{M}_1(\Lambda)$ of $\mathcal{M}_{\mathbb{I}_{(\Lambda)}(\mathbb{R}^2)}$ which is associated to the subcategory of objects in $\mathbb{I}_{(\Lambda)}(\mathbb{R}^2)$ whose image under $\mathfrak{m}_\Lambda$ is a local system (on Λ) of locally free $\Bbbk$-modules of rank one supported in degree zero. In the case that Λ admits a binary Maslov index, the stack $\mathcal{M}_1(\Lambda)$ is equal to its truncation $t_0(\mathcal{M}_1(\Lambda))$, which is an Artin stack.

Now, given an embedded exact Lagrangian filling $L \subset T^*M$ of Λ, the derived stack Loc_L of local systems on L is also equivalent to its truncation and the open inclusion ι_L described gives an inclusion

$$\iota_L : \mathrm{Loc}_1(L) \longrightarrow \mathcal{M}_1(\Lambda)$$

of Artin stacks, where $\mathrm{Loc}_1(L)$ are local systems (on L) of locally free $\Bbbk$-modules of rank one supported in degree zero. Since Abelian local systems on L can be parametrized by $H^1(L,\Bbbk^\times)$, the inclusion ι_L provides a toric chart $\iota_L(Loc_1(L))$ in the moduli stack $\mathcal{M}_1(\Lambda)$. In the course of the article, we typically consider the ground ring $\Bbbk = \mathbb{C}$. If we are given a Legendrian link for which the stabilizers of $\mathcal{M}_1(\Lambda)$ are trivial and $\mathcal{M}_1(\Lambda)$ is smooth, then $\mathcal{M}_1(\Lambda)$ is (represented by) a smooth affine variety and an embedded exact Lagrangian filling L of Λ yields a toric chart

$$\iota_L : (\mathbb{C}^\times)^{2g(L)} \longrightarrow \mathcal{M}_1(\Lambda),$$

where $g(L)$ is the topological genus of the surface L.

Remark 2.36. In a sense, this was a first hint that symplectic geometry might yield a cluster $\mathcal{X}$-structure in $\mathcal{M}_1(\Lambda)$, for we at least have a Hamiltonian invariant way to produce toric charts from Lagrangian fillings. $\qquad\square$

Third, both the inclusions $\iota_L : \mathrm{Loc}_1(L) \longrightarrow \mathcal{M}_1(\Lambda)$ and the microlocal functors $\mathfrak{m}_\Lambda : \mathcal{M}_1(\Lambda) \longrightarrow \mathrm{Loc}_1(\Lambda)$ can be computed explicitly from the front via cones of maps between stalks (of the sheaves parametrized by $\mathcal{M}_1(\Lambda)$). Indeed, we shall use the combinatorial model in [STZ17, Section 3.3], where the points of $\mathcal{M}_1(\Lambda)$ are parametrized by functors from the poset category associated to the stratification induced by the Legendrian front to the Abelian category of $\Bbbk$-modules (modulo acyclic complexes). Note that in the case of Legendrian weaves, this combinatorial model is explained in detail in [CZ22, Section 5].

2.8.3. *A decorated moduli space.* Let $T = \{t_1,\ldots,t_s\}$, $t_i \subset \Lambda$, $i \in [1,s]$, be a set of distinct points in a Legendrian link $\Lambda \subset T^*_\infty\mathbb{R}^2$. The elements of T will be referred to as *marked points*. The moduli stack $\mathcal{M}_1(\Lambda)$ can be decorated with additional trivializing information once a set of marked points T has been fixed, as follows.

Definition 2.37. Let $\Lambda \subset T^*_\infty\mathbb{R}^2$ be a Legendrian link with a fixed choice of Maslov potential and spin structure. Consider a set of marked points $T = \{t_1,\ldots,t_s\}$ and label the components of $\Lambda \setminus T$ by Λ_i, $i \in \pi_0(\Lambda \setminus T)$. The moduli stack $\mathfrak{M}(\Lambda,T)$ is

$$\mathfrak{M}(\Lambda,T) := \{\big(F;\phi_1,\ldots,\phi_{|\pi_0(\Lambda\setminus T)|}\big) : F \in \mathcal{M}_1(\Lambda), \phi_i \text{ trivialization of } \mathfrak{m}_\Lambda(F) \text{ on } \Lambda_i\}.$$

Note that an Abelian local system can always be trivialized over Λ_i if $\Bbbk = \mathbb{C}^*$. For a general ground ring $\Bbbk$, we require that there exist at least one marked point per component of Λ.

---

[9]A choice of spin structure on Λ, and correspondingly choices of spin structures for the Lagrangian fillings we consider, allow a further identification to actual (untwisted) local systems. We implicit have these choices on the background and translate them combinatorially in Section 4 when they are needed to assign signs.

In Definition 2.37, the identification of (global sections of) the codomain of $\mathfrak{m}_\Lambda$ with the stack of local systems is fixed by the choice of Maslov potential and spin structure on Λ, as explained in Subsection 2.8.1 above.

There are at least two advantages to decorating the moduli stack of sheaves $\mathcal{M}_1(\Lambda)$ to $\mathfrak{M}(\Lambda, T)$. First, introducing the data of the trivializations in $\mathfrak{M}(\Lambda, T)$ often results in a smooth affine variety, even if $\mathcal{M}_1(\Lambda)$ was singular; this is similar to the classical setup with character varieties [FG06b]. Second, the trivializations in $\mathfrak{M}(\Lambda, T)$ can be used to define global regular functions. In fact, we will show that $\mathfrak{M}(\Lambda, T)$ admits a cluster $\mathcal{A}$-structure, and our construction of the cluster $\mathcal{A}$-variables crucially relies on the existence of these decorations. Finally, the moduli space $\mathcal{M}_1(\Lambda, T)$ is defined similarly, by considering sheaves in $\mathcal{M}_1(\Lambda)$ with the additional data of trivializations of the stalks of the associated microlocal local systems at each of the marked points in T.
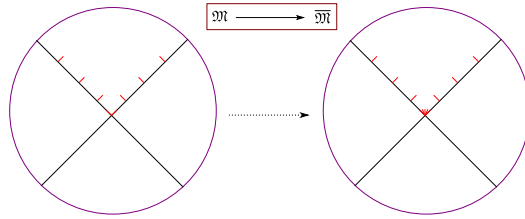
FIGURE 20. The local piece of the front on the left is adding singular support to the piece of front on the right. (Topologically adding a connected interval within the Legendrian fiber.) Such exchange at a given crossing partially compactifies the moduli spaces $\mathfrak{M}$ being discussed in this section.

Remark 2.38. From the perspective of birational geometry and mirror symmetry [HK18] (see also cluster duality [FG06a]) it might be interesting to note that a *front* for a Legendrian link $\Lambda \subset T_\infty^* \mathbb{R}^2$, gives a natural compactification of $\mathfrak{M}(\Lambda, T)$. Indeed, for any subset of the crossing of the front, one can expand the allowed singular support as in Figure 20. This has the effect of partially compactifying the type of moduli spaces $\mathfrak{M}$ being discussed in this section, see e.g. [KS90, STZ17]. In particular, this addition of the allowed singular support can be performed at all the crossings of a given front. For many examples in this manuscript, there are fronts such that this operation results into a projective variety $\overline{\mathfrak{M}}$. We underscore that the resulting compactification *does* depend on the front, see [CGGS21, Section 4], and it is of interest to understand the different (log Calabi-Yau) pairs $(\overline{\mathfrak{M}}, D)$, D a compactifying divisor, associated to the moduli space $\mathfrak{M}(\Lambda, T)$ obtained in this manner.

3. Diagrammatic Weave Calculus and Initial Cycles

The new machinery from contact topology that allows us to construct cluster structures is the study of Legendrian weaves, as initiated in [CZ22]. In this manuscript, we continue to develop techniques for Legendrian weaves so as to prove Theorem 1.1. These weave techniques now relate to GP-graphs $\mathbb{G}$ and their associated Legendrian links $\Lambda(\mathbb{G})$. Among many central facts, the construction of a weave $\mathfrak{w}(\mathbb{G})$ associated to $\mathbb{G}$ yields a canonical embedded exact Lagrangian filling for $\Lambda(\mathbb{G})$, a sheaf quantization, and the flag moduli of the weaves $\mathfrak{w}(\mathbb{G})$ shall provide the initial seeds for our cluster structures. In addition, as explained in Section 4, the weave $\mathfrak{w}(\mathbb{G})$ is used also to carry the explicit computations necessary for the study of cluster $\mathcal{A}$-variables and the proof of Theorem 1.1.

3.1. **Preliminaries on Weaves.** The reader is referred to [CZ22] for the details and background on Legendrian weaves. That said, we provide here a quick primer on the basics. By

definition, an embedded planar graph $G \subset \mathbb{R}^2$ is trivalent if all its vertices have degree three, see Figure 21 (left).

Let $J, K \subset \mathbb{R}^2$ be two trivalent planar graphs having an isolated intersection point at a common vertex $v \in J \cap K$. By definition, 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. Figure 21 (right) depicts such a hexagonal vertex.
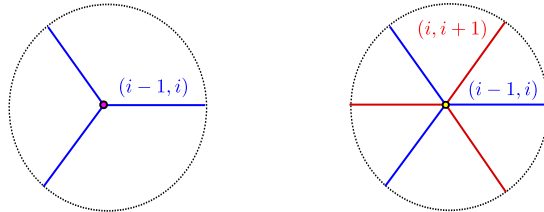

FIGURE 21. Trivalent vertex (left) and Hexagonal Point (right).

Definition 3.1. Given $n \in \mathbb{N}$, an N-weave $\mathfrak{w} \subset \mathbb{R}^2$ is a set $\mathfrak{w} = \{G_i\}_{1 \le i \le N-1}$ of $(N-1)$ embedded trivalent planar graphs $G_i \subset \mathbb{R}^2$, possibly empty or disconnected, such that G_i is allowed to intersect G_{i+1} only at hexagonal points, $1 \le i \le N-2$. By definition, a weave $\mathfrak{w} \subset \mathbb{R}^2$ is an N-weave for some $N \in \mathbb{N}$. $\qquad\square$

We also refer to the image of a weave in the plane as a weave $\mathfrak{w}$, as no confusion arises. The edges of the graphs that constitute a weave $\mathfrak{w}$ are often referred to as weave lines. We note that two graphs $G_i, G_j \subset \mathbb{R}^2$ are allowed to intersect (anywhere) as long as $j \ne i, i \pm 1$, and we always assume that the intersection is transverse. Through its image, we also think of an N-weave as an immersed graph in $\mathbb{R}^2$ with colored (or labeled) edges, the color i corresponding to the graph G_i, $1 \le i \le N-1$. 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. As a graph in the plane, and N-weave has trivalent, tetravalent and hexagonal vertices.

Let $\{s_i\}_{i=1}^{N-1}$ be the set of Coxeter generators of the symmetric group S_N. Instead of colors, we can equivalently label the edges of an N-weave $\mathfrak{w} = \{G_i\}$ which belong to the graph G_i with the transposition s_i: these labeled edges will also be referred to as s_i-edges, or i-edges.

Remark 3.2. By definition, the trivalent vertices belonging to the graph G_i have three incident s_i-edges. Similarly, the hexagonal points in $G_i \cap G_{i+1}$ have six edges incident to it, alternately labeled by s_i and s_{i+1}. Figure 21 depicts the local model for the trivalent vertices of the cubic graph G_{i-1} and a hexagonal intersection point in $G_i \cap G_{i-1}$.

The theory of weaves as developed in [CZ22] is grounded on the theory of Legendrian surfaces in $(\mathbb{R}^5, \xi_{\mathrm{st}})$ and their spatial wavefronts. In brief, a weave $\mathfrak{w} \subset \mathbb{R}^2$ gives rise to a spatial Legendrian wavefront $\Sigma(\mathfrak{w}) \subset \mathbb{R}^3$, which itself lifts to an embedded Legendrian surface $\Lambda(\mathfrak{w})$ in $(\mathbb{R}^5, \xi_{\mathrm{st}})$. In the present manuscript, the main property of the surface $\Lambda(\mathfrak{w})$ that we use is that its image $L(\mathfrak{w}) := \pi(\Lambda(\mathfrak{w})) \subset (\mathbb{R}^4, \omega_{\mathrm{st}})$ is an exact Lagrangian surface in the standard symplectic Darboux ball, where $\pi : (\mathbb{R}^5, \xi_{\mathrm{st}}) \longrightarrow (\mathbb{R}^4, \omega_{\mathrm{st}})$ is the projection along the α_{st}-Reeb direction. A schematic diagram for these objects is as follows:

$$\Lambda(\mathfrak{w}) \subset (\mathbb{R}^5, \xi_{\mathrm{st}})$$

$$\overset{\pi_{front}}{\swarrow} \qquad \overset{\pi}{\searrow}$$

$$\Sigma(\mathfrak{w}) \subset \mathbb{R}^3 \qquad\qquad L(\mathfrak{w}) \subset (\mathbb{R}^4, \lambda_{\mathrm{st}})$$

In this article, unless it is stated otherwise, all the weaves that we construct are free, see [CZ22, Section 7.1.2], which translates into the fact that $L(\mathfrak{w}) \subset (\mathbb{R}^4, \omega_{st})$ will always be an embedded exact Lagrangian surface, and not just immersed. In particular, this implies that $L(\mathfrak{w})$ must have boundary, which it always will. Moreover, when $\mathfrak{w}$ is free, the Lagrangian projection map π is a *homeomorphism*, and hence $H_1(L(\mathfrak{w})) \cong H_1(\Lambda(\mathfrak{w}))$.

In the current combinatorial presentation, the underlying contact geometry dictates that certain weaves ought to be considered equivalent. This leads to the following.

Definition 3.3. The moves depicted in Figure 22 are referred to as *weave equivalences*. By definition, two weaves $G, G' \subset \mathbb{R}^2$ are said to be equivalent if they differ by a sequence of weave equivalences or diffeomorphisms of the plane. We interchangeably refer to a weave and its weave equivalence class when the context permits.
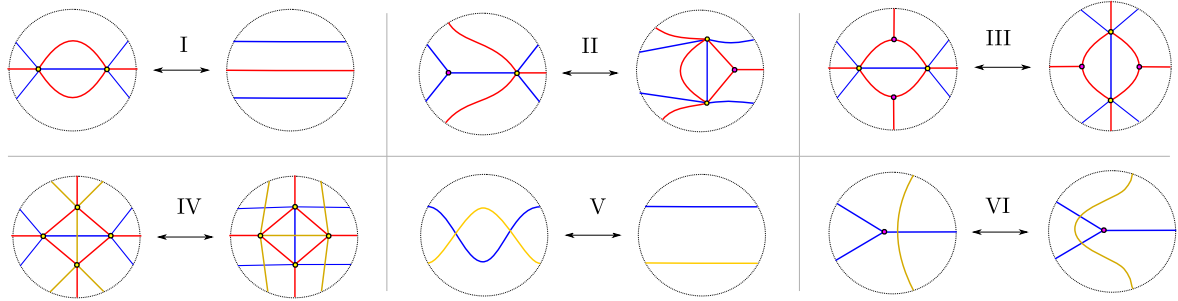
FIGURE 22. Six weave equivalences.

As noted in [CZ22, Section 4], these moves are not entirely independent, as Move III can be deduced from Move I and Move II; it is nevertheless useful to underscore Move III when working with weaves. The results in [CZ22], using the underlying contact geometry, imply that all the constructions that we will associate to a weave are invariant under weave equivalences. (It would to be possible to verify this combinatorially as well, see for instance the computations in [CGGS20].) At this stage, a first goal will be to construct a weave for each GP-graph $\mathbb{G} \subset \mathbb{R}^2$, which we momentarily do in Subsection 3.3, once we have reviewed the necessary material on Y-cycles and weave mutation.

3.2. Y-cycles and weave mutation. Let $\mathfrak{w} \subset \mathbb{R}^2$ be a free weave. The homology group $H_1(L(\mathfrak{w})) \cong H_1(\Lambda(\mathfrak{w}))$ has a central role in our article, as it is a sub-lattice of the defining lattice for the initial seed. In [CZ22, Section 2] we devised a method to describe absolute cycles in $L(\mathfrak{w})$ in terms of $\mathfrak{w} \subset \mathbb{R}^2$. The main concept that is relevant for our purposes is that of a Y-cycle on a weave $\mathfrak{w}$, which is defined as follows.

Definition 3.4. Let $\mathfrak{w} \subset \mathbb{R}^2$ be a weave. An absolute 1-cycle $\gamma \subset \Lambda(\mathfrak{w})$ is said to be Y-*cycle* if its projection onto $\mathbb{R}^2$ consists of weave lines, i.e. it is contained in $\mathfrak{w}$. A Y-cycle is said to be a Y-*tree* if its projection image is a tree, considered as a planar embedded graph in $\mathbb{R}^2$. A Y-tree is a I-*cycle* if its projection onto $\mathbb{R}^2$ does not have any trivalent vertices. Finally, an I-cycle is *short* if it does not pass through any hexagonal vertex of the weave $\mathfrak{w}$. Figure 23 depicts the four possible local models for a Y-tree near a trivalent, tetravalent, and hexagonal vertex of the weave.
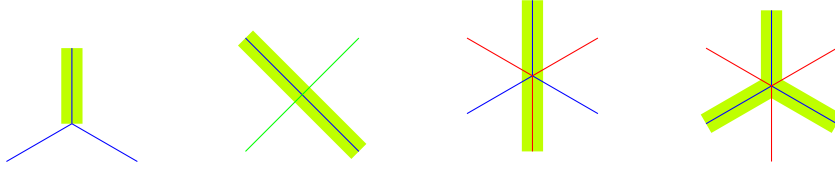
FIGURE 23. The local models for a Y-tree. The Y-tree is highlighted in light green.

Definition 3.4 allows us to associate a unique absolute cycle on $\Lambda(\mathfrak{w})$ (and hence on $L(\mathfrak{w})$) to each Y-tree in a weave $\mathfrak{w}$, as explained in [CZ22, Section 2]. (There are two conventions regarding orientations and choice of sheet at which to lift, but once these conventions are fixed, the absolute cycle is defined uniquely.) Note that a Y-cycle can stack multiple copies of the above patterns at the same vertex: when this happens at a trivalent or hexagonal vertex, the stacking creates self-intersections of the absolute cycle it represents. The distinction between *embedded* and *immersed* representatives of absolute homology classes is at the core of the distinction between *mutable* and *frozen* variables for the cluster structures we construct. In fact, the outstanding role of Y-trees is justified by the following fact.

Proposition 3.5. *Let $\mathfrak{w} \subset \mathbb{R}^2$ be a free weave and δ be an absolute 1-cycle representing a homology class in $H_1(L(\mathfrak{w})) \cong H_1(\Lambda(\mathfrak{w}))$ which is obtained from a Y-tree in $\mathfrak{w}$. Then, there exists a weave equivalence $\mathfrak{w} \sim \mathfrak{w}'$ such that the cycle $\delta \subset \mathfrak{w}$ becomes a short I-cycle in $\mathfrak{w}'$.*

In consequence, any homology class in $H_1(L(\mathfrak{w}))$ represented by a Y-tree admits an embedded representative $\gamma \subset L(\mathfrak{w})$ which bounds an embedded exact Lagrangian disk in $\mathbb{R}^4 \setminus L(\mathfrak{w})$.

Proof. This readily follows from [CZ22], by applying the equivalence moves in Figure 22 and keeping track of the change of a Y-tree under these moves. In fact, it suffices to use of the following two local weave equivalences:
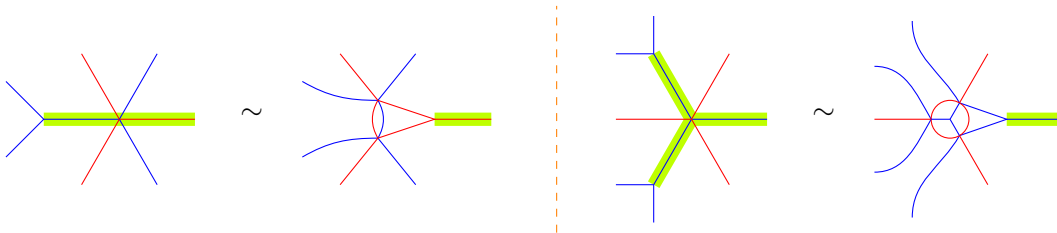
FIGURE 24. Local weave equivalences to turn a Y-tree into a short I-cycle. Note that the first row is just Move II from Figure 22, where we kept track of the Y-tree – highlighted in light green – before and after the equivalence.

By using the two weave equivalences in Figure 24, we can work outside in on the Y-tree δ and replace each weave line of δ with a double track, and shorten δ to a short I-cycle somewhere along the original Y-tree. The double tracks that appear in this shortening process are schematically depicted in Figure 25.

33

FIGURE 25. (Left) A Y-tree cycle highlighted in light green. (Right) The double tracks that remain on the (equivalent) weave after the shortening process, where the Y-cycle has now become the short I-cycle drawn in light green.

The second half of the proposition follows from the description of a short I-cycle in [CZ22]. □

The existence of the embedded Lagrangian disks from Proposition 3.5, i.e. $\mathbb{L}$-compressible curves in $L(\mathfrak{w})$, allows us to perform Lagrangian disk surgeries along Y-trees and produce new exact Lagrangian fillings. In [CZ22, Section 4.8] we proved that it is possible to describe this symplectic geometric operation via a diagrammatic change in a piece of the weave called "weave mutations". This leads to the following definition.

Definition 3.6. Let $\gamma \subset \mathfrak{w}$ be the short I-cycle – a monochromatic blue edge – depicted at the left of Figure 26 (left). Then, the local move illustrated in Figure 26 (left) is said to be a *weave mutation* at the short I-cycle γ. This is the standard Whitehead move for trivalent graphs, dual to a flip in a triangulation. Note that a weave mutation replaces the short I-cycle γ with a new short I-cycle, which we often denote it by γ'; we call γ' the *image* of γ under the weave mutation.

Definition 3.7. For a Y-tree in general, one can apply Proposition 3.5 to turn it into a short I-cycle, perform a weave mutation, and then apply some other weave equivalences. Thus, a *weave mutation* at a Y-tree γ in general means a weave mutation at its short I-cycle counterpart conjugated by sequences of weave equivalences. By following the sequences of weave equivalences, the weave mutation replaces γ by its *image*, which is a new Y-tree γ'.
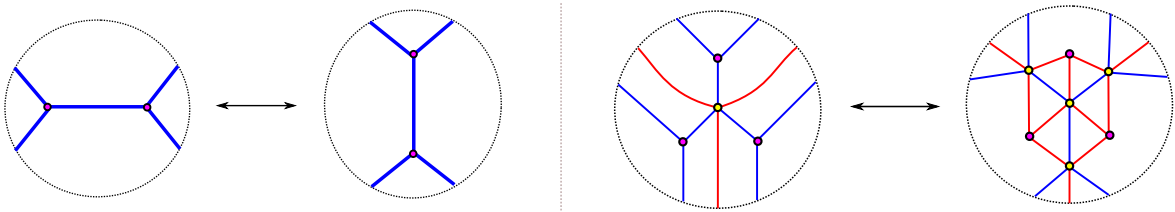
FIGURE 26. (Left) Weave mutation along the short I-cycle given by the unique monochromatic blue edge. (Right) A mweave mutation along monochromatic Y-trees.

Definition 3.8. Two weaves $\mathfrak{w}, \mathfrak{w}'$ are said to be (weave) *mutation equivalent* if they can be connected by a sequence of moves consisting of weave equivalences and weave mutations.

Remark 3.9. We observe that weave mutations generalize square moves, a.k.a. urban renewal, depicted in Figure 28 (left), for plabic graphs. Namely, if $\mathbb{G}, \mathbb{G}'$ are two GP-graphs related by a sequence of square moves. Then the weave surfaces $\mathfrak{w}(\mathbb{G}), \mathfrak{w}(\mathbb{G}')$, as constructed in Subsection 3.3, are weave mutation equivalent. Indeed, it suffices to argue in the 2-weave case for the local models depicted in Figure 28 (left). The associated weaves are drawn in depicted in Figure 28 (right), which is indeed a weave mutation. □

34
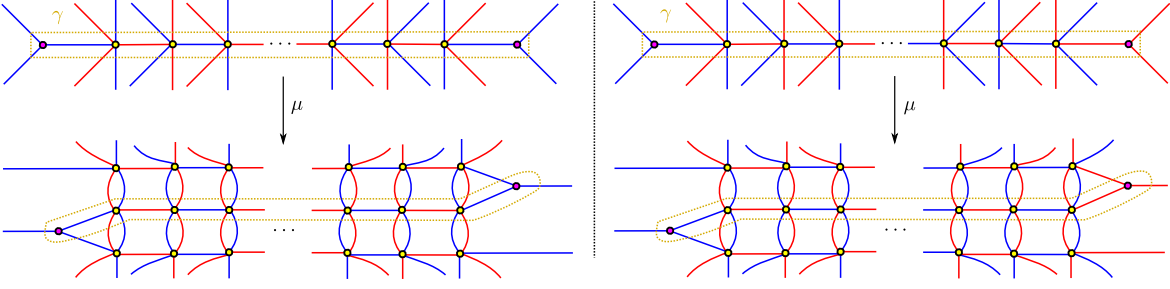
FIGURE 27. Both diagrams, on the left and on the right, depict weave muta-tions along long I-cycles, which are highlighted in the dashed yellow lines. The diagrams also illustrate the image of those long I-cycles after weave mutation.
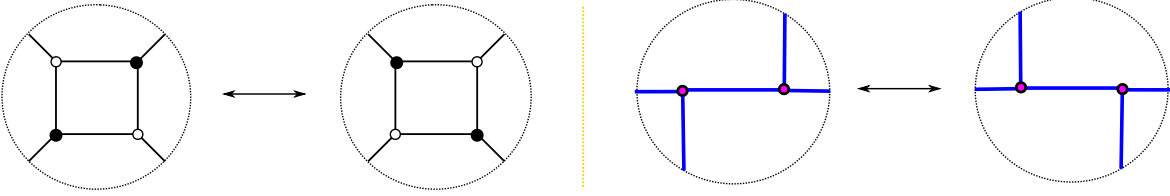
FIGURE 28. (Left) The square move for plabic graphs. (Right) The associated move for the corresponding weaves, whose construction is the content of the upcoming Subsection 3.3.

Finally, we emphasize that weave mutations will allow us to mutate at Y-trees of $\mathfrak{w}(\mathbb{G})$ corresponding to faces and regions of $\mathbb{G}$, even if they might not be square. The resulting weave, typically, will not be of the form $\mathfrak{w}(\mathbb{G}')$ for any GP-graph $\mathbb{G}'$, but all the diagrammatic and symplectic geometric results developed in this article and [CZ22] can still be applied.

3.3. **Initial Weave for a GP-Graph.** In this section we construct the initial weave $\mathfrak{w}(\mathbb{G})$ associated to a GP-graph $\mathbb{G}$. This is done by breaking $\mathbb{G}$ into elementary columns and assigning a local weave associated to each such column. Recall the standard reduced word $\mathtt{w}_{0,n} = s_{[1,1]}s_{[2,1]}s_{[3,1]}\cdots s_{[n-1,1]}$ for the longest element $w_{0,n} \in S_n$. Let us define $\ell = \ell(\mathtt{w}_{0,n}) = n(n-1)/2$. The first three local weaves $\mathfrak{n}(w), \mathfrak{c}^\uparrow(w)$ and $\mathfrak{c}^\downarrow(w)$ are defined as follows.

Definition 3.10. Let $w = s_{i_1}\cdots s_{i_\ell}$ be a reduced expression for $w_{0,n} \in S_n$. By definition, the weave $\mathfrak{n}(w)$ is given by n horizontal parallel weave lines such that the j-th strand, counting from the bottom, is labeled by the transposition s_{i_j}, $j \in [1,\ell]$.

Definition 3.11. The weave $\mathfrak{c}^\uparrow(w)$ is given by the weave $\mathfrak{n}(w)$ where a trivalent vertex is added at the top strand – labelled by s_{i_ℓ} – such that the third leg of this trivalent vertex is a vertical ray starting at the top strand and continuing upwards.

Definition 3.12. Similarly, the weave $\mathfrak{c}^\downarrow(w)$ is given by the weave $\mathfrak{n}(w)$ where a trivalent vertex is added at the bottom strand – labelled by s_{i_1} – such that the third leg of this trivalent vertex is a vertical ray starting at the bottom strand and continuing downwards.

As explained above, the weave $\mathfrak{w}(\mathbb{G})$ associated to $\mathbb{G}$ is built by horizontally concatenating weaves local models: each local model is associated to one of the three types of elementary columns. The corresponding weaves for each of these occurrences are described as follows.

3.3.1. *Local weaves for Type 1 Columns.*

Definition 3.13 (Weave for Type 1). The weave associated to a Type 1 column of a GP-graph $\mathbb{G}$, which consists of n parallel horizontal lines, is $\mathfrak{n}(\mathtt{w}_{0,n})$, where $\mathtt{w}_{0,n}$ is the standard reduced expression for the longest element in a symmetric group S_k.

It is important to note that the s_i-transpositions labeling the strands of $\mathfrak{n}(\mathtt{w}_{0,n})$ depend on the simple transpositions that generate the symmetric group $S_{[a,b]}$ associated to that specific region (see Subsection 2.1). Due to the appearance of lollipops in the GP-graph, the different symmetric groups $S_{[a,b]}$ that we encounter (as we read the GP-graph left to right) have varying discrete intervals $[a,b]$. This was illustrated in Figure 4.

3.3.2. *Local weaves for Type 2 Columns.* It is a well-known property of the symmetric group that any two reduced word expressions for the same element can be transformed into each other via finite sequences of the following two moves (see for instance [BB05, Section 3.3]):

- $s_i s_j \sim s_j s_i$, if $|i - j| > 1$;
- $s_i s_j s_i \sim s_j s_i s_j$, if $|i - j| = 1$.

Now consider the weave $\mathfrak{n}(\mathtt{w}_{0,n})$ and the s_1-strand labelled by the i-th appearance of s_1 in the standard reduced expression $\mathtt{w}_{0,n}$. In order to construct the weave for a Type 2 column, in Definition 3.17 below, we need the following auxiliary local weaves:

Definition 3.14. The weave $\mathfrak{n}_i^\uparrow(\mathtt{w}_{0,n})$ is the unique horizontal weave that coincides with $\mathfrak{n}(\mathtt{w}_{0,n})$ at the left, contains only tetravalent and hexagonal vertices, and brings the i-th s_1-strand of $\mathtt{w}_{0,n}$ to the top level, using a minimal number of weave vertices.

Similarly, the weave $\mathfrak{n}_i^\downarrow(\mathtt{w}_{0,n})$ is the unique weave that coincides with $\mathfrak{n}(\mathtt{w}_{0,n})$ at the left, contains only tetravalent and hexagonal vertices, and brings the i-th s_1-strand of $\mathtt{w}_{0,n}$ to the bottom level, using a minimal number of weave vertices.

Finally, we denote by $\mathfrak{n}_i^\uparrow(\mathtt{w}_{0,n})^{op}$ and $\mathfrak{n}_i^\downarrow(\mathtt{w}_{0,n})^{op}$ the weaves obtained by reflecting $\mathfrak{n}_i^\uparrow(\mathtt{w}_{0,n})$ and $\mathfrak{n}_i^\downarrow(\mathtt{w}_{0,n})$ along a (disjoint) vertical axis. □

In Definition 3.14, bringing the i-th s_1-strand of $\mathtt{w}_{0,n}$ to the top level means considering a horizontal weave which starts at $\mathfrak{n}(\mathtt{w}_{0,n})$ on the left hand side and contains a sequence of tetravalent and hexagonal vertices (no trivalent vertices) such that following the i-th s_1-strand of $\mathtt{w}_{0,n}$ under these vertices (passing through them straight) ends up at the top strand at the right hand side. There are many weaves that verify this property, but by the Zamolodchikov relation proven in [CZ22], they are all equivalent and we might as well take the one with a minimal number of vertices. Figure 29 illustrates several examples of the weaves $\mathfrak{n}_i^\uparrow(\mathtt{w}_{0,n})$ and $\mathfrak{n}_i^\downarrow(\mathtt{w}_{0,n})$ in Definition 3.14 in the cases of $n = 3, 4$. Note that $\mathfrak{n}_{n-1}^\uparrow(\mathtt{w}_{0,n}) = \mathfrak{n}_{n-1}^\uparrow(\mathtt{w}_{0,n})^{op} = \mathfrak{n}_1^\downarrow(\mathtt{w}_{0,n}) = \mathfrak{n}_1^\downarrow(\mathtt{w}_{0,n})^{op} = \mathfrak{n}(\mathtt{w}_{0,n})$ for any $n \in \mathbb{N}$.
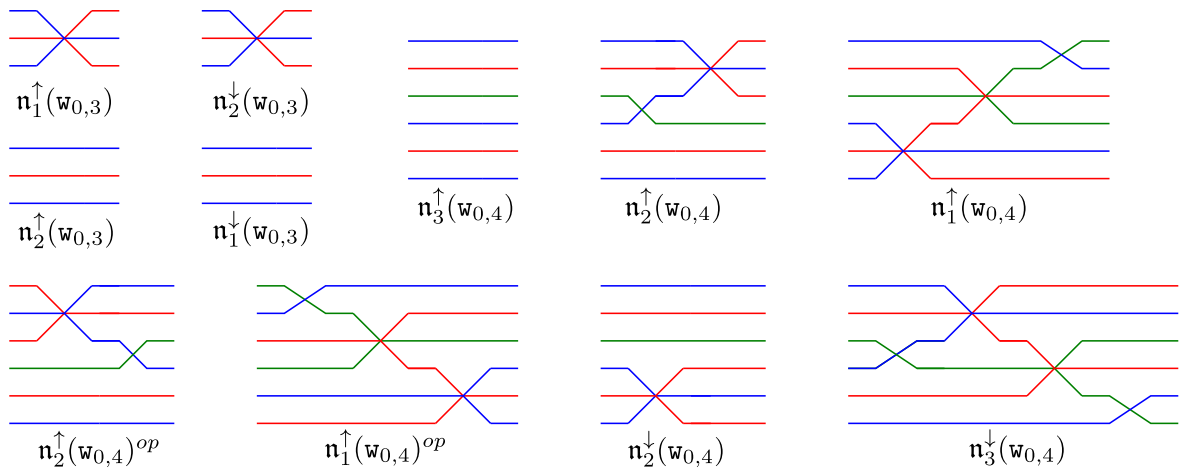

FIGURE 29. Instances of the weaves $\mathfrak{n}_i^\uparrow(\mathtt{w}_{0,n})$ and $\mathfrak{n}_i^\downarrow(\mathtt{w}_{0,n})$, and their opposites, from Definition 3.14, in some of the cases for $n = 3, 4$.

Definition 3.15. The weave $\mathfrak{c}_i^{\uparrow}(\mathtt{w}_{0,n})$ is the weave obtained by horizontally concatenating the three weaves $\mathfrak{n}_i^{\uparrow}(\mathtt{w}_{0,n})$, $\mathfrak{c}^{\uparrow}(w_i)$, and $\mathfrak{n}_i^{\uparrow}(\mathtt{w}_{0,n})^{op}$, left to right, where w_i denotes the reduced expression for $\mathtt{w}_{0,n}$ found at the right of the weave $\mathfrak{n}_i^{\uparrow}(\mathtt{w}_{0,n})$.

Definition 3.16. Similarly, the weave $\mathfrak{c}_i^{\downarrow}(\mathtt{w}_{0,n})$ is the weave obtained by horizontally concatenating the three weaves $\mathfrak{n}_i^{\downarrow}(\mathtt{w}_{0,n})$, $\mathfrak{c}^{\downarrow}(w_i)$ and $\mathfrak{n}_i^{\downarrow}(\mathtt{w}_{0,n})^{op}$, left to right, where w_i denotes the reduced expression for $\mathtt{w}_{0,n}$ found at the right of the weave $\mathfrak{n}_i^{\downarrow}(\mathtt{w}_{0,n})$.

Figures 31 and 32 illustrate examples of the weaves $\mathfrak{c}_i^{\uparrow}(\mathtt{w}_{0,n})$ and $\mathfrak{c}_i^{\downarrow}(\mathtt{w}_{0,n})$ in Definitions 3.15 and 3.16 in the cases of $n = 3, 4$. For the next definition, recall that we always label the horizontal lines, in a Type 1 and Type 2 column of the GP-graph $\mathbb{G}$, with consecutive natural numbers, from bottom to top. For the moment, let us assume that these labels are in $[1, n]$.

Definition 3.17 (Weave for Type 2). For $i \in [1, n-1]$, the weave associated to a Type 2 column of a GP-graph $\mathbb{G}$ whose vertical edge has a white vertex at the ith horizontal line and a black vertex at the $(i+1)$st horizontal line is the weave $\mathfrak{c}_i^{\uparrow}(\mathtt{w}_{0,n})$, and the weave associated to a Type 2 column of a GP-graph $\mathbb{G}$ whose vertical edge has a black vertex at the ith horizontal line and a white vertex at the $(i+1)$st horizontal line is the weave $\mathfrak{c}_{n-i}^{\downarrow}(\mathtt{w}_{0,n})$.

Remark 3.18. It is possible to horizontally concatenate the weaves $\mathfrak{c}_i^{\uparrow}(\mathtt{w}_{0,n})$ and $\mathfrak{c}_j^{\downarrow}(\mathtt{w}_{0,n})$, for any pair $i, j \in [1, n-1]$, in two ways: $\mathfrak{c}_i^{\uparrow}(\mathtt{w}_{0,n})\mathfrak{c}_j^{\downarrow}(\mathtt{w}_{0,n})$ and $\mathfrak{c}_j^{\downarrow}(\mathtt{w}_{0,n})\mathfrak{c}_i^{\uparrow}(\mathtt{w}_{0,n})$. These resulting two weaves are weave equivalent if and only if $i \neq j$. Furthermore, $\mathfrak{c}_i^{\uparrow}(\mathtt{w}_{0,n})\mathfrak{c}_i^{\downarrow}(\mathtt{w}_{0,n})$ and $\mathfrak{c}_i^{\downarrow}(\mathtt{w}_{0,n})\mathfrak{c}_i^{\uparrow}(\mathtt{w}_{0,n})$ differ exactly by one weave mutation. This is the weave analogue of one of the mutation rules for double wiring diagrams [FZ99]. $\square$

3.3.3. *Local weaves for Type 3 Columns.* Let us consider a column of Type 3 with labels $1, 2, \ldots, n$ for the horizontal lines on the right (counting from botom to top) and with a white lollipop attached to the ith horizontal line with $i \in [1, n]$. The case of a black lollipop is similar, and discussed later.

By construction, the weaves associated with the two Type 1 columns sandwiching this Type 3 column are $\mathfrak{n}(\mathtt{w}_{0,n-1})$ and $\mathfrak{n}(\mathtt{w}_{0,n})$, respectively. Hence, the weave we associate to such a Type 3 column must have these boundary conditions. Let us start with the following weave:
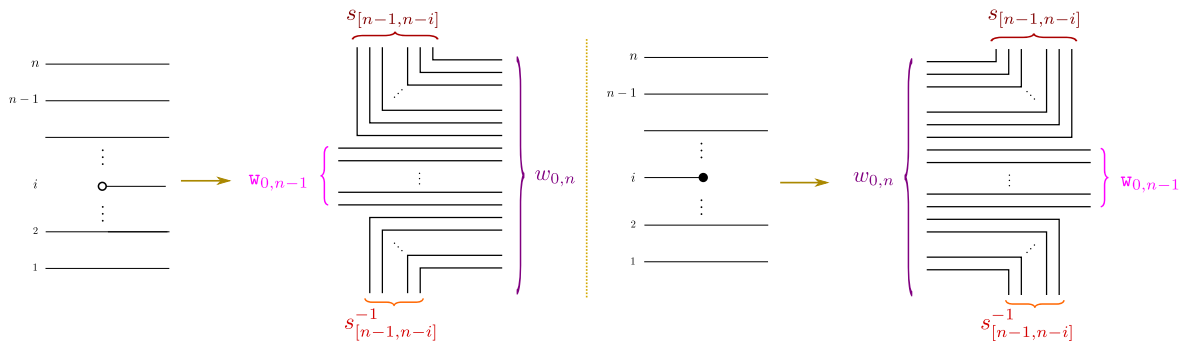
FIGURE 30. The weaves $\mathfrak{i}_i^w$ (left) and $\mathfrak{i}_i^b$ (right) from Definitions 3.19 and 3.22.

Definition 3.19. The weave $\mathfrak{i}_i^w$ is the unique weave with no weave vertices, satisfying:

(i) At its left, $\mathfrak{i}_i^w$ coincides with the horizontal weave $\mathfrak{n}(\mathtt{w}_{0,n-1})$, and at its right, $\mathfrak{i}_i^w$ coincides with the horizontal weave $\mathfrak{n}(s_{[n-1,i+1]}^{-1}\mathtt{w}_{0,n-1}s_{[n-1,n-i]})$.

37
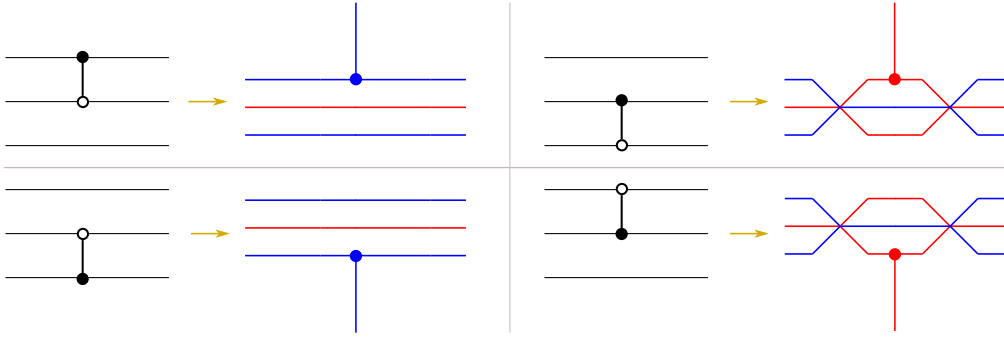
FIGURE 31. All possible Type 2 columns for $n = 3$ and their corresponding weaves. In detail, $\mathfrak{c}_2^{\uparrow}(\mathtt{w}_{0,3})$ (upper left), $\mathfrak{c}_1^{\uparrow}(\mathtt{w}_{0,3})$ (upper right), $\mathfrak{c}_1^{\downarrow}(\mathtt{w}_{0,3})$ (lower left) and $\mathfrak{c}_2^{\downarrow}(\mathtt{w}_{0,3})$ (lower right).
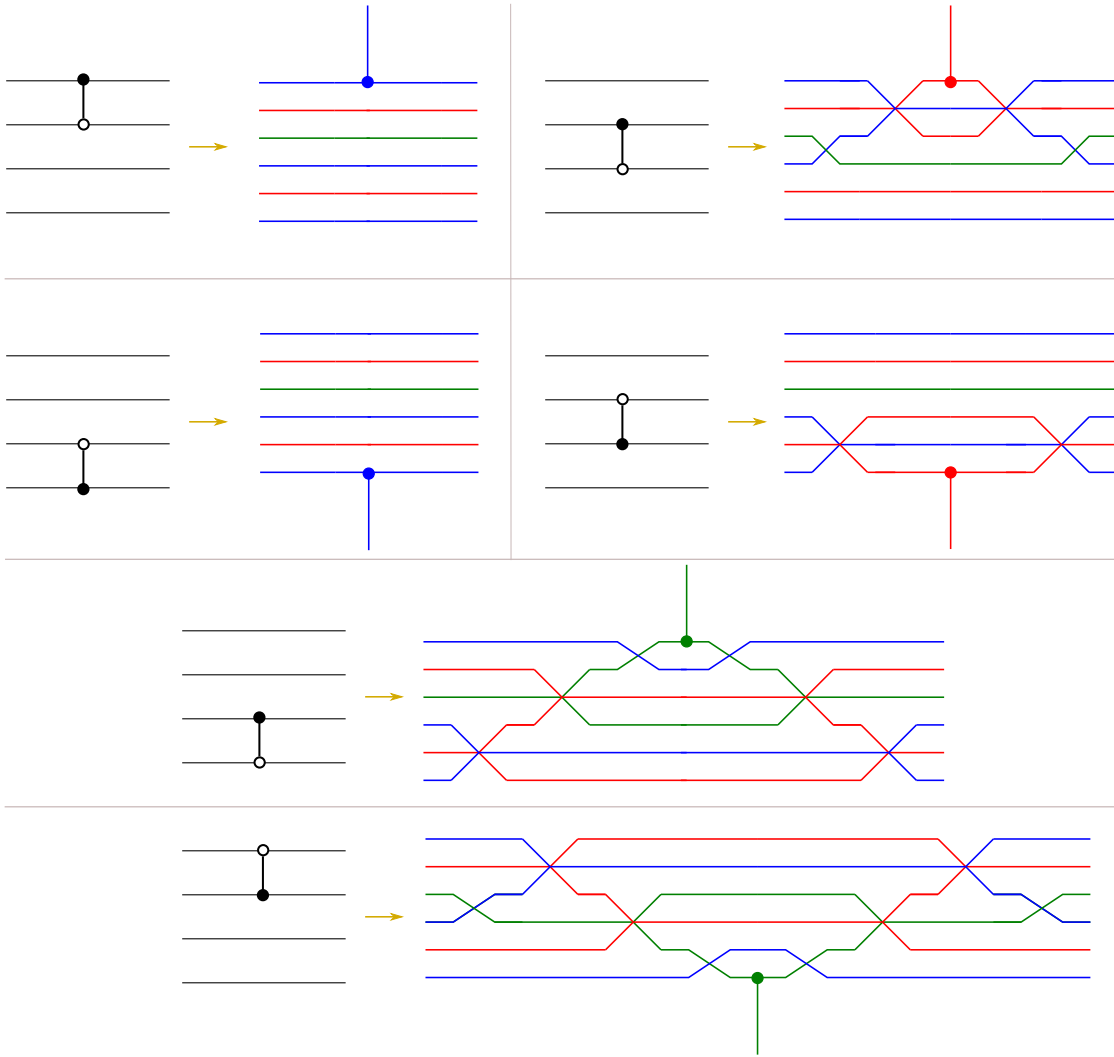
FIGURE 32. Some of the Type 2 columns for $n = 4$ and their corresponding weaves. In detail, $\mathfrak{c}_3^{\uparrow}(\mathtt{w}_{0,4})$ (1st row left), $\mathfrak{c}_2^{\uparrow}(\mathtt{w}_{0,4})$ (1st row right), $\mathfrak{c}_1^{\downarrow}(\mathtt{w}_{0,4})$ (2nd row left), $\mathfrak{c}_2^{\downarrow}(\mathtt{w}_{0,4})$ (2nd row right), $\mathfrak{c}_1^{\uparrow}(\mathtt{w}_{0,4})$ (3rd row) and $\mathfrak{c}_3^{\downarrow}(\mathtt{w}_{0,4})$ (4th row).

(ii) The weave lines in $\mathfrak{n}(s^{-1}_{[n-1,i+1]}\mathsf{w}_{0,n-1}s_{[n-1,n-i]})$ labeled by the transpositions in the reduced expression $s_{[n-1,n-i]}$ diverge upwards to vertical rays.

(iii) The weave lines in $\mathfrak{n}(s^{-1}_{[n-1,i+1]}\mathsf{w}_{0,n-1}s_{[n-1,n-i]})$ labeled by the transpositions in the reduced expression $s^{-1}_{[n-1,n-i]}$ diverge downwards to vertical rays.

See Figure 30 (left) for a depiction of $\mathfrak{i}^w_i$, illustrating what is meant by diverging upwards and downwards to vertical rays.

Note that the word $\mathfrak{n}(s^{-1}_{[n-1,i+1]}\mathsf{w}_{0,n-1}s_{[n-1,n-i]})$ in Definition 3.19 is a reduced expression for the half-twist $w_{0,n}$.

Remark 3.20. There is a geometric reason behind $\mathfrak{i}^w_i$, related to the front $\mathfrak{f}(\mathbb{G})$ of $\Lambda(\mathbb{G})$. Indeed, the appearance of both $s^{-1}_{[n-1,i+1]}$ and $s_{[n-1,n-i]}$ in $\mathfrak{n}(s^{-1}_{[n-1,i+1]}\mathsf{w}_{0,n-1}s_{[n-1,n-i]})$ is caused by the new emerging crossings when one performs the sequence of Reidemeister II moves as indicated in Figure 33 (left). In this sense, a white lollipop introduces new crossings, which translate into new weave lines.
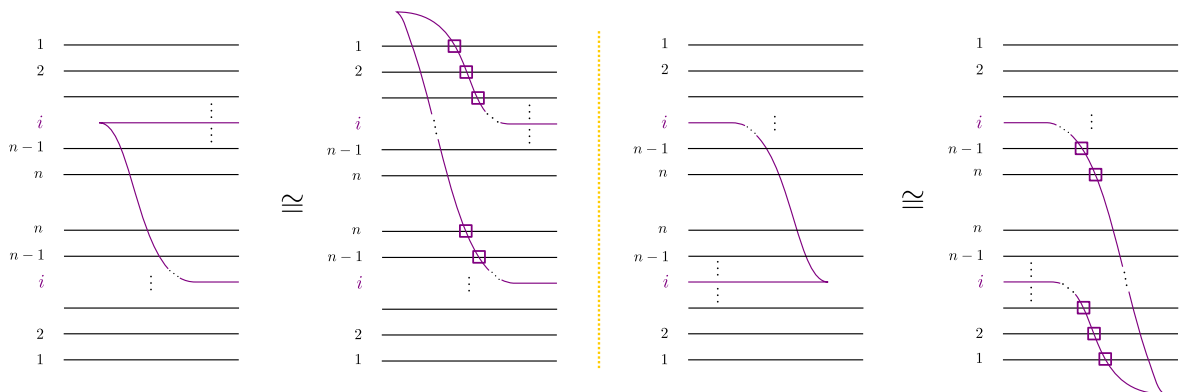
FIGURE 33. A sequence of RII moves in the front projection, which explain the incoming (resp. outcoming) strands in a weave when a white (resp. black) lollipop appears. Note that the level at which the crossing occurs dictated the label of the weave strand, according to Definition 3.21 (resp. 3.23).

