



Geometry & Topology

Volume 28 (2024)

Microlocal theory of Legendrian links and cluster algebras

ROGER CASALS

DAPING WENG



Microlocal theory of Legendrian links and cluster algebras

ROGER CASALS

DAPING WENG

We show the existence of quasicluster \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, including the initial weave associated to a grid plabic graph, cluster mutation along nonsquare faces of a plabic graph, possibly including lollipops, the concept of sugar-free hull, and the notion of microlocal merodromy. Finally, we prove the existence of the cluster DT transformation for shuffle graphs, constructing a contact-geometric realization and an explicit reddening sequence, and establish cluster duality for the cluster ensembles.

13F60, 53D12

1. Introduction	901
2. Grid plabic graphs and Legendrian links	909
3. Diagrammatic weave calculus and initial cycles	925
4. Construction of quasicluster structures on sheaf moduli	951
5. Cluster DT transformations for shuffle graphs	987
References	997

1 Introduction

The object of this article will be to show the existence of intrinsically symplectic quasicluster K_2 -structures and quasicluster 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 quasicluster 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 quasicluster 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




Figure 1: The quasicluster 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.

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 the initial weave of a grid plabic graph, cluster mutations along nonsquare faces, possibly with lollipops, 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, embedded Lagrangian disks whose boundaries lie on embedded exact Lagrangian fillings 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 quasicluster structures.

1.1 Scientific context

Cluster algebras, first introduced by S Fomin and A Zelevinsky [2002; 2003; Berenstein et al. 2005] 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 [2006b; 2006a], are affine schemes equipped with an atlas of torus charts whose transition maps obey certain combinatorial rules. Cluster varieties come in 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 and Zelevinsky [Berenstein et al. 2005].

Since their introduction, cluster algebras and cluster varieties have appeared in many contexts, such as Teichmüller theory [Fock and Goncharov 2006b; Fomin et al. 2008; Gekhtman et al. 2005], birational geometry [Gross et al. 2015; 2018; Hacking and Keel 2018], the Riemann–Hilbert correspondence [Allegretti 2021; Neitzke 2014; Gaiotto et al. 2010], exact WKB analysis [Iwaki and Nakanishi 2014; 2016], and the study of positroid and Richardson varieties [Galashin and Lam 2023; Serhiyenko et al. 2019]. 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; see eg [Auroux 2007; 2009; Pascaleff and Tonkonog 2020]. We also thank Goncharov for pointing out to us his recent work with Kontsevich [2021] focusing on noncommutative 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 [Shende et al. 2019], where it was computed how certain absolute monodromies around a square plabic face change under a square move in a plabic fence. See also the generalization presented in [Shende et al. 2016]. In conjunction, [Shende et al. 2019; 2016] should imply the existence of partial \mathcal{X} -structures for certain moduli stacks of sheaves singularly supported in the Legendrian lifts of the alternating strand diagrams of plabic fences. Nevertheless, they do not imply the existence of the full cluster \mathcal{X} -structures, nor the full cluster \mathcal{A} -structures and certainly not the fact that the rings of regular functions are cluster algebras. (See Section 2.8.) These stronger statements are proven here.

There are two obstacles to proving 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 variable to them and allows for a geometric mutation to be performed. Second, more fundamental, is the regularity problem: even if all faces are square, the 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 restricted to either plabic graphs or absolute cycles, both of which are limiting constraints in that approach.

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; see Section 2.8. 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 used ideas from the theory of cluster algebras for new applications to contact and symplectic geometry — see eg [Casals 2022; 2020; 2021; Casals and Zaslow 2022; Gao et al. 2020a] — including the discovery of infinitely many Lagrangian fillings for many Legendrian links [Casals and Gao 2022]. This article builds in the converse direction, using contact and symplectic topology to construct (upper) cluster algebras, and symplectic topological results to deduce algebraic properties. In fact, we also know $\mathcal{A} = \mathcal{U}$ by [Casals et al. 2022], which builds on the present manuscript.

Note that what can be deduced from [Casals et al. 2020; Casals and Gao 2022; Casals and Zaslow 2022; Gao et al. 2020a; 2020b] is that certain moduli spaces that appear in contact topology are sometimes abstractly isomorphic to certain affine varieties, which themselves can independently be endowed¹ 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 (eg even after a Reidemeister I or II move) and no symplectic computation or interpretation of cluster \mathcal{A} -variables existed. This work finally

¹Explicitly, double Bott–Samelson cells for [Gao et al. 2020a], and positroids for [Casals and Gao 2022; Shende et al. 2019]. These instances are, in any case, particular cases of the moduli stacks that we associate to grid plabic graphs.

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 $\Lambda \subset (T_\infty^*\mathbb{R}^2, \xi_{\text{st}})$ 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 $L \subset (T^*\mathbb{R}^2, \lambda_{\text{st}})$ 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 $\partial \bar{D} \cap L = \gamma \subset \mathbb{R}^4$. A collection $\{\gamma_1, \dots, \gamma_\ell\}$ 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 [Polterovich 1991; Yau 2017].

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 Section 2.7.3, following [Kashiwara and Schapira 1990; Guillermou et al. 2012], 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 Section 2.3 for the definition of the sugar-free hull \mathbb{S}_f of a face f in \mathbb{G} and Section 4.8 for completeness. Note that the concept of sugar-free hulls, and whether a region is sugar-free, only depends on the behavior at nonconvex corners; see Definition 2.2.

Our main result, stated in Theorem 1.1, is the existence and explicit symplectic construction of a full quasicluster \mathcal{A} -structure on $\mathfrak{M}(\Lambda, T)$. In particular, the cluster \mathcal{A} -variables of the initial seed and 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 filling, following [Guillermou et al. 2012; Casals and Zaslow 2022], and we refer to it as a *microlocal merodromy*; see Section 4.

Theorem 1.1 (main result) *Let $\mathbb{G} \subset \mathbb{R}^2$ be a complete grid plabic graph, $\Lambda = \Lambda(\mathbb{G}) \subset (\mathbb{R}^3, \xi_{\text{st}})$ 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 $L = L(\mathbb{G}) \subset (\mathbb{R}^4, \omega_{\text{st}})$ of Λ and a canonical \mathbb{L} -compressing system $\mathfrak{S} = \{\gamma_1, \dots, \gamma_\ell\}$ 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*

$$A_{\eta_i} : \mathfrak{M}(\Lambda, T) \rightarrow \mathbb{C}, \quad \text{ie } A_{\eta_i} \in \mathbb{C}(\mathfrak{M}(\Lambda, T)),$$

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

- (ii) The microlocal merodromies $\{A_{\eta_1}, \dots, A_{\eta_\ell}\}$ associated to the relative cycles that are dual to an \mathbb{L} -compressible absolute cycle in \mathfrak{S} are irreducible functions in $\mathbb{C}(\mathfrak{M}(\Lambda, T))$, whereas the merodromies $\{A_{\eta_{\ell+1}}, \dots, A_{\eta_{b_1(L,T)}}\}$ are nonvanishing functions, ie units in $\mathbb{C}(\mathfrak{M}(\Lambda, T))$.
- (iii) Let $L'_k \subset (\mathbb{R}^4, \omega_{\text{st}})$ 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

$$A_{\eta'_k} : \mathfrak{M}(\Lambda, T) \rightarrow \mathbb{C}, \quad \text{ie } A_{\eta'_k} \in \mathbb{C}(\mathfrak{M}(\Lambda, T)),$$

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

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

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 quasiequivalent 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)–(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 quasicluster \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 $\mathbb{C}(\mathfrak{M}(\Lambda, T))$ 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), eg 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 $\mathbb{C}(\mathfrak{M}(\Lambda, T))$. Table 1 schematically relates different ingredients involved in the proof of Theorem 1.1.

There are several items from Theorem 1.1 that can be helpful to unpack. First, by a modification of the Guillermou–Jin–Treumann map — see [Jin and Treumann 2017] — the Lagrangian filling L yields

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 $\text{Sh}_\Lambda(\mathbb{R}^2)$
Goncharov–Kenyon conjugate surface associated to \mathbb{G}	weave for Lagrangian filling L (\implies sheaf quantization $\mathcal{F}(L)$)	open toric chart $T_L = (\mathbb{C}^\times)^{b_1(L, T)} \subseteq \mathfrak{M}(\Lambda, T)$
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 <i>global regular</i> function $A_\eta: \mathfrak{M}(\Lambda, T) \rightarrow \mathbb{C}$
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} (eg 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 <i>nonvanishing</i> global regular function $A_\phi: \mathfrak{M}(\Lambda, T) \rightarrow \mathbb{C}$
subset of non-square-free 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 (quasicluster 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 a plabic graph but represented by a weave)	<i>Lagrangian surgery</i> $L' = \mu_\gamma(L)$ and relative cycle $\eta' = \mu_\gamma(\eta)$ in L' ; sheaf quantization $\mathcal{F}(L')$ via <i>weave mutation</i> at Y -tree for γ	$T_{L'}$ -coordinate extending to a <i>global regular</i> function $A_{\eta'}: \mathfrak{M}(\Lambda, T) \rightarrow \mathbb{C}$ given by <i>cluster \mathcal{A}-mutation</i> at η

Table 1: Ingredients in the symplectic construction of upper cluster algebra for $\mathcal{O}(\mathfrak{M}(\Lambda, T))$.

an open toric chart $(\mathbb{C}^\times)^{b_1} \subset \mathfrak{M}(\Lambda, T)$, where $b_1 = \text{rk}(H_1(L \setminus T, \Lambda \setminus T)) = \text{rk}(H_1(L, T))$. The group $H^1(L; \mathbb{C}^\times) = \text{Hom}(H_1(L; \mathbb{Z}), \text{GL}_1(\mathbb{C}))$ 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 2.7 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 nontrivial 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. (For example, 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 frozen. The existence of a cluster structure on $\mathfrak{M}(\Lambda, T)$ with a particular quiver Q has neat applications to symplectic geometry, eg studying the possible relative Lagrangian skeleta containing L for the Weinstein relative pair (\mathbb{C}^2, Λ) ; see below for more.

Third, item (iii) in Theorem 1.1 geometrically keeps 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 μSh_Λ , 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 in order 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 quasicluster \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 quasicluster ambiguity. In particular cases, such as \mathbb{G} being a plabic fence, the \mathbb{L} -compressing system already gives a 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 quasicluster 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 [Muller 2013; 2014; 2022], the existence of a full cluster \mathcal{A} -structure on the moduli space² $\mathfrak{M} = \mathfrak{M}(\Lambda, T)$, as proven in Theorem 1.1, leads to:

- (1) The computation of its de Rham cohomology ring $H^*(\mathfrak{M}, \mathbb{C})$, including the refinement of its mixed Hodge structure. These computations are done in [Lam and Speyer 2022] for the locally acyclic cases.
- (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 $\mathbb{C}(\mathfrak{M})$; see [Gekhtman et al. 2010]. 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 [Casals et al. 2020] by different means (using the Cartan 3-form and Bott–Shulman forms), and see work of P Boalch

²If not made explicitly, the set of marked points T is taken to have one marked point per component of Λ .

[2014a; 2014b]. Upcoming work with our collaborators will show that these holomorphic symplectic structures coincide whenever they can be compared.

- (3) In the Louise case [Lam and Speyer 2022], it is possible to compute the eigenvalues of the Frobenius automorphism on ℓ -adic cohomology and perform finite point counts $\#\mathfrak{M}(\mathbb{F}_q)$ over finite fields \mathbb{F}_q for $q = p^k$ and p large enough. These ought to be compared with the contact and symplectic results in [Henry and Rutherford 2015; Ng et al. 2017].

Another byproduct of our result, thinking in terms of cluster ensembles [Fock and Goncharov 2006a], is that there also exists a (full) cluster \mathcal{X} -structure. Let $\mathcal{M}_1(\Lambda)$ be the undecorated stack associated to $\mathfrak{M}(\Lambda, T)$ and $\mathcal{M}_1(\Lambda, T)$ its enhancement 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, $\Lambda = \Lambda(\mathbb{G}) \subset (\mathbb{R}^3, \xi_{\text{st}})$ its associated Legendrian link and $T \subset \Lambda$ marked points. Then there exists a quasicluster \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 quasiequivalent cluster \mathcal{X} -structures on $\mathcal{M}_1(\Lambda, T)$.

Corollary 1.2 is a new result and establishes the existence of a (full) cluster \mathcal{X} -structure. 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 full cluster \mathcal{A} -structure in Theorem 1.1; two mathematical reasons are that the results used from [Berenstein et al. 2005] are only applicable to cluster \mathcal{A} -structures and that the codimension 2 arguments in Section 4 require an explicit understanding of the \mathcal{A} -variables, including their irreducibility; see also Section 2.8.

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 varieties always admit a Donaldson–Thomas (DT) transformation. See [Kontsevich and Soibelman 2010; Goncharov and Shen 2018] 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 $t: (x, y, z) \mapsto (-x, y, -z)$ of $(\mathbb{R}^3, \ker\{dz - y dx\})$. In particular, we conclude the following result:

Corollary 1.3 *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 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 [Casals et al. 2020; 2021; Casals and Zaslow 2022; Gao et al. 2020a], where even the initial seed could not typically be defined (nor computed) after a Legendrian isotopy.³

Notation We denote by $[a, b]$ the discrete interval $\{k \in \mathbb{N} \mid a \leq k \leq b\}$ if $a \leq b$ with $a, b \in \mathbb{N}$. In this article, S_n denotes the group of permutations of n elements for $n \in \mathbb{N}$, and s_i its i^{th} simple transposition for $i \in [1, n-1]$. We abbreviate $s_{[b,a]} := s_b s_{b-1} \dots s_{a+1} s_a$ and $s_{[b,a]}^{-1} := s_a s_{a+1} \dots s_{b-1} s_b$ for $a < b$ with $a, b \in \mathbb{N}$, 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 $w_{0,n}$ for $w_{0,n}$ is defined to be the reduced expression $w_{0,n} := s_{[1,1]} s_{[2,1]} s_{[3,1]} \dots s_{[n-1,1]}$.

Extended version This paper is a condensed account of arXiv 2204.13244, also available on our research websites, which contains a series of additional examples and figures, as well as more detail in some of the proofs and motivation and context in parts of the construction. The interested reader might benefit from the more inviting extended version, as it is more comprehensive and builds the proofs in a more self-contained manner. That said, we believe experts will also appreciate this streamlined version, where only the logically necessary steps for our main results are included.

Acknowledgements We are grateful to C Fraser, H Gao, A Goncharov, E Gorsky, M Gorsky, I Le, W Li, J Simental-Rodriguez, L Shen, M Sherman-Bennett and E Zaslow for their interest, comments on the draft and useful conversations. We also thank the referee for their valuable comments. Casals is supported by the NSF CAREER DMS-1942363 and a Sloan Research Fellowship of the Alfred P Sloan Foundation.

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 Section 2.1, and that of sugar-free hulls in Section 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 Section 2.4, and set up the necessary moduli spaces from the microlocal theory of sheaves in Section 2.7. Several explicit examples are provided in Section 2.5.

2.1 Grid plabic graphs

The input object in our results is the following type of graphs:

³The pullback structures from [Shende et al. 2019, Section 3] had the same issue.

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*.
- (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 endpoints of a vertical edge must be contained in the interior of a horizontal line.

In Definition 2.1, it is fine to allow for bivalent vertices. The Legendrian isotopy type of the zigzag diagram, as introduced in Section 2.4, does not change when inserting such vertices, nor does the Hamiltonian isotopy type of the Lagrangian filling associated to the conjugate surface.

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

- A column is said to be Type 1 if it solely consists of parallel horizontal lines, ie it contains no vertices.
- 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.
- 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.




Figure 2: The three types of elementary column in a GP graph.

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 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 $s_i = (i, i + 1)$ with $i \in \mathbb{N}$. Within $S_{\mathbb{N}}$, we define $S_{[a,b]} \cong S_{b-a+1}$ to be the 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 via these 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).

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 areas 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).

The *boundary* ∂R of a region R is the topological (PL-smooth) boundary of its closure $\bar{R} \subset \mathbb{R}^2$. 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 Γ or \perp , and a 270° corner is said to be *right-pointing* if it is of the form \sqsupset or \sqsubset . Equipped with this terminology, we introduce the following notion:




Figure 3: 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.

Definition 2.2 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 3 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.

The boundary of a sugar-free region has the following characterization, which follows immediately from the fact that all vertical bars must be of different colors at the two ends:

Lemma 2.3 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 four types illustrated in Figure 4.

Lemma 2.4 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 2.3 and therefore $R \cap C$ can have at most one connected component. \square



Figure 4: 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



Figure 5: The local models for an alternating strand diagram associated to a GP graph \mathbb{G} . The small hairs indicate the coorienting direction, which is needed to specify a Legendrian lift.

Note that if a region R is not simply connected, then there must exist a column C such that $R \cap C$ has more than one connected component. Thus, Lemma 2.4 has the following consequence, despite the fact that there may exist non–simply connected faces in the GP graph:

Corollary 2.5 *Sugar-free regions are simply connected.*

2.4 Legendrian links

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

Definition 2.6 Let $\mathbb{G} \subset \mathbb{R}^2$ be a GP graph. The Legendrian link $\Lambda(\mathbb{G}) \subset (\mathbb{R}^3, \xi_{st})$ is the Legendrian lift of the alternating strand diagram of \mathbb{G} , understood as a cooriented front in \mathbb{R}^2 , considered inside a Darboux ball in $(T^*\mathbb{R}^2, \xi_{st})$.

Alternating strand diagrams were introduced in [Postnikov 2006, 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 local models shown in Figure 5, where the hairs indicate the coorientation. The alternating strand diagram near a lollipop (or a bivalent vertex) is the same as in [Postnikov 2006], and the coorientation in these pieces is implied by the coorientations above.

By definition, the Legendrian lift of a cooriented immersed curve on the plane \mathbb{R}^2 is a Legendrian link inside the ideal contact boundary $(T^*\mathbb{R}^2, \xi_{st})$. 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, but, for a 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 $T^*\mathbb{R}^2_{u,v}$ can be identified as the kernel of the contact 1–form $\alpha_{st} := \cos \theta du + \sin \theta dv$, 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 $T^*\mathbb{R}^2$ is diffeomorphic to $\mathbb{R}^2 \times S^1$ and $\theta \in S^1$ records that circle coordinate. In fact, we can consider the 1–jet space (J^1S^1, ξ_{st}) with its standard contact structure $\ker\{\beta_{st}\}$, $\beta_{st} := dz - y d\theta$, where $y \in \mathbb{R}$ is the coordinate along the cotangent fiber and $z \in \mathbb{R}$ the Reeb coordinate, as $J^1S^1 := T^*S^1 \times \mathbb{R}$. Then there exists a strict contactomorphism $\varphi: (T^*\mathbb{R}^2, \alpha_{st}) \rightarrow (J^1S^1, \beta_{st})$ given by

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




Figure 6: The rules to construct the front $f(\mathbb{G})$ 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 $f(\mathbb{G})$ 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.

For any open interval $I \subset S^1$, $(J^1 I, \xi_{st})$ is contactomorphic to a standard Darboux ball (\mathbb{R}^3, ξ_{st}) . In consequence, if a cooriented immersed curve $f \subset \mathbb{R}^2$ in \mathbb{R}^2 has a Gauss map that misses one given angle θ_0 , the Legendrian lift of $f \subset \mathbb{R}^2$ is contained in $(J^1(S^1 \setminus \theta_0), \xi_{st})$, which is contactomorphic to a Darboux ball. This happens for the alternating strand diagram of a GP graph $\mathbb{G} \subset \mathbb{R}^2$ and thus $\Lambda(\mathbb{G})$ naturally lives inside a Darboux ball.

Let us now construct a particular type of (wave)front for the Legendrian link $\Lambda(\mathbb{G})$, which is useful to describe our moduli spaces in Lie-theoretic terms. For that, we consider the front $f(\mathbb{G}) \subset \mathbb{R}^2$ obtained by dividing the GP graph \mathbb{G} into elementary columns and then use the assignments depicted in Figures 6 and 7. 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 i^{th} position, we assign a front consisting of $2n$ parallel horizontal strands with a crossing at the i^{th} position either at the top n strands or the bottom n strands, depending on whether the vertical bar has a white vertex at the top or at the bottom. Figure 6 depicts these three cases for Types 1 and 2. The case of an




Figure 7: The rules to construct the front $f(\mathbb{G})$ 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 $f(\mathbb{G})$ 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.

elementary column of Type 3 involves inserting a right (resp. left) cusp at the i^{th} position (plus some additional crossings) if there is a black (resp. white) lollipop inserted at the i^{th} position. Figure 7 depicts the two possible cases for a Type 3 column.

Note that the $2n$ strands in the front are labeled in a specific manner in Figures 6 and 7, 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 i^{th} top strand (in orange) and the i^{th} bottom strand (in blue) coincide. Now, the front $f(\mathbb{G}) \subset \mathbb{R}^2$ lifts to a Legendrian link $\Lambda(f(\mathbb{G})) \subset (\mathbb{R}^3, \xi_{\text{st}})$. We observe that, in this case, the lift can be considered directly into \mathbb{R}^3 , as the front is cooriented upwards and there are no vertical tangencies. The following proposition follows by applying the above contactomorphism $\varphi: (T^\infty \mathbb{R}^2, \alpha_{\text{st}}) \rightarrow (J^1 S^1, \beta_{\text{st}})$:

Proposition 2.7 *Let $\mathbb{G} \subset \mathbb{R}^2$ be a GP graph. Then the two Legendrian links $\Lambda(\mathbb{G}) \subset (\mathbb{R}^3, \xi_{\text{st}})$ and $\Lambda(f(\mathbb{G})) \subset (\mathbb{R}^3, \xi_{\text{st}})$ are Legendrian isotopic.*

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.

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 $\{1\} \times \mathbb{R}$. Figure 8 depicts instances of such GP graphs. These GP graphs are called plabic fences in [Fomin et al. 2022, Section 12], following L Rudolph’s fence terminology. It follows from Proposition 2.7 and the rules from Figures 6 and 7 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 $\beta \in \text{Br}_n^+$ whose k^{th} crossing is σ_j if and only if the k^{th} vertical edge *with black on bottom* of \mathbb{G} (starting from the left) is between the j^{th} and $(j+1)^{\text{st}}$ horizontal strands. Similarly, consider the positive braid word δ whose m^{th} crossing is σ_{n-j} if and only if the m^{th} vertical edge *with white on*




Figure 8: A front for the Legendrian link associated to the GP graph on the left is drawn on the right, where $\beta, \delta \in \text{Br}_6^+$ 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 9: GP graphs whose alternating strand diagrams have the smooth type of (mirrors of) twist knots. The GP graph on the left (resp. middle) yields a max-tb Legendrian representative of $m(5_2)$ (resp. $m(7_2)$). The right illustrates 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.

bottom of \mathbb{G} is between the j^{th} and $(j+1)^{\text{st}}$ horizontal strands. Then Figure 8, 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 .

