Geometry & Topology

Volume 27 (2023)

Duality between Lagrangian and Legendrian invariants

Tobias Ekholm
Yanki Lekili

msp

msp

Geometry & Topology **27**:6 (2023) 2049–2179
DOI: 10.2140/gt.2023.27.2049
Published: 25 August 2023

Duality between Lagrangian and Legendrian invariants

TOBIAS EKHOLM
YANKI LEKILI

Consider a pair (X, L) of a Weinstein manifold X with an exact Lagrangian submanifold L, with ideal contact boundary (Y, Λ), where Y is a contact manifold and $\Lambda \subset Y$ is a Legendrian submanifold. We introduce the Chekanov–Eliashberg DG–algebra, $\mathrm{CE}^*(\Lambda)$, with coefficients in chains of the based loop space of Λ, and study its relation to the Floer cohomology $\mathrm{CF}^*(L)$ of L. Using the augmentation induced by L, $\mathrm{CE}^*(\Lambda)$ can be expressed as the Adams cobar construction Ω applied to a Legendrian coalgebra, $\mathrm{LC}_*(\Lambda)$. We define a twisting cochain $\mathfrak{t}\colon \mathrm{LC}_*(\Lambda) \to \mathrm{B}(\mathrm{CF}^*(L))^{\#}$ via holomorphic curve counts, where B denotes the bar construction and # the graded linear dual. We show under simple-connectedness assumptions that the corresponding Koszul complex is acyclic, which then implies that $\mathrm{CE}^*(\Lambda)$ and $\mathrm{CF}^*(L)$ are Koszul dual. In particular, $\mathfrak{t}$ induces a quasi-isomorphism between $\mathrm{CE}^*(\Lambda)$ and $\Omega\mathrm{CF}_*(L)$, the cobar of the Floer homology of L.

This generalizes the classical Koszul duality result between $C^*(L)$ and $C_{-*}(\Omega L)$ for L a simply connected manifold, where ΩL is the based loop space of L, and provides the geometric ingredient explaining the computations given by Etgü and Lekili (2017) in the case when X is a plumbing of cotangent bundles of 2–spheres (where an additional weight grading ensured Koszulity of $\mathfrak{t}$).

We use the duality result to show that under certain connectivity and local-finiteness assumptions, $\mathrm{CE}^*(\Lambda)$ is quasi-isomorphic to $C_{-*}(\Omega L)$ for any Lagrangian filling L of Λ.

Our constructions have interpretations in terms of wrapped Floer cohomology after versions of Lagrangian handle attachments. In particular, we outline a proof that $\mathrm{CE}^*(\Lambda)$ is quasi-isomorphic to the wrapped Floer cohomology of a fiber disk C in the Weinstein domain obtained by attaching $T^*(\Lambda \times [0,\infty))$ to X along Λ (or, in the terminology of Sylvan (2019), the wrapped Floer cohomology of C in X with wrapping stopped by Λ). Along the way, we give a definition of wrapped Floer cohomology via holomorphic buildings that avoids the use of Hamiltonian perturbations, which might be of independent interest.

57R17

1 Introduction

In this introduction we first give an overview of our results. The overview starts with a review of well-known counterparts of our constructions in algebraic topology. We then introduce our Legendrian and Lagrangian invariants in Sections 1.1 and 1.2, respectively, and discuss the connection between them and applications thereof in Section 1.3. Among these the most central role is played by the *Chekanov–Eliashberg algebra with based loop space coefficients*, denoted as CE^*. As we show, any other invariant can be obtained from CE^* by algebraic manipulation. Finally, in Section 1.4 we give detailed calculations of the invariants introduced, in the simple yet illustrative example of the Legendrian Hopf link filled by two Lagrangian disks intersecting transversely in one point.

The starting point for our study is a construction in classical topology. Consider the following augmented DG–algebras over a field $\mathbb{K}$ associated to a based, connected, topological space (M, pt):

$$C^*(M) \to \mathbb{K}, \quad C_{-*}(\Omega M) \to \mathbb{K},$$

where $C^*(M)$ is the singular cochain complex equipped with the cup product and $C_{-*}(\Omega M)$ is the singular chain complex of the based (Moore) loop space of M equipped with the Pontryagin product. (We use cohomologically graded complexes throughout the paper so that all differentials increase the grading by 1.) In the case of singular cohomology, the inclusion $i\colon \mathrm{pt} \to M$ gives the augmentation $i^*\colon C^*(M) \to C^*(\mathrm{pt}) = \mathbb{K}$ and in the case of the based loop space, the augmentation is given by the trivial local system $\pi_1(M, \mathrm{pt}) \to \mathbb{K}$.

If M is of finite-type (for example, a finite CW–complex), then it is well known that one can recover the augmented DG–algebra $C^*(M)$ from the augmented DG–algebra $C_{-*}(\Omega M)$ by the Eilenberg–Moore equivalence

$$C^*(M) \simeq \mathrm{RHom}_{C_{-*}(\Omega M)}(\mathbb{K}, \mathbb{K}).$$

In the other direction, if M is *simply connected*, then the Adams construction gives a quasi-isomorphism

$$C_{-*}(\Omega M) \simeq \mathrm{RHom}_{C^*(M)}(\mathbb{K}, \mathbb{K}),$$

and in this case $C^*(M)$ and $C_{-*}(\Omega M)$ are said to be *Koszul dual* DG–algebras. Koszul duality is sometimes abbreviated and simply called *duality*. For more general M, using the method of acyclic models, Brown [13] constructed a twisting cochain

$$\mathfrak{t}\colon C_{-*}(M) \to C_{-*}(\Omega M).$$

This is a degree 1 map that induces a DG–algebra map $\Omega C_{-*}(M) \to C_{-*}(\Omega M)$, where $\Omega C_{-*}(M)$ is the cobar construction applied to chains on M; see Section 2.2.1. By definition, $\mathfrak{t}$ is a quasi-isomorphism when duality holds, and this can be detected by an associated *Koszul complex*, which is acyclic if and only if duality holds. In the general case, $\Omega C_{-*}(M)$ is a certain completion of $C_{-*}(\Omega M)$ and consequently $C_{-*}(\Omega M)$ is a more refined invariant of M than $\Omega C_{-*}(M)$.

In this paper, we pursue this idea in the context of invariants associated to Lagrangian and Legendrian submanifolds. Here the role played by simple connectedness in the above discussion has two natural counterparts: one corresponds to a generalized notion of simple connectedness for intersecting Lagrangian submanifolds and the other is the usual notion of simple connectedness for Legendrian submanifolds.

We start with the geometric data of a Liouville domain X with convex boundary Y and an exact Lagrangian submanifold $L \subset X$ with Legendrian boundary $\Lambda \subset Y$. We assume that $c_1(X) = 0$, that the Maslov class of L vanishes (for grading purposes) and that L is relatively spin (to orient certain moduli spaces of holomorphic disks). Assume that L is subdivided into embedded components intersecting transversely $L = \bigcup_{v \in \Gamma} L_v$, and that Λ is subdivided into connected components $\Lambda = \bigsqcup_{v \in \Gamma} \Lambda_v$. To avoid notational complications, we take both parametrized by the same finite set Γ and assume that the boundary of L_v is Λ_v. We use a base field $\mathbb{K}$ and define the semisimple ring

$$\boldsymbol{k} = \bigoplus_{v \in \Gamma} \mathbb{K} e_v,$$

generated by mutually orthogonal idempotents e_v. Also, we fix a partition

$$\Gamma = \Gamma^+ \cup \Gamma^-$$

into two disjoint sets, and choose a basepoint $p_v \in \Lambda_v$ for each $v \in \Gamma^+$.

For simplicity, let us restrict, in this introduction, to the following situation:

- X is a subcritical Liouville domain.
- If $v \in \Gamma^-$, then the corresponding Legendrian Λ_v is an embedded *sphere*.

From a technical point of view, these restrictions are unnecessary. We make them in order to facilitate the explanation of our constructions from the perspective of Legendrian surgery. (Note that the topology of Λ_v is unrestricted when $v \in \Gamma^+$.)

We write X_Λ for the completion of the Liouville sector obtained from X by attaching critical Weinstein handles along Λ_v for each $v \in \Gamma^-$, and attaching cotangent cones $T^*(\Lambda_v \times [0,\infty))$ along Λ_v for each $v \in \Gamma^+$. If $\Gamma^+ = \varnothing$, X_Λ is an ordinary Liouville manifold. In this case Gromov compactness is ensured by convexity of the boundary. When $\Gamma^+ \neq \varnothing$, we also have part of the boundary that can be identified with a neighborhood of the zero section in the cotangent bundle $\bigcup_{v\in\Gamma^+} T^*(\Lambda_v \times [T,\infty))$, for some $T > 0$. This is a geometrically bounded manifold, hence Gromov compactness [38] still holds, and holomorphic curve theory is well behaved.

In X_Λ, for $v \in \Gamma^-$ there is a closed exact Lagrangian submanifold $S_v = L_v \cup D_v$, the union of the Lagrangian L_v in X and the *Lagrangian core disk* D_v of the Weinstein handle attached to Λ_v, and for $v \in \Gamma^+$ there is a noncompact Lagrangian obtained by attaching the cylindrical boundary $\Lambda_v \times [0,\infty)$ to L_v for $v \in \Gamma^+$, which we will still denote by L_v, by abuse of notation, even when we view them now in X_Λ. Dually, for each $v \in \Gamma^-$ we obtain (noncompact) exact Lagrangian disks C_v, the *Lagrangian cocore disks* of the Weinstein handles attached to Λ_v on X, and for each $v \in \Gamma^+$ we construct *dual* Lagrangians disks C_v intersecting L_v once and asymptotic to a Legendrian meridian of L_v — these can be constructed as the cotangent fiber at the point (p_v, t), $t > 0$, in $T^*(\Lambda_v \times [0,\infty)) \subset X_\Lambda$, where p_v is the basepoint on Λ_v.

The invariants we will construct are associated to the unions of Lagrangian submanifolds

$$L_\Lambda := \bigcup_{v\in\Gamma^+} L_v \;\cup\; \bigcup_{v\in\Gamma^-} S_v \quad \text{and} \quad C_\Lambda := \bigcup_{v\in\Gamma} C_v.$$

The Lagrangian L_Λ will be referred to as a *Lagrangian skeleton* of X_Λ; it is a union of Lagrangian submanifolds which intersect transversely. The dual Lagrangian C_Λ is the union of Lagrangian disks which can be locally identified with cotangent fibers to irreducible components of L_Λ.

We will study two algebraic invariants associated to $(X_\Lambda, L_\Lambda, C_\Lambda)$. The first one is the *Legendrian A_∞–algebra*, LA^*. It corresponds to the endomorphism algebra of L_Λ considered in the infinitesimal Fukaya category of X_Λ (Theorem 63). The second one is the *Chekanov–Eliashberg DG–algebra*, CE^*. It corresponds to the endomorphism algebra of C_Λ considered in the partially wrapped Fukaya category of X_Λ (Theorem 83). However, we will take the pre-surgery perspective as in Bourgeois, Ekholm and Eliashberg [11] and construct all these invariants by studying Legendrian invariants of $\Lambda \subset X$ rather than Floer cohomology in X_Λ. From this perspective, the case $\Gamma^+ \neq \varnothing$ is a new construction, which generalizes the theory from [11] in a way analogous to how the partially wrapped Fukaya categories of Sylvan [62] generalize the wrapped Fukaya categories of Abouzaid and Seidel [3].

The invariants LA^* and CE^* come equipped with canonical augmentations to the semisimple ring $\boldsymbol{k}$, and it is easy to see by construction that LA^* is determined by CE^* via the equivalence

$$\mathrm{LA}^* \simeq \mathrm{RHom}_{\mathrm{CE}^*}(\boldsymbol{k}, \boldsymbol{k}).$$

The duality which would recover CE^* from LA^* holds in the "simply connected" case; see Section 2.3. In the topological case discussed above, this is analogous to the simple-connectedness assumption on M, which makes the augmented algebras $C^*(M)$ and $C_{-*}(\Omega M)$ Koszul dual. In fact, the topological case is a special case of our study for the Weinstein manifold T^*M, with the Lagrangian skeleton $L_\Lambda = M$ given by the zero section, and the dual Lagrangian C_Λ given by a cotangent fiber T_p^*M. This is because the wrapped Floer cohomology complex of a cotangent fiber is quasi-isomorphic to $C_{-*}(\Omega M)$ by Abouzaid [2] and the Floer cohomology complex of the zero section is quasi-isomorphic to $C^*(M)$ (Fukaya and Oh [35]) as augmented A_∞–algebras.

We next sketch the definition of our version of the Chekanov–Eliashberg DG–algebra without any assumption of simple connectedness; see Section 3 for details. This is the DG–algebra over $\boldsymbol{k}$ called CE^* above. Its underlying $\boldsymbol{k}$–bimodule is the unital $\boldsymbol{k}$–algebra generated by Reeb chords between components of Λ and chains in $C_{-*}(\Omega_{p_v}\Lambda_v)$ for $v \in \Gamma^+$. (This is the crucial distinction between Γ^+ and Γ^-.)

We use the cubical chain complex (cf Serre [60]) $C_{-*}(\Omega_{p_v}\Lambda_v)$ for $v \in \Gamma^+$ — see Section 3.1 for a discussion of other possible choices of chain models — to express CE^* as a free algebra over $\boldsymbol{k}$ generated by Reeb chords c and generators of $C_{-*}(\Omega_{p_v}\Lambda_v)$ for $v \in \Gamma^+$. The differential on CE^* is determined by its action on generators. On a generator of $C_{-*}(\Omega_{p_v}\Lambda_v)$ we simply apply the usual differential. On a generator c_0 which is a Reeb chord, the differential is determined by moduli spaces of holomorphic disks in the symplectization $\mathbb{R} \times Y$ which asymptotically converge to c_0 on the positive end and chords $c_1, \dots, c_i$ at the negative end as follows. Consider the moduli space of J–holomorphic maps $u\colon D \to \mathbb{R} \times Y$, where D is a disk with $k+1$ boundary punctures $z_j \in \partial D = S^1$ that are mutually distinct with $(z_0, z_1, \dots z_k)$ respecting the counterclockwise cyclic order of S^1, and u sends the boundary component (z_{j-1}, z_j) of $S^1 \setminus \{z_0, \dots, z_k\}$ to $\mathbb{R} \times \Lambda$ and is asymptotic to c_j near the puncture at z_j for $j = 1, \dots, k$ and to c_0 near the puncture at z_0 (as usual these disks may be anchored

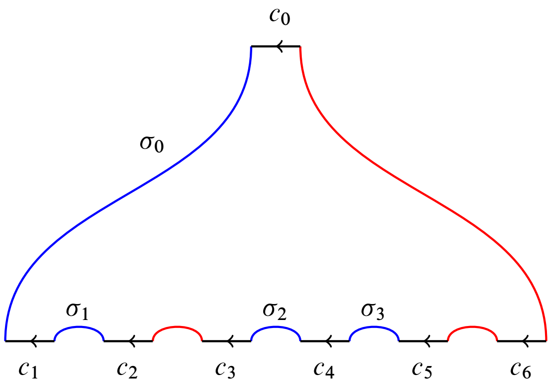

Figure 1: The differential in CE^*: the word $\sigma_0 c_1 \sigma_1 c_2 c_3 \sigma_2 c_4 \sigma_3 c_5 c_6$ appears in dc_0.

in X). The moduli space, which is naturally a stratified space with manifold strata that carries a fundamental chain, comes with evaluation maps to $\Omega_{p_v}\Lambda_v$ for $v \in \Gamma^+$. The image of the fundamental chain determines a word in our chain model of $C_{-*}(\Omega_{p_v}\Lambda_v)$. Reading these together with the Reeb chords in order gives the differential of c_0.

We remark that loop space coefficients have been used in the context of Lagrangian Floer cohomology before; see Barraud and Cornea [7] and Fukaya [34]. See also Abouzaid [2] and Cieliebak and Latschev [19] for uses of high-dimensional moduli spaces in Floer theory.

While CE^* with loop space coefficients is a powerful invariant, it is in general hard to compute as it involves high-dimensional moduli spaces of disks. As mentioned above, duality in the Legendrian Λ will also play a role. More precisely, we define another DG–algebra $\mathrm{CE}^*_{\parallel}$ related to CE^* via a Morse-theoretic version of Adams cobar construction whose definition involves taking parallel copies of Λ but uses only 0–dimensional moduli spaces; see Section 3.4. In fact, we prove that the two DG–algebras are quasi-isomorphic when all Λ_v for $v \in \Gamma^+$ are simply connected.

Theorem 1 *There exists a DG–algebra map*

$$\mathrm{CE}^* \to \mathrm{CE}^*_{\parallel},$$

which is a quasi-isomorphism when the Λ_v are simply connected for all $v \in \Gamma^+$.

Theorem 1 is restated and proved as Theorem 51 in Section 3.4.

1.1 Partially wrapped Fukaya categories by surgery

Let $\Lambda = \bigsqcup_{v\in\Gamma} \Lambda_v$ be a Legendrian submanifold and $\Gamma = \Gamma^+ \cup \Gamma^-$ be as above. Furthermore, we use the notation above for cocore disks and write $\mathrm{CE}^* = \mathrm{CE}^*(\Lambda)$. An important result that is implicit in [11, Remark 5.9] is the following:

Theorem 2 *Suppose $\Gamma = \Gamma^-$. Then there exists a surgery map defined via a holomorphic disk count that gives an A_∞–quasi-isomorphism between the wrapped Floer cochain complex $\mathrm{CW}^* := \bigoplus_{v,w\in\Gamma^-} \mathrm{CW}^*(C_v, C_w)$ and the Legendrian DG–algebra CE^*.*

We prove Theorem 2 in Section B.2 following [11], referring to Ekholm [25] for the necessary technical results omitted there. Section B.1 also contains a construction of wrapped Floer A_∞–algebras that uses only purely holomorphic disks (without Hamiltonian perturbation), and a proof that this agrees with the more standard version defined in Abouzaid and Seidel [3], which uses Hamiltonian perturbations.

One of the main guiding principles for the results in this paper is that Theorem 2 remains true when Γ^+ is nonempty, provided the Lagrangians C_v are considered as objects of the *partially wrapped* Fukaya category of X_Λ, where the noncapped Legendrians Λ_v for $v \in \Gamma^+$ serve as *stops*; cf Sylvan [62]. The full proof of this result when Γ^+ is nonempty can be reduced to the standard surgery result, Theorem 2, and will appear elsewhere. Here we give an outline of a somewhat different and more topological proof; see Section B.3. We will use the geometric intuition provided by this viewpoint, and our constructions of Legendrian invariants provide a rigorous "working definition" of CE^* even in the case that Γ^+ is nonempty, as well as a starting point for the study of "partially wrapped Fukaya categories" via Legendrian surgery (extending the scope of [11] considerably). For future reference, we state this result as a conjecture:

Conjecture 3 *There exists a surgery map defined via moduli spaces of holomorphic disks which gives an A_∞–quasi-isomorphism between the partially wrapped Floer cochain complex $\mathrm{CW}^* := \bigoplus_{v,w\in\Gamma} \mathrm{CW}^*(C_v, C_w)$ and the DG–algebra CE^*.*

While writing this paper, we learned that Sylvan [61] independently considered a similar conjecture in relation with his theory of partially wrapped Fukaya categories [62].

1.2 Augmentations and infinitesimal Fukaya categories

We keep the notation above and now consider an exact Lagrangian filling L in X of Λ. Such a filling gives an augmentation

$$\epsilon_L \colon \mathrm{CE}^* \to \boldsymbol{k}.$$

For chords on components Λ_v with $v \in \Gamma^-$, this is well known and given by a count of rigid disks with one positive puncture and boundary on L_v.

For components Λ_v with $v \in \Gamma^+$, we define an augmentation using the same disks. More formally, we define a chain map

$$\beta_L \colon \mathrm{CE}^* \to \bigoplus_{v \in \Gamma^+} C_{-*}(\Omega L_v),$$

which acts on chains in $\bigoplus_{v\in\Gamma^+} C_{-*}(\Omega\Lambda_v)$ by the inclusion and on Reeb chords c as the chain of loops carried by the moduli spaces of holomorphic disks with boundary on L_v (for each v) and a positive puncture at c. The augmentation ϵ_L is then this map followed by the augmentation on $\bigoplus_{v\in\Gamma^+} C_{-*}(\Omega L_v) \to \boldsymbol{k}$ that takes higher-dimensional chains to zero and takes any loop in L_v to e_v.

This allows us to write

$$\mathrm{CE}^* = \Omega\mathrm{LC}_*$$

for an A_∞–coalgebra $\mathrm{LC}_* = \mathrm{LC}_*(\Lambda)$ that we call the *Legendrian A_∞–coalgebra* (which depends on ϵ_L). Here Ω is the Adams cobar construction. Writing $\mathrm{LA}^* := (\mathrm{LC}_*)^\#$ for the A_∞–algebra which is the linear dual of LC_*, the following result recovers the Floer cochain complex of L in X_Λ:

Theorem 4 *There is an A_∞–quasi-isomorphism between $\mathrm{CF}^* := \mathrm{CF}^*(L_\Lambda)$, the Floer cochain complex in the infinitesimal Fukaya category of X_Λ, and the A_∞–algebra LA^*.*

By the general properties of bar–cobar constructions (see Section 2.2.1), the algebra $\mathrm{RHom}_{\Omega\mathrm{LC}_*}(\boldsymbol{k}, \boldsymbol{k})$ is quasi-isomorphic to the graded $\boldsymbol{k}$–dual of the bar construction on the algebra $\Omega\mathrm{LC}_*$, which can be computed as

$$(\mathrm{B}\Omega\mathrm{LC}_*)^\# \cong (\mathrm{LC}_*)^\# = \mathrm{LA}^*. \tag{1}$$

Remark 5 If Γ^+ is empty, the A_∞–algebra LA^* is obtained from the construction in Civan, Koprowski, Etnyre, Sabloff and Walker [20] and Bourgeois and Chantraine [10], known as the Aug_ category, by adding a copy of $\boldsymbol{k}$, making it unital.

If Γ^- is empty, the A_∞–algebra $\mathrm{LA}^* \approx (\mathrm{BCE}^*)^\#$ — see (1) — is the endomorphism algebra of Λ with the augmentation ϵ_L in the Aug_+ category of Ng, Rutherford, Shende, Sivek and Zaslow [55]. In the setting of microlocal sheaves, a related result was obtained by Nadler [53, Theorem 1.6].

1.3 Duality between partially wrapped and infinitesimal Fukaya categories

We study duality in the setting of the two categories described above: the partially wrapped Fukaya category and the infinitesimal Fukaya category of X_Λ (after surgery) or, equivalently, the augmented DG–algebra $\Omega\mathrm{LC}_*$ and the augmented A_∞–algebra LA^* (before surgery).

As we have seen in Theorem 4, the augmented DG–algebra $\Omega\mathrm{LC}_*$ determines the augmented (unital) A_∞–algebra CF^*. Now, a natural question is to what extent the quasi-isomorphism type of the A_∞–algebra CF^* determines the quasi-isomorphism type of the augmented Legendrian DG–algebra $\Omega\mathrm{LC}_*$.

We emphasize here the phrase "quasi-isomorphism type": even though it is possible to construct chain models of the A_∞–algebra LA^* (which is A_∞–quasi-isomorphic to CF^*) and the DG–algebra $\Omega\mathrm{LC}_*$ by counting "the same" holomorphic disks interpreted in different ways, the two algebras are considered with respect to different equivalence relations, and the resulting equivalence classes can be very different. In particular, it is *not* generally true that $\mathfrak{f}\colon \mathscr{C} \to \mathscr{D}$ being a quasi-isomorphism of A_∞–coalgebras implies that $\Omega\mathfrak{f}\colon \Omega\mathscr{C} \to \Omega\mathscr{D}$ is a quasi-isomorphism.

We will study this question by (geometrically) constructing a *twisting cochain*

$$\mathfrak{t}\colon \mathrm{LC}_* \to (\mathrm{BCF}^*)^\#,$$

where B stands for the bar construction and # is the graded $\boldsymbol{k}$–dual. See Section 2.1.4. This twisting cochain induces a map of DG–algebras,

$$\Omega\mathrm{LC}_* \to \mathrm{RHom}_{\mathrm{CF}^*}(\boldsymbol{k}, \boldsymbol{k}),$$

which is a quasi-isomorphism if and only if $\mathfrak{t}$ is a *Koszul* twisting cochain. For example, we will prove the following result:

Theorem 6 *Suppose that* LC_* *is a locally finite, simply connected* $\boldsymbol{k}$*–bimodule. Then the natural map* $\Omega\mathrm{LC}_* \to \mathrm{RHom}_{\mathrm{CF}^*}(\boldsymbol{k}, \boldsymbol{k})$ *is a quasi-isomorphism.*

This is an instance of *Koszul duality* between the A_∞–algebras $\Omega\mathrm{LC}_*$ and CF^*. It has many useful implications; for example, it implies an isomorphism between Hochschild cohomologies,

$$\mathrm{HH}^*(\Omega\mathrm{LC}_*, \Omega\mathrm{LC}_*) \cong \mathrm{HH}^*(\mathrm{CF}^*, \mathrm{CF}^*).$$

When $\Gamma^+ = \varnothing$, an isomorphism defined via a surgery map [11] was described between *symplectic cohomology*, $\mathrm{SH}^* = \mathrm{SH}^*(X_\Lambda)$, and the Hochschild cohomology $\mathrm{HH}^*(\Omega\mathrm{LC}_*, \Omega\mathrm{LC}_*)$. Therefore, when duality holds (ie $\mathfrak{t}$ induces an isomorphism), we obtain a more economical way of computing SH^*.

In the case of cotangent bundles T^*M of simply connected manifolds M, this recovers a classical result due to Jones [42], which gives

$$H_{n-*}(\mathcal{L}M) \cong \mathrm{HH}^*(C_{-*}(\Omega M), C_{-*}(\Omega M)) \cong \mathrm{HH}^*(\mathrm{CF}^*(M), \mathrm{CF}^*(M)),$$

where M is a simply connected manifold of dimension n and $\mathcal{L}M$ denotes the free loop space of M.

In Section 6, we give several concrete examples where the duality holds beyond the case of cotangent bundles. For example, the duality holds for plumbings of simply connected cotangent bundles according to an arbitrary plumbing tree; see Theorem 68.

In another direction, combining duality and Floer cohomology with local coefficients, we establish the following result for relatively spin exact Lagrangian fillings $L \subset X$ with vanishing Maslov class of a Legendrian submanifold $\Lambda \subset Y$.

Theorem 7 *Let $\Gamma = \Gamma^-$ and assume that $\mathrm{SH}^*(X) = 0$ and that Λ is simply connected. If $\mathrm{CE}^*(\Lambda)$ is supported in degrees < 0, then L is simply connected. Moreover, if Λ is a sphere, then $\mathrm{CE}^*(\Lambda)$ is isomorphic to $C_{-*}(\Omega \bar{L})$, where $\bar{L} = L \cup_\Lambda D$, for a disk D with boundary $\partial D = \Lambda$.*

In general, duality between $\Omega\mathrm{LC}_*$ and CF^* does not hold — as can be seen for example by looking at cotangent bundles of non-simply-connected manifolds, or letting Λ be the standard Legendrian trefoil knot in S^3 filled by a punctured torus. However, there are cases when duality holds even if LC_* is not simply connected, for instance because of the existence of an auxiliary weight grading (see Etgü and Lekili [32]), or, for an example in the 1–dimensional case, see Lekili and Polishchuk [48]. It is a very interesting open question to find a geometric characterization of when duality holds.

Remark 8 Constructions of Legendrian and Lagrangian holomorphic curve invariants require the use of perturbations to achieve transversely cut out moduli spaces. For our main invariant CE^*, all moduli spaces used can be shown to be transverse by classical techniques (see Theorem 74) except for the rigid holomorphic planes in X_Λ with a single positive end that are used to anchor the disks (in the terminology of [11]). These are also relevant for defining the wrapped Floer cochain complex CW^* without Hamiltonian perturbations and for constructing the surgery map. In all cases, there is a distinguished boundary puncture in the main disk that determines asymptotic markers on the split-off planes. Taking this marker into account removes symmetries of the planes, and a specific perturbation scheme for transversality of the resulting moduli spaces was constructed in [25]. We will use that perturbation scheme here; see Section A.2 for details.

1.4 An example: the Hopf link

In this section, we study the example of the Hopf link in order to illustrate our results in a simple and computable example. Some of the algebraic constructions used here are explained in detail only later; see Section 2.

Let $\Lambda \subset S^3$ be the standard Legendrian Hopf link. We work over $\boldsymbol{k} = \mathbb{K}e_1 \oplus \mathbb{K}e_2$ and with the Lagrangian filling L given by two disks in D^4 that intersect transversely in a single point. We choose the partition $\Lambda = \Lambda^+ \cup \Lambda^-$. This means that after attaching a Weinstein 2–handle to Λ^- and $T^*(S^1 \times [0,\infty))$ to Λ^+, we obtain the symplectic manifold X_Λ with Lagrangian skeleton

$$L_\Lambda = S^2 \cup T^*_{\mathrm{pt}} S^2 \subset T^* S^2,$$

or, in the terminology of [62], X_Λ is T^*S^2 with wrapping stopped by a Legendrian fiber sphere. The DG–algebra $\mathrm{CE}^* = \mathrm{CE}^*(\Lambda)$ of Λ has coefficients in

$$C_{-*}(\Omega\Lambda^+)e_1 \oplus \mathbb{K}e_2 \cong \mathbb{K}[t,t^{-1}]e_1 \oplus \mathbb{K}e_2.$$

A free model for $\mathbb{K}[t,t^{-1}]$ is given by the tensor algebra $\mathbb{K}\langle s_1,t_1,k_1,l_1,u_1\rangle$ with $|s_1| = |t_1| = 0$, $|k_1| = |l_1| = -1$, $|u_1| = -2$ and the differential

$$dk_1 = e_1 - s_1t_1, \quad dl_1 = e_1 - t_1s_1, \quad du_1 = k_1s_1 - s_1l_1.$$

The natural map $\mathbb{K}\langle s_1,t_1,k_1,l_1,u_1\rangle \to \mathbb{K}[t,t^{-1}]$ sending $t_1 \to t$ and $s_1 \to t^{-1}$ is a quasi-isomorphism. The subscripts indicate that as $\boldsymbol{k}$–module generators, s_1, t_1, k_1, l_1 and u_1 are annihilated by the idempotent e_2.

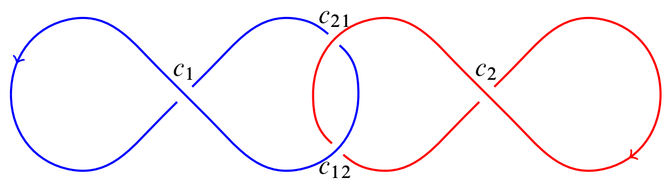

Figure 2: Hopf link when both Γ^+ and Γ^- are nonempty: the blue component lies in Γ^+ and the red in Γ^-.

Next, incorporating the Reeb chords, with notation as in Figure 2, we get the free algebra

$$\boldsymbol{k}\langle c_{12}, c_{21}, c_1, c_2, s_1, t_1, k_1, l_1, u_1\rangle$$

with gradings

$$|u_1| = -2, \quad |c_1| = |c_2| = |k_1| = |l_1| = -1, \quad |c_{12}| = |c_{21}| = |s_1| = |t_1| = 0$$

and differential

$$dc_1 = e_1 + s_1 + c_{12}c_{21}, \quad dc_2 = c_{21}c_{12},$$
$$dk_1 = e_1 - s_1 t_1, \quad dl_1 = e_1 - t_1 s_1, \quad du_1 = k_1 s_1 - s_1 l_1.$$

The only augmentation to $\boldsymbol{k}$ is given by $\epsilon(s_1) = \epsilon(t_1) = -e_1$ and $\epsilon(c_1) = \epsilon(c_2) = \epsilon(c_{12}) = \epsilon(c_{21}) = \epsilon(k_1) = \epsilon(l_1) = \epsilon(u_1) = 0$. After change of variables, $s_1 \to s_1 - e_1$ and $t_1 \to t_1 - e_1$, we obtain the free algebra

$$\boldsymbol{k}\langle c_{12}, c_{21}, c_1, c_2, s_1, t_1, k_1, l_1, u_1\rangle$$

with nonzero differential on generators

$$\begin{gathered} dc_1 = s_1 + c_{12}c_{21}, \quad dc_2 = c_{21}c_{12}, \\ dk_1 = s_1 + t_1 - s_1 t_1, \quad dl_1 = s_1 + t_1 - t_1 s_1, \quad du_1 = l_1 - k_1 + k_1 s_1 - s_1 l_1. \end{gathered} \tag{2}$$

On the other hand, we can compute the Floer cochains $\mathrm{CF}^* = \mathrm{CF}^*(L_\Lambda)$ of L_Λ as

$$\mathrm{CF}^* = \boldsymbol{k} \oplus \mathbb{K}a_{12} \oplus \mathbb{K}a_{21} \oplus \mathbb{K}a_2, \quad \text{where } |a_2| = 2,\ |a_{12}| = |a_{21}| = 1.$$

The cohomology level computation follows easily from the geometric picture and general properties of Floer cohomology: L_Λ is a union of a disk D^2 and a sphere S^2 that intersect transversely in one point, and we have

$$\mathrm{HF}^*(D^2, D^2) = \mathbb{K}e_1, \qquad \mathrm{HF}^*(S^2, S^2) = \mathbb{K}e_2 \oplus \mathbb{K}a_2,$$
$$\mathrm{HF}^*(D^2, S^2) = \mathbb{K}a_{12}, \quad \mathrm{HF}^*(S^2, D^2) = \mathbb{K}a_{21}.$$

The only nontrivial product that does not involve idempotents is $\mathfrak{m}_2(a_{12}, a_{21}) = a_2$. For degree reasons, the only possible nontrivial higher products are

$$\mathfrak{m}_{2k}(a_{12}, a_{21}, \dots, a_{12}, a_{21}) = ca_2 \quad \text{for some } k > 1 \text{ and } c \in \mathbb{K}.$$

It turns out that one can take $c = 0$ for all $k > 1$. Indeed, assuming that the A_∞–structure is strictly unital (which can be arranged up to quasi-isomorphism), consider the A_∞–relation that involves the term

$$\mathfrak{m}_2(m_{2k}(a_{12}, a_{21}, \dots, a_{12}, a_{21}), e_2).$$

By induction on $k > 1$, this term has to vanish, implying $m_{2k}(a_{12}, a_{21}, \dots, a_{12}, a_{21})$ has to vanish for all $k > 1$. Let us confirm this by using the quasi-isomorphism

$$\mathrm{CF}^* \cong \mathrm{RHom}_{\mathrm{CE}^*}(\boldsymbol{k}, \boldsymbol{k}).$$

We introduce the counital A_∞–coalgebra

$$\mathrm{LC}_* = \boldsymbol{k} \oplus \mathbb{K}c_{12} \oplus \mathbb{K}c_{21} \oplus \mathbb{K}c_1 \oplus \mathbb{K}c_2 \oplus \mathbb{K}s_1 \oplus \mathbb{K}t_1 \oplus \mathbb{K}k_1 \oplus \mathbb{K}l_1 \oplus \mathbb{K}u_1$$

with $|u_1| = -3$, $|c_1| = |c_2| = |k_1| = |l_1| = -2$ and $|c_{12}| = |c_{21}| = |s_1| = |t_1| = -1$, for which $\Delta_i = 0$ except for $i = 1$ or 2, where there are the nonzero terms

$$\Delta_1(c_1) = s_1, \quad \Delta_1(k_1) = s_1 + t_1, \quad \Delta_1(l_1) = s_1 + t_1, \quad \Delta_1(u_1) = l_1 - k_1.$$

Write $\Delta_2(x) = 1 \otimes_{\boldsymbol{k}} x + x \otimes_{\boldsymbol{k}} 1 + \bar{\Delta}_2(x)$. Then

$$\bar{\Delta}_2(c_1) = c_{12}c_{21}, \quad \bar{\Delta}_2(c_2) = c_{21}c_{12},$$
$$\bar{\Delta}_2(k_1) = -s_1t_1, \quad \bar{\Delta}_2(l_1) = -t_1s_1, \quad \bar{\Delta}_2(u_1) = k_1s_1 - s_1l_1,$$

where the A_∞ coalgebra operations on LC_* are defined so that $\Omega\mathrm{LC}_*$ is isomorphic to CE^*. Thus, $\mathrm{RHom}_{\mathrm{CE}^*}(\boldsymbol{k}, \boldsymbol{k})$ can be computed as the graded dual of LC_* which is the A_∞–algebra

$$\mathrm{LA}^* = \boldsymbol{k} \oplus \mathbb{K}c_{12}^\vee \oplus \mathbb{K}c_{21}^\vee \oplus \mathbb{K}c_1^\vee \oplus \mathbb{K}c_2^\vee \oplus \mathbb{K}s_1^\vee \oplus \mathbb{K}t_1^\vee \oplus \mathbb{K}k_1^\vee \oplus \mathbb{K}l_1^\vee \oplus \mathbb{K}u_1^\vee,$$

with gradings

$$|u_1^\vee| = 3, \quad |c_1^\vee| = |c_2^\vee| = |k_1^\vee| = |l_1^\vee| = 2, \quad |c_{12}^\vee| = |c_{21}^\vee| = |s_1^\vee| = |t_1^\vee| = 1,$$

where $c^\vee$ is the linear dual of the generator c of LC_*. The differential is

$$\mathfrak{m}_1(s_1^\vee) = c_1^\vee + k_1^\vee + l_1^\vee, \quad \mathfrak{m}_1(t_1^\vee) = k_1^\vee + l_1^\vee, \quad \mathfrak{m}_1(k_1^\vee) = -u_1, \quad \mathfrak{m}_1(l_1^\vee) = u_1,$$

and the products that do not involve idempotents are

$$\mathfrak{m}_2(c_{12}^\vee, c_{21}^\vee) = c_2^\vee, \quad \mathfrak{m}_2(c_{21}^\vee, c_{12}^\vee) = c_1^\vee, \quad \mathfrak{m}_2(t_1^\vee, s_1^\vee) = -k_1^\vee,$$
$$\mathfrak{m}_2(s_1^\vee, t_1^\vee) = -l_1^\vee, \quad \mathfrak{m}_2(k_1^\vee, s_1^\vee) = u_1^\vee, \quad \mathfrak{m}_2(s_1^\vee, l_1^\vee) = -u_1^\vee.$$

All the higher products vanish. We claim that this A_∞–algebra is quasi-isomorphic to the algebra

$$\boldsymbol{k} \oplus \mathbb{K}a_{12} \oplus \mathbb{K}a_{21} \oplus \mathbb{K}a_2, \quad \text{where } |a_2| = 2,\ |a_{12}| = |a_{21}| = 1,$$

with the only nontrivial product (not involving idempotents) given by

$$\mathfrak{m}_2(a_{12}, a_{21}) = a_2.$$

Indeed, it is easy to show that the map defined by

$$c_{12}^\vee \to a_{12}, \quad c_{21}^\vee \to a_{21}, \quad c_2^\vee \to a_2 \quad \text{and} \quad c_1^\vee, s_1^\vee, t_1^\vee, k_1^\vee, l_1^\vee, u_1^\vee \to 0$$

is a DG–algebra (hence also an A_∞–algebra) map, which induces an isomorphism at the level of cohomology.

Dually, we can construct a DG–algebra map

$$\mathrm{CE}^* \to \mathrm{RHom}_{\mathrm{CF}^*}(\boldsymbol{k}, \boldsymbol{k}).$$

The Floer cochain complex CF^* has a unique augmentation given by projection to $\boldsymbol{k}$, and we compute

$$\mathrm{RHom}_{\mathrm{CF}^*}(\boldsymbol{k}, \boldsymbol{k}) \cong \Omega\mathrm{CF}_*,$$

where CF_* is the coalgebra dual to CF^*. This is the free coalgebra

$$\boldsymbol{k}\langle a_{12}^\vee, a_{21}^\vee, a_2^\vee\rangle$$

with $|a_{12}^\vee| = |a_{21}^\vee| = 0$ and $|a_2^\vee| = -1$, and the only nontrivial differential not involving counits is

$$\Delta_2(a_2^\vee) = a_{21}^\vee a_{12}^\vee.$$

We have a twisting cochain

$$\mathfrak{t}\colon \mathrm{LC}_* \to \Omega\mathrm{CF}_*$$

given by

$$\mathfrak{t}(c_2) = a_2^\vee, \qquad \mathfrak{t}(c_{12}) = a_{12}^\vee, \qquad \mathfrak{t}(c_{21}) = a_{21}^\vee,$$
$$\mathfrak{t}(c_1) = 0, \qquad \mathfrak{t}(s_1) = -a_{12}^\vee a_{21}^\vee, \qquad \mathfrak{t}(t_1) = a_{12}^\vee a_{21}^\vee,$$
$$\mathfrak{t}(l_1) = \mathfrak{t}(k_1) = a_{12}^\vee a_2^\vee a_{21}^\vee, \qquad \mathfrak{t}(u_1) = a_{12}^\vee a_2^\vee a_2^\vee a_{21}^\vee.$$

This means that $\mathfrak{t}$ satisfies the equations

$$\begin{aligned}
d\mathfrak{t}(c_1) &= \mathfrak{t}(s_1) + \mathfrak{t}(c_{12})\mathfrak{t}(c_{21}),\\
d\mathfrak{t}(c_2) &= \mathfrak{t}(c_{21})\mathfrak{t}(c_{12}),\\
d\mathfrak{t}(k_1) &= \mathfrak{t}(s_1) + \mathfrak{t}(t_1) - \mathfrak{t}(s_1)\mathfrak{t}(t_1),\\
d\mathfrak{t}(l_1) &= \mathfrak{t}(s_1) + \mathfrak{t}(t_1) - \mathfrak{t}(t_1)\mathfrak{t}(s_1),\\
d\mathfrak{t}(u_1) &= \mathfrak{t}(l_1) - \mathfrak{t}(k_1) + \mathfrak{t}(k_1)\mathfrak{t}(s_1) - \mathfrak{t}(s_1)\mathfrak{t}(l_1).
\end{aligned}$$

Hence, it induces a DG–algebra map

$$\Omega \mathrm{LC}_* \to \Omega \mathrm{CF}_*.$$

We have not checked whether this is a quasi-isomorphism, or equivalently whether $\mathfrak{t}$ is a *Koszul* twisting cochain. Note, however, that the DG–algebra map $\Omega\mathrm{CF}_* \to \Omega\mathrm{LC}_*$ defined by

$$a_2^\vee \to c_2, \quad a_{12}^\vee \to c_{12}, \quad a_{21}^\vee \to c_{21}$$

shows that $\mathfrak{t}$ is a retraction, and $\Omega\mathrm{CF}_*$ is a retract of $\Omega\mathrm{LC}_*$.

Acknowledgements Ekholm is supported by the Knut and Alice Wallenberg Foundation as a Wallenberg scholar KAW2020.0307 and by the Swedish Research Council VR2020-04535. Lekili is supported in part by the Royal Society (URF) and the NSF grant DMS-1509141. Both authors would like to thank the Mittag-Leffler Institute for hospitality and excellent working conditions. We also thank Lenny Ng for providing the example in Section 6.1.6, and Zack Sylvan and Paolo Ghiggini for helpful comments.

2 Algebraic preliminaries

In this section, we review the homological algebra we use in our study of various invariants associated to Legendrian submanifolds and their Lagrangian fillings. Most of this material is well established; see [47] and also [45; 57; 56; 39; 49; 50]. Note though that our sign conventions follow [59]; see Remark 9.

2.1 A_∞–algebras and A_∞–coalgebras

In this section we will discuss the basic algebraic objects we use. These are modules over a ground ring $\boldsymbol{k}$ of the following form. Fix a coefficient field $\mathbb{K}$ (of arbitrary characteristic) and let $\boldsymbol{k}$ be a semisimple ring of the form:

$$\boldsymbol{k} = \bigoplus_{v \in \Gamma} \mathbb{K} e_v,$$

where $e_v^2 = e_v$ and $e_v e_w = 0$ for $v \neq w$, and where the index set Γ is finite.

We will use $\mathbb{Z}$–graded $\boldsymbol{k}$–bimodules. If $M = \bigoplus_i M^i$ is such a module then we call M *connected* if $M^0 \cong \boldsymbol{k}$ and either $M^i = 0$ for all $i > 0$, or $M^i = 0$ for all $i < 0$. We call M *simply connected* if, in addition, in the former case $M^{-1} = 0$, and in the latter $M^1 = 0$. Further, we say that M is *locally finite* if each M^i is finitely generated as a $\boldsymbol{k}$–bimodule.

We have the usual shifting and tensor product operations on modules. If $M = \bigoplus_i M^i$ is a graded $\boldsymbol{k}$–bimodule and s is an integer, then we let the corresponding shifted module $M[s] = \bigoplus_i M[s]^i$ be the module with graded components

$$M[s]^i = M^{i+s}.$$

If $N = \bigoplus_i N^i$ is another graded $\boldsymbol{k}$–bimodule, then $M \otimes_{\boldsymbol{k}} N = \bigoplus_k (M \otimes_{\boldsymbol{k}} N)^k$ is naturally a graded $\boldsymbol{k}$–bimodule with

$$(M \otimes_{\boldsymbol{k}} N)^k = \bigoplus_{i+j=k} M^i \otimes_{\boldsymbol{k}} N^j.$$

For iterated tensor products we write

$$M^{\otimes_{\boldsymbol{k}} r} = \underbrace{M \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} M}_{r}.$$

Our modules will often have further structure as $\mathbb{Z}$–graded A_∞–algebras and A_∞–coalgebras over $\boldsymbol{k}$; see Sections 2.1.1 and 2.1.2. The modules are then in particular chain complexes with a differential, and we will use *cohomological* grading throughout; that is, the differential *increases* the grading by 1. For example, if L is a topological space then its cohomology complex $C^*(L)$ is supported in nonnegative grading, while the homology complex $C_{-*}(L)$ is supported in nonpositive degrees. To be consistent with this, we denote the grading as a subscript (resp. superscript) when the underlying chain complex has a coalgebra (resp. algebra) structure.

2.1.1 A_∞–algebras An A_∞*–algebra* over $\boldsymbol{k}$ is a $\mathbb{Z}$–graded $\boldsymbol{k}$–module $\mathscr{A}$ with a collection of grading-preserving $\boldsymbol{k}$–linear maps

$$\mathfrak{m}_i : \mathscr{A}^{\otimes_{\boldsymbol{k}} i} \to \mathscr{A}[2-i]$$

for all integers $i \geq 1$ satisfying the A_∞–relations

$$\sum_{i,j}(-1)^{|a_1|+\cdots+|a_j|-j}\mathfrak{m}_{d-i+1}(a_d,\ldots,a_{j+i+1},\mathfrak{m}_i(a_{j+i},\ldots,a_{j+1}),a_j,\ldots,a_1) = 0 \tag{3}$$

for all d.

Remark 9 We follow the sign conventions of [59]. Even though $\mathfrak{m}_i$ is written on the left of $(a_{j+i},\ldots,a_{j+1})$, the sign convention is so that $\mathfrak{m}_i$ acts from the right. To be consistent, we will insist that all our operators act on the right independently of how they are written. This convention and the usual Koszul sign exchange rule applied with respect to the shifted grading $\mathscr{A}[1]$ determine the signs that appear in our formulas.

A *DG–algebra* over $\boldsymbol{k}$ is an A_∞–algebra $\mathscr{A}$ such that $\mathfrak{m}_i = 0$ for $i \geq 3$. In this case, we call the first two operations the *differential* and the *product*, respectively, and use the following adjustments to obtain an (ordinary) differential graded algebra:

$$da = (-1)^{|a|}\mathfrak{m}_1(a) \quad \text{and} \quad a_2a_1 = (-1)^{|a_1|}\mathfrak{m}_2(a_2,a_1). \tag{4}$$

In particular, the product is then associative and the graded Leibniz rule for d holds:

$$d(a_2a_1) = (da_2)a_1 + (-1)^{|a_2|}a_2(da_1). \tag{5}$$

An *A_∞–map* $\mathfrak{e}\colon \mathscr{A} \to \mathscr{B}$ between A_∞–algebras $\mathscr{A}$ and $\mathscr{B}$ over $\boldsymbol{k}$, with operations $\mathfrak{m}_i$ and $\mathfrak{n}_i$ for $i \geq 1$, respectively, is a collection of $\boldsymbol{k}$–linear grading-preserving maps

$$\mathfrak{e}_i\colon \mathscr{A}^{\otimes_{\boldsymbol{k}} i} \to \mathscr{B}[1-i], \quad i \geq 1,$$

satisfying the relations

$$\begin{aligned}\sum_{i,j}&(-1)^{|a_1|+\cdots+|a_j|-j}\mathfrak{e}_{d-i+1}(a_d,\ldots,a_{j+i+1},\mathfrak{m}_i(a_{j+i},\ldots,a_{j+1}),a_j,\ldots,a_1)\\ &= \sum_{\substack{1\leq j\leq d\\ 0<i_1<\cdots<i_j<d}} \mathfrak{n}_j(\mathfrak{e}_{d-i_j}(a_d,\ldots,a_{d-i_j}),\ldots,\mathfrak{e}_{i_2-i_1}(a_{i_2},\ldots,a_{i_1+1}),\mathfrak{e}_{i_1}(a_{i_1},\ldots,a_1)).\end{aligned}$$

An A_∞–map $\mathfrak{e}\colon \mathscr{A} \to \mathscr{B}$ is called an *A_∞–quasi-isomorphism* if the map on cohomology $H^*(\mathscr{A}) \to H^*(\mathscr{B})$ induced by $\mathfrak{e}^1$ is an isomorphism.

