

Duality for Legendrian contact homology

JOSHUA M SABLOFF

The main result of this paper is that, off of a “fundamental class” in degree 1, the linearized Legendrian contact homology obeys a version of Poincaré duality between homology groups in degrees k and $-k$. Not only does the result itself simplify calculations, but its proof also establishes a framework for analyzing cohomology operations on the linearized Legendrian contact homology.

57R17; 53D12, 53D40, 57M25

1 Introduction

1.1 Legendrian contact homology

As in smooth knot theory, a fundamental problem in Legendrian knot theory is to find effective invariants and to understand their structure and meaning. Bennequin [1] initiated the modern study of Legendrian knots by introducing two “classical” invariants: the *Thurston–Bennequin number* (which measures the difference between the framing of a knot coming from the contact planes and the Seifert surface framing) and the *rotation number* (which measures the twisting of the tangent to the knot inside the contact planes with respect to a suitable trivialization). These two invariants suffice to classify Legendrian knots in the standard contact structure on \mathbb{R}^3 when the underlying smooth knot type is the unknot (Eliashberg and Fraser [8]), a torus knot or the figure eight knot (Etnyre and Honda [11]) or a cable link (Ding and Geiges [5]).

The first non-classical invariant of Legendrian knots was *Legendrian contact homology*, a Floer-type theory that comes from geometric ideas of Eliashberg and Hofer [9; 7] and was rendered combinatorially computable by Chekanov [4] for knots in the standard contact \mathbb{R}^3 . The Legendrian contact homology of a Legendrian knot is the homology of a freely-generated differential graded algebra (DGA) (\mathcal{A}, ∂) , which we shall refer to as the *Chekanov–Eliashberg DGA*, itself an invariant up to “stable tame isomorphism.” It is difficult to extract information from the stable tame isomorphism class of the Chekanov–Eliashberg DGA, but Chekanov defined a linearized version which was sufficient to distinguish the first examples of Legendrian knots with the same smooth knot type and classical invariants [4]. The homology of the linearized DGA is usually encoded in a

Poincaré–Chekanov polynomial, in which the coefficient of t^k denotes the dimension of the linearized homology in degree k . The set of Poincaré–Chekanov polynomials taken over all possible linearizations of the DGA is invariant under Legendrian isotopy. Though other more powerful — but less computable — methods of distilling information from the Chekanov–Eliashberg DGA have been developed (Ng [19]), there is still much to be learned about Chekanov’s linearized theory.

Many recent advances in Floer-type theories, and contact homology in particular, have come from importing classical Morse-theoretic ideas. For example, Kálmán’s analysis of nontrivial loops of Legendrian knots was motivated by continuation maps in Morse theory [16]. Another example is the extension of the combinatorial definition of the Chekanov–Eliashberg DGA to Legendrian knots in circle bundles with contact structures transverse to the fiber, which was achieved by transporting Morse–Bott methods into the contact homology picture (Sabloff [21]); see also Bourgeois’ work on Morse–Bott methods for the non-relative version of contact homology [3].

This paper translates ideas of Betz and Cohen [2] and Fukaya and Oh [15] on gradient flow trees into the construction of a Poincaré duality map, a cap product, and a fundamental class for the linearized DGA. The main result is the following duality theorem:

Theorem 1.1 *If h_n denotes the dimension in degree n of the homology of a linearization of the Chekanov–Eliashberg DGA of a Legendrian knot, then:*

$$\begin{aligned} h_n &= h_{-n} & |n| > 1, \\ h_1 &= h_{-1} + 1. \end{aligned}$$

Said another way, the Poincaré–Chekanov polynomial $P(t)$ of a linearization satisfies:

$$P(t) = P(t^{-1}) + (t - t^{-1}).$$

The extra class in degree 1 will turn out to be the fundamental class, which is closely related to the fundamental class of the circle. This theorem can greatly simplify calculations (see Melvin and Shrestha [17], for example), but of greater interest is its proof, which introduces an algebraic and geometric framework for understanding the “algebraic topology” of the linearized DGA.

1.2 Morse-theoretic motivation

While the proof of Theorem 1.1 will be combinatorial in nature, it is motivated by geometric ideas from Morse theory. Classically, the Poincaré duality map caps a

cohomology class with the fundamental class. Using gradient flow trees as in [2; 15], the cap product and fundamental class can be reinterpreted in the setting of Morse–Witten theory. Recall that the Morse–Witten complex $C_*(M, f)$ of a manifold M with Morse function f is generated by the critical points of f , while the differential comes from counting rigid negative gradient flow lines between critical points; see Schwarz [23]. The cochain complex $C^*(M, f)$ is generated by the same critical points, but the codifferential counts *upward* gradient flows.

In this language, the fundamental class is represented by a sum of the maxima of f . More interestingly, the cap product between a homology class B and a cohomology class γ is computed by counting certain rigid “Y”-shaped gradient flow trees. Choose three Morse functions f_1, f_2 , and f_3 , and choose representatives b for B in $C_*(M, f_1)$ and c for γ in $C^*(M, f_2)$. The “Y”-shaped tree follows the negative gradient flow of f_1 out of b and splits at the vertex into negative gradient flows for f_2 and f_3 . These two flowlines end at c and at a critical point of f_3 that represents $\gamma \cap B$ in $C_*(M, f_3)$, respectively; see Figure 1. More generally, gradient flow trees give rise to cohomology operations, with homology inputs and cohomology outputs at the outward-flowing, or positive, ends of a negative gradient flow line, and homology outputs and cohomology inputs at the negative ends. Of particular interest in this paper is the Poincaré duality isomorphism, which, as Betz and Cohen note, comes from a tree with two positive ends and one vertex, where one of the positive ends is a homology input and the other a cohomology output.




Figure 1: A gradient flow tree that gives the cap product $[d] = \gamma \cap B$ between $B = [b]$ and $\gamma = [c]$. Notice that there is one homological input at b , one cohomological input at c , and a homological output at d .

Legendrian contact homology fits into this picture as the Morse–Witten–Floer theory of the action functional on the space of paths that begin and end on a Legendrian knot K . In this setting, Reeb chords are the critical points that generate the Morse–Witten–Floer complex, while holomorphic disks in the symplectization $(\mathbb{R}^3 \times \mathbb{R}, d(e^t \alpha))$ take the

place of the gradient flow lines that define the differential ∂ . In order to ensure that $\partial^2 = 0$, the holomorphic disks must be allowed to have multiple negative ends. The end result is a non-commutative DGA generated by the Reeb chords. Chekanov's combinatorial formulation of the DGA arises from the correspondence between Reeb chords and crossings in the xy projection of a Legendrian knot — the Reeb direction in the standard contact \mathbb{R}^3 is the z direction — and between holomorphic disks in the symplectization and immersed disks in the xy projection; see Section 2 for the combinatorics and Etnyre, Ng and Sabloff [12] for more details of the translation between geometry and combinatorics.

The Morse-theoretic basis for Legendrian contact homology would lead one to expect that the longest Reeb chords should give a fundamental class in Legendrian contact homology, and the cap product should come from disks with one positive and two negative ends. The analogy between Legendrian contact homology and Morse theory breaks down in the case of the Poincaré duality map: disks with multiple positive ends do not appear in the contact homology theory. To access disks with two positive ends, it is necessary to expand the algebraic framework of the Chekanov–Eliashberg DGA. The natural expansion is a relative version of Eliashberg, Givental, and Hofer's symplectic field theory (SFT) [9]. The definition of a “Legendrian SFT,” however, runs into some subtle issues regarding compactness of the moduli space of curves used to define the differential and has yet to be rigorously defined. Instead, the expansion can be achieved by an appropriate interpretation of the generators and differentials of the DGA of a link formed by several vertical translates of the original knot.

The remainder of the paper is organized as follows: background notions — including descriptions of Legendrian knot diagrams, the Chekanov–Eliashberg DGA, and the linearization procedure — are described in Section 2. Next, Section 3 describes the structure of the “expanded DGA.” Section 4 describes the relationship between the original linearized DGA and the linearized expanded DGA, resulting in the definition of the Poincaré duality map and the proof of Theorem 1.1. Finally, Section 5 explores the “algebraic topology” of the linearized DGA, defining both the fundamental class and the cap product and showing that a “capping with the fundamental class” construction inverts the duality isomorphism.

Acknowledgments This paper greatly benefited from discussions with John Etnyre, Paul Melvin, Lenny Ng, and Lisa Traynor. The referee's insightful comments spurred me to clean up some ideas surrounding the fundamental class, which, in turn, led to a greatly improved proof of duality. The research behind this paper started out in a rather different direction, and Yasha Eliashberg and Frédéric Bourgeois were instrumental in helping me understand the ideas of symplectic field theory.

2 Background notions

2.1 Legendrian knots and their diagrams

This section briefly reviews the basic notions of Legendrian knot theory; for a more comprehensive introduction, see Etnyre [10] or Sabloff [20].

The *standard contact structure* on \mathbb{R}^3 is the completely non-integrable 2-plane field given by the kernel of $\alpha = dz - y dx$. A *Legendrian knot* is a smooth embedding $K: S^1 \rightarrow \mathbb{R}^3$ that is everywhere tangent to the contact planes. That is, the embedding satisfies $\alpha(K') = 0$. An ambient isotopy of K through other Legendrian knots is a *Legendrian isotopy*. Legendrian knots are plentiful; for example, any smooth knot can be continuously approximated by a Legendrian knot.