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 9. 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 9 draws \mathbb{G}_1 (left) and \mathbb{G}_2 (middle). By using Figures 6 and 7, fronts for the associated Legendrian knots $\Lambda(\mathbb{G}_n)$ are readily drawn: Figure 10 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 representative of a twist knot, with zero rotation number. Note that Legendrian twist knots are classified in [Etnyre et al. 2013]. In particular, the Legendrian knot $\Lambda(\mathbb{G}_1)$ associated to the GP graph depicted in Figure 9, left, 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.⁴




Figure 10: The two Legendrian fronts obtained from the left and middle GP graphs of Figure 9, according to our recipe translating from GP graphs to Legendrian front diagrams. The corresponding GP graphs are drawn in a small box below each front. The front diagram for the general case (Figure 9, right) is readily inferred from these two pictures: a knotted spiraling pattern is iteratively added to the center region of the front.

Shuffle graphs Let us introduce a class of GP graphs which leads to interesting examples.

⁴The other max-tb representative of $m(5_2)$ is not isotopic to $\Lambda(\mathbb{G})$ for any GP plabic graph \mathbb{G} .

Definition 2.8 A GP graph $\mathbb{G} \subset \mathbb{R}^2$ with n horizontal lines is said to be a *shuffle graph* if:

- (1) There exist $M \in \mathbb{N}$ and $\sigma \in S_n$ such that each horizontal line goes from $(-M\sigma(i), i)$ to $(M\sigma(i), i)$.
- (2) Vertical edges are all of the same pattern, ie they either all have a black vertex on top or they all have a white vertex on top.

The two families above, plabic fences and the GP graphs in Figure 9 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 $\text{Dem}(\beta) = w_0 = w_{0,n} \in S_n$. (By [Casals et al. 2020; Casals and Ng 2022], the condition $\text{Dem}(\beta) = w_0 \in S_n$ is necessary.) For example, 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} . In view of this and [Casals et al. 2020; 2021; Casals and Ng 2022], we refer to a positive braid $\beta \in Br_n^+$ as Δ -complete if it is cyclically equivalent to a positive braid of the form $\Delta\gamma$, where Δ is the half twist and $\text{Dem}(\gamma) = w_{0,n}$.


Let us point out two properties of the Legendrian links $\Lambda(\mathbb{G})$ that are useful. 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, ξ_{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 Section 2.4, especially Figures 6 and 7, 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 thus possible to parametrize the moduli of objects of the appropriate dg category by an affine variety (or algebraic quotient thereof). Section 2.7 sets up the necessary ingredients on 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 Section 4.7, in the proof of Proposition 4.23. Let us start with the definition of a single *lollipop reaction*:

Definition 2.9 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 to the immediate left of w (in between which the lollipop appears). A vertical line segment between h_1 and h_2 is said to be a *wall*. By definition, the *lollipop reaction* initiated from the lollipop w pushes this wall to the right along \mathbb{G} with the following rules: the wall shrinks or expands

according to the following five pictures and otherwise the wall stays between the same \mathbb{G} -edges:



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 five pictures and otherwise stays between the same \mathbb{G} -edges:



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 the selection of faces of the GP graph which it has scanned through. If the length of the wall becomes infinite (ie 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 2.10 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. First, select the face f . Then, 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*.

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 2.11 *Let \mathbb{G} be a shuffle graph and f a face of \mathbb{G} for which S_f is nonempty. Then the lollipop chain reaction initiated from f is complete and 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 $S_f = f$, which is also equal to the outcome of the lollipop chain 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} -edges of the type $\begin{array}{c} \circ \\ | \\ \bullet \end{array}$ appear in \mathbb{G} and suppose that f contains a white lollipop. Consider the leftmost 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 situations 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 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. So we have



Now we start scanning leftward from the rightmost 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 leftward 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 leftward 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 nonshuffle GP graphs for which lollipop chain reactions also yield sugar-free hulls, and thus the hypothesis in Proposition 2.11 is sufficient but not necessary.

2.7 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 Section 2.7.1, based on [Kashiwara and Schapira 1985; 1990; Guillermou et al. 2012; Guillermou and Schapira 2014; Shende et al. 2017].⁵ Section 2.7.2 discusses the specific simplifications that occur for the Legendrian links $\Lambda(\mathbb{G})$ and Section 2.7.3 introduces the necessary decorated version of the moduli stacks being discussed. Let \mathbb{k} be a commutative coefficient ring, for us either $\mathbb{k} = \mathbb{Z}$ or $\mathbb{k} = \mathbb{C}$. Consider a smooth manifold M , $\pi_M : T^*M \rightarrow M$ its cotangent bundle and $T_\infty^*M \rightarrow M$ its ideal contact boundary; we will only need $M = \mathbb{R}^2$ and \mathbb{R}^3 .

2.7.1 The general setup The general results on the microlocal theory of constructible sheaves were pioneered by M Kashiwara and P Schapira [1990] and, more recently, in collaboration with S Guillermou in [Guillermou et al. 2012]. The first category that we need is defined as follows:

Definition 2.12 The category $\underline{\mathbb{I}}(\mathbb{k}_M)$ is the full dg subcategory of the dg category of locally bounded complexes of sheaves of \mathbb{k} -modules on M which consist of h -injective complexes of injective sheaves. The homotopy category of $\underline{\mathbb{I}}(\mathbb{k}_M)$ is denoted by $[\underline{\mathbb{I}}(\mathbb{k}_M)]$.

The dg category $\underline{\mathbb{I}}(\mathbb{k}_M)$ is a strongly pretriangulated dg category and the six-functor formalism lifts to this dg enhancement $\underline{\mathbb{I}}(\mathbb{k}_M)$; see [Schnürer 2018]. The homotopy category $[\underline{\mathbb{I}}(\mathbb{k}_M)]$ is triangulated equivalent to the locally bounded derived category of sheaves on M , often denoted by $D^{\text{lb}}(\mathbb{k}_M)$. For an object $F \in \underline{\mathbb{I}}(\mathbb{k}_M)$, we denote by $\mu\text{supp}(F) \subset T^*M$ its singular support understood as an object in $[\underline{\mathbb{I}}(\mathbb{k}_M)] \simeq D^{\text{lb}}(\mathbb{k}_M)$. The notion of singular support leads to defining the following dg categories:

⁵See also Guillermou's notes for his lecture series at the conference *Symplectic topology, sheaves and mirror symmetry* at the IMJ-PRG in Paris (2016) and [Shende et al. 2019].

Definition 2.13 Let $S \subset T^*M$ be a subset. The category $\mathbb{I}_S(\mathbb{k}_M)$ is the subcategory of $\mathbb{I}(\mathbb{k}_M)$ consisting of objects $F \in \mathbb{I}(\mathbb{k}_M)$ such that $\mu\text{supp}(F) \subset S$. The category $\mathbb{I}_{(S)}(\mathbb{k}_M)$ is the subcategory $\mathbb{I}(\mathbb{k}_M)$ consisting of objects $F \in \mathbb{I}(\mathbb{k}_M)$ for which there exists an open neighborhood Ω such that $\mu\text{supp}(F) \cap \Omega \subset S$.

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

The assignment $U \mapsto \mathbb{I}(\mathbb{k}_U)$ to each open subset $U \subset M$ is a stack of dg categories. Similarly, the prestack $\mathbb{I}_{(\Lambda)}$ defined by

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

is a stack. This is an advantage of using the injective dg enhancements instead of derived categories. A central result in symplectic topology [Guillermou et al. 2012] is that the stack $\mathbb{I}_{(\Lambda)}$ on M is a Legendrian isotopy invariant of the Legendrian $\Lambda \subset T^*_\infty M$. There are two constructions associated to the stack $\mathbb{I}_{(\Lambda)}$, as follows:

- (i) **The microlocal functor \mathfrak{m}_Λ** The Kashiwara–Schapira stack $\mu\text{Sh}(\mathbb{k}_\Lambda)$ is the stack on Λ associated to the prestack

$$V \mapsto \mathbb{I}_{(V)}(\mathbb{k}_M; V), \quad V \subset \Lambda \text{ open,}$$

where $\mathbb{I}_{(V)}(\mathbb{k}_M; V)$ is the Drinfeld dg quotient of $\mathbb{I}_{(V)}(\mathbb{k}_M)$ by $\mathbb{I}_{T^*M \setminus (V \cup M)}(\mathbb{k}_M)$. See [Kashiwara and Schapira 1990; Guillermou et al. 2012]. The quotient functor gives a functor of stacks

$$\mathfrak{m}_\Lambda : \mathbb{I}_{(\Lambda)} \rightarrow (\pi_M|_\Lambda)_*(\mu\text{Sh}(\mathbb{k}_\Lambda)).$$

Our use of this functor is twofold: 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 [Nadler 2016, 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 $\text{Sh}^w_\Lambda(M)$ defined in [Nadler 2016, Definition 3.17]. Then the main result of [Toën and Vaquié 2007] 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 $\mathbb{I}_{(\Lambda)}(M)$.

Finally, as explained in [Jin and Treumann 2017, 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}_L applied to the Legendrian lift $\bar{L} \subset J^1(M)$ yields an equivalence of categories between (a subcategory of) $\mathbb{I}_{(\bar{L})}(M)$ and the dg derived category of local systems on L . This induces an open inclusion $\iota_L : \mathbb{R}\text{Loc}(L) \rightarrow \mathcal{M}_{\mathbb{I}_{(\Lambda)}}(M)$, where $\mathbb{R}\text{Loc}(L)$ denotes the derived moduli space of local systems on L .

2.7.2 The concrete models For a Legendrian link $\Lambda \subset T_\infty^*\mathbb{R}^2$ with vanishing rotation number, the category of global sections of the Kashiwara–Schapira stack $\mu\text{Sh}(k_\Lambda)$ admits a simple object Ξ by [Guillermou 2023, Part 10]. In addition, the functor $\mu\text{hom}(\Xi, \cdot)$ is an explicit equivalence between $\mu\text{Sh}(k_\Lambda)$ and the (twisted) stack Loc_Λ of (twisted) local systems on Λ .⁶ In consequence, the microlocal functor \mathfrak{m}_Λ described above yields a functor

$$\mathfrak{m}_{\Lambda, \Xi} : \mathbb{I}(\Lambda)(M) \rightarrow \text{Loc}_\Lambda(\Lambda),$$

where we have considered global sections and identified the codomain of \mathfrak{m}_Λ with Loc_Λ via $\mu\text{hom}(\Xi, \cdot)$ and a choice of spin structure. In addition, given 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}_Λ is a local system (on Λ) of locally free \mathbb{k} -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 $\iota_L : \text{Loc}_1(L) \rightarrow \mathcal{M}_1(\Lambda)$ of Artin stacks, where $\text{Loc}_1(L)$ are local systems (on L) of locally free \mathbb{k} -modules of rank one supported in degree zero. Since abelian local systems on L can be parametrized by $H^1(L, \mathbb{k}^\times)$, the inclusion ι_L provides a toric chart $\iota_L(\text{Loc}_1(L))$ in the moduli stack $\mathcal{M}_1(\Lambda)$. In this article, we typically consider the ground ring $\mathbb{k} = \mathbb{C}$. 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 $\iota_L : (\mathbb{C}^\times)^{2g(L)} \rightarrow \mathcal{M}_1(\Lambda)$, where $g(L)$ is the topological genus of the surface L .

Finally, both the inclusions $\iota_L : \text{Loc}_1(L) \rightarrow \mathcal{M}_1(\Lambda)$ and the microlocal functors $\mathfrak{m}_\Lambda : \mathcal{M}_1(\Lambda) \rightarrow \text{Loc}_1(\Lambda)$ 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 [Shende et al. 2017, 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 \mathbb{k} -modules (modulo acyclic complexes). In the case of Legendrian weaves, this combinatorial model is explained in [Casals and Zaslow 2022, Section 5].

2.7.3 A decorated moduli space Let $T = \{t_1, \dots, t_s\}$ with $t_i \subset \Lambda$ for $i \in [1, s]$ 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 2.14 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 $T = \{t_1, \dots, t_s\}$ and label the components of $\Lambda \setminus T$ by Λ_i for

⁶A choice of spin structure on Λ and corresponding choices of spin structures for the Lagrangian fillings we consider allow a further identification to actual (untwisted) local systems. We implicitly have these choices in the background and translate them combinatorially in Section 4, through sign curves, when they are needed to assign signs.

$i \in \pi_0(\Lambda \setminus T)$. The moduli stack $\mathfrak{M}(\Lambda, T)$ is

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

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

In Definition 2.14, the identification of (global sections of) the codomain of \mathfrak{m}_Λ with the stack of local systems is fixed by the choice of Maslov potential and spin structure on Λ . 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 [Fock and Goncharov 2006b]. 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 .

2.8 A clarification on the notion of cluster structures

In the literature, the sentence “a space Y has a cluster structure” has different meanings. We record here the precise definitions that have been used, implicitly or explicitly, and clarify the type of results we obtain. Let Q be a quiver, or more generally a skew-symmetrizable matrix. Consider the following concepts:

- The cluster algebra \mathbb{A}_Q . This is a commutative algebra and it comes endowed with a (typically infinite) system of generators $A_i \in \mathbb{A}_Q$, called the *cluster variables*. The vertices of the quiver give some of these cluster variables, and the other cluster variables are produced by the process of mutation. Cluster algebras were first introduced and studied by Fomin and Zelevinsky [1999; 2002; 2003]. The affine scheme associated to \mathbb{A}_Q is $\text{Spec}(\mathbb{A}_Q)$.
- The space \mathcal{A}_Q , called the cluster \mathcal{A} -space or cluster K_2 -space, is a scheme obtained by birationally gluing certain tori according to Q . The ring of regular functions $\mathbb{O}(\mathcal{A}_Q) = \Gamma(\mathcal{A}_Q, \mathbb{O}_{\mathcal{A}_Q})$ is often referred to as the upper cluster algebra, due to its connection to [Berenstein et al. 2005]. This scheme is typically not finitely generated, but it is separated by [Gross et al. 2015, Theorem 3.14].
- The space \mathcal{X}_Q , called the cluster \mathcal{X} -space or cluster Poisson space, is also a scheme obtained by birationally gluing certain tori according to Q ; the gluing maps are different than for \mathcal{A}_Q above. The ring of regular functions $\mathbb{O}(\mathcal{X}_Q) = \Gamma(\mathcal{X}_Q, \mathbb{O}_{\mathcal{X}_Q})$ does not have a name. This scheme is typically not separated; see [Gross et al. 2015, Remark 2.6].
- The subset $\mu_{\leq 1}^{\mathcal{A}}(Q) \subset \mathcal{A}_Q$ consisting of the union of \mathcal{A} -tori associated to the initial seed for Q and its adjacent seeds, ie those obtained by performing *one* cluster \mathcal{A} -mutation. The ring of functions $\mathbb{O}(\mu_{\leq 1}^{\mathcal{A}}(Q))$ is known as the upper bound.

- The subset $\mu_{\leq 1}^{\mathcal{X}}(Q) \subset \mathcal{X}_Q$ consisting of the union of \mathcal{X} -tori associated to the initial seed for Q and its adjacent seeds, ie those obtained by performing *one* cluster \mathcal{X} -mutation.

The \mathcal{A} and \mathcal{X} -schemes were first introduced and studied by Fock and Goncharov [2006b; 2006a] and subsequently featured in [Gross et al. 2015; 2018]. There is also the notion of a partial \mathcal{X} -structure (and partial \mathcal{A} -structure), as introduced in [Shende et al. 2019, Definition 5.11], where only some tori in $\mu_{\leq 1}^{\mathcal{X}}(Q)$ are considered. After studying the literature and discussing with experts, our conclusion is that “a space Y has a cluster structure” might mean that Y is equal to either of the (often quite different) spaces above, or even that it is equal up to codimension 2, ie $\mathbb{O}(Y)$ equals any of the (often quite different) rings of functions above. In certain cases, such as [Goncharov and Kontsevich 2021], it might also mean having a partial \mathcal{A} - or partial \mathcal{X} -structure for what the authors referred to as a noncommutative stack.

Remark 2.15 It is crucial to have a rigorous definition of the “space” Y and its “ring of functions” $\mathbb{O}(Y)$ so as to give precise meaning to the notion of admitting a cluster structure. If Y is a scheme, the sheaf of regular functions is well understood [Hartshorne 1977]. In our case, $\mathfrak{M}(\Lambda, T)$ are always affine schemes and $\mathcal{M}_1(\Lambda)$ and $\mathcal{M}_1(\Lambda, T)$ are always algebraic quotients of affine schemes.

Now, the spaces $\text{Spec}(\mathbb{A}_Q)$, \mathcal{A}_Q , \mathcal{X}_Q , $\mu_{\leq 1}^{\mathcal{A}}(Q)$ and $\mu_{\leq 1}^{\mathcal{X}}(Q)$ are often quite different from each other, but the following facts hold:

- (1) The inclusion $\mathbb{A}_Q \subset \mathbb{O}(\mathcal{A}_Q)$ always holds. This is a nontrivial fact known as the Laurent phenomenon. In particular, *all* cluster \mathcal{A} -variables $A_i \in \mathbb{A}_Q$ belong to $\mathbb{O}(\mathcal{A}_Q)$. In fact, \mathbb{A}_Q can be defined to be the subalgebra of $\mathbb{O}(\mathcal{A}_Q)$ generated by the cluster \mathcal{A} -variables. In stark contrast, the cluster \mathcal{X} -variables X_i are almost never elements of $\mathbb{O}(\mathcal{X}_Q)$.
- (2) The inclusion $\mathbb{A}_Q \subset \mathbb{O}(\mathcal{A}_Q)$ of the cluster algebra into its upper cluster algebra may or may not be an equality. This is referred to as the $\mathcal{A} = \mathcal{U}$ problem; see eg [Berenstein et al. 2005; Muller 2014]. The inclusion $\mathbb{O}(\mathcal{A}_Q) \subset \mathbb{O}(\mu_{\leq 1}^{\mathcal{A}}(Q))$ of the upper cluster algebra into its upper bound may or may not be an equality. It is known to be an equality for the case of full rank. In general, the equality $\mathbb{O}(\mathcal{X}_Q) = \mathbb{O}(\mu_{\leq 1}^{\mathcal{X}}(Q))$ always holds.

The main spaces Y we study here are the affine schemes $\mathfrak{M}(\Lambda, T)$. The results we prove imply that $\mathbb{O}(\mathfrak{M}(\Lambda, T))$ equals $\mathbb{O}(\mathcal{A}_Q)$, where Q is the quiver geometrically defined in Section 3. That is, we construct an inclusion $\mathbb{O}(\mathfrak{M}(\Lambda, T)) \subseteq \mathbb{O}(\mathcal{A}_Q)$ and show that it is an equality. We also provide symplectic geometric meaning to the \mathcal{A} -variables in Section 4. In conjunction with [Casals et al. 2022] — which logically depends on [Shen and Weng 2021] and the present manuscript — we know that $\mathbb{O}(\mathcal{A}_Q) = \mathbb{A}_Q$. Therefore, $\mathfrak{M}(\Lambda, T)$ admits a cluster structure in the strongest possible sense: it is an affine scheme whose ring of regular functions *equals* the upper cluster algebra $\mathbb{O}(\mathcal{A}_Q)$, and also the cluster algebra \mathbb{A}_Q .

Remark 2.16 Similarly, a consequence of our results is that $\mathbb{O}(\mathcal{M}_1(\Lambda))$ equals $\mathbb{O}(\mathcal{X}_Q)$, which was also an open question. That is, we show that the inclusion $\mathbb{O}(\mathcal{M}_1(\Lambda)) \subseteq \mathbb{O}(\mathcal{X}_Q)$ is an equality. The results of [Shende et al. 2019], when combined with their later work [Shende et al. 2016], would likely imply that

for Λ associated to plabic fence (no lollipops) one has the inclusion $\mathbb{O}(\mathcal{M}_1(\Lambda)) \subset \mathbb{O}(\mathcal{X}_Q)$. That said, even [Shende et al. 2016; 2019] combined do not prove the equality in these cases.

Finally, we emphasize that the geometric description of the cluster \mathcal{A} -variables $A_i \in \mathbb{A}_Q$ and the particular algebraic geometric description of $\mathfrak{M}(\Lambda, T)$ is what allows for the equalities $\mathbb{O}(\mathfrak{M}(\Lambda, T)) = \mathbb{O}(\mathcal{A}_Q)$ to be proven in Section 4. In particular, the fact that $\mathbb{O}(\mathfrak{M}(\Lambda, T))$ is a unique factorization domain (Section 4.2) and the fact that the (candidate) cluster \mathcal{A} -variables are irreducible in $\mathbb{O}(\mathfrak{M}(\Lambda, T))$ (Section 4.9) are key to deduce $\mathbb{O}(\mathfrak{M}(\Lambda, T)) = \mathbb{O}(\mathcal{A}_Q)$.

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 [Casals and Zaslow 2022]. We continue to develop techniques for Legendrian weaves so as to prove Theorem 1.1. These new 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 [Casals and Zaslow 2022] for the details and background on Legendrian weaves, but we provide here a quick primer on the basics. 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, ie alternately belong to J and K . Figure 11, right, depicts such a hexagonal vertex.

Definition 3.1 Given $n \in \mathbb{N}$, an N -weave $\mathfrak{w} \subset \mathbb{R}^2$ is a set $\mathfrak{w} = \{G_i\}_{1 \leq i \leq N-1}$ 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 for $1 \leq i \leq N-2$. By definition, a weave $\mathfrak{w} \subset \mathbb{R}^2$ is an N -weave for some $N \in \mathbb{N}$.




Figure 11: Trivalent vertex (left) and hexagonal point (right).

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 $j \neq i, i \pm 1$, 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 for $1 \leq i \leq N - 1$. Edges labeled by numbers differing by two or more may pass through one another (hence the immersed property, which is met generically), but not at a vertex. As a graph in the plane, an 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. The theory of weaves as developed in [Casals and Zaslow 2022] is grounded on the theory of Legendrian surfaces in (\mathbb{R}^5, ξ_{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, ξ_{st}) . The main property of the surface $\Lambda(\mathfrak{w})$ that we use here is that its image $L(\mathfrak{w}) := \pi(\Lambda(\mathfrak{w})) \subset (\mathbb{R}^4, \omega_{st})$ is an exact Lagrangian surface in the standard symplectic Darboux ball, where $\pi : (\mathbb{R}^5, \xi_{st}) \rightarrow (\mathbb{R}^4, \omega_{st})$ is the projection along the α_{st} -Reeb direction.

Unless it is stated otherwise, all the weaves that we construct are free — see [Casals and Zaslow 2022, 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 $H_1(L(\mathfrak{w})) \cong H_1(\Lambda(\mathfrak{w}))$. The underlying contact geometry dictates that certain weaves ought to be considered equivalent. This leads to the following:

Definition 3.2 The moves depicted in Figure 12 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 12: Six weave equivalences.




Figure 13: The local models for a Y-tree. The Y-tree is highlighted in light green.