We say that an A_∞–algebra $\mathscr{A}$ is *strictly unital* if there is an element $1_{\mathscr{A}} \in \mathscr{A}$ such that $\mathfrak{m}_1(1_{\mathscr{A}}) = 0$, $\mathfrak{m}_2(1_{\mathscr{A}},a) = \mathfrak{m}_2(a,1_{\mathscr{A}}) = a$ for any $a \in \mathscr{A}$, and $\mathfrak{m}_i$ for $i > 2$ annihilates any monomial containing $1_{\mathscr{A}}$ as a factor. Any A_∞–algebra $\mathscr{A}$ which has a cohomological unit, ie a cocycle representing the identity element in $H^*(\mathscr{A})$, is quasi-isomorphic to a strictly unital A_∞–algebra [56, Section 7.2].

An *augmentation* of a strictly unital A_∞–algebra is an A_∞–map $\epsilon\colon \mathscr{A} \to \boldsymbol{k}$, where $\boldsymbol{k}$ is considered as a strictly unital A_∞–algebra in degree 0 with trivial differential and higher A_∞–products, and is such that $\epsilon_1(1_{\mathscr{A}}) = 1_{\boldsymbol{k}}$ and ϵ_i for $i > 1$ annihilates any monomial containing $1_{\mathscr{A}}$. An augmentation is called *strict* if $\epsilon_i = 0$ for $i > 1$. The category of augmented, strictly unital A_∞–algebras is equivalent to the category of strictly augmented, strictly unital A_∞–algebras; see [56, Section 7.2].

2.1.2 A_∞–coalgebras An A_∞*–coalgebra* $\mathscr{C}$ over $\boldsymbol{k}$ is a $\mathbb{Z}$–graded $\boldsymbol{k}$–module with a collection of $\boldsymbol{k}$–linear grading-preserving maps

$$\Delta_i\colon \mathscr{C} \to \mathscr{C}^{\otimes_{\boldsymbol{k}} i}[2-i]$$

for all integers $i \geq 1$, with the following properties. The maps satisfy the co-A_∞–relations

$$\sum_{i=1}^{d}\sum_{j=0}^{d-i}(\mathbf{1}^{\otimes_{\boldsymbol{k}}(d-i-j)} \otimes_{\boldsymbol{k}} \Delta_i \otimes_{\boldsymbol{k}} \mathbf{1}^{\otimes_{\boldsymbol{k}} j})\Delta_{d-i+1} = 0, \tag{6}$$

where

$$\begin{aligned}\mathbf{1}^{\otimes_{\boldsymbol{k}}(d-i-j)} \otimes_{\boldsymbol{k}} \Delta_i \otimes_{\boldsymbol{k}} \mathbf{1}^{\otimes_{\boldsymbol{k}} j}\,&(c_{d-i+1}, \dots, c_1)\\ &= (-1)^{|c_1|+\cdots+|c_j|-j}(c_{d-i+1}, \dots, c_{j+2}) \otimes_{\boldsymbol{k}} \Delta_i(c_{j+1}) \otimes_{\boldsymbol{k}} (c_j, \dots, c_1)\\ &\qquad\qquad \in \mathscr{C}^{\otimes_{\boldsymbol{k}}(d-i-j)} \otimes_{\boldsymbol{k}} \mathscr{C}^{\otimes_{\boldsymbol{k}} i} \otimes_{\boldsymbol{k}} \mathscr{C}^{\otimes_{\boldsymbol{k}} j}.\end{aligned}$$

Furthermore, the degree 1 map

$$\mathscr{C}[-1] \to \prod_{i=1}^{\infty} \mathscr{C}[-1]^{\otimes_{\boldsymbol{k}} i},$$

with i^{th} component equal to Δ_i, factorizes through the natural inclusion

$$\bigoplus_{i=1}^{\infty} \mathscr{C}[-1]^{\otimes_{\boldsymbol{k}} i} \to \prod_{i=1}^{\infty} \mathscr{C}[-1]^{\otimes_{\boldsymbol{k}} i}$$

of the direct sum into the direct product.

A *DG–coalgebra* over $\boldsymbol{k}$ is an A_∞–coalgebra such that $\Delta_i = 0$ for $i \geq 3$. In this case, we call the first two operations the *differential* and the *coproduct*, respectively, and use the following adjustments to obtain an (ordinary) differential graded coalgebra:

$$\theta c = (-1)^{|c|}\Delta_1(c) \quad\text{and}\quad \Delta(c) = \sum (-1)^{|c_{(2)}|} c_{(1)} \otimes_{\boldsymbol{k}} c_{(2)}, \tag{7}$$

where we write $\Delta_2(c) = \sum c_{(1)} \otimes_{\boldsymbol{k}} c_{(2)}$.

In particular, the coproduct is coassociative, ie $(\Delta \otimes_{\boldsymbol{k}} \mathbf{1}) \circ \Delta = (\mathbf{1} \otimes_{\boldsymbol{k}} \Delta) \circ \Delta$, and the graded co-Leibniz rule holds:

$$\Delta\theta(c) = \sum (-1)^{|c_{(1)}|} c_{(1)} \otimes_{\boldsymbol{k}} \theta(c_{(2)}) + \theta(c_{(1)}) \otimes_{\boldsymbol{k}} c_{(2)}. \tag{8}$$

An A_∞*–comap* $\mathfrak{f}\colon \mathscr{C} \to \mathscr{D}$ between A_∞–coalgebras $\mathscr{C}$ and $\mathscr{D}$ over $\boldsymbol{k}$, with operations Δ_i and Θ_i for $i \geq 1$, respectively, is a collection of $\boldsymbol{k}$–linear grading-preserving maps

$$\mathfrak{f}_i : \mathscr{C} \to \mathscr{D}^{\otimes_{\boldsymbol{k}} i}[1-i], \quad i \geq 1,$$

satisfying the relations

$$\sum_{i=1}^{d} \sum_{j=0}^{d-i} (\mathbf{1}^{\otimes_{\boldsymbol{k}} (d-i-j)} \otimes_{\boldsymbol{k}} \Theta_i \otimes_{\boldsymbol{k}} \mathbf{1}^{\otimes_{\boldsymbol{k}} j}) \mathfrak{f}_{d-i+1} = \sum_{\substack{1 \leq j \leq d \\ 0 < i_1 < i_2 < \cdots < i_j < d}} (\mathfrak{f}_{d-i_j} \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} \mathfrak{f}_{i_2 - i_1} \otimes_{\boldsymbol{k}} \mathfrak{f}_{i_1}) \Delta_j,$$

where

$$\begin{aligned} &\mathbf{1}^{\otimes_{\boldsymbol{k}} (d-i-j)} \otimes_{\boldsymbol{k}} \Theta_i \otimes_{\boldsymbol{k}} \mathbf{1}^{\otimes_{\boldsymbol{k}} j} \, (d_{d-i+1}, \dots, d_1) \\ &\quad = (-1)^{|d_1| + \cdots + |d_j| - j} (d_{d-i+1}, \dots, d_{j+2}) \otimes_{\boldsymbol{k}} \Theta_i(d_{j+1}) \otimes_{\boldsymbol{k}} (d_j, \dots, d_1) \\ &\qquad \in \mathscr{D}^{\otimes_{\boldsymbol{k}} (d-i-j)} \otimes_{\boldsymbol{k}} \mathscr{D}^{\otimes_{\boldsymbol{k}} i} \otimes_{\boldsymbol{k}} \mathscr{D}^{\otimes_{\boldsymbol{k}} j}. \end{aligned}$$

Furthermore, the degree 0 map

$$\mathscr{C}[-1] \to \prod_{i=1}^{\infty} \mathscr{D}[-1]^{\otimes_{\boldsymbol{k}} i},$$

with i^{th} component equal to $\mathfrak{f}_i$, factorizes through the natural inclusion

$$\bigoplus_{i=1}^{\infty} \mathscr{D}[-1]^{\otimes_{\boldsymbol{k}} i} \to \prod_{i=1}^{\infty} \mathscr{D}[-1]^{\otimes_{\boldsymbol{k}} i} \tag{9}$$

of the direct sum into the direct product.

An A_∞–comap $\mathfrak{f}\colon \mathscr{C} \to \mathscr{D}$ is called an A_∞*–quasi-isomorphism* if the map on cohomology $H^*(\mathscr{C}) \to H^*(\mathscr{D})$ induced by $\mathfrak{f}^1$ is an isomorphism.

We say that an A_∞–coalgebra is *strictly counital* if there exists a $\boldsymbol{k}$–linear map $\epsilon : \mathscr{C} \to \boldsymbol{k}$ such that $(\epsilon \otimes \mathbf{1})\Delta_2 = (\mathbf{1} \otimes \epsilon)\Delta_2 = \mathbf{1}$ and $(\mathbf{1}^{\otimes_{\boldsymbol{k}} (i-j)} \otimes_{\boldsymbol{k}} \epsilon \otimes_{\boldsymbol{k}} \mathbf{1}^{\otimes_{\boldsymbol{k}} j-1})\Delta_i = 0$ for all $i \neq 2$ and j. Any A_∞–coalgebra $\mathscr{C}$ which has a cohomological counit, ie a cocycle representing the counit in $H^*(\mathscr{C})$, is quasi-isomorphic to a strictly counital A_∞–coalgebra; see [56, Section 7.5].

A *coaugmentation* of a strictly counital A_∞–coalgebra $\mathscr{C}$ is an A_∞–comap $\eta\colon \boldsymbol{k} \to \mathscr{C}$, where $\boldsymbol{k}$ is considered as a vector space in degree 0 with the trivial A_∞–coalgebra structure, and is such that $\epsilon\eta_1 = 1_{\boldsymbol{k}}$ and $(\mathbf{1}^{\otimes_{\boldsymbol{k}}(i-j)} \otimes_{\boldsymbol{k}} \epsilon \otimes_{\boldsymbol{k}} \mathbf{1}^{\otimes_{\boldsymbol{k}} j-1})\eta_i = 0$ for all $i > 1$ and j. The coaugmentation is called *strict* if $\eta_i = 0$ for $i \geq 2$.

A DG–coalgebra $\mathscr{C}$ is called *conilpotent* (also called *cocomplete*) if for any $c \in \mathscr{C}$, there exists an $n \geq 2$ such that c is in the kernel of the iterated comultiplication map defined recursively by $\Delta^{(2)} = \Delta$, and $\Delta^{(n)} = (\mathbf{1}^{\otimes_{\boldsymbol{k}}(n-2)} \otimes_{\boldsymbol{k}} \Delta) \circ \Delta^{(n-1)}$ for $n > 2$. When considering coaugmented DG–coalgebras, conilpotency is enforced only on the coaugmentation ideal $\operatorname{coker}(\eta)$.

2.1.3 Graded dual We next discuss the graded dual of a graded $\boldsymbol{k}$–module. Since we are working with bimodules over the ring $\boldsymbol{k}$, there are two $\boldsymbol{k}$–linear duals [8].

If $\mathscr{A}$ is a graded $\boldsymbol{k}$–bimodule, $\mathscr{A} = \bigoplus_i \mathscr{A}_i$, then the *graded duals* $\mathscr{A}^\# = \bigoplus_i (\mathscr{A}^\#)_i$ and ${}^\#\mathscr{A} = \bigoplus_i ({}^\#\mathscr{A})_i$ are defined as follows. The graded components $(\mathscr{A}^\#)_i$ of $\mathscr{A}^\#$ are left $\boldsymbol{k}$–module maps

$$\hom_{\boldsymbol{k}-}(\mathscr{A}_{-i}, \boldsymbol{k}),$$

and the $\boldsymbol{k}$–bimodule structure on $\mathscr{A}^\#$ is given as follows: if $e_v, e_w \in \boldsymbol{k}$, $a \in (\mathscr{A}^\#)_i$ and $c \in \mathscr{A}_{-i}$, then

$$(e_v \cdot a \cdot e_w)(c) = a(ce_v)e_w. \tag{10}$$

The graded components $({}^\#\mathscr{A})_i$ of ${}^\#\mathscr{A}$ in degree i are right $\boldsymbol{k}$–module maps, which we write as

$$\hom_{-\boldsymbol{k}}(\mathscr{A}_{-i}, \boldsymbol{k}),$$

and the $\boldsymbol{k}$–bimodule structure is given by: if $e_v, e_w \in \boldsymbol{k}$, $a \in ({}^\#\mathscr{A})_i$ and $c \in \mathscr{A}_{-i}$, then

$$(e_v \cdot a \cdot e_w)(c) = e_v a(e_w c). \tag{11}$$

Both canonical maps $\mathscr{A} \to {}^\#(\mathscr{A}^\#)$ and $\mathscr{A} \to ({}^\#\mathscr{A})^\#$ are $\boldsymbol{k}$–bimodule maps, which are isomorphisms if $\mathscr{A}$ is locally finite.

If $V_1, V_2, \dots, V_n$ are $\boldsymbol{k}$–bimodules, there is a natural map

$$V_n^\# \otimes_{\boldsymbol{k}} V_{n-1}^\# \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} V_1^\# \to (V_1 \otimes_{\boldsymbol{k}} V_2 \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} V_n)^\#$$

given by

$$(a_n \otimes a_{n-1} \otimes \cdots \otimes a_1)(c_1 \otimes c_2 \otimes \cdots \otimes c_n) := a_1(c_1 a_2(c_2 \cdots a_n(c_n) \cdots)). \tag{12}$$

Similarly, there is a natural map

$$^{\#}V_n \otimes_{\boldsymbol{k}} {}^{\#}V_{n-1} \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} {}^{\#}V_1 \to {}^{\#}(V_1 \otimes_{\boldsymbol{k}} V_2 \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} V_n)$$

given by

$$(a_n \otimes a_{n-1} \otimes \cdots \otimes a_1)(c_1 \otimes c_2 \otimes \cdots \otimes c_n) := a_n(\cdots a_2(a_1(c_1)c_2)\cdots c_n). \tag{13}$$

These give the graded duals $\mathscr{C}^{\#}$ and ${}^{\#}\mathscr{C}$ of a coaugmented A_∞–coalgebra $\mathscr{C}$ the structure of augmented A_∞–algebras, with structure maps defined by

$$\mathfrak{m}_i(a_i, \dots, a_1)(c) := (-1)^{|c|}(a_i \otimes \cdots \otimes a_1)\Delta_i(c). \tag{14}$$

Note that to get a nonzero product, we must have $|\mathfrak{m}_i(a_i, \dots, a_1)| = |c|$, hence the sign $(-1)^{|c|}$ equals the sign $(-1)^{|a_1|+\cdots+|a_i|-i}$.

In general, there is no natural way of equipping the graded dual of an augmented A_∞–algebra with an A_∞–coalgebra structure. However, if the grading on $\mathscr{A}$ is locally finite (ie $\mathscr{A}_i$ are finitely generated as $\boldsymbol{k}$–bimodules), it follows that

$$\mathscr{A}^{\#} \otimes_{\boldsymbol{k}} \mathscr{A}^{\#} \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} \mathscr{A}^{\#} \cong (\mathscr{A} \otimes_{\boldsymbol{k}} \mathscr{A} \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} \mathscr{A})^{\#},$$
$$^{\#}\mathscr{A} \otimes_{\boldsymbol{k}} {}^{\#}\mathscr{A} \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} {}^{\#}\mathscr{A} \cong {}^{\#}(\mathscr{A} \otimes_{\boldsymbol{k}} \mathscr{A} \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} \mathscr{A}).$$

Using these isomorphisms, the graded duals $\mathscr{A}^{\#}$ and ${}^{\#}\mathscr{A}$ of an augmented A_∞–algebra $\mathscr{A}$ with locally finite grading can be naturally equipped with the structure of a coaugmented A_∞–coalgebra by using the formulas

$$\Delta_i(c)(a_i \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} a_1) = (-1)^{|c|}c(\mathfrak{m}_i(a_i, \dots, a_1)).$$

2.1.4 Twisting cochains Let $(\mathscr{C}, \Delta_\bullet)$ be an A_∞–coalgebra and let $(\mathscr{A}, \mathfrak{m}_1, \mathfrak{m}_2)$ be a DG–algebra. A *twisting cochain* is a $\boldsymbol{k}$–linear map $\mathfrak{t}\colon \mathscr{C} \to \mathscr{A}$ of degree 1 that satisfies

$$\mathfrak{m}_1 \circ \mathfrak{t} - \mathfrak{t} \circ \Delta_1 + \sum_{d \geq 2} (-1)^d \mathfrak{m}_2^{(d)} \circ \mathfrak{t}^{\otimes_{\boldsymbol{k}} d} \circ \Delta_d = 0, \tag{15}$$

where $\mathfrak{m}_2^{(2)} := \mathfrak{m}_2$ and $\mathfrak{m}_2^{(d)} := \mathfrak{m}_2 \circ (\mathrm{Id}_{\mathscr{A}} \otimes_{\boldsymbol{k}} \mathfrak{m}_2^{(d-1)})$. For $c \in \mathscr{C}$, note that $\Delta_i(c) \neq 0$ for only finitely many i, and hence the potentially infinite sum in (15) is actually finite when it acts on c.

If the coalgebra $\mathscr{C}$ is coaugmented by $\eta\colon \boldsymbol{k} \to \mathscr{C}$ and the algebra $\mathscr{A}$ is augmented $\epsilon\colon \mathscr{A} \to \boldsymbol{k}$, we require in addition that its twisting cochains $\mathfrak{t}$ are compatible in the sense that

$$\mathfrak{t} \circ \eta = \epsilon \circ \mathfrak{t} = 0. \tag{16}$$

We denote the set of twisting cochains from $\mathscr{C}$ to $\mathscr{A}$ by $\mathrm{Tw}(\mathscr{C}, \mathscr{A})$.

Let $\mathfrak{t} \in \mathrm{Tw}(\mathscr{C}, \mathscr{A})$ be a twisting cochain. Consider the twisted tensor product $\mathscr{A} \otimes_{\boldsymbol{k}}^{\mathfrak{t}} \mathscr{C}$ as a chain complex with differential $d^{\mathfrak{t}} \colon \mathscr{A} \otimes_{\boldsymbol{k}}^{\mathfrak{t}} \mathscr{C} \to \mathscr{A} \otimes_{\boldsymbol{k}}^{\mathfrak{t}} \mathscr{C}$ defined by

$$d^{\mathfrak{t}} = \mathfrak{m}_1 \otimes_{\boldsymbol{k}} \mathrm{Id}_{\mathscr{C}} + \mathrm{Id}_{\mathscr{A}} \otimes_{\boldsymbol{k}} \Delta_1 + \sum_{d \geq 2} (\mathfrak{m}_2^{(d)} \otimes \mathrm{Id}_{\mathscr{C}}) \circ (\mathrm{Id}_{\mathscr{A}} \otimes_{\boldsymbol{k}} \mathfrak{t}^{\otimes_{\boldsymbol{k}} d-1} \otimes_{\boldsymbol{k}} \mathrm{Id}_{\mathscr{C}}) \circ (\mathrm{Id}_{\mathscr{A}} \otimes_{\boldsymbol{k}} \Delta_d). \tag{17}$$

Here the differential squares to zero, $d^{\mathfrak{t}} \circ d^{\mathfrak{t}} = 0$, since $\mathfrak{t}$ satisfies (15). This complex is the *Koszul complex* associated with $\mathfrak{t}$. It is called acyclic if the projection to $\boldsymbol{k}$ is a quasi-isomorphism.

One also has an analogous complex of the form $\mathscr{C} \otimes_{\boldsymbol{k}}^{\mathfrak{t}} \mathscr{A}$.

The $\mathbb{K}$–vector space of $\boldsymbol{k}$–bimodule morphisms $\hom_{\boldsymbol{k}\text{-}\boldsymbol{k}}(\mathscr{C}, \mathscr{A})$ carries an A_∞–algebra structure with operations $\mathfrak{n}_d$ for $d \geq 1$, given by

$$\mathfrak{n}_1(t) = \mathfrak{m}_1 \circ t + (-1)^{|t|} t \circ \Delta_1,$$
$$\mathfrak{n}_d(t_d, t_{d-1}, \dots, t_1) = (-1)^{d(|t_d|+\cdots+|t_1|)} \mathfrak{m}_2^{(d)} \circ (t_d \otimes t_{d-1} \otimes \cdots \otimes t_1) \circ \Delta_d \quad \text{for } d \geq 2,$$

where the composition $(t_d \otimes t_{d-1} \otimes \cdots \otimes t_1) \circ \Delta_d$ is defined componentwise. Thus, if $\Delta_d(c) = c_d \otimes \cdots \otimes c_1$, then

$$(t_d \otimes_{\boldsymbol{k}} t_{d-1} \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} t_1) \Delta_d(c) = (-1)^{\dagger} \mathfrak{t}_d(c_d) \otimes_{\boldsymbol{k}} \mathfrak{t}_{d-1}(c_{d-1}) \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} t_1(c_1),$$

where $\dagger = \sum_{j=2}^{d} \sum_{i=1}^{j-1} |c_i||t_j|$. In this setting, a twisting cochain $\mathfrak{t} \colon \mathscr{C} \to \mathscr{A}$ corresponds to a solution of the Maurer–Cartan equation

$$\sum_{i \geq 1} \mathfrak{n}_i(\mathfrak{t}, \mathfrak{t}, \dots, \mathfrak{t}) = 0. \tag{18}$$

(As before, this sum is effectively finite since, for any $c \in \mathscr{C}$, $\Delta_i(c) \neq 0$ only for finitely many i.)

A twisting cochain $\mathfrak{t} \colon \mathscr{C} \to \mathscr{A}$ defines a twisted A_∞–structure on $\hom_{\boldsymbol{k}}(\mathscr{C}, \mathscr{A})$, with operations $\mathfrak{n}_d^{\mathfrak{t}}$ given by

$$\mathfrak{n}_d^{\mathfrak{t}}(t_d, t_{d-1}, \dots, t_1) = \sum_{l_i \geq 0} \mathfrak{n}_{d+l_0+l_1+\cdots+l_d}(\overbrace{\mathfrak{t}, \dots, \mathfrak{t}}^{l_d}, t_d, \overbrace{\mathfrak{t}, \dots, \mathfrak{t}}^{l_{d-1}}, t_{d-1}, \dots, t_1, \overbrace{\mathfrak{t}, \dots, \mathfrak{t}}^{l_0}).$$

We will denote this twisted A_∞–structure by $\hom_{\boldsymbol{k}}^{\mathfrak{t}}(\mathscr{C}, \mathscr{A})$.

There are direct analogues of the above construction if we instead consider a DG–coalgebra $(\mathscr{C}, \Delta_1, \Delta_2)$ and an A_∞–algebra $\mathscr{A}$ with operations $\mathfrak{m}_i$. The module

$\hom_{\boldsymbol{k}\text{-}\boldsymbol{k}}(\mathscr{C},\mathscr{A})$ has the structure of an A_∞–algebra with operations $\mathfrak{n}_d$ given by

$$\mathfrak{n}_1(t) = \mathfrak{m}_1 \circ t + (-1)^{|t|} t \circ \Delta_1,$$
$$\mathfrak{n}_d(t_d, t_{d-1}, \dots, t_1) = \mathfrak{m}_d \circ (t_d \otimes t_{d-1} \otimes \cdots \otimes t_1) \circ \Delta_2^{(d)} \quad \text{for } d \geq 2.$$

To make sense of the twisting cochain (18), one needs to make additional assumptions to ensure the convergence of the infinite sum. This holds, for example, if $\mathscr{C}$ is conilpotent.

We remark that if both $\mathscr{C}$ and $\mathscr{A}$ are A_∞–(co)algebras, then defining a twisting cochain is a more complicated matter; cf [57, Introduction]. We will not need this here.

2.2 Bar–cobar duality for A_∞–(co)algebras

In this section we first introduce the bar and cobar constructions and then discuss basic relations between them.

2.2.1 Bar and cobar constructions Let $(\mathscr{A}, \{\mathfrak{m}_j\}_{j\geq 1})$ be a strictly unital A_∞–algebra with a strict augmentation $\epsilon\colon \mathscr{A} \to \boldsymbol{k}$. Define the *augmentation ideal* $\overline{\mathscr{A}} = \ker(\epsilon)$. If we are given a nonunital A_∞–algebra $\overline{\mathscr{A}}$, we can turn it into a strictly unital A_∞–algebra $\mathscr{A} := \boldsymbol{k} \oplus \overline{\mathscr{A}}$ with an augmentation given by projection to $\boldsymbol{k}$.

We next recall the construction of the (reduced) *bar construction* $\mathrm{B}\mathscr{A}$. For any augmented A_∞–algebra $\mathscr{A}$, $\mathrm{B}\mathscr{A}$ is a coaugmented conilpotent DG–coalgebra. As a coaugmented coalgebra, $\mathrm{B}\mathscr{A}$ is defined as

$$\mathrm{B}\mathscr{A} = \boldsymbol{k} \oplus \overline{\mathscr{A}}[1] \oplus \overline{\mathscr{A}}[1]^{\otimes_{\boldsymbol{k}} 2} \oplus \cdots,$$

where [1] denotes the downwards shift by 1. We write a typical monomial using Eilenberg and Mac Lane's notation

$$[a_d|a_{d-1}|\cdots|a_1] = sa_d \otimes_{\boldsymbol{k}} sa_{d-1} \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} sa_1,$$

where for $a \in \overline{\mathscr{A}}$, $sa \in \overline{\mathscr{A}}[1]$ denotes the corresponding element in $\overline{\mathscr{A}}[1]$ with degree shifted down by 1.

The differential $b\colon \mathrm{B}\mathscr{A} \to \mathrm{B}\mathscr{A}$ is defined to vanish on $\boldsymbol{k} \subset \mathscr{A}$, so $b_{|\boldsymbol{k}} = 0$, and defined on monomials by

$$b([a_d|a_{d-1}|\cdots|a_1]) = \sum_{i,j} (-1)^{|a_1|+\cdots+|a_j|-j} [a_d|\cdots|a_{j+i+1}|\mathfrak{m}_i(a_{j+i},\dots,a_{j+1})|a_j|\cdots|a_1].$$

The coproduct $\Delta_2\colon \mathrm{B}\mathscr{A} \to \mathrm{B}\mathscr{A} \otimes_{\boldsymbol{k}} \mathrm{B}\mathscr{A}$ is defined by

$$\Delta_2([a_d|a_{d-1}|\cdots|a_1]) = \sum_{i=0}^{d} (-1)^{|a_i|+\cdots+|a_1|-i}[a_d|a_{d-1}|\cdots|a_{i+1}] \otimes_{\boldsymbol{k}} [a_i|a_{i-1}|\cdots|a_1].$$

The slightly unusual sign $(-1)^{|a_i|+\cdots+|a_1|-i}$ appears as a consequence of the following two facts:

(i) The equation $b^2 = 0$ is equivalent to the A_∞–relations (3) for $(\mathfrak{m}_i)_{i\geq 1}$.

(ii) The pair (b, Δ_2) satisfies the co-A_∞–relations (6).

Redefining (b, Δ_2) to (θ, Δ) using (7) removes the sign in Δ_2, and (θ, Δ) becomes a (usual) coassociative DG–coalgebra, where the co-Leibniz rule (8) holds. The coaugmentation $\eta\colon \boldsymbol{k} \to \mathrm{B}\mathscr{A}$ is defined by letting η_1 be the inclusion of $\boldsymbol{k}$ and $\eta_i = 0$ for $i > 0$.

There is an increasing, exhaustive and bounded below (hence, complete Hausdorff) filtration on the complex $\mathrm{B}\mathscr{A}$,

$$\boldsymbol{k} = \mathcal{F}^0\mathrm{B}\mathscr{A} \subset \mathcal{F}^1\mathrm{B}\mathscr{A} \subset \cdots \subset \mathrm{B}\mathscr{A}, \quad \text{where } \mathcal{F}^p\mathrm{B}\mathscr{A} := \boldsymbol{k} \oplus \overline{\mathscr{A}}[1] \oplus \cdots \oplus \overline{\mathscr{A}}[1]^{\otimes_{\boldsymbol{k}} p}.$$

This induces the *word-length spectral sequence* with

$$E_1^{p,q} = H^{p+q}(\mathcal{F}^p\mathrm{B}\mathscr{A}/\mathcal{F}^{p-1}\mathrm{B}\mathscr{A})$$

converging strongly to

$$E_\infty^{p,q} = \mathcal{F}^p H^{p+q}(\mathrm{B}\mathscr{A})/\mathcal{F}^{p-1}H^{p+q}(\mathrm{B}\mathscr{A})$$

by the classical convergence theorem [64, Theorem 5.5.1]. It can be proved using this spectral sequence that if an A_∞–map $\mathfrak{e}\colon \mathscr{A} \to \mathscr{B}$ is a quasi-isomorphism, then the naturally induced DG–coalgebra map $\mathrm{B}\mathfrak{e}\colon \mathrm{B}\mathscr{A} \to \mathrm{B}\mathscr{B}$ is a quasi-isomorphism; see [49, Proposition 2.2.3].

There is a universal twisting cochain $\mathfrak{t}_{\mathscr{A}}\colon \mathrm{B}\mathscr{A} \to \mathscr{A}$ which is nonzero only on $\overline{\mathscr{A}}[1] \subset \mathrm{B}\mathscr{A}$ and is given by the inclusion map $\overline{\mathscr{A}}[1] \to \mathscr{A}$. The twisting cochain $\mathfrak{t}_{\mathscr{A}}$ gives rise to a free $\mathscr{A}$–bimodule resolution of $\mathscr{A}$ obtained as a twisted tensor product

$$\mathscr{A} \otimes_{\boldsymbol{k}}^{\mathfrak{t}_{\mathscr{A}}} \mathrm{B}\mathscr{A} \otimes_{\boldsymbol{k}}^{\mathfrak{t}_{\mathscr{A}}} \mathscr{A},$$

with the differential d given by the formula

$$\begin{aligned}(19)\quad d = {}& \mathfrak{m}_1 \otimes_{\boldsymbol{k}} \mathrm{Id}_{\mathrm{B}\mathscr{A}} \otimes_{\boldsymbol{k}} \mathrm{Id}_{\mathscr{A}} + \mathrm{Id}_{\mathscr{A}} \otimes_{\boldsymbol{k}} b \otimes_{\boldsymbol{k}} \mathrm{Id}_{\mathscr{A}} + \mathrm{Id}_{\mathscr{A}} \otimes_{\boldsymbol{k}} \mathrm{Id}_{\mathrm{B}\mathscr{A}} \otimes_{\boldsymbol{k}} \mathfrak{m}_1 \\ & + \Big(\sum_{d\geq 2} (\mathfrak{m}_d \otimes_{\boldsymbol{k}} \mathrm{Id}_{\mathrm{B}\mathscr{A}}) \circ (\mathrm{Id}_{\mathscr{A}} \otimes_{\boldsymbol{k}} \mathfrak{t}^{\otimes_{\boldsymbol{k}} d-1} \otimes_{\boldsymbol{k}} \mathrm{Id}_{\mathrm{B}\mathscr{A}}) \circ (\mathrm{Id}_{\mathscr{A}} \otimes_{\boldsymbol{k}} \Delta_2^{(d)})\Big) \otimes_k \mathrm{Id}_{\mathscr{A}} \\ & + \mathrm{Id}_{\mathscr{A}} \otimes_k \Big(\sum_{d\geq 2} (\mathrm{Id}_{\mathrm{B}\mathscr{A}} \otimes_{\boldsymbol{k}} \mathfrak{m}_d) \circ (\mathrm{Id}_{\mathrm{B}\mathscr{A}} \otimes_{\boldsymbol{k}} \mathfrak{t}^{\otimes_{\boldsymbol{k}} d-1} \otimes_{\boldsymbol{k}} \mathrm{Id}_{\mathscr{A}}) \circ (\Delta_2^{(d)} \otimes_{\boldsymbol{k}} \mathrm{Id}_{\mathscr{A}})\Big).\end{aligned}$$

This can be used to compute Hochschild homology and cohomology of $\mathscr{A}$ with coefficients in an $\mathscr{A}$–bimodule $\mathscr{M}$.

Consider instead a strictly counital A_∞–coalgebra $\mathscr{C}$ with operations Δ_i and with a strict coaugmentation $\eta\colon \boldsymbol{k} \to \mathscr{C}$. Let $\overline{\mathscr{C}} = \operatorname{coker}(\eta)$ be the coaugmentation ideal. We next recall the *cobar construction*, which associates a DG–algebra $\Omega\mathscr{C}$ to $\mathscr{C}$. As an augmented algebra, $\Omega\mathscr{C}$ is

$$\Omega\mathscr{C} = \boldsymbol{k} \oplus \overline{\mathscr{C}}[-1] \oplus \overline{\mathscr{C}}[-1]^{\otimes_{\boldsymbol{k}} 2} \oplus \cdots . \tag{20}$$

As before, we write a typical monomial as

$$[c_d|c_{d-1}|\cdots|c_1] = s^{-1}c_d \otimes_{\boldsymbol{k}} s^{-1}c_{d-1} \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} s^{-1}c_1,$$

where for $c \in \overline{\mathscr{C}}$, $s^{-1}c \in \overline{\mathscr{C}}[-1]$ denotes the corresponding element in $\overline{\mathscr{C}}[-1]$ with degree shifted up by 1. The differential $\mathfrak{m}_1$ on $\Omega\mathscr{C}$ vanishes on $\boldsymbol{k}$, so $\mathfrak{m}_1|_{\boldsymbol{k}} = 0$, and acts on monomials as

$$\mathfrak{m}_1([c_m|\cdots|c_1]) = \sum_{i,j} (-1)^{|c_1|+\cdots+|c_i|-i}[c_m|\cdots|c_{i+2}|\Delta_j(c_{i+1})|c_i|\cdots|c_1].$$

Here, by abuse of notation, we write Δ_j for the induced coproduct $\overline{\mathscr{C}}[-1] \to \overline{\mathscr{C}}[-1]^{\otimes j}$.

The product $\mathfrak{m}_2\colon \Omega\mathscr{C} \otimes \Omega\mathscr{C} \to \Omega\mathscr{C}$ is given by

$$\mathfrak{m}_2([c_m|\cdots|c_{i+1}],[c_i|\cdots|c_1]) = (-1)^{|c_1|+\cdots|c_i|-i}[c_m|\cdots|c_{i+1}|c_i|\cdots|c_1].$$

The slightly unusual sign $(-1)^{|c_1|+\cdots+|c_i|-i}$ appears as a consequence of the following two facts:

(i) The equation $\mathfrak{m}_1^2 = 0$ is equivalent to co-A_∞–relations (6) for $(\Delta_j)_{j\geq 1}$.

(ii) The pair $(\mathfrak{m}_1,\mathfrak{m}_2)$ satisfies the A_∞–relations (3).

Redefining $(\mathfrak{m}_1,\mathfrak{m}_2)$ to $(d,\cdot)$ using (4) removes the sign in $\mathfrak{m}_2$, and $(d,\cdot)$ becomes a (usual) associative DG–algebra, where the Leibniz rule (5) holds. The augmentation $\epsilon\colon \Omega\mathscr{C} \to \boldsymbol{k}$ is given by letting ϵ_1 be the projection to $\boldsymbol{k}$ and $\epsilon_i = 0$ for $i > 0$.

There is a decreasing, exhaustive, bounded above filtration on the complex $\Omega\mathscr{C}$,

$$\Omega\mathscr{C} = \mathcal{F}^0\Omega\mathscr{C} \supset \mathcal{F}^1\Omega\mathscr{C} \supset \cdots ,$$

given by

$$\mathcal{F}^p\Omega\mathscr{C} := \overline{\mathscr{C}}[-1]^{\otimes_{\boldsymbol{k}} p} \oplus \overline{\mathscr{C}}[-1]^{\otimes_{\boldsymbol{k}} (p+1)} \oplus \cdots .$$

This gives the *word-length spectral sequence* with

$$E_1^{p,q} = H^{p+q}(\mathcal{F}^p\Omega\mathscr{C}/\mathcal{F}^{p+1}\Omega\mathscr{C}).$$

Unlike the case of the word-length filtration on the bar construction, for the cobar construction, in general, convergence may fail. Thus, we introduce completions. We define the *completed cobar construction* to be

$$\widehat{\Omega}\mathscr{C} = \varprojlim_{s} (\Omega\mathscr{C})/(\mathcal{F}^s\Omega\mathscr{C}).$$

The length filtration on $\Omega\mathscr{C}$ induces a filtration $\widehat{\mathcal{F}}$ on $\widehat{\Omega}\mathscr{C}$ defined by

$$\widehat{\mathcal{F}}^p\widehat{\Omega}\mathscr{C} = \varprojlim_{s} (\mathcal{F}^p\Omega\mathscr{C})/(\mathcal{F}^s\Omega\mathscr{C}),$$

which is decreasing, exhaustive, bounded above and complete Hausdorff. The spectral sequence associated to the filtration $\widehat{\mathcal{F}}$ on $\widehat{\Omega}\mathscr{C}$ is isomorphic to the length spectral sequence associated with the filtration $\widehat{\mathcal{F}}$ on $\Omega\mathscr{C}$ and converges conditionally to $H^*(\widehat{\Omega}\mathscr{C})$; see [9, Theorem 9.2]. It converges strongly to $H^*(\widehat{\Omega}\mathscr{C})$ if the spectral sequence is regular, ie only finitely many of the differentials $d_r^{p,q}$ are nonzero for each p and q; see [9, Theorem 7.1]. This holds, for example, if $\Omega\mathscr{C}$ is locally finite.

We say that $\Omega\mathscr{C}$ is complete if the natural map $\Omega\mathscr{C} \to \widehat{\Omega}\mathscr{C}$ is a quasi-isomorphism. For example, it is easy to see that this is the case if $\mathscr{C}$ is locally finite and simply connected.

If $\mathfrak{f}\colon \mathscr{C} \to \mathscr{D}$ is an A_∞–comap which is a quasi-isomorphism of A_∞–coalgebras, and if $\Omega\mathscr{C}$ and $\mathscr{D}$ are complete, then $\Omega\mathfrak{f}$ is a quasi-isomorphism. (This follows from [21, Theorem 7.4]; see also [64, Theorem 5.5.11].) The completeness assumptions are necessary and are related to the completeness of the word-length filtration. A counterexample when the completeness assumptions are dropped can be found in [49, Section 2.4.1].

There is a universal twisting cochain $\mathfrak{t}^{\mathscr{C}}\colon \mathscr{C} \to \Omega\mathscr{C}$ given by the composition of canonical projection $\mathscr{C} \to \overline{\mathscr{C}}[-1]$ and the canonical inclusion $\overline{\mathscr{C}}[-1] \to \Omega\mathscr{C}$.

2.2.2 Bar–cobar adjunction Suppose that $\mathscr{C}$ is a coaugmented A_∞–coalgebra and $\mathscr{A}$ is an augmented DG–algebra. Then we have a canonical bijection

$$\mathrm{hom}_{\mathrm{DG}}(\Omega\mathscr{C}, \mathscr{A}) \to \mathrm{Tw}(\mathscr{C}, \mathscr{A}) \tag{21}$$

given by $\phi \mapsto \phi \circ \mathfrak{t}^{\mathscr{C}}$. Similarly, if $\mathscr{C}$ is a coaugmented conilpotent DG–coalgebra and $\mathscr{A}$ is an augmented A_∞–algebra, then we have a canonical bijection

$$\mathrm{hom}_{\mathrm{coDG}}(\mathscr{C}, \mathrm{B}\mathscr{A}) \to \mathrm{Tw}(\mathscr{C}, \mathscr{A}), \tag{22}$$

given by $\phi \mapsto \mathfrak{t}_{\mathscr{A}} \circ \phi$; see [57, lemme 3.17].

Therefore, when $\mathscr{C}$ is a coaugmented conilpotent DG–coalgebra, and $\mathscr{A}$ is an augmented DG–algebra, we have the bar–cobar adjunction

$$\hom_{\mathrm{DG}}(\Omega\mathscr{C},\mathscr{A}) \cong \hom_{\mathrm{coDG}}(\mathscr{C},\mathrm{B}\mathscr{A}).$$

Moreover, the natural DG–maps

$$\Omega\mathrm{B}\mathscr{A} \to \mathscr{A} \quad \text{and} \quad \mathscr{C} \to \mathrm{B}\Omega\mathscr{C} \tag{23}$$

are quasi-isomorphisms for any DG–algebra $\mathscr{A}$ and conilpotent DG–coalgebra $\mathscr{C}$; see [56, Section 6.10]. It is also true that for any A_∞–algebra $\mathscr{A}$, the A_∞–algebra map

$$\mathscr{A} \to \Omega\mathrm{B}\mathscr{A}$$

given by the adjunction map $\mathrm{B}\mathscr{A} \to \mathrm{B}\Omega\mathrm{B}\mathscr{A}$ is an A_∞–quasi-isomorphism; see [47, lemme 2.3.4.3]. Note that any A_∞–quasi-isomorphism is invertible up to homotopy [59, Corollary 1.4].

Similarly, for any A_∞–coalgebra $\mathscr{C}$, the A_∞–comap

$$\mathrm{B}\Omega\mathscr{C} \to \mathscr{C}$$

given by the adjunction map $\Omega\mathrm{B}\Omega\mathscr{C} \to \Omega\mathscr{C}$ is an A_∞–quasi-isomorphism.

However, an A_∞–quasi-isomorphism for a general A_∞–coalgebra is not usually a convenient notion since, as we remarked above, a quasi-isomorphism of A_∞–coalgebras between $\mathscr{C}$ and $\mathbb{C}'$ does not necessarily induce a quasi-isomorphism of DG–algebras $\Omega\mathscr{C}$ and $\Omega\mathbb{C}'$.

For this reason, one considers the category of *conilpotent* A_∞–coalgebras. Let $\mathscr{C}$ be a coaugmented A_∞–coalgebra generated over $\boldsymbol{k}$ by variables $(c_i)_{i\in I}$, with I some countable index set, such that there exists a total ordering

$$c_{\sigma(1)} < c_{\sigma(2)} < \cdots,$$

where $\sigma\colon I \to I$ is a bijection. This produces an increasing filtration

$$\mathcal{F}^0 = \boldsymbol{k} \subset \mathcal{F}^1 \subset \cdots \subset \Omega\mathscr{C}$$

by setting $\mathcal{F}^p = \boldsymbol{k}\langle c_{\sigma(1)},\dots,c_{\sigma(p)}\rangle$. Suppose that the structure maps $(\Delta_i)_{i\geq 1}$ are compatible with this filtration, in the sense that $\Delta_i(c_{\sigma(p)}) \subset \mathcal{F}^{p-1}$ for all i and p. Then we call $\mathscr{C}$ a conilpotent A_∞–coalgebra. (More generally, homotopy retracts of such A_∞–coalgebras are called conilpotent [56, Sections 6.10 and 9]. This notion is called *finite type* in [46].). Given two such A_∞–coalgebras $\mathscr{C}$ and $\mathbb{C}'$, one considers filtered A_∞–comaps between them. In the case of a conilpotent DG–coalgebra $\mathscr{C}$ there exists an increasing filtration on $\Omega\mathscr{C}$ given by the subalgebras $\mathrm{Ker}(\Delta^{(n)})$ that plays the same role; see [47, lemme 1.3.2.3].

We next state the following elementary lemma for later convenience.

Lemma 10 *Let $\mathscr{A}$ be an augmented A_∞–algebra such that the $\boldsymbol{k}$–bimodule structures on $\mathscr{A}$ and $\mathrm{B}\mathscr{A}$ are locally finite. Then there are quasi-isomorphisms of augmented DG–algebras*

$$\Omega(\mathscr{A}^{\#}) \to (\mathrm{B}\mathscr{A})^{\#} \quad \textit{and} \quad \Omega({}^{\#}\mathscr{A}) \to {}^{\#}(\mathrm{B}\mathscr{A}).$$

Note that the assumption is satisfied when $\mathscr{A}$ is locally finite and simply connected. We shall briefly consider the case when $\mathscr{A}$ is only assumed to be locally finite and connected, in which case we have:

Lemma 11 *Let $\mathscr{A} = \bigoplus_i \mathscr{A}^i$ be a connected, locally finite $\boldsymbol{k}$–bimodule equipped with an augmented A_∞–algebra structure. Then there are maps of DG–algebras*

$$\Omega(\mathscr{A}^{\#}) \to (\mathrm{B}\mathscr{A})^{\#} \quad \textit{and} \quad \Omega({}^{\#}\mathscr{A}) \to {}^{\#}(\mathrm{B}\mathscr{A})$$

which become quasi-isomorphisms, after completion,

$$\widehat{\Omega}(\mathscr{A}^{\#}) \to (\mathrm{B}\mathscr{A})^{\#} \quad \textit{and} \quad \widehat{\Omega}({}^{\#}\mathscr{A}) \to {}^{\#}(\mathrm{B}\mathscr{A}).$$

2.3 Koszul duality

Suppose $\mathscr{C}$ is a coaugmented conilpotent A_∞–coalgebra and $\mathscr{A}$ is an augmented DG–algebra. Via the bijection (21), any twisting cochain $\mathfrak{t} \in \mathrm{Tw}(\mathscr{C}, \mathscr{A})$ is of the form $\mathfrak{t} = \phi \circ \mathfrak{t}^{\mathscr{C}}$ for some unique $\phi \in \hom_{\mathrm{DG}}(\Omega\mathscr{C}, \mathscr{A})$. Similarly, if $\mathscr{C}$ is a coaugmented conilpotent DG–coalgebra and $\mathscr{A}$ is an augmented A_∞–algebra, any twisting cochain $\mathfrak{t} \in \mathrm{Tw}(\mathscr{C}, \mathscr{A})$ is of the form $\mathfrak{t} = \mathfrak{t}_{\mathscr{A}} \circ \phi$ for some $\phi \in \hom_{\mathrm{coDG}}(\mathscr{C}, \mathrm{B}\mathscr{A})$.

Definition 12 In either case above we call $\mathfrak{t}$ a *Koszul twisting cochain* if ϕ is a quasi-isomorphism, and we denote the set of Koszul twisting cochains by $\mathrm{Kos}(\mathscr{C}, \mathscr{A})$.

The terminology of Koszul twisting cochains is taken from [49]. They are also called *acyclic twisting cochains* in other sources [47; 56]. This terminology is due to the well-known fact that, under various local-finiteness assumptions, a twisting cochain $\mathfrak{t}$ is Koszul if and only if the Koszul complex (17) associated to $\mathfrak{t}$ is acyclic; see [56, Appendix A].

Informally, if $\mathfrak{t} \in \mathrm{Kos}(\mathscr{C}, \mathscr{A})$, then, depending on whether we write $\mathfrak{t} = \phi \circ \mathfrak{t}^{\mathscr{C}}$ or $\mathfrak{t} = \mathfrak{t}_{\mathscr{A}} \circ \phi$, either $\mathscr{A}$ can be used in place of $\Omega\mathscr{C}$, or $\mathscr{C}$ can be used in place of $\mathrm{B}\mathscr{A}$ in various resolutions. This, in turn, may lead to smaller complexes to compute with. For example, one can compute Hochschild homology and cohomology of $\mathscr{A}$ and $\Omega\mathscr{C}$ using the $\mathscr{A}$–bimodule resolution of $\mathscr{A}$ given by the complex

$$\mathscr{A} \otimes_{\boldsymbol{k}}^{\mathfrak{t}} \mathscr{C} \otimes_{\boldsymbol{k}}^{\mathfrak{t}} \mathscr{A}$$

with the differential as in (19); see [39].

Suppose that $\mathscr{A}$ is an A_∞–algebra with an augmentation $\epsilon\colon \mathscr{A} \to \boldsymbol{k}$. The augmentation ϵ makes $\boldsymbol{k}$ into a left $\mathscr{A}$–module, or equivalently, a right $\mathscr{A}^{\mathrm{op}}$–module.

Definition 13 The *Koszul dual* of an augmented A_∞–algebra $\mathscr{A}$ is the DG–algebra of left $\mathscr{A}$–module maps from $\boldsymbol{k}$ to itself,

$$E(\mathscr{A}) := \mathrm{RHom}_{\mathscr{A}}(\boldsymbol{k}, \boldsymbol{k}).$$

Recall that for a unital A_∞–algebra $\mathscr{A}$ over a field $\mathbb{K}$ (or a semisimple ring such as $\boldsymbol{k}$), any A_∞–module is both h–projective and h–injective; that is, if M is an A_∞–module over $\mathscr{A}$ and N is an acyclic A_∞–module over $\mathscr{A}$, then the complexes $\mathrm{RHom}_{\mathscr{A}}(M, N)$ and $\mathrm{RHom}_{\mathscr{A}}(N, M)$ are acyclic [59, Lemma 1.16]. Hence, the DG–algebra $\mathrm{RHom}_{\mathscr{A}}(\boldsymbol{k}, \boldsymbol{k})$ can be computed as the A_∞–module homomorphisms from $\boldsymbol{k}$ to itself. (More generally, this holds if $\mathscr{A}$ is h–projective as a complex of $\boldsymbol{k}$–modules, which implies that $\boldsymbol{k}$ is h–projective as an A_∞–module over $\mathscr{A}$.) Therefore, we have the following:

Proposition 14 *If $\mathscr{A}$ (resp. $\mathscr{A}^{\mathrm{op}}$) is an augmented unital A_∞–algebra, then*

$$\mathrm{RHom}_{\mathscr{A}}(\boldsymbol{k}, \boldsymbol{k}) \cong (\mathrm{B}\mathscr{A})^{\#} \quad (\textit{resp. } {}^{\#}(\mathrm{B}\mathscr{A})).$$

Proof Recall that $\mathscr{A} \otimes_{\boldsymbol{k}} \mathrm{B}\mathscr{A}$ is quasi-isomorphic to $\boldsymbol{k}$ as an $\mathscr{A}$–module. Hence, by the hom-tensor adjunction, we have $\mathrm{RHom}_{\mathscr{A}}(\mathscr{A} \otimes_{\boldsymbol{k}} \mathrm{B}\mathscr{A}, \boldsymbol{k}) \cong \mathrm{RHom}_{\boldsymbol{k}}(\mathrm{B}\mathscr{A}, \boldsymbol{k})$. Since $\mathscr{A}$ is h–projective as a complex of $\boldsymbol{k}$–modules, so is $\mathrm{B}\mathscr{A}$; hence the latter is computed by $(\mathrm{B}\mathscr{A})^{\#}$. □

In this model of $E(\mathscr{A})$, the $\boldsymbol{k}$–bimodule structure on $E(\mathscr{A})$ can be seen as in (10), since $\boldsymbol{k}$ is viewed as a left $\boldsymbol{k}$–module induced from its structure as a left $\mathscr{A}$–module. If, instead, we have an augmentation of $\mathscr{A}^{\mathrm{op}}$, then we view $\boldsymbol{k}$ as a right $\mathscr{A}$–module, and the $\boldsymbol{k}$–bimodule structure on $\mathrm{RHom}_{\mathscr{A}}(\boldsymbol{k}, \boldsymbol{k})$ would be given by (11).