There are two useful projections of Legendrian knots. The first is the *front projection* π_f to the xz plane. In the front projection, the y coordinate of a knot may be recovered from the slope of its projection. As a result, the projection can have no self or vertical tangencies; it has semi-cubical cusps instead. Further, the crossing information is completely determined: the strand with lesser slope will always pass in front of the strand with greater slope. Any circle in the xz plane that has no vertical tangencies, has no self-tangencies, and that is immersed except at finitely many cusps lifts to a Legendrian knot.

Though the front projection is easier to use, it is more natural to define Legendrian contact homology using the *Lagrangian projection* π_l to the xy plane. Unlike the front projection, not every immersion into the xy plane is the Lagrangian projection of a Legendrian knot K : a system of inequalities involving the areas of the connected components $\mathbb{R}^2 \setminus \pi_l(K)$ must be satisfied (see [4]). It is simpler (and sufficient) to work with *Lagrangian diagrams* of K , ie, immersions D of the circle into the xy plane, together with crossing information, for which there is an orientation-preserving diffeomorphism of the plane carrying D to $\pi_l(K)$. See Figure 2(b), for example.

Ng's *resolution* procedure (see [19]) gives a canonical translation from a front projection to a Lagrangian diagram. Combinatorially, there are three steps:

- (1) smooth the left cusps;
- (2) replace the right cusps with loops (see the right side of the Lagrangian projection in Figure 2); and
- (3) resolve the crossings so that the overcrossing is the one with lesser slope.

A key feature of the resolution procedure is that the heights of the crossings in the Lagrangian diagram strictly increase from left to right, with the jumps in height between

crossings as large as desired. In particular, the crossings in a resolved Lagrangian diagram that come from the rightmost cusps have the greatest height among all crossings.




Figure 2: The (a) front and (b) Lagrangian diagrams of a Legendrian 5_2 knot. The meaning of the numbers will become clear in Section 2.2.

2.2 The Chekanov–Eliashberg DGA

The DGA was originally defined by Chekanov in [4] for Lagrangian diagrams; see also [12; 19; 20]. This section contains a brief review of the definition.

Let K be an oriented Legendrian knot in the standard contact \mathbb{R}^3 with a generic Lagrangian diagram D . Label the crossings with $\{q_1, \dots, q_n\}$. Let A be the graded vector space over $\mathbb{Z}/2$ generated by q_1, \dots, q_n , and let \mathcal{A} be the graded free unital tensor algebra TA .

The grading is determined by the assignment of a *capping path* to each crossing. A capping path is one of the two immersed paths that start at the overcrossing of q_i , trace out a portion of D , and end when D first returns to q_i , necessarily at an undercrossing. Assume, without loss of generality, that the strands of D at each crossing are orthogonal. The *grading* of q_i is defined to be:

$$|q_i| \equiv 2r(\gamma_i) - \frac{1}{2} \pmod{2r(K)}.$$

Extend the grading to all words in \mathcal{A} by letting the grading of a word be the sum of the gradings of its constituent generators.

The next step is to define a differential on \mathcal{A} by counting immersions of the 2-disk B^2 with an ordered set of marked points on its boundary. The number of marked points must be at least one, but otherwise can vary from disk to disk; if there are $k + 1$ of marked points, label them $\{z_0, \dots, z_k\}$. Further, label the corners of D as in Figure 3(a). The immersions of interest are the following:

Definition 2.1 Given an ordered set of generators $q_i, q_{j_1}, \dots, q_{j_k}$, define the set $\Delta(q_i; q_{j_1}, \dots, q_{j_k})$ to consist of orientation-preserving immersions

$$f: B^2 \rightarrow \mathbb{R}^2$$

up to smooth reparametrization that map ∂B^2 to the image of D subject to the following conditions:

- (1) The restriction of f to the boundary is an immersion away from the marked points z_i .
- (2) The map f has the property that $f(z_0) = q_i$ and $f(z_l) = q_{j_l}$, such that the points $q_i, q_{j_1}, \dots, q_{j_k}$ are encountered in counter-clockwise order along the boundary.
- (3) In a neighborhood of the points $q_i, q_{j_1}, \dots, q_{j_k}$, the image of the disk under f has the form indicated in Figure 3(b) and (c).




Figure 3: (a) A labeling of the quadrants surrounding a crossing; (b) The image of $f \in \Delta$ near the crossing q_i and (c) the crossings q_{j_l} . The other positive (resp. negative) corner in (b) (resp. (c)) is also possible.

Note the analogy between the positive and negative corners in Figure 3 and positive and negative ends of a gradient flow tree. With this definition in hand, define the differential as follows:

Definition 2.2
$$\partial(q_i) = \sum_{\text{Words } q_{j_1} \dots q_{j_k}} \sum_{f \in \Delta(q_i; q_{j_1}, \dots, q_{j_k})} q_{j_1} \dots q_{j_k}$$

Extend ∂ to all of \mathcal{A} via linearity and the Leibniz rule.

The fact that the sum in the definition of ∂ is finite comes from the following lemma, which is essentially an application of Stokes' Theorem:

Lemma 2.3 *Let $h(x)$ be the height of the crossing x . If there is an immersed disk that satisfies all but the third condition of Definition 2.1 and has positive corners at x_1, \dots, x_k and negative corners at y_1, \dots, y_l , then*

$$\sum_i h(x_i) - \sum_i h(y_i) > 0.$$

The central results in the theory are:

Theorem 2.4 (Chekanov [4])

- (1) The differential ∂ has degree -1 .
- (2) The differential satisfies $\partial^2 = 0$.
- (3) The “stable tame isomorphism class” (and hence the homology) of the DGA is invariant under Legendrian isotopy.

Here, a *stabilization* of \mathcal{A} is a DGA $S(\mathcal{A})$ with two new generators a and b such that $\partial a = b$ and $\partial b = 0$. Two DGAs \mathcal{A} and \mathcal{A}' are *stable tame isomorphic* if there exist (possibly multiple) stabilizations of each that are tame isomorphic. See [4] for a description of the technical condition “tame;” this condition allows for the invariance of *based* algebras and is essential for the proof of the invariance of the linearized contact homology described in the next section, though its exact technical definition is not necessary for the arguments in this paper.

Example 2.5 The knot in Figure 2 has nine crossings. The generators q_1 and q_2 have grading 0; q_3, q_4, q_8 , and q_9 have grading 1; q_5 and q_7 have grading 2; and q_6 has grading -2 . The differential is given by:

$$\begin{aligned}
 \partial q_1 &= \partial q_2 = \partial q_6 = \partial q_7 = 0 \\
 \partial q_3 &= 1 + q_1 q_2 \\
 \partial q_4 &= 1 + q_2 q_1 \\
 \partial q_5 &= q_3 q_1 + q_1 q_4 \\
 \partial q_8 &= 1 + q_1 + q_1 q_6 q_7 \\
 \partial q_9 &= 1 + q_1 + q_7 q_6 q_1.
 \end{aligned}
 \tag{2-1}$$

2.3 Linearized contact homology

Chekanov introduced an important computable invariant of the stable tame isomorphism class of the Chekanov–Eliashberg DGA called the *linearized contact homology*. The differential ∂ on \mathcal{A} may be split into a sum of differentials $\partial = \sum_{l=0}^{\infty} \partial_l$, where the image of ∂_0 lies in the ground ring $\mathbb{Z}/2$ and ∂_l maps a generator of \mathcal{A} into $\mathcal{A}^{\otimes l}$. If $\partial_0 = 0$, then the equation $\partial^2 = 0$ implies $\partial_1^2 = 0$. In this case, it follows that $(\mathcal{A}, \partial_1)$ is an honest chain complex whose homology $H_*(\mathcal{A}, \partial_1)$ may be easily computed.

Rarely does the DGA have the property that $\partial_0 = 0$. Suppose, however, that there exists an algebra map $\varepsilon: \mathcal{A} \rightarrow \mathbb{Z}/2$ that satisfies:

- (1) $\varepsilon \circ \partial = 0$, and
- (2) $\varepsilon(q_i) = 0$ if $|q_i| \neq 0$.

Such a map is called an *augmentation*; if $\varepsilon(q_i) = 1$, say that q_i is *augmented*. Augmentations are not uncommon, but they do not always exist; it turns out that their existence is equivalent to the existence of a “ruling” of a front diagram of a Legendrian knot (Fuchs [13], Fuchs and Ishkhanov [14] and Sabloff [22]).

Given an augmentation ε , define an automorphism $\Phi^\varepsilon: \mathcal{A} \rightarrow \mathcal{A}$ by:

$$\Phi^\varepsilon(q_i) = q_i + \varepsilon(q_i).$$

Let ∂^ε be the differential induced by Φ^ε :

$$\partial^\varepsilon = \Phi^\varepsilon \partial (\Phi^\varepsilon)^{-1}.$$

It is straightforward to verify that the DGA $(\mathcal{A}, \partial^\varepsilon)$ satisfies $\partial_0^\varepsilon = 0$.

