

Knot contact homology

TOBIAS EKHOLM

JOHN B ETNYRE

LENHARD NG

MICHAEL G SULLIVAN

The conormal lift of a link K in \mathbb{R}^3 is a Legendrian submanifold Λ_K in the unit cotangent bundle $U^*\mathbb{R}^3$ of \mathbb{R}^3 with contact structure equal to the kernel of the Liouville form. Knot contact homology, a topological link invariant of K , is defined as the Legendrian homology of Λ_K , the homology of a differential graded algebra generated by Reeb chords whose differential counts holomorphic disks in the symplectization $\mathbb{R} \times U^*\mathbb{R}^3$ with Lagrangian boundary condition $\mathbb{R} \times \Lambda_K$.

We perform an explicit and complete computation of the Legendrian homology of Λ_K for arbitrary links K in terms of a braid presentation of K , confirming a conjecture that this invariant agrees with a previously defined combinatorial version of knot contact homology. The computation uses a double degeneration: the braid degenerates toward a multiple cover of the unknot, which in turn degenerates to a point. Under the first degeneration, holomorphic disks converge to gradient flow trees with quantum corrections. The combined degenerations give rise to a new generalization of flow trees called multiscale flow trees. The theory of multiscale flow trees is the key tool in our computation and is already proving to be useful for other computations as well.

53D42; 57R17, 57M27

1 Introduction

1.1 Knot contact homology

Let $K \subset \mathbb{R}^3$ be a link. The *conormal lift* $\Lambda_K \subset U^*\mathbb{R}^3$ of K in the unit cotangent bundle $U^*\mathbb{R}^3$ of \mathbb{R}^3 is the sub-bundle of $U^*\mathbb{R}^3$ over K consisting of covectors that vanish on TK . The submanifold Λ_K is topologically a union of 2-tori, one for each component of K . The unit cotangent bundle carries a natural contact 1-form α : if $p dq$ denotes the Liouville form on $T^*\mathbb{R}^3$ then α is the restriction of $p dq$ to $U^*\mathbb{R}^3$. The conormal lift Λ_K is a Legendrian submanifold with respect to the contact structure induced by α . Furthermore, if K_t , $0 \leq t \leq 1$, is a smooth isotopy of links then Λ_{K_t}

is a Legendrian isotopy of tori. Consequently, Legendrian isotopy invariants of Λ_K give topological isotopy invariants of K itself.

Contact homology, and Legendrian (contact) homology, is a rich source of deformation invariants in contact topology. Legendrian homology associates a differential graded algebra (DGA) to a Legendrian submanifold $\Lambda \subset Y$ of a contact manifold Y , in which the algebra is generated by the Reeb chords of Λ and the Reeb orbits in Y , and the differential is defined by a count of holomorphic curves in the symplectization $\mathbb{R} \times Y$ with Lagrangian boundary condition $\mathbb{R} \times \Lambda$. The main result of this paper is a complete description of the DGA of the conormal lift Λ_K of a link K in terms of a braid presentation of K ; see Theorem 1.1 below.

Before presenting Theorem 1.1, we discuss some consequences of it. Theorem 1.1 confirms the conjecture that the combinatorial knot invariant defined and studied by the third author in [26; 27; 29], called “knot contact homology” in those works, indeed equals the Legendrian homology of the conormal lift. (In fact, the Legendrian homology described here is a non-trivial extension of the previously defined version of knot contact homology; see Section 1.2.) We note that Legendrian homology can be expressed combinatorially in other circumstances, for example by Chekanov [7] and Eliashberg [19] for 1–dimensional Legendrian knots in \mathbb{R}^3 with the standard contact structure. Also, computations of Legendrian homology in higher dimensions have been carried out in some circumstances; see eg Dimitroglou Rizell [8], Ekholm, Etnyre and Sullivan [15], and Ekholm and Kálmán [18].

The results of the present paper constitute one of the first complete and reasonably involved computations of Legendrian homology in higher dimensions (in the language of Lagrangian Floer theory, the computation roughly corresponds to the calculation of the differential and all higher product operations). Extensions of our techniques have already found applications, see eg Bourgeois, Ekholm and Eliashberg [5], and we expect that further extensions may be used in other higher-dimensional situations in the future.

A more concrete consequence of Theorem 1.1 is its application to an interesting general question in symplectic topology: to what extent do symplectic- and contact-geometric objects naturally associated to objects in smooth topology remember the underlying smooth topology? More specifically, how much of the smooth topology is encoded by holomorphic-curve techniques on the symplectic/contact side? The construction that associates the (symplectic) cotangent bundle to any (smooth) manifold has been much studied recently in this regard; see eg Abouzaid [1; 2], Fukaya, Seidel and Smith [21], and Nadler [24]. In our setting, the general question specializes to the following.

Question *How much does the Legendrian-isotopy class of the conormal Λ_K remember about the smooth-isotopy class of K ? In particular, if Λ_{K_1} and Λ_{K_2} are Legendrian isotopic, are K_1 and K_2 necessarily smoothly isotopic?*

At this writing, it is possible that the answer to the second question is “yes” in general. One consequence of Theorem 1.1 is that the answer is “yes” if K_1 is the unknot: the conormal lift detects the unknot. See Corollary 1.4 below.

Another geometric application of our techniques is the development of a filtered version of the Legendrian DGA associated to Λ_K when K is a link transverse to the standard contact structure $\ker(dz - y dx)$ in \mathbb{R}^3 . This is carried out in Ekholm, Etnyre, Ng and Sullivan [12], which relies heavily on the computations from the present paper. There is a related combinatorial treatment in Ng [30], where it is shown that this filtered version (“transverse homology”) constitutes a very effective invariant of transverse knots. We remark that although [30] can be read as a stand-alone paper without reference to Legendrian homology or holomorphic disks, the combinatorial structure of transverse homology presented in [12; 30] was crucially motivated by the holomorphic-disk enumerations we present here.

1.2 Main result

We now turn to a more precise description of our main result. To this end we first need to introduce some notation. Let K be an arbitrary r -component link in \mathbb{R}^3 , given by the closure of a braid B with n strands. Then Λ_K is a disjoint union of r 2-tori, which (after identifying the tangent bundle with the cotangent bundle and the normal bundle to K with its tubular neighborhood) we may view as the boundaries of tubular neighborhoods of each component of K . We can then choose a set of generators $\lambda_1, \mu_1, \dots, \lambda_r, \mu_r$ of $H_1(\Lambda_K)$, with λ_i and μ_i corresponding to the longitude (running along the component of K) and meridian (running along a fiber of Λ_K over a point in K) of the i^{th} torus.

We proceed to an algebraic definition of the DGA that computes the Legendrian homology of Λ_K . The DGA is $(\mathcal{A}_n, \partial)$, where the algebra \mathcal{A}_n is the unital graded algebra over \mathbb{Z} generated by the group ring $\mathbb{Z}[H_1(\Lambda_K)] = \mathbb{Z}[\lambda_1^{\pm 1}, \mu_1^{\pm 1}, \dots, \lambda_r^{\pm 1}, \mu_r^{\pm 1}]$ in degree 0, along with the following generators:

$$\begin{aligned} \{a_{ij}\}_{1 \leq i, j \leq n; i \neq j} & \text{ in degree 0,} \\ \{b_{ij}\}_{1 \leq i, j \leq n; i \neq j} & \text{ in degree 1,} \\ \{c_{ij}\}_{1 \leq i, j \leq n} & \text{ in degree 1,} \\ \{e_{ij}\}_{1 \leq i, j \leq n} & \text{ in degree 2.} \end{aligned}$$

Note that in \mathcal{A}_n , the generators $a_{ij}, b_{ij}, c_{ij}, e_{ij}$ do not commute with each other or with non-trivial elements of $\mathbb{Z}[H_1(\Lambda_K)]$.

We next define a differential ∂ on \mathcal{A}_n . Introduce variables $\tilde{\mu}_1, \dots, \tilde{\mu}_n$ of degree 0, and write $\tilde{\mathcal{A}}_n^0$ for the free unital algebra over \mathbb{Z} generated by the ring $\mathbb{Z}[\tilde{\mu}_1^{\pm 1}, \dots, \tilde{\mu}_n^{\pm 1}]$ and the a_{ij} . Now if σ_k is a standard generator of the braid group B_n on n strands, then define the automorphism $\phi_{\sigma_k}: \tilde{\mathcal{A}}_n^0 \rightarrow \tilde{\mathcal{A}}_n^0$ by:

$$\begin{aligned} \phi_{\sigma_k}(a_{ij}) &= a_{ij} & i, j \neq k, k+1 \\ \phi_{\sigma_k}(a_{kk+1}) &= -a_{k+1k} \\ \phi_{\sigma_k}(a_{k+1k}) &= -\tilde{\mu}_k a_{kk+1} \tilde{\mu}_{k+1}^{-1} \\ \phi_{\sigma_k}(a_{ik+1}) &= a_{ik} & i \neq k, k+1 \\ \phi_{\sigma_k}(a_{k+1i}) &= a_{ki} & i \neq k, k+1 \\ \phi_{\sigma_k}(a_{ik}) &= a_{ik+1} - a_{ik} a_{kk+1} & i < k \\ \phi_{\sigma_k}(a_{ik}) &= a_{ik+1} - a_{ik} \tilde{\mu}_k a_{kk+1} \tilde{\mu}_{k+1}^{-1} & i > k+1 \\ \phi_{\sigma_k}(a_{ki}) &= a_{k+1i} - a_{k+1k} a_{ki} & i \neq k, k+1 \\ \phi_{\sigma_k}(\tilde{\mu}_k^{\pm 1}) &= \tilde{\mu}_{k+1}^{\pm 1} \\ \phi_{\sigma_k}(\tilde{\mu}_{k+1}^{\pm 1}) &= \tilde{\mu}_k^{\pm 1} \\ \phi_{\sigma_k}(\tilde{\mu}_i^{\pm 1}) &= \tilde{\mu}_i^{\pm 1} & i \neq k, k+1 \end{aligned}$$

The map ϕ induces a homomorphism from B_n to the automorphism group of $\tilde{\mathcal{A}}_n^0$: if $B = \sigma_{k_1}^{\epsilon_1} \cdots \sigma_{k_m}^{\epsilon_m} \in B_n$, with $\epsilon_l = \pm 1$ for $l = 1, \dots, m$, then $\phi_B = \phi_{\sigma_{k_1}^{\epsilon_1}} \circ \cdots \circ \phi_{\sigma_{k_m}^{\epsilon_m}}$. See Proposition 2.9.

Now let \mathcal{A}_n^0 denote the subalgebra of \mathcal{A}_n of elements of degree 0, generated by the μ_j, λ_j and a_{ij} . For $1 \leq j \leq n$, let $\alpha(j) \in \{1, \dots, r\}$ be the number of the link component of K corresponding to the j^{th} strand of B . Then ϕ_B descends to an automorphism of \mathcal{A}_n^0 , which we also denote by ϕ_B , by setting $\tilde{\mu}_j = \mu_{\alpha(j)}$ for all j , and having ϕ_B act as the identity on λ_j for all j .

For convenient notation we assemble the generators of \mathcal{A}_n into $n \times n$ matrices. Writing \mathbb{M}_{ij} for the element in position ij in the $n \times n$ matrix \mathbb{M} , we define the $n \times n$ matrices

$$\begin{aligned} \mathbb{A}: & \begin{cases} \mathbb{A}_{ij} = a_{ij} & \text{if } i < j, \\ \mathbb{A}_{ij} = a_{ij} \mu_{\alpha(j)} & \text{if } i > j, \\ \mathbb{A}_{ii} = 1 + \mu_{\alpha(i)}, \end{cases} & \mathbb{B}: & \begin{cases} \mathbb{B}_{ij} = b_{ij} & \text{if } i < j, \\ \mathbb{B}_{ij} = b_{ij} \mu_{\alpha(j)} & \text{if } i > j, \\ \mathbb{B}_{ii} = 0, \end{cases} \\ \mathbb{C}: & \{\mathbb{C}_{ij} = c_{ij}, & \mathbb{E}: & \{\mathbb{E}_{ij} = e_{ij}. \end{aligned}$$

To the braid B , we now associate two $n \times n$ matrices Φ_B^L, Φ_B^R with coefficients in \mathcal{A}_n^0 as follows. Set \widehat{B} to be the $(n + 1)$ -strand braid obtained by adding to B an extra strand labeled 0 that does not interact with the other strands. Then ϕ gives a map $\phi_{\widehat{B}}$ acting on the algebra generated by homology classes and $\{a_{ij}\}_{0 \leq i, j \leq n; i \neq j}$, and we can define $\Phi_{\widehat{B}}^L, \Phi_{\widehat{B}}^R$ by

$$\phi_{\widehat{B}}(a_{i0}) = \sum_{j=1}^n (\Phi_B^L)_{ij} a_{j0} \quad \text{and} \quad \phi_{\widehat{B}}(a_{0j}) = \sum_{i=1}^n a_{0i} (\Phi_B^R)_{ij}.$$

(Note that by the above formula for ϕ_{σ_k} , any monomial contributing to $\phi_B(a_{i0})$ (respectively $\phi_B(a_{0j})$) must begin with a generator of the form a_{l0} (respectively, end with a generator a_{0l}). Also, since the 0th strand does not interact with the others, $\tilde{\mu}_0$ does not appear anywhere in the expressions for $\phi_{\widehat{B}}(a_{i0}), \phi_{\widehat{B}}(a_{0j})$, and so $\Phi_{\widehat{B}}^L, \Phi_{\widehat{B}}^R$ have coefficients in \mathcal{A}_n^0 .)

Finally, define a diagonal $n \times n$ matrix λ as follows. Consider the strands $1, \dots, n$ of the braid. Call a strand *leading* if it is the first strand in this ordering belonging to its component. Define

$$\lambda: \begin{cases} \lambda_{ij} = 0 & \text{if } i \neq j, \\ \lambda_{ii} = \lambda_{\alpha(i)} & \text{if the } i^{\text{th}} \text{ strand is leading,} \\ \lambda_{ii} = 1 & \text{otherwise.} \end{cases}$$

We can now state our main result.

Theorem 1.1 *Let $K \subset \mathbb{R}^3$ be an oriented link given by the closure of a braid B with n strands. After Legendrian isotopy, the conormal lift $\Lambda_K \subset U^*\mathbb{R}^3$ has Reeb chords in graded one-to-one correspondence with the generators of \mathcal{A}_n . Consequently, the Legendrian algebra of*

$$\Lambda_K \subset J^1(S^2) \approx U^*\mathbb{R}^3$$

is identified with \mathcal{A}_n . Under this identification, the differential of the Legendrian DGA is the map $\partial: \mathcal{A}_n \rightarrow \mathcal{A}_n$ determined by the following matrix equations:

$$\begin{aligned} \partial \mathbb{A} &= 0, \\ \partial \mathbb{B} &= -\lambda^{-1} \cdot \mathbb{A} \cdot \lambda + \Phi_B^L \cdot \mathbb{A} \cdot \Phi_B^R, \\ \partial \mathbb{C} &= \mathbb{A} \cdot \lambda + \mathbb{A} \cdot \Phi_B^R, \\ \partial \mathbb{E} &= \mathbb{B} \cdot (\Phi_B^R)^{-1} + \mathbb{B} \cdot \lambda^{-1} - \Phi_B^L \cdot \mathbb{C} \cdot \lambda^{-1} + \lambda^{-1} \cdot \mathbb{C} \cdot (\Phi_B^R)^{-1}, \end{aligned}$$

where if \mathbb{M} is a matrix, the matrix $\partial \mathbb{M}$ is defined by $(\partial \mathbb{M})_{ij} = \partial(\mathbb{M}_{ij})$.

Theorem 1.1 has the following corollary, which establishes a conjecture of the third author [28; 29].

Corollary 1.2 *For a knot K , the combinatorial framed knot DGA of K , as defined in [29], is isomorphic to the Legendrian DGA of Λ_K .*

In [29], the framed knot DGA is shown to be a knot invariant via a purely algebraic but somewhat involved argument that shows its invariance under Markov moves. We note that Theorem 1.1 provides another proof of invariance. Here the relevant equivalence relation for invariance is that of stable tame isomorphism of DGAs introduced in [7].

Corollary 1.3 *The DGA given in Theorem 1.1, and consequently the framed knot DGA of [29], is a link invariant: two braids whose closures are isotopic links produce the same DGA up to stable tame isomorphism.*

Proof This is a direct consequence of invariance properties of the Legendrian DGA under Legendrian isotopy; see Theorem 2.6. \square

Quite a bit is known about the behavior of combinatorial knot contact homology, and Corollary 1.2 allows us to use this knowledge to deduce results about the geometry of conormal lifts. For instance, we have the following results.

Corollary 1.4 *A knot K is isotopic to the unknot U if and only if the conormal lift Λ_K of K is Legendrian isotopic to the conormal lift Λ_U of U .*

Proof The degree-0 homology of the framed knot DGA detects the unknot [29, Proposition 5.10]. \square

Corollary 1.5 *If K_1 and K_2 are knots such that Λ_{K_1} and Λ_{K_2} are Legendrian isotopic, then K_1 and K_2 have the same Alexander polynomial.*

Proof The degree-1 linearized homology of the framed knot DGA with respect to a certain canonical augmentation encodes the Alexander polynomial [29, Corollary 4.5]. \square

We close this subsection by discussing a subtlety hidden in the statement of Corollary 1.2. The DGA $(\mathcal{A}_\eta, \partial)$ defined in Theorem 1.1 is actually an *extension* of the combinatorial knot DGA introduced by the third author in [26; 27; 29], in two significant ways.¹ First,

¹There are also some inconsequential sign differences between our formulation and the knot DGA of [26; 27; 29]; see [30] for a proof that the different sign conventions yield isomorphic DGAs.

the original combinatorial knot DGA assumed that K is a one-component knot, whereas our $(\mathcal{A}_n, \partial)$ works for general links, associating separate homology variables λ_j, μ_j to each component. Second, in the original combinatorial knot DGA as presented in [29], and in the filtered version for transverse knots presented in [12], homology variables commute with the generators a_{ij} , etc, while they do not commute here. We may thus think of the original knot DGA as a “homology-commutative” quotient of our “full” DGA; see Section 2.3.2 for further discussion.

Although we do not pursue this point here, it seems quite possible that the full DGA introduced here constitutes a stronger link invariant, or otherwise encodes more information, than the homology-commutative knot DGA from [29]. For example, the proof in [29] that the framed knot DGA detects the unknot uses work of Dunfield and Garoufalidis [9], building on some deep gauge-theoretic results of Kronheimer and Mrowka [22]. However, if we use the full DGA rather than the homology-commutative quotient, then there is an alternate proof that knot contact homology detects the unknot, in unpublished work of the first and third authors along with Cieliebak and Latschev, which uses nothing more than the Loop Theorem.

1.3 Strategy and outline of paper

We conclude this introduction with a description of the strategy of our proof of Theorem 1.1. A loose sketch of this approach has previously been summarized in Ekholm and Etnyre [11].

The unit cotangent bundle $U^*\mathbb{R}^3$ with contact form the restriction of the Liouville form $p dq$ is contactomorphic to the 1-jet space $J^1(S^2) \approx T^*S^2 \times \mathbb{R}$ of S^2 , with contact form $dz - p dq$ where z is the coordinate along the \mathbb{R} -factor. To find the rigid holomorphic disks that contribute to Legendrian homology for a Legendrian submanifold of any 1-jet space, one can use gradient flow trees (Ekholm [10]). In our case, rather than directly examining the gradient flow trees for the conormal lift $\Lambda_K \subset J^1(S^2)$ of a link K , we break the computation down into three steps. First, we compute the differential for the conormal lift $\Lambda = \Lambda_U$ of the unknot U , which we represent as a planar circle. This is done by calculating gradient flow trees for a particular small perturbation of Λ .

Second, given an arbitrary link K , let B be a braid whose closure is K . We can view the closure of B as lying in the solid torus given as a small tubular neighborhood of U , and we can thus realize K as a braid that is C^1 -close to U . Then Λ_K lies in a small neighborhood of Λ , and by the Legendrian neighborhood theorem we can choose this neighborhood to be contactomorphic to the 1-jet space $J^1(\Lambda) = J^1(T^2)$:

$$\Lambda_K \subset J^1(\Lambda) = J^1(T^2) \subset J^1(S^2).$$

We can use gradient flow trees to find the rigid holomorphic disks in $J^1(T^2)$ with boundary on Λ_K . This turns out to be easier than directly calculating the analogous disks in $J^1(S^2)$ because Λ_K is everywhere transverse to the fibers of $J^1(T^2)$ (ie, its caustic is empty). We can assemble the result, which is computed in terms of the braid B , as the Legendrian DGA of $\Lambda_K \subset J^1(T^2)$, which is a subalgebra of the Legendrian DGA of $\Lambda_K \subset J^1(S^2)$.

Finally, we prove that there is a one-to-one correspondence between rigid holomorphic disks in $J^1(S^2)$ with boundary on Λ_K , and certain objects that we call *rigid multiscale flow trees* determined by Λ and Λ_K , which arise as follows. As we let K approach U , Λ_K approaches Λ and each holomorphic disk with boundary on Λ_K approaches a holomorphic disk with boundary on Λ with partial flow trees of $\Lambda_K \subset J^1(\Lambda)$ attached along its boundary. Here the flow trees correspond to the thin parts of the holomorphic disks before the limit; in these parts, the energy approaches zero. In a multiscale flow tree we substitute the holomorphic disk with boundary on Λ in this limit by the corresponding flow tree computed in the first step, and obtain a flow tree of $\Lambda \subset J^1(S^2)$ with portions of flow trees of $\Lambda_K \subset J^1(\Lambda)$ attached along its boundary.

Remark 1.6 An alternative proof strategy would be to apply the relation between flow trees and holomorphic disks directly for the front of $\Lambda_K \subset J^1(S^2)$ and find the relation to multiscale flow trees from there. The current approach has advantages over the direct approach in that it keeps more symmetry, reminiscent of phenomena in Morse theory in Bott degenerate situations where it is sometimes an advantage to keep the Bott symmetry rather than perturbing out.

From a more computational perspective, the actual flow trees on the front of $\Lambda_{K;\eta}$ pass through front singularities numerous times and their study is complicated by tree vertices related to the cusp edge (called Y_1 -vertices, ends, and switches in [10]). In any pure flow tree approach one would have to deal with these. Our current setup, by contrast, introduces some amount of new technical analytical problems, but allows us to push all these complications to the study of the front of the unknot, where they are readily solved. The remaining study of flow trees is then carried out near the conormal lift Λ of the unknot; the front of Λ_K is non-singular in $J^0(\Lambda)$, and the flow trees in our approach do not see the singularities of the front of Λ_K as a subset of $J^0(S^2)$. The relative simplicity of the flow trees for this non-singular front and their rather direct connection to the braid group is what allows us to find a closed formula for the differential.

Here is the plan for the rest of the paper. In Section 2, we present background material, including the definitions of the conormal construction, Legendrian homology, and

gradient flow trees. In Section 3, we use gradient flow trees to accomplish the first two steps in the three-step process outlined above: calculating holomorphic disks for $\Lambda \subset J^1(S^2)$ and $\Lambda_K \subset J^1(T^2)$. We extend these calculations to multiscale flow trees in Section 4 to count holomorphic disks for $\Lambda_K \subset J^1(S^2)$, completing the proof of Theorem 1.1.

The computations in Section 4 rely on some technical results about multiscale flow trees whose proofs are deferred to the final two sections of the paper. In Section 5, we establish Theorem 4.3, which gives a one-to-one correspondence between holomorphic disks and multiscale flow trees. In Section 6, we prove Theorems 4.5 and 4.6, which deal with combinatorial signs arising from choices of orientations of the relevant moduli spaces of flow trees and multiscale flow trees.

Acknowledgements

TE was partially supported by the Göran Gustafsson Foundation for Research in Natural Sciences and Medicine. JBE was partially supported by NSF grant DMS-0804820. LLN was partially supported by NSF grant DMS-0706777 and NSF CAREER grant DMS-0846346. MGS was partially supported by NSF grants DMS-0707091 and DMS-1007260. The authors would also like to thank MSRI for hosting them in spring 2010 during part of this collaboration. Finally, we thank Cecilia Karlsson for useful remarks on coherent orientations, and the referees for many helpful comments.

2 Conormal lifts, Legendrian homology, and flow trees

In this section we review background material. We begin by discussing the conormal lift construction for links in \mathbb{R}^3 in Section 2.1, and place it in the context of 1-jet spaces of surfaces in Section 2.2. In Section 2.3, we review the definition of Legendrian homology. For our purposes, the holomorphic disks counted in Legendrian homology are replaced by flow trees, which we discuss in Section 2.4. Vector splitting along flow trees, which is needed later when assigning signs to rigid flow trees, is discussed in Section 2.5. We end with a compilation in Section 2.6 of algebraic results about the map ϕ_B and the matrices Φ_B^L and Φ_R^L that are used in the proof of our main result, Theorem 1.1.

2.1 The conormal construction

Fixing the standard flat metric on \mathbb{R}^3 we consider the unit cotangent bundle $U^*\mathbb{R}^3 \approx \mathbb{R}^3 \times S^2$ of \mathbb{R}^3 . The Liouville form on the cotangent bundle $T^*\mathbb{R}^3$ is

$$\theta = \sum_{i=1}^3 p_i dq_i,$$

where $q = (q_1, q_2, q_3)$ are coordinates on \mathbb{R}^3 and $p = (p_1, p_2, p_3)$ are coordinates in the fibers of $T^*\mathbb{R}^3$. The restriction $\theta|_{U^*\mathbb{R}^3}$ is a contact 1-form. We denote its associated contact structure $\xi = \ker \theta$.

The standard contact 1-form on the 1-jet space $J^1(S^2) = T^*S^2 \times \mathbb{R}$ of S^2 is given by

$$\alpha = dz - \theta,$$

where z is a coordinate in the \mathbb{R} -factor and where θ is the Liouville form on T^*S^2 . Using the standard inner product $\langle \cdot, \cdot \rangle$ on \mathbb{R}^3 to identify vectors and covectors we may write T^*S^2 as follows:

$$T^*S^2 = \{(x, y) \in \mathbb{R}^3 \times \mathbb{R}^3 : |x| = 1, \langle x, y \rangle = 0\}.$$

We define the map $\phi: U^*\mathbb{R}^3 \rightarrow J^1(S^2) = T^*S^2 \times \mathbb{R}$ by

$$(2-1) \quad \phi(q, p) = (p, q - \langle q, p \rangle p, \langle q, p \rangle)$$

and notice that ϕ is a diffeomorphism such that $\phi^*\alpha = \theta$. Thus $U^*\mathbb{R}^3$ and $J^1(S^2)$ are contactomorphic.

Let K be a knot or link in \mathbb{R}^3 . We associate to K its conormal lift

$$\Lambda_K = \{u \in U^*\mathbb{R}^3|_K : u(v) = 0 \text{ for all } v \in TK\},$$

which topologically is a union of tori, one for each component of K . By definition $\theta|_{\Lambda_K} = 0$, ie, Λ_K is Legendrian. Furthermore, smooth isotopies of K induce Legendrian isotopies of Λ_K . In particular the Legendrian isotopy class of Λ_K (and consequently any invariant thereof) is an invariant of the isotopy class of K . For more on this construction, see [11].

2.2 Legendrian submanifolds in the 1-jet space of a surface

In this section we review well-known material on fronts associated to Legendrian submanifolds in a 1-jet space. For general treatments of this subject, see eg Arnold and Givental [3], and Arnold, Gusein-Zade and Varchenko [4].

Let S be a surface and consider a Legendrian submanifold $\Lambda \subset J^1(S) = T^*S \times \mathbb{R}$. After a small perturbation of Λ we can assume that Λ is in general position with respect to the *Lagrangian projection* $\Pi: J^1(S) \rightarrow T^*S$, the *front projection* $\Pi_F: J^1(S) \rightarrow S \times \mathbb{R}$ and the *base projection* $\pi: J^1(S) \rightarrow S$.

General position for Π means that the self-intersections of the Lagrangian immersion $\Pi(L)$ consists of transverse double points. General position for Π_F means that singularities of $\Pi_F|_\Lambda$ are of two types: *cusplike edges* and *swallow tails*. For a more precise description we first introduce some notation: the *caustic* $\Sigma \subset \Lambda$ is the critical set of $\Pi_F|_\Lambda$. General position for Π_F first implies that Σ is a closed 1–submanifold along which the rank of $\Pi_F|_\Lambda$ equals 1. The kernel field $\ker(d\Pi_F|_\Lambda)$ is then a line field l along Σ and general position for Π_F implies that l has only transverse tangencies with $T\Sigma$ (ie, all tangencies have order one). This gives a stratification of the caustic: the 1–dimensional top stratum, called the *cusplike edge*, consisting of points where l is transverse to Σ and the 0–dimensional complement, called the set of *swallow tail points*, where l is tangent to $T\Sigma$ of order 1. Finally, general position for π means that the image of the caustic under π is stratified: $\pi(\Sigma) = \Sigma_1 \supset (\Sigma_2 \cup \Sigma_2^{sw})$, where Σ_2 is the set of transverse self-intersections of $\pi(\Sigma)$, and Σ_2^{sw} is the image under π of the swallow tail points around which $\pi(\Sigma)$ has the form of a semi-cubical cusp.

For Λ in general position we obtain the following local descriptions:

- If $p \in \pi(\Lambda) - \Sigma_1$ then p has a neighborhood $U \subset S$ such that $\pi^{-1}(U) \cap \Lambda \subset J^1(S)$ is the union of a finite number of 1–jet graphs of functions f_1, \dots, f_n on U . We call such functions *local defining functions of Λ* .




Figure 1: A local model for the cusplike edge is shown on the left and a local model for the swallow tail singularity is shown on the right.

- If $p \in \Sigma_1 - (\Sigma_2 \cup \Sigma_2^{sw})$ then there is a neighborhood $U \subset S$ of p that is diffeomorphic to the open unit disk, $U = U_+ \cup U_-$, where U_+ (respectively U_-) corresponds to the upper (respectively lower) half disk and where $\pi(\Sigma) \cap U$

corresponds to $U_+ \cap U_-$, with the following properties. The intersection $\pi^{-1}(U) \cap \Lambda \subset J^1(S)$ consists of $n \geq 0$ smooth sheets given by the 1-jet graphs of local defining functions and one cusped sheet given by the 1-jet graph of two functions $h_0, h_1: U_+ \rightarrow \mathbb{R}$. We can choose coordinates (x_1, x_2) near ∂U_+ so that h_0 and h_1 have the normal form

$$h_0(x_1, x_2) = \frac{1}{3}(2x_1)^{3/2} + \beta x_1 + \alpha x_2$$

$$h_1(x_1, x_2) = -\frac{1}{3}(2x_1)^{3/2} + \beta x_1 + \alpha x_2$$

for some constants α and β ; see [10, Equation (2–1)]. In particular dh_0 and dh_1 agree near the boundary; see the left diagram in Figure 1. We also call h_0 and h_1 local defining functions (for a cusped sheet).

- If $p \in \Sigma_2^{\text{sw}}$ then there is a neighborhood U in which $n \geq 0$ sheets are smooth and one sheet is a standard swallow tail sheet; see [10, Remark 2.5] and the right diagram in Figure 1.
- If $p \in \pi(\Sigma_2)$ there is a neighborhood U diffeomorphic to the unit disk over which the Legendrian consists of $n \geq 0$ smooth sheets and two cusped sheets, one defined over the upper half-plane and one over the right half-plane.

The following simple observation concerning double points of Π will be important in the definition of the Legendrian homology algebra.

Lemma 2.1 *A point $s \in \Pi(\Lambda)$ is a double point if and only if $p = \pi(\Pi^{-1}(s))$ is a critical point of the difference of two local defining functions for Λ at p .*

2.3 Legendrian homology in the 1-jet space of a surface

In this subsection we briefly review Legendrian (contact) homology in $J^1(S)$, where S is an orientable surface. In fact, for our applications it is sufficient to consider $S \approx S^2$ and $S \approx T^2$. We refer the reader to Ekholm, Etnyre and Sullivan [14] for details on the material presented here.

2.3.1 Geometric preliminaries As in Section 2.1, we consider $J^1(S) = T^*S \times \mathbb{R}$, with contact form $\alpha = dz - \theta$, where z is a coordinate in the \mathbb{R} -factor and θ is the Liouville form on T^*S . The Reeb vector field R_α of α is given by $R_\alpha = \partial_z$ and thus Reeb flow lines are simply $\{p\} \times \mathbb{R}$ for any $p \in T^*S$. In particular, if $\Lambda \subset J^1(S)$ is a Legendrian submanifold then Reeb chords of Λ (Reeb flow segments which begin and end on Λ) correspond to pairs of distinct points $y_1, y_2 \in \Lambda$ such that $\Pi(y_1) = \Pi(y_2)$, where $\Pi: J^1(S) \rightarrow T^*S$ is the Lagrangian projection. As noted in

Lemma 2.1, for such Legendrian submanifolds there is thus a bijective correspondence between Reeb chords of $\Lambda \subset T^*S \times \mathbb{R}$ and double points of $\Pi(\Lambda)$ and in this case we will sometimes use the phrase “Reeb chord” and “double point” interchangeably. As mentioned in Section 2.2, after small perturbation of Λ , $\Pi|_{\Lambda}$ is a selftransverse Lagrangian immersion with a finite number of double points.

Consider an almost complex structure on T^*S that is compatible with $d\theta$, by which we can view $T(T^*S)$ as a complex vector bundle. Since $T(T^*S)|_S$ is the complexification of T^*S we find that $c_1(T(T^*S)) = 0$ and hence there is a complex trivialization of $T(T^*S)$. The orientation of S induces a trivialization of the real determinant line bundle (second exterior power) $\Lambda^2 T^*M$, which in turn gives a trivialization, called the *orientation trivialization*, of the complex determinant line bundle $\Lambda^2 T(T^*S)$ (second exterior power over \mathbb{C}). The complex trivialization of $T(T^*S)$ is determined uniquely up to homotopy by the condition that it induces a complex trivialization of $\Lambda^2 T(T^*S)$, which is homotopic to the orientation trivialization. We will work throughout with complex trivializations of $T(T^*S)$ that satisfy this condition.

Let $\Lambda \subset T^*S \times \mathbb{R}$ be a Legendrian submanifold and let γ be a loop in Λ . Then the tangent planes of $\Pi(\Lambda)$ along $\Pi(\gamma)$ constitute a loop of Lagrangian subspaces. Using the trivialization of $T(T^*S)$ we get a loop of Lagrangian subspaces in \mathbb{C}^2 ; associating to each loop γ the Maslov index of this loop of Lagrangian subspaces gives a cohomology class $\mu \in H^1(\Lambda; \mathbb{Z})$ called the *Maslov class* of Λ . We assume throughout that the Maslov class of Λ vanishes so that the Maslov index of any such loop equals 0; this holds in the situation considered by this paper by Lemma 3.1. (Note that the vanishing Maslov index implies that Λ is orientable.) In order to orient the moduli spaces of holomorphic disks with boundary on Λ we will further assume that Λ is spin and equipped with a spin structure. See Ekholm, Etnyre and Sullivan [16] for more details. For future use, we note that the Maslov index of the loop of Lagrangian tangent spaces as described can be computed by first making the loop generic with respect to fibers of T^*S and then counting (with signs) the number of instances where the tangent space has a 1–dimensional intersection with the tangent space of some fiber. In terms of the front projection, once γ is generic (in particular, transverse to the cusp edges and disjoint from the swallow tail points), one counts the number of times the curve transversely intersects cusp edges going down (that is, with the \mathbb{R} coordinate of $S \times \mathbb{R}$ decreasing) minus the number of times it transversely intersects cusp edges going up.

2.3.2 The Legendrian algebra In the remainder of the present subsection, Section 2.3, we describe the differential graded algebra $(LA(\Lambda), \partial)$, which we call the *Legendrian DGA*, associated to a Legendrian submanifold $\Lambda \subset J^1(S)$, whose homology is

the *Legendrian homology* of Λ . This description is divided into three parts: the algebra, the grading and the differential.

The algebra $LA(\Lambda)$ is simple to describe; in particular, the Legendrian algebra is simpler in our setting than for general contact manifolds since $J^1(S)$ has no closed Reeb orbits. Assume that Λ is in general position so that $\Pi|_\Lambda$ is a Lagrangian immersion with transverse self-intersections. Let \mathcal{Q} denote the set of Reeb chords of Λ . Then $LA(\Lambda)$ is the non-commutative unital algebra over \mathbb{Z} generated by

- elements of \mathcal{Q} (Reeb chords) and
- $\mathbb{Z}[H_1(\Lambda)]$ (homology classes).

Thus a typical generator of $LA(\Lambda)$ viewed as a \mathbb{Z} -module is a monomial of the form

$$\gamma_0 q_1 \gamma_1 q_2 \cdots q_m \gamma_m,$$

where $q_j \in \mathcal{Q}$ and $\gamma_j \in H_1(\Lambda)$, and multiplication of generators is the obvious multiplication (with $H_1(\Lambda)$ viewed as a multiplicative group). Note that homology classes do not commute with Reeb chords; $LA(\Lambda)$ is more precisely defined as the tensor algebra over \mathbb{Z} generated by elements of \mathcal{Q} and elements of $H_1(\Lambda)$, modulo the relations given by the relations in $H_1(\Lambda)$.

To simplify notation, for the remainder of the paper we will assume that Λ is a disjoint union of oriented 2-tori $\Lambda_1, \dots, \Lambda_r$. We will further assume that each component Λ_j is equipped with a fixed symplectic basis (λ_j, μ_j) of $H_1(\Lambda_j)$ (for conormal lifts, these correspond to the longitude and meridian of the link component in \mathbb{R}^3). Then $\mathbb{Z}[H_1(\Lambda)] = \mathbb{Z}[\lambda_1^{\pm 1}, \mu_1^{\pm 1}, \dots, \lambda_r^{\pm 1}, \mu_r^{\pm 1}]$.

Remark 2.2 In the subject of Legendrian homology, it is often customary to quotient by commutators between Reeb chords and homology classes to obtain the *homology-commutative algebra*. This quotient can also be described as the tensor algebra over $\mathbb{Z}[H_1(\Lambda)]$ freely generated by elements of \mathcal{Q} . The homology-commutative algebra is the version of the Legendrian algebra considered in many sources, in particular the combinatorial formulation of knot contact homology in [29] and the transverse version in [12; 30].

From the geometric viewpoint of the present paper, there is no reason to pass to the homology-commutative quotient, and we will adhere to the rule that homology classes do not commute with Reeb chords in the Legendrian algebra. There are indications that the fully non-commutative Legendrian DGA may be a stronger invariant than the homology-commutative quotient; see the last paragraph of Section 1.2.

Remark 2.3 When Λ has more than one component, there is an algebra related to the Legendrian algebra $LA(\Lambda)$, called the *composable algebra*, which is sometimes a more useful object to consider than $LA(\Lambda)$. This was introduced in Bourgeois, Ekholm and Eliashberg [6; 5]; see also Mishachev [23] and Ng [25] for precursors to this notion. We do not need the composable algebra in this paper, but we briefly describe its definition here and note that one can certainly modify our definition of knot contact homology to the composable setting; that is, the Legendrian DGA for knot contact homology descends to a differential on the composable algebra.

Suppose Λ has r components $\Lambda_1, \dots, \Lambda_r$. Let \mathcal{R} denote the ring

$$\mathcal{R} = \bigoplus_{j=1}^r \mathbb{Z}[H_1(\Lambda_j)]$$

with multiplication given as follows: if $\gamma_1 \in \mathbb{Z}[H_1(\Lambda_{j_1})]$ and $\gamma_2 \in \mathbb{Z}[H_1(\Lambda_{j_2})]$, then $\gamma_1 \cdot \gamma_2$ is 0 if $j_1 \neq j_2$, or $\gamma_1 \gamma_2 \in \mathbb{Z}[H_1(\Lambda_{j_1})] \subset \mathcal{R}$ if $j_1 = j_2$. (Note that \mathcal{R} is nearly but not quite a quotient of $\mathbb{Z}[H_1(\Lambda)]$: the identity in $\mathbb{Z}[H_1(\Lambda_j)]$ is an idempotent in \mathcal{R} distinct from 1.) Let Q_{ij} denote the set of Reeb chords beginning on Λ_i and ending on Λ_j . A *composable monomial* in $LA(\Lambda)$ is a monomial of the form

$$\gamma_0 q_1 \gamma_1 q_2 \cdots q_l \gamma_l$$

for some $l \geq 0$, such that there exist $i_0, \dots, i_l \in \{1, \dots, m\}$ for which $\gamma_{i_j} \in H_1(\Lambda_{i_j})$, viewed as an element of \mathcal{R} , and $q_j \in Q_{i_{j-1}i_j}$ for all j .

The *composable algebra* is the \mathbb{Z} -module generated by composable monomials, with multiplication given in the obvious way; note that the product of two composable monomials is either 0 or another composable monomial. This algebra, which is naturally the path algebra of a quiver with vertices given by components of Λ and edges given by Reeb chords, is almost but not quite a quotient of $LA(\Lambda)$: if we replace $\mathbb{Z}[H_1(\Lambda)]$ by \mathcal{R} in the definition of $LA(\Lambda)$, then the composable algebra is the quotient setting non-composable monomials to 0.

2.3.3 Grading in the Legendrian algebra In order to define the grading on $LA(\Lambda)$ we fix a point $p_j \in \Lambda_j$ on each component Λ_j of Λ , and for each Reeb chord endpoint in Λ_j we choose an *endpoint path* connecting the endpoint to p_j . Furthermore, for $j = 1, \dots, r$, we choose paths γ_{1j} in T^*S connecting $\Pi(p_1)$ to $\Pi(p_j)$ and symplectic trivializations of $\gamma_{1j}^* T(T^*S)$ in which the tangent space $\Pi(T_{p_1}\Lambda)$ corresponds to the tangent space $\Pi(T_{p_j}\Lambda)$; for $j = 1$, γ_{11} is the trivial path. For any $i, j \in \{1, \dots, r\}$, we can then define γ_{ij} to be the path $-\gamma_{1i} \cup \gamma_{1j}$ joining $\Pi(p_i)$ to $\Pi(p_j)$, and $\gamma_{ij}^* T(T^*S)$ inherits a symplectic trivialization from the trivializations for γ_{1i} and γ_{1j} .

The grading $|\cdot|$ in $LA(\Lambda)$ is now the following. First, homology variables have degree 0:

$$|\lambda_j| = |\mu_j| = 0 \quad \text{for } j = 1, \dots, r.$$

(Recall that the Maslov class μ of Λ is assumed to vanish; in general a homology variable $\tau \in H_1(\Lambda)$ would be graded by $-\mu(\tau)$.) Second, if q is a Reeb chord we define the grading by associating a path of Lagrangian subspaces to q . We need to consider two cases according to whether the endpoints of the chord lie on the same component of Λ or not. In the case of equal components, both equal to Λ_j , consider the path of tangent planes along the endpoint path from the final point of q to p_j followed by the reverse endpoint path from p_j to the initial point of q . The endpoint of this path is the tangent space of $\Pi(\Lambda)$ at initial point of q . We close this path of Lagrangian subspaces to a loop $\hat{\gamma}_q$ by a “positive rotation” along the complex angle between the endpoints of the path (see [15] for details). In the case of different components, we associate a loop of Lagrangian subspaces $\hat{\gamma}_q$ to q in the same way except that we insert the path of Lagrangian subspaces induced by the trivialization along the chosen path, γ_{ij} , connecting the components of the endpoints in order to connect the two paths from Reeb chord endpoints to chosen points. The grading of q is then

$$|q| = \mu(\hat{\gamma}_q) - 1,$$

where μ denotes the Maslov index.

Remark 2.4 Note that the grading of a pure Reeb chord (a chord whose start and endpoint lie on the same component of Λ) is well-defined (ie, independent of choice of paths to the base point and symplectic trivializations) because the Maslov class vanishes. The grading of mixed chords however depends on the choice of trivializations along the paths γ_{ij} connecting base points. Changing the trivializations changes the gradings as follows: for some fixed $(n_1, \dots, n_r) \in \mathbb{Z}^r$, for all i, j , $|q|$ is replaced by $|q| + n_i - n_j$ for all Reeb chords q beginning on component i and ending on component j . If we choose orientations on each component of Λ , and stipulate that the base points p_j are chosen such that $\pi: \Lambda \rightarrow S$ is an orientation-preserving local diffeomorphism at each p_j (and that the symplectic trivializations on γ_{ij} preserve orientation), then the mod 2 grading of the mixed chords is well-defined, independent of the choice of base points.

For computational purposes we mention that the grading can be computed in terms of the front projection $\Pi_F: J^1(S) \rightarrow S \times \mathbb{R}$; compare [14]. By Lemma 2.1, a Reeb chord q corresponds to a critical point x of the difference of two local defining functions f_1 and f_2 for $\Lambda \subset J^1(S)$. We make the following assumptions, which hold generically: Λ is in general position with respect to the front projection, the critical point is non-degenerate with index denoted by $\text{index}_x(f_1 - f_2)$, where f_1 defines the upper local

sheet (the sheet with the larger z -coordinate) of Λ , and the base points p_j do not lie on cusp edges.

Lemma 2.5 *There is a choice of trivializations for which the following grading formula holds for all Reeb chords. Let q be a Reeb chord with final point (respectively initial point) in component Λ_j (respectively Λ_i). Let γ be the union of the endpoint path from the chord's final point to the base point p_j , and the reverse endpoint path from base point p_i to the chord's initial point. Assume γ is in general position with respect to the stratified caustic $\Sigma \subset \Lambda$. Then*

$$(2-2) \quad |q| = D(\gamma) - U(\gamma) + \text{index}_x(f_1 - f_2) - 1,$$

where $D(\gamma)$ (respectively $U(\gamma)$) is the number of cusps that γ traverses in the downward (respectively upward) z -direction.

2.3.4 Differential in the Legendrian algebra In general, the Legendrian algebra differential of a Legendrian submanifold Λ in a contact manifold Y is defined using moduli spaces of holomorphic curves in the symplectization $\mathbb{R} \times Y$ of Y with Lagrangian boundary condition $\mathbb{R} \times \Lambda$. In our case, $Y = J^1(S)$, one can instead use holomorphic disks in T^*S with boundary on $\Pi(\Lambda)$. We give a brief description. See Ekholm, Etnyre and Sullivan [17] for details and [14] for the relation to curves in the symplectization.

The differential $\partial: LA(\Lambda) \rightarrow LA(\Lambda)$ is defined on generators and then extended by linearity over \mathbb{Z} and the signed Leibniz rule,

$$\partial(vw) = (\partial v)w + (-1)^{|v|}v(\partial w).$$

We set $\partial\lambda_j = \partial\mu_j = 0$ for $j = 1, \dots, m$. It thus remains to define the differential on Reeb chords.

To do this, we begin by fixing an almost complex structure J on $T(T^*S)$ that is compatible with $d\theta$. Let q_0, q_1, \dots, q_k be Reeb chord generators of $LA(\Lambda)$. Let D_{k+1} be the unit disk in \mathbb{C} with $k + 1$ boundary punctures z_0, \dots, z_k listed in counterclockwise order. We consider maps

$$u: (D_{k+1}, \partial D_{k+1} - \{z_0, \dots, z_k\}) \rightarrow (T^*S, \Pi(\Lambda))$$

such that $u|_{\partial D_{k+1} - \{z_0, \dots, z_m\}}$ lifts to a continuous map \tilde{u} into $\Lambda \subset T^*S$. Call a puncture z mapping to the double point q *positive* (respectively *negative*) if the lift of the arc just clockwise of z in ∂D_{k+1} is a path in Λ approaching the upper Reeb chord endpoint q^+ (respectively the lower endpoint q^-) and the arc just counterclockwise of z lifts to

a path approaching q^- (respectively q^+). For a $(k + 1)$ -tuple of homology classes $\bar{A} = (A_0, \dots, A_k)$, $A_j \in H_1(\Lambda; \mathbb{Z})$, $j = 0, \dots, k$, we let

$$\mathcal{M}_{\bar{A}}(q_0; q_1, \dots, q_k)$$

denote the moduli space of J -holomorphic maps

$$u: (D_{k+1}, \partial D_{k+1} - \{z_0, \dots, z_k\}) \rightarrow (T^*S, \Pi(\Lambda))$$

with the following properties: $u|_{\partial D_{k+1} - \{z_0, \dots, z_k\}}$ lifts to a continuous map \tilde{u} into Λ , z_0 is a positive puncture mapping to q_0 , z_j is a negative puncture mapping to q_j , $j = 1, \dots, k$, and when $\tilde{u}|_{(z_j, z_{j+1})}$ is completed to a loop using the endpoint paths then it represents the homology class A_j . Here (z_j, z_{j+1}) denotes the boundary interval in ∂D_{k+1} between z_j and z_{j+1} and we use the convention $z_{k+1} = z_0$. Note that the moduli space is obtained from the space of maps by dividing by the action of the 3-dimensional group of conformal automorphisms of the unit disk D .