The cohomology of $E(\mathscr{A})$ is a graded algebra,

$$\mathrm{Ext}_{\mathscr{A}}(\boldsymbol{k},\boldsymbol{k}) := H^*(\mathrm{RHom}_{\mathscr{A}}(\boldsymbol{k},\boldsymbol{k})) \cong H^*((\mathrm{B}\mathscr{A})^{\#}).$$

Dually, we also have the derived tensor product $\boldsymbol{k}\widehat{\otimes}_{\mathscr{A}}\boldsymbol{k}$, which can be computed by the complex $\mathrm{B}\mathscr{A}$. The cohomology is a graded coalgebra

$$\mathrm{Tor}_{\mathscr{A}}(\boldsymbol{k},\boldsymbol{k}) := H^*(\boldsymbol{k}\widehat{\otimes}_{\mathscr{A}}\boldsymbol{k}) \cong H^*(\mathrm{B}\mathscr{A}).$$

In particular, if $\boldsymbol{k}$ is a field, we have that $\mathrm{Ext}_{\mathscr{A}}(\boldsymbol{k},\boldsymbol{k}) \cong (\mathrm{Tor}_{\mathscr{A}}(\boldsymbol{k},\boldsymbol{k}))^{\#}$ by the universal coefficient theorem.

Remark 15 If $\mathscr{A}$ is a *commutative* algebra (or more generally an E_2–algebra), then $\mathrm{Tor}_{\mathscr{A}}(\boldsymbol{k},\boldsymbol{k})$ also has a graded algebra structure, defined via

$$\mathrm{Tor}_{\mathscr{A}}(\boldsymbol{k},\boldsymbol{k}) \otimes \mathrm{Tor}_{\mathscr{A}}(\boldsymbol{k},\boldsymbol{k}) \to \mathrm{Tor}_{\mathscr{A}\otimes\mathscr{A}}(\boldsymbol{k}\otimes\boldsymbol{k},\boldsymbol{k}\otimes\boldsymbol{k}) \to \mathrm{Tor}_{\mathscr{A}}(\boldsymbol{k},\boldsymbol{k}),$$

induced by the algebra map $\mathscr{A}\otimes\mathscr{A}\to\mathscr{A}$ (which exists since $\mathscr{A}$ is commutative). This should not be confused with the natural coalgebra structure above.

Note that $\mathscr{A}$ itself can be viewed as a left $\mathscr{A}$–module and the map $\epsilon\colon \mathscr{A}\to\boldsymbol{k}$ is a map of left $\mathscr{A}$–modules; hence, it induces a map of left $E(\mathscr{A})^{\mathrm{op}}$–modules

$$\widetilde{\epsilon}\colon \mathrm{RHom}_{\mathscr{A}}(\boldsymbol{k},\boldsymbol{k})^{\mathrm{op}} \to \mathrm{RHom}_{\mathscr{A}}(\mathscr{A},\boldsymbol{k}),$$

which can in turn be viewed as an augmentation of $E(\mathscr{A})^{\mathrm{op}} = \mathrm{RHom}_{\mathscr{A}}(\boldsymbol{k},\boldsymbol{k})^{\mathrm{op}}$, since $\mathrm{RHom}_{\mathscr{A}}(\mathscr{A},\boldsymbol{k})$ can again be identified with $\boldsymbol{k}$ as it is the Yoneda image of $\boldsymbol{k}$ as an $\mathscr{A}$–module. Hence, $\boldsymbol{k}$ can be viewed as a right $E(\mathscr{A})$–module.

Definition 16 The double dual of $\mathscr{A}$ is defined to be $E(E(\mathscr{A})) := \mathrm{RHom}_{E(\mathscr{A})}(\boldsymbol{k},\boldsymbol{k})$.

There is a natural map from $\mathscr{A}$ to its double dual,

$$\Phi\colon \mathscr{A}\to\mathrm{RHom}_{E(\mathscr{A})}(\boldsymbol{k},\boldsymbol{k}),$$

defined via viewing the right $E(\mathscr{A})$–module $\boldsymbol{k}$ as $\mathrm{RHom}_{\mathscr{A}}(\mathscr{A},\boldsymbol{k})$ and acting on the left by $\mathscr{A}\cong\mathrm{RHom}_{\mathscr{A}}(\mathscr{A},\mathscr{A})$.

Definition 17 We say that $\mathscr{A}$ and $E(\mathscr{A})$ are Koszul dual if $\Phi\colon \mathscr{A}\to\mathrm{RHom}_{E(\mathscr{A})}(\boldsymbol{k},\boldsymbol{k})$ is a quasi-isomorphism.

One standard situation in which Koszul duality holds is the following:

Theorem 18 *Suppose $\mathscr{C} = \bigoplus_{i\leq 0}\mathscr{C}^i$ is a locally finite, simply connected $\boldsymbol{k}$–bimodule equipped with an A_∞–coalgebra structure and the coaugmentation $\boldsymbol{k} \cong \mathscr{C}^0 \to \mathscr{C}$. Let $\mathscr{A} = \Omega\mathscr{C}$, which is an augmented connected DG–algebra. Then $E(\mathscr{A}) \cong \mathscr{C}^{\#}$, and $\mathscr{A}$ and $\mathscr{C}^{\#}$ are Koszul dual. In other words, the natural morphism*

$$\Omega\mathscr{C} \to \mathrm{RHom}_{E(\mathscr{A})}(\boldsymbol{k}, \boldsymbol{k})$$

is a quasi-isomorphism.

Proof First, observe that indeed $E(\mathscr{A}) \cong (\mathrm{B}\mathscr{A})^{\#} \cong (\mathrm{B}\Omega\mathscr{C})^{\#} \cong \mathscr{C}^{\#}$ by (23) and because $\mathrm{Hom}_{\boldsymbol{k}-}(-, \boldsymbol{k})$ preserves quasi-isomorphisms. Next, we have that

$$\mathrm{RHom}_{E(\mathscr{A})}(\boldsymbol{k}, \boldsymbol{k}) \cong {}^{\#}(\mathrm{B}(\mathscr{C}^{\#})) \cong \Omega\mathscr{C},$$

where we applied Lemma 10 to $\mathscr{C}^{\#}$ and used the fact that ${}^{\#}(\mathscr{C}^{\#}) \cong \mathscr{C}$ since $\mathscr{C}$ is locally finite. □

Rather than making the grading assumptions on $\mathscr{C}$ as in Theorem 18, which guarantee that $\mathrm{B}\mathscr{C}^{\#}$ is locally finite, one can directly assume that the grading on the cohomology $H^*(\Omega\mathscr{C})$ is locally finite. This assumption is harder to check in practice but Koszul duality still holds under this assumption, which one can prove by combining the above argument with the homological perturbation lemma; see for example [43, Theorem 2.8].

In the case that $\mathscr{C} = \bigoplus_{i\leq 0}\mathscr{C}^i$ is a locally finite, connected (but not simply connected) $\boldsymbol{k}$–bimodule, Lemma 10 no longer applies. We instead use Lemma 11 to deduce the following weaker duality result:

Proposition 19 *Let $\mathscr{C} = \bigoplus_{i\leq 0}\mathscr{C}^i$ be a connected, locally finite $\boldsymbol{k}$–bimodule, equipped with an A_∞–coalgebra structure and coaugmentation $\boldsymbol{k} \cong \mathscr{C}^0 \to \mathscr{C}$, and let $\mathscr{A} = \Omega\mathscr{C} = \boldsymbol{k} \oplus \bigoplus_{j\geq 1}(\overline{\mathscr{C}}[-1])^{\otimes_{\boldsymbol{k}} j}$, which is an augmented DG–algebra where augmentation is given by projection to $\boldsymbol{k}$. Then $E(\mathscr{A}) \cong \mathscr{C}^{\#}$ and there is a quasi-isomorphism*

$$\widehat{\Omega}\mathscr{C} \to \mathrm{RHom}_{E(\mathscr{A})}(\boldsymbol{k}, \boldsymbol{k}).$$

Note that in Proposition 19, $\mathscr{A} = \Omega\mathscr{C}$ is not connected, and may admit other augmentations $\epsilon\colon \mathscr{A} \to \boldsymbol{k}$ than that induced by the cobar construction. Such augmentations will be considered below. For example, suppose that $\mathscr{C} \cong \boldsymbol{k} \oplus \overline{\mathscr{C}}$ is a coaugmented A_∞–coalgebra such that $\overline{\mathscr{C}} = \mathbb{K}\langle c | c \in \mathcal{R}\rangle$ is generated by elements c from an indexing set $\mathcal{R}$ and that $\epsilon\colon \Omega\mathscr{C} \to \boldsymbol{k}$ is an augmentation, which is induced by a map $\mathscr{C} \to \boldsymbol{k}$ since $\Omega\mathscr{C}$ is free. Now we can consider the coaugmented A_∞–coalgebra $\mathscr{C}^\epsilon = \boldsymbol{k} \oplus \overline{\mathscr{C}}^\epsilon$ such that

$$\overline{\mathscr{C}}^\epsilon = \mathbb{K}\langle c - \epsilon(c)1_{\boldsymbol{k}} \mid c \in \mathcal{R}\rangle.$$

Then $\Omega\mathscr{C}$ and $\Omega\mathscr{C}^\epsilon$ are quasi-isomorphic as nonaugmented DG–algebras, and the augmentation on $\Omega\mathscr{C}^\epsilon$ induced by the cobar construction coincides with the given augmentation ϵ on Ω.

Remark 20 When $\mathscr{C}$ is not simply connected, the proof of duality fails precisely because $\mathrm{B}\mathscr{C}^\#$ is not locally finite. Nevertheless, the duality result can still be proved in certain cases where an extra *weight* grading (internal degree, or Adams degree) is available; see [50; 56, Appendix A.2; 39]. We will not study this situation systematically in this paper, but it is important as it extends the range of applicability of Koszul duality theory. In the setting of Chekanov–Eliashberg DG–algebras, such a situation was considered in [32].

3 Legendrian (co)algebra

In this section we introduce our Legendrian invariants. We start by discussing a model for loop space coefficients in Section 3.1. In Section 3.2 we define the Chekanov–Eliashberg algebra with loop space coefficients using moduli spaces of disks of all dimensions, and in Section 3.4 we give a more computable version, which uses only rigid disks and which carries the same information if the Legendrian submanifold is simply connected.

3.1 Coefficients

Before defining our Legendrian invariants, we describe chain models for their coefficients $C_{-*}(\Omega_{p_v}\Lambda_v)$ for $v \in \Gamma^+$. (Notation is as above, Λ_v is a $+$ decorated connected component of the Legendrian Λ.) We work over a field $\mathbb{K}$.

Let $\Omega_{p_v}\Lambda_v$ denote the topological monoid of Moore loops based at p_v, where the monoid structure comes from concatenation of loops; see [5]. Write $C_{-*}(\Omega_{p_v}\Lambda_v)$ for the cubical chain complex (graded cohomologically). Since $\Omega_{p_v}\Lambda_v$ is a topological monoid, the complex $C_{-*}(\Omega_{p_v}\Lambda_v)$ becomes a DG–algebra using the natural product map $\times$ on cubical chains, where the DG–algebra product is given as

$$C_{-*}(\Omega_{p_v}\Lambda_v) \otimes C_{-*}(\Omega_{p_v}\Lambda_v) \xrightarrow{\times} C_{-*}(\Omega_{p_v}\Lambda_v \times \Omega_{p_v}\Lambda_v) \xrightarrow{\circ} C_{-*}(\Omega_{p_v}\Lambda_v).$$

We point out that the $\times$–map

$$\times\colon C_{-*}(\Omega_{p_v}\Lambda_v) \otimes C_{-*}(\Omega_{p_v}\Lambda_v) \to C_{-*}(\Omega_{p_v}\Lambda_v \times \Omega_{p_v}\Lambda_v),$$

when both sides are equipped with the Pontryagin product, is a DG–algebra map.

In what follows, we shall also make use of an inverse to the $\times$, known as the Serre diagonal [60], and the cubical analogue of the Alexander–Whitney map,

$$\eta\colon C_{-*}(\Omega_{p_v}\Lambda_v \times \Omega_{p_w}\Lambda_w) \to C_{-*}(\Omega_{p_v}\Lambda_v) \otimes C_{-*}(\Omega_{p_w}\Lambda_w). \tag{24}$$

To define this map consider the n–cube I^n with coordinates $(x_1,\dots,x_n)$. For an ordered j–element subset $J \subset \{1,2,\dots,n\}$, $J=(i_1,\dots,i_j)$ with $i_1<\dots<i_j$, and for $\epsilon \in \{0,1\}$, let $\iota_J^\epsilon\colon I^j \to I^n$ be the map given in coordinates $y=(y_1,\dots,y_j)$ by

$$x_{i_r}(\iota_J(y)) = y_r \qquad \text{and} \qquad x_m(\iota_J(y)) = \epsilon \quad \text{if } m \notin J.$$

Consider a cubical chain $(\sigma,\tau)\colon I^n \to \Omega_{p_v}\Lambda_v \times \Omega_{p_w}\Lambda_w$. If J is an ordered subset of $\{1,\dots,n\}$, let J' denote its complement ordered in the natural way. Define η by

$$\eta(\sigma,\tau) = \sum_J (-1)^{JJ'}(\sigma \circ \iota_J^0) \otimes (\tau \circ \iota_{J'}^1),$$

where the sum ranges over all ordered subsets J, and $(-1)^{JJ'}$ is the sign of the permutation JJ'. This is a strictly associative chain map inducing a quasi-isomorphism. Note also that there are obvious extensions of η to several products of loop spaces.

As the cubical chain complex $C_{-*}(\Omega_{p_v}\Lambda_v)$ is very large, it is not the most effective complex for computation. We next discuss smaller models. Starting with a 0–reduced simplicial set X with geometric realization $|X| = \Lambda_v$, an explicit economical model for $C_{-*}(\Omega_{p_v}\Lambda_v)$ is obtained by taking normalized chains on the Kan loop group GX; see [44]. We will not say much about this, but point out that GX is a free simplicial group, whose geometric realization $|GX|$ is homotopy equivalent to $\Omega|X|$; see [37, Corollary 5.11]. Hence, by the monoidal Dold–Kan correspondence [58], the normalized chains on GX give a (weakly) equivalent model of $C_{-*}(\Omega_{p_v}\Lambda_v)$. (Another similar construction is sketched in [46], and leads to a free model.)

Alternatively, one can work with CW–complexes. We start with the simply connected case: for a 1–reduced (unique 0–cell and no 1–cells) CW–structure on Λ_v, the Adams–Hilton construction [5] gives a free DG–algebra model for $C_{-*}(\Omega_{p_v}\Lambda_v)$ as follows. Denote the k–cells of Λ_v by e_k^i for $k \geq 2$ and $i = 1,\dots,m_k$. The Adams–Hilton construction gives a CW–monoid with a single 0–cell, and generating cells $\bar{e}_k^i$ in dimension $k-1$, which is quasi-isomorphic to $\Omega_{p_v}(\Lambda_v)$ as a monoid; see [14]. This gives a DG–algebra structure on the free algebra,

$$A(\Lambda_v) := \mathbb{K}\langle \bar{e}_2^1,\dots,\bar{e}_2^{m_2},\bar{e}_3^1,\dots,\bar{e}_3^{m_3},\dots,\dots,\dots\rangle, \quad \text{with } |\bar{e}_k^i| = 1-k,$$

and a DG–algebra map

$$A(\Lambda_v) \xrightarrow{\Psi} C_{-*}(\Omega_{p_v}\Lambda_v),$$

which is a quasi-isomorphism. The differential d on $A(\Lambda_v)$ is generally not explicit. It is defined recursively as follows. For every 2–cell e_2^i, we have $d(\bar{e}_2^i) = 0$. In general, assuming that d_{k-1} and Ψ_{k-1} have been defined on the k–skeleton $\Lambda_v^{(k)}$ of Λ_v, then for each $(k+1)$–cell e, with attaching map $f\colon S^k \to \Lambda_v^{(k)}$, define $d_k\bar{e} = c$ so that $(\Psi_{k-1})(c) = (\Omega f)_*(\xi)$, where ξ a generator of $H_{k-1}(\Omega S^k)$, and define $\Psi_k(\bar{e})$ to be the k–chain of loops in e (which then depends on earlier choices along the boundary of $\bar{e}$). We remark that $A(\Lambda_v)$ can be identified isomorphically with $\Omega C_*^{\mathrm{CW}}(\Lambda_v)$ for a suitable A_∞–coalgebra structure on the cellular chain complex $C_*^{\mathrm{CW}}(\Lambda)$.

This construction can be generalized to the non-simply-connected case as follows.[1] Begin with a 0–reduced CW–structure on Λ_v. Denote the k–cells by e_k^i for $i = 1,\dots,m_k$. For each k–cell e_i^k with $k \ge 2$, we have a free variable in degree $1-k$, which we again denote by $\bar{e}_k^i$. For each 1–cell e_1^j with $j = 1,\dots,m_1$, we have two variables t_j and t_j^{-1} in degree 0 such that $t_j t_j^{-1} = 1 = t_j^{-1}t_j$. Thus, the underlying algebra is the "almost free" algebra of the form

$$A(\Lambda_v) := \mathbb{K}\langle t_1^{\pm 1},\dots,t_{m_1}^{\pm 1}, \bar{e}_2^1,\dots,\bar{e}_2^{m_2}, \bar{e}_3^1,\dots,\bar{e}_3^{m_3},\dots\rangle.$$

This presentation is often more efficient than the presentation one gets from the Kan loop group construction using a simplicial set presentation of Λ_v. However, the differential in the Adams–Hilton model is not easy to describe explicitly. Note that we have

$$d(t_j) = d(t_j^{-1}) = 0$$

for degree reasons. For every 2–cell e_2^i, we have

$$d(\bar{e}_2^i) = 1 - c_i,$$

where $c_i \in \mathbb{K}\langle t_j^{\pm 1} \mid j = 1,\dots,m_1\rangle$ represents the class of the attaching map of e_2^i. The differential on higher-dimensional cells is generally harder to compute and is exactly as in the simply connected case discussed above.

Augmentations $\epsilon\colon A(\Lambda_v) \to \mathbb{K}$ correspond to solutions of the equations

$$\begin{cases} \epsilon(t_j)\epsilon(t_j^{-1}) = 1 & \text{for } j = 1,\dots,m_1, \\ \epsilon(d\bar{e}_2^i) = 0 & \text{for } i = 1,\dots,m_2. \end{cases}$$

[1] See [40; 41]: a generalization was given earlier in [33], however that paper contains an error.

Since $\mathbb{K}\langle t_j^{\pm 1}, j = 1, \dots, m_1 \mid d\bar{e}_2^i, i = 1, \dots, m_2\rangle$ is a presentation of the fundamental group algebra $\mathbb{K}[\pi_1(\Lambda_v, p_v)]$, augmentations correspond exactly to local systems $\pi_1(\Lambda_v, p_v) \to \mathbb{K}$.

We will use the cubical chain complex $C_{-*}(\Omega_{p_v}\Lambda_v)$ to define Legendrian invariants below. Cubical chains work uniformly for all spaces Λ_v and are convenient for showing that the fundamental classes of moduli spaces of pseudoholomorphic disks $\mathcal{M}^{\mathrm{sy}}$, via evaluation maps, take values in the chain complex. The Legendrian invariants can also be studied using any of the smaller models discussed above. It is however important to note that in the non-simply-connected case, we only have either weak equivalence in the homotopy category of DG–algebras, or Morita equivalence [40; 41] of these models and the cubical chain complex $C_{-*}(\Omega_{p_v}\Lambda_v)$.

In the case that Λ_v is simply connected, we can use a DG–algebra map

$$\Phi\colon C_{-*}(\Omega_{p_v}\Lambda_v) \to A(\Lambda_v)$$

that goes in the opposite direction to the Adams–Hilton map to pass to a more economical quasi-isomorphic model. Such a homotopy equivalence Φ is constructed in two steps: first construct, as in [54] using Eilenberg–Moore methods, a DG–algebra quasi-isomorphism

$$C_{-*}(\Omega_{p_v}\Lambda_v) \to \Omega C_*(\Lambda_v), \tag{25}$$

where in both instances C_* refers to the normalized singular chains. Second, using the standard A_∞–coalgebra quasi-isomorphism between the DG–coalgebra of singular chains $C_*(\Lambda_v)$ and the A_∞–coalgebra $C_*^{\mathrm{CW}}(\Lambda_v)$ of normalized cellular chains, one obtains a DG–algebra quasi-isomorphism

$$\Omega C_*(\Lambda_v) \to \Omega C_*^{\mathrm{CW}}(\Lambda_v) = A(\Lambda_v),$$

since we assumed that the complexes C_* and C_*^{CW} are simply connected. (In Section 3.5, we also give a more geometric construction of a DG–algebra quasi-isomorphism Φ corresponding to (25) landing in Morse chains, using Morse flow trees.)

Similarly, if Λ_v is homotopy equivalent to an Eilenberg–Mac Lane space $\mathrm{K}(\pi_1, 1)$, then the singular chains can be replaced with the group algebra $\mathbb{K}[\pi_1]$: there exists a quasi-isomorphism of DG–algebras

$$C_{-*}(\Omega_{p_v}\Lambda_v) \to \mathbb{K}[\pi_1]$$

given by sending a 0–chain to its homology class, and sending all higher-dimensional chains to 0. Note that this DG–algebra map exists for any space Λ_v, but is a quasi-isomorphism only in the case that Λ_v is homotopy equivalent to $\mathrm{K}(\pi_1, 1)$.

It is often convenient to use a cofibrant (or free) replacement for $\mathbb{K}[\pi_1]$. For example, if $\Lambda_v = S^1$, then $\mathbb{K}[\pi_1] \cong \mathbb{K}[t, t^{-1}]$ and a cofibrant replacement is given by the free graded algebra

$$\mathbb{K}\langle s_1, t_1, k_1, l_1, u_1\rangle, \quad \text{where } |s_1| = |t_1| = 0,\ |k_1| = |l_1| = -1,\ |u_1| = -2,$$

with the differential

$$dk_1 = 1 - s_1 t_1, \quad dl_1 = 1 - t_1 s_1, \quad du_1 = k_1 s_1 - s_1 l_1.$$

A DG–algebra defined over $\mathbb{K}[t, t^{-1}]$ can be pulled back to a weakly equivalent DG–algebra over this cofibrant replacement. (See [63] for background in model categories on DG–algebras that we are using in a very simple case here.)

3.2 Construction of Legendrian invariants

As above, let X be a Liouville domain with $c_1(X) = 0$ (for $\mathbb{Z}$–grading) and $\partial X = Y$ its contact boundary. Let $\Lambda = \bigsqcup_{v \in \pi_0(\Lambda)} \Lambda_v$ be a Legendrian submanifold in Y, where Λ_v is a connected component of Λ. Assume that Λ is relatively spin and that its Maslov class vanishes. Let each connected component Λ_v be decorated with a sign and write Λ^+ and Λ^- for the union of the components decorated accordingly. (Our different treatment of Λ^+ and Λ^- is natural from the point of view of handle attachments; recall from the introduction that when Λ^- is a union of spheres, we attach usual Lagrangian disk-handles to Λ^- and handles with cotangent ends to Λ^+.) When we have an exact Lagrangian filling L of Λ (relatively spin and with vanishing Maslov class), L can also be decomposed into embedded components $L = \bigcup_{v \in \Gamma} L_v$. These embedded components are not disjoint: they are allowed to intersect transversely at finitely many points. There is a bijection between Γ and the embedded components of L.

We require that if two components Λ_{w_1} and Λ_{w_2} are boundary components of the same embedded component L_v, then either both belong to Λ^- or both to Λ^+. Using this property, we get a decomposition $\Gamma = \Gamma^+ \sqcup \Gamma^-$, corresponding to the decomposition $\Lambda = \Lambda^+ \sqcup \Lambda^-$.

Let $\boldsymbol{k}$ be the semisimple ring generated by mutually orthogonal idempotents $\{e_v\}_{v \in \Gamma}$. If we are not given a filling of Λ, then the index set Γ is taken to be the connected components, $\pi_0(\Lambda)$, instead. If we need to distinguish between the two choices, we will denote them as $\boldsymbol{k}_\Lambda$ and $\boldsymbol{k}_L$. Note that there is an injective ring map $\boldsymbol{k}_L \to \boldsymbol{k}_\Lambda$ which takes the idempotent e_v corresponding to an embedded component L_v to the sum $e_{w_1} + \cdots + e_{w_r}$ of idempotents of its boundary components Λ_{w_j}. In particular, this map turns any $\boldsymbol{k}_\Lambda$–bimodule into a $\boldsymbol{k}_L$–bimodule.

Let $\mathcal{R}$ denote the set of nonempty Reeb chords of Λ. This is a graded set: the grading of a chord $c \in \mathcal{R}$ is given by $|c| = -\mathrm{CZ}(c)$, where $\mathrm{CZ}(c)$ is the Conley–Zehnder grading; see Appendix A. (With this convention, the unique chord c of the standard Legendrian unknot in $\mathbb{R}^3$ has $|c| = -2$ and for the corresponding Legendrian unknot in $\mathbb{R}^{2n-1}$ with one Reeb chord c, we have $|c| = -n$. See also Remark 30.)

Note that the vector space generated by $\mathcal{R}$ is a $\boldsymbol{k}$–bimodule, where $e_v \mathcal{R} e_w$ corresponds to the set of Reeb chords from Λ_v to Λ_w. The underlying algebra of the standard Chekanov–Eliashberg DG–algebra is generated freely by $\mathcal{R}$ over $\boldsymbol{k}$. We need to modify this in the case that Λ^+ is nonempty to incorporate chains in the based loop space of Λ_v for $v \in \Gamma^+$. Let us first do this using cubical chains.

For each $v \in \Gamma^+$, consider the cubical chains $C_{-*}(\Omega_{p_v} \Lambda_v)$ as a $\boldsymbol{k}$–algebra by requiring that the left or right action of e_w is trivial except if $w = v$, when it acts as identity. Let CE^* be the algebra over $\boldsymbol{k}$ given by adjoining elements of $\mathcal{R}$ to the union of $C_{-*}(\Omega_{p_v} \Lambda_v)$ for $v \in \Gamma^+$. Thus an element of CE^* is a sum of alternating words in Reeb chords, $\sigma_1 c_1 \sigma_2 c_2 \cdots \sigma_m c_m \sigma_{m+1}$, where c_j are Reeb chords and σ_j chains of based loops in the component of the Legendrian where the adjacent Reeb chord lies.

Now the differential on CE^* is defined by extending the differential on the cubical complexes $C_{-*}(\Omega_{p_v} \Lambda_v)$ for $v \in \Gamma_+$. We describe the differential on a single Reeb chord and extend it by the graded Leibniz rule. The differential d on a Reeb chord decomposes to a sum

$$d = \sum_{i \geq 0} \Delta_i,$$

where for any Reeb chord c_0 only finitely many $\Delta_i(c_0)$ are nonzero. The operations $\Delta_i(c_0)$ are defined as follows.

Consider moduli spaces of holomorphic disks with positive puncture at c_0; for definitions and notation see Appendix A. More precisely, consider Reeb chords $c_i, \dots, c_1$ such that $c_0 c_i \cdots c_1$ is a composable word and let $\boldsymbol{c} = c_0^+ c_i^- \cdots c_1^-$. Consider the space of disks D_{i+1} with one distinguished positive puncture and i negative punctures (across which the boundary numbering is constant, in the terminology of Appendix A).

Consider the moduli space $\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})$. As we use a translation-invariant almost complex structure on the symplectization, $\mathbb{R}$ acts by this moduli space by translation. Write

$$\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}) := \mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})/\mathbb{R} \tag{26}$$

for the quotient. Theorems 74 and 75 imply that $\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})$ is a smooth orientable manifold, with a natural compactification as a stratified space that carries a fundamental chain. It follows, via the evaluation map at a point in the boundary arcs of D_{i+1}, that $\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})$ parametrizes a chain of paths in the $(i+1)$–fold product $\Lambda^{\times(i+1)}$.

We transform these chains of paths to chains of based loops as follows. On each component Λ_v pick reference arcs connecting all Reeb chord endpoints to the basepoint. Let $U_v \subset \Lambda_v$ be a disk which is a regular neighborhood of these arcs. For convenience we take the disk to be smooth. Then a collar neighborhood on its boundary gives a smooth map $\theta_v\colon (\Lambda_v, *_v) \to (\Lambda_v, *_v)$ such that $\theta_v(D_v) = *_v$ and $\theta_v|_{\Lambda_v\setminus D_v}\colon \Lambda_v \setminus D_v \to \Lambda_v \setminus \{*_v\}$ is a diffeomorphism. To get a chain of loops parametrized by $\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})$ we compose its chains of paths with the maps θ_v. The resulting chain of paths then takes all Reeb chord endpoints in component Λ_v to the basepoint $*_v$. Thus by composition with θ_v, the moduli space parametrizes a chain of loops in $(\Omega_p\Lambda)^{\times(i+1)}$.

We treat two cases separately. First, if all boundary components of D_{i+1} map to components in Λ^-, then we let

$$[\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})] = \begin{cases} n c_i \cdots c_1 & \text{if } \dim(\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})) = 1, \\ 0 & \text{if } \dim(\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})) \neq 1, \end{cases} \tag{27}$$

where n is the algebraic number of $\mathbb{R}$ components in the moduli space. Second, if some boundary component maps to a component in Λ^+, then we write $[\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})]$ for the chain of paths in $(\Omega_p\Lambda)^{\times(i+1)}$, where we separate the components in the product by the Reeb chords $\boldsymbol{c}' = c_i \cdots c_1$:

$$[\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})] = \sigma_{i+1} c_i \sigma_i \cdots \sigma_2 c_1 \sigma_1,$$

where σ_j are the components of the fundamental chain $\sigma\colon \mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}) \to (\Omega\Lambda)^{\times(i+1)}$. Further, we write e_v for each boundary component that maps to a component in Λ^- in between the Reeb chords $c_i \cdots c_1$ as above.

A subtle point here is that the moduli space $\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})$ naturally gives rise to a chain σ in $C_{-*}(\Omega\Lambda^{\times(i+1)})$ rather than in $C_{-*}(\Omega\Lambda)^{\otimes(i+1)}$. Note that σ_i are simply components of σ, they are not considered as chains. To separate these out we apply the cubical Alexander–Whitney map

$$\eta\colon C_{-*}((\Omega\Lambda)^{\times(i+1)}) \to C_{-*}(\Omega\Lambda)^{\otimes(i+1)},$$

recalled in Section 3.1. With these conventions we then define for $i \geq 0$,

$$\Delta_i(c_0) := \sum_{\boldsymbol{c}=c_0^+ c_i^- \cdots c_1^-} \eta[\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})],$$

where we separate the components of the tensor product by the Reeb chords in $\boldsymbol{c}'$, in analogy with the notation for the product chain and where η is the Serre diagonal from equation (24). The output of $\Delta_i(c_0)$ is thus a sum of alternating words of chains of loops in $C_{-*}(\Omega\Lambda)$ and Reeb chords, and Δ_i is an operation of degree $2-i$ on $\mathrm{LC}_*(\Lambda)$. We point out that if there are Λ^+ components, then higher-dimensional moduli spaces contribute to the differential (unlike the case when $\Lambda = \Lambda^-$). Note also that it is possible to have holomorphic disks contributing to Δ_0, which means that the chord c_0 is the positive puncture of a disk without negative punctures.

Our next result shows that the operations Δ_i give a differential on CE^*. The proof uses boundaries of moduli spaces of holomorphic disks. By SFT compactness [12] and standard gluing results — see eg [31, Appendix A; 23, Appendix B] — the boundary of a moduli space $\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})$ consists of several level holomorphic buildings of curves with top level in $\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c}')$ and lower levels in $\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c}'')$, where the positive puncture of a curve in a lower level is attached at a negative puncture of a curve above it. In terms of $\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})$, standard gluing results imply that in a neighborhood of several-level curves where positive and negative punctures are joined at d Reeb chords, the moduli space $\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})$ is C^1–diffeomorphic to

$$[0,1)^d \times \prod_{j=1}^{d} \mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}_j), \tag{28}$$

where the product runs over positive punctures in the holomorphic building which are not the positive puncture of the curve in $\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})$.

We will use the compact notation $\star$ to denote all such broken configurations and write simply

$$\partial \mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}) = \mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}') \star \mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}'').$$

We next need to consider the fundamental chain of loops $[\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}') \star \mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}'')]$ carried by $\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}') \star \mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}'')$, or in other words the codimension 1 boundary of $[\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})]$. If the dimension of $\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})$ is d then its boundary gives $(d-1)$–dimensional chains of loops in Λ. Consider a several-level building with moduli space components $\mathcal{M}^{\mathrm{sy}}_j$ of dimension $d_j \geq 1$, $j = 1, \dots, m$. Then, by SFT compactness, $d + 1 = \sum_{j=1}^m d_j$. A boundary component of a several-level disk that consists of boundary segments from k disks in $\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}_{j_1}, \dots, \mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}_{j_k}$ will then carry a chain of loops in Λ of dimension

$\sum_{r=1}^{k}(d_{j_r}-1)\le d-1$, with equality only if the broken configuration consists of only two levels. It follows that only two level curves contribute to $[\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}')\star\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}'')]$. More precisely, the codimension 1 boundary of $\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(c_0^+c_m^-\cdots c_1^-)$ corresponding to curves joined at only one Reeb chord contributes with top-dimensional stratum of the boundary in the form of a product,

$$\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(c_0^+c_m^-\cdots c_j^-b^-c_{j-k}^-\cdots c_1^-)\times\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(b^+c_{j-1}^-\cdots c_{j-k+1}^-).$$

In particular, the chains of loops along the two-level boundary segments of the two-level curve are given by the Pontryagin product of the two adjacent chains of one level segments that form the two-level segment. In the two-level moduli space above, if $k>1$ there are two two-level boundary segments: the segments between c_j^- and b^- in the upper-level curve joined to the segment between b^+ and c_{j-1}^-, and the segment between c_{j-k+1}^- and b^+ in the lower level is joined to the segment between b^- and c_{j-k}^- in the upper level. If the upper-level moduli space parametrizes the chains of loops in $C_{-*}(\Omega\Lambda^{\times(m-k)})$ with components given by

$$\sigma_m c_m\sigma_{m-1}\cdots c_j\beta'_{j-1}b\beta'_{j-k}c_{j-k}\cdots\sigma_1c_1\sigma_0,$$

and the lower-level the chain in $C_{-*}(\Omega\Lambda^{\times(k-1)})$ with components given by

$$\beta''_{j-1}c_{j-1}\sigma_{j-2}\cdots\sigma_{j-k+1}c_{j-k+1}\beta''_{j-k},$$

and if $*$ denotes the constant chain and $\cdot$ the Pontryagin product, then the chain in $C_{-*}(\Omega\Lambda^{\times(m+1)})$ that contributes to the boundary has components

$$\begin{aligned}(29)\quad &(\sigma_m\cdot *)\,c_m\,(\sigma_{m-1}\cdot *)\cdots(\sigma_j\cdot *)\,c_j\,(\beta'_{j-1}\cdot\beta''_{j-1})\,c_{j-1}\,(*\cdot\sigma_{j-2})\\ &\cdots(*\cdot\sigma_{j-k+1})\,c_{j-k+1}\,(\beta'_{j-k}\cdot\beta''_{j-k})\,c_{j-k}\,(\sigma_{j-k-1}\cdot *)\cdots(\sigma_1\cdot *)\,c_1\,(\sigma_0\cdot *).\end{aligned}$$

In the case that $k=1$, the lower-level curve lies in $\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(b^+)$ and has no negative punctures. In this case the boundary contribution is

$$(30)\quad (\sigma_m\cdot *)\,c_m\cdots(\sigma_j\cdot *)\,c_j\,(\alpha'_{j-1}\cdot\beta''\cdot\gamma'_{j-1})\,c_{j-1}\,(\sigma_{j-2}\cdot *)\cdots(\sigma_1\cdot *)\,c_1\,(\sigma_0\cdot *),$$

where $\alpha'\cdot\beta''\cdot\gamma'$ denotes the chain of loops parametrized by

$$\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(c_0^+c_m\cdots b\cdots c_1)\times\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(b^+),$$

which at $(s,t)\in\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(c_0^+c_m\cdots b\cdots c_1)\times\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(b^+)$ is the loop $\alpha'(s)\cdot\beta''(t)\cdot\gamma'(s)$, where $\cdot$ denotes concatenation.

Proposition 21 *Let $d\colon \mathrm{CE}^*\to\mathrm{CE}^*$ be the map extended to CE^* by the graded Leibniz rule. Then d is a differential, $d^2=0$. We call CE^* with the differential d the* ***Chekanov–Eliashberg DG–algebra****.*

Remark 22 When $\Lambda = \Lambda^-$, CE^* was called LCA^* in [32]; this is the cohomologically graded version of the usual Legendrian homology algebra LHA_* in [11]. By definition, we have $\mathrm{LHA}_* = \mathrm{CE}^{-*}$.

Proof When there are only components in Λ^- involved, the result follows from standard arguments involving the boundary of 1–dimensional moduli spaces; see eg [28; 23; 11]. Consider therefore the case when there are chains in the loop space involved.

Let $\boldsymbol{c} = c_0^+ c_m^- \cdots c_1^-$. The d–dimensional moduli space $\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})$ contributes to dc_0. The codimension 1 strata of its boundary consists of broken curves with one level of dimension $d-k$ and one of dimension k for $0<k<d$. We find, with ∂ denoting the natural tensor extension of the boundary operator in singular homology over boundary components involved in Λ^+, that

$$\partial[\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})] = [\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}') \star \mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}'')],$$

where $\star$ is as explained above and

$$\boldsymbol{c}' = c_0^+ c_m^- \cdots c_j^- b^- c_{j-k}^- \cdots c_1^- \quad \text{and} \quad \boldsymbol{c}'' = b^+ c_{j-1}^- \cdots c_{j-k+1}^-.$$

We next apply the cubical Alexander–Whitney map η to this formula to deduce

$$\begin{aligned}\partial \circ \eta[\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})] = \eta \circ \partial[\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c})] &= \eta[\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}') \star \mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}'')] \\ &= \eta[\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}')] \cdot \eta[\mathcal{M}^{\mathrm{sy}_{\mathbb{R}}}(\boldsymbol{c}'')],\end{aligned}$$

where $\cdot$ is the Pontryagin product (see (29) and (30)) and we used that η is a chain map and is compatible with the product. The fact that η is a chain map is well known. We verify that it is compatible with the product below. It follows that the terms contributing to d^2 which arise from the differential acting on chains and acting on Reeb chords cancel.

It remains to check that η is compatible with the product. By the explicit product formulas for boundary contributions (29) and (30), we need to check that the compositions

$$\begin{aligned}&C_{-*}(\Omega_{p_u}\Lambda_u \times \Omega_{p_v}\Lambda_v) \otimes C_{-*}(\Omega_{p_v}\Lambda_v \times \Omega_{p_w}\Lambda_w) \\ &\qquad \xrightarrow{\times} C_{-*}(\Omega_{p_u}\Lambda_u \times \Omega_{p_v}\Lambda_v \times \Omega_{p_v}\Lambda_v \times \Omega_{p_w}\Lambda_w) \\ &\qquad\qquad \xrightarrow{1\times\cdot\times 1} C_{-*}(\Omega_{p_u}\Lambda_u \times \Omega_{p_v}\Lambda_v \times \Omega_{p_w}\Lambda_w) \\ &\qquad\qquad\qquad \xrightarrow{\eta} C_{-*}(\Omega_{p_u}\Lambda_u) \otimes C_{-*}(\Omega_{p_v}\Lambda_v \times \Omega_{p_w}\Lambda_w) \\ &\qquad\qquad\qquad\qquad \xrightarrow{1\otimes\eta} C_{-*}(\Omega_{p_u}\Lambda_u) \otimes C_{-*}(\Omega_{p_v}\Lambda_v) \otimes C_{-*}(\Omega_{p_w}\Lambda_w)\end{aligned}$$

and

$$\begin{aligned}
&C_{-*}(\Omega_{p_u}\Lambda_u\times\Omega_{p_v}\Lambda_v)\otimes C_{-*}(\Omega_{p_v}\Lambda_v\times\Omega_{p_w}\Lambda_w)\\
&\qquad\xrightarrow{\eta\otimes\eta} C_{-*}(\Omega_{p_u}\Lambda_u)\otimes C_{-*}(\Omega_{p_v}\Lambda_v)\otimes C_{-*}(\Omega_{p_v}\Lambda_v)\otimes C_{-*}(\Omega_{p_w}\Lambda_w)\\
&\qquad\qquad\xrightarrow{\mathbf{1}\otimes\times\otimes\mathbf{1}} C_{-*}(\Omega_{p_u}\Lambda_u)\otimes C_{-*}(\Omega_{p_v}\Lambda_v\times\Omega_{p_v}\Lambda_v)\otimes C_{-*}(\Omega_{p_w}\Lambda_w)\\
&\qquad\qquad\qquad\xrightarrow{\mathbf{1}\otimes\cdot\otimes\mathbf{1}} C_{-*}(\Omega_{p_u}\Lambda_u)\otimes C_{-*}(\Omega_{p_v}\Lambda_v)\otimes C_{-*}(\Omega_{p_w}\Lambda_w)
\end{aligned}$$

agree. This is easily checked by evaluating them on a test chain (σ,τ). Note that this uses the fact that the cubical chain complex is a quotient; namely, degenerate cubical chains are divided out. □

Remark 23 As discussed in Appendix A, the moduli spaces in the definition of the differential on CE^* above are defined in terms of anchored moduli spaces $\mathcal{M}^{\mathrm{sy}}$, ie moduli spaces of disks with additional interior punctures where holomorphic planes in the filling with asymptotic markers are attached. We point out that in order to calculate the differential one need only take into account rigid such holomorphic planes of dimension zero. For higher-dimensional moduli spaces of planes of dimension $d_0>0$, the dimension of the curves in the symplectization is $d+1-d_0$ and does not contribute to the d–dimensional chain $[\mathcal{M}^{\mathrm{sy}}]$.

Remark 24 As mentioned in the introduction, we relate CE^* as defined above to a parallel copies version of the same algebra, which is defined solely in terms of rigid moduli spaces. In order to do so it is convenient to use a topologically simpler but algebraically more complicated model of CE^*, defined as follows. The generating set of our algebra is extended to chains in the product $(\Omega\Lambda)^{\times(i+1)}$, where we separate the coordinate functions by Reeb chords. We define the product of two such chains by taking the Pontryagin product of the chains at adjacent factors, giving an operation

$$C_{-*}((\Omega\Lambda)^{\times(i+1)})\otimes C_{-*}((\Omega\Lambda)^{\times(j+1)})\to C_{-*}((\Omega\Lambda)^{\times(i+j+1)}).$$

See (29) and (30) for explicit formulas. The differential on this version of CE^* is then defined by the singular differential on the chain, and as

$$\Delta_i(c_0):=\sum_{\boldsymbol{c}=c_0^+c_i^-\cdots c_1^-}[\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})]$$

on Reeb chord generators. (In other words, we define it as above but disregard the diagonal approximation.) It follows from the Künneth formula that the two versions of CE^* are quasi-isomorphic.

Although the definition of CE^* given above works generally, from a computational perspective it is hard to get our hands on, as the cubical chain complexes $C_{-*}(\Omega_{p_v}\Lambda_v)$ have uncountably many elements.

Next, we provide a modification of the definition, which gives a quasi-isomorphic DG–algebra under the assumption that for each $v \in \Gamma^+$, there exists a DG–algebra quasi-isomorphism

$$\Phi\colon C_{-*}(\Omega_{p_v}\Lambda_v) \to \mathbb{K}\langle \mathcal{E}_v\rangle,$$

where $\mathbb{K}\langle \mathcal{E}_v\rangle$ is a DG–algebra structure on a free algebra generated by a graded finite set $\mathcal{E}_v$. For example, as discussed in Section 3.1, such a DG–algebra map exists when Λ_v is simply connected. (Or, if Λ_v is a $K(\pi_1, 1)$ space, one can first work with the group ring $\mathbb{K}[\pi_1]$ and base-change to a cofibrant replacement of it.)

We define a graded quiver $\mathcal{Q}_\Lambda$ with vertex set $\mathcal{Q}_0 = \Gamma$ and arrows in correspondence with

$$\mathcal{Q} := \mathcal{R} \cup \bigcup_{v\in\Gamma^+} \mathcal{E}_v.$$

More precisely, there are arrows from vertex v to w corresponding to the set of Reeb chords from Λ_v to Λ_w. In addition, for each $v \in \Gamma^+$, there are arrows from v to v corresponding to the elements in $\mathcal{E}_v$. Let $\mathrm{LC}_*(\Lambda)$ be the graded $\boldsymbol{k}$–bimodule generated by $\mathcal{Q}$. Thus, there is one generator for each arrow in $\mathcal{Q}$ and an idempotent e_v for each vertex $v \in \mathcal{Q}_0$. We write $\overline{\mathrm{LC}}_*(\Lambda)$ for the submodule without the idempotents.

Let $\mathrm{CE}^*(\Lambda)$ be $\boldsymbol{k}$–algebra given by the tensor algebra

$$\mathrm{CE}^*(\Lambda) = \boldsymbol{k} \oplus \bigoplus_{i=1}^{\infty} \overline{\mathrm{LC}}_*(\Lambda)[-1]^{\otimes_{\boldsymbol{k}} i}.$$

Recall that the path algebra of a quiver is defined as a vector space having all paths in the quiver as basis (including, for each vertex v an idempotent e_v), and multiplication given by concatenation of paths. Thus, the $\boldsymbol{k}$–bimodule $\mathrm{CE}^*(\Lambda)$ is the path algebra of the quiver $\mathcal{Q}_\Lambda$, where the grading of each arrow is shifted up by 1. Just like in the cobar construction, we write elements in $\mathrm{CE}^*(\Lambda)$ as

$$[x_m|\cdots|x_1] = s^{-1}x_m \otimes_{\boldsymbol{k}} \cdots \otimes_{\boldsymbol{k}} s^{-1}x_1 \in \mathrm{CE}^*(\Lambda),$$

where $x_j \in \overline{\mathrm{LC}}_*(\Lambda)$.

Next, we equip the $\boldsymbol{k}$–algebra $\mathrm{CE}^*(\Lambda)$ with a differential using the moduli spaces of holomorphic disks defined in Appendix A. This differential is induced by operations

$$\Delta_i\colon \overline{\mathrm{LC}}_*(\Lambda) \to \overline{\mathrm{LC}}_*(\Lambda)^{\otimes_{\boldsymbol{k}} i}, \quad i = 0, 1, \ldots,$$

where $\overline{\mathrm{LC}}_*(\Lambda)^{\otimes_{\boldsymbol{k}} 0} = \boldsymbol{k}$, which give $\mathrm{LC}_*(\Lambda)$ the structure of an A_∞–coalgebra if $\Delta_0 = 0$.

Consider a generator of $\overline{\mathrm{LC}}_*(\Lambda)$. If it is a generator $\sigma \in \mathcal{E}_v$ of the free model of $C_{-*}(\Omega_{p_v}\Lambda_v)$ for some component $\Lambda_v \subset \Lambda^+$, then we define

$$\Delta_i \sigma = d_i \sigma, \tag{31}$$

where d_i is the coproduct that corresponds to i^{th} homogeneous piece of the differential in the free model $\mathbb{K}\langle \mathcal{E}_v\rangle$ of $C_{-*}(\Omega_{p_v}\Lambda_v)$. If it is a Reeb chord c_0 then we define $\Delta_i(c_0)$ as before using moduli spaces $\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})$ but now take the image of all the singular chains in $C_{-*}(\Omega_{p_v}\Lambda_v)$ under the map $\Phi\colon C_{-*}(\Omega_{p_v}\Lambda_v) \to \mathbb{K}\langle \mathcal{E}_v\rangle$. Since the map Φ is a DG–algebra map, the proof that d is a differential on CE^* is the same. Furthermore, since Φ is a quasi-isomorphism, we get a quasi-isomorphic chain complex CE^* if we use $\mathbb{K}\langle \mathcal{E}_v\rangle$ coefficients instead of $C_{-*}(\Omega_{p_v}\Lambda_v)$.

From now on, unless otherwise specified, we will always assume that we work with a free (over $\boldsymbol{k}$) model of CE^*.