Remark 2.6 Once the augmentation is known, it is possible to read off the linearized differential directly from the Lagrangian diagram. An immersed disk that contributes q_j to the linearized differential of q_i has a positive corner at q_i , a negative corner at q_j , and possibly other negative corners at augmented crossings. For example, if $\partial q_1 = q_2 q_3$ and only q_2 is augmented, then $\partial_1^\varepsilon q_1 = q_3$. If both q_2 and q_3 are augmented, then $\partial_1^\varepsilon q_1 = q_2 + q_3$.

For each augmentation, there is a *Poincaré–Chekanov Polynomial*:

$$P_\varepsilon(t) = \sum_{n=-\infty}^{\infty} \dim H_n(\mathcal{A}, \partial_1^\varepsilon) \cdot t^n.$$

Chekanov proved in [4] that the set $\mathcal{P} = \{P_\varepsilon(t)\}_{\varepsilon \in \mathcal{E}}$ is invariant under Legendrian isotopy, where \mathcal{E} is the set of all possible augmentations of (\mathcal{A}, ∂) . The proof of invariance splits into two parts: the first is the somewhat trivial verification that \mathcal{P} does not change under stabilization. The second part, which involves the analysis of a tame isomorphism $\Psi: (\mathcal{A}, \partial) \rightarrow (\mathcal{A}', \partial')$, will become important in Section 4.3, so a rough sketch of the proof is in order. Let an augmentation ε' on $(\mathcal{A}', \partial')$ be given. The goal is to find a corresponding augmentation ε on (\mathcal{A}, ∂) and an isomorphism between the chain complexes $(\mathcal{A}, \partial_1^\varepsilon)$ and $(\mathcal{A}', (\partial')_1^{\varepsilon'})$. Chekanov proved that any tame isomorphism, such as $\Phi^{\varepsilon'} \circ \Psi$, can be factored into a composition $\bar{\Psi} \circ \Psi^0$, where the image of a generator q under $\bar{\Psi}$ has no constant terms and Ψ^0 is of the form $\Psi^0(q_i) = q_i + c_i$ for some $c \in \mathbb{Z}/2$. It is possible to prove that the constants c_i come from an augmentation ε of (\mathcal{A}, ∂) , and that $\bar{\Psi}$ conjugates ∂^ε and $(\partial')^{\varepsilon'}$. Since $\bar{\Psi}$ is

nondecreasing in the tensor powers of \mathcal{A} , it restricts to a chain isomorphism between the linearized complexes $(A, \partial_1^\varepsilon)$ and $(A', (\partial')_1^{\varepsilon'})$. This argument can clearly be run in the other direction, thus producing a bijection between \mathcal{P} and \mathcal{P}' . For more details, see [4] or [20].

Remark 2.7 The set \mathcal{P} of Poincaré–Chekanov polynomials is not necessarily a one-element set: Melvin and Shrestha [17] found examples of Legendrian knots with arbitrarily large sets.

Example 2.8 Referring back to Example 2.5, it is not hard to check that there is a unique augmentation of (\mathcal{A}, ∂) in which both of the degree 0 generators q_1 and q_2 are sent to 1 by the augmentation. The resulting linearized differential is:

$$(2-2) \quad \begin{aligned} \partial_1^\varepsilon q_1 &= \partial_1^\varepsilon q_2 = \partial_1^\varepsilon q_6 = \partial_1^\varepsilon q_7 = 0 \\ \partial_1^\varepsilon q_3 &= q_1 + q_2 \\ \partial_1^\varepsilon q_4 &= q_1 + q_2 \\ \partial_1^\varepsilon q_5 &= q_3 + q_4 \\ \partial_1^\varepsilon q_8 &= q_1 \\ \partial_1^\varepsilon q_9 &= q_1. \end{aligned}$$

An easy computation shows that the linearized homology is generated by $[q_6]$, $[q_7]$, and $[q_8 + q_9]$, so the set of Poincaré–Chekanov polynomials is

$$\{t^{-2} + t + t^2\}.$$

Notice that the Poincaré polynomial in the example is symmetric about degree 0, with the exception of a class in $H_1(A, \partial_1^\varepsilon)$. Theorem 1.1 asserts that this symmetry holds in general.

3 The expanded algebra

As suggested in the introduction, the proof of duality for the linearized DGA requires disks with two positive corners rather than the single positive corner of the disks in Definition 2.1. One method to access these disks comes from the following construction: given a Legendrian knot K , shrink it so that its Reeb chords have length at most $\delta > 0$. Let $f: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be the vertical translation by 1, and let K^n be the link with components $\{K, f(K), \dots, f^{n-1}(K)\}$. The Reeb chords for this link are far from isolated, but after a perturbation using a Morse function on the knot K , the chords fall into three families:

q chords Chords that start on translates of the bottom strand of a chord for K and end on translates of the top strand (this includes the original chords of K),

p chords Chords that start on translates of the top strand of a chord for K and end on translates of the bottom strand, and

c and d chords Chords corresponding to critical points of the Morse function.

The perturbation procedure and the families of chords will be described in more detail in Sections 3.1 and 3.2. After identifying chords that differ by a translation, these chords will generate the n -fold expanded algebra $\widehat{\mathcal{A}}^n$ of the knot. For the proof of duality, only $\widehat{\mathcal{A}}^2$ will be necessary. The higher-order algebras can be used to define higher cohomology operations; for example, $\widehat{\mathcal{A}}^3$ is required to define the inverse of the duality map.

The p chords have the following useful property: a disk with a negative corner at a p chord locally projects to the same quadrant as a disk that has a positive corner at a q chord. In other words, an output of p is the same as an input of q . This indicates that, at least at the linear level, p chords could be used to represent generators for the cochain complex of the original knot. It follows that a disk in the Lagrangian diagram of K^n with, say, a positive corner at a q chord and a negative corner at a p chord corresponds to a disk with two positive corners in the Lagrangian diagram of the original knot K ; see Figure 6. This property will be exploited to define the duality map in Section 4.

Remark 3.1 This construction was extracted from embedding a “small” Legendrian knot in the contact circle bundle $\mathbb{R}^2 \times S^1$, which is equivalent to taking infinitely many vertical translates of the original knot in \mathbb{R}^3 . The full “Morse–Bott” theory for Legendrian knots in circle bundles, which inspired the language and notation of this section, was fully worked out by Sabloff in [21]. See also the last section of Ekholm, Etnyre and Sullivan [6] for a similar construction.

3.1 Generators of the algebra

In order to define the expanded algebra based on the Chekanov–Eliashberg DGA of K^n , a more precise description of the perturbation is necessary. Let g be a function on the xy plane, supported in a neighborhood of the Lagrangian diagram $\pi_1(K)$, and let \tilde{g} be a z -invariant lift to \mathbb{R}^3 . Choose g so that it has no critical points in a neighborhood of $\pi_1(K)$ and so that $\tilde{g}|_K$ is a Morse function whose critical points do not lie above the crossings of $\pi_1(K)$. The perturbation of K^n is a sequence of small, but progressively larger, shifts of $f^j(K)$ in the direction $i\nabla g$, where i is the usual complex structure

on \mathbb{R}^2 . Figure 4 shows that effects of this perturbation near the critical points and crossings.




Figure 4: The effect of the perturbation on the crossings of K and $f^k(K)$ near (a) a maximum, (b) a minimum, and (c) a crossing of K . For this particular choice of g , the critical points of $\tilde{g}|_K$ to the left of the crossing are maxima and are minima to the right. A different choice of g would give a different configuration of signs and of the position of the surrounding maxima and minima.

As indicated above, the perturbation introduces a host of new generators. For example, Figure 4 shows that every critical point of $\tilde{g}|_K$ gives rise to a new chord between any two components of K^n . Label the maxima of $\tilde{g}|_K$ by $\{c_1, \dots, c_r\}$ and the minima by $\{d_1, \dots, d_r\}$. At a crossing of K , there are four chords for every pair of components: the original chord (q^0 in the figure), its vertical translate (\tilde{q}^0), a chord from the bottom strand of K to the top strand of $f^k(K)$ (q^k), and a chord from the top strand of K to the bottom strand of $f^k(K)$ (p^k). In fact, each critical point and crossing of K gives rise to a half-lattice or lattice of crossings of K^n . In the expanded algebra, unlike the Chekanov–Eliashberg DGA of K^n , the original chord q^0 and its translate \tilde{q}^0 will be identified. In general, *the expanded algebra identifies any two chords that, modulo perturbation, are vertical translates of one another*. This is the only difference between the expanded algebra and the full Chekanov–Eliashberg DGA of K^n .

More formally, the algebra is defined as follows:

Definition 3.2 The n -fold expanded algebra $\widehat{\mathcal{A}}^n$ is the graded free unital tensor algebra over $\mathbb{Z}/2$ generated by the following:

- To the i^{th} double point, associate generators $\{q_i^k, p_i^k\}_{k=0,1,\dots,n-1}$. The generator q_i^k (respectively, p_i^k) corresponds to the chord that starts on the bottom

(resp. top) strand of K and ends on the top (resp. bottom) strand of $f^k(K)$.¹ The number k is the *level* of the generator.