Remark 3.3 As noted in [Casals and Zaslow 2022, Section 4], these moves are not entirely independent, and 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 [Casals and Zaslow 2022], using the underlying contact geometry, imply that all the constructions that we associate to a weave are invariant under weave equivalences. (It would be possible to verify this combinatorially as well; see for instance the computations in [Casals et al. 2020].)

A first goal is constructing a weave for each GP graph $\mathbb{G} \subset \mathbb{R}^2$; it is achieved in Section 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 sublattice of the defining lattice for the initial seed. Casals and Zaslow [2022, Section 2] 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 a Y-cycle if its projection onto \mathbb{R}^2 consists of weave lines, ie 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 1-cycle if its projection onto \mathbb{R}^2 does not have any trivalent vertices. Finally, an 1-cycle is *short* if it does not pass through any hexagonal vertex of the weave \mathfrak{w} . Figure 13 depicts the four possible local models for a Y-tree near a trivalent, tetravalent, and hexagonal vertex of the weave.

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 [Casals and Zaslow 2022, 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. The outstanding role of Y-trees is justified by the following fact:



Figure 14: Local weave equivalences to turn a Y-tree into a short l-cycle. Note that the first row is just Move II from Figure 12, where we kept track of the Y-tree — highlighted in light green — before and after the equivalence.

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 l-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 [Casals and Zaslow 2022], by applying the equivalence moves in Figure 12 and keeping track of the change of a Y-tree under these moves. In fact, it suffices to use of the two local weave equivalences shown in Figure 14.

By using the two weave equivalences in Figure 14, 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 l-cycle somewhere along the original Y-tree. The double tracks that appear in this shortening process are schematically depicted in Figure 15. The second half of the proposition follows from the description of a short l-cycle in [Casals and Zaslow 2022]. □

Remark 3.6 Proposition 3.5 implies that any one Y-tree can be turned into a short l-cycle after weave equivalences. It is not the case that a Y-cycle, which is not necessarily a Y-tree, can always be turned into a short l-cycle. It is also not the case that Proposition 3.5 works for more than one Y-tree at once, in the following sense. If two Y-trees $Y_1, Y_2 \subset \mathfrak{w}$ are given in a weave \mathfrak{w} , then there exists a sequence of weave equivalences from \mathfrak{w} to a weave \mathfrak{w}_1 such that Y_1 becomes a short l-cycle in \mathfrak{w}_1 . There is no guarantee that Y_2 will be a short l-cycle in \mathfrak{w}_1 ; Y_2 will be a short l-cycle in another weave \mathfrak{w}_2 equivalent to \mathfrak{w} , a priori different from \mathfrak{w}_1 . More generally, there are collections of Y-trees in a weave \mathfrak{w} such that




Figure 15: 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 l-cycle drawn in light green.




Figure 16: Left: weave mutation along the short l -cycle given by the blue edge. Right: a weave mutation along a monochromatic Y -tree.

there is no sequence of weave equivalences that would turn at once all those Y -trees in \mathfrak{w} into short l -cycles in any one weave \mathfrak{w}' equivalent to \mathfrak{w} .

The existence of the embedded Lagrangian disks from Proposition 3.5, ie \mathbb{L} -compressible curves in $L(\mathfrak{w})$, allows us to perform Lagrangian disk surgeries along Y -trees and produce new exact Lagrangian fillings. Casals and Zaslow [2022, Section 4.8] 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 3.7 Let $\gamma \subset \mathfrak{w}$ be the short l -cycle — a monochromatic blue edge — depicted in Figure 16, left. Then the local move illustrated in Figure 16, left, is said to be a *weave mutation* at the short l -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 l -cycle γ with a new short l -cycle, which we often denote by γ' ; we call γ' the *image* of γ under the weave mutation.

Definition 3.8 For a Y -tree in general, one can apply Proposition 3.5 to turn it into a short l -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 l -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' .

Definition 3.9 Two weaves \mathfrak{w} and \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.

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 [Casals and Zaslow 2022] 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 $w_{0,n} = s_{[1,1]}s_{[2,1]}s_{[3,1]} \cdots s_{[n-1,1]}$ for the longest element $w_{0,n} \in S_n$. Let us define $\ell = \ell(w_{0,n}) = \frac{1}{2}n(n-1)$. The first three local weaves $n(w)$, $c^\uparrow(w)$ and $c^\downarrow(w)$ 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 $n(w)$ 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} for $j \in [1, \ell]$.

The weave $c^\uparrow(w)$ is given by the weave $n(w)$ where a trivalent vertex is added at the top strand — labeled 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.

Similarly, the weave $c^\downarrow(w)$ is given by the weave $n(w)$ where a trivalent vertex is added at the bottom strand — labeled 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 3.11 (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 $n(w_{0,n})$, where $w_{0,n}$ is the standard reduced expression for the longest element in a symmetric group S_k .

It is important to note that the transpositions s_i labeling the strands of $n(w_{0,n})$ depend on the simple transpositions that generate the symmetric group $S_{[a,b]}$ associated to that specific region (see Section 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 $[a, b]$.

3.3.2 Local weaves for Type 2 columns It is a well-known property of the symmetric group — see for instance [Björner and Brenti 2005, Section 3.3] — that any two reduced word expressions for the same element can be transformed into each other via finite sequences of the following two moves:

- $s_i s_j \sim s_j s_i$ 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 $n(w_{0,n})$ and the s_1 -strand labeled by the i^{th} appearance of s_1 in the standard reduced expression $w_{0,n}$. In order to construct the weave for a Type 2 column, in Definition 3.15 below, we need the following auxiliary local weaves:

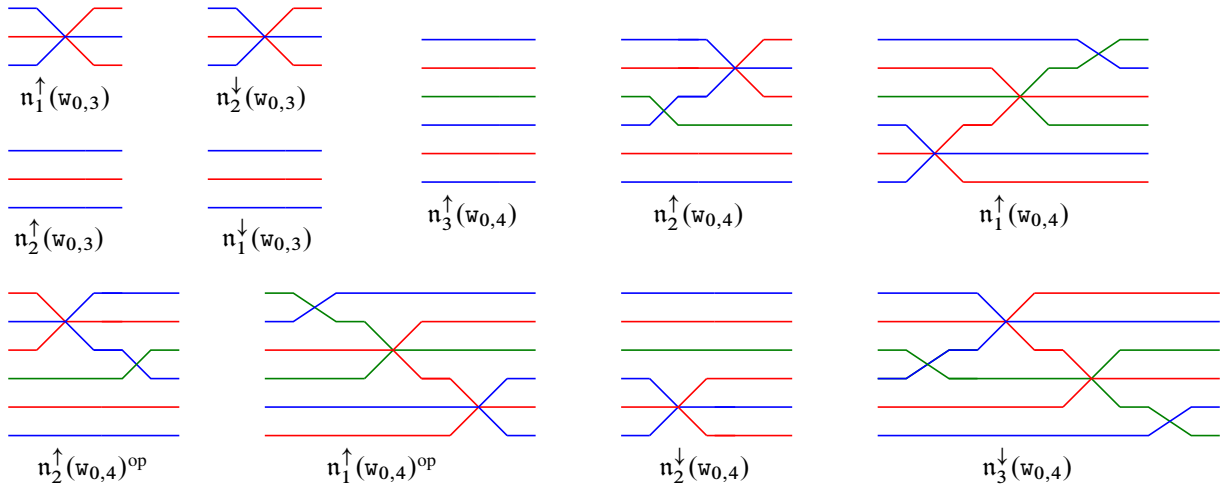


Figure 17: Instances of the weaves $n_i^\uparrow(w_{0,n})$ and $n_i^\downarrow(w_{0,n})$ and their opposites, from Definition 3.12, in some of the cases for $n = 3, 4$.

Definition 3.12 The weave $n_i^\uparrow(w_{0,n})$ is the unique horizontal weave that coincides with $n(w_{0,n})$ at the left, contains only tetravalent and hexagonal vertices, and brings the i^{th} s_1 -strand of $w_{0,n}$ to the top level, using a minimal number of weave vertices.

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

Finally, we denote by $n_i^\uparrow(w_{0,n})^{\text{op}}$ and $n_i^\downarrow(w_{0,n})^{\text{op}}$ the weaves obtained by reflecting $n_i^\uparrow(w_{0,n})$ and $n_i^\downarrow(w_{0,n})$ along a (disjoint) vertical axis.

In Definition 3.12, bringing the i^{th} s_1 -strand of $w_{0,n}$ to the top level means considering a horizontal weave which starts at $n(w_{0,n})$ 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 $w_{0,n}$ 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 [Casals and Zaslow 2022], they are all equivalent and we might as well take the one with a minimal number of vertices. Figure 17 illustrates several examples of the weaves $n_i^\uparrow(w_{0,n})$ and $n_i^\downarrow(w_{0,n})$ in Definition 3.12 for $n = 3, 4$. Note that $n_{n-1}^\uparrow(w_{0,n}) = n_{n-1}^\uparrow(w_{0,n})^{\text{op}} = n_1^\downarrow(w_{0,n}) = n_1^\downarrow(w_{0,n})^{\text{op}} = n(w_{0,n})$ for any $n \in \mathbb{N}$.

Definition 3.13 The weave $c_i^\uparrow(w_{0,n})$ is the weave obtained by horizontally concatenating the three weaves $n_i^\uparrow(w_{0,n})$, $c^\uparrow(w_i)$, and $n_i^\uparrow(w_{0,n})^{\text{op}}$, left to right, where w_i denotes the reduced expression for $w_{0,n}$ found at the right of the weave $n_i^\uparrow(w_{0,n})$.

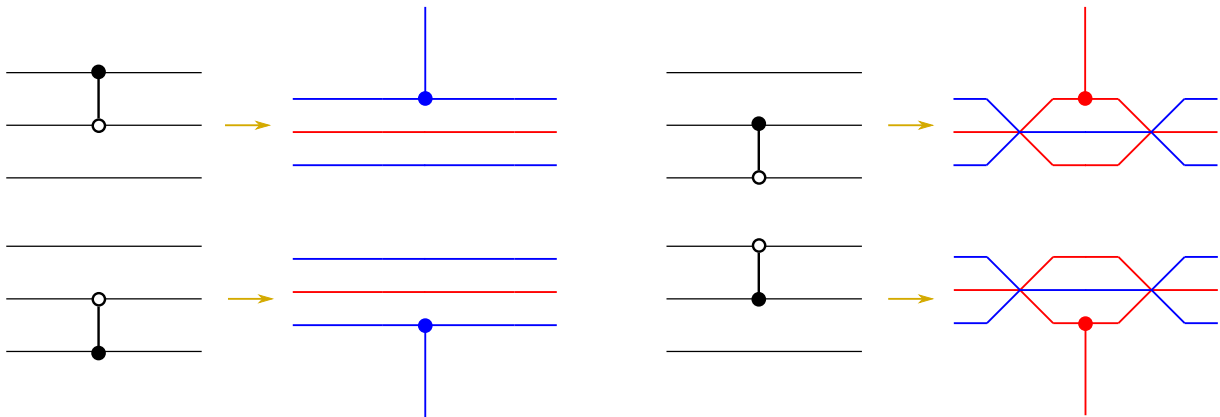


Figure 18: All possible Type 2 columns for $n = 3$ and their corresponding weaves. In detail, $c_2^\uparrow(w_{0,3})$ (upper left), $c_1^\uparrow(w_{0,3})$ (upper right), $c_1^\downarrow(w_{0,3})$ (lower left) and $c_2^\downarrow(w_{0,3})$ (lower right).

Definition 3.14 Similarly, the weave $c_i^\downarrow(w_{0,n})$ is the weave obtained by horizontally concatenating the three weaves $n_i^\downarrow(w_{0,n})$, $c^\downarrow(w_i)$ and $n_i^\downarrow(w_{0,n})^{op}$, left to right, where w_i denotes the reduced expression for $w_{0,n}$ found at the right of the weave $n_i^\downarrow(w_{0,n})$.

Figure 18 illustrates examples of the weaves $c_i^\uparrow(w_{0,n})$ and $c_i^\downarrow(w_{0,n})$ in Definitions 3.13 and 3.14 for $n = 3$. For the next definition, recall that we always label the horizontal lines, in a Type 1 or 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 3.15 (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 i^{th} horizontal line and a black vertex at the $(i + 1)^{\text{st}}$ horizontal line is the weave $c_i^\uparrow(w_{0,n})$, and the weave associated to a Type 2 column of a GP graph \mathbb{G} whose vertical edge has a black vertex at the i^{th} horizontal line and a white vertex at the $(i + 1)^{\text{st}}$ horizontal line is the weave $c_{n-i}^\downarrow(w_{0,n})$.

3.3.3 Local weaves for Type 3 columns Let us consider a column of Type 3 with labels $1, 2, \dots, n$ for the horizontal lines on the right (counting from bottom to top) and with a white lollipop attached to the i^{th} 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 $n(w_{0,n-1})$ and $n(w_{0,n})$, respectively. Hence, the weave we associate to such a Type 3 column must have these boundary conditions. Let us start with the following weave:

Definition 3.16 The weave i_i^w is the unique weave with no weave vertices, satisfying:

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




Figure 19: The weaves i_i^w (left) and i_i^b (right) from Definitions 3.16 and 3.18.

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

See Figure 19, left, for a depiction of i_i^w , illustrating what is meant by diverging upwards and downwards to vertical rays.

Note that the word $n(s_{[n-1, i]}^{-1} w_{0, n-1} s_{[n-1, n-i+1]})$ in Definition 3.16 is a reduced expression for the half-twist $w_{0, n}$. Now, the weaves i_i^w in Definition 3.16 cannot quite be the weaves for the Type 3 column yet because the labeling on the right-hand side is not $w_{0, n}$, but rather $n(s_{[n-1, i]}^{-1} w_{0, n-1} s_{[n-1, n-i+1]})$. To fix this, let n_i^w be any horizontal weave that coincides with $n(s_{[n-1, i]}^{-1} w_{0, n-1} s_{[n-1, n-i+1]})$ on the left, coincides with $n(w_{0, n})$ on the right, and with no trivalent weave vertices in the middle. Any choice of $n(w_{0, n})$ would yield an equivalent weave.

Definition 3.17 (weave for white lollipop) The weave l_i^w associated to a Type 3 column with a white lollipop at the i^{th} horizontal line is the horizontal concatenation of i_i^w and n_i^w .

Figures 20 and 21 illustrate the weaves l_1^w, l_2^w, l_3^w and l_4^w 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 i_i^w pieces, and the yellow boxes contain the n_i^w pieces. The figures also draw the corresponding pieces of the fronts $f(\mathbb{G})$, explaining the contact-geometric origin of these weaves.

The case of a column of Type 3 with labels $1, 2, \dots, n$ for the horizontal lines on the right, and a black lollipop at the i^{th} horizontal line is similar. The necessary definitions are as follows:

Definition 3.18 The weave i_i^b is the unique weave with no weave vertices, satisfying:

- (i) At its left, i_i^b coincides with the horizontal weave $n(s_{[i-1, 1]} w_{0, n-1} s_{[n-i, 1]}^{-1})$ and, at its right, i_i^b coincides with the horizontal weave $n(w_{0, n-1})$.

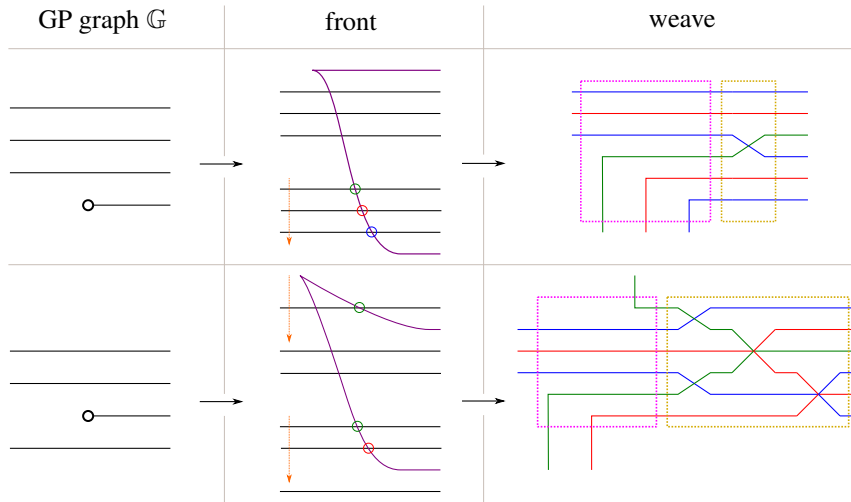


Figure 20: Weaves I_i^w associated to a white lollipop with $n = 4$, as in Definition 3.17. The first row depicts the case $i = 1$ and the second row the case $i = 2$. The weaves n_i^w are drawn within the yellow boxes. The weaves i_i^w , with the incoming weave strands, are depicted within the pink boxes.

- (ii) The weave lines in $n(s_{[i-1,1]}w_{0,n-1}s_{[n-i,1]}^{-1})$ labeled by the transpositions in the reduced expression $s_{[i-1,1]}$ diverge downwards to vertical rays.
- (iii) The weave lines in $n(s_{[i-1,1]}w_{0,n-1}s_{[n-i,1]}^{-1})$ labeled by the transpositions in the reduced expression $s_{[n-i,1]}^{-1}$ diverge upwards to vertical rays.

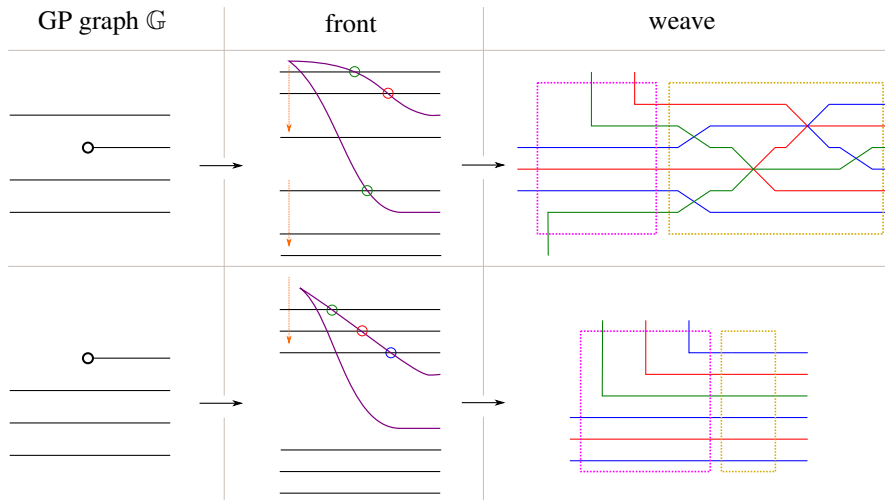


Figure 21: Weaves I_i^w associated to a white lollipop with $n = 4$. The first row depicts the case $i = 3$ and the second row the case $i = 4$. The weaves n_i^w are drawn within the yellow boxes, and the weaves i_i^w are depicted in the pink boxes.




Figure 22: The weave l_2^b associated to a black lollipop when $n = 4$, in accordance with Definition 3.19. The weave n_2^b is drawn in the yellow box, and 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.

See Figure 19, right, for a depiction of i_i^b . Similarly, we denote by n_i^b any horizontal weave which coincides with the horizontal weave $n(s_{[i-1,1]}w_{0,n-1}s_{[n-i,1]}^{-1})$ on the right, coincides with $n(w_{0,n})$ on the left, and with no trivalent weave vertices in the middle.

Definition 3.19 (weave for black lollipop) The weave l_i^b associated to a Type 3 column with a black lollipop at the i^{th} horizontal line is the horizontal concatenation of n_i^b and i_i^b .

See Figure 22 for an example of l_i^b in the case $i = 2$ and $n = 4$.

3.3.4 The initial weave Let G be a GP graph. Sections 3.3.1, 3.3.2 and 3.3.3 have explained how to obtain a weave from each of the three types of elementary columns. Note that each of these weaves coincides with $n(w_{0,k})$ and with $n(w_{0,m})$ 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 G , the associated weaves will coincide at the common side, and thus can be horizontally concatenated.

Definition 3.20 (initial weave) Let G be a GP graph. The initial weave $w(G)$ associated to G is the weave obtained by subdividing G into elementary columns and then concatenating the weaves associated with each elementary column, in the order dictated by the columns.

3.4 Topology of the initial weave

Let G be a GP graph and $w(G) \subset \mathbb{R}^2$ its initial weave. In this subsection we show how to obtain a Legendrian link $\Lambda(G) \subset (\mathbb{R}^3, \xi_{st})$ and an embedded exact Lagrangian filling $L(w) \subset (\mathbb{R}^4, \lambda_{st})$ from the weave $w = w(G)$.

3.4.1 The braid of an initial weave Suppose $w(G)$ is an N -weave. Let $K \subset \mathbb{R}^2$ be a compact subset such that $(\mathbb{R}^2 \setminus K) \cap w(G)$ contains no weave vertices. Note that the number of weave lines in $(\mathbb{R}^2 \setminus K) \cap w(G)$ and their labeling is independent of any K with this 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 3.21 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 leftmost white lollipop in \mathbb{G} .

Part of the usefulness of Definition 3.21 is the following simple lemma:

Lemma 3.22 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 $\Lambda(\mathbb{G})$.

Lemma 3.22 can be phrased as follows. Consider the Legendrian link $\Lambda(\beta(\mathbb{G})) \subset (\mathbb{R}^3, \xi_{\text{st}})$ 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 $(\mathbb{R}^3, \xi_{\text{st}})$.

3.4.2 The surface of the initial weave Let $\mathfrak{w} \subset \mathbb{R}^2$ be a weave and $\Lambda(\mathfrak{w}) \subset (\mathbb{R}^5, \xi_{\text{st}})$ the Legendrian represented by its front. By definition, the Lagrangian $L(\mathfrak{w}) \subset (\mathbb{R}^4, \lambda_{\text{st}})$ is the Lagrangian projection of $\Lambda(\mathfrak{w})$. We refer to [Casals and Zaslow 2022, Section 7.1] for details on how weaves yield exact Lagrangian fillings of Legendrian links in $(\mathbb{R}^3, \xi_{\text{st}})$, and recall that \mathfrak{w} is said to be free if $L(\mathfrak{w}) \subset (\mathbb{R}^4, \lambda_{\text{st}})$ is embedded. The following lemma is readily proven:

Lemma 3.23 Let \mathbb{G} be a GP graph and suppose its initial weave $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$ is an N -weave. Then $L = L(\mathfrak{w}) \subset (\mathbb{R}^4, \lambda_{\text{st}})$ is an embedded exact Lagrangian filling of $\Lambda(\mathbb{G})$ with Euler characteristic

$$\chi(L) = N - \#(\text{trivalent vertices of } \mathfrak{w}) = \#(\text{horizontal lines in } \mathbb{G}) - \#(\text{vertical edges in } \mathbb{G}).$$

The number of boundary components of $L(\mathfrak{w}(\mathbb{G}))$ is readily computed from $\beta(\mathbb{G})$: it is given by the number of cycles in the cycle decomposition of the Coxeter projection of $\beta(\mathbb{G})$. 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 Section 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 $(\mathbb{R}^4, \lambda_{\text{st}})$ 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 nonsquare 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 [Casals and Zaslow 2022], we can manipulate it efficiently and use it to prove the results here. 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.