Now, the weaves $\mathfrak{i}^w_i$ in Definition 3.19 cannot quite be the weaves for the Type 3 column yet because the labeling on the right hand side is not $\mathsf{w}_{0,n}$, but rather $\mathfrak{n}(s^{-1}_{[n-1,i+1]}\mathsf{w}_{0,n-1}s_{[n-1,n-i]})$. To fix this, let $\mathfrak{n}^w_i$ be any horizontal weave that coincides with $\mathfrak{n}(s^{-1}_{[n-1,n-i+1]}\mathsf{w}_{0,n-1}s_{[n-1,i]})$ on the left, coincides with $\mathfrak{n}(\mathsf{w}_{0,n})$ on the right, and with no trivalent weave vertices in the middle. Any choice of $\mathfrak{n}(\mathsf{w}_{0,n})$ would yield and equivalent weave.

Definition 3.21 (Weave for White Lollipop). The weave $\mathfrak{l}^w_i$ associated to a Type 3 column with a white lollipop at the ith horizontal line is the horizontal concatenation of $\mathfrak{i}^w_i$ and $\mathfrak{n}^w_i$.

Figures 34 and 35 illustrate the weaves $\mathfrak{l}^w_1, \mathfrak{l}^w_2, \mathfrak{l}^w_3$, and $\mathfrak{l}^w_4$ for $n = 4$, with the coloring convention that s_1 is blue, s_2 is red and s_3 is green. The pink boxes in the figures contain the $\mathfrak{i}^w_i$ pieces, and the yellow boxes contain the $\mathfrak{n}^w_i$ pieces. The figures also draw the corresponding pieces of the fronts $\mathfrak{f}(\mathbb{G})$, explaining the contact geometric origin of these weaves.

The case of a column of Type 3 with labels $1, 2, \ldots, n$ for the horizontal lines on the right, and a black lollipop at the ith horizontal line is similar. The necessary definitions are as follows.

Definition 3.22. The weave $\mathfrak{i}^b_i$ is the unique weave with no weave vertices, satisfying:

39
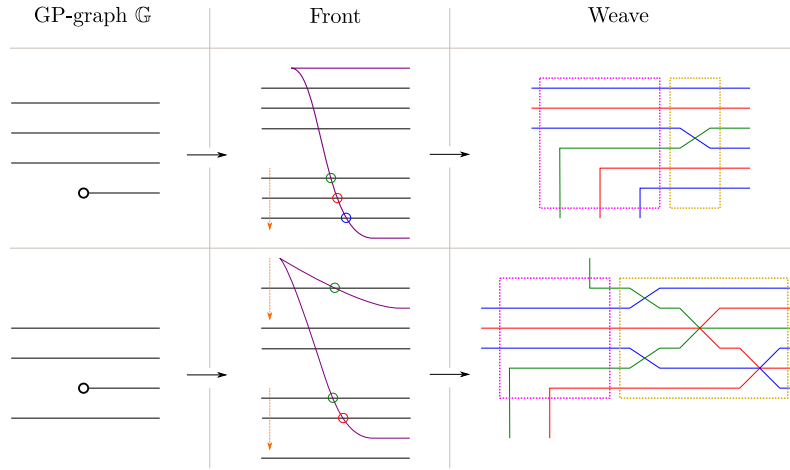
FIGURE 34. Weaves $\mathfrak{l}_i^w$ associated to a white lollipop with $n = 4$, as in Definition 3.21. The first row depicts the case $i = 1$ and the second row the case $i = 2$. The weaves $\mathfrak{n}_i^w$ are drawn within the yellow boxes. The weaves $\mathfrak{i}_i^w$, with the incoming weave strands, are depicted within the pink boxes.
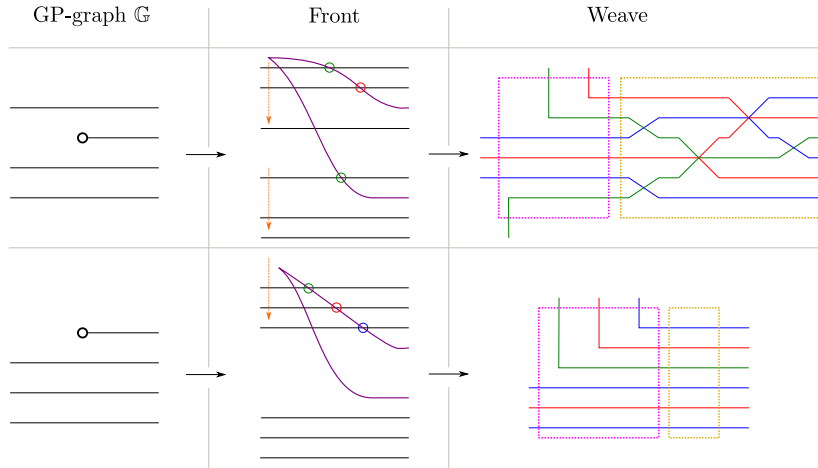
FIGURE 35. Weaves $\mathfrak{l}_i^w$ associated to a white lollipop with $n = 4$. The first row depicts the case $i = 3$ and the second row the case $i = 4$. The weaves $\mathfrak{n}_i^w$ are drawn within the yellow boxes, and the weaves $\mathfrak{i}_i^w$ are depicted in the pink boxes.
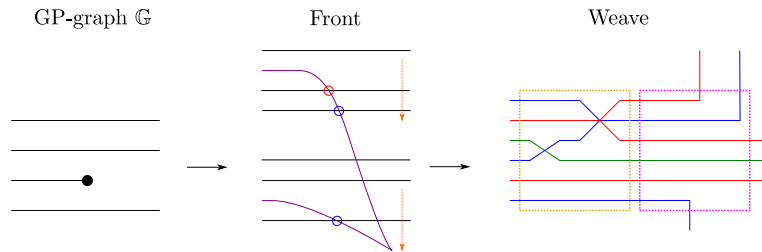
FIGURE 36. The weave $\mathfrak{l}_2^b$ associated to a black lollipop when $n = 4$, in accordance with Definition 3.23. The weave $\mathfrak{n}_2^b$ is drawn in the yellow box, and $\mathfrak{i}_2^b$ in the pink box. Note that the departing strands of the weave (in the pink box) are in bijection with the crossings of the front, as indicated.

(i) At its left, $\mathfrak{i}_i^b$ coincides with the horizontal weave $\mathfrak{n}(s_{[i-1,1]}^{-1}\mathsf{w}_{0,n-1}s_{[n-i+1,1]})$, and at its right, $\mathfrak{i}_i^b$ coincides with the horizontal weave $\mathfrak{n}(\mathsf{w}_{0,n-1})$.

(ii) The weave lines in $\mathfrak{n}(s_{[i-1,1]}^{-1}\mathsf{w}_{0,n-1}s_{[n-i+1,1]})$ labeled by the transpositions in the reduced expression $s_{[i-1,1]}^{-1}$ diverge downwards to vertical rays.

(iii) The weave lines in $\mathfrak{n}(s_{[i-1,1]}^{-1}\mathsf{w}_{0,n-1}s_{[n-i+1,1]})$ labeled by the transpositions in the reduced expression $s_{[n-i+1,1]}$ diverge upwards to vertical rays.

See Figure 30 (right) for a depiction of $\mathfrak{i}_i^b$. Similarly, we denote by $\mathfrak{n}_i^b$ any horizontal weave which coincides with the horizontal weave $\mathfrak{n}(s_{[i-1,1]}^{-1}\mathsf{w}_{0,n-1}s_{[n-i+1,1]})$ on the right, coincides with $\mathfrak{n}(\mathsf{w}_{0,n})$ on the left, and with no trivalent weave vertices in the middle.

Definition 3.23 (Weave for Black Lollipop). The weave $\mathfrak{l}_i^b$ associated to a Type 3 column with a black lollipop at the ith horizontal line is the horizontal concatenation of $\mathfrak{n}_i^b$ and $\mathfrak{i}_i^b$.

See Figure 36 for an example of $\mathfrak{l}_i^b$ in the case $i = 3$ and $n = 4$.

3.3.4. *The initial weave.* Let $\mathbb{G}$ be a GP-graph. The previous three Subsections 3.3.1, 3.3.2, 3.3.3 have explained how to obtain a weave from each of the three types of elementary columns. Note that each these weaves coincides with $\mathfrak{n}(\mathsf{w}_{0,k})$ and with $\mathfrak{n}(\mathsf{w}_{0,m})$ for some k and m at its two ends. For Type 1 and Type 2, $k = m$, and for Type 3 $|k - m| = 1$. Note that, if we consider two adjacent elementary columns in $\mathbb{G}$, the associated weaves will coincide at the common side, and thus can be horizontally concatenated.

Definition 3.24 (Initial weave). Let $\mathbb{G}$ be a GP-graph. The initial weave $\mathfrak{w}(\mathbb{G})$ associated weave $\mathbb{G}$ is the weave obtained by subdividing $\mathbb{G}$ into elementary columns and then concatenating the weaves associated with each elementary column, in the order dictated by the columns.
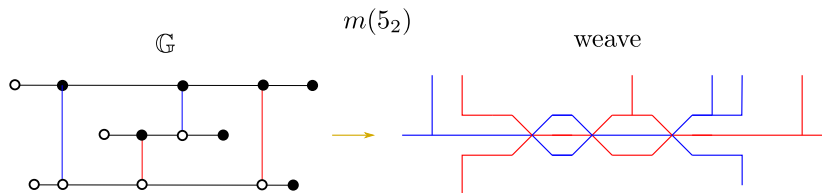
FIGURE 37. A GP-graph for a Legendrian $m(5_2)$ and its associated initial weave.

Example 3.25. *Consider the GP-graphs in the first row of Figure 16, whose Legendrian fronts are depicted in Figure 17. Then, Figure 37 illustrates the initial weave associated to the GP-graph $\mathbb{G}_1$ for $m(5_2)$ in Figure 16 (upper left). Note that the braid word associated to this initial weave is $\beta = (\sigma_2\sigma_1^2\sigma_2^2\sigma_1^2\sigma_2\sigma_1\sigma_2)\Delta_2^{-2} = \sigma_1^3\sigma_2\sigma_1^{-1}\sigma_2$, which is indeed a braid expression for the smooth type $m(5_2)$. Observe that, despite β being a strongly quasi-positive braid, $m(5_2)$ is not representable by a positive braid. Similarly, Figure 38 illustrates the initial weave associated to the GP-graph $\mathbb{G}_2$ for $m(7_2)$ in Figure 16 (upper right).* $\square$

3.4. **Topology of the Initial Weave.** Let $\mathbb{G}$ be a GP-graph and $\mathfrak{w}(\mathbb{G}) \subset \mathbb{R}^2$ its initial weave. In this subsection we show how to obtain a Legendrian link $\Lambda(\mathbb{G}) \subset (\mathbb{R}^3, \xi_{\mathrm{st}})$ and an embedded exact Lagrangian filling $L(\mathfrak{w}) \subset (\mathbb{R}^4, \lambda_{\mathrm{st}})$ from the weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$.
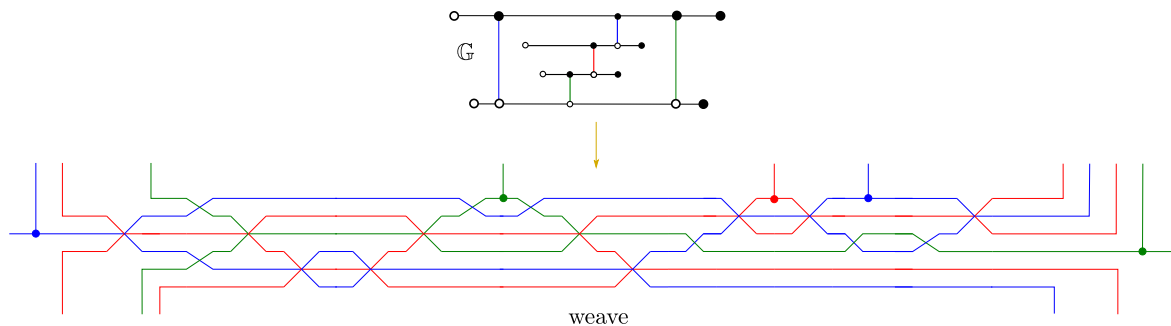
FIGURE 38. A GP-graph for a Legendrian $m(7_2)$ and its associated initial weave.

3.4.1. *The braid of an initial weave.* Suppose $\mathfrak{w}(\mathbb{G})$ is an N-weave. Let $K \subset \mathbb{R}^2$ be a compact subset such that $(\mathbb{R}^2 \setminus K) \cap \mathfrak{w}(\mathbb{G})$ contains no weave vertices. Note that the number of weave lines in $(\mathbb{R}^2 \setminus K) \cap \mathfrak{w}(\mathbb{G})$ and their labeling is independent of any K with such property. The weave lines are labeled by simple transpositions $s_i \in S_N$, which can be lifted to unique positive generators σ_i in the Artin braid group Br_N.

Definition 3.26. Let $\beta(\mathbb{G})$ be the positive braid word obtained by reading the positive braid generators associated with the weave lines of $(\mathbb{R}^2 \setminus K) \cap \mathfrak{w}(\mathbb{G})$ in a counterclockwise manner, starting at the unique strand the corresponds to the left-most white lollipop in $\mathbb{G}$.

Part of the usefulness of Definition 3.26 is the following lemma, which explains that the braid word $\beta(\mathbb{G})$ is actually a braid presentation for the alternating strand diagram of $\Lambda(\mathbb{G})$, understood as a Legendrian link in $(\mathbb{R}^3, \xi_{\mathrm{st}})$. See [CN21, Section 2] for details on the (-1)-framed closure of a positive braid word.

Lemma 3.27. *Let $\mathbb{G}$ be a GP-graph, $\mathfrak{w}(\mathbb{G}) \subset \mathbb{R}^2$ its initial weave and $\beta(\mathbb{G})$ its positive braid word. Then, the (-1)-framed closure of $\beta(\mathfrak{w}(\mathbb{G}))$ is a front for the Legendrian link $\Lambda(\mathbb{G})$.*

Proof. A front for $\Lambda(\mathbb{G})$ is provided in Subsection 2.4, it suffices to argue that it is homotopic to the (-1)-framed closure of $\beta(\mathfrak{w}(\mathbb{G}))$. The desired Legendrian isotopy is obtained by first performing the front homotopies in Figure 33 at each (internal) cusp corresponding to a lollipop, white or black. Then, we use a sequence of Reidemeister II and Reidemeister III moves to push the cusps to the extremes: the left cusps, coming from white lollipops, are sent to the left of the front, and the right cusps, coming from black lollipops, are sent to the right of the front. Once all cusps are pushed to the corresponding extreme of the front, one obtained the (-1)-framed closure of a positive braid. The fact that this positive braid is indeed $\beta(\mathfrak{w}(\mathbb{G}))$ then follows from the construction of the initial weave in Subsection 3.3. □

Lemma 3.27 can be phrased as follows. Consider the Legendrian link $\Lambda(\beta(\mathbb{G})) \subset (\mathbb{R}^3, \xi_{\mathrm{st}})$ whose front is the (-1)-framed closure of the braid word $\beta(\mathbb{G})$. Then the Legendrian links $\Lambda(\beta(\mathbb{G}))$ and $\Lambda(\mathbb{G})$ are Legendrian isotopic in $(\mathbb{R}^3, \xi_{\mathrm{st}})$.

3.4.2. *The surface of the initial weave.* Let $\mathfrak{w} \subset \mathbb{R}^2$ be a weave and $\Lambda(\mathfrak{w}) \subset (\mathbb{R}^5, \xi_{\mathrm{st}})$ the Legendrian represented by its front. By definition, the Lagrangian $L(\mathfrak{w}) \subset (\mathbb{R}^4, \lambda_{\mathrm{st}})$ is the Lagrangian projection of $\Lambda(\mathfrak{w})$. We refer to [CZ22, Section 7.1] for details on how weaves yield exact Lagrangian fillings of Legendrian links in $(\mathbb{R}^3, \xi_{\mathrm{st}})$, and recall that $\mathfrak{w}$ is said to be free if $L(\mathfrak{w}) \subset (\mathbb{R}^4, \lambda_{\mathrm{st}})$ is embedded.

Lemma 3.28. *Let $\mathbb{G}$ be a GP-graph and suppose its initial weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$ is an N-weave. Then $L = L(\mathfrak{w}) \subset (\mathbb{R}^4, \lambda_{st})$ is an embedded exact Lagrangian filling of $\Lambda(\mathbb{G})$ and*

$$\chi(L) = N - \#(\text{trivalent vertices of } \mathfrak{w})$$

42

$$= \#(horizontal\ lines\ in\ \mathbb{G}) - \#(vertical\ edges\ in\ \mathbb{G}).$$

where χ denotes the Euler characteristic.

Proof. Note that the first equality for $\chi(L)$ follows from the Riemann-Hurwitz formula, see [CZ22, Section 2.4]; the second equality follows from the facts that the number of trivalent vertices of $\mathfrak{w}$ equals the number of vertical edges of $\mathbb{G}$, and that N equals the number of horizontal lines in $\mathbb{G}$; exactness is guaranteed by virtue of being the Lagrangian projection of a Legendrian surface.

Thus, it remains to verify embeddedness for each of the elementary columns. For that, we note that $\mathfrak{n}(\mathtt{w}_{0,n})$ is a free weave, for all $n \in \mathbb{N}$. Indeed, the front for $\mathfrak{n}(\mathtt{w}_{0,n})$ can be chosen to have constant slope in the horizontal direction; any slice in the vertical direction yields a front for $\mathfrak{n}(\mathtt{w}_{0,n})$ which is a positive braid coinciding with the reduced expression $\mathtt{w}_{0,n}$. Thus, any pair of strands intersects exactly once, and the slopes between the corresponding sheets can be chosen to increase as they move away from the crossing of intersection. In this manner, no Reeb chords are created and $\mathfrak{n}(\mathtt{w}_{0,n})$ is free. Since $\mathfrak{n}(\mathtt{w}_{0,n})$ is a free weave, for all $n \in \mathbb{N}$, the weaves associated to Type 1 and Type 3 columns are free. For a Type 2 elementary column, we note that adding a trivalent vertex to any strand of a weave $\mathfrak{w}$ which has an open end (i.e. such that the strand goes off to infinity in the plane) gives an embedded Lagrangian cobordism from $\mathfrak{w}$ to a new weave. Thus, the new weave obtained by adding a trivalent vertex is free if $\mathfrak{w}$ is free. Hence, given that a Type 2 column is obtained by adding a trivalent vertex to $\mathfrak{n}(\mathtt{w}_{0,n})$, for some $n \in \mathbb{N}$, its associated weave is always free, as required. $\square$

The number of boundary components of $L(\mathfrak{w}(\mathbb{G}))$ is not as easy to read from $\mathbb{G}$, but it is simply computed from $\beta(\mathbb{G})$: it is given by the number of cycles in the cycle decomposition of the Coxeter projection of $\beta(\mathbb{G})$. That is, we apply the natural projection map $\mathrm{Br}_N \to S_N$ to $\beta(\mathbb{G})$ and compute the number of cycles associated to the resulting permutation. Note that the Euler characteristic $\chi(L(\mathfrak{w}(\mathbb{G})))$ and the number of boundary components of $L(\mathfrak{w}(\mathbb{G}))$ uniquely determine the smooth topology of $L(\mathfrak{w}(\mathbb{G}))$ as an abstract surface.[10]

Remark 3.29. In upcoming work of the first author with W. Li, we prove that $L(\mathfrak{w}(\mathbb{G}))$ is Hamiltonian isotopic to an exact Lagrangian representative of the Goncharov-Kenyon conjugate surface [GK13] associated to the graph $\mathbb{G}$, embedded as in [STWZ19, Proposition 4.9]. $\square$

Finally, a central feature of weaves is the following: it is possible to draw many weaves which coincide with $\mathfrak{w}(\mathbb{G})$, outside a large enough compact set $K \subset \mathbb{R}^2$, and which represent embedded exact Lagrangian fillings of $\Lambda(\mathbb{G})$. In fact, as explained in Subsection 3.2, there are some local modifications that we can perform to the weave – *weave mutations* – such that the smooth embedded class of the associated (Lagrangian) surface in $(\mathbb{R}^4, \lambda_{\mathrm{st}})$ remains the same but the Hamiltonian isotopy class typically changes. Square face mutation of a GP-graph $\mathbb{G}$ is recovered by weave mutations but, importantly, weave mutations allow for more general mutations, including mutations at non-square faces of $\mathbb{G}$ and sugar-free regions. The result of such *weave mutations* applied to $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$ is again another weave $\mu(\mathfrak{w})$: it may no longer be of the form $\mathfrak{w}(\mathbb{G}')$ for any GP-graph $\mathbb{G}'$, but it is a weave and thus, using the calculus in [CZ22], we can manipulate it efficiently and use it to prove the results in this article. In this process, we need explicit geometric cycles representing generators of the absolute homology $H_1(L(\mathfrak{w}))$. In fact, such geometric cycles lead to the quiver for the initial seed. Thus, we now gear towards understanding how to construct geometric representatives of homology classes using weaves.

---

[10]Its embedded topology, especially up to Hamiltonian isotopy, is a different matter. That requires a deeper understanding of the weave $\mathfrak{w}(\mathbb{G})$, see [CZ22].

3.5. **Naive Absolute Cycles in** $L(\mathfrak{w}(\mathbb{G}))$. In this section we explain how to find a set of geometric (absolute) cycles on $L = L(\mathfrak{w}(\mathbb{G}))$ which generate $H_1(L)$.

Since the genus of an embedded exact Lagrangian filling is determined by the (maximal) Thurston-Bennequin invariant of its Legendrian boundary, all embedded exact Lagrangian fillings of a given Legendrian link are topologically equivalent as abstract surfaces, i.e. they have the same genus. In the case of a Legendrian link $\Lambda(\mathbb{G})$ associated with a GP-graph, it is readily seen that this is the same abstract topological type as that of the Goncharov-Kenyon conjugate surface $S = S(\mathbb{G})$ [GK13]. Since the conjugate surface S deformation retracts back to the GP-graph $\mathbb{G}$, it follows that the boundaries of the faces of $\mathbb{G}$ form a basis for the absolute homology groups $H_1(\mathbb{G}) \cong H_1(S) \cong H_1(L)$. This basis, indexed by the faces of $\mathbb{G}$, will be referred to as the *naive basis* of $H_1(L)$.

Remark 3.30. We emphasize that this set of generating absolute cycles is *not* good enough in order to construct cluster structures, nor its intersection quiver gives the correct initial quiver. Thus, these cycles will be referred to as the set of *naive* absolute cycles, and we will perform the necessary corrections in Subsection 3.7 below.

In order to proceed geometrically, we would like identify the naive basis elements of $H_1(L)$ as lifts of a specific collection of absolute cycles on the weave front $\Sigma = \Sigma(\mathfrak{w}(\mathbb{G}))$, ideally a collection of Y-cycles on $\mathfrak{w}(\mathbb{G})$ (Definition 3.4). Since any GP-graph $\mathbb{G}$ can be decomposed into elementary columns, we can try to build these absolute cycles by concatenating appropriate relative cycles associated with each elementary column.

3.5.1. *Local representatives of naive absolute cycles in a Type 1 Column.* In an elementary Type 1 column of $\mathbb{G}$ with n horizontal lines, there are $n - 1$ faces, i.e. gaps, between these n horizontal lines. For each of these $n - 1$ gaps, we identify a unique weave line as follows. First, we observe that a cross-section of the weave front Σ associated to a Type 1 column is, by construction, the reduced expression $\mathtt{w}_{0,n}$ of $w_{0,n}$. In this reduced expression, the lowest Coxeter generator (s_i with the smallest i) appears exactly $(n - 1)$ times. Second, there is a geometric bijection between these $(n - 1)$ faces and the $(n - 1)$ appearances of the lowest Coxeter generator in the reduced expression $\mathtt{w}_{0,k}$. Indeed, for a face f at a Type 1 column, the intersection of ∂f with a Type 1 column has two connected components, which go along two neighboring horizontal lines, say the jth and the $(j + 1)$st. Since each horizontal line is the deformation retract of a sheet in the weave front Σ, a natural choice of the local weave line representative will be the intersection of the two corresponding sheets in Σ, which in turn corresponds to the jth appearance of the lowest Coxeter generator. Figure 39 illustrates a cross-section of the weave front for $n = 4$. Figure 40 illustrates all the possible cases for $n = 3, 4$.
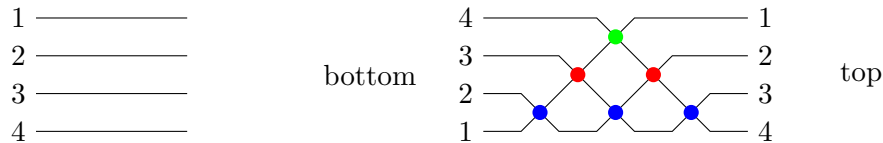
FIGURE 39. (Left) An elementary column with four horizontal lines. (Right) The corresponding cross-section for its associated weave surface $\Sigma(\mathbb{G})$.

3.5.2. *Local representatives of naive absolute cycles in a Type 2 column.* Consider a Type 2 column with n horizontal lines and a single vertical edge between the jth and $(j + 1)$st horizontal lines. First, for the faces bounded by any other pair of consecutive horizontal lines, say kth and $(k + 1)$st with $k \neq j$, the associated naive absolute cycle in the weave is
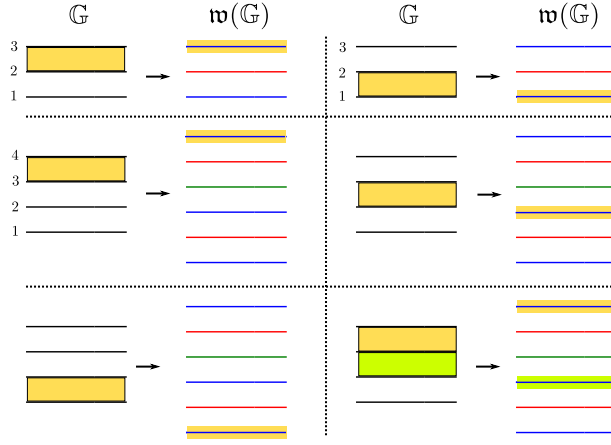
FIGURE 40. Associating a (piece of a) cycle in the weave for Type 1 columns. The first row depicts the two cases for $n = 3$ strands, and the second and third rows depict the three cases for $n = 4$, and an example with a union, in the lower right corner. In all cases, the face $F \subset \mathbb{G}$ is highlighted in yellow, and the associated s_1-edge in the weave $\mathfrak{w}(\mathbb{G})$ is also highlighted in the same color. In the last case of the union, in the lower right, one of the faces and its cycle are highligthed in green.

the *unique* long l-cycle connecting the corresponding weave cycles on the two adjacent Type 1 columns. In other words, one starts at the kth appearance (counting from below) of the lowest Coxeter generator on the left (see Subsection 3.5.1) and follows that weave line straight through any hexagonal vertices. By the construction of the weave $\mathfrak{w}$ in Subsection 3.3.2, this process will go through the weave until it reaches its right hand side at the kth appearance of the lowest Coxeter generator. Figure 41 depicts examples of such faces in purple. Note that, as depicted on the right of the second row in that figure, the l-cycle might go through hexagonal vertices but shall always has the kth lowest Coxeter generator in $\mathfrak{w}_{0,n}$ at the two ends.

Second, for the two faces that involve with the unique vertical edge, the associated absolute cycle in $\mathfrak{w}$ is the unique l-cycle that starts with the jth appearance of the lowest Coxeter generator at its boundary end (see Subsection 3.5.1) and has the other end at the unique trivalent vertex of $\mathfrak{w}$. Figure 41 depicts examples of such faces in yellow and green. Observe that, in general, these l-cycles will also go through hexagonal vertices but always have the jth lowest Coxeter generator at its boundary end.

45
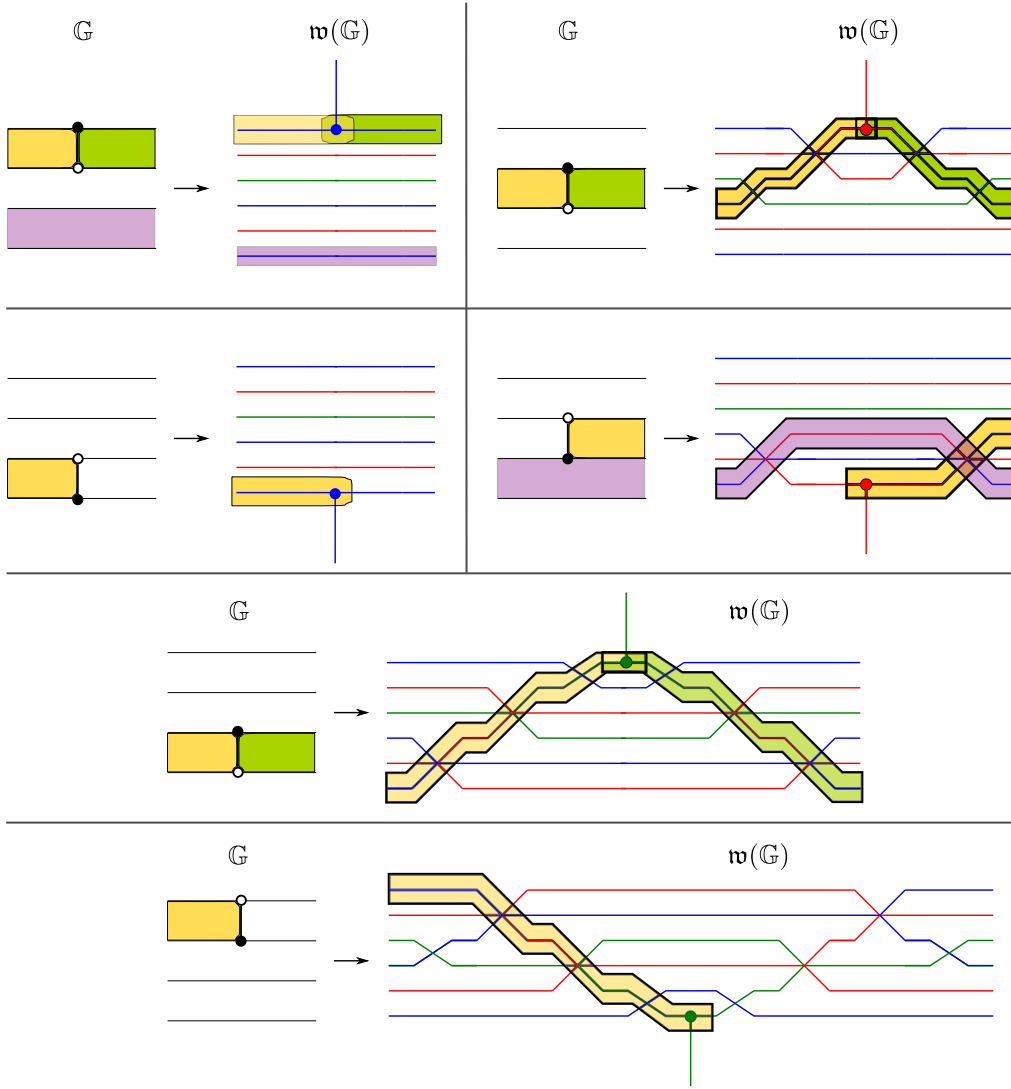
FIGURE 41. Several examples of faces in Type 2 elementary columns with $n = 4$ $\mathbb{G}$-strands, and their associated I-cycle in the weaves $\mathfrak{w}(\mathbb{G})$.

3.5.3. *Local representatives of naive absolute cycles in a Type 3 column.* In a Type 3 column, the majority of faces are similar to a face in a Type 1 column. In the weave, their boundaries are represented by weave lines going across the weave as the unique long I-cycles with the correct boundary conditions. The only exceptional face in a Type 3 column is the face f which contains a lollipop, which we will discuss in detail in this subsection.

Let us first consider the case of a white lollipop attaching to the jth horizontal line on the right. For simplicity let us assume that the horizontal lines on the right are indexed by $1, 2, \ldots, n$ starting from the bottom. Following Subsection 3.5.1, the leftmost and rightmost ends of the cycle γ_f are determined by the Type 1 rules. Namely, given that the face f restricts to one gap on the left and two gaps, on the right, the ends of the cycle γ_f in $\mathfrak{w}$ must be the unique I-cycle associated to those gaps. Thus, the cycle γ_f will start at a blue s_1 edge of the weave on the left and finish at two blue s_1 edges on the right. Now, in general, there does not exist a I-cycle (nor a Y-cycle) with these boundary conditions in $\mathfrak{w}$. This requires introducing a *bident*, as follows. Consider the middle slice of $\mathfrak{w}$ where all the newly emerged weave line have become horizontal, i.e., the right boundary of the weave building

block $\mathfrak{i}_j^w$ (Definition 3.19). Reading the weave lines from bottom to top at this slice yields an expression for the half-twist $w_{0,n} \in S_n$ (note that it is not $\mathbf{w}_{0,n}$). Let us draw the weave slice as the positive braid $s^{-1}_{[n-1,i+1]}\mathbf{w}_{0,n-1}s_{n-1,n-i}$ and mark the s_1 edge in $\mathbf{w}_{0,n-1}$ that corresponds to the gap on the left within which the white lollipop emerges. Then, starting at this marked s_1 edge, we go along the upper-left and upper-right strands until we reach the highest (and last) possible crossing in each of the strands. These two crossings correspond to two weave lines on the right boundary of $\mathfrak{i}_j^w$. These two weave lines are said to be obtained from the (left) s_1 edge by a *bifurcation*.

Type 3 Column

Cross-section of weave front Σ

Weave associated to this column
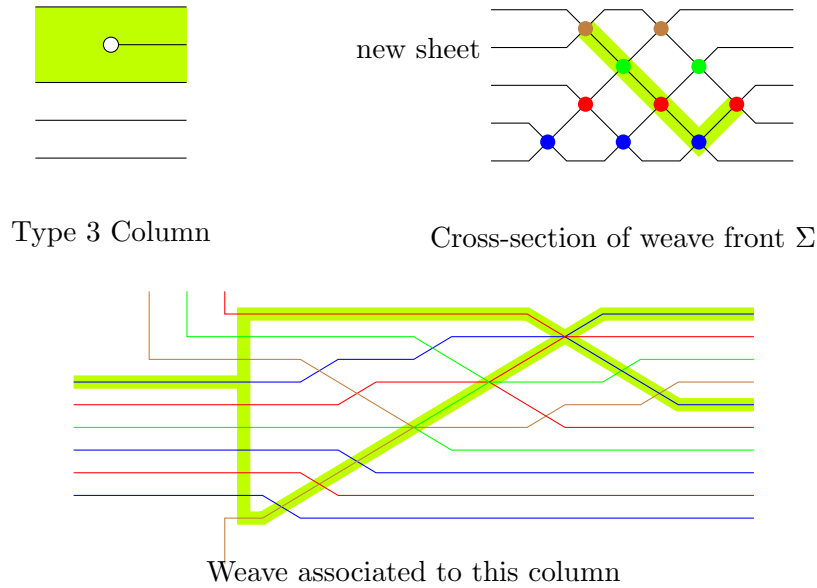
FIGURE 42. The top left is an elementary Type 3 column with a white lollipop. Its associated weave $\mathfrak{w}$ is depicted on the second row. At the top right, a vertical slice of the weave front highlighting the two directions of bifurcation that come up from the s_1-crossing (in blue) at the bottom.

Definition 3.31. A *bident* is an PL-embedding of a T-shape domain into the plane containing the weave such that on the left it coincides with an s_1-edge and on the right it coincides with the two crossings obtained by bifurcation on this s_1-edge.
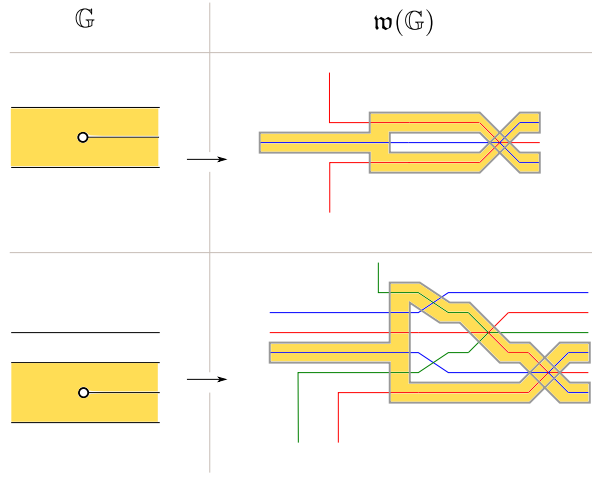
FIGURE 43. (Left) Examples of faces in Type 3 elementary columns with $n = 3, 4$ $\mathbb{G}$-strands. (Right) The associated naive cycles, with the bidents, in the corresponding weaves $\mathfrak{w}(\mathbb{G})$.

Remark 3.32 (Geometric explanation of bidents)**.** In line with the techniques used in [CZ22, Section 2.4], such a bident represents a relative cycle in $L = L(\mathfrak{w})$ for the Type 3 column weave $\mathfrak{w}$. This relative cycle can be described by picturing its intersection with vertical slices of the weave front $\Sigma = \Sigma(\mathfrak{w})$. Figure 44 depicts such slices in two cases, one with $n = 3$ and the other with $n = 4$. Indeed, all that occurs as the vertical slices sweeps from left to right is that a pair of points is created in a strand of the braid (in the vertical slice), and then they move along the strands of the braid until they reach the desired position near the two crossings determined by the bifurcation. Meanwhile, the initial two points that originated from the left s_1 edge also move near those two crossings. Once the points have all moved near the two bifurcation crossings, this cycle can continued to be represented by two l-cycles in the weave. This description allows for the computation of the homology class of this relative cycle and thus the computation of its intersections with other weave cycles.
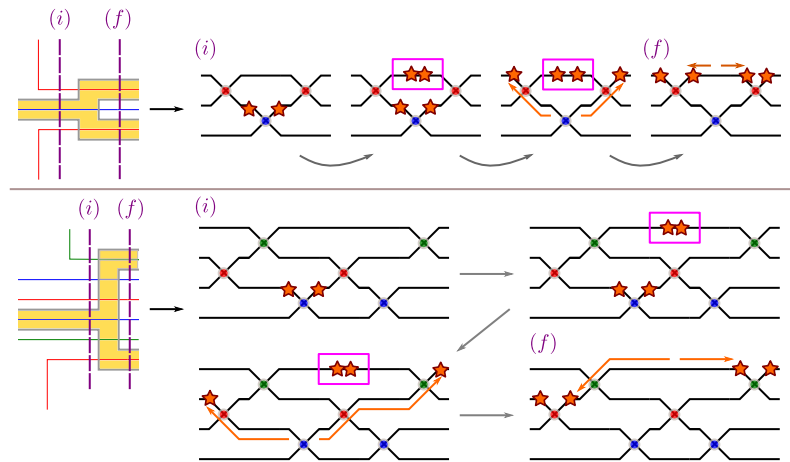
FIGURE 44. The dissection of the bident in two cases. On the left, two weaves $\mathfrak{w}$ are depicted (associated with Type 3 columns), with a piece of a cycle drawn in yellow, and purple dashed vertical lines indicate two locations, the initial (i) and the final (f), where we draw slices. On the right, a movie of slices from (i) to (f) where the cycle in yellow intersects the slices in the star-shaped orange marks. The pink boxes on the right indicate which pair of star-shaped orange marks have been created in the bident process.

48

Finally, the case of an elementary column of Type 3 with a black lollipop is treated in exactly the same manner as for a white lollipop, with the roles vertically reversed. For a face f containing a black lollipop in a Type 3 column, the boundary conditions being l-cycles on the weave and having a unique bident determine the cycle $\gamma_f \subset \mathbb{R}^2$ in the same manner as in the white lollipop case, except now the bident is left-pointing.

A Type 3 column $\mathbb{G}$ Associated weave $\mathfrak{w}(\mathbb{G})$
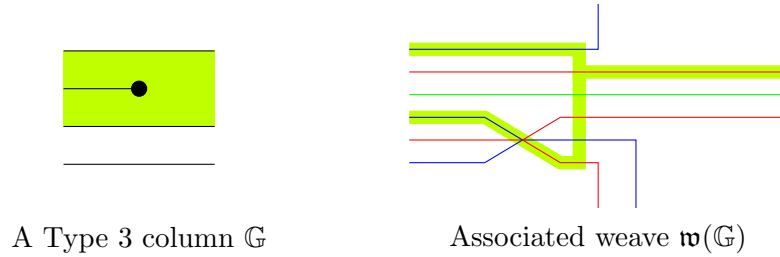
FIGURE 45. A case for the naive absolute cycle in a Type 3 column with a black lollipop. The face $f \in \mathbb{G}$ and its associated cycle $\gamma_f \subset \mathbb{R}^2$ are both highlighted in light green.

Example 3.33. *Figure 46 (upper left) depicts the weave $\mathfrak{w}(\mathbb{G})$ associated to the GP-graph $\mathbb{G}$ given by Figure 16 (upper left), with its two naive absolute cycles γ_1, γ_2. Geometrically, this weave $\mathfrak{w}(\mathbb{G})$ is an embedded exact Lagrangian once-punctured torus bounding one of the two max-tb Legendrian representatives of the $m(5_2)$ twist knot. Note that, since both faces of $\mathbb{G}$ contain lollipops, both cycles γ_1, γ_2 are drawn with bidents. The upper right and the second row of Figure 46 depict the weaves and their naive absolute cycles associated to two plabic fences, which correspond to rainbow closures of braid words $\beta = \sigma_1^5 \in Br_2^+$ and $\beta = (\sigma_1 \sigma_2)^4 \in Br_3^+$, respectively. These plabic fences were drawn in Figure 14. Note that the naive absolute cycles associated to (the weave of) a plabic fence are all long l-cycles.*
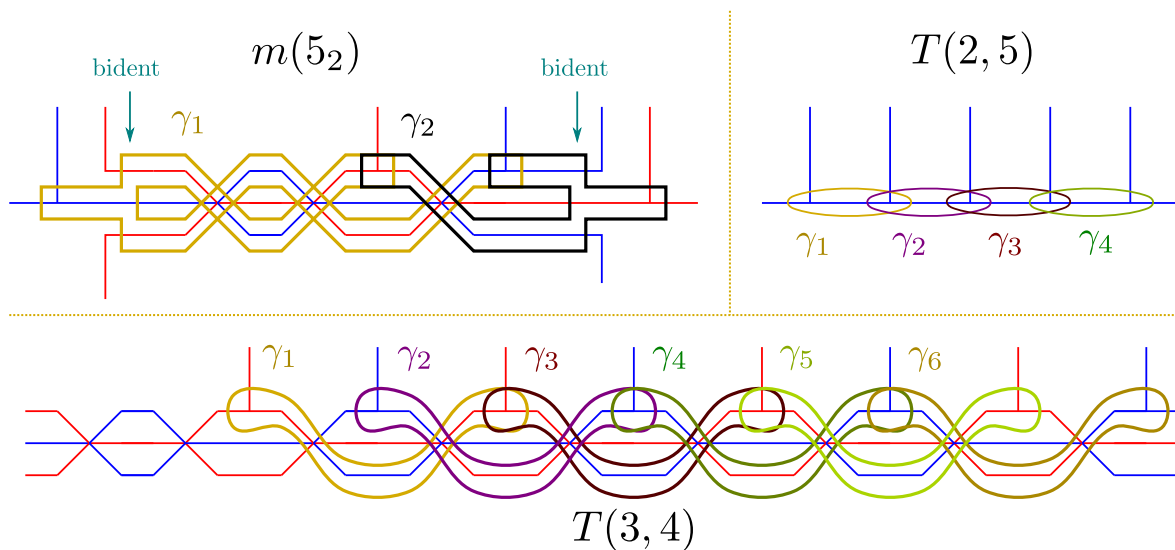
FIGURE 46. (Upper left) The initial weave for the GP-graph in Figure 16 (upper left) and its naive absolute cycles. (Upper right) The initial weave and naive absolute cycles for the plabic fence for the case $n = 5$ for Figure 14 (left). (Second Row) The initial weave and its naive absolute cycles for the plabic fence for the GP-graph drawn in Figure 14 (right).

Due to the possible existence of bidents, it is hard to tell whether a naive absolute cycle has self intersections (and hence it is not an embedded absolute cycle or $\mathbb{L}$-compressible) or not. It is easier if we can represent the naive absolute cycles are Y-cycles (Definition 3.4). Thus, we prove the following:

Proposition 3.34. *Let $\mathbb{G} \subset \mathbb{R}^2$ be a GP-graph and $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$ its associated weave. Then, there exist a weave $\mathfrak{w}'$ and an equivalence $\mathfrak{w}' \sim \mathfrak{w}$ such that, under the isotopy[11] between $\Sigma(\mathfrak{w})$ and $\Sigma(\mathfrak{w}')$, the image of each naive absolute cycle on $L(\mathfrak{w})$ is homologous to a Y-cycle on $L(\mathfrak{w}')$.*

Proof. For Type 1 and Type 2 elementary columns, the associated (pieces of) naive absolute cycles are already I-cycles, and hence Y-cycles. In particular, for $\mathbb{G}$ with no (internal) lollipops, we can take $\mathfrak{w}' = \mathfrak{w}$. It thus suffices to study the case of a Type 3 column, where a bident appears: it suffices to show that there exists a weave equivalence that allows us to replace a bident by a Y-cycle.
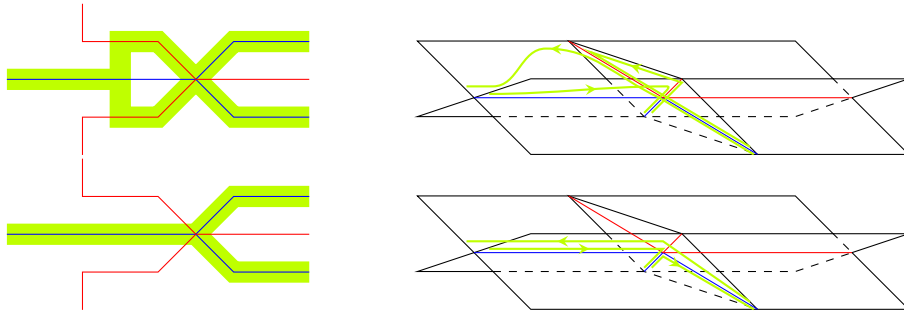
FIGURE 47. Two homologous cycles γ_f and $\widetilde{\gamma}_f$ depicted on the left. The upper left cycle γ_f contains a bident and is not a Y-cycle, the lower left cycle $\widetilde{\gamma}_f$ is Y-cycle. The right hand side of each row depicts these cycles in their spatial Legendrian fronts.

For a Type 3 column with a face f containing a white lollipop, this is done as follows. Consider the two weave lines to the right of the bident where the cycle γ_f propagates. By construction, these two weave lines intersect (to the right) at a unique hexagonal weave vertex of $\mathfrak{w}(\mathbb{G})$. In addition, the horizontal weave line entering from the left at this hexagonal vertex connects with an I-cycle to the s_1-edge on the left of the bident where γ_f starts. Therefore, we can consider the Y-cycle $\widetilde{\gamma}_f$ which starts with this s_1-edge at the left, propagates to the left (as an I-cycle) until the hexagonal vertex and then contains a unique Y-vertex at the hexagonal vertex. Figure 47 depicts both cycles γ_f, at the left of the first row, and $\widetilde{\gamma}_f$, at the left of the second row, in the case of the Type 3 elementary column draw in Figure 43 (upper left). By considering the description of γ_f via vertical slices – as in Figure 44 – it is readily seen that γ_f is homologous to $\widetilde{\gamma}_f$; alternatively, at the right of each row, Figure 44 depcits the cycles in the spatial front and one can visualize the homology between γ_f and $\widetilde{\gamma}_f$ there as well. In consequence, in the case of a white lollipop we can consider the same weave $\mathfrak{w}' = \mathfrak{w}$ and have the naive absolute cycle γ_f with a bident be homologous to the Y-cycle $\widetilde{\gamma}_f$.

For a Type 3 column $\mathbb{G}$ with a black lollipop, the situation is similar, with the exception that a hexagonal vertex might not exist in $\mathfrak{w}(\mathbb{G})$ and thus the bident cannot readily be substituted by a Y-cycle. Nevertheless, we can insert two consecutive hexagonal vertices with a candy twist – Move I in Figure 22 – and then apply the same argument as above. $\square$

Example 3.35. *Let us continue with the example from Figure 46 (upper left). In that figure, neither γ_1 or γ_2 are Y-cycles. By implementing the argument in the proof of Proposition 3.34,*

---

[11]This isotopy naturally induces to a Hamiltonian isotopy between $L(\mathfrak{w})$ and $L(\mathfrak{w}')$.

we obtain a Y*-cycle* $\widetilde{\gamma}_1$ *homologous to* γ_1, *as depicted in Figure 48 (upper left) in the same weave. Similarly, after performing a Move I to the initial weave* $\mathfrak{w}(\mathbb{G})$ *(for the first GP-graph* $\mathbb{G}$ *in Figure 16), we are able to draw a* Y*-cycle* $\widetilde{\gamma}_2$ *homologous to* γ_2, *as shown in Figure 48 (upper right). The second row of Figure 48 draws an absolute cycle* γ_3 *that is homologous to the sum* $\gamma_1 + \gamma_2 = \widetilde{\gamma}_1 + \widetilde{\gamma}_2$.
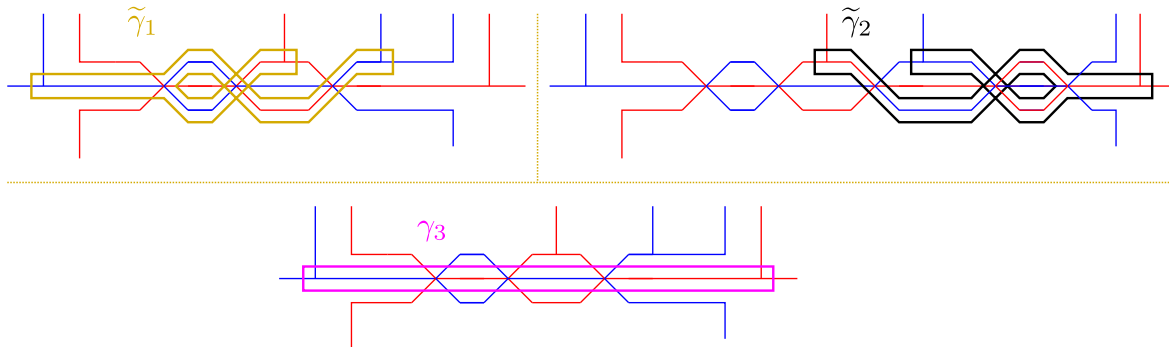
FIGURE 48. (First row) The Y-cycles $\widetilde{\gamma}_1$, on the left, and $\widetilde{\gamma}_2$, on the right, which are homologous to γ_1, and γ_2, following (the proof of) Proposition 3.34. Note that the weave for $\widetilde{\gamma}_2$ has a Move I (candy twist) inserted so as to modify the bident in γ_2 to the Y-cycle $\widetilde{\gamma}_2$. (Second row) An additional absolute I-cycle γ_3 in this weave.