- To each point c_i , associate generators $\{c_i^k\}_{k=1,2,\dots,n-1}$, and to each point d_i , associate generators $\{d_i^k\}_{k=1,2,\dots,n-1}$. The generators c_i^k (resp. d_i^k) represent the Reeb chords that start at a maximum (resp. minimum) of $\tilde{g}|_K$ and end at the corresponding point on $f^k(K)$.

Note that, by construction, we have:

$$(3-1) \quad \begin{aligned} h(q_i^k) &> k, & h(c_i^k) &\approx k, \\ h(p_i^k) &< k, & h(d_i^k) &\approx k. \end{aligned}$$

The gradings of the new generators are easy to define. Let q_i^0 inherit its grading from the DGA (\mathcal{A}, ∂) ; the others are graded as follows:

$$(3-2) \quad \begin{aligned} |q_i^k| &= |q_i^0|, & |c_i^k| &= 0, \\ |p_i^k| &= -1 - |q_i^0|, & |d_i^k| &= -1. \end{aligned}$$

3.2 The differential

The definition of the differential $\widehat{\partial}$ on the n -fold expanded algebra is straightforward: it is simply the contact homology differential for the link K^n , up to identifications of chords that differ by a vertical translation. To show that the differential is well-defined, it is useful to introduce Mishachev’s relative homotopy splitting of a link DGA [18], as interpreted by Ng [19]. The motivation for this language is the following simple observation: while traversing counter-clockwise the boundary of an immersed disk involved in the definition of ∂ , the upper strand of the first negative corner and the lower strand of the last negative corner agree with the upper and lower strands, respectively, of the positive corner. On the other hand, the lower strand of the first of two consecutive negative corners coincides with the upper strand of the second.

More formally, label the components of K^n from bottom to top by $\{0, \dots, n-1\}$. For $j \neq k$, let Γ_{jk} be the submodule of $\widehat{\mathcal{A}}$ generated by words of the form $x_1 \cdots x_m$ so that:

- The upper strand of x_1 is j ,
- The lower strand of x_m is k , and
- The lower strand of x_i coincides with the upper strand of x_{i+1} for $1 \leq i \leq m-1$.

¹The generators p^k would be labeled as p^{k-1} in [21]; the slight change in notation greatly improves readability in the current context.

If $j = k$, include an indeterminate e_j as a generator of Γ_{jj} as well. For K^n , the component j lies above the component k if and only if $j > k$, so it is easy to check that Γ_{jk} is empty for $j < k$ and that for $j \geq k$, the upper (resp. lower) strand of x_i in a generator $x_1 \cdots x_m$ is at most j (resp. at least k).

The key fact about this construction, whose proof is essentially outlined above, is that the module Γ_{jk} is closed under the action of the differential ∂ :

Lemma 3.3 [19, Lemma 2.19] *If the upper strand of a chord x is j and its lower strand is k , then $\partial x \in \Gamma_{jk}$, where a 1 in the differential of $x \in \Gamma_{jj}$ is replaced by e_j .*

More is true: the differential module (Γ_{jk}, ∂) is invariant under Legendrian isotopy up to stable tame isomorphism. Further, if the set of augmentations is restricted to ones that vanish on chords between components, then the set of linearized homologies of (Γ_{jk}, ∂) is also invariant.

To prove that $\widehat{\partial}$ is well-defined, suppose x is an arbitrary generator for the Chekanov–Eliashberg DGA of K^n and let $f^k(x)$ be a vertical translate of x (up to perturbation). Lemma 3.3 and the observation before it regarding the upper and lower strands of the x_i that make up the generators of Γ_{jk} together imply that any disk contributing to ∂x only involves the link consisting of components of K^n that lie between the components of the upper and lower strands of x (inclusive). Similarly, any disk contributing to $\partial f^k(x)$ only involves the link lying between the the components of the upper and lower strands of $f^k(x)$ (inclusive). The Lagrangian diagrams of these links are the same (up to a small shift), so the differentials of x and $f^k(x)$ are the same up to identification of vertically-translated generators. In particular, the differential $\widehat{\partial}$ on $\widehat{\mathcal{A}}$ is well-defined.

Before moving on, a few words are in order about the form of the differential. Following the language of [18],² define a smooth “stick-together” map $s: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ that retracts the Lagrangian diagram of the perturbed link K^n onto the Lagrangian diagram of K . The disks that define the differential split into two classes: the first class consists of *thick disks* that the stick-together maps sends to immersed disks in the original diagram. As can be seen in Figure 4, a disk in K^n with a positive corner at a q crossing will have a positive corner in the original diagram, while a disk in K^n with a positive p corner will have a negative corner in the original diagram; a parallel correspondence holds at negative corners.

²Despite the similarities in language and diagrammatics, the N -copy of [18] and the perturbed link K^n are quite different objects. In particular, K^n is unlinked.

The second class of disks consists of *thin disks* whose images under the stick-together map lie in the original diagram of K . A thin disk has one of the following forms, as can be seen from examining Figure 4:

- (1) A disk that flows down from a positive corner at a c crossing to a negative corner at one of the two adjacent d crossings,
- (2) A disk that flows down from a positive corner at a c crossing to negative p and q corners at a crossing of $\pi_1(K)$ that lies between the c and an adjacent d ,
- (3) A disk that flows down from a positive corner at a q (resp. p) crossing and a negative corner at another q (resp. p) crossing in the same lattice to a negative corner at the next d crossing,³
- (4) A disk that flows down from a positive corner at a q (resp. p) crossing and a negative corner at another q (resp. p) crossing in the same lattice to negative p and q corners in another lattice (of course, there can be no intervening c or d crossings), or
- (5) A disk that lies entirely inside the lattice of crossings created by the perturbation.

These disks have a close relationship with the negative gradient flow of $\tilde{g}|_K$: disks of type (1) correspond to flowlines from maximum to minimum, while disks of types (2–4) coincide with “partial flowlines.” In each case, the positive corner lies at the end of the flowline or partial flowline with the largest value of g and flows down to a negative corner (possibly with other negative corners at each end).

It is easy to see that the differential of a c generator comes entirely from thin disks of types (1) and (2), while the differential of a d generator comes entirely from thick disks and thin disks of type (5). There are exactly two thin disks of type (3) that contribute to the differential of each q and each p generator, and possibly more of types (4) or (5).

Example 3.4 Consider the 2-fold expanded algebra of the knot in Figure 2 given by the perturbation in Figure 5 (the perturbation comes from shifting the entire diagram down). Consider the differential of q_6^1 . The thick disks contribute the following terms:

$$(3-3) \quad p_5^1 + p_7^1.$$

Thin disks of types (3)–(5) contribute quite a few more:

$$(3-4) \quad q_6^0(q_7^0p_7^1 + q_9^0p_9^1 + d_2) + (p_7^1q_7^0 + p_8^1q_8^0 + d_1)q_6^0 + q_6^0p_6^1q_6^0.$$

³More than the two copies of K depicted in Figure 4 are necessary to see the disk with p corners.




Figure 5: A perturbation of the knot in Figure 2 that gives a 2-fold expanded algebra.

To connect this example to the upcoming proof of duality, recall the observation at the beginning of this section that a disk in the Lagrangian diagram of K^n with, say, a positive corner at a q chord and a negative corner at a p chord corresponds to a disk with two positive corners in the Lagrangian diagram of the original knot K . Figure 6 shows this explicitly in the example above: a disk with two positive corners in the Lagrangian diagram of the original knot K yields a p_5^1 term in $\widehat{\partial}q_6^1$ in the expanded algebra of K^2 .




Figure 6: A disk with two positive corners in the Lagrangian diagram of the original knot K corresponds to a disk with a positive corner at a q chord and a negative corner at a p chord in a Lagrangian diagram of K^2 .

4 The proof of duality

When linearized, the expanded algebra carries a lot of structure. In this section, this structure will be uncovered and exploited to prove Theorem 1.1.

4.1 Structure of the linearized expanded algebra

This section examines the following decomposition of the linear pieces of $\widehat{\mathcal{A}}^n$: let Q^m be the vector space generated by the level m chords $\{q_i^m, c_i^m, d_i^m\}$ and let P^m

be the vector space generated by $\{p_i^m\}$. The goal of this section is to understand the relationships between the Q and P vector spaces once $\widehat{\mathcal{A}}$ has been linearized; the primary tools will be Lemmas 2.3 and 3.3.

4.1.1 Extending the Augmentation The first step will be to extend an augmentation for \mathcal{A} to all of $\widehat{\mathcal{A}}^n$. Given an augmentation ε of \mathcal{A} , there is a trivial extension to a map $\widehat{\varepsilon}$ on $\widehat{\mathcal{A}}^n$: simply set $\widehat{\varepsilon}$ to be zero on all generators not in Q^0 . This map turns out to be an augmentation for the expanded algebra:

Proposition 4.1 $\widehat{\varepsilon} \circ \widehat{\partial} = 0$.

Proof The first step is to show that every term in the differential of a generator not in Q^0 contains at least one generator not in Q^0 . This is simple: since a generator x^m , $m \geq 1$, corresponds to a crossing between two different components of K^n , Lemma 3.3 implies that there must be at least one generator in each term of $\widehat{\partial}x^m$ whose upper and lower strands lie on different components of K^n . This implies that $\widehat{\varepsilon} \circ \widehat{\partial}x^m = 0$ for $m \geq 1$ since $\widehat{\varepsilon}$ vanishes on generators of level $m \geq 1$.