3.5 Naive absolute cycles in $L(\mathfrak{w}(\mathbb{G}))$

In this subsection 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, ie 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 $S = S(\mathbb{G})$ [2013]. 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 3.24 This set of generating absolute cycles is *not* good enough in order to construct cluster structures, nor does its intersection quiver give 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 Section 3.7.

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, ie 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 $w_{0,n}$ 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 $w_{0,k}$. 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 j^{th} and the $(j + 1)^{\text{st}}$. 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




Figure 23: Left: an elementary column with four horizontal lines. Right: the corresponding cross-section for its associated weave surface $\Sigma(\mathbb{G})$.

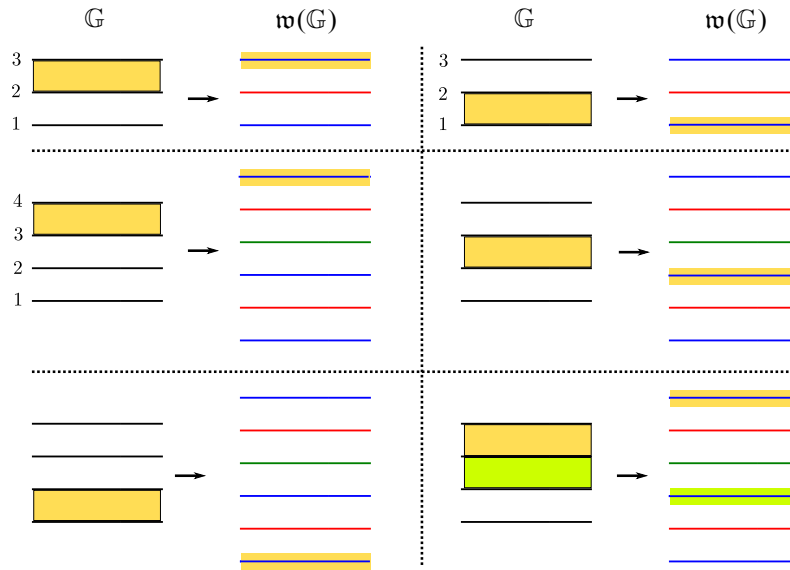


Figure 24: Associating a (piece of a) cycle in the weave for Type 1 columns. The first row depicts the two cases for $n = 3$ strands, and the second and third rows depict the three cases for $n = 4$, and an example with a union, in the lower-right corner. In all cases, the face $f \subset \mathbb{G}$ 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 highlighted in green.

of the two corresponding sheets in Σ , which in turn corresponds to the j^{th} appearance of the lowest Coxeter generator. Figure 23 illustrates a cross-section of the weave front for $n = 4$. Figure 24 illustrates all the possible cases for $n = 3, 4$.

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 j^{th} and $(j + 1)^{\text{st}}$ horizontal lines. First, for the faces bounded by any other pair of consecutive horizontal lines, say the k^{th} and $(k + 1)^{\text{st}}$ with $k \neq j$, the associated naive absolute cycle in the weave is the *unique* long l -cycle connecting the corresponding weave cycles on the two adjacent Type 1 columns. In other words, one starts at the k^{th} appearance (counting from below) of the lowest Coxeter generator on the left (see Section 3.5.1) and follows that weave line straight through any hexagonal vertices. By the construction of the weave \mathfrak{w} in Section 3.3.2, this process will go through the weave until it reaches its right-hand side at the k^{th} appearance of the lowest Coxeter generator. Figure 25 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 have the k^{th} lowest Coxeter generator in $w_{0,n}$ at the two ends.

Second, for the two faces that involve the unique vertical edge, the associated absolute cycle in \mathfrak{w} is the unique l -cycle that starts with the j^{th} appearance of the lowest Coxeter generator at its boundary end (see Section 3.5.1) and has the other end at the unique trivalent vertex of \mathfrak{w} . Figure 25 depicts examples

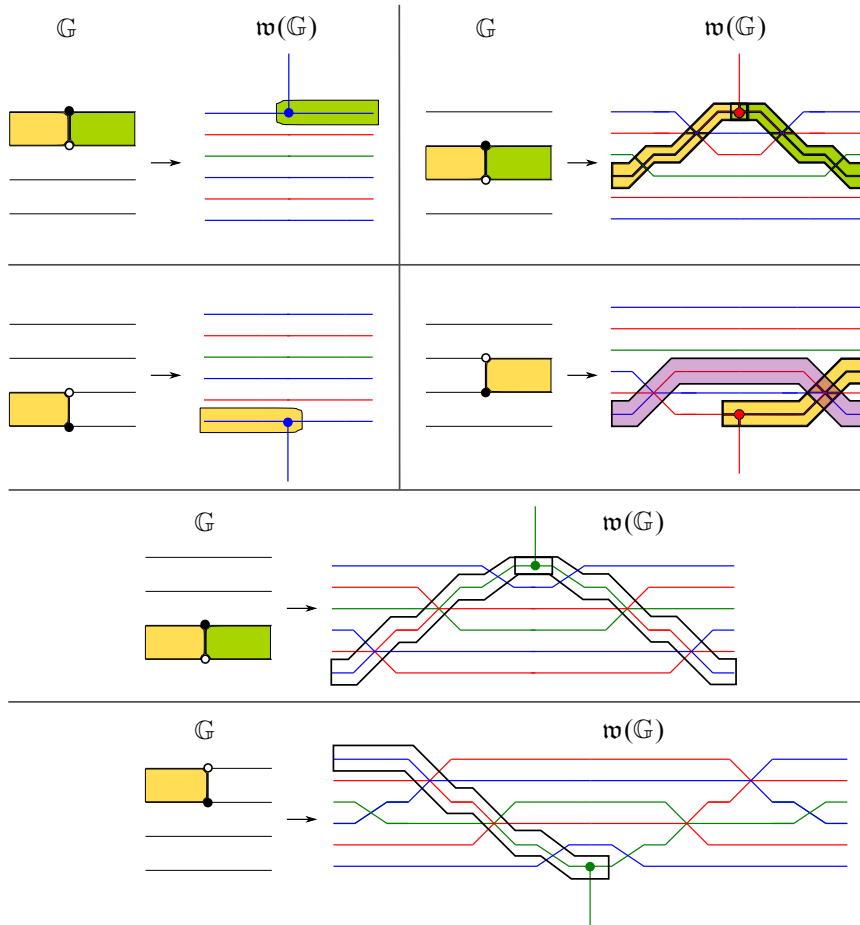


Figure 25: Several examples of faces in Type 2 elementary columns with $n = 4$ \mathbb{G} -strands, and their associated l -cycle in the weaves $w(\mathbb{G})$.

of such faces in yellow and green. Observe that, in general, these l -cycles will also go through hexagonal vertices but always have the j^{th} lowest Coxeter generator at its boundary end.

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 l -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 j^{th} horizontal line on the right. For simplicity let us assume that the horizontal lines on the right are indexed by $1, 2, \dots, n$ starting from the bottom. Following Section 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,




Figure 26: An elementary Type 3 column with a white lollipop (left) and its associated weave \mathfrak{w} (right). Center: 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.

the ends of the cycle γ_f in \mathfrak{w} must be the unique l -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 l -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 lines have become horizontal, ie the right boundary of the weave building block i_j^w (Definition 3.16). 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 $w_{0,n}$). Let us draw the weave slice as the positive braid $s_{[n-1,i]}^{-1} w_{0,n-1} s_{[n-1,n-i+1]}$ and mark the s_1 -edge in $w_{0,n-1}$ 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 i_j^w . These two weave lines are said to be obtained from the (left) s_1 -edge by a *bifurcation*.

Definition 3.25 A *bident* is a 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.

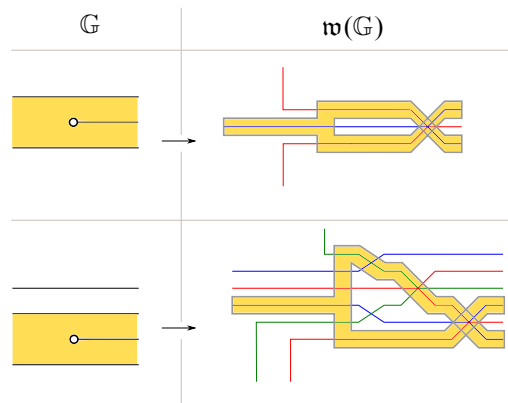


Figure 27: Left: examples of faces in Type 3 elementary columns with $n = 3, 4$ \mathbb{G} -strands. Right: the associated naive cycles, with the bidents, in the corresponding weaves $\mathfrak{w}(\mathbb{G})$.




Figure 28: 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.

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; see Figure 28. 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.

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 as Y -cycles (Definition 3.4). Thus, we prove the following:

Proposition 3.26 *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 isotopy⁷ 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 l -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.

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 l -cycle to the s_1 -edge on the left of the bident where γ_f starts. Therefore, we can consider the Y -cycle $\tilde{\gamma}_f$ which starts with this s_1 -edge at the left, propagates to the left (as an l -cycle) until the hexagonal vertex, and then contains a unique Y -vertex at the hexagonal vertex. Figure 29 depicts both cycles γ_f , at the left of the first row, and $\tilde{\gamma}_f$, at the left of the second row, in the case of the Type 3 elementary column drawn in Figure 27, upper left. By considering the description of γ_f via vertical slices, it is readily seen that γ_f is homologous to $\tilde{\gamma}_f$. 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 $\tilde{\gamma}_f$.

⁷This isotopy naturally induces a Hamiltonian isotopy between $L(\mathfrak{w})$ and $L(\mathfrak{w}')$.




Figure 29: Two homologous cycles γ_f and $\tilde{\gamma}_f$ depicted on the left. The upper-left cycle γ_f contains a bident and is not a Y -cycle; the lower-left cycle $\tilde{\gamma}_f$ is a Y -cycle. The right-hand side of each row depicts these cycles in their spatial Legendrian fronts.

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 12 — and then apply the same argument as above. □

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 $\Lambda = \Lambda(\mathbb{G})$. Section 3.5 constructed an explicit set of generators for a basis of $H_1(L)$ in terms of Y -cycles. 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 nondegenerate pairing between the absolute homology group $H_1(L)$ and the relative homology group $H_1(L, \Lambda)$:

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

Let $\{\gamma_f\}$ be the basis of naive absolute cycles constructed in Section 3.5, where the index f runs over all faces of \mathbb{G} . Consider the Poincaré dual basis $\{\eta_f\}$ on $H_1(L, \Lambda)$.⁸ 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 30, left, depicts two possible oriented lifts of the (dashed) yellow curve κ drawn to its right.

⁸The dual of an absolute cycle γ_f is constructed from the entire basis of naive absolute cycles, not just γ_f .




Figure 30: Left: in yellow, two lifts of the dashed curve κ 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 nondegeneracy 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 nontrivially 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 3.27 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 30, right, depicts a labeled dashed curve κ . From a diagrammatic perspective, it is desirable to be able to manipulate labeled dashed curves in a weave diagram in the same manner that [Casals and Zaslow 2022] explained how to combinatorially manipulate absolute cycles. For that, we have depicted in Figure 31 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.

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 Section 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




Figure 31: Four equivalence moves for labeled dashed curves.




Figure 32: Left: a face $f \in \mathbb{G}$ chosen at an elementary Type 1 column. Center: a slice of the weave $\mathfrak{w}(\mathbb{G})$ 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 $\mathfrak{w}(\mathbb{G})$, with 1 in the weave lines for the crossings inside the diamond, and 0 otherwise.

piece of the weave $\mathfrak{w}(\mathbb{G})$ associated to this column — weave lines arranged according to $w_{0,n}$ — and draw a vertical dashed curve κ_f transverse to 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 $w_{0,n}$ and the gaps in this Type 1 column (see Section 3.5.1). Locate the appearance of s_i that corresponds to a gap belonging to f and consider the 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 32, 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 $\mathfrak{w}(\mathbb{G})$ which are associated to crossings in the braid *inside* the diamond, and by assigning the intersection number 0 for all the remaining weave lines. Figure 32, right, depicts the corresponding curve κ_f with its intersection labels for the face $f \subset \mathbb{G}$.

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 $L(\mathfrak{w}(\mathbb{G}))$

Sections 3.5 and 3.6 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.⁹ A key idea in this manuscript is the consideration and study of *sugar-free* hulls, as introduced in Section 2.3. In this subsection, these two parts, sugar-free hulls and the study of homology cycles compatible with the weave $\mathfrak{w}(\mathbb{G})$, converge: we show that it is

⁹In 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.

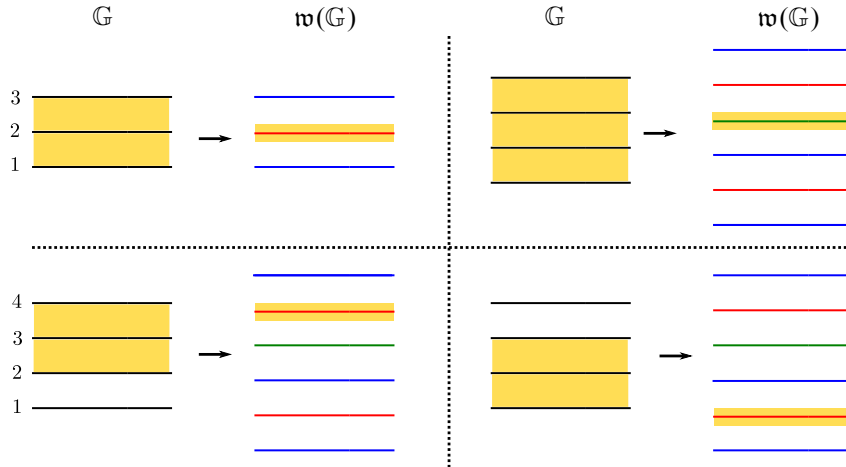


Figure 33: Associating a (piece of a) cycle in the weave for regions on a Type 1 column. The first row depicts the case where the region is the entire column for $n = 3$ and 4 strands. The second row depicts the remaining two cases for $n = 4$ strands.

possible to associate a Y -tree absolute cycle on $\mathfrak{w}(\mathbb{G})$ 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 $\mathfrak{n}(\mathfrak{w}_{0,n})$ introduced in Definition 3.11:

Lemma 3.28 *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 k^{th} level.*

Proof 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. For a union R of k consecutive gaps, the two sheets associated with ∂R intersect at the k^{th} level. □

Figure 33 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.

Proposition 3.26 showed that it is possible, up to possibly performing a weave equivalence, to represent the naive absolute cycles with Y -cycles. But these are typically immersed: it is not always possible,




Figure 34: 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.

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 3.29 (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 $n_k^\uparrow(w_{n,0})n_k^\uparrow(w_{n,0})^{\text{op}}$ and $n_k^\downarrow(w_{n,0})n_k^\downarrow(w_{n,0})^{\text{op}}$ to $\mathfrak{w}(\mathbb{G})$.*

Proof Let C be an elementary column of \mathbb{G} . By Lemma 2.4, the intersection $R \cap C$ has at most one connected component. Therefore, if R intersects C nontrivially, R must be a union of consecutive gaps in the column C . By Lemma 3.28, for a Type 1 or Type 3 elementary column C , we can represent the boundary $\partial(R \cap C)$ by a single \lrcorner -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 2.3 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 $c_k^\uparrow(w_{0,n})$, as introduced in Definition 3.13. In the construction of this weave pattern, we first bring the k^{th} strand in the bottom level upward, using the weave pattern $n_k^\uparrow(w_{0,n})$, subsequently insert a trivalent weave vertex at the top strand, and then insert the weave pattern $n_k^\uparrow(w_{0,n})^{\text{op}}$. Figure 34 depicts a case with four horizontal lines and $k = 2$. (See also Figures 17 and 18.) Since the k^{th} horizontal line is the bottom boundary of ∂R , there must be a unique hexavalent weave vertex in the $n_k^\uparrow(w_{0,n})^{\text{op}}$ that connects to the weave line corresponding to the union of all gaps in R at the right boundary. Therefore,




Figure 35: Example of the first staircase.

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 34, 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.

Next let us study the leftmost of the four staircase patterns. In this case, the chosen weave pattern $c_k^\downarrow(w_{0,n})$, as assigned in Section 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 $n_k^\uparrow(w_{0,n})$ and $n_k^\uparrow(w_{0,n})^{\text{op}}$ 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 $n_k^\uparrow(w_{0,n})^{\text{op}}$, 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 $c_k^\downarrow(w_{0,n})$, to the right. The Y–tree, again in a tripod shape, represents ∂R locally, as desired. Figure 35 depicts an example of such a tripod with $n = 4$ and $k = 2$. An analogous argument also resolves the case of the rightmost staircase pattern. Finally, by combining the local pictures for all three types of column, we conclude that there exists a representative for ∂R which is a Y–tree. \square

The Y–representability lemma allows us to introduce the following definition:

Definition 3.30 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 (nonempty) 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 3.31 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 maintaining 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 still spans 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

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 quasicluster 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, ie a quasicluster structure. 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 latter spans $H_1(L(\mathfrak{w}(\mathbb{G})))$. In a general situation, we can at least consider the following concept:

Definition 3.32 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 3.31. Namely, we start with the naive absolute basis $\{\gamma_f\}$, and then swap some of the basis elements γ_f with their corresponding initial absolute cycles $\partial\mathbb{S}_f$. In the 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 $\mathbb{S}_f = \mathbb{S}_g = \dots$ 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 $\mathbb{S}_f \subsetneq \mathbb{S}_h$, so that the pairing of $\eta_f + N$ with $\partial\mathbb{S}_h$ vanishes for all such h .

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 [Gross et al. 2018], 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 3.33 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 $\bullet \dashrightarrow \bullet$.
- (3) For each pair of quiver vertices, sum up the arrows between them.

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

Lemma 3.34 [Goncharov and Kenyon 2013, 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 3.34 uniquely determines the intersection skew-symmetric form. This intersection form, ie 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 3.35 In the case of a plabic fence \mathbb{G} , all naive absolute cycles are 1-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 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 3.36 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 $f(\mathbb{G})$ 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 Section 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

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

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

Proposition 3.37 *Let L be a connected smooth surface with boundary $\Lambda = \partial L$, and $i : T \rightarrow \Lambda$ an inclusion of a set of marked points with $\pi_0(i)$ surjective. Then*

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

and the intersection pairing $\langle \cdot, \cdot \rangle$ is nondegenerate.

It is possible 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 $H_1(L \setminus T, \Lambda \setminus T)$, as in the following diagram, where all horizontal arrows are dual lattices:

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

4 Construction of quasicluster structures on 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 Section 4.6 which, as we will prove, become the cluster \mathcal{A} -variables. The construction of the cluster structures is obtained purely by symplectic geometric means, using the results for Legendrian weaves from Section 3 above and [Casals and Zaslow 2022] and the microlocal theory of sheaves [Guillermou et al. 2012; Kashiwara and Schapira 1990; Shende et al. 2017]. 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-theoretical moduli spaces and also deduce the factoriality of its ring of regular functions $\mathbb{C}(\mathfrak{M}(\Lambda, T))$. Next, on the moduli space $\mathfrak{M}(\Lambda, T)$, we construct a new family of 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 prove that, within these special microlocal merodromies, 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 nonfree weaves. These results together with [Berenstein et al. 2005] allow us to conclude the 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 $\Lambda \subset (\mathbb{R}^3_{\xi_{st}})$ be a Legendrian link and $T \subset \Lambda$ a set of marked points, and consider the moduli stacks $\mathcal{M}_1(\Lambda)$ and $\mathfrak{M}(\Lambda, T)$ discussed in Section 2.7. These stacks classify (complexes of) constructible sheaves on \mathbb{R}^2 with a singular support condition. In this subsection, we provide Lie-theoretical descriptions for $\mathcal{M}_1(\Lambda)$ and $\mathfrak{M}(\Lambda, T)$ which are suited for our computations, using [Kashiwara and Schapira 1990] and closely following [Shende et al. 2017, Sections 3.3 and 5].¹⁰ 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.

Given a cooriented 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 $\mathbb{R}^2 \setminus \pi_F(\Lambda)$,
- 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 cells (that contain that stratum in their closure). The direction of the arrow is opposite to the coorientation of the front $\pi_F(\Lambda)$.

The following is then proven in [Shende et al. 2017, 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 $\mathcal{M}(Q_F(\Lambda))$ 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 36, 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 36, right.*

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

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 [Shende et al. 2017, Section 5], and we briefly summarize as follows. Let S be the set of singular points (crossings and cusps) in $\pi_F(\Lambda)$ and note that 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:

¹⁰An expert in the results of [Kashiwara and Schapira 1990; Shende et al. 2017] might be able to quickly move forward to Section 4.2.




Figure 36: Identity condition at cusps and exactness condition at crossings.

- At a cusp, as in Figure 36, left, the condition $g \circ f = \text{id}_V$ forces the composition $\ker g \hookrightarrow W \twoheadrightarrow \text{coker } f$ to be an isomorphism. By definition, we glue $\ker g$ and $\text{coker } f$ using this isomorphism.
- At a crossing, as in Figure 36, 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) $\text{coker } f_{sw} \hookrightarrow \frac{V_n}{V_s} \twoheadrightarrow \text{coker } f_{en},$
 $\text{coker } f_{wn} \longleftarrow \frac{V_n}{V_s} \longleftarrow \text{coker } f_{se},$

(4-2) $\ker f_{sw} \hookrightarrow \ker(f_{wn} \circ f_{sw}) = \ker(f_{en} \circ f_{se}) \twoheadrightarrow \ker f_{en},$
 $\ker f_{wn} \longleftarrow \ker(f_{wn} \circ f_{sw}) = \ker(f_{en} \circ f_{se}) \longleftarrow \ker f_{se},$

(4-3) $\text{coker } f_{sw} \xrightarrow{f_{wn}} \text{coker } f_{en},$
 $\ker f_{wn} \xleftarrow{f_{sw}} \ker f_{se}.$

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 [Shende et al. 2017, 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 Section 2.7.3, is a point in $\mathcal{M}_1(\Lambda)$, which is combinatorialized via Proposition 4.1, together with a framing for the local system Φ , ie a trivialization of the (trivial) local system $\Phi|_{\Lambda \setminus T}$. Here two framings are considered *equivalent* if they differ by a global scaling \mathbb{C}^\times factor, and thus $\dim \mathfrak{M}(\Lambda, T) = \dim \mathcal{M}_1(\Lambda) + |T| - 1$.