3.6. **Naive Relative Cycles in** $L(\mathfrak{w}(\mathbb{G}))$. Let $L = L(\mathfrak{w}(\mathbb{G}))$ be the initial filling of the GP-link $\Lambda = \Lambda(\mathbb{G})$. Subsection 3.5 constructed an explicit set of generators for a basis of $H_1(L)$ in terms of Y-cycles. Nevertheless, in order to construct cluster $\mathcal{A}$-variables, we also need access to the lattice given by the relative homology group $H_1(L, \Lambda) = H_1(L, \partial L)$. Recall that, by Poincaré duality, there exists a non-degenerate pairing between the absolute homology group $H_1(L)$ and the relative homology group $H_1(L, \Lambda)$:

$$\langle \cdot, \cdot \rangle : H_1(L) \otimes H_1(L, \Lambda) \longrightarrow \mathbb{Z}.$$

Let $\{\gamma_f\}$ be the basis of naive absolute cycles constructed in Subsection 3.5, where the index f runs over all faces of $\mathbb{G}$. Consider the Poincaré dual basis $\{\eta_f\}$ on $H_1(L, \Lambda)$.[12] In order to perform computations in the moduli stack of sheaves, we also want to describe the relative cycles in $\{\eta_f\}$ combinatorially in terms of the weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$. This is done according to the following discussion.

In general, given an N-weave $\mathfrak{w} \subset \mathbb{R}^2$, we can consider an (unoriented) curve $\kappa \subset \mathbb{R}^2$ that ends at unbounded regions in the complement of $\mathfrak{w} \subset \mathbb{R}^2$ and intersect weave lines of $\mathfrak{w}$ transversely and generically – in particular, away from the weave vertices. There are N natural ways to lift κ to the weave front $\Sigma(\mathfrak{w})$, which in turn correspond to N unoriented curves on L. In consequence, any subset of these N lifts, together with any orientation we choose for each of element of such a subset, defines a relative homology cycle $\eta \in H_1(L, \Lambda)$. Figure 49 (left) depicts two possible oriented lifts of the (dashed) yellow curve κ drawn to its right.

---

[12]Note that the dual of a given absolute cycle γ_f is not constructed merely from γ_f but rather from the entire basis of naive absolute cycles.
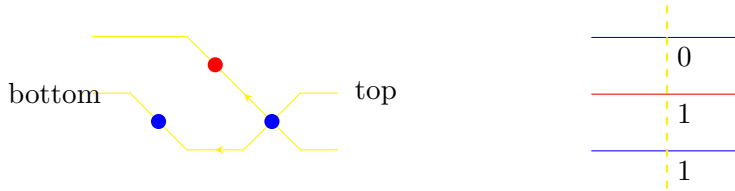
FIGURE 49. (Left) In yellow, two lifts of the κ dashed curve depicted on the right. (Right) A 3-weave with a labeled dashed curve κ: the labels $0, 1, 1$ indicate the intersection number with each of the (pieces of cycles associated to) the weave lines.

Back to the case with $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$, where we have the basis of naive absolute cycles available, we can compute the intersection numbers between such relative cycles η and the naive absolute cycles. This leads to a tuple of integers $I(\eta) := (\langle \eta, \gamma_f \rangle)_{\text{faces } f}$. By the non-degeneracy of the Poincaré pairing $\langle \cdot, \cdot \rangle$ and the fact that $\{\gamma_f\}_{\text{faces } f}$ is a basis, the tuple of intersection numbers $I(\eta)$ uniquely determines the relative homology class of η. In fact, given that all naive absolute cycles that η intersects non-trivially must pass through weave lines, in order to describe the relative homology class of η it suffices to draw the unoriented curve $\kappa \subset \mathbb{R}^2$ and record the collection of the intersection number of its lift η with each of the weave lines. In order to distinguish such curves from weave lines, we will use dashed lines to depict such a curve κ.

Definition 3.36. A dashed curve $\kappa \subset \mathbb{R}^2$ as above, with the data of intersection numbers for each weave line it crosses, is called a *labeled dashed curve*.

Figure 49 (right) depicts a labeled dashed curve κ. From a diagrammatic perspective, it is desirable to be manipulate labeled dashed curves in a weave diagram in the same manner that [CZ22] explained how to combinatorially manipulate absolute cycles. For that, we have depicted in Figure 50 the key moves on labeled dashed curves: these are all equivalences, in that these moves do not change the relative homology classes that the labeled dashed curves represent.
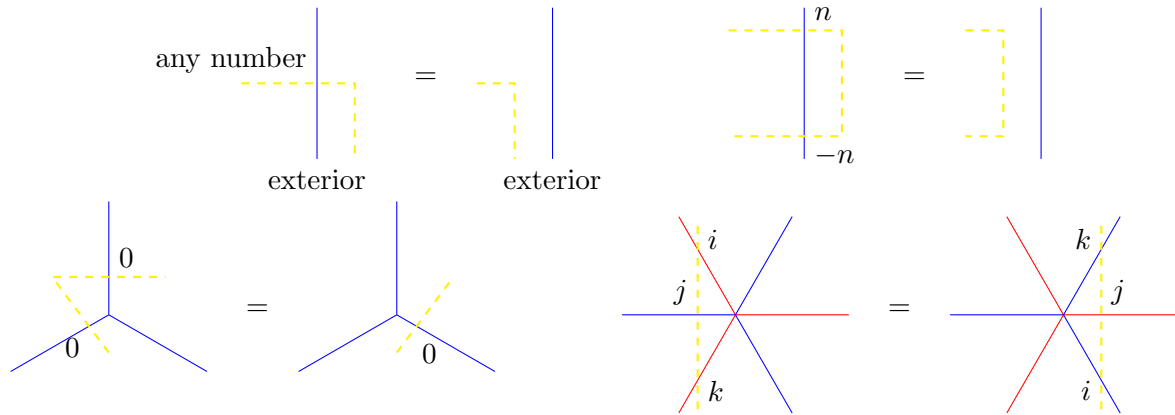
FIGURE 50. Four equivalence moves for labeled dashed curves.

Finally, we can now diagrammatically describe a collection of labeled dashed curves that is a basis of the relative homology group $H_1(\Sigma, \Lambda)$, dual to the naive basis of $H_1(\Sigma)$ built in Subsection 3.5, as follows. First, for each face $f \subset \mathbb{G}$, we select a Type 1 elementary column inside of f. Second, consider the piece of the weave $\mathfrak{w}(\mathbb{G})$ associated to this column – weave lines arranged according to $\mathfrak{w}_{0,n}$ – and draw a vertical dashed curve κ_f transverse to

52

this piece of the weave. It now suffices to specify the correct labels encoding the intersection between the cycle for κ and the naive absolute cycle associated to this face $f \in \mathbb{G}$. For that, consider the braid given by slicing the weave front Σ along κ. Recall that there is a natural bijection between the appearances of the lowest Coxeter generator s_i in $\mathtt{w}_{0,n}$ and the gaps in this Type 1 column (see Subsection 3.5.1). Locate the appearance of s_i that corresponds to a gap belonging to f and consider set of crossings in this braid which are contained within the rhomboid diamond whose unique lowest vertex is at this appearance of s_i. Figure 51 (center) draws an example of such a diamond for the face $f \subset \mathbb{G}$ depicted to its left. Then we label the curve κ, to a labeled curve κ_f, by assigning the intersection number 1 for all the weave lines in $\mathtt{w}(\mathbb{G})$ which are associated to crossings in the braid *inside* of the diamond, and by assigning the intersection number 0 for all the remaining weave lines. Figure 51 (right) depicts the corresponding curve κ_f with its intersection labels for the face $f \subset \mathbb{G}$.
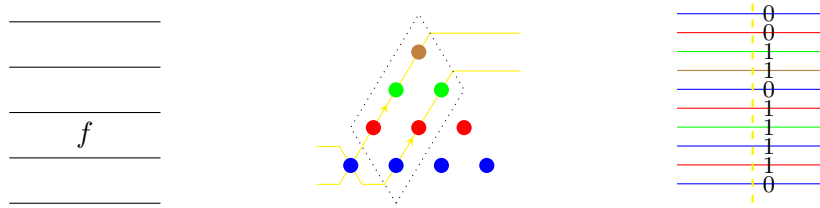
FIGURE 51. (Left) A face $f \in \mathbb{G}$ chosen at an elementary Type 1 column. (Center) A slice of the weave $\mathtt{w}(\mathbb{G})$ with the rhomboid diamond associated to the s_1-crossing, in blue, for the face $f \in \mathbb{G}$. (Right) The labeled dashed curve κ_f, in yellow, in the weave $\mathtt{w}(\mathbb{G})$, with 1 in the weave lines for the crossings inside the diamond, and 0 otherwise.

By construction, the intersection pairing between the labeled dashed curve κ_f and the naive absolute cycles is given by $\langle \eta_f, \gamma_f \rangle = 1$ and $\langle \eta_f, \gamma_g \rangle = 0$ for $g \neq f$. Thus, the collection of relative cycles $\{\eta_f\}$, associated to these particular labeled dashed curves κ_f, are representatives of a dual naive basis of $H_1(L, \Lambda)$. We call this collection of relative cycles the *naive basis of relative cycles* of the relative homology group $H_1(L, \Lambda)$.

3.7. **Initial Absolute Cycles and Initial Relative Cycles in** $L(\mathtt{w}(\mathbb{G}))$. Subsections 3.5 and 3.6 above explain the construction of the naive basis of absolute cycles and the corresponding naive basis of relative cycles. The generators of these basis are not geometrically appropriate: despite being Y-cycles in the weave (or dual to them), they are often represented by *immersed* cycles and it is a priori unclear whether it is possible to mutate at them.[13] A key idea in this manuscript is the consideration and study of *sugar-free* hulls, as introduced in Subsection 2.3. In this section, these two parts, sugar-free hulls and the study of homology cycles compatible with the weave $\mathtt{w}(\mathbb{G})$, converge together: we show that it is possible to associate a Y-tree absolute cycle on $\mathtt{w}(\mathbb{G})$ to every sugar-free hull of $\mathbb{G}$. In consequence, given that Y-trees are *embedded*, it will be possible to perform a weave mutation at every sugar-free hull of $\mathbb{G}$. This leads to the notion of *initial basis*, which eventually give rise to the initial seeds for our cluster structures.

In the study of sugar-free hulls, we must consider cycles which are associated to regions of $\mathbb{G}$ – namely, the sugar-free hulls – and not just faces $f \subset \mathbb{G}$. The simplest case is that of a region in an elementary column of Type 1, which is considered in the following simple lemma, where we use the weave $\mathtt{n}(\mathtt{w}_{0,n})$ introduced in Definition 3.13).

---

[13]In any sense of the word mutation: geometrically, through a Lagrangian surgery, diagrammatically, via a weave mutation, or cluster-theoretically, mutating at the naive vertex representing them in the naive quiver.

Lemma 3.37. *Let $\mathbb{G}$ be a GP-graph and $C \subset \mathbb{G}$ a Type 1 elementary column. Consider the region $R \subset C$ given by the union of k consecutive gaps in C. Then the boundary of ∂R is homologous to the lift of a unique weave line on the kth level.*

Proof. Indeed, the boundary of a single gap has two connected components, and each of them is a deformation retract of a sheet of the spatial wavefront $\Sigma(C)$. For a single gap, the two sheets associated with its boundary intersect at a unique weave line at the bottom level. Now for a union R of k consecutive gaps, the two sheets associated with ∂R intersect at the kth level and the conclusion follows. $\qquad\square$

Figure 52 depicts four cases illustrating how to associate a cycle on a weave line for a region on a Type 1 column. The case of arbitrary strands can be readily imagined by examining these few cases.
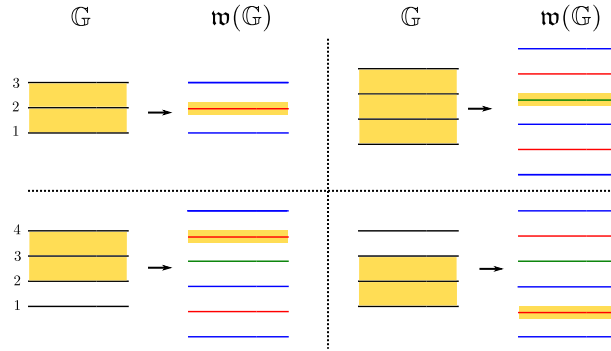
FIGURE 52. Associating a (piece of a) cycle in the weave for regions on a Type 1 column. The first row depicts the case where the region is the entire column, for $n = 3$ and 4 strands. The second row depicts the remaining two cases for $n = 4$ strands.

Proposition 3.34 showed that it is possible, up to possibly performing a weave equivalence, to represent the naive absolute cycles with Y-cycles. Nevertheless, these are typically immersed, e.g. see both $\widetilde{\gamma}_1$ and $\widetilde{\gamma}_2$ in Figure 48. It is not possible, in general, to find embedded Y-cycles representing these homology classes. Now, the following result, which we refer to as the Y-representability Lemma, shows that it *is* possible to represent the boundary cycle ∂R by a Y-tree if the region R is sugar-free.

Lemma 3.38 (Y-Representability Lemma)**.** *Let $\mathbb{G}$ be a GP-graph and $R \subset \mathbb{G}$ a sugar-free region. Then the boundary cycle ∂R is homologous to a Y-tree on $\mathfrak{w}(\mathbb{G})$, up possibly performing a weave equivalence that adds $\mathfrak{n}_k^\uparrow(\mathfrak{w}_{n,0})\mathfrak{n}_k^\uparrow(\mathfrak{w}_{n,0})^{op}$ and $\mathfrak{n}_k^\downarrow(\mathfrak{w}_{n,0})\mathfrak{n}_k^\downarrow(\mathfrak{w}_{n,0})^{op}$ to $\mathfrak{w}(\mathbb{G})$.*

Proof. Let C be an elementary column of $\mathbb{G}$. By Lemma 2.11, the intersection $R \cap C$ has at most one connected component. Therefore, if R intersects C non-trivially, R must be a union of consecutive gaps in the column C. By Lemma 3.37, for a Type 1 or a Type 3 elementary column C, we can represent the boundary $\partial(R \cap C)$ by a single I-cycle weave line going from left to right; note that the weave line color may change within a Type 3 column. It thus remains to treat the cases of elementary columns of Type 2.

By Lemma 2.10 applied to a Type 2 column, we conclude that the following four cases – in correspondence with the four staircase patterns – are to be analyzed.
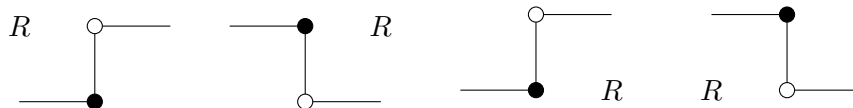
First, let us consider the staircase pattern which is second from the left. The corresponding local weave pattern is $\mathfrak{c}_k^\uparrow(\mathtt{w}_{0,n})$, as introduced in Definition 3.15. In the construction of this weave pattern, we first bring the kth strand in the bottom level upward, using the weave pattern $\mathfrak{n}_k^\uparrow(\mathtt{w}_{0,n})$, subsequently insert a trivalent weave vertex at the top strand, and then insert the weave pattern $\mathfrak{n}_k^\uparrow(\mathtt{w}_{0,n})^{op}$. Figure 53 depicts a case with 4 horizontal lines and $k = 2$. (See also Figures 29 and 31.) Since the kth horizontal line is the bottom boundary of ∂R, there must be a unique hexavalent weave vertex in the $\mathfrak{n}_k^\uparrow(\mathtt{w}_{0,n})^{op}$ that connects to the weave line corresponding to the union of all gaps in R at the right boundary. Therefore, we can create a Y-tree in this region, with the required boundary conditions, by inserting a tripod leg that connects to the trivalent weave vertex (at the top) and continues to the left towards whichever weave line is required by the boundary condition of R. This resulting Y-tree, in the shape of a tripod, then represents ∂R locally, as desired. Figure 53 (right) depicts this Y-tree, highlighted in light green, for the region R on the left, also drawn in the same color. This concludes the second case among the four staircase patterns; the third case, which also contains the region R to the right of the crossing, can be resolved analogously.
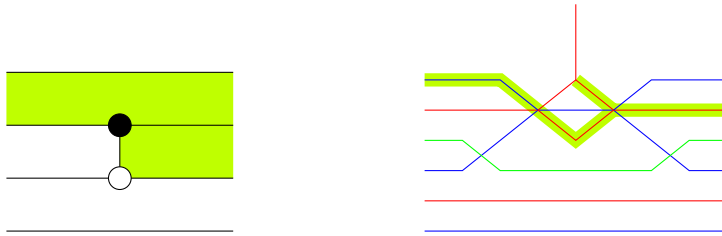
FIGURE 53. (Left) Type 2 column with region R highlighted in light green. (Right) The associated local weave and the Y-tree, in the shape of a tripod.

Next let us study the first (leftmost) of the four staircase patterns. In this case, the chosen weave pattern $\mathfrak{c}_k^\downarrow(\mathtt{w}_{0,n})$, as assigned in Subsection 3.3, does not already have a hexavalent weave vertex that meets our need. Nevertheless, we can create it by concatenating the local weave pieces $\mathfrak{n}_k^\uparrow(\mathtt{w}_{0,n})$ and $\mathfrak{n}_k^\uparrow(\mathtt{w}_{0,n})^{op}$ on the left first. Note that this can be achieved by inserting a series of Moves I and V, and thus the equivalence class of the weave remains the same. Now, inside of the weave piece $\mathfrak{n}_k^\uparrow(\mathtt{w}_{0,n})^{op}$, there exists a unique hexavalent weave vertex that connects to both the weave line representative of ∂R, to the left, and the trivalent weave vertex in $\mathfrak{c}_k^\downarrow(\mathtt{w}_{0,n})$, to the right. The Y-tree, again in a tripod shape, represents ∂R locally, as desired. Figure 54 depicts an example of such tripod with $n = 4$ and $k = 2$. An analogous argument also resolves the case of the fourth (rightmost) staircase pattern.
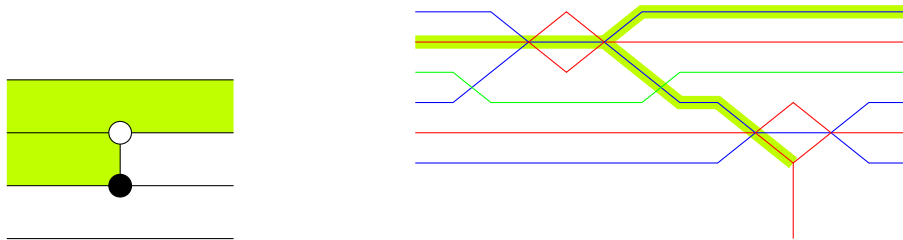
FIGURE 54. Example of the first staircase.

Finally, by combining the local pictures for all three types of columns, we conclude that there exists a representative for ∂R which is a Y-tree. □

The Y-representability Lemma allows us to introduce the following definition.

Definition 3.39. Let $\mathbb{G}$ be a GP-graph. The set of *initial absolute cycles* $\mathfrak{S}(\mathbb{G})$ is the set of all Y-trees on $L(\mathfrak{w}(\mathbb{G}))$ which are associated to the (non-empty) sugar-free hulls in $\mathbb{G}$.

The inclusion relation between sugar-free hulls naturally puts a partial order on the set $\mathfrak{S}(\mathbb{G})$: by definition, $\partial R \leq \partial R'$ if $R \subset R'$ as sugar-free hulls. Also, given a face $f \subset \mathbb{G}$, we note that a face $g \in \mathbb{S}_f$ in its sugar-free hall, must satisfy $\mathbb{S}_g \subset \mathbb{S}_f$ and hence $\partial \mathbb{S}_g \leq \partial \mathbb{S}_f$. Finally, note that multiple faces in $\mathbb{G}$ can share the same sugar-free hull, and there may exist faces with an empty sugar-free hull. Thus, in general, the number of sugar-free hulls may be smaller than the number of faces in $\mathbb{G}$. Nevertheless, we can prove that the set of initial absolute cycles $\mathfrak{S}(\mathbb{G})$ is always linearly independent.

Proposition 3.40. *Let $\mathbb{G}$ be a GP-graph and let $\mathfrak{S}(\mathbb{G})$ be its set of initial absolute cycles. Then $\mathfrak{S}(\mathbb{G})$ is a linearly independent subset of $H_1(L(\mathfrak{w}(\mathbb{G})))$. In addition, it is possible to add naive absolute cycles to $\mathfrak{S}(\mathbb{G})$ to complete $\mathfrak{S}(\mathbb{G})$ into a basis of $H_1(L(\mathfrak{w}(\mathbb{G})))$.*

Proof. Let $L = L(\mathfrak{w}(\mathbb{G}))$ be the initial filling. Consider the naive basis $\{\gamma_f\}_{\text{faces } f}$ of $H_1(L)$. It suffices to show that we can replace $|\mathfrak{S}(\mathbb{G})|$ many naive basis elements with elements in $\mathfrak{S}(\mathbb{G})$ while maintain the spanning property of the set. This is can be done by using the partial order on $\mathfrak{S}(\mathbb{G})$ as follows.

Let us start with the minimal elements in $\mathfrak{S}(\mathbb{G})$ and work our way up, replacing the appropriate naive basis elements in $\{\gamma_f\}$ with elements in $\mathfrak{S}(\mathbb{G})$. At each turn, we select a face $f \subset \mathbb{G}$ whose sugar-free hull $\mathbb{S}_f$ defines a Y-tree $\partial \mathbb{S}_f \in \mathfrak{S}(\mathbb{G})$ and then replace the naive absolute cycle γ_f with the Y-tree $\partial \mathbb{S}_f$. Note that it is always possible to find such a face $f \subset \mathbb{G}$ in this process: if no such face f were available, any faces within the sugar-free hull would not have this particular sugar-free region as their sugar-free hull, which is tautologically absurd.

In the process of implementing each of these replacements, we must argue that the resulting set is still spanning the homology group $H_1(L)$. For that, we observe that if a replacement of γ_g by $\partial \mathbb{S}_g$ is done before the replacement of γ_f by $\mathbb{S}_f$, the sugar-free hull $\mathbb{S}_g$ must not contain the face f inside. This follows from the fact that $f \in \mathbb{S}_g$ implies $\mathbb{S}_f \subset \mathbb{S}_g$. Therefore, we can use the set $\{\partial \mathbb{S}_g \mid g \in \mathbb{S}_f\}$, which is already in the basis set due to the partial order, and $\partial \mathbb{S}_f$ to recover $\gamma_f = \partial f$. This shows that the resulting set after the replacement of γ_f by $\partial \mathbb{S}_f$ as above, is still a basis for $H_1(L)$. $\square$

Example 3.42 below partially illustrates this replacement procedure in a particular Hasse diagram for the initial absolute cycles $\mathfrak{S}(\mathbb{G})$. The choice of completion of $\mathfrak{S}(\mathbb{G})$ to a basis in $H_1(L(\mathfrak{w}(\mathbb{G}))$ is a neat instance of the natural appearance of quasi-cluster structures in (symplectic) geometry. There is no particular canonical manner by which we can typically choose a basis for this complement but, as we shall explain, the different choices all lead to the same cluster structure *up to* monomials in the frozen variables, i.e. a quasi-cluster structure, see Appendix I.

The dualization process, via the Poincaré pairing, requires a choice of basis. Therefore, there is no canonical choice of *initial relative cycles* associated to the set $\mathfrak{S}(\mathbb{G})$ unless the later spans $H_1(L(\mathfrak{w}(\mathbb{G}))$. In a general situation, we can at least consider the following concept.

Definition 3.41. Let $\mathbb{G}$ be a GP-graph, $L = L(\mathfrak{w}(\mathbb{G}))$, $\Lambda = \Lambda(\mathbb{G})$ and $\mathfrak{S}(\mathbb{G})$ its set of initial absolute cycles on L. Let $\mathfrak{B}$ be a basis of $H_1(L)$ which is obtained by adding naive absolute cycles to the set $\mathfrak{S}(\mathbb{G})$. Then the *initial relative cycles* associated to $\mathfrak{B}$ is the collection $\mathfrak{B}^\vee$ of linear combinations of naive relative cycles whose relative homology classes form a basis of $H_1(L, \Lambda)$ dual to the basis $\mathfrak{B}$ of $H_1(L)$.

We emphasize that the basis $\mathfrak{B}^\vee$ depends not only on $\mathfrak{S}(\mathbb{G})$, but also on the chosen basis completion $\mathfrak{B}$. Note that a few simple choices of $\mathfrak{B}$ are available as a result of the replacement construction in the proof of Proposition 3.40. Namely, we start with the naive absolute basis

$\{\gamma_f\}$, and then swaps some of the basis elements γ_f with their corresponding initial absolute cycles $\partial\mathbb{S}_f$. In case that there are multiple faces sharing the same sugar-free hull, only one of the naive absolute cycles gets replaced, and the rest remain in the basis, which will become frozen basis elements.

Furthermore, if the basis $\mathfrak{B}$ is chosen via such a basis replacement process, then the corresponding dual relative cycle basis $\mathfrak{B}^\vee$ can be described with respect to the partial order on sugar-free hulls as well. Indeed, suppose that we have an equality $\mathbb{S}_f = \mathbb{S}_g = \cdots$ of sugar-free hulls, for some faces $f, g \subset \mathbb{G}$, among others, and γ_f was chosen to be replaced by $\partial\mathbb{S}_f$ in the replacement process. Then, the naive relative cycle η_g needs to be replaced by $\eta_g - \eta_f$ for each g with $\mathbb{S}_g = \mathbb{S}_f$. Similarly, η_f would need to be replaced by $\eta_f + N$, where the N summand is a linear combination of the naive relative cycles η_h associated to the chosen faces h with $\mathbb{S}_f \subsetneq \mathbb{S}_h$, such that the pairing of $(\eta_f + N)$ with $\partial\mathbb{S}_h$ vanishes for all such h.

Example 3.42. *Let us suppose that a set of initial absolute cycles $\mathfrak{S}(\mathbb{G})$ conforms to the Hasse diagram below – i.e. this is the inclusion partial order on sugar-free hulls for some GP-graph $\mathbb{G}$ with 16 faces $\{f_1, \ldots, f_{16}\}$, where we have denoted $\mathbb{S}_i := \mathbb{S}_{f_i}$, $i \in [1, 16]$.*
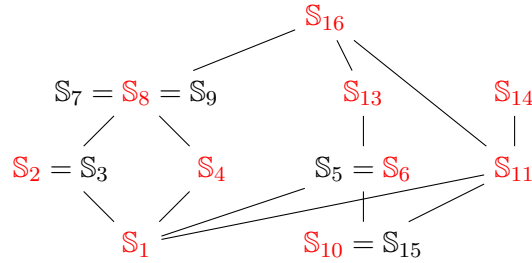
FIGURE 55. The Hasse diagram in Example 3.42, where the sets are the sugar-free hulls and the partial order relation is given by inclusion.

We have marked in red a possible choice of naive absolute cycles for replacement, e.g. the cycles $\gamma_1, \gamma_2, \gamma_4, \gamma_6, \gamma_8\gamma_{10}, \gamma_{11}, \gamma_{13}, \gamma_{14}, \gamma_{16}$ will be modified accordingly, whereas $\gamma_3, \gamma_5, \gamma_7, \gamma_9\gamma_{15}$ will remain part of the basis. Note that, as drawn, there is an additional face f_{12} whose sugar-free hull is empty, and thus γ_{12} also remains in the basis. Given this choice, the replacement construction applied to the basis of naive absolute cycles $\{\gamma_i\}_{i=1}^{16}$, leads to the following basis $\mathfrak{B}$ for $H_1(L)$:

$$\{\gamma_1, \ \gamma_1 + \gamma_2 + \gamma_3, \ \gamma_3, \ \gamma_1 + \gamma_4, \ \gamma_5, \ \gamma_1 + \gamma_5 + \gamma_6 + \gamma_{10} + \gamma_{15}, \ \gamma_7, \ \gamma_1 + \gamma_2 + \gamma_3 + \gamma_4 + \gamma_7 + \gamma_8 + \gamma_9,$$

$$\gamma_9, \ \gamma_{10} + \gamma_{15}, \ \gamma_1 + \gamma_{10} + \gamma_{11} + \gamma_{15}, \ \gamma_{12}, \ \gamma_1 + \gamma_5 + \gamma_6 + \gamma_{10} + \gamma_{13} + \gamma_{15}, \ \gamma_1 + \gamma_{10} + \gamma_{11} + \gamma_{14} + \gamma_{15},$$

$$\gamma_{15}, \ \gamma_1 + \gamma_2 + \gamma_3 + \gamma_4 + \gamma_5 + \gamma_6 + \gamma_7 + \gamma_8 + \gamma_9 + \gamma_{10} + \gamma_{11} + \gamma_{13} + \gamma_{15} + \gamma_{16}\},$$

We have highlighted in red the naive absolute cycles in $\{\gamma_i\}$ that correspond to the red sugar-free hulls in the Hasse diagram. For completeness, we have also highlighted in blue the basis elements that will become frozen variables. Similar, the basis $\mathfrak{B}^\vee$ of $H_1(L, \Lambda)$ dual to $\mathfrak{B}$ is

$$\{\eta_1 - \eta_2 - \eta_4 - \eta_6 + \eta_8 - \eta_{11} + 2\eta_{16}, \ \eta_2 - \eta_8, \ \eta_3 - \eta_2, \ \eta_4 - \eta_8, \ \eta_5 - \eta_6, \ \eta_6 - \eta_{13}, \ \eta_7 - \eta_8,$$

$$\eta_8 - \eta_{16}, \ \eta_9 - \eta_8, \ \eta_{10} - \eta_6 - \eta_{11} + \eta_{16}, \ \eta_{11} - \eta_{14} - \eta_{16}, \ \eta_{12}, \ \eta_{13} - \eta_{16}, \ \eta_{14}, \ \eta_{15} - \eta_{10}, \ \eta_{16}\}.$$
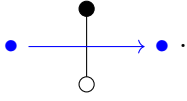
This illustrates that there is a rather simple algorithmic procedure to complete $\mathfrak{S}(\mathbb{G})$ to a basis $\mathfrak{B}$ of $H_1(L)$ and to find the dual basis $\mathfrak{B}^\vee$ of $H_1(L, \Lambda)$.

3.8. **The naive quiver of a GP-graph.** Let us start by emphasizing that the quiver of the initial seed for the cluster structure we construct is *not* always the dual quiver of the GP-graph $\mathbb{G}$. Nevertheless, that dual quiver is useful in order to construct the actual quiver of the initial seed because it can be used to compute the intersection form on the initial filling, which is needed to define the initial quiver. Let us provide the details.

Following the description of M. Gross, P. Hacking, S. Keel, and M. Kontsevich [GHKK18], in order to define a cluster structure, we first fix an integer lattice with a skew-symmetric form on it. In the context of a GP-graph $\mathbb{G}$, the integer lattice is $H_1(L(\mathfrak{w}(\mathbb{G})))$, and a natural skew-symmetric form on it is given by the intersection pairing between absolute homology classes. By using the *naive* basis $\{\gamma_f\}$ of naive absolute cycles and the GP-graph $\mathbb{G}$, we can describe the intersection pairing form $\{\cdot,\cdot\}$ combinatorially using a quiver.

Definition 3.43. Let $\mathbb{G} \subset \mathbb{R}^2$ be a GP-graph. The *naive quiver* $Q_0(\mathbb{G})$, or *dual quiver*, associated to $\mathbb{G}$ is the quiver constructed as follows:

    (1) A quiver vertex is associated to each face $f \subset \mathbb{G}$.

    (2) For every bipartite edge in $\mathbb{G}$, we draw an arrow according to .

    (3) For each pair of quiver vertices, sum up the arrows between them.

Note that in the third step (3) there might be cancellations.

The following lemma is proved by A. Goncharov and R. Kenyon:

Lemma 3.44 ([GK13, Definition 8.2 and Proposition 8.3])**.** *Let ϵ_{fg} be the exchange matrix of the quiver $Q_0(\mathbb{G})$. Then the intersection pairing between γ_f and γ_g is given by $\{\gamma_f, \gamma_g\} = \epsilon_{fg}$.*

Since $\{\gamma_f\}_{\text{faces } f}$ is a basis of $H_1(L(\mathfrak{w}(\mathbb{G})))$, Lemma 3.44 uniquely determines the intersection skew-symmetric form. This intersection form, i.e. the quiver $Q_0(\mathbb{G})$, is then used to compute the correct quiver $Q(\mathbb{G})$ for the initial seed. The unfrozen vertices of the correct initial quiver $Q(\mathbb{G})$ will be indexed by $\mathfrak{S}(\mathbb{G})$, the set of initial absolute cycles, and the remaining frozen vertices are determined by the choice of completion of $\mathfrak{S}(\mathbb{G})$ to a basis of $H_1(L)$. Since we will elaborate more on this in Section 4, we conclude this discussion for now and revisit $Q(\mathbb{G})$ then.

Remark 3.45. In the case of a plabic fence $\mathbb{G}$, all naive absolute cycles are I-cycles and thus they also are initial absolute cycles. Thus, for a plabic fence, $Q(\mathbb{G})$ coincides with $Q_0(\mathbb{G})$.

3.9. **A few examples.** For the purpose of clarity, let us explicitly apply the constructions in this section to three examples of GP-graphs, drawing the associated fronts, weaves and associated bases. In the first and second examples below show, not all naive absolute cycles are (homologous to) initial absolute cycles.[14] In these two examples, the set of initial absolute cycles will form a basis of $H_1(L(\mathfrak{w}(\mathbb{G})))$. The third example shows that the set of initial absolute cycles might not span $H_1(L(\mathfrak{w}(\mathbb{G})))$ (even if the set of naive cycles does), and that there are different natural choices for extending the set of initial absolute cycles to a basis.

Example 1. Let us consider the GP-graph $\mathbb{G}_1$ in Figure 56 (left). The faces f_2, f_3, f_4, f_5 and f_6 are their own sugar-free hulls, whereas the sugar-free hull of f_1 is $f_1 \cup f_2$. The associated Legendrian front is depicted to in Figure 56 (right). This is a front for the unique max-tb Legendrian $(3, 4)$-torus knot, which determines the Legendrian isotopy type of $\Lambda(\mathbb{G}_1)$. Note that the Legendrian knot $\Lambda(\mathbb{G}_1)$ is isotopic to the Legendrian knot associated to the

---

[14]Only in particular cases, such as for plabic fences, all naive absolute cycles are initial absolute cycles. These two sets do not coincide for general examples of GP-graphs $\mathbb{G}$.
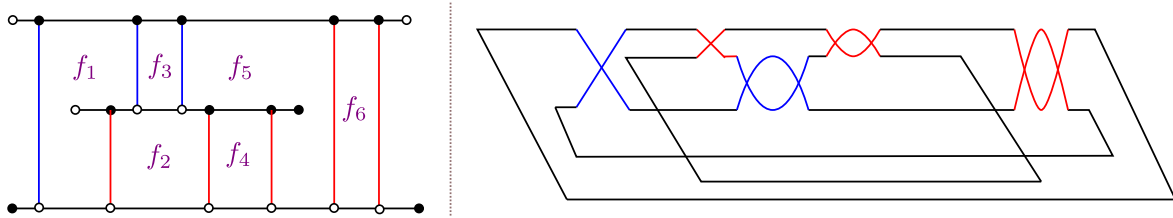
FIGURE 56. The GP-graph $\mathbb{G}_1$, on the left, and its associated Legendrian front, on the right. The Legendrian knot $\Lambda(\mathbb{G}_1)$ is a max-tb Legendrian representative of the $(3,4)$-torus knot.

alternating strand diagram of the plabic fence in Figure 15 (right), but that the GP-graph from Figure 56 (left) is itself not a plabic fence.
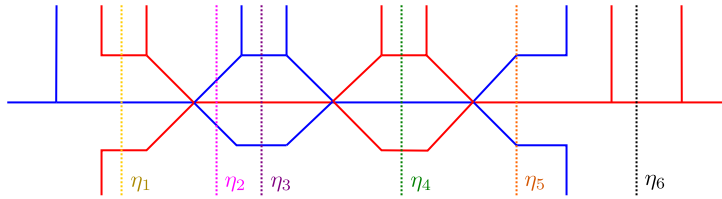
FIGURE 57. The Legendrian weave $\mathfrak{w}(\mathbb{G}_1)$ associated to the GP-graph $\mathbb{G}_1$. Supports of relative cycles η_j, $j \in [1,6]$, are also drawn.

The weave $\mathfrak{w}(\mathbb{G}_1)$ associated to this GP-graph $\mathbb{G}_1$ is drawn in Figure 57. It represents a genus 3 embedded exact Lagrangian filling $L = L(\mathfrak{w}(\mathbb{G}_1))$ for $\Lambda = \Lambda(\mathbb{G}_1)$. A basis of $H_1(L)$ is given by the absolute cycles $\{\gamma_1 + \gamma_2\} \cup \{\gamma_i\}_{i \in [2,6]}$ depicted in Figure 58, where γ_i is the naive cycle associated to f_i, $i \in [1,6]$. In this example, each naive absolute cycles $\gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6$ is also an initial absolute cycle. The naive absolute cycle γ_1 is immersed, and its associated initial absolute cycle is $\gamma_1 + \gamma_2$, as drawn in Figure 58; this is coherent with the fact that the sugar-hull of f_1 is $f_1 \cup f_2$, as noted above.
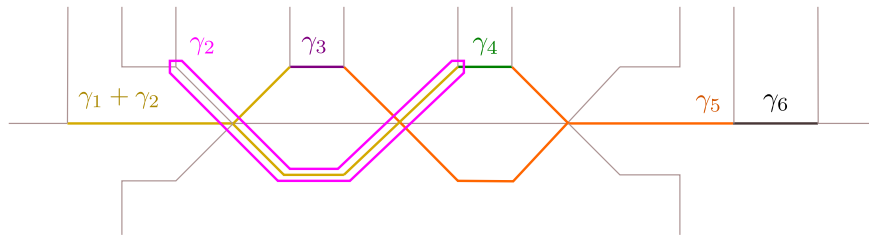
FIGURE 58. A basis of absolute cycles $\{\gamma_i\}_{i \in [1,6]}$ for the lattice $H_1(L(\mathfrak{w}(\mathbb{G}_1)))$.

A basis of relative cycles η_j, $j \in [1,6]$, through its support, is also drawn in Figure 57. The intersection matrix between these two basis is

$$\langle \widetilde{\gamma}_i, \eta_j \rangle = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \quad \widetilde{\gamma}_1 := \gamma_1 + \gamma_2, \quad \widetilde{\gamma}_i := \gamma_i, i \in [2,6],$$

59

and thus the dual basis of $\{\gamma_1 + \gamma_2\} \cup \{\gamma_i\}_{i \in [2,6]}$ in terms of $\{\eta_j\}_{j \in [1,6]}$ is

$$\{\eta_1, \eta_2 - \eta_1, \eta_3, \eta_4, \eta_5, \eta_6\}$$

and the associated quiver $Q(\mathbb{G}_1)$ is depicted in Figure 59. All vertices are in fact mutable, due to the fact that they are represented by initial cycles, as we show in Section 4. It is readily verified that the mutation class of this quiver is that of the simple type E_6, though it is *not* the standard E_6-Dynkin diagram (which could have obtained from a plabic fence).
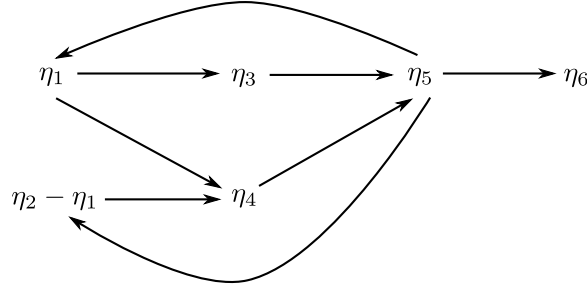
FIGURE 59. The quiver associated to the intersection form and basis for $\mathbb{G}_1$.

Example 2. Consider the GP-graph $\mathbb{G}_2$ in Figure 60 (left). We note that the sugar-free hull of f_1 is $f_1 \cup f_3 \cup f_4 \cup f_5$, that of f_5 is $f_3 \cup f_5$, and f_2, f_3, f_4, f_6 are their own sugar-free hulls. The corresponding front is drawn in 60 (right). The associated Legendrian knot $\Lambda(\mathbb{G}_2)$ is a max-tb representative of the $(2,7)$-torus knot, and we thus would expect the quiver to be in the mutation class of the A_6-Dynkin diagram (with any orientation), with all variables being mutable.
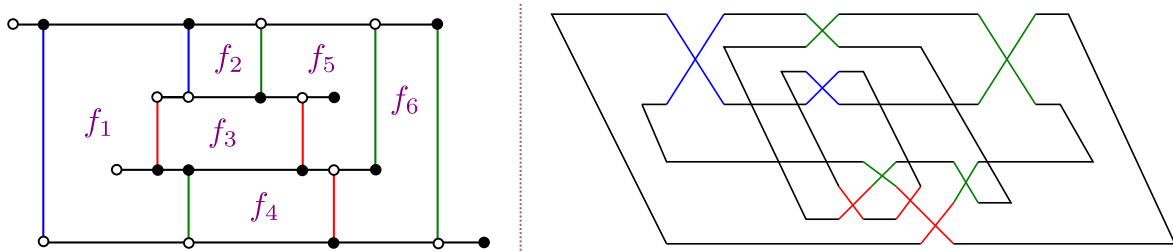
FIGURE 60. A GP-graph $\mathbb{G}_2$ (left) and its associated Legendrian front (right).

The first step is to draw the weave $\mathfrak{w}(\mathbb{G}_2)$, which is done in Figure 61. This weave defines a Lagrangian filling $L = L(\mathfrak{w}(\mathbb{G}_2))$ of $\Lambda(\mathbb{G}_2)$, in this case a once-punctured genus 3 surface, and the second step is finding the initial basis of $H_1(L) \cong \mathbb{Z}^6$. Figure 62 draws the set of naive absolute cycles $\{\gamma_i\}_{i \in [1,6]}$, which indeed span this lattice.

With respect to the naive absolute basis, the intersection pairing has matrix

$$\langle \gamma_i, \gamma_j \rangle = \begin{pmatrix} 0 & 1 & -1 & 1 & 0 & 0 \\ -1 & 0 & 1 & 0 & -1 & 0 \\ 1 & -1 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 & 1 & -1 \\ 0 & 1 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}$$
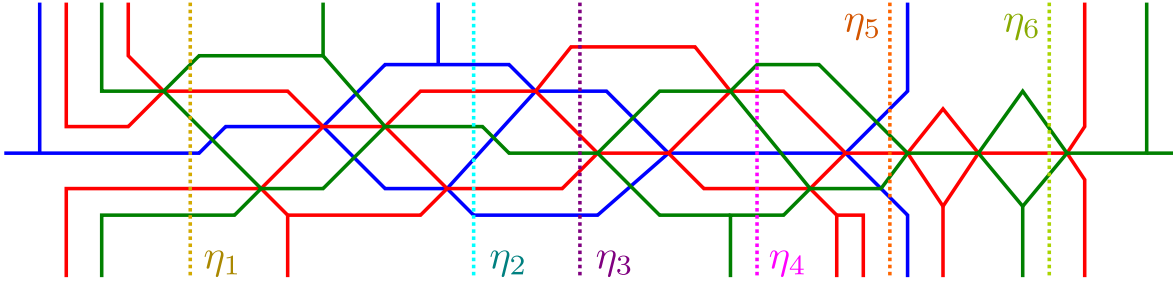
FIGURE 61. The weave $\mathfrak{w}(\mathbb{G}_2)$ associated to the GP-graph of Figure 60, which represents a genus 3 embedded exact Lagrangian filling of $\Lambda(\mathbb{G}_2)$.

In this example, the naive cycles $\gamma_2, \gamma_3, \gamma_4$ are initial absolute cycles, as they are short I-cycles. Nevertheless, the naive cycles $\gamma_1, \gamma_5, \gamma_6$ are not initial absolute cycles. In order to obtain an initial absolute cycle from γ_1, resp. γ_5, we can add $\gamma_3 + \gamma_4 + \gamma_5$ to it, resp. γ_3. Again, this is coherent with the sugar-free hulls of f_1 and f_5. These two initial cycles $\gamma_1 + \gamma_3 + \gamma_4 + \gamma_5$ and $\gamma_3 + \gamma_5$ are depicted in the weave $\mathfrak{w}(\mathbb{G}_2)$ drawn in Figure 63. The cycle γ_6 is in fact homologous to a Y-cycle (using the hexagonal vertex exactly at the left of its bident), and thus we can take that representative, also denoted γ_6, as an initial absolute cycle.
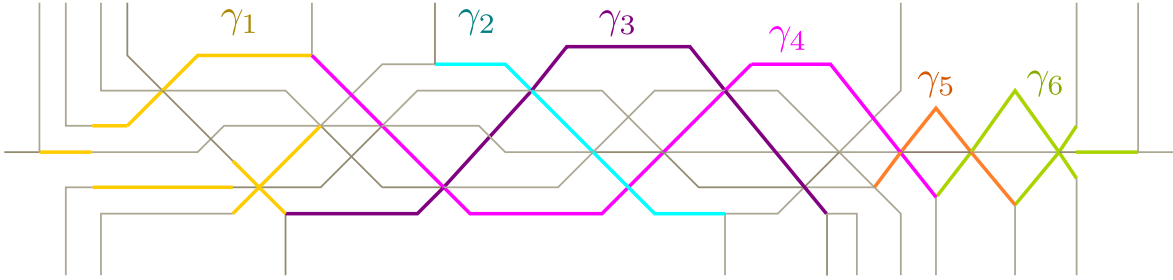
FIGURE 62. The set of naive absolute cycles $\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6$ for $\mathfrak{w}(\mathbb{G}_2)$.
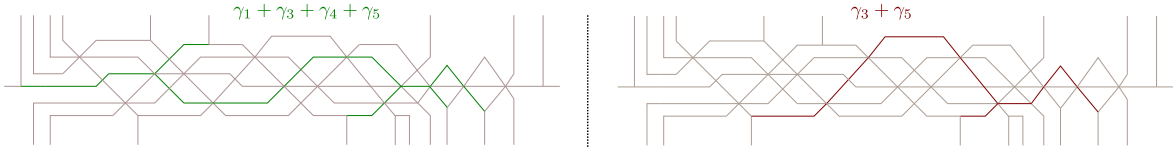
FIGURE 63. The two initial absolute cycles $\gamma_1 + \gamma_3 + \gamma_4 + \gamma_5$ and $\gamma_3 + \gamma_5$ in the weave $\mathfrak{w}(\mathbb{G}_2)$. Note that both homology classes are represented by Y-cycles.

In summary, the set of initial absolute cycles $\mathfrak{S}(\mathbb{G}_2)$ is

$$\mathfrak{S}(\mathbb{G}_2) = \{\gamma_1 + \gamma_3 + \gamma_4 + \gamma_5, \gamma_2, \gamma_3, \gamma_4, \gamma_3 + \gamma_5, \gamma_6\},$$

which is a basis of $H_1(L(\mathfrak{w}(\mathbb{G}_2)))$. Its dual basis, in terms of the relative cycles associated to the dashed curves $\{\eta_i\}_{i \in [1,6]}$ in Figure 57, then reads

$$\langle \eta_1, \eta_2, \eta_3 - \eta_5, \eta_4 - \eta_1, \eta_5 - \eta_1, \eta_6 \rangle = H_1(L, \Lambda),$$

as $\langle \gamma_i, \eta_j \rangle = \delta_{ij}$ for $i, j \in [1, 6]$. The quiver $Q(\mathbb{G}_2)$ is drawn in Figure 64, and all vertices will be mutable (because they are initial cycles). This quiver is indeed in the mutation equivalence of the simple A_6-type.

61
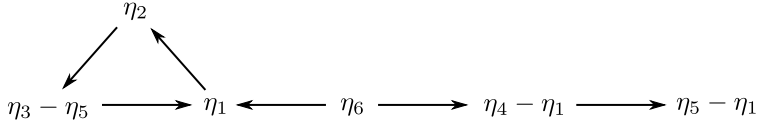
FIGURE 64. The quiver $Q(\mathbb{G}_2)$ associated to the GP-graph $\mathbb{G}_2$.