This leaves Q^0 . When restricted to Q^0 , it turns out that the expanded differential $\widehat{\partial}$ agrees with the original differential ∂ on K . Let $q \in Q^0$. Since the Reeb chords that generate Q^0 are shorter than any other chords, Lemma 2.3 implies that $\widehat{\partial}q$ contains only other generators from Q^0 . In particular, the boundary of any disk contributing to $\widehat{\partial}q$ lies on a single component of K^n , and hence contributes to the original differential of q on K . Thus, on Q^0 , $\widehat{\partial} = \partial$ and the proposition follows from the definition of $\widehat{\varepsilon}$. \square

4.1.2 Splitting the linearized differential The linearized differential $\widehat{\partial}_1^{\widehat{\varepsilon}}$ has a rich structure that will be exploited in the proof of duality. For convenience, the superscript $\widehat{\varepsilon}$ will be dropped from the notation henceforth, ie, all differentials will be assumed to be augmented. The first important property of the linearized differential is that it preserves level:

Lemma 4.2 *The images of the linearized differential obey $\widehat{\partial}_1(Q^0) \subset Q^0$ and $\widehat{\partial}_1(Q^m \oplus P^m) \subset Q^m \oplus P^m$ for all $m \geq 1$. Further, $\widehat{\partial}_1(P^m) \subset P^m$.*

Proof A summand $y_1 \cdots y_k$ of $\widehat{\partial}x^m$ will contribute y_i to the linearized differential if $\widehat{\varepsilon}(y_j) = 1$ for all $j \neq i$. By the definition of $\widehat{\varepsilon}$, all such y_j must be in Q^0 . Hence, the only place that the boundary of the disk giving rise to the summand $y_1 \cdots y_k$ can switch from the upper strand of x^m to the lower strand is at y_i . Thus, y_i must have

the same upper and lower strands as x^m , so the image of the linearized differential of x^m must also have level m .

The second statement follows from the first, Lemma 2.3, and Equation (3–1), which show that a term in $\widehat{\partial}_1 p^m$ must have level m , have height less than p^m , and hence must be an element of P^m . \square

On any given level $m \geq 1$, the linearized differential splits into three pieces by restricting domains and ranges:

- (1) $\widehat{\partial}_Q: Q^m \rightarrow Q^m$,
- (2) $\widehat{\partial}_P: P^m \rightarrow P^m$, and
- (3) $\eta: Q^m \rightarrow P^m$. This is the map that appears schematically in Figure 6.

That there is no piece of the linearized differential with domain P^m and range Q^m is a consequence of the previous lemma.

Example 4.3 As noted in Example 2.8, the only generators augmented in the case of the knot in Figure 2 are q_1^0 and q_2^0 . Thus, for the generators q_6^1 in Figure 5, the thin disks computed in Example 3.4 contribute nothing to the linearized differential. Thus, it is straightforward to read off that $\widehat{\partial}_Q(q_6^1) = 0$. The η map at this crossing, however, is nontrivial:

$$(4-1) \quad \eta(q_6^1) = p_5^1 + p_7^1.$$

The components of the linearized differential give rise to two families of chain complexes:

- Proposition 4.4** (1) $(Q^m, \widehat{\partial}_Q)$ and $(P^m, \widehat{\partial}_P)$ are chain complexes, and
- (2) $\eta: (Q^m, \widehat{\partial}_Q) \rightarrow (P^m, \widehat{\partial}_P)$ is a chain map of degree -1 for $m \geq 1$.

Proof The fact that the degrees of $\widehat{\partial}_Q$, $\widehat{\partial}_P$, and η are all -1 follows from the fact that they are components of the degree -1 differential.

To prove the remainder of the proposition, represent $\widehat{\partial}_1: Q^m \oplus P^m \rightarrow Q^m \oplus P^m$ by the following matrix:

$$\begin{bmatrix} \widehat{\partial}_Q & 0 \\ \eta & \widehat{\partial}_P \end{bmatrix}.$$

Since $\widehat{\partial}^2 = 0$, every component of the following matrix is also zero:

$$\begin{bmatrix} \widehat{\partial}_Q & 0 \\ \eta & \widehat{\partial}_P \end{bmatrix}^2 = \begin{bmatrix} \widehat{\partial}_Q^2 & 0 \\ \eta\widehat{\partial}_Q + \widehat{\partial}_P\eta & \widehat{\partial}_P^2 \end{bmatrix}.$$

This finishes the proof. □

The proof of Proposition 4.1 shows that one of the new chain complexes in the proposition above is familiar:

Corollary 4.5 $(A, \partial_1) \simeq (Q^0, \widehat{\partial}_Q)$

Here is one more corollary of Proposition 4.4:

Corollary 4.6 For $m \geq 1$, $(Q^m \oplus P^m, \widehat{\partial}_1)$ is the mapping cone for the chain map η .

The differential $\widehat{\partial}_Q$ splits further in a manner similar to the original $\widehat{\partial}_1$. For $m \geq 1$, let \overline{Q}^m be the vector space generated by $\{q_i^m\}$ and let C^m be generated by $\{c_i^m, d_i^m\}$. As above, split $\widehat{\partial}_Q$ by restricting domains and ranges:

- (1) $\widehat{\partial}_{\overline{Q}}: \overline{Q}^m \rightarrow \overline{Q}^m$,
- (2) $\widehat{\partial}_C: C^m \rightarrow C^m$, and
- (3) $\rho: \overline{Q}^m \rightarrow C^m$.

By the same argument as in the proof of Lemma 4.2, there is no component of $\widehat{\partial}_Q$ with domain C^m and range Q^m . As in Proposition 4.4, it follows that $(\overline{Q}^m, \widehat{\partial}_{\overline{Q}})$ and $(C^m, \widehat{\partial}_C)$ are both chain complexes and that ρ is a chain map between them. Interestingly, though perhaps not surprisingly given the Morse–Bott motivation for the circle bundle theory from which the expanded algebra is derived, both of these chain complexes are familiar:

Lemma 4.7 $(\overline{Q}^m, \widehat{\partial}_{\overline{Q}}) \simeq (Q^0, \widehat{\partial}_Q)$

Proof There is an obvious correspondence between generators and their degrees, so it suffices to extend this correspondence to the differentials. As demonstrated in the proof of Lemma 4.2, a disk contributing to $\widehat{\partial}_{\overline{Q}}q^m$ must have a positive corner at q^m and negative corners at a q crossing of level m and possibly others at augmented q^0 corners. Thus, under the stick-together map, a thick disk that contributes to $\widehat{\partial}_{\overline{Q}}q^m$

becomes a disk contributing to $\widehat{\partial} Q q^0$. Further, no thin disk can contribute to $\widehat{\partial} \overline{Q}$ since any thin disk with a positive q corner must have either a negative corner at a d (type (3)) or at a p (types (4) and (5)). Overall, then, the differentials on \overline{Q}^m and on Q^0 are constructed out of the same disks, and hence are the same under the correspondence between generators. \square

Lemma 4.8 *Up to shifting degree down by 1, $(C^m, \widehat{\partial}_C)$ is isomorphic to the Morse–Witten complex (C^{MW}, ∂^{MW}) of the circle with respect to $\tilde{g}|_K$.*

Proof It is clear from the construction that there is a one-to-one correspondence between the generators of C^m and those of C^{MW} . As noted in Section 3.2, the differential of a c chord comes completely from thin disks of types (1) or (2). Disks of type (2) have a negative p corner, so cannot contribute to $\widehat{\partial}_C$. Disks of type (1) are in one-to-one correspondence with negative gradient flowlines for g , and hence $\widehat{\partial}_C c^m = \partial^{MW} c^m$. For degree reasons, $\widehat{\partial}_C d^m = 0 = \partial^{MW} d^m$. This completes the proof. \square

4.1.3 “Vertical” relationships There is a “vertical” relationship between any pair of complexes $(Q^m, \widehat{\partial}_Q)$ and $(Q^{m+1}, \widehat{\partial}_Q)$. This relationship can be encoded by a translation map $\tau: Q^m \rightarrow Q^{m+1}$ that raises the level of a generator by 1; a similar map exists for P^m . The maps and spaces defined thus far are pictured in Figure 7. The translation map interacts nicely with the linearized differential:




Figure 7: Relationships between the maps and spaces defined thus far. By Lemma 4.9, the diagram commutes.

Lemma 4.9 *The translation map τ is a chain isomorphism on $(P^m, \widehat{\partial}_P)$ and on $(Q^m, \widehat{\partial}_Q)$ for all $1 \leq m \leq n - 2$. Further, τ commutes with η .*

Proof For all levels in P and for levels greater than zero in Q , the translation map τ is obviously bijective.