If there exists an augmentation $\epsilon\colon \mathrm{CE}^*(\Lambda) \to \boldsymbol{k}$, then there is a change of coordinates which turns $\mathrm{LC}_*(\Lambda)$ into a A_∞–coalgebra. More precisely, consider the restriction of ϵ, $\epsilon_1\colon \mathrm{LC}_*(\Lambda) \to \boldsymbol{k}$, where we think of $\mathrm{LC}_*(\Lambda)$ as the degree 1 polynomials in $\mathrm{CE}^*(\Lambda)$. Define

$$\mathrm{LC}_*^{\epsilon} = \boldsymbol{k} \oplus \ker(\epsilon_1).$$

Note that $\ker(\epsilon_1)$ is generated by idempotents e_v and by $c - \epsilon(c)$, where c ranges over the generators of $\overline{\mathrm{LC}}_*(\Lambda)$. Let

$$\phi_\epsilon\colon \bigoplus_{i\ge 0} \mathrm{LC}_*^{\otimes_{\boldsymbol{k}} i} \to \bigoplus_{i\ge 0} \mathrm{LC}_*^{\otimes_{\boldsymbol{k}} i}$$

be the $\boldsymbol{k}$–algebra automorphism defined on generators as $\phi_\epsilon(c) = c + \epsilon(c)$. Define the operations $\Delta_i^\epsilon\colon \mathrm{LC}_*^\epsilon(\Lambda) \to \mathrm{LC}_*^\epsilon(\Lambda)^{\otimes_{\boldsymbol{k}} i}$ by

$$\Delta_i^\epsilon = \phi_\epsilon \circ \Delta_i \circ \phi_\epsilon^{-1}.$$

Theorem 25 *The operations $(\Delta_i)_{i\ge 1}$ satisfy the A_∞–coalgebra relations, and with these operations, $\mathrm{LC}_*^\epsilon(\Lambda)$ is a coaugmented conilpotent coalgebra.*

Proof Let d denote the differential on $\mathrm{CE}^*(\Lambda)$ and let c be a generator of $\overline{\mathrm{LC}}_*(\Lambda)$. Since ϵ is a augmentation, $\epsilon(dc) = 0$ and it follows that $\Delta_0^\epsilon = 0$. The A_∞–coalgebra relations then follow by combining the equation $d^2 = 0$ from Theorem 25 with the automorphism ϕ_ϵ.

The coaugmentation is simply the inclusion of $\boldsymbol{k}$. The fact that LC^ϵ_* is conilpotent follows from Stokes' theorem: the sum of the actions of the Reeb chords at the negative end of a disk contributing to the differential is bounded above by the action of the Reeb chord at the positive end. This gives the desired finiteness. □

Remark 26 If the original operation Δ_0 on $\mathrm{LC}_*(\Lambda)$ equals 0, then the map ϵ which takes all generators of $\overline{\mathrm{LC}}_*(\Lambda)$ to 0 is an augmentation. In this case $\mathrm{LC}_*(\Lambda)^\epsilon = \mathrm{LC}_*(\Lambda)$ by construction.

Remark 27 If there is an augmentation $\epsilon : \mathrm{CE}^*(\Lambda) \to \boldsymbol{k}$, then $\mathrm{CE}^*(\Lambda)$ can be expressed as the cobar construction of a coalgebra: by construction,

$$\mathrm{CE}^*(\Lambda) = \Omega(\mathrm{LC}^\epsilon_*(\Lambda)).$$

We next consider the $\boldsymbol{k}$–linear dual $\mathrm{LA}^*_\epsilon(\Lambda) := (\mathrm{LC}^\epsilon_*(\Lambda))^\#$ of $\mathrm{LC}^\epsilon_*(\Lambda)$. It follows from Section 2.1.3 that this is an augmented A_∞–algebra. We call it the *Legendrian A_∞–algebra*.

Remark 28 In the case $\Lambda = \Lambda^-$, it can be shown that this A_∞–algebra can be obtained from the endomorphism algebra of the augmentation ϵ in the Aug_- category of [10] by adjoining a unit to it, but is, in general, different from the endomorphism algebra in the Aug_+ category of [55].

Definition 29 Given an augmentation $\epsilon \colon \mathrm{CE}^*(\Lambda) \to \boldsymbol{k}$, we define the *completed Chekanov–Eliashberg DG–algebra* to be $\widehat{\mathrm{CE}}^*_\epsilon := \mathrm{B}(\mathrm{LA}^*_\epsilon)^\#$. The underlying $\boldsymbol{k}$–algebra is the completed tensor algebra

$$\widehat{\mathrm{CE}}^*(\Lambda) = \varprojlim_i \mathrm{CE}^*(\Lambda)/I^i = \boldsymbol{k}\langle\langle \overline{\mathrm{LC}}^\epsilon_*(\Lambda)[-1]\rangle\rangle,$$

where $\overline{\mathrm{LC}}^\epsilon_*(\Lambda)$ is the ideal determined by the natural augmentation.

Note that there is a natural chain map

$$\Phi \colon \mathrm{CE}^*(\Lambda) \to \widehat{\mathrm{CE}}^*_\epsilon(\Lambda).$$

Remark 30 To illustrate the various gradings, the unique Reeb chord for the standard Legendrian unknot Λ in $\mathbb{R}^{2n-1}$ has degree $-n$ in LC_*, n in LA^* and $-(n-1)$ in CE^* (while it is $n-1$ in LHA_*). Therefore, we have the graded isomorphisms $H_*(\mathrm{LC}) \cong H_{-*}(S^n)$, $H^*(\mathrm{LA}) \cong H^*(S^n)$ and $H^*(\mathrm{CE}) \cong H_{-*}(\Omega S^n)$.

3.3 Parallel copies

In this section we will describe the perturbation scheme we will use to define various versions of Lagrangian Floer cohomology. For exact Lagrangian submanifolds with Legendrian boundary we get an induced perturbation scheme for the Legendrian boundary that will allow us to define a simpler version of the Chekanov–Eliashberg algebra which is isomorphic to it when the Legendrian is simply connected.

Let X be a Weinstein manifold and let $L \subset X$ be an exact Lagrangian submanifold with Legendrian boundary Λ. We assume that Λ is embedded but allow L to be a several-component Lagrangian with components that intersect transversely. Assume that the components of L are decorated with signs and write L^+ and L^- for the union of the components decorated with $+$ and $-$, respectively. We will use specific families of Morse functions to shift Lagrangian and Legendrian submanifolds off of themselves in order to relate holomorphic curve theory to Morse theory, and to perform Floer cohomology calculations without Hamiltonian perturbations. Before we discuss the details of this we recall some general results for Morse flow trees.

3.3.1 General results for flow trees In this section we recall several basic results for Morse flow trees from [22]. Morse flow trees live in a neighborhood of a given Lagrangian or Legendrian and are thus defined in the corresponding cotangent bundle or the 1–jet space. In this paper, we will consider only the case of graphical Lagrangians and Legendrians, in the cotangent bundle and 1–jet space, respectively. That corresponds to a simple special case of the more general situation considered in [22], where the nearby Lagrangians and Legendrians are allowed to have singularities when projected to the zero section.

Let M be a smooth manifold with cylindrical ends of the form $\partial M \times [0,\infty)$. Consider the cotangent bundle T^*M and the 1–jet space J^1M. We consider graphical Lagrangians and associated Legendrians $\Gamma_{dF} \subset T^*M$ and $\Gamma_{j^1F} \subset J^1M$. At the ends our functions will have the form $F = e^t f(q) + c$, where $t \in [0,\infty)$, $f\colon \partial M \to \mathbb{R}$ and c is a constant. Let $L_1, \dots, L_m$ be a collection of graphical Lagrangians in T^*M and $\widetilde{L}_j$ be a Legendrian lift of L_j. As in [22, Section 2.2.2], $\widetilde{L}_j$ defines local gradients as well as cotangent and 1–jet lifts of paths in M. Furthermore, [22, Lemma 2.8] shows that there are maximal flow lines, and as in [22, Definition 2.9] we define their flow orientation. We define flow trees of $\widetilde{L} = \bigcup_j \widetilde{L}_j$ as in [22, Definition 2.10] and we will also use partial flow trees, which are flow trees with "free" 1–valent vertices, not necessarily at a critical point.

We next discuss transversality for flow trees following [22, Section 3]. There are two concepts of dimension of a flow tree Γ involved here: the formal dimension $\dim(\Gamma)$ (see [22, Definition 3.4]) and the geometric dimension $\mathrm{gdim}(\Gamma)$ (see [22, Definition 3.5]). In the graphical case considered here these can be described as follows. The formal dimension $\dim(\Gamma)$ is the dimension of the space of flow trees around a tree Γ without degeneracies (ie only trivalent internal vertices and nonzero length flow lines at positive punctures not at a minimum and negative punctures not at a maximum) assuming transverse intersections of flow manifolds at each vertex. The geometric dimension, on the other hand, is the dimension of a flow trees near Γ with fixed degeneracies (higher-valence vertices, etc). It is then clear that $\mathrm{gdim}(\Gamma) \le \dim(\Gamma)$.

We will use a transversality result that says that for generic geometric data, we have:

(FT) *Every flow tree* Γ *comes in a smooth family of dimension* $\mathrm{gdim}(\Gamma)$. *If* Γ *is degenerate then there is a natural Whitney stratification of the* $\dim(\Gamma)$*–dimensional space of flow trees around* Γ *with strata of dimension* $\mathrm{gdim}(\Gamma)$.

This result follows from [22, Proposition 3.14]. We next discuss the adaption (simplification, actually) in the current set-up of the results from [22, Section 3] that lead to [22, Proposition 3.14]. First, since all Lagrangians considered here are graphical, their front projections are smooth with empty singular locus of the front, and the preliminary transversality conditions of [22, Section 3.1.1] hold trivially. This absence of singularities also means that all the results [22, Lemmas 3.9–12] guaranteeing finitely many vertices for trees in the presence of front singularities hold automatically. Then [22, Proposition 3.14] follows readily and shows that for an open dense set of graphical Lagrangians or Legendrians, (FT) holds.

We say that a finite collection of functions $F_1, \dots, F_k$ on M is flow-tree generic provided (FT) holds. It is a consequence of [22, Proposition 3.14] that any collection of functions can be made flow-tree generic by an arbitrarily small perturbation and furthermore that if $F_1, \dots, F_{k-1}$ is already flow-tree generic then $F_1, \dots, F_{k-1}, F_k$ can be made flow-tree generic by an arbitrarily small perturbation of F_k.

3.3.2 Systems of parallel copies In this section we describe how to choose systems of parallel copies for Lagrangians and Legendrians in such a way that higher product and coproduct operations on the Morse complexes can be directly defined (without mapping telescopes of continuation maps, typically used in Hamiltonian Floer cohomology).

Let L be a Lagrangian with Legendrian boundary Λ. Then a neighborhood of L in X looks like T^*L, and along the cylindrical end $[0,\infty) \times \Lambda$, the vector tangent to T^*L in the direction of the dual of the $[0,\infty)$–direction corresponds to the Reeb direction. We consider a collection of parallel copies L_j for $j=0,1,2,\dots$, with $L=L_0$. Here L_j is the graph in T^*L of the differentials dF_j of a Morse function $F_j\colon L\to\mathbb{R}$. The Morse functions F_j will have critical points in the compact part of L and in the cylindrical ends they will look like Morsifications of the Reeb push-off; see below for details.

We next discuss the main strategy without all technicalities: the first Morse function $H_1=F_1$ gives the first parallel copy at small distance $\epsilon>0$ from $L_0=L$. We define L_1 as the graph of the differential of ϵF_1. We want all other copies to be good approximations of L_1 as seen from L_0, so that flow lines between L_j and L_0 and between L_1 and L_0 are sufficiently close that the corresponding spaces of flow lines can be canonically identified. Let $L_j^1=L_1$ for $j>1$.

We next construct $L_2=L_2^2$ as the graph of the differential of a function $\epsilon^2 H_2$ over $L_2^1=L_1$. For small $\epsilon>0$, L_2 is then well approximated by L_1 as seen from L_0, and flow lines between L_0 and L_1 can be identified with flow lines between L_0 and L_2. We also want spaces of flow lines between L_1 and L_2 to be identified with flow lines between L_0 and L_1. This holds provided H_2 is a sufficiently good approximation of H_1. Thus we take

$$H_2=F_1+\epsilon H_2=H_1+\epsilon H_2,$$

where H_2 is sufficiently close to H_1 that the following further condition holds. The Lagrangians L_0, L_1 and L_2 together also define flow trees with three punctures. We take H_2 so that (FT) holds for L_0, L_1 and L_2. It follows from [22, Proposition 3.14] that this can be achieved by an arbitrarily small perturbation of H_2.

The construction now proceeds in the same manner. First, preliminarily, set $L_j^2=L_2$, $j>2$. Then let $L_3=L_3^3$ be the graph of the function $\epsilon^3 H_3$ over $L_3^2=L_2$. In order for L_3 to look like L_1 from the point of view of L_2 and like L_2 from the point of view of L_1, we take

$$H_3=F_2+\epsilon^2 H_3=H_1+\epsilon H_2+\epsilon^2 H_3.$$

For $\epsilon>0$ sufficiently small we may then identify flow lines and flow trees of any three of the functions, and after an arbitrarily small perturbation of H_3, condition (FT) holds for L_j for $j=0,1,2,3$. Continuing like this we get

$$H_k=F_{k-1}+\epsilon^{k-1}H_k=\sum_{k=1}^{k}\epsilon^{k-1}H_k. \tag{32}$$

The corresponding collection $L_0, \dots, L_k$ of parallel copies then has the following properties: flow trees with boundary on any increasing collection $L_{i_1}, \dots, L_{i_k}$ are arbitrarily close to flow trees of $L_0, \dots, L_{k-1}$, and condition (FT) holds for $L_0, \dots, L_k$.

In order to get the system of parallel copies, we need first to set conventions for the description of the ends. Along the ends our Lagrangians L look like cylinders over Legendrians Λ. A small neighborhood of Λ in the contact boundary can be identified with a small neighborhood of the zero section in the 1–jet space of Λ. We think of this as the intersection of $(-\delta, \delta) \times T^*\Lambda$ with a small neighborhood of the zero section in the cotangent bundle factor, and the contact form is $ds - p\,dq$, where s is a coordinate on $(-\epsilon, \epsilon)$. Along this end the Lagrangian is $[0, \infty) \times \Lambda$ and the corresponding neighborhood is $(-\epsilon, \epsilon) \times [0, \infty) \times T^*\Lambda \subset T^*[0, \infty) \times \Lambda$. We observe then that the result of moving Λ ϵ units along the Reeb flow is the graph of the differential of the function $B(t, q) = \epsilon t$, $(t, q) \in [0, \infty) \times \Lambda$. We will Morsify this Bott situation by considering graphs of $F(t, q) = \epsilon(t + f(q))$, where $f(q)$ is a Morse function. Then the Reeb chords inside the neighborhood of L at infinity between the graph of F and $[0, \infty) \times \Lambda$ are in natural one-to-one correspondence with critical points of f. In this set-up, with infinite ends, there are also flow trees with positive punctures asymptotic to Reeb chords at infinity. In the compactification of the space of flow trees there are flow trees entirely in the $\mathbb{R}$–invariant end, $T^*\mathbb{R} \times \Lambda$. Along the end we have $dF = \epsilon(dt + (\partial f/\partial q)\,dq)$ and the results about Morse flow trees from Section 3.3.1 follow readily from the corresponding results for flow trees of f on Λ.

We now turn to a more detailed description of the construction of parallel copies such that flow trees of ordered subcollections of parallel copies can be identified as discussed above. Write $[0, \infty) \times Y$ and $[0, \infty) \times \Lambda$ for the ends of X and L, respectively. We use coordinates $(t, q) \in [0, \infty) \times \Lambda$. Consider a collection of pairs of Morse functions (F_j, f_j) such that $F_j : L \to \mathbb{R}$ and $f_j : \Lambda \to \mathbb{R}$, $j = 1, 2, \dots$, are related at the ends by

$$F_j(t, q) = \epsilon(t + f_j(q)) + b \quad \text{for } 1 \gg \epsilon > 0 \text{ and } b > 0,$$

and F_j does not have any local maxima. We next discuss further restrictions related to critical points.

Let (F_1, f_1) be any pair of positive Morse functions as above. Let $z_1^1, \dots, z_1^m$ be the critical points of F_1 and let $x_1^1, \dots, x_1^l$ be the critical points of f_1. Fix disjoint coordinate balls $B_j^1 \subset L$ around z_j^1 and $D_j^1 \subset \Lambda$ of x_j^1 such that F_1 and f_1 are given by quadratic polynomials in these coordinates. Fix small $\sigma > 0$ such that

$$|dF_1| > \sigma_1 = \sigma \quad \text{on } L - \bigcup_{j=1}^{m} B_1^j \qquad \text{and} \qquad |df_1| > \sigma_1 = \sigma \quad \text{on } \Lambda - \bigcup_{j=1}^{l} D_1^j.$$

Let (F_2, f_2) be another pair of positive Morse functions with m and l critical points $z_2^1, \dots, z_2^m$ and $x_1^1, \dots, x_2^l$, respectively, where

$$\begin{aligned} z_2^j &\in B_1^j \quad \text{with } \operatorname{index}(z_2^j) = \operatorname{index}(z_1^j), \\ x_2^j &\in D_1^j \quad \text{with } \operatorname{index}(x_2^j) = \operatorname{index}(x_1^j). \end{aligned}$$

Let $\sigma_2 < \sigma\sigma_1$ and fix coordinate balls $B_2^j \subset B_1^j$ and $D_2^j \subset D_2^j$ such that

$$|dF_2| > \sigma_2 \quad \text{on } L - \bigcup_{j=1}^{m} B_2^j \qquad \text{and} \qquad |df_2| > \sigma_2 \quad \text{on } \Lambda - \bigcup_{j=1}^{l} D_2^j.$$

Finally, we make sure that $F_2 < \sigma F_1$ and $f_2 < \sigma f_1$, which we obtain by overall scaling. Note that we might have to shrink σ_2 after scaling.

We continue inductively and construct a family of pairs (F_k, f_k), $k = 1, 2, \dots$ of positive Morse functions with the following properties. Each F_k has m critical points $z_k^1, \dots, z_k^m$, each f_k has l critical points $x_k^1, \dots, x_k^l$. There are $\sigma_k > 0$ and disjoint coordinate balls B_k^j around z_k^j and D_k^j around x_k^j such that

$$|dF_k| > \sigma_k \quad \text{on } L - \bigcup_{j=1}^{m} B_k^j \qquad \text{and} \qquad |df_k| > \sigma_k \quad \text{on } L - \bigcup_{j=1}^{l} D_k^j.$$

Furthermore, the following hold:

- $B_k^j \subset B_{k-1}^j$ and $D_k^j \subset D_{k-1}^j$.
- $\sigma_k \le \sigma\sigma_{k-1}$.
- $F_k < \sigma F_{k-1}$ and $f_k \le \sigma f_{k-1}$.

We next take into account the sign decoration. Assume that L^- is nonempty. In this case we first construct functions $\{(G_j, g_j)\}_{j=1}^{\infty}$ exactly as above on all components of L. The actual functions $\{(F_j, f_j)\}$ on (L, Λ) are then $(F_j, f_j) = (G_j, g_j)$ on L^+ and $(F_j, f_j) = (-G_j, -g_j)$ on L^-.

Consider now the Morse functions

$$\left(H_k = \sum_{j=1}^{k} F_j, h_k = \sum_{j=1}^{k} f_j\right).$$

(Compare (F_j, f_j) to the function $\epsilon^k F_k$ in the discussion preceding (32).) Define the system of parallel copies $L_0, \dots, L_k, \dots$ by letting L_k be the graph Γ_{dH_k} of the differential of H_k. Then we have the following:

Lemma 31 *For generic choice of functions (F_j, f_j) and all sufficiently small $\sigma > 0$ in the construction above, the resulting system of parallel copies $\{L_j\}_{j=0}^{\infty}$ has the following properties*:

- *Intersection points $L_{k_0} \cap L_{k_1}$ are transverse and are in natural one-to-one correspondence with intersection points of $L_0 \cap L_1$ (or, in terms of L only, critical points of F_1 and self-intersection points of $L = L_0$).*
- *On L_+, if $k_0 < k_1$, then Reeb chords from Λ_{k_0} to Λ_{k_1} are in natural one-to-one correspondence with Reeb chords from Λ_0 to Λ_1 (or, in terms of Λ only, critical points of f_1 and Reeb chords of $\Lambda = \Lambda_0$).*
- *For all ordered finite subcollections $L_{k_0}, L_{k_1}, \dots, L_{k_m}$ with $k_0 < k_1 < \dots < k_m$ of parallel copies, flow-tree transversality* (FT) *holds. Furthermore, for any two such ordered collections $L_{k_0}, L_{k_1}, \dots L_{k_m}$ and $L_{j_0}, L_{j_1}, \dots L_{j_m}$, the spaces of flow trees are canonically isomorphic.*

Proof Intersection points $L_{k_0} \cap L_{k_1}$ correspond to intersections between the components of L and critical points of $H_{k_1} - H_{k_0}$. Since

$$H_{k_1} - H_{k_0} = F_{k_0+1} + \mathcal{O}(\sigma F_{k_0+1}),$$

we find that, from the point of view of L_{k_0}, L_{k_1} can be viewed as a small perturbation of L_{k_0+1}. In particular, if $k_0 \neq k_1$, then $L_{k_0} \cap L_{k_1}$ is transverse and there is a unique intersection point near each intersection point in $L \cap L_1$ which corresponds to critical points of F_1 and self-intersections of L. The statement on Reeb chords follows similarly. The last statement follows from the special case of [22, Proposition 3.14] as described above. $\square$

Remark 32 In Sections B.2 and B.3, we will also apply this construction to Lagrangian submanifolds $C \subset W$ where W is a Weinstein cobordism with both positive and negative ends. Our Lagrangian submanifolds C that will be equipped with systems of parallel copies will, however, have only positive ends in that case, and the above discussion applies without change. See Remark 65 for a version when both positive and negative ends are perturbed.

3.3.3 Holomorphic disks and flow trees In this section we discuss results relating holomorphic disks and Morse flow trees that are used in computations with parallel copies.

Let $L \subset X$ be a Lagrangian with cylindrical end $\mathbb{R} \times \Lambda \subset \mathbb{R} \times Y$ and let $\overline{L}(\sigma) = \{L_j(\sigma)\}_{j=0}^{\infty}$ be a system of parallel copies for L constructed as in Section 3.3, where $\sigma > 0$ is the scaling parameter. (Roughly, $L_0 = L$, and for $k = 1, 2, 3, \dots$ $L_k(\sigma)$ is at distance σ^k from $L_{k-1}(\sigma)$.)

We first consider the relation between local holomorphic disks and Morse flow trees. We have the following result for words $\boldsymbol{a}$ and $\boldsymbol{c}$ of Reeb chords and intersection points corresponding to critical points of the shifting functions (F_1, f_1); see Section 3.3.

Lemma 33 *If $\overline{L}$ is flow-tree generic and $\kappa = (\kappa_0, \kappa_1, \dots, \kappa_m)$ is an increasing (or decreasing) boundary numbering, then for all $\sigma > 0$ sufficiently small there is a natural one-to-one correspondence between rigid holomorphic disks in $\mathcal{M}^{\mathrm{fi}}(\boldsymbol{a}, \kappa)$ and rigid flow trees of $L_{\kappa_0}, \dots, L_{\kappa_m}$ with asymptotics according to $\boldsymbol{a}$: there is a neighborhood of the cotangent lift of each rigid tree that contains the boundary of a unique rigid holomorphic disk that is transversely cut out, and each rigid disk has boundary in the neighborhood of some rigid tree. Similarly, there are natural one-to-one correspondences between rigid disks in $\mathcal{M}^{\mathrm{co}}(\boldsymbol{c}; \kappa)$ and $\mathcal{M}^{\mathrm{sy}}_{\|}(\boldsymbol{c}, \kappa)$ and corresponding rigid flow trees determined by $L_{\kappa_0}, \dots, L_{\kappa_m}$.*

Proof This is a consequence of the main results in [22], namely Theorems 1.2 and 1.3, which show, for compact L, as $\sigma \to 0$, that any sequence of rigid disks converges to a rigid flow tree ("compactness") and also that near any rigid flow tree in the limit there is a unique rigid holomorphic disk for all sufficiently small $\sigma > 0$ ("gluing"). It is essential for this one-to-one correspondence to hold that there be no multiply covered disks. In the present case the increasing (or decreasing) condition guarantees no disk is multiply covered. The modifications necessary for the case of cylindrical ends and corresponding Lagrangians of the form $\Lambda \times \mathbb{R} \subset T^*\Lambda \times \mathbb{R}$ are straightforward; see eg [29] for flow tree results in a related setting. □

The second result concerns a mixed picture where the disks do not lie entirely in the cotangent bundle. In this case holomorphic disks on a system of parallel copies admit a description with holomorphic disks on the underlying Lagrangian with flow trees attached along their boundaries. Such configurations were considered and the main correspondence was worked out in the setting of knot contact homology in [26].

Consider a Lagrangian L and a system of parallel copies $\overline{L} = \{L_j\}_{j=0}^{\infty}$ as above, and let κ be an increasing (or decreasing) boundary decoration. A quantum flow tree of $\overline{L}$ is a finite collection of holomorphic disks $D_1, \dots, D_m$ with boundaries on subdecorations $\kappa^1 \subset \kappa^m$ of κ with flow trees emanating from their boundaries with boundary subdecorations $\theta^1, \dots, \theta^n$, such that inserting flow-tree domains in the disks gives a disk, and such that inserting the cotangent lifts of the flow tree at the insertion points we get a boundary condition respecting the decoration κ.

In order to establish the desired correspondence between rigid quantum flow trees and rigid holomorphic disks, we need additional transversality properties of the shifting Morse function that controls the interface of holomorphic disks and flow trees. The argument is the following. Start with a system of parallel copies that satisfies flow-tree transversality, and perturb the almost complex structure so that moduli spaces of holomorphic disks with decreasing (or increasing) boundary decoration (that cannot be multiply covered) are transversely cut out. Then perturb the shifting Morse functions slightly so that partial flow trees are transverse to the boundary evaluation maps of the transversely cut out holomorphic curves. We say that parallel copies and almost complex structures with this transversality property are quantum flow-tree transverse. Arguing as for flow-tree transversality it is straightforward to show that flow-tree transversality holds after arbitrarily small perturbation of the shifting Morse functions.

Remark 34 One feature of using parallel copies that all approximate a single push off is that, for a rigid configuration of quantum disks, disk components can be connected only by Morse flow lines (not trees). To see this, consider a rigid configuration with three disks connected by a tree with a trivalent vertex. As $\sigma \to 0$ the tree converges to a flow line and all three disks have to intersect it. This will generically not happen for a rigid configuration, ie such configurations can appear only when the formal dimension is at least 1.

Lemma 35 *If $\overline{L}$ is quantum flow-tree generic and $\kappa = (\kappa_0, \kappa_1, \dots, \kappa_m)$ is an increasing (or decreasing) boundary numbering, then for all $\sigma > 0$ sufficiently small there is a natural one-to-one correspondence between rigid holomorphic disks in $\mathcal{M}^{\mathrm{fi}}(\boldsymbol{a}, \kappa)$ and rigid quantum flow trees of $L_{\kappa_0}, \dots, L_{\kappa_m}$ with asymptotics according to $\boldsymbol{a}$: there is a neighborhood of the cotangent lift of each rigid quantum tree that contains the boundary of a unique rigid holomorphic disk that is transversely cut out, and each rigid disk has boundary in the neighborhood of some rigid quantum tree. Similarly, there are natural one-to-one correspondences between rigid disks in $\mathcal{M}^{\mathrm{co}}(\boldsymbol{c}; \kappa)$ and $\mathcal{M}^{\mathrm{sy}}_{\parallel}(\boldsymbol{c}, \kappa)$ and corresponding rigid quantum flow trees determined by $L_{\kappa_0}, \dots, L_{\kappa_m}$.*

Proof As for flow trees there are two main ingredients: "compactness", ie as $\sigma \to 0$, any sequence of rigid disks converges to a rigid quantum flow tree, and "gluing", ie near any rigid flow tree in the limit there is a unique rigid holomorphic disk for all sufficiently small $\sigma > 0$. The technically most difficult point is gluing. It is present already in the flow tree case in [22] discussed above. The full correspondence was worked out with all details in the case of knot conormals in [26, Sections 5.3 and 5.4]. The case considered here can be established in the same way, as follows.

For compactness, the first step is straightforward: if the positive puncture of a disk maps to a Morse Reeb chord of length $\mathcal{O}(\eta)$, then the whole disk lies in a small neighborhood of the Lagrangian, there is no holomorphic disk part, and the correspondence between disks and flow trees in [22] applies; compare [26, Lemma 5.7]. The second step is to show, via an action/area argument, that for any sequence of disks there are neighborhoods of the punctures mapping to short Reeb chords where the disks converge to flow trees. This argument follows the usual steps in flow-tree convergence once segments of action $\mathcal{O}(\eta)$ near such punctures have been found; see [26, Lemma 5.8 and Corollary 5.9]. The third step uses the fact that there is only one positive puncture to show that there can be only one big disk component in the limit; see [26, Lemma 5.11]. The final step is to show that the flow-tree limits that end at punctures meet the big disk in the limit. This follows from an action argument; see [26, Lemma 5.13]. This establishes flow-tree convergence. (Although the arguments in [26, Section 5.3] are written in the case where the Lagrangian is 2–dimensional, the arguments work unchanged in any dimension n.)

For gluing, the first step is to arrange, by standard transversality arguments for curves and perturbation of Morse functions, transversality of the Morse flow data and evaluation maps of holomorphic curves; see [26, Section 5.4.1]. Then there are finitely many rigid flow-tree configurations. The metric and Lagrangian is adapted near the flow tree parts and the points where they meet the boundary of the disk so that there are explicit holomorphic curves near large parts of the flow trees. Flow tree parts and big disk parts are joined over finite regions in the domain and we construct a weight function which is of size 1 in the finite joining regions and exponentially growing along the parts of the domain where we have solutions. The Floer gluing scheme is applied in this setting and surjectivity of gluing is established; see [26, Section 5.4.3]. Again, [26] works with a 2–dimensional Lagrangian, but most of the arguments above are dimension independent and local models generalize to general dimension in a straightforward way. □

Remark 36 In the setting of Lemma 35, the action integral $\int d\alpha$ of a holomorphic disk in $\mathcal{M}^{\mathrm{sy}}_{\|}(\boldsymbol{c},\kappa)$ is positive since the almost complex structure J is compatible with the contact form. By Stokes' theorem the action integral equals the difference between the action of the Reeb chord at the positive puncture and the sum of the actions at the negative punctures. It follows in particular that if the positive puncture is a Morse chord, then all negative punctures are Morse chords as well and the moduli spaces are controlled already by Lemma 33.

3.4 Chekanov–Eliashberg algebra with parallel copies

We will relate the Legendrian invariants LA^* and LC_* to Floer (co)homology of exact Lagrangian submanifolds. When studying Lagrangian Floer (co)homology we employ the technique of parallel copies. In this setup no holomorphic disk under consideration is multiply covered, and transversality is achieved by perturbation near punctures as in [28, Lemma 4.5]; see Theorem 74 for an outline of the argument. It will be convenient to express the Legendrian invariants in the same language. As it turns out, in the case that the Legendrians are simply connected this technique leads to a simpler formulation of the theory which incorporates a model of chains on the based loop space automatically. Recall that Lagrangian fillings of Legendrian submanifolds induce augmentations, which after a change of variables lead to noncurved Legendrian A_∞–coalgebras. Geometrically, this means that one uses anchored holomorphic disks. We will assume that an augmentation of Λ has been fixed in this section and all disks considered will be anchored with respect to this augmentation. We now turn to the description of this theory in the Legendrian setting.

Let Λ be as above with decomposition $\Lambda = \Lambda^+ \sqcup \Lambda^-$. Fix a Morse function $f\colon \Lambda \to \mathbb{R}$ which is positive on Λ^+ and negative on Λ^-. Use it as described in Section 3.3 to construct a system $\overline{\Lambda} = \{\Lambda_j\}_{j=1}^{\infty}$ of parallel copies of $\Lambda = \Lambda_0$.

Let $\mathcal{Q}_0$, $\boldsymbol{k}$ and $\mathcal{R}$ be as above. Let $\mathcal{R}^+$ denote the Reeb chords connecting Λ_0 to Λ_1 that lie in a small neighborhood of Λ_0. By construction there is then a natural one-to-one correspondence between $\mathcal{R}^+$ and the set of critical points of f on Λ^+. (Since f shifts Λ^- in the negative Reeb direction there are no such short chords near Λ^-.) Write

$$\mathcal{R}_{\|} = \mathcal{R} \cup \mathcal{R}^+,$$

and think of chords in $\mathcal{R}^+$ as connecting a component Λ_v to itself. Again the set $\mathcal{R}_{\|}$ is a graded set: Reeb chords $c \in \mathcal{R}$ are graded as above, $|c| = -\mathrm{CZ}(c)$, and the grading of a short chord $c \in \mathcal{R}^+$ equals the negative of the Morse index of the critical point of f corresponding to c.

We define a graded quiver $\mathcal{Q}_{\parallel,\Lambda}$ with vertex set $\mathcal{Q}_0 = \Gamma$ and arrows in correspondence with

$$\mathcal{Q}_{\parallel} := \mathcal{R}_{\parallel}.$$

More precisely, there are arrows from vertex v to w corresponding to the set of Reeb chords from Λ_v to Λ_w if $v \neq w$, and corresponding to short Reeb chords from $\Lambda_v = \Lambda_{v0}$ to Λ_{v1} if $v = w$.

Let $\mathrm{LC}^{\parallel}_*(\Lambda)$ be the graded $\boldsymbol{k}$–bimodule generated by $\mathcal{Q}_{\parallel}$. We define an A_∞–coalgebra structure on $\mathrm{LC}^{\parallel}_*(\Lambda)$ given by operations Δ_i as follows. Given a chord c_0 (input) and chords $c_i, \dots, c_1$ (outputs), we consider the disk D_{i+1} with distinguished puncture at c_0 and a *strictly decreasing* boundary decoration κ. Let $\boldsymbol{c} = c_0^+ c_i^- \cdots c_1^-$. Consider the moduli space $\mathcal{M}^{\mathrm{sy}}_{\parallel}(\boldsymbol{c};\kappa)$. We write $|\mathcal{M}^{\mathrm{sy}}_{\parallel}(\boldsymbol{c};\kappa)|$ for the algebraic number of $\mathbb{R}$ components in this moduli space provided $\dim(\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c};\kappa)) = 1$, and $|\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c};\kappa)| = 0$ otherwise. Define, for $i > 0$,

$$\Delta_i(c_0) := \sum_{\boldsymbol{c}=c_0^+ c_i^- \cdots c_1^-} |\mathcal{M}^{\mathrm{sy}}_{\parallel}(\boldsymbol{c};\kappa)|\boldsymbol{c}',$$

where $\boldsymbol{c}' = c_i \cdots c_1$. This gives an operation of degree $2-i$ on $\mathrm{LC}^{\parallel}_*(\Lambda)$. Note that $\Delta_0 = 0$ trivially, since the decoration κ is strictly increasing.

It is not a priori clear that the maps Δ_i are well-defined, as the sum may involve infinitely many terms. This can be avoided easily if Λ_v are simply connected for $v \in \Gamma^+$. Namely, the simple connectedness of Λ_v guarantees that the short chords cost index (recall the grading of a Morse chord of Morse index p is graded by $-p$, which means that its input in the dimension formula is $-p+1$, and in the simply connected case $-1 \le -p+1 \le -\dim(\Lambda)+1$), whereas the long chords cost energy. In the non-simply-connected case, we will only consider a completed version of $\mathrm{LC}^{\parallel}_*$ where infinite expressions are allowed. We define the parallel copies algebra first in the simply connected case and turn to the non-simply-connected case in Section 3.4.2.

3.4.1 The parallel copies DG–algebra in the simply connected case When the operations Δ_i defined above have suitable finiteness properties they define a coalgebra. More precisely, we have the following:

Lemma 37 *Suppose Λ is simply connected or, more generally, that Δ_i factorizes through the natural inclusion*

$$\bigoplus_{i=1}^{\infty} \mathrm{LC}_*^{\parallel}[-1]^{\otimes_\kappa i} \to \prod_{i=1}^{\infty} \mathrm{LC}_*^{\parallel}[-1]^{\otimes_\kappa}$$

for all $i \geq 1$. Then $\mathrm{LC}_*^{\parallel}$ equipped with the operations $(\Delta_i)_{i\geq 1}$ is an A_∞–coalgebra.

Proof Recall that the moduli spaces involved in the definition of the operations Δ_i are moduli spaces of anchored disks, where we use the augmentation induced by the filling L. Consider a 1–dimensional moduli space $\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c};\kappa)$ and note that its boundary consists of two-level rigid disks by Theorems 75 and 74. There are two cases, either the levels are joined at a chord connecting the same copy of Λ to itself, or the chord connects distinct copies. Since we count anchored disks, the first type of breaking cancels algebraically; compare with [23, Lemma B.6; 29, Section 1.2; 24, Section 3.4; 6, Section 6.1]. (In the usual treatment of Chekanov–Eliashberg algebras this is the statement that augmentations are chain maps.) Theorem 76 then shows that there are canonical identifications between moduli spaces for different increasing numberings and the proof follows since the breakings of the second type are exactly what contribute to the coalgebra relations and are then in algebraic one-to-one correspondence with the endpoints of an oriented compact 1–manifold. □

We will define the parallel copies version of the Chekanov–Eliashberg algebra as the reduced bar construction of the coalgebra above. To simplify matters we show that for a generic system of parallel copies, the coalgebra has a strict counit. (For more general systems of parallel copies one can instead use the definition in Section 3.4.2.) Consider a generator of $\mathrm{LC}_*^{\parallel}$ which is a small Reeb chord z that corresponds to a critical point of the Morse function f. The coalgebra operations $\Delta_i(z)$ then count holomorphic disks with positive puncture at z which, by an action argument, must lie in a small neighborhood of $\mathbb{R}\times\Lambda$ and rigid such holomorphic disks are in natural correspondence with rigid Morse flow trees; see Remark 36. In the simply connected case, ie for Morse functions without critical points of index 1 and $n-1$, we have the following result:

Lemma 38 *Suppose Λ is simply connected. If the Morse functions for parallel copies described above are sufficiently close to the first function (ie if $\epsilon > 0$ in the construction of shifting copies is sufficiently small) then, if x_v is the minimum of the Morse function on the component Λ_v, the following holds: $\Delta_1(x_v) = 0$, $\Delta_2(x_v) = x_v \otimes x_v$*

and $\Delta_i(x_v) = 0$ for $i > 2$. Furthermore, if c is any other generator corresponding to an arrow from v to w, then $\Delta_2(c) = c \otimes x_v + (-1)^{|c|} x_w \otimes c + D_2(c)$, where $D_2(c)$ does not contain any x_v factor, and $\Delta_i(c)$ does not contain any factor x_v for $i \neq 2$.

Proof Consider a system $\overline{\Lambda}(\sigma) = \{\Lambda_j(\sigma)\}$ of parallel copies as in Section 3.3.2. Note that x_v is a Morse chord of action $\mathcal{O}(\sigma)$. Hence, disks with one positive puncture at x_v can have only other Morse chords as negative punctures. For sufficiently small $\sigma > 0$, Lemma 33 then shows we can compute $\Delta_i(x_v)$ by counting all flow trees with a positive puncture at x_v and Lemma 31 shows that the flow trees are independent of increasing boundary decoration. The equation $\Delta_1(x_v)$ follows since there is no (negative) gradient flow line emanating from a minimum. For the equation $\Delta(x_v) = x_v \otimes x_v$ we consider three copies L_0, L_1 and L_2 and observe that there is a unique flow tree with positive puncture at x_v and two negative punctures at x_v, this flow tree consists simply of two flow lines starting at the minimum chord connecting L_0 to L_2 and ending at the minimum chords connecting L_0 to L_1 and L_1 to L_2, respectively.

To see the equations $\Delta_i(x_v) = 0$ for $i > 2$, we start from a general limiting argument for flow trees of parallel copies. Consider a flow tree with positive puncture at a Morse chord a and negative punctures at Morse chords $b_1, \dots, b_m$. As we take the limit $\sigma \to 0$, all shifting functions approach multiples of the same Morse function and the flow tree limits to a broken flow line starting at a connecting to b_{i_1}, then from b_{i_1} to b_{i_2}, continuing in this way until all negative punctures have been met.

Consider now a tree with positive puncture at x_v. In the limit this converges to a flow line emanating from x_v, which must then be constant. This shows that all negative punctures must be x_v as well. Since the dimension of a tree with positive puncture at x_v and i negative punctures at x_v is $i-2$, it follows that $\Delta_i(x_v) = 0$ if $i > 2$.

We next consider the properties of $\Delta_i(c)$ for $c \neq x_v$. The flow trees contributing $c \otimes x_v + (-1)^{|c|} x_w \otimes c$ to $\Delta_2(c)$ are easily found. Considering three parallel copies L_0, L_1 and L_2, the flow trees consist of a single flow line from either one of the endpoints of the chord c to the minimum x_v or x_w of the corresponding component.

We next show that these are the only contributions with negative punctures at minima. We start in the case when c is a Morse chord. Consider a tree with positive puncture at some Morse chord c which does not limit to a single flow line to the minimum as $\sigma \to 0$, and assume that one of the negative punctures is x_v. Consider first the case when x_v is at the last puncture corresponding to the smallest function difference. Assume that the negative punctures are x_v and a word of punctures b. Then $\|c\| - \|b\| - \|x_v\| - 1 = \|c\| - \|b\| = 0$, where $\|y\| = \operatorname{index}(y) - 1$. Consider the limit when this function difference goes to zero. Then the flow tree goes to a flow tree with a flow line to x_v attached. The remaining flow tree has dimension $\|c\| - \|b\| - 1 = -1$ and hence does not exist by the flow-tree transversality condition (FT) in Section 3.3.1 for the subset of parallel copies obtained by forgetting the last copy.

Consider next the case that x_v is at some other function difference. Then we have negative punctures b before x_v and a after x_v and thus a are Morse chords of smaller function differences. Consider the limit when all these smaller function differences shrink. In the limit we find a flow tree with negative punctures at (b, x_v) with a partial flow tree with negative punctures at a attached. The evaluation dimension of the latter tree (ie the dimension of the partial tree with a free positive puncture) is

$$(n-1) - \|a\| < n-1,$$

where we use the simple connectedness to get strict inequality. Applying the degeneration above to the remaining tree we get a tree with a flow line to x_v attached, and its evaluation dimension is

$$\|c\| - \|b\|.$$

Now, $\|c\| - \|b\| - \|a\| - 1 = -1$ so these two trees do not meet by condition (FT) in Section 3.3.1.

The remaining possibility is that the small tree with negative punctures a intersects the flow line towards x_v. However, such a tree can be viewed as the original partial tree merging with a flow line from the minimum and then continuing. For $\sigma > 0$, at the scale of the tree with punctures c and b (ie σ^k for some k) the flow line from the minimum and the flow at the positive puncture of the partial tree attached are very close to parallel (nonparallel only at order σ^{k+l} for $l > 0$), therefore the evaluation map at the positive puncture of the partial tree with negative punctures (x_v, a) is arbitrarily close to the evaluation map of the original partial tree with negative punctures a and taking the limit $\sigma \to 0$, the dimension count $\|c\| - \|b\| - \|a\| - 1 = -1$ above shows that these do not intersect if (FT) in Section 3.3.1 holds and $\sigma > 0$ is sufficiently small.

We finally consider the case when c is not a Morse chord, and a disk which in the limit $\sigma \to 0$ does not converge to a constant disk with a flow line attached and which has a puncture at x_v. Such a disk must have a nonconstant disk component in the limit and by Lemma 35 it converges to this disk with flow trees attached in the limit. Let b denote the negative punctures of the disk in the limit and a all Morse chord negative punctures except x_v. If a flow line to x_v is directly attached to the disk then the dimension of the quantum flow tree obtained by removing this flow line is $\|c\| - \|b\| - \|a\| - 1 = -1$ and hence it does not exist by quantum flow-tree transversality; see Lemma 35. If this is not the case then x_v is one of the negative punctures in a flow tree attached to the disk. Now that partial flow tree with positive puncture constrained to the evaluation map of the disk must be rigid, and arguing exactly as for the trees above, we see that quantum flow-tree transversality shows that no such configuration exists for sufficiently small $\sigma > 0$. □

Remark 39 The simple connectedness is used in the above proof to ensure that cutting with a small tree really reduces dimension. Here, cutting means intersecting and starting a flow from the intersection locus. In the case that there are index 1 critical points one could have $|a| = 0$ in the above, and indeed there are trees with arbitrarily many punctures at index 1 critical points and then a puncture at x_v.

Lemma 38 shows that, in the simply connected case, there is a strict coaugmentation

$$\eta\colon \boldsymbol{k} \to \boldsymbol{k}_- \oplus \mathrm{LC}_*^{\|}, \quad \text{with } \eta(e_v) = x_v, \tag{33}$$

where η is defined by

$$\eta(e_v) = \begin{cases} x_v & \text{if } \Lambda_v \subset \Lambda^+, \\ e_v & \text{if } \Lambda_v \subset \Lambda^-. \end{cases} \tag{34}$$

Definition 40 If Λ is simply connected, the parallel copies Chekanov–Eliashberg DG–algebra is

$$\mathrm{CE}_{\|}^* = \Omega(\boldsymbol{k}_- \oplus \mathrm{LC}_*^{\|}).$$

3.4.2 The parallel copies DG–algebra in the non-simply-connected case In the non-simply-connected case, the operations Δ_i defined counting holomorphic curves are not necessarily finite. To get a workable definition we will instead start from an algebra structure on the dual LA^* of LC_*. More precisely, we proceed as follows.

Let $\mathrm{LA}_{\|}^*(\Lambda)$ be the graded $\boldsymbol{k}$–bimodule generated by $\mathcal{Q}_{\|}$. We define an A_∞–algebra structure on $\mathrm{LA}_{\|}^*(\Lambda)$ given by operations Δ_i' as follows. Given chords $c_i, \dots, c_1$ (inputs) and a chord c_0 (output), we consider the disk D_{i+1} with distinguished puncture at c_0 and a *strictly increasing* boundary decoration κ. As above, let $\boldsymbol{c} = c_0^+ c_i^- \cdots c_1^-$ and consider $\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c};\kappa)$. Define, for $i > 0$,

$$\Delta_i'(\boldsymbol{c}') := \sum_{\boldsymbol{c} = c_0^+ c_i^- \cdots c_1^-} |\mathcal{M}_{\|}^{\mathrm{sy}}(\boldsymbol{c};\kappa)| c_0,$$

where $\boldsymbol{c}' = c_i \cdots c_1$. This gives an operation of degree $2-i$ on $\mathrm{LA}^*_{\|}(\Lambda)$. Note that $\Delta'_0 = 0$ trivially, since the decoration κ is strictly increasing.

Lemma 41 $\mathrm{LA}^*_{\|}$ *equipped with the operations* $(\Delta_i)'_{i\geq 1}$ *is an* A_∞*–algebra.*

Proof This is identical to the proof of Lemma 37. □

In order to define the parallel copies DG–algebra, consider the minimum x of a Morse function on a component of Λ_+. Write $u_x \in \mathrm{LA}^*_{\|}$ for the corresponding generator. Add idempotents e_v to $\mathrm{LA}^*_{\|}$, one for each component $\Lambda_v \subset \Lambda^-$. We then get the algebra $\boldsymbol{k}_- \oplus \mathrm{LA}^*_{\|}$. Equip it with the trivial augmentation ϵ' which is the projection to $\boldsymbol{k}$. Define

$$\widetilde{\mathrm{CE}}^*_{\|} = (\mathrm{B}(\boldsymbol{k}_- \oplus \mathrm{LA}^*_{\|}))^{\#},$$

and let $\mathscr{I}$ denote the subalgebra of $\widetilde{\mathrm{CE}}^*_{\|}$ defined as the space of functionals which vanish on monomials not containing u_x for some minimum chord $x \in \Lambda$.

Lemma 42 *The subalgebra* $\mathscr{I}$ *is closed under the differential.*

Proof To see that $\mathscr{I}$ is closed under the differential we check that if u_x is an output puncture of the differential then u_x is also an input puncture. Here the output puncture is the positive puncture of the holomorphic disk. The minimum corresponds to a Morse chord which is of smaller action than any Reeb chord of Λ. This means that all negative punctures in a disk with positive puncture at x must also be Morse chords and that the corresponding disks correspond to Morse flow trees. Since there is no negative gradient flow line that starts at a minimum, any gradient flow tree with positive puncture at x must also have a negative puncture at x. This shows that u_x is an output in a differential disk only when it is also an input. □

Remark 43 In the simply connected case, $\mathrm{LA}^*_{\|} = (\mathrm{LC}^{\|}_*)^{\#}$ and there is a natural restriction map $\rho\colon \widetilde{\mathrm{CE}}^*_{\|} \to \Omega\mathrm{LC}^{\|}_*$. Since u_x are strict idempotents by Lemma 38, this is a chain map. The kernel of ρ is $\mathscr{I}$ and, consequently, $\Omega\mathrm{LC}^{\|}_* = \widetilde{\mathrm{CE}}^*_{\|}/\mathscr{I}$.

Guided by Remark 43 we define the parallel copies DG–algebra as follows, in the non-simply-connected case.

Definition 44 If Λ is not simply connected, then we define, with notation as above, the completed DG–algebra

$$\widehat{\mathrm{CE}}^*_{\|}(\Lambda) = \widetilde{\mathrm{CE}}^*_{\|}/\mathscr{I}.$$

3.5 Isomorphism between Chekanov–Eliashberg algebras in the simply connected case

We next show that if Λ is simply connected then $\mathrm{CE}^*_{\parallel}(\Lambda)$ is in fact isomorphic to $\mathrm{CE}^*(\Lambda)$. To this end we first establish a Morse-theoretic version of the Adams result mentioned in the introduction, which here corresponds to the purely local situation of the zero section in a 1–jet space.