Example 3. Consider the GP-graph $\mathbb{G}_3$ in Figure 37 (right) (equivalently the upper left of Figure 16), a front for which is drawn in Figure 17 (left). Denote the two faces of $\mathbb{G}_3$ by f_1, f_2, and note that $f_1 \cup f_2$ is the sugar-free hull of both of them. The associated initial weave $\mathfrak{w}(\mathbb{G}_3)$ is depicted in Figure 37 (right), where the Lagrangian filling $L = L(\mathfrak{w}(\mathbb{G}_3))$ is a once-punctured 2-torus. There are two naive cycles γ_1 and γ_2, drawn in Figure 48 (top row), respectively associated to the two faces f_1 (left) and f_2 (right). The naive quiver $Q(\mathbb{G})$ is

$$\gamma_1 \longrightarrow \gamma_2,$$

and it defines the intersection pairing on the lattice $H_1(L) \cong \mathbb{Z}\gamma_1 \oplus \mathbb{Z}\gamma_2$. However, this is not the quiver needed to define the cluster structure associated with the Legendrian knot $\Lambda(\mathbb{G}_3)$. In order to obtain the correct quiver, we note that there is a unique (non-empty) sugar-free hull $f_1 \cup f_2$ available. In consequence, there is only one initial cycle $\delta := \gamma_1 + \gamma_2$, which does not span $H_1(L)$. Now, to complete the set $\{\delta\}$ into a basis of this lattice, we can throw in either γ_1 or γ_2 (as a frozen vertex). In each such ways, we can produce a quiver using the intersection pairing form. The resulting two possible quivers are

$$\delta \longrightarrow \gamma_2, \qquad\qquad \delta \longleftarrow \gamma_1.$$

As shown in Section 4, either of these two quiver will define a cluster structure for this max-tb Legendrian representative $\Lambda(\mathbb{G}_3)$ of the smooth knot $m(5_2)$, and the resulting cluster structures are quasi-equivalent. We remark that the existence of freedom in completing the set of initial cycles to a basis, is what leads naturally to a quasi-cluster structure, rather than a cluster structure.

Remark 3.46. As we will see in the next section, in general we need to decorate GP-links $\Lambda = \Lambda(\mathbb{G})$ with a set T of marked points subject to the requirement of at least one marked point per link component, and the general lattice we need is $H_1(L, T)$, where $L = L(\mathfrak{w}(\mathbb{G}))$ is the initial filling. However, since all three examples given in this subsection are knots, we can take T to be a singleton set, and hence it is enough to just consider $H_1(L, T) = H_1(L)$.

3.10. **Bases and homology lattices in the presence of marked points.** The construction of the cluster structures in Theorem 1.1, and the definition of the moduli space $\mathfrak{M}(\Lambda, T)$, in general require an additional piece of data: a set T of marked points on $\Lambda(\mathbb{G})$.

Definition 3.47. Let $\mathbb{G}$ be a GP-graph and let $\Lambda = \Lambda(\mathbb{G})$ be its GP-link. A set of *marked points* $T \subset \Lambda$ is a subset of distinct points in Λ, where we require that there is at least one marked point on each link component of Λ and, without loss of generality, the set T is disjoint from all crossings and all cusps in the front $\mathfrak{f}(\mathbb{G})$ of Λ.

All prior statements in Section 3 remain unchanged by the addition of marked points, as they do not affect the associated weaves or the Hamiltonian isotopy class of exact Lagrangian fillings. Therefore, we can still consider the initial embedded exact filling $L = L(\mathfrak{w}(\mathbb{G}))$ of the GP-link Λ. As before, we select the collection of initial absolute cycles $\mathfrak{S}(\mathbb{G})$ associated with sugar-free hulls, and they form a linearly independent subset of $H_1(L)$. The addition of marked points affects only the cluster-theoretic constructions: we need to replace the lattice of absolute homology $H_1(L)$ by the lattice of relative homology $H_1(L, T)$. The natural inclusion $H_1(L) \subset H_1(L, T)$ induced by the inclusions $T \subset \Lambda = \partial L \subset L$, allows us to include

the initial absolute cycles $\mathfrak{S}(\mathbb{G})$ as a linearly independent subset of $H_1(L, T)$. The only difference is that, in order to fix a cluster structure, we must expand $\mathfrak{S}(\mathbb{G})$ further to a basis $\mathfrak{B}$ of $H_1(L, T)$. This expansion can be done in two steps: we first expand $\mathfrak{S}(\mathbb{G})$ to a basis of $H_1(L)$, as done via the replacement process in Subsection 3.7, and then expand this basis of $H_1(L)$ to a basis of $H_1(L, T)$.

As was the case for $H_1(L)$ and its dual $H_1(L, \Lambda)$, we shall need a dual space of $H_1(L, T)$ together with a basis dual to a chosen basis $\mathfrak{B}$ of $H_1(L, T)$. In fact, there is a natural intersection pairing

$$\langle \cdot, \cdot \rangle : H_1(L, T) \otimes H_1(L \setminus T, \Lambda \setminus T) \longrightarrow \mathbb{Z}$$

obtained by algebraically counting geometric intersections of relative cycles in generic position. In the same manner that Poincaré duality was used in Subsection 3.7, a duality also exists in the setting with marked points. We record the precise statement in the following.

Proposition 3.48. *Let L be a connected smooth surface with boundary $\Lambda = \partial L$, and $i : T \longrightarrow \Lambda$ an inclusion of a set of marked points with $\pi_0(i)$ being surjective. Then*

$$\mathrm{rk}(H_1(L, T)) = \mathrm{rk}(H_1(L \setminus T, \Lambda \setminus T)),$$

and the intersection pairing $\langle \cdot, \cdot \rangle$ is non-degenerate.

There are cases in which marked points are not needed, such as the case where Λ is a knot. Nevertheless, marked points are always needed if Λ is a link. Let us illustrate a simple example where marked points will be used, using the the simplest link.

Example 3.49. *Consider the Legendrian representative $\Lambda \subset (\mathbb{R}^3, \xi_{st})$ of the Hopf link with both max-tb components, that admits an embedded exact Lagrangian annulus filling L. Let us consider the set $T = \{t_1, t_2\} \subset \Lambda$ of marked points, with t_1 in one component of Λ and t_2 in the other. Then the relative homology group is $H_1(L, T) \cong \mathbb{Z}^2$, and Figure 65 depicts two blue paths which are a basis of $H_1(L, T)$. Note that exactly one of the generators of the basis naturally lives in $H_1(L) \subset H_1(L, T)$. With respect to the intersection pairing $\langle \cdot, \cdot \rangle$, the blue paths give rise to a dual basis consisting of the red paths. By applying the duality from Proposition 3.48, we obtain a corresponding basis for $H_1(L \setminus T, \Lambda \setminus T)$, which we have depicted in red in the same figure.*
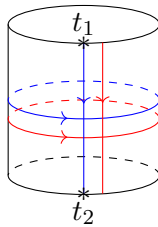
FIGURE 65. The two blue paths are cycles that represent a basis of $H_1(L, T)$. The two red paths are relative cycles in $H_1(L \setminus T, \Lambda \setminus T)$ that represent the dual bases with respect to the intersection pairing.

Remark 3.50. We emphasize that the lattice $H_1(L \setminus T, \Lambda \setminus T)$, or from a cluster theoretical viewpoint, the character lattice for the cluster $\mathcal{A}$-structure, does not carry a natural skew-symmetric form. It is the lattice $H_1(L, T)$ (or more generally, the character lattice for a cluster $\mathcal{X}$-structure) that carries a natural skew-symmetric form, which in turn defines the quiver for the (quasi-)cluster structure. This is why every cluster $\mathcal{X}$-variety is equipped with a natural log-canonical Poisson structure, which further admits a canonical quantization [FG09].

Finally, we conclude this section by noting that it is possible, and in many cases useful, to consider intermediate lattices M and N in between the lattices discussed above. Namely, we can consider sublattices N of $H_1(L,T)$ which include $H_1(L)$, and dually quotients M of $H_1(L \setminus T, \Lambda \setminus T)$:

$$
\begin{array}{ccc}
H_1(L,T) & \longleftrightarrow & H_1(L \setminus T, \Lambda \setminus T) \\
\uparrow & & \downarrow \\
N & \longleftrightarrow & M \\
\uparrow & & \downarrow \\
H_1(L) & \longleftrightarrow & H_1(L, \Lambda)
\end{array}
$$

In the diagram above all horizontal arrows are dual lattices. The intermediate lattices are only interesting if T is large enough; e.g. if Λ is a knot and T consists of a unique marked point, these lattices are all isomorphic. To wit, already in the study of max-tb Legendrian $(2,n)$-torus links $\Lambda \subset (\mathbb{R}^3, \xi_{\mathrm{st}})$, i.e. a plabic fence $\mathbb{G}$ with two $\mathbb{G}$-strands and n vertical edges, it can be desirable to consider the intermediate lattices M, N. Indeed, there are natural cluster structures arising from the study of triangulations of $(n+3)$-gons [FZ03, Section 12.2], with a varying number of frozen vertices. In the literature on cluster algebras, a rather standard choice of frozens is the set of *boundary frozens*, one per boundary edge of the $(n+3)$-gon. This choice is not as intuitive from the perspective of contact topology, where assigning one marked point per component appears to be more natural; thus having a unique frozen vertex if n is even, and none if n is odd. Nevertheless, the addition of such frozen variables often leads to more tractable combinatorics and novel results for cluster variables, e.g. the saturation property of Newton polytopes [Kal14, MSB20]. It can thus be of interest to be able to discuss such cluster structures (with the boundary frozens) from a contact and symplectic topology perspective. With the flexibility in choosing intermediate lattices, this can be achieved. We illustrate this in the following example.
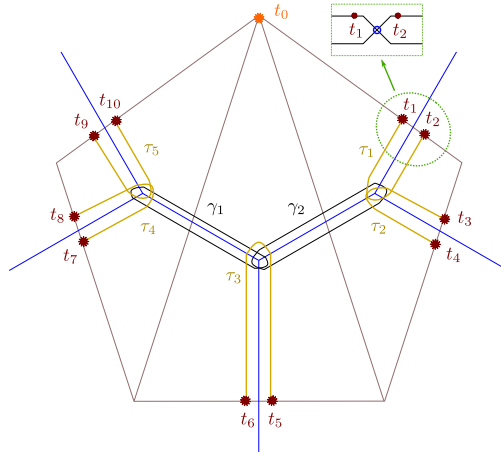
FIGURE 66. Adding marked points T to a front of the trefoil in order to incorporate boundary frozen variables. The weave represents an embedded exact Lagrangian filling of this Legendrian. In this case, the lattices M, N are intermediate lattices.

Example 3.51. *Consider the weave depicted in Figure 66, which represents an embedded exact Lagrangian once-punctured 2-torus filling of the max-tb Legendrian (right-handed) trefoil $\Lambda(3_1)$. It is associated to a plabic fence, and dual to the triangulation drawn in the figure. If we choose the set $T = \{t_0\}$ to be a unique marked point, the $\mathcal{X}$–lattice is*

$H_1(L, T) \cong H_1(L) \cong \mathbb{Z}^2 \langle \gamma_1, \gamma_2 \rangle$, and the $\mathcal{A}$–lattice is generated by their duals. In particular, there are no frozen variables, and the intersection form yields the A_2-quiver with two mutable vertices. In order to incorporate the boundary frozen variables, we enlarge the set of marked point to $T = \{t_0, t_1, t_2, \ldots, t_{10}\}$, adding 10 marked points in the front, as depicted in Figure 66: near each of the five crossings of the front associated to the weave lines we incorporate two marked points, one on the left and one on the right, always in the upper strand. This enlargement is logical: we are incorporating the relative cycles $\tau_1, \ldots, \tau_5$, which now belong to $H_1(L, T)$, and appear naturally in the weave.

The lattice $H_1(L, T)$ has rank 11, and it contains the sublattices $H_1(L) = \mathbb{Z}^2 \langle \gamma_1, \gamma_2 \rangle$ and $N_f := \mathbb{Z}^5 \langle \tau_1, \tau_2, \tau_3, \tau_4, \tau_5 \rangle$. Notice that the marked point t_0 is kept so as to avoid acquiring the relation $\tau_1 + \tau_2 + \tau_3 + \tau_4 + \tau_5 = 0$, which would make N_f of rank 4 instead. At this point we can then consider the sublattice $N := H_1(L) \oplus N_f$ of rank 7, which precisely captures the two mutable variables (dual to the diagonal edges of the triangulation) and the five boundary frozen variables. By construction, the lattice N intermediately (and strictly) lies as $H_1(L) \subset N \subset H_1(L, T)$, and similarly for its dual lattice M.

4. Construction of Quasi-Cluster Structures in Sheaf Moduli

In this section we develop the necessary results to study the geometry of the moduli stack $\mathfrak{M}(\Lambda, T)$ associated to $\Lambda = \Lambda(\mathbb{G})$ and prove Theorem 1.1. In particular, we introduce microlocal merodromies in Subsection 4.6 which, as we will prove, become the cluster $\mathcal{A}$-variables. which. The construction of the cluster structures is obtained purely by symplectic geometric means, using the results for Legendrian weaves from Section 3 above and [CZ22] and the microlocal theory of sheaves [GKS12b, KS90, STZ17]. The proof of its existence also requires a couple of results from algebraic geometry, which we present in detail as well.

Let us review what we have developed in Sections 2 and 3 thus far. Given a GP-graph $\mathbb{G}$, we constructed the following list of objects:

(i) A Legendrian link $\Lambda = \Lambda(\mathbb{G})$, which is a (-1) closure of a positive braid $\beta(\mathbb{G})$.

(ii) An exact Lagrangian filling $L = L(\mathfrak{w})$ of Λ called the *initial filling*. This exact Lagrangian filling L is obtained as the Lagrangian projection of the Legendrian lift associated with the spatial front defined by the *initial weave* $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$.

(iii) A collection of *initial absolute cycles* $\mathfrak{S}(\mathbb{G})$, which form an $\mathbb{L}$-compressing system for L and can be described by Y-trees on $\mathfrak{w}$.

(iv) A skew-symmetric intersection pairing on the lattice $H_1(L)$. This intersection pairing can be computed directly from the GP-graph $\mathbb{G}$.

By specifying an additional generic set of *marked points* $T \subset \Lambda$ with at least one marked point per component, we also obtain the lattice $H_1(L, T)$, which contains $H_1(L)$ and hence the linearly independent subset $\mathfrak{S}(\mathbb{G})$. The skew-symmetric pairing on $H_1(L)$ extends naturally to a skew-symmetric pairing on $H_1(L, T)$. By Poincaré duality, we can identify the dual lattice of $H_1(L, T)$ with the relative homology $H_1(L \setminus T, \Lambda \setminus T)$. Any completion of $\mathfrak{S}(\mathbb{G})$ to a basis $\mathfrak{B}$ of $H_1(L, T)$ gives rise to a unique dual basis $\mathfrak{B}^\vee$ of $H_1(L \setminus T, \Lambda \setminus T)$.

The outline for this section is as follows. First, we give working definitions of the moduli space $\mathfrak{M}(\Lambda, T)$, which allows us to draw connections to Lie-theortical moduli spaces and also deduce the factoriality of $\mathcal{O}(\mathfrak{M}(\Lambda, T))$. Next, on the moduli space $\mathfrak{M}(\Lambda, T)$, we construct a new family of $\mathbb{C}^\times$-valued rational functions called *microlocal merodromies*, which are associated with relative cycles in $H_1(L \setminus T, \Lambda \setminus T)$. Although the definition of microlocal merodromies depends on the initial filling L, we show that for elements in the dual basis $\mathfrak{B}^\vee$, their microlocal merodromies actually extend to $\mathbb{C}$-valued regular functions on the entire

moduli space $\mathfrak{M}(\Lambda, T)$. Moreover, we also prove that within these special microlocal mero-dromies, those dual to $\mathfrak{S}(\mathbb{G})$ can be mutated according to the cluster $\mathcal{A}$-mutation formula as the initial weave $\mathfrak{w}$ undergoes weave mutation, corresponding to a Lagrangian disk surgery on $L(\mathfrak{w})$. Then, we show that the codimension-2 argument in cluster varieties can be applied by studying immersed Lagrangian fillings represented by non-free weaves. Finally, these results together with a theorem of Berenstein-Fomin-Zelevinsky allow us to conclude existence of a cluster $\mathcal{A}$-structure on $\mathfrak{M}(\Lambda, T)$, where the initial and adjacent seeds are constructed via the Lagrangian filling $L(\mathbb{G})$, its Lagrangian surgeries and the associated microlocal merodromies.

4.1. **Descriptions of sheaves with singular support on the Legendrian $\Lambda(\mathbb{G})$.** Let $\Lambda \subset (\mathbb{R}^3 \xi_{\mathrm{st}})$ be a Legendrian link, $T \subset \Lambda$ a set of marked points, and consider the moduli stacks $\mathcal{M}_1(\Lambda)$ and $\mathfrak{M}(\Lambda, T)$ discussed in Section 2.8. These stacks classify (complexes of) constructible sheaves on $\mathbb{R}^2$ with a singular support condition. In this section, we provide Lie-theoretical descriptions for $\mathcal{M}_1(\Lambda)$ and $\mathfrak{M}(\Lambda, T)$ which are suited for our computations, using [KS90] and closely following [STZ17, Section 3.3] and Section 5 in ibid.[15] These are more combinatorial presentations of these stacks, as the constructible and microlocal aspects of the original definition are translated into explicit quiver representations satisfying certain conditions.

First, it suffices to focus on sheaves, rather than complexes of sheaves, as explained in Section 2.8.2, since $\Lambda(\mathbb{G})$ admits a binary Maslov index [STZ17, Section 5.3]. By [KS90, Theorem 8.4.2], the sheaves parametrized by $\mathcal{M}_1(\Lambda)$ and $\mathfrak{M}(\Lambda, T)$ are constructible. A rather combinatorial presentation of (the category of) constructible sheaves is then provided by work of R. MacPherson: for a stratified topological space (M, S), the category of S-constructible sheaves of sets on M is equivalent to the category of set-valued functors on the exit-path category. See [Tre09, Theorem 1.2] and [Cur14, Chapter 11]. In our case, the space being stratified is $M = \mathbb{R}^2$ and the strafication S is provided by a front $\pi_F(\Lambda) \subset \mathbb{R}^2$ of $\Lambda \subset (\mathbb{R}^3 \xi_{\mathrm{st}})$, where $\pi_F : \mathbb{R}^3 \longrightarrow \mathbb{R}^2$ denotes a front (Legendrian) projection. Since we can ensure that this particular fronts are stratifying $\mathbb{R}^2$ as a regular cell-complex, and that the strate are simply-connected, the exit-path category is equivalent to the poset category of (M, S), whose objects are cells and morphisms are inclusions. Thus, by virtue of being constructible, sheaves being parametrized can be described as quiver representations (with certain conditions), where the quiver is the one associated to the (cell) poset category of M. Second, the condition of having singular support contained in Λ translates into constraints, which further simplifies the description, as beautifully explained in [STZ17, Theorem 3.12]. Indeed, it follows from that result that it suffices to consider only maps between the top-dimensional 2-cells $\mathbb{R}^2 \setminus \pi_F(\Lambda)$, only in the direction opposite to co-orientation, and certain condition must be imposed for the maps near cusps and crossings. Following [STZ17, Section 3.3], the combination of these results of R. MacPerson, [KS90] and [STZ17] lead to the following description for $\mathcal{M}_1(\Lambda)$ and $\mathfrak{M}(\Lambda, T)$.

Given a co-oriented front projection $\pi_F(\Lambda)$, consider the following quiver $Q_F(\Lambda)$:

- A vertex of $Q_F(\Lambda)$ is placed at each connected component of the $\mathbb{R}^2 \setminus \pi_F(\Lambda)$,
- For each (1-dimensional) connected component of $\pi_F(\Lambda) \setminus S_0$, where S_0 denotes the set of crossings and cusps in $\pi_F(\Lambda)$, draw an arrow connecting the two vertices associated to the two adjacent 2-dimensional cell (that contain that stratum in their closure). The direction of the arrow is opposite to the co-orientation of the front $\pi_F(\Lambda)$.

The following is then proven in [STZ17, Section 3]:

Proposition 4.1. *Let $\Lambda \subset (\mathbb{R}^3, \xi)$ be a Legendrian and $\pi_F(\Lambda) \subset \mathbb{R}^2$ a front, with a binary Maslov potential, such that $(\mathbb{R}^2, \pi_F(\Lambda))$ is a regular stratification. Consider the stack*

---

[15]An expert in the results of [KS90, STZ17] might be able to quickly move forward to Subsection 4.2.

$\mathcal{M}(Q_F(\Lambda))$ *classifying linear representations of the quiver $Q_F(\Lambda)$ that satisfy the following conditions:*

(1) *The vector space associated with the unbounded region in $\mathbb{R}^2$ is 0.*
(2) *Any two vector spaces associated with neighboring vertices differ in dimension by 1.*
(3) *At each cusp, the composition depicted in Figure 67 (left) is the identity map.*
(4) *At each crossing, the four linear maps involved form a commuting square which is exact, as precised in Figure 67 (right).*

$$g \circ f = \mathrm{id}_V \qquad 0 \longrightarrow V_s \xrightarrow{(f_{sw}, f_{se})} V_w \oplus V_e \xrightarrow{f_{wn} - f_{en}} V_n \longrightarrow 0$$
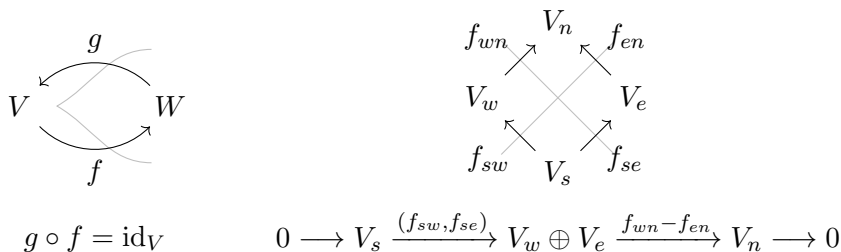
FIGURE 67. Identity condition at cusps and exactness condition at crossings.

Then the stack $\mathcal{M}_1(\Lambda)$ is isomorphic to $\mathcal{M}(Q_F(\Lambda))$.

Remark 4.2. The dimension vector of such quiver representations is uniquely determined by $\pi_F(\Lambda)$. Moreover, the identity condition at cusps and the exactness condition at crossings determine that each homomorphism is one of the following:

- injective with a 1-dimensional cokernel if it crosses a strand with Maslov potential 0;
- surjective with a 1-dimensional kernel if it crosses a strand with Maslov potential 1.

Example 4.3. *The left picture in Figure 68 draws the front from Figure 17 (left) for a Legendrian $m(5_2)$, and the vector spaces assigned to each top-dimensional cell of the complement of the front, according to the quiver $Q(\pi_F(\Lambda))$ and Proposition 4.1.*
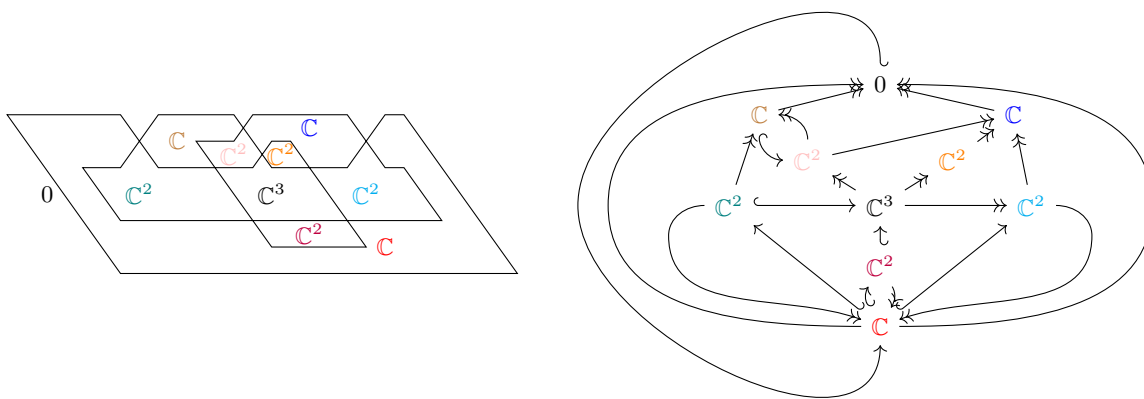
FIGURE 68. (Left) Defining the vertices of the quiver $Q(\pi_F(\Lambda))$ from the front, by assigning a vector space to each top-dimensional strata. In this example $\Lambda = \Lambda(\mathbb{G})$ from Figure 17, a Legendrian representative of $m(5_2)$. (Right) The (representations of the) quiver associated to the front. The representations consist of injective and surjective homomorphisms between vector spaces going across strands, as indicated.

Figure 68 (right) depicts the associated quiver $Q(\pi_F(\Lambda))$. Each quadrilateral piece in the quiver forms a short exact sequence; each bigon consists of one injective map and a surjective map, and the composition satifies (surjective map) ∘ (injective map) = identity.

Proposition 4.1 describes $\mathcal{M}_1(\Lambda)$. Now we gear towards the decorated moduli $\mathfrak{M}(\Lambda, T)$. First, we need a description of microlocal monodromy in terms of these quiver representations, which is provided in [STZ17, Section 5] and we briefly summarize as follows. Let S be the set of singular points (crossings and cusps) in $\pi_F(\Lambda)$. By Remark 4.2, each connected components of $\pi_F(\Lambda) \setminus S$ is associated with a 1-dimensional kernel or a 1-dimensional cokernel. These kernels and cokernels can be glued together along strands of Λ using the identity condition at cusps and the exactness condition at crossings, as follows:

- At a cusp, as in Figure 67 (left), the condition $g \circ f = \mathrm{id}_V$ forces the composition $\ker g \hookrightarrow W \twoheadrightarrow \operatorname{coker} f$ to be an isomorphism. By definition, we glue $\ker g$ and $\operatorname{coker} f$ using this isomorphism.

- At a crossing, as in Figure 67 (right), there are three cases depending, on the injectivity or surjectivity of the four maps: the four maps can be all injective, all surjective, or two injective with two surjective. In each of the three cases, we have the following isomorphisms from the exactness condition:

(4.1)
$$\operatorname{coker} f_{sw} \hookrightarrow \frac{V_n}{V_s} \twoheadrightarrow \operatorname{coker} f_{en}$$

$$\operatorname{coker} f_{wn} \twoheadleftarrow \frac{V_n}{V_s} \hookleftarrow \operatorname{coker} f_{se}$$

(4.2)
$$\ker f_{sw} \hookrightarrow \ker(f_{wn} \circ f_{sw}) = \ker(f_{en} \circ f_{se}) \twoheadrightarrow \ker f_{en}$$

$$\ker f_{wn} \twoheadleftarrow \ker(f_{wn} \circ f_{sw}) = \ker(f_{en} \circ f_{se}) \hookleftarrow \ker f_{se}$$

(4.3)
$$\operatorname{coker} f_{sw} \xrightarrow{f_{wn}} \operatorname{coker} f_{en}$$

$$\ker f_{wn} \xleftarrow{f_{sw}} \ker f_{se}$$

The result of gluing these 1-dimensional vector spaces is a rank-1 local system Φ on Λ. In fact, it coincides with the microlocal monodromy functor, see [STZ17, Section 5.1] for more details. Given the set T of marked points on Λ, with at least one marked point per link component, $\Lambda \setminus T$ is a collection of open intervals. Thus, along each such open interval I, we can trivialize the rank-1 local system Φ by specifying an isomorphism $\phi_I : I \times \mathbb{C} \xrightarrow{\cong} \Phi|_I$. By definition, a collection of such maps $\{\phi_I\}$ are said to be a *framing* for the local system Φ.

In conclusion, a point in the decorated moduli space $\mathfrak{M}(\Lambda, T)$, as defined in Subsection 2.8.3, is a point in $\mathcal{M}_1(\Lambda)$, which is combinatorialized via Proposition 4.1, together with a framing for the local system Φ, i.e. a trivialization of the (trivial) local system $\Phi|_{\Lambda \setminus T}$. In this description, wwo framings are considered *equivalent* if they differ by a global scaling $\mathbb{C}^\times$-factor. It follows that $\dim \mathfrak{M}(\Lambda, T) = \dim \mathcal{M}_1(\Lambda) + |T| - 1$.

4.1.1. *Descriptions for GP-graph $\mathbb{G}$.* In the case that $\Lambda = \Lambda(\mathbb{G})$ comes from a GP-graph $\mathbb{G}$, Section 2.4 provides a specific front $\mathfrak{f}(\mathbb{G}) \subset \mathbb{R}^2$. For this front, the description from Proposition 4.1 can be translated in terms of configurations of flags, as follows. The front projection

$\pi_F(\Lambda) = \mathfrak{f}(\mathbb{G})$ can be sliced into the three types of elementary columns.

For a Type 1 column, there are n strands in the bottom region, with Maslov potential 0, and n strands in the top region, with Maslov potential 1. By Proposition 4.1, the vector space associated with the central region must be $\mathbb{C}^n$. From the quiver representation data, we can construct the following pair of flags in $\mathbb{C}^n$. Since all of the linear maps in the bottom region are injective, their images in the middle $\mathbb{C}^n$ naturally form a first flag. Similarly, since all of the linear maps in the top region are surjective, their kernels in the middle vector space $\mathbb{C}^n$ form a second flag. See Figure 69 for a depiction of the front in Type 1 and its associated pair of flags. We adopt the convention of indexing flags from the bottom region with a subscript, and indexing flags from the top region with a superscript, so as to distinguish between these two types of two flags.

$$\mathscr{F}^0 = \left(\begin{array}{c} 0 \subset \ker g_n \subset \ker(g_{n-1} \circ g_n) \subset \\ \cdots \subset \ker(g_2 \circ \cdots \circ g_n) \subset \mathbb{C}^n \end{array} \right)$$

$$\mathscr{F}_0 = \left(\begin{array}{c} 0 \subset \mathrm{img}(f_2 \circ \cdots \circ f_n) \subset \mathrm{img}(f_3 \circ \cdots \circ f_n) \\ \subset \cdots \subset \mathrm{img}\, f_n \subset \mathbb{C}^n \end{array} \right)$$
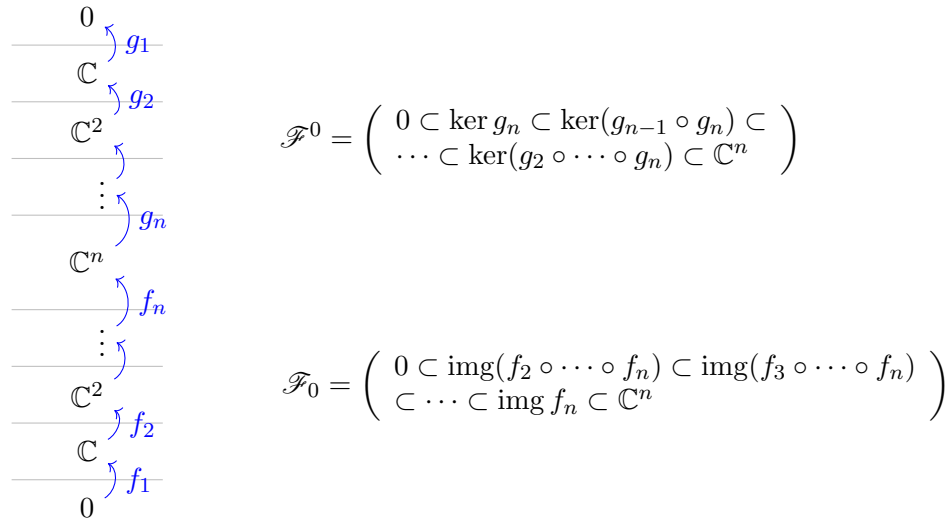
FIGURE 69. The pair of flags associated to a Type 1 column.

Before discussing Type 2 and 3 columns, we recall that the relative positions relations between two flags in $\mathbb{C}^n$ are classified by elements of the symmetric group S_n, which is a Coxeter group with Coxeter generators $\{s_i\}_{i=1}^{n-1}$. For two flags $\mathscr{F} = (0 \subset \mathscr{F}_1 \subset \cdots \subset \mathbb{C}^n)$ and $\mathscr{F}' = (0 \subset \mathscr{F}'_1 \subset \cdots \subset \mathbb{C}^n)$, we denote

- $\mathscr{F} \overset{s_i}{\sim} \mathscr{F}'$ if $F_i \neq F'_i$ but $F_j = F'_j$ for all $j \neq i$;
- $\mathscr{F} \overset{w}{\sim} \mathscr{F}'$ if there exists a sequence of flags $\mathscr{G}_0, \mathscr{G}_1, \ldots, \mathscr{G}_l$ such that

$$\mathscr{F} = \mathscr{G}_0 \overset{s_{i_1}}{\sim} \mathscr{G}_1 \overset{s_{i_2}}{\sim} \mathscr{G}_2 \overset{s_{i_3}}{\sim} \cdots \overset{s_{i_l}}{\sim} \mathscr{G}_l = \mathscr{F}'$$

  and $s_{i_1} s_{i_2} \cdots s_{i_l}$ is a reduced word of w.

This classification can be identified with the Tits distance obtained from a Bruhat decomposition of GL_n. In particular, being in w relative position does not depend on the choice of reduced word of w. If $\mathscr{F} \overset{w}{\sim} \mathscr{F}'$, then for each choice of reduced word $(i_1, \ldots, i_l)$ for w, there exists a unique sequence of flags $(\mathscr{G}_k)_{k=0}^l$ that relate the two flags $\mathscr{F}$ and $\mathscr{F}'$.

We can now translate the local quiver representation data associated with a Type 2 column into relative position relations between flags. Suppose the pair of flags to the left of a Type 2 column is $(\mathscr{L}_0, \mathscr{L}^0)$, and the pair of flags to the right of a Type 2 column is $(\mathscr{R}_0, \mathscr{R}^0)$. If there is a crossing in the bottom region at the ith gap, counting from the bottom in both the front and in the GP-graph, then from the exactness condition at the crossing obtain the

constraints

(4.4)
$$\mathcal{L}_0 \overset{s_i}{\sim} \mathcal{R}_0 \qquad \text{and} \qquad \mathcal{L}^0 = \mathcal{R}^0.$$

Similarly, if there is a crossing in the top region at the ith gap, counting from the top in the front projection or counting from the bottom in the GP-graph, then the exactness condition at the crossing yields

(4.5)
$$\mathcal{L}_0 = \mathcal{R}_0 \qquad \text{and} \qquad \mathcal{L}^0 \overset{s_{n-i}}{\sim} \mathcal{R}^0.$$

Since crossings in Type 2 columns correspond to vertical edges in the GP-graph, we can infer the relative position relation between pairs of flags from the GP-graph as well.

Example 4.4. *The following four Type 2 columns of a GP-graph give rise to the following relative position relations between pairs of flags:*

$$\mathcal{L}^0 = \mathcal{R}^0 \qquad \mathcal{L}^0 = \mathcal{R}^0 \qquad \mathcal{L}^0 \overset{s_2}{\sim} \mathcal{R}^0 \qquad \mathcal{L}^0 \overset{s_1}{\sim} \mathcal{R}^0$$
$$\mathcal{L}_0 \overset{s_1}{\sim} \mathcal{R}_0 \qquad \mathcal{L}_0 \overset{s_2}{\sim} \mathcal{R}_0 \qquad \mathcal{L}_0 = \mathcal{R}_0 \qquad \mathcal{L}_0 = \mathcal{R}_0$$

For a Type 3 column, the pairs of flags on the (Type 1 column on the) left and on the (Type 1 column on the) right are not in the same ambient vector space, as the dimensions of the two vector spaces differ by one. Instead, there is a linear map *from* the ambient vector space for the pair of flags on the left *to* the ambient vector space for the pair of flags on the right. This linear map is injective if the lollipop is white and it is surjective if the lollipop is black. Let us investigate how the two pairs of flags are related.

Suppose first that the lollipop is white, so that the linear map $h : \mathbb{C}^{n-1} \to \mathbb{C}^n$ between the two (middle) adjacent ambient vector spaces is injective. Given any flag $\mathscr{F} = (0 \subset \mathscr{F}_1 \subset \cdots \subset \mathscr{F}_{n-1} = \mathbb{C}^{n-1})$ in $\mathbb{C}^{n-1}$, we can use h to naturally extend it to a flag $h(\mathscr{F})$ in $\mathbb{C}^n$ by defining
$$h(\mathscr{F}) := (0 \subset h(\mathscr{F}_1) \subset h(\mathscr{F}_2) \subset \cdots \subset h(\mathscr{F}_{n-1}) \subset \mathbb{C}^n).$$

This extension from $(\mathcal{L}_0, \mathcal{L}^0)$ to $(h(\mathcal{L}_0), h(\mathcal{L}^0))$ can be achieved geometrically by a sequence of RII moves that pulls the left cusp upward in the front projection. Indeed, consider the following local example:

FIGURE 70. Pulling up a left cusp

The green maps in the bottom region define the extension $h(\mathscr{L}_0)$. By the exactness of the quadrilaterals in the top region, the red maps define the extension $h(\mathscr{L}^0)$. In particular, the extensions $h(\mathscr{L}_0)$ and $h(\mathscr{L}^0)$ are completely determined by the original data of the quiver representations.

Now, with these extensions defined, it follows that if there is a white lollipop emerging in the ith gap with $0 \leq i \leq n-1$, counting from below in the GP-graph[16], then the corresponding relative position conditions are

$$(4.6) \qquad h(\mathscr{L}_0) \overset{s_{n-1}\cdots s_{i+1}}{\sim} \mathscr{R}_0 \quad \text{and} \quad h(\mathscr{L}^0) \overset{s_{n-1}\cdots s_{n-i}}{\sim} \mathscr{R}^0.$$

Suppose that there is a black lollipop, and thus the linear map between the two ambient vector spaces $h : \mathbb{C}^n \to \mathbb{C}^{n-1}$ is surjective. Then, given any flag $\mathscr{F} = (0 \subset \mathscr{F}_1 \subset \cdots \subset \mathscr{F}_{n-1} = \mathbb{C}^{n-1})$ in $\mathbb{C}^{n-1}$, we consider $h^{-1}(\mathscr{F}_i)$ and insert $\ker(h)$ in front of it so as to form a flag in $\mathbb{C}^n$:

$$h^{-1}(\mathscr{F}) := (0 \subset \ker(h) \subset h^{-1}(\mathscr{F}_1) \subset \cdots \subset h^{-1}(\mathscr{F}_{n-1}) = \mathbb{C}^n).$$

FIGURE 71. Pulling down a right cusp

Similar to the white lollipop case, the extension of $(\mathscr{R}_0, \mathscr{R}^0)$ to $(h^{-1}(\mathscr{R}_0), h^{-1}(\mathscr{R}^0))$ can be achieved geometrically by a sequence of RII moves that pulls the right cusp downward in the front projection, as depicted in Figure 71. Note that the green maps in bottom region define the extension $h^{-1}(\mathscr{R}_0)$, whereas the red maps in the top region define the extension $h^{-1}(\mathscr{R}^0)$.

It follows that if there is a black lollipop occurring in the ith gap with $0 \leq i \leq n$, counting from below in the GP-graph, then the corresponding relative position conditions are

$$(4.7) \qquad \mathscr{L}_0 \overset{s_{i-1}\cdots s_1}{\sim} h^{-1}(\mathscr{R}_0) \quad \text{and} \quad \mathscr{L}^0 \overset{s_{n-i+1}\cdots s_1}{\sim} h^{-1}(\mathscr{R}^0).$$

In summary, given a GP-graph $\mathbb{G}$, we can divide $\mathbb{G}$ into columns of three types such that every consecutive pair of non-Type 1 columns is separated by a Type 1 column and every consecutive pair of Type 1 columns is separated by a non-Type 1 column. In between Type 3 columns there is a unique ambient vector space $V_i = \mathbb{C}^n$ for some n, and they are linked by linear maps $h_i : V_{i-1} \to V_i$ that are either injective with a 1-dimensional cokernel or surjective with a 1-dimensional kernel. The above discussion proves the following lemma:

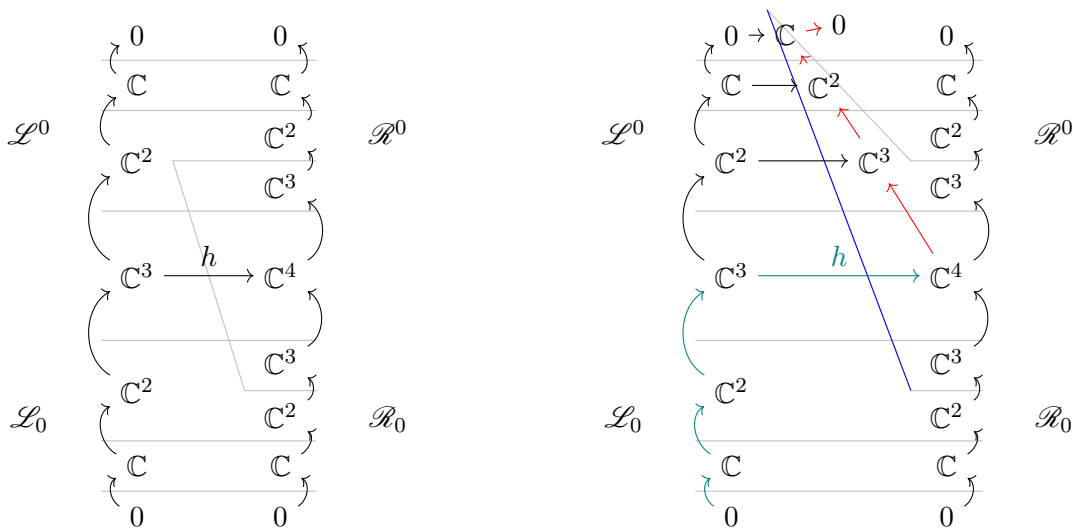[16]The case $i = 0$ is a lollipop at the bottom, and $i = n-1$ is a lollipop at the top.

Lemma 4.5. *For a GP-graph $\mathbb{G}$ with a decomposition into columns as above, the moduli space $\mathcal{M}_1(\Lambda)$ can be described by the following data:*

(1) *a pair of flags in V_i for each Type 1 column contained in the V_i part of $\mathbb{G}$;*
(2) *for each Type 2 column, the neighboring flags satisfy the relative position condition according to (4.4) and (4.5);*
(3) *for each Type 3 column, the neighboring flags satisfy the relative position condition according to (4.6) and (4.7);*

where we quotient this data by the equivalence relation $(\mathscr{F}, h) \sim (\mathscr{F}', h')$ for a collection of elements $g_i \in \mathrm{GL}(V_i)$ such that $h_i \circ g_{i-1} = g_i \circ h'_i$.

In the flag description of Lemma 4.5, the rank-1 local system Φ on Λ can be constructed by taking quotients of consecutive vector subspaces in each flag and then gluing them along strands of Λ at crossings and cusps in the same manner before. Note that in this context, only (4.1) is used when gluing these rank-1 local systems at crossings because all linear maps near a crossing are now inclusions of vector subspaces. In particular, the surjective maps in the top region of the front projection are now turned into inclusions of kernels.

Remark 4.6. Note that the trivialization of the rank-1 local system Φ on Λ can be translated into the flags as well, e.g. by considering affine flags and allowing for a $\mathbb{C}^\times$-rescaling of the (volume form associated to the) flag at the ith position, top or bottom, if a marked point is at the ith strand, top or bottom.

4.1.2. *Descriptions for (-1)-closures.* Finally, there is another description of $\mathcal{M}_1(\Lambda)$ and $\mathfrak{M}(\Lambda, \mathrm{T})$ as moduli space of configurations of flags, which aligns better when comparing with the flag moduli of the weaves $\mathfrak{w}(\mathbb{G})$. In that latter case, there will be only one ambient vector space. The description in Lemma 4.5, which is associated to the specific front $\mathfrak{f}(\mathbb{G})$, after using RII and RIII moves, can be shown to be equivalent to a description with a unique ambient (top-dimensional) vector space. Indeed, rather than using flags from different ambient spaces with varying dimensions, we can perform additional RII and RIII moves to push strands as the blue one in Figure 70 all the way to the left and push strands as the blue one in Figure 71 all the way to the right (see also Lemma 3.27). This will extend all flags from all Type 1 columns to flags in $\mathbb{C}^h$ where h is the total number of horizontal lines in the GP-graph $\mathbb{G}$. Moreover, these flags will satisfy the relative position conditions imposed by the external weave lines of the initial weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$, or equivalently, the cyclic positive braid word $\beta = \beta(\mathbb{G})$ for which Λ is its (-1) closure. In this context, [STZ17, Proposition 1.5], or Lemma 4.1, reads:

Lemma 4.7. *Let $\beta = (i_1, i_2, \ldots, i_l) \in Br_h^+$ be a positive braid word on h strands and Λ be the Legendrian link associated to the front given by the (-1)-closure of β. Then*

$$\mathcal{M}_1(\Lambda) \cong \left\{ (\mathscr{F}_0, \mathscr{F}_1, \mathscr{F}_2, \ldots, \mathscr{F}_l) \; \middle| \; \begin{array}{c} \mathscr{F}_i \text{ is a flag in } \mathbb{C}^h \text{ for all } i \\ \mathscr{F}_0 \overset{s_{i_1}}{\sim} \mathscr{F}_1 \overset{s_{i_2}}{\sim} \ldots \overset{s_{i_l}}{\sim} \mathscr{F}_l = \mathscr{F}_0 \end{array} \right\} \Big/ \mathrm{PGL}_h \,.$$

Example 4.8. *Let us consider the Legendrian representative of $m(5_2)$ from Figure 16 (left). After performing a sequence of RII and RIII moves, we obtain the following front projection:*

72
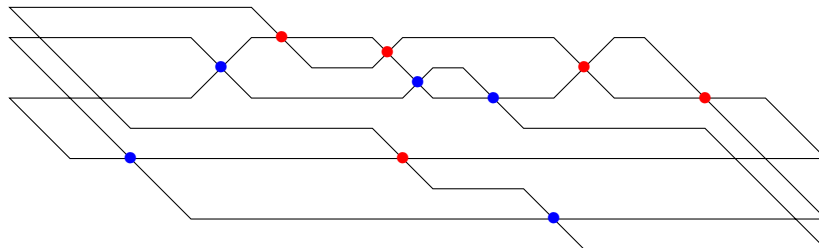
Then Lemma 4.7 implies that $\mathcal{M}_1(\Lambda)$ is isomorphic to the moduli space of configurations of configurations of ten flags in $\mathbb{C}^3$ satisfying the following relative position conditions, modulo the global PGL_3-action:
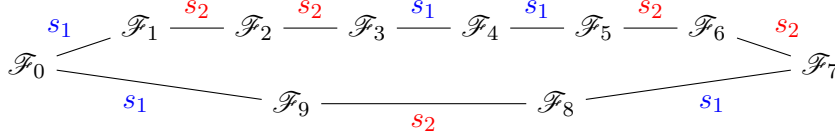
4.2. Factoriality Property.

In the upcoming construction of a cluster $\mathcal{A}$-structure for the moduli space $\mathfrak{M}(\Lambda, T)$, we shall need that the coordinate ring $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ is a unique factorization domain (a.k.a. factorial). This can be a subtle condition to verify and thus we provide in this section an argument that the condition of Δ-completeness of the braid $\beta(\mathbb{G})$, as introduced in Section 2.5, is sufficient for factoriality. Note that all shuffle graphs have $\beta(\mathbb{G})$ be a Δ-complete braid, and thus the rings $\mathcal{O}(\mathfrak{M}(\Lambda(\mathbb{G}), T))$ are factorial if $\mathbb{G}$ is shuffle.

For that we relate $\mathfrak{M}(\Lambda, T)$ to the braid varieties introduced and studied in [CGGS20, CGGS21]. For each simple root α_i, $1 \leq i < n$, there is a natural group homomorphism $\varphi_i : \mathrm{GL}_2 \to \mathrm{GL}_n$, and we denote $B_i(z) := \varphi_i \begin{pmatrix} 0 & 1 \\ 1 & z \end{pmatrix}$. Let $\beta = (i_1, \ldots, i_l)$ be a positive braid (word) and w a permutation matrix. By definition, the braid variety $X(\beta, w)$ is the subvariety in $\mathbb{A}^l = \mathrm{Spec}[z_1, \ldots, z_l]$ defined by the equations that impose that the product

$$B_{i_1}(z_1) B_{i_2}(z_2) \cdots B_{i_l}(z_l) w$$

is an upper triangular matrix. It is proven in [CGGS21, Theorem 1.8] that $X(\beta, w)$ is a smooth affine variety if $\mathrm{Dem}(\beta) = w$. In order to conclude that $\mathcal{O}(X(\beta, w))$, it suffices to show that the ideal class group $\mathrm{Cl}(X(\beta, w)) = 0$ vanishes. This is in fact true, as was proven to us by E. Gorsky, to whom we are very grateful, in a private conversation:

Lemma 4.9 (E. Gorsky). *Let β be a positive braid with Demazure product $\mathrm{Dem}(\beta) = w_0$. Then the coordinate ring $\mathcal{O}(X(\beta, w_0))$ is factorial.*

Proof. The braid variety $X = X(\beta, w_0)$ admits a smooth compactification by the closed brick manifold Y for the same braid word, as discussed in β [CGGS21, Theorem 4.8 (ii)]. It follows from the excision exact sequence of class groups [Vak15, Section 14.2.8], that $\mathcal{O}(X)$ is a unique factorization domain if and only if the irreducible components of $Y \setminus X$ span the Picard group $\mathrm{Pic}(Y)$. Now, the Picard group of the brick manifold Y is generated by the line bundles $\{\mathcal{L}_s\}_{s=1}^l$, one for each crossings in $\beta = (i_1, \ldots, i_l)$, see [LT04, And19]. The divisor associated to $\mathcal{L}_s$ is either the closed brick manifold for $\beta_s := (i_1, \ldots, \hat{i}_s, \ldots, i_l)$, if $\mathrm{Dem}(\beta_s) = w_0$, which is an irreducible component of $Y \setminus X$, or empty otherwise. In the former case this is an irreducible component in $Y \setminus X(\beta, w_0)$, and in the latter case the line bundle is trivial. Thus irreducible components of $Y \setminus X$ span the Picard group $\mathrm{Pic}(Y)$ and $\mathcal{O}(X(\beta, w_0))$ is factorial as required. $\square$

Lemma 4.9 implies the factoriality of $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ as required:

Proposition 4.10. *Let $\mathbb{G}$ be a GP-graph with $\beta(\mathbb{G})$ a Δ-complete braid. Then the moduli space $\mathfrak{M}(\Lambda(\mathbb{G}), T)$ is an affine variety whose coordinate ring is factorial.*

Proof. Since the moduli space $\mathfrak{M}(\Lambda, T)$ is an Legendrian invariant, without loss of generality we can turn Λ into the (-1)-closure of an n-stranded positive braid $\beta(\mathbb{G}) = \Delta\gamma$ and use the description from Section 4.1.2. Let us first consider the case where the set T of marked points can be arranged into a configuration with one marked point per level along a vertical

73

line between Δ and γ. (It follows that $|T| = n$.) This case is depicted as follows:
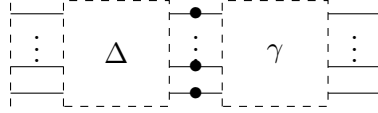
Let B_+ and B_- be the Borel subgroups of PGL_n of upper triangular and lower triangular matrices, respectively. We can exhaust the PGL_n-action on flag configurations by fixing the two flags at the two ends of Δ to be the two unique flags stabilized by B_+ and B_-, respectively, while requiring that the decoration on the flag $\mathscr{F}_l$ at the dashed line (after γ on the right or before Δ on the left) to be the standard one, i.e., mapping $\bar{e}_i$ to 1 for each consecutive quotient $\mathrm{Span}\{e_1, \ldots, e_i\}/\mathrm{Span}\{e_1, \ldots, e_{i-1}\} \cong \mathrm{Span}(\bar{e}_i)$.