To prove that τ is a chain map and commutes with η , consider Lemma 4.2 and the definition of $\widehat{\varepsilon}$. These combine to show that any disk that contributes to the linearized differential of a generator of level m has one negative corner at a crossing between K and $f^m(K)$ and all other negative corners at (augmented) self-crossings of K or of $f^m(K)$. In particular, such disks depend only on the link $K \cup f^m(K)$. For $m \geq 1$, the Lagrangian projections of the links $K \cup f^m(K)$ and $K \cup f^{m+1}(K)$ are combinatorially identical, and hence so are the disks that contribute to $\widehat{\partial}_1$ — and hence to $\widehat{\partial}_Q$, $\widehat{\partial}_P$, and η — at levels m and $m + 1$. \square

More interesting phenomena occur when comparing $(Q^0, \widehat{\partial}_Q)$ to $(Q^1, \widehat{\partial}_Q)$. The chain complex $(Q^1, \widehat{\partial}_Q)$ can be written as follows, with subscripts denoting degrees:

$$\begin{array}{ccccccccccc}
 \xrightarrow{\widehat{\partial}_{\overline{Q}}} & \overline{Q}_2 & \xrightarrow{\widehat{\partial}_{\overline{Q}}} & \overline{Q}_1 & \xrightarrow{\widehat{\partial}_{\overline{Q}}} & \overline{Q}_0 & \xrightarrow{\widehat{\partial}_{\overline{Q}}} & \overline{Q}_{-1} & \xrightarrow{\widehat{\partial}_{\overline{Q}}} & \overline{Q}_{-2} & \xrightarrow{\widehat{\partial}_{\overline{Q}}} & \\
 & & & \searrow \rho & & \searrow \rho & & & & & & \\
 & & & 0 & \longrightarrow & C_0^1 & \xrightarrow{\widehat{\partial}_C} & C_{-1}^1 & \longrightarrow & 0 & &
 \end{array}$$

It follows from Lemma 4.7 that:

$$(4-2) \quad \dim H_k(Q^1, \widehat{\partial}_Q) = \dim H_k(\overline{Q}^1, \widehat{\partial}_{\overline{Q}}) = \dim H_k(Q^0, \widehat{\partial}_Q)$$

for $k > 1$ and $k < -1$. This, together with knowledge of the case $k = -1$, will suffice for the upcoming proof of duality.

In degree -1 , it turns out that

$$(4-3) \quad \dim H_{-1}(Q^1, \widehat{\partial}_Q) = \dim H_{-1}(Q^0, \widehat{\partial}_Q) + 1.$$

On one hand, the kernel of $\widehat{\partial}_Q$ on Q_{-1}^1 is the direct sum of the kernel of $\widehat{\partial}_{\overline{Q}}$ on \overline{Q}_{-1}^1 and all of C_{-1}^1 , the space generated by the d 's. On the other, the image of $\widehat{\partial}_Q$ on Q_0^1 comes from studying the images of $\widehat{\partial}_{\overline{Q}}$, $\widehat{\partial}_C$, and ρ . Lemma 4.8 implies that the c^1 terms have a linearized differential given by

$$(4-4) \quad \widehat{\partial}_Q(c^1) = d^1 + \overline{d}^1,$$

where d and \bar{d} are adjacent to c along the knot. Because of the configuration of positive and negative corners around a d crossing (see Figure 4), only thin disks of type (3) that begin at an augmented crossing contribute to the image of ρ . In fact, by the discussion of how thin disks contribute to the differential $\widehat{\partial}$ in Section 3.2, if q_i^0 is augmented, then $\rho(q_i^1)$ is the sum of the two d^1 generators adjacent to q_i . This is also the image under $\widehat{\partial}_C$ of the sum of the c^1 generators that lie between the two d^1 generators on the knot K .⁴ Thus, it makes sense to consider the following change of basis for Q_0^1 : let β_i be the sum of the c^1 generators that lie between the two d^1 generators whose sum is $\rho(q_i^1)$. Change basis so that if q_i^0 is augmented, then q_i^1 is replaced by $q_i^1 + \beta_i$. In this basis, it is clear that the image of $\widehat{\partial}_Q$ is the direct sum of the image of $\widehat{\partial}_{\overline{Q}}$ on \overline{Q}_0^1 and the image of $\widehat{\partial}_C$. Thus, the degree -1 homology of Q^1 is the direct sum $H_{-1}(\overline{Q}^1, \widehat{\partial}_{\overline{Q}}) \oplus H_0(S^1)$, with the latter summand generated by $[d_1^1] = [d_2^1] = \dots = [d_r^1]$. Equation (4-3) follows.

4.1.4 “Horizontal” relationships There is also a “horizontal” relationship between the chain complexes $(Q^0, \widehat{\partial}_Q)$ and $(P^1, \widehat{\partial}_P)$. Define a degree 1 pairing⁵ on $P^1 \otimes Q^0$ by:

$$\langle p_i, q_j \rangle = \delta_{ij}.$$

This pairing allows us to view $(Q^0, \widehat{\partial}_Q)$ and $(P^1, \widehat{\partial}_P)$ as dual complexes since a disk that contributes q_j to $\widehat{\partial}_Q(q_i)$ also contributes p_i to $\widehat{\partial}_P(p_j)$. Thus, the pairing descends to homology. In particular, by (3-2):

Lemma 4.10 $\dim H_k(Q^0, \widehat{\partial}_Q) = \dim H_{-k-1}(P^1, \widehat{\partial}_P).$

Example 4.11 Combining the lemma with equation (2-2) shows that p_7^1 represents a non-trivial homology class and that p_5^1 is a boundary in $(P^1, \widehat{\partial}_P)$. Thus, since η is a chain map, (4-1) becomes:

(4-5) $\eta_*[q_6^1] = [p_7^1].$

Similarly, it is possible to compute that:

(4-6) $\eta_*[q_7^1] = [p_6^1].$

⁴There are two such sets of c^1 generators, depending on the direction in which the knot K is traversed. Either set works. This is true even if one is the empty set, as this is the case that the two adjacent d^1 generators coincide and hence sum to 0 over $\mathbb{Z}/2$.

⁵This assumes that the base ring lies in degree 0.

Since q_6 has degree -2 and q_7 has degree 2 , these equations indicate that the η map could be the duality map that explains the symmetry of the Poincaré–Chekanov polynomial for the knot in Figure 2.

4.2 Proof of duality

The analogy between contact homology and Morse theory described in Section 1.2 suggests that the duality map for linearized Legendrian contact homology should come from disks with two positive corners. The structure analyzed in the previous section provides the setting in which to make this analogy precise: disks with two positive corners are the defining disks for the η map, as shown in Figure 6. The following proposition demonstrates that the analogy holds in this case:

Proposition 4.12 *The map $\eta: (Q^m, \widehat{\partial}_Q) \rightarrow (P^m, \widehat{\partial}_P)$ is an isomorphism.*

Proof By Corollary 4.6, it suffices to show that $(Q^m \oplus P^m, \widehat{\partial}_1)$ is acyclic. This, in turn, is obvious, since this complex is the same as Γ_{0m} for the unlink consisting of K and the large vertical translate $f^m(K)$. \square

The proof of duality is now a matter of combining Proposition 4.12 with the relationships between the homologies of $(Q^1, \widehat{\partial}_Q)$ and $(Q^0, \widehat{\partial}_Q)$ that were calculated in Section 4.1.

Proof of Theorem 1.1 First suppose that $|k| > 1$:

$$\begin{aligned} \dim H_k(A, \partial_1) &= \dim H_{-k-1}(P^1, \widehat{\partial}_P) && \text{by Lemma 4.10} \\ &= \dim H_{-k}(Q^1, \widehat{\partial}_Q) && \text{by Proposition 4.12} \\ &= \dim H_{-k}(A, \partial_1) && \text{by (4-2)} \end{aligned}$$

The proof of the case where $k = 1$ is exactly the same up until the last step, where (4-3) implies that

$$\dim H_1(A, \partial_1) = \dim H_{-1}(Q^1, \widehat{\partial}_Q) = \dim H_{-1}(A, \partial_1) + 1,$$

as required by the theorem. Finally, $\dim H_0(A, \partial_1)$ does not matter for Theorem 1.1, so the proof is complete. \square

4.3 Invariance

The duality map η was defined using a fixed Lagrangian diagram and perturbing Morse function. Two pairs of Lagrangian diagrams and perturbing functions are related by Reidemeister moves, and it is indeed the case that the duality maps before and after the induced stable tame isomorphism are, in some suitable sense, conjugate. To set notation for a precise statement of invariance, let $(\widehat{\mathcal{A}}, \widehat{\partial})$ and $(\widetilde{\mathcal{A}}, \widetilde{\partial})$ be the expanded algebras of the perturbed Lagrangian diagrams before and after the Reidemeister moves. Assume that both algebras have been appropriately stabilized and that Ψ is a tame isomorphism between them. Let $\widehat{\varepsilon}$ and $\widetilde{\varepsilon}'$ be augmentations that correspond under Ψ in the manner outlined in Section 2.3; recall that there is a map $\overline{\Psi}$ that restricts to a chain isomorphism between the two *linearized* expanded algebras. Further, by Mishachev's work, $\overline{\Psi}$ must respect levels. Finally, write the restriction of $\overline{\Psi}$ to $Q^1 \oplus P^1$ in terms of its components as follows:

$$(4-7) \quad \overline{\Psi} = \begin{bmatrix} \Psi_Q & G \\ H & \Psi_P \end{bmatrix}.$$

With this notation set, the statement of invariance is:

Proposition 4.13 *The maps Ψ_Q and Ψ_P are chain isomorphisms that, on homology, conjugate η_* and η'_* .*

Proof The key to the proof is the fact that the map $G: P^1 \rightarrow (Q')^1$ is zero. This fact comes from examining the tame isomorphisms used to prove invariance under the Reidemeister moves; see [4; 12]. In all cases, the isomorphisms are nondecreasing in height, and hence a generator of $(Q')^m$ cannot appear in the image of a generator of P^m .