4.1.1 Description for a 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

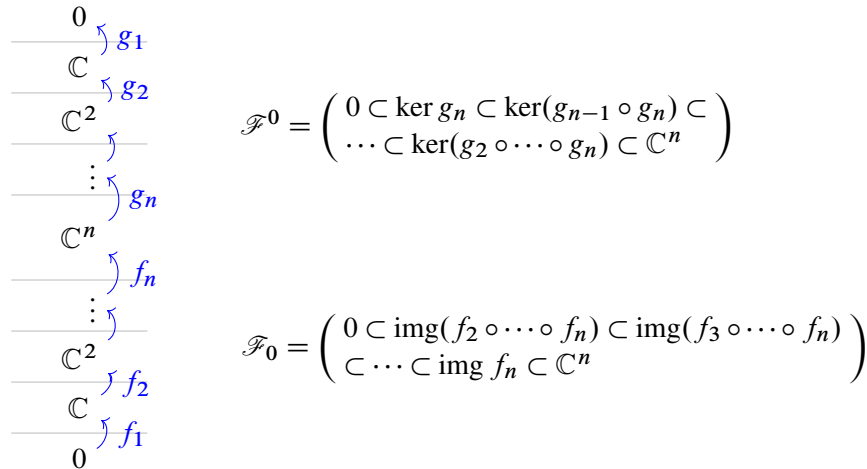


Figure 37: The pair of flags associated to a Type 1 column.

in terms of configurations of flags, as follows. The front projection $\pi_F(\Lambda) = f(\mathbb{G})$ 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 37 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 them.

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 $\mathcal{F} = (0 \subset \mathcal{F}_1 \subset \dots \subset \mathbb{C}^n)$ and $\mathcal{F}' = (0 \subset \mathcal{F}'_1 \subset \dots \subset \mathbb{C}^n)$, we write

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

$$\mathcal{F} = \mathcal{G}_0 \stackrel{s_{i_1}}{\sim} \mathcal{G}_1 \stackrel{s_{i_2}}{\sim} \mathcal{G}_2 \stackrel{s_{i_3}}{\sim} \dots \stackrel{s_{i_l}}{\sim} \mathcal{G}_l = \mathcal{F}'$$

and $s_{i_1}s_{i_2} \dots s_{i_l}$ 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 $\mathcal{F} \stackrel{w}{\sim} \mathcal{F}'$, then, for each choice of reduced word (i_1, \dots, i_l) for w , there exists a unique sequence of flags $(\mathcal{G}_k)_{k=0}^l$ that relate the two flags \mathcal{F} and \mathcal{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 $(\mathcal{L}_0, \mathcal{L}^0)$, and the pair of flags to the right of a Type 2 column is $(\mathcal{R}_0, \mathcal{R}^0)$. If there is a crossing in the bottom region at the i^{th} gap, counting from the bottom in both the front and the GP graph, then, from the exactness condition at the crossing, we obtain the constraints

$$(4-4) \quad \mathcal{L}_0 \stackrel{s_i}{\sim} \mathcal{R}_0 \quad \text{and} \quad \mathcal{L}^0 = \mathcal{R}^0.$$

Similarly, if there is a crossing in the top region at the i^{th} 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

$$(4-5) \quad \mathcal{L}_0 = \mathcal{R}_0 \quad \text{and} \quad \mathcal{L}^0 \stackrel{s_{n-i}}{\sim} \mathcal{R}^0.$$

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.

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 $h: \mathbb{C}^{n-1} \rightarrow \mathbb{C}^n$ between the two (middle) adjacent ambient vector spaces is injective. Given any flag $\mathcal{F} = (0 \subset \mathcal{F}_1 \subset \dots \subset \mathcal{F}_{n-1} = \mathbb{C}^{n-1})$ in \mathbb{C}^{n-1} , we can use h to naturally extend it to a flag $h(\mathcal{F})$ in \mathbb{C}^n by defining

$$h(\mathcal{F}) := (0 \subset h(\mathcal{F}_1) \subset h(\mathcal{F}_2) \subset \dots \subset h(\mathcal{F}_{n-1}) \subset \mathbb{C}^n).$$

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 local example in Figure 38.

The green maps in the bottom region define the extension $h(\mathcal{L}_0)$. By the exactness of the quadrilaterals in the top region, the red maps define the extension $h(\mathcal{L}^0)$. In particular, the extensions $h(\mathcal{L}_0)$ and $h(\mathcal{L}^0)$ 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 i^{th} gap with $0 \leq i \leq n - 1$, counting from below in the GP graph,¹¹ then the corresponding relative position conditions are

$$(4-6) \quad h(\mathcal{L}_0) \stackrel{s_{n-1} \dots s_{i+1}}{\sim} \mathcal{R}_0 \quad \text{and} \quad h(\mathcal{L}^0) \stackrel{s_{n-1} \dots s_{n-i}}{\sim} \mathcal{R}^0.$$

¹¹The case $i = 0$ is a lollipop at the bottom, and $i = n - 1$ is a lollipop at the top.

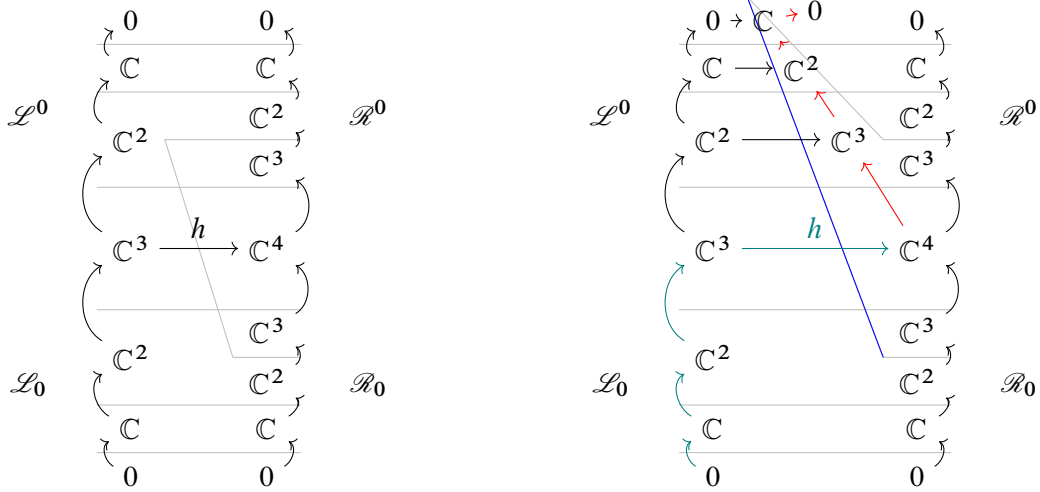


Figure 38: Pulling up a left cusp.

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

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

Similar to the white lollipop case, the extension of $(\mathcal{R}_0, \mathcal{R}^0)$ to $(h^{-1}(\mathcal{R}_0), h^{-1}(\mathcal{R}^0))$ can be achieved geometrically by a sequence of RII moves that pulls the right cusp downward in the front projection, as depicted in Figure 39. Note that the green maps in bottom region define the extension $h^{-1}(\mathcal{R}_0)$, whereas the red maps in the top region define the extension $h^{-1}(\mathcal{R}^0)$.




Figure 39: Pulling down a right cusp.

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

$$(4-7) \quad \mathcal{L}_0^{s_{i-1} \dots s_1} h^{-1}(\mathcal{R}_0) \quad \text{and} \quad \mathcal{L}^0 s_{n-i+1} \dots s_1 h^{-1}(\mathcal{R}^0).$$

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 $h_i: V_{i-1} \rightarrow V_i$ that are either injective with a 1-dimensional cokernel or surjective with a 1-dimensional kernel. The above discussion proves:

Lemma 4.2 *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 $(\mathcal{F}, h) \sim (\mathcal{F}', h')$ for a collection of elements $g_i \in \text{GL}(V_i)$ such that $h_i \circ g_{i-1} = g_i \circ h'_i$.

In the flag description of Lemma 4.2, 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 as 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.

4.1.2 Description for (−1)–closures Finally, there is another description of $\mathcal{M}_1(\Lambda)$ and $\mathfrak{M}(\Lambda, T)$ 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 4.2, 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 like the blue one in Figure 38 all the way to the left and push strands like the blue one in Figure 39 all the way to the right (see also Lemma 3.22). 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, [Shende et al. 2017, Proposition 1.5], or Proposition 4.1, reads:

Lemma 4.3 Let $\beta = (i_1, i_2, \dots, i_l) \in Br_h^+$ 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


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

4.2 Factoriality property

In the upcoming construction of a cluster \mathcal{A} -structure for the moduli space $\mathcal{M}(\Lambda, T)$, we shall need that the coordinate ring $\mathbb{C}(\mathcal{M}(\Lambda, T))$ is a unique factorization domain (aka. 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 $\mathbb{C}(\mathcal{M}(\Lambda(\mathbb{G}), T))$ are factorial if \mathbb{G} is shuffle.

Proposition 4.4 Let \mathbb{G} be a GP graph with $\beta(\mathbb{G})$ a Δ -complete braid. Then the moduli space $\mathcal{M}(\Lambda(\mathbb{G}), T)$ is an affine variety whose coordinate ring is factorial.

Proof Since the moduli space $\mathcal{M}(\Lambda, T)$ is a 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 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 \mathcal{F}_l at the dashed line (after γ on the right or before Δ on the left) be the standard one, ie mapping \bar{e}_i to 1 for each consecutive quotient $\text{Span}\{e_1, \dots, e_i\} / \text{Span}\{e_1, \dots, e_{i-1}\} \cong \text{Span}(\bar{e}_i)$.

Let (i_1, \dots, 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 \mathcal{F}_l can be recorded by the permutation matrix w_0 . Starting from the flag \mathcal{F}_l , the flags $\mathcal{F}_{l-1}, \mathcal{F}_{l-2}, \dots$ to the left of \mathcal{F}_l can then be given by

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

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

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

which is equivalent to requiring that $B_{i_1}(z_1)B_{i_2}(z_2)\cdots B_{i_k}(z_l)w_0$ 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 [Casals et al. 2020], which has a factorial coordinate ring.¹²

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 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

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

If there is an element in $\mathbb{C}(\mathfrak{M}(\Lambda, T'))$ admitting two nonequivalent factorizations, then these two factorizations are still valid and nonequivalent in $\mathbb{C}(\mathfrak{M}(\Lambda, T))$, contradicting the fact that $\mathbb{C}(\mathfrak{M}(\Lambda, T))$ is factorial. Thus, we can conclude that $\mathbb{C}(\mathfrak{M}(\Lambda, T'))$ 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

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

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

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 of [Jin and Treumann 2017], or [Ekhholm et al. 2016; Casals and Ng 2022], imply that there are open embeddings

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

¹²We thank Eugene Gorsky for an explanation of why this is the case. See also upcoming work of the first author with Gorsky and coauthors, where this is written in detail.

whose domains 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 [Casals and Zaslow 2022]; we shall use it in the proof of Theorem 1.1. The definition of $\mathcal{M}_1(\mathfrak{w})$ from [Casals and Zaslow 2022] is as follows:

Definition 4.5 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} comprises tuples of flags, as follows:

- (i) There is a flag $\mathcal{F}^\bullet(F)$ assigned to each face F of the weave \mathfrak{w} , ie 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 $\mathcal{F}^\bullet(F_1)$ and $\mathcal{F}^\bullet(F_2)$ are in relative position $s_i \in S_n$, ie they must satisfy

$$\mathcal{F}_j(F_1) = \mathcal{F}_j(F_2) \quad \text{for } 0 \leq j \leq N \text{ with } j \neq i \quad \text{and} \quad \mathcal{F}_i(F_1) \neq \mathcal{F}_i(F_2).$$

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 $\mathcal{M}_1(\mathfrak{w}) := \tilde{\mathcal{M}}_1(\mathfrak{w}) / \text{PGL}_n$.

By Section 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 [Casals and Zaslow 2022, 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 [Casals and Zaslow 2022] 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–Tremann 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 by $\mathfrak{M}(\mathfrak{w}, T)$, 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 Section 4.6.

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 Section 4.3. We now build a function

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

associated to each Y -cycle γ , generalizing [Casals and Zaslow 2022, Section 7] — see also [Shende et al. 2017, Section 5.1] — to our context. First we observe that the data of $\mathcal{M}_1(\mathfrak{w})$ associates a flag \mathcal{F} in each connected component of the complement of \mathfrak{w} in \mathbb{R}^2 . We associate the 1-dimensional vector space




Figure 40: A weave for the unique filling of the max tb unknot and microlocal parallel transports from its sheaf quantization.

$\mathcal{F}_i / \mathcal{F}_{i-1}$ to the i^{th} sheet in the lift of each connected component. Then, across the lifts of each weave line, we define two linear isomorphisms

$$(4-8) \quad \begin{aligned} \psi_+ : \frac{\mathcal{L}_i}{\mathcal{L}_{i-1}} &\hookrightarrow \frac{\mathcal{L}_{i+1}}{\mathcal{L}_{i-1}} = \frac{\mathcal{R}_{i+1}}{\mathcal{R}_{i-1}} \twoheadrightarrow \frac{\mathcal{R}_{i+1}}{\mathcal{R}_i}, & \mathcal{L}_{i+1}/\mathcal{L}_i &\xrightarrow{\mathcal{L}_{i+1} = \mathcal{R}_{i+1}} \mathcal{R}_{i+1}/\mathcal{R}_i \\ \psi_- : \frac{\mathcal{R}_i}{\mathcal{R}_{i-1}} &\hookrightarrow \frac{\mathcal{R}_{i+1}}{\mathcal{R}_{i-1}} = \frac{\mathcal{L}_{i+1}}{\mathcal{L}_{i-1}} \twoheadrightarrow \frac{\mathcal{L}_{i+1}}{\mathcal{L}_i}. & \mathcal{L}_i/\mathcal{L}_{i-1} &\xrightarrow{\mathcal{L}_i = \mathcal{R}_i} \mathcal{R}_i/\mathcal{R}_{i-1} \end{aligned}$$

Note that ψ_{\pm} are isomorphisms because \mathcal{L} and \mathcal{R} are in s_i -transverse position, as they are separated by a weave line labeled with s_i . Now, given a loop γ on L , we may perturb it so that it intersects with any lifts of weave lines transversely. Then, by composing several of the isomorphisms ψ_{\pm} above and their inverses, 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 nonzero scalar ψ_{γ} . This nonzero 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 ,¹³ as the following illustrates:

Example 4.6 Consider the weave with a unique trivalent vertex, which depicts a Lagrangian 2-disk filling, as drawn in blue in Figure 40, 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. 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 40, left. By definition, the microlocal parallel transport ψ_{γ} should be the map in Figure 40, center, which is the linear map that projects parallel to the line l_2 . 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 in Figure 40, right, 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 $L(\mathfrak{w})$.


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

¹³They give a *twisted* local system as in [Guilleumou 2023, Part 13], or a twisted flat connection as in [Gaiotto et al. 2013, Part 10].

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 3.29, 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 l–cycle, ie an equivalent monochromatic edge, we may use it to compute X_γ explicitly, as follows. In a neighborhood of a short l–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 \mathcal{F}^j in each region for $j \neq i$, and for $j = i$ we additionally require the data in each region of a line l in the 2–dimensional space $V := \mathcal{F}^{i+1}/\mathcal{F}^{i-1}$. This is the data of four lines $a, b, c, d \subset V$. 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 [Fock and Goncharov 2006b; Shende et al. 2019; Casals and Zaslow 2022], is not particularly new. It is also possible to define X_γ directly and combinatorially from the Y–trees, in line with [Casals and Zaslow 2022, 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 [Shende et al. 2019], and under the more general weave mutation due to [Casals and Zaslow 2022]. 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 [Casals and Zaslow 2022, 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 yield equivalent, or even quasiequivalent, 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 do in a moment. By [Berenstein et al. 2005], 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 Sections 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 subsection, we describe a way to add signs to monodromies to obtain a (noncanonical) isomorphism between $\mathcal{M}_1(\mathfrak{w})$ and $\text{Loc}_1(L)$. 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 (noncanonically) also isomorphic to the category of local systems. In terms of weave combinatorics, we proceed as follows:

Definition 4.7 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 $P \sqcup T$. By definition, 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.

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.

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

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

In fact, we can do better than an arbitrary isomorphism $\mathcal{M}_1(\mathfrak{w}) \xrightarrow{\cong} \text{Loc}_1(L)$. From Section 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 4.8 A coherent set C of sign curves on L is said to be *compatible* if, for all initial absolute cycles $\gamma \in \mathfrak{S}(\mathbb{G})$,

$$\Phi_C(p)(\gamma) = X_\gamma(p) := -\psi_\gamma(p) \quad \text{for all } p \in \mathcal{M}_1(\mathfrak{w}).$$

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 *noncritical* otherwise. By construction, each critical Type 2 column has a unique initial cycle γ that ends there.

If a Type 2 column is noncritical, 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 nontrivially. 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 $\text{Loc}_1(L)$ 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.

Proposition 4.9 *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 (noncanonical) compatible set of sign curves.*

Proof Both weave equivalences and weave mutations are local operations on the weave. Therefore, it suffices to verify that compatible sets of sign curves can be constructed locally, before and after such local operations. That is, locally in a neighborhood where the weave equivalence or mutation is going to be performed, we want to argue that any given compatible set of sign curves on that piece of the initial weave — before an equivalence or mutation — we can construct a compatible set of sign curves afterwards, locally on that piece of the weave after the operation.

For weave equivalences, Moves I, IV and V in Definition 3.2 do not involve any trivalent weave vertices. Thus, sign curves that pass through any of these local pictures can be carried through these equivalences using planar homotopies. In contrast, Move II (the push-through move), III and VI do involve trivalent weave vertices. In the case of Move VI, the weave lines lift to sheets that are not adjacent to each other; therefore, the weave line with no trivalent vertex (yellow in Figure 12) can be ignored when studying the set of compatible sign curves, reducing to the constant case of a trivalent vertex. By Remark 3.3, Move III is a concatenation of Moves I and II. Thus, the only weave equivalence move that remains to be studied is Move II, which will be discussed in a moment. For weave mutations, it suffices to check mutations along short l -cycles, since any Y -tree is weave equivalent to a short l -cycle by Proposition 3.5. In conclusion, we need to study compatible sets of sign curves locally near a push-through and a weave mutation.

A priori, we must study sign curves that arrive at the trivalent vertices from different faces of the weave; a face being any connected component of the complement of the weave lines. That said, if a sign curve is incident to a trivalent weave vertex, we can apply a planar homotopy to the sign curve so that it arrives at the trivalent vertex from any of the another faces near the trivalent vertex. This is depicted in Figure 41. Therefore, it suffices to study the case that a sign curve arrives at a trivalent only from one of the three




Figure 41: Applying a planar homotopy to a sign curve so that it arrives at the trivalent vertex from a different face. (The index i can be either 1 or 2.)

faces near the trivalent vertex. (If the curve arrived from another face, we could homotope the curve and, locally in a small neighborhood around the vertex, have it arrive from another face.)

With this reduction, it suffices to study compatible sets of sign curves locally near a push-through and a weave mutation which arrive from one face (of our choice) at each trivalent vertex. Up to symmetry, there are only three cases to check, shown in Figure 42. The figure illustrates how to resolve the problem at hand: for each set of compatible sign curves before the equivalence or mutation (on the left of each diagram), we can construct a set of compatible sign curves afterwards (on the right of each diagram). \square

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 subsection, 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 $A_\eta: \mathfrak{M}(\Lambda, T) \rightarrow \mathbb{C}$ 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 $\Phi_C: \mathcal{M}_1(\mathfrak{w}) \xrightarrow{\cong} \text{Loc}_1(L)$ in Section 4.5, together with a framing (trivialization) of the rank-1 local system Φ on the connected components of $\Lambda \setminus T = (\partial L) \setminus T$.




Figure 42: Existence of compatible sets of sign curves before and after weave equivalences (left) and weave mutations (center and right). (The index i can be either 1 or 2.)




Figure 43: 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 .

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 nonzero vector in $\eta(\phi_s) \in \Phi_t$. The ratio $\eta(\phi_s)/\phi_t$ is a nonzero 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 4.10 The function $A_\eta: \mathfrak{M}(\mathfrak{w}, T) \rightarrow \mathbb{C}^\times$ 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 extend to a regular function on $\mathfrak{M}(\Lambda, T)$. We emphasize that the decorations in $\mathfrak{M}(\Lambda, T)$ are needed in order to define A_η , and thus microlocal merodromies cannot be defined in $\mathcal{M}_1(\Lambda)$.

The microlocal merodromies associated to relative cycles coming from marked points are nonvanishing. Indeed, for each marked point $t \in T$, we pick a small half-disk neighborhood U_t of t , as in Figure 43, and define $\xi_t := \partial U_t$. By definition,

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

In particular, by using the ratio $\xi_t(\lambda)/\rho$, 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 exact sequence of lattices

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

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

Corollary 4.11 $\prod_{t \in T} A_t = 1$ and hence A_t is a unit in $\mathbb{O}(\mathfrak{M}(\Lambda, T))$ for every $t \in T$.


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

Corollary 4.12 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 for $t \in T$.

Proof This follows from the fact that $\ker \pi = \text{Span}\{\xi_t \mid t \in T\}$. □

The latter corollary starts to hint at the quasicluster 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 for $t \in T$ 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 commutative diagram of vector space inclusions



and assume that $0 \rightarrow V_s \rightarrow V_w \oplus V_e \rightarrow V_n \rightarrow 0$ is exact. Let $\alpha_s, \alpha_w, \alpha_e$ and α_n be nonzero top-dimensional (volume) forms in V_s, V_w, V_e and 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 4.13 In the context of a diagram as above, we define

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

The top form $\alpha_w \wedge^{\alpha_s} \alpha_e$ is nonzero on V_n and does not depend on the choice of β_w or β_e . By definition, the ratio $(\alpha_w \wedge^{\alpha_s} \alpha_e)/\alpha_n$ 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. Consider a flag $\mathcal{F} = (0 \subset \mathcal{F}_1 \subset \mathcal{F}_2 \subset \dots \subset \mathcal{F}_n = \mathbb{C}^n)$ with a choice of $\phi_i \neq 0 \in \mathcal{F}_i/\mathcal{F}_{i-1}$ for all $i \in [1, n]$. The choice of such $\phi = (\phi_i)$ for $i \in [1, n]$ is said to be a *framing* for the flag \mathcal{F} . Given such *framed flag* (\mathcal{F}, ϕ) , we can construct top forms $\alpha_i \in \bigwedge^i \mathcal{F}_i$ for $i \in [1, n]$, by first lifting each ϕ_j to a vector in $\tilde{\phi}_j \in \mathcal{F}_j$ for $j \in [1, n]$, and then taking ordered wedges, leading to the forms

$$(4-9) \quad \alpha_i := \tilde{\phi}_i \wedge \tilde{\phi}_{i-1} \wedge \dots \wedge \tilde{\phi}_1 \quad \text{for } i \in [1, n].$$

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

Definition 4.14 Given a flag \mathcal{F} , a collection $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$ of nonvanishing forms $\alpha_i \in \bigwedge^i \mathcal{F}_i$ for $i \in [1, n]$ is said to be a *decoration* on the flag \mathcal{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 (ϕ_1, \dots, ϕ_n) and decorations $(\alpha_1, \alpha_2, \dots, \alpha_n)$ of a flag \mathcal{F} are equivalent pieces of data. By definition, two decorated (or framed) flags (\mathcal{L}, α) and (\mathcal{R}, β) are in s_i -transverse position if the underlying flags $\mathcal{L} \stackrel{s_i}{\sim} \mathcal{R}$ are in s_i -transverse position, ie s_i -transversality does not see decorations (or framings).