Let Q be a *simply connected* smooth manifold with a basepoint $q \in Q$. Fix a system of positive Morse functions $\bar{f} = \{f_j\}_{j=1}^{\infty}$ as in Section 3.3, and assume that the functions have only one minimum and no index 1 critical points. (This can always be arranged by handle cancellation if $\dim(Q) \geq 5$; see Remarks 50 and 52 for the lower-dimensional case.) We will first discuss a Morse flow tree model for chains on Q, which we denote by $\mathrm{CM}_{-*}(Q)$. Our treatment of Morse flow trees follows [22]; see Sections 3.3.1 and 3.3.2. We first recall the details of the flow tree definitions from [22, Section 2] in the special case needed here.

Consider a strip $\mathbb{R}\times[0,m]$ or half-strip $[T,\infty)\times[0,m]$ with coordinates $s+i\tau$ and with $m-1$ slits along $[a_j,\infty)\times j$ for $j=1,\dots,m-1$, and $T \leq a_j$ in the half-strip case. In the half-strip case the vertical segment $T\times[0,m]$ is a finite end that will be used as an input, and we do not consider it as a part of the boundary of the strip with slits. In the strip case, the input is at the puncture $-\infty\times[0,m]$, and in both cases we call punctures at $+\infty$ "output". Order the boundary components according to the positive boundary orientation of the disk with punctures starting from the input and decorate its boundary components by a strictly increasing sequence of positive integers $\kappa_1<\kappa_2<\cdots<\kappa_m$. Let $\kappa=\{\kappa_i\}_{i=1}^m$ denote this decoration. Cutting the strip by line segments $a_j\times[0,m]$ for $j=1,\dots,m-1$, subdivides it into *strip regions* of the form $[s_0,s_1]\times[\tau_0,\tau_1]$, where $s_0\in\{-\infty,a_1,\dots,a_m\}$, $s_1\in\{a_1,\dots,a_m,\infty\}$ and $\tau_0,\tau_1\in\{0,1,\dots,m\}$, and with a numbering κ_j on each boundary component $[s_0,s_1]\times\{\tau_0\}$ and $[s_0,s_1]\times\{\tau_1\}$.

Definition 45 [22, Definition 2.10] A flow tree is a continuous map from a strip with slits into Q which in each strip region $[s_0,s_1]\times[\tau_0,\tau_1]$ depends only on the first coordinate $s\in[s_0,s_1]$, and there satisfies the gradient equation

$$\dot{x}(s) = -\nabla(f_{\kappa_i}-f_{\kappa_j})(x(s)),$$

where κ_i is the numbering of the upper horizontal boundary of the strip region and κ_j that of the lower.

A partial flow tree is defined analogously except that the domain is a half-strip with slits $[T,\infty)\times[0,m]$.

If y is a critical point of f_1 then we let $|y| = -\mathrm{index}(y)$ denote the negative Morse index of y. If $\boldsymbol{y} = y_0 y_1 \cdots y_m$ is a word of critical points of f_1 then the space of flow trees $\mathcal{T}(\boldsymbol{y})$ with input puncture at y_0 and output punctures $\boldsymbol{y}' = y_1 \cdots y_m$, in the order induced by the boundary orientation, has dimension (formal dimension in the language of Section 3.3.1)

$$\dim(\mathcal{T}(\boldsymbol{y})) = |\boldsymbol{y}'| - |y_0| + (m-2). \tag{35}$$

For a sufficiently small perturbing system of Morse functions $\bar{f}$, the space of flow trees is independent of the increasing boundary decoration κ; see Lemma 31.

Let $\mathrm{CM}_{-*}(Q)$ denote the $\mathbb{K}$–module generated by critical points of f_1 and equip $\mathrm{CM}_{-*}(Q)$ with the structure of a coalgebra with operations Δ_i given by

$$\Delta_i(y_0) = \sum_{|\boldsymbol{y}'| = |y_0| - (i-2)} |\mathcal{T}(\boldsymbol{y})|\boldsymbol{y}',$$

where the sum ranges over all $\boldsymbol{y}'$ of word length i. It is not hard to see that the boundary of a 1–dimensional space of flow trees consists of broken rigid flow trees from which it follows that the operations Δ_i satisfy the coalgebra relations; compare Lemma 37. Furthermore, by Lemma 38, the coalgebra has a natural counit, the critical point which is the minimum of f_1. We will call this critical point the *counit critical point*. We say that a flow tree with no puncture mapping to the counit critical point is *counit-free*.

The coalgebra $\mathrm{CM}_{-*}(Q)$ agrees with the Floer coalgebra $\mathrm{CF}_*(Q)$, where we view Q as the zero section in its own cotangent bundle T^*Q as follows. Let $\bar{Q}(\eta) = \{Q_j(\eta)\}_{j=0}^{\infty}$ be the system of parallel copies of the zero section $Q \subset T^*Q$ corresponding to the system of functions $\bar{f}$, where $\eta > 0$ gives the size of the perturbation, $|f_{k+1} - f_k|_{C^s} = \mathcal{O}(\eta^{k+1})$; see Section 3.3.2. Then, by Lemma 33, there is, for all sufficiently small shifts, a natural one-to-one correspondence between rigid holomorphic disks with boundary on $\bar{Q}$ and Morse flow trees in Q determined by $\bar{f}$. This gives a chain isomorphism $\mathrm{CM}_{-*}(Q) \to \mathrm{CF}_*(Q)$.

We now return to the Morse-theoretic approach to Adams' result. We will define a map

$$\phi\colon C_{-*}(\Omega Q) \to \Omega\mathrm{CM}_{-*}(Q)$$

in terms of the operation of attaching counit-free partial flow trees to Moore loops $\sigma_v : [0, r_v] \to Q$, $r_v \geq 0$, based at $q \in Q$, parametrized by a simplex $v \in \Delta$. To define this

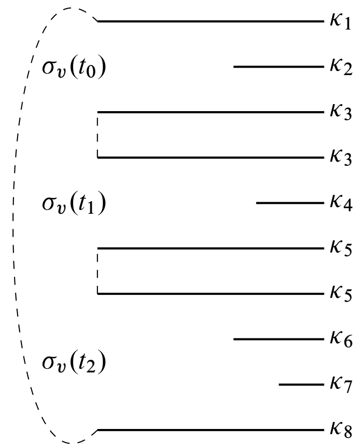

Figure 3: A configuration with three partial flow trees attached to σ_v at the points $\sigma_v(t_0)$, $\sigma_v(t_1)$ and $\sigma_v(t_2)$. The numbers on the right determine the gradient equation at that end. The dashed part represents the loop σ_v.

operation we will use the following notion: we say that a partial flow tree parametrized by a half-strip $\gamma\colon [T,\infty)\times[0,m]\to Q$ starts at a point $p\in Q$ if its input puncture maps to p, where $\gamma(T\times[0,m])=p$.

Attaching partial flow trees to $\sigma_v\colon[0,r_v]\to Q$ then means fixing points $0\le t_0\le t_1\cdots\le t_m\le r_v$ and partial flow trees Γ_j for $j=1,\dots,m$ that start at $\sigma_v(t_j)$. Our map ϕ takes values in $\Omega\mathrm{CM}_{-*}(Q)$ and will accordingly be defined by attaching flow trees which have no output at the counit.

Take the system of parallel copies $\bar{Q}(\eta)$ to be flow-tree generic (to satisfy (FT) of Section 3.3.1). Then the set of positive punctures of minimum free partial flow trees for a fixed increasing boundary numbering constitutes a codimension-two subset in Q and that the corresponding subset for any numbering lies in an $\mathcal{O}(\eta^2)$–neighborhood of it. We pick our Morse functions for the flow trees so that the basepoint q does not lie on any minimum free partial flow tree.

If $\sigma_v\colon[0,r_v]\to Q$ is a loop with flow trees attached at $0\le t_0\le t_1\le\cdots\le t_m\le r_v$, we also introduce a numbering of the components of $[0,r_v]-\{t_1,\dots,t_m\}$ induced by the flow trees attached as follows. The rightmost interval $(t_m,r_v]$ is numbered by κ_0. The right boundary segment of the strip with slits attached at $\sigma_v(t_m)$ is numbered by κ_0 as well, whereas the left boundary segment of its domain is numbered by κ_{k_m}. We number the boundary segment (t_{m-1},t_m) and the right boundary segment of the domain of the flow tree attached at $\sigma_v(t_{m-1})$ by κ_{k_m} as well. The left boundary segment of the flow tree attached at $\sigma_v(t_{m-1})$ is then numbered by $\kappa_{k_{m-1}}$, which determines the numbering of the segment (t_{m-2}, t_{m-1}) as well as the right boundary segment in the flow tree at $\sigma_v(t_{m-2})$, etc. We view the end result of this process as the domain for a loop with flow trees attached with numbering κ that decreases; see Figure 3.

Note next that if $\sigma\colon I^d \to \Omega Q$ is a d–dimensional cube in general position with respect to $\bar{Q}$ (ie transverse to the stratified space of the partial puncture of all minimum free partial flow trees), the set of σ_v with $v \in I^d$ for which a single partial flow tree can be attached is at most $(d-1)$–dimensional. Attaching more partial flow trees, the dimension decreases further, by at least one for each flow tree. We say that the loops in σ with flow trees attached which form a 0–dimensional family are the rigid loops with flow trees in σ. If σ is a cubical simplex in ΩQ and if $\boldsymbol{y}$ is a word of critical points, then we let $\mathcal{T}(\sigma; \boldsymbol{y})$ denote the space of loops with flow trees in σ, where the critical points at punctures of the flow trees read in order give the word $\boldsymbol{y}$. The formal dimension of $\mathcal{T}(\sigma; \boldsymbol{y})$ is then

$$\dim(\mathcal{T}(\sigma; \boldsymbol{y})) = |\boldsymbol{y}| + \dim(\sigma) + (\ell(\boldsymbol{y}) - 1),$$

where ℓ is the word length, and for chains transverse to the system of parallel copies the formal dimension equals the actual dimension.

Note that if the set of loops in σ with flow trees is transversely cut out, then, by construction of the system of parallel copies, loops with flow trees corresponding to different decreasing numberings are canonically diffeomorphic; see the proof of Theorem 76. We define the map ϕ by counting rigid loops with flow trees in cubical simplices σ:

$$\phi(\sigma) = \sum_{\dim(\mathcal{T}(\sigma;\boldsymbol{y}))=0} |\mathcal{T}(\sigma; \boldsymbol{y})|\boldsymbol{y}. \tag{36}$$

Remark 46 We sketch technical aspects of the definition of the map ϕ in (36). In order to get a suitable chain model for ΩQ on which the map ϕ is defined we equip Q with a Riemannian metric and consider piecewise smooth loops in Q. It is shown in [52, Section 17] that the inclusion of piecewise smooth loops into all loops is a homotopy equivalence and we will work with piecewise smooth loops. In [52, Section 16] it is shown that if E is the energy functional on the space of piecewise smooth loops, then the preimage $E^{-1}(b)$ for any noncritical value b is compact and is a deformation retract of the corresponding subset of the space of piecewise geodesic loops, which has a natural structure of a finite-dimensional manifold. Furthermore, as we increase the energy level the spaces of piecewise geodesics are naturally included in piecewise geodesics with finer subdivision. In this way we get a directed system of inclusions

$$M_{E_0} \subset M_{E_1} \subset \cdots \subset M_{E_j} \subset \cdots$$

such that $E^{-1}(b) \cap M_{E_j} \subset E^{-1}(b) \cap M_{E_k}$ for $b < E_j < E_k$ is a deformation retract. To define the map ϕ we can, for example, use chains of simplices in a sufficiently fine triangulation of the finite-dimensional manifolds M_{E_j} that are suitably transverse to the system of parallel copies $\overline{Q}(\eta)$.

Since the shifting Morse functions do not have any index 1 critical points a partial flow tree has at most $\dim(M)$ punctures. Consider the natural evaluation map on partial flow trees that takes a flow tree to the location of its positive puncture discussed above. The image of this map for partial flow trees not involving the minimum is a stratified space of codimension two and by construction of parallel copies, the corresponding set for partial flow trees defined by distinct boundary numberings lie $\mathcal{O}(\eta^2)$–close to each other. The map ϕ above is defined for chains in M_{E_j} (chains of piecewise smooth loops) with evaluation maps that are transverse to this stratified subset. It is straightforward to see that the chains of simplices in M_{E_j} can be made transverse without destroying transversality for chains at earlier energy levels. This means that we can define the map on the direct limit of chains which is a chain model for the based loop space.

In order to connect this to the path loop fibration we consider a similar map

$$\widehat{\phi}\colon C_{-*}(\mathrm{P}Q) \to \Omega\mathrm{CM}_{-*}(Q) \otimes^{\mathfrak{t}} \mathrm{CM}_{-*}(Q),$$

where $\mathfrak{t}$ denotes the canonical twisting cochain of the cobar construction and $\mathrm{P}Q$ the based path space. This map can be described geometrically as follows. The chain complex $C_{-*}(\mathrm{P}Q) \to \Omega\mathrm{CM}_{-*}(Q) \otimes^{\mathfrak{t}} \mathrm{CM}_{-*}(Q)$ can be thought of as generated by words of critical points in which the last letter is distinguished and may be the minimum x; in other words, the words are either minimum free, or the last letter (and only the last) is the minimum. The differential counts rigid flow trees as usual and also here only the last letter may be the minimum. To define the map $\widehat{\phi}$ we consider chains of paths. As above we attach counit-free partial flow trees to paths in such a chain at interior points and also attach a partial flow tree with last puncture distinguished at the endpoint of the path. Also here, only the distinguished (ie the last puncture in the tree at the endpoint of the path) may be the minimum. The map $\widehat{\phi}$ then counts rigid paths with flow trees attached as described.

Lemma 47 *The maps ϕ and $\widehat{\phi}$ are chain maps.*

Proof For the chain map property of $\widehat{\phi}$, we consider 1–dimensional moduli spaces of chains of paths with flow trees attached as described above, including flow trees at the endpoint. This moduli spaces form oriented 1–manifolds with a natural compactification consisting of the following 0–dimensional configurations:

(i) spaces of paths with flow trees attached which is obtained by attaching trees to the boundary of the original chain of paths, and

(ii) spaces of paths with flow trees attached where one of the flow trees is broken.

The configuration (i) contributes to the composition $\widehat{\phi} \circ d$ and the configuration (ii) to $d \circ \widehat{\phi}$. Since the algebraic number of boundary points of a 1–dimensional oriented manifold equals zero we conclude the chain map equation, $d \circ \widehat{\phi} = \widehat{\phi} \circ d = 0$.

The chain-map property of ϕ is proved applying the same argument to 1–dimensional spaces of chains of loops with flow trees attached. □

Remark 48 The codimension-one boundary of $\mathcal{T}(\sigma; \boldsymbol{y})$ corresponds either to the loop or path moving to the boundary of σ or to the breaking of a flow tree at a critical point. Instances when two trees are attached at the same point are naturally interior points of the moduli space where the disks with slits join to a new disk with a slit of width equal to the sum of the widths. See Figure 4.

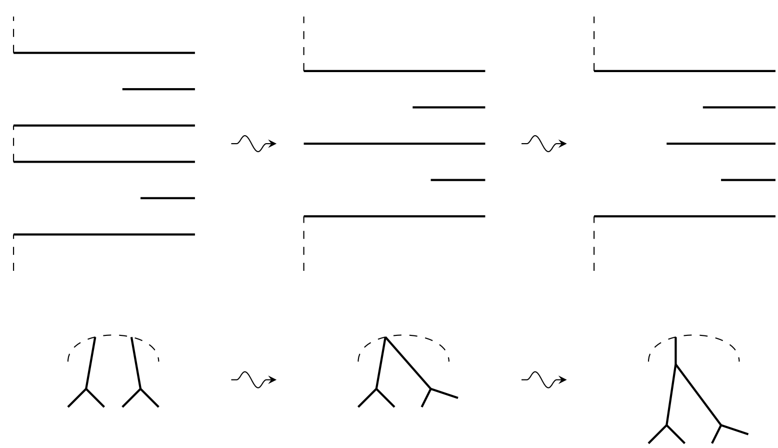

Figure 4: Flow trees attached at the same point are interior points: in the source, top, and in the target, bottom.

With this established we can now prove Adams' result:

Theorem 49 *The flow tree map $\phi\colon C_{-*}(\Omega Q) \to \Omega\mathrm{CM}_{-*}(Q)$ is a quasi-isomorphism.*

Proof The first observation is that the chain complex $\Omega\mathrm{CM}_{-*}(Q) \otimes^{\mathrm{t}} \mathrm{CM}_{-*}(Q)$ is acyclic. To see that, note that for each critical point y there are exactly two flow trees with positive puncture y and two negative punctures, one at the counit x and one at y; see Lemma 38. Exactly one of these — the tree in which the negative punctures are in the order y followed by x — contributes to the differential on $\Omega\mathrm{CM}_{-*}(Q) \otimes^{\mathrm{t}} \mathrm{CM}_{-*}(Q)$.

Add a constant to the Morse function f used to build the parallel copies so that the minimum x lies at level 0 and all other critical points at positive levels. We then filter by action of the parallel copies $\bar{f} = \{f_j\}$; more precisely, we associate to a word $y_1 \cdots y_m$ of critical points the action $\sum_j (f_j - f_{j-1})(y_j)$. Then, by definition of a flow tree for the flow-tree map, the differential does not increase action. (To see this, recall that the flow-tree map uses flow trees with only one positive puncture which decrease action since the value of a Morse function decreases along negative gradient flow; see [22, Lemma 2.3, equations (2-2)–(2-3)] for the calculation.) Since all flow trees except those involving the counit decrease action, we find that the differential on the associated graded complex acts only on the last (distinguished) letter in the word and it acts there as $y \mapsto yx$ if $y \neq x$ and $x \mapsto 0$. Since this is an isomorphism from words not ending with the counit x to those which end with x, the desired acyclicity follows. Clearly, $C_{-*}(\mathrm{P}Q)$ is also acyclic.

Consider next the stratification of Q induced by the stable manifolds of the Morse function f and the corresponding filtration on $C_{-*}(\mathrm{P}Q)$ induced by evaluation at the endpoint. The corresponding filtration on $\Omega\mathrm{CM}_{-*}(Q) \otimes^{\mathrm{t}} \mathrm{CM}_{-*}(Q)$ is filtration by degree of the distinguished (the last) critical point and the map $\hat{\phi}$ respects these filtrations.

The corresponding E_1–terms with induced maps are

$$C_{-*}(Q; H_{-*}(\Omega Q)) \to H_*(\Omega\mathrm{CM}_{-*}(Q)) \otimes \mathrm{CM}_{-*}(Q).$$

To see this, consider first the left-hand side. The associated graded complex can be represented by chains of paths in Q with endpoint in a Morse stratum of fixed dimension, divided by such chains of paths with endpoints in Morse strata of lower dimension, and the differential in the associated graded complex acts as the singular differential on such paths. The resulting homology is then the homology of the chain with endpoints in cells of fixed dimension modulo their boundary, which gives $C_{-*}(Q; H_*(\Omega Q))$.

Consider the right-hand side: here the associated graded complex can thought of as the direct sum of complexes with fixed distinguished last critical point. The differential attaches minimum-free flow trees at all nondistinguished critical points. The resulting homology is clearly $H_*(\Omega\mathrm{CM}_{-*}(Q)) \otimes \mathrm{CM}_{-*}(Q)$.

The E_2–terms are then

$$H_{-*}(Q; H_{-*}(\Omega Q)) \to H_{-*}\big(Q; H_*(\Omega\mathrm{CM}_{-*}(Q))\big). \tag{37}$$

To see this, on the left-hand side the part of the boundary operator on a chain of loops that remains after passing to the E_1–level corresponds to the endpoint going to the boundary of the Morse chain in Q. On the right-hand side, it remains to add flow trees at the last distinguished critical points. The homologies of these differentials are then clearly as stated. Equation (37) together with Zeeman's comparison theorem [51, Section 3.3] then establishes the result. □

Remark 50 If $\dim(Q) = 4$ then the Morse function may have critical points of index one. In this case we use stabilization as follows. Multiply Q by $\mathbb{R}^N$ for any $N \geq 2$, and consider the function $F(q,x) = f(q) + x^2$. Then F has the same critical points as f and $-\nabla F$ is inward pointing at infinity. In $Q \times \mathbb{R}^N$ there is room to cancel 1–handles and the above applies. In this case we define $\mathrm{CM}_{-*}(Q)$ to be $\mathrm{CM}_{-*}(Q \times \mathbb{R}^N)$, which is a 1–reduced version of the original complex. Noting that $C_{-*}(Q \times \mathbb{R}^N)$ and $C_{-*}(Q)$ are canonically isomorphic, the result follows also in this case.

We next show that $\mathrm{CE}^*_{\parallel}(\Lambda)$ and $\mathrm{CE}^*(\Lambda)$ are isomorphic if Λ is simply connected. This is a more or less a direct consequence of the description of rigid disks on a Legendrian with parallel copies in Lemma 35 and the isomorphism in Theorem 49. Since components in Λ_- are not affected by this choice of $\mathrm{CE}^*_{\parallel}$ versus CE^*, we disregard them and assume that $\Lambda = \Lambda_+$ in what follows.

Recall the definition of $\mathrm{CE}^*(\Lambda)$ given in Remark 24, which is generated by chains $C_{-*}((\Omega_p\Lambda)^{\times(i+1)})$ in the product of the based loop space of Λ with factors separated by Reeb chords. Here the differential of a chain is just the usual differential of the chain whereas the differential of a Reeb chord is a sum over all moduli spaces of disks with one positive puncture at the chord and any number of negative punctures. Such a moduli space carries a fundamental chain and the contribution to the differential is an alternating word of chains in the based loop space corresponding to the boundary arcs of the disk carried by the fundamental chain and Reeb chords at the negative puncture,

$$dc_0 = \sum_{\boldsymbol{c}'} [\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})],$$

where $\boldsymbol{c} = c_0\boldsymbol{c}'$ and $\boldsymbol{c}'$ is a word of Reeb chords. Here we use the diagonal in the product of loop spaces; see Remark 24.

We next consider a system of parallel copies $\overline{\Lambda}(\eta) = \{\Lambda_j(\eta)\}_{j=0}^{\infty}$ defined by a system of positive Morse functions (Section 3.3.2), where $\Lambda_0 = \Lambda$. Recall that the generators of the algebra $\mathrm{CE}^*_{\|}(\Lambda)$ are the Reeb chords connecting Λ_0 to Λ_1, and that these can be long, corresponding to Reeb chords of Λ, and short, corresponding to critical points of f_1 except for the minimum. The differential counts rigid disks with one positive puncture in $\mathcal{M}^{\mathrm{sy}}_{\|}(\boldsymbol{b};\kappa)$ where κ is a decreasing boundary numbering, $\boldsymbol{b} = b_0\boldsymbol{b}'$.

We next consider the map

$$\phi\colon \mathrm{CE}^*(\Lambda) \to \mathrm{CE}^*_{\|}(\Lambda)$$

which takes every Reeb chord to itself and takes a chain σ in the based loop space to $\phi(\sigma)$, where ϕ is as in (36) and where we identify the critical points of f_1 with the corresponding Reeb chords connecting Λ_0 to Λ_1.

Theorem 51 *The map ϕ is a DG–algebra map and if Λ^+ is simply connected then ϕ is a quasi-isomorphism.*

Proof The fact that ϕ is a chain map follows from Lemma 35. Filter the algebras by action of Reeb chords on the left-hand side and actions of long Reeb chords on the right-hand side. The map respects this filtration. The E_2–pages are obtained by acting by the differential on the chains on the based loop space only on the left-hand side and on Morse chords only on the right-hand side. The result is words of Reeb chords separated by homology classes in the based loop space and by homology classes in the (reduced) bar construction on the Morse coalgebra on the left- and right-hand sides, respectively. On these E_2–pages the map ϕ induces an isomorphism by Theorem 49. Since the sum of actions of the Reeb chords at the negative end is bounded by that at the positive end, the spectral sequences converge. The theorem follows. □

Remark 52 The isomorphism in Theorem 51 is compatible with the stabilization of Remark 50. To see this we multiply the ambient contact manifold Y with contact form α by $T^*\mathbb{R}^N$ and consider $\Lambda\times\mathbb{R}^N \subset Y\times T^*\mathbb{R}^N$ with contact form $\theta = (\alpha - y\,dx)$. The Reeb chords of $\Lambda\times\mathbb{R}^N$ then come in $\mathbb{R}^N$–families, one for each Reeb chord of Λ. Consider the contact form $e^{x^2}\theta$ and note that with respect to this contact form the Reeb chords of $\Lambda\times\mathbb{R}^N$ are in natural one-to-one correspondence with those of Λ and there is a canonical isomorphism between $\mathrm{CE}^*(\Lambda)$ and $\mathrm{CE}^*(\Lambda\times\mathbb{R}^N)$. In fact the disks in the differential are canonically identified. It follows in particular that Theorem 51 holds also if $\dim(Q) \le 4$.

4 Lagrangian (co)algebra

As before, X is a Liouville manifold with $c_1(X) = 0$ and L is an exact relatively spin Lagrangian in X with vanishing Maslov class and ideal boundary given by the Legendrian Λ.

We will associate several chain-level structures to L. To begin with, let us first assume that L is an embedded Lagrangian. Since L has boundary, in classical topology, one can consider either $C^*(L)$ or $C^*(L, \partial L)$. In our case, these two choices are reflected in the choices of $+$ or $-$ decorations on L, respectively. More generally, let L^v, $v \in \Gamma$, be the (irreducible) components of L. As with the Legendrian submanifolds in Section 3.2, we assume these components of L are decorated by signs and we write $L = L^+ \cup L^-$ for the corresponding decomposition. Let $F\colon L \to \mathbb{R}$ be a Morse function with prescribed behavior at infinity (depending on the $+$ or $-$ decoration) as explained in Section 3.3. We use this to construct a system of parallel copies $\bar{L} = \{L_j\}_{j=1}^{\infty}$, as in Section 3.3, shifted at infinity along the Reeb flow either in the positive or negative direction on L^+ and L^-, respectively.

Now, using the parallel copies, $\{L_j\}_{j=1}^{\infty}$, we define a graded quiver $\mathcal{Q}_L$ as follows. The parallel copies $\{L_j\}_{j=1}^{\infty}$ give rise to following sets, for fixed $i_1 < i_2$ positive integers, and $v, w \in \Gamma$ with $v \neq w$:

- Intersection points $L^v_{i_1} \cap L^v_{i_2}$ in bijection with the union of critical points of $F|_{L_v}$. These critical points may depend on the $+$ or $-$ decoration on L_v, one can for example turn a $-$ decorated component L^v into a $+$ decorated one, by introducing critical points corresponding to the topology of ∂L_v; see Figure 5.
- Intersection points $L^v_{i_1} \cap L^w_{i_2}$ in bijection with $L^v \cap L^w$.

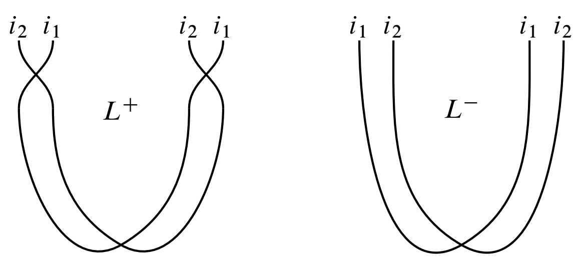

Figure 5: Difference between $+$ and $-$ generators for $i_1 < i_2$. Both the left- and the right-hand sides depict shifts corresponding to Morse functions with a maximum. One of the intersection points in L_+ is the minimum and corresponds to the unit for the Floer cohomology product.

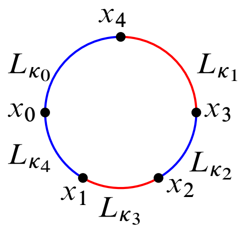

Figure 6: An example of a disk with labelings. The blue labeled Lagrangians are perturbed with $+$ perturbations and red labeled Lagrangians are perturbed with $-$ perturbation.

Furthermore, by the construction in Section 3.3, there are canonical bijections between the above sets associated to any pairs (i_1, i_2) and (i_1', i_2') with $i_1 < i_2$ and $i_1' < i_2'$. So, fix a pair (i_1, i_2) such that $i_1 < i_2$, and define a graded quiver $\mathcal{Q}_L$ with vertex set Γ and with an arrow connecting v to w (possibly equal to v) for each element of the above sets. Let $\mathcal{I}$ denote the set of arrows.

Alternatively, one can describe the generators as the set of intersection points in $L_0 \cap L_1$, between the original L and the first shifted copy.

Let $\mathrm{CF}^*(L)$ be the graded $\boldsymbol{k}$–bimodule generated by $\mathcal{I}$. Thus, there is one generator x_{vw} in degree $|x_{vw}|$ for each arrow in $\mathcal{Q}_L$ from v to w. We endow $\mathrm{CF}^*(L)$ with the structure of an A_∞–algebra. Let x_0 be an intersection point generator and let $\boldsymbol{x}' = x_i \cdots x_1$ be a word of intersection points. Consider the disk D_{i+1} with $i+1$ boundary punctures and with a decreasing numbering of its boundary arcs κ. Let $\boldsymbol{x} = x_0 x_i \cdots x_1$. Consider the moduli space $\mathcal{M}^{\mathrm{fi}}(\boldsymbol{x};\kappa)$; see Appendix A for notation. Define the operations $\mathfrak{m}_i$ by

$$\mathfrak{m}_i(\boldsymbol{x}') = \sum_{\substack{\ell(\boldsymbol{x}')=i \\ |x_0|=|\boldsymbol{x}'|+(2-i)}} |\mathcal{M}^{\mathrm{fi}}(\boldsymbol{x};\kappa)| x_0,$$

where $\ell(\boldsymbol{y})$ denotes the word length of $\boldsymbol{y}$ and $|\mathcal{M}^{\mathrm{fi}}(\boldsymbol{x};\kappa)|$ denotes the algebraic number of points in the oriented 0–dimensional manifold.

Lemma 53 *The operations $\mathfrak{m}_i$ satisfy the A_∞–algebra relations and are independent of the decreasing boundary labeling κ.*

Proof This follows by the usual argument: after noting that the decreasing boundary numbering ensures that there is no boundary bubbling, one observes that the terms in the A_∞–algebra relations count the ends of a 1–dimensional oriented compact manifold by Theorems 74 and 75. The operations compose because of Theorem 76. □

We call $\mathrm{CF}^*(L)$ the Lagrangian Floer cohomology algebra of L. Let u_v denote the generator corresponding to the minimum on the component $L_v \subset L^+$. If L_v is simply connected, then, by Lemma 38, u_v is a strict idempotent. We write $\boldsymbol{k}_- \oplus \mathrm{CF}^*(L)$ for the augmented algebra where we adjoined an idempotent e_w for each component $L_w \subset L^-$. (On these components the shifting function is decreasing at infinity and has a maximum rather than a minimum in the compact part.) This is a connected algebra over $\boldsymbol{k}$.

Remark 54 The two different choices of perturbations at infinity corresponding to $+$ and $-$ are the two extremal constructions, where one pushes the copies either always in the positive direction or always in the negative direction. One can also choose perturbations at infinity to depend on the topology of the manifold at infinity; see, for example, Section 4 of [1]. All our constructions should extend meaningfully to this more general setting, but we have not pursued this direction in this paper.

We next consider various linear duals of $\mathrm{CF}^*(L)$ and associated algebraic objects. The simplest case occurs when $\mathrm{CF}^*(L)$ is simply connected. In this case the linear dual $\mathrm{CF}_*(L)$ is a coalgebra with operations Δ_i dual to $\mathfrak{m}_i$, and as before we can adjoin $\boldsymbol{k}_-$ so that $\boldsymbol{k}_- \oplus \mathrm{CF}_*(L)$ is coaugmented over $\boldsymbol{k}$. Then we define the *Adams–Floer DG–algebra*

$$\Omega(\boldsymbol{k}_- \oplus \mathrm{CF}_*(L))$$

by applying the cobar construction.

In the non-simply-connected case, we replace this object by the *completed Adams–Floer DG–algebra*

$$\big(\mathrm{B}(\boldsymbol{k}_- \oplus \mathrm{CF}^*(L))\big)^{\#}.$$

Example 55 Let L be the standard Lagrangian D^n filling of the standard Legendrian unknot $\Lambda \subset S^{2n-1}$. The Floer cohomology can be computed as

$$\mathrm{CF}^*(L) = \begin{cases} \mathbb{K}x & \text{with } |x| = 0 \text{ if } L \text{ is decorated with } +, \\ \mathbb{K}c & \text{with } |c| = n \text{ if } L \text{ is decorated with } -. \end{cases}$$

Alternatively, if we want compatibility with the inclusion $C^*(D^n, \partial D^n) \to C^*(D^n)$, it can be computed as

$$\mathrm{CF}^*(L) = \begin{cases} \mathbb{K}c \oplus \mathbb{K}y \oplus \mathbb{K}x & \text{with } |c| = n,\ |y| = n-1,\ |x| = 0 \text{ and } dy = c \\ & \hfill \text{if } L \text{ is decorated with } +, \\ \qquad \mathbb{K}c & \text{with } |c| = n \text{ if } L \text{ is decorated with } -. \end{cases}$$

In Section B.1, we introduce a model for wrapped Floer cohomology without a Hamiltonian term and prove it is quasi-isomorphic to the usual wrapped Floer cohomology.

We refer there for details and give only a short description here. The chain complex underlying $\mathrm{CW}^*(L)$ is the following. Let $L = L_0$ and shift L off itself to L_1 by a Morse function that is positive at infinity (as in the definition of parallel copies when L is decorated $+$). The generators of $\mathrm{CW}^*(L)$ are then Reeb chords connecting L_1 to L_0 and intersection points in $L_0 \cap L_1$.

There is an A_∞–functor, often called the *acceleration functor*,

$$\mathrm{CF}^*(L) \to \mathrm{CW}^*(L).$$

If L is decorated $+$, it can be shown that this functor is unital.

5 Maps relating Legendrian and Lagrangian (co)algebras

We continue with our usual set-up, where X is a Liouville manifold with $c_1(X) = 0$ and L is an exact Lagrangian in X with vanishing Maslov class and ideal boundary given by the Legendrian Λ. Let Γ be the set of embedded components of L subdivided into $\Gamma^+ \cup \Gamma^-$. Let Θ be the set of components of Λ with induced subdivision $\Theta^+ \cup \Theta^-$.

In this section we will define twisting cochains and associated DG–algebra maps relating the parallel copies version $\mathrm{CE}^*_{\parallel}(\Lambda)$ and the Floer cohomology $\mathrm{CF}^*(L)$. Since L is an exact filling, we have an augmentation $\epsilon_L\colon \mathrm{CE}^*_{\parallel}(\Lambda) \to \boldsymbol{k}$. As in Section 3.2 we use this augmentation throughout to change coordinates in such a way that $\Delta_0 = 0$.

As explained in Section 3.4, the definition of $\mathrm{CE}^*_{\parallel}$ differs depending on whether or not the components of Λ in Θ^+ are simply connected. We will start in the simply connected case and turn later to the non-simply-connected case, using the definitions in Section 3.4.2.

Assume thus that all components of Λ in Θ^+ are simply connected. As usual, let $\boldsymbol{k}_- \oplus \mathrm{LC}^{\parallel}_*(\Lambda)$ denote the coalgebra corresponding to $\mathrm{CE}^*_{\parallel}(\Lambda)$ augmented by the Lagrangian filling, with counits e_v adjoined to all components Λ_v in Θ^-. As Λ^+ has simply connected components, by Lemma 38, this is a counital coalgebra with counit

$$\sum_{v \in \Theta^+} x_v + \sum_{v \in \Theta^-} e_v.$$

Let $\eta\colon \boldsymbol{k} \to \boldsymbol{k}_- \oplus \mathrm{LC}^{\parallel}_*(\Lambda)$ denote the coaugmentation

(38) $$\eta(e_v) = \begin{cases} x_v & \text{if } v \in \Theta^+, \\ e_v & \text{if } v \in \Theta^-, \end{cases}$$

(see (33)), so that $\mathrm{CE}^*_{\parallel} = \Omega(\boldsymbol{k}_- \oplus \mathrm{LC}^{\parallel}_*)$. This means that x_v is traded for e_v for $v \in \Theta^+$.

Consider the Floer cohomology A_∞–algebra $\mathrm{CF}^*(L)$. If all components of L^+ are simply connected there exists a strict idempotent $u_v \in \mathrm{CF}^*(L)$ for each $v \in \Gamma^+$ corresponding to the minimum of the shifting Morse function, and we make $\mathrm{CF}^*(L)$ unital by adding an idempotent e_w for each $w \in \Gamma^-$. We write the strictly unital algebra as $\boldsymbol{k}_- \oplus \mathrm{CF}^*(L)$. Let $\epsilon\colon \boldsymbol{k}_- \oplus \mathrm{CF}^*(L) \to \boldsymbol{k}$ be the augmentation that maps u_v to e_v for $v \in \Gamma^+$ and e_w to e_w for $w \in \Gamma^-$. Consider the dual of the bar construction,

$$\mathscr{A} = \big(\mathrm{B}(\boldsymbol{k}_- \oplus \mathrm{CF}^*(L))\big)^\#, \tag{39}$$

or in other words the completed Adams–Floer DG–algebra. In what follows we will represent $\mathscr{A}$ as a quotient in way that generalizes to the non-simply-connected case in analogy with the construction in Section 3.4.2. In the non-simply-connected case we introduce strict idempotents by hand as follows.

Consider adding extra idempotents e_v for $v \in \Gamma^+$ to $\boldsymbol{k}_- \oplus \mathrm{CF}^*(L)$. This gives $\boldsymbol{k} \oplus \mathrm{CF}^*(L)$ and we equip it with the trivial augmentation ϵ' which is the projection to $\boldsymbol{k}$. Let

$$\mathscr{A}' = \big(\mathrm{B}(\boldsymbol{k} \oplus \mathrm{CF}^*(L))\big)^\#,$$

and let $\mathscr{I}$ denote the subalgebra of $\mathscr{A}'$ given by the space of functionals which vanish on monomials not containing u_v for some $v \in \Gamma^+$. Let $\rho\colon \mathscr{A}' \to \mathscr{A}$ denote the restriction map induced by the inclusion $\boldsymbol{k}_- \oplus \mathrm{CF}^*(L) \to \boldsymbol{k} \oplus \mathrm{CF}^*(L)$.

Lemma 56 *The subalgebra $\mathscr{I}$ is closed under the differential. In the simply connected case, the map ρ is a chain map with kernel $\mathscr{I}$ and consequently $\mathscr{A}$ is quasi-isomorphic to $\mathscr{A}'/\mathscr{I}$.*

Proof Similar to Lemma 42. To see that $\mathscr{I}$ is closed under the differential we check that if u_v is an output of the differential, then u_v is also an input. Here the output is the positive puncture of the holomorphic disk, or equivalently flow tree; see Lemma 33. Since there is no negative-gradient flow line that starts at a minimum, any gradient flow tree with positive puncture at a minimum must also have a negative puncture at that minimum. This shows that if u_v is an output then it is an input as well.

In the simply connected case, monomials not containing u_v come from $\mathrm{B}(\boldsymbol{k}_- \oplus \mathrm{CF}^*)$, therefore the kernel of ρ is contained in $\mathscr{I}$ and, conversely, any element in $\mathscr{I}$ restricts to zero on $\mathrm{B}(\boldsymbol{k}_- \oplus \mathrm{CF}^*)$. Thus, in this case $\mathscr{I}$ is the kernel of ρ. □

In the general case we define $\mathscr{A} = \mathscr{A}'/\mathscr{I}$. Lemma 56 shows that in the simply connected case this definition agrees with the alternative definition of $\mathscr{A}$ given in (39).

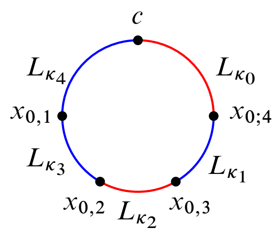

Figure 7: An example of a disk contributing to $\mathfrak{t}'$. The blue labeled Lagrangians are perturbed with $+$ perturbations and red labeled Lagrangians are perturbed with $-$ perturbation.

Remark 57 The somewhat artificial construction of $\mathscr{A}$ as $\mathscr{A} = \mathscr{A}'/\mathscr{I}$ is used to adapt the bar construction to a not necessarily strictly unital algebra.

We next define a map $\mathfrak{t}'$ on generators of $\mathrm{CE}^*_{\|}(\Lambda)$ which then gives a map

$$\mathfrak{t}' \colon \mathrm{LC}^{\|}_*(\Lambda) \to \mathscr{A}'$$

in the simply connected case, and in that case it will induce a twisting cochain

$$\mathfrak{t} \colon \boldsymbol{k}_- \oplus \mathrm{LC}^{\|}_*(\Lambda) \to \mathscr{A}.$$

The map $\mathfrak{t}'$ is defined by the following curve count for generators of $\mathrm{LC}^{\|}_*(\Lambda)$. Fix systems of parallel copies $\overline{L}$ of L. Recall that the components labeled with a $+$ sign are shifted by a positive Morse function and the components labeled with a $-$ sign are shifted by a negative Morse function.

Let c be a Reeb chord of $\overline{\Lambda}$ and let $\boldsymbol{x}_0 = x_{0;1} \cdots x_{0;j}$, $j > 0$, be a nonempty word of intersection points of $\overline{L}$. Let

$$\boldsymbol{c} = c\boldsymbol{x}_0$$

and define

$$\mathfrak{t}'(c) = \sum_{|\boldsymbol{x}_0| = |c| + (1-j)} |\mathcal{M}^{\mathrm{fi}}(\boldsymbol{c})|\boldsymbol{x}_0, \tag{40}$$

where we interpret $\boldsymbol{x}_0$ as an element in $\mathscr{A}'$.

Remark 58 In the non-simply-connected case we use the same formula to define $\mathfrak{t}'(c)$ and note that the sum in the definition may be infinite.

We have the following:

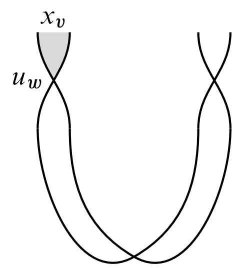

Figure 8: Minimum x_v is sent to the minimum u_w.

Proposition 59 *If $v \in \Theta^+$ is such that Λ_v is a boundary component of L_w for $w \in \Gamma^+$, then*

$$\mathfrak{t}'(x_v) = u'_w,$$

where u'_w is the dual of the minimum u_w. Furthermore, $\mathfrak{t}'$ satisfies the equation of a twisting cochain.

Proof For the first property we need to understand rigid holomorphic disks with positive puncture at the Reeb chord x_v. By small action such a holomorphic disk must lie in a neighborhood of L_w and is hence given by Morse flow trees. There is only one flow line emanating from the minimum in Λ_v and the flow line generically ends at the minimum of the shifting of L_w; see Figure 8. The first equation follows.

To see the twisting cochain equation, we need to check that

$$\mathfrak{m}_1 \circ \mathfrak{t}' - \mathfrak{t}' \circ \Delta_1 + \sum_{d \geq 2} (-1)^d \mathfrak{m}_2^{(d)} \circ \mathfrak{t}'^{\otimes_{\boldsymbol{k}} d} \circ \Delta_d = 0,$$

where $\mathfrak{m}_2^{(2)} := \mathfrak{m}_2$ and $\mathfrak{m}_2^{(i)} := \mathfrak{m}_2 \circ (\mathrm{Id}_{\mathscr{A}} \otimes_{\boldsymbol{k}} \mathfrak{m}_2^{(d-1)})$. To this end, we consider the boundary of the 1–dimensional moduli space $\mathcal{M}^{\mathrm{fi}}(c_0\boldsymbol{x})$. By Theorem 75 this corresponds to two-level curves which by Theorems 74 and 76 form the boundary of an oriented compact 1–manifold. $\square$

Proposition 59 shows that $\mathfrak{t}'$ maps the submodule generated by x_v for $v \in \Theta^+$ into $\mathscr{I} \subset \mathscr{A}'$. Hence, by letting $\mathfrak{t}(e_v) = 0$, $\mathfrak{t}'$ induces a map

$$\mathfrak{t}\colon \boldsymbol{k}_- \oplus \mathrm{LC}^{\parallel}_*(\Lambda) \to \mathscr{A}'/\mathscr{I} = \mathscr{A}.$$

Note that, if $\eta\colon \boldsymbol{k} \to \boldsymbol{k}_- \oplus \mathrm{LC}^{\parallel}_*(\Lambda)$ is the coaugmentation in (38) and $\epsilon\colon \mathscr{A} \to \boldsymbol{k}$ is the trivial augmentation, then $\epsilon \circ \mathfrak{t} = \mathfrak{t} \circ \eta = 0$.

Corollary 60 *The map $\mathfrak{t}$ is a twisting cochain.*

Proof Since $\mathfrak{t}'$ satisfies the twisting cochain equation, so does $\mathfrak{t}$. $\square$

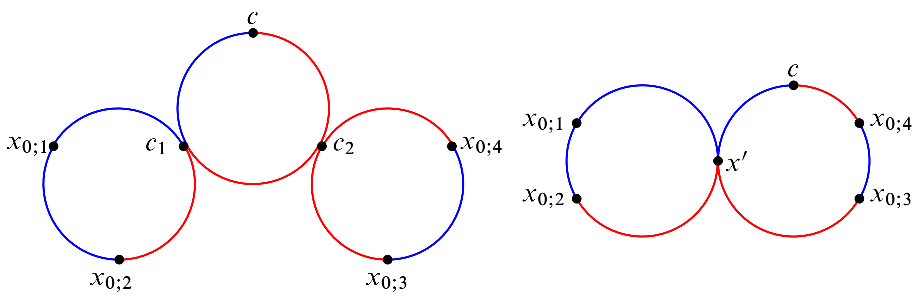

Figure 9: Left: a two-story disk contributing to the term $\mathfrak{m}_2 \circ \mathfrak{t}^{\otimes 2} \circ \Delta_2$ applied to c. There are similar disks in the compactification of the 1–dimensional moduli space, with 2 replaced by n. Right: a two-story disk which contributes to the term $\mathfrak{m}_1 \circ \mathfrak{t}$ applied to c.

This twisting cochain is always defined, and determines a map

$$\mathrm{CE}^*_{\parallel}(\Lambda) \to \mathscr{A}. \tag{41}$$

We next consider the question whether $\mathfrak{t} \in \mathrm{Kos}(\mathrm{LC}^{\parallel}_*(\Lambda), \mathscr{A})$. The following theorem gives a sufficient condition for $\mathfrak{t}$ to be Koszul:

Theorem 61 *Suppose that $\mathrm{LC}^{\parallel}_*(\Lambda)$ is a locally finite, simply connected $\boldsymbol{k}$–bimodule. Suppose further that $\mathrm{HW}^*(L) = 0$. Then $\mathscr{A}$ is quasi-isomorphic to $\Omega(\boldsymbol{k}_- \oplus \mathrm{CF}_*(L))$ and $\mathfrak{t}\colon \mathrm{LC}^{\parallel}_*(\Lambda) \to \mathscr{A}$ is a Koszul twisted cochain. In other words, the induced DG–algebra map*

$$\mathrm{CE}^*_{\parallel}(\Lambda) \to \mathscr{A} \approx \Omega(\boldsymbol{k}_- \oplus \mathrm{CF}_*(L))$$

is a quasi-isomorphism.

Corollary 62 *In the situation of Theorem 61, suppose that L is connected and decorated by $-$ and that $\partial L = \Lambda$ is diffeomorphic to a sphere S^{n-1}. Writing $\overline{L} = L \cup_\partial D^n$, there exists a quasi-isomorphism of DG–algebras*

$$\mathrm{CE}^*(\Lambda) \to C_{-*}(\Omega \overline{L}),$$

where $\Omega \overline{L}$ is the based loop space of $\overline{L}$.

Proof We first note that there exists a quasi-isomorphism $\boldsymbol{k} \oplus \mathrm{CF}^*(L) \to C^*(\overline{L})$ since L is an exact Lagrangian; this is well known and can be deduced eg from Lemma 33. We next use Theorem 61 and Adams' cobar equivalence [4]

$$\Omega(C_*(\overline{L})) \simeq C_{-*}(\Omega \overline{L}),$$

which holds for the simply connected topological space $\overline{L}$. □

Theorem 61 is obtained as a corollary of Theorem 18 and the following result. For the notion of wrapped Floer cohomology, see [3] and Section B.1.

Theorem 63 *Suppose that* $\mathrm{HW}^*(L) = 0$. *Then there exists a quasi-isomorphism of augmented* A_∞*–algebras*

$$\mathfrak{e}\colon \mathrm{CF}^*(L) \to \mathrm{LA}^*_{\|}(\Lambda)$$

such that

$$\mathfrak{e}(u_v) = \sum_{\Lambda_w \subset \partial L_v} x_w \quad \textit{for } v \in \Gamma^+.$$

Note that $\mathrm{HW}^*(L) = 0$ if X is subcritical [3], or more generally a flexible Weinstein manifold. (A Weinstein manifold is flexible if the attaching spheres of all top handles are loose and hence have trivial Chekanov–Eliashberg algebras [18]; see Section B.2 for the vanishing of wrapped Floer cohomology.)

Note also that if there is no bijection between connected components of L and connected components of Λ, then we have to work over the semisimple ring of idempotents determined by the Lagrangian.

Proof We construct an A_∞–map $\mathfrak{e} = (\mathfrak{e}_i)_{i\ge 1}$, where the maps

$$\mathfrak{e}_i\colon (\mathrm{CF}^*(L))^{\otimes_{\Bbbk} i} \to \mathrm{LA}^*_{\|}(\Lambda)$$

are constructed by dualizing the components of the map $\mathfrak{k}'$. More explicitly, given c_0 and $\boldsymbol{x}_0 = x_{0;n}, \dots, x_{0;1}$, write $\boldsymbol{c} = c_0\boldsymbol{x}_0$ as in (40), and define

$$\mathfrak{e}_i(\boldsymbol{x}_0) = \sum_{c_0 \in \mathcal{R}} |\mathcal{M}^{\mathrm{fi}}(\boldsymbol{c})| c_0.$$

The proof that $(\mathfrak{e}_i)_{i\ge i}$ is an A_∞–map follows as in the proof of Proposition 59.