For a generic Λ and J the following holds (see Ekholm, Etnyre and Sabloff [13], and Ekholm, Etnyre and Sullivan [15] for details):

- $\mathcal{M}_{\bar{A}}(q_0; q_1, \dots, q_k)$ is a manifold of dimension (recall that the Maslov class of Λ is assumed to vanish)

$$\dim(\mathcal{M}_{\bar{A}}(q_0; q_1, \dots, q_k)) = |q_0| - \sum_{i=1}^k |q_i| - 1,$$

which is transversely cut out by the $\bar{\partial}_J$ -operator. Furthermore it admits a compactification as a manifold with boundary with corners in which the boundary consists of broken disks. Consequently, if the dimension equals 0 then the manifold is compact.

- The moduli spaces $\mathcal{M}_{\bar{A}}(q_0; q_1, \dots, q_k)$ determined by Λ and J can be “coherently” oriented; see Section 6. (Note that the assumption that the components of Λ are tori and thus admit spin structures is used here. The moduli space orientations depend on the choice of spin structure on Λ .)

The differential of a Reeb chord generator q_0 is then defined as follows:

$$\partial q_0 = \sum_{\substack{\bar{A}=(A_0, \dots, A_k), \\ \sum_{j=1}^k |q_j|=|q_0|-1}} |\mathcal{M}_{\bar{A}}(q_0; q_1 \dots q_k)| A_0 q_1 A_1 q_2 A_2 \dots A_{k-1} q_k A_k,$$

where $|\mathcal{M}|$ denotes the algebraic number of points in the oriented compact 0-manifold \mathcal{M} .

2.3.5 Invariance of Legendrian homology The main properties of $LA(\Lambda)$ are summarized in the following theorem.

Theorem 2.6 (Ekholm, Etnyre and Sullivan [17]) *The map $\partial: LA(\Lambda) \rightarrow LA(\Lambda)$ is a differential, that is, $\partial^2 = 0$. The stable tame isomorphism class of $LA(\Lambda)$ is an invariant of Λ up to Legendrian isotopy; in particular, its homology $LH(\Lambda)$ is a Legendrian isotopy invariant. (For the notion of stable tame isomorphism in this setting, see [14].)*

This result is stated and proven in [17] for the homology-commutative quotient, but the proof there extends verbatim to the full non-commutative algebra.

2.4 Flow trees

Consider a Legendrian submanifold $\Lambda \subset J^1(S) = T^*S \times \mathbb{R}$ as in Section 2.2. There is a Morse-theoretic description of the differential in the Legendrian DGA of Λ via flow trees, as developed in [10], which we will describe in this subsection (in less generality than [10]). The motivation is as follows. For $\sigma > 0$, the map

$$\phi_\sigma: T^*S \times \mathbb{R} \rightarrow T^*S \times \mathbb{R}, \quad (q, p, z) \mapsto (q, \sigma p, \sigma z),$$

satisfies $\phi_\sigma^*(dz - \theta) = \sigma(dz - \theta)$. Hence $\Lambda_\sigma = \phi_\sigma(\Lambda)$ is a Legendrian submanifold that is Legendrian isotopic to Λ . For $\sigma > 0$ small enough there are regular almost complex structures for which there is a one-to-one correspondence between rigid holomorphic disks with boundary on Λ_σ with one positive puncture and rigid flow trees determined by Λ ; see [10].

We now define flow trees. Fix a metric g on S . Then two local defining functions f_0 and f_1 for Λ defined on the same open set in S define a local vector field on S :

$$-\nabla(f_1 - f_0),$$

where ∇ denotes the g -gradient. The 1-jet lift of a path $\gamma: (-\epsilon, \epsilon) \rightarrow S$ is a pair of continuous paths $\gamma_i: (-\epsilon, \epsilon) \rightarrow \Lambda$, $i = 0, 1$ with the following properties: $\pi \circ \gamma_i = \gamma$ and either $\gamma_0(t) \neq \gamma_1(t)$ or $\gamma_0(t) = \gamma_1(t)$ is a point in Σ . A path $\gamma: (a, b) \rightarrow S$ is called a *flow line of Λ* if it has a 1-jet lift γ_i , $i = 0, 1$, such that for each $t \in (a, b)$ there are local defining functions f_0, f_1 defined near $\gamma(t)$ such that γ_i lies in the sheet determined by f_i , $i = 0, 1$, and

$$\dot{\gamma}(t) = -\nabla(f_1 - f_0)(\gamma(t)).$$

See Figure 2. If γ is a flow line with 1-jet lift γ_0, γ_1 we define its *cotangent lift* as $\Pi \circ \gamma_0, \Pi \circ \gamma_1$.




Figure 2: Graphs Γ_{f_i} of the local defining functions f_i for two sheets of $\Pi_F(\Lambda)$. A flow line γ and critical point s in S and their lifts to $\Pi_F(\Lambda)$. Notice s is a positive puncture.

Let γ be a flow line of Λ with 1-jet lift γ_0, γ_1 in sheets with local defining functions f_0, f_1 . Then the *flow orientation* of γ_0 (respectively γ_1) is given by the lift of the vector

$$-\nabla(f_0 - f_1) \quad (\text{respectively } -\nabla(f_1 - f_0)).$$

If $\gamma: (-\infty, b) \rightarrow S$ (respectively $\gamma: (a, \infty) \rightarrow S$) is a flow line as above such that

$$\lim_{t \rightarrow -\infty} \gamma(t) = s \in S \quad (\text{respectively } \lim_{t \rightarrow \infty} \gamma(t) = s)$$

is a critical point of $f_1 - f_0$, then we say s is a *puncture* of the flow line γ . Let s_i be the point in the sheet of f_i with $\pi(s_i) = s$, $i = 0, 1$. Choose notation so that $f_1(s_1) > f_0(s_0)$; then s is a

- $\left\{ \begin{array}{l} \text{positive puncture} \\ \text{negative puncture} \end{array} \right.$ if the flow orientation of γ_0 points toward s_0 and that of γ_1 points away from s_1 ,
- if the flow orientation of γ_1 points toward s_1 and that of γ_0 points away from s_0 .

If s is a puncture of a flow tree then the *chord at s* is the vertical line segment oriented in the direction of increasing z that connects s_0 and s_1 .

A *flow tree* of $\Lambda \subset J^1(S)$ is a map into S with domain a finite tree Γ , which may have semi-infinite edges, with extra structure consisting of a cyclic ordering of the edges at each vertex and with the following three properties.

- (1) The restriction of the map to each edge is a flow line of Λ .
- (2) If v is a k -valent vertex with cyclically ordered edges e_1, \dots, e_k and \bar{e}_j^0, \bar{e}_j^1 is the cotangent lift of e_j , $1 \leq j \leq k$, then there exists a pairing of lift components such that for every $1 \leq j \leq k$ (with $k + 1 = 1$)

$$\bar{e}_j^1(v) = \bar{e}_{j+1}^0(v) = p \in \Pi(\Lambda) \subset T^*S,$$

and such that the flow orientation of \bar{e}_j^1 at p is directed toward p if and only if the flow orientation of \bar{e}_{j+1}^0 at p is directed away from p . Thus the cotangent lifts of the edges of Γ then fit together as an oriented curve $\bar{\Gamma}$ in $\Pi(\Lambda)$.

- (3) This curve $\bar{\Gamma}$ is closed.

We first notice that vertices may contain punctures. We will be interested in flow trees with only one positive puncture. Such flow trees can have only one puncture above each vertex; see [10, Section 2]. Thus for such flow trees Γ we divide the vertices into three sets: the set of positive punctures $P(\Gamma)$, the set of negative punctures $N(\Gamma)$ and the set of other vertices $R(\Gamma)$. Recall at a (non-degenerate) puncture v the corresponding difference between the local defining functions has a non-degenerate critical point. Denote its index by $I(v)$.

If Γ is a flow tree as above then its *formal dimension* (see [10, Definition 3.4]), which measures the dimension of the space of flow trees with 1-jet lift near the 1-jet lift of Γ , is

$$(2-3) \quad \dim(\Gamma) = \left(\sum_{v \in P(\Gamma)} (I(v) - 1) - \sum_{v \in N(\Gamma)} (I(v) - 1) + \sum_{v \in R(\Gamma)} \mu(v) \right) - 1,$$

where $\mu(v)$ is the *Maslov content* of v and is defined as follows. For a vertex $v \in R(\Gamma)$ let $x \in \pi^{-1}(v)$ be a cusp point that lies in the 1-jet lift of Γ (if such a point exists). If γ_0 and γ_1 are two 1-jet lifts of an edge of Γ adjacent to v that contain x and for which the flow orientation of γ_0 is pointed towards x and that of γ_1 is pointed away from x , then we set $\tilde{\mu}(x) = +1$ (respectively -1) if γ_0 is on the upper (respectively lower) local sheet of Λ near x and γ_1 is on the lower (respectively upper) local sheet. Otherwise define $\tilde{\mu}(x) = 0$. We can now define $\mu(v) = \sum \tilde{\mu}(x)$, where the sum is taken over all $x \in \pi^{-1}(v)$ that are cusp points in the 1-jet lift of Γ .

There is also a notion of *geometric dimension* of Γ , $\text{gdim}(\Gamma)$ (see [10, Definition 3.5]), which measures the dimension of the space of flow trees near Γ that have the exact same

geometric properties as Γ . In [10, Lemma 3.7] it is shown that $\text{gdim}(\Gamma) \leq \dim(\Gamma)$ for any flow tree Γ , and a characterization of the vertices of trees for which equality holds is given. In combination with transversality arguments for Morse flows, this leads to the following result; see [10, Lemma 3.7 and Proposition 3.14].

Theorem 2.7 (Ekholm 2007 [10]) *Given a number $n > 0$, after a small perturbation of Λ and the metric g on S we may assume that for any flow tree Γ having one positive puncture and $\dim(\Gamma) \leq n$, the space of flow trees with the same geometric properties in a neighborhood of Γ is a transversely cut out manifold of dimension $\text{gdim}(\Gamma)$. In particular, trees Γ with $\dim(\Gamma) = 0$ form a transversely cut out 0-manifold. Furthermore, such rigid trees satisfy $\text{gdim}(\Gamma) = \dim(\Gamma) = 0$ and have vertices only of the following types; see Figure 3.*

- (1) Valency one vertices that
 - (a) are positive punctures with Morse index $\neq 0$,
 - (b) are negative punctures with Morse index $\neq 2$, or
 - (c) lift to a cusp edge and have Maslov content $+1$, called a cusp end
- (2) Valency two vertices that
 - (a) are positive punctures with Morse index 0,
 - (b) are negative punctures with Morse index 2, or
 - (c) have order one tangencies with a cusp edge and Maslov content -1 , called a switch
- (3) Valency three vertices that
 - (a) are disjoint from the projection $\pi(\Sigma)$ of the singular locus to S , called a Y_0 vertex, or
 - (b) lie on the image of the cusp locus and have Maslov content -1 , called a Y_1 vertex

Notice that we may endow a flow tree Γ that has exactly one positive puncture with a *flow orientation*: any edge e is oriented by the negative gradient of the positive difference of its defining functions.

When working with flow trees it will also be useful to consider the symplectic area of a flow tree Γ . Given a flow tree Γ , let $\hat{\Gamma}$ denote its 1-jet lift (which, in previous notation, projects to $\bar{\Gamma} \subset \Pi(\Lambda)$), and define the *symplectic area* of Γ to be

$$A(\Gamma) = - \int_{\hat{\Gamma}} p \, dq = - \int_{\hat{\Gamma}} dz.$$

The name comes from the connection between flow trees and holomorphic curves. The important features of the symplectic area are summarized in the following lemma.

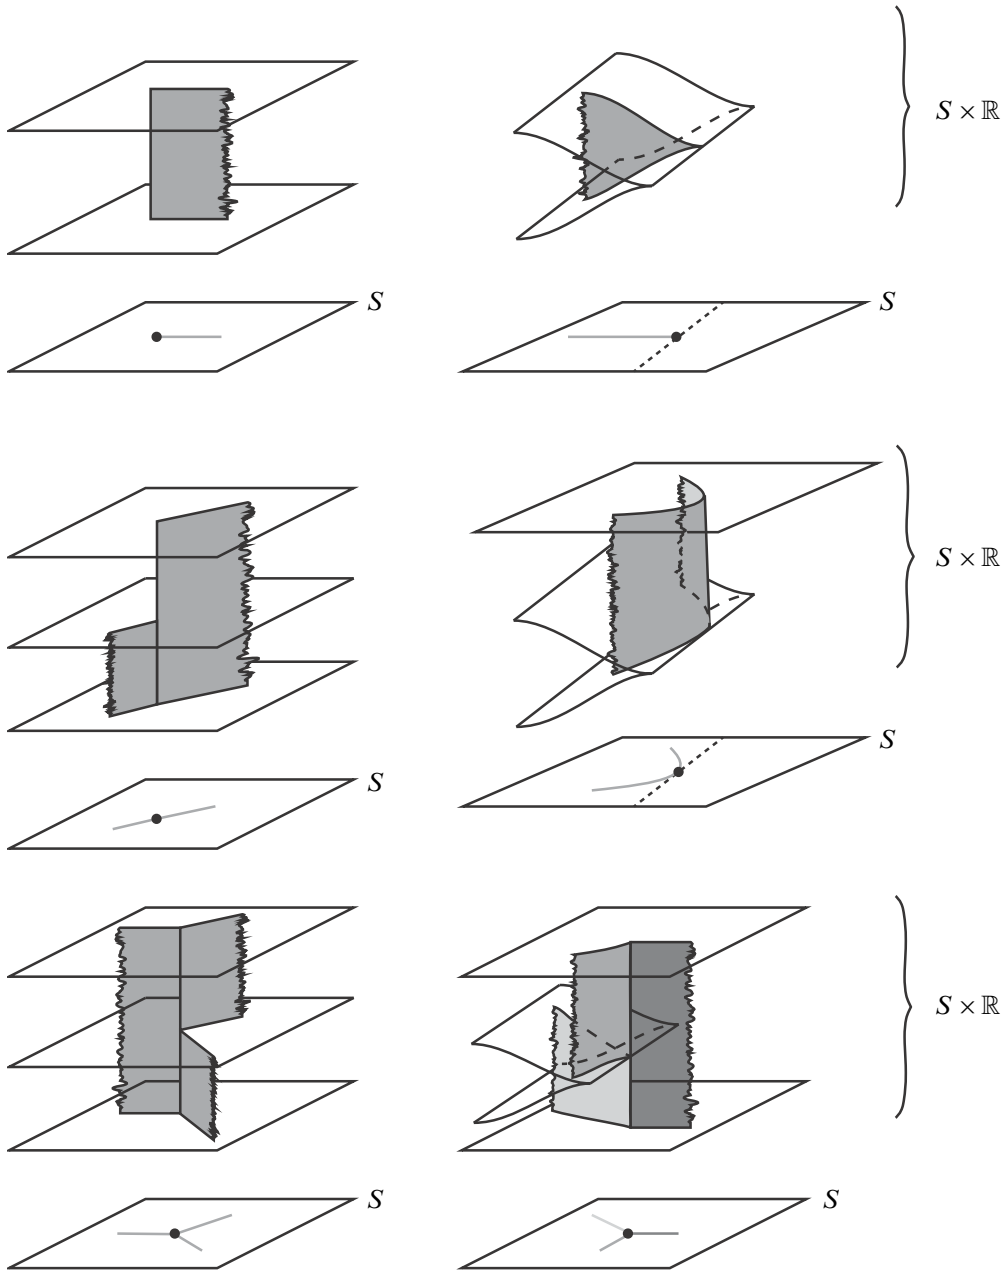


Figure 3: Lifts of neighborhoods, in Γ , of vertices into the front space $S \times \mathbb{R}$. Dashed lines in S are cusp edges.

Before stating it we introduce some notation. A puncture a of a flow tree lifts to a double point in the Lagrangian projection $\Pi(\Lambda)$ and hence corresponds to a Reeb chord. Thus in the 1-jet lift of Γ there will be two points that project to a . We denote them a^+ and a^- where a^+ has the larger z -coordinate.

Lemma 2.8 (Ekholm 2007 [10]) *For any flow tree Γ the symplectic action is positive: $A(\Gamma) > 0$. The symplectic action can be computed by the formula*

$$A(\Gamma) = \sum_p (z(p^+) - z(p^-)) - \sum_q (z(q^+) - z(q^-)),$$

where the first sum is over positive punctures of Γ , the second sum is over negative punctures of Γ and $z(a)$ denotes the z -coordinate of the point a .

For our applications two further types of flow trees will be needed. First, a *partial flow tree* is a flow tree Γ for which we drop the condition that the cotangent lift $\bar{\Gamma}$ is closed and allow Γ to have 1-valent vertices v such that $\bar{\Gamma}$ intersects the fiber over v in two points. We call such vertices *special punctures*. In the dimension formula (2-3), $I(v) = 3$ for a positive special puncture. Theorem 2.7 holds as well for partial flow trees with at least one special puncture; see [10]. Second, in Section 5.1 we consider *constrained flow trees*: if p_1, \dots, p_r are points in Λ then a flow tree constrained by p_1, \dots, p_r is a flow tree Γ with 1-jet lift which passes through the points p_1, \dots, p_r .

2.5 Vector splitting along flow trees

In this subsection we describe a combinatorial algorithm for transporting normal vectors in a flow tree to all its vertices which will combinatorially determine the sign of a flow tree; see Section 3.4. Specifically we will be concerned only with flow trees that do not involve cusp edges and only have punctures at critical points of index 1 and 2.

Suppose Λ is a Legendrian submanifold in $J^1(S)$ that does not have cusp edges (see Remark 1.6) and only has Reeb chords corresponding to critical points of index 1 and 2. Let Γ be a partial flow tree with positive special puncture determined by $\Lambda \subset J^1(S)$.

Consider the local situation at a Y_0 -vertex of Γ at $t \in S$. (Note that because of the lack of cusp edges, Γ has no Y_1 -vertices.) In the flow orientation of Γ one edge adjacent to t is pointing toward it (we call this edge *incoming*) and the other two edges are pointing away from it (we call them *outgoing*). Furthermore, the natural orientation of the 1-jet lift of the tree induces a cyclic order on the three edges adjacent to t and thus an order of the two outgoing edges. (Specifically, the edge e_1 will have the same upper sheet as the incoming edge while the edge e_2 will have the same lower sheet.) If v_0 denotes the negative gradient of the incoming edge at t and v_1, v_2 the negative

gradients of the two outgoing edges, all pointing according to the flow orientation of Γ , then v_1 and v_2 are linearly independent and the following balance equation holds: $v_0 = v_1 + v_2$. (This follows from the fact that the difference between the local defining functions along the incoming edge is the sum of the function differences along the outgoing edges.)

We next define vector splitting. Let p denote the special positive puncture of Γ , let t_1, \dots, t_{k-1} denote its trivalent vertices, and let q_1, \dots, q_k denote its negative punctures. Vector splitting along Γ is a function

$$N_p\Gamma \rightarrow \prod_{j=1}^{k-1} F_{t_j}S \times \prod_{i=1}^k N_{q_i}\Gamma,$$

where $N_x\Gamma$ denotes the normal bundle of the edge of Γ containing the point $x \in \Gamma$ (which is assumed not to be a trivalent vertex) and where FS denotes the frame bundle of S . It is defined as follows.

Translating $n \in N_x\Gamma$ as a normal vector along the edge in the direction of the flow orientation of the tree, it eventually arrives at a trivalent vertex t . At this vertex, the translated vector $n(t)$ is perpendicular to the incoming edge at t and determines two unique vectors $w_1(t)$ and $w_2(t)$ in T_tS perpendicular to the first and second outgoing edges at t , respectively, by requiring that $n(t) = w_1(t) + w_2(t)$; see Figure 4. The frame at t is $(w_1(t), w_2(t))$.




Figure 4: Vector splitting at a trivalent Y_0 -vertex

Regarding $w_1(t)$ and $w_2(t)$ as normal vectors of their respective edges, the above construction can be applied with n replaced by $w_j(t)$, $j = 1, 2$, and Γ replaced by the partial tree Γ_j that is obtained by cutting Γ in the j^{th} outgoing edge at t and taking the component that does not contain t . When a subtree Γ_j contains no trivalent vertex we translate $w_j(t)$ along its respective edge to the negative puncture at the vertex. Continuing in this way until the cut-off partial trees do not have any trivalent vertices, we get:

- two vectors $w_1(t)$ and $w_2(t)$ perpendicular to the first and second outgoing edge at t , for each trivalent vertex t , and
- a vector $w(q)$ perpendicular to the edge ending at q , for each negative puncture q .

The vector splitting of n along Γ is

$$n \mapsto ((w_1(t_1), w_2(t_1)), \dots, (w_1(t_{k-1}), w_2(t_{k-1}))); w(q_1), \dots, w(q_k)).$$

2.6 Algebraic results

In this subsection, we collect for convenience some algebraic results about the map ϕ and the matrices Φ^L and Φ^R , which were defined in Section 1.2 and play an essential role in the combinatorial formulation of the Legendrian algebra for knot contact homology, Theorem 1.1. These results, which will occasionally be needed in the remainder of the paper, have essentially been established in previous work of the third author [26; 27; 29; 30]. However, we repeat the caveat from Section 1.2 that in our present context, homology classes (μ and λ) do not commute with Reeb chords (a_{ij}), while the previous papers deal with the homology-commutative quotient (see Section 2.3.2). Nevertheless, all existing proofs extend in an obvious way to the present setting.

Proposition 2.9 *The map $\phi: \tilde{\mathcal{A}}_n^0 \rightarrow \tilde{\mathcal{A}}_n^0$, defined in Section 1.2 for braid generators σ_k of the braid B_n , respects the relations in the B_n and thus extends to a homomorphism from B_n to $\text{Aut } \tilde{\mathcal{A}}_n^0$.*

Proof The proof is by direct computation; cf [26, Proposition 2.5]. □

Proposition 2.10 *Let $B \in B_n$, and let $\phi_B(\mathbb{A})$ be the $n \times n$ matrix defined by $(\phi_B(\mathbb{A}))_{ij} = \phi_B(\mathbb{A}_{ij})$. Then we have the matrix identity:*

$$\phi_B(\mathbb{A}) = \Phi_B^L \cdot \mathbb{A} \cdot \Phi_B^R$$

Proof The proof is by induction on the length of the braid word representing B ; cf [26, Proposition 4.7; 30, Lemma 2.8]. The latter reference proves the result stated here ($\hat{\mathbb{A}}$ and $\check{\mathbb{A}}$ there correspond to \mathbb{A} here once we set $U = V = 1$), but in the homology-commutative quotient, for the case of a single-component knot, and with slightly different sign conventions. Nevertheless, the inductive proof given there works here as well; we omit the details. □

We remark that Proposition 2.10 can be given a more natural, geometric proof via the language of “cords” [27]. This approach also provides an explanation for the precise placement of the homology classes $\tilde{\mu}$ in the definition of ϕ_{σ_k} from the introduction. We refer the interested reader to [29, Section 3.2], which treats the homology-commutative single-component case, and leave the straightforward extension to the general case to the reader.

One consequence of Proposition 2.10 is that the differential for knot contact homology presented in Theorem 1.1 is well-defined. More precisely, the differential for the b_{ij} generators is given in matrix form by:

$$\partial\mathbb{B} = -\lambda^{-1} \cdot \mathbb{A} \cdot \lambda + \Phi_B^L \cdot \mathbb{A} \cdot \Phi_B^R$$

However, since \mathbb{B} has 0’s along the main diagonal, it is necessary that the right hand side has 0’s along the diagonal as well. This is indeed the case: the (i, i) entry of $\lambda^{-1} \cdot \mathbb{A} \cdot \lambda$ is $1 + \mu_{\alpha(i)}$, while the (i, i) entry of $\Phi_B^L \cdot \mathbb{A} \cdot \Phi_B^R$ is $\phi_B(1 + \mu_{\alpha(i)}) = 1 + \mu_{\alpha(i)}$.

3 The differential and flow trees

In Section 4, we will prove Theorem 1.1 by using multiscale flow trees to compute the differential of Λ_K in $J^1(S^2) = U^*\mathbb{R}^3$. These multiscale flow trees combine two types of flow trees, which are the focus of this section. Specifically we will see that if Λ is the conormal lift of the unknot U then by thinking of K as a braid about U we can isotope Λ_K into an arbitrarily small neighborhood of $\Lambda = \Lambda_U$, which can be identified with a neighborhood of the zero section in $J^1(\Lambda)$. Thus we may think of Λ_K as a subset of $J^1(\Lambda)$.

To use multiscale flow trees to compute the differential of Λ_K in $J^1(S^2)$ we will combine flow trees of $\Lambda \subset J^1(S^2)$ and $\Lambda_K \subset J^1(\Lambda)$. The content of this section is a computation of these flow trees; in Section 4, we then combine the flow trees to complete the computation of the Legendrian DGA for $\Lambda_K \subset J^1(S^2)$.

Here is a more detailed summary of this section. In Section 3.1, we discuss the Legendrian torus Λ and describe a generic front projection for it. In Section 3.2, we compute the rigid flow trees for Λ as well as the 1-parameter families of flow trees. In Section 3.3, we give an explicit identification of a neighborhood of the zero section in $J^1(\Lambda)$ with a neighborhood of Λ in $J^1(S^2)$, use this identification to explicitly describe Λ_K in $J^1(\Lambda)$, and find all the Reeb chords of $\Lambda_K \subset J^1(\Lambda)$. Section 3.4 computes the rigid flow trees of Λ_K in $J^1(\Lambda)$ modulo some technical considerations concerning “twist regions” that are handled in Section 3.5. We comment that Section 3.4 produces an invariant of braids (cf [26]) using only existing flow tree technology and not multiscale flow trees.

3.1 A generic front for Λ

Let U be the round unknot given by the unit circle in the xy -plane in \mathbb{R}^3 . We give a description of the conormal lift $\Lambda = \Lambda_U$ in $J^1(S^2) = U^*\mathbb{R}^3$ by describing its front projection in $S^2 \times \mathbb{R}$. In the figures below, we draw $S^2 \times \mathbb{R}$ as $\mathbb{R}^3 - \{0\}$ and identify the zero section $S^2 \times \{0\}$ with the unit sphere. If C is the circle in the conormal bundle of U lying over a point $x \in U$ then its image in $S^2 \times \{0\} \subset S^2 \times \mathbb{R}$ is a great circle running through the north and south poles of S^2 . See Figure 5. By the contactomorphism ϕ between $U^*\mathbb{R}^3$ and $J^1(S^2)$ from Section 2.1, the image of C in $S^2 \times \mathbb{R}$ is the graph of $\langle x, y \rangle$, where $y \in C$. This is shown in the leftmost picture in Figure 6.




Figure 5: On the left is the unknot U in the xy -plane with a point p on U labeled and its unit (co-)normal bundle shown. On the right is the tangent space $\mathbb{R}^3 = T_p\mathbb{R}^3$ at p with the unit sphere indicated along with the image of the unit (co-)normal bundle to U at p .




Figure 6: The front of the unknot. On the left is the image of the circle shown in Figure 5 in the front projection in $J^0(S^2) = S^2 \times \mathbb{R} = \mathbb{R}^3 - \{0\}$. The center image shows the entire front projection (with the unit sphere $S^2 \times \{0\}$ shown in light grey) and the image on the right shows the two Reeb chords after perturbation.

By symmetry we get the entire front projection by simply rotating this image of C about the axis through the north and south pole as shown in the middle picture of Figure 6. Using Lemma 2.1 we see that this representative of Λ has an S^1 's worth of Reeb chords over the equator. We perturb Λ using a Morse function on the equator with one maximum and one minimum, so that only two Reeb chords c and e , as indicated in the rightmost picture in Figure 6, remain. Here e and c correspond to the maximum and minimum, respectively, of the perturbing function, and both correspond to transverse double points of $\Pi(\Lambda)$.

This perturbation does not suffice to make the front of Λ generic with respect to fibers of T^*S^2 : over the poles of S^2 we see that Λ consists of Lagrangian cones. We first describe these in local coordinates and then show how to perturb them to become front generic. Let $x = (x_1, x_2)$ be local coordinates near the pole in S^2 and let $(x, y) = (x_1, y_1, x_2, y_2)$ be corresponding Darboux coordinates in T^*S^2 , with symplectic form given by $dx \wedge dy = dx_1 \wedge dy_1 + dx_2 \wedge dy_2$.

Consider $S^1 = \{\xi \in \mathbb{R}^2 : |\xi| = 1\}$. The *Lagrangian cone* is the exact Lagrangian embedding $C: S^1 \times \mathbb{R} \rightarrow \mathbb{R}^4$ given by

$$C(\xi, r) = (r\xi, \xi).$$

As mentioned above, the Lagrangian cone is not front generic: the front projection $\Pi_F \circ C$ is regular for $r \neq 0$ but maps all of $S^1 \times \{0\}$ to the origin. In order to describe perturbations of C we first find a cotangent neighborhood of it. See Figure 7.

Consider $T^*(S^1 \times \mathbb{R}) = T^*S^1 \times T^*\mathbb{R}$. Let (ξ, η, r, ρ) be coordinates on this space where ρ is dual to r and η is dual to ξ . Here we think of η as a (co)vector perpendicular to $\xi \in S^1 \subset \mathbb{R}^2$, that is, $\eta \in \mathbb{R}^2$ and $\eta \cdot \xi = 0$. Consider the map

$$\Phi: T^*S^1 \times T^*\mathbb{R} \rightarrow T^*\mathbb{R}^2 = \mathbb{R}^4$$

given by

$$\Phi(\xi, \eta, r, \rho) = \left(r\xi - \frac{1}{1+\rho} \eta, (1+\rho)\xi \right).$$

Then $\Phi^*(dx \wedge dy) = d\xi \wedge d\eta + dr \wedge d\rho$. To see this we compute

$$\begin{aligned} dx &= (dr)\xi + r d\xi - \frac{1}{1+\rho} d\eta + \frac{d\rho}{(1+\rho)^2} \eta, \\ dy &= (d\rho)\xi + (1+\rho) d\xi, \end{aligned}$$




Figure 7: A perturbed Lagrangian cone. Along the top of the figure is an annular neighborhood of the circle in $\Lambda = \Lambda_U$ that maps to the fiber above the north pole. On the middle left, we see the image of this annulus near the north pole in the front projection, a cone whose boundary is two circles. On the bottom left is the image of this annulus near the north pole in S^2 (that is, the top view of the cone where we have slightly offset the circles so that they are both visible). On the middle right, we see the top view of the cone after it has been perturbed to have a generic front projection. More specifically, the lighter outer curve is the image of the cusp curves, the dotted lines are the image of double points in the front projection and the darkest inner curve is the image of the circle that mapped to the cone point before the perturbation. On the bottom right, we see the image in S^2 of the cusp curve and the two boundary circles on Λ_U .

and hence

$$\begin{aligned}
 dx \wedge dy &= (dr \wedge d\rho)(\xi \cdot \xi) + d\xi \wedge d\eta + ((1 + \rho)dr - rd\rho) \wedge (\xi \cdot d\xi) \\
 &\quad + \frac{d\rho}{1 + \rho} \wedge (\xi \cdot d\eta + \eta \cdot d\xi) \\
 &= dr \wedge d\rho + d\xi \wedge d\eta,
 \end{aligned}$$

since $\xi \cdot \xi = 1$ and $\xi \cdot \eta = 0$.

Thus Φ is a symplectic neighborhood map extending the Lagrangian cone. It follows that exact Lagrangian submanifolds C^1 -near C can be described by $\Phi(\Gamma_{df})$, where f is a smooth function on $S^1 \times \mathbb{R}$ and where Γ_{df} denotes the graph of its differential.

We will consider specific functions of the form

$$f(\xi, r) = \alpha(r)g(\xi),$$

where $\alpha(r)$ is a cut off function equal to 0 for $|r| > 2\delta$ and equal to 1 for $|r| < \delta$ for some small $\delta > 0$. Write $C_f = \Phi(\Gamma_{df})$.

We first give a description of the caustic of C_f . Note that $\Pi_F \circ C$ is an immersion for $|r| > \delta$. The same therefore holds true for C_f provided f is small enough. In order to describe the caustic we thus focus on the region where $|r| < \delta$ and hence $\alpha(r) = 1$. Here

$$C_f(\xi, r) = (r\xi + dg, \xi) = (C_f^1(\xi, r), C_f^2(\xi, r)).$$

The caustic is the set where C_f has tangent lines in common with the fiber. Consequently a point $p \in S^1 \times \mathbb{R}$ belongs to the fiber provided the differential of the first component C_f^1 of the map C_f has rank less than 2. Write $\xi = (\cos \theta, \sin \theta) \in S^1$ and take $g(\xi) = g_1(\xi) = \epsilon \cos 2\theta$. The caustic is then the image of the locus $r = 4\epsilon \cos 2\theta$ under the first component of the map

$$\Phi(\theta, r) = (r(\cos \theta, \sin \theta) - 2\epsilon \sin \theta(-\sin \theta, \cos \theta), (\cos \theta, \sin \theta));$$

see Figure 7.

Lemma 3.1 *The Maslov class of Λ vanishes and consequently the grading of any Reeb chord of Λ is independent of choice of capping path. Let e and c denote the Reeb chords of Λ , as described above; then*

$$|e| = 2 \quad \text{and} \quad |c| = 1.$$

Proof To see that the Maslov class vanishes we need to check that the Maslov index of any generator of $H_1(\Lambda)$ vanishes. We compute the Maslov index of a curve as described at the end of Section 2.3.1. Take one generator as a curve in Λ over the equator; since this curve does not intersect any cusp edge its Maslov index vanishes. Take the other generator as a curve perpendicular to the equator going to the poles and then back; since such a curve has two cusp edge intersections, one up-cusp and one down-cusp, its Maslov index vanishes as well.

Finally choose the capping path of e and c which goes up to the north pole and then back. This capping path has one down-cusp and the Morse indices of e and c are 2 and 1, respectively. The index assertions now follow from Lemma 2.5. \square

3.2 Flow trees of Λ

We next determine all flow trees of Λ .

Lemma 3.2 *There are exactly six rigid flow trees of Λ : four with positive puncture at c (I_N, Y_N, I_S and Y_S) and two with positive puncture at e and negative puncture at c (E_1 and E_2). Furthermore, if p is a point of Λ lying over a point where the front of Λ has 2 sheets then there are exactly two constrained rigid flow trees with positive puncture at e with 1-jet lift passing through p .*

Before proving Lemma 3.2, we make a couple of remarks.

Remark 3.3 In fact one can show the following: There are exactly four 1-parameter families of trees with positive puncture at e : $\tilde{I}_N, \tilde{Y}_N, \tilde{I}_S$ and \tilde{Y}_S . The boundaries of these 1-parameter families are as follows:

$$\begin{aligned} \partial\tilde{I}_N &= (E_1 \# I_N) \cup (E_2 \# I_N), & \partial\tilde{Y}_N &= (E_1 \# Y_N) \cup (E_2 \# Y_N), \\ \partial\tilde{I}_S &= (E_1 \# I_S) \cup (E_2 \# I_S), & \partial\tilde{Y}_S &= (E_1 \# Y_S) \cup (E_2 \# Y_S). \end{aligned}$$

Here $E_1 \# I_N$ denotes the broken tree obtained by adjoining I_N to E_1 , etc. See Figures 8 and 9. Furthermore, the 1-jet lifts of the flow trees in each of these families sweep the part of the torus lying over the corresponding hemisphere (N or S) once.

The formal proof of this result about 1-parameter families would require a more thorough study of flow trees in particular including a description of all vertices of flow trees that appear in generic 1-parameter family. This is fairly straightforward; see [10, Section 7]. For the purposes of this paper it suffices to work with constrained rigid trees rather than 1-parameter families so details about 1-parameter families of flow trees will be omitted.




Figure 8: Rigid flow trees for Λ on the northern hemisphere of S^2

Remark 3.4 We will not need a precise expression of the 1–dimensional families \tilde{I}_N , \tilde{Y}_N , \tilde{I}_S , and \tilde{Y}_S in our computations, but we do need a rough understanding of them. To see the family of disks, start with the symmetric picture of Λ coming from the conormal lift of U . Now make a small perturbation of the north and south poles as shown in Figure 7. Then we see an I and Y flow tree from each point on the equator into the northern hemisphere and another into the southern hemisphere. Now perturbing slightly so the equator is no longer a circle of critical points but contains only the critical points c and e and two flow lines between them, we will see that each of the I and Y disks will become part of one of the 1–dimensional families of disks \tilde{I}_N , \tilde{Y}_N , \tilde{I}_S , and \tilde{Y}_S . See Figure 9.

It is also useful to see these trees as arising from the Bott-degenerate conormal lift of the round unknot. Here there are four holomorphic disks emanating from each Reeb chord. The corresponding trees are just flow lines from the equator to the pole. The 1–jet lift of such a flow line can then be completed by one of the two half circles of the circle in Λ which is the preimage of the pole. The Bott-degenerate 1–parameter family then consists of a flow segment starting at the maximum in the Bott-family and ending at some point where a disk emanating at that point (corresponding to a flow line from that point to a pole) is attached.




Figure 9: One of the 1–dimensional flow tree families for Λ lying over the northern hemisphere of S^2

Proof of Lemma 3.2 We are only considering flow trees with exactly one positive puncture. First consider a rigid flow tree Γ with a positive puncture at c . Using

Lemma 2.8 we see that there can be no negative punctures since the symplectic area must be positive. Thus all vertices of Γ are Y_0 or Y_1 vertices, switches or cusp ends.

We can rule out switches as follows. Since the front is a small perturbation of a Lagrangian cone, it is possible to arrange the following: along a cusp edge in the image $\pi(\Sigma)$ of the caustic, the gradient vector fields for the function difference between a sheet meeting the cusp edge and a sheet not meeting the cusp edge are transverse to the caustic except at the swallow tail points, where they are tangent to the caustic (or zero). This is a non-generic situation, generically tangencies occur only at smooth points of the caustic, and after small perturbation tangency points lie close to the swallowtails. Switches of flow trees lie at tangency points in the caustic and thus a switch could only occur near a swallow tail point. Now, using again that the front is a small perturbation of a Lagrangian cone, we can also arrange that near a swallow tail, the function difference between any one of the three sheets involved in the swallow tail and the fourth sheet is larger than any function difference in the small region bounded by $\pi(\Sigma)$ away from swallow tails. Then since the unique flow line that leaves c hits $\pi(\Sigma)$ away from the swallow tails, by positivity of symplectic area (in this case, the fact that the function difference decreases along a flow line), there is a neighborhood of the swallow tail points that the flow tree cannot reach and in particular it cannot have any switches.

Thus the vertices of Γ are all Y_0 or Y_1 vertices or cusp ends. By the dimension formula (2-3), we see that for the flow tree to be rigid there is some $n \geq 0$ such that there are n type Y_1 vertices and $n + 1$ cusp ends (each with $\mu = +1$). To have a Y_1 vertex a flow line must intersect a cusp edge so that when traveling along the flow the number of sheets used describing Λ increases (as one passes the cusp). Around the north and south pole of S^2 the cusp edges are arranged so that only flow lines traveling towards the poles could possibly have a Y_1 vertex. Thus for $n > 1$ one of the edges in the flow tree will be a flow line that travels from near the north pole to near the south pole. Since this clearly does not exist, as there are two flow lines connecting e to c along the equator and Λ has only two sheets along the equator, we must have $n = 0$ or 1 .

There are only two flow lines leaving c (that is, two flow lines that could have c as a positive puncture), one heading towards the north pole and one heading towards the south pole. When $n = 0$ we clearly get I_N and I_S from these flow lines when they do not split at a Y_1 vertex and when $n = 1$ we get Y_N and Y_S when they do split at a Y_1 vertex. Conversely, this argument shows that I_N , I_S , Y_N , Y_S are indeed rigid flow trees with positive puncture at c (refer to Figure 8 for I_N and Y_N).

Notice for future reference that the above analysis shows that there are no 1-parameter families of flow trees with c as a positive puncture.

Now consider a flow tree Γ with e as a positive puncture. Noting that the Reeb chord above c is only slightly shorter than the chord above e we see, using Lemma 2.8, that Γ can have either no negative punctures or just one negative puncture at c . If $\dim(\Gamma) = 0$, then by the dimension formula there must be a negative puncture at c . Then, since Λ is defined by only two functions away from neighborhoods of the poles, the only vertices of Γ are the punctures e and c . Thus Γ is simply a flow line from e to c and there are precisely two: E_1 and E_2 .

The argument for constrained rigid trees with positive puncture at e follows from the argument used for rigid flow trees with positive puncture at c above. \square

3.3 Conormal lifts of general links

In order to describe the conormal lift of a general link $K \subset \mathbb{R}^3$ we first represent it as the closure of a braid around the unknot U . More precisely, K lies in a small neighborhood $N = S^1 \times D^2 \subset \mathbb{R}^3$ of U and is transverse to the fiber disks $\{\theta\} \times D^2$ for all $\theta \in S^1$. Note that B is a braid on n strands if and only if K intersects any fiber disk n times. We write B for the closed braid corresponding to K considered as lying in $S^1 \times D^2$.

We represent a closed braid B on n strands as the graph Γ_{f_B} of a multi-section $f_B: S^1 \rightarrow D^2$, where $f_B(s)$ consists of n distinct points in D^2 varying smoothly with s , so that Γ_{f_B} is a smooth submanifold. Representing S^1 as $[0, 2\pi]$ with endpoints identified, we can express f_B as a collection of n functions $\{f_1(s), \dots, f_n(s)\}$, where $f_i: [0, 2\pi] \rightarrow D^2$ are smooth functions, $i = 1, \dots, n$. (The sets $\{f_1(0), \dots, f_n(0)\}$ and $\{f_1(2\pi), \dots, f_n(2\pi)\}$ are equal but it does not necessarily hold that $f_i(0) = f_i(2\pi)$.) Note that the distance between Λ and Λ_K is controlled by the C^1 -distance from f_B to the trivial multi-section, which consists of n points at the origin. In particular if N is a fixed neighborhood of Λ then $\Lambda_K \subset N$ provided that f_B is sufficiently C^1 -small.

3.3.1 A 1-jet neighborhood of Λ In order to describe the multiscale flow trees on Λ determined by Λ_K we need to identify some neighborhood of Λ with $J^1(\Lambda) \approx T^*T^2 \times \mathbb{R}$. Thinking of T^2 as $S^1 \times S^1$ we use two different versions of the cotangent bundle of S^1 to describe T^*T^2 . First, we think of S^1 as $[0, 2\pi]$ with endpoints identified, we let s be a coordinate on $[0, 2\pi]$ and $\sigma \in \mathbb{R}$ be a fiber coordinate in $T^*S^1 = S^1 \times \mathbb{R}$. We write $T^*S^1_\lambda$ for the cotangent bundle with these coordinates:

$$T^*S^1_\lambda = \{(s, \sigma) : s \in [0, 2\pi], \sigma \in \mathbb{R}\}.$$

We denote the second version $T^*S^1_\mu$ and define it as

$$T^*S^1_\mu = \{(\xi, \eta) \in \mathbb{R}^2 \times \mathbb{R}^2 : |\xi| = 1, \xi \cdot \eta = 0\}.$$

Let $r(s) = (\cos s, \sin s, 0)$, $0 \leq s \leq 2\pi$ denote a unit vector in the x_1x_2 -plane at an angle s from the x_1 -axis and $\theta(s) = (-\sin s, \cos s, 0)$ the vector $r(s)$ rotated $\pi/2$ counter-clockwise in the x_1x_2 -plane (we can think of it as the standard angular vector at $r(s)$ translated back to the origin). Let ζ be the coordinate on \mathbb{R} . Consider the map

$$\Phi: T^*S_\lambda^1 \times T^*S_\mu^1 \times \mathbb{R} \rightarrow U^*\mathbb{R}^3 = \mathbb{R}^3 \times S^2,$$

where $\Phi = (\Phi_1, \Phi_2)$ is defined by

$$(3-1) \quad \begin{aligned} \Phi_1(s, \sigma, \xi, \eta, \zeta) &= r(s) + \frac{1}{\sqrt{1-\sigma^2}} (\eta_1 r(s) + \eta_2(0, 0, 1)) \\ &\quad + \zeta (-\sigma \theta(s) + \sqrt{1-\sigma^2} (\xi_1 r(s) + \xi_2(0, 0, 1))), \\ \Phi_2(s, \sigma, \xi, \eta, \zeta) &= -\sigma \theta(s) + \sqrt{1-\sigma^2} (\xi_1 r(s) + \xi_2(0, 0, 1)). \end{aligned}$$

Then $\Phi|_{S_\lambda^1 \times 0 \times S_\mu^1 \times 0 \times 0}$ is a parametrization of Λ and its restriction to a small neighborhood of the 0-section is an embedding. Furthermore, since $\xi \cdot \eta = 0$, we know that $\eta \cdot d\xi = -\xi \cdot d\eta$. Using this and the fact that $r(s) \cdot \theta(s) = 0$ and $r'(s) = \theta(s)$ we can compute

$$\Phi^*(p dq) = d\zeta - \sigma ds - \xi \cdot d\eta - \frac{1}{\sqrt{1-\sigma^2}} \sigma \eta_1 ds.$$

We introduce the following notation:

$$(3-2) \quad \beta_0 = \sigma ds + \xi \cdot d\eta,$$

$$(3-3) \quad \alpha = \frac{1}{\sqrt{1-\sigma^2}} \sigma \eta_1 ds,$$

$$(3-4) \quad \beta_t = \beta_0 - t\alpha, \quad 0 \leq t \leq 1.$$

Note that $d\beta_t$ is symplectic in a neighborhood of the 0-section, for all t . Using Moser's trick we define a time-dependent vector field X_t by

$$(3-5) \quad -\alpha = d\beta_t(X_t, \cdot)$$

and find that if ψ_t denotes the time t flow of X_t then

$$\psi_t^* d\beta_t = d\beta_0.$$

In particular,

$$d(\psi_t^* \beta_t - \beta_0) = 0.$$

Equation (3-5) and the definition of α imply that $X_t = 0$ along the 0-section and thus $\psi_t^* \beta_t - \beta_0 = 0$ along the 0-section. By the homotopy invariance of de Rham cohomology, the closed form $\psi_t^* \beta_t - \beta_0$ is exact. Let the function h be such that $\beta_0 = \psi_1^* \beta_1 + dh$ and such that $h = 0$ on the 0-section.

We can bound the growth of α, X_t, h in terms of distance $r := (|\eta|^2 + |\sigma|^2)^{1/2}$ from the origin. From the explicit expression for α in (3-2), we have $\alpha = \mathcal{O}(r^2)$ and thus from (3-5),

$$(3-6) \quad |X_t| = \mathcal{O}(|\eta|^2 + |\sigma|^2) \quad \text{and} \quad |dX_t| = \mathcal{O}((|\eta|^2 + |\sigma|^2)^{1/2}).$$

Then from the definition of $h, |dh| = \mathcal{O}(r^2)$ and so

$$(3-7) \quad |d^{(k)}(h)| = \mathcal{O}((|\eta|^2 + |\sigma|^2)^{(3-k)/2}), \quad k = 0, 1, 2.$$

We will use these estimates in the proof of Lemma 3.6 below.

Finally, if Ψ is the diffeomorphism of $T^*S_\lambda^1 \times T^*S_\mu^1 \times \mathbb{R}$ given by

$$\Psi((s, \sigma, \xi, \eta), \zeta) = (\psi_1(s, \sigma, \xi, \eta), \zeta + h(s, \sigma, \xi, \eta)),$$

then

$$\Psi^*\Phi^*(p dq) = d\zeta - \sigma ds - \xi \cdot d\eta.$$

The map $\Theta = \Phi \circ \Psi$ is the 1-jet neighborhood map we will use. It is an embedding from a neighborhood of the 0-section in $J^1(T^2) = T^*T^2 \times \mathbb{R}$ to a neighborhood of $\Lambda \subset U^*\mathbb{R}^3 = J^1(S^2)$ such that $\Theta^*(p dq) = d\zeta - \theta$, where ζ is a coordinate in the \mathbb{R} -factor and where θ is the Liouville form on T^*T^2 .