Suppose (\mathcal{L}, λ) and (\mathcal{R}, ρ) are two framed flags such that $\mathcal{L} \stackrel{s_i}{\sim} \mathcal{R}$. Let α and β be the decorations constructed from λ and ρ , respectively. Consider the parallel transport maps ψ_\pm defined in (4-8). The images $\psi_+(\lambda_i)$ and $\psi_-(\rho_i)$ are readily computed in terms of decorations as follows:

Lemma 4.15 We have

$$\psi_+(\lambda_i) = \frac{\alpha_i \wedge^{\alpha_{i-1}} \beta_i}{\beta_{i+1}} \rho_{i+1} \quad \text{and} \quad \psi_-(\rho_i) = \frac{\beta_i \wedge^{\beta_{i-1}} \alpha_i}{\alpha_{i+1}} \lambda_{i+1}.$$

Proof Let $\tilde{\lambda}_i$ and $\tilde{\rho}_{i+1}$ be lifts of λ_i and ρ_{i+1} . By construction, the framing $\psi_+(\lambda_i)$ can be obtained as follows. First, lift $\lambda_i \in \mathcal{L}_i / \mathcal{L}_{i-1}$ to a vector $\tilde{\lambda}_i \in \mathcal{L}_i$ and consider this vector as $\tilde{\lambda}_i \in \mathcal{L}_{i+1}$ via the inclusion $\mathcal{L}_i \subset \mathcal{L}_{i+1}$. Then, using $\mathcal{L}_{i+1} = \mathcal{R}_{i+1}$, we can view $\tilde{\lambda}_i \in \mathcal{R}_{i+1}$ and thus finally $\psi_+(\lambda_i) = \pi(\tilde{\lambda}_i)$, where $\pi : \mathcal{R}_{i+1} \rightarrow \mathcal{R}_{i+1} / \mathcal{R}_i$ is the quotient map. Each of $\psi_+(\lambda_i)$ and ρ_{i+1} is a (volume) 1-form on $\mathcal{R}_{i+1} / \mathcal{R}_i$, and can be pulled back via π to 1-forms in \mathcal{R}_{i+1} . By wedging these forms with (any) top form in \mathcal{R}_i , such as β_i , we obtain the top forms $\tilde{\lambda}_i \wedge \beta_i$ and $\tilde{\rho}_{i+1} \wedge \beta_i$. Since we wedged both $\psi_+(\lambda_i)$ and ρ_{i+1} with the same form β_i , their ratios are equal:

$$\frac{\psi_+(\lambda_i)}{\rho_{i+1}} = \frac{\tilde{\lambda}_i \wedge \beta_i}{\tilde{\rho}_{i+1} \wedge \beta_i}.$$

By (4-9), we also have $\alpha_i = \tilde{\lambda}_i \wedge \alpha_{i-1}$ and $\beta_{i+1} = \tilde{\rho}_{i+1} \wedge \beta_i$. Therefore,

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

The equality for ψ_- is obtained similarly. □

Remark 4.16 By Lemma 4.15, the inverses of ψ_{\pm} are computed analogously. Namely,

$$\psi_+^{-1}(\rho_{i+1}) = \frac{\beta_{i+1}}{\alpha_i \wedge^{\alpha_{i-1}} \beta_i} \lambda_i \quad \text{and} \quad \psi_-^{-1}(\lambda_{i+1}) = \frac{\alpha_{i+1}}{\beta_i \wedge^{\beta_{i-1}} \alpha_i} \rho_i.$$

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 $\mathcal{L} = \mathcal{F}_0, \mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_l = \mathcal{R}$. For each flag \mathcal{F}_i for $0 < i < l$ we choose a sequence of top forms $\alpha_j = (\alpha_{i,j})$. Since the parallel transport along η consists of compositions of linear isomorphisms like the maps ψ_{\pm} in Lemma 4.15, or their inverses, Lemma 4.15 allows us to compute A_{η} .

Example 4.17 Consider the cross-section of a weave surface depicted in Figure 44, and let η be the blue relative cycle. The sequences of top forms α and δ are determined by the decorations λ and ρ , respectively. The tuples of top forms β and γ are chosen arbitrarily. Note that $\alpha_0, \beta_0, \gamma_0$ and δ_0 are trivial. Let us denote the ψ_{\pm} maps associated to each of the three crossings by ${}_i\psi_{\pm}$, where $i \in [1, 3]$, $i = 1$ being associated to the leftmost crossing, $i = 2$ to the center crossing and $i = 3$ to the rightmost crossing.

By definition, the microlocal merodromy along η is

$$A_{\eta} = \frac{\psi_{\eta}(\lambda_2)}{\rho_3} = \frac{({}_3\psi_+ \circ {}_2\psi_+ \circ {}_1\psi_-^{-1})(\lambda_2)}{\rho_3}.$$

By Lemma 4.15, each of the microlocal merodromies ${}_i\psi_{\pm}$ and their inverses are computed as

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




Figure 44: Computation of a merodromy.

Two observations based on this computation: First, the right-hand side of this expression shows that A_η depends on the underlying undecorated flag associated with the β , but it is invariant under any nonzero rescaling of the decoration β . This is a general fact. Namely, the function A_η does depend on the intermediate flags between the two flags at the endpoints; nevertheless, it does not depend on the decorations of these intermediate flags.

Second, a reason for the decoration γ not appearing in the computation of A_η above is that the flag associated to γ is uniquely determined by the flags associated to β and δ . Observe that in the case that the slice along η yields a reduced braid word, the intermediate flags are uniquely determined by the flags at the endpoints and thus the microlocal merodromy only depends on the decorated flags at the endpoints. For more general computations, the study of microlocal merodromies involves understanding properties, such as regularity, of products of crossing values and inverses thereof.

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* functions. Indeed, let us consider the following special family of merodromies. By Section 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 $H_1(L, \Lambda)$, where $\pi: H_1(L \setminus T, \Lambda \setminus T) \rightarrow H_1(L, \Lambda)$ is the natural projection map. By Corollary 4.12, we conclude that different choices of η_f only change A_f by a multiple of units.

Proposition 4.18 *Let f be a face in a GP graph \mathbb{G} . Then the microlocal merodromy*

$$A_f: \mathfrak{M}(\Lambda(\mathbb{G}), T) \rightarrow \mathbb{C}$$


is a regular function.

Proof Suppose that the face $f \in \mathbb{G}$ corresponds to the k^{th} gap in \mathbb{G} , counting from the bottom. Then the associated relative cycle η_f can be written as

$$\eta_f = \sum_{i=1}^k \xi_i,$$

where the relative cycles ξ_i are described as follows. Consider the braid word $\mathfrak{w}_{0,n}$; then ξ_i is the relative cycle given by the i^{th} strand in $\mathfrak{w}_{0,n}$, counting from the bottom on the left, when considered as a slice of

the weave $\mathfrak{w}(\mathbb{G})$ along η_f . The following figure depicts such ξ_i for $n = 4$:



Now, the microlocal parallel transport along ξ_i , from the i^{th} strand on the left to the $(n-i+1)^{\text{st}}$ strand on the right, can be computed via Lemma 4.15. In particular, each microlocal merodromy A_{ξ_i} has contributions from $i - 1$ crossings because $\mathfrak{w}_{0,n}$ is a half-twist, and thus it is obtained after composing $i - 1$ instances of ψ_{\pm}^{-1} . In such a slice spelling $\mathfrak{w}_{0,n}$, let (\mathcal{L}, α) be the decorated flag at the left endpoints and (\mathcal{R}, β) the decorated flag at the right endpoints. In line with Example 4.17, we obtain

$$A_{\xi_i} = \frac{\gamma \wedge \beta_{n-i}}{\beta_{n-i+1}} \cdot \frac{\alpha_i}{\gamma \wedge \alpha_{i-1}} \quad \text{for } i \in [1, k],$$

where γ is a nonzero vector in the line $\mathcal{L}_i \cap \mathcal{R}_{n-i+1}$. Note that the formula for A_{ξ_i} does not actually depend on γ . In fact, since $\mathfrak{w}_{0,n}$ is a reduced word, we have $\alpha_{i-1} \wedge \beta_{n-i+1} \neq 0$. Thus, after wedging both $\gamma \wedge \beta_{n-i}$ and β_{n-i+1} with α_{i-1} , the expression for A_{ξ_i} above reads

$$\begin{aligned} A_{\xi_i} &= \frac{\alpha_{i-1} \wedge \gamma \wedge \beta_{n-i}}{\alpha_{i-1} \wedge \beta_{n-i+1}} \cdot \frac{\alpha_i}{\gamma \wedge \alpha_{i-1}} = (-1)^{i-1} \frac{\alpha_{i-1} \wedge \gamma \wedge \beta_{n-i}}{\alpha_{i-1} \wedge \beta_{n-i+1}} \cdot \frac{\alpha_i}{\alpha_{i-1} \wedge \gamma} \\ &= (-1)^{i-1} \frac{\alpha_{i-1} \wedge \gamma \wedge \beta_{n-i}}{\alpha_{i-1} \wedge \beta_{n-i+1}} \cdot \vartheta = (-1)^{i-1} \frac{\vartheta \cdot (\alpha_{i-1} \wedge \gamma) \wedge \beta_{n-i}}{\alpha_{i-1} \wedge \beta_{n-i+1}} \\ &= (-1)^{i-1} \frac{\alpha_i \wedge \beta_{n-i}}{\alpha_{i-1} \wedge \beta_{n-i+1}} \quad \text{for } i \in [1, k], \end{aligned}$$

where we have denoted by $\vartheta \in \mathbb{C}^\times$ the unique nonzero scalar such that $\alpha_i = \vartheta \cdot (\alpha_{i-1} \wedge \gamma)$. In conclusion, we obtain

$$A_f = \prod_{i=1}^k A_{\xi_i} = (-1)^{k(k-1)/2} \cdot \frac{\alpha_k \wedge \beta_{n-k}}{\beta_n}.$$

By definition, $\beta_n \neq 0$ is nonzero and therefore A_f is a regular function for each face f . □

Remark 4.19 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 subsection we study the vanishing loci of the microlocal merodromies $A_f : \mathfrak{M}(\Lambda(\mathbb{G}), T) \rightarrow \mathbb{C}$ associated to faces $f \subset \mathbb{G}$ of a GP graph \mathbb{G} . The key technical result, Proposition 4.23, 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 , ie $u \leq v$ if $\mathbb{O}_u \subset \overline{\mathbb{O}_v}$. Combinatorially, the Bruhat order can be computed through set comparison.

Definition 4.20 For two equal-sized subsets $I = \{i_1 < i_2 < \dots < i_m\}$ and $J = \{j_1 < j_2 < \dots < j_m\}$ of $\{1, \dots, n\}$, 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 $\{u(1), \dots, u(m)\} \leq \{v(1), \dots, v(m)\}$ for all $1 \leq m < n$.

By Section 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 \mathcal{F}_0 and \mathcal{F}^0 ; see Figure 37. In the points of the open torus chart $\mathfrak{M}(\mathfrak{w}, T) \subset \mathfrak{M}(\Lambda(\mathbb{G}), T)$, these two flags \mathcal{F}_0 and \mathcal{F}^0 are in w_0 -relative position, but in general the relative position between \mathcal{F}_0 and \mathcal{F}^0 at another point of $\mathfrak{M}(\Lambda(\mathbb{G}), T)$ might vary. The dependence of \mathcal{F}_0 and \mathcal{F}^0 on the point $p \in \mathfrak{M}(\Lambda(\mathbb{G}), T)$ will be denoted by $\mathcal{F}_0(p)$ and $\mathcal{F}^0(p)$.

By Proposition 4.18, the microlocal merodromy A_i associated to the i^{th} 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 $(\alpha_i \wedge \beta_{n-i})/\beta_n$, up to a multiple of units, where α and β are decorations on the pair of flags \mathcal{F}_0 and \mathcal{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 nonvanishing function. The following lemma shows that we can describe the vanishing locus of this microlocal merodromy $A_i: \mathfrak{M}(\Lambda(\mathbb{G}), T) \rightarrow \mathbb{C}$ in terms of the relative position between the two flags \mathcal{F}_0 and \mathcal{F}^0 :

Lemma 4.21 Let \mathbb{G} be a GP graph, $(\mathcal{F}_0, \alpha), (\mathcal{F}^0, \beta)$ the pair of decorated flags associated with a Type 1 column C and A_i the i^{th} microlocal merodromy associated to C for $i \in [1, n]$. Consider a point $p \in \mathfrak{M}(\Lambda(\mathbb{G}), T)$ and the permutation $w \in S_n$ such that $\mathcal{F}_0(p) \stackrel{w}{\sim} \mathcal{F}^0(p)$. 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 decorations α and β are proportional to

$$(e_{w(1)}, e_{w(1)} \wedge e_{w(2)}, \dots, e_{w(1)} \wedge e_{w(2)} \wedge \dots \wedge e_{w(n)}) \quad \text{and} \quad (e_1, e_1 \wedge e_2, \dots, e_1 \wedge e_2 \wedge \dots \wedge e_n),$$

respectively. From this we know that $A_i = 0$ if and only if

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

which is equivalent to saying that

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

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

$$(4-10) \quad \{u(1), u(2), \dots, u(l)\} \leq \{i \leq \dots\},$$

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

Finally, the core of this subsection is the following result, which states that we can use lollipop reactions (Definition 2.9) to keep track of the relative position conditions on flags and, in turn, understand vanishing conditions for microlocal merodromies associated to faces:


Proposition 4.23 *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 Section 2.6, the scanning wall in a lollipop reaction moves to the right if the lollipop is white and to the left if the lollipop is black. By Lemma 4.21, at the starting point, $A_f = 0$ implies that the flags at the two ends of the wall are at most $w_{[i,j]} = \bar{w}_{[i,i]}$ 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 as



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. $\bar{w}_{[i,j]}$) 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 4.23 for white lollipop reactions; the proof for the case of black lollipop reactions is completely symmetric. Let us start with the following lemma:

Lemma 4.24 *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 $(\mathcal{L}_0, \mathcal{L}^0)$ be the pair of flags to the left and $(\mathcal{R}_0, \mathcal{R}^0)$ be the pair of flags to the right. Suppose that $\mathcal{L}_0 \stackrel{u}{\sim} \mathcal{L}^0$ and $\mathcal{R}_0 \stackrel{v}{\sim} \mathcal{R}^0$. (Hence, $h^{-1}(\mathcal{R}_0) \stackrel{v}{\sim} h^{-1}(\mathcal{R}^0)$.) Then, in the Bruhat order,

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

Proof By construction, $h^{-1}(\mathcal{R}_0)$ and $h^{-1}(\mathcal{R}^0)$ share the same 1-dimensional subspace. Therefore, $v(1) = 1$, which implies that $vs_1s_2 \dots s_{n-k}$ is reduced. Lemma 4.22(1) implies that

$$\mathcal{L}_0 \stackrel{s_{k-1} \dots s_1}{\sim} h^{-1}(\mathcal{R}_0) \stackrel{vs_1 \dots s_{n-k}}{\sim} \mathcal{L}^0,$$

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 $(w(1), w(2), \dots, w(n))$. Since $v(1) = 1$, we may assume that $v = (1, v(2), v(3), \dots, v(n))$. Then

$$vs_1s_2 \dots s_{n-k} = (v(2), v(3), \dots, v(n-k+1), 1, v(n-k+2), \dots, v(n)).$$

Note that left multiplication by s_i interchanges the entries i and $i + 1$. From (the proof of) Lemma 4.22(2), we know that, when multiplying s_i on the left of a permutation w , there is only one possible relative position s_iw if i is on the left of $i + 1$, and there can be two possible relative positions w and s_iw if i is on the right of $i + 1$, in which case $s_iw < w$. Thus, performing all the left multiplications s_1, \dots, s_{k-1} in turn on $vs_1 \dots s_{n-k}$ yields the smallest relative position relation, and hence $u \geq s_{k-1} \dots s_1vs_1 \dots s_{n-k}$, as claimed. \square

Proof of Proposition 4.23 It suffices to argue in 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 remains at most $w_{[i,j]}$.

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

Let us start with the hardest case, namely a Type 3 column with a black lollipop at the k^{th} gap with $i < k \leq j$, as depicted in Figure 45, 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 $[2, n] = \{2, 3, \dots, n\}$. We want to prove that $v \leq w_{[i,j-1]} \in S_{[2,n]}$. By shifting the indices of the Coxeter generators in (4-10), $v \leq w_{[i,j-1]} \in S_{[2,n]}$ if and only if, for all $n - j \leq l \leq n - i - 1$,

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

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

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

where $a > i + 1$ is the smallest element in the set on the right. To deduce a contradiction, it suffices to prove that $\tilde{v} := s_{k-1} \dots s_1vs_1 \dots s_{n-k} \not\leq w_{[i,j]} \in S_{[1,n]}$, since Lemma 4.24 states that \tilde{v} is the smallest possible relative position between \mathcal{L}_0 and \mathcal{L}^0 . Direct computation yields that

$$\tilde{v} = (v(2)', v(3)', \dots, v(n-k+1)', k, v(n-k+2)', \dots, v(n)'),$$




Figure 45: Three of the cases in the proof of Proposition 4.23.

where

$$m' := \begin{cases} m - 1 & \text{if } m \leq k, \\ m & \text{if } m > k. \end{cases}$$

If $l \leq n - k$, then

$$\{\tilde{v}(1), \tilde{v}(2), \dots, \tilde{v}(l)\} = \{v(2)', v(3)', \dots, v(l+1)'\} = \{a', a' + \dots, \dots\} \not\leq \{i \leq \dots\},$$

where the last $\not\leq$ relation is because $a' \geq a - 1 > i$. This shows that $\tilde{v} \not\leq w_{[i,j]}$.

Otherwise, if $l > n - k$, then

$$\begin{aligned} \{\tilde{v}(1), \tilde{v}(2), \dots, \tilde{v}(l+1)\} &= \{v(2)', v(3)', \dots, v(n-k+1)', k, v(n-k+2)', v(l+1)'\} \\ &= \{k, a', a' + \dots, \dots\} \not\leq \{i \leq \dots\}, \end{aligned}$$

where the last $\not\leq$ relation is because both a' and k are greater than i .

There are three more special cases to consider for Type 3 columns with black lollipops, two of them depicted in Figure 45, center and right, as well as the case where neither $k - 1$ nor k is contained in $[i, j]$.

For the case of Figure 45, center, we want to prove that, if $u \leq w_{[i,k-1]} \in S_{[1,n]}$, then $v \leq w_{[i,k-1]} \in S_{[2,n]}$. We proceed by contradiction again. Suppose $v \not\leq w_{[i,k-1]}$; then there must exist some l with $n - k \leq l \leq n - i - 1$ such that

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

where $a > i + 1$ is the smallest element on the right. By the same argument, we see that

$$\begin{aligned} \{\tilde{v}(1), \tilde{v}(2), \dots, \tilde{v}(l+1)\} &= \{v(2)', v(3)', \dots, v(n-k+1)', k, v(n-k+2)', v(l+1)'\} \\ &= \{k, a', a' + \dots, \dots\} \not\leq \{i \leq \dots\}. \end{aligned}$$

This shows that $\tilde{v} \not\leq w_{[i,k-1]} \in S_{[1,n]}$.

For the case of Figure 45, right, we want to prove that, if $u \leq w_{[k,j]} \in S_{[1,n]}$, then $v \leq w_{[k-1,j-1]} \in S_{[2,n]}$; a proof by contradiction works again, as follows. Suppose $v \not\leq w_{[k-1,j-1]}$; then there must exist some l with $n - j \leq l \leq n - k$ such that

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


where $a > k$ is the smallest element on the right. Since $a > k$, we know that $a' = a$, and hence

$$\{\tilde{v}(1), \tilde{v}(2), \dots, \tilde{v}(l)\} = \{a', a' + \dots, \dots\} = \{a, a + \dots, \dots\} \not\subseteq \{k \leq \dots\}.$$

This shows that $\tilde{v} \not\leq w_{[k,j]}$.

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 k^{th} gap, counting from below in the GP graph, as depicted below:



By the Type 3 column requirement, we know that

$$\mathcal{R}_0 \stackrel{s_k s_{k+1} \dots s_n}{\sim} h(\mathcal{L}_0) \stackrel{u}{\sim} h(\mathcal{L}^0) \stackrel{s_n s_{n-1} \dots s_{n-k}}{\sim} \mathcal{R}^0,$$

where, by assumption, $u \leq w_{[i,j]}$. By Lemma 4.22(2) and a direct computation, we conclude that $\mathcal{R}_0 \stackrel{v}{\sim} \mathcal{R}^0$ for

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

Lastly, we consider Type 2 columns. By using Lemma 4.22, we compute that, unless we are in one of the two wall shrinking situations



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

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 Section 4.2, the factoriality of the coordinate ring $\mathbb{O}(\mathfrak{M}(\Lambda(\mathbb{G}), T))$ is a requirement. Following the results from Section 4.7, we add an additional hypothesis, as follows:

Definition 4.25 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 $\mathbb{C}(\mathfrak{M})$ 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 direct computation (eg using Gröbner bases), whereas the factoriality condition is, to our knowledge, more subtle. That said, as explained in Section 4.2, we have developed combinatorial criteria to ensure the first condition, eg 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 4.26 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 4.4 and 4.23, \mathbb{S} -complete GP graphs satisfy the second condition in Definition 4.26. By Proposition 2.11, shuffle graphs are \mathbb{S} -complete, and therefore complete. The schematic of implications is: shuffle graphs $\mathbb{G} \implies (\beta(\mathbb{G}) \Delta\text{-complete}) + (\mathbb{S}\text{-complete } \mathbb{G}) \implies \text{complete } \mathbb{G}$.

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 (eg all 3-stranded ones), and more.

4.9 Proof of the main theorem

In this subsection, 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 Sections 3 and 4.3, with the candidate cluster \mathcal{A} -variables being the microlocal merodromies (constructed in Section 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 $\{\eta_1, \dots, \eta_s\}$ for the pair $(L(\mathbb{G}), T)$ associated to $\mathfrak{w} = \mathfrak{w}(\mathbb{G})$, we have an isomorphism $\mathbb{C}[A_{\eta_1}^{\pm 1}, \dots, A_{\eta_s}^{\pm 1}] \cong \mathbb{C}(\mathfrak{M}(\mathfrak{w}, T))$, and thus $\{A_{\eta_1}, \dots, A_{\eta_s}\}$ 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 $\{\eta_1, \dots, \eta_s\}$.

We shall now show that the algebra $\mathbb{C}(\mathfrak{M}(\Lambda, T))$ 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 cover the moduli space $\mathfrak{M}(\Lambda, T)$ up to codimension 2, ie

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




Figure 46: The local models for erasing Y-trees.