Let $(i_1, \ldots, i_l)$ be a positive word for the positive braid γ such that $\beta(\mathbb{G}) = \Delta\gamma$. Let us record a flag as a matrix with row vectors, such that the span of the last k row vectors give the k-dimensional subspace in the flag. Then $\mathscr{F}_l$ can be recorded by the permutation matrix w_0. Starting from the flag $\mathscr{F}_l$, the flags $\mathscr{F}_{l-1}, \mathscr{F}_{l-2}, \ldots$ to the left of $\mathscr{F}_l$ can then be given by

$$\mathscr{F}_k = B_{i_{k+1}}(z_{k+1})B_{i_{k+2}}(z_{k+2}) \cdots B_{i_l}(z_l)w_0.$$

In the end, we need $\mathscr{F}_0$ to be the standard flag

$$0 \subset \mathrm{Span}\{e_n\} \subset \mathrm{Span}\{e_{n-1}, e_n\} \subset \cdots \subset \mathrm{Span}\{e_2, \ldots, e_n\} \subset \mathbb{C}^n,$$

which is equivalent to requiring that $B_{i_1}(z_1)B_{i_2}(z_2) \cdots B_{i_k}(z_l)w_0$ to be upper triangular. This shows that $\mathfrak{M}(\Lambda(\mathbb{G}), T)$ is isomorphic to the braid variety $X(\beta(\mathbb{G}), w_0)$ from [CGGS20], which has a factorial coordinate ring.[17]

Now let us consider the case of an arbitrary number of marked points. Let us start with the set T having one marked point per level, as in the case above. Suppose m of the marked points share the same link component, then we can move these marked points along that link component until they get inside an horizontal interval with no crossings or cusps. Then these marked points are just changing decorations on the same underlying 1-dimensional quotient of consecutive vector spaces of the same flag. Thus we can extract out a $(\mathbb{C}^\times)^{m-1}$-torus factor and replace these marked points with one marked point. By doing this for each link component, we can reduce T to a set T' with one marked point per link component, and conclude that

$$\mathfrak{M}(\Lambda, T) \cong \mathfrak{M}(\Lambda, T') \times (\mathbb{C}^\times)^{n-N}$$

as affine varieties, where N is the number of link components in $\Lambda = \Lambda(\mathbb{G})$. This implies that

$$\mathcal{O}(\mathfrak{M}(\Lambda, T)) \cong \mathcal{O}(\mathfrak{M}(\Lambda, T')) \otimes \mathbb{C}[t_i^{\pm 1}]_{i=1}^{n-N}.$$

If there is an element in $\mathcal{O}(\mathfrak{M}(\Lambda, T'))$ admitting two non-equivalent factorizations, then these two factorizations are still valid and non-equivalent in $\mathcal{O}(\mathfrak{M}(\Lambda, T))$, contradicting the fact that $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ is factorial. Thus, we can conclude that $\mathcal{O}(\mathfrak{M}(\Lambda, T'))$ is factorial when T' consists of one marked point per link component. In general, for any set T'' with at least one marked point per link component, we can implement the same argument above and write

$$\mathfrak{M}(\Lambda, T'') \cong \mathfrak{M}(\Lambda, T') \times (\mathbb{C}^\times)^{|T''|-N}.$$

Algebraically this implies that

$$\mathcal{O}(\mathfrak{M}(\Lambda, T'')) \cong \mathcal{O}(\mathfrak{M}(\Lambda, T')) \otimes \mathbb{C}[t_i^{\pm 1}]_{i=1}^{|T''|-N}.$$

Again, since $\mathcal{O}(\mathfrak{M}(\Lambda, T'))$ is factorial, so is the tensor product $\mathcal{O}(\mathfrak{M}(\Lambda, T')) \otimes \mathbb{C}[t_i]_{i=1}^{|T''|-N}$. Given that $\mathcal{O}(\mathfrak{M}(\Lambda, T')) \otimes \mathbb{C}[t_i^{\pm 1}]_{i=1}^{|T''|-N}$ is a localization of this factorial tensor product, it is factorial as well. $\qquad \square$

---

[17]We thank Eugene Gorsky for an explanation of why this is the case. See also upcoming work of the first author with E. Gorsky and co-authors where this is written in detail.

4.3. **Moduli spaces for the Lagrangian** $L(\mathfrak{w}(\mathbb{G}))$. The moduli spaces $\mathcal{M}_1(\Lambda)$ and $\mathfrak{M}(\Lambda, T)$ depend only on the Legendrian isotopy type of Λ. In particular, if $\Lambda = \Lambda(\mathbb{G})$ is a GP-link, then these moduli spaces are invariant under square moves and other combinatorial equivalences of the GP-graph $\mathbb{G}$ which preserve the Legendrian isotopy class of Λ. The GP-graph also provides the information of an embedded exact Lagrangian filling for $\Lambda(\mathbb{G})$. Namely, the exact Lagrangian filling $L = L(\mathbb{G})$ described by the initial weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$. The Guillermou-Jin-Treumann map [JT17], or [EHK16, CN21], imply that there are open embeddings

$$H^1(L; \mathbb{C}^\times) \longrightarrow \mathcal{M}_1(\Lambda), \quad H^1(L, T; \mathbb{C}^\times) \longrightarrow \mathfrak{M}(\Lambda, T),$$

where the domains of these embeddings parametrize (decorated) $\mathbb{C}$-local systems on L (with decoration T), and the map is essentially the microlocalization functor. These open torus charts $(\mathbb{C}^\times)^{b_1(L)}$ and $(\mathbb{C}^\times)^{b_1(L,T)}$ can be described in terms of flags if the Lagrangian filling L is obtained from a weave, as explained in [CZ22]; we shall use it in the proof of Theorem 1.1. The definition of $\mathcal{M}_1(\mathfrak{w})$ from [CZ22] is as follows:

Definition 4.11. Let $\mathfrak{w} \subset \mathbb{R}^2$ be a weave. By definition, the *total flag moduli space* $\tilde{\mathcal{M}}_1(\mathfrak{w})$ associated to $\mathfrak{w}$ is comprised of tuples of flags, as follows:

 i) There is a flag $\mathscr{F}^\bullet(F)$ assigned to each face F of the weave $\mathfrak{w}$, i.e. to each connected component of $\mathbb{R}^2 \setminus \mathfrak{w}$.

 ii) For each pair of adjacent faces $F_1, F_2 \subset \mathbb{R}^2 \setminus \mathfrak{w}$, sharing an s_i-edge, their two associated flags $\mathscr{F}^\bullet(F_1), \mathscr{F}^\bullet(F_2)$ are in relative position $s_i \in S_n$, i.e. they must satisfy

$$\mathscr{F}_j(F_1) = \mathscr{F}_j(F_2), \quad 0 \le j \le N, j \ne i, \text{ and } \mathscr{F}_i(F_1) \ne \mathscr{F}_i(F_2).$$

The group PGL_n acts on the space $\tilde{\mathcal{M}}_1(\mathfrak{w})$ simultaneously. By definition, the *flag moduli space* of the weave $\mathfrak{w}$ is the quotient stack $\mathcal{M}_1(\mathfrak{w}) := \tilde{\mathcal{M}}(\mathfrak{w})/\mathrm{PGL}_n$.

By Subsection 4.1.1, $\mathcal{M}_1(\mathfrak{w})$ is an open subspace of $\mathcal{M}_1(\Lambda)$ via restriction to the boundary. Indeed, since the weaves $\mathfrak{w}$ are free weaves [CZ22, Section 7.1.2], the data of flags at the boundary of the initial weave, uniquely determines the flags at each face of $\mathfrak{w}$. (This fact can also be verified combinatorially.) It follows from [CZ22] that $\mathcal{M}_1(\mathfrak{w})$ are complex tori $\mathcal{M}_1(\mathfrak{w}) \cong (\mathbb{C}^\times)^{\dim \mathcal{M}_1(\Lambda)}$, and thus these moduli spaces of flags associated to the initial weave $\mathfrak{w}$ are natural candidates for an initial cluster chart in the moduli space $\mathcal{M}_1(\Lambda)$ for a GP-link Λ. (These complex tori are indeed the images of the Guillermou-Jin-Treumann maps.) The definition of candidate cluster $\mathcal{X}$-variables will be the subject of the next subsection.

The decorated version of the flag moduli $\mathcal{M}_1(\mathfrak{w})$, which we denote as $\mathfrak{M}(\mathfrak{w}, T)$, is naturally defined by adding a framing away from T along the boundary $\partial L(\mathfrak{w}) = \Lambda$. It also follows that $\mathfrak{M}(\mathfrak{w}, T)$ is naturally an open torus chart in $\mathfrak{M}(\Lambda, T)$. The corresponding definition of the candidate cluster $\mathcal{A}$-variables is undertaken in Subsection 4.6 below.

4.4. **Microlocal monodromies: unsigned candidate $\mathcal{X}$-variables.** Let us consider the open toric chart $\mathcal{M}_1(\mathfrak{w}) \subset \mathcal{M}_1(\Lambda)$ from Subsection 4.3. We now build a function

$$X_\gamma : \mathcal{M}_1(\mathfrak{w}) \longrightarrow \mathbb{C}$$

associated to each Y-cycle γ, generalizing our previous work in [CZ22, Section 7], see also [STZ17, Section 5.1], to our context. First we observe that the data of $\mathcal{M}_1(\mathfrak{w})$ associates a flag $\mathscr{F}$ in each connected component of the complement of $\mathfrak{w}$ in $\mathbb{R}^2$. We associate the 1-dimensional vector space $\mathscr{F}_i/\mathscr{F}_{i-1}$ to the ith sheet in the lift of each connected component.

Then across the lifts of each weave line, we define two linear isomorphisms

$$\mathscr{L}_{i+1} = \mathscr{R}_{i+1}$$

$$\mathscr{L}_{i+1}/\mathscr{L}_i \qquad \mathscr{R}_{i+1}/\mathscr{R}_i$$

$$\mathscr{L}_i \qquad \mathscr{R}_i$$

$$\mathscr{L}_i/\mathscr{L}_{i-1} \qquad \mathscr{R}_i/\mathscr{R}_{i-1}$$

$$\mathscr{L}_{i-1} = \mathscr{R}_{i-1}$$

$$\psi_+ : \frac{\mathscr{L}_i}{\mathscr{L}_{i-1}} \hookrightarrow \frac{\mathscr{L}_{i+1}}{\mathscr{L}_{i-1}} = \frac{\mathscr{R}_{i+1}}{\mathscr{R}_{i-1}} \twoheadrightarrow \frac{\mathscr{R}_{i+1}}{\mathscr{R}_i}.$$

$$\psi_- : \frac{\mathscr{L}_i}{\mathscr{L}_{i-1}} \hookrightarrow \frac{\mathscr{L}_{i+1}}{\mathscr{L}_{i-1}} = \frac{\mathscr{R}_{i+1}}{\mathscr{R}_{i-1}} \twoheadrightarrow \frac{\mathscr{R}_{i+1}}{\mathscr{R}_i}.$$

Given a loop γ on L, we may perturb it so that it intersects with any lifts of weave lines transversely. Then by composing isomorphisms like $\psi_\pm$ above, we obtain a linear automorphism for each generic fiber along γ. Since each generic fiber is a 1-dimensional vector space, we can represent this linear automorphism by a non-zero scalar ψ_γ. This non-zero scalar ψ_γ is also known as the *microlocal monodromy* of the sheaf moduli space $\mathcal{M}_1(\mathfrak{w})$ along γ. However, the microlocal monodromies ψ_γ do not naturally give rise to a local system on L[18], as the following illustrate:

Example 4.12. *Consider the weave with a unique trivalent vertex, which depicts a Lagrangian 2-disk filling, as drawn in blue in Figure 72 (left). According to the definition of $\mathcal{M}_1(\mathfrak{w})$, there is a flag $l_i \subset \mathbb{C}^2$ in each of the three sectors, and they are pairwise transverse from each other. Let γ be a curve on $L(\mathfrak{w}) \cong \Lambda(\mathfrak{w})$, which, under the front projection, goes from the lower sheet to the upper sheet and then back to the lower sheet; see again Figure 72 (left). By definition, the microlocal parallel transport ψ_γ should be the map drawn in Figure 72 (center), which is the linear map that projects parallel to the line l_2.*
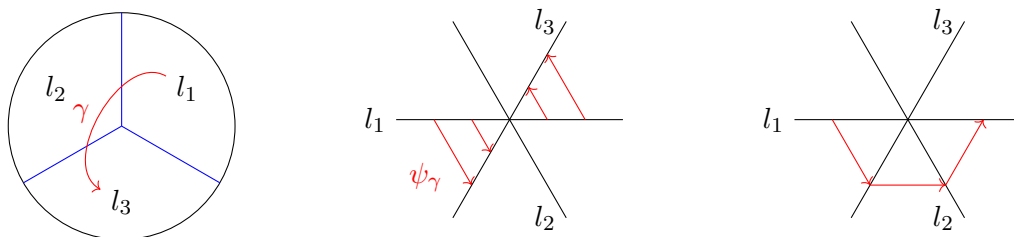
FIGURE 72. A weave for the unique filling of the max tb unknot and microlocal parallel transports from its sheaf quantization.

Consider the lift ξ of a loop that goes around a trivalent weave vertex in $\mathbb{R}^2$, which is a double cover for the projection onto the weave plane. Without loss of generality, let us suppose ξ starts at the lower sheet in the sector containing l_1. The parallel transport along ξ is then the composition of the three linear projections, as drawn in Figure 72 (center), which is equal to the linear map $v \mapsto -v$ on l_1. In other words, $\psi_\xi = -1$. However, ξ is a contractible cycle on $L(\mathfrak{w})$ and thus the microlocal monodromy assignment $\xi \mapsto \psi_\xi$ cannot be a local system on $L(\mathfrak{w})$.

Remark 4.13. This (-1) appearing in the microlocal monodromy along a contractible loop around a trivalent weave vertex is the only obstruction to obtain a local system on L. If we replace $\mathbb{C}$ with a field $\mathbb{F}$ of characteristic 2, then $\gamma \mapsto \psi_\gamma$ naturally defines a rank 1 local

---

[18]They naturally give a *twisted* local system as in [Gui19, Part 13], or a twisted flat connection as in [GMN13, Part 10].

system on L with fiber $\mathbb{F}$. In Subsection 4.5 we explain how to resolve this sign matter in our weave context.

Let us now specialize to our situation, with $\mathbb{G}$ a GP-graph and $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$ its initial weave. By Section 3, there is a distinguished linearly independent subset $\mathfrak{S}(\mathbb{G}) \subset H_1(L(\mathfrak{w}(\mathbb{G})))$ of $\mathbb{L}$-compressible cycles parametrized by the sugar-free hulls in the GP-graph. For each element in $\mathfrak{S}(\mathbb{G})$, we choose a Y-tree representative γ, which exists by Lemma 3.38, and define

$$X_\gamma := -\psi_\gamma.$$

These functions shall become our cluster $\mathcal{X}$-variables, once signs are fixed and Theorem 1.1 is proven. Note that, since we can isotope the Y-tree γ to a short I-cycle, i.e. an equivalent monochromatic edge, we may use it to compute X_γ explicitly, as follows. In a neighborhood of a short I-cycle, labeled with the permutation s_i, a point in the flag moduli $\mathcal{M}_1(\mathfrak{w})$ is specified by the data of a quadruple of flags. Each of these flags has the same subspaces $\mathscr{F}^j$ 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 $V := \mathscr{F}^{i+1}/\mathscr{F}^{i-1}$. This is the data of four lines $a, b, c, d \subset V$.

$$
\begin{array}{c}
b \\
a \,\diagup\!\!\!\!\diagdown\, c \\
d
\end{array}
$$
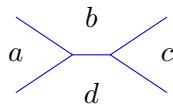
The function X_γ is then equal to the cross-ratio

$$X_\gamma = \langle a, b, c, d \rangle = -\frac{a \wedge b}{b \wedge c} \cdot \frac{c \wedge d}{d \wedge a}.$$

The definition of X_γ, following [FG06b] and [STWZ19, CZ22], is not particularly new. It is also possible to define X_γ directly and combinatorially from the Y-trees, in line with [CZ22, Section 7]. The fact that these functions $\{X_\gamma\}$ transform according to an $\mathcal{X}$-mutation formula under a square-face mutation is due to [STWZ19], and under the more general weave mutation due to [CZ22]. Indeed, let $\Gamma = \{\gamma_i\}$ be a maximal collection of Y-trees in $\mathfrak{w}(\mathbb{G})$ which are linearly independent in $H_1(L(\mathfrak{w}(\mathbb{G})))$, $Q(\Gamma)$ be their (algebraic) intersection quiver, and $X_\Gamma = \{X_{\gamma_i}\}$ be a labeling of each vertex of the quiver. Then, it is shown in [CZ22, Section 7.2.2] that weave mutation at one such Y-tree $\gamma \in \Gamma$ induces a quiver mutation of $Q(\Gamma)$ at the vertex associated to γ, and the set of variables X_Γ changes according to a cluster $\mathcal{X}$-mutation.

Defining these candidate cluster $\mathcal{X}$-variables is relatively useless for the purpose of proving existence of cluster structures: the variables X_γ do *not* extend to global in $\mathcal{M}_1(\Lambda)$ in general and we cannot deduce the existence of a cluster $\mathcal{X}$-structure merely from constructing this initial seed $(Q(\Gamma), X_\Gamma)$. Moreover, in general there could be many choices of Γ for a fixed general weave $\mathfrak{w}$, and it is not known whether different choices yields equivalent, or even quasi-equivalent, cluster seeds. It thus becomes crucial to construct cluster $\mathcal{A}$-variables for $\mathfrak{M}(\Lambda, T)$, ideally in a symplectic invariant manner, as we will momentarily do. By [BFZ05], a cluster $\mathcal{A}$-structure can be shown to exist, once the necessary properties of the candidate $\mathcal{A}$-variables are proven. As a byproduct, Corollary 1.2 then deduces the existence of the cluster $\mathcal{X}$-structure on $\mathcal{M}_1(\Lambda)$ where the variables are microlocal monodromies.

4.5. **Collections of sign curves: fixing signs.** Let $\mathbb{G}$ be a GP-graph, $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$ its initial weave and $L := L(\mathfrak{w})$ its initial filling, and T a set of marked points in $\Lambda(\mathbb{G}) = \partial L$. Let us denote the set of lifts of trivalent weave vertices on L by $P \subset L$. It follows from Subsections 4.3 and 4.4 that each point of the flag moduli $\mathcal{M}_1(\mathfrak{w})$ defines a rank 1 local system on $L \setminus P$ with -1 monodromy around each point in P. In this section, we describe a way to add signs to monodromies to obtain a (non-canonical) isomorphism between $\mathcal{M}_1(\mathfrak{w})$ and $\mathrm{Loc}_1(L)$. This is a combinatorial expression of the fact that, in our case, global sections of the Kashiwara-Schapira stack are (canonically) isomorphic to the category of twisted local

systems and (non-canonically) also isomorphic to the category of local systems. In terms of weave combinatorics, we proceed as follows:

Definition 4.14. A *sign curve* is an unoriented curve on the weave surface L that intersects the lifts of weave lines transversely and whose endpoints lie in the set $P \sqcup T$.

We record sign curves on L by drawing dotted curves on $\mathbb{R}^2$ in juxtaposition with the weave $\mathfrak{w}$ and labeling the indices of the sheets they are on.

Definition 4.15. A collection C of sign curves on L is *coherent* if each point in P is incident to one and only one sign curve in C, and all curves in C intersect transversely.

Example 4.16. *Here are three examples of a coherent set of sign curves in a Lagrangian annulus:*
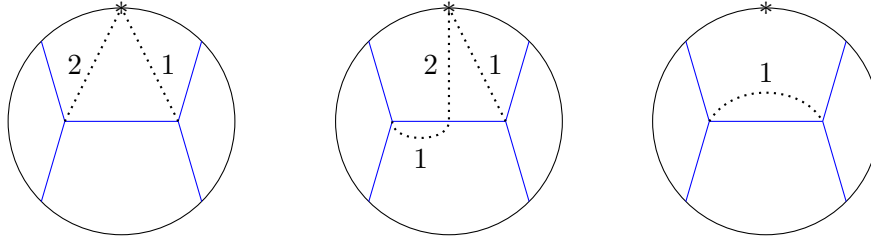
FIGURE 73. Three examples of coherent sets of sign curves. The $*$ at the boundary of the disk is lifted to two marked points on Λ. The left picture and the middle picture are compatible (Definition 4.17), the right picture is not.

Fix a coherent set C of sign curves on L. For any path γ on L, we may perturb γ so that it intersects elements of C transversely. Then, we redefine the parallel transport along γ to be the microlocal parallel transport ψ_γ multiplied by a factor of -1 whenever the curve γ passes through a sign curve in C. Since each branch point of L is incident to one and only one sign curve, this new parallel transport corrects the monodromy around each point in P to be 1, defining an isomorphism

$$\Phi_C : \mathcal{M}_1(\mathfrak{w}) \longrightarrow \mathrm{Loc}_1(L) \cong H^1(L; \mathbb{C}^\times) \cong (\mathbb{C}^\times)^{b_1(L)}.$$

In fact, we can construct better than an arbitrary isomorphism $\mathcal{M}_1(\mathfrak{w}) \xrightarrow{\cong} \mathrm{Loc}_1(L)$. From Subsection 4.4, our candidates for cluster $\mathcal{X}$-variables are of the form $-\psi_\gamma$ for initial absolute cycles $\gamma \in \mathfrak{S}(\mathbb{G})$, and we can in fact incorporate this extra sign in front of ψ_γ into the set of coherent sign curves.

Definition 4.17. A coherent set C of sign curves on L is said to be *compatible* if

$$\Phi_C(p)(\gamma) = X_\gamma(p) := -\psi_\gamma(p), \quad \forall p \in \mathcal{M}_1(\mathfrak{w}),$$

and for all initial absolute cycles $\gamma \in \mathfrak{S}(\mathbb{G})$.

For the initial free weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$ constructed from a GP-graph $\mathbb{G}$, we can find a compatible set of sign curves as follows. First, we observe that all trivalent weave vertices of $\mathfrak{w}$ occur near the boundary of the weave. Thus, at each trivalent weave vertex, two of the three adjacent sectors are facing away from the weave: we will draw our sign curves inside these two sectors. Next, we break the weave $\mathfrak{w}$ down into weave columns, and by Section 3 trivalent weave vertices only occur inside Type 2 columns.

Let us further classify Type 2 columns into two types: a Type 2 column is said to be *critical* if it is the rightmost Type 2 column that contains part of an initial cycle $\gamma \in \mathfrak{S}(\mathbb{G})$; it is said to be *non-critical* otherwise. Note that by construction, each critical Type 2 column has a unique initial cycle γ that ends there.

If a Type 2 column is non-critical, we draw a sign curve in either sector on either sheet, and then lead it towards the boundary of L; once it gets within a collar neighborhood of the boundary $\partial L = \Lambda$, the sign curve will follow along Λ until it reaches a marked point. (Such a marked point exists because we have at least one marked point per link component.)

If a Type 2 column is critical, we consider the unique initial cycle γ that ends at this Type 2 column. We compute the product of all the signs γ has picked up along all the previous trivalent weave vertices. If the product is 1, we add a sign curve c on the appropriate sheet of either of the two sectors so that γ intersects with c non-trivially. If the product is -1, we add a sign curve c on the appropriate sheet of either of the two sectors so that γ intersects with c trivially. By doing so, we guarantee that Φ_C maps γ to $X_\gamma = -\psi_\gamma$, as desired.

Thus, a compatible set C of sign curves exists in our setting, and we can explicitly identify $\mathcal{M}_1(\mathfrak{w})$ with the moduli space $\mathrm{Loc}_1(L)$ of rank 1 local systems on L. This identification allows us to interpret the cluster $\mathcal{X}$-variables X_γ as actual monodromies of local systems along the initial cycles γ, and also fixing the necessary signs for the upcoming constructions.

Example 4.18. *Consider the GP-graph $\mathbb{G}$ in Figure 74 (left), whose GP-link is a Legendrian trefoil. The weave $\mathfrak{w}(\mathbb{G})$ is drawn in Figure 74 (right), together with a compatible set of sign curves obtained from the recipe described above. Note that among the three Type 2 columns, the first one (counting from the left) is non-critical, but the last two are both critical.*
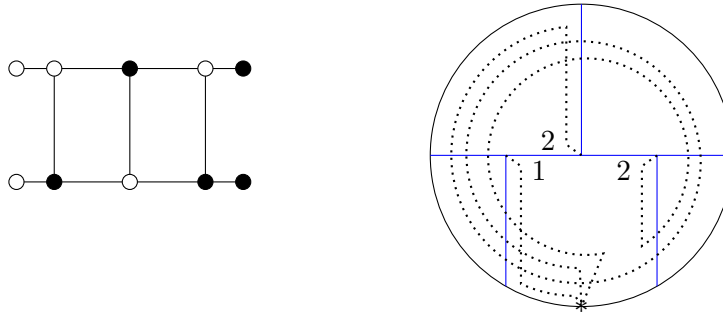
FIGURE 74. A GP-link with a compatible set of sign curves. The marked point is located on the lower sheet (sheet 1).

Proposition 4.19. *If $\mathfrak{w}$ admits a compatible set of sign curves and $\mathfrak{w}'$ is mutation equivalent to $\mathfrak{w}$, then $\mathfrak{w}'$ also admits a (non-canonical) compatible set of sign curves.*

Proof. Since weave equivalence and weave mutation can both be done locally, it suffices to verify that given a compatible set of sign curves locally before any weave equivalence or weave mutation, we can find a compatible set of sign curves locally after the weave equivalence or weave mutation as well. In particular, it suffices to check the local pictures in Figure 22 and the left picture in Figure 26. Below are two such cases. We leave it as an exercise for the readers to check all remaining cases. $\qquad\square$
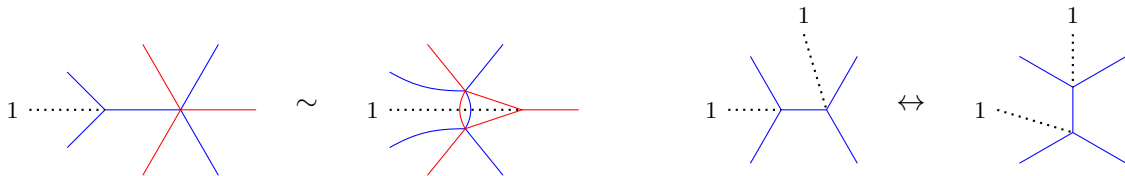
FIGURE 75. Existence of compatible sets of sign curves before and after weave equivalences (Left) and weave mutations (Right).

4.6. **Microlocal merodromies: candidate cluster $\mathcal{A}$-variables.** This subsection addresses the construction of what shall become the cluster $\mathcal{A}$-variables on the moduli space $\mathfrak{M}(\Lambda, T)$ for a GP-link $\Lambda = \Lambda(\mathbb{G})$. In the previous section, we explained that cluster $\mathcal{X}$-variables were indexed by certain *absolute* cycles $\gamma \in H_1(L)$ in the Lagrangian filling $L = L(\mathfrak{w}(\mathbb{G}))$ and X_γ was a natural rational function with a symplectic origin: the microlocal monodromy along γ of the sheaf associated with the weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$.

Now, the new idea is that cluster $\mathcal{A}$-variables $\{A_\eta\}$ will be indexed by certain *relative* cycles $\eta \in H_1(L \setminus T, \Lambda \setminus T)$ and the functions $A_\eta : \mathfrak{M}(\Lambda, T) \longrightarrow \mathbb{C}$ will be defined by what we call the *microlocal merodromy* along η. Intuitively, this merodromy along η is constructed as a microlocal parallel transport along η. Here are the details.

4.6.1. *Microlocal merodromy.* Let $\mathbb{G}$ be an GP-graph, $\Lambda = \Lambda(\mathbb{G})$ be its GP-link, and T be a collection of marked points on Λ with at least one marked point per link component. Fix a compatible set C of sign curves and let $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$ be the initial weave. The flag moduli $\mathfrak{M}(\mathfrak{w}, T)$ is an open subset of the moduli space $\mathfrak{M}(\Lambda, T)$, and every point in this open subset defines a local system, via the identification $\Phi_C : \mathcal{M}_1(\mathfrak{w}) \xrightarrow{\cong} \mathrm{Loc}_1(L)$ in Subsection 4.5, together with a framing (trivialization) of the rank-1 local system Φ on the connected components of $\Lambda \setminus T = (\partial L) \setminus T$.

The framing data defines a special vector $\phi_x \in \Phi_x$ at any point $x \in \Lambda \setminus T$. Given an oriented curve $\eta \subset L$ with both the source point $s = \partial_-\eta$ and the target point $t = \partial_+\eta$ contained inside $\Lambda \setminus T$, we can parallel transport ϕ_s from the source s to the target t along η, obtaining a non-zero vector in $\eta(\phi_s) \in \Phi_t$. The ratio $\frac{\eta(\phi_s)}{\phi_t}$ is a non-zero number A_η, which defines a $\mathbb{C}^\times$-valued function on $\mathfrak{M}(\mathfrak{w}, T)$. This can be naturally generalized to relative 1-cycles $\eta \in H_1(L \setminus T, \Lambda \setminus T)$.

Definition 4.20. The function $A_\eta : \mathfrak{M}(\mathfrak{w}, T) \longrightarrow \mathbb{C}^\times$ is said to be the *microlocal merodromy* along the oriented curve η.

Since $\mathfrak{M}(\mathfrak{w}, T)$ is an open subset of $\mathfrak{M}(\Lambda, T)$, A_η can also be viewed as a rational function on $\mathfrak{M}(\Lambda, T)$. Note that, a priori, A_η might not be extend to a regular function on $\mathfrak{M}(\Lambda, T)$. We emphasize that the decoration in $\mathfrak{M}(\Lambda, T)$ are needed in order to define A_η, and thus microlocal merodromies cannot be defined in $\mathcal{M}_1(\Lambda)$.

Remark 4.21. In Greek, *mono* essentially translates to *one*, and thus *monodromy* can be translated to "a path that goes around once". Similarly, *mero* in Greek can be understood as *a partial, a fraction*, and thus *merodromy*, translating to "a path that goes around only partially", seems a natural nomenclature for a relative cycle. The use of the adjective *microlocal* is also fitting since the construction, which factors through the Kashiwara-Schapira stack, involves the use of the μhom functor. $\square$

The microlocal merodromies associated to relative cycles coming from marked points are non-vanishing. Indeed, for each marked point $t \in T$, we pick a small half-disk neighborhood U_t of t, as drawn in Figure 76, and define $\xi_t := \partial U_t$. By definition, we have that

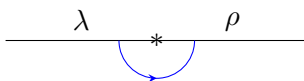$$A_t := A_{\xi_t} = \frac{\xi_t(\lambda)}{\rho} \neq 0.$$

FIGURE 76. A marked point $t \in T$, with decorations λ and ρ to the left and right, and the boundary of a half-disk neighborhood U_t.

In particular, by using the ratio $\frac{\xi_t(\lambda)}{\rho}$, we can extend A_t to a global invertible function on the entire moduli space $\mathfrak{M}(\Lambda, T)$. (This property does not in general hold for A_η if η is an arbitrary relative cycle in $H_1(L \setminus T, \Lambda \setminus T)$.) Now consider the following exact sequence of lattices:

$$0 \longrightarrow \mathbb{Z} \overset{i}{\longrightarrow} \bigoplus_{t \in T} \mathbb{Z}\xi_t \longrightarrow H_1(L \setminus T, \Lambda \setminus T) \overset{\pi}{\longrightarrow} H_1(L, \Lambda) \longrightarrow 0,$$

where $i(1) := \sum_{t \in T} \xi_t$. This exact sequence implies the following two corollaries.

Corollary 4.22. $\prod_{t \in T} A_t = 1$ and hence A_t is a unit in $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ for every $t \in T$.

Proof. It follows from the fact that $\sum_{t \in T} \xi_t = 0$ in $H_1(L \setminus T, \Lambda \setminus T)$. $\qquad\square$

Corollary 4.23. If $\eta_1, \eta_2 \in H_1(\Sigma \setminus T, \Lambda \setminus T)$ satisfy $\pi(\eta_1) = \pi(\eta_2)$, then A_{η_1} and A_{η_2} are related to each other by a Laurent monomial in the variables A_t, $t \in T$.

Proof. It follows from the fact that $\ker \pi = \operatorname{Span}\{\xi_t \mid t \in T\}$. $\qquad\square$

The later corollary starts to hint at the quasi-cluster equivalence that appears if different basis completions in $H_1(\Sigma \setminus T, \Lambda \setminus T)$ are chosen, as the former corollary indeed hints at the fact that A_t, $t \in T$ are a type of frozen variables.

4.6.2. *Crossing values.* The next aim is to compute A_η for curves whose support is transverse to a weave $\mathfrak{w}$. Consider the following commutative diagram of vector space inclusions

$$\begin{array}{ccc}
 & V_n & \\
 \nearrow & & \nwarrow \\
V_w & & V_e \\
 \nwarrow & & \nearrow \\
 & V_s &
\end{array}$$

and assume that $0 \to V_s \to V_w \oplus V_e \to V_n \to 0$ is exact. Let $\alpha_s, \alpha_w, \alpha_e, \alpha_n$ be non-zero top-dimensional (volume) forms in V_s, V_w, V_e, V_n, respectively. Then we write $\alpha_w = \beta_w \wedge \alpha_s$ and $\alpha_e = \beta_e \wedge \alpha_s$ for some forms β_w and β_e.

Definition 4.24. In the context of a diagram as above, we define

$$\alpha_w \overset{\alpha_s}{\wedge} \alpha_e := \beta_w \wedge \beta_e \wedge \alpha_s.$$

The top form $\alpha_w \overset{\alpha_s}{\wedge} \alpha_e$ is a non-zero on V_n and does not depend on the choice of β_w or β_e. By definition, the ratio $\frac{\alpha_w \overset{\alpha_s}{\wedge} \alpha_e}{\alpha_n}$ is said to be the *crossing value* of the quadruple of top forms $\alpha_s, \alpha_w, \alpha_e, \alpha_n$.

Let us describe how to use crossing values to compute merodromies along planar relative cycles. Given a flag $\mathscr{F} = (0 \subset \mathscr{F}_1 \subset \mathscr{F}_2 \subset \cdots \subset \mathscr{F}_n = \mathbb{C}^n)$ with framing $\phi_i \neq 0 \in \mathscr{F}_i/\mathscr{F}_{i-1}$ for all $i \in [1, n]$, we can construct a top form $\alpha_i \in \bigwedge^i \mathscr{F}_i$ by first lifting each ϕ_i to a vector in $\tilde{\phi}_i \in \mathscr{F}_i$ and then taking the wedge

$$\alpha_i := \tilde{\phi}_i \wedge \tilde{\phi}_{i-1} \wedge \cdots \wedge \tilde{\phi}_1.$$

Note that each form α_i is independent of the choice of lifts.

Definition 4.25. Given a flag $\mathscr{F}$, such a collection $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_n)$ of forms is said to be a *decoration* on the flag $\mathscr{F}$. A flag with a decoration is referred to as a *decorated flag*.

Note that we can reverse the construction above and recover a framing from a decoration. Thus, framings $(\phi_1, \ldots, \phi_n)$, and decorations $(\alpha_1, \alpha_2, \ldots, \alpha_n)$ of a flag $\mathscr{F}$ are equivalent in the sense that there is a natural construction of one from the data of the other.

Now suppose $(\mathscr{L}, \lambda)$ and $(\mathscr{R}, \rho)$ are two framed flags such that $\mathscr{L} \overset{s_i}{\sim} \mathscr{R}$. Let α and β be the (associated) decoration on the framed flags $(\mathscr{L}, \lambda)$ and $(\mathscr{R}, \rho)$, respectively. By construction, we have

(4.8)
$$\begin{cases} \alpha_j = \beta_j & \text{if } j < i, \\ \alpha_j = -\beta_j & \text{if } j > i. \end{cases}$$

Since $\mathscr{L}_{i\pm 1} = \mathscr{R}_{i\pm 1}$, we can also define the linear isomorphism

$$\psi : \frac{\mathscr{L}_i}{\mathscr{L}_{i-1}} \hookrightarrow \frac{\mathscr{L}_{i+1}}{\mathscr{L}_{i-1}} = \frac{\mathscr{R}_{i+1}}{\mathscr{R}_{i-1}} \twoheadrightarrow \frac{\mathscr{R}_{i+1}}{\mathscr{R}_i}.$$

The image under ψ of a framing is readily computed in terms of decorations as follows:

Lemma 4.26. $\psi(\lambda_i) = \dfrac{\alpha_i \overset{\alpha_{i-1}}{\wedge} \beta_i}{\beta_{i+1}} \rho_{i+1}.$

Proof. Let $\tilde{\lambda}_i$ and $\tilde{\rho}_{i+1}$ be lifts of λ_i and ρ_{i+1}. Then by definition, $\alpha_i = \tilde{\lambda}_i \wedge \alpha_{i-1}$ and $\beta_{i+1} = \tilde{\rho}_{i+1} \wedge \beta_i$. By construction,

$$\frac{\alpha_i \overset{\alpha_{i-1}}{\wedge} \beta_i}{\beta_{i+1}} = \frac{\tilde{\lambda}_i \wedge \beta_i}{\tilde{\rho}_{i+1} \wedge \beta_i} = \frac{\psi(\lambda_i)}{\rho_{i+1}}. \qquad \square$$

Now suppose $\eta \subset \Sigma(\mathfrak{w})$ is a lift of a planar curve in $\mathbb{R}^2$ to the weave front. Then, it defines a partial cross-section of the weave surface, where η passes through a collection of (framed) flags $\mathscr{L} = \mathscr{F}_0, \mathscr{F}_1, \mathscr{F}_2, \ldots, \mathscr{F}_l = \mathscr{R}$. For each flag $\mathscr{F}_i$, $0 < i < l$, we choose a sequence of top forms $\alpha_i = (\alpha_{i,j})$. Since the parallel transport along η consists of compositions of linear isomorphisms like the map ψ in Lemma 4.26, or its inverse, Lemma 4.26 allows us to compute A_η.

Example 4.27. *Consider the cross-section of a weave surface depicted in Figure 77, and let η be the blue relative cycle. The sequences of top forms α and δ are determined by the decoration λ and ρ, respectively, but the sequences of top forms β and γ are chosen arbitrary. By Lemma 4.26, the microlocal merodromy along η can be obtained as*

$$A_\eta = \frac{\eta(\lambda_2)}{\rho_3} = \frac{\alpha_2}{\beta_1 \wedge \alpha_1} \frac{\beta_1 \wedge \gamma_1}{\gamma_2} \frac{\gamma_2 \overset{\gamma_1}{\wedge} \delta_2}{\delta_3} = \frac{\alpha_2}{\beta_1 \wedge \alpha_1} \frac{\beta_1 \wedge \delta_2}{\delta_3}.$$

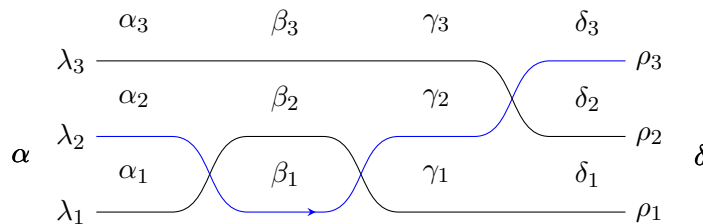FIGURE 77. Computation of a merodromy

The right-hand side this expression shows that A_η is invariant under any non-zero rescaling of β_1. In general, the study of microlocal merodromies involves understanding properties (such as regularity) of products of crossing values and inverses thereof.

Example 4.28. *For consistency, let us use crossing values to verify that the introduction of sign curves resolve the −1 microlocal monodromy problem around trivalent weave vertices. For that, we revisit the unique filling of the max tb unknot (Figure 72), and let ξ be the double lift of a loop that goes around the trivalent weave vertex counterclockwise. A cross section of a circle around the trivalent vertex, being traveled twice, is depicted in Figure 78.*
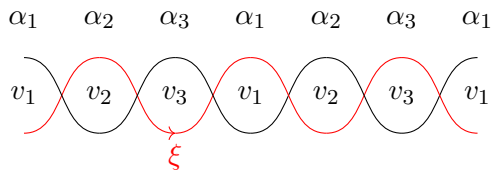
FIGURE 78. A cross-section of a weave, where we can use crossing values to verify that the signs introduced in Subsection 4.5 indeed yield a local system.

Without loss of generality let us suppose it starts at the lower sheet in the sector containing l_1. Let us fix decorations (v_i, α_i) on each flag. Without any sign curves, the microlocal parallel transport is

$$\frac{\psi_\xi(v_1)}{v_1} = \frac{v_1 \wedge v_2}{\alpha_2} \frac{\alpha_2}{v_3 \wedge v_2} \frac{v_3 \wedge v_1}{\alpha_1} \frac{\alpha_1}{v_2 \wedge v_1} \frac{v_2 \wedge v_3}{\alpha_3} \frac{\alpha_3}{v_1 \wedge v_3} = -1.$$

By introducing a sign curve as in[19] Figure 79 (left), we are adding in an extra factor of −1 to any parallel transport that crosses the sign curve. If we use crossing values to compute microlocal parallel transportation, all decorations (forms) associated to regions above the sign curve will need to be multiplied by −1, as we pass through a sign curve, as illustrated in Figure Figure 79 (right).
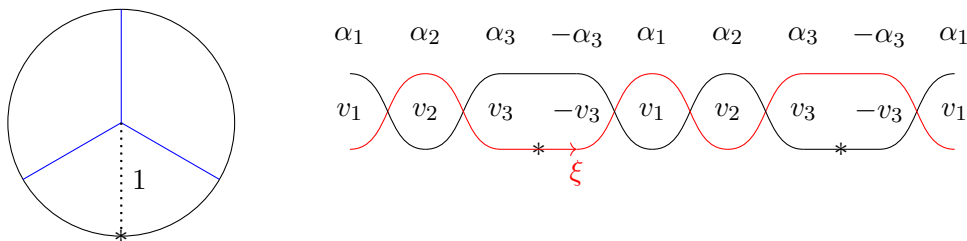
FIGURE 79. Sign curve correction to the microlocal monodromy (left) and the boundary cross-section for the computation of the corrected microlocal parallel transport (right).

The corrected microlocal parallel transport along ξ computed by crossing values is

$$\frac{\psi_\xi(v_1)}{v_1} = \frac{v_1 \wedge v_2}{\alpha_2} \frac{\alpha_2}{v_3 \wedge v_2} \frac{-v_3 \wedge v_1}{\alpha_1} \frac{\alpha_1}{v_2 \wedge v_1} \frac{v_2 \wedge v_3}{\alpha_3} \frac{-\alpha_3}{v_1 \wedge (-v_3)} = 1.$$

In general, a microlocal merodromy A_η will be expressed in terms of ratios of crossing values, and is only a *rational* function. Nevertheless, certain choices of η within the initial weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$ yield *regular* funtions. Indeed, let us consider the following special family of merodromies. By Subsection 3.6, each face f of $\mathbb{G}$ has an associated naive relative cycle η_f in $H_1(L \setminus T, \Lambda \setminus T)$. Let $A_f := A_{\eta_f}$ be the microlocal merodromy of this naive relative cycle. Since $\{\partial f\}$ is a basis of $H_1(L)$, it follows from Poincaré duality that $\{\pi(\eta_f)\}$ is a basis of

---

[19]Here we chose to place the marked point on the lower sheet, but placing it on the upper sheet would work as well.

$H_1(L, \Lambda)$, where $\pi : H_1(L \setminus T, \Lambda \setminus T) \to H_1(L, \Lambda)$ is the natural projection map. By Corollary 4.23, we conclude that different choices of η_f only change A_f by a multiple of units.

Proposition 4.29. *Let $\mathbb{G}$ be GP-graph and $f \in \mathbb{G}$ a face. Then the microlocal merodromy $A_f : \mathfrak{M}(\Lambda(\mathbb{G}), T) \longrightarrow \mathbb{C}$ is a regular function.*

Proof. Let us consider the parallel transport along ξ_i, whose support is the connected component in a cross-section of L that goes from the ith level on the left to the $(n-i)$th level on the right. For example, below is such a cross-section picture when $n = 4$.
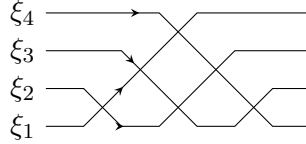
By using the crossing value technique, we see that

$$A_{\xi_i} = \frac{\alpha_i}{\alpha_{i-1} \wedge \gamma} \frac{\gamma \wedge \beta_{n-i}}{\beta_{n-i+1}},$$

where γ is a non-zero vector in the line $\mathscr{L}_i \cap \mathscr{R}_{n-i+1}$. By the definition of crossing value, the right hand side of this expression is also equal to

$$\frac{\alpha_i \wedge \beta_{n-i}}{\alpha_{i-1} \wedge \beta_{n-i+1}}$$

By construction, if f is the ith gap counting from the bottom, which is left in the cross-section picture above, then η_f can be chosen to be $\sum_{k=1}^{i} \xi_k$. As a result,

$$A_f = \prod_{k=1}^{i} A_{\xi_i} = \frac{\alpha_i \wedge \beta_{n-i}}{\beta_n}.$$

Since $\beta_n \neq 0$ by definition, we conclude that A_f is a regular function for each face f. $\quad\square$

Remark 4.30. Microlocal merodromies can also be used to define the frozen cluster $\mathcal{X}$-variables associated to the relative cycles in $H_1(L, T)$ that are not in the image of $H_1(L)$. In the moduli space $\mathcal{M}_1(\Lambda, T)$, the microlocal merodromy allows one to compare framings at the endpoints of the relative cycles, which are marked points T where the (stalk of the microlocal) local system has been trivialized.

4.7. **Vanishing of microlocal merodromies and flag relative positions.** In this section we study the vanishing loci of the microlocal merodromies $A_f : \mathfrak{M}(\Lambda(\mathbb{G}), T) \longrightarrow \mathbb{C}$ associated to faces $f \subset \mathbb{G}$ of a GP-graph $\mathbb{G}$. The key technical result, Proposition 4.34, relates the vanishing loci of microlocal merodromies associated to different faces of the graph $\mathbb{G}$. This result is crucial to deduce the necessary properties of these candidate cluster $\mathcal{A}$-variables, such as regularity, and conclude Theorem 1.1.

Thus far, we have parametrized the relative position of a pair of flags in $\mathbb{C}^n$ by the symmetric group S_n. This relative position is invariant under the diagonal GL_n action, and hence is also in bijection with GL_n-orbits in $\mathcal{B}(n) \times \mathcal{B}(n)$. The inclusion relation on closures of these orbits defines a partial order called the *Bruhat order* on S_n, i.e. $u \leq v$ if $\mathcal{O}_u \subset \overline{\mathcal{O}}_v$. Combinatorially, the Bruhat order can be computed through set comparison.

Definition 4.31. For two equal-sized subsets $I = \{i_1 < i_2 < \cdots < i_m\}$ and $J = \{j_1 < j_2 < \cdots < j_m\}$ of $\{1, \ldots, n\}$, we define $I \leq J$ if $i_k \leq j_k$ for all $1 \leq k \leq m$. By definition, for two permutations u and v of S_n, $u \leq v$ in the *Bruhat order* if and only if $\{u(1), \ldots, u(m)\} \leq \{v(1), \ldots, v(m)\}$ for all $1 \leq m < n$.

By Subsection 4.1, the moduli $\mathfrak{M}(\Lambda(\mathbb{G}), T)$ can be understood in terms of tuples of flags, with maps between them and incidence constraints. The flags can be read directly from the front $\mathfrak{G}$. In particular, for a Type 1 column of $\mathbb{G}$, there exists a unique pair of (decorated) flags $\mathscr{F}_0$ and $\mathscr{F}^0$, see Figure 69. In the points of the open torus chart $\mathfrak{M}(\mathfrak{w}, T) \subset \mathfrak{M}(\Lambda(\mathbb{G}), T)$, these two flags $\mathscr{F}_0$ and $\mathscr{F}^0$ and in w_0-relative position, but in general the relative position between $\mathscr{F}_0$ and $\mathscr{F}_0$, at another point of $\mathfrak{M}(\Lambda(\mathbb{G}), T)$ might vary. The dependence of $\mathscr{F}_0, \mathscr{F}_0$ on the point $p \in \mathfrak{M}(\Lambda(\mathbb{G}), T)$ will be denoted by $\mathscr{F}_0(p), \mathscr{F}_0(p)$.