With this fact in hand, it follows that the matrix for $\overline{\Psi}$ is lower triangular, which immediately shows that Ψ_Q and Ψ_P are isomorphisms. Further, the fact that $\overline{\Psi}$ is a chain map implies the following three equations:

$$(4-8) \quad \begin{aligned} \widetilde{\partial}'_Q \Psi_Q &= \Psi_Q \widehat{\partial}_Q, \\ \widetilde{\partial}'_P \Psi_P &= \Psi_P \widehat{\partial}_P, \text{ and} \\ \eta' \Psi_Q + \Psi_P \eta &= \widetilde{\partial}'_P H + H \widehat{\partial}_Q. \end{aligned}$$

The first equation shows that Ψ_Q is a chain isomorphism from $(Q^1, \widehat{\partial}_Q)$ to $(Q'^1, \widetilde{\partial}'_Q)$; a similar statement holds for Ψ_P . The last equation shows that, on homology, $\eta'_*(\Psi_Q)_* = (\Psi_P)_* \eta_*$, as desired. \square

5 Duality and the cap product

The expanded algebra defined in the previous sections not only provides the structure with which to prove duality, but also to begin to examine the “algebraic topology” of the linearized Legendrian contact homology. In this section, the proof of duality will be reinterpreted in terms of the classical “capping with the fundamental class” construction.

5.1 The fundamental class

In Morse theory, the fundamental class of a compact manifold is represented by the sum of the maxima of a Morse function. Analogously, the longest Reeb chords — for example, those coming from the rightmost cusps in a resolved front diagram — should give the fundamental class in degree 1 in the linearized DGA. This is not quite accurate, however, as the change of variables given by an augmentation can introduce new terms. One way to overcome this difficulty — and to produce a nontrivial class in degree 1 — was discovered by Chekanov [4, Section 12]: closing a long Legendrian knot induces a one-to-one correspondence between Legendrian isotopy classes of long Legendrian knots and those of Legendrian knots. Chekanov also proved that the mapping from the set of Poincaré–Chekanov polynomials for a long Legendrian knot to the set for its closure defined by

$$P(t) \mapsto P(t) + t$$

is bijective. In particular, there is always a nontrivial homology class in degree 1 for the closure. Chekanov’s construction, however, is not precise enough for the needs of this paper.

The necessary precision is derived from the calculations that established the relationship between $H_{-1}(Q^0, \widehat{\partial}_Q)$ and $H_{-1}(Q^1, \widehat{\partial}_Q)$ in Section 4.1. Recall that the proof of equation (4–3) decomposed $H_{-1}(Q^1, \widehat{\partial}_Q)$ into $H_{-1}(\overline{Q}^1, \widehat{\partial}_{\overline{Q}}) \oplus H_0(S^1)$, with the second factor generated by the class $[d]$. In essence, the following theorem asserts that the fundamental class is dual (with respect to the pairing $\langle \cdot, \cdot \rangle$) to the image of the generator of $H_0(S^1)$ under the Poincaré duality map η .

Theorem 5.1 *There exists a unique class κ , called the fundamental class, in $H_1(A, \partial_1)$ that pairs to 1 with $\eta_*([d])$ and to 0 on the image of $H_{-1}(\overline{Q}^1, \widehat{\partial}_{\overline{Q}})$ under η_* .*

Proof The theorem follows immediately from (4–3) and Proposition 4.12. □

The statement of invariance in Section 4.3 only shows that the fundamental class is defined up to an isomorphism of the expanded algebra for a given Lagrangian diagram of K . The fundamental class, however, lies in the original linearized contact homology of the diagram, so it would be nice if the restriction of the map η to \overline{Q}_{-1}^1 were to be independent of the perturbation function. To see that this is indeed true, consider a q^1 generator of degree -1 . No thin disk can contribute to $\eta(q^1)$: a thin disk of type 3 (resp. 4 or 5) would have corners at a d crossing (resp. a p crossing) and a q^0 crossing of degree -1 , and hence could not contribute to the linearized differential. Since the thick disks depend only on the original Lagrangian diagram, the image of q^1 under η does not depend on the perturbation. As the coefficient ring is $\mathbb{Z}/2$, the image of $H_{-1}(\overline{Q}_{-1}^1, \widehat{\partial}_{\overline{Q}})$ determines the fundamental class. Thus, the restriction of the map η to \overline{Q}_{-1}^1 does not depend on the perturbation and can be read off of the original Lagrangian diagram using thick disks.

The definition of the fundamental class in Theorem 5.1 captures some of the original Morse-theoretic motivation. If a Lagrangian diagram is the resolution of a front diagram, then the longest chords are those coming from the rightmost cusps; see Figure 2. Fix a chord coming from a right cusp. Choose a perturbation so that there is a d chord on the loop to the right of this crossing. The image of d under η is precisely the right cusp chord, leading to the conclusion that every rightmost cusp chord is a summand of any representative of the fundamental class.

Example 5.2 The knot in Figure 2 has a fundamental class consisting purely of the right cusps: $[\kappa] = [q_8 + q_9]$. This is not always the case for resolutions of front diagrams: it is possible to calculate that the figure 8 knot in Figure 8 has a fundamental class of $[\kappa] = [q_5 + q_7 + q_8 + q_9]$ for either of its two augmentations.




Figure 8: The resolution of a front diagram of a Figure 8 knot.

5.2 The cap product

The goal of this section is to use the cap product and the fundamental class to exhibit an inverse to the duality map η_* . The first step in constructing an inverse to η_* is

to define the cap product. As argued in Section 1.2, the cap product for linearized Legendrian contact homology should come from disks with one positive corner and two negative corners. As with the fundamental class, the augmentation interferes with this geometric description, but the length 2 differential, $\widehat{\partial}_2$, encodes the relevant algebraic information.⁶ Geometrically, the length 2 differential counts disks with one positive corner, two negative corners, and possibly other augmented negative corners.

More technically, the length 2 differential $\widehat{\partial}_2$ can be split into components just as $\widehat{\partial}_1$ was in Section 4.1. Write:

$$(5-1) \quad P = \bigoplus_{1 \leq m \leq n-2} P^m \quad \text{and} \quad Q = \bigoplus_{0 \leq m \leq n-2} Q^m.$$

For now, only the components of $\widehat{\partial}_2$ with domain P are of interest; they are:

- (1) $\Phi_{QP}: P \rightarrow Q \otimes P$,
- (2) $\Phi_{PQ}: P \rightarrow P \otimes Q$, and
- (3) $\Phi_{PP}: P \rightarrow P \otimes P$.

That these preserve total level and are the only components follows from the same arguments as for Lemma 4.2. The first of these maps is a chain map into the tensor product $(Q, \widehat{\partial}_Q) \otimes (P, \widehat{\partial}_P)$. To see why, look at the length two component of the equation $\widehat{\partial}^2 = 0$:

$$\widehat{\partial}_1 \widehat{\partial}_2 + \widehat{\partial}_2 \widehat{\partial}_1 = 0,$$

and more specifically, the component of the left hand side that maps P to $Q \otimes P$:

$$(5-2) \quad \Phi_{QP} \widehat{\partial}_P + (\widehat{\partial}_Q \otimes 1 + 1 \otimes \widehat{\partial}_P) \Phi_{QP} = 0.$$

The second map is similarly a chain map.

Pairing the second tensor factor of Φ_{QP} with κ turns this map into the cap product with the fundamental class. More precisely:

Definition 5.3 Let k be a representative of the fundamental class κ . The *evaluation map* $ev: P \rightarrow \mathbb{Z}/2$ is defined by

$$ev(p) = \langle p, k \rangle$$

on P^1 and is extended to the rest of P by 0.

The *cap product map* $\phi: P \rightarrow Q$ is defined by

$$\phi = (1 \otimes ev) \Phi_{QP}.$$

⁶Recall that the superscript $\widehat{\varepsilon}$ has been dropped from the notation for the augmented differential.

The only component of Φ_{QP} involved in the definition of ϕ is $Q^{m+1} \rightarrow Q^m \otimes P^1$, so ϕ reduces level by 1. The maps and spaces defined thus far can be organized into the diagram in Figure 9, an expansion of Figure 7.




Figure 9: Relationships between the maps and spaces defined in Sections 4.1 and 5.2.