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 [Berenstein et al. 2005, 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 nonfree weaves that are obtained by erasing Y-cycles 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 with 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 46, 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 46, 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 46, left.

Alternatively, it is possible to first shorten the Y-tree to a short l-cycle, and erase the l-cycle, which only requires applying Figure 46, left, twice. The following lemma verifies that this resulting weave is equivalent to the weave obtained by erasing the Y-tree directly:

Lemma 4.27 (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 l-cycle (see Proposition 3.5), which we still denote by $\gamma \subset \mathfrak{w}'$. 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 l-cycle γ in \mathfrak{w}' . Erasing γ in \mathfrak{w}' leaves two weave lines with a Reeb chord in the middle; see Figure 47, left. Follow the rest of the double tracks and contract these two weave




Figure 47: Three types of removals of (pieces of) Y-tree cycles.

lines inductively. 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 47, center. At the part where the double tracks branch off, the contracting weave line can be pulled through it by using weave equivalences, thus becoming two contracting weave lines, as illustrated Figure 47, right. In the end, we recover the weave $\mathfrak{w}_{\hat{\gamma}}$, as required. \square

Let us now study the vanishing loci of microlocal merodromies A_f associated to faces $f \subset \mathbb{G}$ by using Lemma 4.27. 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 $\text{Loc}_1(L)$. For each relative 1-cycle $\eta \in H_1(L \setminus T, \Lambda \setminus T)$, its microlocal merodromy A_η (Definition 4.10) is well defined and, in particular, we can associate a naive microlocal merodromy A_f with each face f of \mathbb{G} . By Corollary 4.12 and Proposition 4.18, these A_f 's are regular functions on $\mathfrak{M}(\Lambda, T)$ and they are unique up to multiples of units.

Proposition 4.28 *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 nonempty.*

Proof By assumption, f admits a sugar-free hull \mathbb{S}_f , which, by Lemma 3.29, 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 l-cycle, and place this short l-cycle near the end of the original Y-tree, so that it lies inside some Type 1 weave column. If we delete this short l-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 l-cycle. We claim that its associated stratum $\mathcal{M}_1(\mathfrak{w}')$ in $\mathfrak{M}(\Lambda(\mathbb{G}), T)$ is nonempty.

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 [Jin and Treumann 2017], or [Ekholm et al. 2016; Ng et al. 2020], the two strata $\mathcal{M}_1(\mathfrak{w}_1)$ and $\mathcal{M}_1(\mathfrak{w}_2)$ are nonempty. 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 (nonunique) 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 nonempty.

Since $\mathfrak{M}(\mathfrak{w}', T)$ fibers over $\mathcal{M}_1(\mathfrak{w}')$, it follows that $\mathfrak{M}(\mathfrak{w}', T)$ is nonempty as well. Let p be a point in $\mathfrak{M}(\mathfrak{w}', T)$. By Lemma 4.27, 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 i^{th} 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 4.21. \square

We also establish the converse of Proposition 4.28, ie if the lollipop chain reaction initiated from f is incomplete, then the microlocal merodromy A_f must be a unit in $\mathbb{C}(\mathfrak{M}(\Lambda, T))$.

Proposition 4.29 *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 4.28 implies that $\{A_f = 0\}$ is nonempty and hence A_f cannot be 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 4.23, $A_g(p) = 0$ as well. But this is impossible because any relative 1-chain η_g associated with g must map to the identity under the projection map $H_1(L \setminus T, \Lambda \setminus T) \rightarrow H_1(L, \Lambda)$, and hence A_g is a unit by Corollary 4.12. Therefore, $\{A_f = 0\}$ is empty and A_f is a unit. \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 there can be multiple faces sharing the same sugar-free hull, some faces (naive absolute cycles) are set aside during the replacement process. 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 4.30 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 4.23, $A_f(p) = 0$ if and only if $A_g(p) = 0$ for any $p \in \mathfrak{M}(\Lambda, T)$. Since $\mathbb{C}(\mathfrak{M}(\Lambda, T))$ 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. \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 by $\{A_f\}$, where f runs through the faces of \mathbb{G} , and the microlocal monodromies associated with marked points by $\{A_t\}_{t \in T}$. 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,

$$\mathbb{C}(U_0) = \mathbb{C}[A_{f_1}^{\pm 1}, \dots, A_{f_{b_1(L)}}^{\pm 1}, A_t^{\pm 1}] = \mathbb{C}[A_1^{\pm 1}, \dots, A_{b_1(L)}^{\pm 1}, A_t^{\pm 1}],$$

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 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 :

$$\begin{array}{c} \text{---} \star \text{---} \\ \text{---} \star \text{---} \end{array} = \frac{1}{2} \begin{array}{c} \text{---} \star \text{---} \\ \text{---} \star \text{---} \end{array}$$

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

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

Proposition 4.31 Consider the forgetful map $p: \mathfrak{M}(\mathfrak{w}, T) \rightarrow \mathcal{M}_1(\mathfrak{w})$. For an initial absolute cycle $\gamma_i \in \mathfrak{G}(\mathbb{G})$,

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




Figure 48: Mutation of initial relative cycles.

Proof Let us denote the relative 1-chain dual to $\gamma_i \in \mathfrak{B}$ 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 $\theta = \sum_{k \in \mathfrak{B}} c_k \gamma_k \in H_1(L, T)$ that is a lift of a relative 1-cycle in $H_1(L, \Lambda)$,

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

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

Corollary 4.32 *The rectangular exchange matrix $\epsilon|_{\mathfrak{S}(\mathbb{G}) \times \mathfrak{B}}$ is full-ranked.*

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

Let us now focus 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 3.29, we represent 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 48, 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.

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 48. 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 48, we have also drawn a relative 1-chain η' , which connects to the rest of $-\eta_k$ outside of the local picture. But 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 nontrivial 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

$$(4-11) \quad \eta'_k = \eta' + \sum_{j \in \tilde{\mathfrak{B}}} [-\epsilon_{kj}]_+ \eta_j.$$

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 k^{th} 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 4.33 *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 $p: U_0 \rightarrow \mathcal{M}_1(\mathfrak{w})$ 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 1-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 nonzero vector in each line l_i and let \det denote the dual of the nonzero 2-form associated with \mathbb{C}^2 . Following the crossing-value formula, we obtain

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

Therefore, the product can be computed as

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

Remark 4.34 Instead of the relative 1-chain η' depicted in Figure 48, right, we could have chosen η' to have support the other zigzag  with appropriate intersection numbers. In this other choice, equation (4-11) would need to be modified to $\eta'_k = \eta' + \sum_{j \in \tilde{\mathfrak{B}}} [\epsilon_{kj}]_+ \eta_j$ and the equation in Proposition 4.33 would be modified to $A'_k = A_k^{-1} (1 + p^*(X_k^{-1})) \prod_{j \in \tilde{\mathfrak{B}}} A_j^{[\epsilon_{kj}]_+}$. These compatible changes of signs $\epsilon_{kj} \rightarrow -\epsilon_{kj}$ define the chiral dual cluster structure on the same variety, which is the cluster structure defined by the opposite quiver. This chiral dual is discussed in [Fock and Goncharov 2009, Section 1.2]. Since the quiver is given by the intersection pairing between absolute cycles, the choice in Figure 48 is, in a sense, naturally dictated by the chosen orientation on the filling $L(\mathfrak{w})$.

Propositions 4.33 and 4.31 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 4.35 *Let \mathbb{G} be a GP graph, $\{\eta_1, \dots, \eta_s\}$ the set of naive relative cycles and $\{A_i\}$ the associated set of naive microlocal merodromies. Consider 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 4.18, 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 $i \in \mathfrak{B}$, the 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 4.36 *Let $U_0 \subset \mathfrak{M}(\Lambda, T)$ be the initial open toric chart and f a unit in $\mathbb{C}(U_0)$ (resp. $\mathbb{C}(U_k)$). Suppose that $f = gh$ in $\mathbb{C}(\mathfrak{M}(\Lambda, T))$ for some $g, h \in \mathbb{C}(\mathfrak{M}(\Lambda, T))$. Then g and h are also units in $\mathbb{C}(U_0)$ (resp. $\mathbb{C}(U_k)$).*

Proof Indeed, if $f = gh$ in $\mathbb{C}(\mathfrak{M}(\Lambda, T))$, then $f = gh$ in $\mathbb{C}(U_0)$ as well, and, if f is a Laurent monomial in $\mathbb{C}(U_0)$, then each of g and h must be a Laurent monomial, too. The proof for the case where f is a unit in $\mathbb{C}(U_k)$ is analogous. □

We first show that the initial merodromies are irreducible assuming they are regular functions:

Lemma 4.37 *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, ie an element of $\mathbb{C}(\mathfrak{M}(\Lambda, T))$. Then A_k is irreducible.*

Proof Suppose $A_k = gh$ in $\mathbb{C}(\mathfrak{M}(\Lambda, T))$ with neither g nor h being a unit. Then Lemma 4.36 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 $\mathbb{C}(\mathfrak{M}(\Lambda, T))$. 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 4.29, 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 4.36, h is also a Laurent monomial in the new chart $\mathbb{C}(U_j)$ associated to \mathfrak{w}_j and hence we can write g and h as

$$g = A_j^{p_j} \prod_{i \neq j} A_i^{p_i} \quad \text{and} \quad h = A_j^{q_j} \prod_{i \neq j} A_i^{q_i}.$$

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 nonzero, then one of them must be positive; suppose $p_j > 0$. By Corollary 4.35, $A'_j = M_1 + M_2$, where M_1 and M_2 are two algebraically independent Laurent monomials in $\{A_i\}_{i \in \tilde{\mathfrak{G}}}$, 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 $\{A_i\}_{i \in \tilde{\mathfrak{G}}}$. This is again a contradiction, and therefore A_k must be an irreducible element in $\mathbb{C}(\mathfrak{M}(\Lambda, T))$. □

We are ready to conclude regularity, and thus irreducibility, of initial merodromies:

Proposition 4.38 *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 $\mathbb{C}(\mathfrak{M}(\Lambda, T))$.*

Proof By Lemma 4.37, 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 4.18 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 4.37, A_i are irreducible elements in $\mathbb{C}(\mathfrak{M}(\Lambda, T))$ as well. Let f_i be the face selected for each vertex i of \mathcal{H} . Then, for each vertex i of \mathcal{H} ,

$$A_{f_i} = \prod_{j \geq i} A_j.$$

In particular, if $i > k$, then the above is the unique factorization of the naive microlocal merodromies A_{f_i} in $\mathbb{C}(\mathfrak{M}(\Lambda, T))$, up to multiple 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 4.12.

In addition to the above, if $i > k$, then f_i is contained in the sugar-free hull $S_{f_i} = S_i$. Proposition 4.23 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 $\mathbb{C}(\mathfrak{M}(\Lambda, T))$, 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. □

4.9.4 Conclusion of the argument We finalize the proof of the covering of $\mathbb{C}(\mathfrak{M}(\Lambda, T))$ 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 associated microlocal merodromy by $D_k := \{A_k = 0\} \subset \mathfrak{M}(\Lambda, T)$. Since A_k is an irreducible element of $\mathbb{C}(\mathfrak{M}(\Lambda, T))$, D_k is irreducible as a codimension 1 subvariety in $\mathfrak{M}(\Lambda, T)$.

Proposition 4.39 *Let \mathbb{G} be a GP graph, $\gamma_k \in \mathfrak{S}(\mathbb{G})$ the initial absolute cycles, $U_k \subset \mathfrak{M}(\Lambda, T)$ 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 nonempty open subset of the vanishing locus D_k .*

Proof It suffices to prove that $U_k \cap D_k$ is nonempty. Similar to the proof of Proposition 4.28, 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 is clear: For instance, in the case where the weave looks like the one on the left in



we first fix a point in $\mathcal{M}_1(\mathfrak{w}_1)$ and then, based on the flags L_0, L_1, L_2 and L_3 , we choose a point in $\mathcal{M}_1(\mathfrak{w}_2)$ with flags $\mathcal{R}_0 = L_0, \mathcal{R}_1 = L_1$ and $\mathcal{R}_2 = L_2$, but $\mathcal{R}_3 \neq L_3$, 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. □

At this stage, the covering property, up to codimension 2, readily follows:

Proposition 4.40 *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*

$$\text{codim}\left(U_0 \cup \bigcup_{k \in \mathfrak{S}(\mathbb{G})} U_k\right) \geq 2,$$

ie the union of U_0 and all the adjacent charts U_k covers $\mathfrak{M}(\Lambda, T)$ up to codimension 2.

Proof By Proposition 4.39, the intersection $U_j \cap D_j$ is open in D_j for all j , and thus $\text{codim}(D_j \cap U_j^c) \geq 2$ for each j . Thus, $\text{codim}(U_0 \cup \bigcup_k U_k) \geq 2$ by the inclusions

$$\left(U_0 \cup \bigcup_k U_k\right)^c = U_0^c \cap \bigcap_k U_k^c = \left(\bigcup_j D_j\right) \cap \bigcap_k U_k^c = \bigcup_j \left(D_j \cap \bigcap_k U_k^c\right) \subset \bigcup_j (D_j \cap U_j^c). \quad \square$$

Theorem 1.1 and Corollary 1.2 are now concluded as follows:

Theorem 4.41 *Let \mathbb{G} be a complete GP graph. The coordinate ring of regular functions $\mathbb{O}(\mathfrak{M}(\Lambda(\mathbb{G}), T))$ has the structure of an upper cluster algebra.*

Proof Consider the open subset $U_0 \cup \bigcup_{k \in \mathbb{S}(\mathbb{G})} U_k \subset \mathfrak{M}(\Lambda, T)$. Proposition 4.40 shows the equality of coordinate rings $\mathbb{O}(\mathfrak{M}(\Lambda, T)) = \mathbb{O}(U_0 \cup \bigcup_{k \in \mathbb{S}(\mathbb{G})} U_k)$. Corollary 4.35 implies that $\mathbb{O}(U_0 \cup \bigcup_{k \in \mathbb{S}(\mathbb{G})} U_k)$ is an upper bound of a cluster algebra. In addition, since T has at least one marked point per link component, Corollary 4.32 shows that the rectangular exchange matrix $\epsilon|_{\mathbb{S}(\mathbb{G}) \times \tilde{\mathfrak{B}}}$ is full-ranked. Then [Berenstein et al. 2005, Corollary 1.9] implies that this upper bound coincides with its upper cluster algebra and therefore we conclude that $\mathbb{O}(\mathfrak{M}(\Lambda, T))$ is an upper cluster algebra. \square

Corollary 4.42 *Let \mathbb{G} be a complete GP graph. Then $\mathbb{O}(\mathcal{M}_1(\Lambda(\mathbb{G})))$ 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 $p: \mathfrak{M}(\Lambda, T) \rightarrow \mathcal{M}_1(\Lambda)$, the cluster-theoretical map $p: \mathcal{A} \rightarrow \mathcal{X}$ is also surjective. Both $\mathbb{O}(\mathcal{A})$ and $\mathbb{O}(\mathcal{X})$ 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 [Shen and Weng 2020, Lemma A.1]. That said, given that $p: \mathfrak{M}(\Lambda, T) \rightarrow \mathcal{M}_1(\Lambda)$ 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 commutative diagram

$$\begin{array}{ccc} \mathfrak{M}(\Lambda, T) & \xrightarrow[\cong]{\alpha} & \mathcal{A} \\ p \downarrow & & \downarrow p \\ \mathcal{M}_1(\Lambda) & \xrightarrow{\chi} & \mathcal{X} \end{array}$$

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 4.41 implies that the top map induces an isomorphism between $\mathbb{O}(\mathfrak{M}(\Lambda, T))$ and $\mathbb{O}(\mathcal{A})$. Now, given a regular function $f \in \mathbb{O}(\mathcal{X})$, the pullback $\chi^*(f)$ 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 $f \in \mathbb{O}(\mathcal{M}_1(\Lambda))$, 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 $\chi^*: \mathbb{O}(\mathcal{X}) \rightarrow \mathbb{O}(\mathcal{M}_1(\Lambda))$, and hence $\mathbb{O}(\mathcal{M}_1(\Lambda))$ is a cluster Poisson algebra. \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 [Kontsevich and Soibelman 2010; Keller 2017; Goncharov and Shen 2018]. In this section we prove Corollary 1.3,


ie 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, ξ_{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 2.8. □

Corollary 5.2 *Let \mathbb{G} be a shuffle graph. Then all its sugar-free hulls must be one of the 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 $Q(\mathbb{G})$, ie $\langle \partial\mathbb{S}_f, \partial\mathbb{S}_g \rangle = 0$.*

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 endpoints, 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 endpoints 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 endpoints of $(\partial\mathbb{S}_f) \setminus (\partial\mathbb{S}_g)$ are the same as the two endpoints 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. □

By Definition 2.8, 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 vertex set of the initial quiver $Q(\mathbb{G})$ as follows:




Figure 49: Left: the relative lengths of horizontal lines in a shuffle graph \mathbb{G} associated with σ . Right: 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.

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 i^{th} and the j^{th} 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 Figures 49, right, and 50. 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 50 illustrates this for a shuffle graph associated with the permutation $\sigma = [4\ 1\ 2\ 3]$, with two branchings on each side.

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 [Shen and Weng 2021].

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




Figure 50: 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.




Figure 51: 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.

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.

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 i^{th} horizontal line in a grid plabic graph \mathbb{G} . Take the first vertical edge e' with a black vertex on the i^{th} horizontal line as we search from right to left starting from b . Suppose the other vertex of e' lies on the j^{th} horizontal line whose right endpoint lies to the right of b , and suppose there are no more vertical edges (of either pattern) between the i^{th} and j^{th} horizontal lines to the right of e' . 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 i^{th} and j^{th} horizontal lines on the right side of e' , resulting in a possibly nonplanar bicolor graph. Figure 52 gives 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 .

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

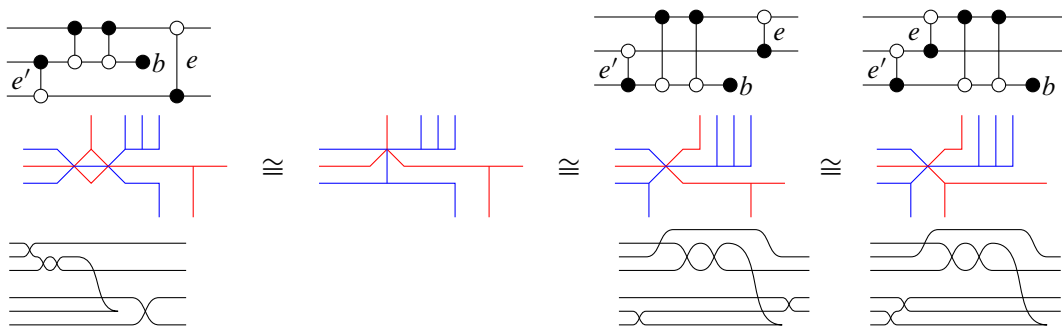


Figure 52: Example of a reflection move near a black lollipop.

obstructing lollipop. In Figure 52, 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 [Shen and Weng 2021].

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 52, then this reflection move recovers the right reflection move in [Shen and Weng 2021]. (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 $K^{1/2}: \mathcal{M}_1(\Lambda) \rightarrow \mathcal{M}_1(\Lambda')$, 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 reverses 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 $t: \mathcal{M}_1(\Lambda') \rightarrow \mathcal{M}_1(\Lambda)$ in [Shen and Weng 2021]. 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 define $DT := t \circ K^{1/2} = K^{1/2} \circ t$ 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 results 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

$$\tilde{\epsilon} = \begin{pmatrix} \epsilon & \text{id} \\ -\text{id} & 0 \end{pmatrix},$$

where ϵ is the exchange matrix of Q .




Figure 53: Left: the square move in a plabic fence. Center and right: sliding vertical edges of opposite patterns on different levels through each other.

For any mutation sequence μ_i on the quiver Q , we can apply the same mutation sequence μ_i to \tilde{Q} and get a new quiver $\tilde{Q}' := \mu_i(Q)$. Note that the unfrozen part of \tilde{Q}' is identical to $Q' := \mu_i(Q)$. A remarkable property of \tilde{Q}' is that, for any unfrozen vertex i , $\tilde{c}'_{ij'}$ is either nonnegative or nonpositive for all framing frozen vertices j' ; this is known as the *sign coherence* phenomenon of c -vectors in cluster theory [Derksen et al. 2010; Gross et al. 2018]. We say a vertex i in Q' is *green* if $\tilde{c}'_{ij'}$ is nonnegative for all framing frozen vertices j' and a vertex i in Q' is *red* if $\tilde{c}'_{ij'}$ is nonpositive for all framing frozen vertices j' . For a given initial quiver Q (all of whose vertices are green), if a mutation sequence μ_i turns every quiver vertex red, then we say μ_i is a *reddening sequence*. If additionally a reddening sequence μ_i only mutates at green vertices, then we say that μ_i is a *maximal green sequence*. For any fixed initial seed, the cluster Donaldson–Thomas transformation can be captured combinatorially by a reddening sequence [Keller 2017; Goncharov and Shen 2018]. 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 53.

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 l -cycle, which can be locally described by the movie in Figure 54.

A maximal green sequence on $Q(\mathbb{G})$ can be constructed recursively as follows:

- Take the rightmost vertical edge e of the \circlearrowleft pattern and change it to the opposite pattern.
- Move this newly changed vertical edge e to the left, passing all remaining \circlearrowleft edges.




Figure 54: Square move in terms of Legendrian weaves: the first and last moves are weave equivalences; the middle move is a weave mutation.


This iterative process terminates when we run out of vertical edges with a black vertex on top. Note that all mutations 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 [Shen and Weng 2021, 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 leftmost 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 2.8, if a vertical edge can be placed between the i^{th} and j^{th} 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 i^{th} or $(i + 1)^{\text{st}}$ 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 i^{th} and the j^{th} horizontal lines, and through all  edges incident to the i^{th} or j^{th} 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 $j - i - 1$ black lollipops. Each time we encounter a black lollipop, we will try to apply a move similar to Figure 52. 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 52; see Figure 55.

Note that all moves in Figure 55 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




Figure 55: Two reflection moves to cover the vertical edge e .

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 nonplanarity caused by the earlier reflection move (Figure 52, right) will be canceled, 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 52. 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 56.

(III) If e lies in the left region between the i^{th} and the j^{th} horizontal lines with $|i - j| > 1$:

(III.1) Move e all the way to the right region between the i^{th} and the j^{th} horizontal lines so that it becomes the rightmost vertical edge in the plabic graph.

(III.2) Apply a reflection move to change the pattern of e to \circlearrowleft .

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 \circlearrowleft pattern already. Thus, moving e , which is of the \circlearrowleft pattern, to the right through \circlearrowleft edges, is completely mirror to step (II.1) before. Figure 57 depicts the reflection moves we need to perform during this process.