By Proposition 4.29 the microlocal merodromy A_i associated to the ith gap of a Type 1 column, counting from below in the GP-graph, is a regular function on $\mathfrak{M}(\Lambda, T)$. Moreover, it can be expressed as $\frac{\alpha_i \wedge \beta_{n-i}}{\beta_n}$, up to a multiple of units, where α and β are decorations on the pair of flags $\mathscr{F}_0$ and $\mathscr{F}^0$ placed at the bottom and the top of that Type 1 column, respectively. In particular, the restriction of $A_i|_{\mathfrak{M}(\mathfrak{w}, T)}$ to the open torus chart $\mathfrak{M}(\mathfrak{w}, T) \subset \mathfrak{M}(\Lambda(\mathbb{G}), T)$ is a non-vanishing function. The following lemma shows that we can describe the vanishing locus of this microlocal merodromy $A_i : \mathfrak{M}(\Lambda(\mathbb{G}), T) \longrightarrow \mathbb{C}$ in terms of the relative position between the two flags $\mathscr{F}_0$ and $\mathscr{F}^0$:

Lemma 4.32. *Let $\mathbb{G}$ be a GP-graph, $(\mathscr{F}_0, \alpha), (\mathscr{F}^0, \beta)$ the pair of decorated flags associated with a Type 1 column C and A_i the ith microlocal merodromy associated to C, $i \in [1, n]$. Consider a point $p \in \mathfrak{M}(\Lambda(\mathbb{G}), T)$ and the permutation $w \in S_n$ such that $\mathscr{F}_0(p) \overset{w}{\sim} \mathscr{F}^0(p)$. Then $A_i(p) = 0$ if and only if $w \leq s_i w_0$ in the Bruhat order.*

Proof. Without loss of generality, we may assume that the decoration α is proportional to $(e_{w(1)}, e_{w(1)} \wedge e_{w(2)}, \ldots, e_{w(1)} \wedge e_{w(2)} \wedge \cdots e_{w(n)})$ and the decoration β is proportional to $(e_1, e_1 \wedge e_2, \ldots, e_1 \wedge e_2 \wedge \cdots e_n)$. From this we know that $A_i = 0$ if and only if

$$e_{w(1)} \wedge e_{w(2)} \wedge \cdots \wedge e_{w(i)} \wedge e_1 \wedge e_2 \wedge \cdots \wedge e_{n-i} = 0,$$

which is equivalent to saying that the following intersection is non-empty:

$$\{w(1), w(2), \ldots, w(i)\} \cap \{1, 2, \ldots, n - i\} \neq \emptyset.$$

Now note that for the permutation $v = s_i w_0$, all $\{v(1), \ldots, v(m)\}$ are maximal sets with respect to the linear order on $\{1, \ldots, n\}$ except when $m = i$, where

$$\{v(1), \ldots, v(i)\} = \{n, n - 1, \ldots, n - i + 2, n - i\}.$$

If $w \leq v$, then $\{w(1), w(2), \ldots, w(i)\} \leq \{n, n - 1, \ldots, n - i + 2, n - i\}$. This implies that among $w(1), w(2), \ldots, w(i)$, some index no greater than $n - i$ must have appeared. Therefore we have $\{w(1), w(2), \ldots, w(i)\} \cap \{1, 2, \ldots, n - i\} \neq \emptyset$ and hence $A_i = 0$.

Conversely, if $w \not\leq v$, then we must have

$$\{w(1), w(2), \ldots, w(i)\} = \{n, n - 1, \ldots, n - i + 1\},$$

which implies that $\{w(1), w(2), \ldots, w(i)\} \cap \{1, 2, \ldots, n - i\} = \emptyset$ and hence $A_i \neq 0$. $\qquad\square$

Lemma 4.32 shows that in order to study whether the microlocal merodromies A_f vanish or not, it suffices to consider the relative position between the pair of flags in a Type 1 column that contains part of the face f. We use the following simple lemma in the proof of Proposition 4.34, through Lemma 4.35:

Lemma 4.33. *Suppose $\mathscr{F} \overset{u}{\sim} \mathscr{F}' \overset{v}{\sim} \mathscr{F}''$ with $u, v \in S_n$. Let l denote the length function on S_n. For any $w \in S_n$, let $\underline{w}$ denote the positive braid represented by a (equivalently any) reduced word of w.[20]*

(1) If $l(uv) = l(u) + l(v)$, then $\mathscr{F} \overset{uv}{\sim} \mathscr{F}''$.

(2) In general, if $\mathscr{F} \overset{w}{\sim} \mathscr{F}'$, then $w \leq \mathrm{Dem}(\underline{u}\ \underline{v})$ in the Bruhat order.

---

[20] As before, Dem denotes the Demazure product on positive braids.

Proof. (1) Follows from the standard fact that, for Bruhat cells, if $l(uv) = l(u) + l(v)$ then $(BuB)(BvB) = BuvB$. For (2), at the level of Bruhat cells we know that $(Bs_iB)(Bs_iB) = Bs_iB \sqcup B$; this is equivalent to saying that if $\mathscr{F} \overset{s_i}{\sim} \mathscr{F}' \overset{s_i}{\sim} \mathscr{F}''$, then there are two possible relative position relations between $\mathscr{F}$ and $\mathscr{F}''$: either $\mathscr{F} \overset{s_i}{\sim} \mathscr{F}''$ or $\mathscr{F} = \mathscr{F}''$. For both cases, the relative position is at most $\mathrm{Dem}(\underline{s_i}\ \underline{s_i}) = s_i$. The general statement follows from the well-definedness of the Demazure product. $\qquad\square$

For notation convenience, for $1 \leq i \leq j < n$, let us define the following permutation:

$$
w_{[i,j]} := \quad
\begin{array}{c}
s_{n-1} \\
s_{n-2} \quad s_{n-2} \\
\reflectbox{$\ddots$} \qquad\qquad \ddots \\
s_2 \quad s_2 \quad s_2 \quad \cdots \quad s_2 \\
s_1 \quad s_1 \quad s_1 \qquad\qquad s_1 \\
\underbrace{\qquad\qquad\qquad}_{} \\
\text{delete all } s_1\text{'s from the } i\text{th to } j\text{th copy}
\end{array}
$$

and denote

$$\overline{w} := w_0 w^{-1} w_0.$$

In this notation, $s_i w_0 = w_{[i,i]} = \overline{w}_{[i,i]}$. In terms of set comparison, $u \leq w_{[i,j]}$ in the Bruhat order if and only if for all $n - j \leq l \leq n - i$,

(4.9)
$$\{u(1), u(2), \ldots, u(l)\} \leq \{i \leq \cdots\},$$

where $\{i \leq \cdots\}$ means the set of the appropriate size (say of size l) consisting of the greatest $l - 1$ elements in $\{1, \ldots, n\}$ together with the element i.

Finally, the core of this subsection is the following result, which states that we can use lollipop reactions (Definition 2.22) to keep track of the relative position conditions on flags and, in turn, understand vanishing conditions for microlocal merodromies associated to faces. The goal of the remaining part of this subsection is to prove:

Proposition 4.34. *Let $\mathbb{G}$ be a GP-graph and $f, g \subset \mathbb{G}$ two faces. Suppose that g is selected in a lollipop reaction initiated from a lollipop in f. Then $A_f = 0$ implies $A_g = 0$.*

As discussed in Subsection 2.6, the scanning wall in in a lollipop reaction moves to the right if the lollipop is white and to the left if the lollipop is black. By Lemma 4.32, at the starting point, $A_f = 0$ implies that the flags at the two ends of the wall are at most $w_{[i,i]} = \overline{w}_{[i,i]}$ apart, where i is the index of the gap between the two adjacent horizontal lines, counting from below in the GP-graph. This is schematically illustrated in the following picture:

The heart of the argument is proving that in the case of a white (resp. black) lollipop reaction, as the wall scans to the right (resp. left), the flags at the two ends of the wall will be at most $w_{[i,j]}$ (resp. $\overline{w}_{[i,j]}$) apart, where $[i,j]$ is the interval containing the indices of the gaps that the wall crosses (counting from below in the GP-graph). Due to symmetry, we will only prove Proposition 4.34 for white lollipop reactions; the proof for the case of black lollipop reactions is completely symmetric. Let us start with the following lemma:

Lemma 4.35. *Let $\mathbb{G}$ be a GP-graph and consider a Type 3 column with a black lollipop (we shift the indices on the right to match the Coxeter generators), as follows:*
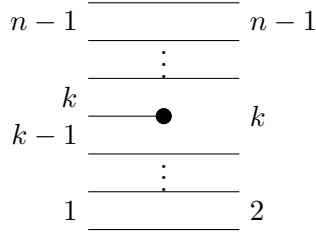
Let $(\mathscr{L}_0, \mathscr{L}^0)$ be the pair of flags to the left and $(\mathscr{R}_0, \mathscr{R}^0)$ be the pair of flags to the right. Suppose that $\mathscr{L}_0 \overset{u}{\sim} \mathscr{L}^0$ and $\mathscr{R}_0 \overset{v}{\sim} \mathscr{R}^0$. (Hence $h^{-1}(\mathscr{R}_0) \overset{v}{\sim} h^{-1}(\mathscr{R}^0)$.) Then, in the Bruhat order we have that

$$u \geq s_{k-1}s_{k-2} \cdots s_1 v s_1 s_2 \cdots s_{n-k}.$$

Proof. By construction, $h^{-1}(\mathscr{R}_0)$ and $h^{-1}(\mathscr{R}^0)$ share the same 1-dimensional subspace. Therefore $v(1) = 1$, which implies that $v s_1 s_2 \cdots s_{n-k}$ is reduced. Lemma 4.33 (1) implies that

$$\mathscr{L}_0 \overset{s_{k-1} \cdots s_1}{\sim} h^{-1}(\mathscr{R}_0) \overset{v s_1 \cdots s_{n-k}}{\sim} \mathscr{L}^0,$$

where the first relative position is given by the Type 3 column requirement.

Let us record a permutation $w \in S_n$ as an n-tuple $(w(1), w(2), \ldots, w(n))$. Since $v(1) = 1$, we may assume that $v = (1, v(2), v(3), \ldots, v(n))$. Then

$$v s_1 s_2 \cdots s_{n-k} = (v(2), v(3), \ldots, v(n-k+1), 1, v(n-k+2), \ldots, v(n)).$$

Note that left multiplication by s_i interchanges the entries i and $i+1$. From (the proof of) Lemma 4.33 (2), we know that when multiplying s_i on the left of a permutation w, there is only one possible relative position $s_i w$ if i is on the left of $i+1$, and there can be two possible relative positions w and $s_i w$ if i is on the right of $i+1$, in which case $s_i w < w$. Thus, performing all the left multiplications $s_{k-1} \cdots s_1$ to $v s_1 \cdots s_{n-k}$ yield the smallest relative position relation, and hence $u \geq s_{k-1} \cdots s_1 v s_1 \cdots s_{n-k}$ as claimed. $\qquad\square$

Proof of Proposition 4.34. It suffices to argue the case of a white lollipop reaction, by symmetry. We inductively verify the claim that, as the wall scans from left to right, the relative position between the pair of flags is kept to be at most $w_{[i,j]}$.

Suppose that the wall scanning is passing through a column of Type 2 or 3. Let $(\mathscr{L}_0, \mathscr{L}^0)$ be the pair of flags on the left of this column and $(\mathscr{R}_0, \mathscr{R}^0)$ be the pair of flags on the right of this Type 3 column. Suppose the that wall on the left goes across the interval $[i, j]$. Then by assumption, $\mathscr{L}_0 \overset{u}{\sim} \mathscr{L}^0$ for $u \leq w_{[i,j]}$. Let v be the permutation such that $\mathscr{R}_0 \overset{v}{\sim} \mathscr{R}^0$.

Let us start with the hardest case, namely a Type 3 column with a black lollipop at the kth gap with $i < k \leq j$, as depicted in Figure 80 (left). By a shift of indices, we can view v as an element in $S_{[2,n]}$, the permutation group that acts on the set $[2, n] = \{2, 3, \ldots, n\}$. We want to prove that $v \leq w_{[i,j-1]} \in S_{[2,n]}$. By shifting the indices of the Coxeter generators in (4.9), we have that $v \leq w_{[i,j-1]} \in S_{[2,n]}$ if and only if for all $n - j \leq l \leq n - i - 1$,

$$\{v(2), v(3), \ldots, v(l+1)\} \leq \{i+1 \leq \cdots \}.$$

Let us proceed by contradiction: suppose $v \not\leq w_{[i,j-1]} \in S_{[2,n]}$; then there must exist some l with $n - j \leq l \leq n - i - 1$ such that

$$\{v(2), v(3), \ldots, v(l+1)\} = \{a, a + \cdots, \ldots \},$$

where $a > i + 1$ is the smallest element in the set on the right. To deduce a contradiction, it suffices to prove that $\tilde{v} := s_{k-1} \cdots s_1 v s_1 \cdots s_{n-k} \not\leq w_{[i,j]} \in S_{[1,n]}$, since Lemma 4.35 states

87

that $\tilde{v}$ is the smallest possible relative position between $\mathscr{L}_0$ and $\mathscr{L}^0$. Direction computation yields that
$$\tilde{v} = (v(2)', v(3)', \ldots, v(n-k+1)', k, v(n-k+2)', \ldots, v(n)'),$$
where
$$m' := \begin{cases} m-1 & \text{if } m \le k, \\ m & \text{if } m > k. \end{cases}$$
If $l \le n-k$, then
$$\{\tilde{v}(1), \tilde{v}(2), \ldots, \tilde{v}(l)\} = \{v(2)', v(3)', \ldots, v(l+1)'\} = \{a', a' + \cdots, \ldots\} \not\le \{i \le \cdots\},$$
where the last $\not\le$ relation is because $a' \ge a - 1 > i$. This shows that $\tilde{v} \not\le w_{[i,j]}$.
Else, if $l > n-k$, then
$$\{\tilde{v}(1), \tilde{v}(2), \ldots, \tilde{v}(l+1)\} = \{v(2)', v(3)', \ldots, v(n-k+1)', k, v(n-k+2)', v(l+1)'\}$$
$$= \{k, a', a' + \cdots, \ldots\} \not\le \{i \le \cdots\},$$
where the last $\not\le$ relation is because both a' and k are greater than i.
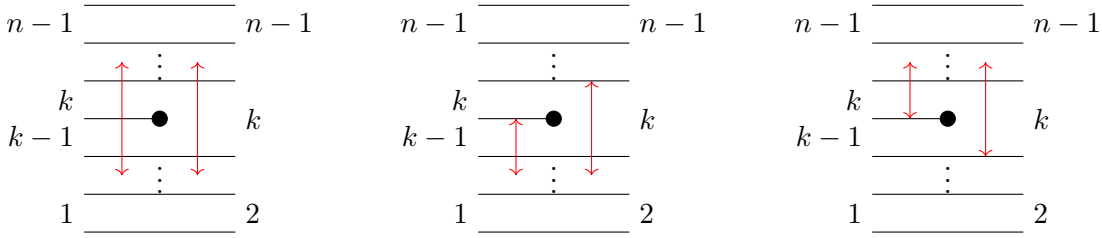
FIGURE 80. Three of the cases in the proof of Proposition 4.34.

There are three more special cases to consider for Type 3 columns with black lollipops, two of them depicted at the center and right of Figure 80, as well as the case where neither $k-1$ nor k is not contained in $[i,j]$.

For the case of Figure 80 (center), we want to prove that if $u \le w_{[i,k-1]} \in S_{[1,n]}$, then $v \le w_{[i,k-1]} \in S_{[2,n]}$. We proceed by contradiction again. Suppose $v \not\le w_{[i,k-1]}$; then there must exist some l with $n-k \le l \le n-i-1$ such that
$$\{v(2), v(3), \ldots, v(l+1)\} = \{a, a + \cdots, \ldots\},$$
where $a > i+1$ is the smallest element on the right. By the same argument, we see that
$$\{\tilde{v}(1), \tilde{v}(2), \ldots, \tilde{v}(l+1)\} = \{v(2)', v(3)', \ldots, v(n-k+1)', k, v(n-k+2)', v(l+1)'\}$$
$$= \{k, a', a' + \cdots, \ldots\} \not\le \{i \le \cdots\}.$$
This shows that $\tilde{v} \not\le w_{[i,k-1]} \in S_{[1,n]}$.

For the case of Figure 80 (right), we want to prove that if $u \le w_{[k,j]} \in S_{[1,n]}$, then $v \le w_{[k-1,j-1]} \in S_{[2,n]}$; a proof by contradiction works again, as follows. Suppose $v \not\le w_{[k-1,j-1]}$; then there must exist some l with $n-j \le l \le n-k$ such that
$$\{v(2), v(3), \ldots, v(l+1)\} = \{a, a + \cdots, \ldots\},$$
where $a > k$ is the smallest element on the right. Since $a > k$, we know that $a' = a$, and hence
$$\{\tilde{v}(1), \tilde{v}(2), \ldots, \tilde{v}(l)\} = \{a', a' + \cdots, \ldots\} = \{a, a + \cdots, \ldots\} \not\le \{k \le \cdots\}.$$
This shows that $\tilde{v} \not\le w_{[k,j]}$.

The remaining case, where neither $k-1$ nor k is contained in $[i,j]$, can be proved similarly, and it is left as an exercise. This covers all the cases with a Type 3 column with a black lollipop.

Next we consider the case of a Type 3 column with a white lollipop in the kth gap, counting from below in the GP-graph, as depicted below. By the Type 3 column requirement, we know that

$$\mathscr{R}_0 \overset{s_k s_{k+1}\cdots s_n}{\sim} h(\mathscr{L}_0) \overset{u}{\sim} h(\mathscr{L}^0) \overset{s_n s_{n-1}\cdots s_{n-k}}{\sim} \mathscr{R}^0,$$

where by assumption $u \leq w_{[i,j]}$. By Lemma 4.33 (2) and a direct computation, we conclude that $\mathscr{R}_0 \overset{v}{\sim} \mathscr{R}^0$ for

$$v \leq \mathrm{Dem}(\underline{s_k \cdots s_n}\, \underline{u}\, \underline{s_n \cdots s_{n-k}}) \leq \mathrm{Dem}(\underline{s_k \cdots s_n}\, \underline{w_{[i,j]}}\, \underline{s_n \cdots s_{n-k}}) = \begin{cases} w_{[i,j]} & \text{if } j < k-1, \\ w_{[i,j+1]} & \text{if } i \leq k-1 \leq j, \\ w_{[i+1,j+1]} & \text{if } i \geq k. \end{cases}$$

Lastly, we consider Type 2 columns. By using Lemma 4.33, we compute that, unless we are in one of the two wall shrinking depicted below, we have the logical implication $u \leq w_{[i,j]} \implies v \leq w_{[i,j]}$.

For the leftmost of the two cases depicted above, we directly have $u \leq w_{[i,j]} \implies v \leq w_{[i,j-1]}$, and for the rightmost one, we obtain $u \leq w_{[i,j]} \implies v \leq w_{[i+1,j]}$, as required. $\square$

Example 4.36. *Consider the GP-graph $\mathbb{G}$ in Figure 81, which we have discussed throughout the article and yields a Legendrian representative of the $m(5_2)$ knot. There are two faces in $\mathbb{G}$, both of which are non-square. Let A_1 and A_2 be the microlocal merodromies associated with the naive relative cycles of the left face and the right face, respectively. Figure 81 (left) records the relative position conditions that need to be satisfied along the vanishing locus $\{A_1 = 0\} \subset \mathfrak{M}(\Lambda(\mathbb{G}), T)$, as the result from the single lollipop reaction initiated from the left face. Figure 81 does the same thing for the vanishing locus $\{A_2 = 0\} \subset \mathfrak{M}(\Lambda(\mathbb{G}), T)$ as a result of a single lollipop reaction initiated from the right face.*
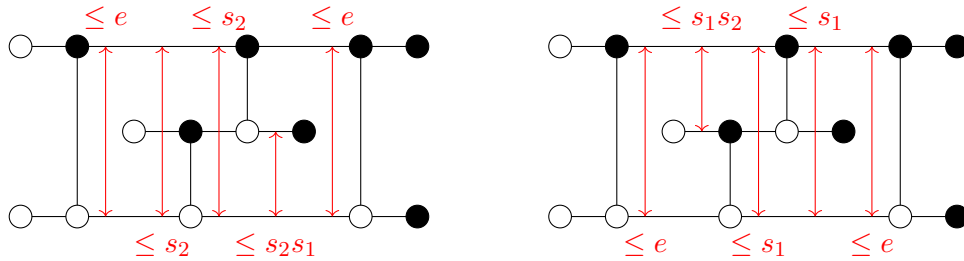

FIGURE 81. Relative position conditions from lollipop reactions.

Thus, it follows from these two lollipop reactions that A_1 vanishes at a point in $\mathscr{F}(\Lambda(\mathbb{G}), T)$ if and only if A_2 vanishes at that point as well. Moreover, by combining the two lollipop reactions together, we deduce that along their common vanishing locus, the flags must satisfy

*the relative position conditions depicted in Figure 82 (left) which, non-coincidentally, can
be depicted by the (non-free) weave shown in Figure 82 (right). This weave can be obtained
from the initial weave* $\mathfrak{w}(\mathbb{G})$ *by deleting the unique long* l-*cycle. These cycle deletions will
now have a central role in the remaining part of the argument for Theorem 1.1.*
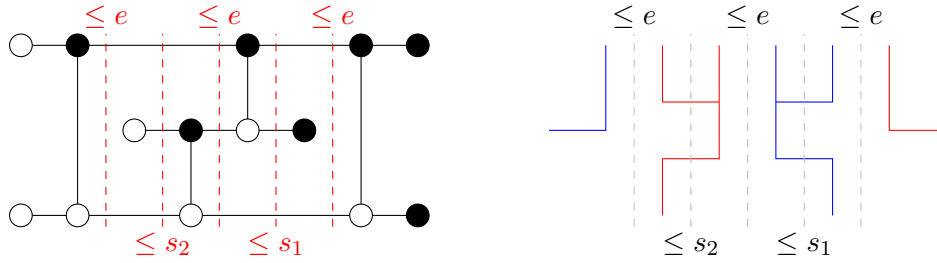
FIGURE 82. Relative position conditions from a lollipop chain reaction (left)
and a comparison with a non-free weave (right).

4.8. **Completeness of GP-graphs.** This brief subsection discusses the concept of *complete*
GP-graphs, for which the argument we present gives a complete proof of Theorem 1.1. As
prefaced in Subsection 4.2, the factoriality of the coordinate ring $\mathcal{O}(\mathfrak{M}(\Lambda(\mathbb{G}), T))$ is a re-
quirement. Following the results from Subsection 4.7, we add an additional hypothesis, as
follows:

Definition 4.37. A grid plabic graph $\mathbb{G}$ is said to be *complete* if the moduli stack $\mathfrak{M} = \mathfrak{M}(\Lambda(\mathbb{G}), T)$ satisfies the following properties

- The coordinate ring $\mathcal{O}(\mathfrak{M})$ is a unique factorization domain (UFD).

- For any face $f \subset \mathbb{G}$ with a sugar-free hull $\mathbb{S}_f$, the vanishing locus of the microlocal
  merodromy A_f is contained in the vanishing loci of A_g, for all faces $g \in \mathbb{S}_f$.

These two conditions are technical, and are only trying to capture the most general type
of GP-graph $\mathbb{G}$ for which the argument works. In practice, if a reasonable example of a $\mathbb{G}$
is given, it is possible to verify the second condition by direction computation (e.g. using
Gröbner basis), whereas the factoriality condition is, to our knowledge, more subtle. That
said, as explained in Subsection 4.2, we have developed combinatorial criteria to ensure the
first condition, e.g. Δ-completeness of $\beta(\mathbb{G})$, or even more combinatorially, $\mathbb{G}$ being a shuffle
graph. In fact, shuffle graphs $\mathbb{G}$ also satisfy the second condition, as can be seen by examining
the following combinatorial property:

Definition 4.38. A GP-graph $\mathbb{G}$ is said to be $\mathbb{S}$-complete if every sugar-free hull of $\mathbb{G}$ can
be obtained via some lollipop chain reaction.

By Propositions 4.10 and 4.34, $\mathbb{S}$-complete GP-graphs satisfy the second condition in Defini-
tion 4.38. By Proposition 2.26, shuffle graphs are $\mathbb{S}$-complete, and therefore complete. The
schematics of implications are:

$$\text{Shuffle graphs } \mathbb{G} \Longrightarrow (\beta(\mathbb{G}) \ \Delta\text{-complete}) + (\mathbb{S}\text{-complete } \mathbb{G}) \implies \text{complete } \mathbb{G}$$

In summary, though Theorem 1.1 is proven for complete grid plabic graphs, there are large
classes of $\mathbb{G}$-graphs that can be proven combinatorially to be complete, either because they
are shuffle or because $\mathbb{S}$-completeness and Δ-completeness are directly verified. Note that
shuffle graphs include all plabic fences, so all open Bott-Samelson varieties at the level of
$\mathfrak{M}(\Lambda, T)$, several families of interesting links (such as the twist knots), many braid varieties
(e.g. all 3-stranded ones), and more.

4.9. **Proof of the Main Theorem.** In this section, we conclude the proof Theorem 1.1. At this stage, we can consider the initial open torus chart in $\mathfrak{M}(\Lambda, T)$ given by $\mathfrak{M}(\mathfrak{w}(\mathbb{G}), T)$, as built in Section 3 and Subsection 4.3, with the candidate cluster $\mathcal{A}$-variables being the microlocal merodromies (constructed in Subsection 4.6) along a set of initial relative cycles (built in Section 3). Namely, given a GP-graph $\mathbb{G}$ and an initial set of relative cycles $\{\eta_1, \ldots, \eta_s\}$ for the pair $(L(\mathbb{G}), T)$ associated to $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$, we have an isomorphism $\mathbb{C}[A_{\eta_1}^{\pm 1}, \ldots, A_{\eta_s}^{\pm 1}] \cong \mathcal{O}(\mathfrak{M}(\mathfrak{w}, T))$, and thus $\{A_{\eta_1}, \ldots, A_{\eta_s}\}$ and their inverses naturally coordinatize the open torus chart $\mathfrak{M}(\mathfrak{w}(\mathbb{G}), T) \subset \mathfrak{M}(\Lambda, T)$. This defines an initial open toric chart U_0, candidate for an initial cluster $\mathcal{A}$-chart, with the quiver $Q(\mathbb{G}, \eta)$ being the intersection quiver associated to the (duals of the) relative cycles $\{\eta_1, \ldots, \eta_s\}$.

We shall now show that the algebra $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ coincides with an upper cluster algebra, along with the remaining items of Theorem 1.1. In geometric terms, the key ingredient for the former claim will be to prove that the initial cluster chart $U_0 \subset \mathfrak{M}(\Lambda, T)$, together with all the once-mutated charts, covers the moduli space $\mathfrak{M}(\Lambda, T)$ up to codimension 2, i.e.

$$U_0 \bigcup_{\eta \text{ mutable}} \mu_\eta(U_0) \overset{\text{up to codim. 2}}{=} \mathfrak{M}(\Lambda, T).$$

By Hartogs's extension theorem, any two normal varieties that differ at most in codimension 2 have the same algebra of regular functions. It thus follows from this codimension 2 isomorphism that $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ is equal to the coordinate ring of the union of the initial chart U_0 and all the once-mutated charts. Then [BFZ05, Corollary 1.9] is used to conclude that the coordinate ring of such a union is an upper cluster algebra.

4.9.1. *Erasing Y-trees on weaves and vanishing loci of face merodromies.* In order to establish the above covering of $\mathfrak{M}(\Lambda, T)$, up to codimension 2, by U_0 and the once-mutated charts $\mu_\eta(U_0)$, we need to gain understanding of the codimension 1 strata in $\mathfrak{M}(\Lambda, T)$ that appear as vanishing loci of certain microlocal merodromies. These loci can be explicitly described via non-free weaves that are obtained by erasing Y-cycle in $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$.

First, we begin with the diagrammatic process on the weave that erases Y-trees. Let $\mathfrak{w}$ be a free weave and Y-tree $\gamma \subset \mathfrak{w}$. By definition, the weave $\mathfrak{w}_{\hat{\gamma}}$ is the weave obtained by erasing $\gamma \subset \mathfrak{w}$ from $\mathfrak{w}$ according to the following local models:

- If γ ends at a trivalent vertex, then we simply erase the weave line contained in γ together with the trivalent vertex, turning the local picture into a single weave line. This is depicted in Figure 83 (left).

- If γ goes straight through a hexavalent vertex, then we erase the two weave lines contained in γ and pull the remaining weave lines apart according to their colors, turning the local picture into two weave lines. We draw this in Figure 83 (center).

- If γ branches off at a hexavalent vertex, erase the three weave lines contained in γ, turning the local picture into a trivalent weave vertex picture. See Figure 83 (left).
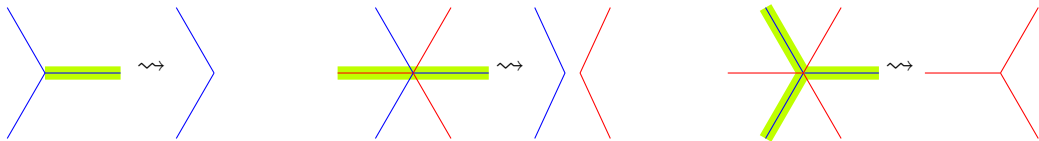
FIGURE 83. The local models for erasing Y-trees.

Alternatively, it is possible to first shorten the Y-tree to a short I-cycle, and erase the I-cycle, which only requires applying Figure 83 (left) twice. The following lemma verifies that this resulting weave is equivalent to the weave obtained by erasing the Y-tree directly:

Lemma 4.39 (Y-tree erasing)**.** *Let $\mathfrak{w}$ be a free weave and $\gamma \subset \mathfrak{w}$ a Y-tree. Let $\mathfrak{w}' \sim \mathfrak{w}$ be a weave obtained from $\mathfrak{w}$ by the double track trick that shortens γ into a short I-cycle (see Proposition 3.5), which we still denote $\gamma \subset \mathfrak{w}'$. Then there exists a weave equivalence $\mathfrak{w}'_{\hat\gamma} \sim \mathfrak{w}_{\hat\gamma}$. In particular, since $\mathfrak{w}'_{\hat\gamma}$ is not free, neither is $\mathfrak{w}_{\hat\gamma}$.*

Proof. Let us start with the short I-cycle γ in $\mathfrak{w}'$. Erasing γ in $\mathfrak{w}'$ leaves two weave lines with a Reeb chord in the middle, see Figure 84 (left). Follow the rest of the double tracks and contract these two weave lines inductively.
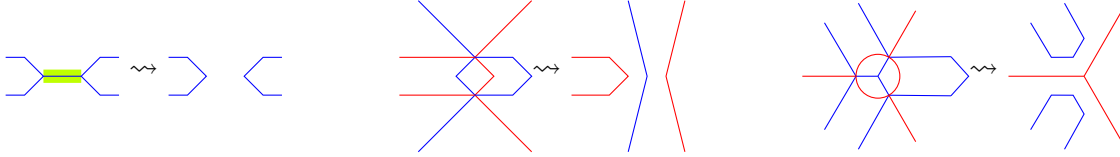

FIGURE 84. Three types of removals of (pieces of) Y-tree cycles.

At the part where the double track goes straight though, the contracting weave line can be pulled through this part by undoing a candy twist, as depicted in Figure 84 (center). At the part where the double track branches off, the contracting weave line can be pulled through it by using weave equivalences, thus becoming two contracting weave lines, as illustrated Figure 84 (right). In the end, we recover the weave $\mathfrak{w}_{\hat\gamma}$, as required. $\qquad\square$

Let us now study the vanishing loci of microlocal merodromies A_f associated to faces $f \subset \mathbb{G}$ by using Lemma 4.39. For that, fix a compatible set C of sign curves on the initial weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$ so that the chart $\mathcal{M}_1(\mathfrak{w})$ can be identified with $\mathrm{Loc}_1(L)$. For each relative 1-cycle $\eta \in H_1(L \setminus T, \Lambda \setminus T)$, its microlocal merodromy A_η (Definition 4.20) is well-defined and, in particular, we can associate a naive microlocal merodromy A_f with each face f of $\mathbb{G}$. By Corollary 4.23 and Proposition 4.29, these A_f's are regular functions on $\mathfrak{M}(\Lambda, T)$ and they are unique up to multiples of units.

Proposition 4.40. *Let $\mathbb{G}$ be a complete GP-graph and $f \subset \mathbb{G}$ a face. Suppose that the lollipop chain reaction initiated from f is complete. Then for any microlocal merodromy A_f associated with f, the vanishing locus $\{A_f = 0\} \subset \mathfrak{M}(\Lambda(\mathbb{G}), T)$ is non-empty.*

Proof. By assumption, f admits a sugar-free hull $\mathbb{S}_f$, which, by Lemma 3.38, gives rise to a Y-tree $\gamma \subset \mathfrak{w}$ in the initial weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$. Apply Proposition 3.5 to γ, making it a short I-cycle, and place this short I-cycle near the end of the original Y-tree, so that it lies inside some Type 1 weave column. If we delete this short I-cycle, we obtain a weave $\mathfrak{w}' = \mathfrak{w}_{\hat\gamma}$ whose associated weave surface is immersed, with a single interior Reeb chord at the midpoint of the short I-cycle. We claim that its associated stratum $\mathcal{M}_1(\mathfrak{w}')$ in $\mathfrak{M}(\Lambda(\mathbb{G}), T)$ is non-empty.

To prove this claim, cut the weave $\mathfrak{w}'$ open vertically across the column into two weaves $\mathfrak{w}_1$ and $\mathfrak{w}_2$, so that both $\mathfrak{w}_1$ and $\mathfrak{w}_2$ are free weaves again. By [JT17], or [EHK16, NRS$^+$20], the two strata $\mathcal{M}_1(\mathfrak{w}_1)$ and $\mathcal{M}_1(\mathfrak{w}_2)$ are non-empty. Now, since this column used to be a Type 1 column, the vertical slice along the cut is a reduced word of $w = s_i w_0$ for some i. Given that any two pairs of flags of relative position w are related to each other by a (non-unique) general linear group element, we use this action to line up the two pairs of flags of relative position w and glue any point in $\mathcal{M}_1(\mathfrak{w}_1)$ with any point in $\mathcal{M}_1(\mathfrak{w}_2)$ and get a point in $\mathcal{M}_1(\mathfrak{w}')$. This shows that $\mathcal{M}_1(\mathfrak{w}')$ is non-empty.

Since $\mathfrak{M}(\mathfrak{w}', T)$ fibers over $\mathcal{M}_1(\mathfrak{w}')$, it follows that $\mathfrak{M}(\mathfrak{w}', T)$ is non-empty as well. Let p be a point in $\mathfrak{M}(\mathfrak{w}', T)$. By Lemma 4.39, the weave $\mathfrak{w}'$ is equivalent to the weave $\mathfrak{w}_{\hat\gamma}$. Let η_f be a relative 1-chain associated with f and, without loss of generality, we may assume that η_f is contained in some Type 1 column. By construction, the cross-section of the initial weave

$\mathfrak{w}$ at a Type 1 column is always the positive (half-twist) braid Δ, the positive lift of w_0. The erasing of the Y-tree γ turns the cross-sectional positive braid into a positive braid β whose Demazure product satisfies $D(\beta) \leq s_i w_0$, for all i such that the ith gap, counting from below in the GP-graph $\mathbb{G}$, is contained in the face f. Moreover, the two flags at the two ends of η_f are of relative position $w \leq D(\beta)$. Thus we conclude that $w \leq s_i w_0$ and hence $A_f(p) = 0$ by Lemma 4.32. $\qquad\square$

We also establish the converse of Proposition 4.40, i.e. if the lollipop chain reaction initiated from f is incomplete, then the microlocal merodromy A_f must a unit in $\mathcal{O}(\mathfrak{M}(\Lambda, T))$.

Proposition 4.41. *Let $\mathbb{G}$ be a GP-graph and $f \subset \mathbb{G}$ a face. Then the microlocal merodromy A_f is a unit if and only if the lollipop chain reaction initiated at f is incomplete.*

Proof. Indeed, if the lollipop chain reaction initiated from f is complete, then Proposition 4.40 implies that $\{A_f = 0\}$ is non-empty and hence A_f cannot a unit. For the converse implication, note that the lollipop chain reaction initiated from f being incomplete implies that during a certain lollipop reaction in the chain, the selection process runs out of faces to select. This is equivalent to saying that the selection process selects an unbounded face g of the GP-graph. Now, if there exists a point $p \in \mathfrak{M}(\Lambda, T)$ with $A_f(p) = 0$, then by Proposition 4.34, $A_g(p) = 0$ as well. Nevertheless, this is impossible because any relative 1-chain η_g associated with g must map to the identity under the projection map $H_1(L \setminus T, \Lambda \setminus T) \to H_1(L, \Lambda)$, and hence A_g is a unit by Corollary 4.23. Therefore $\{A_f = 0\}$ is empty and A_f is a unit. $\qquad\square$

Finally, let us discuss the ratios of microlocal merodromies that appear when two faces share a sugar-free hull . Recall that the dual basis $\mathfrak{B}^\vee$ of $H_1(L, \Lambda)$ is constructed by starting with the set $\mathfrak{S}(\mathbb{G})$ of initial absolute cycles, which is a linearly independent subset of $H_1(L)$ and is in bijection with the sugar-free hulls of the GP-graph $\mathbb{G}$. Each element of $\mathfrak{S}(\mathbb{G})$ is a linear combination of the naive absolute cycles γ_f, which are in bijection with the faces of $\mathbb{G}$. Then, we complete $\mathfrak{S}(\mathbb{G})$ to a basis $\mathfrak{B}$ of $H_1(L)$ via a replacement process from bottom to top along the Hasse diagram $\mathcal{H}$ of the sugar-free hulls with respect to inclusion. On the dual side, this replacement process is performed from the top down along the Hasse diagram $\mathcal{H}$, replacing the naive relative cycles η_f one-by-one and thus obtaining a basis $\mathfrak{B}^\vee$ of $H_1(L, \Lambda)$ dual to $\mathfrak{B}$. By choosing a representative A_f for each naive relative cycle η_f, we construct a microlocal merodromy function A_i for each $i \in \mathfrak{B}^\vee$.

As illustrated in Example 3.42, there can be multiple faces sharing the same sugar-free hull. Thus, during the replacement process, some faces (naive absolute cycles) are set aside. Suppose $\mathbb{S}_f = \mathbb{S}_g$ for two different faces $f, g \subset \mathbb{G}$ and suppose that we set aside γ_f, while selecting γ_g. Then, on the dual side, the set-aside naive absolute cycle γ_f will correspond to the relative cycle $\eta_f - \eta_g$, which in turn gives rise to the microlocal merodromy function A_f/A_g.

Proposition 4.42. *Let $\mathbb{G}$ be a GP-graph and $f, g \subset \mathbb{G}$ two faces with equal sugar-free hulls. Then, the microlocal merodromy A_f/A_g corresponding to a set-aside naive absolute cycle γ_f is a unit.*

Proof. Since the faces f and g have the same sugar-free hull, the lollipop chain reaction initiated from one of them must contain the other. Therefore, by Proposition 4.34, $A_f(p) = 0$ if and only if $A_g(p) = 0$ for any $p \in \mathfrak{M}(\Lambda, T)$. Since $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ is a UFD, this implies that A_f and A_g are associates of each other, and thus A_f/A_g is a unit. $\qquad\square$

4.9.2. *Rank of exchange matrix and mutation formulae for Lagrangian surgeries.* Recall the notation $U_0 = \mathfrak{M}(\mathfrak{w}, T) \subset \mathfrak{M}(\Lambda, T)$ for the open toric chart associated to the (Lagrangian filling for the) weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$. Let us denote the naive microlocal monodromies as $\{A_f\}$, where f runs through the faces of $\mathbb{G}$, and the microlocal monodromies associated with marked

points as $\{A_t\}_{t \in T}$. The microlocal merodromies associated with the dual basis $\mathfrak{B}^\vee$ will be denoted by $\{A_i\}$, where $i \in [1, b_1(L)]$. By construction,

$$\mathcal{O}(U_0) = \mathbb{C}\left[A_{f_1}^{\pm 1}, \ldots, A_{f_{b_1(L)}}^{\pm 1}, A_t^{\pm 1}\right] = \mathbb{C}\left[A_1^{\pm 1}, \ldots, A_{b_1(L)}^{\pm 1}, A_t^{\pm 1}\right],$$

where the variable t runs through the set of marked points T. Let us also fix $\tilde{\mathfrak{B}}$ to be a completion of the basis $\mathfrak{B} \subset H_1(L)$ to a basis of $H_1(L, T)$. Then, by possibly adding relative 1-chains $\{\xi_t\}_{t \in T}$, we can modify elements of $\mathfrak{B}^\vee$ so that $\mathfrak{B}^\vee \sqcup \{\xi_t\}_{t \in T}$ becomes a dual basis of $\tilde{\mathfrak{B}}$. In this modification, each microlocal merodromy A_i is multiplied by some Laurent monomial in the A_t's. To ease notation, we rename the microlocal merodromies A_i's to include these Laurent monomials as well. Since $\tilde{\mathfrak{B}}$ and $\tilde{\mathfrak{B}}^\vee$ are dual bases, there is a natural bijection between them, and we will use them interchangeably as index sets for microlocal monodromies and microlocal merodromies.

Now, the intersection form $\{\cdot, \cdot\}$ on $H_1(L)$ can be extended to a skew-symmetric form on $H_1(L, T)$ by imposing a half integer value for intersections at T.
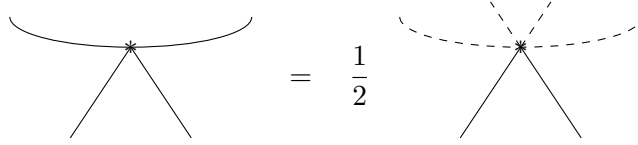
With respect to the basis $\tilde{\mathfrak{B}}$, the intersection form on $H_1(L, T)$ is then encoded by a $\tilde{\mathfrak{B}} \times \tilde{\mathfrak{B}}$ skew-symmetric matrix ϵ, where

$$\epsilon_{ij} = \{\gamma_i, \gamma_j\}$$

for any pair $\gamma_i, \gamma_j \in \tilde{\mathfrak{B}}$. For any absolute, Subsection 4.4 constructs the microlocal monodromy function ψ_γ on $\mathcal{M}_1(\mathfrak{w})$ for each 1-chain γ in $H_1(L)$. After the correction by sign curves, we have $X_i := \psi_{\gamma_i}$, and th collection $\{X_i\}_{i \in \mathfrak{B}}$ are our candidates for the initial cluster $\mathcal{X}$-variables.

Let $p : \mathfrak{M}(\Lambda, T) \longrightarrow \mathcal{M}_1(\Lambda)$ be the forgetful map. Then, by restricting to the respective tori supported on the initial weave $\mathfrak{w}$, we also obtain $p : \mathfrak{M}(\mathfrak{w}, T) \to \mathcal{M}_1(\mathfrak{w})$.

Proposition 4.43. *Consider the forgetful map* $p : \mathfrak{M}(\mathfrak{w}, T) \to \mathcal{M}_1(\mathfrak{w})$. *For an initial absolute cycle* $\gamma_i \in \mathfrak{S}(\mathbb{G})$,

$$p^*(X_i) = \prod_{j \in \tilde{\mathfrak{B}}} A_j^{\epsilon_{ij}}.$$

Proof. Let us denote the relative 1-chain dual to $\gamma_i \in \tilde{\mathfrak{B}}$ by η_i. It suffices to prove that

$$\gamma_i = \sum_j \epsilon_{ij} \eta_j$$

under the inclusion map $H_1(L) \cong H_1(L \setminus T) \hookrightarrow H_1(L \setminus T, \Lambda \setminus T)$. Note that since we have at least one marked point per link component in Λ, we can lift elements from $H_1(L, \Lambda)$ to $H_1(L, T)$. Now, for any element $\theta = \sum_{k \in \tilde{\mathfrak{B}}} c_k \gamma_k \in H_1(L, T)$ that is a lift of a relative 1-cycle in $H_1(L, \Lambda)$, we have

$$\left\langle \sum_j \epsilon_{ij} \eta_j, \theta \right\rangle = \left\langle \sum_j \epsilon_{ij} \eta_j, \sum_k c_k \gamma_k \right\rangle = \sum_j \epsilon_{ik} c_k = \langle \gamma, \theta \rangle.$$

Since the intersection form is non-degenerate on the tensor product $H_1(L) \otimes H_1(L, \Lambda)$, it indeed follows that $\gamma = \sum_j \epsilon_{ij} \eta_j$. $\qquad \square$

Corollary 4.44. *The rectangular exchange matrix* $\epsilon|_{\mathfrak{S}(\mathbb{G}) \times \tilde{\mathfrak{B}}}$ *is full-ranked.*

Proof. Since $\epsilon|_{\mathfrak{S}(\mathbb{G}) \times \tilde{\mathfrak{B}}}$ is a submatrix of $\epsilon|_{\mathfrak{B} \times \tilde{\mathfrak{B}}}$, it suffices to prove that $\epsilon|_{\mathfrak{B} \times \tilde{\mathfrak{B}}}$ is full-ranked. The latter follows from the surjectivity of the map $p : \mathfrak{M}(\mathfrak{w}, T) \longrightarrow \mathcal{M}_1(\mathfrak{w})$. $\qquad \square$

Let us now focus towards on the effect that a Lagrangian surgery, in the form of weave mutation, has on microlocal merodromies. For that we need to understand how relative cycles change under such an operation, as follows. Let γ_k be an initial absolute cycle which, by Lemma 3.38, we represent γ_k as a Y-tree on the initial weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$. By Proposition 3.5, there exists a weave equivalence that isotopes γ_k to a short l-cycle. In this weave equivalence, the dual basis element η_k of γ_k must also be isotoped to a curve that cuts through this short l-cycle in the middle. Thus, near the short l-cycle γ_k, the local model is the one depicted in Figure 85 (left). Note that each of the four weave lines extending out of this local picture may be part of multiple initial absolute cycles. However, we may assume without loss of generality that all other initial relative cycles are outside this local picture.
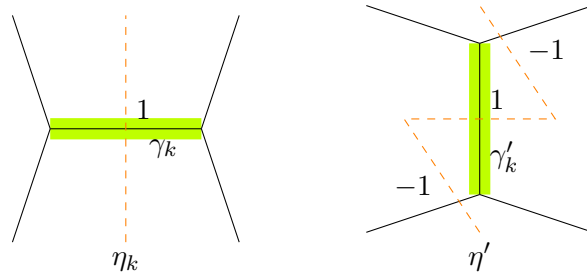
FIGURE 85. Mutation of initial relative cycles.

By performing a weave mutation at γ_k, we obtain a new weave $\mathfrak{w}_k$, which is mostly identical to $\mathfrak{w}$ except the local picture is replaced by Figure 85. Note that the initial absolute cycle γ_k in $\mathfrak{w}$ is replaced by the new absolute cycle γ'_k in $\mathfrak{w}_k$.

In the local model in Figure 85, we have also drawn a relative 1-chain η', which connects to the rest of $-\eta_k$ outside of the local picture. Nevertheless, this relative 1-chain η' is not the correct replacement for η_k, after the surgery, because in addition to having intersection 1 with γ'_k, the new relative cycle η'_k also needs to have trivial intersection with all other absolute cycles in $\mathfrak{B}$. At this stage, η' could possibly have non-trivial intersections with absolute cycles that come into the local picture from the northeast and the southeast. Thus, the correct replacement for the relative cycle η_k after the weave mutation is the linear combination

$$(4.10) \qquad \eta'_k = \eta' + \sum_{j \in \tilde{\mathfrak{B}}} [-\epsilon_{kj}]_+ \eta_j.$$

This explains how to keep track of relative cycles after a weave mutation, and thus a Lagrangian surgery in our context. Note that the moduli space $\mathfrak{M}(\mathfrak{w}_k, T)$ also defines an open toric chart $\mathfrak{M}(\mathfrak{w}_k, T) \subset \mathfrak{M}(\Lambda, T)$, as $\mathfrak{w}_k$ defines an embedded exact Lagrangian filling as well. Let us denote this chart, where we have performed a Lagrangian surgery at the kth disk of the $\mathbb{L}$-compressing system, by $U_k \subset \mathfrak{M}(\Lambda, T)$, and denote the microlocal merodromy associated with the relative cycle η'_k by A'_k. In order to understand the change of the microlocal merodromies under surgery, we have the following result.

Proposition 4.45. *At any point* $u \in U_0 \cap U_k$,

$$A'_k = A_k^{-1}(1 + p^*(X_k)) \prod_{j \in \tilde{\mathfrak{B}}} A_j^{[-\epsilon_{kj}]_+},$$

where $p : U_0 \to \mathcal{M}_1(\mathfrak{w})$ *is the forgetful map restricted to* $U_0 \cap U_k$.