Remark 3.5 Notice that the contactomorphism Θ sends the Reeb flow of $J^1(T^2)$ to the Reeb flow of $J^1(S^2)$. (Here, as throughout the rest of the paper, we are identifying $J^1(S^2)$ with $U^*\mathbb{R}^3$ using the contactomorphism in (2-1).) Thus any Reeb chord of Λ_K in $J^1(T^2)$ corresponds to a Reeb chord of Λ_K in $J^1(S^2)$ and any Reeb chord of Λ_K in $J^1(S^2)$ that lies entirely in N corresponds to a Reeb chord of Λ_K in $J^1(T^2)$.

3.3.2 Conormal lifts of closed braids as multi-sections Consider $T^2 = S_\lambda^1 \times S_\mu^1$ as above. A multi-section of $J^0(T^2)$ is a smooth map $F: T^2 \rightarrow J^0(T^2)$ such that $\pi \circ F$ is an immersion (ie, a covering map). In particular, a multi-section can be thought of as the graph of a multi-function $F: T^2 \rightarrow \mathbb{R}$. (Context will designate whether F refers to a (multi)function or its graph, a (multi)section.) The 1-jet extension of a generic multi-section is a Legendrian submanifold. We denote it $\Gamma_{j^1(F)}$. Let $K \subset \mathbb{R}^3$ be a link and let B be a closed braid representing K with corresponding multi-section $f_B = \{f_1, \dots, f_n\}, f_j: [0, 2\pi] \rightarrow D^2$, as described at the beginning of Section 3.3.

Lemma 3.6 *The conormal lift Λ_K is Legendrian isotopic to the Legendrian submanifold $\Theta(\Gamma_{j^1(F_B)})$, where F_B is the multi-section given by the functions*

$$F_j: [0, 2\pi] \times S_\mu^1 \rightarrow \mathbb{R}, \quad j = 1, \dots, n, \quad \text{where } F_j(s, \xi) = f_j(s) \cdot \xi.$$

Here we think of S_λ^1 as $[0, 2\pi]$ with endpoints identified and we identify S_μ^1 with the unit circle in the plane of the disk where $f_j: [0, 2\pi] \rightarrow D^2$ takes values.

Proof Let $N \subset U^*\mathbb{R}^3$ denote a δ -neighborhood of Λ in which Θ and Φ give local coordinates. Take the C^2 -norm of f_B sufficiently small so that Λ_K will be in N . We first show that $\Phi^{-1}(\Lambda_K)$ is given by the 1-jet lift of F_B . To see this notice that the image of the braid B is given by the image of the maps

$$(f_j(s))_x r(s) + (f_j(s))_y (0, 0, 1),$$

where $(f_j)_x$ and $(f_j)_y$ are the x and y -coordinates of f_j , and $r(s)$ is as in (3-1). Thus the normal component in $J^1(S^2) = T^*S^2 \times \mathbb{R}$ (ie, the \mathbb{R} -component) of the braid at the point $(f_j(s))_x r(s) + (f_j(s))_y (0, 0, 1)$ is given by

$$((f_j(s))_x r(s) + (f_j(s))_y (0, 0, 1)) \cdot (\xi_1 r(s) + \xi_2 (0, 0, 1)) = f_j(s) \cdot \xi.$$

According to the definition of Φ we see that the \mathbb{R} -factor of $T^*T^2 \times \mathbb{R} = J^1(T^2)$ maps to the \mathbb{R} -factor in $T^*S^2 \times \mathbb{R}$ by $\zeta \mapsto \zeta \sqrt{1 - \sigma^2}$. Thus the multi-section of $J^0(T^2)$ corresponding to $\Phi^{-1}(\Lambda_K)$ is given by $(1/\sqrt{1 - \sigma^2}) f_j(s) \cdot \xi$, but in $J^0(T^2)$ the σ -coordinate is always equal to zero. Thus the multi-function F_B does indeed describe $\Phi^{-1}(\Lambda_K)$ as claimed.

Now $\Theta^{-1}(\Lambda_K)$ is the 1-jet graph $\Gamma_{j^1(G_B)}$ of some multi-section G_B . In general, Legendrian submanifolds of $J^1(T^2)$ will be given by cusped multi-sections, but since each point in $\Theta^{-1}(\Lambda_K)$ has a neighborhood in $\Theta^{-1}(\Lambda_K)$ that can be made C^1 -close to the zero section in $J^1(T^2)$, we see that $\Theta^{-1}(\Lambda_K)$ has empty caustic and hence is the 1-jet extension of a multi-section.

From the above discussion we see that $\Gamma_{j^1(G_B)}$ is the same as $\Psi^{-1}(\Gamma_{j^1(F_B)})$. The estimates (3-6) and (3-7) then imply that the C^1 -distance between F_B and G_B is $\mathcal{O}(\delta^2)$. Consequently, for $\delta > 0$ sufficiently small, $\Gamma_{j^1(G_B)} = \Lambda_K$ and $\Gamma_{j^1(F_B)}$ are Legendrian isotopic. \square

3.3.3 Reeb chords and grading Let $K \subset \mathbb{R}^3$ be a link. We assume that K is braided around the unknot as a braid on n strands and we represent Λ_K as $\Gamma_{j^1(F_B)} \subset J^1(T^2)$ as in Lemma 3.6. Then the Reeb chords of Λ_K in $J^1(S^2)$ are of two types: *short chords*, which are entirely contained in the neighborhood N of Λ , and *long chords*, which are not. According to Remark 3.5 we see that the short chords of Λ_K correspond to chords of Λ_K in $J^1(T^2)$. As with Λ , one can use the techniques of Section 2.3.1 to conclude that the Maslov class of Λ_K vanishes, and thus the grading of any Reeb chord of Λ_K with both endpoints on the same component is independent of capping path. For Reeb chords with endpoints on distinct components the grading depends on

the chosen transport along the paths connecting the components; cf Remark 2.4. Here we use the following choice throughout. Fix a point $x \in \Lambda$ and let the base points of the components of Λ_K all lie in the intersection $J_x^1(\Lambda) \cap \Lambda_K$ of Λ_K and the fiber of $J^1(\Lambda)$ at x . Take the connecting paths as straight lines in $T_{\Pi^*\Lambda(x)}^*\Lambda$ and use parallel translation in the flat metric followed by a rotation along the complex angle or the complementary complex angle, whichever is less than $\frac{\pi}{2}$, as transport. Note that as K gets closer to U the angle of this rotation approaches 0.

Lemma 3.7 *Up to smooth isotopy, we can choose the link $K \subset \mathbb{R}^3$ so that Λ_K has exactly $2n(n - 1)$ short Reeb chords,*

$$\begin{aligned} \{a_{ij}\}_{1 \leq i, j \leq n, i \neq j}, \quad |a_{ij}| = 0, \\ \{b_{ij}\}_{1 \leq i, j \leq n, i \neq j}, \quad |b_{ij}| = 1, \end{aligned}$$

and exactly $2n^2$ long Reeb chords,

$$\begin{aligned} \{c_{ij}\}_{1 \leq i, j \leq n}, \quad |c_{ij}| = 1, \\ \{e_{ij}\}_{1 \leq i, j \leq n}, \quad |e_{ij}| = 2. \end{aligned}$$

Here c_{ij} (respectively e_{ij}) lie in small neighborhoods of the Reeb chords c (respectively e) of Λ for all i, j . Furthermore all Reeb chords can be taken to correspond to transverse intersection points in T^*S^2 .

Remark 3.8 We will use the following notation for the Reeb chords of Λ_K . We say that the short Reeb chords are of type **S** and the long of type **L**. We sometimes specify further and say that a short (long) Reeb chord of grading j is of type **S_j** (**L_j**).

In a neighborhood of the short chords, the Legendrian consists of the conormal lifts of n distinct arcs. The notation for short Reeb chords in Lemma 3.7 is chosen so that a_{ij} begins on the conormal lift of the i^{th} strand and ends on that of the j^{th} strand and similarly for b_{ij} .

Proof of Lemma 3.7 The statement in the lemma for long chords is immediate from the fact that Λ_K is the 1-jet graph of a multi-section with n sheets over Λ which has two Reeb chords: e with $|e| = 2$ and c with $|c| = 1$. To prove the statement on short chords we note that we may choose the multi-section f_B so that for any $i \neq j$, $|f_i - f_j|$ has a maximum at $2\pi - \delta \in [0, 2\pi]$, a minimum at $2\pi - 2\delta$, and no other critical points. Parameterizing S_μ^1 by $\xi = (\cos t, \sin t)$, $t \in [0, 2\pi]$, we find that the difference between two local functions of Λ_K is

$$F_{ij}(s, t) = (f_i(s) - f_j(s)) \cdot (\cos t, \sin t).$$

Now $dF_{ij} = 0$ if and only if

$$(3-8) \quad (f'_i(s) - f'_j(s)) \cdot (\cos t, \sin t) = 0 \quad \text{and}$$

$$(3-9) \quad (f_i(s) - f_j(s)) \cdot (-\sin t, \cos t) = 0,$$

which in turn happen if and only if s is critical for $|f_i(s) - f_j(s)|$ and t takes one of the two values, say t_0 and $t_0 + \pi$, for which (3-9) holds. We take a_{ij} to be the chord corresponding to the minimal distance between strands and b_{ij} to the maximal distance.

In order to compute the gradings of a_{ij} and b_{ij} we note that the front of Λ_K in $J^0(\Lambda) = \Lambda \times \mathbb{R}$ has no singularities and that the chords a_{ij} (respectively b_{ij}) correspond to saddle points (respectively maxima) of positive function differences of local defining functions. The grading statement then follows from (2-2). \square

3.4 Counting flow trees of Λ_K in $J^1(\Lambda)$

In this subsection we determine all the flow trees for $\Lambda_K \subset J^1(\Lambda)$. Our enumeration relies on a particular and fairly technical choice of position for Λ_K over the regions where the braid twists, whose details we defer to Section 3.5. For our current computational purposes, we only need a few qualitative features of these twist regions, as described in Sections 3.4.1 through 3.4.3 below, which will serve to motivate the more technical parts of our discussion of twist regions in Section 3.5.

Given these qualitative features, we perform the actual combinatorial computation of flow trees in Sections 3.4.5 through 3.4.8 (after first presenting a scheme for calculating signs for flow trees in Section 3.4.4), culminating in Lemma 3.21, which presents a purely algebraic formula for the Legendrian DGA of $\Lambda_K \subset J^1(\Lambda)$. This comprises an important subalgebra of the Legendrian DGA of $\Lambda_K \subset J^1(S^2)$, the rest of which is computed in Section 4.

3.4.1 Basic setup Recall that K is the closure of a braid in $S^1 \times D^2$ given by a collection $\{f_1, \dots, f_n\}$ of functions $f_j: [0, 2\pi] \rightarrow D^2$, $j = 1, \dots, n$. We use Lemma 3.6 to represent $\Lambda_K \subset J^1(\Lambda)$ as the 1-jet graph of the functions $F_j: [0, 2\pi]^2 \rightarrow \mathbb{R}$ given by

$$F_j(s, t) = f_j(s) \cdot (\cos t, \sin t).$$

Here $(s, t) \in [0, 2\pi]^2$ are coordinates on Λ , with s corresponding to the parameter along the unknot (represented by the unit circle in the x_1x_2 -plane) and t to the parameter along a unit circle in the normal fiber of the unknot with 0 corresponding to the positive outward normal of the unit circle in the x_1x_2 -plane. Furthermore, recall from

Lemma 3.7 that Reeb chords of $\Lambda_K \subset J^1(\Lambda)$ correspond to points $(s, t) \in \Lambda$ where s is a critical point of $|f_i(s) - f_j(s)|$ and where t is such that the vector $(\cos t, \sin t)$ is parallel to the vector $f_i(s) - f_j(s)$.

Let $[s_0, s_1] \subset [0, 2\pi]$. We refer to the part of Λ_K lying over an interval $[s_0, s_1]$,

$$\Lambda_K \cap J^1([s_0, s_1] \times S^1) \subset J^1(\Lambda),$$

as the $[s_0, s_1]$ -slice of Λ_K . We will represent the braid as follows: the actual twists will take place in an $[s_0^{br}, s_1^{br}]$ -slice, where $[s_0^{br}, s_1^{br}] \subset (0, \frac{\pi}{2})$. Inside the $[s_0^{br}, s_1^{br}]$ -slice the braid is given by sub-slices where it twists so that two strands are interchanged, separated by slices where the braid is trivial. Outside the $[s_0^{br}, s_1^{br}]$ -slice the braid is trivial.

We will choose perturbations so that the following holds: Reeb chord endpoints of the Reeb chords e and c of Λ have the following coordinates:

$$e^+ = (0, 0), \quad e^- = (\pi, \pi), \quad c^+ = (\pi, 0) \quad \text{and} \quad c^- = (0, \pi);$$

see Lemma 3.1. Consequently, all Reeb chords e_{ij} and c_{ij} , $1 \leq i, j \leq n$, of type **L** are located near these points. The Reeb chords a_{ij} and b_{ij} of type **S** are located near $s = \underline{s}$ and $s = \bar{s}$, where $\frac{\pi}{2} < \underline{s} < \bar{s} < 2\pi$; see the proof of Lemma 3.7.

3.4.2 Reeb chords and trivial slices Consider an interval $[s_0, s_1]$ where the braid is trivial. We will describe a model for the $[s_0, s_1]$ -slice of the conormal lift of a trivial braid which we will use in two ways: to control Reeb chords and to define a normal form of a slice of the trivial braid in which there are no Reeb chords. We first describe a somewhat degenerate $[s_0, s_1]$ -slice of Λ_K : we represent the trivial braid by

$$f_j(s) = (0, j \psi(s)), \quad j = 1, \dots, n,$$

where $\psi(s)$ is a positive function that has a non-degenerate local minimum at \underline{s} and a non-degenerate local maximum at \bar{s} . Here $s_0 < \underline{s} < \bar{s} < s_1$ and ψ has no other critical points.

We call the 1-jet graph of the function

$$F_j(s, t) = f_j(s) \cdot (\cos t, \sin t)$$

the j^{th} sheet of Λ_K and denote it by S_j , $j = 1, \dots, n$. Writing

$$\begin{aligned} F_{ij}(s, t) &= F_i(s, t) - F_j(s, t) = (f_i(s) - f_j(s)) \cdot (\cos t, \sin t) \\ &= (i - j)\psi(s) \sin t, \end{aligned}$$

Reeb chords correspond to critical points of F_{ij} and are all located in the fibers over the following points in Λ :

$$\left(\underline{s}, \frac{\pi}{2}\right), \left(\underline{s}, \frac{3\pi}{2}\right), \left(\bar{s}, \frac{\pi}{2}\right), \text{ and } \left(\bar{s}, \frac{3\pi}{2}\right).$$

The Reeb chords lying in the fibers over the first (respectively second) two points correspond to saddle points (respectively extrema) of the functions F_{ij} . We denote them by a_{ij} (respectively b_{ij}). Here the labeling is such that a_{ij} and b_{ij} begin on S_j and end on S_i . Thus, if $i > j$ then the t -coordinate of a_{ij} and b_{ij} equals $\frac{\pi}{2}$ whereas if $i < j$ it equals $\frac{3\pi}{2}$.

The flow trees that we study all have their negative punctures at a_{ij} and thus we must understand the unstable manifold $W^u(a_{ij})$ of a_{ij} as a critical point of F_{ij} . (Recall that we are using the positive gradient flow when discussing stable and unstable manifolds.) It is straightforward to check that

$$W^u(a_{ij}) = \begin{cases} \{(s, t) : t = \frac{\pi}{2}\} & \text{if } i > j, \\ \{(s, t) : t = \frac{3\pi}{2}\} & \text{if } i < j. \end{cases}$$

Note that the $[s_0, s_1]$ -slice of Λ_K as defined above is degenerate: Reeb chords are not disjoint, the unstable manifolds $W^u(a_{ij})$ are not mutually transverse, and their stable counterparts are not mutually transverse either. The following lemma describes the $[s_0, s_1]$ -slice of Λ_K after a small perturbation which makes it generic. In particular, the perturbation is so small that there is a natural one-to-one correspondence between Reeb chords before and after perturbation and we will keep the notation a_{ij} and b_{ij} from above. There are of course many perturbations which make Λ_K generic. The particular choice studied here is designed to make counting flow trees as simple as possible.

Lemma 3.9 *For $\epsilon_0 > 0$ arbitrarily small, there exist a Legendrian isotopy of Λ_K and a collection of functions $\epsilon_{ij}, \epsilon'_{ij}, \epsilon''_{ij} : [s_0, s_1] \rightarrow (0, \epsilon_0)$ for each $n \geq i > j \geq 1$, so that the following conditions hold:*

- ϵ_{ij} are lexicographically ordered in (i, j) : for any $s, s', \epsilon_{ij}(s) > \epsilon_{i'j'}(s')$ if $i > i'$, or if $i = i'$ and $j > j'$.
- Similarly, ϵ'_{ij} and ϵ''_{ij} are lexicographically ordered in (i, j) .
- The unstable manifolds $W^u(a_{ij})$ are curves of the following form:
 - if $i > j$ then $W^u(a_{ij}) = \{(s, t) : t = \frac{\pi}{2} - \epsilon_0 + \epsilon_{ij}(s)\}$ and
 - if $i < j$ then $W^u(a_{ij}) = \{(s, t) : t = \frac{3\pi}{2} + \epsilon'_{ji}(s)\}$.
- The Reeb chords corresponding to critical points in these unstable manifolds satisfy the following:

- if $i > j$ then the s -coordinate of a_{ij} (respectively b_{ij}) equals $\underline{s} + \epsilon''_{ij}(\underline{s})$ (respectively $\bar{s} + \epsilon''_{ij}(\bar{s})$),
- if $i < j$ then the s -coordinate of a_{ij} (respectively b_{ij}) equals $\underline{s} + \epsilon''_{ji}(\underline{s})$ (respectively $\bar{s} + \epsilon''_{ji}(\bar{s})$).

Remark 3.10 The functions $\epsilon_{ij}, \epsilon'_{ij}, \epsilon''_{ij}$ are chosen so that the unstable manifolds $W^u(a_{ij}), i > j$, appear in the lexicographical order on $\{(i, j)\}_{1 \leq j < i \leq n}$ if read in the increasing t -direction, so that the $W^u(a_{ij}), i < j$, appear in the lexicographical order of $\{(j, i)\}_{1 \leq i < j \leq n}$ if read in the positive t -direction, and so that no $W^u(a_{ij})$ lies in the region $\frac{\pi}{2} \leq t \leq \frac{3\pi}{2}$; see Figure 10.




Figure 10: Ordered unstable manifolds after perturbation

Proof of Lemma 3.9 Throughout the proof we use the notational conventions above for Reeb chords. We will consider two different perturbations of the braid representative for the trivial braid given at the beginning of this subsection and then combine them to give the desired perturbation.

We first choose the scaling function $\psi(s)$ from the beginning of this subsection to additionally satisfy $\psi(s) = \frac{M}{2}(s - \underline{s})^2 + c$ in a $2\epsilon_0$ neighborhood of \underline{s} and $\psi(s) = -\frac{M}{2}(s - \bar{s})^2 + c'$ in a $2\epsilon_0$ neighborhood of \bar{s} , where c and c' are constants, M is some large constant and ϵ_0 is some small positive constant.

Now choose constants $\delta_1, \dots, \delta_n$ so that

$$0 \leq \delta_1 < \delta_2 < 2\delta_2 < \delta_3 < 3\delta_3 < \dots < \delta_n < n\delta_n < \epsilon_0.$$

For $i > j$, define $\epsilon''_{ij} = (i\delta_i - j\delta_j)/(i - j)$. The above inequalities imply that the ϵ''_{ij} are positive, less than ϵ_0 and lexicographically ordered in (i, j) : indeed, we have

$$\epsilon''_{ij} \in \left(\frac{i}{i-j+1}\delta_i, \frac{i}{i-j}\delta_i \right) \subset (\delta_i, i\delta_i).$$

Set $\psi_i(s) = \psi(s - \delta_i)$. Defining the trivial braid by the multi-function determined by $f_j(s) = (0, j\psi_j(s))$, we see that the Reeb chord a_{ij} has s -coordinate determined by the solution near \underline{s} to

$$iM(s - \delta_i - \underline{s}) = i\psi'_i(s) = j\psi'_j(s) = jM(s - \delta_j - \underline{s}),$$

which is precisely $s = \underline{s} + \epsilon''_{ij}$, and similarly for b_{ij} and \bar{s} . Thus the Reeb chords a_{ij} are at

$$\begin{cases} (\underline{s}, \frac{\pi}{2}) + (\epsilon''_{ij}, 0) & \text{if } i > j, \\ (\underline{s}, \frac{3\pi}{2}) + (\epsilon''_{ji}, 0) & \text{if } i < j, \end{cases}$$

and the Reeb chords b_{ij} are at

$$\begin{cases} (\bar{s}, \frac{\pi}{2}) + (\epsilon''_{ij}, 0) & \text{if } i > j, \\ (\bar{s}, \frac{3\pi}{2}) + (\epsilon''_{ji}, 0) & \text{if } i < j. \end{cases}$$

Now consider a different representation of the trivial braid. In particular, returning to the original representation of the trivial braid given at the beginning of this subsection, we can replace the curve $\{x = 0\}$ where the functions $f_i(s)$ take values, with a curve family $\{x = h_t(y)\}$ where $h_t: \mathbb{R} \rightarrow \mathbb{R}$ is a smoothly varying family of functions such that:

- h_t is constant in t near $t = \frac{\pi}{2}$ and $t = \frac{3\pi}{2}$,
- $h_{\pi/2}(j) = j(\delta_n - \delta_j)$ for $j = 1, \dots, n$,
- $h_{3\pi/2}(j) = -j\delta_j$ for $j = 1, \dots, n$,

where the δ_j are as before. (We can assume that the t dependence of h_t is supported in an arbitrarily small neighborhood of π and 0 . While this is not necessary here, it will be important when describing the lift of a non-trivial braid in Section 3.5.) Then let

$$F_j(s, t) = \psi(s)(h_t(j), j) \cdot (\cos t, \sin t)$$

be the function of the j^{th} sheet. The critical points for $F_{ij} = F_i - F_j$ are at $s = \underline{s}$ or $s = \bar{s}$, with t given by

$$\cot(t) = \frac{h_t(i) - h_t(j)}{i - j + (\partial h / \partial t)(i) - (\partial h / \partial t)(j)}.$$

Since h_t is small (and can be made arbitrarily small by the appropriate choice of δ_i), this equation can only hold for t near $\pi/2$ or $3\pi/2$, whence

$$\cot(t) = \frac{h_{\pi/2}(i) - h_{\pi/2}(j)}{i - j} \quad \text{or} \quad \cot(t) = \frac{h_{3\pi/2}(i) - h_{3\pi/2}(j)}{i - j},$$

respectively.

In the former case, the expression for $h_{\pi/2}$ implies that the Reeb chords a_{ij}, b_{ij} for $i > j$ have t -coordinate near $\pi/2$ and are given by the solution to $\cot(t) = \delta_n - \epsilon''_{ij} > 0$. If $t = t_{ij}$ is the solution to this equation, then $t_{ij} < \pi/2$ and t_{ij} is ordered in lexicographic order on (i, j) , and so we can write $t_{ij} = \pi/2 - \epsilon_0 + \epsilon_{ij}$ with ϵ_{ij} ordered lexicographically in (i, j) . Similarly, the expression for $h_{3\pi/2}$ implies that the Reeb chords a_{ij}, b_{ij} for $i < j$ have t -coordinate near $3\pi/2$ and are given by the solution to $\cot(t) = -\epsilon''_{ji} < 0$, which yields $t = 3\pi/2 + \epsilon'_{ji}$ with ϵ'_{ji} ordered lexicographically on (i, j) . Furthermore, the unstable manifolds $W^u(a_{ij})$ and $W^u(a_{ji})$ are horizontal (constant in t) since the ∂_t component of ∇F_{ij} is zero at $t = t_{ij}$.

Combining the perturbations, we find that the location and ordering of the critical points and unstable manifolds of

$$F_j(s, t) = \psi_j(s)(h_t(j), j) \cdot (\cos t, \sin t)$$

is as desired. □

Remark 3.11 Recall that our notation for (un)stable manifolds refers to the positive gradient flow of positive function differences and note that the unstable manifolds $W^u(a_{ij})$ can be characterized as the only flow line determined by sheets S_i and S_j along which the local function difference stays positive for all time under the negative gradient flow: along any other (non-constant) flow lines of the negative gradient, the local function difference eventually becomes negative.

Remark 3.12 Consider next an $[s', s'']$ -slice where the braid is trivial, eg, the slices mentioned above that separate the slices where twists of the braid occur. In each such slice we will take the braid to look much like in the $[s_0, s_0 + \delta]$ -slice of the braid in Lemma 3.9, where $\delta > 0$ is small enough so that the s -coordinate of any Reeb chord a_{ij} is larger than $s_0 + \delta$. More precisely, we require that the functions $f_j(s)$, $j = 1, \dots, n$ take values in a family of graphical curves $\{x = h_t(y)\}$ (with $y = j$ for f_j), where h_t is a small function which is independent of t near $\frac{\pi}{2}$ and $\frac{3\pi}{2}$. In order to make sure that the $[s', s'']$ -slice does not have any Reeb chords we let the points $f_j(s)$ move away from each other along the curves $\{x = h_t(y)\}$ as we go through $[s', s'']$ from right to left, so that $|f_i(s) - f_j(s)|$ decreases with s for all i, j . We call

a braid with these properties a *standard trivial braid* and we number the functions f_1, \dots, f_n according to the order in which they appear along the curve $\{x = h_t(y)\}$ with orientation induced from the positively oriented y -axis.

3.4.3 A model for Λ_K , endpoint paths, and homology indicators Here we describe the qualitative features of the gradient flows associated to Λ_K in $J^1(\Lambda)$ that will be necessary for our computation of the rigid flow trees. Some of these features have been discussed in Sections 3.4.1 through 3.4.3. For the actual construction of gradient flows with these properties, see Section 3.5.

For the braid $B = \sigma_{k_1}^{\epsilon_1} \cdots \sigma_{k_m}^{\epsilon_m}$, $\epsilon_l = \pm 1$, we can build a braiding slice $[s_0^{\text{br}}, s_1^{\text{br}}]$ where B lives, by starting with a twist slice corresponding to $\sigma_{i_m}^{\epsilon_m}$ and then attaching slices corresponding to the other twists $\sigma_{i_l}^{\epsilon_l}$ consecutively, working backwards in l . We then extend B to the complement of the braiding slice by closing it with a trivial braid as in Lemma 3.9. This is the model of Λ_K that we will use below. More precisely, we will construct the braid model to have the following properties:

- A slice $[s', 2\pi]$ of a trivial braid contains all of the Reeb chords of $\Lambda_B \subset J^1(\Lambda)$, as in Lemma 3.9. The Reeb chords a_{ij} (respectively b_{ij}) of Λ_K lie near $s = 2\pi - 4\delta$ (respectively $s = 2\pi - 2\delta$), which is in the slice that is complementary to the braiding region. The manifolds $W^u(a_{ij})$ in this slice lie just below $t = \frac{\pi}{2}$ and are ordered in the t -direction according to the lexicographical order on $\{(i, j)\}_{1 \leq j < i \leq n}$. The manifolds $W^u(a_{ji})$ lie just above $t = \frac{3\pi}{2}$ and are ordered in the t -direction according to the lexicographical order on $\{(i, j)\}_{1 \leq j < i \leq n}$.
- All twists of the braid occur in the braiding region $2\delta < s < 4\delta$, where Λ_K looks like twist slices separated by standard trivial braid slices; see Section 3.5.
- There are points $s_1 = 2\delta < s_2 < \dots < s_{2m} = 4\delta$ so that the $[s_{2l-1}, s_{2l}]$ -slice contains the twist region for $\sigma_{i_l}^{\epsilon_l}$. For a slice corresponding to the braid crossing $\sigma_{i_l}^{\epsilon_l} = \sigma_k^{\pm 1}$, the unstable manifolds for a_{kk+1} and a_{k+1k} are shown in Figure 11. The unstable manifolds for the other a_{ij} are as for the trivial braid.
- For each $i \neq j$ and $l \in \{1, \dots, m\}$, there is an interval neighborhood $J_{ij}^{2l} \subset \{s = s_{2l}\}$ of $W^u(a_{ij}) \cap \{s = s_{2l}\}$ such that the intervals J_{ij}^{2l} are disjoint for fixed l , and if we consider the set of all negative gradient flow lines of F_{ij} in the $[s_{2l-1}, s_{2l}]$ slice that start on J_{ij}^{2l} for all i, j , then any pair of distinct flow lines from this set intersect transversely. See Figure 12.
- On each $\{s = s_{2l-1}\}$ there are two intervals $J_{\pi/2}^{2l-1}$ and $J_{3\pi/2}^{2l-1}$ such that

$$\begin{aligned}
 J_{\pi/2}^{2l-1} &\supset \bigcup_{i>j} (W^u(a_{ij}) \cap \{s = s_{2l-1}\}) \quad \text{and} \\
 J_{3\pi/2}^{2l-1} &\supset \bigcup_{i<j} (W^u(a_{ij}) \cap \{s = s_{2l-1}\}).
 \end{aligned}$$




Figure 11: Flow lines in positive and negative twist slices




Figure 12: The intervals J_{ij}^{2l} , J_{ij}^{2l-1} , and J_{ij}^{2l-2} , and some representative gradient flow lines

Moreover, any negative gradient flow line of F_{ij} ($i > j$) that intersects J_{ij}^{2l-1} also intersects J_{ij}^{2l-2} , and similarly for J_{ij}^{2l-1} for $i < j$.

Assume that Λ_K has r components $\Lambda = \Lambda_{K;1} \cup \dots \cup \Lambda_{K;r}$. We will keep track of homology classes of cycles in $H_1(\Lambda_K)$ by counting intersections with certain fixed cycles. On each component $\Lambda_{K;j}$, fix the curve μ'_j that is the preimage under the base projection map $\pi: \Lambda_{K;j} \rightarrow \Lambda$ of the curve $t = \frac{1}{2}\pi - \epsilon_1$ for ϵ_1 positive and extremely small (for now, the line $t = \frac{\pi}{2}$ will suffice; in Section 4, we will need μ'_j to lie below $t = \frac{\pi}{2}$ but above the unstable manifolds of the a_{ij} for all $i > j$). Also, fix the curve λ'_j which is the preimage of the curve $s = 2\pi - 3\delta$ (a vertical curve between the a_{ij}

and the b_{ij}) in the leading sheet of $\Lambda_{K;j}$, where we recall from the introduction that “leading” refers to the first of the n sheets of Λ_K that belongs to component $\Lambda_{K;j}$. Intersections of a cycle in $H_1(\Lambda_K)$ with μ'_j and λ'_j then count the multiplicity of the j^{th} meridian and longitude class in the cycle. See Figure 13.

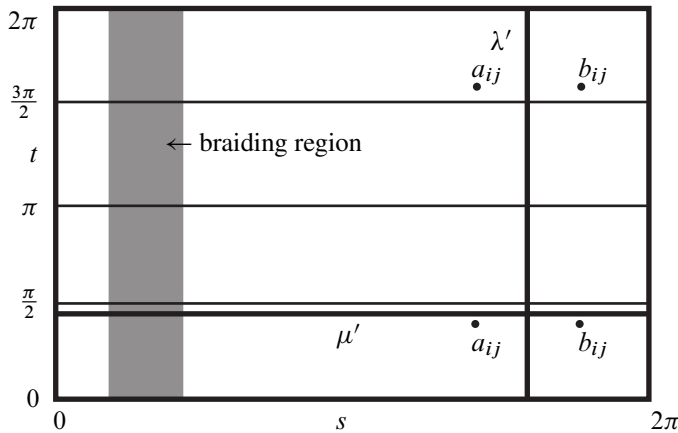


Figure 13: Schematic picture of Λ_K , as projected to the torus $\Lambda \approx T^2$

We choose a base point in each component over $(s, t) = (\frac{\pi}{2}, \frac{9\pi}{8})$ and endpoint paths for each Reeb chord endpoint which are disjoint from λ'_j and μ'_j . (Since the complement of $\lambda'_j \cup \mu'_j$ in $\Lambda_{K;j}$ is a disk such paths exist.)

3.4.4 Orientation choices and sign rules for $\Lambda_K \subset J^1(\Lambda)$ Before we proceed with the computation of flow trees for $\Lambda_K \subset J^1(\Lambda)$, we discuss the general method we use to assign signs to flow trees. These signs come from a fairly elaborate orientation scheme which depend on certain initial choices, some of which are of global nature which we call *basic orientation choices*, and others which are local, more specifically, the choice of orientations of determinant lines of a capping operator associated to each Reeb chord. Here we will simply state a combinatorial rule that comes from one particular set of choices. The derivation of the combinatorial rule and the effect of orientation choices is discussed in detail in Section 6. (For later computations, we will also need signs for multiscale flow trees; this is discussed in Section 4.3.)

We will discuss signs of rigid flow trees with one positive puncture or partial flow trees of dimension 1 with a special positive puncture of $\Lambda_K \subset J^1(\Lambda)$. Cutting a rigid flow tree close to its positive puncture we obtain a partial flow tree of dimension 1 with special positive puncture so it suffices to consider this case.

We first discuss how orientation choices for the “capping operators” corresponding to the Reeb chords a_{ij} and b_{ij} of $\Lambda_K \subset J^1(\Lambda)$ are encoded geometrically:

- Consider a Reeb chord a_{ij} , which is of type S_0 (with notation from Remark 3.8). Let $W^u(a_{ij})$ denote the unstable manifold of the positive local function difference defining a_{ij} . Fix a vector $v^{\ker}(a_{ij})$ perpendicular to $W^u(a_{ij})$; see Figure 14. (This choice corresponds to the choice of an orientation of the capping operator of a_{ij} .)
- Consider a Reeb chord b_{ij} , which is of type S_1 and note that $b_{ij} \in W^u(a_{ij})$. Fix vectors $v^{\ker}(b_{ij})$ parallel to $W^u(a_{ij})$ and $v^{\text{coker}}(b_{ij})$ perpendicular to $W^u(a_{ij})$; see Figure 14. (This choice corresponds to the choice of an orientation of the capping operator of b_{ij} .)
- If t is a trivalent vertex of a partial flow tree Γ then we let $v^{\text{con}}(t)$ be a vector tangent to the incoming edge at t and pointing into this edge; see Figure 15. (This is a reflection of a chosen orientation on the space of conformal structures on the disk with boundary punctures.)




Figure 14: Orientation data at b_{ij} and a_{ij} . (The unstable manifold $W^u(a_{ij})$ is in grey.)




Figure 15: Orientation data at a trivalent Y_0 -vertex. (Gradient flow tree is in grey.)

We next define two functions that are central to our definition of signs of rigid multiscale flow trees. Let $\langle \cdot, \cdot \rangle$ denote a Riemannian metric on S (which we will take to be the flat metric on the torus) and let

$$\text{sign}: \mathbb{R} - \{0\} \rightarrow \mathbb{R}$$

be the function that maps negative numbers to -1 and positive numbers to 1 . First consider a flow tree Γ of $\Lambda \subset J^1(S)$. Let b_{ij} be its positive puncture, $v^{\text{flow}}(\Gamma)$ denote

the vector field of the flow orientation of Γ and define

$$\sigma_{\text{pos}}(\Gamma) = \text{sign}(\langle v^{\text{flow}}(\Gamma), v^{\text{ker}}(b_{ij}) \rangle).$$

Next consider a partial flow tree Γ of $\Lambda \subset J^1(S)$ with positive special puncture p , trivalent vertices t_1, \dots, t_{k-1} , and negative punctures q_1, \dots, q_k and let n be a normal vector of Γ at p . Denote the result of vector splitting of n along Γ by

$$((w_1(t_1), w_2(t_1)), \dots, (w_1(t_{k-1}), w_2(t_{k-1}))); w(q_1), \dots, w(q_k))$$

and define

$$(3-10) \quad \sigma_{n,\Gamma}(t) = \text{sign}(\langle w_2(t) - w_1(t), v^{\text{con}}(t) \rangle), \quad t \in \{t_1, \dots, t_{k-1}\},$$

$$(3-11) \quad \sigma_{n,\Gamma}(q) = \text{sign}(\langle w(q_j), v^{\text{ker}}(q_j) \rangle), \quad q \in \{q_1, \dots, q_k\}.$$

Finally we define

$$\sigma(n, \Gamma) = \prod_{j=1}^k \sigma_{n,\Gamma}(q_j) \prod_{j=1}^{k-1} \sigma_{n,\Gamma}(t_j).$$

We can now assign orientations to trees according to the following theorem.

Theorem 3.13 *There exists a choice of basic orientations and of orientations of capping operators for all Reeb chords of type \mathbf{L} such that for a rigid flow tree Γ of Λ in $J^1(S)$ with positive puncture at b_{ij} of type \mathbf{S}_1 , the sign of Γ is*

$$\sigma_{\text{pos}}(\Gamma) \sigma(v^{\text{coker}}(b_{ij}), \Gamma).$$

Proof Theorem 3.13 is a special case of Theorem 4.6, which is proved in Section 6.6. □

3.4.5 Partial flow trees in twist slices We now begin our enumeration of flow trees for $\Lambda_K \subset J^1(\Lambda)$. Consider a braid whose closure is K , and assume that the braid is in the form given in Section 3.4.3. Then flow trees for Λ_K decompose nicely into pieces in each twist region $[s_{2l-2}, s_{2l}]$, which we call *partial flow trees*. Using the notation from Section 3.4.3, we focus on one of these twist regions, an interval $[s_{2l-2}, s_{2l}] \times [0, 2\pi]$ (for fixed $l \in \{1, \dots, m\}$) containing the l^{th} twist $\sigma_k^{\pm 1}$ from the braid. We first define a special type of partial flow tree in such a slice, which we call a *slice tree*.

Let $n \geq i > j \geq 1$. Fix symbols \bar{a}_{ij} and \bar{a}_{ji} for negative punctures and \bar{b}_{ij} and \bar{b}_{ji} for positive punctures; at the moment, these are just symbols and do not correspond to actual punctures or special punctures. Now for each $i > j$ choose one point $\bar{a}'_{ij} \in \{s_{2l}\} \times J_{ij}^{2l}$ and think of it as a special puncture connecting the sheet S_j to S_i , where the sheets are numbered by the order of the braid strands at $s = s_{2l}$. Similarly choose $\bar{a}'_{ji} \in \{s_{2l}\} \times J_{ji}^{2l}$

and think of it as a special puncture connecting the sheet S_i to S_j . Once this choice is made we will frequently conflate the variables \bar{a}_{ij} and \bar{a}'_{ij} . We will only use the primes when we need to refer to specific special punctures.

Given these choices any flow tree lying entirely in the $[s_{2l-2}, s_{2l}]$ -slice that has a positive special puncture connecting sheets S_j to S_i at any point in J_{ij}^{2l-2} and negative punctures at the \bar{a}'_{pq} will be called a *slice tree with positive special puncture at \bar{b}_{ij}* (think of \bar{b}_{ij} as the intersection of the flow tree with J_{ij}^{2l-2}). We let $\mathcal{T}^\pm(\bar{b}_{ij})$ denote the set of slice trees in the $[s_0, s_1]$ -slice with positive special puncture at \bar{b}_{ij} .

Remark 3.14 Note that the formal dimension of slice trees $\mathcal{T}^\pm(\bar{b}_{ij})$ is 0 (recall we have fixed the locations of all the \bar{a}_{pq}). Note also that the set $\mathcal{T}^\pm(\bar{b}_{ij})$ is independent of the specific choice of the \bar{a}'_{pq} , ie, a different choice of the \bar{a}'_{pq} induces a one-to-one correspondence of $\mathcal{T}^\pm(\bar{b}_{ij})$. This follows from the choice of the intervals J_{ij}^{2l} , $J_{\pi/2}^{2l-1}$, $J_{3\pi/2}^{2l-1}$ in Section 3.4.3.

In order to keep track of orientation signs of flow trees, we will decorate the special punctures with arrows which should be thought of as normal vectors to flow lines through the special punctures. We write \bar{a}_{kl}^\uparrow and \bar{b}_{kl}^\uparrow for the relevant special puncture decorated with a vector v normal to the slice tree at that point with $\langle v, \partial_t \rangle > 0$; similarly, we write \bar{a}_{kl}^\downarrow and \bar{b}_{kl}^\downarrow for the same special puncture decorated with a normal vector v' such that $\langle v', \partial_t \rangle < 0$.

Furthermore, as the twist is part of a larger braid, each strand belongs to a link component and we need to keep track of the corresponding homology coefficients. We write μ_A (respectively μ_B) for the homology variable of the meridian of the component of Λ_K associated to sheet A (respectively B), which is the sheet numbered by k (respectively $k + 1$) at $s = s_{2l}$ and by $k + 1$ (respectively k) at $s = s_{2l-2}$.

Let $\bar{\mathcal{B}}$ denote the set of decorated special chords at $s = s_{2l-2}$:

$$\bar{\mathcal{B}} = \{\bar{b}_{ij}^\uparrow, \bar{b}_{ij}^\downarrow\}_{1 \leq i, j \leq n, i \neq j},$$

and let $\bar{\mathcal{A}}$ denote the \mathbb{Z} -algebra generated by $\mu_A^{\pm 1}$, $\mu_B^{\pm 1}$, and decorated special chords at $s = s_{2l}$:

$$\bar{\mathcal{A}} = \mathbb{Z}\langle \mu_A^{\pm 1}, \mu_B^{\pm 1}, \bar{a}_{ij}^\uparrow, \bar{a}_{ij}^\downarrow \rangle_{1 \leq i, j \leq n, i \neq j}.$$

Given a tree $\Gamma \in \mathcal{T}^\pm(\bar{b}_{ij})$ and a normal v for \bar{b}_{ij} , the 1-jet lift of Γ determines a word in $\bar{\mathcal{A}}$, in a manner that we now describe. Orient the 1-jet lift of Γ by the flow orientation (see Section 2.4). This orientation induces an ordering $a_{i_1 j_1}, \dots, a_{i_q j_q}$ of the negative punctures of Γ so that the 1-jet lift of Γ consists of a union of oriented

paths $\gamma_0, \dots, \gamma_q \subset \Lambda_K$, with the beginning point of γ_0 and the end point of γ_q equal to the ends of \bar{b}_{ij} (the points in Λ_K lying over \bar{b}_{ij} on sheets i and j), and with the beginning point of γ_r and the end point of γ_{r-1} equal to the ends of $\bar{a}_{i_r j_r}$ for $r = 1, \dots, q$.

For $0 \leq r \leq q$, define $n'_A, n'_B \in \mathbb{Z}$ to be the intersection numbers of γ_r with the preimage of $t = \frac{\pi}{2}$ in sheets A, B , respectively. (In the notation of Section 3.4.3, these numbers count intersections with μ'_A and μ'_B .) Also, use vector splitting (see Section 4.3) from the normal vector ν at \bar{b}_{ij} to obtain normal vectors ν_r at each $\bar{a}_{i_r j_r}$. Finally, define the word

$$q(\Gamma, \nu) := \left(\mu_A^{n'_A} \mu_B^{n'_B} \right) \bar{a}_{i_1 j_1}^{\nu_1} \left(\mu_A^{n_1} \mu_B^{n_1} \right) \bar{a}_{i_2 j_2}^{\nu_2} \cdots \bar{a}_{i_q j_q}^{\nu_q} \left(\mu_A^{n_A} \mu_B^{n_B} \right).$$

We can now define maps $\eta_{\sigma_k^{\pm 1}}: \bar{\mathcal{B}} \rightarrow \bar{\mathcal{A}}$ as follows, with $\nu \in \{\uparrow, \downarrow\}$:

$$\eta_{\sigma_k^{\pm 1}}(\bar{b}_{ij}^\nu) = \sum_{\Gamma \in \mathcal{T}^\pm(\bar{b}_{ij})} \tau(\Gamma) q(\Gamma, \nu),$$

where $\tau(\Gamma) = \prod_t \sigma_{n, \Gamma}(t) \in \{1, -1\}$ with the product running over all trivalent punctures of Γ ; see (3-10).

If $b \in \bar{\mathcal{B}}$ then define b^\dagger as the same chord but with the decorating normal reversed. Similarly, for a monomial $q \in \bar{\mathcal{A}}$ define q^\dagger as the same monomial of chords but with all decorating normals reversed. Write

$$\eta_{\sigma_k^{\pm 1}}(b) = \eta_{\sigma_k^{\pm 1}}^{\text{odd}}(b) + \eta_{\sigma_k^{\pm 1}}^{\text{even}}(b),$$

where the two terms on the right hand side are summands containing all monomials with an odd and even number of variables, respectively.

Lemma 3.15 *For any $b \in \bar{\mathcal{B}}$, the map $\eta_{\sigma_k^{\pm 1}}$ satisfies*

$$\eta_{\sigma_k^{\pm 1}}(b^\dagger) = \left(\eta_{\sigma_k^{\pm 1}}^{\text{odd}}(b) \right)^\dagger - \left(\eta_{\sigma_k^{\pm 1}}^{\text{even}}(b) \right)^\dagger.$$

Proof Let n be a normal at the positive puncture. Then the vector splittings of $-n$ and n along Γ differs by an over all sign. Since the number of trivalent vertices in Γ is one less than the number of negative punctures of Γ , the lemma follows. \square

We next turn to the actual calculation of the maps η_{σ_k} and $\eta_{\sigma_k^{-1}}$. By Lemma 3.15, it is sufficient to compute for \bar{b}_{ij}^\uparrow , $1 \leq i, j \leq n$, $i \neq j$.

Lemma 3.16 Let $\eta_{\sigma_k^{\pm 1}}: \bar{\mathcal{B}} \rightarrow \bar{\mathcal{A}}$ denote the maps associated to a twist representing the braid group generator $\sigma_k^{\pm 1}$ in an $[s_0, s_1]$ -slice, as described above. Then

$$\begin{aligned} \eta_{\sigma_k}(\bar{b}_{ij}^\uparrow) &= \bar{a}_{ij}^\uparrow & i, j \neq k, k+1 \\ \eta_{\sigma_k}(\bar{b}_{kk+1}^\uparrow) &= \bar{a}_{k+1k}^\uparrow \\ \eta_{\sigma_k}(\bar{b}_{k+1k}^\uparrow) &= \mu_A \bar{a}_{kk+1}^\uparrow \mu_B^{-1} \\ \eta_{\sigma_k}(\bar{b}_{ik+1}^\uparrow) &= \bar{a}_{ik}^\uparrow & i \neq k, k+1 \\ \eta_{\sigma_k}(\bar{b}_{k+1i}^\uparrow) &= \bar{a}_{ki}^\uparrow & i \neq k, k+1 \\ \eta_{\sigma_k}(\bar{b}_{ik}^\uparrow) &= \bar{a}_{ik+1}^\uparrow + \bar{a}_{ik}^\uparrow \bar{a}_{kk+1}^\uparrow & i < k \\ \eta_{\sigma_k}(\bar{b}_{ik}^\uparrow) &= \bar{a}_{ik+1}^\uparrow + \bar{a}_{ik}^\uparrow \mu_A \bar{a}_{kk+1}^\uparrow \mu_B^{-1} & i > k+1 \\ \eta_{\sigma_k}(\bar{b}_{ki}^\uparrow) &= \bar{a}_{k+1i}^\uparrow - \bar{a}_{k+1k}^\uparrow \bar{a}_{ki}^\uparrow & i \neq k, k+1 \end{aligned}$$

and

$$\begin{aligned} \eta_{\sigma_k^{-1}}(\bar{b}_{ij}^\uparrow) &= \bar{a}_{ij}^\uparrow & i, j \neq k, k+1 \\ \eta_{\sigma_k^{-1}}(\bar{b}_{kk+1}^\uparrow) &= \mu_B^{-1} \bar{a}_{k+1k}^\uparrow \mu_A \\ \eta_{\sigma_k^{-1}}(\bar{b}_{k+1k}^\uparrow) &= \bar{a}_{kk+1}^\uparrow \\ \eta_{\sigma_k^{-1}}(\bar{b}_{ik}^\uparrow) &= \bar{a}_{ik+1}^\uparrow & i \neq k, k+1 \\ \eta_{\sigma_k^{-1}}(\bar{b}_{ki}^\uparrow) &= \bar{a}_{k+1i}^\uparrow & i \neq k, k+1 \\ \eta_{\sigma_k^{-1}}(\bar{b}_{ik+1}^\uparrow) &= \bar{a}_{ik}^\uparrow - \bar{a}_{ik+1}^\uparrow \mu_B^{-1} \bar{a}_{k+1k}^\uparrow \mu_A & i < k \\ \eta_{\sigma_k^{-1}}(\bar{b}_{ik+1}^\uparrow) &= \bar{a}_{ik}^\uparrow - \bar{a}_{ik+1}^\uparrow \bar{a}_{k+1k}^\uparrow & i > k+1 \\ \eta_{\sigma_k^{-1}}(\bar{b}_{k+1i}^\uparrow) &= \bar{a}_{ki}^\uparrow + \bar{a}_{kk+1}^\uparrow \bar{a}_{k+1i}^\uparrow & i \neq k, k+1. \end{aligned}$$

Proof We label the sheets of Λ_K in the slice under consideration by

$$S_1, \dots, S_{k-1}, A, B, S_{k+2}, \dots, S_n.$$

Here the sheet S_j corresponds to the j^{th} strand of the standard trivial braid, and sheets A and B are defined as before.