Lemma 5.4

- (1) *The evaluation map descends to homology and is independent of the choice of the representative of κ .*
- (2) *The cap product map ϕ is a chain map of degree 1.*

Proof The first part of the lemma stems from the following computation, using Lemma 4.10:

$$\begin{aligned}
 \text{ev} \circ \widehat{\partial}_P(p) &= \langle \widehat{\partial}_P(p), \kappa \rangle \\
 (5-3) \qquad \qquad &= \langle p, \widehat{\partial}_Q(\kappa) \rangle \\
 &= 0.
 \end{aligned}$$

The fact that ϕ is a chain map follows from the fact that it is the composition of two chain maps (the evaluation map is a chain map into the trivial complex $(\mathbb{Z}/2, 0)$). The fact that ϕ is a degree 1 map follows from the facts that Φ_{QP} has degree -1 (it is derived from $\widehat{\partial}$) and ev has degree 2 (it is nonzero only on degree -2 generators in P^1 dual to the degree 1 class κ). □

The Morse theoretic motivation suggests that ϕ_* inverts η_* . This is indeed the case, up to an application of τ_* .

Proposition 5.5 The inverse of $\eta_*: H_*(Q^1, \widehat{\partial}_Q) \rightarrow H_*(P^1, \widehat{\partial}_P)$ is $\tau_*\phi_* = \phi_*\tau_*$.

Proof It suffices to find chain homotopies $H: P^1 \rightarrow P^1$ and $K: Q^1 \rightarrow Q^1$ such that

$$(5-4) \quad \eta\phi\tau + \iota_P = H\widehat{\partial}_P + \widehat{\partial}_P H$$

and

$$(5-5) \quad \phi\eta\tau + \iota_Q = K\widehat{\partial}_Q + \widehat{\partial}_Q K,$$

where $\iota_P: P^1 \rightarrow P^1$ and $\iota_Q: Q^1 \rightarrow Q^1$ are the identity maps.

First, define H by:

$$H = (1 \otimes \text{ev})\Phi_{PP}\tau.$$

Consider the $P^2 \rightarrow P^1 \otimes P^1$ component of $\widehat{\partial}^2 = 0$:

$$\Phi_{PP}\widehat{\partial}_P + (\widehat{\partial}_P \otimes 1 + 1 \otimes \widehat{\partial}_P)\Phi_{PP} + (\eta \otimes 1)\Phi_{QP} + (1 \otimes \eta)\Phi_{PQ} = 0.$$

Precomposing with τ , post-composing with $(1 \otimes \text{ev})$, and using Lemmas 4.9 and 5.4 yields

$$H\widehat{\partial}_P + \widehat{\partial}_P H + \eta\phi\tau + (1 \otimes \text{ev} \eta)\Phi_{PQ}\tau = 0.$$

Let $\iota_P = (1 \otimes \text{ev} \eta)\Phi_{PQ}\tau$; this is clearly a chain map since it is a composition of chain maps. Further, since η acts trivially on Q^0 , only the component of Φ_{PQ} with image in $P^1 \otimes Q^1$ contributes to the definition of ι_P . It now suffices to prove that ι_P is the identity map.

Let $p \in P^2$ represent a homology class. Thick disks contribute terms of the form $p' \otimes q$ to $\Phi_{PQ}\tau(p)$, where, by the Künneth theorem, p' represents a homology class in $H_*(P^1, \widehat{\partial}_P)$ and q represents a homology class in $H_*(Q^1, \widehat{\partial}_Q)$. In fact, since a d generator cannot appear as the negative corner of a thick disk, q must represent a class in $H_*(\overline{Q}^1, \widehat{\partial}_{\overline{Q}})$. It follows from Theorem 5.1 that $\text{ev} \eta(q) = 0$, and hence that thick disks not contribute to ι_P . On the other hand, there is always exactly one thin disk that contributes a term of the form $p \otimes d$ to $\Phi_{PQ}\tau(p)$. Since Theorem 5.1 implies that $\text{ev} \eta(d) = 1$, it follows that $\iota_P(p) = p$.

The definition of K and the proof of (5-5) are almost identical to the above. To define K , let $\Psi_{QP}: Q^2 \rightarrow Q^1 \otimes P^1$ be a component of $\widehat{\partial}_2$. Then let

$$(5-6) \quad K = (1 \otimes \text{ev})\Psi_{QP}\tau. \quad \square$$

Example 5.6 To illustrate the theorem, return to the original observation of duality for the knot in Figure 2. As noted at the end of Section 4.1, $\eta_*[q_6^1] = [p_7^1]$, the latter of which is dual to $[q_7^0]$ with respect to the pairing $\langle \cdot, \cdot \rangle$. How does ϕ_* behave here?

A computation similar to that in Example 3.4 shows that

$$\widehat{\partial}_2 p_7^2 = q_6^1 p_9^1 + \text{terms not in } Q^1 \otimes P^1.$$

Since $\kappa = [q_8 + q_9]$, the first term contributes non-trivially to ϕ :

$$(5-7) \quad \phi_*[p_7^1] = [q_6^1].$$

This illustrates that the cap product map ϕ_* is the inverse of the duality map η_* .

References

- [1] **D Bennequin**, *Entrelacements et équations de Pfaff*, from: “Third Schnepfenried geometry conference, Vol. 1 (Schnepfenried, 1982)”, Astérisque 107, Soc. Math. France, Paris (1983) 87–161 MR753131
- [2] **M Betz, R L Cohen**, *Graph moduli spaces and cohomology operations*, Turkish J. Math. 18 (1994) 23–41 MR1270436
- [3] **F Bourgeois**, *A Morse-Bott approach to contact homology*, from: “Symplectic and contact topology: interactions and perspectives (Toronto, ON/Montreal, QC, 2001)”, Fields Inst. Commun. 35, Amer. Math. Soc. (2003) 55–77 MR1969267
- [4] **Y Chekanov**, *Differential algebra of Legendrian links*, Invent. Math. 150 (2002) 441–483 MR1946550
- [5] **F Ding, H Geiges**, *Legendrian knots and links classified by classical invariants* arXiv: math.SG/0503033
- [6] **T Ekhholm, J Etnyre, M Sullivan**, *Orientations in Legendrian contact homology and exact Lagrangian immersions*, Internat. J. Math. 16 (2005) 453–532 MR2141318
- [7] **Y Eliashberg**, *Invariants in contact topology*, from: “Proceedings of the International Congress of Mathematicians, Vol. II (Berlin, 1998)”, Doc. Math. Extra Vol. II (1998) 327–338 MR1648083
- [8] **Y Eliashberg, M Fraser**, *Classification of topologically trivial Legendrian knots*, from: “Geometry, topology, and dynamics (Montreal, PQ, 1995)”, CRM Proc. Lecture Notes 15, Amer. Math. Soc., Providence, RI (1998) 17–51 MR1619122
- [9] **Y Eliashberg, A Givental, H Hofer**, *Introduction to symplectic field theory*, Geom. Funct. Anal. (2000) 560–673 MR1826267GAFA 2000 (Tel Aviv, 1999)
- [10] **J B Etnyre**, *Legendrian and transversal knots*, from: “Handbook of knot theory”, Elsevier (2005) 105–185 MR2179261

- [11] **JB Etnyre, K Honda**, *Knots and contact geometry. I. Torus knots and the figure eight knot*, *J. Symplectic Geom.* 1 (2001) 63–120 MR1959579
- [12] **JB Etnyre, LL Ng, JM Sabloff**, *Invariants of Legendrian knots and coherent orientations*, *J. Symplectic Geom.* 1 (2002) 321–367 MR1959585
- [13] **D Fuchs**, *Chekanov-Eliashberg invariant of Legendrian knots: existence of augmentations*, *J. Geom. Phys.* 47 (2003) 43–65 MR1985483
- [14] **D Fuchs, T Ishkhanov**, *Invariants of Legendrian knots and decompositions of front diagrams*, *Mosc. Math. J.* 4 (2004) 707–717, 783 MR2119145
- [15] **K Fukaya, Y-G Oh**, *Zero-loop open strings in the cotangent bundle and Morse homotopy*, *Asian J. Math.* 1 (1997) 96–180 MR1480992
- [16] **T Kálmán**, *Contact homology and one parameter families of Legendrian knots*, *Geom. Topol.* 9 (2005) 2013–2078 MR2209366
- [17] **P Melvin, S Shrestha**, *The nonuniqueness of Chekanov polynomials of Legendrian knots*, *Geom. Topol.* 9 (2005) 1221–1252 MR2174265
- [18] **K Mishachev**, *The N -copy of a topologically trivial Legendrian knot*, *J. Symplectic Geom.* 1 (2003) 659–682 MR2039159
- [19] **LL Ng**, *Computable Legendrian invariants*, *Topology* 42 (2003) 55–82 MR1928645
- [20] **JM Sabloff**, *Invariants for Legendrian Knots from Contact Homology* in preparation
- [21] **JM Sabloff**, *Invariants of Legendrian knots in circle bundles*, *Commun. Contemp. Math.* 5 (2003) 569–627 MR2003211
- [22] **JM Sabloff**, *Augmentations and rulings of Legendrian knots*, *Int. Math. Res. Not.* (2005) 1157–1180 MR2147057
- [23] **M Schwarz**, *Morse homology*, *Progress in Mathematics* 111, Birkhäuser, Basel (1993) MR1239174

Haverford College, Haverford, PA 19041, USA

jsabloff@haverford.edu

Proposed: Yasha Eliashberg
Seconded: Leonid Polterovich, Joan Birman

Received: 7 September 2005
Revised: 29 July 2006