Proof. It suffices to prove that $A_{\eta'} = A_k^{-1}(1 + p^*(X_k))$. Since γ_k is a short l-cycle, we can assume that the four neighboring flags are four lines l_e, l_s, l_w, and l_n in $\mathbb{C}^2$. Let v_i be a

non-zero vector in each line l_i and let det denote the dual of the non-zero 2-form associated with $\mathbb{C}^2$. Following the crossing-value formula, we obtain

$$A_k = \det(v_s \wedge v_n), \qquad A_{\eta'} = \frac{\det(v_e \wedge v_w)}{\det(v_n \wedge v_e)\det(v_s \wedge v_w)}.$$

Therefore the product can be computed as

$$\begin{aligned}
A_k A_{\eta'} &= \frac{\det(v_s \wedge v_n)\det(v_e \wedge v_w)}{\det(v_n \wedge v_e)\det(v_s \wedge v_w)} \\
&= \frac{\det(v_n \wedge v_e)\det(v_s \wedge v_w) + \det(v_n \wedge v_w)\det(v_e \wedge v_s)}{\det(v_n \wedge v_e)\det(v_s \wedge v_w)} \\
&= 1 + p^*(X_k). \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad \square
\end{aligned}$$

Proposition 4.45 and Proposition 4.43, yield the desired cluster $\mathcal{A}$-mutation formula for the microlocal merodromies under Lagrangian surgeries on the set of initial $\mathbb{L}$-compressing disks:

Corollary 4.46. *Let $\mathbb{G}$ be a GP-graph, $\{\eta_1, \ldots, \eta_s\}$ the set of naive relative cycles and $\{A_i\}$ the associated set of naive microlocal merodromies. Consisder the microlocal merodromy A'_k along the relative 1-chain η'_k obtained from η_k by weave mutation at the dual 1-cycle γ_k. Then*

$$A'_k = \frac{\prod_{j \in \tilde{\mathfrak{B}}} A_j^{[\epsilon_{kj}]_+} + \prod_{j \in \tilde{\mathfrak{B}}} A_j^{[-\epsilon_{kj}]_+}}{A_k}.$$

4.9.3. *Regularity of initial microlocal merodromies.* By Proposition 4.29, the naive microlocal merodromies A_f are regular functions on the moduli space $\mathfrak{M}(\Lambda, T)$. However, since the adjusted microlocal merodromies $\{A_i\}_{i \in \tilde{\mathfrak{B}}}$, corresponding to the initial basis are ratios of the naive microlocal merodromies, the initial merodromies $\{A_i\}$ are only rational functions a priori. Our next goal is to prove that for all $i \in \mathfrak{B}$, A_i are actually global regular functions, and that they are either irreducible if $i \in \mathfrak{S}(\mathbb{G})$, or units otherwise. Let us start with the following lemma:

Lemma 4.47. *Let $U_0 \subset \mathfrak{M}(\Lambda, T)$ be the initial open toric chart and f a unit in $\mathcal{O}(U_0)$, resp. $\mathcal{O}(U_k)$. Suppose that $f = gh$ in $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ for some $g, h \in \mathcal{O}(\mathfrak{M}(\Lambda, T))$. Then g and h are also units in $\mathcal{O}(U_0)$, resp. $\mathcal{O}(U_k)$.*

Proof. Indeed, if $f = gh$ in $\mathcal{O}(\mathfrak{M}(\Lambda, T))$, then $f = gh$ in $\mathcal{O}(U_0)$ as well, and if f is a Laurent monomial in $\mathcal{O}(U_0)$, then each of g and h must be a Laurent monomial, too. The proof for the case where f is a unit in $\mathcal{O}(U_v)$ is analogous. $\qquad\qquad\qquad\qquad\square$

Let us first show that the initial merodromies are irreducible under the assumption that they are regular functions:

Lemma 4.48. *Let $\mathbb{G}$ be a GP-graph, $\gamma_k \in \mathfrak{S}(\mathbb{G})$ an initial absolute cycle and consider $A_k : \mathfrak{M}(\Lambda, T) \dashrightarrow \mathbb{C}$ an associated microlocal merodromy. Suppose that A_k is a regular function, i.e., an element of $\mathcal{O}(\mathfrak{M}(\Lambda, T))$. Then A_k is irreducible.*

Proof. The proof is in line with Geiss-Leclerc-Shröer's argument of irreducibility of cluster variables in [GLS13, Theorem 3.1]. Suppose $A_k = gh$ in $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ with neither g nor h being a unit. Then Lemma 4.47 implies that g and h must be Laurent monomials , and hence can be expressed as $\prod_{i \in \tilde{\mathfrak{B}}} A_i^{m_i}$ and $\prod_{i \in \tilde{\mathfrak{B}}} A_i^{n_i}$, respectively, up to a multiple of units in $\mathcal{O}(\mathfrak{M}(\Lambda, T))$. Since $A_k = \prod_j A_i^{m_i + n_i}$, then at least one of m_k and n_k must be positive. Without loss of generality, let us assume that $m_k > 0$. Then, since h is not a unit, there must be some $j \neq k$ such that A_j is not a unit and $n_j > 0$. If A_j is not a unit, then by Proposition 4.41, j must correspond to an initial absolute cycle γ_j. By mutating along the

initial absolute cycle γ_j, we obtain a new weave $\mathfrak{w}_j$. By Lemma 4.47, h is also a Laurent monomial in the new chart $\mathcal{O}(U_j)$ associated to $\mathfrak{w}_j$ and hence we can write g and h as

$$g = A_j'^{p_j} \prod_{i \neq j} A_i^{p_i} \qquad \text{and} \qquad h = A_j'^{q_j} \prod_{i \neq j} A_i^{q_i}.$$

Note that since $A_k = gh$, we must have $p_j + q_j = 0$. If $p_j = q_j = 0$, then we have a contradiction because $\prod_{i \neq j} A_i^{q_i} = h = \prod_i A_i^{n_i}$ with $n_j > 0$. That said, if p_j and q_j are non-zero, then one of them must be positive; suppose $p_j > 0$. By Corollary 4.46, we have $A_j' = M_1 + M_2$ where M_1 and M_2 are two algebraically independent Laurent monomials in $\{A_i\}_{i \in \tilde{\mathfrak{B}}}$, up to units. It then follows that

$$g = (M_1 + M_2)^{p_j} \prod_{i \neq j} A_i^{p_i},$$

which shows that g is not a Laurent monomial in $\{A_i\}_{i \in \tilde{\mathfrak{B}}}$. This is again a contradiction, and therefore A_k must be an irreducible element in $\mathcal{O}(\mathfrak{M}(\Lambda, T))$. $\qquad \square$

We are ready to conclude regularity, and thus irreducibility, of initial merodromies:

Proposition 4.49. *Let $\mathbb{G}$ be a GP-graph, $\gamma_k \in \mathfrak{S}(\mathbb{G})$ an initial absolute cycle and A_k an associated microlocal merodromy. Then A_k is a regular function and an irreducible element in $\mathcal{O}(\mathfrak{M}(\Lambda, T))$.*

Proof. By Lemma 4.48, it suffices to prove that A_k is a regular function. We proceed by induction from top down along the Hasse diagram $\mathcal{H}$; recall that vertices of $\mathcal{H}$ are sugar-free hulls and hence they are naturally indexed by the set of initial absolute cycles $\mathfrak{S}(\mathbb{G})$. For the base case, suppose k is a maximal vertex in the Hasse diagram. Then $A_k = A_f$ for some naive relative cycle η_f. Then Proposition 4.29 implies that $A_k = A_f$ is a regular function, as required.

Inductively, suppose for all $i > k$ in the Hasse diagram $\mathcal{H}$, A_i is a regular function on $\mathfrak{M}(\Lambda, T)$. By Lemma 4.48, A_i are irreducible elements in $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ as well. Let f_i be the face selected for each vertex i of $\mathcal{H}$. Then, for each vertex i of $\mathcal{H}$, we have

$$A_{f_i} = \prod_{j \geq i} A_j.$$

In particular, if $i > k$, then the above is the unique factorization of the naive microlocal merodromies A_{f_i} in $\mathcal{O}(\mathfrak{M}(\Lambda, T))$, up to mutliple of units. This also implies that the irreducible elements $\{A_i\}_{i > k}$ are not associates of each other because A_{f_i} are not units by Corollary 4.23.

In addition to the above, if $i > k$, then f_i is contained in the sugar-free hull $\mathbb{S}_{f_i} = \mathbb{S}_i$. Proposition 4.34 implies $\{A_{f_i} = 0\} \subset \{A_{f_k} = 0\}$, but we obtain the inclusion $\{A_i = 0\} \subset \{A_{f_i} = 0\}$ as well because A_i is an irreducible factor of A_{f_i}. Therefore $\{A_i = 0\} \subset \{A_{f_k} = 0\}$, which is equivalent to A_{f_k} being divisible by A_i for all $i > k$. Since A_i are distinct irreducible elements of $\mathcal{O}(\mathfrak{M}(\Lambda, T))$, it follows that the quotient

$$A_k = A_{f_k} \prod_{i > k} A_i^{-1}$$

is also a regular function on $\mathfrak{M}(\Lambda, T)$ as well. The induction is now complete. $\qquad \square$

4.9.4. *Conclusion of the argument.* We finalize the proof of the covering of $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ by the initial and adjacent charts, up to codimension 2. For each initial absolute cycle $\gamma_k \in \mathfrak{S}(\mathbb{G})$, we denote the vanishing locus of the associatied microlocal merodromy by $D_k := \{A_k = 0\} \subset \mathfrak{M}(\Lambda, T)$. Since A_k is an irreducible element of $\mathcal{O}(\mathfrak{M}(\Lambda, T))$, D_k is irreducible as a codimension 1 subvariety in $\mathfrak{M}(\Lambda, T)$.

Proposition 4.50. *Let $\mathbb{G}$ be a GP-graph, $\gamma_k \in \mathfrak{S}(\mathbb{G})$ initial absolute cycle, $U_k \subset \mathfrak{M}(\Lambda, T)$ the open torus chart associated to the Lagrangian surgery of $L(\mathfrak{w}(\mathbb{G}))$ at γ_k, and $D_k \subset \mathfrak{M}(\Lambda, T)$ the vanishing locus of its associated microlocal merodromy. Then the intersection $U_k \cap D_k \subset \mathfrak{M}(\Lambda, T)$ is a non-empty open subset of the vanishing locus D_k.*

Proof. It suffices to prove that $U_k \cap D_k$ is non-empty. Similar to the proof of Proposition 4.40, we apply Proposition 3.5 to move γ_k to a short l-cycle near the end of the original Y-tree, so that it lies inside some Type 1 weave column. By deleting this short l-cycle, we obtain a weave $\mathfrak{w}'$ whose moduli space $\mathfrak{M}(\mathfrak{w}', T)$ is a subset of D_k. It suffices to show that $\mathfrak{M}(\mathfrak{w}', T) \cap U_k \neq \emptyset$, but this clear: for instance, in the case where the weave looks like the one on the left below,

we first fix a point in $\mathcal{M}_1(\mathfrak{w}_1)$ and, then based on the flags $\mathscr{L}_0, \mathscr{L}_1, \mathscr{L}_2, \mathscr{L}_3$, we choose a point in $\mathcal{M}_1(\mathfrak{w}_2)$ with flags $\mathscr{R}_0 = \mathscr{L}_0$, $\mathscr{R}_1 = \mathscr{L}_1$, $\mathscr{R}_2 = \mathscr{L}_2$, but $\mathscr{R}_3 \neq \mathscr{L}_3$, and then glue them together. This gives a point in $\mathcal{M}_1(\mathfrak{w}')$ which is also in $\mathcal{M}_1(\mathfrak{w}_k)$, where $\mathfrak{w}_k$ is the mutated weave, which is also shown on the right above). $\qquad\square$

At this stage, the covering property, up to codimension 2, readily follows:

Proposition 4.51. *Let $\mathbb{G}$ be a GP-graph, $\gamma_k \in \mathfrak{S}(\mathbb{G})$ the initial absolute cycles and $U_k \subset \mathfrak{M}(\Lambda, T)$ the open torus charts associated to the Lagrangian surgery of $L(\mathfrak{w}(\mathbb{G}))$ at each γ_k, where U_0 is the initial chart associated to $L(\mathfrak{w}(\mathbb{G}))$. Then*

$$\operatorname{codim}\left(U_0 \cup \bigcup_{k \in \mathfrak{S}(\mathbb{G})} U_k\right) \geq 2,$$

i.e. the union of U_0 and all the adjacent charts U_k covers $\mathfrak{M}(\Lambda, T)$ up to codimension 2.

Proof. Consider the following inclusions:

$$\left(U_0 \cup \bigcup_k U_k\right)^c = U_0^c \cap \bigcap_k U_k^c = \left(\bigcup_j D_j\right) \cap \bigcap_k U_k^c = \bigcup_j \left(D_j \cap \bigcap_k U_k^c\right) \subset \bigcup_j \left(D_j \cap U_j^c\right).$$

By Proposition 4.50, the intersection $U_j \cap D_j$ is open in D_j, for all j, and thus we have $\operatorname{codim}(D_j \cap U_j^c) \geq 2$ for each j. Thus we conclude that $\operatorname{codim}\left(U_0 \cup \bigcup_k U_k\right) \geq 2$ as well. $\qquad\square$

The concluding statement[21] in Theorem 1.1 is now proven as follows:

Theorem 4.52. *Let $\mathbb{G}$ be a complete GP-graph. Then the coordinate ring of regular functions $\mathcal{O}(\mathfrak{M}(\Lambda(\mathbb{G}), T))$ has the structure of an upper cluster algebra.*

---

[21]By construction, and its proof, the following statement implicitly contains the fact that such structure of an upper cluster algebra has the microlocal merodromies as initial cluster $\mathcal{A}$-variables and all other symplectic geometric ingredients developed throughout the manuscript. See the statement of Theorem 1.1.

Proof. Consider the open subset $U_0 \cup \bigcup_{k \in \mathfrak{S}(\mathbb{G})} U_k \subset \mathfrak{M}(\Lambda, T)$. Proposition 4.51 shows the equality of coordinate rings $\mathcal{O}(\mathfrak{M}(\Lambda, T)) = \mathcal{O}(U_0 \cup \bigcup_{k \in \mathfrak{S}(\mathbb{G})} U_k)$. Corollary 4.46 implies that $\mathcal{O}(U_0 \cup \bigcup_{k \in \mathfrak{S}(\mathbb{G})} U_k)$ is an upper bound of a cluster algebra. In addition, since T has at least one marked point per link component, Corollary 4.44 shows that the rectangular exchange matrix $\epsilon|_{\mathfrak{S}(\mathbb{G}) \times \tilde{\mathfrak{B}}}$ is full-ranked. Then [BFZ05, Corollary 1.9] implies that this upper bound coincides with its upper cluster algebra and therefore we conclude that $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ is an upper cluster algebra. $\square$

Corollary 1.2 is deduced as follows:

Corollary 4.53. *Let $\mathbb{G}$ be a complete GP-graph. Then $\mathcal{O}(\mathcal{M}_1(\Lambda(\mathbb{G})))$ has the structure of a cluster Poisson algebra.*

Proof. Let us temporarily denote the cluster $\mathcal{A}$-variety defined by the $\tilde{\mathfrak{B}} \times \tilde{\mathfrak{B}}$ exchange matrix ϵ by $\mathcal{A}$ and denote the cluster $\mathcal{X}$-variety associated with the submatrix $\epsilon|_{B \times B}$ by $\mathcal{X}$. Since $\epsilon|_{B \times \tilde{B}}$ is full-ranked, which follows from the surjectivity of $p : \mathfrak{M}(\Lambda, T) \to \mathcal{M}_1(\Lambda)$, the cluster theoretical map $p : \mathcal{A} \to \mathcal{X}$ is also surjective. Noth $\mathcal{O}(\mathcal{A})$ and $\mathcal{O}(\mathcal{X})$ are intersections of Laurent polynomial rings, and thus a rational function f on $\mathcal{X}$ is regular if and only if $p^*(f)$ is regular on $\mathcal{A}$, see [SW20, Lemma A.1]. That said, given that $p : \mathfrak{M}(\Lambda, T) \to \mathcal{M}_1(\Lambda)$ is surjective, a rational function on $\mathcal{M}_1(\Lambda)$ is regular if and only if $p^*(g)$ is regular on $\mathfrak{M}(\Lambda, T)$. Now consider the following commutative diagram

$$
\begin{array}{ccc}
\mathfrak{M}(\Lambda, T) & \overset{\alpha}{\underset{\cong}{\dashrightarrow}} & \mathcal{A} \\
{\scriptstyle p}\downarrow & & \downarrow{\scriptstyle p} \\
\mathcal{M}_1(\Lambda) & \underset{\chi}{\dashrightarrow} & \mathcal{X}.
\end{array}
$$

Both horizontal maps are birational because $\mathfrak{M}(\Lambda, T)$ (resp. $\mathcal{M}_1(\Lambda)$) and $\mathcal{A}$ (resp. $\mathcal{X}$) share an open torus chart $\mathfrak{M}(\mathfrak{w}, T)$ (resp. $\mathcal{M}_1(\mathfrak{w})$). In addition, Theorem 4.52 implies that the top map induces an isomorphism between $\mathcal{O}(\mathfrak{M}(\Lambda, T))$ and $\mathcal{O}(\mathcal{A})$. Now, given a regular function $f \in \mathcal{O}(\mathcal{X})$, the pull-back $\chi^*(f)$ is a rational function on $\mathcal{M}_1(\Lambda)$ by birationality; but since $p^* \circ \chi^*(f) = \alpha^* \circ p^*(f)$ is regular on $\mathfrak{M}(\Lambda, T)$, it follows that $\chi^*(f)$ is regular on $\mathcal{M}_1(\Lambda)$. Conversely, if we are given a regular function $f \in \mathcal{O}(\mathcal{M}_1(\Lambda))$, we know that $(\chi^{-1})^*(f)$ is a rational function on $\mathcal{X}$ by birationality; but since $p^* \circ (\chi^{-1})^*(f) = (\alpha^{-1})^* \circ p^*(f)$ is regular on $\mathcal{A}$, it follows that $(\chi^{-1})^*(f)$ is regular on $\mathcal{X}$ as well. Therefore we conclude that χ induces an algebra isomorphism between $\chi^* : \mathcal{O}(\mathcal{X}) \to \mathcal{O}(\mathcal{M}_1(\Lambda))$, and hence $\mathcal{O}(\mathcal{M}_1(\Lambda))$ is a cluster Poisson algebra. $\square$

4.10. **Examples and comments.** This manuscript concludes with a few more explicit computations and examples. The two examples that we present illustrate two different points. First, the computations for the Legendrian $m(5_2)$ associated to the GP-graph in Figure 16 (upper-left) are the continuation of this running example, fitting within the framework of Theorem 1.1 and explaining how to implement the ingredients and results developed in Section 4 thus far in this particular instance.

Second, Theorem 1.1 has a grid plabic graph $\mathbb{G}$ as an input. Nevertheless, a major part of the ingredients and results that we have developed throughout the article only involve the weave $\mathfrak{w}(\mathbb{G})$ and the set of initial absolute cycles. From a symplectic geometric viewpoint, it is also natural to start with a Lagrangian filling $L = L(\mathfrak{w})$ given by a weave $\mathfrak{w}$. In case that a basis for $H_1(L, T)$ can be found through Y-trees in the weave $\mathfrak{w}$, with appropriate properties, many (often all) of the conclusions of Theorem 1.1 also follow. In general, finding a basis of Y-tree with the right properties can be challenging, but there are circumstances (and many explicit examples) where this can be done directly, see [CGGS20, Section 4.7] and upcoming work of the first author with I. Le, E. Gorsky, M. Gorsky, J. Simental and L. Shen.

The second example we present is an explicit example that illustrates how the methods of the article work directly starting with a weave filling $L(\mathfrak{w})$ of a Legendrian knot Λ, where no GP-graph is ever mentioned. We are thankful to L. Ng for asking us a question about cluster structures for a certain Legendrian representative of $m(7_2)$, which is *not* Legendrian isotopic to the Legendrian representative of $m(7_2)$ associated to the GP-graph in Figure 16 (upper-right). Here are the two final examples.

The running example of $m(5_2)$. Let us consider the shuffle graph $\mathbb{G}$ depicted at the Figure 86 once (and only once) more. The associated GP-link $\Lambda = \Lambda(\mathbb{G})$ is the Legendrian knot of smooth knot type $m(5_2)$ that we have been using in many examples, see also Figures 37 and 16 (upper-left). The weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$ drawn in Figure 37 simplifies by undoing a candy twist to the weave in Figure 86 (right).

FIGURE 86. A shuffle graph (left) and its initial weave together with a compatible set of sign curves.

Let us denote the naive absolute cycles as $\gamma_i := \partial f_i$ for $i = 1, 2$. There is only one sugar-free hull $f_1 \cup f_2$. Thus, the set of initial absolute cycles consists of only one element, namely $\delta := \gamma_1 + \gamma_2$. It embeds in the initial weave as the unique long I-cycle between the left most and the right most trivalent weave vertices. As $\Lambda(\mathbb{G})$ is smoothly a knot, $m(5_2)$, we only need to put one marked point on its boundary, which we place on the lowest strand at upper left hand corner. Among the four Type 2 columns, only the right most one is critical. Let us fix a compatible set of sign curves as depicted on the right in Figure 86. Since the unique marked point is the end point of four distinct sign curves, there is no need to change the decorations across this marked point. By reading the boundary weave lines, we conclude that $\mathcal{M}(\Lambda(\mathbb{G}))$ is the configuration of 10 flags satisfying the relative position conditions described in Example 4.8.

$$
\begin{pmatrix} \det \\ \alpha_0 \\ v_0 \end{pmatrix} \overset{s_1}{\diagup} \begin{pmatrix} -\det \\ -\alpha_0 \\ v_1 \end{pmatrix} \overset{s_2}{-} \begin{pmatrix} \det \\ \alpha_1 \\ v_1 \end{pmatrix} \overset{s_2}{-} \begin{pmatrix} -\det \\ \alpha_2 \\ v_1 \end{pmatrix} \overset{s_1}{-} \begin{pmatrix} \det \\ -\alpha_2 \\ v_2 \end{pmatrix} \overset{s_1}{-} \begin{pmatrix} -\det \\ \alpha_2 \\ v_3 \end{pmatrix} \overset{s_2}{-} \begin{pmatrix} \det \\ \alpha_3 \\ v_3 \end{pmatrix} \overset{s_2}{\diagdown} \begin{pmatrix} -\det \\ \alpha_4 \\ v_3 \end{pmatrix}
$$

$$
\begin{pmatrix} -\det \\ -\alpha_0 \\ v_4 \end{pmatrix} \overset{s_2}{\rule{4em}{0.4pt}} \begin{pmatrix} \det \\ -\alpha_4 \\ v_4 \end{pmatrix} \overset{s_1}{\diagup}
$$

Now, the framing data decorate each flag by associating a non-zero form to each subspace. Following the cyclic order in Example 4.8, we name them as follows. The leftmost decorated flag is associated with the sector in which the unique marked point is located. Note that the sign changes on the decorations are in accordance with Equation (4.8). In addition, by Lemma 4.26, we have the following identities

$$\alpha_0 = v_4 \wedge v_0 = v_1 \wedge v_0, \qquad \alpha_2 = v_2 \wedge v_1 = v_2 \wedge v_3, \qquad \alpha_4 = v_4 \wedge v_3,$$

$$\det = (-\alpha_0) \overset{v_1}{\wedge} \alpha_1 = \alpha_2 \overset{v_1}{\wedge} \alpha_1 = \alpha_2 \overset{v_3}{\wedge} \alpha_3 = \alpha_4 \overset{v_3}{\wedge} \alpha_3 = (-\alpha_0) \overset{v_4}{\wedge} (-\alpha_4).$$

Suppose $\alpha_1 = u \wedge v_1$ and $\alpha_3 = w \wedge v_3$. Then from the det identity above, we see that

$$\det = v_0 \wedge u \wedge v_1 = v_2 \wedge u \wedge v_1 = v_2 \wedge w \wedge v_3 = v_4 \wedge w \wedge v_3 = v_0 \wedge v_3 \wedge v_4.$$

Next let us define the naive microlocal merodromies. We have a few possible choices for the dual naive relative cycle for each of the two faces. In Figure 87, we draw a few possible choices for the face f_2 in purple.

FIGURE 87. Possible choices for naive relative cycles.

We can now apply Proposition 4.29 to write down expressions for these microlocal merodromies. We have labeled the sheet index for the sign curves to help us keep track of whether these relative 1-cycles cross the sign curves or not. The resulting expressions are as follows:

$$A_2 = \frac{v_4 \wedge \alpha_2}{-\det} = \frac{v_4 \wedge (-\alpha_2)}{\det} = -\frac{v_2 \wedge (-\alpha_4)}{\det} = \frac{\alpha_4 \overset{v_3}{\wedge} \alpha_2}{-\det}.$$

Similarly, by using the relative 1-cycle in green, one can find that

$$A_1 = -\frac{v_4 \wedge v_1}{-\alpha_0} = \frac{v_4 \wedge v_1}{\alpha_0}.$$

Since $\mathfrak{S}(\mathbb{G}) = \{\delta\}$ is a singleton set but $H_1(L, T) \cong H_1(L) \cong \mathbb{Z}^2$ (Example 3 in Subsection 3.9), we need to choose a completion of $\mathfrak{S}(\mathbb{G})$ to a basis of $H_1(L)$. One such choice is by adding γ_1, which gives the following initial quiver, where the vertex γ_1 is frozen:

$$\delta \longleftarrow \gamma_1$$

The dual basis to $\{\delta, \gamma_1\}$ is $\{\eta_2, \eta_1 - \eta_2\}$. Therefore the initial cluster is

$$A_2 \longleftarrow \frac{A_1}{A_2}$$

In particular, Proposition 4.42 implies that

$$\frac{A_1}{A_2} = \frac{v_4 \wedge v_1}{\alpha_0} \frac{\det}{v_4 \wedge v_1 \wedge v_2} = \frac{\det}{\alpha_0 \wedge v_2}$$

is a unit. In other words, $\alpha_0 \wedge v_2 \neq 0$ up to codimension 2 and hence when computing with functions on $\mathfrak{M}(\Lambda(\mathbb{G}), T)$, we may assume that the line $\mathrm{Span}(v_2)$ and the plane $\mathrm{Span}(\alpha_0)$ are always transverse to each other.

Let us now compute the cluster $\mathcal{A}$-mutation at A_2. According to the cluster $\mathcal{A}$-mutation formula, we should get

$$A_2' = \frac{1 + \frac{A_1}{A_2}}{A_2} = \frac{\det + \frac{\det}{\alpha_0 \wedge v_2} \det}{v_4 \wedge v_3 \wedge v_2} = \frac{(v_1 \wedge v_0 \wedge v_2) \det + \det^2}{(v_4 \wedge v_3 \wedge v_2)(\alpha_0 \wedge v_2)}$$

$$=\frac{(v_3 \wedge v_0 \wedge v_2)(v_4 \wedge w \wedge v_3) + (v_2 \wedge w \wedge v_3)(v_0 \wedge v_3 \wedge v_4)}{(v_4 \wedge v_3 \wedge v_2)(\alpha_0 \wedge v_2)}$$

$$=\frac{(v_4 \wedge v_2 \wedge v_3)(v_0 \wedge w \wedge v_3)}{(v_4 \wedge v_3 \wedge v_2)(\alpha_0 \wedge v_2)}$$

$$=\frac{v_0 \wedge v_3 \wedge w}{\alpha_0 \wedge v_2}.$$

Note that in the second to the last equality, we used a Plücker relation among the vectors v_4, v_0, v_2, w in $\mathbb{C}^3/\mathrm{Span}(v_3)$. Also, since $\alpha_0 \wedge v_2$ never vanishes, we can deduce that A_2' is a regular function on $\mathfrak{M}(\Lambda(\mathbb{G}), T)$.

Finally, let us also demonstrate how we can compute A_2' via a weave mutation. We first implement a weave equivalence to the initial weave to obtain the weave on the top left of Figure 88, in which the unique long I-cycle δ becomes a short I-cycle. We draw a representative of the dual initial relative cycle η_2 in purple.
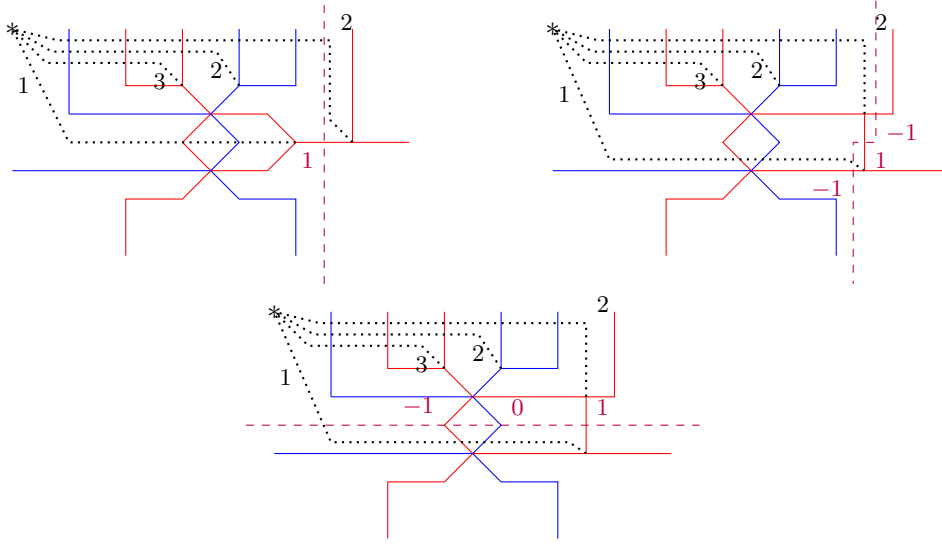
FIGURE 88. Mutating the initial weave for $m(5_2)$.

According to Figure 85, the weave mutation at the short I-cycle δ now produces the weave on the top right of Figure 88, together with a new relative cycle η' as drawn. Now we can homotope η' to the purple relative cycle representative, as drawn in the second row of Figure 88. By applying Proposition 4.29 to the relative cycle representative in the last picture, we conclude that

$$A_{\eta'} = -\frac{v_0 \wedge \alpha_3}{\det} \Big/ \frac{v_3 \wedge \alpha_0}{\det} = -\frac{v_0 \wedge w \wedge v_3}{v_3 \wedge v_4 \wedge v_0} = \frac{v_0 \wedge v_3 \wedge w}{\det}.$$

In addition, Equation (4.10) implies that

$$A_2' = A_{\eta_2'} = A_{\eta'} A_{\eta_1 - \eta_2} = \frac{v_0 \wedge v_3 \wedge w}{\det} \frac{\det}{\alpha_0 \wedge v_2} = \frac{v_0 \wedge v_3 \wedge w}{\alpha_0 \wedge v_2},$$

which is precisely the result obtained in the cluster computation above. $\qquad\square$

The 4th Legendrian $m(7_2)$. Let us consider the (-1)-closure of the following 4-stranded positive braid word

$$\beta = \sigma_1\sigma_3\sigma_3\sigma_3\sigma_2\sigma_2\sigma_2\sigma_3\sigma_1\sigma_2\sigma_2\sigma_2\sigma_3(\sigma_3\sigma_2\sigma_2\sigma_3\sigma_1\sigma_2\sigma_2\sigma_3) = \sigma_1\sigma_3\sigma_3\sigma_3\sigma_2\sigma_2\sigma_2\sigma_3\sigma_1\sigma_2\sigma_2\sigma_2\sigma_3\Delta,$$

This (-1)-closure describes a max-tb Legendrian representative $\Lambda_4(m(7_2)) \subset (\mathbb{R}^3, \xi_{st})$ in the smooth type $m(7_2)$, which in fact corresponds to the fourth entry in the part of the Legendrian Knot Atlas [CN13] dedicated to $m(7_2)$. The second entry represents the Legendrian isotopy class of the (lift of the) alternating strand diagram of 16 (upper right). The first and third entries do not admit binary Maslov indices, and thus do not arise as alternating strand diagrams of GP-graphs. We do not know whether this fourth Legendrian representative of $m(7_2)$ is the alternating strand diagram of a GP-graph; notice that it does admit a binary Maslov index, by virtue of being a (-1)-closure of a positive braid. We iterate our thanks L. Ng for fruitful discussions with regards to Legendrian knots. Now, regardless of whether it arises from a GP-graph, the techniques and principles discovered in this article still apply, as follows.

Let us set T to be a unique mark point, since $\Lambda_4(m(7_2))$ is a knot. An embedded exact Lagrangian filling of $\Lambda_4(m(7_2))$ is smoothly a once-punctured torus and since $H_1(L) = H_1(L, T) = \mathbb{Z}^2$, we can forgo keeping track of the unique marked point in T, which we do. Note that the Demazure product

$$Dem(\sigma_1\sigma_3\sigma_3\sigma_3\sigma_2\sigma_2\sigma_3\sigma_1\sigma_2\sigma_2\sigma_3) = Dem(\sigma_1\sigma_3\sigma_2\sigma_3\sigma_1\sigma_2\sigma_3) = \sigma_1\sigma_3\sigma_2\sigma_3\sigma_1\sigma_2 = w_0$$

is indeed w_0, and thus the Legendrian admits an embedded exact Lagrangian filling by [CN21, Proposition 2.8]. In this case, there is no GP-graph $\mathbb{G}$ from which we can build a weave $\mathfrak{w}(\mathbb{G})$ that represents a Lagrangian filling. Nevertheless, it is perfectly possible to do this by hand. In fact, there is an algorithmic procedure to produce weave fillings of (-1)-closures, using Demazure weaves, as explained in [CGGS20, Section 4.7]. In either case, we have drawn a weave filling in Figure 89. Following the ideas and results from the manuscript, this dictates that the (mutation equivalence class of the) quiver ought to be the A_2-quiver, with γ_1 representing a mutable vertex and γ_2 representing a frozen vertex. The corresponding microlocal merodromies are associated to η_1, η_2, which we now compute in coordinates.
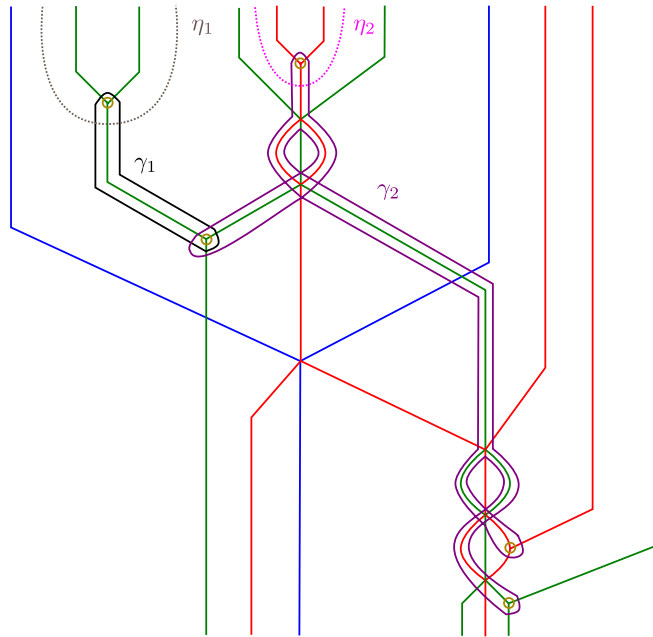
FIGURE 89. A weave $\mathfrak{w}$ representating an embedded exact Lagrangian filling $L(\mathfrak{w})$ of the Legendrian $\Lambda_4(m(7_2))$, which is topologically a once-punctured torus. A basis of $H_1(L(\mathfrak{w})) = \mathbb{Z}^2$ is given by γ_1, γ_2, of which the 1-cycle γ_1 is $\mathbb{L}$-compressible. A dual basis of relative cycles for the dual lattice $H_1(L(\mathfrak{w}) \setminus T, \Lambda_4(m(7_2)) \setminus T) = \mathbb{Z}^2$ is given by the two relative 1-cycles η_1, η_2.

Let us introduce coordinates $z_1, \ldots, z_{17} \in \mathbb{C}$, corresponding to the crossings of β left to right. Following the notation introduced in [CN21, Section 5], we can use the path matrix $P_\beta \in GL(4, \mathbb{C})$ associated to the braid β for $\Lambda_4(m(7_2))$ in order to find implicit equations for the decorated moduli $\mathfrak{M}(\Lambda_4(m(7_2)), T)$, as it coincides with the augmentation variety. Explicitly, by [CN21, Proposition 5.2], this moduli is cut out by the entry-wise equations $P_\beta + \mathrm{Id}_4 = 0$. The microlocal monodromies associated to η_1 and η_2 are

$$A_{\eta_1} = z_2, \quad A_{\eta_2} = -z_5,$$

both of them being understood as the restrictions of the ambient functions $z_2, z_5 \in \mathbb{C}[z_1, \ldots, z_{17}]$ to $\mathfrak{M}(\Lambda_4(m(7_2)), T) \subset \mathbb{C}^{17}$. Still in the notation of [CN21], we set $t_1 = t_2 = t_3 = 1$ and $t_4 = -1$ with regards to the marked point in T. Note that the general theory states that A_{η_2} is a frozen variable, and thus our first computation in this explicit example should verify that the ambient regular function z_5 is non-vanishing on the affine subvariety $\mathfrak{M}(\Lambda_4(m(7_2)), T)$. Indeed, consider the set $\mathbb{P}$ of the 16 polynomial entries of $P_\beta + \mathrm{Id}_4$. Then a Gröbner basis for $\mathbb{C}[z_1, \ldots, z_{17}]/\langle \mathbb{P}, z_5 \rangle$, whose spectrum is the intersection $\mathfrak{M}(\Lambda_4(m(7_2)), T) \cap \{z_5 = 0\}$, can be computed to be $\{1\}$, e.g. using Mathematica. In consequence, the intersection is empty and the function z_5 does not vanish on $\mathfrak{M}(\Lambda_4(m(7_2)), T)$. The second computation is that A_{η_1}, the only mutable variable, does indeed mutate to a regular function. Given the quiver, the mutated variable is

$$A'_{\eta_1} = \frac{1 + A_{\eta_2}}{A_{\eta_1}} = \frac{1 - z_5}{z_2}.$$

Now, $A_{\eta_1} = z_2$ certainly vanishes on $\mathfrak{M}(\Lambda_4(m(7_2)), T)$, but we can verify that, when restricted to $\mathfrak{M}(\Lambda_4(m(7_2)), T)$, if z_2 vanishes, then so does $1 - z_5$. Indeed, a Gröbner basis for the subvariety

$$\mathfrak{M}(\Lambda_4(m(7_2)), T) \cap \{z_2 = 0\} \subset \mathbb{C}^{17}$$

is computed to be

$$\{z_{16} - 1, z_{15} - 1, z_{14} - 1, z_{13}, z_{12} + 1, z_{11} - 1, z_{10} - 1, z_9 + 1, z_8, z_7, z_6 + 1, {\color{blue}z_5 - 1}, z_4, z_3 - z_{17}, z_1 - 1\}.$$

This Gröbner basis contains the element $z_5 - 1$, highlighted in blue above, and we conclude that A'_{η_1} is a regular function on $\mathfrak{M}(\Lambda_4(m(7_2)), T)$, as desired.

As a last verification, let us explicitly compute that the intersection

$$T_{z_2, z_5} := \mathfrak{M}(\Lambda_4(m(7_2)), T) \cap \{z_2 \neq 0, z_5 \neq 0\} \subset \mathbb{C}^{17}$$

is indeed an algebraic torus $(\mathbb{C}^*)^2$ whose coordinate ring can be identified with the Laurent polynomial ring $\mathbb{C}[z_2^{\pm 1}, z_5^{\pm 1}]$. A Gröbner basis for the coordinate ring of this intersection T_{z_2, z_5} is

$$\left\{ \frac{z_{16}}{z_2} + z_{17} - \frac{1}{z_2}, z_{15} - 1, z_{14} - 1, z_{13}, z_{17}z_{12} - \frac{z_{12}}{z_2} - \frac{1}{z_2}, z_{12}z_{16} + 1, z_{11} - z_{16}, \right.$$

$$z_{10} - z_{16}, z_9 - z_{12}, z_8, z_7, z_6 - z_{12}, z_5 - z_{16}, \frac{z_4}{z_2} + z_{16}, -z_{16}^2 + z_{16} + z_4 z_{17},$$

$$\left. z_3 + \frac{z_{12}}{z_2} + \frac{1}{z_2}, z_2 - z_4 z_{12}, z_1 - z_{16} \right\},$$

which readily tells us that $z_7 = z_8 = z_{13} = 0$ and $z_{14} = z_{15} = 1$ when we restrict to T_{z_2, z_5}. Also, we learn that

$$z_1 = -z_6^{-1} = -z_9^{-1} = z_{10} = z_{11} = -z_{12}^{-1} = z_{16} = z_5$$

and thus all these coordinates are graphs on the coordinate z_5. The remaining variables are z_2, z_3, z_4. Since the element $z_4 z_2^{-1} + z_{16}$ is in the basis, or equivalently $z_2 - z_4 z_{12}$, we have that $z_4 = -z_2 z_5$, and thus z_4 is graphical on the pair z_2, z_5. The presence of $z_3 + z_{12} z_2^{-1} + z_2^{-1}$ in the basis implies that $z_3 = -z_2^{-1}(1 - z_5^{-1})$. This concludes that $z_1, z_3, z_4, z_6, z_7, \ldots, z_{17}$ are graphical on $z_2^{\pm 1}, z_5^{\pm 1}$ in the above intersection T_{z_2, z_5}, and that z_2, z_5 have no relations among

104

them: T_{z_2,z_5} is algebraically isomorphic to $\mathbb{C}[A_{\eta_1}^{\pm 1}, A_{\eta_2}^{\pm 1}] \cong \mathbb{C}[z_2^{\pm 1}, z_5^{\pm 1}] \cong (\mathbb{C}^*)^2$ as desired. A similar computation shows that $\mathbb{C}[(A'_{\eta_1})^{\pm 1}, A_{\eta_2}^{\pm 1}]$ is an algebraic torus $(\mathbb{C}^*)^2$ as well. $\square$

5. Cluster DT Transformations for Shuffle Graphs

The cluster Donaldson-Thomas (DT) transformation is a cluster variety automorphism that manifests the Donaldson-Thomas invariants of a 3d Calabi-Yau category associated with the cluster ensemble [Kel17, KS10, GS18]. Combinatorially a cluster DT transformation can be captured by a reddening sequence. In this section we prove Corollary 1.4, i.e. we focus on the cluster varieties $\mathcal{M}_1(\Lambda)$ associated with shuffle graphs and in particular show that their cluster DT transformation is the composition of a Legendrian isotopy and a contactomorphism of $(\mathbb{R}^3, \xi_{\mathrm{st}})$.

5.1. Initial Quivers of Shuffle Graphs.
Let us first prove features of the initial quivers associated with shuffle graphs. From now onward, we assume without loss of generality that shuffle graphs have all vertical edges with a black vertex on top.

Proposition 5.1. *Let $\mathbb{G}$ be a shuffle graph and $f \subset \mathbb{G}$ a face. Then the sugar-free hull $\mathbb{S}_f$ can have a staircase pattern on at most one of its sides.*

Proof. If a sugar-free hull has staircase patterns on more than one of its sides, then somewhere in this sugar-free hull we must have two opposing staircases that look like
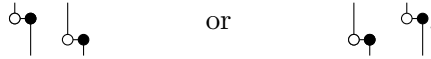
In either case, the horizontal lines containing the horizontal edges above violate Definition 2.18. $\square$

Corollary 5.2. *Let $\mathbb{G}$ be a shuffle graph. Then all its sugar-free hulls must be in one of the following three shapes:*

Proposition 5.3. *Let $\mathbb{G}$ be a shuffle graph. If $\mathbb{S}_f$ and $\mathbb{S}_g$ are two sugar-free hulls and $\mathbb{S}_f \subset \mathbb{S}_g$, then there is no arrow between their corresponding quiver vertices $Q(\mathbb{G})$, i.e., $\langle \partial \mathbb{S}_f, \partial \mathbb{S}_g \rangle = 0$.*

Proof. If $\mathbb{S}_f$ and $\mathbb{S}_g$ do not share boundaries, then $\langle \partial \mathbb{S}_f, \partial \mathbb{S}_g \rangle = 0$. If $(\partial \mathbb{S}_f) \cap (\partial \mathbb{S}_g) \neq \emptyset$, then based on their possible shapes listed in Corollary 5.2, we see that $(\partial \mathbb{S}_f) \cap (\partial \mathbb{S}_g)$ must be the union of a consecutive sequence of edges. By going over all possibilities of having opposite colors at the two end points, we deduce that each possibility will always cut $\mathbb{S}_g$ into smaller sugar-free regions, making $\mathbb{S}_g$ no longer a sugar-free hull. Thus, the two end points of this union must be of the same color. Note that the pairing $\langle \partial \mathbb{S}_f, \partial \mathbb{S}_g \rangle$ can be computed by summing over contributions from the bipartite edges in $(\partial \mathbb{S}_f) \setminus (\partial \mathbb{S}_g)$: since the two end points of $(\partial \mathbb{S}_f) \setminus (\partial \mathbb{S}_g)$ are the same as the two end points of $(\partial \mathbb{S}_f) \cap (\partial \mathbb{S}_g)$, we can conclude that the contributions from the bipartite edges must cancel each other out, leaving $\langle \partial \mathbb{S}_f, \partial \mathbb{S}_g \rangle = 0$ as a result. $\square$

By Definition 2.18, a shuffle graph $\mathbb{G}$ with n horizontal lines is equipped with a permutation $\sigma \in S_n$. Based on the permutation σ, we decompose the initial quiver $Q(\mathbb{G})$ as follows.

Definition 5.4. For each integer m with $1 \leq m < n$, we define $\sigma^{-1}[m,n]$ to be the preimage of the $(n - m + 1)$-element set $[m,n]$. We order elements in $\sigma^{-1}[m,n]$ according to the ordinary linear order on natural numbers. We say that (i,j) form a *level* if $i < j$ in $\sigma^{-1}[m,n]$ for some m and there is no $k \in \sigma^{-1}(m,n)$ such that $i < k < j$. We say a quiver vertex is

on level (i, j) if its corresponding sugar-free hull is sandwiched between the ith and the jth horizontal lines.

If $\sigma^{-1}(m) = k$ and there exists $i < j$ in $\sigma^{-1}[m+1, n]$ such that $i < k < j$, then there can be sugar-free hulls on level (i, j) containing sugar-free hulls on levels (i, k) and (k, j). It is possible to visualize this phenomenon as a branching on the quiver $Q(\mathbb{G})$: the *main branch* contains sugar-free hulls on levels (i, k) and (k, j) and the *side branch* contains sugar-free hulls on level (i, j). See Figure 90 (right) and 91. Note that such a branching may happen multiple times, with the side branch of the former branching becoming the main branch of the next. Figure 91 illustrates this for a shuffle graph associated with the permutation $\sigma = [4123]$, with two branchings on each side.
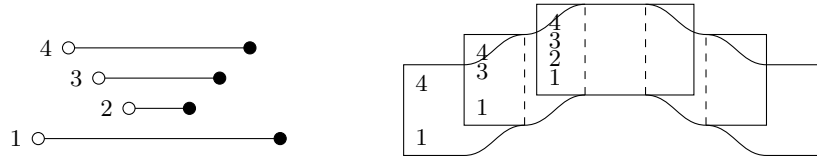
FIGURE 90. The left picture shows the relative lengths of horizontal lines in a shuffle graph $\mathbb{G}$ associated with σ. The right picture describes the branching of the quiver $Q(\mathbb{G})$; the collection of numbers on each branching records the horizontal lines that define the levels for quiver vertices on that branch.
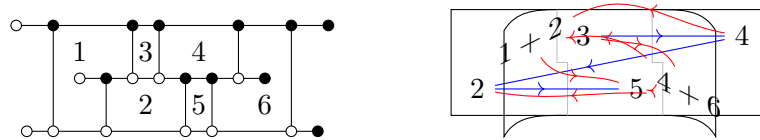
FIGURE 91. Example of the branching phenomenon of the quiver of a shuffle graph. The blue arrows lie on the main branch (the plabic fence part). The red arrows go between the side branches and the main branch

5.2. **Reflection Moves.** In order to geometrically construct the DT-transformations for general shuffle graphs, we now generalize the left and right reflection moves introduced in [SW19].

Consider a Type 2 weave column with an outgoing weave line s_i on one side (top or bottom). By using weave equivalences, we can extend this outgoing weave line inward, penetrating through the weave column and forming a trivalent weave vertex on the other side with color s_{n-i}. If in addition, either of the two horizontal weave lines incident to the new trivalent weave vertex happens to be outgoing as well, then we can homotope the weave locally so that the local picture becomes a Type 2 weave column with an outgoing weave line s_{n-i} on the other side.
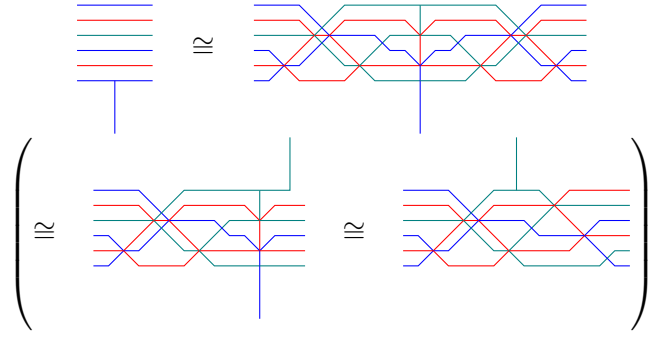
FIGURE 92. The first two pictures are an example of the local penetration move. If there is an additional weave equivalence that turns the 2nd picture into the 3rd one with an outgoing weave line on the top, then we can use it to turn the 3rd one back to a Type 2 weave column again.

These weave equivalences are more general than reflection moves for rainbow closures in *loc. cit.*. In our upcoming construction of cluster DT transformations, we will apply this weave equivalence to lollipops. For example, suppose b is a black lollipop on the ith horizontal line in a grid plabic graph $\mathbb{G}$. Take the first vertical edge e' with a black vertex on the ith horizontal line as we search from right to left starting from b. Suppose the other vertex of e' lies on the jth horizontal line whose right end point lies to the right of b, and suppose there are no more vertical edges (of either pattern) between the ith and jth horizontal lines to the right of e'. Then then the reflection move can be used to turn e' into its opposite pattern; a side-effect is that this move would also switch the portions of the ith and jth horizontal lines on the right side of e', resulting in a possibly non-planar bicolor graph. Below is an example of such a move done on a plabic graph with three horizontal lines. A similar move can be applied to a white lollipop w, and the vertical edge is found by scanning rightward from w.
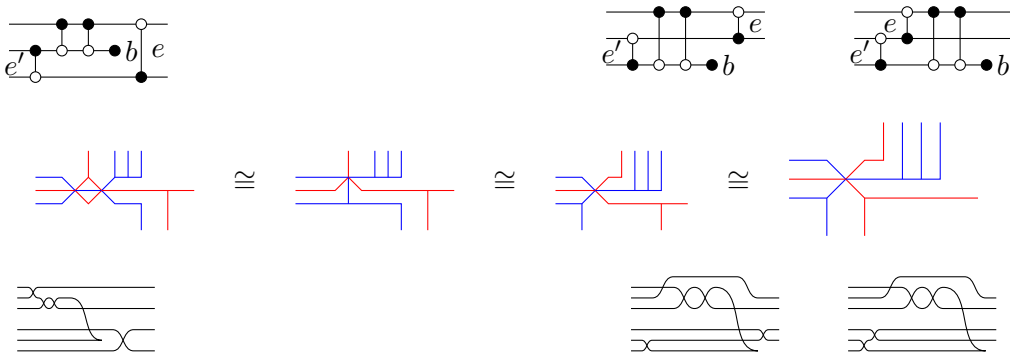
FIGURE 93. Example of a reflection move near a black lollipop.