Consider first the linear terms in the expressions for η_{σ_k} and $\eta_{\sigma_k^{-1}}$. The trees that give these contributions are simply the (negative) gradient flow lines of positive function differences that end at \bar{a}_{ij} . They obviously exist, are unique, and transport normal vectors as claimed. This proves that the linear terms of the equations are correct except

for homology coefficients. To see these, note that only the flow lines between sheets A and B intersect $t = \frac{\pi}{2}$. Further, a flow line between sheets A and B intersects this curve only if the twist is positive and it ends at $\bar{a}_{k k+1}$ or if the twist is negative and it ends at $\bar{a}_{k+1 k}$. Finally, noting that the component in the upper sheet of the 1-jet lift is oriented according to the flow orientation and that the component in the lower sheet is oriented opposite to the flow orientation, it follows that the coefficients are as stated.

We next study higher order terms arising from trees with trivalent vertices. Such a vertex arises as follows: a flow line determined by sheets X and Y splits into two flow lines determined by sheets X and Z , and by Z and Y . Furthermore, since we consider only trees with one positive puncture it is required that all flow lines correspond to positive function differences. In other words, the z -coordinate of the sheet Z must lie between the z -coordinates of the sheets X and Y at the splitting point.

Given a slice tree Γ suppose that \bar{a}_{ij} is a negative special puncture of Γ with $i, j \neq k, k + 1$. Then the negative gradient flow γ_{ij} of F_{ij} ending at \bar{a}'_{ij} in $[s_{2l-1}, s_{2l}]$ is disjoint from all the flow lines starting at all the other $\bar{a}'_{i'j'}$, except the flow line of $F_{k k+1}$ ending at $\bar{a}'_{k k+1}$ or of $F_{k+1 k}$ ending at $\bar{a}'_{k+1 k}$. However we cannot have a Y_0 vertex here since $i, j \neq k, k + 1$. Also, γ_{ij} is disjoint from the other flow lines in the interval $[s_{2l-2}, s_{2l-1}]$. Thus if \bar{a}'_{ij} is a negative special puncture of Γ then Γ has just one edge.

If $\bar{a}'_{i k+1}$ is a negative puncture of Γ and $i < k$ then the negative gradient flow line of $F_{i k+1}$ starting at $\bar{a}'_{i k+1}$ intersects only the flow line of $F_{k k+1}$, but since they have the same upper sheets these flow lines cannot merge at a Y_0 vertex. Thus again Γ just has one edge. Similarly Γ will only have one edge if $i > k + 1$. The same argument shows that a slice tree with negative puncture at $\bar{a}'_{k+1 i}$ must have only one edge.

It remains to consider slice trees with negative punctures among $\bar{a}'_{k k+1}, \bar{a}'_{k+1 k},$ and $\bar{a}'_{i k}, \bar{a}'_{k i}$ for $i \neq k, k + 1$. Besides the flow lines ending at $\bar{a}'_{k k+1}$ or $\bar{a}'_{k+1 k}$, all such slice trees must have some $\bar{a}'_{i k}$ or $\bar{a}'_{k i}, i \neq k, k + 1$, as a negative puncture.

If $\bar{a}'_{i k}$ is a negative puncture of Γ then the negative gradient flow line $\gamma_{i k}$ of $F_{i k}$ ending at $\bar{a}'_{i k}$ intersects the flow line $\gamma_{k k+1}$ of $F_{k k+1}$ ending at $\bar{a}'_{k k+1}$. Thus we can have a Y_0 splitting of the flow line from $F_{i k+1}$ into these flow lines; see Figure 16 for an illustration when $i > k + 1$. Thus there is a point c in $J_{\frac{\pi}{2}}^{2l-1}$ or $J_{\frac{3\pi}{2}}^{2l-1}$ and a flow line γ of $F_{i k+1}$ starting at c so that $\gamma \cup \gamma_{k k+1} \cup \gamma_{i k}$ forms a flow tree. Notice assuming the twist interval is small enough it is clear that γ does not intersect any unstable manifolds of the a_{ij} in $[s_{2l-1}, s_{2l}]$. We can extend γ through the interval $[s_{2l-2}, s_{2l-1}]$ and see that there is a unique point $\bar{b}'_{i k}$ in $J_{i k}^{2l-2}$ that γ runs through. (Recall our labeling conventions from the beginning of the subsection: the sheets on the left side of the interval $[s_{2l-2}, s_{2l}]$ are labeled just as on the right, but with the role




Figure 16: A flow tree with one Y_0 vertex and negative punctures at \bar{a}'_{ik} and \bar{a}'_{k+1} . The case $i < k$ is shown here.

of k and $k + 1$ switched.) Moreover notice that by the lexicographical ordering on the unstable manifolds, γ will not intersect any of the unstable manifolds of the $\bar{a}'_{i'j'}$ over the interval $[s_{2l-2}, s_{2l-1}]$.

Thus we see that if \bar{a}'_{ik} is a negative puncture of Γ then Γ either has one edge (this is the trivial flow line mentioned above) or exactly one Y_0 vertex. A similar argument says the same for Γ having a negative puncture at \bar{a}'_{ki} . See Figure 17 for an illustration of all slice trees with a Y_0 vertex and a negative puncture at \bar{a}'_{ik} or \bar{a}'_{ki} for some $i \neq k, k + 1$.

We conclude that all slice trees in $[s_{2l-2}, s_{2l}]$ either have one edge or, if the positive puncture is \bar{b}_{ik} or \bar{b}_{ki} , are of the type described above (and shown in Figure 17) with three edges and one Y_0 vertex. This shows that the quadratic terms of the equations are correct except for homology coefficients and vector splittings. To see the homology coefficients simply notice that the only flow trees that intersect $t = \frac{\pi}{2}$ are the ones containing a flow line of F_{kk+1} and having $i > k + 1$. Since such a flow tree crosses μ'_A positively between \bar{a}'_{ik} and \bar{a}'_{kk+1} , and crosses μ'_B negatively after \bar{a}'_{kk+1} , it contributes the word $\bar{a}'_{ik}\mu_A\bar{a}'_{kk+1}\mu_B^{-1}$ to $\eta_{\sigma_k}(\bar{b}_{ik})$.

As for vector splittings for the Y_0 trees, note that the upward normal at \bar{b}'_{ik} or \bar{b}'_{ki} is split into upward normals at each of the negative punctures (\bar{a}'_{ik} and \bar{a}'_{kk+1} , or \bar{a}'_{ki} and \bar{a}'_{k+1k}). For a Y_0 tree Γ , an easy application of the definitions from Section 3.4.4 shows that $\tau(\Gamma)$ is $+1$ for the top two trees in Figure 17 and -1 for the bottom two trees. (The difference between these pairs is the relative placement at the trivalent vertex of the flows labeled 1 and 2 from Figure 4.) Thus the arrow decorations and signs are as given in the statement of the lemma.

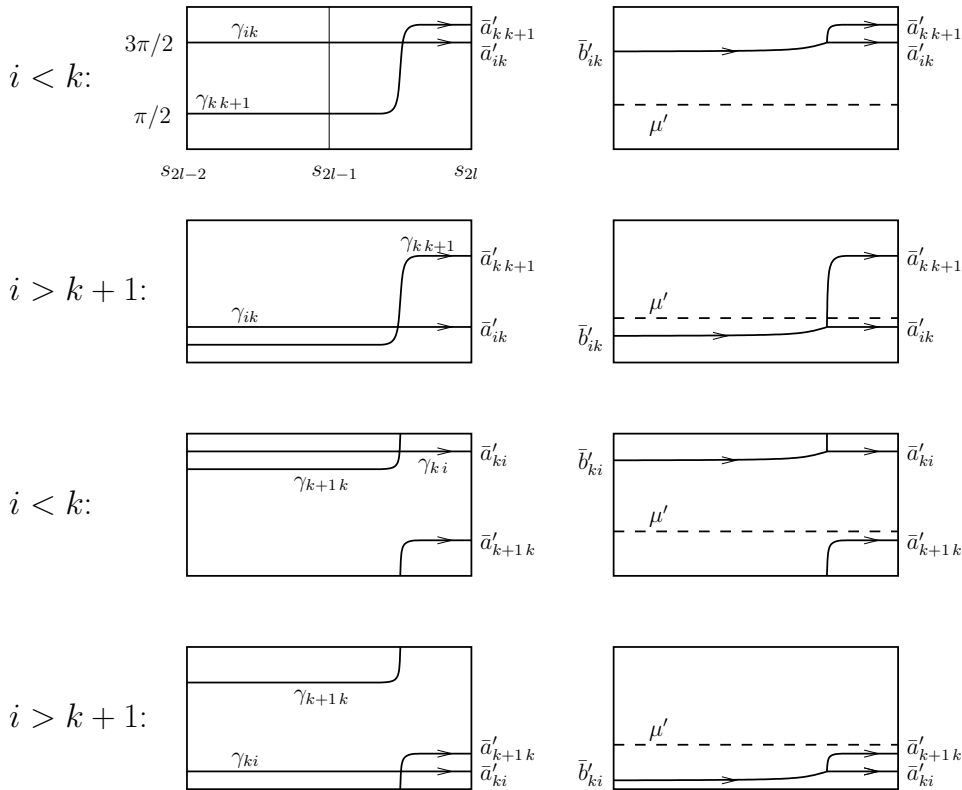


Figure 17: The four slice trees that consist of more than one edge (right), resulting from the four configurations of intersecting flow lines (left). The cycle μ' is also shown on the right, and intersects only the second slice tree. Up to sign, we can read off the words associated to these four slice trees, from top to bottom: $\bar{a}'_{ik}\bar{a}'_{k,k+1}$, $\bar{a}'_{ik}\mu_A\bar{a}'_{k,k+1}\mu_B^{-1}$, $\bar{a}'_{k+1k}\bar{a}'_{ki}$, $\bar{a}'_{k+1k}\bar{a}'_{ki}$.

The case for σ_k^{-1} can be handled similarly. □

3.4.6 Decomposing flow trees Our expression for the differential derives from the following geometric decomposition of flow trees. Consider a braid $B = \sigma_{k_1}^{\epsilon_1} \cdots \sigma_{k_m}^{\epsilon_m}$, $\epsilon_l = \pm 1$, on n strands. Let K denote the closure of B and consider Λ_K . We represent Λ_K as described above as a sequence of twists slices separated by trivial braid slices in the braiding region, one small slice containing all Reeb chords, and the remaining trivial braid slice, which is the complement of these two slices. Consider an $[s', s'']$ -slice as just described in which there are no Reeb chords.

Lemma 3.17 Any rigid flow tree or 1–dimensional partial flow tree of Λ_K with one positive puncture intersects the $[s', s'']$ –slice in a union of slice trees (for appropriately chosen \bar{a}'_{ij}).

Proof Suppose Γ is a flow tree for Λ_K . We see that Γ will intersect the slice immediately preceding the a_{ij} –chords of Λ_K in a union of slice trees, where the $\bar{a}'_{ij} = a_{ij}$, by definition. Notice that the slice trees in this slice define \bar{a}'_{ij} for the next slice. Using the new \bar{a}'_{ij} we see again that Γ will intersect this next slice in a union of slice trees. Continuing by induction we see that Γ will intersect each slice in a union of slice trees determined by appropriately chosen \bar{a}'_{ij} . \square

Remark 3.18 We note that any partial flow tree Γ' of Λ_K , with only the positive puncture p special, can be completed in a unique way to a flow tree Γ by appending a flow line connecting p to a Reeb chord b_{ij} .

3.4.7 Twist morphisms The following algebraic construction makes it possible to apply the result in Lemma 3.16 inductively. Let K be a link with r components and let $\Lambda_K = \Lambda_{K;1} \cup \dots \cup \Lambda_{K;r}$ be the subdivision of its conormal lift into components. Let $\mu_j, \lambda_j \in H_1(\Lambda_{K;j}; \mathbb{Z})$ be as described above. As in the introduction, consider the algebra \mathcal{A}_n^0 over \mathbb{Z} generated by $\mathbb{Z}[H_1(\Lambda_K)]$ along with the Reeb chords a_{ij} , $1 \leq i, j \leq n, i \neq j$. We will define morphisms $\phi_{\sigma_k^{\pm 1}}: \mathcal{A}_n^0 \rightarrow \mathcal{A}_n^0$ associated to braid group generators.

Consider a twist corresponding to $\sigma_k^{\pm 1}$. We use the notation μ_A and μ_B for homology variables exactly like in Lemma 3.16 and in order to connect to that result we make the following identifications:

$$\begin{aligned}
 \bar{b}_{ij}^\uparrow &= +a_{ij} & \text{if } i > j, \\
 \bar{b}_{ij}^\downarrow &= -a_{ij} & \text{if } i > j, \\
 \bar{b}_{ij}^\uparrow &= -a_{ij} & \text{if } i < j, \\
 \bar{b}_{ij}^\downarrow &= +a_{ij} & \text{if } i < j
 \end{aligned}
 \tag{3-12}$$

in the source \mathcal{A}_n^0 , and

$$\begin{aligned}
 \bar{a}_{ij}^\uparrow &= +a_{ij} & \text{if } i > j, \\
 \bar{a}_{ij}^\downarrow &= -a_{ij} & \text{if } i > j, \\
 \bar{a}_{ij}^\uparrow &= -a_{ij} & \text{if } i < j, \\
 \bar{a}_{ij}^\downarrow &= +a_{ij} & \text{if } i < j
 \end{aligned}
 \tag{3-13}$$

in the target \mathcal{A}_n^0 . We define a homomorphism $\phi_{\sigma_k}: \mathcal{A}_n^0 \rightarrow \mathcal{A}_n^0$ using this identification in combination with Lemmas 3.16 and 3.15. That is, we define it as follows on generators:

$$\begin{aligned} \phi_{\sigma_k}(a_{ij}) &= a_{ij} & i, j \neq k, k + 1, \\ \phi_{\sigma_k}(a_{kk+1}) &= -a_{k+1k} \\ \phi_{\sigma_k}(a_{k+1k}) &= -\mu_A a_{kk+1} \mu_B^{-1} \\ \phi_{\sigma_k}(a_{ik+1}) &= a_{ik} & i \neq k, k + 1, \\ \phi_{\sigma_k}(a_{k+1i}) &= a_{ki} & i \neq k, k + 1, \\ \phi_{\sigma_k}(a_{ik}) &= a_{ik+1} - a_{ik} a_{kk+1} & i < k, \\ \phi_{\sigma_k}(a_{ik}) &= a_{ik+1} - a_{ik} \mu_A a_{kk+1} \mu_B^{-1} & i > k + 1, \\ \phi_{\sigma_k}(a_{ki}) &= a_{k+1i} - a_{k+1k} a_{ki} & i \neq k, k + 1. \end{aligned}$$

Similarly, we define $\phi_{\sigma_k^{-1}}: \mathcal{A}_0 \rightarrow \mathcal{A}_0$ as follows on generators:

$$\begin{aligned} \phi_{\sigma_k^{-1}}(a_{ij}) &= a_{ij} & i, j \neq k, k + 1, \\ \phi_{\sigma_k^{-1}}(a_{kk+1}) &= -\mu_B^{-1} a_{k+1k} \mu_A \\ \phi_{\sigma_k^{-1}}(a_{k+1k}) &= -a_{kk+1} \\ \phi_{\sigma_k^{-1}}(a_{ik}) &= a_{ik+1} & i \neq k, k + 1, \\ \phi_{\sigma_k^{-1}}(a_{ki}) &= a_{k+1i} & i \neq k, k + 1, \\ \phi_{\sigma_k^{-1}}(a_{ik+1}) &= a_{ik} - a_{ik+1} \mu_B^{-1} a_{k+1k} \mu_A & i < k, \\ \phi_{\sigma_k^{-1}}(a_{ik+1}) &= a_{ik} - a_{ik+1} a_{k+1k} & i > k + 1, \\ \phi_{\sigma_k^{-1}}(a_{k+1i}) &= a_{ki} - a_{k+1k} a_{k+1i} & i \neq k, k + 1. \end{aligned}$$

Note that $\phi_{\sigma_k^{-1}} \circ \phi_{\sigma_k} = \phi_{\sigma_k} \circ \phi_{\sigma_k^{-1}} = \text{id}$, since A, B switch places between the σ_k and σ_k^{-1} twists.

Remark 3.19 In any equation above where the sign differs from that of the corresponding equation of the formulas in Lemma 3.16 the following holds. If we use (3-12) and (3-13) to substitute decorated variables \bar{b}_{ij}^\uparrow etc, exactly one arrow of a decorated chord in the target monomial is oriented differently than all other arrows in the equation.

For the braid $B = \sigma_{k_1}^{\pm 1} \dots \sigma_{k_m}^{\pm 1}$, define

$$\phi_B = \phi_{\sigma_{k_1}^{\pm 1}} \circ \dots \circ \phi_{\sigma_{k_m}^{\pm 1}}.$$

Note that ϕ then gives a representation of the braid group into the group of automorphisms of \mathcal{A}_n^0 .

Remark 3.20 In order for ϕ_B to respect the order of composition we think of the braid as written in the “operator order,” so that B above should be interpreted as: apply $\sigma_{k_m}^{\pm 1}$ first and $\sigma_{k_1}^{\pm 1}$ last. Thinking of the braid in the opposite order, $B \mapsto \phi_B$ would be an anti-homomorphism. See Figure 18.

3.4.8 Flow trees for Λ_K in $J^1(\Lambda)$ Since there are no cusps of Λ_K in $J^1(\Lambda)$, we know for grading reasons that any rigid flow tree must have its positive puncture at some Reeb chord b_{ij} and its negative punctures at Reeb chords a_{ij} . Consequently, Theorem 2.7 implies that the differential of the Legendrian algebra can be computed as

$$\partial(b_{ij}) = \sum_{\Gamma \in \mathcal{T}(b_{ij})} \epsilon(\Gamma)q(\Gamma),$$

where $\mathcal{T}(b_{ij})$ denotes the set of all flow trees with positive puncture at b_{ij} and where, if Γ is such a tree, $q(\Gamma)$ denotes the monomial of its negative punctures and $\epsilon(\Gamma)$ its sign. To compute this differential we fix orientation choices as follows, see Figure 18:

$$(3-14) \quad v^{\ker}(b_{ij}) = \partial_s, \quad \text{for all } 1 \leq i, j \leq n, i \neq j$$

and

$$(3-15) \quad v^{\text{coker}}(b_{ij}) = v^{\ker}(a_{ij}) = \begin{cases} \partial_t & \text{if } i > j, \\ -\partial_t & \text{if } i < j. \end{cases}$$




Figure 18: Flow trees contributing to $\partial\mathbb{B}$

The following lemma, where we use the matrix notation from Theorem 1.1, determines the differential discussed above.

Lemma 3.21 *With orientation data as in (3-14) and (3-15), the following equation holds:*

$$\partial\mathbb{B} = -\mathbb{L}^{-1} \cdot \mathbb{A} \cdot \mathbb{L} + \phi_{\mathbb{B}}(\mathbb{A}) = -\mathbb{L}^{-1} \cdot \mathbb{A} \cdot \mathbb{L} + \Phi_{\mathbb{B}}^L \cdot \mathbb{A} \cdot \Phi_{\mathbb{B}}^R.$$

Proof We will prove the left equality, as the right equality follows from Proposition 2.10. The first term in the right hand side of the left equality comes from the short flow lines connecting b_{ij} to a_{ij} , with flow orientation given by $-\partial_s$. These flow lines clearly exist and are unique. The sign ϵ of one of these flow lines Γ is given, according to Theorem 3.13, by

$$\begin{aligned} \epsilon(\Gamma) &= \text{sign}(\langle v^{\text{coker}}(b_{ij}), v^{\text{flow}}(\Gamma) \rangle) \text{sign}(\langle v^{\text{coker}}(b_{ij}), v^{\text{ker}}(a_{ij}) \rangle) \\ &= (-1) \cdot 1 = -1. \end{aligned}$$

Consider their homology coefficients. Since no endpoint path intersects μ'_j or λ'_j it is sufficient to consider the intersection between μ'_j and λ'_j and the 1-jet lift of the trees. Clearly, all flow lines under consideration are disjoint from μ'_j , $j = 1, \dots, r$. Furthermore, the 1-jet lift in the upper (respectively lower) sheet of flow line passes λ' in the negative (respectively positive) direction. Consequently a_{ij} comes with homology coefficients if and only if either the sheet S_i or the sheet S_j is the leading sheet for that link component. If S_i (respectively S_j) is the leading sheet, then the intersection with $\lambda'_{\gamma(i)}$ (respectively $\lambda_{\gamma(j)}$) in that sheet contributes the coefficient $\lambda_{\gamma(i)}^{-1}$ on the left of a_{ij} (respectively $\lambda_{\gamma(j)}$ on the right). This corresponds to multiplication by the matrix \mathbb{L}^{-1} from the right and \mathbb{L} from the left as claimed. See the leftmost picture in Figure 18.

The second term in the right hand side comes from the flow trees that end as flow lines in $W^u(a_{ij})$ flowing in the $+\partial_s$ direction. By Lemma 3.17 we find that the intersection of such a flow tree with any twist slice is a twist tree and thus any flow tree starting at b_{ij} and ending at the a_{ij} along unstable manifolds oriented in the ∂_s direction will contribute a term from $\phi_{\mathbb{B}}(a_{ij})$. Moreover any term in $\phi_{\mathbb{B}}(a_{ij})$ gives a flow tree starting at b_{ij} . See the rightmost picture in Figure 18. Note also that no such tree passes λ' . Lemma 3.16 in combination with the composition formula for $\phi_{\mathbb{B}}$ and the sign rule in Theorem 3.13 then shows that the second term is $\phi_{\mathbb{B}}(\mathbb{A})$ (since all the non-trivial terms in the sign rule appear at the Y_0 vertices, which were accounted for in the formula for ϕ). □

3.5 Negative gradient flows in twist slices

Our goal in this subsection is to construct a braid $B = \sigma_{k_1}^{\epsilon_1} \cdots \sigma_{k_m}^{\epsilon_m}$, $\epsilon_l = \pm 1$, in a braiding slice $[s_0^{\text{br}}, s_1^{\text{br}}]$, in such a way that the properties of unstable manifolds detailed in Section 3.4.3 are satisfied. We perform this construction twist by twist going from right (larger s) to left (smaller s). To facilitate the formulas we will change coordinates from (s, t) to (u, t) where $u = s_1^{\text{br}} - s$. So the braid region happens over $u \in [0, U]$ where $U = s_1^{\text{br}} - s_0^{\text{br}}$. Throughout the computation we will also change u by translations, but the key is that u is always $-s$ up to translations. In (u, t) coordinates we will build up the braid region, twist by twist, going from left (smaller u) to right (larger u); moreover in the u coordinates we consider the reversed word $\sigma_{k_m}^{\epsilon_m} \cdots \sigma_{k_1}^{\epsilon_1}$ so that when we switch back to the s coordinates it is K that is represented.

More specifically we will break $[0, U]$ into $3m$ subintervals I_1, \dots, I_{3m} , ordered from left to right. For $l = 1, \dots, m$, the union of the three intervals $I_{3l-2} \cup I_{3l-1} \cup I_{3l}$ is associated to the braid generator $\sigma_{k_{m-l+1}}^{\epsilon_{m-l+1}}$ and will be called the *braid interval* associated to $\sigma_{k_{m-l+1}}^{\epsilon_{m-l+1}}$.

The intervals I_{3l-2} , $l = 1, \dots, m$, will contain trivial braids as discussed in Remark 3.12 and are used to adjust the braid to prepare for a twist between two strands of the braid. These will be called *preparatory intervals*. The intervals I_{3l-1} , $l = 1, \dots, m$, will be the intervals over which two strands of the braid will actually twist. These will be called *twist intervals*. The intervals I_{3l} , $l = 1, \dots, m$, will contain the trivial braid and will be used to adjust the braid so that we can more easily count the flow trees. These will be called *concentration intervals*.

For comparison with Section 3.4.3, we set s_{2l-2} and s_{2l-1} to be the endpoints (in reverse order) of the concentration intervals $I_{3(m-l+1)}$ for $l = 1, \dots, m$, and s_{2m} to be the leftmost (in u) endpoint of I_1 . Then s_0, \dots, s_{2m} are arranged in increasing order and B is trivial in each $s \in [s_{2l-2}, s_{2l-1}]$ slice (which corresponds to a concentration interval) and comprises the braid generator $\sigma_{k_l}^{\epsilon_l}$ in each $s \in [s_{2l-1}, s_{2l}]$ slice (the union of a preparatory and a twist interval).

In the rest of this subsection we will describe the braid interval corresponding to the twist $\sigma_{k_l}^{\epsilon_l}$, but before focusing on this we make an observation and some conventions. First, for convenience, we will think of the function f_j describing the braid as maps $[0, C] \rightarrow \mathbb{R}^2$, for some arbitrarily large constant C , rather than $[0, U] \rightarrow D^2$. Since scaling the variable u and multiplying all the functions by a small constant will not affect the discussion below we will be able to return to the appropriate braid setup once we have constructed our desired functions $[0, C] \rightarrow \mathbb{R}^2$. (This step could be avoided by choosing appropriate scaling constants throughout the argument, but as

these constants would depend on the entire braid it is considerably simpler to proceed as we do.) Moreover we will start by considering the trivial braid over $\mathbb{R}_{\geq 0}$, that is, our functions f_j will be maps $[0, \infty) \rightarrow \mathbb{R}^2$. We will then alter the f_j over some interval $[0, c_1]$, which we call I_1 , then over $[c_1, c_2]$, which we call I_2 , and so on. Once we finish with the interval I_{3m} we let $[0, C]$ be the union of these intervals and our braid will be described by the functions f_j restricted to this interval.

Just as in the proof of Lemma 3.9 we will describe our Legendrian Λ_K in two steps. We begin, using the notation from Lemma 3.9 and its proof, by considering the standard trivial braid $\{f_1, \dots, f_n\}$ given by

$$f_i(u) = \psi_i(u)(h_t(i), i),$$

where $\psi_i(u) = u + k_i$ for some constant $k_i > 0$. Throughout our construction, as we alter the functions ψ_i , we will always assume that $i\psi_i$ has slope in the interval $[i - 1/2, i + 1/2)$. We will also assume the f_i are exactly equal to $(u + k_i)(h_t(i), i)$ near the endpoints of each of the braid intervals, though the k_i will depend on the particular braid interval. Before we perform any braiding the unstable manifolds $W^u(a_{ij})$ of any Reeb chord a_{ij} intersect $\{u\} \times S^1$ near $t = \frac{\pi}{2}$ if $i > j$ or $t = \frac{3\pi}{2}$ if $i < j$ and are lexicographically ordered as in Lemma 3.9. As we inductively build up our braid we will assume that these unstable manifolds have this same property at the boundary of all the I_l . (We will see in the construction that we can make them as near as we like.)

We will now focus on the braid interval associated to the braid generator σ_k ; the case of σ_k^{-1} is completely similar and is treated at the end of this subsection. We reset our coordinate u so that the braid interval for σ_k is $[0, U]$. Recall the braid interval consists of three subintervals: the preparatory interval I_p , the twist interval I_t and the concentration interval I_c . In the interval I_p we will alter the slopes of the curves $i\phi_i(u)$. In particular, we alter the slopes of the strands over the interval so that near the upper endpoint of the interval we have that the difference of the slopes of the k^{th} and $(k + 1)^{\text{st}}$ strands is constant and very small (that is each slope is near $k + 1/2$) and the slope of the i^{th} strand is i . Thus the difference in the slopes of the functions f_i and f_k is greater than 1 whereas the slope of the difference function $f_{k+1} - f_k$ is arbitrarily small. Allowing u to increase sufficiently, we can assume that $|f_j(u) - f_k(u)|$ and $|f_j(u) - f_{k+1}(u)|$ are sufficiently large compared to $|f_{k+1}(u) - f_k(u)|$ for each $j \neq k, k + 1$ so that a certain approximation described below is valid. This completes the description of the braid in I_p . Notice that none of these alterations affect the unstable manifold of the a_{ij} .

(Here, as below, it might be useful to consider the situation when $h_t(y)$ is zero or constant in y , and then notice that perturbing it slightly to another function $h_t(y)$ does

not affect the qualitative behavior of the flow. Here this is clear since the unstable manifolds stay far away from the regions near $t = 0$ and π where $h_t(y)$ actually depends on t . Also keep in mind that $h_t(y)$ can be taken to be arbitrarily small.)

Now consider the twist region I_t for a braid generator σ_k that interchanges the k^{th} and $(k + 1)^{\text{st}}$ strands by a π rotation of the line segment between them around its midpoint in the positive direction as s increases. This means that the strands are also interchanged by a rotation in the positive direction as u increases. For the standard trivial braid under consideration we let $f_k(u)$ and $f_{k+1}(u)$ rotate at a fast rate around the midpoint between $f_k(u)$ and $f_{k+1}(u)$ while still moving slowly away from each other.

For the functions f_k and f_{k+1} we begin by replacing $h_t(k)$ and $h_t(k + 1)$ by 0, though we leave the other f_i as they were. Specifically before we perform the twist we can assume (after possible translations) that there are constants $c < c'$ such that $f_k(u) = f_k^0(u) := (0, (k + 1/2 - \epsilon)u + c)$ and $f_{k+1}(u) = f_{k+1}^0(u) := (0, (k + 1/2 + \epsilon)u + c')$ for some very small ϵ . Now to perform the twist over the interval $[u', u'']$, choose an increasing surjective function $\beta: [u', u''] \rightarrow [0, \pi]$ that is constant near the endpoints, and let

$$f_k(u) = \frac{f_k^0(u) + f_{k+1}^0(u)}{2} + \left| \frac{f_{k+1}^0(u) - f_k^0(u)}{2} \right| (\sin \beta(u), -\cos \beta(u))$$

$$f_{k+1}(u) = \frac{f_k^0(u) + f_{k+1}^0(u)}{2} - \left| \frac{f_{k+1}^0(u) - f_k^0(u)}{2} \right| (\sin \beta(u), -\cos \beta(u)).$$

This describes the half twist between the two strands.

Note that if the distances $|f_j(u) - f_k(u)|$ and $|f_j(u) - f_{k+1}(u)|$, $j \neq k, k + 1$, are sufficiently large compared to $|f_k(u) - f_{k+1}(u)|$ then the gradient flows of $\pm F_{jk}$ and $\pm F_{j k+1}$ can be made arbitrarily close to the corresponding flows for the standard trivial braid, ie, the same braid but with non-rotating $f_k(u)$ and $f_{k+1}(u)$. We assume that in the interval I_p we arranged that the other points are sufficiently far away from $f_k(u)$ and $f_{k+1}(u)$ so that these approximations of the gradient flows of $\pm F_{jk}$ and $\pm F_{j k+1}$ are valid.

Consider next the gradient flow of $\pm F_{k+1 k}$, which, in contrast to the flows just discussed, changes drastically. We take the rotation to be supported in a small subinterval $[u', u'']$ of $[0, U]$. We have

$$F_{k+1 k}(u, t) = (2\epsilon u + c' - c)(\sin(\beta(u)), \cos(\beta(u))) \cdot (\cos t, \sin t)$$

$$= (2\epsilon u + c' - c) \sin(t + \beta(u)).$$

The gradient is

$$\begin{aligned} \nabla F_{k+1k} = & (2\epsilon \sin(t + \beta(u)) + (2\epsilon u + c' - c) \frac{d\beta}{du} \cos(t + \beta(u))) \partial_u \\ & + (2\epsilon u + c' - c) \cos(t + \beta(u)) \partial_t. \end{aligned}$$

In order to understand relevant negative gradient flow lines of this vector field we first note that F_{k+1k} is positive for $-\beta(u) < t < \pi - \beta(u)$ and negative in the complementary region. Moreover, in the limit where $\epsilon = 0$, the gradient flow of F_{k+1k} is perpendicular to the level sets $\{t + \beta(u) = a\}$ and flowing towards the curve of critical points $\{t + \beta(u) = \pi/2\}$.

Now instead let $\epsilon > 0$ be small, and choose $a_m < \pi/2 < a_M$ so that $a_M - a_m$ is small. Notice ∇F_{k+1k} is still transverse to $\{t + \beta(u) = a_m\}$ and $\{t + \beta(u) = a_M\}$ and pointing into the region R bounded by these level sets and containing $\{t + \beta(u) = \pi/2\}$. Along the curve $I = R \cap \{u = u'\}$, the u -component of ∇F_{k+1k} is positive and the flow is into R , while along the curve $I' = R \cap \{u = u''\}$, the u -component of ∇F_{k+1k} is also positive but now flows out of R . Thus since there are no critical points of ∇F_{k+1k} inside R , we see that any flow line starting on I will exit R along I' . It follows that the flow lines in R are approximately equal to the level sets $\{t + \beta(u) = a\}$ inside R .

From this we see that the unstable manifold $W^u(a_{k+1k})$ exits the twist region near $3\pi/2$ and by choosing ϵ small in the region I_p (which can be done without affecting any essential feature mentioned above) we can arrange that the t -coordinate of $W^u(a_{k+1k}) \cap \{u = u''\}$ is as close to $3\pi/2$ as we like. Furthermore, we can choose small intervals J'' and J' in the slices $\{u = u''\}$ and $\{u = u'\}$, containing the respective intersection points of $W^u(a_{k+1k})$ with these slices, such that any flow line of $-\nabla F_{k+1k}$ that starts in J'' leaves through J' and is transverse to the flow lines of the other F_{ij} . We can also then choose intervals J''_{ij} and J'_{ij} in the slices $\{u = u''\}$ and $\{u = u'\}$, containing the respective intersection points of $W^u(a_{ij})$ with these slices, so that all of these intervals are disjoint from each other and from J'' and J' , and so that all flow lines of $-\nabla F_{ij}$ that start in J''_{ij} are disjoint from each other and transversely intersect the flow lines of $-\nabla F_{k+1k}$ that begin in J'' .

The above discussion assumes that $h_t(k) = h_t(k+1) = 0$, whereas in fact $h_t(y)$ is a small function that is constant in t outside of small neighborhoods of 0 and π (see the proof of Lemma 3.9). In the argument above, the same qualitative features of the flow hold if we change $h_t(k)$ and $h_t(k+1)$ from 0 to constants in t , since the level sets of F_{k+1k} only change slightly. If now we choose $h_t(y)$ to be the function from the proof of Lemma 3.9, as is needed to define the trivial braid, then the flow only changes near $t = 0$ and $t = \pi$. But here by taking β to have very large derivative over most




Figure 19: The twist interval. The shaded regions are where F_{k+1k} is negative. The dark lines are level curves of $t + \beta(u)$ and the thinner lines are approximate flow lines of the negative gradient flow of F_{k+1k} . The left-hand end is $u = u'$ and the right-hand end is $u = u''$.

of its support we see that the alteration to the gradient of F_{k+1k} in the t -support of $h_t(y)$ can be thought of as arbitrarily small. Thus again we see that the qualitative features of the flow are unchanged.

Analogously, $W^u(a_{kk+1})$ lies close to the curve $\{t = \frac{3\pi}{2} - \beta(u)\}$ and we can argue for the same intervals J' and J'' here too. This completes the discussion of the twist interval I_t for σ_k .

Now for the concentration interval I_c . The purpose of this interval is to concentrate gradient flow lines near the unstable manifolds. Specifically, in the first part of I_c we alter our ψ_i so they are the standard affine functions again. Notice that in the region where F_{ij} is positive all the flow lines of ∇F_{ij} converge towards a constant t line near $t = \pi/2$ if $i > j$ or $3\pi/2$ if $i < j$. Thus choosing the interval $I_c = [a, b]$ large enough we can find intervals $J_{\pi/2}$ and $J_{3\pi/2}$ on $\{u = a\}$ and for each (i, j) intervals J_{ij} on $\{u = b\}$ that are an arbitrarily small neighborhood of $W^u(a_{ij}) \cap \{u = b\}$, so that all the unstable manifolds intersect $J_{\pi/2}$ or $J_{3\pi/2}$ and any flow line of ∇F_{ij} , $i > j$, that starts on $J_{\pi/2}$ intersects J_{ij} and if $i < j$ then a flow line that starts in $J_{3\pi/2}$ intersects J_{ij} .

Similarly, we consider the inverse σ_k^{-1} of the braid generator σ_k that interchanges the k^{th} and $(k + 1)^{\text{th}}$ strands by a rotation of magnitude π of the line segment between them around its midpoint in the negative direction as s increases. The analysis of this situation is exactly as above, except the unstable manifolds veer up instead of down. See Figure 11 above (where we have returned to s coordinates).

4 Combinatorial computation of the differential

In this section we compute the differential in the Legendrian algebra of Λ_K , where $K \subset \mathbb{R}^3$ is a link braided around the unknot. Our computation heavily uses the results of Section 3, where we determined all flow trees of Λ_K viewed as a Legendrian submanifold of $J^1(\Lambda)$, where $\Lambda \approx T^2$ is the conormal lift of the unknot. These flow trees give the differential for a subalgebra of the Legendrian DGA of $\Lambda_K \subset J^1(S^2)$.

In Section 4.1, we introduce the notion of a multiscale flow tree of $\Lambda_K \subset J^1(S^2)$, which is essentially a collection of partial flow trees for $\Lambda_K \subset J^1(\Lambda)$, glued to a flow tree of $\Lambda \subset J^1(S^2)$. By a result (Theorem 4.3) whose proof is deferred to Section 5, there is a one-to-one correspondence between rigid multiscale flow trees and rigid holomorphic disks with boundary on Λ_K and one positive puncture. This allows us to reduce the computation of the Legendrian DGA of $\Lambda_K \subset J^1(S^2)$ to a combinatorial enumeration of all rigid multiscale flow trees. The enumeration is performed in Sections 4.2 through 4.4 (for signs associated to multiscale flow trees, we use some results whose proofs are postponed to Section 6) and completes the proof of the main theorem of this paper, Theorem 1.1.

4.1 Multiscale flow trees

We begin by discussing multiscale trees. We first recall the basic notation that will be used in this section. Let $U \subset \mathbb{R}^3$ denote the unknot and write $\Lambda = \Lambda_U \subset J^1(S^2)$. Let K be a link given by the closure of an n -strand braid around U such that each (local) strand is C^2 -close to U . Then $\Lambda_K \subset J^1(\Lambda) \subset J^1(S^2)$, and we have the front projection $\Pi_F^\Lambda: J^1(\Lambda) \rightarrow \Lambda \times \mathbb{R}$ and the base projection $\pi^\Lambda: J^1(\Lambda) \rightarrow \Lambda$. The latter induces an n -fold cover $\Lambda_K \rightarrow \Lambda$; if γ is a path in Λ , then there are n distinct lifts $\tilde{\gamma}$ of γ with $\pi^\Lambda \circ \tilde{\gamma} = \gamma$, which we call *neighborhood lifts* of γ .

If Γ is a flow tree of $\Lambda \subset J^1(S^2)$, let $\tilde{\Gamma}$ denote its 1-jet lift.

Definition 4.1 A multiscale flow tree Γ_Δ on Λ determined by Λ_K is a flow tree Γ of $\Lambda \subset J^1(S^2)$ and a finite set of partial flow trees $\Delta = \{\Delta_j\}_{j=1}^m$ of $\Lambda_K \subset J^1(\Lambda)$ each with exactly one special puncture x_j , $j = 1, \dots, m$, such that the following holds.

- $x_j \in \tilde{\gamma}$, $j = 1, \dots, m$.
- For each component of $\tilde{\gamma} - \{x_1, \dots, x_m\}$ a neighborhood lift to Λ_K in $J^1(T^2)$ is specified.
- The union of the 1-jet lifts of the flow trees Δ_j , $j = 1, \dots, m$, and the specified neighborhood lifts, together with their flow orientation, gives a collection of consistently oriented curves $\hat{\Gamma} \subset \Lambda_K$.

- The curve $\widehat{\Gamma}$ is such that $\Pi^\Lambda(\widehat{\Gamma})$ is closed, where $\Pi^\Lambda: J^1(\Lambda) \rightarrow T^*\Lambda$ is the Lagrangian projection.

Γ is called the *big tree* and Δ the *small tree* part of Γ_Δ .

Remark 4.2 As we shall see, cf Section 4.2, the partial trees Δ_j of Γ_Δ are of two types: either Δ_j is constant at one critical point b of Morse index 2, with both its positive special puncture x_j and its negative puncture lying at b , or Δ_j is non-constant with positive special puncture at x_j .

The punctures of a multiscale flow tree Γ_Δ are the punctures of the flow trees Δ_j (not including the special punctures) and the punctures of the tree Γ . We say that the chord at a positive (respectively negative) puncture of Γ connects the sheets determined by the neighborhood lift of the arc oriented toward (respectively away from) the puncture, to the sheet determined by the neighborhood lift of the arc oriented away from (respectively toward) the puncture. A puncture of a multiscale flow tree is positive (respectively negative) if the corresponding puncture of the flow tree Γ or Δ_j is positive (respectively negative). A straightforward application of Stokes' theorem shows that every multiscale flow tree has at least one positive puncture.

Define the formal dimension of a multiscale flow tree Γ_Δ as

$$\dim(\Gamma_\Delta) = \dim(\Gamma) + \sum_{\Delta_j \in \Delta} (\dim(\Delta_j) - 1),$$

where the (formal) dimension of a (partial) flow tree is given in (2-3); see also [10, Definition 3.4].

We say that a multiscale flow tree Γ_Δ is *rigid* if $\dim(\Gamma_\Delta) = 0$ and if it is transversely cut out by its defining equation.

As discussed above in Lemma 3.7, the set of Reeb chords $\mathcal{Q}(\Lambda_K)$ of $\Lambda_K \subset J^1(S^2)$ can be written as follows:

$$\mathcal{Q}(\Lambda_K) = \mathcal{Q}^\Lambda(\Lambda_K) \cup \bigcup_{1 \leq i, j \leq n} \mathcal{Q}(\Lambda)_{ij},$$

where $\mathcal{Q}^\Lambda(\Lambda_K)$ is the set of short Reeb chords (contained in $J^1(\Lambda) \subset J^1(S^2)$) and $\mathcal{Q}(\Lambda)_{ij}$ denotes the set of long Reeb chords of $\Lambda \subset J^1(S^2)$ with endpoint on the i^{th} sheet of Λ_K and beginning point on the j^{th} sheet.

Theorem 4.3 For any $\epsilon > 0$, there exists an almost complex structure J on T^*S^2 , regular with respect to holomorphic disks with one positive puncture of dimension at

most 1, such that: there is a one-to-one correspondence between rigid holomorphic disks with one positive puncture and boundary on Λ_K , and rigid multiscale flow trees on Λ determined by Λ_K with one positive puncture, and the 1-jet lift of a multiscale flow tree lies in an ϵ -neighborhood of the boundary lift of the corresponding holomorphic disk.

Theorem 4.3 is proved in Section 5, and constitutes the basic tool in our calculation of the differential in $LA(\Lambda_K)$.

Above we described the front of Λ , the conormal lift of the unknot, and its flow trees. With this established we will next classify possible multiscale flow trees determined by a braid closure (Section 4.2) and give an algorithm for the sign of such a tree (Section 4.3). In Section 4.4 we then turn to the actual calculation of the Legendrian DGA of Λ_K by explicitly computing all multiscale flow trees for Λ_K with signs.

4.2 Classification of rigid multiscale flow trees of Λ_K

To apply Theorem 4.3 to calculate the Legendrian DGA of Λ_K , we need to classify all possible rigid multiscale flow trees of $\Lambda_K \subset J^1(S^2)$. We give a rough characterization of such trees in this subsection, examine the signs associated to the trees in Section 4.3, and present the full classification in Section 4.4.

Let K be a link and assume that Λ_K satisfies Lemma 3.7. We first consider flow trees and partial flow trees for Λ_K in $J^1(\Lambda)$. Since the front of Λ_K in $J^0(\Lambda)$ does not have any singularities and since all critical points of positive differences of local defining functions are either maxima or saddle points it follows from Section 2.4 that for generic functions a rigid tree for $\Lambda_K \subset J^1(\Lambda)$ must have

- (1) a positive puncture at a Reeb chord b_{ij} of type \mathbf{S}_1 ,
- (2) $k - 1$ Y_0 -vertices (trivalent vertices away from (non-existent) cusps of Λ) and
- (3) k negative punctures at Reeb chords a_{ij} of type \mathbf{S}_0 .

Likewise, a partial flow tree of dimension 1 has the same vertices and punctures, except its positive puncture is a special puncture instead of a maximum. Also the constant partial flow tree with both special positive and negative puncture at b_{ij} of type \mathbf{S}_1 will be of importance.

Since the 1-jet lift of any flow tree of $\Lambda \subset J^1(S^2)$ has codimension 1 in Λ , it follows that any tree in the small tree part of any rigid multiscale flow tree is either a 1-dimensional partial flow tree with positive puncture on the 1-jet lift of the big tree on Λ or it is a constant tree at some b_{ij} of type \mathbf{S}_1 . Furthermore, the flow tree on

Λ corresponding to the big tree part must either be rigid, or rigidified by a constant tree (a point condition at some $\pi(b_{ij})$). In the case where the big tree is rigidified by point conditions, its dimension must be equal to the number of point conditions, ie, the number of constant trees in the multiscale flow tree. Combining this discussion with Lemma 3.2 we find that (after small perturbation) there are the following types of rigid multiscale flow trees for Λ determined by Λ_K :

- MT₀ A rigid flow tree Γ of Λ with constrained rigid partial flow trees of Λ_K attached. Here the constraint says that the special positive puncture of each partial flow tree must lie on the 1–jet lift of Γ .
- MT₁ A constrained rigid flow tree Γ^* with constrained rigid partial flow trees of Λ_K attached. Here the constraint of Γ^* is the requirement that the 1–jet lift passes a point in Λ in the fiber where a Reeb chord b_{ij} lies, and the constraint of the partial flow trees is as above.
- MT_∅ A rigid flow tree of Λ_K .

Remark 4.4 In our setting, we can rule out one other ostensible possibility for a rigid multiscale tree: those with big tree a constant rigid flow tree Γ of Λ and with small tree a constrained rigid flow tree of Λ_K . Here Γ would correspond to a Reeb chord and the constraint would say that the special positive puncture of the partial flow tree must lie on the Reeb chord on the 1–jet lift of Γ . If the location of the Reeb chord is generic with respect to the flow determined by Λ_K then its endpoints do not lie on $W^u(a_{ij})$ for any a_{ij} or on b_{ij} and such a configuration is rigid only if the flow line ends at a minimum. As there are no positive local function differences of Λ_K that are local minima no such trees correspond to disks with one positive puncture. (The rigid configurations with flow lines that end at a negative local minimum correspond to disks with two positive punctures.)

4.3 Signs of rigid multiscale flow trees

In this subsection we describe a combinatorial algorithm for computing the sign of a rigid multiscale flow tree, which determines its contribution to the Legendrian algebra differential. This is the analogue for multiscale flow trees in the discussion in Section 3.4.4. We will discuss the derivation of the combinatorial rule as well as the effect of orientation choices in detail in Section 6.