We can now conclude the following result:

Theorem 5.8 Let \mathbb{G} be a shuffle graph. Then $DT = t \circ K^{1/2}$ 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 50 and 49); we claim that, after each iterative step (I) and step (II) of moving an edge e on level (i, j) , the leftmost green vertex of level (i, j)




Figure 56: Example of a special case where the edge e' is absent.




Figure 57: 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.

turns red, and after each iterative step (III) of moving an edge e on level (i, j) , the rightmost 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 \circ 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 k^{th} horizontal line (with $i < k < j$). If we are in the situation of Remark 5.7, then we can directly jump through the whole k^{th} 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 52 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 nonrectangular before the moves in Figure 52 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 reaches the left white lollipop of the k^{th} 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 55, 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;¹⁵ 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.


¹⁵Note 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.

The case (III) is essentially (II) in reverse. Suppose we are moving a \bullet edge e on level (i, j) , and suppose the first white lollipop it encounters is on the k^{th} horizontal line. Let v be the quiver vertex to the right of e , which is the rightmost green on level (i, j) at the beginning of the iterative step. Under the moves in the top row of Figure 57, we move v into level (k, j) , and obtain another vertical edge e' of the \bullet 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 green¹⁶ 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 it exists) turns from red to green. As a result, when e and e' get to the right end of the k^{th} 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 57 (which do not involve mutations), we need to send e' back to where it was, and that will make the quiver vertices on the level of e' undergo the reverse sequence of color changes 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 \bullet to \circ , all quiver vertices would turn red. Therefore, $K^{1/2}$ is indeed a reddening sequence. □

The statement in Corollary 1.3 about the cluster duality conjecture now follows from [Gross et al. 2018], as our quivers are full-ranked and a DT transformation exists. Finally, we remark that the same argument used in [Gao et al. 2020b] 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, as will be its square, the Legendrian Kálmán loop.

Example 5.9 Consider the following shuffle graph \mathbb{G} , with its quiver depicted and $\Lambda(\mathbb{G})$ at its right:



This is mutation equivalent to an acyclic quiver of infinite type, eg 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

¹⁶These mutations are also not green.

smooth knot type 10_{161} is not a rainbow closure of a positive braid. The reddening sequence realizing DT is

$$(\mu_2 \circ \mu_5 \circ \mu_4 \circ \mu_3 \circ \mu_5 \circ \mu_2) \circ () \circ () \circ (\mu_4) \circ (\mu_5) \circ (\mu_2 \circ \mu_5) \circ (\mu_{4+6} \circ \mu_4) \\ \circ (\mu_5 \circ \mu_2 \circ \mu_{1+2} \circ \mu_3 \circ \mu_{4+6} \circ \mu_2 \circ \mu_5),$$

where we have grouped the mutations of each iterative step inside a pair of parentheses, and empty parentheses mean no mutations at that step. The first mutation is μ_5 , then μ_2 , then μ_{4+6} , and so on until the last three mutations are μ_4 , μ_5 and lastly μ_2 .

References

- [Allegretti 2021] **D G L Allegretti**, *Stability conditions, cluster varieties, and Riemann–Hilbert problems from surfaces*, Adv. Math. 380 (2021) art. id. 107610 MR Zbl
- [Auroux 2007] **D Auroux**, *Mirror symmetry and T–duality in the complement of an anticanonical divisor*, J. Gökova Geom. Topol. 1 (2007) 51–91 MR Zbl
- [Auroux 2009] **D Auroux**, *Special Lagrangian fibrations, wall-crossing, and mirror symmetry*, from “Geometry, analysis, and algebraic geometry: forty years of the Journal of Differential Geometry” (H-D Cao, S-T Yau, editors), Surv. Differ. Geom. 13, International, Somerville, MA (2009) 1–47 MR Zbl
- [Berenstein et al. 2005] **A Berenstein, S Fomin, A Zelevinsky**, *Cluster algebras, III: Upper bounds and double Bruhat cells*, Duke Math. J. 126 (2005) 1–52 MR Zbl
- [Björner and Brenti 2005] **A Björner, F Brenti**, *Combinatorics of Coxeter groups*, Graduate Texts in Math. 231, Springer (2005) MR Zbl
- [Boalch 2014a] **P P Boalch**, *Geometry and braiding of Stokes data; fission and wild character varieties*, Ann. of Math. 179 (2014) 301–365 MR Zbl
- [Boalch 2014b] **P Boalch**, *Poisson varieties from Riemann surfaces*, Indag. Math. 25 (2014) 872–900 MR Zbl
- [Casals 2022] **R Casals**, *Lagrangian skeleta and plane curve singularities*, J. Fixed Point Theory Appl. 24 (2022) art. id. 34 MR Zbl
- [Casals and Gao 2022] **R Casals, H Gao**, *Infinitely many Lagrangian fillings*, Ann. of Math. 195 (2022) 207–249 MR Zbl
- [Casals and Ng 2022] **R Casals, L Ng**, *Braid loops with infinite monodromy on the Legendrian contact DGA*, J. Topol. 15 (2022) 1927–2016 MR Zbl
- [Casals and Zaslow 2022] **R Casals, E Zaslow**, *Legendrian weaves: N–graph calculus, flag moduli and applications*, Geom. Topol. 26 (2022) 3589–3745 MR Zbl
- [Casals et al. 2020] **R Casals, E Gorsky, M Gorsky, J Simental**, *Algebraic weaves and braid varieties*, preprint (2020) arXiv 2012.06931
- [Casals et al. 2021] **R Casals, E Gorsky, M Gorsky, J Simental**, *Positroid links and braid varieties*, preprint (2021) arXiv 2105.13948
- [Casals et al. 2022] **R Casals, E Gorsky, M Gorsky, I Le, L Shen, J Simental**, *Cluster structures on braid varieties*, preprint (2022) arXiv 2207.11607

- [Derksen et al. 2010] **H Derksen, J Weyman, A Zelevinsky**, *Quivers with potentials and their representations, II: Applications to cluster algebras*, J. Amer. Math. Soc. 23 (2010) 749–790 MR Zbl
- [Ekholm et al. 2016] **T Ekholm, K Honda, T Kálmán**, *Legendrian knots and exact Lagrangian cobordisms*, J. Eur. Math. Soc. 18 (2016) 2627–2689 MR Zbl
- [Etnyre et al. 2013] **J B Etnyre, L L Ng, V Vértesi**, *Legendrian and transverse twist knots*, J. Eur. Math. Soc. 15 (2013) 969–995 MR Zbl
- [Fock and Goncharov 2006a] **V V Fock, A B Goncharov**, *Cluster \mathcal{X} -varieties, amalgamation, and Poisson–Lie groups*, from “Algebraic geometry and number theory” (V Ginzburg, editor), Progr. Math. 253, Birkhäuser, Boston, MA (2006) 27–68 MR Zbl
- [Fock and Goncharov 2006b] **V Fock, A Goncharov**, *Moduli spaces of local systems and higher Teichmüller theory*, Publ. Math. Inst. Hautes Études Sci. 103 (2006) 1–211 MR Zbl
- [Fock and Goncharov 2009] **V V Fock, A B Goncharov**, *Cluster ensembles, quantization and the dilogarithm*, Ann. Sci. Éc. Norm. Supér. 42 (2009) 865–930 MR Zbl
- [Fomin and Zelevinsky 1999] **S Fomin, A Zelevinsky**, *Double Bruhat cells and total positivity*, J. Amer. Math. Soc. 12 (1999) 335–380 MR Zbl
- [Fomin and Zelevinsky 2002] **S Fomin, A Zelevinsky**, *Cluster algebras, I: Foundations*, J. Amer. Math. Soc. 15 (2002) 497–529 MR Zbl
- [Fomin and Zelevinsky 2003] **S Fomin, A Zelevinsky**, *Cluster algebras, II: Finite type classification*, Invent. Math. 154 (2003) 63–121 MR Zbl
- [Fomin et al. 2008] **S Fomin, M Shapiro, D Thurston**, *Cluster algebras and triangulated surfaces, I: Cluster complexes*, Acta Math. 201 (2008) 83–146 MR Zbl
- [Fomin et al. 2022] **S Fomin, P Pylyavskyy, E Shustin, D Thurston**, *Morsifications and mutations*, J. Lond. Math. Soc. 105 (2022) 2478–2554 MR Zbl
- [Gaiotto et al. 2010] **D Gaiotto, G W Moore, A Neitzke**, *Four-dimensional wall-crossing via three-dimensional field theory*, Comm. Math. Phys. 299 (2010) 163–224 MR Zbl
- [Gaiotto et al. 2013] **D Gaiotto, G W Moore, A Neitzke**, *Spectral networks*, Ann. Henri Poincaré 14 (2013) 1643–1731 MR Zbl
- [Galashin and Lam 2023] **P Galashin, T Lam**, *Positroid varieties and cluster algebras*, Ann. Sci. Éc. Norm. Supér. 56 (2023) 859–884 MR Zbl
- [Gao et al. 2020a] **H Gao, L Shen, D Weng**, *Augmentations, fillings, and clusters*, preprint (2020) arXiv 2008.10793
- [Gao et al. 2020b] **H Gao, L Shen, D Weng**, *Positive braid links with infinitely many fillings*, preprint (2020) arXiv 2009.00499
- [Gekhtman et al. 2005] **M Gekhtman, M Shapiro, A Vainshtein**, *Cluster algebras and Weil–Petersson forms*, Duke Math. J. 127 (2005) 291–311 MR Zbl
- [Gekhtman et al. 2010] **M Gekhtman, M Shapiro, A Vainshtein**, *Cluster algebras and Poisson geometry*, Mathematical Surveys and Monographs 167, Amer. Math. Soc., Providence, RI (2010) MR Zbl
- [Goncharov and Kenyon 2013] **A B Goncharov, R Kenyon**, *Dimers and cluster integrable systems*, Ann. Sci. Éc. Norm. Supér. 46 (2013) 747–813 MR Zbl

- [Goncharov and Kontsevich 2021] **A Goncharov, M Kontsevich**, *Spectral description of non-commutative local systems on surfaces and non-commutative cluster varieties*, preprint (2021) arXiv 2108.04168
- [Goncharov and Shen 2018] **A Goncharov, L Shen**, *Donaldson–Thomas transformations of moduli spaces of G -local systems*, *Adv. Math.* 327 (2018) 225–348 MR Zbl
- [Gross et al. 2015] **M Gross, P Hacking, S Keel**, *Birational geometry of cluster algebras*, *Algebr. Geom.* 2 (2015) 137–175 MR Zbl
- [Gross et al. 2018] **M Gross, P Hacking, S Keel, M Kontsevich**, *Canonical bases for cluster algebras*, *J. Amer. Math. Soc.* 31 (2018) 497–608 MR Zbl
- [Guillermou 2023] **S Guillermou**, *Sheaves and symplectic geometry of cotangent bundles*, *Astérisque* 440, Soc. Math. France, Paris (2023) MR Zbl
- [Guillermou and Schapira 2014] **S Guillermou, P Schapira**, *Microlocal theory of sheaves and Tamarkin’s non displaceability theorem*, from “Homological mirror symmetry and tropical geometry” (R Castano-Bernard, F Catanese, M Kontsevich, T Pantev, Y Soibelman, I Zharkov, editors), *Lect. Notes Unione Mat. Ital.* 15, Springer (2014) 43–85 MR Zbl
- [Guillermou et al. 2012] **S Guillermou, M Kashiwara, P Schapira**, *Sheaf quantization of Hamiltonian isotopies and applications to nondisplaceability problems*, *Duke Math. J.* 161 (2012) 201–245 MR Zbl
- [Hacking and Keel 2018] **P Hacking, S Keel**, *Mirror symmetry and cluster algebras*, from “Proceedings of the International Congress of Mathematicians, II: Invited lectures” (B Sirakov, P N de Souza, M Viana, editors), *World Sci.*, Hackensack, NJ (2018) 671–697 MR Zbl
- [Hartshorne 1977] **R Hartshorne**, *Algebraic geometry*, *Graduate Texts in Math.* 52, Springer (1977) MR Zbl
- [Henry and Rutherford 2015] **M B Henry, D Rutherford**, *Ruling polynomials and augmentations over finite fields*, *J. Topol.* 8 (2015) 1–37 MR Zbl
- [Iwaki and Nakanishi 2014] **K Iwaki, T Nakanishi**, *Exact WKB analysis and cluster algebras*, *J. Phys. A* 47 (2014) art. id. 474009 MR Zbl
- [Iwaki and Nakanishi 2016] **K Iwaki, T Nakanishi**, *Exact WKB analysis and cluster algebras, II: Simple poles, orbifold points, and generalized cluster algebras*, *Int. Math. Res. Not.* 2016 (2016) 4375–4417 MR Zbl
- [Jin and Treumann 2017] **X Jin, D Treumann**, *Brane structures in microlocal sheaf theory*, preprint (2017) arXiv 1704.04291
- [Kashiwara and Schapira 1985] **M Kashiwara, P Schapira**, *Microlocal study of sheaves*, *Astérisque* 128, Soc. Math. France, Paris (1985) MR Zbl Correction in 130 (1985) 209
- [Kashiwara and Schapira 1990] **M Kashiwara, P Schapira**, *Sheaves on manifolds*, *Grundle. Math. Wissen.* 292, Springer (1990) MR Zbl
- [Keller 2017] **B Keller**, *Quiver mutation and combinatorial DT-invariants*, preprint (2017) arXiv 1709.03143
- [Kontsevich and Soibelman 2010] **M Kontsevich, Y Soibelman**, *Motivic Donaldson–Thomas invariants: summary of results*, from “Mirror symmetry and tropical geometry” (R Castaño Bernard, Y Soibelman, I Zharkov, editors), *Contemp. Math.* 527, Amer. Math. Soc., Providence, RI (2010) 55–89 MR Zbl
- [Lam and Speyer 2022] **T Lam, D E Speyer**, *Cohomology of cluster varieties, I: Locally acyclic case*, *Algebra Number Theory* 16 (2022) 179–230 MR Zbl
- [Muller 2013] **G Muller**, *Locally acyclic cluster algebras*, *Adv. Math.* 233 (2013) 207–247 MR Zbl

- [Muller 2014] **G Muller**, $\mathcal{A} = \mathcal{U}$ for locally acyclic cluster algebras, *Symmetry Integrability Geom. Methods Appl.* 10 (2014) art. id. 094 MR Zbl
- [Nadler 2016] **D Nadler**, *Wrapped microlocal sheaves on pairs of pants*, preprint (2016) arXiv 1604.00114
- [Neitzke 2014] **A Neitzke**, *Cluster-like coordinates in supersymmetric quantum field theory*, *Proc. Natl. Acad. Sci. USA* 111 (2014) 9717–9724 MR Zbl
- [Ng et al. 2017] **L Ng, D Rutherford, V Shende, S Sivek**, *The cardinality of the augmentation category of a Legendrian link*, *Math. Res. Lett.* 24 (2017) 1845–1874 MR Zbl
- [Ng et al. 2020] **L Ng, D Rutherford, V Shende, S Sivek, E Zaslow**, *Augmentations are sheaves*, *Geom. Topol.* 24 (2020) 2149–2286 MR Zbl
- [Pascaleff and Tonkonog 2020] **J Pascaleff, D Tonkonog**, *The wall-crossing formula and Lagrangian mutations*, *Adv. Math.* 361 (2020) art. id. 106850 MR Zbl
- [Polterovich 1991] **L Polterovich**, *The surgery of Lagrange submanifolds*, *Geom. Funct. Anal.* 1 (1991) 198–210 MR Zbl
- [Postnikov 2006] **A Postnikov**, *Total positivity, grassmannians, and networks*, preprint (2006) arXiv math/0609764
- [Schnürer 2018] **O M Schnürer**, *Six operations on dg enhancements of derived categories of sheaves*, *Selecta Math.* 24 (2018) 1805–1911 MR Zbl
- [Serhiyenko et al. 2019] **K Serhiyenko, M Sherman-Bennett, L Williams**, *Cluster structures in Schubert varieties in the Grassmannian*, *Proc. Lond. Math. Soc.* 119 (2019) 1694–1744 MR Zbl
- [Shen and Weng 2020] **L Shen, D Weng**, *Cyclic sieving and cluster duality of Grassmannian*, *Symmetry Integrability Geom. Methods Appl.* 16 (2020) art. id. 067 MR Zbl
- [Shen and Weng 2021] **L Shen, D Weng**, *Cluster structures on double Bott–Samelson cells*, *Forum Math. Sigma* 9 (2021) art. id. e66 MR Zbl
- [Shende et al. 2016] **V Shende, D Treumann, H Williams**, *On the combinatorics of exact Lagrangian surfaces*, preprint (2016) arXiv 1603.07449
- [Shende et al. 2017] **V Shende, D Treumann, E Zaslow**, *Legendrian knots and constructible sheaves*, *Invent. Math.* 207 (2017) 1031–1133 MR Zbl
- [Shende et al. 2019] **V Shende, D Treumann, H Williams, E Zaslow**, *Cluster varieties from Legendrian knots*, *Duke Math. J.* 168 (2019) 2801–2871 MR Zbl
- [Toën and Vaquié 2007] **B Toën, M Vaquié**, *Moduli of objects in dg-categories*, *Ann. Sci. École Norm. Sup.* 40 (2007) 387–444 MR Zbl
- [Yau 2017] **M-L Yau**, *Surgery and isotopy of Lagrangian surfaces*, from “Proceedings of the sixth international congress of Chinese mathematicians, II” (C-S Lin, L Yang, S-T Yau, J Yu, editors), *Adv. Lect. Math.* 37, International, Somerville, MA (2017) 143–162 MR Zbl

*Department of Mathematics, University of California Davis
Davis, CA, United States*

*Department of Mathematics, University of California Davis
Davis, CA, United States*

casals@math.ucdavis.edu, dweng@ucdavis.edu

Proposed: Leonid Polterovich
Seconded: Dmitri Burago, Mark Gross

Received: 1 August 2022
Revised: 19 June 2023

GEOMETRY & TOPOLOGY

msp.org/gt

MANAGING EDITOR

András I Stipsicz Alfréd Rényi Institute of Mathematics
stipsicz@renyi.hu

BOARD OF EDITORS

Mohammed Abouzaid	Stanford University abouzaid@stanford.edu	Mark Gross	University of Cambridge mgross@dpmms.cam.ac.uk
Dan Abramovich	Brown University dan_abramovich@brown.edu	Rob Kirby	University of California, Berkeley kirby@math.berkeley.edu
Ian Agol	University of California, Berkeley ianagol@math.berkeley.edu	Bruce Kleiner	NYU, Courant Institute bkleiner@cims.nyu.edu
Arend Bayer	University of Edinburgh arend.bayer@ed.ac.uk	Sándor Kovács	University of Washington skovacs@uw.edu
Mark Behrens	University of Notre Dame mbehren1@nd.edu	Urs Lang	ETH Zürich urs.lang@math.ethz.ch
Mladen Bestvina	University of Utah bestvina@math.utah.edu	Marc Levine	Universität Duisburg-Essen marc.levine@uni-due.de
Martin R Bridson	University of Oxford bridson@maths.ox.ac.uk	Ciprian Manolescu	University of California, Los Angeles cm@math.ucla.edu
Jim Bryan	University of British Columbia jbryan@math.ubc.ca	Haynes Miller	Massachusetts Institute of Technology hrm@math.mit.edu
Dmitri Burago	Pennsylvania State University burago@math.psu.edu	Tomasz Mrowka	Massachusetts Institute of Technology mrowka@math.mit.edu
Tobias H Colding	Massachusetts Institute of Technology colding@math.mit.edu	Aaron Naber	Northwestern University anaber@math.northwestern.edu
Simon Donaldson	Imperial College, London s.donaldson@ic.ac.uk	Peter Ozsváth	Princeton University petero@math.princeton.edu
Yasha Eliashberg	Stanford University eliash-gt@math.stanford.edu	Leonid Polterovich	Tel Aviv University polterov@post.tau.ac.il
Benson Farb	University of Chicago farb@math.uchicago.edu	Colin Rourke	University of Warwick gt@maths.warwick.ac.uk
David M Fisher	Rice University davidfisher@rice.edu	Roman Sauer	Karlsruhe Institute of Technology roman.sauer@kit.edu
Mike Freedman	Microsoft Research michaelf@microsoft.com	Stefan Schwede	Universität Bonn schwede@math.uni-bonn.de
David Gabai	Princeton University gabai@princeton.edu	Natasa Sesum	Rutgers University natasas@math.rutgers.edu
Stavros Garoufalidis	Southern U. of Sci. and Tech., China stavros@mpim-bonn.mpg.de	Gang Tian	Massachusetts Institute of Technology tian@math.mit.edu
Cameron Gordon	University of Texas gordon@math.utexas.edu	Ulrike Tillmann	Oxford University tillmann@maths.ox.ac.uk
Jesper Grodal	University of Copenhagen jg@math.ku.dk	Nathalie Wahl	University of Copenhagen wahl@math.ku.dk
Misha Gromov	IHÉS and NYU, Courant Institute gromov@ihes.fr	Anna Wienhard	Universität Heidelberg wienhard@mathi.uni-heidelberg.de


See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2024 is US \$805/year for the electronic version, and \$1135/year (+\$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFLOW[®] from MSP.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing
<http://msp.org/>

© 2024 Mathematical Sciences Publishers

GEOMETRY & TOPOLOGY

Volume 28 Issue 2 (pages 497–1003) 2024

On the top-weight rational cohomology of \mathcal{A}_g	497
MADELINE BRANDT, JULIETTE BRUCE, MELODY CHAN, MARGARIDA MELO, GWYNETH MORELAND and COREY WOLFE	
Algebraic uniqueness of Kähler–Ricci flow limits and optimal degenerations of Fano varieties	539
JIYUAN HAN and CHI LI	
Valuations on the character variety: Newton polytopes and residual Poisson bracket	593
JULIEN MARCHÉ and CHRISTOPHER-LLOYD SIMON	
The local (co)homology theorems for equivariant bordism	627
MARCO LA VECCHIA	
Configuration spaces of disks in a strip, twisted algebras, persistence, and other stories	641
HANNAH ALPERT and FEDOR MANIN	
Closed geodesics with prescribed intersection numbers	701
YANN CHAUBET	
On endomorphisms of the de Rham cohomology functor	759
SHIZHANG LI and SHUBHODIP MONDAL	
The nonabelian Brill–Noether divisor on $\overline{\mathcal{M}}_{13}$ and the Kodaira dimension of $\overline{\mathcal{R}}_{13}$	803
GAVRIL FARKAS, DAVID JENSEN and SAM PAYNE	
Orbit equivalences of \mathbb{R} –covered Anosov flows and hyperbolic-like actions on the line	867
THOMAS BARTHELMÉ and KATHRYN MANN	
Microlocal theory of Legendrian links and cluster algebras	901
ROGER CASALS and DAPING WENG	
Correction to the article Bimodules in bordered Heegaard Floer homology	1001
ROBERT LIPSHITZ, PETER OZSVÁTH and DYLAN P THURSTON	