Since such a move can potentially destroy planarity to the right of edge e', sugar-free hulls no longer make sense there. However, if the part of the quiver corresponding to the region on the right of edge e' does not get involved in the current iterative step, then this does not affect the construction of the cluster DT transformations. Moreover, the reflection move can enable us to pass a vertical edge through an obstructing lollipop. In Figure 93, the lollipop b is preventing the vertical edge e from moving to the left in the left picture; after the reflection move, we can now move e through the lollipop b; even better, the edge e is now contained within a subgraph that is a plabic fence, for which we know a recursive procedure to construct the cluster DT transformation [SW19].

In the front projection, the reflection move is a Legendrian RII move that pulls out a cusp, which is a Legendrian isotopy. Since the reflection move is a weave equivalence, the quiver does not change under such a move. Note that if e' is the only the vertical edge present in each of the bicolor graphs in Figure 93, then this reflection move recovers the right reflection move in [SW19]. (A similar picture can be drawn for left reflection moves.)

By realizing the reflection moves as Legendrian RII moves on the front projection, we can define a Legendrian isotopy on Legendrian links associated with shuffle graphs. By construction, between the top region and the bottom region of the initial weave associated with a shuffle graph, the outgoing weave lines inside one of them is always just a half twist. In the front projection, this region can be untangled into a collection of parallel horizontal lines. Thus, we can clockwise rotate every crossing in the other region one-by-one to this region using just reflection moves (and homotopy) on the front projection. We call this Legendrian isotopy the *half Kálmán loop* $K^{1/2}$. It is not a Legendrian loop and $K^{1/2}$ does not automatically give rise to an automorphism on $\mathcal{M}_1(\Lambda)$: it only gives rise to an isomorphism $K^{1/2} : \mathcal{M}_1(\Lambda) \to \mathcal{M}_1(\Lambda')$, where Λ' is the image of Λ under the Legendrian isotopy $K^{1/2}$. In order to make this into an automorphism, we need the involution t induced from the strict contactomorphism $t : (x, y, z) \mapsto (-x, y, -z)$ on $\mathbb{R}^3$. By Proposition 4.1, we see that this strict contactomorphism reverse all maps in the quiver representation, which implies that we need to dualize all vector spaces and take transpositions of all the maps. Note that this coincides with the definition of the transposition map $t : \mathcal{M}_1(\Lambda') \to \mathcal{M}_1(\Lambda)$ in [SW19]. All parallel transportation maps are now dualized as well but the microlocal monodromies and microlocal merodromies remain unchanged and therefore t preserves the cluster structure and is a cluster isomorphism. We denote $\mathrm{DT} := t \circ K^{1/2} = K^{1/2} \circ t$ as our candidate for the cluster Donaldson-Thomas transformation for shuffle graphs.

5.3. **Edge Migration in a Plabic Fence.** Besides the reflection moves, we also need to move vertical edges through regions that locally look like plabic fences, and cluster mutations are needed for this process. In this subsection, we will discuss these moves and prove some basic result about the color change of quiver vertices (green vs. red). We begin with a quick review of the meaning of vertex colors, green and red, in a quiver. Fix an initial quiver Q with no frozen vertices. We construct a framed quiver $\tilde{Q}$ from Q by adding a frozen vertex i' for every vertex i of Q, together with a single arrow pointing from i to i'. Note that by construction, the exchange matrix of $\tilde{Q}$ is $\tilde{\epsilon} = \begin{pmatrix} \epsilon & \mathrm{id} \\ -\mathrm{id} & 0 \end{pmatrix}$ where ϵ is the exchange matrix of Q.

For any mutation sequence $\mu_{\mathbf{i}}$ on the quiver Q, we can apply the same mutation sequence $\mu_{\mathbf{i}}$ to $\tilde{Q}$ and get a new quiver $\tilde{Q}' := \mu_{\mathbf{i}}(Q)$. Note that the unfrozen part of $\tilde{Q}'$ is identical to $Q' := \mu_{\mathbf{i}}(Q)$. A remarkable property of $\tilde{Q}'$ is that for any unfrozen vertex i, $\tilde{\epsilon}'_{ij'}$ is either non-negative or non-positive for all framing frozen vertices j': this is known as the *sign coherence* phenomenon of c-vectors in cluster theory [DWZ10, GHKK18]. We say a vertex i in Q' is *green* if $\tilde{\epsilon}'_{ij'}$ is non-negative for all framing frozen vertices j' and a vertex i in Q' is *red* if $\tilde{\epsilon}'_{ij'}$ is non-positive for all framing frozen vertices j'. For a given initial quiver Q (all of whose vertices are green), if a mutation sequence $\mu_{\mathbf{i}}$ turns every quiver vertex red, then we say $\mu_{\mathbf{i}}$ is a *reddening sequence*. If additionally a reddening sequence $\mu_{\mathbf{i}}$ only mutates at green vertices, then we say that $\mu_{\mathbf{i}}$ is a *maximal green sequence*. For any fixed initial seed, the cluster Donaldson-Thomas transformation can be captured combinatorially by a reddening sequence [Kel17, GS18]. Thus, it is important to keep track of color change of quiver vertices as we perform cluster mutations.

Let us now consider a plabic fence $\mathbb{G}$. By construction, the initial quiver $Q = Q(\mathbb{G})$ is a planar quiver with one unfrozen vertex for each face in $\mathbb{G}$, and the arrows in Q are drawn in a way such that they form a clockwise cycle around a neighboring group of white vertices

and form a counterclockwise cycle around a neighboring group of black vertices. If we have two adjacent vertical edges of opposite patterns on the same level, we can exchange them by doing a mutation at the quiver vertex corresponding to the face they bound: this is just the square move. If we have two adjacent vertical edges of opposite patterns not on the same level, then we can slide them through each other without doing any mutation on the quiver. See Figure 94.
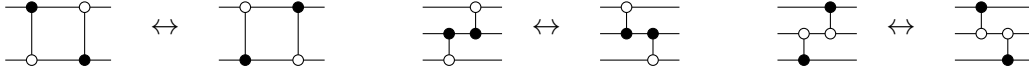

FIGURE 94. Left: the square move in a plabic fence. Middle and right: sliding vertical edges of opposite patterns on different levels through each other.

On a Legendrian weave, the sliding of edges corresponds to a weave equivalence, whereas the square move can be described by a weave mutation along a long I-cycle, which can be locally described by the following movie.
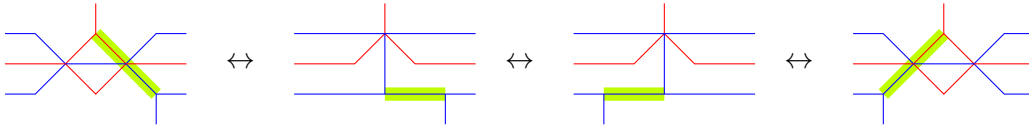
FIGURE 95. Square move in terms of Legendrian weaves: the first and last moves are weave equivalences; the middle move is a weave mutation.

A maximal green sequence on $Q(\mathbb{G})$ can be constructed recursively as follows:

- Take the right most vertical edge e of the ○ pattern and change it to the opposite pattern.
- Move this newly changed vertical edge e to the left, passing all remaining vertical edges of the ○ pattern.

This iterative process terminates when we run out of vertical edges with a black vertex on top. Note that all mutations occur in this maximal green sequence come from square moves; this will not be the case for general shuffle graphs. Lastly, here is a result that we will need in the next subsection:

Lemma 5.5 ([SW19, Proposition 4.6]). *Each square move turns the mutating vertex from green to red and turns the vertex directly to the right of the mutating vertex (if such a vertex exists) from red back to green. As a result, at the end of each iteration of moving a vertical edge e to the left, the left most quiver vertex on the level of e turns red, while the color of every other vertex remain the same.*

5.4. **DT Transformations for Shuffle Graphs.** Let us construct the cluster DT transformations for shuffle graphs. By Condition (1) of Definition 2.18, if a vertical edge can be placed between the ith and jth horizontal lines with $|i - j| > 1$, then there must be disjoint two continuous regions we can place vertical edges between them, with one on the left and the other one on the right. Let us call them the *left region* and the *right region*, respectively. If $|i - j| = 1$, then there is only one continuous region where we can place vertical edges between the two horizontal lines. The main strategy is to go through all vertical edges of $\mathbb{G}$ one by one from right to left. For each vertical edge e we do one of the following depending, on its location in $\mathbb{G}$.

(I) If e lies on level $(i, i+1)$:

(I.1) Apply a reflection move to change the pattern of e to .

(I.2) Move e to the left through all edges incident to the ith or $(i+1)$st horizontal lines. Note that a cluster mutation occurs whenever we exchange e and a edge at the same horizontal level.

(II) If e lies in the right region of level (i,j) with $j - i > 1$:

(II.1) Apply a reflection move to change the pattern of e to .

(II.2) Move e all the way into the left region between the ith and the jth horizontal lines, and through all edges incident to the ith or jth horizontal lines.

Remark 5.6. There is the following subtlety in step (II.2). Since $j - i > 1$, when moving e to the left, we will encounter $j - i - 1$ many black lollipops. Each time when we encounter a black lollipop, we will try to apply a move similar to Figure 93. If such a move can be applied, then we will get another vertical edge of the same pattern as e, and we need to move this newly changed edge along with e as a group to the left. Moreover, this newly changed edge itself may encounter a black lollipop, too, and consequently introduce another vertical edge into the moving group. This moving group eventually will reach the other side, and we need a way to recover e back as a vertical edge on level (i,j). This can be done by a move mirror to that of Figure 93:
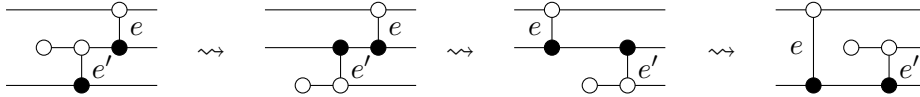
FIGURE 96. Two reflection moves to cover the vertical edge e.

Note that all moves in Figure 96 correspond to weave equivalences and hence no cluster mutations occur. After the vertical edge e is recovered, we can send the auxiliary vertical edge e' back to where it was before, and then do another reflection move to restore the pattern of e'. Note that upon restoring the location and the pattern of e', the non-planarity caused by the earlier reflection move (right picture of Figure 93) will be cancelled, and we return to a grid plabic graph after the iteration.

Remark 5.7. There is also a possibility that although there is a black lollipop b in the way, no vertical edge e' is found and hence it is not possible to perform the move in Figure 93. We claim that in this case we can directly move e through the obstructing horizontal line directly without the need of any weave (cluster) mutations. This follows from the fact that if no vertical edge e' is present, then the incoming weave line corresponding to the gap where e' should have been does not need to be tangled in the weave, and hence we can perform a weave equivalence to move the vertical edge e through. See Figure 97.
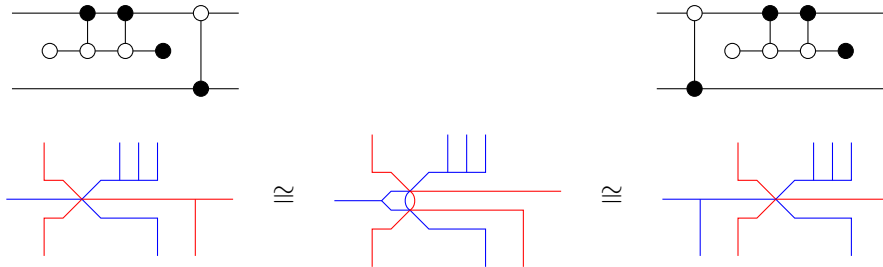

FIGURE 97. Example of a special case where the edge e' is absent

(III) If e lies in the left region between the ith and the jth horizontal lines with $|i - j| > 1$:

(III.1) Move e all the way to the right region between the ith and the jth horizontal lines so that it becomes the right most vertical edge in the plabic graph.

(III.2) Apply a reflection move to change the pattern of e to ⌐.

Note that in this case, we need to move the vertical edge e to the right before changing its pattern. But since we are going through vertical edges in $\mathbb{G}$ one by one from right to left, by the time we get to e, all vertical edges to its right must be of the ⌐ pattern already. Thus, moving e, which is of the ⌐ pattern, to the right through ⌐ edges, is completely mirror to step (II.1) before. Below are pictures depicting the reflection moves we need to perform during this process.
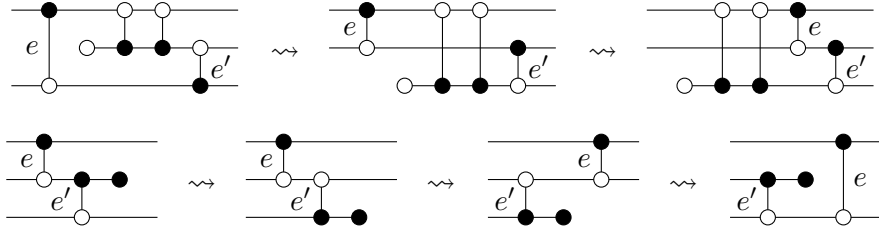
FIGURE 98. Reflection moves in the (III.1) step: the top row is an example of how to shrink e to a shorter vertical edge, and the bottom row is an example of how to recover e after moving through a horizontal line.

We can now conclude the following result:

Theorem 5.8. *Let $\mathbb{G}$ be a shuffle graph. Then $\mathrm{DT} = t \circ K^{1/2}$ is the cluster Donaldson-Thomas transformation on $\mathcal{M}_1(\Lambda)$.*

Proof. Since t is a cluster isomorphism, it suffices to prove that $K^{1/2}$ gives rise to a reddening sequence. The vertices of the quiver $Q(\mathbb{G})$ are grouped into regions (Figures 91 and 90): we claim that after each iterative step (I) and step (II) of moving an edge e on level (i, j), the left most green vertex of level (i, j) turns red, and after each iterative step (III) of moving an edge e on level (i, j), the right most green vertex of level (i, j) turns red. Indeed, the case (I) follows from Lemma 5.5 directly; so it remains to consider (II) and (III).

Let us consider (II) first. In the process of moving a ⌐ edge e on level (i, j) to the left, before we encounter any lollipop, the quiver vertices would change according to Lemma 5.5, turning from green to red when a square move occurs and then turning back to green in the next square move. Let us now consider what happens when we encounter a black lollipop at the kth horizontal line (with $i < k < j$). If we are in the situation of Remark 5.7, then we can directly jump through the whole kth horizontal line without any cluster mutations, and Lemma 5.5 will continue to take care of the rest. It thus remains to consider what happens when we need to do moves according to Remark 5.6. Note that since the moves in Figure 93 do not induce any cluster mutations, there is no change to the quiver itself. However, the way we branch the quiver is different: the sugar-free hull to the left of edge e was non-rectangular before the moves in Figure 93 but it becomes rectangular after the moves; thus, the corresponding quiver vertex was on the side branch before the moves and relocates itself to the main branch after the moves.

After this quiver vertex is relocated to the main branch (which is a quiver of a plabic fence), we can make use of Lemma 5.5 again. Note that we need to move e as well as the auxiliary

edge e' together to the left as a group, and in that process, there is still a possibility of introducing more edges to that left-moving group. Nevertheless, by induction it is enough to consider what happens when the two-member group e' and e reach the left white lollipop of the kth horizontal line. By Lemma 5.5 we know that both the quiver vertex v to the right of e and the quiver vertex v' to the right of e' have turned red. Next, under the moves in Figure 96, we restore e to a vertical edge on level (i, j) without any cluster mutations. Finally, we need to send e' back to where it was before, and this process reverses the mutations we did on the level of e': the quiver vertices on that level will turn from green to red and then back to green again one-by-one[22]; in the end all quiver vertices on the same level as e' are restored back to green. The iterative step can now continue further to the left on level (i, j), and the color change will again follow Lemma 5.5.

The case (III) is essentially (II) in reverse. Suppose we are moving a ⬥ edge e on level (i, j), and suppose the first white lollipop it encounters is on the kth horizontal line. Let v be the quiver vertex to the right of e, which is the right-most green on level (i, j) at the beginning of the iterative step. Under the moves in the top row of Figure 98, we move v into level (k, j), and obtain another vertical edge e' of the ⬦ pattern. Note that the vertex v' to the right of e' is red at this moment. Now we need to move the vertical edges e and e' to the right: first e', then e. For each square move on level (i, k) (the level of e'), the mutating quiver vertex changes from red to green[23] and stays green afterward. On the other hand, for each square move on level (k, j) (the level of e), the mutating quiver vertex turns from green to red and the one in the next mutation (if exists) turns from red to green. As a result, when e and e' gets to the right end of the kth horizontal line, all quiver vertices on the level of e are red and all quiver vertices on the level of e' are green. After the second row of moves in Figure 98 (which do not involve mutations), we need to send e' back to where it was, and that will make the quiver vertices on level of e' undergo another the reverse sequence of color change again, restoring all of them back to red. Of course, if there are move vertices further to the right of e, we need to continue moving e rightward, which will make the remaining quiver vertices to the right on level (i, j) turn from red to green and then back to red again one-by-one. In the end, we see that precisely the quiver vertex v turns red after this iterative step.

In conclusion, since all quiver vertices in $Q(\mathbb{G})$ are either inside the middle region, to the left of a vertical edge in the right region, or to the right of a vertical edge in the left region, we see that after all vertical edges in $\mathbb{G}$ change from ⬥ to ⬦, all quiver vertices would turn red. Therefore $K^{1/2}$ is indeed a reddening sequence. □

Example 5.9. *Let us consider the shuffle graph $\mathbb{G}$ below and follow our procedure to construct its cluster DT transformation. We will keep track of the change on the bicolor graph, the quiver, and the front projection at each step.*
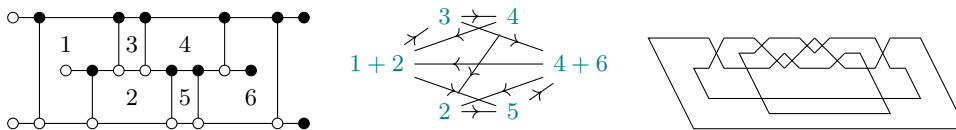

We start by taking the right most vertical edge and changing its pattern to its opposite. This does not do anything to the quiver. In the front projection, this moves the right most crossing

---

[22]Note that we are mutating at red quiver vertices in this process; thus, we are not claiming that the whole mutation sequence is maximal green. On the other hand, such moves are not needed in the case of plabic fences, which is why we can obtain a maximal green sequence.

[23]These mutations are also not green.

from the top to the bottom.
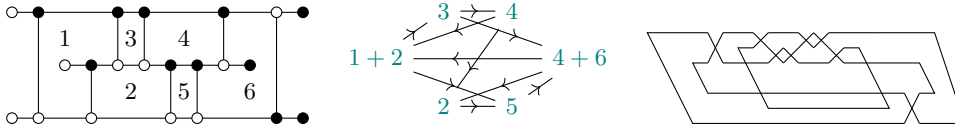
Next we apply a reflection move to the black lollipop of the 2nd horizontal line. This does not cause any mutations, but the quiver vertex $4+6$ now moves to level $(2,3)$ and the quiver vertex 4 moves to level $(1,3)$. In the front projection, there are now two crossings in the bottom region, ready to be moved to the left. From now on, we will label a face only if it is sugar-free, and we label it with its corresponding quiver vertex.
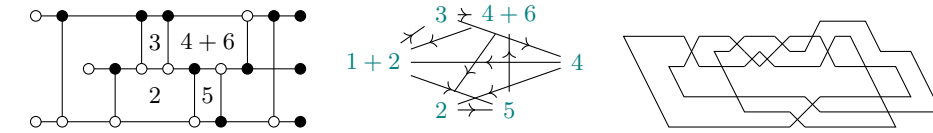

Next we need to move the bottom two crossings to the left through the plabic fence region. The mutation sequence for this is $\mu_3 \circ \mu_{4+6} \circ \mu_2 \circ \mu_5$. This mutation sequence turns the quiver vertices 2 and 3 red.
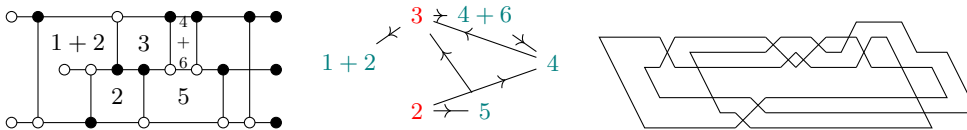

Next we apply two reflection moves similar to those in Figure 96, recovering a vertical edge on level $(1,3)$. No mutation occurs in this step, but the quiver vertex 3 now moves to level $(1,3)$.
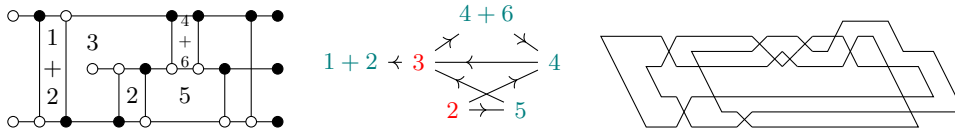
Next we continue moving the ⊶ edge on level $(1,3)$ further to the left and return the ⊶ edge on level $(1,2)$ back to its original place. The corresponding mutation sequence is $(\mu_5 \circ \mu_2) \circ (\mu_{1+2})$. Note that the two mutation subsequences in parentheses commute with each other. As a result, the quiver vertex $1+2$ turns red, but quiver vertices 2 and 3 go back to being green.
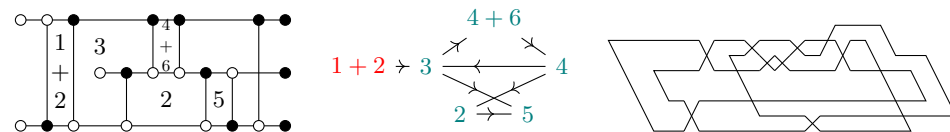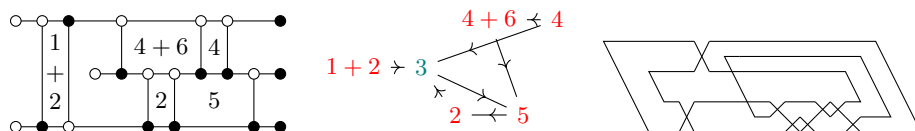

Next we perform a reflection move at the black lollipop on the 1st horizontal line. This does not cause any mutations, but it moves the quiver vertex 4 from level $(1,3)$ back to level $(2,3)$. Note that the bicolor graph becomes planar again after this move.
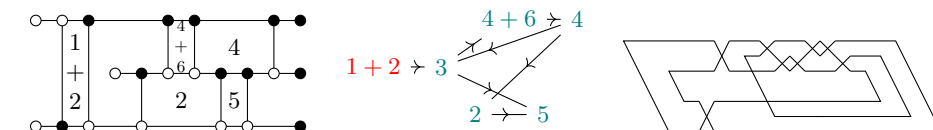

Next we deal with quiver vertices inside the plabic fence region, and turn each of them from green to red by following Lemma 5.5. The mutation sequence is $() \circ () \circ (\mu_4) \circ (\mu_5) \circ (\mu_2 \circ \mu_5) \circ (\mu_{4+6} \circ \mu_4)$, where each pair of parentheses corresponds to an iterative step (the last

113

two iterative steps do not require any mutations). In the front projection, this rotates all crossings on levels $(1,2)$ and $(2,3)$ from the top region to the bottom region.


Next we need to move the right most *edge to the right before flipping it. In order to do that, we need to first apply a reflection move to the white lollipop on the 2nd horizontal line. This does not cause any mutations, but it moves the quiver vertices $1+2$ and 3 to level $(2,3)$ and moves the quiver vertex $4+6$ to level $(1,3)$.*
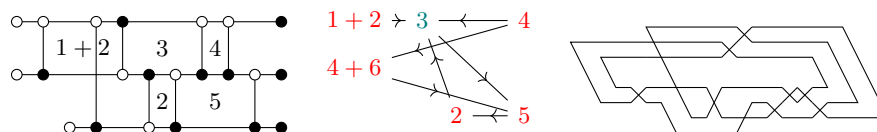

Next we need to move the right most two *edges to the far right. The mutation sequence is $\mu_4 \circ \mu_3 \circ \mu_5 \circ \mu_2$. After this mutation sequence, the quiver vertex 3 turns red, but the quiver vertices 2 and 5 turn back to green.*
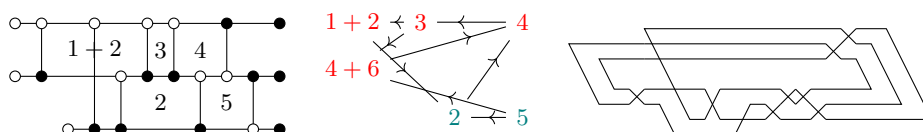
Next we need to do a couple reflection moves similar to those of the second row in Figure 98. This does not cause any mutations, but it moves the quiver vertex 4 to level $(1,3)$.
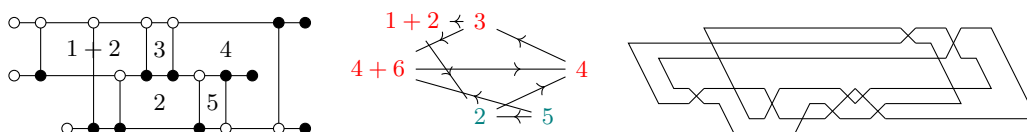
Next we flip the right most *edge and send the auxiliary vertical edge back to where it was, restoring the planarity of the bicolor graph. The mutation sequence is $\mu_2 \circ \mu_5$, which turns both vertices 2 and 5 red. The reflection move after the mutation sequence puts quiver vertices $1+2$ and 4 back to level $(1,3)$ and the quiver vertex $4+6$ back to level $(2,3)$.*
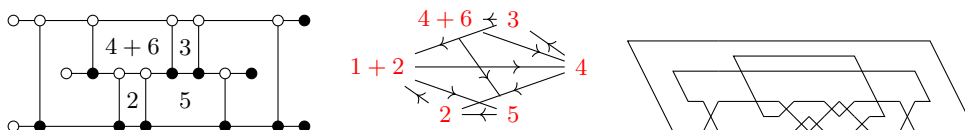
In total, the reddening sequence is

$$\circ \, (\mu_2 \circ \mu_5 \circ \mu_4 \circ \mu_3 \circ \mu_5 \circ \mu_2) \circ () \circ () \circ (\mu_4) \circ (\mu_5) \circ (\mu_2 \circ \mu_5) \circ (\mu_{4+6} \circ \mu_4)$$

$$\circ \, (\mu_5 \circ \mu_2 \circ \mu_{1+2} \circ \mu_3 \circ \mu_{4+6} \circ \mu_2 \circ \mu_5),$$

where we have grouped the mutations of each iterative step inside a pair of parentheses. By depicting the action on the front projection we conclude that this reddening sequence corresponds to the half Kálmán loop $K^{1/2}$. Composing it with the cluster isomorphism t gives the cluster DT transformation on $\mathcal{M}_1(\Lambda(\mathbb{G}))$.

The statement in Corollary 1.4 about the cluster duality conjecture now follows from [GHKK18], as our quivers are full-ranked and a DT-transformation exists. Finally, we remark that the same argument used in [GSW20b] to distinguish infinitely many Lagrangian fillings also works for any shuffle graph whose quiver is mutation equivalent to an acyclic quiver

of infinite type. Indeed, the DT-transformation will be of infinite order and so will be its square, the Legendrian Kálmán loop.

Example 5.10. *Consider the shuffle graph* $\mathbb{G}$ *from Example 5.9. The quiver* $Q(\mathbb{G})$ *is mutation equivalent to an acyclic quiver of infinite type, e.g. consider the mutation sequence* $\mu_5 \circ \mu_{4+6} \circ \mu_2 \circ \mu_4 \circ \mu_3 \circ \mu_{4+6}$. *Thus* $\Lambda(\mathbb{G})$, *which is a max-tb representative in the smooth knot type* 10_{161}, *admits infinitely many non-Hamiltonian isotopic embedded exact fillings. Note that the smooth knot type* 10_{161} *is not a rainbow closure of a positive braid.*

A. Appendix: Quasi-Cluster Structures

This appendix contains the necessary definitions regarding quasi-cluster structures and their associated objects, as needed for the purposes of this article. See also C. Fraser's [Fra16] and M. Sherman-Bennett's [SB21, Section 4.3].

Let us consider a lattice N of finite rank, a saturated sublattice $N^{\mathrm{uf}} \subset N$, and a $\mathbb{Z}$-valued skew-symmetric form $\{\cdot, \cdot\}$ on the lattice N. The triple $(N, N^{uf}, \{\cdot, \cdot\})$ is considered as input data, and it is fixed throughout. By definition, we set $M := N^*$ and define the linear map

$$p^* : N^{\mathrm{uf}} \longrightarrow M$$
$$n \longmapsto \{n, -\}.$$

For instance, in the present article we use $N = H_1(L, T)$, for L a Lagrangian filling and $T \subset \partial L$ a set of marked point at the boundary, $\{\cdot, \cdot\}$ the intersection form, and N^{uf} the sub-lattice spanned by an $\mathbb{L}$-compressing system.

Definition A.1. Consider a triple $(N, N^{\mathrm{uf}}, \{\cdot, \cdot\})$ as above. A basis $\{e_i\}$ of N^{uf} is said to be a *seed* $\mathbf{s}$ associated to $(N, N^{\mathrm{uf}}, \{\cdot, \cdot\})$.

Definition A.2. Given a seed $\mathbf{s}$ and an element $e_k \in \mathbf{s}$, the seed $\mathbf{s}' := \mu_{e_k}\mathbf{s}$ which consists of vectors

$$e_i' = \begin{cases} -e_k & \text{if } i = k, \\ e_i + \{e_i, e_k\}_+ e_k & \text{if } i \neq k, \end{cases}$$

where $\{e_i, e_k\}_+ := \max\left(\{e_i, e_k\}, 0\right)$, is said to be obtained by a *mutation* of the seed $\mathbf{s}$ at the element e_k. Finally, two seeds of the same given fixed data are said to be *mutation equivalent* if they can be obtained from one another via a sequence of mutations, and we denote the family of seeds that are mutation equivalent to $\mathbf{s}$ by $|\mathbf{s}|$.

Consider a mutation equivalence family of seeds $|\mathbf{s}_0|$. For each seed $\mathbf{s} \in |\mathbf{s}_0|$, we consider the two algebraic tori

$$\mathcal{A}_{\mathbf{s}} := \mathrm{Hom}(M, \mathbb{C}^\times) \quad \text{and} \quad \mathcal{X}_{\mathbf{s}} := \mathrm{Hom}(N, \mathbb{C}^\times).$$

which are referred to as the *seed tori* associated to $\mathbf{s}$; the former is said to be of type A and the latter of type X. By construction. M and N are the character lattices of $\mathcal{A}_{\mathbf{s}}$ and $\mathcal{X}_{\mathbf{s}}$ respectively, and elements of M and N respectively define Laurent monomial functions on $\mathcal{A}_{\mathbf{s}}$ and $\mathcal{X}_{\mathbf{s}}$. We adopt the notation of writing these functions as A^m and X^n for $m \in M$ and $n \in N$. In particular, multiplication of these Laurent monomial functions corresponds to addition of elements in the corresponding lattices.

Given any two seeds $\mathbf{s} = \{e_i\}$ and $\mathbf{s}' = \mu_{e_k}\mathbf{s}$ which differ by exactly one mutation, we define the two birational maps

$$\mu_{e_k} : \mathcal{A}_{\mathbf{s}} \dashrightarrow \mathcal{A}_{\mathbf{s}'} \qquad\qquad \mu_{e_k} : \mathcal{X}_{\mathbf{s}} \dashrightarrow \mathcal{X}_{\mathbf{s}'}$$

$$\mu_{e_k}^*(A^m) := A^m \left(1 + A^{p^*\{e_k\}}\right)^{-\langle m, e_k \rangle} \qquad \mu_{e_k}^*(X^n) := X^n \left(1 + X^{e_k}\right)^{-\{n, e_k\}}.$$

By definition, the *quasi-cluster varieties* $\mathcal{A}$ and $\mathcal{X}$ are the algebraic spaces obtained by gluing the (algebraic) seed tori, of the respective types A and X, via the birational maps μ_{e_k} as above, for every pair of mutation adjacent seeds.

Remark A.3. Intuitively, the *quasi* in *quasi*-cluster allows for a controlled ambiguity coming from the frozen variables: one allows certain rescalings of the cluster variables by Laurent monomials in frozen variables. (Scaling preserving exchange ratios.) It is possible to make additional choices so that a *quasi*-cluster variety becomes a cluster variety. Namely, in order to obtain a seed in the theory of cluster algebras (without the *quasi*), we need to extend our seed $\{e_i\}$, which was a basis for N^{uf} to a basis of the entire lattice N; in turn this yields a dual basis $\{f_i\}$ of M. The mutation of the dual basis is given by

$$
f_i' = \begin{cases} -f_k + \sum_j \{e_k, -e_j\}_+ f_j & \text{if } i = k, \\ f_i & \text{if } i \neq k. \end{cases}
$$

For simplicity, let us denote A^{f_i} by A_i and $A^{f_i'}$ by A_i'. Then the first mutation map becomes

$$
\mu_{e_k}^*(A_k') = A_k^{-1} \left(\prod_{j:\{e_k, -e_j\}>0} A_j^{\{e_k, -e_j\}} \right) \left(1 + \prod_j A_j^{\{e_k, e_j\}} \right)^{-\langle -f_k + \sum_j \{e_k, -e_j\}_+ f_j, e_k \rangle}
$$

$$
= A_k^{-1} \left(\prod_{j:\{e_k, e_j\}<0} A_j^{-\{e_k, e_j\}} + \prod_{j:\{e_k, e_j\}>0} A_j^{\{e_k, e_j\}} \right),
$$

and for $i \neq k$,

$$
\mu_{e_k}^*(A_i') = A_i \left(1 + \prod_j A_j^{\{e_k, e_j\}} \right)^{-\langle f_j, e_k \rangle} = A_i,
$$

which are precisely the classical cluster $\mathcal{A}$-mutation formulas. For the cluster $\mathcal{X}$-mutations, just note that

$$
\mu_{e_k}^*(X_k') = X_k^{-1}(1 + X_k)^{-\{e_k, e_k\}} = X_k^{-1}
$$

and for $i \neq k$,

$$
\mu_{e_k}^*(X_i') = X_i X_k^{\{e_i, e_k\}_+} (1 + X_k)^{-\{e_i, e_k\}},
$$

which are the classical cluster $\mathcal{X}$-mutation formulas.

Consider the lattice $M^{\mathrm{fr}} := (N^{\mathrm{uf}})^\perp$, which is the saturated sublattice of M corresponding to frozen cluster $\mathcal{A}$-variables. In particular, M^{fr} defines a frozen subtorus $\mathcal{A}^{\mathrm{fr}} \subset \mathcal{A}_{\mathbf{s}}$ for each seed $\mathbf{s}$, and these frozen subtori can be (and are) glued via the identity map under mutations. Since M^{fr} is saturated, the quotient M/M^{fr} is also a lattice and any seed $\mathbf{s} = \{e_i\}$ gives rise to a dual basis $\{\overline{f}_i\}$ of M/M^{fr}. Let us choose a lift $f_i \in M$ for each $\overline{f}_i$. The subset

$$
M_+ := M^{\mathrm{fr}} + \sum \mathbb{Z}_{\geq 0} f_i
$$

then defines a monoid of global functions on $\mathcal{A}$, which we refer to as the monoid of *cluster monomials* associated with the seed $\mathbf{s}$.

Finally, *quasi*-cluster transformations arise as follows. Suppose $\mathbf{s} = \{e_i\}$ and $\mathbf{s}' = \{e_i'\}$ are two seeds (not necessarily mutation adjacent) in the same mutation equivalent family of seeds and that $\sigma^* : N \to N$ is a lattice isomorphism preserving the skew-symmetric form $\{\cdot, \cdot\}$ and such that $\sigma^*(\mathbf{s}') = \mathbf{s}$. (Let us denote the induced map $M \to M$ also by σ^*.) Then these two dual maps induce the following isomorphisms between the respective seed tori

$$
\sigma : \mathcal{A}_{\mathbf{s}} \longrightarrow \mathcal{A}_{\mathbf{s}'} \qquad\qquad \sigma : \mathcal{X}_{\mathbf{s}} \longrightarrow \mathcal{X}_{\mathbf{s}'}
$$

$$
\sigma^*(A^m) := A^{\sigma^*(m)} \qquad\qquad \sigma^*(X^n) := X^{\sigma^*(n)}.
$$

Since σ^* preserves the skew-symmetric form, σ^* commutes with seed mutation as well, i.e., $\sigma^*(\mu_{e_k}\mathbf{s}') = \mu_{\sigma^*(e_k)}\mathbf{s}$. Thus, the above isomorphisms of algebraic tori also commute with the birational mutation maps and, as a result, we obtain a pair of automorphisms $\sigma : \mathcal{A} \to \mathcal{A}$ and $\sigma : \mathcal{X} \to \mathcal{X}$ between the quasi-cluster varieties. By definition, we refer to any such automorphisms of the quasi-cluster varieties as *quasi-cluster transformations*.

References

[All21]   Dylan G. L. Allegretti. Stability conditions, cluster varieties, and Riemann-Hilbert problems from surfaces. *Adv. Math.*, 380:Paper No. 107610, 62, 2021.

[And19]   Dave Anderson. Effective divisors on Bott-Samelson varieties. *Transform. Groups*, 24(3):691–711, 2019.

[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.

[BB05]   Anders Björner and Francesco Brenti. *Combinatorics of Coxeter groups*, volume 231 of *Graduate Texts in Mathematics*. Springer, New York, 2005.

[BFZ05]   Arkady Berenstein, Sergey Fomin, and Andrei Zelevinsky. Cluster algebras. III. Upper bounds and double Bruhat cells. *Duke Math. J.*, 126(1):1–52, 2005.

[Cas22]   Roger Casals. Lagrangian skeleta and plane curve singularities. to appear *J. Fixed Point Theory App. (Viterbo 60 Volume)*, 2022.

[CG22]   Roger Casals and Honghao Gao. Infinitely many Lagrangian fillings. *Ann. of Math. (2)*, 195(1):207–249, 2022.

[CGGS20]   Roger Casals, Eugene Gorsky, Mikhail Gorsky, and José Simental. Algebraic weaves and braid varieties. arXiv:2012.06931, 2020.

[CGGS21]   Roger Casals, Eugene Gorsky, Mikhail Gorsky, and José Simental. Positroid links and braid varieties. arXiv:2105.13948, 2021.

[CN13]   Wutichai Chongchitmate and Lenhard Ng. An atlas of Legendrian knots. *Exp. Math.*, 22(1):26–37, 2013.

[CN21]   Roger Casals and Lenhard Ng. Braid loops with infinite monodromy on the legendrian contact dga. preprint, 2021.

[Cur14]   Justin Michael Curry. *Sheaves, cosheaves and applications*. ProQuest LLC, Ann Arbor, MI, 2014. Thesis (Ph.D.)–University of Pennsylvania.

[CZ22]   Roger Casals and Eric Zaslow. Legendrian weaves. to appear in *Geometry&Topology*, 2022.

[DWZ10]   Harm Derksen, Jerzy Weyman, and Andrei Zelevinsky. Quivers with potentials and their representations II: applications to cluster algebras. *J. Amer. Math. Soc.*, 23(3):749–790, 2010.

[EHK16]   Tobias Ekholm, Ko Honda, and Tamás Kálmán. Legendrian knots and exact Lagrangian cobordisms. *J. Eur. Math. Soc. (JEMS)*, 18(11):2627–2689, 2016.

[ENV13]   John B. Etnyre, Lenhard L. Ng, and Vera Vértesi. Legendrian and transverse twist knots. *J. Eur. Math. Soc. (JEMS)*, 15(3):969–995, 2013.

[Fao17]   Giovanni Faonte. Simplicial nerve of an $\mathcal{A}_\infty$-category. *Theory Appl. Categ.*, 32:Paper No. 2, 31–52, 2017.

[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.

[FG09]   Vladimir V. Fock and Alexander B. Goncharov. Cluster ensembles, quantization and the dilogarithm. *Ann. Sci. Éc. Norm. Supér. (4)*, 42(6):865–930, 2009.

[FPST17]   Sergey Fomin, Pavlo Pylyavskyy, Eugenii Shustin, and Dylan Thurston. Morsifications and mutations. preprint, 2017.

[Fra16]   Chris Fraser. Quasi-homomorphisms of cluster algebras. *Adv. in Appl. Math.*, 81:40–77, 2016.

[FST08]   Sergey Fomin, Michael Shapiro, and Dylan Thurston. Cluster algebras and triangulated surfaces. I. Cluster complexes. *Acta Math.*, 201(1):83–146, 2008.

[FZ99]   Sergey Fomin and Andrei Zelevinsky. Double Bruhat cells and total positivity. *J. Amer. Math. Soc.*, 12(2):335–380, 1999.

[FZ02]   Sergey Fomin and Andrei Zelevinsky. Cluster algebras. I. Foundations. *J. Amer. Math. Soc.*, 15(2):497–529, 2002.

[FZ03]   Sergey Fomin and Andrei Zelevinsky. Cluster algebras. II. Finite type classification. *Invent. Math.*, 154(1):63–121, 2003.

[GHK15]   Mark Gross, Paul Hacking, and Sean Keel. Birational geometry of cluster algebras. *Algebr. Geom.*, 2(2):137–175, 2015.

[GHKK18]   Mark Gross, Paul Hacking, Sean Keel, and Maxim Kontsevich. Canonical bases for cluster algebras. *J. Amer. Math. Soc.*, 31(2):497–608, 2018.

[GK13]   Alexander B. Goncharov and Richard Kenyon. Dimers and cluster integrable systems. *Ann. Sci. Éc. Norm. Supér. (4)*, 46(5):747–813, 2013.

[GK21]   Alexander Goncharov and Maxim Kontsevich. Spectral description of non-commutative local systems on surfaces and non-commutative cluster varieties. *ArXiv e-prints*, 2021.

[GKS12a]   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.

[GKS12b]   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.

[GL19]   Pavel Galashin and Thomas Lam. Positroid varieties and cluster algebras. to appear Ann. Sci. Ec. Norm. Super., 2019.

[GL21]   Pavel Galashin and Thomas Lam. Positroids, knots, and q, t-Catalan numbers. *Sém. Lothar. Combin.*, 85B:Art. 54, 12, 2021.

[GLS13]   Christof Geiss, Bernard Leclerc, and Jan Schröer. Factorial cluster algebras. *Doc. Math.*, 18:249–274, 2013.

[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.

[GS14]   Stéphane Guillermou and Pierre Schapira. Microlocal theory of sheaves and Tamarkin's non displaceability 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.

[GSV05]   Michael Gekhtman, Michael Shapiro, and Alek Vainshtein. Cluster algebras and Weil-Petersson forms. *Duke Math. J.*, 127(2):291–311, 2005.

[GSV10]   Michael Gekhtman, Michael Shapiro, and Alek Vainshtein. *Cluster algebras and Poisson geometry*, volume 167 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2010.

[GSW20a]   Honghao Gao, Linhui Shen, and Daping Weng. Augmentations, fillings, and clusters. arXiv:2008.10793, 2020.

[GSW20b]   Honghao Gao, Linhui Shen, and Daping Weng. Positive braid links with infinitely many fillings. arXiv:2009.00499, 2020.

[Gui19]   Stephane Guillermou. Sheaves and symplectic geometry of cotangent bundles. *ArXiv e-prints*, 2019.

[HK18]   Paul Hacking and Sean Keel. Mirror symmetry and cluster algebras. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. II. Invited lectures*, pages 671–697. World Sci. Publ., Hackensack, NJ, 2018.

[HR15]   Michael B. Henry and Dan Rutherford. Ruling polynomials and augmentations over finite fields. *J. Topol.*, 8(1):1–37, 2015.

[IN14]   Kohei Iwaki and Tomoki Nakanishi. Exact WKB analysis and cluster algebras. *J. Phys. A*, 47(47):474009, 98, 2014.

[IN16]   Kohei Iwaki and Tomoki Nakanishi. Exact WKB analysis and cluster algebras II: Simple poles, orbifold points, and generalized cluster algebras. *Int. Math. Res. Not. IMRN*, (14):4375–4417, 2016.

[JT17]   Xin Jin and David Treumann. Brane structures in microlocal sheaf theory. arXiv:1704.04291, 2017.

[Kal14]   Adam Kalman. Newton polytopes of cluster variables of type A_n. In *26th International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2014)*, Discrete Math. Theor. Comput. Sci. Proc., AT, pages 137–147. Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2014.

[Kel17]   Bernhard Keller. Quiver mutation and combinatorial DT-invariants. *Discrete Mathematics and Theoretical Computer Science*, 2017.

[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 re-
           sults. In *Mirror symmetry and tropical geometry*, volume 527 of *Contemp. Math.*, pages 55–89.
           Amer. Math. Soc., Providence, RI, 2010.
[LS22]   Thomas Lam and David E. Speyer. Cohomology of cluster varieties, I: Locally acyclic case. *Algebra
           Number Theory*, 16(1):179–230, 2022.
[LT04]   Niels Lauritzen and Jesper Funch Thomsen. Line bundles on Bott-Samelson varieties. *J. Algebraic
           Geom.*, 13(3):461–473, 2004.
[MSB20]   Amal Mattoo and Melissa Sherman-Bennett. Saturation of newton polytopes of type a and d
           cluster variables. e-print at arxiv: 2012.07500, 2020.
[Mul13]   Greg Muller. Locally acyclic cluster algebras. *Adv. Math.*, 233:207–247, 2013.
[Mul14]   Greg Muller. $\mathscr{A} = \mathscr{U}$ for locally acyclic cluster algebras. *SIGMA Symmetry Integrability Geom.
           Methods Appl.*, 10:Paper 094, 8, 2014.
[Nad16]   David Nadler. Wrapped microlocal sheaves on pairs of pants. arXiv:1604.00114, 2016.
[Nei14]   Andrew Neitzke. Cluster-like coordinates in supersymmetric quantum field theory. *Proc. Natl.
           Acad. Sci. USA*, 111(27):9717–9724, 2014.
[NRS+20]   Lenhard Ng, Dan Rutherford, Vivek Shende, Steven Sivek, and Eric Zaslow. Augmentations are
           Sheaves. *Geom. Topol.*, 24(5):2149–2286, 2020.
[NRSS17]   Lenhard Ng, Dan Rutherford, Vivek Shende, and Steven Sivek. The cardinality of the augmenta-
           tion category of a Legendrian link. *Math. Res. Lett.*, 24(6):1845–1874, 2017.
[Pan17]   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.
[Pos06]   Alexander Postnikov. Total positivity, grassmannians, and networks. preprint, 2006.
[PT20]   James Pascaleff and Dmitry Tonkonog. The wall-crossing formula and Lagrangian mutations. *Adv.
           Math.*, 361:106850, 67, 2020.
[SB21]   Melissa Sherman-Bennett. *On the Combinatorics of Cluster Structures on Positroid Varieties.*
           ProQuest LLC, Ann Arbor, MI, 2021. Thesis (Ph.D.)–University of California, Berkeley.
[Sch18]   Olaf M. Schnürer. Six operations on dg enhancements of derived categories of sheaves. *Selecta
           Math. (N.S.)*, 24(3):1805–1911, 2018.
[SSBW19]   K. Serhiyenko, M. Sherman-Bennett, and L. Williams. Cluster structures in Schubert varieties in
           the Grassmannian. *Proc. Lond. Math. Soc. (3)*, 119(6):1694–1744, 2019.
[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 Legen-
           drian 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.
[SW19]   Linhui Shen and Daping Weng. Cluster structures on double bott-samelson cells. preprint, 2019.
[SW20]   Linhui Shen and Daping Weng. Cyclic sieving and cluster duality for Grassmannian. *SIGMA,
           Special Issue on Cluster Algebras*, 16, 2020.
[Tre09]   David Treumann. Exit paths and constructible stacks. *Compos. Math.*, 145(6):1504–1532, 2009.
[TV07]   Bertrand Toën and Michel Vaquié. Moduli of objects in dg-categories. *Ann. Sci. École Norm. Sup.
           (4)*, 40(3):387–444, 2007.
[TV08]   Bertrand Toën and Gabriele Vezzosi. Homotopical algebraic geometry. II. Geometric stacks and
           applications. *Mem. Amer. Math. Soc.*, 193(902):x+224, 2008.
[Vak15]   Ravi Vakil. *THE RISING SEA, Foundations of Algebraic Geometry.* 2015.
[Yau17]   Mei-Lin Yau. Surgery and isotopy of Lagrangian surfaces. In *Proceedings of the Sixth International
           Congress of Chinese Mathematicians. Vol. II*, volume 37 of *Adv. Lect. Math. (ALM)*, pages 143–
           162. Int. Press, Somerville, MA, 2017.

UNIVERSITY OF CALIFORNIA DAVIS, DEPT. OF MATHEMATICS, USA

Email address: casals@math.ucdavis.edu

UNIVERSITY OF CALIFORNIA DAVIS, DEPT. OF MATHEMATICS, USA

Email address: dweng@ucdavis.edu