We will use the notation established in Section 2.5 for vector splitting along flow trees and signs associated to rigid flow trees of Λ_K as well as partial flow trees of Λ_K of dimension 1 with special positive puncture.

Before we can state the combinatorial rule for orienting rigid multiscale flow trees, we need to discuss signs of rigid trees determined by Λ ; see Sections 3.1 and 3.2 for the notation for these rigid trees. Except for basic orientation choices the signs depend on orientations of determinants of capping operators. We call such choices *capping orientations*. Recall that there are two Reeb chords of Λ , e and c , and that if $K \subset \mathbb{R}^3$ is a link represented as a closed braid on n strands, then the long Reeb chords of Λ_K are e_{ij} and c_{ij} , $1 \leq i, j \leq n$, where c_{ij} lies very close to c and e_{ij} lies very close to e . In particular, capping orientations for c and e induce capping orientations for c_{ij} and e_{ij} , respectively.

Theorem 4.5 *There is a basic orientation choice and choice of capping orientation for c so that the sign $\epsilon(T)$, for T a rigid flow tree of Λ , satisfies*

$$\epsilon(I_N) = \epsilon(Y_N) = \epsilon(I_S) = \epsilon(Y_S) = 1 \quad \text{and} \quad \epsilon(E_1) = -\epsilon(E_2).$$

Proof This is a consequence of Theorem 6.4 below. \square

Furthermore, the choice of capping orientation of e induces an orientation of the 1-dimensional moduli spaces of flow trees such that the induced orientation at the broken disk $E_j \# T$ is $\epsilon(E_j)\epsilon(T)$ for $T \in \{I_N, Y_N, I_S, Y_S\}$. If Γ is a flow tree in such a 1-dimensional moduli space we consider the orientation as a normal vector field v along the 1-jet lift of Γ .

In order to state the sign rule for multiscale flow trees, we first make some definitions.

MT₁ Consider a multiscale rigid flow tree Θ of type **MT₁** with 1-dimensional big tree Γ and a negative puncture at b_{ij} . Let $v^{\text{flow}}(\Gamma)$ denote the vector field along the 1-jet lift oriented in the positive direction (ie, the 1-jet lift of each edge is equipped with the flow-orientation, defined in Section 2.4). The sign $\epsilon(\Theta)$ of the rigid tree constrained by b_{ij} of type **S₁** is defined to be

$$\epsilon(\Theta) = \text{sign}(\langle v, v^{\text{ker}}(b_{ij}) \rangle \langle v^{\text{flow}}(\Gamma), v^{\text{coker}}(b_{ij}) \rangle).$$

If non-constant flow trees $\Gamma_1, \dots, \Gamma_k$ are attached to Θ then define n_j to be the normal vector at the special puncture of Γ_j with positive inner product with $v^{\text{flow}}(\Gamma)$.

MT₀ Consider a multiscale rigid flow tree Θ of type **MT₀** with big tree Γ . Let $\epsilon(\Theta)$ equal the sign of Γ . If Γ has two punctures (positive at e , negative at c), then we define $v^{\text{flow}}(\Gamma)$ as the vector field along the boundary pointing toward the positive puncture. If Γ has only one puncture (positive at c), let $v^{\text{flow}}(\Gamma)$ point in the positive direction along the boundary. Then take n_j as in **MT₁**.

MT_∅ Consider a flow tree Γ of Λ_K in $J^1(\Lambda)$. Let $n = v^{\text{coker}}(b_{ij})$ be the normal vector of Γ at its positive puncture.

Theorem 4.6 *There exists a choice of basic orientations and of orientations of capping operators for all long Reeb chords such that Theorem 4.5 holds and such that the sign of a rigid multiscale flow tree is as follows.*

MT₁ *Let Γ be a multiscale rigid flow tree of type MT_1 with constrained rigid flow tree Θ (ie, Θ has only one negative puncture at some b_{ij}) and attached flow trees $\Gamma_1, \dots, \Gamma_k$. Then the sign of Γ is*

$$\epsilon(\Gamma) = \epsilon(\Theta) \prod_{j=1}^k \sigma(n_j, \Gamma_j).$$

MT₀ *Let Γ be a multiscale rigid flow tree of type MT_0 with rigid flow tree Θ and attached flow trees $\Gamma_1, \dots, \Gamma_k$. Then the sign of Γ is*

$$\epsilon(\Gamma) = \epsilon(\Theta) \prod_{j=1}^k \sigma(n_j, \Gamma_j).$$

MT_∅ *Let Γ be a flow tree of type MT_\emptyset . Then the sign of Γ is*

$$\sigma_{\text{pos}}(\Gamma) \sigma(n, \Gamma).$$

Proof Theorem 4.6 is proved in Section 6.6. □

4.4 Counting multiscale flow trees

In this section we complete the computation of the Legendrian algebra differential of $\Lambda_K \subset J^1(S^2)$, and thereby obtain a proof of Theorem 1.1, by counting all multiscale flow trees determined by Λ_K and Λ . In Section 3, we counted flow trees of $\Lambda_K \subset J^1(\Lambda)$. This leads to the expression for $\partial\mathbb{B}$ in Theorem 1.1. In this subsection we derive the expression for $\partial\mathbb{C}$ and $\partial\mathbb{E}$ in Theorem 1.1 by counting multiscale flow trees with non-empty big tree part. Our technique relates these multiscale trees to ordinary flow trees of a stabilized braid obtained by adding a trivial non-interacting strand to the given braid.

For notation used throughout this section, see Section 3.4.

4.4.1 Multiscale flow trees of type MT_\emptyset We first consider the part of the differential in the Legendrian algebra of Λ_K that accounts for multiscale flow trees of type MT_\emptyset , ie, the parts which count only trees of the braid localized near Λ . Such a tree has its positive puncture at some Reeb chord b_{ij} and its negative punctures at Reeb chords a_{ij} . Furthermore, a straightforward action argument shows that any multiscale flow tree with its positive puncture at a Reeb chord b_{ij} must lie inside the 1-jet neighborhood

of Λ . Consequently, flow trees of type MT_\emptyset account for the boundary of the Reeb chords b_{ij} of type S_1 :

$$\partial(b_{ij}) = \sum_{\Gamma \in \mathcal{T}(b_{ij})} \epsilon(\Gamma)q(\Gamma),$$

where $\mathcal{T}(b_{ij})$ denotes the set of all flow trees with positive puncture at b_{ij} and where, if Γ is such a tree, $q(\Gamma)$ denotes the monomial of its negative punctures and $\epsilon(\Gamma)$ its sign.

In Section 3.4.8 above, orientation conventions were picked for the Reeb chords a_{ij} and b_{ij} and the differential was computed in Lemma 3.21.

4.4.2 Multiscale flow trees with positive puncture at c_{ij} of type L_1 In this section, we compute the differentials of the c_{ij} Reeb chords. This involves counting rigid multiscale flow trees of type MT_0 , with big-tree component given by one of the four rigid big trees with positive puncture at c .

We introduce the following notation for 1–jet lifts of flow trees with one positive puncture. Each point in the 1–jet lift that is neither a 1–valent vertex nor a trivalent vertex belongs to either the upper or the lower sheet of its flow line. We call points on the 1–jet lift *head-points* if they belong to the upper sheet and *tail-points* if they belong to the lower. Because of positivity of local function differences of trees with one positive puncture, points in a component of the complement of preimages of 1–valent and trivalent vertices in the 1–jet lifts are either all head-points or all tail-points.

There are four rigid flow trees with positive puncture at c , denoted in Section 3.2 by I_N, Y_N, I_S, Y_S . When we project the 1–jet lifts of these trees to the torus Λ , we obtain the four curves $\Gamma_{\alpha\beta}(c)$, $\alpha, \beta \in \{0, 1\}$, shown in Figure 20. We decompose $\Gamma_{\alpha\beta}(c)$ as follows:

$$\Gamma_{\alpha\beta}(c) = \Gamma_{\alpha\beta}^{v+}(c) \cup \Gamma_{\alpha\beta}^h(c) \cup \Gamma_{\alpha\beta}^{v-}(c)$$

See Figure 20 for $\Gamma_{00}(c)$. More precisely, $\Gamma_{\alpha\beta}^{v+}(c)$ consists of head-points, $\Gamma_{\alpha\beta}^{v-}(c)$ consists of tail-points, and $\Gamma_{\alpha\beta}^h(c)$ is the portion near the cusp edge. For our purposes, we will assume that $\Gamma_{\alpha\beta}^h(c)$ lies at $t = \frac{\pi}{2}$ (for $\beta = 0$) or $t = \frac{3\pi}{2}$ (for $\beta = 1$), as is the case in the degenerate picture where U is the unperturbed round unknot (see Section 3.1).

The following lemma determines the differential acting on Reeb chords c_{ij} . We use the matrix notation of Theorem 1.1.

Lemma 4.7 *With capping operator of Reeb chord c of Λ as in Theorem 4.6 and with orientation choices as in (3-14) and (3-15), the following equation holds:*

$$\partial\mathbb{C} = \mathbb{A} \cdot \mathbb{A} + \mathbb{A} \cdot \Phi_B^R.$$




Figure 20: The 1-jet lifts of flow trees with positive puncture at c

Proof We count multiscale flow trees contributing to ∂c_{ij} , divided into four cases based on their big tree part, which must be one of the $\Gamma_{\alpha\beta}(c)$. For ease of reference, all of these flow trees are pictorially represented in Table 1.

Before we proceed, note that the braiding region is disjoint from $\Gamma_{10}(c)$ and from $\Gamma_{11}(c)$ and intersects $\Gamma_{00}(c)$ and $\Gamma_{01}(c)$ in tail-points, since the braiding region has small s -coordinate values. Also, from Section 3.4, the unstable manifolds $W^u(a_{ij})$ are essentially horizontal and have t -coordinate as follows: just less than $\frac{\pi}{2}$ for $i > j$, ordered lexicographically by (i, j) , and just over $\frac{3\pi}{2}$ for $i < j$, ordered lexicographically by (j, i) . Finally, we recall from Section 3.4.3 and Figure 13 that we have cycles λ' and μ' in Λ for the purposes of counting homology classes, where λ' is a vertical line between the a_{ij} and the b_{ij} , and μ' is a horizontal line just below $t = \frac{\pi}{2}$. We can in particular choose μ' to lie above all of the $W^u(a_{ij})$ for $i > j$.

Case 1 Big tree $\Gamma_{10}(c)$.

If $i = j$, there is a “trivial” multiscale flow tree with boundary on sheet S_i that projects to $\Gamma_{10}(c)$. To count other multiscale flow trees corresponding to $\Gamma_{10}(c)$, note that the curve $\Gamma_{10}^{v+}(c)$ intersects $W^u(a_{ij})$ for all $i > j$, while the curves $\Gamma_{10}^h(c)$ and $\Gamma_{10}^{v-}(c)$ are disjoint from $W^u(a_{ij})$ for all i, j . As we move along $\Gamma_{10}^{v+}(c)$ in the sheet S_i , at the intersection between $\Gamma_{10}^{v+}(c)$ and $W^u(a_{ij})$, $i > j$, a flow line to a_{ij} can split off and the 1-jet lift of $\Gamma_{10}(c)$ continues to move along sheet S_j . Note that no other flow









Big tree	Multiscale flow tree	Term
$\Gamma_{10}(c)$	 $(i = j)$	$\mu_{\alpha(i)}$
	 $(i > j)$	$a_{ij} \mu_{\alpha(j)}$
$\Gamma_{11}(c)$	 $(i = j)$	1
	 $(i < j)$	a_{ij}
$\Gamma_{00}(c)$	 $\mu_{\alpha(i)} (\Phi_B^R)_{ij}$	
	 $\sum_{k < i} a_{ik} \mu_{\alpha(k)} (\Phi_B^R)_{kj}$	
$\Gamma_{01}(c)$	 $(\Phi_B^R)_{ij}$	
	 $\sum_{k > i} a_{ik} (\Phi_B^R)_{kj}$	

Table 1: Contributions to ∂c_{ij} . For each of the four big disks, schematic diagrams for corresponding multiscale flow trees are shown, along with the algebraic contribution to ∂c_{ij} (with powers of the longitudinal homology classes λ suppressed for simplicity). In the diagrams, along the boundary of the big disk, the index of the sheet (one of i, j, k) is labeled.

line can split off after this event since, according to the orientation requirement in the definition of multiscale flow tree, such a flow line could split off only at an intersection with $W^u(a_{jk})$, $j > k$; however, by our choice of ordering, such an intersection precedes the intersection with $W^u(a_{ij})$ and hence no further splitting is possible.

Thus for each c_{ij} with $i \geq j$, there is a multiscale flow tree with big tree $\Gamma_{10}(c)$ and a single negative puncture at a_{ij} (or no negative puncture if $i = j$). We next consider

homology coefficients and signs. Note that each tree which is lifted to a leading sheet $S_{\gamma(i)}$ intersects $\lambda'_{\gamma(i)}$ once positively and that each lifted tree intersects $\mu'_{\alpha(i)}$, where $\gamma(i)$ and $\alpha(i)$ are the indices of the components of the sheets considered. Furthermore, by Theorem 4.6, each tree has sign $+1$: at the splitting point (for $i > j$), the tangent vector to the positively oriented 1-jet lift is ∂_t , which transports to $\partial_t = v^{\ker}(a_{ij})$ at a_{ij} , $i > j$, and the big tree has sign $+1$. Finally, we conclude that the contribution to ∂c_{ij} of Case 1 is:

$$\begin{cases} \lambda_{\gamma(i)}\mu_{\alpha(i)} & \text{if } i = j \text{ and } S_i \text{ is leading,} \\ \mu_{\alpha(i)} & \text{if } i = j \text{ and } S_i \text{ is not leading,} \\ a_{ij}\lambda_{\gamma(j)}\mu_{\alpha(j)} & \text{if } i > j \text{ and } S_j \text{ is leading,} \\ a_{ij}\mu_{\alpha(j)} & \text{if } i > j \text{ and } S_j \text{ is not leading.} \end{cases}$$

Case 2 Big tree $\Gamma_{11}(c)$.

As with $\Gamma_{10}(c)$, if $i = j$ then there is a trivial multiscale tree with boundary on sheet S_i that projects to $\Gamma_{10}(c)$. To count other multiscale trees, notice that $\Gamma_{11}^{v+}(c)$ intersects $W^u(a_{ji})$, for all $i > j$. As above we find that a 1-jet lift in sheet S_j can split off a flow line to a_{jk} for $k > j$, then continue along S_k , and that no further splittings are possible. We thus find multiscale flow trees with positive puncture at c_{ji} and negative puncture at a_{ji} for all $i > j$. As above the (signed) coefficient equals $+1$ since the normal at the special puncture of the partial tree attached is $-\partial_t$, which agrees with $v^{\ker}(a_{ij})$, $i < j$. To see homology coefficients, we note that all 1-jet lifts are disjoint from μ' and calculate as above. We conclude that the contribution to ∂c_{ij} from Case 2 is:

$$\begin{cases} \lambda_{\gamma(j)} & \text{if } i = j \text{ and } S_j \text{ is leading,} \\ 1 & \text{if } i = j \text{ and } S_j \text{ is not leading,} \\ a_{ij}\lambda_{\gamma(j)} & \text{if } i < j \text{ and } S_j \text{ is leading,} \\ a_{ij} & \text{if } i < j \text{ and } S_j \text{ is not leading.} \end{cases}$$

When we combine the contributions of Cases 1 and 2 to ∂c_{ij} , we obtain precisely the (i, j) entry of the matrix $\mathbb{A} \cdot \mathbb{L}$.

Case 3 Big tree $\Gamma_{00}(c)$.

Here we will make use of the stabilized braid \widehat{B} , which is B along with one non-interacting strand labeled 0. Any multiscale flow tree begins on sheet S_i along $\Gamma_{00}^{v+}(c)$. As above it is possible to split off either no flows and arrive on S_i at the initial point of $\Gamma_{00}^h(c)$, or one flow to a_{ik} (for any $k < i$) and arrive on S_k at the initial point of $\Gamma_{00}^h(c)$. Note further that $\Gamma_{00}^h(c)$ intersects the braiding region in tail-points since the

braiding region is to the left of $s = \pi/2$; thus, its flow orientation is the same as the oriented lift of the flow line of \widehat{B} in $W^u(a_{0k})$ for any k . Consequently, the flow trees that can split off from $\Gamma_{00}^h(c)$ on sheet S_k as it passes the braiding region agree with the flow trees that can split off from $W^u(a_{0k})$. It follows that the contribution to ∂c_{ij} from Case 3, up to sign and homology coefficients, is

$$\sum_{k < i} a_{ik} (\Phi_B^R)_{kj}.$$

We next consider the homology coefficients and signs of these trees. For homology, note that μ' is a horizontal line that lies just below $t = \frac{\pi}{2}$, where $\Gamma_{00}^h(c)$ sits, but just above $W^u(a_{ik})$, where a flow line to a_{ik} can split off. Thus none of the trees pass μ' ; a tree that splits off a flow to a_{ik} intersects $\mu'_{\alpha(k)}$ once positively, and a tree that does not split off such a flow intersects $\mu'_{\alpha(j)}$ once positively.

In order to compute the sign we first note that the sign contribution from the flow line splitting off to a_{ik} is positive: ∂_t is the vector of the boundary orientation as well as $v^{\ker}(a_{ik})$. Second, we consider sign contributions from trees in the braiding region. At a positive twist the induced normal is $-\partial_s$, which corresponds to a vector splitting of the normal ∂_t along the incoming edge and the sign at the corresponding trivalent vertex of the tree of \widehat{B} is $+1$. At a negative twist the induced normal is still $-\partial_s$, which now corresponds to a splitting of the normal $-\partial_t$ along the incoming edge and the sign at the trivalent vertex of the tree of \widehat{B} equals -1 . Thus also in the case of a negative twist the total sign contribution is $(-1)^2 = 1$, which shows that all multiscale flow trees from $\Gamma_{00}(c)$ have sign $+1$.

In sum, we find that the total contribution to ∂c_{ij} from Case 3 is

$$\mu_{\alpha(j)}(\Phi_B^R)_{ij} + \sum_{k < i} a_{ik} \mu_{\alpha(k)}(\Phi_B^R)_{kj}.$$

Case 4 Big tree $\Gamma_{01}(c)$.

As in Case 3, we find that either one or zero flow lines split off along $\Gamma_{01}^{v+}(c)$. Again Γ_{01}^h intersects the braiding region in tail points and is oriented isotopic to $W^u(a_{0i})$. An argument similar to Case 3 shows that the contribution to ∂c_{ij} from Case 4 is

$$(\Phi_B^R)_{ij} + \sum_{k > i} a_{ik} (\Phi_B^R)_{kj},$$

where the difference from the previous case arises because the 1-jet lift of $\Gamma_{01}(c)$ is disjoint from μ'_j .

When we combine the contributions of Cases 3 and 4 to ∂c_{ij} , we obtain precisely the (i, j) entry of the matrix $\mathbb{A} \cdot \Phi_B^R$. The lemma follows. \square

4.4.3 Multiscale flow trees of type MT_0 with positive puncture at e_{ij} of type L_2

To complete the computation of the Legendrian homology of Λ_K , we need to compute the differentials of the e_{ij} Reeb chords. These have contributions from two types of multiscale rigid flow trees: trees of type MT_0 and trees of type MT_1 . In this subsection, we compute the first type; the second type is computed in the following subsection, Section 4.4.4.

There are two rigid flow trees with positive puncture at e and negative puncture at c , corresponding to E_1 and E_2 in the language of Section 3.2. Denote their 1-jet lifts by $\Gamma_\alpha(e; c)$, $\alpha \in \{0, 1\}$. We decompose $\Gamma_\alpha(c)$ as follows:

$$\Gamma_\alpha(e; c) = \Gamma_\alpha^{he}(e; c) \cup \Gamma_\alpha^{ta}(e; c),$$

where $\Gamma_\alpha^{he}(e; c)$ consists of head-points and $\Gamma_\alpha^{ta}(e; c)$ of tail-points; see Figure 21.




Figure 21: The 1-jet lifts of flow trees with positive puncture at e and negative puncture at c

Writing the contribution to ∂e_{ij} from trees of type MT_0 as $\partial_0 e_{ij}$ and using the matrix notation of Theorem 1.1, we have the following result.

Lemma 4.8 *With capping operator of c so that Theorem 4.6 holds and with orientation choices as in (3-14) and (3-15), there is a choice of capping operator for e so that the*

following equation holds:

$$\partial_0 \mathbb{E} = -\Phi_B^L \cdot \mathbb{C} \cdot \lambda^{-1} + \lambda^{-1} \cdot \mathbb{C} \cdot (\Phi_B^R)^{-1}.$$

Proof First consider the contributions from $\Gamma_0(e, c)$: $\Gamma_0^{\text{ta}}(e, c)$ is disjoint from the braiding region and from all $W^u(a_{ij})$. Consequently, no flow tree can split off from $\Gamma_0^{\text{ta}}(e, c)$. Also, if $\Gamma_0^{\text{ta}}(e, c)$ is lifted to S_i then it intersects $\lambda'_{\gamma(i)}$ once negatively if S_i is leading and not at all otherwise. The curve of head-points $\Gamma_0^{\text{hc}}(e, c)$, oriented as in Figure 21, is isotopic without crossing the cycle μ' to the oriented lift of the flow line $W^u(a_{i0})$ of the stabilized braid \hat{B} which lies in the sheet S_i . As in the proof of Lemma 4.7 we conclude that the trees that split off along Γ correspond to the trees which split off from $W^u(a_{0i})$. As with $\Gamma_{00}(c)$ (Case 3) in the proof of Lemma 4.7, we find that the sign contribution from the split-off trees agrees with the sign of ϕ_B . Thus the contribution from $\Gamma_0(e, c)$ to $\partial_0 \mathbb{E}$ is

$$\epsilon(\Gamma_0(e; c)) \Phi_B^L \cdot \mathbb{C} \cdot \lambda^{-1}.$$

Choose the capping operator of e so that $\epsilon(\Gamma_0(e; c)) = -1$, and note that by Theorem 4.5 this implies $\epsilon(\Gamma_1(e; c)) = +1$.

Next consider the contributions from $\Gamma_1(e, c)$: $\Gamma_1^{\text{hc}}(e, c)$ is disjoint from the braiding region and from all $W^u(a_{ij})$. Consequently, no flow tree can split off from $\Gamma_1^{\text{hc}}(e, c)$. Also, if $\Gamma_1^{\text{hc}}(e, c)$ is lifted to S_i then it intersects $\lambda'_{\gamma(i)}$ once with negative intersection number if S_i is leading and not at all otherwise. The tail-points curve $\Gamma_1^{\text{ta}}(e, c)$ is isotopic to the lift of the flow line $W^u(a_{i0})$ of the stabilized braid \hat{B} with orientation reversed. Switching the roles of e and c in the Morse–Bott perturbation we would get the following contribution to $\partial'_0 \mathbb{C}$ from this disk:

$$\lambda \cdot \mathbb{E} \cdot \Phi_B^R,$$

where ∂'_0 denotes the analogue of ∂_0 with the alternative Morse–Bott perturbation, by a repetition of the argument above for $\Gamma_0(e; c)$.

Observing that the multiscale flow trees we are interested in are exactly the same as those for the alternative Morse–Bott perturbation except for the big disk having the opposite orientation, we can view the equation above as a linear system of equations with coefficients in \mathcal{A}_n^0 and invert it to get the contribution to $\partial_0 \mathbb{E}$. Thus we find that the contribution from $\Gamma_1(e, c)$ is

$$\epsilon(\Gamma_0(e; c)) \lambda^{-1} \cdot \mathbb{C} \cdot (\Phi_B^R)^{-1} = \lambda^{-1} \cdot \mathbb{C} \cdot (\Phi_B^R)^{-1}$$

and the lemma holds. □

4.4.4 Multiscale flow trees of type MT_1 with positive puncture at e_{ij} Finally, we enumerate multiscale flow trees of type MT_1 . Recalling Lemma 3.2 and Remark 3.4, we see there are four constrained rigid flow trees with positive puncture at e and no negative punctures that are constrained to pass through some b_{ij} . We denote their 1-jet lifts $\Gamma_{\alpha\beta}(e)$, $\alpha, \beta \in \{0, 1\}$; see Figures 22 and 23. We note that each constrained tree is a deformation of a broken tree, with $\Gamma_{\alpha\beta}(c)$ as defined in Section 4.4.2 and $\Gamma_\alpha(e; c)$ as defined in Section 4.4.3:

$$\begin{aligned} \Gamma_{00}(e) &\simeq \Gamma_1(e; c) \# \Gamma_{00}(c), \\ \Gamma_{10}(e) &\simeq \Gamma_1(e; c) \# \Gamma_{10}(c), \\ \Gamma_{01}(e) &\simeq \Gamma_1(e; c) \# \Gamma_{01}(c), \\ \Gamma_{11}(e) &\simeq \Gamma_1(e; c) \# \Gamma_{11}(c). \end{aligned}$$

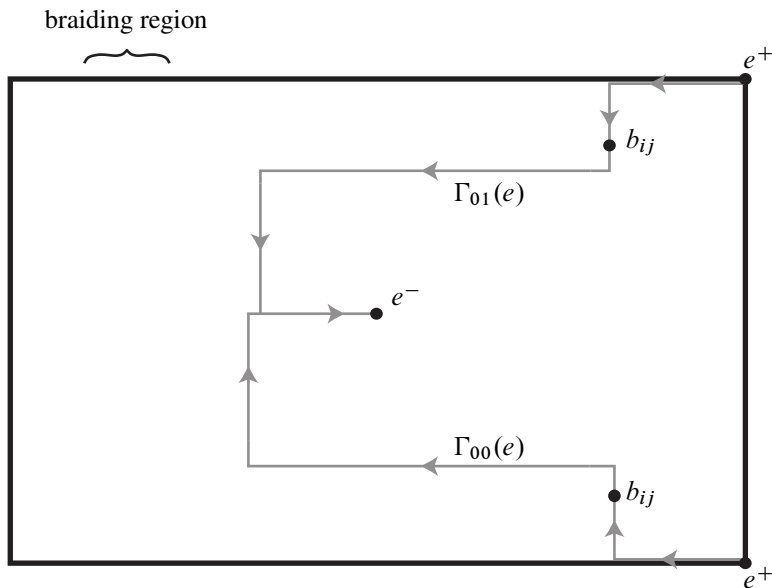


Figure 22: 1-jet lifts of constrained flow trees with positive puncture at e that are disjoint from the braiding region

Writing the contribution to ∂e_{ij} from trees of type MT_1 as $\partial_1 e_{ij}$ and using the matrix notation of Theorem 1.1, we have the following result.

Lemma 4.9 *With capping operators and orientation data as in Lemma 4.8, the following equation holds:*

$$\partial_1 \mathbb{E} = \mathbb{B} \cdot (\Phi_B^R)^{-1} + \mathbb{B} \cdot \lambda^{-1}.$$

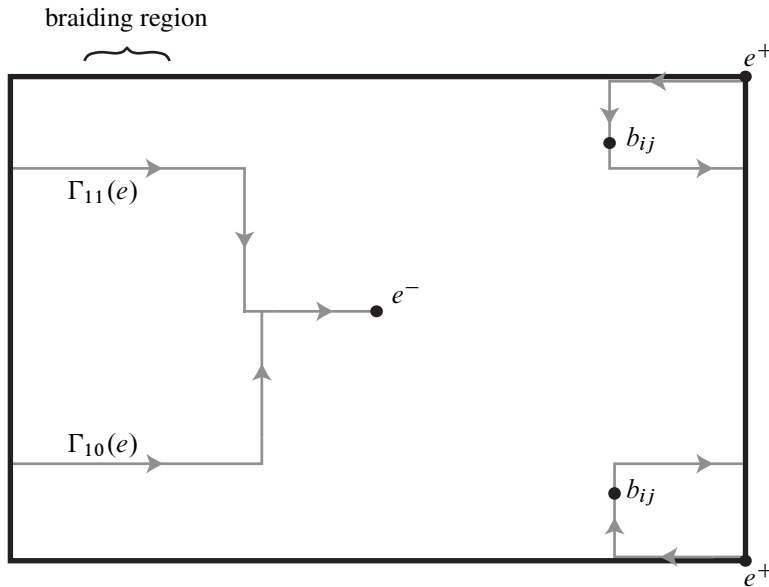


Figure 23: 1–jet lifts of constrained flow trees with positive puncture at e that intersect the braiding region. As with $\Gamma_{\alpha\beta}(c)$, the horizontal segments in the middle of $\Gamma_{10}(e)$ and $\Gamma_{11}(e)$ are at $t = \pi/2$ and $t = 3\pi/2$, respectively.

Proof Consider first the contributions of the big trees $\Gamma_{00}(e)$ and $\Gamma_{01}(c)$. Note that these trees do not intersect the braiding region and that because of the lexicographic order, exactly as in the proof of Lemma 4.7, a disk that is constrained at b_{ij} cannot split off flow lines to any a_{ik} . Thus there is exactly one multiscale flow tree contributing to $\partial_1 e_{ij}$ that arises from $\Gamma_{00}(e)$ or $\Gamma_{01}(e)$ if $i \neq j$ (from $\Gamma_{00}(e)$ if $i > j$ or $\Gamma_{01}(e)$ if $i < j$), and it begins at e^+ on sheet i , jumps at b_{ij} to sheet j , and remains there until e^- . (If $i = j$, then there is no such multiscale tree.) Up to sign and homology classes, we conclude that $\Gamma_{00}(e)$ and $\Gamma_{01}(c)$ combined contribute the term b_{ij} (or 0 if $i = j$) to ∂e_{ij} .

As for homology classes, the 1–jet lift of this multiscale tree is disjoint from all λ' cycles, with the exception of $\lambda'_{\alpha(j)}$ (one negative intersection) if j is leading; furthermore, it is disjoint from all μ' cycles, with the exception of $\mu'_{\alpha(j)}$ (one positive intersection) if $i > j$. Next we determine the signs. The sign of the endpoints of the moduli spaces of $\Gamma_{00}(e)$ and $\Gamma_{01}(e)$ listed above are both $+1$ by our choice of capping operators. Thus the orienting vector field v of the moduli space at b_{ij} satisfies $v = \partial_s$ in both cases. Furthermore, $v^{\ker}(\Gamma_{00}(e))$ equals ∂_t at b_{ij} , $i > j$, and $v^{\ker}(\Gamma_{01}(e))$ equals $-\partial_t$ at b_{ij} , $i < j$. This shows that the signs are positive. Collecting homology classes and signs, we conclude that the big trees $\Gamma_{00}(e)$ and $\Gamma_{01}(e)$ contribute $\mathbb{B} \cdot \lambda^{-1}$ to $\partial_1 \mathbb{E}$ (recall that the (i, j) entry of \mathbb{B} is $b_{ij}\mu_j$ if $i > j$, 0 if $i = j$, and b_{ij} if $i < j$).

Next consider the contributions to ∂e_{ij} from big trees $\Gamma_{10}(e)$ and $\Gamma_{11}(e)$. The 1–jet lifts of the multiscale flow trees corresponding to these big trees begin at e^+ on sheet i , switch to sheet k for some $k < i$ (for $\Gamma_{10}(e)$) or $k > i$ (for $\Gamma_{11}(e)$) at the constraint b_{ik} , and then pass through the braiding region and end at e^- on sheet j . The (horizontal) portion passing through the braiding region is located at $t = \frac{\pi}{2}$ or $t = \frac{3\pi}{2}$ and is thus isotopic to $\Gamma_1^{\text{ta}}(e; c)$ (which is at $t = \pi$) in the complement of all the unstable manifolds $W^u(a_{ij})$ (which are below $t = \frac{\pi}{2}$ or above $t = \frac{3\pi}{2}$) and the cycles μ' (which is below $t = \frac{\pi}{2}$). Thus repeating the computation in Lemma 4.8, we find that $\Gamma_{10}(e)$ and $\Gamma_{11}(e)$ combined give a contribution to $\partial_1 \mathbb{E}$ of

$$\mathbb{B} \cdot (\Phi_B^R)^{-1}.$$

Here we find that the sign of this term is $+1$ after observing that the sign of the underlying restricted rigid disk is again $+1$ by repeating the calculation above. The lemma follows. □

4.4.5 Proof of Theorem 1.1 By Theorem 4.3, the differential for the Legendrian DGA $LA(\Lambda_K)$ can be computed in terms of multiscale flow trees determined by Λ_K and Λ . The contribution from multiscale flow trees of types MT_\emptyset , MT_0 and MT_1 are calculated in Lemma 3.21, 4.7, 4.8 and 4.9. □

5 Multiscale flow trees and holomorphic disks

The main purpose of this section is to establish Theorem 4.3. The proof has several steps. In Section 5.1 we establish the connection between holomorphic disks and flow trees for Λ_U following [10]. In Section 5.2 we discuss the slightly stronger disk/flow tree correspondence needed here. Finally, in Sections 5.3 and 5.4 we establish the correspondence between holomorphic disks of Λ_K in $J^1(S^2)$ and multiscale flow trees, ie, holomorphic disks with boundary on Λ_U with flow trees of $\Lambda_K \subset J^1(\Lambda_U)$ attached along its boundary.

5.1 Basic results on constrained flow trees and disks

In this subsection we give a slight modification of results from [10] in the case of a Legendrian surface Λ in the 1–jet space of a surface S , $\Lambda \subset J^1(S)$. We will apply these results in two cases relevant to this paper; namely when $\Lambda = \Lambda_U$ and $S = S^2$, and $\Lambda = \Lambda_K$ and $S = \Lambda_U$.

For the notion of a flow tree of Λ we refer to Section 2.4 and for a more thorough account [10, Section 2.2]. As explained in [10, Section 3.1], associated to each flow

tree Γ is its formal dimension $\dim(\Gamma)$, see (2-3), which is the dimension of the manifold of nearby flow trees for sufficiently generic Λ . The main result of [10] is concerned with the relation between rigid holomorphic disks and rigid flow trees. Here we will need a slight generalization. To this end we introduce *constrained* rigid holomorphic disks and *constrained* flow trees. If p_1, \dots, p_r are distinct points in Λ then a holomorphic disk *constrained by* p_1, \dots, p_r is a holomorphic disk with r extra boundary punctures at which the evaluation map hits p_1, \dots, p_r . Similarly, a flow tree of Λ *constrained by* p_1, \dots, p_r is a flow tree with 1-jet lift with r extra marked points where the evaluation map hits p_1, \dots, p_r . If Γ' is a flow tree (or holomorphic disk) constrained by $p_1, \dots, p_r \in \Lambda$ and if Γ is that flow tree (or holomorphic disk) with the constraining conditions forgotten then the formal dimension of the constrained flow tree (or holomorphic disk) Γ' satisfies

$$\dim(\Gamma') = \dim(\Gamma) - r,$$

where $\dim(\Gamma)$ is the formal dimension of Γ . We say that a flow tree (or holomorphic disk) is *constrained rigid* if its formal dimension equals 0 and it is transversely cut out by its defining equations.

For $0 < \sigma \leq 1$ consider the map

$$s_\sigma: J^1(S) \rightarrow J^1(S), \quad s_\sigma(q, p, z) = (q, \sigma p, \sigma z)$$

where $q \in S$, $p \in T_q^*S$ and $z \in \mathbb{R}$ and write

$$\Lambda_\sigma = s_\sigma(\Lambda).$$

Notice that since s_σ preserves the contact structure, Λ_σ is still Legendrian and clearly Legendrian isotopic to Λ . In order to state the correspondence theorem relating constrained rigid disks and trees we recall the following notation: D_m denotes the unit disk with m boundary punctures and $\Pi: J^1(S) \rightarrow T^*S$ denotes the Lagrangian projection.

Theorem 5.1 *Given $\Lambda \subset J^1(S)$ as above there is a small perturbation of Λ so that for a generic metric g on S there exist $\sigma_0 > 0$, almost complex structures J_σ , $0 < \sigma < \sigma_0$, and perturbations $\tilde{\Lambda}_\sigma$ of Λ_σ with the following properties.*

- *The Legendrian submanifold $\tilde{\Lambda}_\sigma$ is obtained from Λ_σ by a C^0 -deformation supported near the cusp edges of Λ_σ .*
- *The constrained rigid flow trees defined by $\tilde{\Lambda}_\sigma$ have well-defined limits as $\sigma \rightarrow 0$.*

- The (constrained) rigid J_σ -holomorphic disks with boundary on $\tilde{\Lambda}_\sigma$ with one positive puncture are in one-to-one correspondence with the (constrained) rigid flow trees of $\tilde{\Lambda}_\sigma$ with one positive puncture. In particular, the following holds for all sufficiently small $\sigma > 0$: if $u_\sigma: D_m \rightarrow T^*S$ is a (constrained) rigid J_σ -holomorphic disk then there exists a (constrained) rigid flow tree Γ of $\tilde{\Lambda}_\sigma$ such that $u_\sigma(\partial D_m)$ lies in an $\mathcal{O}(\sigma \log(\sigma^{-1}))$ neighborhood of the Lagrangian lift $\bar{\Gamma}$ of Γ . Moreover, outside $\mathcal{O}(\sigma \log(\sigma^{-1}))$ -neighborhoods of the Y_0, Y_1 -vertices and switches of Γ , the curve $u_\sigma(\partial D_m)$ lies at C^1 -distance $\mathcal{O}(\sigma \log(\sigma^{-1}))$ from $\bar{\Gamma}$.

Proof The proof is an adaption of results from [10], where Theorems 1.2 and 1.3 give a version of Theorem 5.1 for unconstrained rigid disks and trees. (See the proof of Theorem 1.3 and Lemma 5.13 in [10] for the $\mathcal{O}(\sigma \log(\sigma^{-1}))$ -estimate). We briefly recall the construction in order to adapt it to the constrained rigid case.

The first step is to fix a Riemannian metric on S such that there are only a finite number of (constrained) flow trees of formal dimension 0 determined by Λ and such that all such flow trees are transversely cut out. A straightforward modification of the unconstrained case, [10, Proposition 3.14], shows that the set of such metrics is open and dense.

The second step is to change the metric to g , to introduce almost complex structures J_σ , and to isotope the Legendrian Λ_σ to a new Legendrian $\tilde{\Lambda}_\sigma$. The main features of these objects are the following. The rigid flow trees determined by g and $\tilde{\Lambda}_\sigma$ are in one-to-one correspondence with the rigid flow trees of Λ and corresponding trees lie very close to each other; see [10, Lemma 4.4]. The submanifold $\tilde{\Lambda}_\sigma$ is rounded near its cusps and changed accordingly near its swallowtails. The metric g is flat in a neighborhood of any rigid flow tree and $\Pi(\tilde{\Lambda}_\sigma)$ is affine in this neighborhood outside a finite number of regions of diameter $\mathcal{O}(\sigma)$ where it is curved in only one direction; see [10, Subsection 4.2] and Remark 5.2 below for details. The almost complex structure J_σ agrees with the almost complex structure J_g induced by the metric in a neighborhood of all rigid flow trees and swallowtails and outside a neighborhood of the caustic (the locus of fiber tangencies) of $\Pi(\tilde{\Lambda}_\sigma)$. Near the points in the caustic outside a neighborhood of the swallowtails, J_σ is constructed so that both J_σ itself and the Lagrangian boundary condition given by $\Pi(\Lambda)$ splits as products with one direction along the cusp edge and one perpendicular to it; see [10, Section 4.2.3].

In fact, $\Pi(\tilde{\Lambda}_\sigma) = \Phi_\sigma^{-1}(\hat{L}_\sigma)$ where \hat{L}_σ is a totally real immersed submanifold and where Φ_σ is a diffeomorphism with $d_{C^0}(\Phi_\sigma, \text{id}) = \mathcal{O}(\sigma)$ and with $d_{C^1}(\Phi_\sigma, \text{id})$ arbitrarily small (but finite), which is supported in a small neighborhood of the caustic of $\Pi(\tilde{\Lambda}_\sigma)$ and which is equal to the identity in a neighborhood of all rigid flow trees of $\tilde{\Lambda}_\sigma$. We

use the almost complex structure $J_\sigma = d\Phi_\sigma \circ J \circ d\Phi_\sigma^{-1}$ on T^*S , where J is the almost complex structure induced by a metric on S . Then J_σ -holomorphic disks with boundary on $\tilde{\Lambda}_\sigma$ correspond to J -holomorphic disks with boundary on \hat{L}_σ .

We modify the deformations of the metric and of the Legendrian discussed above, in order to deal also with constrained rigid trees. In the presence of point conditions, repeat the construction in [10, Section 4.2], deforming the metric to g and constructing \hat{L}_σ in the exact same way as along rigid flow trees. Furthermore, this should be done in such a way that no constraining point is an edge point; see [10, Section 4.2.H]. Then take $J_\sigma = d\Phi_\sigma \circ J \circ d\Phi_\sigma^{-1}$, and let $\tilde{\Lambda}_\sigma$ be such that $\Phi_\sigma(\Pi\tilde{\Lambda}_\sigma) = \hat{L}_\sigma$, where J is the almost complex structure on T^*S induced by the special metric. We use the notation from [10, Remark 3.8] for vertices of (constrained) rigid flow trees.

After the modifications of the constructions in [10, Section 4] (first and second steps) described above, the theorem follows from [10, Theorem 1.2, Theorem 1.3] with the following additions. Theorem 1.2 shows that any rigid holomorphic disk with one positive puncture converges to a flow tree. Since the condition that the boundary of a disk passes through a constraining point is closed, it follows that constrained disks converge to constrained trees. In Theorem 1.3, rigid J_σ -holomorphic disks with boundary on \hat{L}_σ near any rigid flow tree are constructed and proved to be unique. The corresponding construction and uniqueness proof in the case of constrained rigid flow trees is completely analogous after the following alteration. If Γ is a constrained rigid flow tree with constraining point m then take the preimage of m (which is not an edge point) to be a marked point in the domain $\Delta_{p,m}^{(0,0)}(\Gamma, \sigma)$ of the approximately holomorphic disks; see [10, Section 6.2.A], and let $V_{\text{sol}}(m) = 0$ instead of $V_{\text{sol}}(m) \approx \mathbb{R}$, and see [10, Section 6.3.B, Definition 6.15]. \square

Remark 5.2 Below we will use the following special features of the metric g and the Legendrian submanifold $\tilde{\Lambda}_\sigma$ in Theorem 5.1. (For the construction of the metric and $\tilde{\Lambda}_\sigma$ with these properties we refer to [10, Section 4].)

- (1) There exists a neighborhood $X \subset S$ that contains all (constrained) rigid flow trees in which the metric g on S is flat. We write $X_{\tilde{\Lambda}_\sigma} \subset \Pi_F^{-1}(X) \cap \tilde{\Lambda}_\sigma$.
- (2) The image of any flow segment in a (constrained) rigid flow tree is a geodesic of the metric g .
- (3) The two sheets of the Lagrangian projection $\Pi(\tilde{\Lambda}_\sigma)$ near each double point consist of two transverse affine Lagrangian subspaces. The Lagrangian projection $\Pi(\tilde{\Lambda}_\sigma)$ is parallel to the 0-section (ie the graph of a polynomial function in flat local coordinates of degree at most 1) in neighborhoods of the following points: trivalent vertices of any (constrained) rigid flow tree, and intersection points of any (constrained) rigid flow tree.

- (4) In Fermi coordinates along any edge the Lagrangian is parallel to the 0–section in the coordinate perpendicular to the edge (ie the Lagrangian is the differential of the graph of a function of the form $f(x_1) + cx_2$, where x_1 is the coordinate along the edge, x_2 perpendicular to it, and c a constant).
- (5) Outside $\mathcal{O}(\sigma)$ –neighborhoods of a finite number of *edge points* in each (constrained) rigid flow tree the Lagrangian projection $\Pi(\tilde{\Lambda}_\sigma)$ is affine (ie the function f in (4) has the form $f(x_1) = a_2x_1^2 + a_1x_1 + a_0$). For our study of multiscale flow trees below we will also assume that the edge point regions are disjoint from the junction points (ie, the points where, using the notation of Definition 4.1, small trees are attached to the big tree).
- (6) We will also assume that the following extra condition is met: at a finite number of fixed *extra points* the Lagrangian projection has the form mentioned in (3). (These extra points are the junction points mentioned above and points where two distinct rigid multiscale flow trees intersect or one such tree self-intersects transversely.)

5.2 Refined results on constrained flow trees and disks

Recall that Theorem 4.3 relates holomorphic disks to multiscale flow trees. This relation is the result of a double degeneration: first the conormal lift of the unknot is pushed to the 0–section in $J^1(S^2)$ and then the conormal lift of a more general closed braid is pushed toward the almost degenerate conormal lift of the unknot. To deal with this we will stop the first degeneration close to the limit where actual holomorphic disks on the almost degenerate conormal lift of the unknot are close to flow trees. Then we degenerate the conormal lift of a general braid toward the almost degenerate conormal lift of the unknot and show that holomorphic disks near the limit admit a description in terms of quantum flow trees, ie, holomorphic disks with flow trees attached along their boundaries. Since both quantum flow trees and multiscale trees are defined as intersection loci of evaluation maps of holomorphic disks and flow trees, respectively, we need to show that the disks and the flow trees are arbitrarily C^1 –close in order to get the desired relation between quantum flow trees and multiscale flow trees. However the relation between quantum trees and holomorphic disks holds only for almost complex structures with special properties near $\Pi(\Lambda_U)$. In this section we show that there exist almost complex structures with these special properties for which the holomorphic disks are still close to flow trees.

5.2.1 Definition of quantum flow trees As already mentioned quantum flow trees will be central to establishing the relation between multiscale flow trees and holomorphic disks. We define them as follows. Consider $\Lambda_K \subset J^1(\Lambda_U)$ as above. A *quantum flow*

tree Ξ of Λ_K is a holomorphic disk $u: D_m \rightarrow T^*S^2$ with boundary on $\Pi(\Lambda_U)$ and a collection Γ of partial flow trees $\Gamma = \{\Gamma_1, \dots, \Gamma_m\}$ with one special positive puncture on the lift $\tilde{u}: \partial D_m \rightarrow \Lambda_U$ of the boundary of u . Note that Γ_j could be a constant flow tree at a Reeb chord of Λ_K . Then the positive special punctures subdivide the boundary of D_m into arcs and we require that there is a lift of these arcs to Λ_K which together with the 1-jet lifts of the trees in Γ from a closed curve when projected to T^*S^2 . As for multiscale flow trees we call the points where flow trees are attached to u *junction points*.

5.2.2 Modifying the almost complex structure — metric around flow trees Consider the degeneration $\sigma \rightarrow 0$ in Theorem 5.1. Fix a small $\sigma = \sigma_0 > 0$ so that the boundaries of all (constrained) rigid holomorphic disks are close to the cotangent lifts of their corresponding (constrained) rigid flow trees. More precisely, we take σ_0 so that the boundaries of all (constrained) rigid J_{σ_0} -holomorphic disks lie well inside the finite neighborhood X of the tree where the metric is flat and where $\tilde{\Lambda}_{\sigma_0}$ is as described in Remark 5.2. For simpler notation we write

$$\Lambda = \tilde{\Lambda}_{\sigma_0} = \Lambda_U,$$

where U is the unknot. We continue to use the subscript σ_0 in J_{σ_0} from Theorem 5.1 since we will modify the almost complex structure some more.

Consider an arbitrary closed braid $K \subset \mathbb{R}^3$ lying in a tubular neighborhood of U . If K is sufficiently close to U then Λ_K lies in a tubular neighborhood of $\Lambda = \Lambda_U$ that is symplectomorphic to $J^1(\Lambda)$. Furthermore, the front projection $\Pi_F^\Lambda: \Lambda_K \rightarrow \Lambda$ is an immersion. Since these properties are preserved under the global scaling by σ we consider $\Lambda_K \subset N \subset J^1(\Lambda)$, where the neighborhood N of the 0-section in $J^1(\Lambda)$ is identified with a neighborhood of Λ in $J^1(S^2)$.