Now, we need to check that $\mathfrak{e}_1$ is a quasi-isomorphism. We point out that $\mathfrak{e}_1$ concerns only strips and is defined using only two parallel copies. In the case that $L = L^-$, this is a consequence of the exact sequence for wrapped Floer cohomology induced by the subdivision of the complex into high- and low-energy generators,

$$0 \to \mathrm{CW}^*_0(L) \to \mathrm{CW}^*(L) \to \mathrm{CW}^*_+(L) \to 0,$$

as follows. In terms of the version of wrapped Floer cohomology presented in Section B.1.1, the low-energy subcomplex $\mathrm{CW}^*_0(L)$ is generated by Lagrangian intersection points in $L_0 \cap L_1$, where $L = L_0$ and L_1 is the first parallel copy of L, shifted by a positive Morse function f that increases at the end; see Section 3.3. The differential on $\mathrm{CW}_0^*(L)$ counts holomorphic strips which are incoming along L_1 and outgoing along L_0 at the output puncture. Similarly, the high-energy quotient $\mathrm{CW}_+^*(L)$ is generated by Reeb chords connecting Λ_1 to Λ_0, and the differential counts holomorphic strips interpolating between such Reeb chords. Since $\mathrm{CW}^*(L)$ is acyclic it follows that the connecting homomorphism $\mathrm{HW}_+^*(L) \to \mathrm{HW}_0^{*+1}(L)$ is an isomorphism. In order to connect this to $\mathrm{CF}^*(L)$ and $\mathrm{LA}_{\|}^*(\Lambda)$, renumber the parallel copies so that L_1 now lies in the negative Reeb direction of L_0 at infinity and the shifting function f is replaced by $-f$. Then note that since L is labeled by $-$ it holds that:

- The linear dual of $\mathrm{CW}_+^{*-1}(L)$ is canonically identified with $\mathrm{LA}_{\|}^*(\Lambda)$ as a chain complex. Note that $\mathrm{CW}_+^{*-1}(L)$ also has an A_∞–coalgebra structure as defined in [30] and this should dualize to the A_∞–algebra structure on LA^*, but we do not need that here.
- The linear dual of $\mathrm{CW}_0^*(L)$ is canonically identified with $\mathrm{CF}^*(L)$ as a chain complex.
- The linear dual of the connecting homomorphism can be canonically identified with the map $\mathfrak{e}_1 : \boldsymbol{k}_- \oplus \mathrm{CF}^*(L) \to \boldsymbol{k}_- \oplus \mathrm{LA}_{\|}^*(\Lambda)$ on critical points which counts strips with an input puncture at $L_0 \cap L_1$ and output puncture at a Reeb chord and is the canonical map on $\boldsymbol{k}_-$.

Since the connecting homomorphism is an isomorphism so is its linear dual. (The argument here is originally due to Seidel; compare [29, Theorem 7.2].)

In the case that $L^+ \neq \varnothing$, the argument just given applies after a certain deformation, which we describe next. For components in L^+, we define $\mathrm{CF}^*(L)$ and $\mathrm{LA}_{\|}^*(\Lambda)$ via parallel copies shifted in the positive Reeb direction at infinity. To connect to the previous case we consider a Lagrangian cobordism of two cylinders: $\mathbb{R} \times \Lambda_0$, which is constant, and $\mathbb{R} \times \Lambda_1$, which is the trace of an isotopy pushing Λ_1 across Λ_0 in the negative Reeb direction. This can be arranged so that the intersection points of the two cylinders are in natural one-to-one correspondence between the short Reeb chords between Λ_0 and Λ_1. Furthermore, there is exactly one transverse holomorphic strip connecting each intersection point to the corresponding Reeb chord at the negative end of the cobordism. To see this, note that $\mathbb{R} \times \Lambda_1$ can be obtained from a graphical Lagrangian in a cotangent neighborhood of $\mathbb{R} \times \Lambda_0$ that intersects the zero section $\mathbb{R} \times \Lambda_0$ cleanly along $\Lambda_0 \times \{0\}$. More precisely, it is the graph of the pullback of a 1–form on the $\mathbb{R}$–factor with a single transverse zero at $0 \in \mathbb{R}$. The corresponding

Lagrangians are then a product of the zero section in $T^*\Lambda_0$ and a 1–dimensional Lagrangian in $T^*\mathbb{R}$, and of the zero section and a curve that intersects the zero section once transversely at 0. The transverse holomorphic strips after perturbation by a Morse function on Λ_0 are products of constants at the critical points of the Morse function and the obvious strips for the 1–dimensional Lagrangian. Transversality is a consequence of transversality for the components of the product. We call such curves *basic strips*.

Adding these cylinders to L^+ we get a 1–parameter family of pairs of Lagrangian submanifolds $\hat{L}_0^\rho$ and $\hat{L}_1^\rho$, where $\rho > 0$ is a gluing parameter that measures the length of the trivial cobordism between L^+ and the added cylinders. The wrapped Floer cohomology $\mathrm{CW}^*(\hat{L}_0^\rho, \hat{L}_1^\rho)$ between Lagrangians $\hat{L}_0^\rho$ and $\hat{L}_1^\rho$ vanishes: it is isomorphic to the wrapped Floer cohomology $\mathrm{CW}^*(L)$ by Hamiltonian deformation invariance. Write $\mathrm{CW}^*(\hat{L}^\rho) := \mathrm{CW}^*(L_0^\rho, L_1^\rho)$. This complex is then acyclic and is generated by the set of long Reeb chords C^+ from L_0 to L_1, the set of intersection points I between the cylinders, and intersection points P in L. Let C^- denote the short Reeb chords connecting L_0 to L_1 and recall the natural one-to-one correspondence $C^- \approx I$ above. Let $\rho > 0$. We claim that the following sets are in natural one-to-one correspondence for all sufficiently large ρ:

(i) Rigid strips of $\hat{L}^\rho$ with input puncture at $c \in C^+$ and output puncture at $p \in P$ and rigid strips of L with input puncture at $c \in C^+$ and output puncture at $p \in P$.

(ii) Rigid strips of $\hat{L}^\rho$ with input puncture at $c \in C^+$ and output puncture at $q \in I$ and rigid-up-to-translation strips of $\mathbb{R} \times \Lambda$ with input puncture in $c \in C^+$ and output puncture at $q \in C^-$.

(iii) Rigid strips of $\hat{L}^\rho$ with input puncture at $p \in P$ and output puncture at $q \in I$ and rigid strips of L with input puncture at $p \in P$ and positive puncture at $q \in C^-$.

To see this note first that the strips in (i) are unaffected by adding the almost trivial cobordism: the strips are transversely cut out and therefore solutions for small variations of the boundary data are canonically identified. Taking ρ sufficiently large the boundary data of the disks can be made arbitrarily close.

For (ii), in the limit $\rho \to \infty$ the disks limit to an anchored disk with a positive puncture and a puncture at the intersection point. Gluing to it the basic strip (see above for this notion) connecting the intersection point to the short Reeb chord gives a 1–dimensional moduli space. The other boundary component of this moduli space consists of a rigid strip in the cobordism and a disk or plane of dimension one at either symplectization end. We next argue that the other boundary component of the moduli space must be a trivial strip followed by a strip connecting c to q at the negative symplectization end. To see this we note that the cobordism is obtained by a very small perturbation of the trivial cobordism $\mathbb{R}\times\Lambda$. It is well known that the only rigid strips of the trivial cobordism are the trivial strips; nontrivial curves have dimension ≥ 1. Therefore, a rigid strip limits either to a trivial strip or does not leave a small neighborhood of the trivial cobordism $\mathbb{R}\times\Lambda$. In this neighborhood holomorphic strips correspond to Morse flow lines of the shifting function, and thus the only rigid strips in the cobordism with negative ends at Reeb chords are either close to trivial strips or a basic strip. Our assertion follows. For (iii) we note that every rigid strip must break under stretching into two rigid strips. Since the only rigid strips in the upper part are the basic strips, the claim follows.

Observe that the strips in (i) and (iii) contribute to $\mathfrak{k}'$, the strips in (ii) to the differential on $\mathrm{LC}^{\parallel}_*(\Lambda)$. The vanishing of the wrapped Floer cohomology of $\widehat{L}^\rho$ then implies that $\mathfrak{e}_1$ is a quasi-isomorphism. The last statement follows from Proposition 59. $\square$

Proof of Theorem 61 The A_∞–quasi-isomorphism given in Theorem 63 induces a quasi-isomorphism of DG–algebras

$$\Phi\colon \mathrm{B}(\boldsymbol{k}_-\oplus \mathrm{CF}^*(L)) \to \mathrm{B}(\mathrm{LA}^*_{\parallel}(\Lambda))$$

by an application of [21, Theorem 7.4] with respect to length filtrations on the bar construction.

By the local-finiteness and simple-connectedness assumptions, each of these bar constructions is locally finite. So we can apply the linear dual operation to get a quasi-isomorphism of DG–algebras

$$\Phi^\#\colon \mathrm{B}(\mathrm{LA}^*_{\parallel}(\Lambda))^\# \to \mathrm{B}(\boldsymbol{k}_-\oplus \mathrm{CF}^*(L))^\#. \tag{42}$$

The result then follows as in Theorem 18, where local-finiteness of the grading enabled us to appeal to Lemma 10. Therefore, the quasi-isomorphism given in (42) induces the required quasi-isomorphism

$$\Omega(\mathrm{LC}^{\parallel}_*(\Lambda)) \to \Omega(\boldsymbol{k}_-\oplus \mathrm{CF}_*(L)). \qquad \square$$

We next turn to the non-simply-connected case, where we use $\mathrm{CE}^*_{\parallel}(\Lambda)$ as defined in Section 3.4.2 directly without using the corresponding coalgebra. Note that the A_∞–algebras $\boldsymbol{k}_-\oplus \mathrm{CF}^*(L)$ and $\mathrm{LA}^*_{\parallel}(\Lambda)$ are finite-rank $\boldsymbol{k}$–bimodules (in particular, they are locally finite), thus we can consider their $\boldsymbol{k}$–duals, which are, by definition, the A_∞–coalgebras $\boldsymbol{k}_- \oplus \mathrm{CF}_*(L)$ and $\mathrm{LC}_*^{\|}(\Lambda)$. However, unless we have the simple-connectedness assumption, the A_∞–quasi-isomorphism

$$\mathfrak{e}\colon \boldsymbol{k}_- \oplus \mathrm{CF}^*(L) \to \mathrm{LA}_{\|}^*(\Lambda)$$

does not necessarily dualize to an A_∞–comap

$$\mathfrak{f}\colon \mathrm{LC}_*^{\|}(\Lambda) \to \boldsymbol{k}_- \oplus \mathrm{CF}_*(L),$$

because A_∞–comaps are required to factorize through inclusion of the corresponding direct sum into the direct product as in (9). This is to ensure that a A_∞–coalgebra map $\mathfrak{f}$ induces a DG–algebra map $\Omega\mathfrak{f}$ on the cobar construction.

If we drop this condition, the A_∞–quasi isomorphism dualizes to DG–algebra map

$$\widetilde{\mathrm{CE}}_{\|}^*(\Lambda) \to \mathscr{A},$$

and this is just the map (41) induced by the twisting cochain $\mathfrak{t}$. Furthermore, by Proposition 59 this gives a DG–algebra map

$$\widehat{\Omega}(\mathfrak{f})\colon \widehat{\mathrm{CE}_{\|}^*}(\Lambda) \to \widehat{\Omega}(\boldsymbol{k}_- \oplus \mathrm{CF}_*(L)).$$

Now, since $\mathfrak{f}$ is a quasi-isomorphism, by using the length filtration on $\widehat{\Omega}$, and appealing to [21, Theorem 7.4], we can conclude the following:

Theorem 64 *Suppose that* $\mathrm{HW}^*(L) = 0$. *Then there exists a quasi-isomorphism of DG–algebras*

$$\widehat{\mathrm{CE}_{\|}^*}(\Lambda) \to \widehat{\Omega}(\boldsymbol{k}_- \oplus \mathrm{CF}_*(L)).$$

Note that the completion $\widehat{\mathrm{CE}_{\|}^*}$ is in general a cruder invariant than both $\mathrm{CE}_{\|}^*(\Lambda)$ and $\mathrm{CE}^*(\Lambda)$. Nevertheless, we always have a map

$$\mathrm{CE}^*(\Lambda) \to \widehat{\mathrm{CE}_{\|}^*}(\Lambda).$$

Theorem 64 can be used to compute $\widehat{\mathrm{CE}_{\|}^*}$ in a variety of cases. For example, if L is a connected Lagrangian filling decorated by $-$ of a Legendrian Λ diffeomorphic to a sphere S^{n-1}, then writing $\overline{L} = L \cup_\partial D^n$, we have a quasi-isomorphism $\boldsymbol{k} \oplus \mathrm{CF}_*(L) \simeq C_*(\overline{L})$ since $\overline{L}$ is an exact Lagrangian. Hence, we have a map

$$\mathrm{CE}^*(\Lambda) \to \widehat{\Omega}(C_*(\overline{L})).$$

Here the right-hand side can often be computed; in particular, $H^0\big(\widehat{\Omega}(C^*(L))\big)$ is the group ring of the unipotent completion of the fundamental group $\pi_1(L)$; see [16].

In particular, any information on the completion map $\mathrm{CE}^*(\Lambda) \to \widehat{\mathrm{CE}}^*_{\|}(\Lambda)$ can help to obtain information about $\mathrm{CE}^*(\Lambda)$. We will see an application of this idea in the next section.

We end this section with a discussion of the twisting cochains constructed above from an after-surgery perspective. Assume that all components of Λ^- are spheres and recall the surgery isomorphism

$$\Theta\colon \mathrm{CW}^*(C) \to \mathrm{CE}^*(\Lambda)$$

of Conjecture 89. Let $\widetilde{\Theta} = \phi \circ \Theta$, where ϕ is the identity map on components in Λ_- and the map ϕ of Theorem 51. We next note that there is a natural A_∞–algebra map

$$\Psi\colon \mathrm{CW}^*(C) \to \mathrm{B}(\boldsymbol{k} \oplus \mathrm{CF}^*(L))^\# = \mathscr{A}' \to \mathscr{A} = \mathscr{A}'/\mathscr{I},$$

where we identify $\boldsymbol{k}_- \oplus \mathrm{CF}^*(L)$ with the Floer cohomology $\mathrm{CF}^*(L')$ of the manifold after surgery obtained by capping off all boundary spheres in Λ^- by disks. (Note that in the simply connected case, the shifting Morse function then extends with a unique minimum in each disk which gives an idempotent corresponding to e_v.)

The map Ψ is defined by a curve count. Fix systems of parallel copies $\overline{C}$ of C and $\overline{L}'$ of L'. Let $\boldsymbol{c}' = c_1 \cdots c_i$ be a word of Reeb chords of C and let $\boldsymbol{x}_0 = x_{0;1} \cdots x_{0;j}$ be a word of intersection points of L'. Let

$$\boldsymbol{c} = \boldsymbol{c}' z_v \boldsymbol{x}_0 z_w$$

and define

$$\Psi(\boldsymbol{c}') = \sum_{|\boldsymbol{x}_0| = |\boldsymbol{c}'| + (1-i)} |\mathcal{M}^{\overline{\mathrm{co}}}(\boldsymbol{c})| \boldsymbol{x}_0.$$

Remark 65 We require here that the parallel copies $\overline{L}'$ give a system of parallel copies of Λ near the surgery region in such a way that, for the components of Λ_v labeled by $-$ (resp. $+$), the induced parallel copies $\overline{\Lambda}_v$ lie in the negative (resp. positive) Reeb direction. Compare with Figure 5.

Theorem 66 *The map Ψ is a map of A_∞–algebras and $\Psi(z_v) = u_v'$, where z_v is the strict unit in $\mathrm{CW}^*(C_v)$ and u_v' is the dual of the unit in $\mathbb{K}e_v \oplus \mathrm{CF}^*(L_v)$ for each v.*

Proof This follows as usual by identifying terms contributing to the A_∞–relations with the oriented boundary of an oriented 1–manifold and Theorems 74–76.

To compute $\Psi(z_v)$ note that we can represent z_v as the minimum of the shifting Morse function of C and there is a unique flow line from this minimum to the intersection point $C_v \cap L'_v$, and a unique flow line in L'_v from the intersection point to the minimum in L_v. The corresponding holomorphic disk starts at the intersection point between C_0 and C_1, has two corners at $C_0 \cap L'_0$ and $C_1 \cap L'_1$, and ends at the intersection point in $L'_0 \cap L'_1$ corresponding to the minimum of the shifting Morse function. □

The pre-twisting cochain $\mathfrak{t}' \colon \mathrm{LC}^{\|}_*(\Lambda) \to \mathscr{A}$ can now be seen to arise via SFT stretching as follows. Consider the first component Ψ_1 of the A_∞–map above and a holomorphic disk contributing to it. Now stretch the lower end of the cobordism. Then by SFT compactness each curve contributing to Ψ_1 breaks up into a curve contributing to the map $\phi \circ \Theta$ followed by the twisting cochain at each negative puncture. The z_v are the only low-energy generators, so it follows that the map induces a map of the high-energy quotient into $\operatorname{coker}(\eta)$, and we can write the induced map Ψ_1^+ as $\Psi_1^+ = (\Omega\mathfrak{t}) \circ \phi \circ \Theta^+$.

6 Examples and applications

6.1 Concrete calculations

In this section we compute Legendrian and Lagrangian invariants in a number of concrete examples.

6.1.1 The unknot Consider the Legendrian unknot $\Lambda \subset S^{2n-1}$ for $n > 1$ with its standard filling $L = D^n \subset D^{2n}$. Then Λ can be represented as a standard unknot in a small Darboux chart which has effectively only one Reeb chord with respect to the standard Reeb flow on S^{2n-1}; see [11, Section 7.1].

We work over a field $\mathbb{K}$. Consider first the case when L is decorated by $-$. Then

$$\mathrm{LC}_*(\Lambda) = \mathbb{K} \cdot 1 \oplus \mathbb{K} \cdot c \quad \text{with } |c| = -n,$$

with all coalgebra maps $(\Delta_i)_{i \geq 1}$ trivial, except Δ_2, for which we have

$$\Delta_2(1) = 1 \otimes 1 \quad \text{and} \quad \Delta_2(c) = 1 \otimes c + c \otimes 1$$

by the counitality. Using a Morse function on D^{2n} with a unique local maximum a and which decreases along the end corresponding to a shift in the negative Reeb direction, we have

$$\mathbb{K} \oplus \mathrm{CF}_*(L) = \mathbb{K} \cdot 1 \oplus \mathbb{K} \cdot a \quad \text{with } |a| = -n.$$

Let $\mathscr{A} = \Omega(\mathbb{K} \oplus \mathrm{CF}_*(L))$ be the Adams–Floer algebra, where the degree of a is now shifted up by 1. Then we have the twisting cochain

$$\mathfrak{t} \colon \mathrm{LC}_*(\Lambda) \to \mathscr{A}, \quad \mathfrak{t}(c) = a.$$

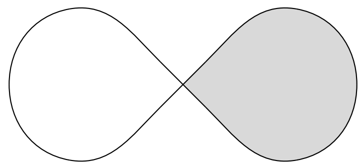

Figure 10: Computation of $\mathfrak{t}$ for the Legendrian unknot. The disk drawn lies in an $(n-1)$–dimensional family that sweeps the filling once.

Here the disk with input c and output a corresponds to the family of disks with positive puncture at c that sweeps L once.

The case of $n=2$ is drawn in Figure 10.

The Koszul complex is generated over $\mathbb{K}$ by a^k, $a^k c$ with $k \geq 0$. We can compute the nontrivial part of the differential to be

$$d^{\mathfrak{t}}(a^k c) = a^{k+1} \quad \text{for all } k \geq 0;$$

hence, $\mathfrak{t}$ is acyclic. This is consistent with the classical Koszul duality between the algebras $C^*(S^n)$ and $C_{-*}(\Omega S^n)$ for $n>1$.

Consider next the case when L is decorated by $+$. Since S^{n-1} is simply connected, $\mathrm{CE}^*(\Lambda) \approx \mathrm{CE}^*_{\parallel}(\Lambda)$ and we will use the parallel copies version in our calculation. Choose a Morse function on Λ with a single minimum and a single maximum. Denote the corresponding Reeb chords by x (the counit chord) and y. Then

$$\mathrm{LC}^{\parallel}_*(\Lambda) = \mathbb{K}\cdot x \oplus \mathbb{K}\cdot y \oplus \mathbb{K}\cdot c, \quad \text{with } |x|=0,\ |y|=-(n-1),\ |c|=-n.$$

Here

$$\Delta_1(c) = y, \quad \Delta_2(x) = x\otimes x + +(-1)^{n-1} x\otimes y + y\otimes x + (-1)^n x\otimes c + c\otimes x,$$

and all other operations are trivial. It follows that

$$\mathrm{CE}^*_{\parallel}(\Lambda) = \Omega(\mathrm{LC}^{\parallel}_*(\Lambda)) \simeq \mathbb{K}.$$

This is in line with Conjecture 3, which says that $\mathrm{CE}^*(\Lambda) \simeq \mathrm{CE}^*_{\parallel}(\Lambda)$ is isomorphic to $\mathrm{CW}^*(C)$, where C is the cotangent fiber in the manifold obtained by attaching a cotangent end $T^*(S^{n-1}\times[0,\infty))$ to the ball along Λ. The manifold that results from this attachment is simply $T^*\mathbb{R}^n$, and the wrapped Floer cohomology of the cotangent fiber C has rank 1 and is generated by the minimum in the disk C, in accordance with the above calculations.

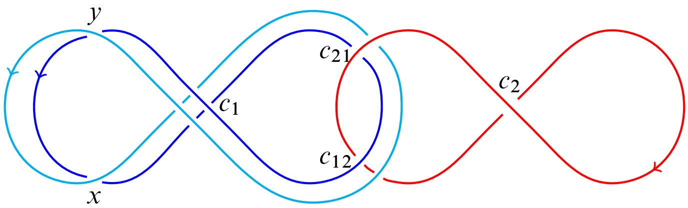

Figure 11: Hopf link with one (blue) marked $+$ and one (red) marked $-$ component.

Finally, the twisting cochain $\mathfrak{t}$ in the $+$ case is the canonical map $\mathbb{K} \to \mathbb{K}$, and again the Koszul complex is acyclic. As explained in Section 5 this map is induced by a map $\mathfrak{t}' \colon \mathrm{LC}^{\parallel}_{*}(\Lambda) \to (\mathrm{BCF}^{*}(L))^{\#}$. To define $\mathrm{CF}^{*}(L)$ we use a Morse function on L with a single local minimum and which is increasing along the end corresponding to a shift in the positive Reeb direction. Denote the generator of $\mathrm{CF}^{*}(L)$ corresponding to the minimum by u, $|u| = 0$. Then

$$\mathfrak{t}'(x) = u',$$

where u' is the dual of u and the holomorphic strip is the thin strip corresponding to a rigid Morse flow line from the minimum u to the minimum y in the boundary.

6.1.2 Geometric twisting cochain for the Hopf link In this section we carry out the geometric calculation of the twisting cochain for the Hopf link. As explained we cannot directly calculate the twisting cochain into the Legendrian coalgebra with coefficients in chains of the based loop space. We can however calculate the corresponding twisting cochain when we replace chains on the based loop space with the Morse theory of parallel copies for the components decorated by a positive sign. To carry out the calculation we pick a Morse function on the component Λ^{+} with on minimum x and a maximum y. We place them on the circle and choose parallel copies as shown in Figure 11. The parallel copies algebra $\mathrm{CE}^{*}_{\parallel}(\Lambda)$ is

$$\boldsymbol{k}\langle x, y, c_1, c_{12}, c_{21}, c_2\rangle,$$

where we use notation for Reeb chords and Floer cohomology generators as in Section 1.4. The differential is

$$\begin{aligned} dc_1 &= xc_1 + c_1x + y + c_{12}c_{21}, & dx &= xx, & dy &= xy - yx, \\ dc_{12} &= -xc_{12}, & dc_{21} &= c_{21}x, & dc_2 &= c_{21}c_{12}. \end{aligned}$$

Passing to $\mathrm{CE}^{*}_{\parallel}(\Lambda)$ means dividing out by the cokernel of the coaugmentation that takes e_1 to x. This gives the algebra

$$\boldsymbol{k}\langle y, c_1, c_{12}, c_{21}, c_2\rangle,$$

and the differential becomes

$$dc_1 = y + c_{12}c_{21}, \quad dc_2 = c_{21}c_{12}.$$

The twisting cochain $\mathfrak{t}$ is induced from a map $\mathfrak{t}'\colon \mathrm{LC}^{\parallel}_*(\Lambda) \to \big(\mathrm{B}(\boldsymbol{k}_- \oplus \mathrm{CF}^*(L))\big)^{\#}$ that counts holomorphic disks with one positive puncture and boundary on L, and with several punctures at Lagrangian intersection points in the compact part; see (40). In the current example it is straightforward to find these disks. Note first that, by general properties (see Proposition 59),

$$\mathfrak{t}'(x) = a_1^{\vee},$$

where a_1 is the idempotent corresponding to the minimum of the shifting Morse function on L_1. The holomorphic disk corresponds to a Morse flow line connecting x to u_1. We next consider $\mathfrak{t}'(c_1)$ and $\mathfrak{t}'(c_2)$. Consider first the moduli spaces $\mathcal{M}^{\mathrm{fi}}(c_j)$ of holomorphic disks with a positive puncture at c_j and boundary on L_j. As in the case of the unknot this moduli space sweeps L_j. On both L_1 and L_2 the shifting functions have one critical point, on L_1 it is a minimum and on L_2 a maximum. The maximum is constraining for the map into the linear dual of $\mathrm{CF}^*(L_j)$, whereas the minimum is not. We find that

$$\mathfrak{t}'(c_1) = 0 \quad \text{and} \quad \mathfrak{t}'(c_2) = a_2^{\vee}.$$

The spaces $\mathcal{M}^{\mathrm{fi}}(c_j)$ give further information of the twisting cochain as follows. Note that as the evaluation map hits the double point one can glue on a constant disk. These broken disks are also the boundary of the 1–dimensional moduli spaces $\mathcal{M}^{\mathrm{fi}}(c_1a_{12}a_{21})$ and $\mathcal{M}^{\mathrm{fi}}(c_2a_{21}a_{12})$. The other end of these moduli spaces correspond to broken disks with one level in the symplectization, a disk in $\mathcal{M}^{\mathrm{sy}}(c_1c_{12}c_{21})$ in the former case and in $\mathcal{M}^{\mathrm{sy}}(c_2c_{21}c_{12})$ in the latter, and two disks one in $\mathcal{M}^{\mathrm{fi}}(c_{12}a_{12})$ and one in $\mathcal{M}^{\mathrm{fi}}(c_{21}a_{21})$ attached at its negative end. The last disks contribute to $\mathfrak{t}'$ and we conclude that

$$\mathfrak{t}'(c_{12}) = a_{12}^{\vee} \quad \text{and} \quad \mathfrak{t}'(c_{21}) = a_{21}^{\vee}.$$

Finally, we compute $\mathfrak{t}'(y)$. Since y is a small chord corresponding to a critical point at infinity of the shifting Morse function the only contributions to $\mathfrak{t}'(y)$ come from small holomorphic disks that are controlled by the Morse theory. It is straightforward to check that the only rigid disk corresponds to a flow line in L_1 connecting y to the intersection point and that this flow line corresponds to a holomorphic triangle in $\mathcal{M}^{\mathrm{fi}}(ya_{12}a_{21})$. It follows that

$$\mathfrak{t}'(y) = a_{21}^{\vee}a_{12}^{\vee}.$$

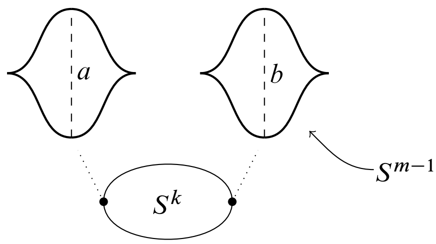

Figure 12: Front of products of spheres.

The actual twisting cochain takes the cokernel $\overline{\mathrm{LC}}_*(\Lambda)$ of the coaugmentation into the cokernel of the counits. Concretely, this means disregarding x and $a_1^\vee$, and we get

$$\mathfrak{t}(c_1)=0,\quad \mathfrak{t}(c_2)=a_2^\vee,\quad \mathfrak{t}(c_{12})=a_{12}^\vee,\quad \mathfrak{t}(c_{21})=a_{21}^\vee,\quad \mathfrak{t}(y)=a_{12}^\vee a_{21}^\vee.$$

Remark 67 The parallel copies algebra $\mathrm{CE}^*_{\|}(\Lambda)$ is defined using a fixed augmentation (in the current example the zero augmentation) on the one copy version of $\mathrm{CE}^*(\Lambda)$. Here this is reflected in the change of variables $t-e_1=y$.

6.1.3 Products of spheres We consider a Legendrian embedding $\Lambda\subset\mathbb{R}^{2(m+k)-1}$, where the ambient space is standard contact $(2(m–k)–1)$–space with coordinates $(x,y,z)\in\mathbb{R}^{m+k-1}\times\mathbb{R}^{m+k-1}\times\mathbb{R}$ and contact form $dz-y\,dx$. We will define it by describing its front in $\mathbb{R}^{m+k-1}\times\mathbb{R}$. To this end consider first the following construction of the front of the Legendrian unknot in $\mathbb{R}^n\times\mathbb{R}$. Take a disk D^n lying in $\mathbb{R}^n$. Think of it as having multiplicity two and lift one of the sheets up in the auxiliary $\mathbb{R}$–direction (with coordinate z) keeping it fixed along the boundary. In this way we construct the front of the standard unknot with Reeb chord at the maximum distance between the two sheets and a cusp edge along the boundary of D^n. Consider now instead the base $\mathbb{R}^{m+k-1}$ and the standard embedding of the k–dimensional sphere S^k into this space. A tubular neighborhood of this embedding has the form $S^k\times D^{m-1}$ with fibers D^{m-1}. Now take two copies of this tubular neighborhood and repeat the above construction for the D^{m-1} in each fiber. The result is the front of a Legendrian $S^k\times S^{m-1}$ with an S^k Bott family of Reeb chords corresponding to the maxima in fibers. Figure 12 shows this front after Morse perturbation. The resulting Legendrian Λ has only two Reeb chords. We denote them by a, with grading $|a|=-(m+k)$, and b, with grading $|b|=-m$. Note also that Λ has an exact Lagrangian filling $L\approx D^m\times S^k$.

Consider first the case when L is decorated by $-$. If d is the differential in $\mathrm{CE}^*(\Lambda)$, we have

$$da=0\quad\text{and}\quad db=0.$$

The Floer cohomology of L is defined by choosing a shifting function which is decreasing at infinity and we find that CF^* has two generators M, with $|M| = m+k$, and S, with $|S| = m$. As in the case of the unknot $\mathcal{M}^{\mathrm{fi}}(a)$ sweeps L and $\mathcal{M}^{\mathrm{fi}}(b)$ sweeps $D^m \times \mathrm{pt}$. It follows that the twisting cochain satisfies

$$\mathfrak{t}(a) = M^\vee \quad \text{and} \quad \mathfrak{t}(b) = S^\vee,$$

and duality holds.

Consider second the case when L is decorated by $+$. In this case there are additional generators of $\mathrm{LC}_*^{\parallel}(\Lambda)$ corresponding to the Morse theory of Λ. We have in addition to a and b above also

$$x,\ s_1,\ s_2,\ y, \quad \text{with} \quad |x| = 0,\ |s_1| = -k,\ |s_2| = -(m-1),\ |y| = -(m+k-1).$$

It follows that $\mathrm{CE}_{\parallel}^*(\Lambda)$ is generated by s_1, s_2, y, a and b. Using the flow tree description of moduli spaces $\mathcal{M}^{\mathrm{sy}}$ one verifies that if d is the differential on $\mathrm{CE}_{\parallel}^*(\Lambda)$ then

$$da = y - ((-1)^{km} b s_1 + s_1 b), \quad dy = s_1 s_2 + (-1)^{k(m-1)} s_2 s_1,$$
$$db = s_2, \quad ds_1 = 0, \quad ds_2 = 0.$$

The Floer cohomology of L is defined by choosing a shifting function which is increasing at infinity and we find that CF^* has two generators M, with $|M| = 0$, and S, with $|S| = k$, where M is the unit. It follows that $(\mathrm{BCF}^*(L))^\# \simeq \Omega\mathrm{CF}_*(L))$ is generated by $S^\vee$ with $|S^\vee| = -k$ and the twisting cochain is

$$\mathfrak{t}(a) = 0, \quad \mathfrak{t}(y) = 0, \quad \mathfrak{t}(b) = 0, \quad \mathfrak{t}(s_2) = 0, \quad \mathfrak{t}(s_1) = S^\vee,$$

and duality holds.

6.1.4 Plumbings of simply connected cotangent bundles Let $\mathcal{T}$ be a tree with vertex set Γ. For each $v \in \Gamma$, let M_v be a compact simply connected manifold of dimension $n \geq 3$. We will see that duality holds between the wrapped and the compact Fukaya categories of the symplectic manifold $X_{\mathcal{T}}$ obtained by plumbing the collection of $T^* M_v$ according to the tree $\mathcal{T}$.

As usual, we take the pre-surgery perspective. Hence, consider a handle decomposition of each M_v with a unique top-dimensional n–handle. Removing this handle, we get manifolds L_v with spherical boundary Λ_v. Let $W_{\mathcal{T}}$ be the subcritical Weinstein manifold obtained by plumbing the cotangent bundles $T^* L_v$ according to the tree $\mathcal{T}$. We take the plumbing region to be away from the boundary of L_v. Write $\Lambda = \bigsqcup_v \Lambda_v$ for the Legendrian in the boundary of the subcritical Weinstein manifold $W_{\mathcal{T}}$ which is filled by the Lagrangian $L = \bigcup_v L_v$. Equip the components of Λ with either $+$ or $-$ labeling.

Theorem 68 *If $n \geq 3$ and L_v is simply connected for each $v \in \Gamma$, then* $\mathrm{CE}^*(\Lambda)$ *and* $\mathrm{CF}^*(L)$ *are Koszul dual.*

Proof Consider first the case $n \geq 5$. Pick a handle decomposition of L_v without 1– and $(n-1)$–handles. The existence of such a handle decomposition is equivalent to simple connectedness in high dimensions by the work of Smale. Consider the corresponding Weinstein handle decomposition of T^*L_v. Attaching a k–handle alters the Legendrian boundary by surgery and adds Reeb chords in the cocore sphere of the handle, of grading $\leq -(n-k)$. It follows that all Reeb chords of Λ_v have grading ≤ -2. To construct $\Lambda \subset \partial W_{\mathcal{T}}$, we perform a version of boundary connected sum as follows. For each edge in $\mathcal{T}$ we pick a $2n$–ball B with two transversely intersecting Lagrangian disks $D \subset B$ that intersect the boundary sphere ∂B in a standard Legendrian Hopf link Δ. We then make the boundary connected sum adding (B, Δ) to join the T^*L_v according to the tree. This adds Reeb chords of index $\leq -(n-1)$ in the boundary connected sum handles and further Reeb chords corresponding to the Reeb chords of each Hopf link, which effectively has four Reeb chords: two Reeb chords connecting the unknot components to themselves of grading $-n$, and two mixed Reeb chords connecting distinct components. We can pick gradings so that the gradings of these two mixed chords are $-d$ and $-(n-d)$ for any d, where the first Reeb chord goes from the component closest to the a priori fixed root of the tree $\mathcal{T}$ to the one further from it and the other in the opposite direction. Taking d between 2 and $n-2$, we find that $\mathrm{LC}_*(\Lambda)$ is simply connected, as is Λ. The result then follows from Theorem 61.

For $n = 4$, we can stabilize by multiplying by $\mathbb{R}^N$ as described in Remarks 50 and 52 to get 1–reduced versions of the Legendrian and Lagrangian algebras, then use the above argument.

Finally, for $n = 3$, the assumptions of Theorem 61 do not hold, but we recall from Remark 20 that to apply the duality result from Theorem 61 all we really need is that $\mathrm{B}(\mathrm{LA}^*)$ is locally finite. This is easily seen to be the case in our case, since the plumbing is according to a tree (which by definition has no cycles) and for any word of Reeb chords that do not consist of connecting Reeb chords going away from the root of the tree, the grading has to increase with the size of the word. Alternatively, in this case we know by Perelman's theorem that we can take Λ as a link of standard

Legendrian spheres linked according to the tree as Hopf links, and $\mathrm{B}(\mathrm{LA}^*)$ can be seen to be locally finite directly. □

Remark 69 The case $n = 2$ corresponds to plumbing of copies of T^*S^2. This case was studied in [32] and a version of the duality result still holds, at least when $\mathrm{char}(\mathbb{K}) = 0$. However, the above argument fails in that case and a more complicated argument using an additional grading is used in [32]. Also a set of examples for $n = 1$ are studied in [48], where the plumbing tree is a star and the corresponding symplectic manifold is a punctured torus. The duality still holds in this case, even though this is a plumbing of T^*S^1 copies (not simply connected). The proof of duality given in [48] uses homological mirror symmetry.

6.1.5 The trefoil Consider the standard Legendrian trefoil $\Lambda \subset S^3$ described in Figure 13. Let us first consider the case when Λ is marked $-$. With respect to the standard choice of orientation datum, the Chekanov–Eliashberg DG–algebra CE^* is then given by the free algebra

$$\mathbb{K}\langle c_1, c_2, b_1, b_2, b_3\rangle, \quad \text{with } |c_1| = |c_2| = -1 \text{ and } |b_1| = |b_2| = |b_3| = 0,$$

and the nontrivial part of the differential can be read from Figure 13 as

$$dc_1 = 1 + b_1 + b_3 + b_3 b_2 b_1, \quad dc_2 = -1 - b_1 - b_3 - b_1 b_2 b_3.$$

It is well known that Λ has an exact Lagrangian torus filling. (In fact, there are at least five of them; see [29].) Any of these can be obtained by doing surgery (pinch move) at b_1, b_2 and b_3 in some order. Corresponding augmentations $\epsilon_L \colon \mathrm{CE}^* \to \mathbb{K}$ are given by their values $\epsilon(b_1), \epsilon(b_2), \epsilon(b_3) \in \mathbb{K}$ subject to the condition

$$1 + \epsilon(b_1) + \epsilon(b_3) + \epsilon(b_1)\epsilon(b_2)\epsilon(b_3) = 0.$$

Note that since CE^* is not simply connected, Theorem 61 does not apply here. In fact, duality does not hold in this example. However, Theorem 64 shows that there is a

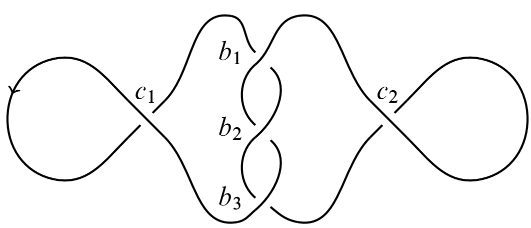

Figure 13: Trefoil.

quasi-isomorphism of completions

$$\widehat{\mathrm{CE}}^* \cong \hat{\Omega} C_*(T^2) \cong \mathbb{K}[\![u,v]\!],$$

where the latter is a power series ring in two commuting variables concentrated in degree 0.

Thus, the completion map composed with the twisting cochain gives an algebra map

$$H^0(\mathrm{CE}^*) = \mathbb{K}\langle b_1, b_2, b_3\rangle / \langle 1 + b_1 + b_3 + b_3 b_2 b_1, 1 + b_1 + b_3 + b_1 b_2 b_3\rangle \to \mathbb{K}[\![u,v]\!].$$

We claim that this map is injective. Indeed, observe that $H^0(\mathrm{CE}^*)$ is a commutative algebra,

$$d(c_1 + c_2 + b_3 b_2 c_1 + c_1 b_2 b_3) = b_3 b_2 - b_2 b_3,$$

and thus, b_2 commutes with b_3 in homology. Using this, one shows similarly that b_1 commutes with b_2 and b_3; see [15]. Hence, we have a completion map from a commutative ring to its completion,

$$H^0(\mathrm{CE}^*) = \mathbb{K}[b_1, b_2, b_3]/(1 + b_1 + b_3 + b_1 b_2 b_3) \to \mathbb{K}[\![u,v]\!].$$

It is a well-known theorem in commutative algebra, the Krull intersection theorem, that, for any commutative Noetherian ring which is an integral domain, the completion map is injective. Thus, even though duality fails in this example, partial information can still be obtained by considering completions.

We next describe the twisting cochain $\mathfrak{t}\colon \mathrm{LC}_*(\Lambda) \to \mathrm{BCF}^*(L)^\#$ for one of the Lagrangian fillings L of Λ. More precisely, we choose the filling which is obtained by pinching first at b_1, then at b_2, then at b_3, and finally filling the resulting unknots with Lagrangian disks; see [29, Section 8.1]. We think of the Lagrangian filling as two disks connected by three twisted bands corresponding to the three pinchings. We next consider moduli spaces of holomorphic disks with boundary in L. Here [29, Sections 4–5] gives a description in terms of Morse flow trees which gives the following:

- $\mathcal{M}^{\mathrm{fi}}(b_1)$ consists of one disk, δ_1^1.
- $\mathcal{M}^{\mathrm{fi}}(b_2)$ consists of two disks, δ_2^1 and δ_2^2.
- $\mathcal{M}^{\mathrm{fi}}(b_3)$ consists of two disks, δ_3^1 and δ_3^2.

The boundaries of these disks are as follows:

- $\partial\delta_1^1$ is a fiber in the first twisted band.

- $\partial\delta_2^1$ is a fiber in the second twisted band, and $\partial\delta_2^2$ runs across the first and the second twisted band.
- $\partial\delta_3^1$ is a fiber in the third twisted band, and $\partial\delta_3^2$ runs across the second and the third twisted bands.

We next describe the moduli spaces $\mathcal{M}^{\mathrm{fi}}(c_2)$ and $\mathcal{M}^{\mathrm{fi}}(c_1)$ in a completely analogous manner. The space has four components, all diffeomorphic to intervals, as follows:

- θ_1 with one boundary point the disk in $\mathcal{M}^{\mathrm{sy}}(c_2b_1)$ with δ_1^1 attached, and the other the disk in $\mathcal{M}^{\mathrm{sy}}(c_2b_1b_2b_3)$ with δ_1^1, δ_2^1 and δ_3^2 attached.
- θ_2 with one boundary point the disk in $\mathcal{M}^{\mathrm{sy}}(c_2b_3)$ with δ_3^1 attached, and the other the disk in $\mathcal{M}^{\mathrm{sy}}(c_2b_1b_2b_3)$ with δ_1^1, δ_2^2 and δ_3^1 attached.
- θ_3 with one boundary point the disk in $\mathcal{M}^{\mathrm{sy}}(c_2b_3)$ with δ_3^2 attached, and the other the disk in $\mathcal{M}^{\mathrm{sy}}(c_2b_1b_2b_3)$ with δ_1^1, δ_2^2 and δ_3^2 attached.
- θ_4 with one boundary point the disk in $\mathcal{M}^{\mathrm{sy}}(c_2)$ and the other the disk in $\mathcal{M}^{\mathrm{sy}}(c_2b_1b_2b_3)$ with δ_1^1, δ_2^1 and δ_3^1 attached.

Here the disks in θ_4 sweep the right-hand disk of L, whereas the evaluation maps of θ_j for $1 \le j \le 3$ do not map surjectively onto either disk. The moduli space $\mathcal{M}^{\mathrm{fi}}(c_1)$ also has four components, only one of which sweeps the left-hand disk of L.

In order to compute the Floer cohomology $\mathrm{CF}^*(L)$ we pick a Morse function with one maximum M and two saddle points S_1 and S_2. Letting the maximum lie in the right-hand disk we find

$$\mathfrak{t}(c_1) = 0,$$

since the only way to rigidify a disk of dimension one is that its boundary passes the maximum M.

The twisting cochain can now in principle be computed from the moduli spaces $\mathcal{M}^{\mathrm{fi}}$ described above by attaching flow trees. To get an algebraically feasible twisting cochain we first arrange the perturbation scheme so that

$$\mathfrak{t}(b_1) = S_1^\vee + S_1^\vee S_2^\vee,$$

where the first term is δ_1^1 with a flow line and the second with a flow tree, and next so that

$$\mathfrak{t}(b_3) = -S_1^\vee,$$

the disk δ_2^2 with a flow line contributes, while other contributions cancel (the disk δ_2^2 with two flow lines and the same disk with a flow tree, the two distinct disks with one flow line). Remaining parts of the twisting cochain are now determined from the coproduct in the Floer homology

$$dM^\vee = S_1^\vee S_2^\vee - S_2^\vee S_1^\vee,$$

the twisting cochain equation, and the sweeping property of θ_4, as

$$\mathfrak{t}(c_2) = M^\vee\left(\frac{1+S_1^\vee}{1+S_1^\vee+S_1^\vee S_2^\vee}\right) \quad\text{and}\quad \mathfrak{t}(b_2) = \frac{S_2^\vee}{(1+S_1^\vee+S_1^\vee S_2^\vee)}.$$

It is possible to check that these power series agree with the geometric count. We leave out the details but describe the mechanism. In order to arrange that only one flow line can be attached to δ_1^1 we order the stable manifolds of the flow line of the parallel copies so that if a flow line (or flow tree) between copy j and $j-l$ is attached then following δ_1^1 we already passed all intersections with stable manifolds between copy $j-l$ and k for $k<j-l$. Now, if the disk intersects the collection of stable manifolds in the opposite direction this means that we can jump down in all ways, which then gives the desired power series.

We finish this section with a brief discussion of the case of L decorated by $+$ and parallel copies. As usual this introduces two extra generators in addition to the Reeb chords above in LC_*, namely x, with $|x|=0$, and y, with $|y|=-1$, where x is the counit corresponding to the minimum of the shifting function and y is the maximum. In the reduced coalgebra (disregarding x) we get the new differential (using the augmentation ϵ_L above to change coordinates)

$$dc_1 = b_1+b_3-b_3b_2+b_3b_2b_1, \quad dc_2 = -y-b_1-b_3+b_2b_3-b_1b_2b_3.$$

In this case CF^* is defined instead by choosing a shifting function that increases at infinity, and CF^* is generated by the unit u, with $|u|=0$, and s_1 and s_2, with $|s_j|=1$. The new twisting cochain is

$$\mathfrak{t}(c_1)=\mathfrak{t}(c_2)=0,\ \mathfrak{t}(y)=(s_1^\vee s_2^\vee - s_2^\vee s_1^\vee)\left(\frac{1+s_1^\vee}{1+s_1^\vee+s_1^\vee s_2^\vee}\right),$$

and $\mathfrak{t}(b_j)$ is exactly as above, after the substitutions $s_j\to S_j$ for $j=1,2$.

6.1.6 Mirror of 7_2 We next discuss an example where Koszulity fails and also the completion map fails to be injective, even though CE^* is supported in nonpositive degrees. This example was shown to us by Lenhard Ng.

Consider the Legendrian knot Λ drawn in Figure 14 decorated $-$. It is easy to see that two pinch moves indicated by dashed lines in Figure 14 give a Legendrian unknot,

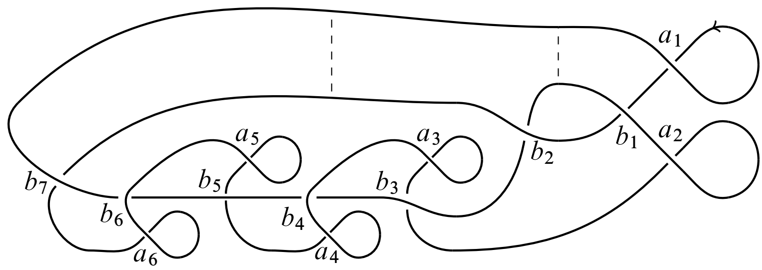

Figure 14: Mirror of 7_2.

hence Λ has a Lagrangian torus filling, call it L. Thus, we again have a completion map

$$H^0(\mathrm{CE}^*) \to \mathbb{K}[\![u, v]\!],$$

where $K[\![u, v]\!]$ is the commutative power-series algebra in two variables.

CE^* is given by the free algebra

$$\mathbb{K}\langle a_1, a_2, a_3, a_4, a_5, a_6, b_1, b_2, b_3, b_4, b_5, b_6, b_7\rangle,$$

with $|a_i| = -1$ and $|b_i| = 0$.

The differential is given by

$$da_1 = -1 + (1 + b_1b_2)b_7 + b_1(1 + b_4b_3)(1 + b_6b_5), \quad da_2 = 1 - b_3(1 + b_2b_1),$$
$$da_3 = 1 + b_3b_4, \quad da_4 = 1 + b_5b_4, \quad da_5 = 1 + b_5b_6, \quad da_6 = 1 + b_7b_6.$$

Taking the quotient of $H^0(\mathrm{CE}^*)$ by letting $b_4 = b_6$, $b_3 = b_5 = b_7$, $b_1 = 1$ and $b_2 = -1 - b_4$ gives

$$\langle b_3, b_4\rangle / \langle 1 + b_3b_4\rangle,$$

which is a noncommutative algebra.

Thus, the completion map cannot be injective in this case. Otherwise, $H^0(\mathrm{CE}^*)$, and thus any quotient of it, would have been commutative.