When we compute the Legendrian homology of Λ_K , we will use an almost complex structure J_η (to be defined after Lemma 5.3 for small $\eta > 0$) on T^*S^2 , which differs from J_{σ_0} . In particular, to relate holomorphic disks with boundary on Λ_K with quantum flow trees of Λ_K , it will be important that J_η agrees with the almost complex structure induced by a metric on Λ in a neighborhood N_η of $\Pi(\Lambda) \subset T^*S^2$. Here N_η is the image under a symplectic immersion of a small neighborhood of the 0-section in $T^*\Lambda$ that extends $\Pi|_\Lambda$. Lemma 5.3, below, establishes the existence of such a metric on Λ . Specifically, the metric induces an almost complex structure on N_η , which (has a push-forward under an immersion which) agrees up to first order with J_{σ_0} in the fixed size neighborhood $\Pi(X_\Lambda) \subset \Pi(\Lambda)$ of the union of all boundaries of constrained rigid flow trees in Λ . Our desired almost complex structure J_η will interpolate between this push-forward and J_{σ_0} .

Recall the special form of $\Pi(\Lambda)$ near its double points; see Remark 5.2(3). Choose a metric g on Λ flat near its Reeb chord endpoints. Let J_g be the almost complex structure induced by g . Note that for such a metric we can find an immersion ϕ defined on the cotangent bundle of the neighborhood of the Reeb chord endpoint that is (J_g, J_{σ_0}) holomorphic. In particular, if $\phi: T^*\Lambda \rightarrow T^*S^2$ is an immersion that extends $\Pi|_\Lambda$ and that has these properties near Reeb chord endpoints, then the push-forward of J_g, ϕ_*J_g , is well-defined. Furthermore, we assume that (g, ϕ) satisfies this (J_g, J_{σ_0}) -holomorphicity condition also near the other points mentioned in Remark 5.2(3) and (6), where we take the extra points to be the junction points of the multiscale trees. We call the points in (3) and the extra points the *distinguished* points. We call a pair consisting of a metric and a symplectic immersion (g, ϕ) with properties as above *adapted to Λ* .

Lemma 5.3 *There exists a neighborhood N of the 0-section in $T^*\Lambda$ and a pair (g, ϕ) , consisting of a metric g and an immersion $\phi: N \rightarrow T^*S^2$ that extends $\Pi|_\Lambda$, that is adapted to Λ and such that the following hold on $\phi(N)$.*

- (1) J_{σ_0} and ϕ_*J_g agree along $\Pi(\Lambda)$.
- (2) J_{σ_0} and ϕ_*J_g agree in neighborhoods of distinguished points.
- (3) J_{σ_0} and ϕ_*J_g agree to first order in $\Pi(X_\Lambda)$.

Proof It is straightforward to check that statement (1) can be achieved and statement (2) follows by the definition of adapted pair (where the fact that the metric is flat and the Lagrangian affine near distinguished points readily implies existence).

We turn our attention to statement (3). Let p be a point on the 1-jet lift of a rigid flow tree. Pick normal coordinates $x = (x_1, x_2)$ on Λ around p with the 1-jet lift corresponding to $\{x_2 = 0\}$. Since the metric on S^2 is flat we can identify it locally with \mathbb{C}^2 with coordinates (u_1, v_1, u_2, v_2) and we can choose these coordinates so that the flow tree under consideration lies along $\{u_2 = 0\}$.

Since $\Pi(\Lambda)$ is a product of a curve in the u_1 -direction and a line segment parallel to the 0-section in the u_2 -direction we have the following local parametrization of $\Pi(\Lambda)$:

$$f(x_1, x_2) = (x_1, f(x_1), x_2, c_2).$$

Let $y = (y_1, y_2)$ denote the fiber coordinate. Defining the local immersion

$$\psi(x, y) = f(x_1, x_2) + y_1(-f'(x_1)\partial_{x_1} + \partial_{y_1}) + y_2\partial_{y_2}$$

we find that condition (1) holds and that the Taylor expansion of ψ^*J (ie, the complex structure of the flat metric on S^2 that corresponds to the standard complex structure in $u + iv$ -coordinates pulled back by ψ) with respect to y is

$$(\psi^*J)(x, y) = J_0 + B(x_1)y_1 + \mathcal{O}(2),$$

where J_0 is the standard complex structure on \mathbb{C}^2 with $x + iy$ -coordinates which is the complex structure induced by the flat metric on Λ . Here $B(x_1)J_0 + J_0B(x_1) = 0$, and $B(x_1)\partial_{x_2} = 0$. A straightforward calculation shows that if we change the immersion ψ by pre-composing with a diffeomorphism Φ with the Taylor expansion

$$\Phi(x, y) = (x, y) + \frac{1}{2}(B_1(x_1)\partial_{x_1})y_1^2,$$

then the almost complex structures agree to first order, ie, if $\phi = \psi \circ \Phi$ then $\phi^*J_{\sigma_0}$ and J_0 agree to first order along $\{y_1 = y_2 = 0\}$. □

Remark 5.4 For a general ambient almost complex structure J it is not possible to make the push forward agree up to first order. The Taylor expansion above with respect to (y_1, y_2) for a general J is

$$(\psi^*J_{\sigma_0})(x, y) = J_0 + B_1(x)y_1 + B_2(x)y_2 + \mathcal{O}(2),$$

where B_j anti-commutes with J_0 . One would then look for a map with Taylor expansion of the form

$$\Phi(x, y) = (x, y) + \frac{1}{2}(B_1\partial_{x_1})y_1^2 + \frac{1}{2}(B_2\partial_{x_2})y_2^2 + Cy_1y_2.$$

However, in order for the almost complex structures to agree up to first order one needs both $C = B_1\partial_{x_2}$ and $C = B_2\partial_{x_1}$. In general $B_1\partial_{x_2} \neq B_2\partial_{x_1}$ so no solution Φ exists.

Let (g, ϕ) be as in Lemma 5.3. For small $\eta > 0$, write N_η for the image under ϕ of an η -neighborhood of the 0-section in $T^*\Lambda$ and let Ξ_η denote the image of an η -neighborhood of T^*X_Λ . Theorem 5.1 implies that (for σ_0 small enough) every (constrained) rigid J_{σ_0} -holomorphic disk intersects a neighborhood of $\Pi(\Lambda)$ inside $\Xi_{\eta'}$ for some $\eta' > 0$. Write M_η for a η -neighborhood of all (constrained) rigid holomorphic disks. Let J_η denote an almost complex structure on T^*S^2 that equals ϕ_*J_g on N_η , equals J_{σ_0} outside $N_{2\eta}$, and that interpolates between the two in the remaining region in such a way that the following hold:

- (5-1) $|J_{\sigma_0} - J_\eta|_{C^0} \rightarrow 0$ as $\eta \rightarrow 0$
- (5-2) $|J_{\sigma_0} - J_\eta|_{C^1} \rightarrow 0$ in $\Xi_{\eta_0} \cup M_{\eta_0}$ as $\eta \rightarrow 0$, for fixed $\eta_0 > 0$
- (5-3) $|J_{\sigma_0} - J_\eta|_{C^2} \leq K_1$ in $\Xi_{\eta_0} \cup M_{\eta_0}$, for fixed η_0, K_1 , and $\eta < \eta_0$
- (5-4) $|J_{\sigma_0} - J_\eta|_{C^1} \leq K_2$, for fixed η_0, K_2 , and for $\eta < \eta_0$

We make two remarks. As a consequence of the fact that in general J_{σ_0} and $\phi_* J_g$ do not agree up to first order outside X we typically have $|J_{\sigma_0} - J_\eta|_{C^2} \rightarrow \infty$ as $\eta \rightarrow 0$. As shown in the C^0 -convergence portion of the proof of Lemma 5.5 below, we may assume that for sufficiently small η , any rigid J_η -holomorphic disk lies in M_{η_0} .

We next show that rigid J_η -holomorphic disks are C^1 -close to rigid J_{σ_0} -holomorphic disks. Note that it follows from Theorem 5.1 that (for σ_0 sufficiently small) the almost complex structure J_{σ_0} is regular in the sense that all (constrained) J_{σ_0} -holomorphic disks in T^*S^2 with boundary on $\Pi(\Lambda)$ of formal dimension at most 0 are transversely cut out.

Lemma 5.5 *Let $u_\eta: D_m \rightarrow T^*S^2$, $\eta \rightarrow 0$ be a sequence of (constrained) rigid J_η -holomorphic disks with boundary on Λ . Then some subsequence of u_η C^1 -converges on compact subsets of D_m to a (constrained) rigid J_{σ_0} -holomorphic disk. Moreover, for all $\eta > 0$ small enough there is a unique (constrained) rigid J_η -holomorphic disk in a neighborhood of each (constrained) rigid J_{σ_0} -holomorphic disk.*

Proof For the first statement we use Gromov compactness (for $|J_\eta - J_{\sigma_0}|_{C^0} \rightarrow 0$) to conclude that either $|Du_\eta|$ is uniformly bounded or there is bubbling in the limit, see eg Sikorav [32], and the fact that point constraints are closed. The case that u_η is a sequence of rigid disks bubbling is not possible: all bubbles of the limit have dimension at least 0, since J_{σ_0} is regular. This implies $\dim(u_\eta) > 0$ in contradiction to u_η being (constrained) rigid. We conclude thus that $|Du_\eta|$ is uniformly bounded; thus u_η converges uniformly to a (constrained) rigid J_{σ_0} -holomorphic disk u . We must show that it converges with one derivative as well.

Consider a point $z \in D$ and its image $u(z)$ under u . Since $u_\eta(z) \rightarrow u(z)$ and since $|Du_\eta|$ and $|Du|$ are bounded we can find a coordinate neighborhood $W \subset M_{\eta_0}$ of $u(z)$ and a small disk E around z in D so that $u_\eta(E) \subset W$ for all $\eta > 0$ and $u(E) \subset W$.

We pick \mathbb{C}^2 -coordinates on W so that $\Pi(\Lambda)$ corresponds to the totally real $\mathbb{R}^2 \subset \mathbb{C}^2$. The neighborhood E is either a disk or a half disk, with complex structure j . We find as in [32, Section 2.3] that in local coordinates u and u_η satisfy the equations

$$(5-5) \quad \bar{\partial}u + q \partial u = 0,$$

$$(5-6) \quad \bar{\partial}u_\eta + q_\eta \partial u_\eta = 0,$$

where

$$(5-7) \quad q(z) = (i + J_{\sigma_0}(u))^{-1}(i - J_{\sigma_0}(u)),$$

$$(5-8) \quad q_\eta(z) = (i + J_\eta(u_\eta))^{-1}(i - J_\eta(u_\eta)).$$

Letting $u - u_\eta = h_\eta$ we conclude that

$$(5-9) \quad \bar{\partial}h_\eta + q_\eta \partial h_\eta = (q - q_\eta)\partial u.$$

By scaling we may take $|q_\eta|_{C^2} \leq \epsilon \ll 1$; see [32]. Moreover, by C^0 -convergence u_η lies in M_{η_0} and therefore $|q - q_\eta|_{C^1} \rightarrow 0$. A standard bootstrap argument for h_η now shows that $|h_\eta|_{C^1} \rightarrow 0$ as $\eta \rightarrow 0$.

Uniqueness of u_η near u_{σ_0} for small η follows from transversality at u_σ and the C^1 -convergence (5-3) of complex structures near u_{σ_0} . \square

Remark 5.6 A general sequence of J_η -holomorphic disks that does not blow up would C^0 -converge to a J_{σ_0} -holomorphic disk but not necessarily C^1 -converge.

5.2.3 Metric and perturbations—metric near flow trees of Λ_K Note that the metric g in Lemma 5.3 has arbitrary form outside a small neighborhood of the 1-jet lifts of constrained rigid flow trees of Λ and that furthermore it is flat near all distinguished points. In this section we will impose further conditions on g outside this region in order to adapt it to the (partial) flow trees of $\Lambda_K \subset J^1(\Lambda)$ that are parts of rigid multiscale flow trees of Λ_K . Furthermore, we will also deform Λ_K itself in a way analogous to how Λ_σ was deformed into $\tilde{\Lambda}_\sigma$. The construction is completely analogous to the construction in [10, Section 4], although it is simpler in the present situation since Λ_K has no front singularities. The construction in Section 5.2.2 gives a metric in Λ , which is flat in a neighborhood of the 1-jet lifts of all (constrained) rigid flow trees whose Π -projections contain the boundaries of all (constrained) rigid holomorphic disks. Furthermore, we take the distinguished points to include all junction points as well as points where multiscale flow trees intersect constrained rigid disks. We next extend the region where the metric is flat to contain all rigid flow trees of $\Lambda_K \subset J^1(\Lambda)$ as well as all partial flow trees of Λ_K that are parts of rigid generalized disks. Note that these regions include the projection of any Reeb chord of $\Lambda_K \subset J^1(\Lambda)$. We next deform Λ_K slightly so that it has the form described in Remark 5.2 over all the flow trees just mentioned; see [10, Section 4.2]. (Here we treat junction points corresponding to positive punctures of special trees like the 3-valent vertices of flow trees in [10] and treat the other junction points like the 2-valent punctures in [10].) In particular, Λ_K is affine at Reeb chord endpoints, lifts of flow trees are geodesics in the flat metric, and the sheets of Λ_K near a junction point that is a special Reeb chord will be parallel to the 0-section in $J^1(\Lambda)$ (which in turn is parallel to the 0-section in $J^1(S^2)$ over a subset $U \subset S^2$ where the metric on S^2 is flat). Similarly, the metric on S^2 is flat near junction points that are Reeb chords, where the sheet of Λ is parallel to the 0 section, and where the sheets of Λ_K are affine (and almost parallel to the 0-section).

Let $0 \leq \eta \leq 1$ and define $\Lambda_{K;\eta} = s_\eta(\Lambda_K)$ to be the image under fiber scaling by η in $J^1(\Lambda)$. Then as above, along 1-jet lifts of flow lines that are part of rigid generalized disks, the Lagrangian $\Pi(\Lambda_{K;\eta})$ is a product of a horizontal line segment and a curve over the distinguished curve. Thus, as in [10] and [13], the regions where this curve is not affine have diameters $\mathcal{O}(\eta)$ as $\eta \rightarrow 0$. With this metric we construct the almost complex structure J_η of Lemma 5.5 for some sufficiently small but fixed $\eta > 0$. Note that J_η then agrees with the complex structure induced by the metric on Λ in a neighborhood of $\Pi(\Lambda)$ that is the image under an immersion of a small neighborhood of the 0-section in $T^*\Lambda$.

5.3 From disks to quantum flow trees

Now that we have finished modifying our almost complex structures to achieve J_η with the desired properties from the previous subsection, we simplify notation and let

$$J = J_\eta.$$

The main result of this section is that any sequence of rigid J -holomorphic disks with boundary on $\Lambda_{K,\eta} = s_\eta(\Lambda_K)$ has a subsequence that converges to a rigid quantum flow tree of Λ_K and Λ . In Section 5.3.2 we characterize certain subsets of the domains of any sequence of J -holomorphic disks with boundary on $\Lambda_{K,\eta}$ such that the restrictions of the maps to these subsets converge to a (partial) flow tree of $\Lambda_K \subset J^1(\Lambda)$. In particular, in case these subsets constitute the whole domains of the members in the sequence we find that the J -holomorphic disks converge to a flow tree. In Section 5.3.3 we show that there can be at most one disk bubbling off in the limit of a sequence of rigid J -holomorphic disks with boundary on $\Lambda_{K,\eta}$ as $\eta \rightarrow 0$ and that by adding a puncture in the domain near the point where the bubble forms we ensure that the maps in the sequence satisfy a uniform derivative bound. In Section 5.3.4 we prove the main result of the section. After the previous subsections there are two main points that must be demonstrated. First, we show that the limit has only one holomorphic disk part, which must be a (constrained) rigid disk. Second, we show that our analysis of the two separate parts (the flow tree- and the disk part) gives a complete description of the limit objects.

5.3.1 Notation Consider $\Lambda_K \subset N \subset J^1(\Lambda)$ where N is a neighborhood of the 0-section that is identified with a neighborhood of $\Lambda \subset J^1(S^2)$. In particular, if $0 < \eta < 1$ and $s_\eta: J^1(\Lambda) \rightarrow J^1(\Lambda)$ denotes the fiber scaling then $s_\eta(N) \subset N$ and hence $\Lambda_{K,\eta}$ is a Legendrian submanifold in N that is Legendrian isotopic to Λ_K .

Consider the Reeb chords of $\Lambda_{K;\eta}$. Recall that these are of two types: short and long chords. We will use the notation for chords introduced in Remark 3.8. The *action* $a(c)$

of a Reeb chord c is the positive difference of the z -coordinates of its two endpoints. Stokes' Theorem implies the area of a holomorphic disk is the signed sum of the actions of the Reeb chords at its punctures; see [15, Lemma 2.1], for example. Recall the two Reeb chords e and c of $\Lambda = \Lambda_U$ introduced in Section 3.1. Chords of $\Lambda_{K;\eta}$ satisfy the following: chords of type L_2 lie close to e and have action $\alpha(e) + \mathcal{O}(\eta)$, chords of type L_1 lie close to c and have action $\alpha(c) + \mathcal{O}(\eta)$, and chords of types S_0 and S_1 have action $\mathcal{O}(\eta)$.

For the fixed almost complex structure $J = J_\eta$, Lemma 5.5 holds and J has properties as in Section 5.2.3. Then, in particular, boundaries of (constrained) rigid disks of Λ lie C^1 -close to 1-jet lifts of its corresponding (constrained) rigid trees. It follows, in particular, that there is a natural one-to-one correspondence between rigid quantum trees of Λ_K and rigid generalized trees of Λ_K . Furthermore, in a neighborhood of $\Pi(\Lambda)$ the almost complex structure J agrees with the one induced by the metric on Λ by Lemma 5.3.

Below we will discuss J -holomorphic disks in T^*S^2 with boundary punctures. Throughout we will think of these as maps $u: \Delta_m \rightarrow T^*S^2$, where the source is a standard domain. For details on standard domains we refer to [10, Section 2.2.1]; here we give a brief description. Consider \mathbb{R}^{m-2} with coordinates $\underline{\tau} = (\tau_1, \dots, \tau_{m-2})$. Let $t \in \mathbb{R}$ act on \mathbb{R}^{m-2} by $t \cdot \underline{\tau} = (\tau_1 + t, \dots, \tau_{m-2} + t)$. The orbit space of this action is the space of conformal structures of the disk with m boundary punctures, one of which is distinguished, $\mathcal{C}_m \approx \mathbb{R}^{m-3}$. Define a *standard domain* $\Delta_m(\underline{\tau})$ as the subset of $\mathbb{R} \times [0, m]$ obtained by removing $m - 2$ horizontal slits of width ϵ , $0 < \epsilon \ll 1$, starting at (τ_j, j) , $j = 1, \dots, m - 2$ and going to $+\infty$. All slits have the same shape, ending in a half-circle; see Figure 24. The points (τ_j, j) are called the *boundary minima*.




Figure 24: The standard domain $\Delta_5(\tau_1, \tau_2, \tau_3)$ with three boundary minima

In the case that u has less than two negative punctures we add marked points at intersections with a small circle around the positive Reeb chord endpoint and puncture the domain there, so that it admits a description as a standard domain. We will often write Δ_m dropping the precise information about the conformal structure from the notation.

Consider a standard domain $\Delta_m \subset \mathbb{C}$. Let $V_\beta = \{x + iy : x = \beta\}$. A connected component of the closure of $V_\beta \cap (\Delta_m - \partial\Delta_m)$ in Δ_m will be called a *vertical segment* in Δ_m .

5.3.2 Flow tree convergence Consider a sequence $u_\eta: \Delta_m \rightarrow T^*S^2$ of rigid J -holomorphic disks with boundary on $\Lambda_{K;\eta}$, $\eta \rightarrow 0$. As $\eta \rightarrow 0$, the actions of Reeb chords of $\Lambda_{K;\eta}$ of type \mathbf{S} satisfy a $\mathcal{O}(\eta)$ bound. Hence, by Stokes' Theorem and the dimension formula for holomorphic disks, if $\eta > 0$ is sufficiently small then a moduli space of J -holomorphic disks with one positive puncture, with boundary on $\Lambda_{K;\eta}$, and of formal dimension 0 can be non-empty only if it contains disks with punctures of the following types:

- (QT $_\emptyset$) The positive puncture is of type \mathbf{S}_1 and all negative punctures are type \mathbf{S}_0 .
- (QT $_0$) The positive puncture is of type \mathbf{L}_1 and all negative punctures are of type \mathbf{S}_0 .
- (QT' $_0$) The positive puncture is of type \mathbf{L}_2 , one negative puncture is of type \mathbf{L}_1 , and all other negative punctures are of type \mathbf{S}_0 .
- (QT $_1$) The positive puncture is of type \mathbf{L}_2 , one negative puncture is of type \mathbf{S}_1 , and all other negative punctures are of type \mathbf{S}_0 .

Lemma 5.7 *If $u_\eta: \Delta_m \rightarrow T^*S^2$ is a sequence of J -holomorphic disks of type (QT $_\emptyset$) then u_η has a subsequence that converges to a flow tree of $\Lambda_K \subset N \subset J^1(\Lambda)$ as $\eta \rightarrow 0$.*

Proof The actions of Reeb chords in \mathbf{S}_1 are $\mathcal{O}(\eta)$. Thus the area of u_η is $\mathcal{O}(\eta)$ as well. Monotonicity then implies that $u_\eta(\Delta_m)$ must stay inside an $\mathcal{O}(\eta^{1/2})$ -neighborhood of $\Pi(\Lambda)$. Since J agrees with the complex structure coming from the metric on Λ in a finite neighborhood of $\Pi(\Lambda)$ and since disks lift to the symplectization of $J^1(S^2)$ where $\Lambda \times \mathbb{R}$ is embedded, the lemma follows from [10, Theorem 1.2]. □

We next show that for any sequence of J -holomorphic disks $u_\eta: \Delta_m \rightarrow T^*S^2$ there are neighborhoods of each negative puncture of type \mathbf{S} , where the disk converges to a flow tree (which may be constant). The key to establishing this convergence is an $\mathcal{O}(\eta)$ derivative bound on neighborhoods of the punctures. Consider the inclusion

$\Lambda_{K;\eta} \subset N \subset J^1(\Lambda)$ and let z denote a coordinate in the \mathbb{R} -direction of $J^1(\Lambda) \approx T^*\Lambda \times \mathbb{R}$. If γ is a curve in T^*S^2 then let $\ell(\gamma)$ denote the length of γ in the metric induced by ω and J . Fix $M > 0$ larger than the maximum of the function $|z|$ on $\Lambda_K = \Lambda_{K;1}$. A vertical segment $l_\eta \approx [0, 1]$ in the domain Δ_m of u_η such that

$$(5-10) \quad \ell(u_\eta|_{l_\eta}) \leq M\eta,$$

$$(5-11) \quad |z(u_\eta|_{l_\eta}(1)) - z(u_\eta|_{l_\eta}(0))| \leq M\eta,$$

will be called an η -short vertical segment.

By elementary Fourier analysis, near punctures, holomorphic disks u_η have expansions of the form

$$(5-12) \quad u_\eta(\tau + it) = e^{i\theta_0(\eta)} \sum_{n>0} c_n e^{\theta(\eta)(\tau+it)}, \quad \tau + it \in [0, \infty) \times [0, 1],$$

(see [14, Section 6], and Robbin and Salamon [31]), where $\theta(\eta) = \mathcal{O}(\eta)$ is the angle between two sheets of the double point. This implies that there exist η -short vertical segments in a neighborhood of each puncture of u_η of type **S**. Note that a vertical segment l_η subdivides Δ_m into two components: $\Delta_m = \Delta_m^+(l_\eta) \cup \Delta_m^-(l_\eta)$, where $\Delta^+(l_\eta)$ contains the positive puncture of u_η . For $d > 0$, let $\Delta^\pm(l_\eta, d)$ denote the subset of points in $\Delta^\pm(l_\eta)$ that are at a distance at least d from l_η .

Lemma 5.8 *For all sufficiently small $\eta > 0$ the following derivative bound holds: if l_η is an η -short vertical segment in the domain of $u_\eta: \Delta_m \rightarrow T^*S^2$ then*

$$|du_\eta(z)| = \mathcal{O}(\eta), \quad z \in \Delta^-(l_\eta, 1).$$

Proof Let b_1, \dots, b_r denote the negative punctures of u_η that lie in $\Delta_m^-(l_\eta)$. Then the area A_η^- of $u_\eta(\Delta^-(l_\eta))$ satisfies

$$\begin{aligned} 0 \leq A_\eta^- &= \int_{u_\eta(\partial\Delta^-(l_\eta))} p \, dq \\ &= \int_{u_\eta(l_\eta)} p \, dq + (z(u_\eta|_{l_\eta}(1)) - z(u_\eta|_{l_\eta}(0))) - \sum_{j=1}^r \ell(b_j) = \mathcal{O}(\eta). \end{aligned}$$

We conclude by monotonicity that $u_\eta(\Delta^-(l_\eta))$ must lie in an $\mathcal{O}(\eta^{1/2})$ -neighborhood of Λ (in particular all chords b_j are of type **S**). Since J agrees with the almost complex structure induced by the metric on Λ in such a neighborhood, [10, Lemma 5.4] shows that the function $|p|^2$, where p is the fiber coordinate in $T^*\Lambda$, composed with u_η is subharmonic on $\Delta^-(l_\eta)$ and therefore attains its maximum on the boundary. The lemma then follows from [10, Lemma 5.6]. □

The derivative bound of Lemma 5.8 leads to flow tree convergence on $\Delta^-(l_\eta)$. Consider a sequence of η -short vertical segments l_η such that each domain $\Delta^-(l_\eta)$ contains a puncture mapping to a Reeb chord $\eta \cdot b$ of $\Lambda_{K;\eta}$ for some Reeb chord b of Λ_K with $\Pi(b) \in N$.

Corollary 5.9 *There exists a constant $C > 0$ such that the sequence of restrictions $u_\eta|_{\Delta^-(l_\eta, C \log(\eta^{-1}))}$ has a subsequence that converges to a flow tree of $\Lambda_K \subset J^1(\Lambda)$. (Note that the flow tree in the limit may be constant.)*

Proof By Lemma 5.8, the image under u_η of any region in $\Delta_m^-(l_\eta, 1)$ of diameter $\log(\eta^{-1})$ lies inside a disk of radius $\mathcal{O}(\eta \log(\eta^{-1}))$. Furthermore, along any strip region in $\Delta^-(l_\eta, \log(\eta^{-1}))$ outside an $\mathcal{O}(\log(\eta^{-1}))$ -neighborhood of the boundary minima in $\Delta^-(l_\eta)$, the map converges to a flow tree by the proof of [10, Theorem 1.2]; see in particular Lemmas 5.12, 5.16 and 5.17 in [10]. \square

Remark 5.10 If the limiting flow tree in Corollary 5.9 is constant then it lies at $\Pi(b) \in N$, where b is the Reeb chord of Λ_K above. To see this note that $\Delta_m^-(l_\eta, C \log(\eta^{-1}))$ always contains a half infinite strip that is a neighborhood of the puncture mapping to b . If the vertical segment bounding this strip does not converge to $\Pi(c)$ for some Reeb chord c , then the limiting tree would be non-constant. Thus, each such segment converges to the projection of the Reeb chord and since $\Pi(b)$ is in the image we find that the whole tree must lie at $\Pi(b)$.

5.3.3 Blow-up analysis We next show that the limit of any sequence of J -holomorphic disks $u_\eta: \Delta_m \rightarrow T^*S^2$ with boundary on $\Lambda_{K;\eta}$ can contain at most one bubble. We also show how to add one puncture consistently to each domain so that this forming bubble corresponds to some coordinate of the domains of u_η , which give points in the space of conformal structures on the disk with m boundary punctures, one distinguished, $\mathcal{C}_m \approx \mathbb{R}^{m-3}$, approaching ∞ (rather than the derivative of u_η blowing up).

Lemma 5.11 *If $u_\eta: \Delta_m \rightarrow T^*S^2$ is a sequence of J -holomorphic disks with boundary on $\Lambda_{K;\eta}$ and with one positive puncture such that $\sup_{\Delta_m} |du_\eta|$ is unbounded as $\eta \rightarrow 0$ then, after adding one puncture in the domains Δ_m of u_η , we get an induced sequence $u_\eta: \Delta_{m+1} \rightarrow T^*S^2$ for which $|du_\eta|$ is uniformly bounded from above.*

Proof The proof uses standard blow-up arguments; see eg [32]. Assume that $M_\eta = \sup_{\Delta_m} |du_\eta|$ is not bounded as $\eta \rightarrow 0$. Using asymptotic properties of J -holomorphic disks near the punctures of u_η , see (5-12), we find that for $\eta > 0$ there exist points

$p_\eta \in \Delta_m$ such that $|du_\eta(p_\eta)| = M_\eta$. View Δ_m as a subset of \mathbb{C} and consider the sequence of maps

$$g_\eta(z) = u_\eta\left(p_\eta + \frac{z}{M_\eta}\right) \quad \text{defined on } U = \left\{z \in \mathbb{C} : p_\eta + \frac{z}{M_\eta} \in \Delta_m\right\},$$

where Δ_m refers to the domain of u_η . Note that the derivative $|dg_\eta|$ of g_η is uniformly bounded as $\eta \rightarrow 0$. We can thus extract a convergent subsequence, which in the limit $\eta = 0$ gives a non-constant holomorphic disk $v: H \rightarrow T^*S^2$, where H is the upper half plane, with boundary on Λ , with a positive puncture at infinity, and no other puncture. (There may be other bubble disks in the limiting bubble tree, but since there are no bubbles forming near punctures and since the total topology of the domain is recoverable from the limit, at least one of them has a single puncture. Note also that in the limit, only the chords of type **L** exist.) Fix an arc A in Λ that intersects $v(\partial H)$ transversely at a point far from all Reeb chord endpoints. It follows from the convergence $g_\eta \rightarrow v$ that there exists a point q_η in the domain Δ_m for u_η with $|p_\eta - q_\eta| \rightarrow 0$ as $\eta \rightarrow 0$ such that $u_\eta(q_\eta) \in \Pi(\Lambda)$. Adding a puncture at q_η in the domain Δ_m of u_η gives a new sequence of maps $u_\eta^1: \Delta_{m+1} \rightarrow T^*S^2$. Assume now that $\sup_{\Delta_{m+1}} |du_\eta^1|$ is unbounded. Repeating the blow up argument sketched above, we would again find a bubble disk $v^1: H \rightarrow T^*S^2$ in the limit with one positive puncture and no other punctures.

Since the area contributions of u_η^1 in a neighborhood of the added puncture q_η are uniformly bounded from below and since this neighborhood can be taken to map to a region far from all Reeb chords, it follows that there are at least two non-constant disks in the limit, both with one positive puncture and no other punctures. The sum of the areas of these two disks is bounded from below by $2a(c) + \mathcal{O}(\eta)$, where c is the shorter of the two Reeb chords of Λ . This however contradicts u_η having one positive puncture since the lengths of Reeb chords then imply $\text{Area}(u_\eta) \leq a(e) + \mathcal{O}(\eta)$ and $2a(c) > a(e)$. The lemma follows. \square

5.3.4 Quantum flow tree convergence Consider a sequence $u_\eta: \Delta_m \rightarrow T^*S^2$ of rigid J -holomorphic disks with boundary on $\Lambda_{K,\eta}$. Assume, without loss of generality (see Lemma 5.11), that $|du_\eta|$ is uniformly bounded.

Lemma 5.12 *If each of the disks u_η has a positive puncture at a Reeb chord of type **L** then $\sup_{\Delta_m} |du_\eta|$ is uniformly bounded from below.*

Proof Consider the neighborhood N of $\Pi(\Lambda)$ where J is induced by the metric on Λ . Since u_η maps $\partial\Delta_m$ to an $\mathcal{O}(\eta)$ -neighborhood of $\Pi(\Lambda)$ there exists $\epsilon > 0$ such that if $|du_\eta| \leq \epsilon$ then $u_\eta(\Delta_m) \subset N$ for all sufficiently small $\eta > 0$. [10, Lemma 5.4]

then shows that the function $|p \circ u_\eta|^2$ is subharmonic and, consequently, that the sequence u_η has a subsequence $u_{\eta'}$ that converges to a flow tree; see [10, Lemma 5.6]. In particular the area of $u_{\eta'}$ is $\mathcal{O}(\eta')$. But since $u_{\eta'}$ has a positive puncture of type **L**, there is a uniform bound from below on the area of $u_{\eta'}$ restricted to a neighborhood of this puncture. See the proof of [14, Lemma 9.3], for example. \square

Lemma 5.12 leads to a description of the J -holomorphic components in the limit of the sequence u_η with positive puncture at a chord of type **L** as follows. Consider a sequence of points p_η in the domains Δ_m of u_η such that $|du_\eta(p_\eta)| \geq \epsilon > 0$ for all $\eta > 0$. After passing to a subsequence, we may assume that $u_\eta(p_\eta)$ converges in T^*S^2 . Consider coordinates $\tau + it$ on \mathbb{C} and represent the domain $\Delta_m \subset \mathbb{C}$ of u_η by letting the τ -coordinate of p_η equal 0. Then as $\eta \rightarrow 0$ in the sequence of domains of u_η , some boundary minima of Δ_m stay at finite distance from p_η and others do not. After passing to a subsequence we may assume that every sequence of boundary minima on a fixed height has a limit, which may be finite or infinite, and we find a limiting conformal structure on a domain Δ_{m_0} that contains $p_0 = \lim_\eta p_\eta$. It follows in a straightforward way that u_η converges (uniformly on compacts) to a non-constant J -holomorphic disk $v: \Delta_{m_0} \rightarrow T^*S^2$ with boundary on Λ . We say that such a disk is a *non-constant J -holomorphic component* of the limit.

We next consider the role of the choice of p_η . If the sequence of domains Δ_m converges and if they contain some point q_η such that $|du_\eta(q_\eta)| \geq \delta > 0$ and if $|p_\eta - q_\eta| \rightarrow \infty$ as $\eta \rightarrow 0$ then, repeating the above argument, we extract another non-constant J -holomorphic component v' containing $q_0 = \lim_{\eta \rightarrow 0} q_\eta$ of the limit, which is distinct from the component v that contains p_0 . Furthermore, if the τ -coordinate of q_η approaches $\pm\infty$ in coordinates where the τ -coordinate of p_η equals 0 and if a is the Reeb chord at the positive puncture of v' then a is also the negative puncture of some non-constant J -holomorphic component in the limit. Arguing by action it is easy to see that the number of non-constant components in the limiting configuration is finite. And since Λ has only two Reeb chords of almost equal actions the number of such components is at most two.

We next show that the flow trees in Lemma 5.8 fit together with the non-constant components to form a quantum flow tree. Consider a puncture ζ_r at $+\infty$ in the domains Δ_m of u_η that maps to a Reeb chord r of $\Lambda_{K;\eta}$. By asymptotic properties of holomorphic disks at punctures there are η -short vertical segments that separate ζ_r from the positive puncture ζ_+ at $-\infty$ of Δ_m . In particular, there is an η -short vertical segment $l_\eta(\zeta_r)$ of minimal τ -coordinate that separates ζ_r from ζ_+ ; we call it the *extremal η -small* vertical segment of ζ_r . If $v: \Delta_{m'} \rightarrow T^*S^2$ is a non-constant J -holomorphic component in the limit configuration then we write $\partial v = v(\partial\Delta_{m'})$ for the image of v restricted to the boundary.

Lemma 5.13 *If ζ_r is any puncture at $+\infty$ in the domains Δ_m of u_η that maps to a Reeb chord r and if l_η is the extremal η -small vertical segment of ζ then there exists a non-constant component v of the limit such that $u_\eta(l_\eta)$ converges to a point in ∂v .*

Proof We prove this lemma by contradiction: if the statement of the lemma does not hold, then the area difference between a limit disk and the disks before the limit violates an $\mathcal{O}(\eta)$ area bound. In this proof we let Area refer to area defined by the Riemannian metrics $\omega(J_\eta, 0)$, $\eta > 0$.

Assume the lemma does not hold. Then there exists $\epsilon > 0$ such that for any sequence of l_η which satisfies the inequalities (5-10) and (5-11), some point on l_η maps a distance at least $\epsilon > 0$ from ∂v . Consider a strip region between two slits, $[-d, d] \times [0, 1] \subset \Delta_m$, for which some point converges to a point at distance δ from ∂v , where $\frac{\epsilon}{4} < \delta < \frac{\epsilon}{2}$. Let $\sup_{[-d, d] \times [0, 1]} |du_\eta| = K$. Then K is not bounded by $M\eta$ for any $M > 0$. Since the difference between the sum of areas of the non-constant components in the limit and that of u_η is $\mathcal{O}(\eta)$, it follows that $|du_\eta| = \mathcal{O}(\eta^{1/2})$ (by monotonicity and a standard bootstrap estimate). Thus $K = \mathcal{O}(\eta^{1/2})$. Consider next the scaling of the target by K^{-1} at the image of $(0, 0) \in [-d, d] \times [0, 1]$. We get a sequence of maps \hat{u}_η from $[-d, d] \times [0, 1]$ with bounded derivative. Note moreover that the scaled boundary condition is $\mathcal{O}(\eta^{1/2})$ from the 0-section. Changing coordinates to the standard $(\mathbb{C}^n, \mathbb{R}^n)$ respecting the complex structure at the limit point, we find that there are maps $f_\eta: [-d, d] \times [0, 1] \rightarrow \mathbb{C}^n$ with the following properties:

- $\sup_{[-d, d] \times [0, 1]} |D^k f_\eta| = \mathcal{O}(\eta^{1/2})$, $k = 0, 1$,
- $\hat{u}_\eta + f_\eta$ satisfies \mathbb{R}^n boundary conditions, and
- $\bar{\partial}(\hat{u}_\eta + f_\eta) = \mathcal{O}(\eta^{1/2})$.

It follows that $\hat{u}_\eta + f_\eta$ converges to a holomorphic map with boundary on \mathbb{R}^n , which takes 0 to 0 and which has derivative of magnitude 1 at 0. Using solvability of the $\bar{\partial}$ -equation in combination with L^2 -estimates in terms of area we find that the area of \hat{u}_η must be uniformly bounded from below by a constant C . The area contribution to the original disks near the limit is thus at least $K^2 C$. Since the image of $[-d, d] \times [0, 1]$ in $\Pi(\Lambda)$ has diameter at most $2Kd$, we may repeat the argument with many disjoint finite strips with maximal derivatives K_j and with the sum of diameters bounded below by $\frac{\epsilon}{100}$. We find that the area contribution is bounded from below by $C \sum K_j^2$. Since the length contribution is bounded below, we get:

$$2d \sum K_j \geq \frac{\epsilon}{100}.$$

Now,

$$C \sum K_j^2 \geq C \inf_j \{K_j\} \sum K_j \geq C' \inf_j \{K_j\}.$$

For any $M > 0$, $\inf_j \{K_j\} \geq M\eta$. To see this assume that it does not hold true. Then there is a sequence of vertical segments l_η such that $|du_\eta| \leq 2M\eta$ with the property that the distance between $u(l_\eta)$ and ∂v is at most $\frac{3}{4}\epsilon$. This however contradicts our hypothesis. Consequently, the area contribution from the remaining part of the disk is not $\mathcal{O}(\eta)$.

Consider now the limit disk v . (For simplicity we explain how to deal with the case of one extremal vertical segment, the case of many is completely analogous.) On its boundary there is a point p such that for all η sufficiently small there are points on the boundary of u_η lying on different sheets of $\Pi(\Lambda_{K;\eta})$ near $v(p)$, such that both limit to p , and which in the boundary of the domain of u_η separate the positive puncture from the strips considered above. Fix a small sphere of radius of order of magnitude $T\eta$ around p , where T is a large constant. Then for generic spheres there is an arc γ_η in the intersection of this sphere and the image of u_η that connects the two sheets of $\Pi(\Lambda_{K;\eta})$. Write E_η and I_η for the two components of the complement of $u_\eta^{-1}(\gamma_\eta)$ in the domain of u_η , where I_η contains the positive puncture.

By the above area estimate we have

$$\text{Area}(u_\eta(E_\eta)) = \int_{\gamma_\eta} p \, dq - \mathcal{O}(\eta)$$

by Stokes' Theorem, since u_η is J_η -holomorphic. Now consider the shortest geodesic arc γ_η^0 in the sphere connecting the two endpoints of γ_η . Then there is a 2-chain in the sphere with boundary $\gamma_\eta^0 \cup \gamma_\eta$ of area $\mathcal{O}(\eta^2)$, and clearly $|\int_{\gamma_\eta^0} p \, dq| = \mathcal{O}(\eta)$.

We next construct a 2-chain A with boundary on $\Pi(\Lambda_{K;\eta})$ as follows. The first piece is the union of $u_\eta(I_\eta)$ and the above 2-chain in the sphere. The resulting chain has area such that

$$\eta^{-1}(\text{Area}(u_\eta) - \text{Area}(A)) \geq \eta^{-1}(\text{Area}(E_\eta) + \mathcal{O}(\eta)) \rightarrow \infty \quad \text{as } \eta \rightarrow 0.$$

Now connect p to the start point of the flow tree corresponding to the extremal η -small vertical segments (that converge to a point at distance ϵ from p). Take the chain that has boundary the cotangent lift of this curve and that consists of vertical fiber segments. This chain has area $\mathcal{O}(\eta)$. Finally, consider the analogous chain associated to the flow tree starting at the end point of this arc. We interpolate between these chains and have constructed a 2-chain A' with boundary homotopic to the boundary of u_η and such that

$$\text{Area}(A') < \text{Area}(u_\eta).$$

This, however, contradicts the fact that holomorphic curves minimize area in their homology class. The lemma follows. □

As a consequence of the preceding lemmas, we get the following result.

Corollary 5.14 *Any sequence of rigid holomorphic disks u_η with boundary on $\Lambda_{K;\eta}$ has a subsequence that converges to a quantum flow tree. The quantum flow trees that arise as limits of rigid disks are of the following types:*

- (QT $_\emptyset$) *No non-constant J -holomorphic components (ie, flow trees), positive puncture of type \mathbf{S}_1 and negative punctures of type \mathbf{S}_0 .*
- (QT $_0$) *Rigid non-constant J -holomorphic component, positive puncture of type \mathbf{L}_1 and negative punctures of type \mathbf{S}_0 .*
- (QT' $_0$) *Rigid non-constant J -holomorphic component, positive puncture of type \mathbf{L}_2 , one negative puncture is of type \mathbf{L}_1 and all other negative punctures are of type \mathbf{S}_0 .*
- (QT $_1$) *Constrained rigid non-constant J -holomorphic component, positive puncture of type \mathbf{L}_2 , one negative puncture of type \mathbf{S}_1 and remaining negative punctures of type \mathbf{S}_0 .*

Proof Consider (QT $_\emptyset$): if the positive puncture of u_η is of type \mathbf{S}_1 , flow tree convergence was established in Lemma 5.7. For (QT $_0$) and (QT $_1$), assume that the positive puncture has type \mathbf{L} . Convergence to a quantum flow tree follows from the above: Lemma 5.12 implies that there is some non-constant component in the limit and the discussion following that lemma shows that the non-constant component is a broken disk with at most two levels, and Lemma 5.13 then implies that the flow tree pieces of Corollary 5.9 are attached to the non-constant components. If the positive puncture is of type \mathbf{L}_1 then by the dimension formula there can be only one non-constant component in the limit, and this component must be rigid. Since no rigid disk passes through any chord of type \mathbf{S} and since flow trees with negative punctures at chords of type \mathbf{S}_1 are constant we conclude that (QT $_0$) holds in this case. A similar argument shows that (QT' $_0$) holds also when the positive puncture is of type \mathbf{L}_2 and there is a negative puncture of type \mathbf{L}_1 . In the case that the positive puncture is of type \mathbf{L}_2 and all negative punctures are of type \mathbf{S} it follows from the dimension formula that exactly one negative puncture maps to a chord b of type \mathbf{S}_1 . Since all flow trees with a negative puncture at b are constant it follows that some non-constant component in the limit passes $\Pi(b)$. Since no rigid disk passes through $\Pi(b)$ it follows that the non-constant component is un-broken and hence constrained rigid. We conclude that (QT $_1$) holds in this case. \square

Remark 5.15 As in [10], the proof of Corollary 5.14 allows us to control the conformal structures of the sources of a sequence of rigid disks $u_\eta: \Delta_m \rightarrow T^*S^2$ with boundary on $\Lambda_{K;\eta}$ in the following way. The distance from a boundary minimum that maps

near a trivalent puncture of the tree to its nearest boundary minimum equals $c\eta^{-1} + \mathcal{O}(\log(\eta^{-1}))$, where c is a constant determined by the quantum flow tree. The distance between other boundary minima equals $c' + \mathfrak{o}(1)$, where c' depends only on the disk component (and its marked points) of the quantum limit tree. Thus, the conformal structure of the big disk part converges to that of the limiting disk and the conformal structures (represented as truncated standard domains) of the flow tree parts converge after rescaling by η^{-1} .

5.4 From quantum flow trees to disks

In this section we construct rigid J -holomorphic disks near any rigid quantum flow tree. Technical results needed for this were already developed in [10] and [13]. Here we will thus present the main steps together with detailed references to these two papers. The construction follows the standard gluing scheme often used in Floer theory: we associate an approximately holomorphic disk to each rigid quantum flow tree and use Floer's Picard Lemma to show that near each such disk there is a unique actual holomorphic disk. We begin with the following observations.

5.4.1 Properties of rigid quantum flow trees We start by recalling some of the properties of rigid quantum flow trees of $\Lambda_K \subset N \subset T^*(S^2)$ that will be used below. We will write (u, Γ) for a quantum flow tree where $u: \Delta_m \rightarrow T^*S^2$ is the holomorphic disk with boundary on Λ part of the quantum flow tree and where Γ denotes the flow tree part.

- Rigid quantum flow trees (u, Γ) of Λ_K are of two main types: those with u constant and those with u non-constant.
- If (u, Γ) is a rigid quantum flow tree with u non-constant then u is either rigid or constrained rigid.
- If (u, Γ) is a rigid quantum flow tree then it consists of a (constrained) rigid disk with a finite number of partial flow trees Γ attached along its boundary at junction points. In the case that u is constrained rigid then the constraint is at the image in $\Pi(\Lambda)$ of a Reeb chord of type \mathbf{S}_1 . The partial flow trees attached to the holomorphic disk have trivalent Y_0 -vertices, 1-valent vertices at critical points of index 1, and no other vertices.

Rigid quantum flow trees (u, Γ) with u constant are rigid flow trees of $\Lambda_{K;\eta} \subset J^1(\Lambda)$ in the sense of [10]. By Lemma 5.7, holomorphic disks with positive puncture at a Reeb chord of type \mathbf{S} lie in an $\mathcal{O}(\eta)$ -neighborhood of $\Pi(\Lambda)$. Consequently existence and uniqueness of rigid holomorphic disks near each local rigid quantum flow tree

follows from [10, Theorem 1.2]. (In fact, the convergence rate is $\mathcal{O}(\eta \log(\eta^{-1}))$; see [10].) Because of this we will mainly focus on quantum flow trees with non-constant u below.

Consider a rigid quantum flow tree (u, Γ) and let Γ' denote one of the partial flow trees attached to u at the junction point p , which is then also the special puncture of Γ' . Recall that we choose a metric g on Λ and an almost complex structure on T^*S^2 that agrees with that in a neighborhood of $\Pi(\Lambda)$ and that has the following additional properties:

- The metric is flat in a neighborhood of Γ' .
- $\Pi(\Lambda_{K;\eta}) \subset T^*\Lambda$ is affine outside neighborhoods of a finite number of edge points, which are points on the edges of Γ' , at least one on each edge. We call these neighborhoods *edge point regions*. Furthermore, along any edge $\Pi(\Lambda_{K;\eta})$ is a product of a curve in the direction of the edge and horizontal line segments perpendicular to it.
- Near each trivalent vertex $\Pi(\Lambda_{K;\eta}) \subset T^*\Lambda$ is parallel to the 0-section, and near each critical point $\Pi(\Lambda_{K;\eta}) \subset T^*\Lambda$ is affine.

5.4.2 Local solutions The approximately holomorphic disks near the rigid quantum flow tree (u, Γ) will be constructed by patching local solutions. The local solutions are as follows:

- The map $u: \Delta_m \rightarrow T^*S^2$. Here the standard domain has punctures mapping to double points of $\Pi(\Lambda)$ as well as to each junction point and constraining critical points. See [13, Equations (6.9) and (6.10)] for explicit forms of u near junction points and double points. The normal form at a constraining puncture is the same as that at a junction point; see [13, Section 6.5.1 (gd.2)].
- Along each part of an edge between edge points where $\Pi(\Lambda_K)$ is affine we have a local solution $s_\eta: [T_1, T_2] \times [0, 1] \rightarrow T^*S^2$, which is a holomorphic strip with image in the strip region that consists of straight line segments in the fibers connecting the two sheets of the tree. See [10, Section 6.1.1] for details.
- At each 1-valent puncture we choose coordinates (x_1, x_2) along $\Pi(\Lambda)$ with the critical point at 0 and corresponding holomorphic coordinates $(z_1, z_2) = (x_1 + iy_1, x_2 + iy_2)$ so that the flow line of the tree lies along the x_1 -direction. The local solution $s_\eta: [0, \infty) \times [0, 1]$ is then

$$s_\eta(\tau + it) = (\eta ic_1 + c_\eta e^{-\theta_\eta(\tau + it)}, 0),$$

where $\pi - \theta_\eta$ is the largest complex (Kähler) angle of the intersection point and where c_η is chosen so that the distance from $s_\eta(0 + it)$ to the nearest edge point is $\mathcal{O}(\eta)$. See [10, Section 6.1.2] for details.

- At each trivalent vertex we chose \mathbb{C}^2 -coordinates as above with the three sheets corresponding to constant sections with values iv_j , $j = 1, 2, 3$, where a flow line of $v_1 - v_2$ breaks into flow lines of $v_1 - v_3$ and $v_3 - v_2$. Let p_1, p_2, p_3 denote the punctures of a standard domain Δ_3 . Consider the biholomorphic maps

$$U_j: \Delta_3 \rightarrow \mathbb{R} \times [0, 1], \quad j = 2, 3,$$

with $U_2(p_1) = -\infty$, $U_2(p_3) = i$, $U_2(p_2) = \infty$, $U_3(p_1) = -\infty$, $U_3(p_2) = 0$ and $U_3(p_3) = \infty$. Let $a_j: \mathbb{R} \times [0, 1] \rightarrow \mathbb{C}^2$, $j = 2, 3$ be the maps $a_2(z) = (v_3 - v_2)z + iv_2$, $a_3(z) = (v_1 - v_2)z$. Then s_η is a restriction to a subdomain of Δ_3 cut off by vertical segments of the map

$$\tilde{s}_\eta = \eta(a_2 \circ U_2 + a_3 \circ U_3)$$

such that the vertical segments lie at distance $\mathcal{O}(\eta \log(\eta^{-1}))$ from an ϵ -sphere around 0. See [10, Section 6.1.5] for details.

- At each junction point we choose \mathbb{C}^2 -coordinates so that the two sheets of $\Pi(\Lambda_{K;\eta})$ correspond to \mathbb{R}^2 , and to the section $(i\eta, 0)$, respectively. We take

$$w_\eta^{\text{jun}}(z) = (\eta z, 0) + \sum_n c_n e^{n\pi z}, \quad c_n \in \mathbb{R}^2,$$

where the latter sum agrees with the Fourier expansion of u in the strip neighborhood of the junction point.

5.4.3 The domain of a rigid quantum flow tree and approximately holomorphic disks Consider a rigid flow tree (u, Γ) . The local solutions discussed above are associated with domains that are determined by requiring that their vertical boundary segments map close to edge points. There are in particular finite strip regions of length bounded below by $c\eta^{-1}$ and above by $C\eta^{-1}$ associated to each segment of an edge between edge points, finite neighborhoods of the boundary minimum in Δ_3 of the same size around each trivalent vertex, as well as half-infinite strips associated to critical points. These regions are patched together over uniformly finite size rectangles corresponding to the edge points where we also interpolate between local solutions. We thus get a domain $\bar{\Delta}_{p,m}(\Gamma', \eta)$ for each tree Γ' in Γ . We construct the domains $\Delta_{p,m}(u, \Gamma, \eta)$ by gluing the tree domains to the disk domain in the strip regions corresponding to the junction points; see Figure 25.

In this way we obtain the desired domain with a map $w_\eta: \Delta_{p,m}(u, \Gamma, \eta) \rightarrow T^*S^2$ obtained by patching local solutions, which is approximately holomorphic in a sense that we will next make precise.

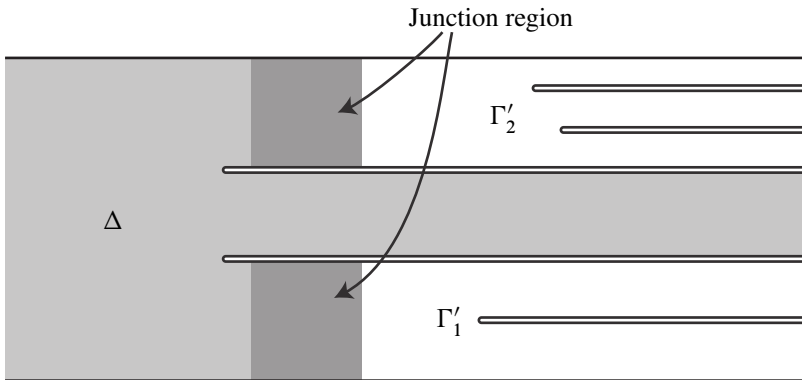


Figure 25: The domain $\Delta_{p,m}(u, \Gamma, \eta)$ for a quantum flow tree. The light shaded region is the domain for the disk Δ with boundary on Λ , while the regions Γ'_1 and Γ'_2 are domains for two gradient flow trees that are glued to the disk Γ .

In order to prove existence and uniqueness of holomorphic disks near rigid flow trees we need an appropriate functional analytic setting for Fredholm theory. Here one cannot use standard Sobolev spaces because the domains are degenerating near the limit and derivatives of maps go to 0 accordingly. For this reason we use weighted Sobolev spaces. The norms of the natural vector fields associated to shifting the local solutions are then unbounded as $\eta \rightarrow 0$. To correct this we use instead a subspace of the Sobolev space determined by a vanishing condition at a marked point in the middle of each strip region between edge points and add a finite-dimensional space of shifts endowed with the supremum norm. The total configuration space is then obtained by adding conformal variations of the target, which corresponds to moving boundary minima and the marked points of the above mentioned vanishing conditions. More precisely, the functional analytic spaces are constructed as follows:

- The domain $\Delta_{p,m}$ of u is cut off as described in [13, Section 6.5.1], at vertical segments corresponding to junction points and to punctures. In the finite strip regions near the cut-offs corresponding to junction points of the form $[0, d\eta^{-1}] \times [0, 1]$ we use an exponential weight peaked in the middle of this strip; see [13, Equation (6.12)].
- For each flow tree Γ' in Γ we take the domain $\bar{\Delta}_{p',m'}(\Gamma', \eta)$ associated to a partial flow tree as in [10, Section 6.4.1], cut off at a vertical segment corresponding to the edge point closest to its special puncture, with the weight function constructed there. We attach these domains at the vertical segments of the corresponding junction points. Note that the weight functions match.

We denote the resulting standard domain $\Delta_{p,m}^\eta$ and the weight function $h: \Delta_{p,m}^\eta \rightarrow [1, \infty)$. Associated to the junction points are spaces of cut off local solutions. We denote the direct sum of these spaces with the supremum norm

$$V_{\text{sol}}^{\text{jun}}.$$

See [13, Section 6.5.1, page 66]. Associated to each partial flow tree there is a space of cut off solutions; we denote the sum of these spaces with the supremum norm

$$V_{\text{sol}}^\Gamma.$$

See [10, Section 5.4.1, page 1209]. Furthermore, conformal variations corresponding to moving boundary minima in the domain of u give conformal variations of $\Delta_{p,m}$; we denote the sum of these spaces with the supremum norm

$$V_{\text{con}}^u.$$

See [10, Sections 6.2.3 and 6.3.5]. Finally there are also conformal variations in the domains of the trees corresponding to moving boundary minima and marked points. They form a space $V_{\text{con}}(\Gamma', \eta)$ and we denote the sum over all trees in Γ

$$V_{\text{con}}^\Gamma.$$

Finally, we denote the space of vector fields that vanish at the marked points in $\Delta_{p,m}^\eta$ and that satisfy Lagrangian boundary conditions and are holomorphic along the boundary, and that have two derivatives in L^2 weighted by h

$$\mathcal{H}_{2,\delta},$$

where δ denotes the small positive exponential weight that controls the size of h . Then, as in [13], we view the $\bar{\partial}$ -operator on function in a neighborhood of a tree as a Fredholm map

$$\bar{\partial}_J: \mathcal{H}_{2,\delta} \oplus V_{\text{sol}} \oplus V_{\text{con}} \rightarrow \mathcal{H}_{1,\delta},$$

where $\mathcal{H}_{1,\delta}$ is the Sobolev space of vector fields with one derivative in L^2 weighted by h . We denote the norm in $\mathcal{H}_{1,\delta}$ by $\|\cdot\|_{1,\delta}$ and that in $\mathcal{H}_{2,\delta} \oplus V_{\text{sol}} \oplus V_{\text{con}}$ by $\|\cdot\|_{2,\delta}$.

The proof now follows the same steps as in [13]. First we estimate the approximate solution:

Lemma 5.16 *The function w_η satisfies*

$$\|\bar{\partial}_J w_\eta\|_{1,\delta} = \mathcal{O}(\eta^{3/4-\delta} \log(\eta^{-1})).$$

Proof The restriction of $\bar{\partial}_J w_\eta$ to the part of the domain corresponding to flow trees is controlled by [10, Remark 6.16]. The proof is then a repetition of the proof of [13, Lemma 6.20]. \square

Second we show that the differential of $\bar{\partial}_J$ is invertible.

Lemma 5.17 *The differential*

$$L\bar{\partial}: \mathcal{H}_{2,\delta} \oplus V_{\text{sol}} \oplus V_{\text{con}} \rightarrow \mathcal{H}_{1,\delta}$$

is uniformly invertible.

Proof After replacing the vector field v_η^{mo} in the proof of [13, Lemma 6.21] with a sum $v^{\text{mo}} = a^\Gamma + b^\Gamma$ where a^Γ is a vector field in $\mathcal{H}_{2,\delta}$ supported in the part of the domain corresponding to Γ and $b^\Gamma \in V_{\text{sol}}^\Gamma \oplus V_{\text{con}}^\Gamma$, and using the inductive procedure for obtaining estimates in rigid flow trees from corresponding flow subtrees [10, Proposition 6.20, pages 1211–1213] to control v^{mo} , the proof is a word by word repetition of the proof of [13, Lemma 6.21]. \square

Third we establish a quadratic estimate for the $\bar{\partial}_J$ -map. We let w_η correspond to $0 \in \mathcal{H}_{2,\delta} \oplus V_{\text{sol}} \oplus V_{\text{con}}$.

Lemma 5.18 *There exists a constant C so that*

$$\bar{\partial}_J(v) = \bar{\partial}_J(0) + L\bar{\partial}_J(v) + N(v),$$

where

$$\|N(v_1) - N(v_2)\|_{2,\delta} = C(\|v_1\|_{2,\delta} + \|v_2\|_{2,\delta})\|v_1 - v_2\|_{2,\delta}.$$

Proof See [13, Lemma 6.22]. \square

With these results established we get the following result as a consequence of Floer's Picard Lemma:

Corollary 5.19 *For all sufficiently small $\eta > 0$ there exists a unique rigid holomorphic disk in a finite $\|\cdot\|_{2,\delta}$ -neighborhood of w_η .*

The last lemma needed to show the correspondence is the following.

Lemma 5.20 *For sufficiently small $\eta > 0$, if a holomorphic disk lies in a sufficiently small C^0 -neighborhood of w_η , then it lies in a $\mathcal{O}(\eta^{1/2})$ $\|\cdot\|_{2,\delta}$ -neighborhood of it.*

Proof See [13, Lemma 6.24]. \square

5.5 Proof of Theorem 4.3

Consider a rigid holomorphic disk with boundary on $\Lambda_{K;\eta}$. By Corollary 5.14, for all sufficiently small η , u_η lies in a neighborhood of a rigid quantum flow tree and we map u_η to that tree. Consider next a rigid quantum flow tree. By Corollary 5.19, there exists a holomorphic disk u_η in a small functional analytic neighborhood of its approximately holomorphic map w_η and we map the quantum flow tree to u_η . Composing these maps starting with a quantum tree it is clear that we get the identity. Composing them starting with a holomorphic disk we also get the identity, by Lemma 5.20: the disk being C^0 -close to the approximate solution w_η , which follows by definition, in fact implies that it is also close functional analytically, and therefore unique by Corollary 5.19.

Finally, applying the correspondence between (constrained) rigid flow trees and holomorphic disks for the conormal of the unknot in Theorem 5.1 we get a one-to-one correspondence between rigid quantum flow trees and rigid multiscale trees. \square

6 Orientations

The main purpose of this section is to prove Theorems 4.5 and 4.6. To this end we first give an overview of the general orientation scheme constructed in [16] and then interpret this scheme in geometric terms for rigid holomorphic disks near quantum flow trees of Λ_K .

6.1 The general orientation scheme

We first give a rough outline of the orientation scheme that we will employ below. For simplicity we restrict attention to closed orientable Legendrian surfaces Λ inside the 1-jet space $J^1(S)$ of some orientable surface S . Let $\bar{\partial}$ be the operator on functions $v: D \rightarrow \mathbb{C}^n$, where D is the unit disk in \mathbb{C} , such that $v(e^{i\theta}) \in L(\theta)$, where $L(\theta)$ is a *trivialized* Lagrangian boundary condition. The starting point for constructing orientations of moduli spaces of holomorphic disks with boundary on Λ is the fact that the index bundle of $\bar{\partial}$ is orientable, and that a choice of orientation on \mathbb{C} and on \mathbb{R}^n determines an orientation on this index bundle. We call this induced orientation the *canonical orientation*; see [20, Proposition 8.1.4] (or [16, Section 3.3]).

Consider now a holomorphic disk $u: \Delta_m \rightarrow T^*S$ with boundary on Λ . In order to parameterize the moduli space of holomorphic disks near u , we look at the $\bar{\partial}_J$ -operator as a section of the bundle of complex anti-linear maps over the configuration space with local chart at u given by $\mathcal{H} \oplus \mathcal{C}_m$, where \mathcal{H} is a Sobolev space of vector fields along u and where \mathcal{C}_m denotes the space of conformal structures on the source Δ_m

of u . In this setting the tangent space of the moduli space can be identified with the kernel of the linearized operator $L\bar{\partial}_J$ acting on $\mathcal{H} \oplus \mathcal{C}_m$. In order to orient the moduli space we choose capping operators at all Reeb chords of u with oriented determinant bundles. Then closing up the boundary condition of u by these capping operators gives a Lagrangian boundary condition on the closed disk and, provided the Lagrangian boundary conditions are compatibly trivialized, we may use the canonical orientation to orient the determinant of the resulting operator. The orientation of the glued problem and orientations on all capping operators induce an orientation of the determinant of the original linearized problem. For Legendrian homology in general, the exact sequence which relates all the orientations of the different operators was introduced in [16, Equation 3.17]. Together with an orientation of the space of conformal structures this then gives an orientation on the tangent space to the moduli spaces, as described in [16, Remark 3.18].

Appropriate trivializations on the boundary condition of u can be defined provided Λ is spin. In order to have the above scheme compatible with disk breaking at the boundary of the compactified moduli space one must choose oriented capping operators at Reeb chords as positive and negative punctures that add to the trivialized boundary condition with the canonical orientation. We note also that when discussing orientations we can stabilize the operator and add oriented finite-dimensional spaces to the source or target of an operator or take direct sums with other oriented Fredholm problems as long as we keep track of the orientations that these extra directions carry.

6.2 Basic choices for the canonical orientation

As mentioned in Section 6.1 the orientation scheme that we use derives from orientation properties of the index bundle over trivialized Lagrangian boundary conditions on the disk. In this section we study some of the details of this construction. Let L be an n -dimensional Lagrangian boundary condition on the unit disk $D \subset \mathbb{C}$, which is trivialized. Consider the $\bar{\partial}$ -operator acting on vector fields $v: D \rightarrow \mathbb{C}^n$ that satisfy the boundary condition given by L , $v(e^{i\theta}) \in L(e^{i\theta})$. Denote this operator $\bar{\partial}_L$. Then $\bar{\partial}_L$ is a Fredholm operator of index

$$\text{index}(\bar{\partial}_L) = n + \mu(L),$$

where μ is the Maslov index. Since the boundary condition is trivialized, the monodromy of L is orientation-preserving and $\mu(L)$ is even.

As mentioned above, an orientation of \mathbb{R}^n together with a choice of complex orientation in \mathbb{C} induces a canonical orientation on the determinant of the operator $\bar{\partial}_L$. This canonical orientation is obtained by trivializing the complex bundle over almost all

of D , splitting off a complex vector bundle over $\mathbb{C}P^1$ at the center of the disk in the limit. The split problem in the limit consists of the operator $\bar{\partial}_{\mathbb{R}^n}$ on D , where \mathbb{R}^n denotes the constant trivialized Lagrangian boundary condition given by $\mathbb{R}^n \subset \mathbb{C}^n$ with the standard basis, and the $\bar{\partial}$ -operator on a complex vector bundle over $\mathbb{C}P^1$. The determinant of the latter operator is a wedge product of complex vector spaces and hence has an orientation induced from the choice of complex orientation on \mathbb{C} . The former operator has trivial cokernel and kernel spanned by constant sections with values in \mathbb{R}^n ; hence, the orientation of \mathbb{R}^n gives an orientation on its determinant. For operators near the split limit the kernel and cokernel are isomorphic to the sum of kernel and cokernels of the split problem and we get an induced canonical orientation of the determinant of the original problem by transporting this orientation along the path of the deformation. We call the choice of orientation of \mathbb{R}^n and \mathbb{C} *basic orientation choices*.

Remark 6.1 Changing the choice of basic orientation on \mathbb{R}^n clearly changes the canonical orientation of $\det(\bar{\partial}_L)$ for every L . Changing the choice of basic orientation of \mathbb{C} preserves the canonical orientation of $\det(\bar{\partial}_L)$ for all L with $\text{index}(\bar{\partial}_L) - n = \mu(L)$ divisible by 4 and reverses the canonical orientation of $\det(\bar{\partial}_L)$ for all L with $\text{index}(\bar{\partial}_L) - n = \mu(L)$ not divisible by 4.

In our calculations below the orientation on \mathbb{R}^n above will correspond to the choice of an orientation on the conormal lift of the unknot, which we take as fixed once and for all. We will denote the chosen basic orientation on \mathbb{C} by $\sigma_{\mathbb{C}}$.

Furthermore, the trivialization of the boundary conditions of the linearized operators at holomorphic disks with boundary on Λ_K are induced from a trivialization of the tangent bundle $T\Lambda_K$. Note that the orientation above depends only on the trivialization modulo 2, ie, on the corresponding spin structure. In the calculations below we will fix a spin structure on the conormal lift Λ of the unknot and pull it back to Λ_K under the natural projection in the 1-jet neighborhood of Λ .

6.3 Capping operators and orientation data at Reeb chords

As mentioned in Section 6.1 orientations of moduli spaces are constructed using capping operators at Reeb chords. In this section we discuss capping operators and their orientations for the conormal lift Λ_K of a closed braid K by applying [16, Section 3.3] to Λ_K .

6.3.1 Auxiliary directions Before we start this discussion we note that the construction of coherent orientations in [16] uses auxiliary directions. More precisely, the

boundary condition of a punctured disk with boundary on Λ is stabilized, ie, multiplied by boundary conditions for vector fields in \mathbb{C}^2 with boundary conditions close to constant \mathbb{R}^2 boundary conditions; see [16, Section 3.4.2]. The resulting problem is split and the operator in the auxiliary direction is an isomorphism. The reason for introducing these auxiliary directions is that they enable a continuous choice of capping operators for varying Legendrian submanifolds and thereby simplify the proof of invariance. When studying kernels and cokernels of linearized operators below we need not consider the auxiliary directions. We do include the auxiliary directions in order to use the capping operators in [16], which we discuss next.

We first describe capping operators of the Reeb chords of Λ_K and then connect the operators to geometry. In Tables 2–5 below we represent the Legendrian submanifold near its Reeb chords. For detailed properties of the capping operators, we refer to [16, Section 3.3.6]. The actual Legendrian $\Lambda_K \subset J^1(S^2)$ is stabilized and appears as the restriction of an embedding $\Lambda_K \times (-\epsilon, \epsilon)^2 \rightarrow J^1(S^2 \times (-\epsilon, \epsilon)^2)$ to $(0, 0) \in (-\epsilon, \epsilon)^2$. Reeb chords are generic and hence locally their fronts are determined by the Hessian for the function difference of the local sheets of Reeb chord endpoints. In [16], the stabilization is constructed in such a way that the two eigenvalues of the Hessian of smallest norm are positive and lie in the auxiliary direction. Here we take these eigenvalues to be of largest norm for simpler combinatorics; see Remark 6.2. In Tables 2–5 the corresponding eigendirections are denoted Aux_j , $j = 1, 2$, the two remaining eigendirections are denoted Real_j , $j = 1, 2$, and we will use the following notation:

- $\delta, \delta', \delta'', \delta_0$ are numbers such that $0 < \delta' < \delta'' < \delta < \delta_0$, and such that δ', δ'' , and δ all approach 0 as the conormal lift of the link approaches the 0–section.
- If $\lambda = (\lambda_1, \dots, \lambda_m)$ is a collection of paths of Lagrangian subspaces such that the endpoint of λ_j is transverse to the start point of λ_{j+1} then $\widehat{\mu}(\lambda)$ denotes the Maslov index of the loop of Lagrangian subspaces obtained by closing up the collection of paths by rotating the incoming subspace to the outgoing one in the negative direction. Thus $n + \widehat{\mu}(\lambda)$ is the index of the $\bar{\partial}$ –operator on the $(m + 1)$ –punctured disk with boundary conditions given by λ , where n is the dimension of the Lagrangian subspaces.
- The expressions “Coker, const” and “Ker, const” indicate that the cokernel and the kernel of some operator can be represented by constant functions, ie, the actual kernel or cokernel functions are approximately constant in the sense that they converge to constant functions on any compact subset as the conormal of the link approaches the 0–section.

Remark 6.2 As mentioned above, our choice of eigenvalues in the auxiliary directions differs from that in [16] and therefore the capping operators differ as well. It is easy to

see that this does not affect the DGA: there is a DGA isomorphism relating the two choices that take each Reeb chord generator c of the algebra to $\pm c$, where the sign depends on the choice of orientation of the capping operator. In fact, the same argument shows that any capping operator would do as long as the negative and positive capping operators glue to an operator with the canonical orientation. Our particular choice of auxiliary directions here was made in order to allow for a uniform treatment of all kernel and cokernel elements in the auxiliary directions; see Remark 6.3.

Remark 6.3 As mentioned above, after the capping operators have been defined, we may disregard the auxiliary directions when studying orientations. The reason for this is that the stabilized boundary conditions of the linearized operator split into real and auxiliary directions. Here the boundary conditions in the auxiliary directions give an operator that is an isomorphism over any disk with one positive puncture and boundary on Λ_K , the capping operators in the auxiliary directions at all negative punctures are isomorphisms as well, and the capping operator in the auxiliary directions at the positive puncture is independent of the particular puncture and has index -2 and trivial kernel. Thus using the canonical orientation for the isomorphisms and fixing an oriented basis in the cokernel in the auxiliary directions of the capping operators at positive punctures we get induced orientations of moduli spaces of holomorphic disks with boundary on Λ_K . If the orientation of the cokernel of the positive capping operator is changed then the orientations of moduli spaces change by an overall sign. Thus, auxiliary directions affect the differential in the Legendrian contact homology algebra of Λ_K only by an overall sign. We therefore suppress auxiliary directions in our calculations below.

6.3.2 Details for the orientation data at a chord of type S_1 As indicated in Corollary 5.14 a Reeb chord b of Λ_K of type S_1 can appear as a positive puncture for a disk of type (QT_\emptyset) and a negative puncture for a disk of type (QT_1) . Let $\bar{\partial}_{b+}$ (respectively, $\bar{\partial}_{b-}$) denote the capping operator associated to the Reeb chord b of Λ_K when it appears as a positive (respectively, negative) puncture of the J_η -holomorphic (see (5-1)) disk u_η . Although as Lemma 3.7 indicates, there are many such chords of type S_1 , each is a parallel translate of the other; thus, we can consider their capping operators simultaneously.

At a Reeb chord b of Λ_K of type S_1 , $\bar{\partial}_{b-}$ splits into two 1-dimensional problems. Recall that auxiliary directions are disregarded. Because the grading of b is odd, the conventions set in [16, Section 3.3.6] imply one 1-dimensional problem has index 1 with 1-dimensional kernel and the other has index -1 with 1-dimensional cokernel. This is indicated in the two left “Real” columns at the bottom row of Table 2.

Consider first the index 1 component. This operator is an operator on a Sobolev space

	Aux ₁	Aux ₂	Real ₁	Real ₂
Front				
Lagrangian				
Complex angle, positive puncture	 δ	 δ	 $\pi - \delta'$	 $\pi - \delta''$
Complex angle, negative puncture	 $\pi - \delta$	 $\pi - \delta$	 δ'	 δ''
Closing rotation, positive puncture	 $-(\pi - \delta)$	 $-(\pi - \delta)$	 $-\delta'$	 $-\delta''$
Closing rotation, negative puncture	 $-\delta$	 $-\delta$	 $-(\pi - \delta')$	 $-(\pi - \delta'')$
Capping operator, positive puncture	 $e^{-i(2\pi-\delta)s}$ $\hat{\mu} = -2$ index = -1 Coker, const	 $e^{-i(2\pi-\delta)s}$ $\hat{\mu} = -2$ index = -1 Coker, const	 $e^{-i\delta's}$ $\hat{\mu} = -1$ index = 0 Isomorphism	 $e^{-i\delta''s}$ $\hat{\mu} = -1$ index = 0 Isomorphism
Capping operator, negative puncture	 $e^{-i\delta s}$ $\hat{\mu} = -1$ index = 0 Isomorphism	 $e^{-i\delta s}$ $\hat{\mu} = -1$ index = 0 Isomorphism	 $e^{i\delta's}$ $\hat{\mu} = 0$ index = 1 Ker, const	 $e^{-i(2\pi-\delta'')s}$ $\hat{\mu} = -2$ index = -1 Coker, const

Table 2: Capping operator at a chord of type S_1 . The direction corresponding to the angle δ' is along the flow line in the direction of the unknot parameter. The direction corresponding to the angle δ'' is the direction of the fiber of the conormal bundle.

	Aux ₁	Aux ₂	Real ₁	Real ₂
Front				
Lagrangian				
Complex angle, negative puncture	 $\pi - \delta$	 $\pi - \delta$	 $\pi - \delta'$	 δ''
Closing rotation, negative puncture	 $-\delta$	 $-\delta$	 $-\delta'$	 $-(\pi - \delta'')$
Capping operator, negative puncture	 $e^{-i\delta s}$ $\hat{\mu} = -1$ index = 0 Isomorphism	 $e^{-i\delta s}$ $\hat{\mu} = -1$ index = 0 Isomorphism	 $e^{-i\delta' s}$ $\hat{\mu} = -1$ index = 0 Isomorphism	 $e^{i\delta'' s}$ $\hat{\mu} = 0$ index = 1 Ker, const

Table 3: Local data at a chord of type S_0

of complex-valued functions on the disk with one boundary puncture. As the parameter $\delta' \rightarrow 0$ in Table 2 we continue the operator family continuously to the limit by introducing a small negative exponential weight in a strip neighborhood of the puncture. In the limit, the kernel is spanned by a constant real-valued function. By continuity, it follows that solutions near the limit are close to constant functions, and in particular, the L^2 -pairing with a kernel function is close to the L^2 -pairing of the corresponding constant function in the limit. We thus fix an orientation of the kernel of the index 1 component of the capping operator at b by fixing a vector

$$v^{\ker}(b) \in T_b \Lambda$$

parallel to the direction of the knot. To see that $v^{\ker}(b)$ should be chosen parallel to the knot, note that the direction of the knot is the direction of the Bott manifold of Reeb chords of Λ_U and hence corresponds to the smaller of the two eigenvalues of the Hessian.

Similarly, elements in the cokernel of the index -1 component are solutions to a dual boundary value problem for the ∂ -operator. The cokernel functions converge to constant real valued functions as $\delta'' \rightarrow 0$ and we fix an orientation in the cokernel by

	Aux ₁	Aux ₂	Real ₁	Real ₂
Front				
Lagrangian				
Complex angle, positive puncture	 δ	 δ	 δ'	 $\pi - \delta_0$
Complex angle, negative puncture	 $\pi - \delta$	 $\pi - \delta$	 $\pi - \delta'$	 δ_0
Closing rotation, positive puncture	 $-(\pi - \delta)$	 $-(\pi - \delta)$	 $-(\pi - \delta')$	 $-\delta_0$
Closing rotation, negative puncture	 $-\delta$	 $-\delta$	 $-\delta'$	 $-(\pi - \delta_0)$
Capping operator, positive puncture	 $e^{-i(2\pi-\delta)s}$ $\hat{\mu} = -2$ index = -1 Coker, const	 $e^{-i(2\pi-\delta)s}$ $\hat{\mu} = -2$ index = -1 Coker, const	 $e^{-i(\pi-\delta')s}$ $\hat{\mu} = -1$ index = 0 Isomorphism	 $e^{-i\delta_0s}$ $\hat{\mu} = -1$ index = 0 Isomorphism
Capping operator, negative puncture	 $e^{-i\delta s}$ $\hat{\mu} = -1$ index = 0 Isomorphism	 $e^{-i\delta s}$ $\hat{\mu} = -1$ index = 0 Isomorphism	 $e^{-i(\pi+\delta')s}$ $\hat{\mu} = -2$ index = -1 Coker, const	 $e^{i\delta_0s}$ $\hat{\mu} = 0$ index = 1 Ker, const

Table 4: Capping operator at a chord of type L_1 . The direction corresponding to δ' is along the equator (along the parameter of the unknot). The direction corresponding to δ_0 is perpendicular to the equator (along the fiber).

fixing a vector

$$v^{\text{coker}}(b) \in T_b\Lambda,$$

perpendicular to the direction of the knot, which is the direction corresponding to the positive eigenvalue of the Hessian along the Bott manifold of Reeb chords of Λ_U and

	Aux ₁	Aux ₂	Real ₁	Real ₂
Front				
Lagrangian				
Complex angle, positive puncture	 δ	 δ	 $\pi - \delta'$	 $\pi - \delta_0$
Closing rotation, positive puncture	 $-(\pi - \delta)$	 $-(\pi - \delta)$	 $-\delta'$	 $-\delta_0$
Capping operator, positive puncture	 $e^{-i(2\pi-\delta)s}$ $\hat{\mu} = -2$ index = -1 Coker, const	 $e^{-i(2\pi-\delta)s}$ $\hat{\mu} = -2$ index = -1 Coker, const	 $e^{-i\delta's}$ $\hat{\mu} = -1$ index = 0 Isomorphism	 $e^{-i(\pi+\delta_0)s}$ $\hat{\mu} = -2$ index = -1 Coker, const

Table 5: Capping operator at a chord of type L_2 . The direction of δ' is along the equator. The direction of δ_0 perpendicular to it.

hence corresponds to the largest eigenvalue after perturbation. The basis

$$(v^{\text{coker}}(b), v^{\text{ker}}(b))$$

determines the orientation of the operator $\bar{\partial}_{b-}$, and constant functions with values in the lines spanned by the basis vectors are approximate solutions.

When b appears as a negative puncture, we see from its parity and the conventions of [16, Section 3.3.6] that both the kernel and cokernel are trivial.

6.3.3 Orientation data at a chord of type S_0, L_1 or L_2 The discussion for the other types of chords is similar to that of chords of type S_1 in Section 6.3.2, so we discuss them only briefly here. Recall that while a chord of type L_1 can appear as a positive or negative puncture, chords of the other two types only occur as punctures of one sign.

A Reeb chord a of Λ_K of type S_0 has even parity and only appears as a negative puncture. Thus, as indicated in Table 3, the capping operator $\bar{\partial}_{a-}$ splits into two 1-dimensional problems, one of index 0, which is an isomorphism and one of index 1 with 1-dimensional kernel. As in Section 6.3.2 the kernel functions of the index 1

component are approximately constant and we fix an orientation of the capping operator by fixing a vector

$$v^{\ker}(a) \in T_a \Lambda$$

perpendicular to the direction of the knot. We can assume $v^{\ker}(a)$ is perpendicular to the knot because the second real coordinate (the last column in Table 3) represents the unstable manifold of a , thought of as a braid saddle point. Constant functions with values in the line spanned by this vector are approximate solutions of $\bar{\partial}_{a^-}$.

Noting the parity of the grading of a chord c of type \mathbf{L}_1 and a chord e of type \mathbf{L}_2 , the conventions of [16, Section 3.3.6] imply the two capping operators each split into two 1-dimensional problems with kernels and cokernels as indicated in Tables 4 and 5. As above, we fix orientations on these problems by fixing vectors $v^{\ker}(c)$, $v^{\text{coker}}(c)$ and $v^{\text{coker}}(e)$.

6.4 Signs in the unknot differential — proof of Theorem 4.6

Fix the Lie group spin structure on Λ . The corresponding trivialization of $T\Lambda$ is then the canonical trivialization of the tangent bundle of the 2-torus coming from the identification $T^2 = \mathbb{R}^2 / \mathbb{Z}^2$. Here, we take coordinates on Λ as described in Section 4.

Recall for the unknot there are four rigid flow trees I_N , I_S , Y_N and Y_S with one puncture, which is at c , that contribute to ∂c , and there are two rigid strips with positive puncture at e and negative puncture at c . Consequently, by Theorem 5.1, for $\sigma > 0$ sufficiently small there are four corresponding rigid holomorphic disks with positive puncture at c and two corresponding rigid holomorphic strips E_1 and E_2 connecting e to c . We next compute their signs.

Theorem 6.4 *For any choice of basic orientations there is a choice of capping operator at c such that the signs of the rigid disks satisfy*

$$\begin{aligned} \epsilon(I_S) &= \epsilon(I_N) = \epsilon(Y_S) = \epsilon(Y_N) = 1, \\ \epsilon(E_1) &= -\epsilon(E_2). \end{aligned}$$

Proof We first show that $\epsilon(I_S) = \epsilon(I_N)$. To this end consider a geodesic arc in S^2 that passes through $\Pi(c)$ and that contains both poles. Let x_1 be a coordinate along this arc with $\Pi(c)$ corresponding to 0 and let x_2 be a coordinate perpendicular to the arc. Then $\{x = (x_1, x_2) : |x_1| \leq \pi + \delta, |x_2| < \delta\}$ parametrizes a disk D in S^2 and we find a complex trivialization of the tangent bundle of $T(T^*S^2)$ over D by noting that the metric is flat and using coordinates $(x + iy)$.

Let $\bar{\partial}_c$ denote the positive capping operator at c , let $\bar{\partial}_I$ denote the linearized boundary condition of any one of the four disks with positive puncture at c , and let \hat{I} denote the problem on the closed disk obtained by gluing these two. Then the following gluing sequence is used to orient moduli spaces:

$$0 \rightarrow \ker(\partial_{\hat{I}}) \rightarrow \ker(\bar{\partial}_I) \rightarrow 0$$

See [16, Equation (3.17)] and Table 4. Note first that the trivialized boundary conditions of I_N and I_S agree. The signs of the disks are then obtained by comparing the oriented kernel of $\bar{\partial}_I$ with the orientation induced by conformal automorphism. Note furthermore that if $u: \Delta_1 \rightarrow T^*S^2$ is a parametrization of I_N then $-u$ parametrizes I_S , where $-(x + iy) = (-x - iy)$ in the coordinates discussed above. Since the automorphism group of Δ_1 is 2-dimensional the signs of the two disks agree, $\epsilon(I_N) = \epsilon(I_S)$. An identical argument shows that $\epsilon(Y_N) = \epsilon(Y_S)$.

After noting that the orientation of the capping operator at c determines the sign in the orienting isomorphism above, it remains only to show that $\epsilon(I_N) = \epsilon(Y_N)$ to complete the proof of the first equation. To this end we compare the boundary conditions of $\bar{\partial}_{I_N}$ and $\bar{\partial}_{Y_N}$. Note that the boundary conditions of the disks are arbitrarily close to the boundary conditions of the corresponding trees and that the trees I_N and Y_N are identical except near the north pole, $(x_1, x_2) = (\pi, 0)$. Using the trivialization (over the disk D above) of the $(x + iy)$ -coordinates around the north pole, the Lie group spin of the torus Λ is (∂_s, ∂_t) and induces the trivialized boundary condition

$$(\cos t \partial_{x_1} + \sin t \partial_{x_2}, -\sin t \partial_{y_1} + \cos t \partial_{y_2})$$

on the I_N -disk and

$$(\cos t \partial_{x_1} - \sin t \partial_{x_2}, \sin t \partial_{y_1} + \cos t \partial_{y_2})$$

on the Y_N -disk.

Now, the homotopy of Lagrangian boundary conditions that are given by acting by the complex matrix

$$\begin{pmatrix} 1 & 0 \\ 0 & e^{i\theta} \end{pmatrix}, \quad 0 \leq \theta \leq \pi,$$

takes one trivialization to the other and we conclude that multiplication by the matrix at $\theta = \pi$ takes the positively oriented kernel of $\bar{\partial}_{I_N}$ to that of $\bar{\partial}_{Y_N}$. Comparing this orientation to the orientation induced by source isomorphisms, for example by evaluation at a point where the disks agree, we find that the signs of the two disks agree.

The argument for finding the signs of the two disks E_1 and E_2 is similar to the above arguments: both disks come from flow lines and their boundary conditions are identical. Again the disks are related by multiplication by -1 in suitable coordinates. Here however, the kernel of the linearized operator and the automorphism group are both 1-dimensional and it follows that multiplication by -1 reverses orientation. The lemma follows. \square

In our computations of the signs for the differential in the Legendrian algebra of Λ_K we will use the capping operators of chords of type \mathbf{L}_2 and \mathbf{L}_1 , which correspond to the capping operators of e and c , respectively, for which Theorem 6.4 holds.

6.5 Conformal structures

Conformal parameters for flow trees are best represented as moving boundary minima in standard domains; see Section 5.3.1. In the general orientation scheme of [16] the space of conformal structures was represented as the location of boundary punctures on the unit disk in the complex plane. The main purpose of this section is to relate these two representations in order to allow for the representation best adapted to trees to be used in computations.

Consider first the representation of conformal structures \mathcal{C}_m , used in [16], on the (unit) disk D_m in \mathbb{C} with $m \geq 3$ boundary punctures p_0, \dots, p_{m-1} . Recall punctures are ordered counter-clockwise. Fix the (distinguished) puncture p_0 at 1, p_1 at i , and p_{m-1} at $-i$. Then the locations of the remaining punctures in the boundary arc between i and $-i$ determine the conformal structure uniquely. Thus the space \mathcal{C}_m of conformal structures on the disk with m boundary punctures, one of which is distinguished, is an $(m-3)$ -dimensional simplex. We write b_j for the tangent vector that corresponds to moving the j^{th} puncture p_j in the positive direction and keeping all other punctures fixed. Then b_2, \dots, b_{m-2} is a basis in $T\mathcal{C}_m$.

Consider second the representation of \mathcal{C}_m using standard domains Δ_m . Recall that a standard domain is a strip with slits of fixed width, that a standard domain determines a conformal structure on the disk with m boundary punctures, and that two standard domains determine the same conformal structure if and only if they differ by an overall translation. Assume that $m > 3$ and let $t_j \in T\mathcal{C}_m$ denote the tangent vector that is the first order variation that corresponds to moving the j^{th} boundary minimum toward $-\infty$ and keeping all other boundary minima fixed.

Lemma 6.5 *Let $m > 3$. Then*

$$t_j = \sum_{k=2}^j \xi_k b_k + \sum_{k=j+1}^{m-2} \eta_k b_k \in T\mathcal{C}_m,$$

where $\xi_k > 0, k = 2, \dots, j$ and $\eta_k < 0, k = j + 1, \dots, m - 2$. In particular, we can represent TC_m as the vector space generated by t_1, \dots, t_{m-2} divided by the 1-dimensional subspace generated by the vector

$$t_1 + t_2 + \dots + t_{m-2},$$

and the orientation given by the basis b_2, \dots, b_m agrees with that induced by t_2, \dots, t_m .

Proof Consider the map which takes a neighborhood of $+\infty$ in an infinite strip to a neighborhood of the origin in the upper half plane:

$$w = -e^{-\pi z}.$$

Under such a change of coordinates the vector field 1 on the w -plane corresponds to the vector field $\pi^{-1}e^{\pi z}$ on the strip since

$$1 = \left(\frac{dw}{dz}\right)\pi^{-1}e^{\pi z}.$$

Let $\mathcal{H}_{-\delta;k}$ denote the Sobolev space of vector fields along Δ_m that are tangent to Δ_m along the boundary and with a small negative weight at each puncture, ie, a weight function of the form $e^{-\delta|\tau|}$ in a strip region, $\tau + it \in [0, \infty) \times [0, 1]$ or $\tau + it \in (-\infty, 0] \times [0, 1]$, and with k derivatives in L^2 . The $\bar{\partial}$ -operator $\bar{\partial}: \mathcal{H}_{-\delta;k} \rightarrow \mathcal{H}_{-\delta;k-1}$ has index

$$1 - (m - 2) = -(m - 3),$$

where $1 = \dim(\mathbb{R})$ and $-(m - 2)$ is the Maslov index of the boundary condition with a negative half turn at each boundary minimum. The exact degree of regularity of the vector fields we use will be of no importance and will be dropped from the notation. Let b'_j denote cut-off versions of the vector fields $e^{\pi z}$ supported in the j^{th} strip end. Then we can think of TC_m as the quotient space

$$\bar{\partial}(\mathcal{H}_{-\delta} \oplus \langle b'_2, \dots, b'_{m-2} \rangle) / \bar{\partial}(\mathcal{H}_{-\delta}).$$

In this setting, we can interpret t_j as follows; see [10, Section 2.1.1]. Consider $\Delta_m \subset \mathbb{C}$ and let $z = x + iy$ be the standard complex coordinate on \mathbb{C} . Let T_j denote a vector field on \mathbb{C} supported in a small ball B_r centered at the j^{th} boundary minimum and equal to $-\partial_x$ in $B_{r/2}$ and tangent to the boundary of Δ_m in $B_r - B_{r/2}$. The conformal variation t_j is then represented by $\bar{\partial}T_j \in \mathcal{H}_{-\delta}$. (To see this, linearize the comparison of conformal structures κ and $d\phi^{-1}\kappa d\phi$, where $\phi: \Delta_m \rightarrow \Delta_m$ is a small diffeomorphism associated to the vector field T_j , To first order, $d\phi^{-1}\kappa d\phi \approx \kappa + \kappa \bar{\partial}T_j$.) Because $\{\bar{\partial}b'_k\}_k$

spans the cokernel of the operator $\bar{\partial}$,

$$\bar{\partial}T_j = \sum_{k=2}^{m-2} \alpha_k (\bar{\partial}b'_k) + \bar{\partial}v,$$

for some $v \in \mathcal{H}_{-\delta}$ and real constants $\alpha_2, \dots, \alpha_{m-2}$. We first show that $\alpha_k \neq 0$, $k = 2, \dots, m - 2$. Define the vector field $w: \Delta_m \rightarrow \mathbb{C}$ as $w = T_j - \sum_k \alpha_k b'_k - v$. Then w satisfies a Lagrangian boundary condition of Maslov index $-(m - 3)$ and lies in a Sobolev space with exponential weights as follows. The weight is $-\delta$ in the strip-like end around the puncture at $-\infty$, at the first and last punctures at $+\infty$, and at all punctures for which $\alpha_k = 0$. It is $-\pi - \delta$ at punctures where $\alpha_k \neq 0$. If the number of punctures with $\alpha_k = 0$ is N then the index of the $\bar{\partial}$ -operator on the Sobolev space with Lagrangian boundary condition and weights as just explained equals

$$\text{index}(\bar{\partial}) = 1 - (m - 3) + (m - 3 - N) = 1 - N.$$

Note that $w \neq 0$ since T_j is not tangent at the boundary while $v + \sum_{k=2}^{m-2} \alpha_k b'_k$ is tangent. Since $\bar{\partial}w = 0$ and $w \neq 0$, it follows by automatic transversality in dimension 1 (ie, the argument principle) that $N = 0$, ie, $\alpha_k \neq 0$, for all k .

To determine the signs in the expression for t_j , consider the limit as the shift of the boundary minimum goes to $-\infty$. In this limit the disk Δ_m splits into three components: a three punctured disk containing the puncture at $-\infty$ and two punctures at $+\infty$ where two standard domains $\Delta_{m'}$ and $\Delta_{m''}$ are attached. We choose notation so that the punctures in Δ_m at $+\infty$ below the moving slit end up in $\Delta_{m'}$ and those above in $\Delta_{m''}$. From the point of view of the representations of conformal structures via boundary punctures on the closed disk, the punctures in $\Delta_{m'}$ collide at i and those in $\Delta_{m''}$ collide at $-i$. As the coefficients α_k , $k = 2, \dots, m - 2$ are non-zero for the infinitesimal deformation t_j at each instance of this total deformation, it follows that $\alpha_k > 0$ for $k < j$ and $\alpha_k < 0$ for $k \geq j$. □

Remark 6.6 Lemma 6.5 also has an intuitive justification using harmonic measure. Suppose the conformal structure changes by slightly decreasing the j^{th} boundary minimum. Then the harmonic measure of the j^{th} slit (the probability of a Brownian motion particle first hitting the boundary of Δ_m at that slit) increases while the measures of all other boundary components decrease. Harmonic measure is preserved under a conformal map from Δ_m to D_m . Thus, the corresponding changes in measures of the boundary arcs of D_m can only occur if puncture p_i moves in the negative direction for $i = 2, \dots, j$ and the positive direction for $i = j + 1, \dots, m - 2$.

6.6 Signs of rigid quantum flow trees — proof of Theorem 4.6

In this subsection we compute the signs of rigid quantum flow trees determined by $\Lambda_K \subset J^1(S^2)$ and thereby prove Theorem 4.6. Recall from Corollary 5.14 that there are four types of rigid quantum flow trees: (QT_\emptyset) , (QT_0) , (QT'_0) and (QT_1) . We will consider each case separately.

Let A denote the Lagrangian boundary condition, suppressing auxiliary directions, on the domain Δ_m of the linearized $\bar{\partial}$ -problem corresponding to the J_η -holomorphic disk u_η . Recall that we think of $\eta > 0$ as small; see (5-1). As in [14, Section 6], we must close up the boundary conditions at each puncture using a “negative J -twist.” This is illustrated in Tables 2–5 under the “Closing rotation” row(s). Note that this may contribute to the index of $\bar{\partial}_A$. Let $\bar{\partial}_A$ denote the linearized problem with this boundary. Let \hat{A} denote the boundary condition after adding the appropriate capping operators. Define $\bar{\partial}_{\hat{A}}$ similarly. For each of the four cases above, we must compute the exact sequence [16, Equation 3.17].

6.6.1 The sign of a quantum flow tree of type (QT_\emptyset) A rigid holomorphic disk near the limit in a neighborhood of a quantum flow tree of type (QT_\emptyset) lies in a small neighborhood of a rigid flow tree in $\Lambda \subset J^1(S^2)$ determined by Λ_K . Let Γ be such a rigid flow tree with positive puncture b and negative punctures a_1, \dots, a_{m-1} . Recall that, since the front of $\Lambda_K \subset J^1(\Lambda)$ has no singularities, all vertices of such a rigid tree are trivalent Y_0 -vertices except for 1-valent vertices at Reeb chords. Let t_1, \dots, t_{m-2} denote the trivalent vertices of Γ . Note that each trivalent vertex corresponds to a boundary minimum in the domain Δ_m of the holomorphic disks u_η , which corresponds to Γ for small η . We number the trivalent vertices according to the order of the corresponding boundary minima in the vertical direction of the complex plane. We write τ_j for the boundary minimum corresponding to the trivalent vertex t_j .