6.2 Simply connected Legendrian submanifolds

Let $\Lambda \subset Y$ be a Legendrian $(n-1)$–submanifold with $\pi_1(\Lambda) = 1$ in the boundary Y of a Weinstein $2n$–manifold X that bounds an exact Lagrangian $L \subset X$. Assume that $c_1(X) = 0$ and that the Maslov class of L vanishes and that L is relatively spin. Decorate L by $-$. Our next result shows that if the symplectic homology of X vanishes and if all Reeb chords of Λ have negative grading as generators of $\mathrm{CE}^*(\Lambda)$, then

$\mathrm{CE}^*(\Lambda)$ is determined by the topology of L, and conversely. If Λ is a sphere then we write $\overline{L} = L \cup_\partial D^n$ for the closed manifold obtained by adding a disk to L along Λ.

Theorem 70 *Suppose that $\Lambda = \Lambda^-$ is simply connected. Assume that $\mathrm{SH}^*(X) = 0$ and that $\mathrm{CE}^*(\Lambda)$ is supported in degrees ≤ -1. Then L is simply connected. Moreover, if Λ is a sphere, then $\mathrm{CE}^*(\Lambda)$ is isomorphic to $C_{-*}(\Omega\overline{L})$.*

Proof Consider the wrapped Floer cohomology $\mathrm{HW}^*_{\pi_1}(L)$ of L with coefficients in $\mathbb{Z}[\pi_1(L)]$. Using our model for wrapped Floer cohomology in Section B.1, a chain complex $\mathrm{CW}^*_{\pi_1}(L)$ which calculates $\mathrm{HW}^*_{\pi_1}(L)$ can be described as follows. Let $L = L_0$ and let L_1 be a parallel copy of L shifted in the negative Reeb direction at infinity. The complex $\mathrm{CW}^*_{\pi_1}(L)$ is then generated over $\mathbb{Z}[\pi_1]$ by the intersection points in $L_0 \cap L_1$, which we call the Morse generators, and the Reeb chords starting on Λ_0 and ending on Λ_1. The differential of a generator counts the usual rigid holomorphic strips, keeping track of the homotopy class of the loop obtained from the boundary component of the disk in L_1 completed by the reference paths connecting Reeb chord endpoints and intersection points to the basepoint. We point out that since there are no Reeb chords of degree 0 the augmentation induced by L is trivial and the high-energy part of the differential on $\mathrm{CW}^*_{\pi_1}$ counts honest holomorphic strips (without extra negative punctures at augmented Reeb chords). As for usual wrapped Floer homology, $\mathrm{HW}^*_{\pi_1}(L)$ is naturally a module over symplectic cohomology $\mathrm{SH}^*(X)$ and hence vanishes.

We next describe a geometric version of the complex $\mathrm{CW}^*_{\pi_1}(L)$ that we call $\mathrm{CW}^*_{\widetilde{p}}$ and that also computes $\mathrm{HW}^{\pi_1}(L)$. Let $\widetilde{p}\colon \widetilde{L} \to L$ denote the universal covering of L and let $\widetilde{\Lambda} = \widetilde{p}^{-1}(\Lambda)$. Pick a Morse function $f\colon \widetilde{L} \to \mathbb{R}$ such that

- f has exactly one local maximum M,
- f has no index $n-1$ critical points,
- f has no local minima,
- f is constant on $\widetilde{\Lambda}$ where it attains its global minimum and if ν is the unit normal vector field along $\widetilde{\Lambda}$ then $df(\nu) = -1$.

The generators of $\mathrm{CW}^*_{\widetilde{p}}$ are of two types:

(i) The preimages of endpoints of Reeb chords $L_0 \to L_1$ in $L \approx L_1$ under $\widetilde{p}$ graded as the corresponding Reeb chord in $\mathrm{CE}^*(\Lambda)$.

(ii) The critical points of the Morse function $f\colon \widetilde{L} \to \mathbb{R}$ graded by the negative of the Morse index.

Let M_{-*} denote the Morse chain complex of f with cohomological grading, with generators as in (ii) and differential δ which counts negative gradient flow lines. Then M_{-*} is supported in degrees d with $-n \le d \le -1$ and $M_{-(n-1)} = 0$. Let C^* denote the complex generated by the generators of type (i) and equip it with the differential ∂ that counts lifts of the boundary of holomorphic strips in the symplectization interpolating between Reeb chords. (This corresponds naturally to the high-energy part of the differential on $\mathrm{CW}^*_{\pi_1}$.) By our assumption on Reeb chord grading, the grading of C^* is supported in degrees d where $d \le -1$.

We now define the complex $\mathrm{CW}^*_{\widetilde{p}} = C^* \oplus M_{-*}$, with differential

$$d = \begin{pmatrix} \partial & \phi \\ 0 & \delta \end{pmatrix},$$

where δ and ∂ are the differentials on M^* and C^*, and where ϕ counts rigid lifts of disks with flow lines of f attached. (This is the linear part of the map ϕ in (36).)

The homology of d is then isomorphic to $\mathrm{HW}^*_{\pi_1}(L)$. To see this note that we can describe $\mathrm{CW}^*_{\pi_1}(L)$ exactly as $\mathrm{CW}^*_{\widetilde{p}}(L)$ just replacing the Morse function f above with a Morse function $h \circ \widetilde{p}$, where $\widetilde{p}$ is a Morse function on L without minimum and with the required boundary behavior. Thus, passing from $\mathrm{CW}^*_{\pi_1}(L)$ to $\mathrm{CW}^*_{\widetilde{p}}(L)$ corresponds to deforming the Morse function on $\widetilde{L}$, and it is well known that this induces a homotopy of complexes. In particular, $\mathrm{CW}^*_{\widetilde{p}}(L)$ is acyclic.

We next want to show that $\pi_1(L) \approx 1$ or, equivalently, that the map $\widetilde{L} \to L$ has degree one. To show this we first observe that since there are no Reeb chords of grading 0 the augmentation of $\mathrm{CE}^*(\Lambda)$ is trivial and the differential on C^* counts honest holomorphic strips in the symplectization. This in turn means that the whole boundary of any holomorphic strip contributing to ∂ actually lies in $\Lambda \times \mathbb{R}$ and therefore cannot pick up any nontrivial $\mathbb{Z}[\pi_1]$–coefficient.

Consider the part of the chain complex $C^* \oplus M_{-*}$ given by

$$\cdots \to C^{-(n+1)} \to C^{-n} \oplus M_{-n} \to C^{-(n-1)} \to \cdots,$$

where we use that $M_{-k} = 0$ for $k = n-1$ and $k > n$. It follows from the above discussion and the vanishing of the wrapped homology $\mathrm{HW}^*(L)$ with trivial coefficients that the cohomology $\mathrm{HW}^{-n}_{\pi_1}(L)$ in degree $-n$ has one generator for each nontrivial element in π_1. On the other hand, $\mathrm{HW}^{-n}_{\pi_1}(L) = 0$, and we conclude that $\pi_1 = 1$.

The statement about the isomorphism class of $\mathrm{CE}^*(\Lambda)$ follows from Corollary 62. □

Appendix A Basic results for moduli spaces

Consider as above a Weinstein manifold with an exact Lagrangian submanifold (X, L), which outside a compact set agrees with the positive part of the symplectization of the contact manifold with Legendrian submanifold (Y, Λ). We assume that the Maslov class of L vanishes and that L is relatively spin. We will consider several versions of punctured holomorphic spheres and disks with boundary on L. The most basic disks we consider will lie either in X or in the symplectization $\mathbb{R} \times Y$. We call the former *filling curves* and the latter *symplectization curves*. We will also consider a more general cobordism setting where, like in the symplectization, disks may have both positive and negative punctures. Here we assume that (W, K) is a Weinstein cobordism with negative end $(-\infty, 0] \times (Y, \Lambda)$ and positive end $[0, \infty) \times (Z, \Gamma)$, where Z is a contact manifold and Γ is a Legendrian submanifold. We call disks in W *cobordism disks*.

When we want to consider the relation of our "pre-surgery" invariants to "post-surgery" invariants, we will consider the case when K decomposes as a Lagrangian $C \subset W$ with positive end Γ and empty negative end and a Lagrangian L with negative end Λ and empty positive end. We further assume that there is a natural one-to-one correspondence between the components of L^v of L and C^v of C for $v \in Q_0$, and that corresponding components L^v and C^v intersect transversely at one point z^v and that $L^v \cap C^w = \varnothing$ if $v \neq w$.

We will describe a geometric setup that covers the cases considered below. Let Y be the contact boundary of the Weinstein manifold X, where $c_1(X) = 0$. Let $\Lambda \subset Y$ be a Legendrian with connected components $\Lambda_1, \dots, \Lambda_m$. Let D denote the unit disk in $\mathscr{C}$ and let $z_1, \dots, z_r$ be boundary punctures and $\zeta_1, \dots, \zeta_k$ interior punctures. Let each component of $\partial D \setminus \{z_1, \dots, z_r\}$ be decorated with a component Λ_j. The boundary punctures come in two types, positive and negative; all interior punctures are negative. Following [23], we make the further requirement that the disk be *admissible*:

> Any arc in D that connects two boundary arcs in $\partial D \setminus \{z_1, \dots, z_r\}$ subdivides the boundary punctures into two subsets. If both these subsets contain positive punctures, then the labels of the two boundary segments at the endpoints of the arc are different.

Remark 71 When we consider parallel copies of Lagrangians and Legendrians, boundary arcs labeled by different numbers in the numbering of parallel copies lie on copies shifted by different Morse functions, and correspond to distinctly labeled boundary conditions in the current discussion.

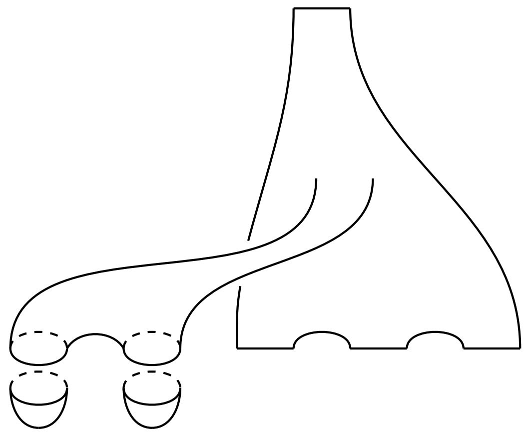

Figure 15: Anchored disks.

We also assume that the disk has one distinguished boundary puncture. Note that using a conformal model where the distinguished positive puncture lies at $1 \in \partial D$ and an interior puncture ζ_j lies at the origin, the positive real axis determines an asymptotic marker at ζ_j for each j. In the conformal model of the upper half-plane with the distinguished puncture at infinity, this marker at any interior puncture is that determined by the vertical axis.

A.1 Moduli spaces of spheres for anchoring, and compactifications of moduli spaces of disks

Following [11], all symplectization and cobordism disks we consider will be anchored. This means that the actual disks we consider have, aside from their boundary punctures, also additional interior punctures, where the maps are asymptotic to Reeb orbits at the negative end. An anchored disk is such a disk completed by holomorphic planes in X at all its negative interior punctures; see Figure 15.

When defining our version of the Chekanov–Eliashberg algebra CE^*, use moduli spaces of anchored disks to parametrize chains of boundary paths. When defining our versions of the Legendrian coalgebra LC_* and the wrapped Floer A_∞–algebra CW^*, we will need to consider disks in the symplectization with additional interior and boundary punctures, completed by rigid planes in X and disks in (X, L), respectively. We will also call such disks "anchored disks".

Although standard arguments using classical methods allow us to prove transversality for the disks with boundary punctures that we consider, anchored disks require transversality and gluing also for holomorphic planes in (X, L), and that requires a more general perturbation scheme. Necessary perturbations for such curves were constructed in [25].

To state the relevant result let Y denote the contact boundary of X and consider a Reeb orbit γ in Y with a marker (ie a point $p \in \gamma$) on it. We write γ' for the orbit γ with a marker. Let $\mathcal{M}(\gamma')$ denote the moduli space of holomorphic planes in X with positive puncture with an asymptotic marker where the curve is asymptotic to γ, with the asymptotic maker mapping to the marker on γ. As in [25, Theorem 1.1], we define perturbation data λ so that $\mathcal{M}^{\lambda}(\gamma')$ is a transversely cut out space of solutions to a perturbed Cauchy–Riemann equation $\bar\partial_{J_\lambda} u = 0$, where J_λ is a domain-dependent almost complex structure that is allowed to depend also on the map u in the neighborhood of $\mathcal{M}(\gamma')$. The moduli space $\mathcal{M}^{\lambda}(\gamma')$ furthermore has a natural compactification $\overline{\mathcal{M}}(\gamma')$ as a manifold with boundary with corners, where boundary strata correspond to several level spheres. Here levels not in moduli spaces of the form $\mathcal{M}^{\lambda}(\beta')$ just discussed lie in a moduli space $\mathcal{M}^{\mathrm{sy},\lambda}(\beta', \boldsymbol{\eta}')$, where β' is a Reeb orbit with marker and $\boldsymbol{\eta}' = \eta_1' \cdots \eta_k'$ is a word of Reeb orbits with markers. Elements in $\mathcal{M}^{\mathrm{sy},\lambda}(\beta', \boldsymbol{\eta}')$ are maps $u\colon S \to \mathbb{R} \times Y$ of a punctured spheres S into the symplectization. There are fixed cylindrical ends $S^1 \times [0, \infty)$ near the punctures in S that are compatible with breaking in the sense of [30, Section 2.1]. The map u takes $(1, \infty)$ at a puncture to the marker of the corresponding Reeb orbit and u again solves a perturbed Cauchy–Riemann equation $\bar\partial_{J_\lambda} u = 0$, where J_λ is domain dependent and only depends on the angular coordinate along the ends near the punctures, has a positive puncture at β' and negative punctures at the orbits in $\boldsymbol{\eta}'$.

We refer to [25, Section 2.4] for more details on $\overline{\mathcal{M}}(\gamma')$. Here we only point out that the asymptotic marker at the positive puncture of a curve in the compactification determines asymptotic markers at all negative punctures and that the level structure is compatible with this in the sense that the asymptotic marker at the positive puncture in a lower-level curve agrees with the asymptotic marker at the negative puncture where it is attached.

We next consider holomorphic disks in a cobordism (Z, K) with positive and negative ends $(\partial_\pm Z, \partial_\pm K)$. We include also the case when the cobordism (Z, K) is trivial, ie the symplectization $(\mathbb{R} \times Y, \mathbb{R} \times \Lambda)$, with $(\partial_\pm Z, \partial_\pm K) = (Y, \Lambda)$. Let $\boldsymbol{c}$ be a word of Reeb chords of $\partial_\pm K$. Let $\boldsymbol{\gamma} = \gamma_1 \cdots \gamma_k$ be a word of Reeb orbits in $\partial_- Z$. We define $\mathcal{M}^{\mathrm{neg}}(\boldsymbol{c}, \boldsymbol{\gamma}')$ to be the moduli space of punctured holomorphic disks in Z with boundary on K, with boundary punctures mapping to Reeb chords in the word $\boldsymbol{c}$ and one distinguished boundary puncture, and with additional negative interior punctures at

$\zeta_1,\dots,\zeta_k$ mapping to Reeb orbits in the word $\boldsymbol{\gamma}$. Note that the distinguished boundary puncture determines an asymptotic marker at each interior puncture ζ_j that determines a marker on the corresponding γ_j. Let $\boldsymbol{\gamma}'$ denote the corresponding word or Reeb orbits with markers. Below we will show that such moduli spaces of disks with interior negative punctures with markers that are relevant to our study cannot contain multiple covers and are transversely cut out for a generic almost complex structure.

Recall the symplectic filling X of the negative end Y of the cobordism above. We will use punctured sphere curves in X in the compactification of the moduli spaces $\overline{\mathcal{M}}^{\lambda}(\gamma')$ to fill the interior punctures of the disks and treat them as disks with only boundary punctures. For this purpose, we define the moduli space of anchored disks $\mathcal{M}^{\mathrm{anc}}(\boldsymbol{c})$ as

$$\mathcal{M}^{\mathrm{anc}}(\boldsymbol{c})=\bigcup_{\boldsymbol{\gamma}'}\biggl(\mathcal{M}^{\mathrm{neg}}(\boldsymbol{c},\boldsymbol{\gamma}')\times\prod_{\gamma_j'\in\boldsymbol{\gamma}'}\mathcal{M}^{\lambda}(\gamma_j')\biggr),$$

where markers on Reeb orbits are induced from the distinguished boundary puncture. Here the topology on the moduli space of anchored curves is the product topology. This means in particular that the dimension of the boundary evaluation map equals the dimension $\dim(\mathcal{M}^{\mathrm{anc}}(\boldsymbol{c}))$ only on components $\mathcal{M}^{\mathrm{neg}}(\boldsymbol{c},\boldsymbol{\gamma}')\times\prod_{\gamma_j'\in\boldsymbol{\gamma}'}\mathcal{M}^{\lambda}(\gamma_j')$ where $\dim(\mathcal{M}^{\lambda}(\gamma_j'))=0$ for all j.

We consider next the case when the cobordism is trivial $(\mathbb{R}\times Y,\mathbb{R}\times\Lambda)$ and when all punctures mapping to chords in $\boldsymbol{c}$ are positive. The above construction then gives a stratification of the moduli space $\mathcal{M}(\boldsymbol{c})$ of holomorphic disks in (X,L) as follows. First, since all moduli spaces $\mathcal{M}^{\mathrm{neg}}(\boldsymbol{c},\boldsymbol{\gamma}')$ are transversely cut out, the corresponding moduli spaces $\mathcal{M}^{\mathrm{neg},\lambda}(\boldsymbol{c},\boldsymbol{\gamma}')$, where a small perturbation near the negative ends corresponding to the perturbation λ of holomorphic planes with asymptotic marker has been turned on, is canonically diffeomorphic to $\mathcal{M}^{\mathrm{neg}}(\boldsymbol{c},\boldsymbol{\gamma}')$. Gluing the curves in $\mathcal{M}^{\mathrm{neg},\lambda}(\boldsymbol{c},\boldsymbol{\gamma}')$ to the curves in $\mathcal{M}^{\lambda}(\gamma_j')$ and extending the perturbation, we get a compactification of the moduli space $\mathcal{M}(\boldsymbol{c})$, with boundary given by tree configurations of anchored disks. Near a broken configuration the moduli space is a manifold with boundary with corners, with corner structure induced by the gluing parameters; compare [31, Sections 6.4–6.6].

For example, if $\boldsymbol{c}=c$ is a single Reeb chord, then $\mathcal{M}(c)$ has a natural compactification with boundary of the form

$$\mathcal{M}^{\mathrm{anc}}(c,\boldsymbol{b})\times\prod_{b_j\in\boldsymbol{b}}\mathcal{M}^{\mathrm{anc}}(b_j).$$

All moduli spaces of disks considered in this paper will be anchored, and we will drop the superscript "anc" from the notation.

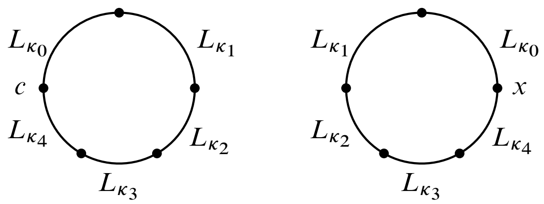

Figure 16: Strictly decreasing (left) and strictly increasing (right).

A.2 Moduli spaces of anchored disks

Consider a system of parallel copies $\overline{L} = \{L_j\}_{j=0}^{\infty}$, where $L_0 = L$ is an embedded Lagrangian, as in Section 3.3. This induces a system $\overline{\Lambda} = \{\Lambda_j\}_{j=0}^{\infty}$ of parallel copies of Λ in Y. We first discuss numberings that determine the boundary conditions of the holomorphic disks that we use in the various theories considered. Let D_m denote the unit disk in the complex plane with m boundary punctures $\zeta_1, \dots, \zeta_m$. One of the boundary punctures is distinguished. We choose notation so that ζ_1 is distinguished. The m punctures subdivide the boundary of D_m into m boundary arcs. We will consider disks with numbered boundary arcs where the numbers correspond to the parallel copies L_j. We will consider two types of numberings, *increasing* and *decreasing*. Traversing the boundary of the disk across a boundary puncture in the positive direction, the numbering increases, remains constant, or decreases as we pass the puncture. We call punctures increasing, constant, and decreasing, accordingly. We call a disk increasing (resp. decreasing) if all its nondistinguished punctures are either increasing (resp. decreasing) or constant. Then its distinguished puncture is decreasing (resp. increasing) or constant.

When defining operations Δ_k for the Legendrian A_∞–coalgebra LC_*, we count anchored increasing disks in the symplectization asymptotic to Reeb chords at all punctures. When defining the operations for the Lagrangian A_∞–algebras CF^* and CW^*, we count decreasing disks in X. When defining the twisting cochain $\mathfrak{t}\colon \mathrm{LC}_*^{\parallel} \to \Omega(\mathbf{k}_- \oplus \mathrm{CF}_*)$, we count increasing disks in X which are asymptotic to a Reeb chord at the distinguished puncture, and to Lagrangian intersection points at the other punctures.

We next consider asymptotic conditions at the boundary punctures. There are two basic forms of asymptotics: a puncture is either asymptotic to a Lagrangian intersection point or to a Reeb chord. The former case is the standard form of asymptotics in Lagrangian Floer theory, and the latter in Legendrian DG–algebras. More precisely, we choose an almost complex structure on X which along the cylindrical end $\mathbb{R} \times Y$ is invariant under

$\mathbb{R}$–translation, leaves the contact planes invariant and is compatible with the symplectic form induced on the contact planes by the contact form. Furthermore, it pairs the $\mathbb{R}$–direction with the Reeb direction. This means in particular that the Reeb chord strip, which is the product of a Reeb chord and $\mathbb{R}$, is holomorphic. "Reeb chord asymptotics" means we study holomorphic disks with boundary punctures that are asymptotic to these Reeb chord solutions, while "Lagrangian intersection point asymptotics" means we study holomorphic disks that are asymptotic to constant strips at intersection points.

First we consider disks in the filling (X, L). Consider a disk D_m as above with strictly increasing or decreasing numbering $\kappa = (\kappa_1, \dots, \kappa_m)$ and let $\boldsymbol{a} = a_1 \cdots a_m$ be a word of Reeb chords and Lagrangian intersection points in $L_0 \cap L_1$. We let $\mathcal{M}^{\mathrm{fi}}(\boldsymbol{a}; \kappa)$ denote the moduli space of holomorphic disks $u\colon (D_m, \partial D_m) \to (X, \bar{L})$ such that

- u takes the boundary component labeled by κ_j to the Lagrangian L_{κ_j}, and
- u is asymptotic to the unique Reeb chord or Lagrangian intersection point of L_{κ_j} and $L_{\kappa_{j+1}}$ near a_j at ζ_j, where we let $\kappa_{m+1} = \kappa_1$.

We next consider disks in the symplectization. Consider the disk D_{m+k} with increasing or decreasing boundary numbering κ' and punctures $\zeta_1, \dots, \zeta_{m+k}$. We next note that in the symplectization there are two possible Reeb chord asymptotics, positive or negative according to the sign of the t–coordinate near the puncture. Let $\boldsymbol{c}' = c_1^{\sigma_1} \cdots c_{m+k}^{\sigma_{m+k}}$ be a word of signed Reeb chords of $\Lambda_0 \cup \Lambda_1$, where $\sigma \in \{+, -\}$ is a sign. If $m > 1$ then we require that the Reeb chords c_r at all constant punctures connect Λ_0 to Λ_0 and that their signs are all negative, $\sigma_r = -1$. (These constant punctures will be capped by augmentation disks.) We let $\mathcal{M}^{\mathrm{sy},\circ}_{\parallel}(\boldsymbol{c}'; \kappa')$ denote the moduli space of anchored holomorphic disks $v\colon (D_m, \partial D_m) \to (\mathbb{R} \times Y, \mathbb{R} \times \Lambda)$ such that

- v takes the boundary components labeled by κ_j to the Lagrangian Λ_{κ_j}, and
- v is asymptotic at positive or negative infinity, according to the sign of σ_j, to the unique Reeb chord between Λ_{κ_j} and $\Lambda_{\kappa_{j+1}}$ near c_j at a puncture ζ_j.

If $\boldsymbol{c}$ is a word of strictly increasing (resp. decreasing) Reeb chords then we define the moduli space $\mathcal{M}^{\mathrm{sy}}_{\parallel}(\boldsymbol{c}; \kappa)$ by anchoring also at constant boundary punctures,

$$\mathcal{M}^{\mathrm{sy}}_{\parallel}(\boldsymbol{c}; \kappa) = \bigcup_{\boldsymbol{c} \subset \boldsymbol{c}'} \biggl(\mathcal{M}^{\mathrm{sy},\circ}_{\parallel}(\boldsymbol{c}'; \kappa') \times \prod_{c_r \in \boldsymbol{c}' \setminus \boldsymbol{c}} \mathcal{M}(c_r) \biggr),$$

where the union runs over all words $\boldsymbol{c}'$ extending $\boldsymbol{c}$ by constant punctures.

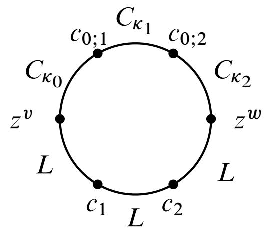

Figure 17: Disk contributing to $\mathcal{M}^{\mathrm{co}}(\boldsymbol{c};\kappa)$.

We will also consider a simpler version with only one copy of the Legendrian $\Lambda=\Lambda_0$ in the case that exactly one puncture in $\boldsymbol{c}$ is positive and all others are negative. In the language above, all punctures of such a disk are constant and we write $\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})$ for the moduli space of such disks. The (constant) negative punctures of disks in this moduli space are typically not filled by augmentation disks.

Finally, we consider disks in the cobordism and in the filled cobordism. We start with the cobordism disks without filling. Recall that we assume a decomposition $K=C\cup L$, where C has no negative end and L has no positive end. Consider a system of parallel copies $\overline{C}=\{C_j\}_{j=0}^{\infty}$. Consider the disk D_{i+j+2} where we fix two punctures that subdivide the boundary of the disk into two arcs, upper and lower. Let κ be a decreasing boundary numbering of the boundary components in the upper arc and extend it to a constant numbering in the lower arc. Let $\boldsymbol{c}_0=c_{0;1}\cdots c_{0;j}$ be a composable word of Reeb chords connecting Λ_v to Λ_w, and let $\boldsymbol{c}'=c_i\cdots c_1$ be a word of Reeb chords of Γ. Consider the word of Reeb chords and intersection points

$$\boldsymbol{c}=c_{0;1}\cdots c_{0;j}z^w c_i\cdots c_1 z^v,$$

and let

$$\mathcal{M}^{\mathrm{co}}(\boldsymbol{c};\kappa)$$

denote the moduli space of holomorphic disks $u\colon(D_{i+j+2},\partial D_{i+j+2})\to(W,\overline{C}\cup L)$ such that:

- u is asymptotic to the Reeb chord $c_{0;r}$ at its r^{th} constant puncture, takes adjacent boundary arcs to L, and neighboring punctures to the unique intersection point near z^w in $L\cap C^w_{\kappa_i}$ and near z^v in $L\cap C^v_{\kappa_i}$, respectively.
- On remaining boundary arcs and punctures, the boundary maps as described by $\boldsymbol{c}$ and the numbering κ, exactly as above.

The disks in the filled cobordism are entirely analogous. Here we assume that L is a Lagrangian submanifold in $X=X_0\cup W$. Consider again a system of parallel copies

$\bar{C} = \{C_j\}_{j=0}^{\infty}$ and also a system of parallel copies $\bar{L} = \{L_j\}_{j=0}^{\infty}$ of L. Consider the disk D_{i+j+2} where we fix two punctures that subdivide the boundary of the disk into two arcs, upper and lower. Let κ be a decreasing boundary numbering of the boundary components in the upper and lower arcs. Let $\boldsymbol{x}_0 = x_{0;1} \cdots x_{0;j}$ be a word of intersection points of L and let $\boldsymbol{c}' = c_i \cdots c_1$ be a word of Reeb chords of Γ. Consider the word of Reeb chords and intersection points

$$\boldsymbol{c} = x_{0;1} \cdots x_{0;j} z^w c_i \cdots c_1 z^v,$$

and let

$$\mathcal{M}^{\overline{\mathrm{co}}}(\boldsymbol{c}; \kappa)$$

denote the moduli space of holomorphic disks $u\colon (D_{i+j+2}, \partial D_{i+j+2}) \to (X, \bar{C} \cup \bar{L})$ such that:

- u is asymptotic to the intersection point $x_{0;r}$ at its r^{th} puncture in the lower arc, takes adjacent boundary arcs to L, and neighboring punctures to the unique intersection point near z^w in $L^w_{\kappa_j} \cap C^w_{\kappa_{j+1}}$ and the unique intersection point near z^v in $K_{\kappa_1} \cap C^v_{\kappa_{i+j}}$.
- On remaining boundary arcs and punctures, the boundary maps as described by $\boldsymbol{c}$ and the numbering κ, exactly as above.

The disks in the cobordism without filling will be used to map into the DG–algebra of the negative end, whereas the disks in the filled cobordism will be used to map into the Floer cohomology of a Lagrangian. This is why we use parallel copies in one case but not the other.

The formal dimension of the moduli spaces above is computed in terms of the negative of a Conley–Zehnder index CZ of the Reeb chords. Recall $\mathrm{CZ}(a)$ of a Reeb chord a of Λ as defined for example in [11, Section 2.1]: we pick paths connecting basepoints in the boundary of the various components of L, and paths connecting Reeb chord endpoints to the basepoints. We define $\mathrm{CZ}(a)$ to be the Maslov index of this path closed up by a positive rotation in the contact plane, and the grading $|a| = -\mathrm{CZ}(a)$. For a Lagrangian intersection x between L^1 and L^2 we similarly pick paths connecting to the basepoints and use these to form a loop γ starting in L^2 and ending in L^1, and define $\mathrm{CZ}(x)$ to be the Maslov index of the loop of Lagrangian planes that results from closing up the path of Lagrangian planes along γ by a positive rotation, and $|x| = -\mathrm{CZ}(x)$. The Conley–Zehnder index is independent of the basepoint paths since the Maslov class vanishes.

Remark 72 The above gradings are related to the grading $|\cdot|_{\mathrm{Leg}}$ in the Legendrian contact homology algebra [28; 11] as

$$|c|_{\mathrm{Leg}} = -|c| - 1.$$

Gradings generally depend on the choice of paths connecting endpoints to the basepoint: two such choices differ by a loop and the grading is shifted by the Maslov index of that loop. In particular, if the Maslov class vanishes the grading is well defined. Also, the paths connecting tangent planes at basepoints in different components are defined only up to choice. Changing the homotopy class shifts the Maslov potential between components and indices of mixed Reeb chords accordingly.

Lemma 73 *The formal dimension of the moduli space $\mathcal{M}^{\mathrm{fi}}(\boldsymbol{a};\kappa)$ is given by*

$$\dim \mathcal{M}^{\mathrm{fi}}(\boldsymbol{a};\kappa) = (n-3) - \sum_{j=1}^{m}(|a_j| - (n-2)).$$

The formal dimension of the moduli spaces $\mathcal{M}^{\mathrm{sy}}_{\|}(\boldsymbol{c};\kappa)$ is given by

$$\dim \mathcal{M}^{\mathrm{sy}}_{\|}(\boldsymbol{c};\kappa) = (n-3) + \sum_{\sigma_j=-1}(|c_j|+1) - \sum_{\sigma_j=+1}(|c_j| - (n-2)).$$

The formal dimension of the moduli spaces $\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})$ is given by

$$\dim \mathcal{M}^{\mathrm{sy}}(\boldsymbol{c}) = (n-3) + \sum_{\sigma_j=-1}(|c_j|+1) - \sum_{\sigma_j=+1}(|c_j| - (n-2)).$$

The formal dimension of the moduli space $\mathcal{M}^{\mathrm{co}}(\boldsymbol{c};\kappa)$ is given by

$$\dim \mathcal{M}^{\mathrm{co}}(\boldsymbol{c};\kappa) = 1 - \sum_{r=1^i}(|c_r| - (n-2)) + \sum_{s=1}^{j}(|c_{0;s}|+1).$$

The formal dimension of the moduli space $\mathcal{M}^{\overline{\mathrm{co}}}(\boldsymbol{c};\kappa)$ is given by

$$\dim \mathcal{M}^{\overline{\mathrm{co}}}(\boldsymbol{c};\kappa) = 1 - \sum_{r=1^i}(|c_r| - (n-2)) - \sum_{s=1}^{j}(|x_{0;s}| - (n-2)).$$

Proof See [17, Theorem A.1]. □

We next study topological properties of the moduli spaces just defined. It turns out to be comparatively simple because of two key features. First, since we require our disks to switch copies at punctures "in the same direction" they cannot be multiply covered, and second, for the same reason there can be no boundary splitting. As in [23], the first property allows us to prove transversality by perturbing the almost complex structure, and the second shows that the moduli spaces admit compactifications consisting only of punctured curves joined at Reeb chords or Lagrangian intersection points. Precise formulations of these results are as follows.

Theorem 74 *For a generic almost complex structure J the moduli spaces $\mathcal{M}^{\mathrm{fi}}(\boldsymbol{a};\kappa)$, $\mathcal{M}^{\mathrm{sy}}_{\|}(\boldsymbol{c};\kappa)$, $\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})$, $\mathcal{M}^{\mathrm{co}}(\boldsymbol{c};\kappa)$ and $\mathcal{M}^{\overline{\mathrm{co}}}(\boldsymbol{c};\kappa)$ are transversely cut out manifolds of respective dimensions* $\dim\mathcal{M}^{\mathrm{fi}}(\boldsymbol{a};\kappa)$, $\dim\mathcal{M}^{\mathrm{sy}}_{\|}(\boldsymbol{c};\kappa)$, $\dim\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})$, $\dim\mathcal{M}^{\mathrm{co}}(\boldsymbol{c};\kappa)$ *and* $\dim\mathcal{M}^{\overline{\mathrm{co}}}(\boldsymbol{c};\kappa)$.

Proof A well-known argument gives transversality for disks that are somewhere injective on the boundary by perturbing the almost complex structure: any element in the cokernel of the linearized operator must be zero on the set of injectivity and then identically zero by unique continuation. Our disks are not necessarily somewhere injective but the disks cannot be multiple covers and there is a region with the property of the region of injectivity above. We briefly recall the argument for this from [28, Lemma 4.5]. Fix a puncture of u.

Consider first the more difficult case when this puncture maps to a Lagrangian intersection. Pick coordinates so that the intersection point lies at the origin in $\mathbb{C}^n$ and so that the two Lagrangians correspond to $\mathbb{R}^n$ and $i\mathbb{R}^n$. Let $(x_1+iy_1,\dots,x_n+iy_n)$ be standard coordinates on $\mathbb{C}^n$. Consider the complex hyperplanes $H_{\pm\epsilon}=\{x_1+iy_1=\pm\epsilon(1+i)\}$. Looking at the Fourier expansion of u near the puncture it is clear that for suitable coordinates (such that the leading Fourier coefficient of u lies in the direction of the first coordinate) the number of intersection points of the image of u and $H_{\pm\epsilon}$ near the puncture have different parities depending on the sign of ϵ. With more details: choose coordinates in $\mathbb{C}^n$ for which the complex structure J agrees with the standard almost complex structure at the origin and coordinates $[0,\infty)\times[0,1]$ around the puncture in the domain. Then the argument in the proof of [25, Lemma 2.1] shows that u is conjugate to a standard holomorphic map $\tilde{u}$ with boundary condition $\mathbb{R}^n$ and $i\mathbb{R}^n$. That map has Fourier expansion

$$\tilde{u}(s+it)=\sum_{k\geq 0}c_k e^{-\left(k\pi+\frac{1}{2}\pi\right)(s+it)},\quad \text{with } c_k\in\mathbb{R}^n.$$

For the intersection with $H_{\pm\epsilon}$ pick coordinates so that the first Fourier coefficient c_k which is nonzero has the form $c_k=(a,0,\dots,0)$, with $a\neq 0$. By analytic continuation, other disks and half-disks mapping there either have images agreeing completely in the ball or there are injective points of the disk near the puncture. In the case where they agree completely, we note that the parity of the number of intersection points in each local sheet not containing the puncture with $H_{\pm\epsilon}$ is independent of the sign of $\pm\epsilon$. We then find that we can achieve transversality by perturbing the complex structure near $H_{\pm\epsilon}$: because the sheet with the puncture intersects only one of $H_{\pm\epsilon}$, if the contributions of the sheets mapping to $H_{-\epsilon}$ cancel then those mapping to $H_{+\epsilon}$ cannot cancel and vice versa, by unique continuation.

Once transversality is achieved the statement that solutions form manifolds follows from a well-known argument; see eg [28, Proposition 2.3]. □

Theorem 75 *The moduli space $\mathcal{M}^{\mathrm{fi}}(\boldsymbol{a};\kappa)$ admits a compactification consisting of several-level disks joined at Reeb chords and intersection points, where some levels may lie in the symplectization.*

The moduli spaces $\mathcal{M}^{\mathrm{sy}}_{\|}(\boldsymbol{c};\kappa)$ and $\mathcal{M}^{\mathrm{sy}}(\boldsymbol{c})$ admit compactifications consisting of several-level disks joined at Reeb chords.

The moduli space $\mathcal{M}^{\mathrm{co}}(\boldsymbol{a};\kappa)$ admits a compactification consisting of several-level disks joined at Reeb chords. There is one level of disks in the cobordisms and the remaining levels are in the symplectization ends.

The moduli space $\mathcal{M}^{\overline{\mathrm{co}}}(\boldsymbol{a};\kappa)$ admits a compactification consisting of several-level disks joined at Reeb chords and intersection points.

Proof The boundary conditions on our punctured disks have the following property: any arc in a disk with more than one positive puncture that subdivides the source into two components with a positive puncture in each must connect boundary components numbered with distinct numbers. This shows that there can be no boundary splitting. The theorem then follows from SFT compactness; see [22, Appendix B.1] for the curve with boundary version of [12]. □

We next discuss orientations of moduli spaces following [27]. We fix capping operators at all Reeb chords and Lagrangian intersection points so that the two capping operators there glue to a disk with the Fukaya orientation; see [36]. Recall that the relative spin structure on the Lagrangian submanifold induces an orientation on the determinant bundle over the space of disks with boundary condition in the Lagrangian; see [27] or

[28, Section 4.4]. As in [27] we see that these choices then induce a system of coherent orientations on the moduli spaces.

We will use one more property of the moduli spaces, which says that they are effectively independent of the increasing or decreasing boundary labeling κ.

Theorem 76 *Let κ and κ' be two increasing (decreasing) boundary numberings. Then there are canonical orientation-preserving diffeomorphisms*

$$\mathcal{M}^{\mathrm{fi}}(\boldsymbol{a};\kappa) \approx \mathcal{M}^{\mathrm{fi}}(\boldsymbol{a};\kappa'), \quad \mathcal{M}^{\mathrm{sy}}_{\|}(\boldsymbol{c};\kappa) \approx \mathcal{M}^{\mathrm{sy}}_{\|}(\boldsymbol{c};\kappa'), \quad \mathcal{M}^{\mathrm{co}}(\boldsymbol{c};\kappa) \approx \mathcal{M}^{\mathrm{co}}(\boldsymbol{c};\kappa').$$

Proof Let $\mathcal{M}(\kappa)$ denote either one of the above moduli spaces. This moduli space is the transverse zero set of a Fredholm section in a Banach bundle. Changing the numbering from κ to κ' corresponds to an arbitrarily small isotopy, which induces an arbitrarily small deformation of the section. The theorem follows. □

Appendix B Wrapped Floer cohomology and Legendrian surgery

In this section we present the argument that establishes the isomorphism between $\mathrm{CE}^*(\Lambda)$ and $\mathrm{CW}^*(C)$, where C is the cocore disk of the surgery. Our proof is a generalization of the corresponding result under Lagrangian handle attachment explained in [11], and uses the technical results on relevant moduli spaces in [25].

We first define a version of wrapped Floer cohomology using only purely holomorphic disks and show that the resulting theory agrees with the usual version defined in terms of holomorphic disks with a Hamiltonian term. Second we discuss the surgery isomorphism in [11], and third we discuss how to generalize that argument to partially wrapped Floer cohomology calculations.

B.1 Wrapped Floer cohomology without Hamiltonian

Let X be a Weinstein manifold and L be an exact Lagrangian. Fix a system of shifting Morse functions that are positive at infinity and let $\overline{L} = \{L_j\}_{j=0}^{\infty}$ be the corresponding family of parallel Lagrangian submanifolds. Define $\mathrm{CW}^*(L)$ to be the chain complex generated by Reeb chords of L and intersection points $L_0 \cap L_1$. We define operations $\mathfrak{m}_i$ on $\mathrm{CW}^*(L)$ using what we call *partial holomorphic buildings*.

We start in the simplest case when the output of $\mathfrak{m}_i$ is an intersection point c_0. Consider i generators $c_i, \dots, c_1$ and consider a disk D_{i+1} with a decreasing boundary numbering κ, distinguished negative (output) puncture and remaining punctures positive (inputs). Let $\boldsymbol{c}' = c_i \cdots c_1$ and $\boldsymbol{c} = c_0 c_i \cdots c_1$. Define

$$\mathfrak{m}_i'(\boldsymbol{c}') = \sum_{|c_0| = |\boldsymbol{c}'| + (2-i)} |\mathcal{M}^{\mathrm{fi}}(\boldsymbol{c};\kappa)| c_0.$$

Here we use the temporary notation $\mathfrak{m}_i'$ to denote the summand of the full operation $\mathfrak{m}_i$ that takes values in intersection points. We next turn to the more complicated definition of the part $\mathfrak{m}_i''$ of the operation that takes values in Reeb chord generators, and to this end we introduce the notion of a partial holomorphic building.

The domain of a partial holomorphic building is a possibly broken disk D_{i+1} with decreasing boundary numbering κ. The partial holomorphic buildings we consider always have exactly one disk in the symplectization. We call it the primary disk of the building. We require that the distinguished puncture is increasing and is a negative puncture of this primary disk. If the distinguished puncture is the only negative puncture of the primary disk then the partial building consists only of its primary component. If on the other hand the primary disk has additional negative punctures then we require that at each additional negative puncture (which is decreasing or constant) there is a disk in the filling with decreasing boundary condition that is attached at its distinguished increasing or constant puncture to the additional negative puncture. We call these disks the secondary disks of the partial building. The resulting partial holomorphic building is then a disk with domain a broken D_{i+1}, with distinguished puncture a negative puncture at a Reeb chord and with remaining i punctures either Reeb chords or intersection points. See Figure 18.

Remark 77 At additional negative punctures there may be holomorphic disks with one positive puncture and boundary on L attached. These are the usual augmentation disks, or disks on L used as anchoring disks in the definition of $\mathcal{M}^{\mathrm{sy}}_{\|}(\boldsymbol{c},\kappa)$.

We write the punctures of the partial holomorphic disk building as $\boldsymbol{c} = c_0 c_i \cdots c_1$, where c_0 is the distinguished puncture. Write

$$\mathcal{M}^{\mathrm{pb}}(\boldsymbol{c};\kappa)$$

for the moduli space of partial holomorphic disk buildings with boundary condition on $\bar{L}$ according to κ. Using this we define for generators $\boldsymbol{c}' = c_i \cdots c_1$ the operation

$$\mathfrak{m}_i''(\boldsymbol{c}') = \sum_{|c_0| = |\boldsymbol{c}'| + (2-i)} |\mathcal{M}^{\mathrm{pb}}(\boldsymbol{c};\kappa)| c_0,$$

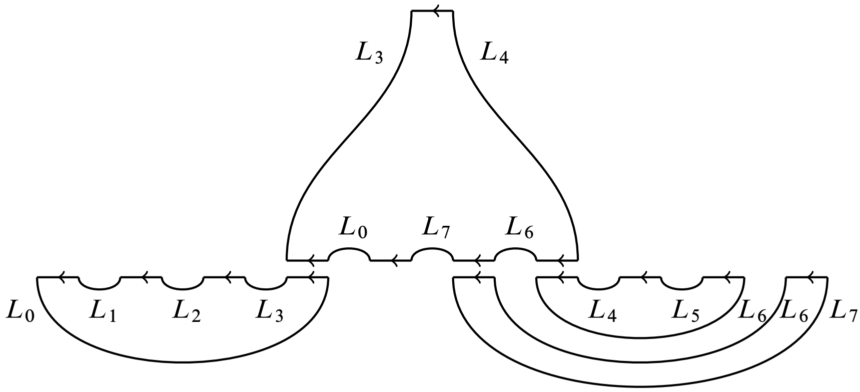

Figure 18: The domain of a partial building contributing to the operation $\mathfrak{m}_7$ of $\mathrm{CW}^*(L)$ with a possible decoration. The map sends the tensor product of chords $L_0 \leftarrow L_1, L_1 \leftarrow L_2, \dots, L_6 \leftarrow L_7$ to a chord $L_0 \leftarrow L_7$.

where the sum ranges over Reeb chords c_0 with grading as indicated. Finally we define the total operation $\mathfrak{m}_i$ as the sum

$$\mathfrak{m}_i(\boldsymbol{c}') = \mathfrak{m}'_i(\boldsymbol{c}') + \mathfrak{m}''_i(\boldsymbol{c}').$$

Lemma 78 *The A_∞–relations hold for the operations $\mathfrak{m}_i$.*

Proof First, Theorem 76 shows that the operations compose and that they are independent of the choice of decreasing boundary numbering. To see that the relations hold, we will as usual identify the terms contributing to them with the boundary of an oriented 1–dimensional compact manifold.

To this end we first consider 1–dimensional moduli spaces $\mathcal{M}'$ of the form $\mathcal{M}' = \mathcal{M}^{\mathrm{fi}}(\boldsymbol{c}; \kappa)$, where the distinguished puncture c_0 is an intersection point. As usual, the boundary numbering precludes boundary bubbling and we find that the boundary consists of broken disks that either break at an intersection point, in which case the holomorphic parts both have dimension zero, or break into a partial holomorphic building with a rigid disk attached at its negative puncture, in which case the primary component of the partial building has dimension one. We find the boundary points of $\mathcal{M}'$ are in one-to-one correspondence with disks contributing to compositions of $\mathfrak{m}'_i$ and $\mathfrak{m}'_j$ (disks breaking at intersection points) and disks contributing to $\mathfrak{m}''_i$ and $\mathfrak{m}'_j$.

The remaining contributions to the A_∞–relations correspond to compositions of $\mathfrak{m}''_i$ and $\mathfrak{m}''_j$. We show that all contributions to this composition constitute the boundary of an oriented 1–manifold. The contributions are of two forms: either the output puncture of the first operation (which lies in the primary disk of the corresponding partially broken configuration) is glued to an input puncture of the primary disk in the partially broken configuration of the second operation, or it is glued to an input puncture in a secondary disk.

The configurations of the former type correspond to a part of the boundary of the moduli space $\mathcal{M}^{\mathrm{sy}}_{\|}(\boldsymbol{b};\kappa)$ of dimension two with distinguished negative puncture and decreasing boundary numbering (after we divide out the natural $\mathbb{R}$–action this is a 1–dimensional space) capped off by rigid disks in $\mathcal{M}^{\mathrm{co}}(\boldsymbol{a})$ at all nondistinguished negative punctures. This is the part of the boundary where the distinguished puncture belongs to the lower level disk.

The configurations of the second type correspond to the part of the boundary of the 1–dimensional moduli space $\mathcal{M}^{\mathrm{co}}(\boldsymbol{a};\kappa)$, with a distinguished increasing positive puncture where a negative puncture in the primary disk of the second operation is attached (other negative punctures in the primary disk of the second operation are capped off as usual). The part of the boundary containing the distinguished positive puncture lies in the rigid disk in the cobordism.

Finally, the remaining part of the boundary in the first case is two-level buildings in $\mathcal{M}^{\mathrm{sy}}_{\|}(\boldsymbol{b};\kappa)$, where the distinguished negative puncture belongs to the top-level curves. These are exactly the configurations that we get from the remaining parts of the boundary in the second case (ie configurations where the distinguished puncture belongs to the component in the symplectization) when we glue to it the primary disk of the second operation.

We conclude that also the composition of $\mathfrak{m}''_i$ and $\mathfrak{m}''_j$ cancels. The lemma follows. □

B.1.1 Isomorphism with the Hamiltonian version In this section we show that the above definition of wrapped Floer cohomology agrees with the standard theory. Similar results can be found in [30; 29]. Here we will give a sketch. We keep the geometric setting as above and write $CW^*_{\mathrm{Ham}}(L)$ for the usual version of Hamiltonian wrapped Floer cohomology. We give a brief recollection of the definition.

We define the wrapped Floer cohomology complex $CW^*_{\mathrm{Ham}}(L)$ of L as follows. Write $X=\overline{X}\cup[0,\infty)\times Y$, where $\overline{X}$ is a compact domain and $[0,\infty)\times Y$ the positive end of the Weinstein manifold X. Consider time-dependent Hamiltonians $H_a\colon X\times[0,1]\to\mathbb{R}$

which are perturbations of functions that equal 0 on $\overline{X}$ and are linear of the form

$$(t, y) \mapsto ae^t + b \quad \text{on } [0, \infty) \times Y,$$

where a is not in the chord and orbit spectrum or the contact form on Λ. We choose these Hamiltonians in such a way that if $a_0 < a_1$ then $H_{a_0} < H_{a_1}$ on X. After a small perturbation, Hamiltonian time 1 chords and Hamiltonian time 1 orbits are nondegenerate.

Define the chain complex $\mathrm{CW}^*_{\mathrm{Ham}}(L, H_a)$ to be generated by Hamiltonian chords $\gamma\colon [0,1] \to X$ of C of action

$$\mathfrak{a}(\gamma) = \int_0^1 \big(\lambda(\dot\gamma(t)) - H_a(\gamma(t))\big)\, dt < a.$$

The differential on $\mathrm{CW}^*(L, H_a)$ is defined by counting solutions of the perturbed Cauchy–Riemann equation over the strip with coordinates $s + it \in \mathbb{R} \times [0,1]$:

$$(du + X_{H_a} \otimes dt)^{0,1} = 0.$$

Choosing an increasing interpolation between H_{a_0} and H_{a_1} we get continuation maps

$$\mathrm{CW}^*_{\mathrm{Ham}}(L, H_{a_0}) \mapsto \mathrm{CW}^*_{\mathrm{Ham}}(L, H_{a_1}),$$

and we define the wrapped Floer cohomology complex as the direct limit

$$\mathrm{CW}^*_{\mathrm{Ham}}(L) = \varinjlim_a \mathrm{CW}^*_{\mathrm{Ham}}(C, H_a).$$

The wrapped Floer cohomology $\mathrm{HW}^*_{\mathrm{Ham}}(C)$ is the homology of this complex. Writing $\mathrm{HW}^*_{\mathrm{Ham}}(L, H_a)$ for the homology of $\mathrm{CW}^*_{\mathrm{Ham}}(L, H_a)$ we then have

$$\mathrm{HW}^*_{\mathrm{Ham}}(L) = \varinjlim_a \mathrm{HW}^*_{\mathrm{Ham}}(L, H_a),$$

by exactness of direct limits.

A well-known argument shows that $\mathrm{CW}^*(L)$ with differential $\mathfrak{m}_1$ is quasi-isomorphic to the wrapped Floer cohomology by a geometrically defined chain map [29]. We recall the argument here.

The filling L of Λ gives an augmentation of $\mathrm{CE}^*(\Lambda)$ and we define $\mathrm{CW}^*(L)$ (as a chain complex disregarding higher product operations) without Hamiltonian as the "Morse extended linearized Chekanov–Eliashberg complex" with respect to this augmentation as generated by Reeb chords and the critical point of a Morse function on L with a unique minimum, and take the differential to count unperturbed augmented and anchored holomorphic strips. We also introduce the subcomplexes $\mathrm{CW}^*(L,a)$ generated by chords of action $< a$. Then, by definition,

$$\mathrm{CW}^*(L) = \varinjlim_{a} \mathrm{CW}^*(L,a).$$

The isomorphism $\mathrm{CW}^*(L) \to \mathrm{CW}^*_{\mathrm{Ham}}(L)$ is now constructed by interpolating exactly as in the continuation maps above from the zero Hamiltonian (ordinary Cauchy–Riemann equation) to the Hamiltonians H_a above. Choosing the interpolations compatibly, we get the commutative diagram

$$\begin{array}{ccccccccc}
\mathrm{CW}^*(L,a_0) & \longrightarrow & \mathrm{CW}^*(L,a_1) & \longrightarrow & \cdots & \longrightarrow & \mathrm{CW}^*(L,a_j) & \longrightarrow & \cdots \\
\downarrow & & \downarrow & & & & \downarrow & & \\
\mathrm{CW}^*_{\mathrm{Ham}}(L, H_{a_0}) & \longrightarrow & \mathrm{CW}^*_{\mathrm{Ham}}(L, H_{a_0}) & \longrightarrow & \cdots & \longrightarrow & \mathrm{CW}^*_{\mathrm{Ham}}(L, H_{a_j}) & \longrightarrow & \cdots
\end{array}$$

Here all vertical arrows are chain isomorphisms by the standard argument — see for instance [30, Section 6] — and taking limits we find a chain isomorphism

$$\mathrm{CW}^*(L) \to \mathrm{CW}^*_{\mathrm{Ham}}(C).$$

We extend this chain map to an A_∞–map, then the standard spectral sequence argument establishes the desired A_∞–quasi-isomorphism.

We follow the approach in [30], where similar isomorphisms between contact and symplectic differential graded algebras were constructed. More precisely, we construct a splitting compatible nonnegative field of 1–forms with values in Hamiltonian vector fields, and further a 1–parameter family of such forms interpolating between the zero Hamiltonian at the positive end and the Hamiltonian used to define wrapped Floer cohomology at the negative end [30, Section 2]. We then define the corresponding moduli spaces over the deformation interval. Keeping the notation from [30] we write

$$\mathcal{F}_{\mathbb{R}}(\boldsymbol{a}, b).$$

In order for the asymptotics at infinity of these maps to make sense we need to include the parallel copies of the Lagrangians according to boundary numbering, and in particular also to incorporate this in the description of wrapped Floer cohomology. More precisely, as in the case above we will have moduli spaces of Floer holomorphic disks with boundary in distinct Lagrangians that are arbitrarily close. The analogue of Theorem 76 holds by the same argument and the corresponding moduli spaces are canonically isomorphic for sufficiently small perturbations. Using these observations we then define $\Phi\colon \mathrm{CW}^*(L) \to \mathrm{CW}^*_{\mathrm{Ham}}(L)$ by

$$\Phi(\boldsymbol{a}) = \sum_{\dim \mathcal{F}_{\mathbb{R}}(\boldsymbol{a},b)=0} |\mathcal{F}_{\mathbb{R}}(\boldsymbol{a},b)|\,b.$$

Lemma 79 *The map Φ is an A_∞–homomorphism.*

Proof To see this we note again that the disks which contributes to the A_∞ relations correspond exactly to the ends of 1–dimensional moduli space. □

Lemma 80 *The map Φ is a quasi-isomorphism.*

Proof The map respects the word-length filtration and is the standard isomorphism from the linearized Legendrian cohomology to the wrapped Floer cohomology, discussed above, on the E_2–page. □

B.2 Wrapped Floer cohomology and Lagrangian handle attachment

In this subsection we prove the results in [11] giving a Legendrian surgery description of the wrapped Floer cohomology of a cocore disk in a Weinstein manifold obtained by Lagrangian handle attachment along a Legendrian sphere, referring to [25] for the results on holomorphic curves missing in [11]. To state this result we first introduce notation.

Suppose that X_0 is a Weinstein $2n$–manifold with ideal boundary the contact $(2n-1)$–manifold Y_0. Let $\Lambda = \Lambda_1 \cup \cdots \cup \Lambda_m$ be a Legendrian submanifold such that all of its components Λ_j are parametrized $(n-1)$–spheres. Let X be the Weinstein manifold that results from attaching Lagrangian handles H to Λ. Here $H = H_1 \cup \cdots \cup H_m$, where each component H_j is a disk subbundle of the cotangent bundle T^*D of the n–disk D, and where H_j is attached to Λ_j. Then X contains m cocore disks corresponding to the cotangent fibers at the center of the disk in each H_j. We let $C_j \subset X$ denote the cocore disk in H_j, let $\Gamma_j \subset Y$ denote its Legendrian boundary inside the contact boundary Y of X, and write $\Gamma = \Gamma_1 \cup \cdots \cup \Gamma_m$.

As a first step in the calculation of the wrapped Floer cohomology of C we describe the generators of the underlying chain complex. By definition — see Section B.1 — generators of $\mathrm{CW}^*(C)$ are of two kinds: Lagrangian intersection points and Reeb chords. Here the Lagrangian intersection points are easily understood: pick the shifting Morse function so that it has one minimum on each component of C and no other critical

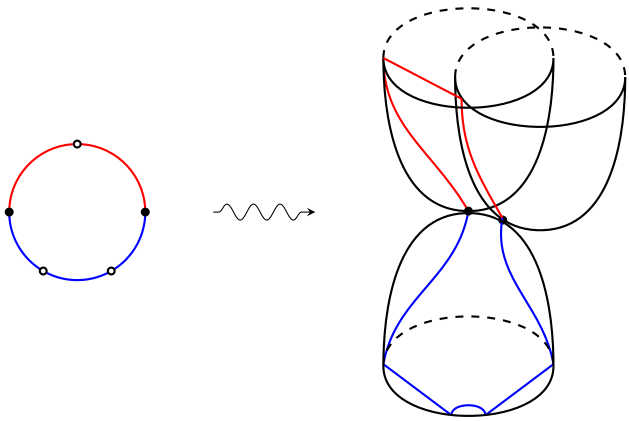

Figure 19: A picture illustrating a curve contributing to Φ^1 of the A_∞–functor Φ.

points, then there is exactly one intersection point for each component of C. We denote the intersection point of C_j by m_j and we denote the subcomplex generated by the m_j by $\mathrm{CW}_0^*(\Gamma)$. Remaining generators are Reeb chords of Γ we write $\mathrm{CW}_+^*(\Gamma)$ for the quotient complex $\mathrm{CW}^*(\Gamma)/\mathrm{CW}_0^*(\Gamma)$ and note that CW_+^* is generated by Reeb chords.

Consider the link Λ and let all components be decorated by minus, $\Lambda^- = \Lambda$. Consider $\mathrm{CE}^*(\Lambda)$ as a chain complex, generated by composable words of Reeb chords with differential d and with product $\cdot$ given by concatenation if the words are composable and zero otherwise. Let $\epsilon > 0$ denote the size of the attaching region, ie the size of the tubular neighborhood of Λ where H is attached. We then have the following:

Lemma 81 *For any $A > 0$ there exists $\epsilon_0 > 0$ such that if $\epsilon < \epsilon_0$ then there is a natural one-to-one correspondence between the generators of $\mathrm{CW}_+^*(\Gamma)$ (Reeb chords of Γ) of action $< A$ and the generators of $\mathrm{CE}^*(\Lambda)$ (words of Reeb chords of Λ) of action $< A$.*

Proof This is [25, Theorem 1.2]. □

We will next define the surgery map, which is an A_∞–morphism

$$\Phi\colon \mathrm{CW}^*(\Gamma) \to \mathrm{CE}^*(\Lambda)$$

that counts certain holomorphic disks. See Figure 19.

As in Appendix A, consider the disk D_{i+j+2} with two special punctures subdividing the boundary into an upper and a lower arc with i and j punctures, respectively, and with a boundary numbering in the upper arc. Let $\boldsymbol{c}_0 = c_{0;1} \cdots c_{0;j}$ be a composable word of Reeb chords connecting Λ_v to Λ_w, and let $c_i \cdots c_1$ be a word of generators of $\mathrm{CW}^*(C)$. Consider the word of Reeb chords and intersection points

$$\boldsymbol{c} = c_{0;1} \cdots c_{0;j} z^w c_i \cdots c_1 z^v.$$

Define $\Phi_i : \mathrm{CW}^*(C)^{\otimes i} \to \mathrm{CE}^*(\Lambda)$ by

$$\Phi_i(\boldsymbol{c}') = \sum_{|\boldsymbol{c}_0| = |\boldsymbol{c}'| + i(n-2)} |\mathcal{M}^{\mathrm{co}}(\boldsymbol{c})| \boldsymbol{c}_0.$$

Remark 82 If m^v is the minimum of the Morse function on C^v as above, then

$$\Phi_1(m^v) = e_v,$$

because of the unique holomorphic disk corresponding to the flow line from the minimum to the intersection point between $C^v \cap L$; for the parallel copies this gives a triangle with corners at $m^v = C_0^v \cap C_1^v$, at $C_0^v \cap L$ and at $C_1^v \cap L$, and since there are no negative punctures the output is e_v. Also, if a word $\boldsymbol{c}'$ of generators of $\mathrm{CW}^*(C)$ contains a generator m^v and has length $i > 1$, then

$$\Phi_i(\boldsymbol{c}') = 0,$$

as this corresponds — see Lemma 35 — to a holomorphic disk with a flow line from the minimum attached, and such a configuration cannot be rigid unless the disk is constant.

Theorem 83 *The maps Φ_i give an A_∞–map $\mathrm{CW}^*(C) \to \mathrm{CE}^*(\Lambda)$, which is an A_∞–quasi-isomorphism.*

Proof In order to see the A_∞–relations we study the boundary of the moduli space $\mathcal{M}^{\mathrm{co}}(\boldsymbol{c})$. As usual the boundary numbering guarantees that there is no boundary splitting on C. The fact that there is no boundary splitting on L follows from Stokes' theorem: such a splitting would give a disk without positive puncture. The boundary of the moduli space thus consists of the following configurations:

(i) Two level curves with one level in the cobordism and one in either symplectization end.

(ii) Curves which split at the intersection point $C \cap L$.

Splitting (i) corresponds to the map followed by the operation d in $\mathrm{CE}^*(\Lambda)$ when the symplectization disk lies in the negative end, and to an operation in $\mathrm{CW}^*(C)$ followed by the map when the symplectization disk lies in the positive end. Splitting (ii) corresponds to the tensor product of the map followed by the product operation $\cdot$ in $\mathrm{CE}^*(\Lambda)$. The A_∞–relations follow.

To see that Φ is a quasi-isomorphism we argue as follows. We first fix an action cut-off $A > 0$ and show that Φ_1 induces an isomorphism on homology below action A by constructing algebraically one holomorphic disk interpolating between a Reeb chord of Γ and the corresponding word of Reeb chords of Λ. For a complete proof we refer to [25, Theorem 1.3]; the argument is roughly as follows. One starts from unique and uniformly transversely cut out such disks for single chord words obtained by a straightforward explicit geometric construction. Gluing such disks at their Lagrangian intersection punctures in $L \cap C$ and using small action to rule out all breakings except one, we find that there is algebraically one disk interpolating between a chord on Γ and the corresponding word of chords of Λ. Together with Remark 82, which shows $\Phi_1(m^v) = e^v$, the existence of such disks implies that the map Φ_1 has a triangular matrix with respect to the action filtration and hence is a chain isomorphism (compare [11, Section 6.2]): since Φ_1 is an isomorphism on the subquotients of the action filtration (and the isomorphism on generators is given by the bijection given in Lemma 81), Φ_1 is an isomorphism below action A for any A. The A_∞–isomorphism below action $A > 0$ then follows from the usual spectral sequence argument.

To see that we get an isomorphism on the full complex we show that the isomorphisms discussed are compatible with action limits. More precisely, in order to increase the action limit $A > 0$ for the one-to-one correspondence between Reeb chords of Γ of action $< A$ and words of Reeb chords of Λ of action $< A$, we must shrink the size $\delta > 0$ of the handle attached. Consider attaching a handle of size $\delta > 0$ to Λ and denote the resulting new Weinstein manifold by X_δ and the cocore disk by $C_\delta \subset X_\delta$.

If $\delta_0 > \delta_1$, then by the isomorphism in Lemma 80 and standard results for wrapped Floer cohomology — see eg [30, Section 5.5] — there is a cobordism map

$$\mathrm{CW}^*_+(C_{\delta_0}) \to \mathrm{CW}^*_+(C_{\delta_1}),$$

which is a quasi-isomorphism. Moreover, by the surgery description of chords for any $A > 0$ there exists $\delta_1 > 0$ such that the above map has ± 1 on the diagonal (with respect to the identification of generators in Lemma 81) for all chords of Γ and words of chords of Λ of action $< A$.

Consider the directed system

$$\mathrm{CW}^*_+(C_{\delta_0}) \to \mathrm{CW}^*_+(C_{\delta_1}) \to \cdots \to \mathrm{CW}^*_+(C_{\delta_j}) \to \cdots , \tag{43}$$

where $\delta_j \to 0$, and let

$$\overline{\mathrm{CW}}^*_+(C) = \varinjlim_{\delta} \mathrm{CW}^*_+(C_\delta).$$

Then the homology $\overline{\mathrm{HW}}^*_+(C)$ of $\overline{\mathrm{CW}}^*_+(C)$ satisfies

$$\overline{\mathrm{HW}}^*_+(C) = \varinjlim_{\delta} \mathrm{HW}^*_+(C_\delta) = \mathrm{HW}^*_+(C_{\delta_j}) \quad \text{for any fixed } j.$$

Here the last equality follows since all the arrows in the directed system of homology groups of (43) are isomorphisms.

Consider next the Chekanov–Eliashberg algebra $\mathrm{CE}^*(\Lambda)$ of Λ. We define the action-truncated subcomplex $\mathrm{CE}^*(\Lambda, a)$ generated by words of chords of total action $< a$. Viewing $\mathrm{CE}^*(\Lambda)$ as a chain complex generated by words of chords, we then have

$$\mathrm{CE}^*(\Lambda) = \varinjlim_{a} \mathrm{CE}^*(\Lambda, a).$$

For each a_j the surgery map gives $\delta_j > 0$ such that the map

$$\overline{\mathrm{CW}}^*_+(C_{\delta_j}) \to \mathrm{CE}^*(\Lambda, a_j)$$

is a chain isomorphism with ± 1 on the diagonal below action a_j. By definition of surgery and cobordism maps, the diagram

$$\begin{tikzcd}
\overline{\mathrm{CW}}^*_+(C_{\delta_0}) \arrow[r] \arrow[d] & \overline{\mathrm{CW}}^*_+(C_{\delta_1}) \arrow[r] \arrow[d] & \cdots \arrow[r] \arrow[d, phantom, "\vdots"] & \overline{\mathrm{CW}}^*_+(C_{\delta_j}) \arrow[r] \arrow[d] & \cdots \\
\mathrm{CE}^*(\Lambda, a_0) \arrow[r] & \mathrm{CE}^*(\Lambda, a_1) \arrow[r] & \cdots \arrow[r] & \mathrm{CE}^*(\Lambda, a_j) \arrow[r] & \cdots
\end{tikzcd} \tag{44}$$

commutes, where $\delta_{j+1} < \delta_j$ and $a_j < a_{j+1}$. Taking limits of the sequences we get a chain map

$$\overline{\mathrm{CW}}^*_+(C) \to \mathrm{CE}^*(\Lambda). \tag{45}$$

Taking the limits of the sequence (44) on the homology level and using that all vertical arrows are homology isomorphisms then gives homology isomorphisms in the limit, and (45) is a quasi-isomorphism inducing an isomorphism

$$\overline{\mathrm{HW}}^*(C) \approx HCE^*(\Lambda).$$

The above gives a homology isomorphism of chain complexes. To consider also products one uses the exact same argument. The product operation on $\overline{\mathrm{CW}}^*(C)$ is induced from the action-truncated version

$$\overline{\mathrm{CW}}^*(C, a_1) \otimes \cdots \otimes \overline{\mathrm{CW}}^*(C, a_m) \to \overline{\mathrm{CW}}^*(C, a_1 + \cdots + a_m),$$

and similarly, on CE^*,

$$\mathrm{CE}^*(C, a_1) \otimes \mathrm{CE}^*(C, a_2) \to \mathrm{CE}^*(C, a_1 + a_2). \qquad \square$$

Remark 84 In the above proof we obtain the isomorphism by taking smaller and smaller handles. To see that such a procedure is necessary note that the correspondence between words of chords and chords is true only below an action limit determined by the size of the handle. For larger actions there are Reeb flows before the surgery that hit the neighborhood of the Legendrian without being close to a chord and which could give chords after the surgery.

Remark 85 There is also an "upside-down" perspective on the surgery just described. Namely, one can start from the contact manifold Y and produce the contact manifold Y_0 by doing so-called $+1$–surgery on Γ. In complete analogy with the above, one shows that Reeb chords on Λ are in natural one-to-one correspondence with words of Reeb chords on Γ, and one can construct an upside-down surgery map of A_∞–coalgebras,

$$\mathrm{BCW}^*(C) \to \mathrm{LC}_*(\Lambda).$$

A similar argument also shows that this map is a quasi-isomorphism. Alternatively, one can prove this from the original surgery map using only algebra as follows. First write $\mathrm{CE}^*(\Lambda) = \Omega\mathrm{LC}_*(\Lambda)$. Then

$$\mathrm{BCW}^*(C) \simeq \mathrm{B}\Omega\mathrm{LC}_*(\Lambda) \simeq \mathrm{LC}_*(\Lambda),$$

since $\mathrm{LC}_*(\Lambda)$ is conilpotent; see Section 2.2.2.

B.3 Legendrian surgery and stopped wrapping

In this section we outline a surgery approach to the computation of wrapped Floer cohomology in a Weinstein manifold X with wrapping stopped by a Legendrian Λ in its boundary. We will use the following model for the ambient manifold. Fix a tubular neighborhood of Λ in the contact boundary Y of X. Attach a disk-bundle neighborhood of the zero section in $T^*([0,\infty)\times\Lambda)$ along the boundary $T^*([0,\infty)\times\Lambda)|_{0\times\Lambda}$, just like in Lagrangian handle attachment. We use a Liouville vector field on this domain that agrees with the standard Liouville vector field pointing outwards along fibers in the cotangent bundle over $[T,\infty)\times\Lambda$ for some $T>0$. Let the components of Λ be denoted by Λ_v for $v\in Q_0$. Fix a basepoint $p_v\in\Lambda_v$ for each v. Let $C^{v;\tau}$ denote the cotangent fiber $T^*_{(p_v,\tau)}([0,\infty)\times\Lambda)$. We compute the wrapped Floer cohomology of $C^\tau=\bigcup_{v\in Q_0}C^{v;\tau}$ for sufficiently large τ using a surgery approach. A straightforward monotonicity argument shows that the noncompactness of the cotangent bundle $T^*\Lambda\times[0,\infty)$ does not interfere with the compactness results for holomorphic curves used in the definition of wrapped Floer cohomology.

We first consider the surgery map into the Chekanov–Eliashberg algebra with loop space coefficients. Consider all components of Λ decorated by a positive sign $\Lambda = \Lambda^+$ and consider $\mathrm{CE}^*(\Lambda)$, which now involves, aside from Reeb chords, also chains $C_{-*}(\Omega\Lambda)$ on the based loop space. We define an A_∞–map

$$\Phi\colon \mathrm{CW}^*(C) \to \mathrm{CE}^*(\Lambda),$$

where the A_∞–structure on the right-hand side is the standard DG–algebra structure induced by concatenation and the Pontryagin product (as defined in this paper). As in Appendix A, consider a disk D_{i+j+2} with two dividing punctures that subdivides the boundary into two arcs, lower and upper. Let the upper arc contain i boundary punctures and be equipped with a decreasing boundary numbering κ, and the lower arc have j boundary punctures. Let $\boldsymbol{c}' = c_i \cdots c_1$ be Reeb chords of C and let $\boldsymbol{c}_0 = c_{0;1} \cdots c_{0;j}$ be Reeb chords of Λ. Let

$$\boldsymbol{c} = c_{0;1} \cdots c_{0;j} z^v c_i \cdots c_1 \cdots z^w,$$

and consider $\mathcal{M}^{\mathrm{co}}(\boldsymbol{c};\kappa)$, again as in Figure 19.

Theorems 74 and 75 show that this moduli space carries a fundamental chain. We view this chain as parametrizing chains of paths in Λ connecting the Reeb chord endpoints in $\boldsymbol{c}_0$. We write $[\mathcal{M}^{\mathrm{co}}(\boldsymbol{c})]$ for the alternating word of chains of loops and Reeb chords and view it as an element in $\mathrm{CE}^*(\Lambda)$. Define the maps

$$\Psi_i\colon \mathrm{CW}^*(C)^{\otimes_\kappa i} \to \mathrm{CE}^*(\Lambda) \quad \text{by} \quad \Psi_i(\boldsymbol{c}') = \sum_{\boldsymbol{c}_0} [\mathcal{M}^{\mathrm{co}}(\boldsymbol{c})].$$

Theorem 86 *The map $\Psi\colon \mathrm{CW}^*(C) \to \mathrm{CE}^*(\Lambda)$ is an A_∞–map.*

Proof To see that the A_∞–relations hold we look at the boundary of the moduli space $\mathcal{M}^{\mathrm{co}}(\boldsymbol{c})$ of dimension d. The codimension-one boundary consists of three splitting types:

(i) A one-dimensional curve splits off in the positive symplectization end.

(ii) A curve splits off at the negative end.

(iii) Splitting at one of the intersection points z^v.

In order for splittings of the form (i) to contribute to the codimension-one boundary of the moduli space, the part of the holomorphic building in W consists of rigid disks with only positive punctures attached at one puncture to a negative puncture of the disk in the positive end and a disk of dimension $d-1$ in $\mathcal{M}^{\mathrm{co}}(\boldsymbol{b})$ attached at the remaining negative puncture. (Splittings where the dimension of the components of the holomorphic building are distributed differently have higher dimension along the disk C and correspond to "hidden faces" from the point of view of $C_*(\Omega\Lambda)$.) Assembling the rigid disks and the one-dimensional disk we get a partial holomorphic disk building that contributes to the A_∞–operations in $\mathrm{CW}^*(L)$ followed by the map Ψ. Splittings of type (ii) correspond to the map Ψ followed by the differential μ_1 in $\mathrm{CE}^*(\Lambda)$. Finally splittings of type (iii) correspond to the map Ψ followed by the product μ_2 on $\mathrm{CE}^*(\Lambda)$. We conclude that the terms contributing to A_∞–relations express the codimension-one boundary of $[\mathcal{M}^{\mathrm{co}}(\boldsymbol{c})]$ in two different ways and hence Ψ is an A_∞–map. □

Remark 87 In the boundary of the moduli space $\mathcal{M}^{\mathrm{co}}(\boldsymbol{c})$ considered in the proof of Theorem 86 there are also higher-dimensional curves splitting off in the positive symplectization end. Such splittings contribute neither to the codimension-one boundary of the chains of loops, nor to the operations in the wrapped Floer cohomology, and hence play no role in the A_∞ chain map equation.

We will use slight generalizations of the map Ψ. More precisely, if p_j for $j = 1, \dots, m$ are points in $[0,\infty) \times \Lambda$ and if F_j is the cotangent fiber at p_j, then we have a similar surgery map

$$\Psi^{p_i p_j} : \mathrm{CW}^*(F_i, F_j) \to \mathrm{CE}^*_{p_i p_j}(\Lambda),$$

which counts holomorphic disks with a positive Reeb chord connecting F_i to F_j, two Lagrangian intersection punctures at p_i and at p_j, and a word of chains of loops in Λ and Reeb chords of Λ as output, and where CE^*_{ij} is directly analogous to CE^* but where the first chain of loops is in a word is replaced by a chain of paths from p_i to the basepoint and the last is replaced by a chain of paths from the basepoint to p_j. In this setup the counterpart of the second component Ψ_2 is

$$\Psi^{p_i p_j p_k} : \mathrm{CW}^*(F_j, F_k) \otimes \mathrm{CW}^*(F_i, F_j) \to \mathrm{CE}^*_{p_i p_k}(\Lambda), \tag{46}$$

and counts disks with two positive punctures at Reeb chords and two Lagrangian intersection punctures at p_i and p_k. The counterpart of the A_∞–equations in this setup is then

$$d \circ \Psi^{p_i p_k} + \Psi^{p_j p_k} \cdot \Psi^{p_i p_j} + \Psi^{p_i p_j p_k} \circ (1 \otimes \mu_1 + \mu_1 \otimes 1) + \Psi^{p_i p_k} \circ \mu_2 = 0, \tag{47}$$

where d is the differential on CE^*_{ik} and $\cdot$ is the (Pontryagin) product

$$\mathrm{CE}^*_{p_j p_k} \otimes \mathrm{CE}^*_{p_i p_j} \to \mathrm{CE}^*_{p_i p_k}.$$

The proofs of these statements are word for word repetitions of the proof of Theorem 86.

We next sketch a proof that the map Ψ in Theorem 86 is in fact a quasi-isomorphism, or, in other words, that its first component Ψ_1 induces an isomorphism on homology. We filter $\mathrm{CW}^*(C)$ by action of its Reeb chord generators. To get a corresponding filtration on $\mathrm{CE}^*(\Lambda)$ we use action on Reeb chords in combination with the energy on the loops. We start with a discussion of the energy of loops, following [52].

Equip Λ with a Riemannian metric and let $\Omega = \Omega(\Lambda)$ denote the space of based loops in Λ with the supremum norm: for two loops $\gamma, \beta \colon [0,1] \to \Lambda$,

$$d^*(\gamma,\beta) = \sup_{t\in[0,1]} \rho(\gamma(t),\beta(t)),$$

where ρ is the metric on Λ induced by the Riemannian structure. Then the metric topology on Ω agrees with the standard compact open topology.

Let $\Omega' = \Omega'(\Lambda)$ denote the space of piecewise smooth paths with metric

$$d(\gamma,\beta) = d^*(\gamma,\beta) + \int_0^1 (|\dot\gamma| - |\dot\beta|)^2 \, dt,$$

where $\dot\gamma$ denotes the derivative of γ. The natural inclusion $\Omega' \to \Omega$ is a homotopy equivalence [52, Theorem 17.1]. We will use finite-dimensional approximations to study Ω'. The energy of a piecewise smooth loop in Λ is

$$E(\gamma) = \int_0^1 |\dot\gamma|^2 \, dt.$$

For $c > 0$, let $\Omega^c \subset \Omega'$ denote the subset of loops of energy $E < c$. The space Ω^c can be approximated by piecewise geodesic loops. More precisely, fixing a subdivision $0 = t_0 < t_1 < \cdots < t_m = 1$ of $[0,1]$ we consider the space B^c of loops of energy $E < c$ that are geodesic on each interval $[t_i, t_{i+1}]$. Then [52, Lemma 16.1] shows that for all sufficiently fine subdivisions, B^c is a finite-dimensional manifold (a submanifold of the product $\Lambda^{\times m}$ in a natural way). Moreover, by [52, Theorem 16.2], all critical points of $E|_{\Omega^c}$ lie in B^c, which is a deformation retract of Ω^c, and for a generic metric $E|_{B^c}$ is a Morse function.

With these preliminaries established we turn to the actual proof. The first step will be to describe the Reeb chords of C^τ. Let g be a Riemannian metric on Λ as above and let $f \colon [0,\infty) \to \mathbb{R}$ be a positive function with $f(0) = 1$, $f'(0) = -1$ and $f'(t) < 0$ monotone increasing. Define the metric h on $\Lambda \times \mathbb{R}$ by

$$h = dt^2 + f(t) g.$$

Then, if x and y are points in Λ and $\gamma\colon [0,s] \to \Lambda$ is a geodesic with $\gamma(0) = x$ and $\gamma(s) = y$, there is a unique geodesic $(\gamma(t), r(t)) \in \Lambda \times [0,\infty)$ such that

- $(\gamma(0), 0) = (x, 0)$ and $(\gamma(s), r(s)) = (y, 0)$;
- $r(t)$ is Morse and has a unique maximum at an interior point $t = t_0$.

Note that the Reeb flow in the unit disk bundle is the natural lift of the geodesic flow. Assume next as above that the metric g on Λ is generic in the sense that the length functional for curves connecting any two Reeb chord endpoints in Λ has only Morse critical points. Concretely, this means that the index form of any geodesic connecting two Reeb chord endpoints is nondegenerate. As in Lemma 81, this allows us to control the Reeb chords of C^τ below a given action for all sufficiently thin handles. More precisely, let ϵ denote the size of the tubular neighborhood of Λ in Y where we attach $T^*(\Lambda \times [0,\infty))$. We introduce the following notion of a geodesic-Reeb chord word. A *geodesic-Reeb chord word* is a word

$$\gamma_1 c_1 \gamma_2 c_2 \cdots c_m \gamma_m,$$

where γ_1 is a geodesic from one of the basepoints p_v to the start point of the Reeb chord c_1, where γ_2 is a geodesic from the endpoint of c_1 to the start point of c_2, etc, until finally γ_m is a geodesic from the endpoint of the Reeb chord c_m to one of the basepoints p_w. We define the action of a geodesic-Reeb chord word to be the sum of actions of its Reeb chords and the energies of its geodesics.

Lemma 88 *For any $A > 0$ there exist $\epsilon_0 > 0$ and $\tau_0 > 0$ such that for any $\epsilon < \epsilon_0$ and any $\tau > \tau_0$ there is a natural one-to-one correspondence between Reeb chords of C^τ of action $< A$, and geodesic-Reeb chord words of Λ of action $< A$.*

Sketch of proof The proof uses the transversality of the Reeb chords and of the geodesics. The basic observation is that the point in the normal fiber of Λ where the Reeb flow hits determines the direction of the geodesic in $\Lambda \times [0,\infty)$. After introducing a concrete smoothing of corners the lemma then follows from the finite-dimensional inverse function theorem. □

To show that Ψ_1 is a quasi-isomorphism we will show that it is represented by a triangular matrix with ones on the diagonal with respect to the action/energy filtration. To this end we will use the Morse-theoretic finite-dimensional model for the chain complex underlying the homology of the based loop space described above. In order to have $[\mathcal{M}^{\mathrm{co}}(\boldsymbol{c})]$ defined as a chain in this model we need to ensure that the paths on the boundary of the holomorphic disk are sufficiently well behaved. We only sketch the construction. On holomorphic disks with unstable domains, we fix gauge using small spheres surrounding a Reeb chord endpoint; compare [31, Section A.2]. As in [31, Section A.1] we use a configuration space for holomorphic curves consisting of maps with two derivatives in L^2. This means that the restriction to the boundary has $\frac{3}{2}$ derivatives in L^2 and in particular the projection to Λ has bounded energy. Since the action of the positive puncture in a holomorphic disk contributing to the differential controls the norm of the solution, it follows that we can use configuration spaces of bounded energy to study the disks in the differential: we approximate the boundary curves uniformly by a piecewise geodesic curve by introducing a uniformly bounded number of subdivision points and straight-line homotopies in small charts.

Conjecture 89 *The chain map* $\Psi_1 \colon \mathrm{CW}^*(C) \to \mathrm{CE}^*(\Lambda)$ *induces an isomorphism on homology.*

Sketch of proof Consider a word of the form

$$\gamma_0 c_1 \gamma_1 c_2 \cdots c_m \gamma_m,$$

where γ_j are geodesics in $\Lambda \times [0,\infty)$ and c_j are Reeb chords. We aim to construct algebraically one disk connecting the Reeb chord a of C, corresponding to this word (see Lemma 88), to the word itself. We use an inductive argument and energy filtration. To start the argument we pick additional fiber disks F_{c^+} and F_{c_-} in $T^*(\Lambda \times [0,\infty))$ at (c^+,ϵ) and (c_-,ϵ) for very small $\epsilon > 0$ near all Reeb chord endpoints c_+ and c_- in Λ. We use the natural counterparts of the correspondence between mixed words of geodesics and Reeb chords before surgery and Reeb chords after, for mixed wrapped Floer cohomologies. For example, there is a straightforward analogue of Lemma 88: Reeb chord generators of $\mathrm{CW}^*(F_{c^+}, C)$ correspond to before-surgery words of the form

$$\gamma_1 c_1 \gamma_2 \cdots c_m \gamma_m,$$

where γ_1 is a geodesic connecting the basepoint of F_{c^+} to the initial point of c_1, and γ_2 connects the endpoint of c_1 to the start point of c_2, etc. To start the argument we note that it is straightforward to construct holomorphic strips corresponding to the short geodesics starting at F_{c^-} followed by the chord c and then the short geodesic to F_{c^+} and to show that they are unique. This corresponds to a generator of $\mathrm{CW}^*(F_{c^-}, F_{c^+})$. Likewise, it is immediate to construct the holomorphic disk connecting a Reeb chord generator of $\mathrm{CW}^*(F_{c^+}, C)$ corresponding to a geodesic, and show that it is unique; compare Theorem 83.

We now use these two to construct algebraically one disk from the Reeb chord generator of $\mathrm{CW}^*(F_{c_-}, C)$ corresponding to the short geodesic, the chord c, and a geodesic connecting the endpoint of c to the basepoint of C. To this end we consider the natural map

$$\Psi^{p_v c^+ c^-} : \mathrm{CW}^*(F_{c^+}, C) \otimes \mathrm{CW}^*(F_{c^-}, F_{c^+}) \to \mathrm{CE}^*(\Lambda);$$

see (47). For the two Reeb chords, a connecting F_{c^-} to F_{c^+} corresponding to the chord c of Λ, and b connecting F_{c^+} to C corresponding to the geodesic, we then have, with p_v denoting the basepoint,

$$d(\Psi^{p_v c^+ c^-}(b,a)) + (\Psi^{p_v c^-}(b)) \cdot (\Psi^{c^+ c^-}(a)) + \Psi^{p_v c^-}(\mathfrak{m}_2(b,a)) \\ + (-1)^{|a|-1} \Psi^{p_v c^+ c^-}(\mathfrak{m}_1(b), a) + \Psi^{p_v c^+ c^-}(b, \mathfrak{m}_1(a)) = 0.$$

Here we know that the terms containing $\mathfrak{m}_1$ and d involve nontrivial holomorphic disks or Morse flows in the finite-dimensional approximation, and hence lower action/energy by an amount bounded below by some $\delta > 0$, which we assume is much larger than $\epsilon > 0$ above. Therefore, if we restrict attention to a small action window, we find

$$(\Psi^{p_v c^-}(b) \cdot \Psi^{c^+ c^-}(a)) + \Psi^{p_v c^-}(\mathfrak{m}_2(b,a)) = 0.$$

Here the first term is simply the Pontryagin product at the common endpoint of the curves, which is homologous to the word $\epsilon' c \gamma$ of the small geodesic, the Reeb chord and then the longer geodesic, by rounding the corner at c^+. It follows that $\mathfrak{m}_2(a,b) = r$, where r is a Reeb chord with action between the sum of the actions of a and b and the action of $\epsilon' c \gamma$, and that $\Psi^{c^- p_v}(r)$ contains this word with coefficient ± 1. Noting that there is only one Reeb chord in the action window studied, we find that the desired coefficient equals ± 1. It is now clear how to continue the induction: at each step we add one more geodesic or Reeb chord to any word. Using already constructed curves and (47) in a small action window, we find that the map Ψ_1 has a triangular action matrix with ± 1 on the diagonal, hence it is a quasi-isomorphism. $\square$

References

[1] **M Abouzaid**, *Morse homology, tropical geometry, and homological mirror symmetry for toric varieties*, Selecta Math. 15 (2009) 189–270 MR Zbl

[2] **M Abouzaid**, *On the wrapped Fukaya category and based loops*, J. Symplectic Geom. 10 (2012) 27–79 MR Zbl

[3] **M Abouzaid**, **P Seidel**, *An open string analogue of Viterbo functoriality*, Geom. Topol. 14 (2010) 627–718 MR Zbl

[4] **J F Adams**, *On the cobar construction*, Proc. Nat. Acad. Sci. U.S.A. 42 (1956) 409–412 MR Zbl

[5] **J F Adams**, **P J Hilton**, *On the chain algebra of a loop space*, Comment. Math. Helv. 30 (1956) 305–330 MR Zbl

[6] **M Aganagic**, **T Ekholm**, **L Ng**, **C Vafa**, *Topological strings, D–model, and knot contact homology*, Adv. Theor. Math. Phys. 18 (2014) 827–956 MR Zbl

[7] **J-F Barraud**, **O Cornea**, *Lagrangian intersections and the Serre spectral sequence*, Ann. of Math. 166 (2007) 657–722 MR Zbl

[8] **A Beilinson**, **V Ginzburg**, **W Soergel**, *Koszul duality patterns in representation theory*, J. Amer. Math. Soc. 9 (1996) 473–527 MR Zbl

[9] **J M Boardman**, *Conditionally convergent spectral sequences*, from "Homotopy invariant algebraic structures" (J-P Meyer, J Morava, W S Wilson, editors), Contemp. Math. 239, Amer. Math. Soc., Providence, RI (1999) 49–84 MR Zbl

[10] **F Bourgeois**, **B Chantraine**, *Bilinearized Legendrian contact homology and the augmentation category*, J. Symplectic Geom. 12 (2014) 553–583 MR Zbl

[11] **F Bourgeois**, **T Ekholm**, **Y Eliashberg**, *Effect of Legendrian surgery*, Geom. Topol. 16 (2012) 301–389 MR Zbl

[12] **F Bourgeois**, **Y Eliashberg**, **H Hofer**, **K Wysocki**, **E Zehnder**, *Compactness results in symplectic field theory*, Geom. Topol. 7 (2003) 799–888 MR Zbl

[13] **E H Brown, Jr**, *Twisted tensor products, I*, Ann. of Math. 69 (1959) 223–246 MR Zbl

[14] **G Carlsson**, **R J Milgram**, *Stable homotopy and iterated loop spaces*, from "Handbook of algebraic topology" (I M James, editor), North-Holland, Amsterdam (1995) 505–583 MR Zbl

[15] **R Casals**, **E Murphy**, *Legendrian fronts for affine varieties*, Duke Math. J. 168 (2019) 225–323 MR Zbl

[16] **K T Chen**, *Iterated path integrals*, Bull. Amer. Math. Soc. 83 (1977) 831–879 MR Zbl

[17] **K Cieliebak**, **T Ekholm**, **J Latschev**, *Compactness for holomorphic curves with switching Lagrangian boundary conditions*, J. Symplectic Geom. 8 (2010) 267–298 MR Zbl

[18] **K Cieliebak**, **Y Eliashberg**, *Flexible Weinstein manifolds*, from "Symplectic, Poisson, and noncommutative geometry" (T Eguchi, Y Eliashberg, editors), Math. Sci. Res. Inst. Publ. 62, Cambridge Univ. Press (2014) 1–42 MR Zbl

[19] **K Cieliebak**, **J Latschev**, *The role of string topology in symplectic field theory*, from "New perspectives and challenges in symplectic field theory" (M Abreu, F Lalonde, L Polterovich, editors), CRM Proc. Lecture Notes 49, Amer. Math. Soc., Providence, RI (2009) 113–146 MR Zbl

[20] **G Civan**, **P Koprowski**, **J Etnyre**, **J M Sabloff**, **A Walker**, *Product structures for Legendrian contact homology*, Math. Proc. Cambridge Philos. Soc. 150 (2011) 291–311 MR Zbl

[21] **S Eilenberg**, **J C Moore**, *Limits and spectral sequences*, Topology 1 (1962) 1–23 MR Zbl

[22] **T Ekholm**, *Morse flow trees and Legendrian contact homology in* 1*–jet spaces*, Geom. Topol. 11 (2007) 1083–1224 MR Zbl

[23] **T Ekholm**, *Rational symplectic field theory over* $\mathbb{Z}_2$ *for exact Lagrangian cobordisms*, J. Eur. Math. Soc. 10 (2008) 641–704 MR Zbl

[24] **T Ekholm**, *A version of rational SFT for exact Lagrangian cobordisms in* 1*–jet spaces*, from "New perspectives and challenges in symplectic field theory" (M Abreu, F Lalonde, L Polterovich, editors), CRM Proc. Lecture Notes 49, Amer. Math. Soc., Providence, RI (2009) 173–199 MR Zbl

[25] **T Ekholm**, *Holomorphic curves for Legendrian surgery*, preprint (2019) arXiv

[26] **T Ekholm**, **J B Etnyre**, **L Ng**, **M G Sullivan**, *Knot contact homology*, Geom. Topol. 17 (2013) 975–1112 MR Zbl

[27] **T Ekholm**, **J Etnyre**, **M Sullivan**, *Orientations in Legendrian contact homology and exact Lagrangian immersions*, Internat. J. Math. 16 (2005) 453–532 MR Zbl

[28] **T Ekholm**, **J Etnyre**, **M Sullivan**, *Legendrian contact homology in* $P\times\mathbb{R}$, Trans. Amer. Math. Soc. 359 (2007) 3301–3335 MR Zbl

[29] **T Ekholm**, **K Honda**, **T Kálmán**, *Legendrian knots and exact Lagrangian cobordisms*, J. Eur. Math. Soc. (JEMS) 18 (2016) 2627–2689 MR Zbl

[30] **T Ekholm**, **A Oancea**, *Symplectic and contact differential graded algebras*, Geom. Topol. 21 (2017) 2161–2230 MR Zbl

[31] **T Ekholm**, **I Smith**, *Exact Lagrangian immersions with one double point revisited*, Math. Ann. 358 (2014) 195–240 MR Zbl

[32] **T Etgü**, **Y Lekili**, *Koszul duality patterns in Floer theory*, Geom. Topol. 21 (2017) 3313–3389 MR Zbl

[33] **Y Félix**, **J C Thomas**, *Extended Adams–Hilton's construction*, Pacific J. Math. 128 (1987) 251–263 MR Zbl

[34] **K Fukaya**, *Application of Floer homology of Langrangian submanifolds to symplectic topology*, from "Morse theoretic methods in nonlinear analysis and in symplectic topology" (P Biran, O Cornea, F Lalonde, editors), NATO Sci. Ser. II Math. Phys. Chem. 217, Springer (2006) 231–276 MR Zbl

[35] **K Fukaya**, **Y-G Oh**, *Zero-loop open strings in the cotangent bundle and Morse homotopy*, Asian J. Math. 1 (1997) 96–180 MR Zbl

[36] **K Fukaya**, **Y-G Oh**, **H Ohta**, **K Ono**, *Lagrangian intersection Floer theory: anomaly and obstruction, I*, AMS/IP Studies in Advanced Mathematics 46, Amer. Math. Soc., Providence, RI (2009) MR Zbl

[37] **P G Goerss**, **J F Jardine**, *Simplicial homotopy theory*, Progr. Math. 174, Birkhäuser, Basel (1999) MR Zbl

[38] **M Gromov**, *Pseudo holomorphic curves in symplectic manifolds*, Invent. Math. 82 (1985) 307–347 MR Zbl

[39] **E Herscovich**, *Hochschild (co)homology and Koszul duality*, preprint (2014) arXiv

[40] **J V S Holstein**, *Morita cohomology*, PhD thesis, University of Cambridge (2014) Available at `https://api.repository.cam.ac.uk/server/api/core/bitstreams/38e69cc9-eb4d-435a-b49b-5a09764ceda5/content`

[41] **J V S Holstein**, *Morita cohomology*, Math. Proc. Cambridge Philos. Soc. 158 (2015) 1–26 MR Zbl

[42] **J D S Jones**, *Cyclic homology and equivariant homology*, Invent. Math. 87 (1987) 403–423 MR Zbl

[43] **M Kalck**, **D Yang**, *Relative singularity categories, I: Auslander resolutions*, Adv. Math. 301 (2016) 973–1021 MR Zbl

[44] **D M Kan**, *Abstract homotopy, IV*, Proc. Nat. Acad. Sci. U.S.A. 42 (1956) 542–544 MR Zbl

[45] **B Keller**, *Deriving DG categories*, Ann. Sci. École Norm. Sup. 27 (1994) 63–102 MR Zbl

[46] **M Kontsevich**, *Symplectic geometry of homological algebra*, preprint (2009) Available at `https://www.ihes.fr/~maxim/TEXTS/Symplectic_AT2009.pdf`

[47] **K Lefèvre-Hasegawa**, *Sur les A_∞–catègories*, PhD thesis, Université Paris 7 (2003) arXiv

[48] **Y Lekili**, **A Polishchuk**, *Arithmetic mirror symmetry for genus* 1 *curves with n marked points*, Selecta Math. 23 (2017) 1851–1907 MR Zbl

[49] **J-L Loday**, **B Vallette**, *Algebraic operads*, Grundl. Math. Wissen. 346, Springer (2012) MR Zbl

[50] **D M Lu**, **J H Palmieri**, **Q S Wu**, **J J Zhang**, *Koszul equivalences in A_∞–algebras*, New York J. Math. 14 (2008) 325–378 MR Zbl

[51] **J McCleary**, *User's guide to spectral sequences*, Mathematics Lecture Series 12, Publish or Perish, Wilmington, DE (1985) MR Zbl

[52] **J Milnor**, *Morse theory*, Annals of Mathematics Studies 51, Princeton Univ. Press (1963) MR Zbl

[53] **D Nadler**, *Wrapped microlocal sheaves on pairs of pants*, preprint (2016) arXiv

[54] **J A Neisendorfer**, *What is loop multiplication anyhow?*, J. Homotopy Relat. Struct. 12 (2017) 659–690 MR Zbl

[55] **L Ng**, **D Rutherford**, **V Shende**, **S Sivek**, **E Zaslow**, *Augmentations are sheaves*, Geom. Topol. 24 (2020) 2149–2286 MR Zbl

[56] **L Positselski**, *Two kinds of derived categories, Koszul duality, and comodule–contramodule correspondence*, Mem. Amer. Math. Soc. 996, Amer. Math. Soc., Providence, RI (2011) MR Zbl

[57] **A Prouté**, *A_∞–structures: modèles minimaux de Baues–Lemaire et Kadeishvili et homologie des fibrations*, Repr. Theory Appl. Categ. (2011) 1–99 MR Zbl

[58] **S Schwede**, **B Shipley**, *Equivalences of monoidal model categories*, Algebr. Geom. Topol. 3 (2003) 287–334 MR Zbl

[59] **P Seidel**, *Fukaya categories and Picard–Lefschetz theory*, Eur. Math. Soc., Zürich (2008) MR Zbl

[60] **J-P Serre**, *Homologie singulière des espaces fibrés: applications*, Ann. of Math. 54 (1951) 425–505 MR Zbl

[61] **Z Sylvan**, *Partially wrapped Floer theory*, talk, IAS workshop "Homological mirror symmetry: methods and structures" (2016) Available at `https://www.ias.edu/video/whms/2016/1110-ZachSylvan`

[62] **Z Sylvan**, *On partially wrapped Fukaya categories*, J. Topol. 12 (2019) 372–441 MR Zbl

[63] **G Tabuada**, *Theorie homotopique des DG–categories*, preprint (2007) arXiv

[64] **C A Weibel**, *An introduction to homological algebra*, Cambridge Studies in Advanced Mathematics 38, Cambridge Univ. Press (1994) MR Zbl

TE: *Department of Mathematics, Uppsala University*
Uppsala, Sweden

Insitut Mittag-Leffler
Djursholm, Sweden

YL: *Department of Mathematics, Imperial College London*
South Kensington, London, United Kingdom

`tobias.ekholm@math.uu.se`, `y.lekili@imperial.ac.uk`

Proposed: Yakov Eliashberg — Received: 3 December 2019
Seconded: Leonid Polterovich, Paul Seidel — Revised: 16 September 2021

Geometry & Topology

msp.org/gt

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

The subscription price for 2023 is US \$740/year for the electronic version, and \$1030/year (+\$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

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

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

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

Geometry & Topology

Volume 27 Issue 6 (pages 2049–2496) 2023