Lemma 6.7 *There exists a choice of basic complex orientation $o_{\mathbb{C}}$ such that if $\eta > 0$ is sufficiently small and if u_η is a rigid holomorphic disk in a neighborhood of the rigid flow tree Γ then the sign of u_η is given by*

$$\epsilon(u_\eta) = \epsilon(\Gamma) = \sigma_{\text{pos}}(\Gamma)\sigma(n, \Gamma),$$

where $n = v^{\text{coker}}(b)$; see Section 3.4.4 for notation.

Proof Since η is small, u_η lies close to Γ . Using the trivialization of $T(T^*S^2)$ in a neighborhood of Λ induced by the trivialization of $T(T^*\Lambda)$, the boundary condition A is very close to constant \mathbb{R}^2 boundary conditions (for \mathbb{C}^2 -valued vector fields) on

Δ_m . For such disks with m punctures, using closing rotation angles from Tables 2 and 3, we compute

$$\text{index}(\bar{\partial}_A) = \mu(A) + 2 = (m - 1) \times (-1) + 1 \times 0 + 2 = -(m - 3).$$

Adding capping operators at the punctures to A — see Table 2 for the positive puncture and Table 3 for the negative punctures — we get a boundary condition \hat{A} on the closed disk, which also has boundary conditions very close to constant. It follows that $\ker(\bar{\partial}_{\hat{A}})$ is 2–dimensional with kernel spanned by almost constant sections, and that $\text{coker}(\bar{\partial}_{\hat{A}})$ is 0–dimensional. By definition of the canonical orientation, see Section 6.2, the positive orientation of the determinant of $\bar{\partial}_{\hat{A}}$ is represented by a basis of its kernel, which converges to constant solutions that form a positively oriented basis of $T\Lambda$ (and a positive sign on its 0–dimensional cokernel).

Consider first the case $m = 2$. In this case the tree is simply a flow line, and the operator $\bar{\partial}_A$ has index 1 and a 1–dimensional kernel spanned by an almost constant solution that converges to $v^{\text{flow}}(\Gamma)$ as $\eta \rightarrow 0$; see Section 3.4.4 for notation. The exact gluing sequence that determines the orientation is then (see [16, Equation (3.17)]):

$$0 \rightarrow \ker(\bar{\partial}_{\hat{A}}) \rightarrow \ker(\bar{\partial}_{a_1-}) \oplus \ker(\bar{\partial}_A) \rightarrow 0$$

Here $\ker(\bar{\partial}_{a_1-})$ is spanned by an approximately constant section that converges to $v^{\text{ker}}(a_1)$, which is perpendicular to Γ ; see Table 3. It follows that the sign of the disk agrees with the orientation sign of the basis $(v^{\text{ker}}(a_1), v^{\text{flow}}(\Gamma))$ of $T\Lambda$, where $v^{\text{flow}}(\Gamma)$ is the vector field induced by the automorphism of the strip, times the sign of the determinant of the capping operator of a positive puncture at b , which is an isomorphism.

Gluing the positive and negative capping operators $\bar{\partial}_{b+}$ and $\bar{\partial}_{b-}$ at b gives an operator $\bar{\partial}_{b0}$ of index 0 with $\dim(\ker \bar{\partial}_{b0}) = \dim(\text{coker}(\bar{\partial}_{b0})) = 1$. Here the kernel is spanned by the constant solution $v^{\text{ker}}(b)$ and the cokernel by the constant solution $v^{\text{coker}}(b)$ of the dual problem. Note that the canonical orientation of $\det(\bar{\partial}_{b0})$ changes with o_C . Choose o_C so that $v^{\text{ker}}(b) \wedge v^{\text{coker}}(b)$ represents the positive orientation of $\det(\bar{\partial}_{b0})$ if $(v^{\text{coker}}(b), v^{\text{ker}}(b))$ is a positively oriented basis of $T\Lambda$. Then the sign of the disk can be expressed as

$$\text{sign}(\langle v^{\text{flow}}(\Gamma), v^{\text{ker}}(b) \rangle \cdot \langle v^{\text{ker}}(a), v^{\text{coker}}(b) \rangle) = \sigma_{\text{pos}}(\Gamma)\sigma(n, \Gamma)$$

as claimed.

Consider the case $m \geq 3$. Here the exact gluing sequence that gives the orientation is

$$(6-1) \quad 0 \rightarrow \ker(\bar{\partial}_{\hat{A}}) \xrightarrow{\alpha} \ker(\bar{\partial}_-) \xrightarrow{\beta} \text{coker}(\bar{\partial}_A) \rightarrow 0,$$

where we write:

$$\bar{\partial}_- = \bigoplus_{j=1}^{m-1} \bar{\partial}_{a_j-}$$

Since u_η is a rigid disk and $\text{index}(\bar{\partial}_A) = -(m-3)$, we get that $\dim(\text{coker}(\bar{\partial}_A)) = m-3$, and that $\text{coker}(\bar{\partial}_A)$ is spanned by linearized conformal variations which we represent as motion of the boundary minima in the domain; see Lemma 6.5. Here $\ker(\bar{\partial}_{\hat{A}})$ is endowed with the canonical orientation and $\ker(\bar{\partial}_-)$ with the orientation from capping operators. The sequence then induces an orientation on $\text{coker}(\bar{\partial}_A)$, which gives a sign when compared to the orientation induced from the space of conformal variations, as indicated in [16, Remark 3.18].

Applying Lemma 3.1 and Remark 3.3 of [16], the map α is defined as follows. An element in $\ker(\bar{\partial}_{a_j-})$ is a solution of the $\bar{\partial}$ -equation with boundary condition given by the negative capping operator at a_j . This solution is cut off and thereby defines an element in the space of sections over the closed disk. In this way, we identify $\ker(\bar{\partial}_-)$ with an $(m-1)$ -dimensional subspace of the domain of the operator $\bar{\partial}_{\hat{A}}$. The map α is then given by L^2 -projection of $\ker(\bar{\partial}_{\hat{A}})$ to this subspace. Likewise, we identify $\text{coker}(\bar{\partial}_A)$ with a subspace of the target space of $\bar{\partial}_{\hat{A}}$ by cutting off solutions of the dual problem and define the map β as $\bar{\partial}_{\hat{A}}$ followed by L^2 -projection.

The boundary condition of u_η is very close to constant \mathbb{R}^2 boundary conditions and the complex (Kähler) angles at the punctures are close to either 0 or π . Thus there is a deformation of $\bar{\partial}_A$ that takes the boundary conditions to constant \mathbb{R}^2 boundary conditions and that introduces a small negative exponential weight where the complex angle is close to 0 and a small positive exponential weight where it is close to π , which is sufficiently small so that the kernel and cokernel of $\bar{\partial}_A$ undergoes a continuous deformation (in particular dimensions of kernel and cokernel do not change). The capping operators can be deformed accordingly. We will use these deformed operators to determine the sign in the gluing sequence above. For simplicity, we will keep the notation $\bar{\partial}_A$ and $\bar{\partial}_{\hat{A}}$ for the deformed operators.

We next introduce notation for parts of the tree Γ as well as for corresponding parts of the domain Δ_m . See Figure 26.

We write E_{ij} for the edge connecting the i^{th} trivalent vertex t_i to the j^{th} , t_j , and R_{ij} for the (finite) strip region in Δ_m corresponding to E_{ij} , where we think of the boundary of R_{ij} as vertical line segments located to the right of the i^{th} boundary minimum τ_i and to the left of the j^{th} , τ_j . We write E_0 for the edge ending at the positive puncture and we take R_0 to be the half strip with a slit in Δ_m with boundary given by the two vertical segments bounding R_{ji} and $R_{ji'}$, where τ_j is the minimal




Figure 26: The graph Γ and domain Δ_m

boundary minimum. We write $E_l, l = 1, \dots, m - 1$ for edges that end at negative punctures and R_l for the corresponding half infinite strip which is a neighborhood of the l^{th} negative puncture in Δ_m .

Then

$$\Delta_m - \left(\bigcup_{E_i \subset \Gamma} R_i \cup \bigcup_{E_{ik} \subset \Gamma, i < k} R_{ik} \right)$$

is a disjoint union

$$\bigcup_{1 \leq i \leq m-2, i \neq j} V_i,$$

where V_i is a neighborhood of τ_i . Consider the vertical segment l through τ_i . The boundary ∂l of l lies in the boundary of Δ_m ; if the boundary segment containing the lower endpoint of l lies in the lower boundary on Δ_m we define $i_- = 0$ and if it lies on a boundary segment containing τ_k we define $i_- = k$. Likewise, if the upper endpoint lies in the upper boundary segment of Δ_m then we define $i_+ = m - 1$ and if it lies in a boundary segment containing $\tau_{k'}$ we define $i_+ = k'$. Note that $i_- < i < i_+$.

In order to deal with $\text{coker}(\bar{\partial}_A)$, we introduce below the space V_{con} of conformal variations of Δ_m . Write \tilde{v}_i^{con} for the conformal variation supported in V_i . Then $\bar{\partial}_A(\tilde{v}_i^{\text{con}}) = du_\eta(\bar{\partial}(\tilde{\partial}_\tau))$, where $\tilde{\partial}_\tau$ is a cut-off of the constant vector field ∂_τ supported in V_i . See Section 6.5. Thus, as $\eta \rightarrow 0$, $\bar{\partial}_A(\tilde{v}_i^{\text{con}})$ is supported in three rectangular regions $R_s^i, s = 0, 1, 2$, containing the vertical segments in the boundary ∂V_i and lying in the strip regions corresponding to the incoming edge E_0^i and the outgoing edges E_1^i and E_2^i , respectively, at t_i . Then in $R_s^i, s = 0, 1, 2$, $\bar{\partial}_A(\tilde{v}_i^{\text{con}})$ approaches $\bar{\partial}T_s^i$, where $\bar{\partial}T_s^i$ is a cut-off constant vector field tangent to $E_s^i, s = 0, 1, 2$, directed towards the positive puncture. For the purpose of calculating signs we thus replace \tilde{v}_i^{con} with v_i^{con} , where $\bar{\partial}_A v_i^{\text{con}} = \bar{\partial}(T_0^i + T_1^i + T_2^i)$ and think of V_{con} as the vector space spanned by the $m - 3$ conformal variations $v_i^{\text{con}}, i \neq j$.

Let \mathcal{H}_A and \mathcal{H}'_A (respectively $\mathcal{H}_{\hat{A}}$ and $\mathcal{H}'_{\hat{A}}$) denote the spaces of vector fields on the closed disk D (respectively the punctured disk Δ_m), which are the domain and target,

respectively, of $\bar{\partial}_A$ (respectively of $\bar{\partial}_{\hat{A}}$). Recall that

$$\bar{\partial}_A: \mathcal{H}_A \oplus V_{\text{con}} \rightarrow \mathcal{H}'_A$$

is an isomorphism because $\ker(\bar{\partial}_A)$ is trivial and $\text{coker}(\bar{\partial}_A)$ is mapped onto by $\bar{\partial}_A(V_{\text{con}})$. Viewing Δ_m as a subset of D , we define $\bar{\partial}_{\hat{A}}(v_i^{\text{con}}) = \bar{\partial}_A(v_i^{\text{con}})$ for $v_i^{\text{con}} \in V_{\text{con}}$ and write

$$\bar{\partial}_{\hat{A}, \text{con}}: \mathcal{H}_{\hat{A}} \oplus V_{\text{con}} \rightarrow \mathcal{H}'_{\hat{A}}$$

for the operator with extended domain. Then $\bar{\partial}_{\hat{A}, \text{con}}$ is an operator of index $2 + (m - 3)$ which has a $(m - 1)$ -dimensional kernel.

We will define a map $\psi: V_{\text{con}} \rightarrow \ker(\bar{\partial}_-)$ such that $\beta \circ \psi: V_{\text{con}} \rightarrow \text{coker}(\bar{\partial}_A)$ (see (6-1)) is an isomorphism that induces the same orientation on $\text{coker}(\bar{\partial}_A)$ as the isomorphism $\bar{\partial}_A: V_{\text{con}} \rightarrow \text{coker}(\bar{\partial}_A)$. It then follows that the sign of the disk u_η equals the sign of the determinant of the isomorphism

$$\ker(\bar{\partial}_{\hat{A}, \text{con}}) \approx \ker(\bar{\partial}_{\hat{A}}) \oplus V_{\text{con}} \xrightarrow{\alpha + \psi} \ker(\bar{\partial}_-)$$

between oriented vector spaces, where α is as in (6-1). To finish the proof we must thus first define ψ and then compute α and the resulting determinant.

We introduce certain vector fields on Δ_m that are supported in neighborhoods of the strip regions of the form R_l or R_{lk} in Δ_m associated to edges in Γ as explained above. We will call these vector fields *edge solutions*. More precisely, we take n_0^1 and n_0^2 to be constant sections supported in R_0 , cut off in a neighborhood of its boundary where n_0^1 is tangent to the second outgoing edge at t_j and n_0^2 tangent to the first; see Figure 27.

Along edges E_{ik} (respectively E_i) we define two cut-off constant vector fields: v_{ik} (respectively w_i) perpendicular to the edge and ζ_{ik} (respectively ζ_i) tangent to the edge. Here ζ_{ik} (respectively ζ_i) has support in a neighborhood of R_{ik} (respectively R_i), whereas v_{ik} (respectively w_i) has support in a neighborhood of $R_i \cup V_i$; see Figure 28. We call v_{ik} , ζ_{ik} and ζ_i *interior edge solutions* and w_i *exterior edge solutions*.

Furthermore, we assume that the elements $v_i^{\text{con}} \in V_{\text{con}}$ are chosen in such a way that the following holds: in any component C of the support of $\bar{\partial}_A v_i^{\text{con}}$ lying in an edge region R_{lk} or R_l the corresponding edge solution ζ_{lk} (or ζ_l) satisfies the matching condition $\bar{\partial}_A \zeta_{lk} = \bar{\partial}_A v_i^{\text{con}}$ (or $\bar{\partial}_A \zeta_l = \bar{\partial}_A v_i^{\text{con}}$).




Figure 27: Vector fields supported in R_0

We can now say that the $(m + 1)$ -dimensional kernel of $\bar{\partial}_{\hat{A}, \text{con}}$ is spanned by $m + 1$ linearly independent sections

$$\begin{aligned} \kappa_0^1 &= n_0^1 + (w_1 + \cdots + w_j) + \mathbf{E}_j^1, \\ \kappa_0^2 &= n_0^2 + (w_{j+1} + \cdots + w_m) + \mathbf{E}_j^2, \\ \kappa_i &= (-(w_{i-+1} + \cdots + w_i) + (w_{i+1} + \cdots + w_{i+})) \\ &\quad - \sigma_{n_0^1, \Gamma^1}(t_i)v_i^{\text{con}} + \mathbf{E}_i, \quad 1 \leq i \leq j - 1, \\ \kappa_i &= (-(w_{i-+1} + \cdots + w_i) + (w_{i+1} + \cdots + w_{i+})) \\ &\quad - \sigma_{n_0^2, \Gamma^2}(t_i)v_i^{\text{con}} + \mathbf{E}_i, \quad j + 1 \leq i \leq m - 1, \end{aligned}$$

where Γ^1 and Γ^2 are the partial flow trees obtained by cutting Γ at the first and second outgoing edge at t_j , respectively, and where $\sigma_{n, \Gamma}$ is as in (3-10). Here \mathbf{E}^i , $i \neq j$ (resp. \mathbf{E}_j^α , $\alpha = 1, 2$) are some linear combinations of interior edge solutions and conformal variations, respectively, that are supported in the component of $\Delta_m - s$, where s the vertical segment through τ_i (resp. τ_j), that contains punctures at $+\infty$; see Figure 29. (The matching conditions for edge solutions imply that linear combinations \mathbf{E}_i and \mathbf{E}_j^α , $\alpha = 1, 2$, exists so that the sections indeed lie in $\ker(\bar{\partial}_{\hat{A}, \text{con}})$.

Using these equations we define the map ψ as follows:

$$\psi(v_i) = \sigma_{n_0^s, \Gamma^s}(t_i)(-(w_{i-+1} + \cdots + w_i) + (w_{i+1} + \cdots + w_{i+}) + \mathbf{E}_i),$$




Figure 28: Supports of edge solutions





Figure 29: Support of additional interior edge solutions

where $s = 1$ if $i < j$ and $s = 2$ if $i > j$. Since $\bar{\partial}\kappa_i = 0$ we have, by construction of κ_i , $\bar{\partial}\psi(v_i) = \bar{\partial}v_i$. By definition, the orientation on $\text{coker}(\bar{\partial}_A)$ induced by conformal variations is given by L^2 -projection of $\bar{\partial}_A(V_{\text{con}})$ (with orientation on V_{con} as in Section 6.5). Thus $\beta \circ \psi$ induces the correct orientation on $\text{coker}(\bar{\partial}_A)$.

In order to compute the sign we first note that the projection of the subspace spanned by $\kappa_1, \dots, \kappa_{j-1}, \kappa_{j+1}, \dots, \kappa_{m-3}$ to V_{con} is an isomorphism and that the map from its complement spanned by κ_0^1, κ_0^2 given by evaluation at the positive puncture gives an isomorphism to $\ker(\bar{\partial}_{\hat{A}})$, which consists of constant solutions. It follows that the sign of the rigid disk u_η is given by

$$\epsilon(u_\eta) = s_1 s_2 s_3,$$

where $s_k, k = 1, 2, 3$ are as follows. First, s_1 equals the sign of the orientation given by

$$\xi_0^1 \wedge \xi_0^2 \wedge \xi_1 \wedge \dots \wedge \xi_{j-1} \wedge \xi_{j+1} \wedge \xi_{m-2}$$

on $\ker(\bar{\partial}_-)$ where

$$\begin{aligned} \xi_0^1 &= w_1 + \dots + w_j, \\ \xi_0^2 &= w_{j+1} + \dots + w_{m-1}, \\ \xi_1 &= -w_1 + (w_2 + \dots + w_{1+}), \\ \xi_2 &= -(w_{2-+1} + w_2) + (w_3 + \dots + w_{2+}), \\ &\vdots \\ \xi_{j-1} &= -(w_{(j-1)-+1} + \dots + w_{j-1}) + w_j, \\ \xi_{j+1} &= -(w_{j+1}) + (w_{j+2} + \dots + w_{(j+1)+}), \\ &\vdots \\ \xi_{m-1} &= -(w_{(m-2)-+1} + \dots + w_{m-2}) + w_{m-1}. \end{aligned}$$

Second, s_2 equals the sign of the orientation on V_{con} given by

$$\begin{aligned} (\sigma_{n_0^1, \Gamma^1}(t_1)v_1^{\text{con}}) \wedge \dots \wedge (\sigma_{n_0^1, \Gamma^1}(t_{j-1})v_{j-1}^{\text{con}}) \\ \wedge (\sigma_{n_0^2, \Gamma^2}(t_{j+1})v_{j+1}^{\text{con}}) \wedge \dots \wedge (\sigma_{n_0^2, \Gamma^2}(t_{m-2})v_{m-2}^{\text{con}}). \end{aligned}$$

Third, s_3 equals the sign of the orientation on $\ker(\partial_{\hat{A}})$ given by

$$\bar{n}_0^1 \wedge \bar{n}_0^2,$$

where \bar{n}_0^s is a constant solution agreeing with n_0^s in the region where n_0^s is constant, $s = 1, 2$.

We have

$$\begin{aligned} \xi_0^1 \wedge \xi_0^2 \wedge \xi_1 \wedge \cdots \wedge \xi_{j-1} \wedge \xi_{j+1} \wedge \xi_{m-2} &= (-1)^{j-1} \xi_0^1 \wedge \xi_1 \wedge \cdots \wedge \xi_{j-1} \wedge \xi_0^2 \wedge \xi_{j+1} \wedge \xi_{m-2} \\ &= (-1)^{j-1} w_1 \wedge \cdots \wedge w_{m-1} \\ &= (-1)^{j-1} \prod_{s=1}^{m-1} \langle w_s, v^{\ker}(a_s) \rangle v^{\ker}(a_1) \wedge \cdots \wedge v^{\ker}(a_{m-1}), \end{aligned}$$

and thus

$$s_1 = (-1)^{j-1} \prod_{s=1}^{m-1} \langle w_s, v^{\ker}(a_s) \rangle.$$

Next,

$$v_1 \wedge \cdots \wedge v_{j-1} \wedge v_{j+1} \wedge \cdots \wedge v_{m-2} = (-1)^{j-1} \partial_{q_3} \wedge \cdots \wedge \partial_{q_m},$$

by Lemma 6.5, where $\partial_{q_3} \wedge \cdots \wedge \partial_{q_m}$ is the standard orientation of \mathcal{C}_m corresponding to moving the last $m - 3$ punctures counter clockwise along the boundary of the disk, and thus

$$s_2 = (-1)^{j-1} \prod_{s=1}^{j-1} \sigma_{n_0^1, \Gamma^1}(t_s) \prod_{s=j+1}^{m-1} \sigma_{n_0^2, \Gamma^2}(t_s).$$

Finally, recall that $o_{\mathbb{C}}$ was chosen so that $v^{\text{coker}}(b) \wedge v^{\ker}(b)$ represents the positive orientation on $T\Lambda$. Since $v^{\ker}(b)$ is parallel to the vector $v^{\text{con}}(t_j)$ related to the conformal variation at τ_j , we find that

$$\begin{aligned} \bar{n}_0^1 \wedge \bar{n}_0^2 &= (n_0^1 + n_0^2) \wedge (-n_0^1 + n_0^2) \\ &= \langle (n_0^1 + n_0^2), v^{\text{coker}}(b) \rangle \langle (-n_0^1 + n_0^2), v^{\ker}(b) \rangle v^{\text{coker}}(b) \wedge v^{\ker}(b) \\ &= \langle (n_0^1 + n_0^2), v^{\text{coker}}(b) \rangle \sigma(t_j) \langle v_{\Gamma}^{\text{flow}}, v^{\ker}(b) \rangle v^{\text{coker}}(b) \wedge v^{\ker}(b). \end{aligned}$$

Thus, if n_0^1 and n_0^2 are the vector splittings of $v^{\text{coker}}(b)$ then $s_3 = \sigma(t_j)$, and

$$s_1 s_2 s_3 = \sigma_{\text{pos}}(\Gamma) \sigma(n, \Gamma). \quad \square$$

6.6.2 The sign of a quantum flow tree of type (QT'_0) Let Ξ be a quantum flow tree whose big disk part is a rigid strip Θ . Assume that Ξ has m punctures. Let e denote its positive puncture (a Reeb chord of type \mathbf{L}_2), let c denote its negative puncture (of type \mathbf{L}_1) and denote the remaining punctures by a_1, \dots, a_{m-2} (all of type \mathbf{S}_0). It then follows that Θ is a strip with positive puncture at a Reeb chord of Λ close to e and negative puncture at a Reeb chord of Λ close to c . Let t_1, \dots, t_{m-1} denote the junction points (ie, the points on the boundary of Θ where trees are attached)

and the trivalent vertices in the trees attached. Then each t_j corresponds to a unique boundary minimum τ_j in the domain Δ_m of a holomorphic disk u_η corresponding to η and we number the points t_j according to the vertical coordinate of the boundary minima τ_j in Δ_m . We write \mathcal{I} for the set of junction points of Ξ and for $t_j \in \mathcal{I}$ we write Γ_j for the tree attached at t_j and n_j for the vector at t_j that is tangent to the boundary of Θ and points toward the positive puncture e .

Lemma 6.8 *The sign of the rigid disk u_η corresponding to Ξ is given by*

$$\epsilon(u_\eta) = \epsilon(\Xi) = \epsilon(\Theta) \prod_{t_j \in \mathcal{I}} \sigma(n_j, \Gamma_j);$$

see Sections 3.4.4 and 4.3 for notation.

Proof Consider first the case $m = 2$, ie, when there are no flow trees attached to the disk. In this case the linearized operator $\bar{\partial}_A$ as well as the capping operators are small deformations of the linearized operator and capping operators of the corresponding disk with boundary on Λ . It follows that the signs of the two disks agree and hence $\epsilon(u_\eta) = \epsilon(\Theta)$ as claimed.

In order to prepare for the case $m > 2$ we write down the gluing sequence for $m = 2$, see [16, Equation (3.17)], that gives the sign explicitly. Using Tables 3, 4 (as a negative puncture) and 5 we have

$$0 \rightarrow \ker(\bar{\partial}_{c-}) \oplus \ker(\bar{\partial}_A) \xrightarrow{\beta} \text{coker}(\bar{\partial}_{e+}) \oplus \text{coker}(\bar{\partial}_{c-}) \rightarrow 0,$$

where we use the fact that the glued operator $\bar{\partial}_{\hat{A}}$ is an isomorphism. (To see this note that $\bar{\partial}_{\hat{A}}$ splits into a direct sum of two 1-dimensional operators both of index 0.)

Noting that the boundary conditions of $\bar{\partial}_A$ are close to \mathbb{R}^2 conditions we deform them to constant \mathbb{R}^2 boundary conditions inserting weights as determined by the complex angles, exactly as in the proof of Lemma 6.7. The solutions in the above sequence can then be thought of as cut-off constant solutions (which may be extended on a sufficiently large domain so that the supports of $v^{\ker}(c-)$ and $v^{\text{coker}}(e+)$ overlap in order for the sequence to be exact) and the sign is given by

$$\epsilon(u_\eta) = \epsilon(\Gamma) = \langle v^{\ker}(c), v^{\text{coker}}(e) \rangle \langle v^{\text{flow}}(\Gamma), v^{\text{coker}}(c) \rangle.$$

Consider next the case $m > 2$. As above we use the associated weighted problem corresponding to A with constant boundary conditions and exponential weights. Again $\bar{\partial}_{\hat{A}}$ is an isomorphism and the gluing sequence which determines the sign is

$$0 \rightarrow \ker(\bar{\partial}_{-}) \oplus \ker(\bar{\partial}_{c-}) \rightarrow \text{coker}(\bar{\partial}_{e+}) \oplus \text{coker}(\bar{\partial}_{c-}) \oplus \text{coker}(\bar{\partial}_A) \rightarrow 0,$$

where $\ker(\bar{\partial}_-) = \bigoplus_{j=1}^{m-2} \ker(\bar{\partial}_{a_j-})$ is spanned by cut off (constant) solutions of the capping operators at the negative punctures a_1, \dots, a_{m-2} , where $\ker(\bar{\partial}_{c-})$ and $\text{coker}(\bar{\partial}_{e+})$ are as above, where the orientation of $\ker(\bar{\partial}_-) \oplus \ker(\bar{\partial}_{c-})$ is induced by the order of the punctures, and where $\text{coker}(\bar{\partial}_A)$ is equipped with the orientation induced by the space of conformal variations of Δ_m .

As in the proof of Lemma 6.7 we stabilize the operator $\bar{\partial}_{\hat{A}}$ by adding the finite-dimensional space V_{con} spanned by conformal variations supported near all boundary minima except τ_r , where we take τ_r to be the boundary minimum corresponding to the junction point t_r immediately following the negative puncture c , or if there is no junction point after c , the junction point t_r immediately preceding c . Here the conformal variation v_j^{con} near a boundary minimum τ_j such that $t_j \notin \mathcal{I}$ is defined exactly as the elements of V_{con} in the proof of Lemma 6.7. Conformal variations supported near boundary minima τ_i with $t_i \in \mathcal{I}$ have the form $\bar{\partial}_A(\tilde{v}_i^{\text{con}}) = du_\eta(\bar{\partial}(\tilde{\partial}_\tau))$, where $\tilde{\partial}_\tau$ is a cut-off of the vector field $\partial\tau$ in the domain of the holomorphic disk part of Γ that is tangent to the boundary and directed towards the positive puncture and continued constantly into the domain corresponding to the tree attached; see Figure 30.




Figure 30: Support of conformal variation at junction point

As $\eta \rightarrow 0$, $\bar{\partial}_A \tilde{v}_j^{\text{con}}$ converges to a $\bar{\partial} \tilde{v}^{\text{flow}}$ in the part of the domain corresponding to the holomorphic disk part of Γ , where \tilde{v}^{flow} is a cut-off of the constant vector field v^{flow} pointing toward the positive puncture, and to \tilde{n}_i in the part of the domain corresponding to the tree (near the junction point), where \tilde{n}_i is a cut-off of the constant vector field n_i . In analogy with v_j^{con} for $t_j \notin \mathcal{I}$, we define v_i^{con} for the deformed boundary conditions so that $\bar{\partial}_A v_i^{\text{con}}$ agrees with the operator acting on these cut off constant vector fields in the two components of its support. This gives the stabilized operator

$$\bar{\partial}_{\hat{A}, \text{con}} : \mathcal{H}_{\hat{A}} \oplus V_{\text{con}} \rightarrow \mathcal{H}'_{\hat{A}}$$

and the new sequence

$$0 \rightarrow \ker(\bar{\partial}_{\hat{A}, \text{con}}) \rightarrow \ker(\bar{\partial}_-) \oplus \ker(\bar{\partial}_{c-}) \rightarrow \text{coker}(\bar{\partial}_{e+}) \oplus \text{coker}(\bar{\partial}_{c-}) \rightarrow 0,$$

which by an argument similar to that used for the stabilized sequence in Lemma 6.7 gives the correct sign.

In addition to conformal variations we also introduce *internal* and *external edge solutions* corresponding to edges of the flow trees Γ_j , $t_j \in \mathcal{I}$ attached exactly as in Lemma 6.7, where we start the construction for the tree Γ_j with the vector n_j at $t_j \in \mathcal{I}$. We also introduce an extra cut-off solution, which compensates for the lack of conformal variation at τ_r : let $v^{\text{sol}} = v^{\text{flow}} + v_r$ where v^{flow} is the constant vector field along Θ cut off near each junction point and near c , and where v_r is a conformal variation of the usual type supported near τ_r ; see Figure 31.




Figure 31: The support of v^{sol} . Here Γ_3 is the first partial tree to follow the puncture c in the positive direction of the boundary of the domain.

As in the proof of Lemma 6.7 we assume that cut-off tangential edge solutions and v^{sol} agree with conformal variations on intersections of supports of their derivatives.

Define V_{sol} as the 1-dimensional space spanned by v^{sol} and stabilize the operator one more time:

$$\bar{\partial}_{\hat{A}, \text{con}, \text{sol}}: \mathcal{H}_{\hat{A}} \oplus V_{\text{con}} \oplus V_{\text{sol}} \rightarrow \mathcal{H}'_{\hat{A}}$$

Noting that the projection of $\bar{\partial}v^{\text{sol}}$ to $\text{coker}(\bar{\partial}_{c-})$ equals the corresponding projection of $\bar{\partial}v^{\text{flow}}$ to $\text{coker}(\bar{\partial}_{c-})$ for the disk Θ we can determine the sign in the exact sequence for the stabilized operator $\bar{\partial}_{\hat{A}, \text{con}, \text{sol}}$, which reads

$$0 \rightarrow \ker(\bar{\partial}_{\hat{A}, \text{con}, \text{sol}}) \rightarrow \ker(\bar{\partial}_{-}) \oplus \ker(\bar{\partial}_{c-}) \rightarrow \text{coker}(\bar{\partial}_{e+}) \rightarrow 0,$$

in combination with the sign of $\langle v^{\text{sol}}, v^{\text{coker}(c-)} \rangle$.

In analogy with the sign calculation in Lemma 6.7, we find a basis in the $(m - 2)$ -dimensional kernel of $\bar{\partial}_{\widehat{A}, \text{con, sol}}$ given by $\kappa_1, \dots, \kappa_{m-2}$, where

$$\begin{aligned} \kappa_i &= -(w_{i-+1} + \dots + w_i) + (w_{i+1} + \dots + w_{i+}) - \sigma_{n_{s(i)}, \Gamma_{s(i)}}(t_i)v_i^{\text{con}} + \mathbf{E}_i \quad \text{if } t_i \notin \mathcal{I}, \\ \kappa_j &= (w_{j-+1} + \dots + w_j) - v_j^{\text{con}} + \mathbf{E}_j \quad \text{if } t_j \in \mathcal{I} \text{ and } j \neq r, \\ \kappa_r &= (w_{r-+1} + \dots + w_r) - v^{\text{sol}} - \sum_{j \neq r} v_j^{\text{con}} + \mathbf{E}_j, \end{aligned}$$

where $\Gamma_{s(i)}$ is the attached flow tree in which $t_i \notin \mathcal{I}$ lies, and where \mathbf{E}_j denotes linear combinations of interior edge solutions of edges below t_j , exactly as in the proof of Lemma 6.7.

It follows that $\epsilon(u_\eta) = s_1 s_2 s_3$, where the signs s_j , $j = 1, 2, 3$ are as follows. First, s_1 equals the orientation on $\ker(\bar{\partial}_-)$ given by

$$(-1)^{r-1} \left(\prod_{t_i \notin \mathcal{I}} \sigma_{n_{s(i)}, \Gamma_{s(i)}}(t_i) \right) \xi_1 \wedge \xi_{r-1} \wedge \xi_{r+1} \wedge \dots \wedge \xi_{m-1},$$

where

$$\begin{aligned} \xi_i &= -(w_{i-+1} + \dots + w_j) + (w_{i+1} + \dots + w_{i+}) \quad \text{if } t_i \notin \mathcal{I}, \\ \xi_j &= (w_{j-+1} + \dots + w_j) \quad \text{if } t_j \in \mathcal{I}. \end{aligned}$$

Second, up to a positive factor,

$$s_2 = \langle v^{\text{sol}}, v^{\text{coker}}(c-) \rangle = \langle v_r + v^{\text{flow}}, v^{\text{coker}}(c-) \rangle = \langle v^{\text{flow}}, v^{\text{coker}}(c-) \rangle.$$

Third, $s_3 = (-1)^{r-1} \langle v^{\text{ker}}(c), v^{\text{coker}}(e) \rangle$, by the orientation conventions for sequences; see [16, Section 3.2.1]. We conclude that, with notation as in Theorem 4.6,

$$\epsilon(u_\eta) = \epsilon(\Xi) = \epsilon(\Theta) \prod_{j=1}^l \sigma(n_j, \Gamma_j). \quad \square$$

6.6.3 The sign of a quantum flow tree of type (QT_0) Let Ξ be a quantum flow tree with holomorphic component Θ , a rigid disk with one puncture. Assume that Ξ has m punctures. Let c denote its positive puncture (a Reeb chord of type \mathbf{L}_1) and a_1, \dots, a_{m-1} denote its negative punctures (all Reeb chords of type \mathbf{S}_0). We use notation as in Section 6.6.2 for junction points and trivalent vertices of Ξ .

We mark two points q'_0 and q_0 on the boundary of the rigid disk Θ right after the positive puncture using two parallel oriented hypersurfaces near the positive puncture. Let \tilde{v}^{flow} denote the holomorphic vector field along the disk that vanishes at the positive puncture and at the second marked point, and that is directed toward the positive puncture along the boundary. At each junction point $t_j \in \mathcal{I}$, let n_j denote the value of \tilde{v}^{flow} at t_j .

Lemma 6.9 *The sign of the rigid disk u_η corresponding to Ξ is given by:*

$$\epsilon(u_\eta) = \epsilon(\Xi) = \epsilon(\Theta) \prod_{j=1}^l \sigma(n_j, \Gamma_j)$$

See Sections 3.4.4 and 4.3 for notation.

Proof Let $\bar{\partial}_A$ denote the linearized operator corresponding to the rigid disk Θ with domain thought of as a strip with punctures at the positive puncture and at q_0 . Then the linearized operator

$$\bar{\partial}_A: \mathcal{H}_A \rightarrow \mathcal{H}'_A$$

with $m - 1 = 0$ negative punctures has index 1. The 1–dimensional kernel for an appropriate choice of orientation of the hypersurface makes \tilde{v}^{flow} give the positive orientation to the moduli space. The proof then follows from an argument similar to the proof of Lemma 6.8 so we just sketch the details.

Assume that u_η has m punctures. Adding capping operators we get an operator $\bar{\partial}_{\hat{A}}$ on the closed disk of index 2 with two–dimensional kernel corresponding to linearized conformal automorphism. Adding the vanishing condition at the marked point q_0 for the vector fields in $\mathcal{H}_{\hat{A}}$ we get a new operator $\bar{\partial}_{\hat{A}^*}$ with domain a codimension-one subspace $\mathcal{H}_{\hat{A}^*} \subset \mathcal{H}_{\hat{A}}$. The restriction of the operator gives an operator

$$\bar{\partial}_{\hat{A}^*}: \mathcal{H}_{\hat{A}^*} \rightarrow \mathcal{H}'_{\hat{A}}$$

of index 1 with 1–dimensional kernel. We then consider the stabilized problem

$$\bar{\partial}_{\hat{A}^*, \text{con}}: \mathcal{H}_{\hat{A}^*} \oplus V_{\text{con}} \rightarrow \mathcal{H}'_{\hat{A}}$$

which has index $(m - 1)$ and find that the sign of u_η is given by the sign of the map

$$0 \rightarrow \ker(\bar{\partial}_{\hat{A}^*, \text{con}}) \rightarrow \ker(\partial_-) \rightarrow 0,$$

where $\ker(\partial_-) = \bigoplus_{j=1}^{m-1} \ker(\bar{\partial}_{a_j-})$ is the sum of cut off kernel functions of the capping operators at all the negative punctures of u_η . The equation

$$\epsilon(u_\eta) = \epsilon(\Theta) \prod_{j=1}^l \sigma(n_j, \Gamma_j)$$

then follows from a slightly simpler version of the analogous calculation in the proof of Lemma 6.9. □

6.6.4 The sign of quantum flow tree of type (QT₁) Consider a quantum flow tree Ξ with m punctures and with big disk part a once punctured constrained rigid disk Θ that is constrained to pass through $\Pi(b)$ where b is a Reeb chord of type S_1 . Let the positive puncture of Ξ be e (a Reeb chord of type L_2), let the negative punctures be b (a Reeb chord of type S_1), and a_1, \dots, a_{m-2} (Reeb chords of type S_0). We use notation for trivalent vertices and junction points of Ξ as in Section 6.6.2. Let \tilde{v}^{flow} denote the holomorphic vector field along Θ that vanishes at e and b and that is directed toward the positive puncture e along the boundary. Let n_j be the value of \tilde{v}^{flow} at $t_j \in \mathcal{I}$.

Lemma 6.10 *The sign of the rigid disk u_η corresponding to Ξ is given by*

$$\epsilon(u_\eta) = \epsilon(\Xi) = \epsilon(\Theta) \prod_{j=1}^l \sigma(n_j, \Gamma_j).$$

See Sections 3.4.4 and 4.3 for notation.

Proof The proof follows along the lines of previous lemmas in the section and detailed calculations will be omitted. We use notation as before. Start with the case of no negative punctures. Consider first the orientation sequence for the moduli space containing the 1-dimensional disk. We write $\bar{\partial}_E$ for the corresponding operator. Then $\bar{\partial}_E$ has 3-dimensional kernel whereas the capped off operator $\bar{\partial}_{\hat{E}}$ has 2-dimensional kernel corresponding to linearized automorphism of the disk with one puncture and we get the gluing sequence

$$0 \rightarrow \ker(\bar{\partial}_{\hat{E}}) \rightarrow \ker(\bar{\partial}_E) \rightarrow \text{coker}(\bar{\partial}_{e+}) \rightarrow 0,$$

thus $\text{coker}(\bar{\partial}_{e+})$ together with conformal automorphisms orient the moduli space and we can identify $\text{coker}(\bar{\partial}_{e+})$ with the vector field v that is the positively oriented tangent vector of the 1-dimensional moduli space; see Section 4.3.

Consider first the case $m = 2$. The operator $\bar{\partial}_{\hat{A}}$ is an isomorphism and the gluing sequence is

$$0 \rightarrow \ker(\bar{\partial}_{b-}) \oplus \ker(\bar{\partial}_A) \rightarrow \text{coker}(\bar{\partial}_{e+}) \oplus \text{coker}(\bar{\partial}_{b-}) \rightarrow 0$$

and we find that

$$\epsilon(u_\eta) = \langle v, v^{\ker(b)} \rangle \langle v^{\text{flow}}(\Gamma), v^{\text{coker}(b)} \rangle = \epsilon(\Theta)$$

as claimed. In the case $m > 2$ the gluing sequence is

$$0 \rightarrow \ker(\bar{\partial}_-) \oplus \ker(\bar{\partial}_{b-}) \rightarrow \text{coker}(\bar{\partial}_{b-}) \oplus \text{coker}(\bar{\partial}_{e+}) \oplus \text{coker}(\bar{\partial}_A) \rightarrow 0.$$

As in Lemma 6.8 we stabilize to an operator

$$\bar{\partial}_{\hat{A}, \text{con}, \text{sol}}: \mathcal{H}_{\hat{A}} \oplus V_{\text{con}} \oplus V_{\text{sol}} \rightarrow \mathcal{H}'_{\hat{A}}$$

of index $m - 2$, where V_{con} is spanned by conformal variations and where V_{sol} is a 1-dimensional space spanned by v^{sol} , which is a sum of an extra conformal variation at the fixed boundary minimum and a cut-off \tilde{v}^{flow} of v^{flow} , and which maps non-trivially to $v^{\text{coker}}(b)$. The resulting map which determines the sign is then

$$0 \rightarrow \ker(\bar{\partial}_{\hat{A}, \text{con}, \text{sol}}) \rightarrow \ker(\bar{\partial}_{-}) \oplus \ker(\bar{\partial}_{b-}) \rightarrow \text{coker}(\bar{\partial}_{e+}) \rightarrow 0,$$

and the lemma follows from computation similar to those in the proof of Lemma 6.7. \square

References

- [1] **M Abouzaid**, *Maslov 0 nearby Lagrangians are homotopy equivalent* arXiv: 1005.0358
- [2] **M Abouzaid**, *Framed bordism and Lagrangian embeddings of exotic spheres*, Ann. of Math. 175 (2012) 71–185 MR2874640
- [3] **V I Arnold, A B Givental**, *Symplectic geometry*, from: “Dynamical systems, IV”, (V I Arnold, S P Novikov, editors), Encyclopaedia Math. Sci. 4, Springer, Berlin (2001) 1–138 MR1866631
- [4] **V I Arnold, S M Gusein-Zade, A N Varchenko**, *Singularities of differentiable maps, I: classification of critical points, caustics and wave fronts*, Modern Birkhäuser Classics, Birkhäuser/Springer, New York (2012) MR2896292
- [5] **F Bourgeois, T Ekholm, Y Eliashberg**, *Symplectic homology product via Legendrian surgery*, Proc. Natl. Acad. Sci. USA 108 (2011) 8114–8121 MR2806647
- [6] **F Bourgeois, T Ekholm, Y Eliashberg**, *Effect of Legendrian surgery*, Geom. Topol. 16 (2012) 301–389 MR2916289
- [7] **Y Chekanov**, *Differential algebra of Legendrian links*, Invent. Math. 150 (2002) 441–483 MR1946550
- [8] **G Dimitroglou Rizell**, *Knotted Legendrian surfaces with few Reeb chords*, Algebr. Geom. Topol. 11 (2011) 2903–2936 MR2846915
- [9] **N M Dunfield, S Garoufalidis**, *Non-triviality of the A -polynomial for knots in S^3* , Algebr. Geom. Topol. 4 (2004) 1145–1153 MR2113900
- [10] **T Ekholm**, *Morse flow trees and Legendrian contact homology in 1-jet spaces*, Geom. Topol. 11 (2007) 1083–1224 MR2326943

- [11] **T Ekholm, J B Etnyre**, *Invariants of knots, embeddings and immersions via contact geometry*, from: “Geometry and topology of manifolds”, (H U Boden, I Hambleton, A J Nicas, B D Park, editors), Fields Inst. Commun. 47, Amer. Math. Soc. (2005) 77–96 MR2189927
- [12] **T Ekholm, J Etnyre, L Ng, M Sullivan**, *Filtrations on the knot contact homology of transverse knots*, Math. Ann. 355 (2013) 1561–1591 MR3037024
- [13] **T Ekholm, J B Etnyre, J M Sabloff**, *A duality exact sequence for Legendrian contact homology*, Duke Math. J. 150 (2009) 1–75 MR2560107
- [14] **T Ekholm, J Etnyre, M Sullivan**, *The contact homology of Legendrian submanifolds in \mathbb{R}^{2n+1}* , J. Differential Geom. 71 (2005) 177–305 MR2197142
- [15] **T Ekholm, J Etnyre, M Sullivan**, *Non-isotopic Legendrian submanifolds in \mathbb{R}^{2n+1}* , J. Differential Geom. 71 (2005) 85–128 MR2191769
- [16] **T Ekholm, J Etnyre, M Sullivan**, *Orientations in Legendrian contact homology and exact Lagrangian immersions*, Internat. J. Math. 16 (2005) 453–532 MR2141318
- [17] **T Ekholm, J Etnyre, M Sullivan**, *Legendrian contact homology in $P \times \mathbb{R}$* , Trans. Amer. Math. Soc. 359 (2007) 3301–3335 MR2299457
- [18] **T Ekholm, T Kálmán**, *Isotopies of Legendrian 1-knots and Legendrian 2-tori*, J. Symplectic Geom. 6 (2008) 407–460 MR2471099
- [19] **Y Eliashberg**, *Invariants in contact topology*, from: “Proceedings of the International Congress of Mathematicians, II” (1998) 327–338 MR1648083
- [20] **K Fukaya, Y-G Oh, H Ohta, K Ono**, *Lagrangian intersection Floer theory: anomaly and obstruction, II*, AMS/IP Studies in Advanced Mathematics 46, Amer. Math. Soc. (2009) MR2548482
- [21] **K Fukaya, P Seidel, I Smith**, *Exact Lagrangian submanifolds in simply-connected cotangent bundles*, Invent. Math. 172 (2008) 1–27 MR2385665
- [22] **P B Kronheimer, T S Mrowka**, *Dehn surgery, the fundamental group and $SU(2)$* , Math. Res. Lett. 11 (2004) 741–754 MR2106239
- [23] **K Mishachev**, *The N -copy of a topologically trivial Legendrian knot*, J. Symplectic Geom. 1 (2003) 659–682 MR2039159
- [24] **D Nadler**, *Microlocal branes are constructible sheaves*, Selecta Math. 15 (2009) 563–619 MR2565051
- [25] **L L Ng**, *Computable Legendrian invariants*, Topology 42 (2003) 55–82 MR1928645
- [26] **L Ng**, *Knot and braid invariants from contact homology, I*, Geom. Topol. 9 (2005) 247–297 MR2116316
- [27] **L Ng**, *Knot and braid invariants from contact homology, II*, Geom. Topol. 9 (2005) 1603–1637 MR2175153

- [28] **L Ng**, *Conormal bundles, contact homology and knot invariants*, from: “The interaction of finite-type and Gromov–Witten invariants”, (D Auckly, J Bryan, editors), *Geom. Topol. Monogr.* 8 (2006) 129–144 MR2402823
- [29] **L Ng**, *Framed knot contact homology*, *Duke Math. J.* 141 (2008) 365–406 MR2376818
- [30] **L Ng**, *Combinatorial knot contact homology and transverse knots*, *Adv. Math.* 227 (2011) 2189–2219 MR2807087
- [31] **J W Robbin, D A Salamon**, *Asymptotic behaviour of holomorphic strips*, *Ann. Inst. H. Poincaré Anal. Non Linéaire* 18 (2001) 573–612 MR1849689
- [32] **J-C Sikorav**, *Some properties of holomorphic curves in almost complex manifolds*, from: “Holomorphic curves in symplectic geometry”, (M Audin, J Lafontaine, editors), *Progr. Math.* 117, Birkhäuser, Basel (1994) 165–189 MR1274929

TE: *Department of Mathematics, Uppsala University*
Box 480, 751 06 Uppsala, Sweden

JE: *School of Mathematics, Georgia Institute of Technology*
686 Cherry Street, Atlanta GA 30332-0160, USA

LN: *Department of Mathematics, Duke University*
Box 90320, Durham NC 27708-0320, USA

MS: *Department of Mathematics and Statistics, University of Massachusetts*
Amherst MA 01003-9305, USA

tobias@math.uu.se, etnyre@math.gatech.edu, ng@math.duke.edu,
sullivan@math.umass.edu

<http://www.math.gatech.edu/~etnyre>, <http://www.math.duke.edu/~ng/>,
<http://www.math.umass.edu/~sullivan/>

Proposed: Leonid Polterovich

Received: 16 January 2012

Seconded: Yasha Eliashberg, Ralph Cohen

Accepted: 5 January 2013